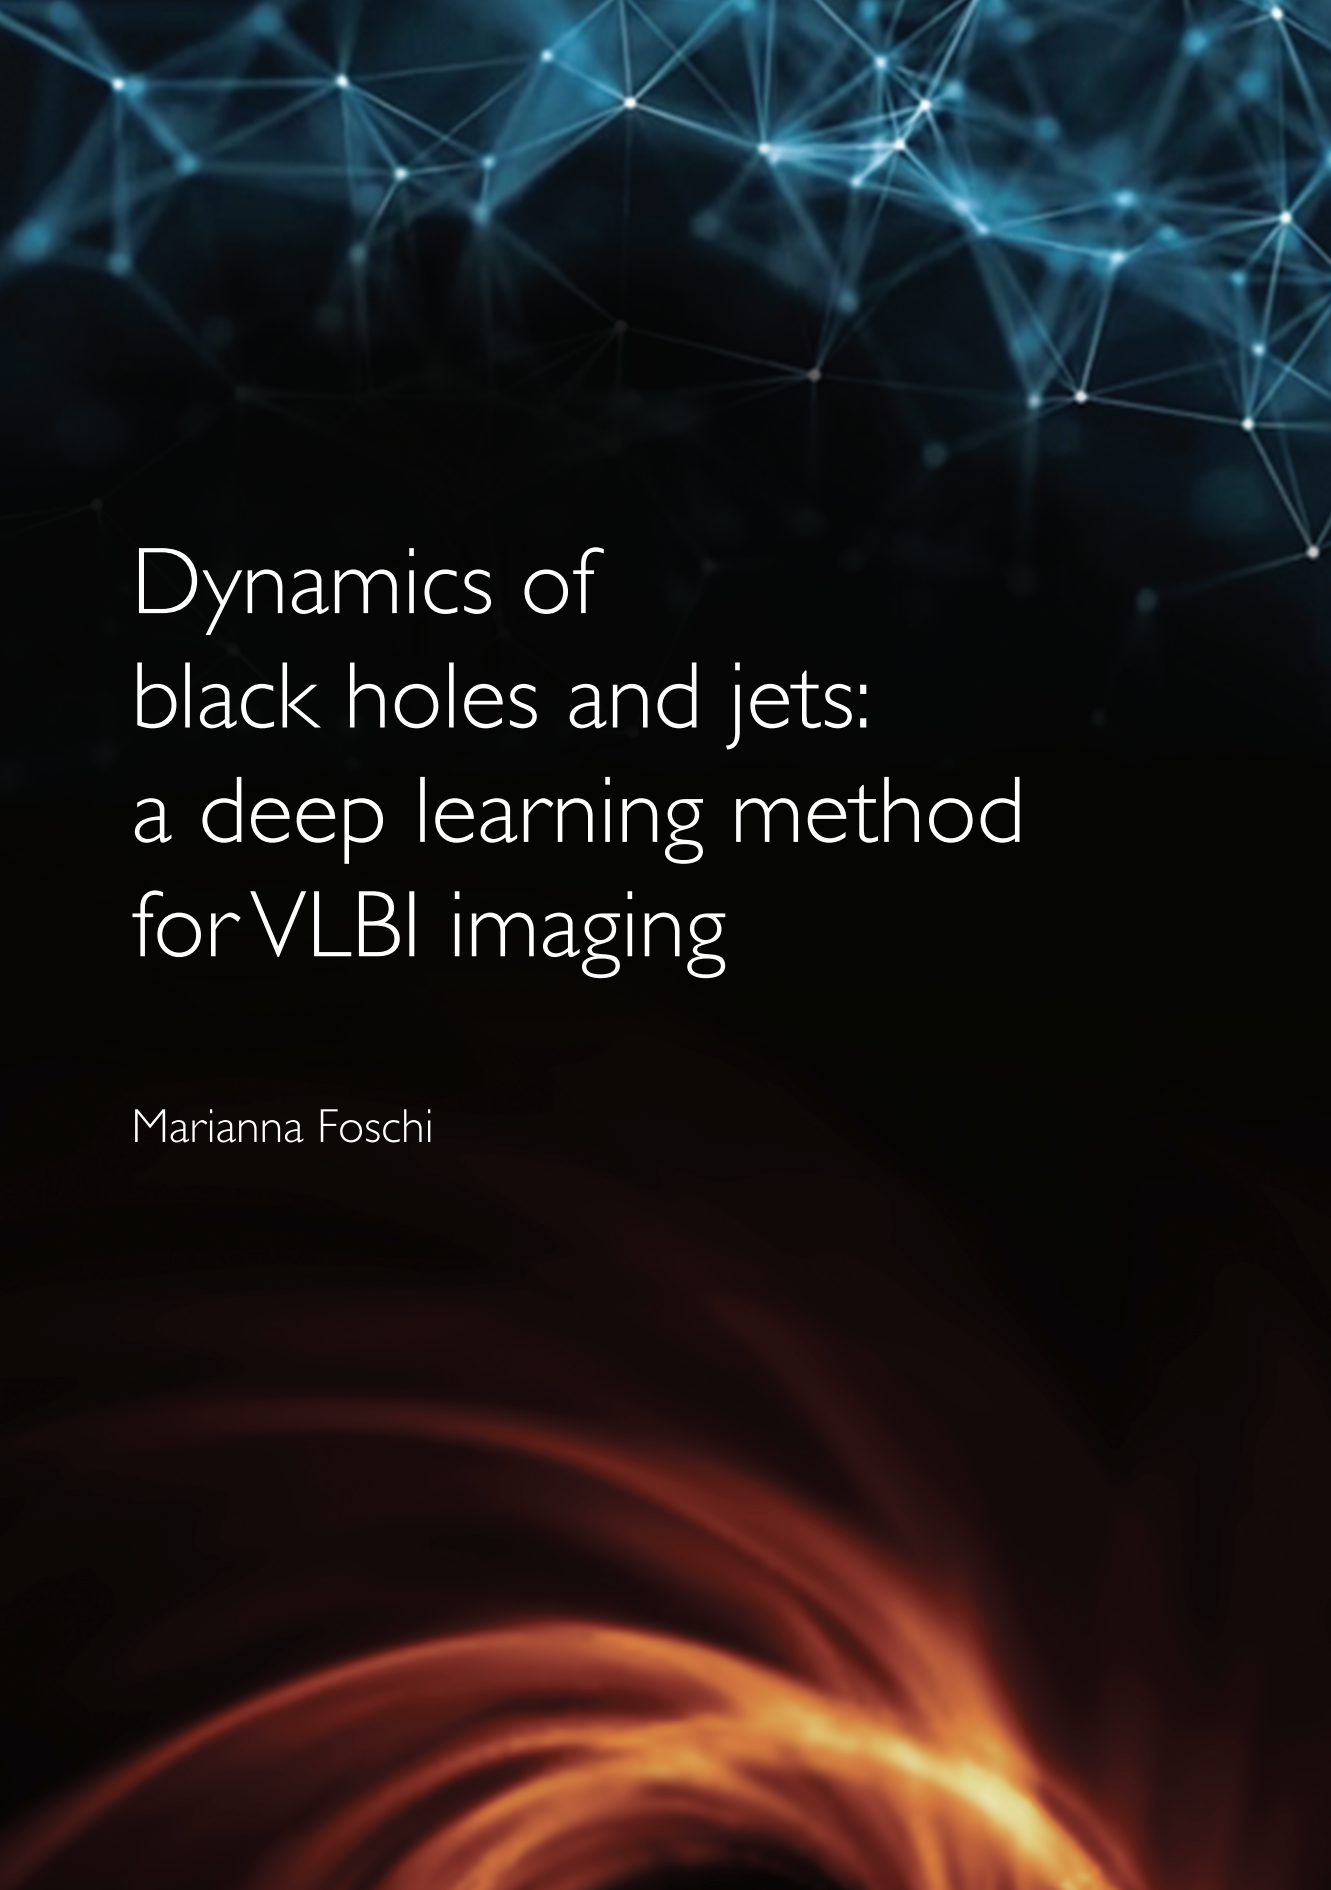

Dynamics of black holes and jets: a deep learning method for VLBI imaging

Marianna Foschi

Dynamics of black holes and jets: a deep learning method for VLBI imaging

Marianna Foschi




Dynamics of black holes and jets:
a deep learning method for VLBI imaging

by
Marianna Foschi

supervised by
Dr. José Luis Gómez Fernández

Institute for Astrophysics of Andalucía
University of Granada
Granada, Spain

September 10, 2025

A thesis submitted in partial fulfillment of the requirements
for the Doctoral Program in Physics and Space Sciences

You taught me the courage of stars before you left
How light carries on endlessly, even after death
With shortness of breath you explained the infinite
And how rare and beautiful it is to even exist
I couldn't help but ask for you to say it all again
I tried to write it down, but I could never find a pen
I'd give anything to hear you say it one more time
That the universe was made just to be seen by my eyes.

— Sleeping at Last, *Jupiter*

Perdere tempo con il cielo, farlo di lavoro,
pagati per immaginare qualcosa che non puoi fotografare.

— Zen Circus, *L'amore è una dittatura*

Abstract

This thesis focuses on the development of imaging and analysis techniques to study the dynamics of black hole accretion and jet launching. At the center of active galactic nuclei, supermassive black holes are surrounded by disks of accreting material and eject collimated, highly relativistic jets of plasma that extend far beyond the host galaxy. Despite extensive observational and theoretical efforts, several fundamental questions about these systems remain unresolved. At horizon scales, uncertainties persist regarding the dominant accretion regime in observed sources, the physical mechanisms responsible for jet launching, the origin of flaring events and variability, and the measurement of specific black hole properties such as the spin. At the larger scales of relativistic jets, we still need to understand the complex details of the plasma fluid propagation through the jet, and its interactions with the interstellar medium.

From an experimental point of view, the answer to these questions can be researched through observations of the radio emission originating from active galactic nuclei. Because of the astronomical distances at which these objects are located, only very long baseline interferometry (VLBI) provides the angular resolution necessary to image the lensed emission around black holes and distinguish the details of jet formation and propagation. Numerous imaging methods have been developed in recent years to achieve robust images at super-resolution, but accretion and jet propagation are intrinsically dynamic processes, whose properties cannot be fully captured by static imaging only. Until recently, there has been a lack of dedicated effective imaging and analysis techniques capable of robustly resolving and measuring the temporal variability of these sources at horizon or intra-jet scales. In this thesis we set out to tackle variability studies of relativistic jets and accreting black holes in three steps of increasing complexity, which coincide with the three main chapters of the thesis.

The first work that we present consists in the application of a novel Regularized Maximum Likelihood (RML) imaging method to a decade of multi-epoch observations of a parsec scale jet. The super-resolution achieved by the imaging method allowed the measurement of the instantaneous jet expansion speed and the irregular precession of its core. However the discontinuities from one image to the next prevented further analysis of the jet dynamics. The second work consists in the development of a dynamic imaging pipeline capable of reconstructing time-continuous videos from multiple repeated observations of a slowly varying source or from a single observation of a source with intra-day variability. The method, named `kine`, was applied to almost three decades of observations of the 3C 345 relativistic jet, producing a smooth video, from which it was possible to compute an instantaneous pixel-by-pixel map of the plasma velocity in the jet. While the `kine` method proved useful for dynamic imaging of multi-epoch observations, it was originally developed to recover a video of the Sgr A* black hole at the galactic center, which presents significant time variability over the course of a single observation. Therefore, the third work presented in this thesis focuses on the validation of the `kine` imaging pipeline on realistic synthetic data,

simulated with the noise and the extremely sparse coverage of the horizon-scale observations of Sgr A* with the Event Horizon Telescope (EHT). A successful validation of the method means that the imaging algorithm can be reliably applied by the EHT Collaboration to obtain the first horizon-scale video of a black hole. In the following paragraphs we provide a more detailed introduction to each work.

In Chapter 3, we employ the forward imaging algorithm `eht-imaging` to image 12 years of roughly monthly observation of the 3C 84 relativistic jet at 43 GHz. Observations were conducted with the Very Long Baseline Array (VLBA) by the BEAM-ME monitoring program between 2010 and 2023, for a total of 121 images. Thanks to the super-resolving power of the imaging algorithm, which reaches an effective resolution a factor of 2-3 better compared to CLEAN, we were able to distinguish fine structures of the jet morphology, such as the limb brightening, the twisting of the core, multiple traveling hot spots, and the propagating edge of the jet. We measured the instantaneous expansion speed of the jet over time, identifying three different expansion regimes, marked by a Fanaroff-Riley I (FR I) to FR II morphological transition, a hot spot frustration phase, and an opposite FR II to FR I transition. For every epoch, we fitted the shape of the two jet limbs, from which we recovered the winding jet axis and the projected jet launching direction. Tracking the latter over time we found an irregular variation of the jet core orientation. Overall the results of the work confirms previous studies of the morphological transition undergone by 3C 84 adding details regarding the transition and the hot spot frustration. More importantly, the analysis provided quantitative instantaneous measurements of the jet's expansion speed and irregular precession over a decade-long timescale.

In Chapter 4, we present `kine`, a newly developed video reconstruction algorithm for VLBI observations. The method is based on a neural representation, which is able to process simultaneously all observations available, while learning and leveraging the spatio-temporal correlations existent in the data in full polarization. We applied `kine` to 27 years of 15 GHz VLBA observations of the 3C 345 blazar from the MOJAVE program, obtaining a full polarization video with resolution and dynamic range significantly greater than what is achievable with frame-by-frames imaging using RML methods or other super-resolving methods. The total intensity video shows a highly variable jet that exhibits non-periodic changes in its launching direction. The polarization field structure indicates the presence of an evolving helical magnetic field threading the jet. The continuity and resolution of the video enabled the recovery of the variable projected velocity field of the jet plasma, using an optical flow method, in contrast to previous kinematic studies which could only measure the pattern speed of broad components through Gaussian model fitting. We find that the speed of the bright components traveling through the jet is of the same order of the average plasma speed, which indicates that the local brightness increases are not moving shocks, as previously proposed, but more likely regions with increased emissivity in a turbulent flow.

In Chapter 5, we present an extended version of the `kine` imaging method, developed to address the specific challenges posed by EHT observations of the supermassive black hole Sgr A*. The imaging pipeline is validated on an extensive suite of synthetic data generated after EHT observations of Sgr A* on 2017 April 11, with the aim of assessing `kine`'s ability to reconstruct the ground-truth dynamics of the models

with the extremely sparse coverage of EHT observations. The synthetic data were generated from static and dynamic geometric models, as well as more complex and realistic general-relativistic magneto-hydrodynamic black hole simulations. Realistic noise was introduced in the data, including thermal noise, gain corruption and interstellar scattering. `kine` passes all the validation tests successfully, proving its ability to correctly reconstruct different morphologies in both total intensity and polarization, and recover a wide range of possible motion and variability. The reconstructions also recover important physical quantities, such as the speed of orbiting features, the brightness position angle of the ring shadow, and the linear polarization orientation. This extensive validation of the imaging method is part of a larger effort by the EHT Collaboration to recover a video of the lensed emission surrounding Sgr A*. Given the successful validation, `kine` is the main pipeline that is being used to reconstruct the first video of a supermassive black hole.

Resumen

Esta tesis se centra en el desarrollo de técnicas de imagen y análisis para estudiar la dinámica de la acreción de agujeros negros y el lanzamiento de chorros relativistas. En el centro de los núcleos activos de galaxia, los agujeros negros supermasivos están rodeados por discos de material en acreción y eyectan chorros de plasma colimados y altamente relativistas que se extienden mucho más allá de la galaxia que los origina. A pesar de los grandes esfuerzos teóricos y observacionales, aún quedan por resolver varias cuestiones fundamentales sobre estos sistemas. En la escala del horizonte de sucesos, persisten incertidumbres sobre el régimen de acreción dominante en las fuentes observadas, los mecanismos físicos responsables del lanzamiento de chorros, el origen de los fenómenos de flare y la variabilidad, y la medición de propiedades específicas de los agujeros negros como el espín. En la escala de los chorros relativistas, aún necesitamos comprender los detalles de la propagación del plasma a través del chorro y sus interacciones con el medio interestelar.

Desde un punto de vista experimental, la respuesta a estas preguntas puede investigarse mediante observaciones de la emisión de radio procedente de núcleos activos de galaxia. Debido a las distancias astronómicas a las que se encuentran estos objetos, sólo la interferometría de muy larga línea de base (VLBI) proporciona la resolución angular necesaria para obtener imágenes de la emisión curvada alrededor de los agujeros negros y distinguir los detalles de la formación y propagación de los chorros. En los últimos años se han desarrollado numerosos métodos de obtención de imágenes para conseguir imágenes robustas con superresolución. Sin embargo la acreción y la propagación de chorros son procesos intrínsecamente dinámicos, cuyas propiedades no pueden captarse por completo sólo con imágenes estáticas. Hasta hace poco, se carecía de técnicas eficaces de imagen y análisis capaces de resolver y medir con solidez la variabilidad temporal de estas fuentes a escalas del horizonte o intrachorro. En esta tesis nos proponemos abordar el estudio de variabilidad de los chorros relativistas y los agujeros negros en acreción en tres pasos de complejidad creciente, que coinciden con los tres capítulos principales de la tesis.

En el primer trabajo que presentamos, aplicamos un novedoso método de imagen basado en probabilidad máxima regularizada (RML, por su siglas en inglés) para obtener imágenes de un chorro a escala de un parsec, a partir de datos tomados a lo largo de una década. La superresolución alcanzada por el método de imagen permitió medir la velocidad de expansión instantánea del chorro y la precesión irregular de su núcleo. Sin embargo, las discontinuidades entre una imagen y la siguiente impidieron un análisis más profundo de la dinámica del chorro. El segundo trabajo consiste en el desarrollo de un algoritmo capaz de reconstruir vídeos continuos en el tiempo a partir de múltiples observaciones de una fuente que varía lentamente o a partir de una única observación de una fuente con variabilidad más rapida que el tiempo de observación. El método, denominado `kine`, se aplicó a casi tres décadas de observaciones del chorro relativista 3C 345, produciendo un vídeo continuo, a partir del cual fue posible calcular

un mapa instantáneo píxel a píxel de la velocidad del plasma en el chorro. Aunque el método `kine` resultó útil para la obtención de imágenes dinámicas de observaciones multi-epoca, se desarrolló originalmente para reconstruir un vídeo del agujero negro Sgr A* en el centro galáctico, que presenta una variabilidad temporal significativa en el transcurso de una única observación. Finalmente, el tercer trabajo presentado en esta tesis se centra en la validación del algoritmo `kine` con datos sintéticos realistas, simulados con el ruido y la cobertura extremadamente escaso de las observaciones a escalas del horizonte de sucesos de Sgr A* con el Event Horizon Telescope (EHT). La validación con éxito del método significa que el algoritmo puede ser aplicado con fiabilidad por la Colaboración EHT para obtener el primer vídeo a escalas del horizonte de sucesos de un agujero negro. En los párrafos siguientes ofrecemos una introducción más detallada de cada trabajo.

En el Capítulo 3, empleamos el algoritmo `eht-imaging` para obtener imágenes de 12 años de observación aproximadamente mensual del chorro relativista 3C 84 a 43 GHz. Las observaciones se llevaron a cabo con el Very Long Baseline Array (VLBA) por el programa de monitorizado BEAM-ME entre 2010 y 2023, resultando en un total de 121 imágenes. Gracias al poder de superresolución del algoritmo, que alcanza una resolución efectiva un factor de 2-3 mejor en comparación con CLEAN, pudimos distinguir estructuras con más detalle de la morfología del chorro, como el brillo del borde, la torsión del núcleo, múltiples puntos brillantes itinerantes y el borde de propagación del chorro. Medimos la velocidad de expansión instantánea del chorro a lo largo del tiempo, identificando tres regímenes de expansión diferentes, marcados por una transición morfológica de Fanaroff-Riley I (FR I) a FR II, una fase de frustración del punto caliente, y una transición opuesta de FR II a FR I. Para cada época, ajustamos la forma de las dos extremidades del chorro, a partir de las cuales recuperamos el eje sinuoso del chorro y la dirección de lanzamiento del chorro proyectada sobre el plano del cielo. Al realizar un seguimiento de esta última a lo largo del tiempo, observamos una variación irregular de la orientación del núcleo del chorro. En conjunto, los resultados del trabajo confirman estudios anteriores sobre la transición morfológica experimentada por 3C 84 , añadiendo detalles sobre la transición y la frustración del punto brillante. Además queremos destacar que este análisis proporcionó medidas cuantitativas instantáneas de la velocidad de expansión del chorro y de su precesión irregular a lo largo de una década.

En el Capítulo 4, presentamos `kine`, un algoritmo de reconstrucción de vídeo que hemos desarrollado para observaciones VLBI. El método se basa en una representación neuronal, que es capaz de procesar simultáneamente todas las observaciones disponibles, a la vez que aprende y aprovecha las correlaciones espacio-temporales existentes en los datos en polarización completa. Hemos aplicado `kine` a 27 años de observaciones con el VLBA a 15 GHz del blazar 3C 345 dentro del programa MOJAVE, obteniendo un vídeo en polarización completa con una resolución y un rango dinámico significativamente superiores a los que se pueden conseguir con imágenes individuales obtenidas con métodos RML u otros métodos capaces de superresolución. El vídeo en intensidad total muestra un chorro altamente variable que exhibe cambios no periódicos en su dirección de lanzamiento. La estructura del campo de polarización indica la presencia de un campo magnético helicoidal en evolución que enhebra el chorro. La continuidad y resolución del vídeo han permitido recuperar el campo de velocidad instántanea

del plasma del chorro, instántenea, proyectada sobre el plano del cielo, utilizando un método de flujo óptico. Esto se contrapone a estudios cinemáticos previos que sólo podían medir la velocidad del patrón de ciertas componentes mediante el ajuste de modelos Gaussianos. Encontramos que la velocidad de las componentes brillantes que viajan a través del chorro es del mismo orden que la velocidad media del plasma, lo que indica que los aumentos locales de brillo no son choques en movimiento, como se había propuesto anteriormente, sino más bien regiones con mayor emisividad en un flujo turbulento.

En el Capítulo 5, presentamos una versión ampliada del algoritmo de reconstrucción de vídeo `kine`, desarrollado para abordar los retos específicos planteados por las observaciones del EHT del agujero negro supermasivo Sgr A*. La pipeline se valida en un amplio conjunto de datos sintéticos generados siguiendo las observaciones del EHT de Sgr A* el 11 de abril de 2017, con el objetivo de evaluar la capacidad de `kine` para reconstruir la dinámica real de los modelos con el cubrimiento extremadamente escaso de las observaciones del EHT. Los datos sintéticos se generaron a partir de modelos geométricos estáticos y dinámicos, así como de simulaciones magnetohidrodinámicas en relatividad general de agujeros negros más complejas y realistas. Se introdujeron condiciones de ruido realistas en los datos, incluyendo ruido térmico, ganancias en amplitud y en fase y scatteringinterestelar. `kine` supera con éxito todas las pruebas de validación, demostrando su capacidad de reconstruir correctamente diferentes morfologías tanto en intensidad total como en polarización, y recuperar una amplia gama de posibles dinámicas y variabilidad. Las reconstrucciones también recuperan importantes cantitades físicas, como la velocidad de los ondas espirales y plasmoides, el ángulo de posición del plasma que rodea el horizonte de sucesos y la orientación de la polarización lineal. Esta extensa validación del algoritmo forma parte de un esfuerzo mayor de la Colaboración EHT para reconstruir un vídeo de la emisión que rodea a Sgr A*. Dado el éxito de la validación, `kine` es el algoritmo principal que se está utilizando para reconstruir el primer vídeo de un agujero negro supermasivo.

Contents

List of Figures

List of Tables

1

Introduction

1.1 Supermassive black holes

In this section, we present both the theoretical foundations and observational evidence related to astrophysical black holes. Section 1.1.1 provides an overview of the fundamental concepts and definitions of black holes as predicted by General Relativity, following and adapting Carroll (2019). Section 1.1.2 discusses the astrophysical processes responsible for black hole formation, along with the current experimental evidence and direct observations supporting their existence.

1.1.1 The Schwarzschild and Kerr metrics

General Relativity (GR) is the theory that describes gravitation and, together with the Standard Model for particle physics, is one of the two major theories that describe the laws of the physical world. In GR, gravity is not described as a force, but rather it arises as a consequence of the spacetime curvature caused by presence of mass and energy. The spacetime manifold is described by the *metric tensor* $g_{\mu\nu}$, while energy and matter are represented in the *stress-energy tensor* $T_{\mu\nu}$. General Relativity relates these two tensors through Einstein's field equations:

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu} \ , \tag{1.1}$$

where the Ricci curvature tensor $R_{\mu\nu}$ and and its trace, the Ricci scalar R, are functions of the metric tensor and its derivatives and they contain information about the spacetime curvature. The other terms appearing in the equations are the gravitational constant G and the speed of light c. For a spherically symmetric and static spacetime, Einstein's field equations are solved in vacuum ($T_{\mu\nu} = 0$) by the Schwarzschild metric:

$$ds^2 = -\left(1 - \frac{2GM}{c^2 r}\right)c^2 dt^2 + \left(1 - \frac{2GM}{c^2 r}\right)^{-1} dr^2 + r^2 d\theta^2 + r^2 \sin^2\theta d\phi^2 \ , \tag{1.2}$$

written here is spherical coordinates (t, r, θ, ϕ), which describes the spacetime outside a spherical object of mass M.

At a radius $R_s = 2GM/c^2$, denominated the *Schwarzschild radius*, the metric presents a coordinate singularity, where its tt component vanishes and the rr component diverges. This is not a physical singularity because it can be eliminated by a change of coordinates, for example employing Eddington–Finkelstein or Kruskal–Szekeres coordinates. However, by choosing the right set of coordinates it can be demonstrated that inside the Schwarzschild radius, all future-directed spacetime trajectories proceed in the direction of diminishing radius. This means that any particle located inside, including photons, can't help but fall towards the central mass. A surface like the $r = R_s$ hypersurface, past which particles can never escape to infinity, is called an *event horizon*. In the case of celestial bodies such as stars or planets, the Schwarzschild radius is located inside the object, where the spacetime is not vacuum and therefore the Schwarzschild metric doesn't hold and no event horizon arises. However, astrophysical events such as the collapse of high-mass star can lead to a compression of the object's mass inside R_s, leading to the formation of an event horizon. Such as object is called a *black hole* (BH).

In addition to the event horizon, a Schwarzschild black hole is characterized by the presence of a singularity at $r = 0$ where the spacetime curvature becomes infinite. This is indeed a physical singularity because, not only metric terms, but also physically meaningful quantities such as the Ricci scalar diverge. General Relativity is not able to describe what happens at a singularity. This is an important long-standing theoretical problem, but from an experimental perspective it doesn't represent an issue because nothing that is inside an black hole event horizon can be observed at an external point of view and, according to the weak cosmic censorship conjecture, singularities cannot appear outside of event horizons (Penrose 1965).

Differently from Newtonian gravity, where massless particles such as photons follow straight trajectories, in GR, photons follow geodesic trajectories that may bend according to the spacetime curvature, an effect called *gravitational lensing*. An object described by a Schwarzschild metric allows bound geodesic orbits for both massive and massless particles. For massless particles such as photons, an unstable circular orbit exists at a radius of $R_{\rm ps} = \frac{3GM}{c^2}$, which is located outside the event horizon. A photon with the right direction of the angular momentum may orbit endlessly around this orbit, but any small deviation may cause the photon to either escape towards infinity or fall towards the event horizon. The sphere defined by $r = R_{\rm ps}$ is called the *photon sphere*. Light emitted inside the photon sphere but outside of the horizon may escape to infinity depending on its direction, while any light crossing the photon sphere from outside will inevitably fall into the event horizon.

If the requirement of spherical symmetry is relaxed into axial symmetry and we consider stationary instead of static metrics, the most general solution to Einstein's field equations is the Kerr metric:

$$\begin{aligned} ds^2 = &- \left(1 - \frac{2GMr}{c^2\rho^2}\right) dt^2 - \frac{4GMar\sin^2\theta}{c^3\rho^2} dt d\phi + \frac{\rho^2}{\Delta} dr^2 + \rho^2 d\theta^2 + \\ &+ \left(r^2 + a^2 + \frac{2GMa^2 r \sin^2\theta}{c^2\rho^2}\right) \sin^2\theta d\phi^2 \,, \end{aligned} \tag{1.3}$$

with:

$$\begin{aligned} a &:= J/Mc \,, \\ \rho^2 &:= r^2 + a^2 \cos^2\theta \,, \\ \Delta &:= r^2 - \frac{2GM}{c^2} r + a^2 \,, \end{aligned} \tag{1.4}$$

which is written here in elliptical (Boyer-Lindquist) coordinates. This metric describes a rotating black hole, parametrized by its mass M and angular momentum J (or specific angular momentum a) and reduces to the Schwarzschild metric for $a \to 0$. The event horizon of the Kerr black hole is the surface for which $\Delta = 0$, which happens at $r_\pm = GM/c^2 \pm \sqrt{G^2M^2/c^4 - a^2}$. This means that the Kerr black holes has two event horizons at r_+ and r_-. When an infalling object crosses the outer horizon it is forced to move towards decreasing radii, similarly to the Schwarzschild case. Inside the inner horizon this is reversed, and an object at $r < r_-$ can move towards increasing r, causing a possible accumulation of matter around the inner horizon. A singularity also arises in the Kerr black hole at $\rho = 0$, which happen when $r = 0$ and $\phi = 0$. The effect of the black hole rotation on the singularity is thus to stretch the singularity point into a disk.

Another important difference with respect to the static black hole is the presence of a non-null metric cross term $g_{t\phi}$ between the time coordinate and a spatial one. This results in the phenomenon of frame dragging, which means that inertial frames are not stationary but they rotate with respect to an observer at infinity. Free falling particles with initial null angular momentum will begin to rotate along with the black hole, with angular speed increasing as the black hole approaches. It can be proved that, at the surface defined by

$$\left(r - \frac{GM}{c^2}\right)^2 = \frac{G^2M^2}{c^4} - a^2 \cos^2\theta \,, \tag{1.5}$$

photons emitted tangentially in the opposite direction of the black hole rotation will have null angular velocity, while inside this surface they will nonetheless move in the same direction of the black hole. This surface is called the *stationary limit surface*, while the region enclosed between it and the outer horizon is called the *ergosphere* (Figure 1.1). Particles inside the ergosphere can't help but be dragged along with the rotating spacetime, but don't necessary fall inside the event horizon, they are instead able to exit the ergosphere depending on their energy and direction. This property leads to a mechanism named *Penrose process*, which enables the extraction of rotational energy from a black hole. We consider the case of an object that enters the ergosphere with positive energy E and then separates in two parts, for example a particle that decays in two other particles. Because of the frame dragging, one of the two separated particles, depending on its angular momentum, may have negative energy $E_1 < 0$ and fall towards the black hole. Because of the conservations of energy the other particle must have an energy $E_2 = E - E_1 > E$, and depending on its angular momentum may exit the ergosphere with more energy than the initial particle. The net energy gain comes at the expense of the black hole's rotational energy, effectively slowing it down. The Penrose process is a powerful energy extraction mechanism that shows how rotating black holes may accelerate particles that travel in the proximity of its ergosphere.

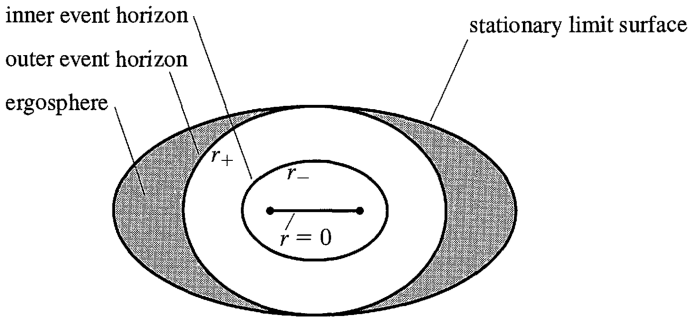

Figure 1.1: **Kerr black hole.** Vertical axial section of the Kerr black hole spacetime in Boyer-Lindquist coordinates. Figure from Carroll (2019).

1.1.2 Supermassive black hole observations

In the previous section we have treated black holes as mathematical metric solutions to the Einstein's field equation of General Relativity. Here we will discuss how black holes may arise in realistic astrophysical processes, what experimental proof we have of their existence and how we can observe them.

The major mechanism that can lead to the formation of a black hole is stellar collapse Figure 1.2. When a star exhausts the fuel for nuclear fusion, without the outward pressure originated by nuclear reactions, the star collapses under it own gravity forming first a neutron star. If the neutron-star mass is higher than the Oppenheimer-Volkoff limit, it will continue to collapse until all the mass is contained within the Schwarzshild radius, giving origin to a black hole. Black holes may also originate in the merger of two compact objects like for a pair of neutron stars. The black holes that originate from the collapse or merging of compact objects are called *stellar-size* black holes and they have masses up to 10^3 $M_\odot$. Another possible mechanism for black hole formation is the direct collapse of gas clouds in the early universe, without fragmentation into stars. The mechanism is believe to form *intermediate-size* black holes, with masses of $10^4 - 10^6$ $M_\odot$. Once a black hole is formed, it may gain mass by merging with other black holes, attract mass from a companion star or accrete nearby interstellar gas and stars. A combination of these mechanisms may lead to the formation of *supermassive* black holes (SMBH), which have masses of $10^6 - 10^9$ $M_\odot$. Regardless of the specific mechanism that leads to black hole formation, the initial material typically possesses nonzero angular momentum. Due to the conservation of angular momentum during collapse, this generally results in the formation of a Kerr black hole rather than a Schwarzschild one. Consequently, the Kerr metric is widely considered to provide an accurate description of astrophysical black holes. Additionally, as matter accretes into black holes, especially in the case of supermassive ones, it tends to settle in a rotating accretion disk, which gradually transfers additional angular momentum to the black hole.

Black holes may form by different mechanisms and from different initial components, but once the horizon is formed and the black hole reaches a stationary state, no information about the initial matter distribution or formation mechanism is accessible to an external observer. This is known as the *no-hair theorem* and it states that station-

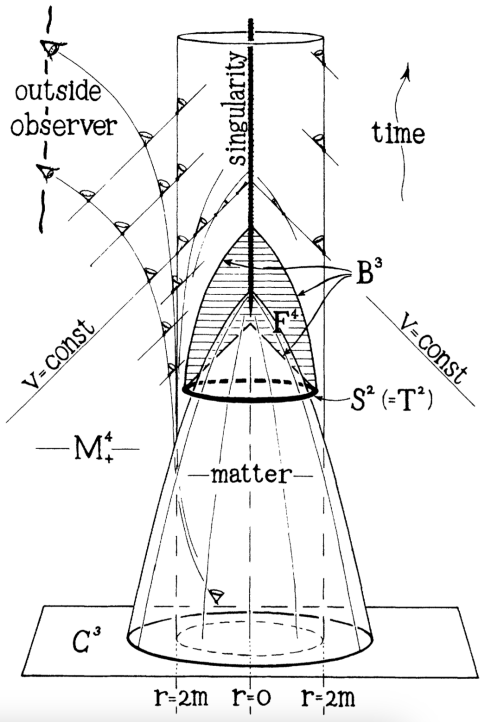

Figure 1.2: **Stellar collapse into a black hole.** The diagram from Penrose (1965) shows how a Schwarzschild event horizon and singularity arise as a consequence of the collapse of a spherical object. The time coordinate in on the vertical axis, while 2D spatial coordinates are on the horizontal plane. In 2+1 dimensions the event horizon is an infinite cylinder, while the singularity is a timelike infinite straight line.

ary, asymptotically flat black hole solutions to general relativity are fully characterized by the parameters of mass and angular momentum. Therefore, a black hole spacetime can be described only by the Schwarzschild or Kerr metric[1]. All other details—such as the composition, shape, or multipole moments of the collapsing matter—are radiated away during the formation process via gravitational and electromagnetic radiation. This means that the only black hole parameters that can be measured are its mass and spin.

Today there are various experimental evidences of the existence of black holes. The first black holes were detected by observing a star orbiting around an invisible binary companion or multiple stars orbiting around an invisible central object.

In 1974, Sgr A*, a highly compact, very bright radio object, was discovered within Sagittarius A (Sgr A), at the very center of the Galactic Center region (Balick and Brown 1974). In the 1990s and early 2000s, Ghez et al. (2008) and Gillessen et al. 2009, tracked stellar orbits around Sgr A* through near-infrared observations. They

[1]Here we neglect the Reissner-Nordström and Kerr–Newman metrics, which arise from coupling GR with electromagnetism. In fact, it is physically unrealistic that an astrophysical black hole may maintain a stable electric or magnetic charge, so these cases are not relevant to the scope of the thesis.

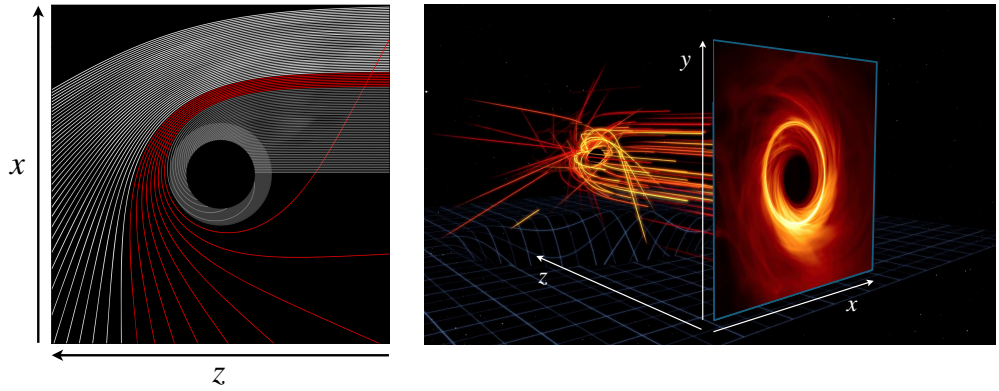

Figure 1.3: **Black hole shadow.** (Left) Null geodesics in the equatorial plane of a Schwarzschild black hole, that reach an observer located to the right, at infinity. Gray geodesics originate inside the photon sphere and have shorter light paths, resulting in a dimmer observed emission compared to the red geodesics, whose light paths are increased by the strong lensing. White geodesics undergo weaker lensing than the red ones, so their brightness is not increased as much as the red ones. Grey path correspond to the black hole shadow, red ones to the photon ring and white ones to the extended emission. Figure adapted from Bronzwaer and Falcke (2021). (Right) Tridimensional representation of the lensed emission around a black holes that results in the black hole shadow and the bright photon ring. Figure adapted from a video produced by the Center for Astrophysics, Harvard & Smithsonian.

observed that the stars followed Keplerian orbits around an invisible central object, whose mass was estimated to be 4 million solar masses. The most likely interpretation of such a compact object, was for it to be a black hole. Another experimental evidence of black holes is the detection of gravitational waves created during the merging of compact objects (LIGO and Collaborations 2016). The existence of black holes is also supported by observations of Active Galactic Nuclei (AGN). AGN are compact, highly luminous galaxy centers, emitting across the whole electromagnetic spectrum, with luminosity so high that it can outshine the whole host galaxy. Such a powerful source of energy can only be provided by supermassive black holes and it is in fact believed that supermassive black holes are present at the center of all galaxies, with some of them producing AGN emission. A more in depth introduction to AGN is presented in section 1.2.1.

However, all of these measurements are indirect observations of black holes. The first direct observation was achieved by the Event Horizon Telescope Collaboration (EHTC), which obtained radio images of the supermassive black hole shadows of M 87*, in the homonymous giant elliptical galaxy, and Sgr A*, at the center of the Milky Way (EHTC 2019a; EHTC 2022a). The black hole shadow is a distinctive observational feature that emerges when emitting material is present near the event horizon, such as in the case of an accretion disk surrounding a supermassive black hole. For the shadow to be observable, the emitting region must also be optically thin, which is the case for both Sgr A* and M 87. More specifically, as depicted in Figure 1.3, the BH shadow is the lensed projection of the black hole's photon sphere on the observer's image plane (Falcke et al. 2000; Bronzwaer and Falcke 2021). This

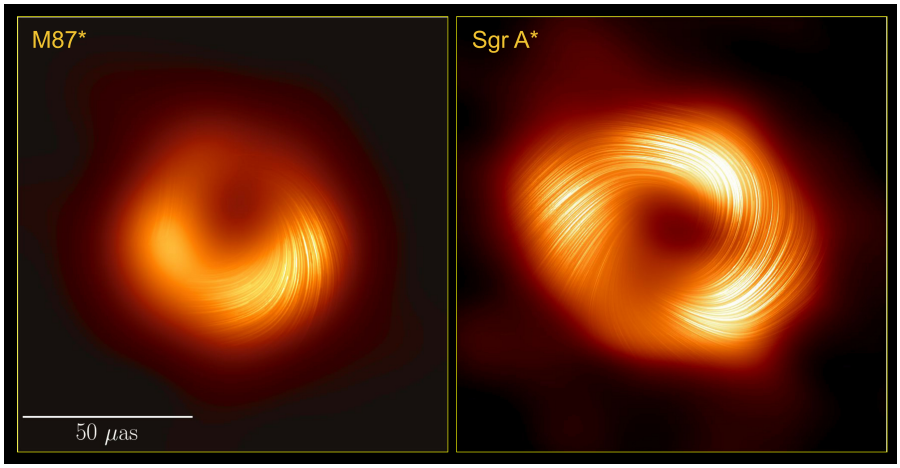

Figure 1.4: **M 87* and Sgr A* EHT images.** Horizon scale polarized images of the black hole shadows in M 87* and Sgr A* from 2017 EHT observations. The images display the total intensity emission with the linear polarimetric emission overplotted as field lines. Figure from the EHT Collaboration outreach material.

feature appears because most of the emission that originates inside the photon sphere is lensed into the black hole horizon, while most of the light emitted just outside of it is strongly lensed around the black hole, resulting in a ring-like shape in the image plane. Light geodesics that are warped around the black hole for more than half of a full rotation converge into a thin bright ring in the image plane, which is called the *photon ring*. The shape of the shadow and the surrounding emission depends on the spacetime metric and black hole orientation and on the characteristics of the accreting flow, meaning that observations of BH shadows can set constrains on both astrophysical processes and General Relativity in the strong gravity regime.

In 2017, the EHT observed the strong radio sources in M 87 and Sgr A with Very Long Baseline Interferometry (VLBI) observations (see section 2.1) and obtained horizon scale images of two supermassive black holes in both total intensity and linear polarization (Figure 1.4). In addition to being the first direct evidence of supermassive black holes, by demonstrating the existence of an event horizon and testing GR predictions in the strong gravity regime, these images led to numerous other scientific results. From the size of the shadows it was possible to measure the black hole masses, and for M 87*, the asymmetry of the ring also allowed to determine the spin direction. Both images were compared to extensive suites of General Relativistic Magneto-Hydro-Dynamic (GRMHD) simulations, leading to constrains on the simulation parameters compatible with the images. In both cases, observations strongly favor models with prograde spin and with a magnetically arrested accretion disk, i.e. a strong ordered magnetic field connecting the disk to the jet. The images also allowed to exclude many alternative scenarios explaining the emission in these radio source and reject some alternative models to Kerr black holes, such as black holes predicted by GR coupled with additional fields and black hole mimickers, which are horizon-less compact objects.

Black holes, however, display dynamic behaviors that are difficult to observe and measure from individual "snapshot" images. A time-resolved video of a black hole

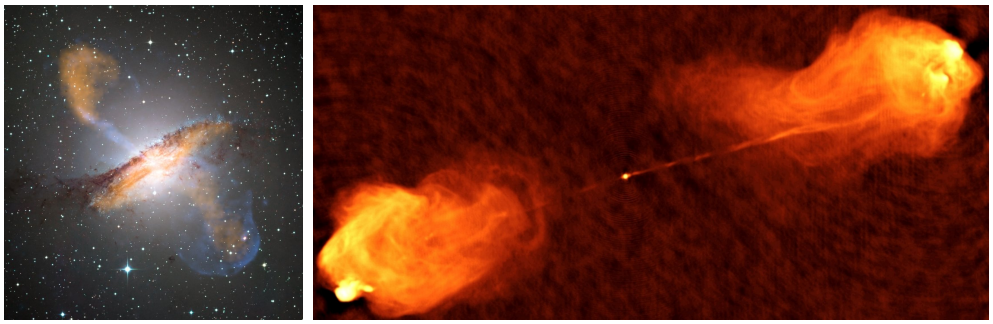

Figure 1.5: **Example of AGN jets.** (Left) Multifrequancy image of the Centaurus A galaxy with its relativistic jets. Credit: ESO/WFI (optical, in gray); MPIfR/ESO/APEX/Weiß et al. 2008 (submillimetre, in blue); NASA/CXC/CfA/Kraft et al. 2002 (X-ray, in blue). (Right) Jets of the Cygnus A radio galaxy at 4.9 GHz (Perley et al. 1984).

shadow, by contrast, would offer a powerful tool to probe the dynamics of the accretion flow. It would enable the measurement of the direction and speed of the plasma in the accretion disk (Conroy et al. 2023), as well as the orientation and evolution of the magnetic fields threading the black hole. These measurements could lead to tighter constraints on parameters such as the spin and orientation of the black hole, a better understanding of the jet launching mechanism, and to the identification of the accretion model and the role of magnetic fields. Additionally, the shape and size of the photon ring is completely determined by the spacetime metric and remains constant in time, while the projection of the lensed extended emission depends on the turbulent, variable plasma in the accretion flow. A video of the BH evolution can help disentangle the effects of the spacetime metric from those of the accretion process, by discriminating the persistent static emission from the variable one. Furthermore, a video of the black hole shadow would also allow to better understand variability in the emission originating around supermassive black holes, which results in flares observed at multiple wavelength from radio to X-rays (e.g. Wielgus et al. 2022).

The timescale over which the emission around a supermassive black hole shows variability depends on the black hole mass, which sets the period of the matter orbiting in the proximity of the horizon. The measured mass of M 87* is $\sim 6{\times}10^9$ $M_\odot$ (EHTC 2019f), resulting in a variability timescale from a few days to a week, while for Sgr A*, whose mass is $\sim 4{\times}10^6$ $M_\odot$ (Gravity Collaboration 2023), the variability timescale ranges from a few minutes to tens of minutes. The strategy adopted to observe a video of the black hole may change significantly depending on the variability timescale. In the case of M 87*, the EHT Collaboration is planning a multi-day observing campaign that will be conducted in 2026. Obtaining a video of Sgr A* is more complicated and requires the development of "dynamic" imaging algorithms, capable of recovering intra-day variability from instantaneous observations, rather than the full day observations (see section 2.1.5). One of the objectives of this thesis was to develop an effective dynamic imaging method for observations of the Sgr A* black hole. The method is presented, validated, and applied in Chapter 4 and Chapter 5 and it is currently being applied to image the first video of a supermassive black hole.

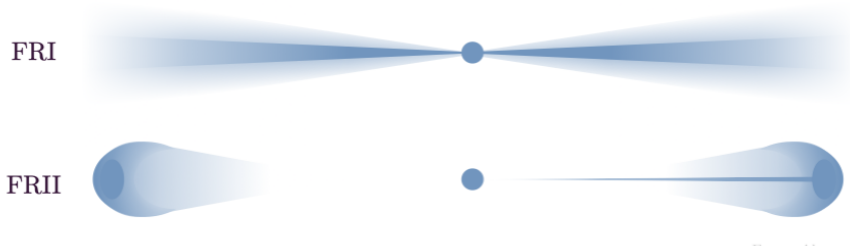

Figure 1.6: **Fanaroff–Riley class I and II.** Morphologies of the two FR classes of relativistic jet. Credit: E. Alexander.

1.2 Active galactic nuclei and relativistic jets

In this section, we describe the early observations and the proposed model of active galactic nuclei (section 1.2.1). We also present blazar jet observations, and describe accretion, jet launching mechanisms, and jet propagation (section 1.2.2).

1.2.1 Active galactic nuclei: observations and unified model

Active galactic nuclei (AGN) are bright and energetic regions located at the center of some massive galaxies. They emit across the whole electromagnetic spectrum with extremely high luminosities (10^{40} to 10^{48} erg/s), often outshining the host galaxy. AGN emission cannot be attributed to stellar activity but is rather the consequence of the accretion of gas and dust into a supermassive black hole. The first signatures of AGN were detected in 1943 by Seyfert, who observed a class of galaxies (now known as Seyfert galaxies) with luminous nuclei and strong broad emission lines (Seyfert 1943). The following decades saw the rise of the recently born field of radio astronomy (Jansky 1933; Reber 1940) with the development of the aperture synthesis technique and the first interferometers (McCready et al. 1947; Ryle 1962). This led to the discovery of the radio galaxies Cygnus A (Jennison and Das Gupta 1953; Baade and Minkowski 1954; Ryle et al. 1965) and Centaurus A (Bolton et al. 1949), with distinctive large-scale radio lobes extending far outside their optical counterparts (Figure 1.5). This indicated the presence of powerful non-thermal processes capable of emitting collimated jets able to reach extragalactic scales.

The first quasar, 3C 273, was discovered in 1963 (Hazard et al. 1963; Schmidt 1963), soon followed by the identification of many others (e.g. Matthews and Sandage 1963). In early observations, quasars appeared as star-like in optical wavelengths but also exhibited strong radio emission. Spectral redshifts indicated that they were located at intergalactic distances, leading to the discovery of compact objects capable of outshining their host galaxies. In the following years, a wide range of AGN were observed and categorized accordingly to their luminosity and spectral properties. In addition to quasars, this included Seyfert I and Seyfert II galaxies, which are nearby radio-quiet AGN with strong spectral lines, blazars, that are AGN whose jets are pointing almost along the line of sight, and radio galaxies, which emit strong radio emission from the pair of large jets and lobes. Radio galaxies are also classified into Fanaroff-Riley I and Fanaroff-Riley II, based on the morphology and brightness distribution of the jets(Fanaroff and Riley 1974). FR I galaxies are brightest near the core, with jets that

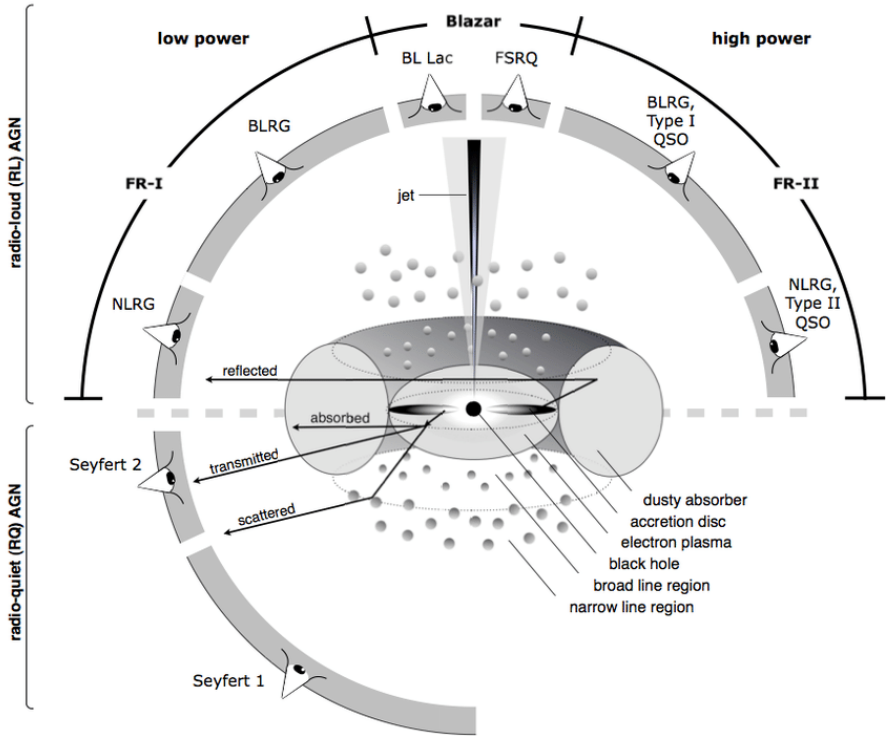

Figure 1.7: **The unified model of AGN.** Different classes of AGN are interpreted as referring to the same kind of object observed under different viewing angles. In addition to the viewing angle AGN are classified according to the emitted power (low or high) and the amount of radio emission (radio-loud or radio quiet). Image from Beckmann and Shrader (2012).

progressively fade and become diffuse as they move outwards, while FR 2 galaxies present powerful, collimated jets that end in bright radio lobes (Figure 1.6).

In the 1990s, unified models were proposed to reconcile different AGN classes into the same astrophysical object (Antonucci 1993; Urry and Padovani 1995). According to the models, the differences in the properties of AGN classes arise only from the power of the nucleus and from orientation and obscuration effects. As shown in Figure 1.7, an AGN is composed by a central supermassive black hole surrounded by a disk of infalling gas and dust. As the material orbits in the accretion disk, it heats up and emits radiation at high frequencies (UV and X-rays). Irregular clouds of gas are located above and below the disk. The inner clouds move at high speeds and emit broad spectral lines because of Doppler effects, while the outer clouds moving a lower speeds emit narrow spectral lines. Surrounding this system, in axis with the disk, is a thick torus of dust which, depending on the observing angle, may obscure the inner regions of the AGN. Some AGN also emit collimated, highly relativistic jets of plasma accelerated by strong magnetic fields in the vicinity of the black hole. If the orientation of the jets is close to the line of sight, like in the case of blazars, Doppler boosting may cause the receding jet to appear significantly dimmer, so that only one

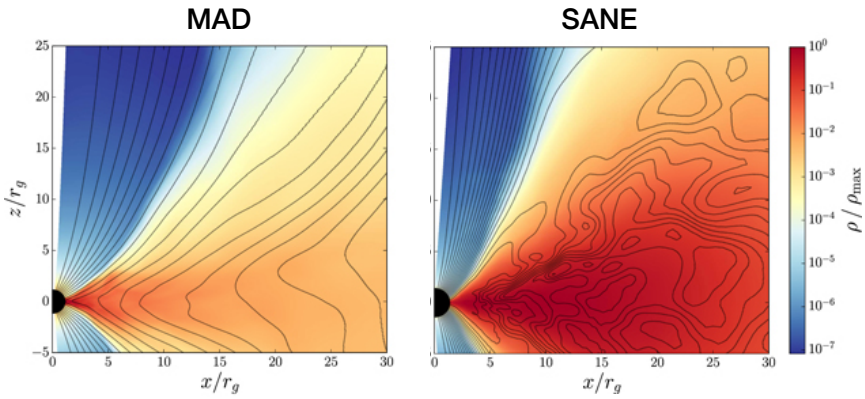

Figure 1.8: **MAD and SANE accretion models.** Simulations of different black hole accretion models. In MAD models the disk magnetic field is strong and coherent while in SANE models it is weak and turbulent. Figure adapted from O' Riordan et al. (2018).

jet is observed.

1.2.2 Accretion models and relativistic jets

Current theoretical models of black hole accretion propose two distinct scenarios, differing in the dynamic role of magnetic fields (Figure 1.8). In the magnetically arrested disk (MAD) model, the accreting gas drags in a strong magnetic field toward the black hole. As the field accumulates, it eventually becomes strong enough to arrest the inflow of matter (Narayan et al. 2003). MAD systems are highly efficient at launching powerful relativistic jets. In contrast, the standard and normal evolution (SANE) model involves accretion with relatively weak magnetic fields, insufficient to disrupt the inflow. This results in a more disordered, chaotic magnetic field within the disk. While relativistic jets can still form in SANE systems, they tend to be less powerful and less efficiently produced than in MAD scenarios.

The models of the accretion processes around supermassive black holes can be explored through GRMHD simulations (e.g. HARM, Gammie et al. 2012, KORAL, Sadowski et al. 2014, BHAC, Porth et al. 2017). A simulation begins by setting initial conditions for the spacetime metric (usually Kerr), the plasma distribution (typically a thick torus) and the magnetic field structure. Then the configuration is evolved according to the GR and magneto-hydrodynamic equations of motion.

The initial conditions of GRMHD simulations do not include jets, which instead arise naturally as the simulations progress. Two models describe the possible mechanisms responsible for the launching and collimation of the jets. In the Blandford-Znajek model (BZ, Blandford and Znajek 1977), rotational energy is transferred directly from the black hole spin to the jet by large scale magnetic fields threaded through the horizon. Charged particles are accelerated from the polar regions of the black hole magnetosphere into a pair of highly-relativistic, collimated jets. A second mechanism is the Blandford-Payne model (BP, Blandford and Payne 1982), in which the energy that powers the jet is extracted from the rotational kinetic energy of the accretion disk via magneto-centrifugal forces. In this model, charged particles are accelerated

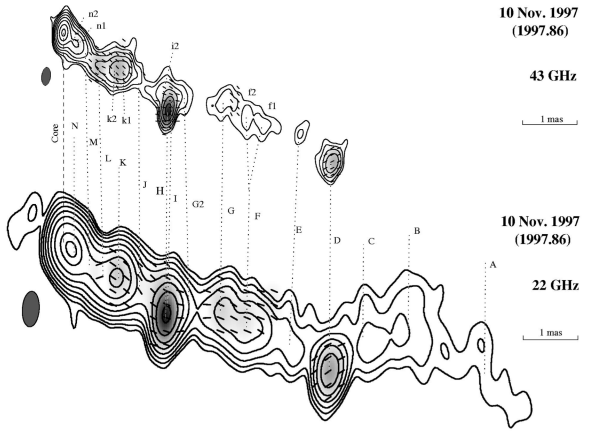

Figure 1.9: **Radio jet in 3C 120.** The image of 3C 120 observed by the VLBA at 43 and 22 GHz shows the typical knots in the jet, i.e. the bright, discrete "blobs", which are usually interpreted as shock perturbations in the jet plasma. Image from Gómez et al. (1999).

along large-scale magnetic field lines that are anchored in the rotating disk. The BP mechanism produces a collimated jet or wind, that is less relativistic than BZ jets. The BZ mechanism is generally associated with MAD accretion flows and is favored by observations (e.g. EHTC 2019e), while the BP mechanism is more predominant in SANE models. The two mechanisms can also coexist in the same system, leading to a jet with higher velocities in the spine and lower ones in the external sheath. As the jet progresses away from the black hole, the magnetic fields responsible for the collimation, remain wrapped around it in a helix structure. The charged particles in the jet plasma are accelerated along the threaded magnetic field, resulting in the characteristic emission of synchrotron radiation (see section 1.3.2).

The flow of plasma along the jet is often turbulent and subject to instabilities. Typical jets images from VLBI observations show the presence of multiple *knots*, i.e. bright, discrete, "blob" features (e.g. Figure 1.9), along the jet. In the last decades, extensive observations and studies have been conducted to track the motion of knots in multiple jet sources. These features are usually interpreted as shock waves in the plasma, following the first shock models by Blandford and Königl (1979) and Marscher and Gear (1985), and are characterized by a sharp increase in density and a discontinuity in velocity with respect to the surrounding, unshocked regions. The pattern speed of the shock perturbations may be higher or lower than the rest of the plasma. Traditionally, the motion of knots has been measured by modeling the jet with simple Gaussian components fitted to the observations, a technique referred to as *model fitting*.

In this thesis, we focus on the parsec-scale relativistic jets observed in blazars, which are primarily observed though VLBI observations. Initial studies of these sources began in the 1990s (Wilkinson 1995; Zensus 1997) and were significantly advanced by dedicated monitoring programs such as MOJAVE (Monitoring of Jets in Active Galactic Nuclei with VLBA Experiments, Lister et al. 2016) and BEAM-ME (Blazars Entering the Astrophysical Multi-Messenger Era, formerly VLBA-BU-BLAZAR, Jorstad

et al. 2017). Monitoring programs have been observing dozen of blazars at multiple frequencies, for decades up to today, using the Very Long Baseline Array (VLBA) interferometer (Napier et al. 1994). We image two of these sources, 3C 84 and 3C 345, in Chapters 3 and 4.

1.3 Radiative processes in relativistic jets

1.3.1 Relativistic effects

In the absence of gravity, General Relativity reduces to Special Relativity (SR), which is the theory that describes the geometry of space and time for observers in inertial frames. Special Relativity must be invoked to describe objects that move at relativistic speeds, while for non-relativistic velocities it reduces to Newtonian physics. The most important effects of special relativity are *time dilation* and *length contraction*. Considering two inertial reference frames moving with relative velocity v, the *proper time* t and the *proper length* s measured at rest in one frame are measured in the comoving frame of a moving observer as:

$$\begin{aligned} t' &= \gamma\, t\ , \\ s' &= \gamma^{-1} s\ , \end{aligned} \tag{1.6}$$

where the *Lorentz factor* γ is:

$$\gamma = \frac{1}{\sqrt{1-\beta}} \quad , \quad \text{with} \quad \beta = \frac{v}{c}\ . \tag{1.7}$$

The particles accelerated in AGN jets are highly relativistic and therefore exhibit relativistic phenomena. Following Ghisellini (2013), we introduce here two of them, superluminal motion and Doppler boosting, whose effects are taken into considerations in the discussions in Chapters 3 and 4.

Superluminal motion Because of the finite velocity of light, objects that are traveling at relativistic velocities close to the line of sight may appear to have velocities higher than the speed of lights. This effect is called *superluminal motion* and is caused by the different light paths covered by the light emitted by a moving object. Let's consider the situation represented in Figure 1.10, where an object is moving from point P_1 to point P_2, covering a distance s in a time $\Delta t = t_2 - t_1$, with velocity $v = s/\Delta t$. The object is emitting light at an angle θ towards an observer at a distance d. The light emitted by the object in P_2 is received by the observer at a time $t_{\text{app, 2}} = t_2 + d/c$, while the light emitted in P_1 is received at a time $t_{\text{app, 1}} = t_1 + d/c + s_\parallel/c$. Therefore, the observer sees the object cover the distance $s_\perp$ in a time interval:

$$\Delta t_{\text{app}} = t_{\text{app, 2}} - t_{\text{app, 1}} = t_1 - t_2 - \frac{s_\parallel}{c} = \Delta t - \frac{s}{c}\cos\theta\ , \tag{1.8}$$

resulting in an apparent velocity:

$$v_{\text{app}} = \frac{s_\perp}{\Delta t_{\text{app}}} = \frac{v\Delta t \sin\theta}{\Delta t - \frac{\Delta s}{c}\cos\theta} = \frac{v\sin\theta}{1 - \frac{v}{c}\cos\theta}\ , \tag{1.9}$$

or equivalently

$$\beta_{\text{app}} = \frac{\beta\sin\theta}{1-\beta\cos\theta}\ , \tag{1.10}$$

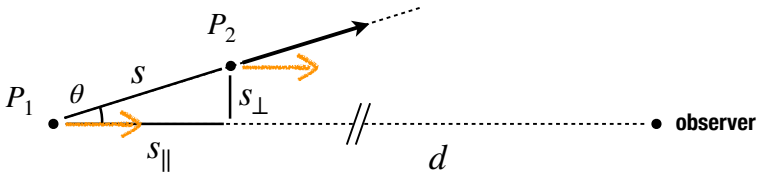

Figure 1.10: **Superluminal motion.**

where $\beta = v/c$ is the velocity expressed in light speed units. As velocities approach the speed of light ($v \to c$), the apparent speed $\beta_{\rm app}$ can exceed the speed of light, increasing for smaller angles with respect to the line of sight. Superluminal motion is a kinematic effect, not a special relativity effect, but it only becomes relevant for objects moving at relativistic speeds. This happens for traveling bright components in blazar, which can reach apparent speeds of a few tens of the speed of light (e.g. Lister et al. 2016).

Doppler boosting Considering again the situation in Figure 1.10, if the object is moving at relativistic speed, the frequency and intensity of the emission measured in the observed frame are different from the ones in the comoving frame. In classical mechanics, a wave emitted by a source moving at velocity v along a trajectory forming an angle θ with the line of sight is observed at a frequency ν', different from the emission frequency ν. This is a kinematic effect, known as Doppler boosting and the two frequencies are related through the equation:

$$\nu' = \frac{c}{c - v}\nu = \frac{1}{1 - \beta}\,\nu\ , \tag{1.11}$$

where c is the wave velocity. In the relativistic case, because of time dilation, the wave period T_o measured in the observer's frame is related to the wave period T_c observed in a frame comoving with the wave source by

$$T_o = \gamma\, T_c\ , \tag{1.12}$$

which, in terms of frequencies, is equivalent to

$$\nu_o = \gamma^{-1}\, \nu_c\ . \tag{1.13}$$

Combining the two effects, we obtain that the received frequency, measured in the observer's frame, is equal to:

$$\nu'_o = \gamma^{-1}\, \nu'_c = \frac{1}{\gamma\,(1 - \beta)}\,\nu_c = \delta\,\nu_c\ , \tag{1.14}$$

where

$$\delta = \frac{1}{\gamma\,(1 - \beta)} \tag{1.15}$$

is the relativistic Doppler factor.

Doppler boosting does not only affect the frequency of the emission but also the intensity. The *monochromatic intensity* I_ν is defined as the energy $dE = h\nu dN$ passing through a unit surface dA perpendicular to the photons arrival direction, per unit time, per unit solid angle, per unit frequency:

$$I_\nu = h\nu \frac{dN}{dt d\nu dA d\Omega} \, . \tag{1.16}$$

Applying the transformations $dt' = \delta^{-1} dt$, $d\Omega' = \delta^{-2} d\Omega$, which can be derived similarly to the frequency transformation in equation (1.14), the observed intensity is

$$I'_\nu = h\nu' \frac{dN'}{dt' d\nu' dA' d\Omega'} = h\delta\nu \frac{dN}{\delta^{-1} dt \, \delta d\nu \, \delta^{-2} d\Omega \, dA} = \delta^3 I_\nu \, , \tag{1.17}$$

where the number of particles dN and the area dA perpendicular to the motion are Lorentz invariants. For particles moving close to the line of sight ($\theta \to 0$) at relativistic speed ($\beta \to 1$), the Doppler effect boosts the intensity by a large factor. This makes the relativistic jets in blazar sources appear significantly brighter than they are. At the same time, the counterjets that are pointing away from the observer are affected by the Doppler effect in the opposite way and appear dimmer, to the point that in many cases they are undetectable and only one of the two jets is observed.

1.3.2 Synchrotron emission and polarization

The primary mechanism that generates the radio emission that we observe in blazars is synchrotron radiation, produced by the acceleration of relativistic electrons along the magnetic fields threading the jet. Here we present the main characteristics of the emission, following Rybicki and Lightman (1979) and Ghisellini (2013).

A charged electron moving with velocity $\vec{v}$ in a magnetic field $\vec{B}$ is subject to the Lorentz force:

$$\vec{F} = e \, (\vec{v} \times \vec{B}) \, , \tag{1.18}$$

where e is the electron charge. For relativistic speeds, Newton's second law takes the form:

$$\frac{d}{dt}(\gamma m \vec{v}) = e(\vec{v} \times \vec{B}) \, , \tag{1.19}$$

which can be decomposed into components parallel and perpendicular to the magnetic field:

$$\gamma m \vec{a}_\parallel = e \, (\vec{v}_\parallel \times \vec{B}) = 0 \, , \tag{1.20}$$

$$\gamma m \vec{a}_\perp = e \, (\vec{v}_\perp \times \vec{B}) \, , \tag{1.21}$$

where $\vec{a}_{\parallel,\perp} = d\vec{v}_{\parallel,\perp}/dt$ and we used $\gamma = \text{const}$ because the Lorenz force is conservative. These equations describe a constant motion in the direction parallel to the magnetic field, composed with a circular motion in the plane perpendicular to it, resulting in a helical motion along the the field lines (a diagram representing the situation is shown in Figure 1.11). The centripetal acceleration is $a_\perp = v_\perp^2 / r_L$, where r_L is the Larmor radius:

$$r_L = \frac{\gamma m v_\perp}{eB} \, . \tag{1.22}$$

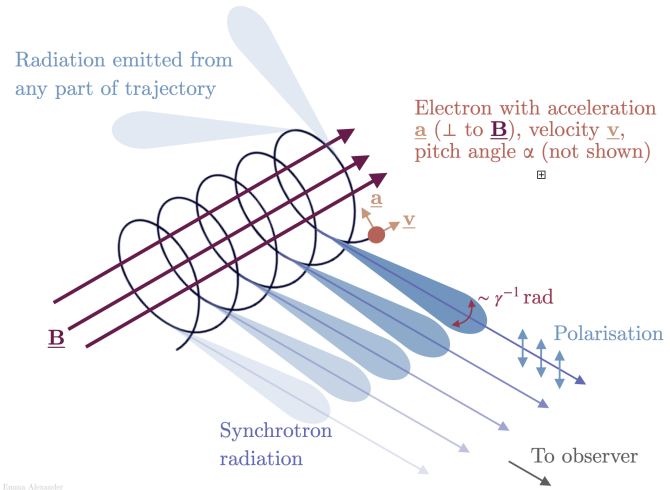

Figure 1.11: **Synchrotron emission.** Credit: E. Alexander.

The angular frequency of the circular motion, known as the *cyclotron frequency* ω_c, is:

$$\omega_c = \frac{v_\perp}{r_L} = \frac{eB}{\gamma m} \, . \tag{1.23}$$

The power emitted by an accelerated particle is Lorentz invariant and is given by the Larmor formula

$$P = \frac{2e^2}{3c^3}\left(a_\parallel'^2 + a_\perp'^2\right) = \frac{2e^2\gamma^4}{3c^3}\left(a_\parallel^2\gamma^2 + a_\perp^2\right) \, , \tag{1.24}$$

where the primed quantities refer to the particle's reference frame, and the unprimed ones to the reference frame of a static observer.

In our case $a_\parallel = 0$ and, considering a pitch angle α between the magnetic field and the electron, we can insert the expressions for the centripetal acceleration and the Larmor radius in the above equation to obtain:

$$P = \frac{2e^4}{3cm^2}\gamma^2\beta^2B^2\sin^2\alpha \, . \tag{1.25}$$

This is the total power radiated by a single relativistic electron moving in the magnetic field.

For a non-relativistic electron in circular motion, the emitted radiation would be at the cyclotron frequency ω_c. However, because of beaming effects, at relativistic velocities the majority of the radiation is focused into a narrow cone of half-angle $\sim 1/\gamma$ along the direction of motion. An observer therefore sees only short pulses of radiation as the cone briefly points in its direction. These short pulses lead to a spectrum that extends to much higher frequencies than ω_c. The spectrum of synchrotron radiation from a single electron is broad and peaks around a critical frequency ν_s, which is given by:

$$\nu_c = \frac{3}{4\pi}\frac{e\gamma^2B\sin\alpha}{mc} \, . \tag{1.26}$$

The resulting spectral distribution is:

$$P(\nu) = \frac{\sqrt{3}\, e^3 B \sin\alpha}{mc^2} F\left(\frac{\nu}{\nu_c}\right) , \tag{1.27}$$

where $F(x)$ is a dimensionless function describing the shape of the synchrotron spectrum (for the derivation and the details of F, see e.g., Rybicki and Lightman 1979.)

To obtain the total power emitted by the plasma, one needs to integrate the power $P(\nu, E)$ emitted by a single electron of energy E, over the whole energy distribution. Typically, in the jet plasma, it is assumed that electrons follow a power law distribution

$$N(E)dE = N_0 E^{-p} dE , \tag{1.28}$$

where N(E) is the number density of electrons with energy in $[E, E + dE]$ and p is the particle population index. Assuming isotropic emission, the synchrotron *emissivity* (or *emission coefficient*) of the plasma is given by:

$$j_v = \frac{1}{4\pi} \int N(E) P(\nu, E) dE , \tag{1.29}$$

which will also result in a power law distribution over frequency. However, the emission that actually exits the source will depend on the amount of radiation re-absorbed by the source itself. This is described by the radiative transfer equation, which relate the monochromatic intensity I_v to the emission and absorption coefficients j_ν and κ_ν by:

$$dI_\nu = j_\nu \, ds - \kappa_\nu I_\nu \, ds \tag{1.30}$$

where s is the linear space coordinate. We define the *optical depth* τ_ν as the line integral of the absorption coefficient:

$$\tau_\nu = \int \kappa_\nu \, ds . \tag{1.31}$$

When $\tau_\nu < 1$ the source is *optically thin*, meaning that a part of the radiation at frequency ν escapes the source. On the contrary, when $\tau \geq 1$ the source is *optically thick* and all the radiation is absorbed. Integrating equation (1.30), under the assumption that j_ν and κ_ν are constant in space, we obtain:

$$I_\nu = \frac{j_\nu}{\kappa_\nu}(1 - e^\tau) \tag{1.32}$$

which is the monochromatic intensity emitted by the source. Then, the intensity received by the observer will be affected by the relativistic effects described in the previous section.

Polarization For a single electron, the emitted radiation is linearly polarized in the plane of acceleration. In the case of synchrotron radiation, the acceleration is perpendicular to the magnetic field, so the observed polarization will be perpendicular to the magnetic field lines projected in the sky plane. The polarization of electromagnetic radiation can be described by the Stokes parameters:

$$\begin{aligned} \mathcal{I} &= |E_x|^2 + |E_y|^2 , \\ \mathcal{Q} &= |E_x|^2 - |E_y|^2 , \\ \mathcal{U} &= 2\mathrm{Re}(E_x E_y^*) , \\ \mathcal{V} &= -2\mathrm{Im}(E_x E_y^*) , \end{aligned} \tag{1.33}$$

which are defined here with respect to the x and y components of the electric field. The linearly polarized intensity P, the fractional linear polarization m_ℓ, the electric vector position angle (EVPA) χ, and the fractional circular polarization m_c are related to the Stokes parameters by:

$$\begin{aligned} P &= \sqrt{\mathcal{Q}^2 + \mathcal{U}^2} \quad , \quad \chi = \frac{1}{2} \arctan\left(\frac{\mathcal{U}}{\mathcal{Q}}\right) \ , \\ m_\ell &= \frac{\sqrt{\mathcal{Q}^2 + \mathcal{U}^2}}{\mathcal{I}} \quad , \quad m_c = \frac{|\mathcal{V}|}{\mathcal{I}} \ . \end{aligned} \tag{1.34}$$

The amount of fractional polarization in a jet depends on the electron energy distribution. The maximum fractional polarization for an optically thin source is around 70% for a perfectly ordered magnetic field. If the magnetic field is disordered, linear polarizations of different orientations cancel out, resulting in typical polarization values of 10-20% or less.

2

Methodology

2.1 VLBI observations and imaging

Astrophysical objects are observed across the full electromagnetic spectrum, from γ-rays to radio frequencies, using different kinds of telescopes. The smallest angular scale that a telescope can resolve is proportional to the observing wavelength λ and inversely proportional to the diameter D of the telescope, through the relationship:

$$\theta \simeq \frac{1.22\,\lambda}{D}\,, \tag{2.1}$$

which sets the diffraction limit of the telescope. For radio frequencies, antennas can reach sizes of $\sim$100 m, while larger antennas are limited by engineering and maintenance considerations.

So, for example, for the Effelsberg 100-m Radio Telescope (Wielebinski et al. 2011) this means that at the observing wavelength of 3.5 mm the maximum achievable resolution is $\sim 10''$. However, numerous interesting radio sources, like AGN relativistic jets and supermassive black holes, have angular sizes that are $10^3 - 10^5$ orders of magnitude smaller than the arcsecond. The resolution limitation of single radio antennas can be overcome through radio interferometry.

An interferometric array consists of a set of different antennas that observe simultaneously the same target in the sky. According to the van Cittert-Zernike theorem (Cittert 1934; Zernike 1938), the time-correlated signal of each pair of antennas is equal to the spatial Fourier transform of the observed source, evaluated at a frequency proportional to the projection of the antennas distance (or *baseline*) in the image plane. Therefore, a set of accurately spaced antennas may sample numerous Fourier components of the observed target, from which it is possible to recover the sky image, with an angular resolution up to:

$$\theta \simeq \frac{\lambda}{B}\,, \tag{2.2}$$

where B is the maximum baseline distance. The field of view (FOV) is instead limited by the minimum baseline distance b and is affected by time and frequency averaging smearing effects as well as the a beam size of each antenna.

The following subsections will introduce the fundamental concepts of radio interferometry and image reconstructions, in the context of Very Long Baseline Interferometry, with a focus on observations of relativistic jets and supermassive black holes. An in-depth explanation of radio interferometry and imaging can be found in Thompson et al. (2017), which served as the main reference for the present section. section 2.1.1 presents the assumptions and proof of the van Cittert-Zernike theorem, section 2.1.3 discusses the sources of noise in interferometric observations and the general strategies to minimize their impact. section 2.1.2 introduces the concept of Earth rotation synthesis on which VLBI is based, while section 2.1.4 and section 2.1.4 discuss different imaging methods for static and variable sources.

2.1.1 Interferometer response

We consider the case of an extended incoherent source, located in a distant plane, whose emission is being measured by two antennas at points P_1 and P_2 in a plane parallel to the source (Figure 2.1). The source is assumed to be in the far field, implying that the line of sight direction from points 1 and 2 is the same, that the source can be approximated as two-dimensional, and that the angular extension of the source is small compared to its distance from the observation points.

A radio antenna records a signal proportional to the received electric field. For the emission of a point element in the source, which is identified with the direction vector $\vec{s}$, the two antennas receive the signals:

$$\begin{aligned} E_1(\vec{s},t) &= A(\vec{s},t)e^{-i2\pi\nu t} \\ E_2(\vec{s},t) &= A(\vec{s},t+\tau)e^{-i2\pi\nu(t+\tau)} \ , \end{aligned} \tag{2.3}$$

where A and ν are the amplitude and frequency of the electromagnetic wave and τ is the *geometric delay* between the arrival of the signal at the two antennas. For a baseline distance $\vec{b}$, the time delay is proportional to the scalar product between the baseline vector and the line of sight vector:

$$\tau = -\frac{\vec{b}\cdot\vec{s}}{c} \ , \tag{2.4}$$

where the minus sign introduced because of the direction of vector $\vec{b}$ chosen is Figure 2.1. For an extended emitting source, the measured electric field is given by integrating equation (2.3) over the solid angle S subtended by the source:

$$\begin{aligned} E_1(t) &= \iint_S d\vec{s}\, A(\vec{s},t)e^{-i2\pi\nu t} \\ E_2(t) &= \iint_S d\vec{s}\, A\left(\vec{s}, t+\frac{\vec{b}\cdot\vec{s}}{c}\right) e^{-i2\pi\nu\left(t-\frac{\vec{b}\cdot\vec{s}}{c}\right)} \ . \end{aligned} \tag{2.5}$$

Assuming that the bandwidth $\Delta\nu$ is small enough that $\frac{\vec{b}\cdot\vec{s}}{c} \ll \frac{1}{\Delta\nu}$, we can assume that the amplitude of the signal is constant over the time delay scale. During correlation,

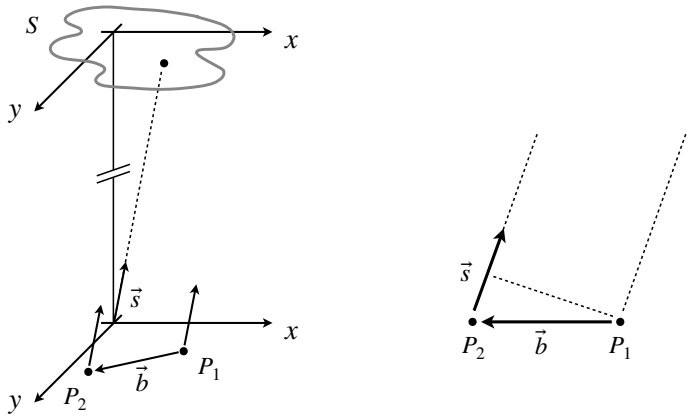

Figure 2.1: **Interferometric pair.** An incoherent extended source S is observed in the far field by two antennas at points P_1 and P_2. The antennas are separated by the baseline vector $\vec{b}$, while points in the source are identified with the direction vector $\vec{s}$. The difference in the light path from the source to P_1 and P_2 is given by $\vec{b} \cdot \vec{s}$.

the conjugate product of the signal recorded by each pair of antennas is averaged in time. The signal exiting the correlator is:

$$R(\vec{b}) := \langle E_1(t) E_2^*(t) \rangle_t = \iint_S d\vec{s_1} \iint_S d\vec{s_2} \, \langle A(\vec{s_1}, t) A^*(\vec{s_2}, t) \rangle_t \, e^{-i2\pi\nu \frac{\vec{b} \cdot \vec{s_2}}{c}} \, . \tag{2.6}$$

For an incoherent source, the radiation from any two points is statistically independent, so:

$$\langle A(\vec{s_1}, t) A^*(\vec{s_2}, t) \rangle_t = \delta(\vec{s_1}, \vec{s_2}) I(\vec{s_1}) \, , \tag{2.7}$$

where $I(\vec{s})$ is the time averaged intensity. In this case, equation (2.6) becomes:

$$R(\vec{b}) = \iint I(\vec{s}) e^{-i2\pi\nu \frac{\vec{b} \cdot \vec{s}}{c}} d\vec{s} \, . \tag{2.8}$$

The correlation product of two antennas is referred to as a *complex visibility* V. The coordinates of the baseline vector projected in the image plane are usually expressed in units of the observing wavelength $\vec{b} = (b_x, b_y) = (u\lambda, v\lambda)$, while the source extension is described by angular sky coordinates $\vec{s} = (x, y)$. With this notation, the van Cittert-Zernike theorem takes the form:

$$V(u, v) = \iint I(x, y) e^{-i2\pi(ux + vy)} dx \, dy \, , \tag{2.9}$$

which states that the interferometric correlation products are equal to the spatial Fourier transform of the sky brightness distribution. Imaging VLBI data consists in inverting the above equation to recover the source image $I(x, y)$.

In practice, the receivers of radio telescopes record the incoming electromagnetic wave under two separate polarization states, either with a linear feed, which is sensitive to the x and y components of the electric field (E^x, E^y), or a circular feed, which is

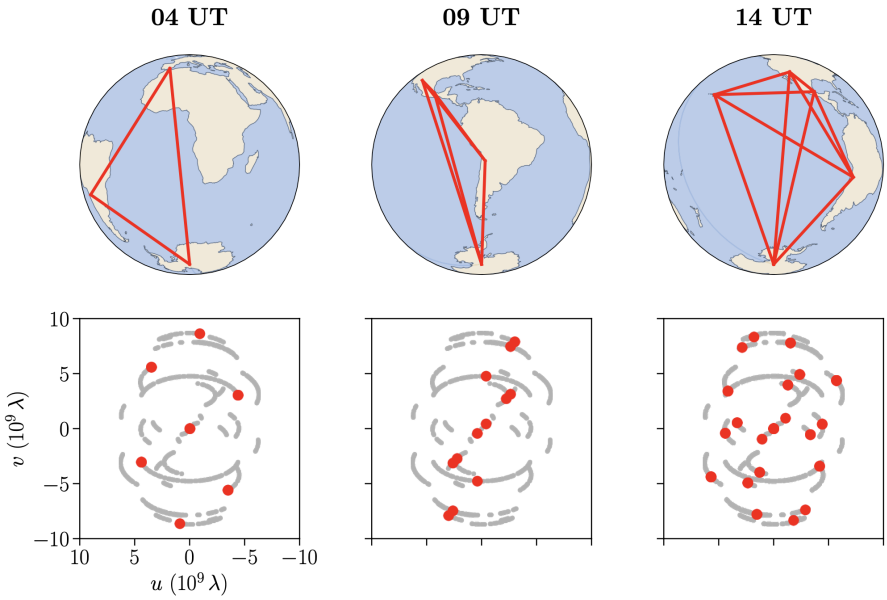

Figure 2.2: **Earth rotation synthesis.** Example from observations of Sgr A* with the EHT array in 2017 (EHTC 2022a). The first row shows the baselines between the observing antennas for three different times, while the second row shows the corresponding instantaneous coverage in red and the total one in gray. Figure from Fuentes (2022).

sensitive to the left and right components of circular polarization (E^r, E^ℓ). Therefore, for a pair of antennas (i, j), there are four possible correlation products, that are described by the *coherency matrix*:

$$\rho_{ij} := \begin{pmatrix} \langle E_i^r E_j^{r\,*} \rangle & \langle E_i^r E_j^{\ell\,*} \rangle \\ \langle E_i^\ell E_j^{r\,*} \rangle & \langle E_i^\ell E_j^{\ell\,*} \rangle \end{pmatrix} , \tag{2.10}$$

which is here written for the case of a circular feed. Assuming that the measurements are only affected by stochastic thermal noise, the coherency matrix is related to the polarimetric visibilities $(V^{\mathcal{I}}, V^{\mathcal{Q}}, V^{\mathcal{U}}, V^{\mathcal{V}})$ by the relationship:

$$\rho_{ij} = \begin{pmatrix} (V^{\mathcal{I}} + V^{\mathcal{V}})_{ij} & (V^{\mathcal{Q}} + iV^{\mathcal{U}})_{ij} \\ (V^{\mathcal{Q}} - iV^{\mathcal{U}})_{ij} & (V^{\mathcal{I}} - V^{\mathcal{V}})_{ij} \end{pmatrix} , \tag{2.11}$$

which can be inverted to obtain the visibilities associated with the Fourier transform of each Stokes parameter $(\mathcal{I}, \mathcal{Q}, \mathcal{U}, \mathcal{V})$. In section 2.1.3, we will discuss how the above equations are affected by realistic noise corruption.

2.1.2 Earth rotation synthesis and VLBI

In order to invert equation (2.9) and image the observed source, the Fourier transform must be sufficiently sampled in the (u, v) plane. This can be achieved with a high number of properly spaced antennas and/or by relying on the rotation of the Earth. Indeed, as the Earth rotates, the projected baseline vectors change length and orientation with respect to the line of sight, allowing to sample multiple spatial frequencies

with the same pair of antennas (Figure 2.2). Observations may be conducted for several hours to allow the maximum exploitation of the Earth rotation. This technique is known as *Earth rotation synthesis* and enables a much denser sampling of the Fourier space, and/or a reduction of the minimum number of antennas required.

To maximize an interferometer's resolving power, the antennas in the array can be placed as far apart from each other as possible. In fact, antennas do not need to be physically connected, but can be placed thousands of kilometers apart, with baselines distances up to the size of the Earth. The technique that combines signals from a network of widely separated antennas that operate independently is called *Very Long Baseline Interferometry*. Connected arrays, like the Atacama Large Millimeter Array (ALMA, Wootten and Thompson 2009), the Very Large Array (VLA, Thompson et al. 1980), or the future Square Kilometer Array (SKA, Dewdney et al. 2009), perform real-time correlation of the observed signals. On the contrary, in VLBI arrays, like the Very Long Baseline Array (VLBA, Napier et al. 1994, Figure 2.3) or the Event Horizon Telescope (EHT, EHTC 2019b, Figure 2.3), the antennas record the incoming radio signals with time stamps measurements by atomic clocks. These recordings are later physically shipped to the same location and correlated offline. VLBI arrays can reach resolutions down to 20 microarcseconds, which enables detailed observations of the inner sections of relativistic jets and horizon scale observations of supermassive black holes.

2.1.3 Measurement noise and error mitigation

Interferometric measurements are affected by various sources of noise. Measurement errors can be modeled into antenna-based errors (or *complex gains*) and baseline dependent errors (i.e. stochastic noise). Accordingly, for a pair of antennas (i, j), the measured visibilities V' are related to the true visibilities V by:

$$V'_{ij} = G_i \, G_j^* \, V_{ij} + \epsilon_{ij} \ , \tag{2.12}$$

where $G_{i,j} := g_{i,j} e^{\psi_{i/j}}$ are the complex antenna-based gains describing phase errors ψ and amplitude gains g associated with antennas i and j, while ϵ_{ij} is the stochastic thermal noise associated with the (i, j) baseline. For a given baseline (i, j), the measured amplitudes A'_{ij} and phases ϕ'_{ij} are related to the true amplitudes A_{ij} and phases ϕ_{ij} by:

$$A'_{ij} = g_i \, g_j \, A_{ij} \tag{2.13}$$

$$\phi'_{ij} = \phi_{ij} + \psi_i - \psi_j \tag{2.14}$$

Phase errors may be caused by an incorrect estimation of the time delay between two antennas. This may be due to errors in the measurement of the antennas' positions, clock errors or to local random fluctuations in the atmosphere's refractive index, which increase the light's path length. To mitigate phase errors, a calibration process called *fringe fitting* is performed to search for the phase correction that yield the highest correlation values. Amplitude gains are determined by the antenna's amplification of the received signal. A pair of amplitude gains is the multiplicative factor that converts correlation coefficients, which are given in units of noise power, to physical flux density

Figure 2.3: **VLBI arrays.** Locations of the antennas in the EHT array (top) and the VLBA array (bottom). For the EHT array, red points mark the antennas present in 2017, while green points denote antennas added between 2018 and 2025.

units (usually Jansky). They are given by the expression:

$$g_i = \frac{1}{\sqrt{\mathrm{SEFD}_i}} = \sqrt{\frac{\eta_{\mathrm{el},i} \cdot \mathrm{DPFU}}{T_{\mathrm{sys}_i}}}\,, \tag{2.15}$$

where the System Equivalent Flux Density (SEFD) is the total system noise represented in units of equivalent incident flux density, Tsys is the effective system noise temperature, the Degrees Per Flux Unit (DPFU) is the conversion factor from temperature units (Kelvin, K) to intensity units (Jansky, Jy) correcting for the antenna's aperture efficiency, and $\eta_{\mathrm{el},i}$ is the gain curve, which models the elevation dependence of the telescope's aperture efficiency. In principle amplitude gains can be computed from the antennas specifications (*a-priori calibration*) but in practice some of these quantities may be incorrectly estimated, resulting in residual amplitude gain errors.

To mitigate this possibility, the observation of the target source is alternated with the observation of a nearby source of known brightness, which is used as a calibrator for the amplitude gains. However, for millimeter VLBI it is very difficult to find suitable calibrator sources with well-known brightness and structure because, at these frequencies, most of them are internally resolved and time variable. In addition to errors in amplitude and phase, interferometric measurements are also affected by thermal noise due to the receiving instrumentation. This source of noise is stochastic and can be well characterized by a Gaussian distribution, whose width determines the uncertainty of the observed visibilities.

Even after careful fringe fitting and a-priori amplitude calibration, residual calibration errors may still be present in the data. One possible mitigation strategy is *self-calibration*, which consists in making an assumption about the true image, computing the visibilities associated with that image, equating the visibilities to the observed ones, and solving for the complex gains. Self-calibration can be performed on phase gains, amplitude gains or both, but is effective only if the assumed image is a good model of the true one, otherwise it carries the risk of biasing the data in favor of the model image.

Another way to address residual calibration uncertainty is to construct data products that are invariant with respect to phase or amplitude gains. Considering a triangle of antennas (i, j, k) and the corresponding baselines, we can define a *closure phase* Φ_{ijk} as the sum of the three visibilities phases associated with the triangle:

$$\Phi_{ijk} := \phi'_{ij} + \phi'_{jk} + \phi'_{ki} \ , \tag{2.16}$$

where primed quantities refer to measured phases, to distinguish them from the true phases. Using equation (2.12), the closure phase can be rewritten as:

$$\begin{aligned} \Phi_{ijk} &= (\phi_{ij} + \psi_i - \psi_j) + (\phi_{jk} + \psi_j - \psi_k) + (\phi_{ki} + \psi_k - \psi_i) \\ &= \phi_{ij} + \phi_{jk} + \phi_{ki} \ , \end{aligned} \tag{2.17}$$

which shows that phase errors are canceled in the triple phase sum. Similarly, considering a set of four antennas $ijkl$, we can define a *closure amplitude* A_{ijkl} as:

$$A_{ijkl} = \frac{|V'_{ij}||V'_{kl}|}{|V'_{ik}||V'_{jl}|} \ , \tag{2.18}$$

where again primed quantities indicate measured values and equation (2.12) allows to show that amplitude gain errors cancel in the double ratio:

$$A_{ijkl} = \frac{g_i g_j |V_{ij}| g_k g_l |V_{kl}|}{g_i g_k |V_{ik}| g_j g_l |V_{jl}|} = \frac{|V_{ij}||V_{kl}|}{|V_{ik}||V_{jl}|} \ . \tag{2.19}$$

For an array of N antennas the number of complex visibilities is $N(N-1)/2$, for a total of $N(N-1)$ degrees of freedom. Instead, the maximum number of independent closure phases and independent closure amplitudes that can be constructed is respectively $(N-1)(N-2)/2$ and $N(N-3)/2$, meaning that, while closure quantities are independent of gain errors, they carry less information compared to the full set of visibilities. In particular, closure phases lose the information about the absolute position of the source, while closure amplitudes lose information about the total flux

of the image. However, as the number of antennas grows, the information contained in closure quantities tends to that of the complex visibilities, reaching, for example, 80% of the total information for arrays of 11 antennas. As discussed in section 2.1.4, the recently developed forward imaging methods are able to make use of closure quantities in the image reconstruction process, thus avoiding mis-calibration biases.

Additional sources of error may affect polarimetric visibilities. Due to hardware imperfections, the receiver may not separate correctly the two orthogonal polarization states, resulting in *polarization leakage*, where part of the signal from one polarization state is mixed with the other. The effect of leakage is modeled and quantified by D-terms (D_R, D_L), which, for a circular feed, are defined by the equations:

$$\begin{aligned} E'_r &= E_r + D_R E_\ell \\ E'_\ell &= E_\ell + D_L E_r \ , \end{aligned} \tag{2.20}$$

in which (E'_r, E'_ℓ) are the measured polarization states, while (E_r, E_ℓ) are the true ones. Polarization leakage and complex gains can be elegantly incorporated in the coherence matrix formalism through the introduction of the *Jones matrix* associated to the antenna:

$$J = \begin{pmatrix} G & 0 \\ 0 & G \end{pmatrix} \begin{pmatrix} 1 & D_R \\ D_L & 1 \end{pmatrix} \ , \tag{2.21}$$

which links the measured coherency matrix ρ' to the true one:

$$\rho'_{i,j} = J_i \, \rho_{ij} \, J_j^\dagger \ . \tag{2.22}$$

2.1.4 Image reconstruction and imaging methods

In VLBI, even with Earth rotation synthesis, the spatial frequencies of the (u, v)-domain remain under-sampled, with the coverage often displaying anisotropy in the distribution of sampled frequencies and different sampling densities from region to region. If this was not the case, an image of the target source could be recovered simply with a discrete inverse Fourier transform of the visibilities. Instead, imaging VLBI observations constitutes an ill-posed inverse problem, meaning that there is an infinite set of possible images that would be compatible with the data. In mathematical terms, the true image source $I_{\text{true}}(x, y)$ is observed through an under-sampled Fourier transform $\mathcal{F}\left[I(x, y)_{\text{true}}\right]$ which yields a set of N visibilities $\boldsymbol{V} = \{V_n(u, v)\}_{n\in[1,N]}$. Imaging consists in reconstructing an image $I(x, y)$ that is consistent with the observed data and is as close as possible to the true underlying image. To retrieve the most likely image, among the infinite set of possible ones, prior assumptions or information about the source need to be incorporated into the imaging process.

A variety of methods have been developed, especially in recent years, to address the VLBI imaging problem. These approaches are generally classified into inverse imaging methods and forward imaging methods. Inverse imaging methods begin by applying a direct Fourier transform to the visibilities, producing the so called *dirty image*, which is the convolution of the real image with the interferometer's point source response function, known as *dirty beam*. The dirty image is then progressively refined to reconstruct the final image. In contrast, forward imaging methods start with a tentative model of the image, which is optimized until its Fourier transform

matches the observed visibilities. Inverse methods amount to variations of the same core algorithm, CLEAN, which was the first method to be developed (Högbom 1974) and remains the most widely used in the VLBI imaging community.

In the last decade, the development of forward imaging methods was driven by the necessity to image the extremely sparse datasets of horizon-scale observations of the M 87* and Sgr A* black holes by the EHT (EHTC 2019a; EHTC 2022a). To reach a resolution of ∼20 μas, the EHT is composed of ∼10 antennas scattered around the globe, which means that many stations cannot observe simultaneously or do so for a short amount of time, making the Fourier coverage extremely sparse. Furthermore, at the observing frequency of 230 GHz, phase gains from atmospheric fluctuations vary on a timescale of minutes, completely corrupting the visibility phases. In the case of Sgr A*, interstellar scattering from the Milky Way and the intrinsic variability of the source further complicate the imaging process. The effort to tackle the unique set of imaging challenging posed by EHT observations resulted in a variety of imaging algorithms, which have in common the ability to produce super-resolved images, improving over the nominal beam resolution limit of CLEAN. The majority of forward modeling methods belongs to the general categories of Regularized Maximum Likelihood (RML) methods and Bayesian methods. In the following paragraphs we provide a brief overview of CLEAN, RML, and Bayesian imaging methods, focusing on the core concepts of each kind of algorithm and highlighting its advantages and disadvantages.

CLEAN. The CLEAN imaging method is based on the assumption that the image can be described by a collection of unresolved point sources, making it especially suitable for compact sources, rather than extended ones. The first step of the algorithm consists in computing the inverse Fourier transform of the observed visibilities to obtain the dirty image $I_{\rm dirty}(x,y) = \mathcal{F}^{-1}[\boldsymbol{V}]$, which is equal to the convolution between the true sky brightness and the dirty beam:

$$I_{\rm dirty}(x,y) = I_{\rm true}(x,y) * B_{\rm dirty}(x,y)\ . \tag{2.23}$$

Then, a scaled version of the dirty beam is subtracted from the dirty image at the location of the maximum brightness point, while the location (x_i, y_i) and brightness I_i values are added to a list of delta function components $\{I_i\delta(x-x_i, y-y_i)\}$. This process is repeated iteratively until the image only consists of the residual brightness $R(x,y)$ that is lower than the noise level. The final set of components is convolved with the *clean beam* $B_{\rm clean}$, an elliptical Gaussian fitted to the dirty beam, which defines the nominal resolution of the array. The resulting image is then added to the residuals to obtain the final image:

$$I(x,y) = \sum_i I_i\delta(x-x_i, y-y_i) * B_{\rm clean} + R(x,y)\ . \tag{2.24}$$

The iterative cleaning process, usually requires the user to restrict the regions in which components can be searched (*cleaning windows*), and this introduces arbitrariness in the imaging procedure. Another limitation of CLEAN is that it relies on complex visibilities to compute the inverse Fourier transform, requiring a good a priori calibration of both amplitudes and phase gains. This is achieved by first self-calibrating the dataset to a simple model, usually a Gaussian, and then alternating CLEAN steps with

self-calibration to the provisional image. The CLEAN method has been extensively employed for decades and has proven effective in imaging a wide range of sources, establishing it as the standard benchmark against which more advanced methods are compared.

RML methods. Regularized Maximum Likelihood methods define the image as a square matrix of flux density pixels $\boldsymbol{I} = \{I(x_i, y_j)\}$ and optimize the pixel values to minimize the loss function:

$$J(\boldsymbol{I}) = \sum_D \alpha_D \chi^2_D(\boldsymbol{I}, \boldsymbol{V}) - \sum_R \beta_R S_R(\boldsymbol{I}) \, . \tag{2.25}$$

The first sum in the loss is the *data term*, which enforces compatibility between the image and the observations through a χ^2 computed for the chosen data products D. The second sum is the *regularizer term*, which enforces correlation among pixels, constraining the possible solutions to the ill-posed imaging problem. The choice of regularizers R may favor different properties of the image, such as sparsity, similarity with a prior image, continuity, or smoothness, or set physical constrains, like the total integrated flux (see EHTC 2019d, for a list of the most common regularizers). The coefficients α_D and β_R that weight the data terms and the regularizers are hyperparameters of the method.

The main advantage of RML methods over CLEAN is the ability to perform imaging using only closure quantities, without requiring complex visibilities, making the imaging process independent from calibration errors. This is especially important for millimeter wavelength observations, like in the case of the EHT array, where complex gains completely corrupt the visibility phases. Additionally, the images produced by RML methods do not require convolution with the nominal beam, since the smoothness in the image is already provided by the regularizers. This allows the algorithm to achieve an effective resolution higher than the nominal one, typically by a factor of 2 or 3. The downside of RML methods is that the choice of regularizers significantly affects the final image and presents some degree of arbitrariness. If the risk of bias from an incorrect regularizer choice is high, a hyperparameter survey can be performed, in order to select combinations of regularizer weights that perform well on simulated data. Parameter surveys are computationally expensive and necessitate the preparation of synthetic data. Even though most RML methods have been developed in the context of black hole observations, they often proved useful also to image less sparse observations of a variety of sources. For instance, the images of the relativistic jet in 3C 84, which are presented in Chapter 3, were imaged with the `eht-imaging` algorithm (Chael et al. 2018). Other methods that are based on RML approaches include `SMILI` (Akiyama et al. 2017), `DoG-HiT` (Müller and Lobanov 2023), and `ngMEM` (Mus and Martí-Vidal 2024). In the past few years, more than a hundred publications relied on RML methods for imaging interferometric observations, from multiple arrays including EHT, VLBA, VLA, and ALMA. This indicates the relevance and effectiveness of this class of algorithms as well as the demand of novel imaging algorithm in the VLBI imaging community.

Bayesian methods. VLBI imaging can also be approached from a statistical inference perspective. Bayesian methods aim to reconstruct not only the most likely image, given the observed data and prior assumptions, but also its posterior distribution, which describes the image's stochastic uncertainty. Most Bayesian methods

provide a simultaneous parametrized model $(I(x,y),\{g_i\})_{\boldsymbol{\theta}}$ of the image and the instrument response (e.g. antenna gains, D-terms, noise distribution), where the parameters $\boldsymbol{\theta} = \{\theta_1, \ldots, \theta_n\}$ are considered as stochastic variables. Prior information about the image is incorporated in the parameter's prior distributions $\mathcal{P}(\boldsymbol{\theta})$ and regularization is implemented through correlation among the distributions of different parameters. A likelihood function $\mathcal{P}(\boldsymbol{V}|\boldsymbol{\theta})$ is defined to quantify the probability distribution of the observed visibilities for given combinations of parameter values. The likelihood function implements the data constraint, which is usually based on a weighted squared difference between the data products resulting from a choice of image parameters and the observed ones. The imaging problem can then be stated as a inference problem, in which the goal is to recover the posterior distribution of the model parameters, from the parameters prior distribution and the likelihood distribution of the observed data. From Bayes theorem:

$$\mathcal{P}(\boldsymbol{\theta}|\boldsymbol{V}) = \frac{\mathcal{P}(\boldsymbol{V}|\boldsymbol{\theta})\mathcal{P}(\boldsymbol{\theta})}{\mathcal{P}(\boldsymbol{V})} \,, \tag{2.26}$$

where the evidence $\mathcal{P}(\boldsymbol{V})$ can be treated as a simple normalizing factor since it has no dependence on the parameters. The core of the imaging process consists in sampling and/or approximating the posterior distribution, by exploring the parameter space with Monte Carlo methods or other sampling techniques.

Similarly to RML methods, Bayesian algorithms are also able to achieve super-resolution and use closure data products. Additionally, they have the advantage of a modular framework, in which it is straightforward to switch between data formats (e.g. coherence matrix rather than Stokes parameters) and add further modeling of the instrument's response and noise corruption. Most importantly, Bayesian methods provide a quantification of the image uncertainty, which is of great importance in the case of extremely sparse datasets. This comes at the cost of computational time, since the sampling of the posterior requires significantly longer run times compared to CLEAN imaging or RML optimization, making Bayesian methods more suitable for small sparse datasets and/or outputs described by smaller sets of parameters. A more subtle limitation of Bayesian methods is that the recovered posterior distribution is limited to the configurations allowed by the parameter priors, which need to be carefully defined, in order to allow for sufficient expressivity of the model and at the same time provide sufficient constrains to the image. State of the art Bayesian methods include `THEMIS` (Broderick et al. 2020), `Comrade` (Tiede 2022), and `Resolve` (Junklewitz et al. 2016).

Machine learning methods. In the past few years, a fourth class of methods based on machine learning (ML) approaches began to emerge, with successful proof of concept demonstrations but still with limited applications to real science cases (e.g. `DPI`, Sun and Bouman 2021, `PRIMO`, Medeiros et al. 2023, Deep Generative Image Priors, Feng et al. 2024, `DIRECT`, Lai et al. 2025). `DPI`, is an unsupervised Bayesian method based on variational inference with a normalizing flow. The method is able to characterize the image uncertainty and does not rely on explicit morphological priors. However, due to the network's architecture, the number of tunable parameters grows linearly with the number of input pixels, severely limiting the scalability of the model. The other methods mentioned above are based on supervised machine learning models, which are trained on a dataset of prior images. They produce high quality results but

Method	Strong morphological priors	Uncertainty quantification	Run time	Scalability
RML	Yes	No	short*	easy
Bayesian	Generally No	Yes	long	difficult
Unsupervised ML	No	Yes	short	difficult
Supervised ML	Yes	No	short†	difficult
`kine`	**No**	**No**	**short**	**easy**

Table 2.1: **Forward imaging methods comparison.** Summary of the most relevant properties of the different classes of forward imaging methods for VLBI. Common to all forward methods is the ability to employ closure quantities and achieve super-resolution. *If a parameter survey is not required. †Run time time is short once the method is trained.

are strongly biased by the choice of the training set. Part of the work of this thesis consisted in developing a new imaging method based on neural radiance fields, which overcomes the present limitations of RML, Bayesian, and machine learning method. The proposed method represents a competitive and more promising alternative to RML algorithms, and we expect it to have a significant impact in VLBI imaging, broadening the application of forward imaging methods. The method is presented in Chapter 4 where it is also applied to the imaging of multi-epoch blazar observations and it is validated on realistic synthetic EHT observations in Chapter 5. In Table 2.1 we provide a summary of the advantages and disadvantages of the different forward imaging methods, in comparison with the new method presented in this thesis.

2.1.5 Dynamic imaging

Because of the sparsity of interferometric arrays, VLBI observations strongly rely on Earth rotation synthesis to sufficiently sample the Fourier space. However, this can only be applied under the assumption that the flux density distribution of the source remains unchanged during the course of one observation. This is not the case for observations of the Sgr A* black hole with the EHT. In fact, the black hole at the center of our galaxy presents light-curve variability and evidence of structural variability on timescales of 5-30 minutes (Wielgus et al. 2022; EHTC 2022b; EHTC 2022d), which is considerably shorter than the ∼10 hours of observations required for a sufficient (u, v)-coverage. Two approaches can be pursued to image the brightness distribution of a source with intra-day variability. The first possibility is to treat time variability as statistical uncertainty and increase the error budget on the observed visibilities to account for it. The imaging process then recovers a "motion-blurred" image, that is representative of the time average source structure. This approach was applied by the EHTC to reconstruct the first horizon-scale image of Sgr A* (EHTC 2022c). However, the method marginalizes over time variability instead of characterizing it, limiting the analysis only to time-invariant properties of the black hole.

The second approach consists in reconstructing a minute-by-minute video of the source, where each frame corresponds to a single instantaneous observation. This is feasible if the imaging algorithm enforces correlations across neighboring frames to compensate for the extreme sparsity of each instantaneous observation, since the instantaneous coverage of the EHT array alone would be insufficient to constrain

an image. This method is called *dynamic imaging* and outputs a sequence of time-regularized frames, by propagating information from observations at different times. Like in the case of spatial regularization in classical *static imaging*, time regularization can be implemented using different frameworks by adding the temporal dimension to the domain of the flux density distribution. In RML methods, dynamic imaging is performed by introducing explicit time regularizers that favor frame-to-frame similarity and motion continuity (e.g. `eht-imaging`, Johnson et al. 2017), while Bayesian methods assume a correlated prior distribution for different frames (e.g. `StarWarps`, Bouman et al. 2018 or `Resolve`, Arras et al. 2022). The new method presented in Chapter 4 relies instead on the spectral bias of coordinate-based networks (see section 2.2.2) to provide space and time correlations.

First attempt at black hole dynamic imaging. Dynamic imaging of sources with intra-day variability is more challenging than static imaging because, with the addition of the time dimension, the imaging inverse problem becomes even more ill-posed. Therefore it is important to assess whether the (u, v)-coverage of a specific dataset is sufficient for a robust reconstruction. A first attempt at dynamic imaging of the Sgr A* black hole was made by the EHT Collaboration using the `StarWarps` pipeline (Bouman et al. 2018), which is based on a Bayesian framework. The method assumes that each frame $\boldsymbol{I}_t = I(x, y, t)$ is sampled from a multivariate Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Lambda}$:

$$\mathcal{P}(\boldsymbol{I}_t) = \mathcal{N}_{\boldsymbol{I}_t}(\boldsymbol{\mu}, \boldsymbol{\Lambda}) \,, \tag{2.27}$$

where $\boldsymbol{\mu}$ sets the image prior and $\boldsymbol{\Lambda}$ determines the amount of spatial correlation. Frames are also correlated in time, with a conditional Gaussian prior distribution

$$\mathcal{P}(\boldsymbol{I}_t|\boldsymbol{I}_{t-1}) = \mathcal{N}_{\boldsymbol{I}_t}(\boldsymbol{I}_{t-1}, \beta^{-1}\mathbb{1}) \,, \tag{2.28}$$

in which the parameter β sets the amount of temporal regularization determining the maximum frame-to-frame variability. The likelihood function is also assumed to be a multivariate Gaussian distribution:

$$\mathcal{P}(\boldsymbol{V}_t|\boldsymbol{I}_t) = \mathcal{N}_{\boldsymbol{V}_t}(\mathcal{F}_t(\boldsymbol{I}_t), \mathrm{diag}[\sigma^2]_t) \,, \tag{2.29}$$

which can be defined for visibilities, as in the above equation, or other data products. Assuming these priors and likelihoods, `StarWarps` computes the most likely image for each time t, taking into account all the observations $\{\boldsymbol{V}_t\}_{t=1,\ldots,N}$, using the elimination algorithm (Koller and Friedman 2009).

`StarWarps` was applied to EHT observations of Sgr A* on April 6 and 7 of the 2017 observing campaign (EHTC 2022c). The observed datasets were restricted to a $\sim$1.5 hour window which contained an instantaneous coverage of at least 5 antennas. The pipeline was tested on synthetic data simulated from static geometric models and GRMHD simulations. Different imaging choices, such as the prior image $\boldsymbol{\mu}$ and the amount of temporal regularization β were explored, in the reconstruction of both synthetic and real data. Dynamic imaging results showed that Sgr A* maintained a stable position angle on April 6, while on April 7 the position angle underwent a $\sim$140° rotation over the course of 100 minutes. However, these results cannot be considered sufficiently constraining because different amount of time regularization and

different image priors (e.g. a ring and a disk prior) resulted in different orientations of the position angle. `StarWarps` was able to reconstruct datasets simulated from some GRMHD models, but, in other cases, the reconstruction of GRMHD synthetic data was not successful in recovering the correct orientation of the position angle. Additionally, the pipeline appeared to introduce more variability than what present in the synthetic data analyzed. This led to the conclusion that the dynamic imaging methods available in 2022 were not able to recover a robust horizon-scale video of Sgr A* with the instantaneous coverage of the EHT array on April 6 and 7, 2017. This motivated, on one hand, the expansion of the EHT array, which progressively increased from 8 antennas in 2017 to 12 in 2025, and, on the other hand, it prompted the development of new dynamic imaging methods. Specifically, an effective dynamic method should be capable of learning the right amount of space and time correlation from the observed data, reconstruct a video with only negligible dependence on priors, hyperparameters, or initialization, and be effective even for extremely sparse (u, v)-coverages.

The major contribution of this thesis consisted in developing such an algorithm, by employing a deep learning approach based on neural fields. The proposed imaging method, named `kine`, is presented in Chapter 4 and successfully validated on EHT data in Chapter 5. Even though dynamic imaging methods were developed to address the imaging challenges of intra-day variable sources, they can be also applied to repeated observations of slowly-varying sources, to image multiple datasets simultaneously in a continuous video. Chapter 4 presents such an application, describing the advantages of multi-epoch dynamic imaging.

2.2 Deep learning and neural representations

This section consists of a general introduction to deep learning, followed by an overview of neural representation models. A comprehensive review on the topic is beyond the scope of the chapter, which rather aims to provide the reader with the notions necessary to understand the `kine` imaging algorithm presented in Chapter 4. In section 2.2.1 we present basic deep learning concepts and introduce the associated terminology. We will describe neural representation models and discuss some of their properties in section 2.2.2. Finally, in section 2.3 we describe the optical flow method for video analysis.

2.2.1 Neural networks basics

Traditional algorithms in computer science consists of a sequence of explicit instructions that are applied to solve a given task. The last decade and a half, however, has seen the rise of the field of machine learning, which aims at building algorithms that learn how to best perform a task from available data, without being explicitly programmed. ML algorithms are implemented in two steps: first a training step, where the algorithm learns how to perform the task, and then an evaluation step, where the algorithm is applied to solve the task. A subcategory of machine learning is deep learning (DL), which includes all ML algorithms that are based on neural networks (NN). Neural networks are information processing structures, whose core functioning was inspired, as the name suggests, by models of biological neurons. In this section we will introduce the basic components of a simple neural network architecture, state one of the general problems that neural networks can solve, and describe the steps of

the NN learning process.

The MLP architecture. The *node* (or *neuron*) is the fundamental unit of a neural network. A node receives an input, performs a sequence of mathematical operations on it, and returns as output the result of the operations. The simplest kind of node is a perceptron (Rosenblatt 1958), which applies the following operations (Figure 2.4):

1. Receives a real-valued vector $x = [x_1, \dots, x_n]$ as input.

2. Makes a linear combination $x \to wx + b$ of the input values with coefficients (*weights*) $w = [w_1, \dots w_n]$ and constant term (*bias*) b.

3. Applies a non-linear transformation (*activation function*) $x \to \sigma(x)$ and returns the obtained scalar value as output.

Nodes can be combined in parallel to output a multi-dimensional array. A sequence of parallel nodes is referred to as a layer. Layers can be combined in sequence to create a fully-connected multi-layer perceptron (MLP), where the output of each layer is used as input to the following layer. An MLP is one of the simplest feed-forward neural

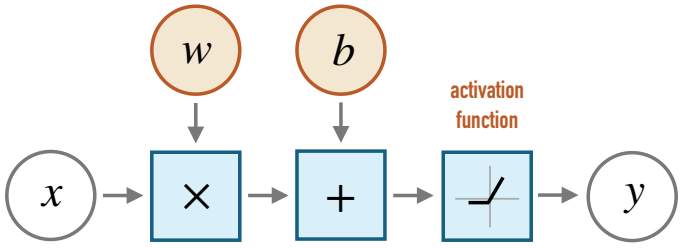

Figure 2.4: **The perceptron.** A perceptron is a node of the MLP network architecture. It performs a linear transformation of the input, followed by the application of a non-linear transformation.

network (FNN) architectures (Popescu et al. 2009), meaning that it does not have any recursive loops in its structure. It is characterized by the numbers of layers (*depth*) D, the number of neurons per layer (*width*) W, and the form of the activation function σ (Figure 2.5), all of which are the *hyperparameters* of the network. The first and last layers are called the input and output layers, while those in between are called hidden layers. Mathematically an MLP is represented by a composition of functions:

$$\mathrm{MLP}(x) = \Psi^{(D)} \circ \Psi^{(D-1)} \circ \dots \Psi^{(1)}(x) \,, \tag{2.30}$$

where each function

$$\Psi^{(i)}(x) = \sigma\left(w^{(i)}x + b^{(i)}\right) \tag{2.31}$$

represents a layer. The collection of all weights and biases $\{w^{(i)}, b^{(i)}\}$ is the set of tunable *parameters* of the MLP.

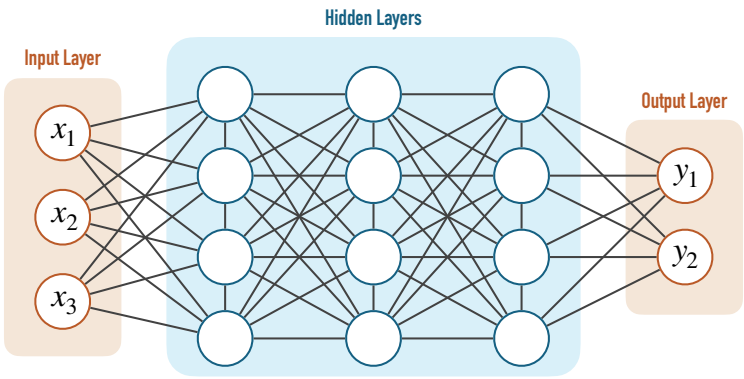

Figure 2.5: **MLP architecture.** An MLP is a simple feed-forward neural network that processes the input through a sequence of fully-connected layers. For each layer, each line represents a weight w_{ij}, while blue circles represent the bias b_j and activation function. Orange circles represent input and output values. In this example, the MLP has a depth of 4 and a width of 4.

Universal function approximators. Neural networks can be employed for a wide variety of tasks, one of them being function approximation. In particular, NN are especially useful to approximate complex functions that are difficult or impossible to define explicitly. In this context, the general problem consists in finding the function f that maps an input $x \in X$ to the corresponding output $f(x) = y \in Y$, without knowing the explicit form of f. A possible approach is to find the best approximation to f by defining a parametrized function f^*_θ and tuning the parameters θ until f^* converges to f. The function f^* can be provided by a neural network, for example an MLP, so that f^* will consists of an iterative composition of linear and non-linear functions, parametrized by the weights and biases of the network. Cybenko (1989) and Hornik (1991) proved that a sufficiently wide FNN, with at least one hidden layer and a non-linear activation function, can approximate any continuous function on a compact domain to an arbitrary degree of accuracy. This is known as the Universal Approximation Theorem and explains why neural networks are optimal tools for function approximation. In particular, the non-linearity of the activation function is the key factor that enables a NN to express complex functions. However, the Universal Approximation Theorem does not provide indications on how to find the best network parameters or what is the most efficient architecture. Therefore, the choice of the network hyperparameters and optimization method are crucial to steer the optimization in the right direction and achieve a good approximation. The VLBI imaging algorithm presented in 4, employs a neural network to approximate the function describing the time and space distribution of the source's flux density. This means that function modeled by the network is the image or video itself, rather than an operation that produces the image or video.

Training procedure. To approximate a desired function, a neural network needs first to be trained, which means optimizing its parameters so that the difference $|f - f^*_\theta|$ is minimized. Since the function f is unknown, in practice one needs to define a

different but equivalent loss function $\mathcal{L}$, which is also minimized when $f^* \to f$. There are two kinds of training procedure, named *supervised* and *unsupervised* learning. Supervised learning is usually applied when the DL algorithm needs to learn the underlying relation between the elements of two datasets and it requires a labeled set of input-output pairs $\{(x^{(i)}, y^{(i)})\}$, called *training set*. The network uses the pairs in the training set as reference examples to optimize its parameters to reproduce the correct output for any given input. In this case, the loss function is usually a function of the distance between the labeled outputs y_i and the estimated network outputs $y_i^* = f^*(x_i)$, averaged over the training set:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^{N} ||y_i - y_i^*|| \ . \tag{2.32}$$

Unsupervised learning, instead, is performed without labeled example data and it is applied when the algorithm is required to learn implicit patterns and correlations in the data. In this case, the loss function will include a term that constrains the output to be compatible with available data and possibly terms that constrains specific properties of the output.

In both cases the loss function depends on the network's output, which in turn depends on the network's parameters. Training the network consists of applying a gradient descent optimization on the loss function. This is based on the computation of the loss function derivatives with respect to the network parameters by applying the chain rule, a procedure referred to as *backpropagation*. At each iteration of the gradient descent, the weights and biases of the network are updated to:

$$\theta \to \theta - \alpha \frac{\partial \mathcal{L}}{\partial \theta} \ , \tag{2.33}$$

where the *learning rate* (LR) α determines the extent of the gradient descent step and the speed of the convergence of the training.

Computing the full loss function gradient is often a computationally expensive operation. In this case, the gradient descent optimization can be replaced by the stochastic gradient descent, which consist in computing the loss function only over a subset of randomly selected training data, called *batch*, which changes at every iteration. This enables faster iterations at the expense of the convergence rate. Different variants of the gradient descent can be used as optimization procedure. One of the most commonly used is Adam (Kingma and Ba 2014), an adaptive moment estimation optimizer, which adapts the learning rate of each parameter independently, depending on the average and variance of the last few gradients. This algorithm is appropriate for optimization problems with a large set of parameters and noisy gradients.

Overfitting, validation, and hyperparameters. During the training phase, the network learns to produce the desired output corresponding to the input data available in the training set. However, the NN must also be able to provide a correct output for input data outside of the training set. The risk indeed is that the network learns a function that overfits the training set and is not able to generalize further. This issue is addressed in the *validation* phase, in which the NN is evaluated on a separate set of data. As opposed to the training phase, during validation, the trained network is applied and evaluated on new data (the *validation set*) without varying the network

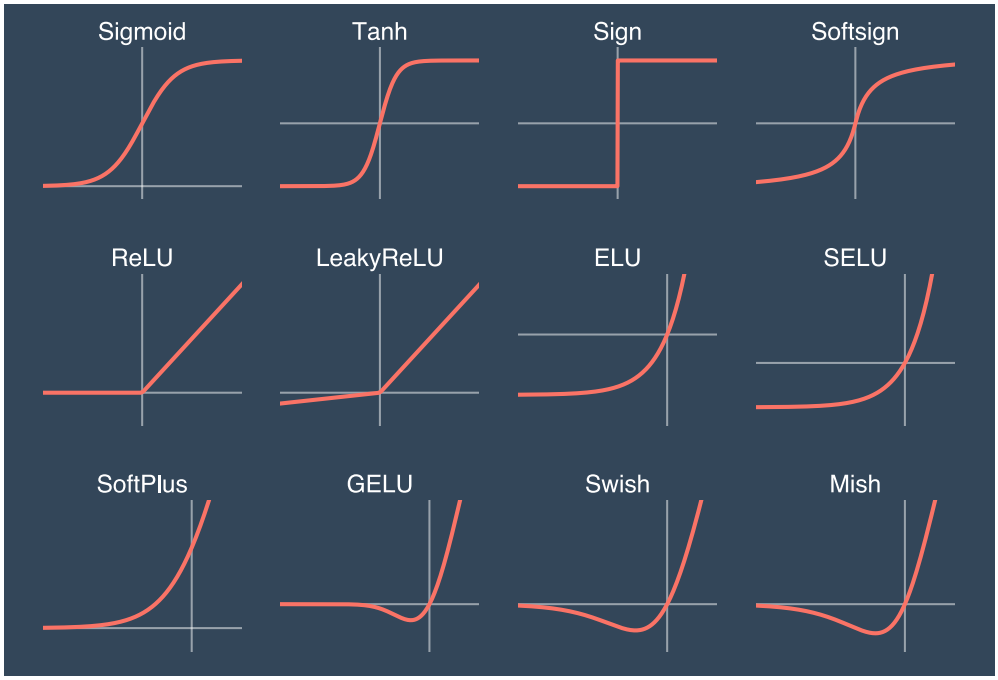

Figure 2.6: **Activation functions.** Examples of the most common activation functions.

parameters. The validation phase is usually alternated to the training phase and the hyperparameters are manually adjusted to find which architecture and settings are the most efficient in minimizing the loss function and provide the best performance. After training and validation, the *testing* phase takes place. In this final step the network is applied to a new set of data (the *testing set*), without varying the parameters nor the hyperparameters, in order to evaluate the final network's performance. In the case of the MLP architecture, the hyperparameters that may be adjusted during validation include: the depth and width of the network, additional connections among nodes (*skip connections*), the activation function (some examples are shown in Figure 2.6), the learning rate, the batch size, the weight of each loss term, and the total number of iterations.

The width and depth determine the amount of information that the network can process and therefore its expressivity. The more complex the function that the NN should learn, the higher the number of necessary layers and nodes. However, an excessive increase of the number of tunable parameters may result in instabilities in the backpropagation and longer training time. Therefore the width and depth should be kept to the minimum value sufficient to provide the required expressivity. Skip connections can help propagate information and structures from the input to further layers. The activation function can affect the stability of the training and smoothness properties of the output. The activation of the output layer may differ from that of the hidden layers in order to constrain properties of the output such as the range of possible values. The learning rate is a critical parameter that affects the convergence of the optimization algorithm. It is possible to schedule changes in the learning rate after a

set number of iterations. A common LR schedule consist in periodically decreasing the learning rate, to guide the optimization into the correct minima of the loss function.

A possible problem that may affect training is vanishing or diverging loss function gradients. To help prevent this, it is important that the values passed to the activation functions belong to an interval close to $[-1, 1]$, which is where the activations are the most non-linear. One way to achieve this is through *batch normalization*, which normalizes the input to each activation with respect to the present batch. Specifically, batch normalization consists in the following steps:

1. Given a set of N training data $\{x_i\}$, and the corresponding inputs $\{z_i\}$ to the activation function of a specific node, compute the average μ and standard deviation σ of z:

$$\mu = \frac{1}{N}\sum_{i=1}^{N} z_i \quad , \quad \sigma^2 = \frac{1}{N}\sum_{i=1}^{N}(\mu - z_i)^2 \quad . \tag{2.34}$$

2. Normalize the activation input:

$$z_i' = \frac{z_i - \mu}{\sqrt{\sigma^2 + \epsilon}} \quad , \quad \text{with } \epsilon \ll 0 \tag{2.35}$$

   so that the new input has zero mean and unit variance over the dataset.

3. Apply the linear transformation:

$$z_i'' = \gamma z_i' + \beta \tag{2.36}$$

   and pass z_i'' to the activation function. γ and β are tunable parameters of the network, which ensure that the normalized values fall in the optimal range.

In addition to avoid gradient issues, since batch normalization maintains activation inputs in a constant range, it allows higher learning rates to be used, making training faster. Batch normalization can be applied even when batching is not used, by considering the whole training set as a batch.

2.2.2 Neural field representation

One possible application of neural networks is to model physical fields, through what is called *neural field representation*. Neural fields have found applications in a variety of tasks in computer vision, from image synthesis, to volume modeling, scene rendering, and dynamics reconstruction. One of the most popular applications is the representation of complex 3D scenes from a small set of 2D projections, using a neural radiance field (NeRF, Mildenhall et al. 2021), which showed that high-resolution 3D scenes can be recovered from partial, compressed information.

Following Xie et al. (2022), we define a field to be an m-dimensional physical quantity defined over a continuous n-dimensional domain. In this context, we restrict the field to be defined over low-dimensional (typically $n \in \{2, 3, 4\}$) spatio-temporal domains. The field is also of low dimensionality, with $m \sim n$. A field can be represented by a function f mapping the coordinates $x \in \mathbb{R}^n$ of the domain to the corresponding value $y \in \mathbb{R}^m$ of the physical quantity. The analytical form of such a mapping may

be unknown, but it may be approximated by a parametrized function $f_\theta(x)$ of known form. We define a *neural field* as a field that is parametrized by a neural network. The neural networks employed in neural fields are called *coordinate-based* networks because their input consists in the continuous coordinates of the field domain, while the output is the field value corresponding to the input coordinates.

Usually coordinate-based networks are MLPs, with differentiable activation functions, whose parameters are optimized through gradient descent. This is opposed to other architectures, such as convolutional neural networks (CNNs), which instead rely on a discretization of the spatial domain into pixels or voxels. The inputs and outputs of CNNs consists of high-dimensional arrays or matrices of sampled field values. An advantage of neural fields over the discrete parametrizations is that they scale well with increasing resolution of the domain size, which is especially important for 3D or 4D domains. Additionally, the continuous parametrization provided by neural fields can be sampled at any point of the domain, instead of being restricted to the pixel grid, granting smooth and correct interpolation.

Training a neural field. The optimization of a neural field generally consists in the iteration of the following steps, which are schematized in Figure 2.7:

1. Sample the coordinates of the input domain and feed them to the network to obtain the corresponding sampled field values.

2. Apply a measurement operator (or *forward map*) that estimates observable quantities from the samples of the reconstructed field.

3. Compute a loss function based on the distance between the estimated observables and the observed data.

4. Optimize the network parameters through backpropagation so to minimize the loss function.

After training, the network can be sampled on a set of domain points, typically denser than those sampled during training, to evaluate the recovered field. Training a neural field is equivalent to the standard optimization of a parametrized function. It differs from traditional training because no training set is involved, making the algorithm completely unsupervised. Therefore in this context the terms "training" and "optimization" can be used interchangeably. The VLBI dynamic imaging algorithm that is presented in Chapter 4, and applied in Chapter 5, relies on a neural field to model the flux density distribution in the video or image of the source. In that case, the input domain consists of two linear spatial coordinates (right ascension and declination) and one time dimension, while the modeled field consists of the total intensity flux, the fractional linear polarization, the electric vector position angle and the fractional circular polarization. The measurement operator consists of a two-dimensional spatial Fourier transform, sampled in correspondence with the (u, v)-coverage of the observed data, while the loss function contains the χ^2 distance between observed and estimated data products.

Spectral bias and positional encoding. Even though, coordinate-based neural networks are, in principle, universal approximators of high expressivity, when initialized to random distributions and optimized with gradient descent, they are biased

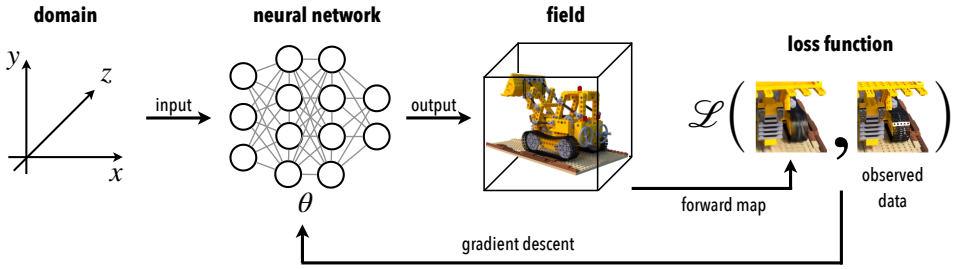

Figure 2.7: **Neural field training.** Diagram of the components of a neural field and the steps required for its training.

towards outputting low-frequency functions (Rahaman et al. 2019). This property is called *spectral bias* and arises because the network learns the low-frequency components of the output faster than the high-frequency ones. This property has been demonstrated in the limit of an infinitely wide, fully-connected, neural network, by proving that training the network with gradient descent is equivalent to performing a kernel regression (Lee et al. 2018; Jacot et al. 2018). In computer vision, where most applications involve high-resolution natural images, spectral bias is often an undesirable property. Instead, in the field of VLBI imaging, where the ill-posed inverse problem of image reconstruction requires additional prior information, spectral bias is often useful to provide implicit regularization of the image.

It is possible to overcome spectral bias and enable MLPs to output high-frequency function. This may be achieved with *positional encoding* of the input coordinates, which consists of mapping each coordinate to a higher-frequency function through a Fourier feature mapping (Vaswani et al. 2017). The higher frequency of the input then enables the possibility to achieve higher frequencies in the output too (Tancik et al. 2020). Mathematically this is performed by applying the transformation:

$$\gamma(x) = \left[\sin(2^0\pi x), \cos(2^0\pi x), \ldots, \ldots, \sin(2^{L-1}\pi x), \cos(2^{L-1}\pi x)\right] \tag{2.37}$$

to every input coordinate x and feeding the result to the first layer of the network. L is the degree of the positional encoding and sets the high-frequency limit of the output field. Figure 2.8 shows the effect of positional encoding of the input coordinates for different neural field applications. The images clearly show that Fourier features allow the output field to achieve higher resolutions.

2.3 Optical flow

The development of imaging algorithms capable of recovering high-resolution videos of variable radio sources prompts the employment of advanced motion analysis techniques to study the dynamics of the evolving plasma. One possible approach is to track kinematics by computing the *optical flow* (OF), that is a two-dimensional vector field which quantifies the projected apparent motion of visual elements between consecutive frames (Gibson 1950).

The optical flow is based on the assumptions of brightness constancy and infinitesimal motion between frames, meaning that the brightness of each pixel is conserved in

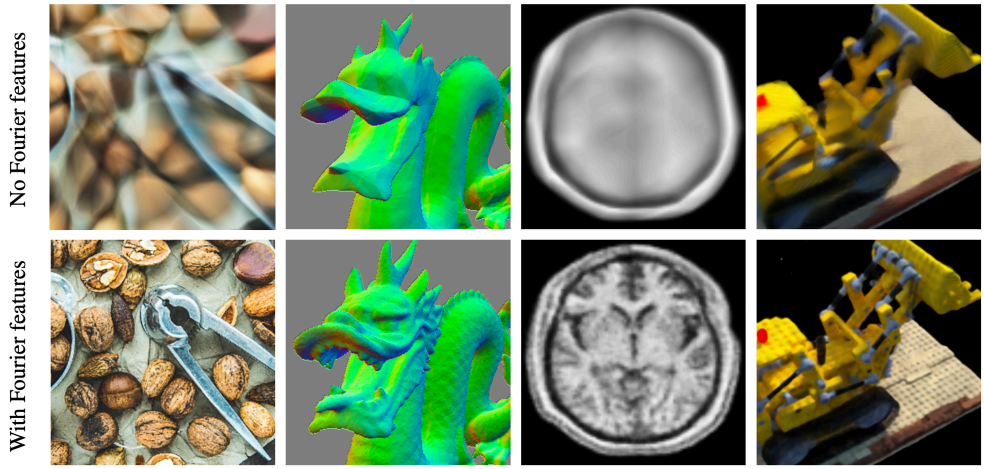

Figure 2.8: **Positional encoding.** The first rows shows the output neural field models when no Fourier features are applied, while the second shows the result of the application of Fourier features. Figure adapted from Tancik et al. (2020).

the infinitesimal displacement undergone from one frame to the other. For two adjacent frames $I(x, y, t)$ and $I(x, y, \delta t)$ and an optical flow velocity field $(u(x, y), v(x, y))$, the first assumption requires the brightness of a displaced pixel to remain constant:

$$I(x + u\delta t, y + v\delta t, \delta t) = I(x, y, t) \ , \tag{2.38}$$

while the second allows to approximate the equation to its first-order Taylor expansion, leading to:

$$I_x u + I_y v + I_t = 0 \ , \tag{2.39}$$

where $I_x = \frac{\partial I}{\partial x}$, $I_y = \frac{\partial I}{\partial y}$, $I_t = \frac{\partial I}{\partial t}$. Solving this equation for every pixel (x, y) of the frame, requires additional constraints since the unknowns are twice as much the equations.

Constraints can be provided by regularization of the optical flow field, for example by assuming a continuous OF from pixel to pixel. This is implemented in the Horn-Schunck (HS) algorithm (Horn and Schunck 1981), which redefines the OF recovery problem in terms of the minimization of the functional:

$$E(u, v) = \sum_{x,y} \left(I_x u + I_y v + I_t\right)^2 + \lambda^2 \left(||\nabla u||^2 + ||\nabla v||^2\right) \ , \tag{2.40}$$

where the first term enforces brightness constancy, while the second, weighted by the positive parameter α^2, favors smoothness, by aiming to minimize the gradients ∇u, ∇v of the optical flow components. The optical flow that minimizes the functional is

found by the iterative solution:

$$
\begin{aligned}
u^{k+1} &= \bar{u}^k - \frac{I_x(I_x\bar{u}^k + I_y\bar{v}^k + I_t)}{\alpha^2 + I_x^2 + I_y^2} \\
v^{k+1} &= \bar{v}^k - \frac{I_x(I_x\bar{u}^k + I_y\bar{v}^k + I_t)}{\alpha^2 + I_x^2 + I_y^2} \,,
\end{aligned}
\tag{2.41}
$$

where $\bar{u}^k(x,y)$ and $\bar{v}^k(x,y)$ are the averages of u and v in the neighborhood of the pixel (x,y).

Thanks to the smoothness assumption, the main advantage of the Horn-Schunck algorithm is that it is able to provide a good estimate of the optical flow for all pixels in the image, even in areas of homogeneous brightness ($\nabla I \sim 0$). However, since it relies on the assumption of small motion between frames, it performs poorly in recovering large displacements.

This can be solved by implementing HS into a multi-scale approach that reconstructs a progressively finer optical flow field (Meinhardt-Llopis et al. 2013). This approach starts by computing a sequence of progressively downsampled versions of the video frames. The optical flow is initialized to 0 and the HS algorithm is applied to the coarsest pair of frames. The resulting optical flow is then used as initialization for the optical flow computation in the next finer pair of frames, until the finest resolution is reached.

A multi-scale implementation of the Horn-Schunck algorithm was employed in Chapter 4 to estimate the optical flow in a video reconstruction of multi-epoch observations of blazar 3C 345. In this application, the optical flow allowed to measure the apparent velocity field of the plasma moving in the relativistic jet.

3

Evolution of a relativistic jet

Adapted from

Evolution, speed, and precession of the parsec-scale jet in the 3C 84 radio galaxy

M. Foschi, J. L. Gómez, A. Fuentes, I. Cho, A. P. Marscher, S. Jorstad

Astronomy and Astrophysics, 696, A17, (2025)

Abstract

We present high-resolution images of the radio source 3C 84 at 43 GHz from 121 observations conducted by the BEAM-ME monitoring program between 2010 and 2023. Imaging was performed using the recent forward modeling imaging method `eht-imaging` and achieved a resolution of 80 μas, which is a factor of $\sim$2-3 better than traditional imaging methods such as CLEAN. The sequence of images depicts the growth and expansion of the parsec-scale relativistic jet in 3C 84; it clearly shows a complex internal structure with bending in the jet and changes in its launching direction and expansion speed. We report measurements of the expansion speed over time, which show that the jet goes through three regimes, marked by the start and end of a hot spot frustration phase. The high resolution of the images also allowed us to measure the projected launching direction as a function of time, and we find an irregular variation pattern. Our results confirm previous studies of the morphological transition undergone by 3C 84 and provide quantitative measurements of the jet's kinematic properties over a decade-long timescale.

3.1 Introduction

Some supermassive black holes at the center of galaxies generate collimated jets of ionized relativistic particles, which are accelerated by the strong magnetic fields surrounding the black hole and the accretion disk. These highly energetic and luminous jets propagate through the host galaxy and beyond, interacting with the interstellar (ISM), intergalactic, and intracluster media (ICM). Temperature, density, and pressure differences between the plasma in the jet and the surrounding medium influence the jet's expansion by altering its shape and profile and by affecting its direction and expansion speed.

The relativistic jet in the radio galaxy 3C 84 (NGC 1275), located in the Perseus Cluster, is a valuable source of information about the interactions between the ICM and the parsec-scale jet. NGC 1275 is at the center of a strong cooling flow cluster, where large amounts of inflowing gas are reheated by the galaxy's active nucleus. The galaxy also displays bright X-ray emission (Forman et al. 1972), and observations by Chandra have shown pairs of opposed bubbles in the ICM located at different distances and directions relative to the 3C 84 radio source (Fabian et al. 2003). A proposed explanation for these structures is that they are inflated by a precessing and restarting of a jet–counterjet pair (Dunn et al. 2006).

3C 84 is also a variable radio source (Dent 1966; Pauliny-Toth and Kellermann 1966) and has been observed since the 1950s through multiple Very Long Baseline Interferometry observations at both millimeter and centimeter wavelengths. Giovannini et al. (2018) and Savolainen et al. (2023) provide a historical overview of the radio observations of 3C 84, of which we give a short summary here. The source presents various lobe-like structures south and north of the core, on parsec (e.g., Walker et al. 2000; Asada et al. 2006) to kiloparsec (Pedlar et al. 1990) scales, which may indicate a repeatedly restarting jet. Observations at the parsec scale from the 2010s have shown the presence of a dim radio lobe (C2; Nagai et al. 2010) and a bright radio lobe (C3), the latter of which is connected to the core (see the top-left panel in our Figure 3.2; Nagai et al. 2014). The C2 component was ejected in the early 1960s during a period of increasing brightness that lasted until the mid 1980s, before dropping in the 1990s and early 2000s (Nesterov et al. 1995). C3 was emitted around 2003 (Suzuki et al. 2012; Nagai et al. 2017) during a second period of high brightness. While C3 propagated from the core through a limb-brightened structure, the propagation of C2 happened through a jet with a centrally peaked morphology. In more recent years, higher-resolution observations have been able to resolve the internal structure of the parsec-scale jet in 3C 84. Giovannini et al. (2018) presented results at 22 GHz from a global array of ground antennas plus the space antenna RadioAstron (Kardashev and Khartov 2013). The reconstructed image clearly shows strong limb-brightening and a wide opening angle near the core, followed by a quasi-cylindrical jet profile. The jet ends with a bright spot with surrounding diffuse emission. Giovannini et al. (2018) suggest that the cylindrical profile may be due to the jet connecting C1 to C3 being embedded in a uniform-pressure cavity carved by past activity of the jet. This is supported by other observations by RadioAstron at 5 GHz (Savolainen et al. 2023), which show that the C2 and C3 components are both surrounded by low-intensity emission from a cocoon-like structure. Savolainen et al. (2023) discuss that, in the interaction between the C3 structure and the ISM, energy is transferred to the ISM, heating the

gas that forms the cocoon. They also suggest that the cocoon-like structure could be caused by the jet moving through a multiphase medium consisting of gas clouds of different sizes and densities. The embedding of the jet in a clumpy medium is supported by results from Nagai et al. (2017) and Kino et al. (2018) and Kino et al. (2021). In particular, Kino et al. (2021) analyzed 43 GHz images of 3C 84 from the Very Long Baseline Array (VLBA) from 2012 to 2020. They tracked the motion of a hot spot and, in 2016-2017, observed a year-long frustration phase, during which the hot spot followed a circular trajectory after reaching the edge of C3, rather than propagating farther through the jet. They attribute this event to a collision between the head of the jet and a compact dense cloud. After the collision, the jet breaks through the cloud, deviating its expansion direction to the west and transitioning from an FR II- to FR I-class radio lobe morphology. However, despite the significant number of studies, observations of 3C 84 have provided either repeated images of the jet at low resolution (Kino et al. 2018; Kino et al. 2021) or hard-to-repeat single-epoch images at high resolution (Giovannini et al. 2018). This hinders a proper study of the kinematics of the plasma in the jets and the dynamics of the jet expansion.

In this work we present a reimaging of all 121 VLBA observations of 3C 84 at 43 GHz obtained by the BEAM-ME monitoring program (Jorstad and Marscher 2016) from 2010 to 2023 obtained using the regularized maximum likelihood (RML) imaging method `eht-imaging` (Chael et al. 2018). With `eht-imaging`, we obtained images of the parsec-scale jet at a resolution of $\sim$80 μas. This is $\sim$2-3 times higher than the nominal beam used to convolve CLEAN images, whose average across different epochs is (280, 150) μas. RML methods produce super-resolved images by incorporating reasonable prior assumptions that regularize the image. These methods have proven to achieve higher fidelity at super-resolution than CLEAN (see, e.g., Fuentes et al. 2023). Thanks to the combination of the super-resolving power of RML methods and the constant monitoring by the BEAM-ME program, we were able to observe, for the first time, the evolution of the overall and internal structure dynamics of the parsec-scale jet, over a 12-year period. The images we present resolve the hot spots and the internal structure of the jet, as well as the connection between the limb-brightened structure and the core. At this resolution, it is also possible to resolve the front of the jet head, discerning the expansion of the jet from the motion of components through the jet. We considered the source redshift, z, to be 0.0176 (Strauss et al. 1992). In continuity with previous publications on 3C 84, we assumed a Λ cold dark matter cosmology with $H_0 = 70.7$ km s^{-1} Mpc^{-1}, $\Omega_M = 0.27$, and $\Omega_\Lambda = 0.73$; this means that 1 mas in the image plane corresponds to 0.35 pc.

The chapter is organized as follows: in section 3.2 we provide details of the observed data and explain the method used to image them. In section 3.3 we present the imaging results, provide a quantitative estimate of the jet's speed and direction, and discuss the evolution of the jet in the context of previous observations of the source. We summarize and discuss our results in section 3.4.

3.2 Data and imaging

3.2.1 Observations

We analyzed data from the BEAM-ME monitoring program conducted by Boston University (previously named VLBA-BU-BLAZAR)[1], which observes multiple gamma-ray blazars and radio sources using the VLBA at 43 and 86 GHz (Jorstad and Marscher 2016). We focused on total intensity observations of the radio source 3C 84 in the Perseus cluster, at 43 GHz, conducted on a roughly monthly basis from late 2010 until early 2023, resulting in a total of 121 individual epochs.

The BEAM-ME data are already fully calibrated and self-calibrated to the CLEAN images provided in the archive. However, to avoid any bias from possible residual calibration errors or from the self-calibration, we chose to run the first imaging step using only closure quantities. The archival data were already time averaged with a 30 s interval, so no additional time averaging was performed before imaging.

3.2.2 Imaging procedure

The data were imaged using `eht-imaging`, a forward modeling imaging method for VLBI observations (Chael et al. 2018). This method defines the image as a discrete square matrix of flux density values, $\boldsymbol{I} = \{I_{ij}\}$, and optimizes these values to minimize the objective function:

$$J(\boldsymbol{I}) = \sum_D \alpha_D \chi^2_D(\boldsymbol{I}, \boldsymbol{V}) - \sum_R \beta_R S_R(\boldsymbol{I}) \,, \tag{3.1}$$

where the first sum runs over the reduced χ^2 of different data products D computed from the image $\boldsymbol{I}$ and the observed data $\boldsymbol{V}$, while the second sum runs over various regularizers R that impose additional correlations among pixel values, thus constraining the possible solutions to the ill-posed problem of VLBI imaging (EHTC 2019d). The coefficients α_D and β_R that weight the data terms and the regularizers are hyperparameters of the method.

The imaging procedure followed these steps:

1. Iteratively apply `eht-imaging`'s optimization step, using only log closure amplitudes and closure phases as data products in the objective function. Stop the optimization when the reduced χ^2 of both data products decreases by less than 2% in a single step.

2. Perform self-calibration of amplitudes and phases to the obtained image, to correct potential residual station-based errors.

3. Reapply `eht-imaging`'s optimization steps, this time using both complex visibilities and closure quantities as data products in the objective function. Stop the optimization when the χ^2 of both data products decreases by less than 1%.

The `eht-imaging` optimization step mentioned in point 1 and 3 consists of alternating between a series of quasi-Newton gradient descent steps and blurring of the resulting image with a Gaussian kernel with a full width at half maximum (FWHM) equal to 1.5 times the nominal resolution of the array ($\sim$150 μas). The blurring step prevents the

[1] https://www.bu.edu/blazars/BEAM-ME.html

optimizer from getting trapped in local minima of the objective function. The images obtained with this procedure provide a good fit to the data, with an average reduced visibility χ^2 of 1.61. We note that, in most cases, the images obtained after step 1 are very similar to those obtained after step 3, meaning that the closure quantities are sufficient to constrain the images and that self-calibration only provides small refinements to the final images. The gains obtained from the self-calibration step were negligible, which was expected since the archival data were already self-calibrated using CLEAN imaging.

Field of view. The field of view was increased linearly from 5.1 mas for the earliest epoch to 13.2 mas for the latest one, following the growth of the emitting region in the jet, as seen in the archival CLEAN images. Accordingly, the number of pixels was increased from 170 to 440, while the pixel size remained constant at 30 μas[2].

Initialization image. The initialization image for the optimization process was chosen to be an elliptical Gaussian with major axis rotated 6° clockwise from the north, with (FWHM$_x$, FWHM$_y$) values ranging linearly from (385, 1375) μas to (910, 3250) μas to match the average direction and increasing dimensions of the jet. For a few epochs (specifically, from April 30, 2022, to December 6, 2022), the optimizer would not converge if initialized to a Gaussian. We attribute this to the combination of a suboptimal coverage and a complex jet morphology, which could not be well approximated by a Gaussian. In these cases, we initialized the optimizer to the CLEAN image provided by the BEAM-ME program, blurred with a Gaussian kernel of FWHM equal to twice the nominal resolution of the array. We tested the effect of initializing with the CLEAN image instead of a Gaussian on the other epochs and found that the final image was not noticeably affected by the choice of prior image.

Regularizers. We made minimal use of regularizers, since the (u, v)-coverage of the array is sufficiently dense. In step 1 of the imaging procedure, the total flux of the image cannot be constrained by closure amplitudes, so we used a flux regularizer to constrain it to the maximum amplitude of the shortest baseline. In both steps 1 and 2, we used the entropy regularizer to constrain the emission in the center of the image and the ℓ_1 regularizer to encourage sparsity, since significant portions of the images were expected to have no emission. The exact regularizers' weights, along with other imaging parameters, are reported in Table 3.1.

Images. Imaging results from all epochs are presented in Figure 3.1 and in the corresponding movie (`https://www.aanda.org/articles/aa/olm/2025/04/aa53406-24/aa53406-24.html`), while Figure 3.2 shows the comparison between CLEAN and `eht-imaging` images for a few selected epochs. The `eht-imaging` images we obtained are consistent with the CLEAN images provided by the BEAM-ME program, but in higher resolution. This enables a more precise detection of the outlines of the bright limbs and the edge of the jet head, as well as the resolution of intra-jet features and the jet orientation at the subparsec scale.

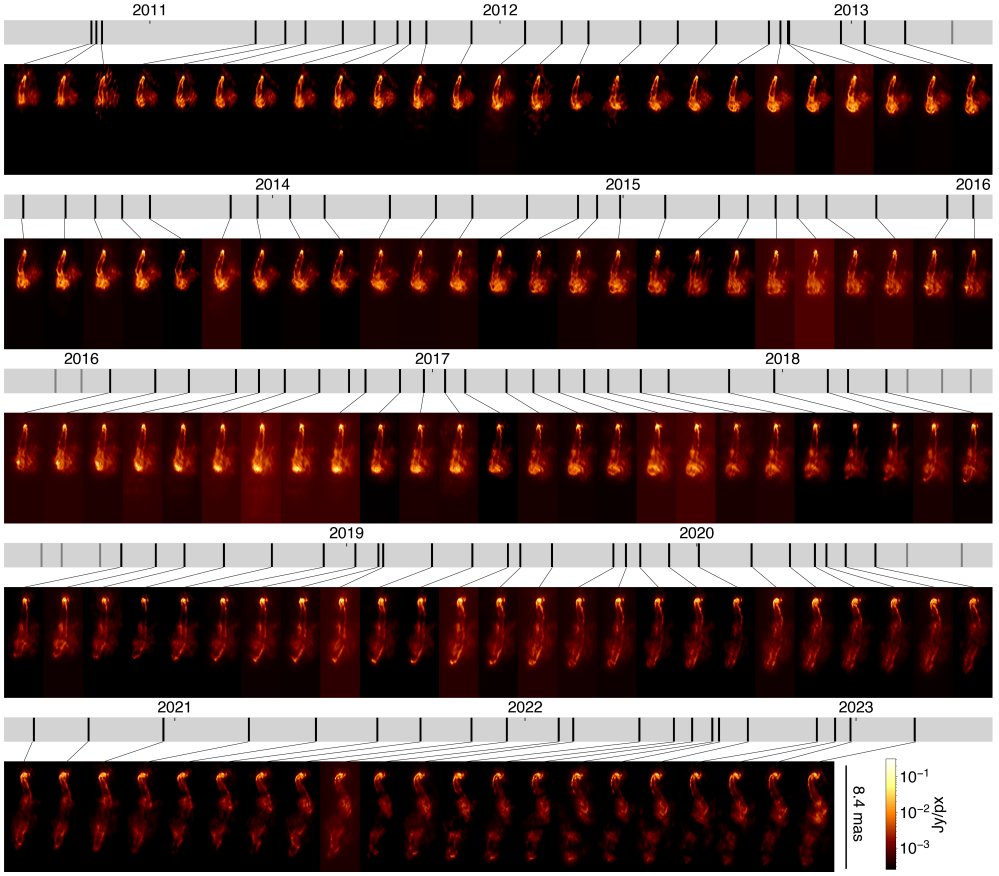

Figure 3.1: **Evolution of 3C 84 over 12 years.** Radio source 3C 84 as observed from 2010 to 2023 by the VLBA at 43 GHz and imaged with `eht-imaging`. The black vertical lines mark the dates of the observation epoch corresponding to each image. The time evolution of the source is shown in the movie available at https://www.aanda.org/articles/aa/olm/2025/04/aa53406-24/aa53406-24.html.

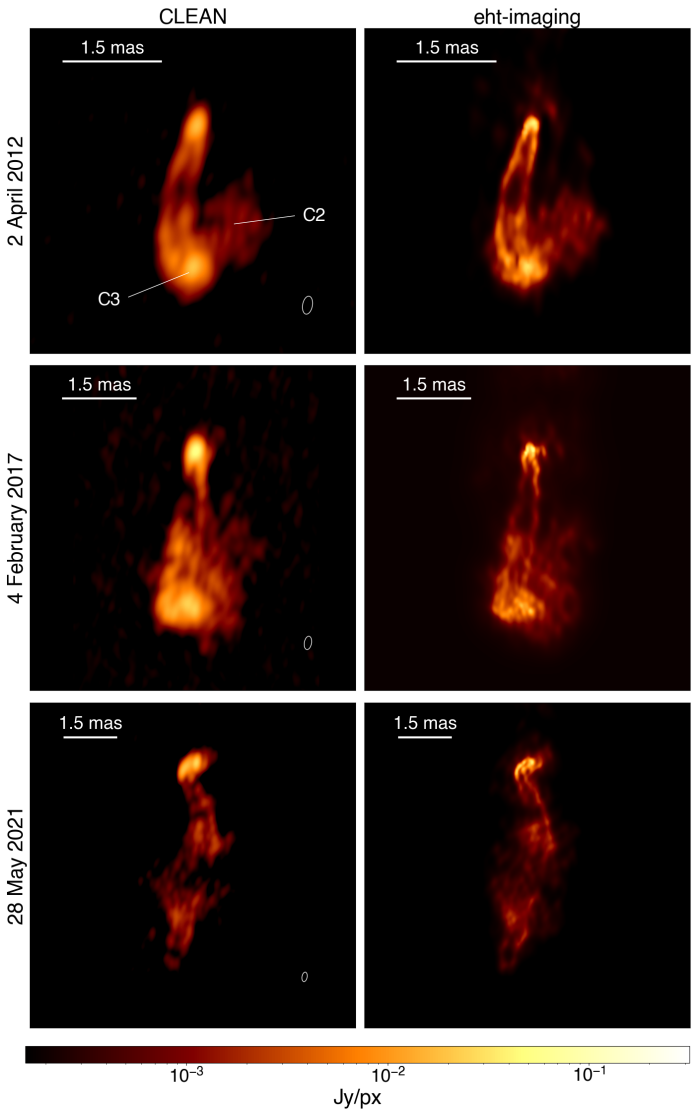

Figure 3.2: **CLEAN and eht-imaging comparison.** Comparison between images of radio source 3C 84 from VLBA observations at 43 GHz, obtained with the CLEAN (left) and `eht-imaging` (right) imaging methods.

Image parameters	**Values**
pixel size	30 μas
number of pixels	170 - 440
field of view	5.1 - 13.2 mas
Observation parameters	
added systematic error	0.5 % (of visibility amplitudes)
time averaging	30 s
Prior parameters	
prior image	Gaussian / CLEAN image
Gaussian fwhm (maj axis)	1375 - 3250 μas
Gaussian fwhm (min axis)	385 - 910 μas
Gaussian orientation	-6°
Optimization parameters	
max N iterations	200
data term weights (steps 1 - 2)	
vis	0 - 1
amp	0 - 1
cphase	1 - 1
logcamp	1 - 1
regularizer weights (steps 1 - 2)	
flux	50 - 10
entropy	0.1 - 1
ℓ_1	0.1 - 10

Table 3.1: **Imaging hyperparameters** used in the `eht-imaging` pipeline.

3.3 Jet analysis

The sequence of jet images in Figure 3.1 reveals various features of the jet evolution. In the first images (corresponding to epochs in late 2010), the jet presents clear limb-brightening, with the two limbs originating from an unresolved core and undergoing a slight counterclockwise bending. The limb-brightened structure extends until the core, to which it connects at a high opening angle, confirming the analysis from Giovannini et al. (2018), which resolved the limb structure at a distance of 0.03 mas (or 350 gravitational radii de-projected). Over the years, the limb-brightening persists as the jet increases in length, reaching $\sim$4 times its initial size. A dim emission feature (C2) is evident in the bottom-right part of the jet in 2011-2013. From 2014, C2 gradually becomes dimmer and more elongated, as if dragged by the expanding C3 component. A single localized bright spot from C2 is still visible in 2015, before the component disappears into the diffuse emission on the right side of the jet, barely visible above the noise level until 2019. A detailed analysis of the C2 component is beyond the scope of the present chapter, but future works may extract information about the energy exchange associated with its disappearance.

The jet launching direction changes over the years, slowly rotating counterclockwise from 2011 to 2015, slowly rotating back to the initial direction from 2015 to 2016,

[2]The images are cropped with respect to the original field of view, to be displayed in the most effective way.

maintaining the same direction during 2017-2018, then drastically rotating clockwise from 2019 to 2022. From 2019, the core undergoes a severe twisting due to the clockwise rotation, which complicates its structure, while in 2017 a secondary component appears west of the core and persists until the end of the considered time window. A complex picture of the core was already described by Punsly et al. (2021), who modeled the core structure with two or three components aligned in the east-west direction. The 2019 twist is also reported by Park et al. (2024b), who describe the ejection of a knot east of the core (Paraschos et al. 2022) and track its motions as it abruptly changes direction toward the south. Park et al. (2024b) suggest that the knot's deflection may be due to the jet colliding with a dense clumpy cloud. In our images, however, we do not see the ejection of an individual knot, but rather we observe emission appearing west of the core because of the twisting of the jet rails. This scenario also explains the apparent southward deflection of the knot, which moves along the preexisting jet structure.

Various components can be tracked moving along the jet (Hodgson et al. 2021), notably a bright spot is seen approaching the head of the jet, "bouncing" against it and then dissipating, in the time span between late 2015 and early 2017. Our results confirm the hot spot's counterclockwise trajectory reported by Kino et al. (2021). However, our images show that the hot spot appears at the end of 2015 and dissipates at the beginning of 2018 as the jet pierces through the lobe. The 2015 hot spot's flip and the 2018 hot spot's breakout reported by Kino et al. (2021), should be attributed to a component mismatch caused by insufficient resolution.

At the end of 2010, the jet presents a straight morphology. However, a lobe begins to form in late 2011 and undergoes a significant expansion from early 2013 to early 2017. The inflation coincides with a slowdown in the jet's expansion velocity, followed, in 2018, by a burst through the inflated bubble and an increase in the expansion speed. This confirms the abrupt morphological transition from an FR II- to FR I-class radio lobe observed by Kino et al. (2021), to which we add a gradual opposite transition from FR I to FR II observed from 2010 to 2013. From late 2020 to the end of the considered epochs, some portions of the jet appear darkened. This could be caused by a lower local emissivity, a change in the viewing angle, or the presence of an absorbing foreground. With respect to this last hypothesis, it should be noted that 3C 84 is likely surrounded by an accretion disk associated with ionized gas, which absorbs and obscures the inner section of the counterjet (Walker et al. 2000; Fujita and Nagai 2017).

Some of our images also show the presence of an emission region north of the core. From 2011 to 2013 we detect a persistent emission 0.5 mas north-west of the core, while from 2017 we detect a more diffuse emission 2-3 mas north of the core, coinciding with the N1 component detected by Fujita and Nagai (2017). However, in some epochs, the emission cannot be detected above the noise level of the image background. Whenever the northern emission is detected in the CLEAN images, it is also visible in the corresponding images from `eht-imaging` which proves the ability of `eht-imaging` to recover diffuse structures. We consider the northern emission to be produced by the counterjet of the radio source, partially obscured by the accretion disk. If both the NW and N emitting regions were to be attributed to the counterjet, it would mean that the latter also presents a winding and irregular profile like its southern counterpart. However, the emission that we detect is diffuse and barely

above the noise level of the images. Because of this, we are not able to use it to constrain the core shift and we chose not to include further analysis of the counterjet in this work.

3.3.1 Feature extraction

Traditionally, VLBI images of relativistic jets have been analyzed by fitting Gaussian components to the bright features of the jet and tracking the motion of these components over time, which is known as *model fitting*. This was the best approach to analyzing jet dynamics when the resolution was not sufficient to resolve features inside the jets. However, new super-resolution imaging methods such as `eht-imaging` now allow us to resolve intra-jet features (see Janssen et al. 2021; Fuentes et al. 2023; Savolainen et al. 2023; Park et al. 2024a), making it unnecessary to approximate and oversimplify jet images using a set of Gaussians. In these cases, model fitting is not an adequate tool. Instead, case-by-case methods should be chosen, depending on the features visible in the image. Because of the winding and evolving jet structure observed in our images of 3C 84, we characterized the jet by measuring its maximum radial expansion from the core (section 3.3.5) and by tracing the profiles of the two bright limbs, from which we computed the overall jet outline and the jet launching direction (section 3.3.4).

3.3.2 Alignment

To compare features extracted from images at different epochs, it is crucial that the images are properly aligned. To align the images with respect to the jet core, we applied the following steps:

1. Locate the brightest pixel in each image. In a few cases the brightest pixel is not located in the core, but in the C3 component. In those cases the brightest pixel is replaced by a randomly chosen pixel from the core.

2. Apply a circular mask of 1 mas radius, centered in the brightest point, setting to zero every pixel outside the mask.

3. Shift each image by the amount that maximizes the cross-correlation between the masked image and the masked image of the previous epoch.

This procedure was effective for correctly aligning all epochs except those between November 3 and 6, 2021. For those epochs, a simple cross-correlation alignment was not effective because of the rapid twisting of the core region, so an additional shift, linearly spaced from 35 μas up to 246 μas, was applied after the cross-correlation shift.

3.3.3 Edge fitting

The most straightforward way to characterize the features present in the jet is to outline its edges by tracking the position of the two bright limbs. For this purpose, as shown in Figure 3.3, we considered circular sections of the image, centered around the jet core. We find the two highest emission peaks along each of these profiles and assign their positions to the two bright edges, creating two sets of points outlining each edge (blue points in Figure 3.3). In some sections, the position of one or both limbs could not be detected using this method. In such cases, the limb position was determined by interpolating between the positions in the preceding and following sections. The

overall jet outline (white points in Figure 3.3) was determined by the set of midpoints of the distance segments between each point of one limb and the other limb.

3.3.4 Jet launching direction

The local jet direction at each section of the jet is defined by the vector tangent to the jet outline curve. We determined the initial jet launching direction by averaging the tangent vectors corresponding to the segments of the jet outline within 90 μas from the core, to avoid resolution-induced biases. The standard deviation associated with the average was assigned as the error on the measured direction. The left panel in Figure 3.4 shows the angle corresponding to the jet launching direction as a function of time. The points in blue represent the measured direction from each epoch, while the orange lines indicate the average and standard deviation of a Gaussian Process (GP) regression to the data. The right panel in Figure 3.4 shows the average jet launching direction from the GP regression with the angle plotted in angular coordinates for a more intuitive representation of the directional shift. The plot shows that the orientation of the jet within the first 0.03 pc from the core undergoes several irregular oscillations. We observe a 20° oscillation from the beginning of 2011 to the end of 2012, a 40° oscillation from early 2013 to early 2017, a 10° shift until early 2018, followed by a clear 60° shift until 2021 and a constant trend in late 2021 and 2022. The overall change in the jet's orientation spans 80 degrees.

A change in the direction of the jet has been proposed by Dunn et al. (2006) to explain the presence of X-ray holes at different orientations with respect to the core on kiloparsec scales. They suggest two possible causes for jet precession: a binary black hole system would make the jet of the primary black hole undergo a regular precession (Katz 1997), while an instability in the accretion disk or a misalignment between the black hole spin and the accretion disk axis could cause the disk to warp, resulting in a stochastic jet precession (Pringle 1997). However, they estimate a precession timescale on the order of 10^7 years. Here instead we observe a drastic irregular variation in the jet's direction over a timescale of a few years. It is possible that the jet in 3C 84 undergoes precession cycles over different timescales. It should also be noted that, since monitoring of 3C 84 began, at least two epochs of jet activity, with different properties and directions, have been observed, one from the 1960s related to the ejection of C2, and the other from the 2000s, associated with the expansion of C3 and the direction oscillations reported above. Therefore, the presence of X-ray holes at different orientations could also be due to different activity phenomena associated with different jet orientations and properties.

3.3.5 Jet length and expansion speed

To measure the jet length, we first computed the longitudinal intensity profile of the jet by identifying the highest intensity value of the image along circular sections centered on the core. We defined the jet length, as projected in the image plane, as the maximum distance from the core reached by the head of the jet. The head of the jet was identified as the point where the longitudinal jet profile dropped below the average noise floor, which was computed as the mean image value, in a portion of the image not covered by the jet. The threshold value was adjusted for epochs from June 2016 to January 2017 to accommodate for a significantly higher noise floor and for epochs from May to September 2021 to take into account the darkening of portions

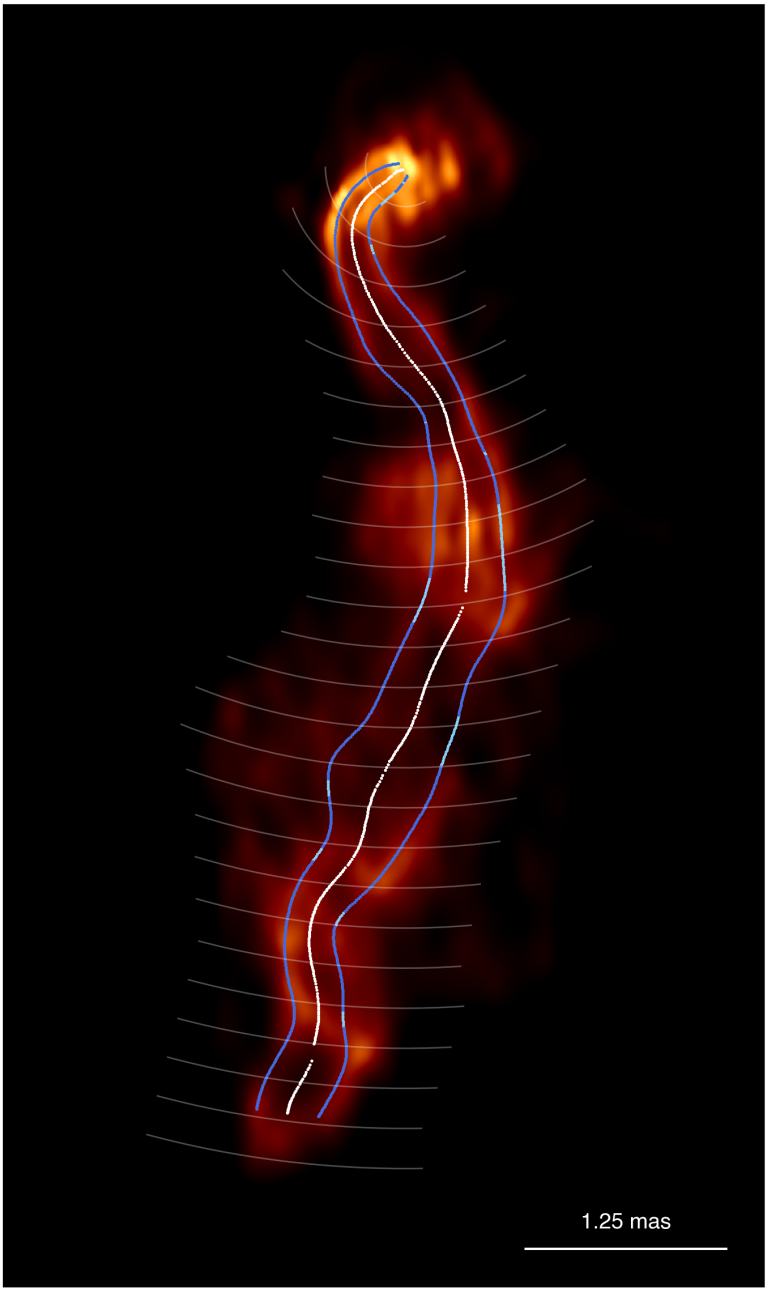

Figure 3.3: **Edge fitting in 3C 84.** Example of the fitted edges and jet direction for the jet in epoch 21-07-2021. Dark blue points represent fitted edge points, while light blue points represent interpolated edge points. White points represent the jet outline. A subset of the circular sections used to detect the edge points is shown as thin white lines.

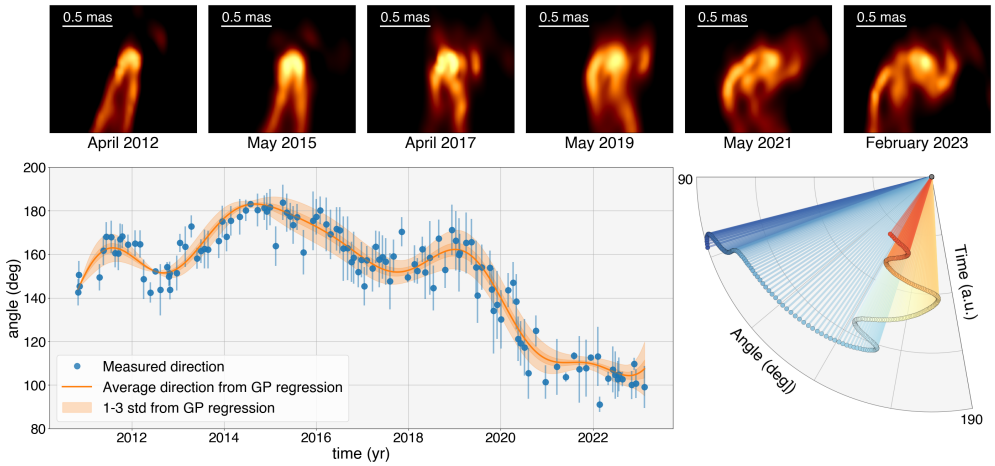

Figure 3.4: **Jet precession in 3C 84.** (Top) Details of the core region in 3C 84 for six selected epochs. (Bottom left) Angle of the jet launching direction as a function of time. The angle is measured east of north. The measured direction from each epoch is shown in blue and the average of a GP regression to the data in orange. The orange shading indicates the 1 sigma (dark shade) and 3 sigma (light shade) uncertainty from the GP regression. (Bottom right) Average direction from the GP regression, plotted in angular coordinates for a more intuitive representation of the direction change. Red corresponds to earlier epochs and blue to later epochs.

of the jet. The darkening appears in various portions of the jet in later epochs, but it affects the measurement of the jet length only in mid 2021 because a darkened portion coincides with the jet head, as shown, for example, in the lower panel of Figure 3.2 or in Figure 3.3. The uncertainty associated with the jet length measurements was taken as the pixel size used in the imaging process (30 μas), under the assumption that the limited resolution is the main source of uncertainty for this measurement.

Figure 3.5 shows the jet length as a function of time. As visible in the online movie, three different trends are evident. Until the end of 2012 the jet length increases linearly, from the beginning of 2013 to the beginning of 2017 the increase occurs at a lower rate, and finally, from 2017, the expansion occurs at a higher speed than the initial one. A possible explanation for the speed change, is that the jet propagates across a medium with different densities, possibly shaped by past activity of the jet. For each of these three expansion regimes, we performed a linear fit to the jet length to compute the expansion velocity. The residuals of the linear fits are shown in the lower panel of Figure 3.5 and do not show significant trends, meaning that the expansion in each regime was indeed occurring at a constant speed. From the measured velocities projected in the image plane, we computed the true de-projected velocities of the jet front, accounting for special relativistic effects and assuming an inclination of $\theta = 18°$ (Tavecchio and Ghisellini 2014).

Table 3.2 reports the values of the apparent speed in the image plane and the de-projected physical speed. The reduced χ^2 values of the linear fits are also reported in Table 3.2 and show that the uncertainties are properly estimated for the first two regimes, while they might be slightly underestimated in the last case. This might be

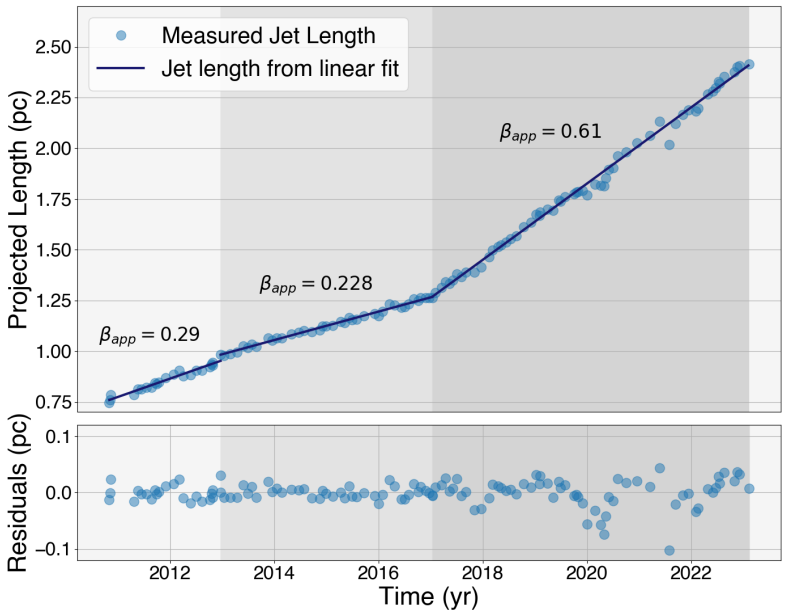

Figure 3.5: **Jet length in 3C 84.** (Top) Measured jet length as a function of time (points) and piece-wise linear fit (continuous line). Uncertainties on the measurements are not shown because the error bars are smaller than the size of the points. (Bottom) Residuals of the linear fits.

.

due to some portions of the jet being obscured in the latest epochs, which causes a higher uncertainty in the detection of the jet head. Previous estimates of the speed of the 3C 84 jet head, in our considered time period, measured an average apparent speed of 0.27 ± 0.02 c between 2007 and 2013 (Hiura et al. 2018), and an average apparent speed of 0.33 between 2003 and 2020 (Kino et al. 2021). Weaver et al. (2022) measured the apparent speed of the components in 3C 84 in the period from 2010 to 2019, through model fitting. The speed of the components is not necessarily the same speed of the jet expansion, but it can be useful to compare the two. In the period between 2010 and 2012, the apparent speed of components C2, C3, and C6 are 0.27 c, 0.21 c, and 0.34 c, which match our estimate of the expansion speed of 0.29 c. In the period between 2013 and 2015, Weaver et al. (2022) reported a speed of 0.84 c for component C9, matching the observation of the hot spot frustration, which moves at a higher speed than the jet expansion. Finally, from 2016 to 2019, they report component C10 moving at a speed of 1.38 c, indicating a drastic increase in velocity, which matches our measured expansion speed increase. Regarding the jet expansion after 2017, Kam et al. (2024) report measurements of the velocity of four subcomponents propagating from C3 after the jet breaks through the inflated bubble. In the time period between 2017 and 2022, they measure apparent speeds between 0.46 c and 1.35 c, indicating that short-lived components inside the jet may travel faster than the jet front itself.

time period	apparent speed (μas/day)	apparent speed (c)	physical speed (c)	χ^2
Nov 2010 - Dec 2012	0.71 ± 0.03	0.29 ± 0.01	0.50 ± 0.02	1.7
Dec 2012 - Jan 2017	0.549 ± 0.009	0.228 ± 0.004	0.434 ± 0.007	0.81
Jan 2017 - Feb 2023	1.47 ± 0.02	0.61 ± 0.01	0.69 ± 0.01	7.0

Table 3.2: **Jet front expansion velocity of 3C 84.** Velocities computed from the linear fit for each of the three expansion regimes. We report both apparent velocities projected in the image plane and the corresponding physical velocities, assuming an inclination angle of $\theta = 18°$.

3.4 Conclusions

In this chapter we present and discuss the parsec-scale structure of 3C 84. We show images from 121 VLBA observations at 43 GHz from late 2010 to early 2023, and we performed a quantitative kinematic analysis of the jet's expansion and precession. Studying the restarted jet in 3C 84 is important to understanding how the ISM interacts with the jet, affecting its direction, morphology, and propagation speed, and how, in return, the ISM is affected by the irregular jet activity.

Thanks to the super-resolution enabled by the `eht-imaging` imaging method, our images resolve the internal structure of the jet, its profile and edges, and different bright components moving inside the jet. The images were obtained using a standard iterative procedure in `eht-imaging`, mostly relying on closure quantities to avoid bias from residual calibration errors and with minimal use of regularizers. The images were later aligned with respect to the jet core by maximizing the cross-correlation between subsequent images.

We confirm the presence of a limb-brightened structure that connects to the unresolved core at a wide opening angle (Giovannini et al. 2018). In the initial observations, a dim component (C2) was still present in the bottom-right side of the jet but disappeared completely around 2019. During the 12-year period of observations, the restarted jet grew in length, expanding toward the south. In the most recent images, starting from 2020, some portions of the far side of the jet appear darkened. The direction at which the jet emerges from the core varies across the years. We measured the jet direction within 0.03 pc of the core and observed it undergoing various irregular oscillations, spanning an overall angle of 80°, including a fast 60° orientation change from 2019 to 2021. The parsec-scale jet gradually transitioned from an FR I- to an FR II-like morphology from 2010 to 2013, while in 2017 it abruptly transitioned back to an FR I-like morphology. We measured the expansion speed of the jet head (C3 component) and observed three separate regimes of linear expansion. The first regime corresponds to the FR I – FR II transition, the second to the inflation of the jet head in the FR II state, and the third to the expansion following the FR II – FR I transition. We measure apparent speeds of 0.29±0.01 c, 0.228±0.004 c, and 0.61±0.01 c, respectively. We confirm previous observations of a hot spot frustration during the epochs preceding the jet's abrupt FR II – FR I transition (Kino et al. 2021). The hot spot is observed following a counterclockwise trajectory around the lobe of the jet head. However, contrary to prior observations, we distinguish the frustrated hot spot from other components present before 2016 and after 2017. In some of our images,

especially from 2017, we detect a diffuse emission north of the core, which we interpret as being produced by the counterjet that is partially obscured by the accretion disk. However, the emission is barely above the noise level of the images, so we did not rely on it for the analysis of the jet features.

Overall, our results indicate that the jet is propagating in an irregular ISM that is characterized by the presence of clumps of denser material, which affects the jet's speed, direction, and morphology. The presence of localized absorbing gas in front of the jet may also be an explanation for its local darkening, which alternatively may be due to an increase in the viewing angle or a change in emissivity. The observed evolution of the parsec-scale jet suggests that the presence of radio lobes may be a temporary stage in the evolution of a jet, caused by density differences in the ambient medium. In 3C 84 this seems to be confirmed by the presence of ancient lobes at multiple scales. The limb-brightening suggests a possible spine-sheath structure, where the inner part of the jet moves at a higher speed than the outer part, and could explain the bright gamma-ray emission observed by Abdo et al. (2009). Limb-brightening could also be caused by a higher number of emitting electrons, which are accelerated in the interaction between the jet and the ISM (Stawarz and Ostrowski 2002). Park et al. (2024a) also suggest that a higher emissivity in the jet boundary layer, due to an interaction with a dense medium, may be at the origin of the observed limb-brightening in NGC315. Possible reasons for the irregular jet precession include an instability in the disk or a misalignment between the angular momentum of the accretion disk and the spin of the black hole (Dunn et al. 2006; Pringle 1997), which may cause a warping of the disk and a stochastic variation in the jet's direction.

Our results show that by using innovative super-resolving imaging methods, it is possible to resolve complex features in the jet structure, which were previously accessible only at higher observing frequencies or with significantly longer baselines. This marks a change in how jet features can be analyzed, shifting from the fitting of simple Gaussian components to more specific analyses adapted to the jet morphology. For example, similar to Park et al. (2024a), our `eht-imaging` images clearly highlight the limb-brightening structure, which allowed us to trace the jet edges and precisely track the changes in the jet direction over time. Future imaging of jet sources with new imaging methods could uncover additional cases of limb-brightening that were previously undetectable due to insufficient resolution.

4

The `kine` imaging method

Adapted from

Turbulent flow in the relativistic jet of 3C 345 from neural video reconstructions of interferometric data

M. Foschi[†], B. Zhao[†], A. Fuentes[†], et al.

Under review, (2025)

[†]These authors contributed equally to this work

Abstract

The ejection of collimated and highly relativistic jets of plasma from the vicinity of supermassive black holes to intergalactic scales is an intrinsically dynamic process shaped by large-scale magnetic fields. A proper characterization of the turbulent, magnetized flow of plasma is fundamental for understanding how relativistic jets form and evolve through the dense intergalactic medium. Until now, the study of jet dynamics has been limited to tracking the motion of unresolved components recovered from radio-interferometric observations, usually interpreted as shocks traveling downstream in the jet. Here we present a comprehensive pixel-by-pixel kinematic analysis of the parsec scale jet in the 3C 345 blazar from 27 years of data. This is enabled by `kine`, our newly developed video reconstruction algorithm based on a neural representation, which is able to process simultaneously all observations available, while learning and leveraging the spatio-temporal correlations present in the data in full polarization. Our video reconstruction reveals, in great detail, a highly variable jet that exhibits non-periodic changes in its launching direction, in contrast to what is suggested by previous work. The polarization field structure indicates the presence of an evolving helical magnetic field threading the jet. Additionally, the continuity and high resolution of our video allows us to characterize the complete two-dimensional velocity field

associated with the jet plasma evolution, in contrast to previous kinematic studies, which could only measure the pattern speed of broad components. We find that the instantaneous speed and degree of polarization of the traveling bright compact features are of the same order of those of the bulk plasma, indicating that these over-densities cannot be produced by strong shocks, but most probably they correspond to regions with increased emissivity in a turbulent flow. Here we present a comprehensive pixel-by-pixel kinematic analysis of the parsec scale jet in the 3C 345 blazar from 27 years of data. This is enabled by `kine`, our newly developed video reconstruction algorithm based on a neural representation, which is able to process simultaneously all observations available, while learning and leveraging the spatio-temporal correlations existent in the data in full polarization. Our video reconstruction reveals, in great detail, a highly variable jet that exhibits non-periodic changes in its launching direction, in contrast to what is suggested by previous work. The polarization field structure indicates the presence of an evolving helical magnetic field threading the jet. Additionally, the continuity and high resolution of our video allows us to characterize the complete two-dimensional velocity field associated with the jet plasma evolution, in contrast to previous kinematic studies which could only measure the pattern speed of broad components. We find that the instantaneous speed and degree of polarization of the traveling bright compact features are of the same order of those of the bulk plasma, indicating that these over-densities are just regions with increased emissivity in a turbulent flow, rather than traveling shocks, as assumed until now. The algorithm and methodology described in this work can be applied to entire jet monitoring programs, providing a complete description of the kinematics of hundreds of sources and possibly leading to a reinterpretation of established models. The algorithm and methodology described in this work can be applied to entire jet monitoring programs, providing a complete description of the kinematics of hundreds of sources and possibly leading to a reinterpretation of established models. Our algorithm is, in addition, capable of reconstructing videos of variable sources from single-epoch observations, as in the case of horizon-scale observations of the black hole Sgr A* by the Event Horizon Telescope.

4.1 Introduction

Supermassive black holes at the center of active galactic nuclei drive the ejection of jets (Blandford and Znajek 1977; Blandford and Payne 1982), highly relativistic and collimated streams of plasma that propagate through the intergalactic medium up to mega-parsecs away from the SMBH (Oei et al. 2024). Early very long baseline interferometric observations of jets reported on the presence of broad, unresolved features, often called components, moving downstream at apparent superluminal speeds due to relativistic effects (e.g., Cohen et al. 1977; Vermeulen and Cohen 1994; Gómez et al. 2000; Jorstad et al. 2005). These bright components are typically the only features apart from the jet core with enough emissivity to be captured in VLBI observations, and the trajectories they follow as they propagate downstream highlight the underlying jet structure (e.g., Steffen et al. 1995). Superluminal components have historically been associated with traveling shock waves, which would compress the magnetized plasma leading to an increase in the emissivity and fractional polarization (Marscher and Gear 1985). Alternative models suggest that plasma instabilities and geometrical effects could also potentially explain the observed components commonly interpreted as shocks (e.g., Raiteri et al. 2017; Fuentes et al. 2023).

Until now, the dynamics of the jet plasma have been studied by tracking the motion of discrete components across multiple observations. Components are identified by modeling the jet structure with Gaussians directly fitted to the data, a process known as model fitting (e.g., Lister et al. 2021; Weaver et al. 2022), or by identifying discrete features in VLBI jet images with a wavelet analysis (Mertens and Lobanov 2015). This approach has been widely used due to the limited resolution and dynamic range of VLBI images reconstructed with traditional imaging methods (e.g., CLEAN; Högbom 1974), as well as irregular time sampling of observations, which complicates continuous tracking of features.

We developed a new dynamic imaging algorithm for VLBI data, named `kine`, which overcomes these limitations and recovers a full polarimetric video of a variable radio source by integrating information from observations at different times. The algorithm, originally intended for dynamic imaging of intra-day variable sources like Sgr A*, relies on a neural representation of the source's video and recovers its frames simultaneously from all observations, by learning and leveraging the spatio-temporal correlations present in the data. The method provides a continuous representation of the brightness density distribution, which can be evaluated at any chosen instant in time. Additionally, thanks to the simultaneous dynamic imaging of multiple datasets, each frame of the video is reconstructed using more information than in traditional individual static imaging. This improves resolution and dynamic range, which is especially important in cases where some observations have more limited coverage, missing radio telescopes, or lower signal-to-noise ratio (SNR). The continuity and high resolution of our video reconstructions allow us to apply post-processing video techniques, such as the optical flow (OF) estimation (Horn and Schunck 1981), and recover the pixel-by-pixel instantaneous velocity field associated to the plasma.

In this work, we applied `kine` to multi-epoch observations of radio quasar 3C 345, obtaining a high-resolution polarimetric video of the relativistic jet. At a redshift of $z = 0.593$, this source has been monitored at centimeter wavelengths since the late 1970s (Unwin et al. 1983; Biretta et al. 1986). The relativistic jet in 3C 345 is highly

dynamic, and characterized by the ejection of numerous superluminal components (Steffen et al. 1995; Lister et al. 2019; Lister et al. 2021; Pötzl et al. 2021; Weaver et al. 2022; Röder et al. 2024), which have been interpreted as shocks moving across the plasma fluid. Past observations have shown that the jet structure presents irregularities and a persistent bend at 4 milliarcseconds (mas) from the core, which may be caused by interactions with the surrounding medium. In order to explain the apparent superluminal speeds measured along bent paths and the periodic changes in the inner jet position angle, Steffen et al. (1995) proposed a helical model for the jet structure, while Lobanov and Roland (2005) proposed a binary system of equal mass supermassive black holes inducing precession of the main accretion disk. Thanks to our newly developed imaging method, paired with motion analysis tools, we are able to recover a high resolution polarimetric video of 3C 345 and measure the 2D velocity field in the jet plasma, shedding light on the nature of jet dynamics and superluminal components.

4.2 Methods

`kine` is a static or dynamic image reconstruction algorithm for interferometric data. It is a forward imaging method, based on a neural field representation of the source's brightness distribution. In this section we describe our methodology, starting with a general overview of VLBI measurements, followed by an explanation of the `kine` imaging algorithm, which is summarized by the diagram in Figure 4.1.

4.2.1 VLBI measurements

In radio interferometric observations, each antenna in the array records a signal proportional to the amplitude of the received electromagnetic flux density. By the Van Cittert–Zernike theorem, the time averaged correlation product of the signals recorded by any pair of antennas (called visibility) is the Fourier transform of the flux density spatial distribution on the sky plane, evaluated at a frequency proportional to the antennas' distance. Specifically, for a pair of antennas A and B, with projected baseline vector $\bar{b} = (b_x, b_y)$, observing at wavelength λ, the ideal complex visibility $V_{AB}(u, v, t)$ is related to the flux density distribution $\mathcal{I}(x, y, t)$ by:

$$V_{AB}^{(\mathcal{I})}(u, v, t) = \int\int e^{-2\pi i(ux+vy)}\, \mathcal{I}(x, y, t) dx dy \ , \tag{4.1}$$

where $(u, v) = (\frac{b_x}{\lambda}, \frac{b_x}{\lambda})$ are the x and y spatial frequencies (refer to Thompson et al. 2017 for a detailed description of the theory of radio interferometric observations and imaging). The above equation is written for the Stokes parameter $\mathcal{I}$, but it holds for all Stokes visibilities $(V_{AB}^{(\mathcal{I})}, V_{AB}^{(\mathcal{Q})}, V_{AB}^{(\mathcal{U})}, V_{AB}^{(\mathcal{V})})$.

In practice, different sources of noise corrupt the measurement of visibilities. They can be classified as baseline-dependent errors and site-dependent errors, so the measured visibilities can be expressed as:

$$V'_{AB} = G_A G_B e^{i(\phi_A - \phi_B)} \left(V_{AB} + \epsilon_{AB}\right) \ , \tag{4.2}$$

where $G_{A,B} e^{i\phi_{A,B}}$ are the site-dependent errors in amplitude and phase, referred to as complex gains, while ϵ_{AB} is the thermal noise, which is Gaussian distributed with baseline-dependent standard deviation (Thompson et al. 2017). The most problematic

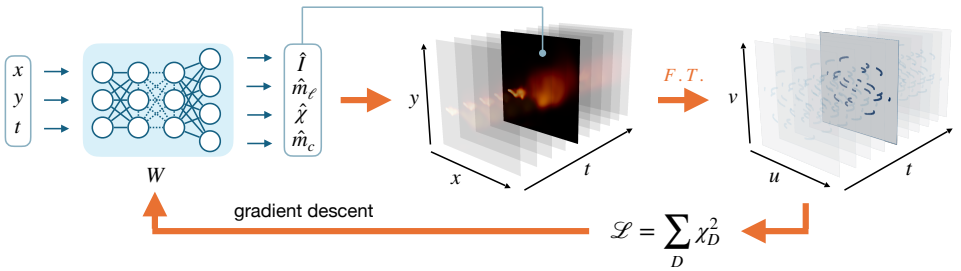

Figure 4.1: **The `kine` algorithm.** The algorithm is composed by a forward model, a data constraint and an optimization method. The forward model consists of a coordinate-based MLP which takes as input space and time coordinates and outputs the values of total intensity, fractional linear polarization, polarization angle and fractional circular polarization, at the location of the input coordinates. A spatial Fourier transform is applied to the image or video estimated by the network to compute the visibilities. The difference between the estimated and observed visibilities is computed through χ^2s of the chosen data products. The loss function is built as the sum over all the data products' χ^2s with a regularization term enforcing consistency with the observed total flux. The parameters of the MLP are optimized to minimize the loss function.

source of error is the complex gains, as they might not be estimated correctly from the a-priori calibration, while thermal noise can be fully characterized and incorporated in the loss function through the visibilities' uncertainties σ_{AB}.

Therefore, even though complex visibilities are the fundamental data product resulting from VLBI observations, the imaging process can use different data products which are constructed to be independent from site-dependent amplitude or phase corruptions. `kine` supports the following data products: complex visibilities V_{AB}, amplitudes of the visibilities $|V_{AB}|$, closure phases $\Phi_{ABC} := \arg(V_{AB}V_{BC}V_{CA})$, closure amplitudes $A_{ABCD} := (V_{AB}V_{CD})/(V_{AC}V_{BD})$, and the logarithm of closure amplitudes. Closure phases are unaffected by station-based phase errors, because they cancel out in the triple visibility product (Jennison 1958; Rogers et al. 1974). For an an array of N_a antennas, a complete and non redundant set of closure phases still contains a fraction $(N_a - 2)/N_a$ of the total information in the visibilities, though the information about the absolute location of the image is lost (Twiss et al. 1960). Closure amplitudes are unaffected by station based amplitude errors, because they cancel out in the double visibility ratio (Twiss et al. 1960). A complete and non redundant set of closure amplitudes still contains a fraction $(N_a - 3)/(N_a - 1)$ of the total visibility information, though the information about the image total flux is lost. Logarithmic closure amplitude contain the same information as closure amplitudes, but have the useful property of having Gaussian-distributed uncertainties in the high SNR limit, which justifies the use of a χ^2 loss for this data product.

4.2.2 Model

`kine` is a forward imaging algorithm for interferometric data, that models the polarized flux density distribution of a continuous image or video with the weights $\mathcal{W}$ of a neural field (Xie et al. 2022). The neural field consists in a coordinate-based neural network, specifically a Multi-Layer Perceptron (MLP) (Popescu et al. 2009), that takes as input both space (right ascension x and declination y) and time (t) coordinates and outputs

an estimated total intensity flux density $\hat{I}$, fractional linear polarization $\hat{m}_\ell$, electric vector position angle $\hat{\chi}$, and fractional circular polarization $\hat{m}_c$ at that location in time and on the sky plane:

$$(\hat{\mathcal{I}}, \hat{m}_\ell, \hat{\chi}, \hat{m}_c)_{\mathcal{W}}(x, y, t) = \mathrm{MLP}_{\mathcal{W}}(x, y, t) \; . \tag{4.3}$$

The hat symbol indicates the estimated quantities as opposed to the true or observed quantities of the source. The Stokes parameters are then computed from the polarization quantities $(\hat{\mathcal{I}}, \hat{m}_\ell, \hat{\chi}, \hat{m}_c)$ using the transformation:

$$\hat{\mathcal{I}} = \hat{\mathcal{I}} \tag{4.4}$$

$$\hat{\mathcal{Q}} = \hat{\mathcal{I}} \cdot \hat{m}_\ell \cdot \cos(2\hat{\chi}) \tag{4.5}$$

$$\hat{\mathcal{U}} = \hat{\mathcal{I}} \cdot \hat{m}_\ell \cdot \sin(2\hat{\chi}) \tag{4.6}$$

$$\hat{\mathcal{V}} = \hat{\mathcal{I}} \cdot \hat{m}_c \; . \tag{4.7}$$

We compute VLBI measurements from the predicted video with a fully differentiable forward model, then optimize the weights of the MLP to minimize a χ^2 data fit loss between the predicted and observed measurements of the source.

Estimated complex visibilities $\hat{V}$, and associated data products, are computed by evaluating the 2D Fourier transform of the network-estimated flux density $\hat{\mathcal{I}}$ for the (u, v, t) points corresponding to the observation tracks. In practice, $\hat{\mathcal{I}}$ is estimated for a discrete set of $\{(x, y, t)\}$ points, so the visibilities are obtained by a discrete Fourier transform $\mathbf{F}$ implemented through a matrix multiplication. For each observation time t_j and u-v point i, the estimated visibilities are:

$$\hat{V}_{\mathcal{W}}((u, v)_i, t_j) = \sum_k \mathbf{F}_{ik} \hat{\mathcal{I}}_{\mathcal{W}}((x, y)_k, t_j) \; , \tag{4.8}$$

where k runs over all spatial points at which the video is estimated and both the estimated total intensity and the estimated visibilities are dependent on the network parameters $\mathcal{W}$.

MLP networks have been shown to suffer from "spectral bias" (Rahaman et al. 2019) causing the network to favor representing only low spatial frequencies when trained with gradient descent. This tendency is advantageous for VLBI reconstructions as it acts as an implicit regularizer, encouraging the network to learn smooth reconstructions without spurious high-frequency features in space or time. In addition, the continuous neural representation allows us to sample the video reconstruction at times where observations are not available, by leveraging the correlations from neighboring frames. The implicit regularization provided by MLP networks has already been applied successfully to other ill-posed inverse problems in astrophysical imaging, where many possible solution fit a given set of measurements (Levis et al. 2024; Zhao et al. 2024; Zhong et al. 2021).

4.2.3 Optimization

The `kine` algorithm optimizes the network's parameters so to minimize the difference $|D - \hat{D}_{\mathcal{W}}|$ between observed and estimated data. We assume that the noise residuals in the data products are Gaussian, which is a valid assumption for visibility and amplitudes, and for closure quantities in a high SNR regime. Therefore, the natural

loss function $\mathcal{L}$ for the optimization process is the sum of the χ^2s of the chosen data products:

$$\mathcal{L} = \sum_D \chi^2_D = \sum_D \left[\frac{1}{N_t} \sum_{j=1}^{N_t} \frac{1}{k_D N_{D,j}} \sum_i^{N_{D,j}} \frac{\left(D_{ij} - \hat{D}_{\mathcal{W}ij} \right)^2}{\sigma^2_{D,ij}} \right] , \tag{4.9}$$

where D indicates a generic data product, the index j runs over all the observed times, from 1 to the total number of observed times N_t, the index i runs over all data corresponding to the observed time t_j, from 1 to the total number of data $N_{D,j}$, and k_D is a normalization factor that takes into account the data product's degrees of freedom.

The optimization process consists in iterating the following steps:

1. Estimate the video, with the current network parameters $\mathcal{W}_k$, on a coordinate grid with x and y regularly spaced and t spaced accordingly to the observations times.

2. Compute the loss function from the estimated video and the observed data.

3. Update the MLP parameters to $\mathcal{W}_{k+1}$, by applying a gradient descent optimization step of the loss function. More specifically, we use the `Adamax` optimizer (Kingma and Ba 2014).

In this work, convergence is considered to be reached when the moving average of the loss function over a window of 100 iterations does not decrease by more than 1% for at least 1000 iterations, provided that all data products in the loss function have reached a value of ~ 1. This convergence was achieved after approximately 2×10^4 iterations, a value that we finally set as the total number of iterations.

4.2.4 Calibration-free imaging details

When imaging a dataset with no gain corruption, the best data product to use are complex visibilities, since they contain the full amount of available information. In such a case, the imaging procedure in `kine` consist of running the optimization algorithm estimating all Stokes parameters simultaneously. However, if the dataset contains amplitude and/or phase gains, it is better to perform imaging using closure quantities, which are error independent. When imaging with closure or log-closure amplitudes, the information about the total flux of the frames is lost. In this case, a total flux regularizer term is added to the loss function, constraining the total flux in each frame to a value either provided by the user or computed from the instantaneous visibility of the shortest baseline. When imaging with closure phases, the information about the absolute location of the source in the frame is lost. Because of this, the reconstructed video may show the source drifting smoothly across the field of view. We solve this issue post-imaging, by re-aligning the frames to a chosen feature, such as a jet core. The alignment process consists of shifting each frame by the amount that maximizes its cross-correlation with the reference frame.

The current version of `kine` doesn't model station gains simultaneously with the imaging. For the total intensity reconstruction, this problem is solved by relying on closure quantities. The closure quantities for the $\mathcal{Q}$ and $\mathcal{U}$ signals, however, are

gain independent only in the absence of polarization leakage (Roelofs et al. 2023), but, even in that case, the low SNR of the $\mathcal{Q}$ and $\mathcal{U}$ data, combined with the loss of information due to the use of closure quantities, may result in insufficient data constraints. Therefore, we found that the most effective procedure to produce a full polarimetric video, in the presence of gain corruption, is to apply the following steps:

1. Image the data in total intensity only, using the appropriate closure quantities.
2. Self-calibrate the data to the total intensity video reconstruction.
3. Image the self-calibrated data in full Stokes, using complex visibilities.

The `kine` imaging method was born out of the need to perform dynamic imaging of a sparse-coverage observation of a variable source or multiple observations of the same evolving source. However, the algorithm can reduce to static imaging by removing the time dimension from the input coordinates.

4.2.5 Architecture and training

The algorithm is implemented in a python code, using the machine learning framework `JAX` (Bradbury et al. 2018). We use an MLP of depth $D = 4$ and width $W = 256$. A width of $W = 256$ is generally considered sufficient for the network to fall in the infinitely wide limit. The depth of the network determines the expressivity of the network. The higher the number of layers, the more complex the functions that the network can approximate. However, a high number of layers increases the number of tunable parameters, making optimization difficult and slower. Therefore D should be set by trial and error to the minimum value that guarantees the expressivity required for the task. The network includes a skip connection (He et al. 2016), which adds the output of the first layer to the input of the last layer. The output of each node, in each layer except the last one, is normalized with respect to the total set of input coordinates, so to have zero average and unit standard deviation, (Ioffe 2015). The activation function used in the hidden layers is the Gaussian Error Linear Unit (GELU) (Hendrycks and Gimpel 2016), while the activation function of the output layer is a SoftPlus function for $\hat{\mathcal{I}}$ and a sigmoid for $\hat{m}_\ell$, $\hat{m}_c$, and $\hat{\chi}$. The code makes use of the `eht-imaging` library (Chael et al. 2018) to deal with interferometric data and standard interferometry functions. During training, the input points have space coordinates sampled on a regular grid, while time coordinates follow the gridding of the observation's timestamps. In the imaging of 3C 345 we used 200x200 spatial input points spaced at 75 μas.

The runtime of the code depends on the dataset, specifically on the number of visibilities, the number of epochs, the ratio between field of view and beam size and the data products used. Training the network to image 116 epochs of 3C 345 data from the VLBA required ∼1.3 hours on four NVIDIA A100 GPUs, while imaging 4 days of M 87* observations by the EHT required ∼20 minutes on one NVIDIA A100. The RAM memory required to run the code is set by the size of the matrices used to compute the Fourier transform and can be determined with the formula:

$$\mathrm{RAM} = N_{\mathrm{vis}} \cdot N_{\mathrm{epochs}} \cdot N_{\mathrm{px}}^2 \cdot \sum_D f_D \ , \tag{4.10}$$

where N_{vis} is the number of visibilities per epoch, N_{epochs} is the number of epochs, N_{px} is the number of spatial input points, and f_D is a multiplying factor that depends

on the data products used, being 1 for visibilities and amplitudes, 3 for closure phases, 4 for closure amplitudes and log closure amplitudes.

4.2.6 Initialization

When the u-v coverage of the VLBI data is sufficiently dense, such as in the case of VLBA observations, `kine` is able to reconstruct the image or video starting from a random initialization of the network weights, specifically we use the He uniform initialization (He et al. 2015), and a zero initialization of the biases. However, when imaging data with sparse coverages, such a the EHT array ones using closure quantities, convergence to a correct reconstruction requires initializing the image or each frame of the video to a simple shape such a disk or a Gaussian which constrains the total flux and helps retain the majority of the flux in a specific area of the image (such as the center), hence preventing the appearance of tiling artifacts (i. e. artifacts consisting of adjacent repetitions of the main structure in the image). Even in cases in which the initialization is not necessary for a correct converge of the network, it may be useful to apply it, to constrain the image to a specific position in the frame. We have found that, while it is important that the initialization image has the correct total flux and covers approximately the area where most of the source's flux should be located, the specific shape of the initialization image does not affect the resulting image or video obtained when convergence is reached.

The initialization is carried out by training the network directly on the initialization image $\mathcal{I}_{\rm init}(x, y)$, with a pixel-to-pixel distance loss function:

$$\mathcal{L}_{\rm init} = \sum_{i,j}, (\mathcal{I}_{\rm init}(x_i, y_i) - \hat{\mathcal{I}}_{\mathcal{W}}(x_i, y_i, t_j))^2 \ , \tag{4.11}$$

without the computation of the image Fourier transform or other data products. After the weights and biases of the network are optimized to output the initialization image, the training on the data products begins.

4.2.7 Optical flow

Traditional imaging methods, such as CLEAN, characterize the dynamics of the jet by tracking the motion of Gaussian components obtained with model fitting (e.g. Lister et al. 2021). With `kine`, we can move beyond this approach and estimate the continuous velocity field of the superluminal jet via video post-processing techniques. In our work we present velocity estimates using the Horn-Schrunk optical flow algorithm (Horn and Schunck 1981), though `kine` provides reconstructions from which this field could be easily estimated with other algorithms. The Horn-Schrunk algorithm estimates the optical flow field, which measures the pixel displacement between neighboring pairs of frames. The pixel displacement, combined with the length of the time interval between frames, gives an estimate of the jet's velocity field at each pixel location of the video. In addition, we integrate the velocity field to obtain the trajectories followed by volume elements in the jet and track significant jet components. Optical flow recovery methods are widely used in many research fields involving video observations, including astrophysics (e.g. Colaninno and Vourlidas 2006).

4.3 Results

The MOJAVE monitoring program (Lister et al. 2018) hosts observational data of 3C 345 since 1995 and continues to observe it regularly with the Very Long Baseline Array (VLBA) up to the present day. We used `kine` (details in section 4.2) to recover the time-resolved polarimetric emission of 3C 345 from the MOJAVE database, consisting of 116 observations at 15 GHz and spanning 27 years, from January 1995 to May 2022. Thanks to `kine`'s ability to integrate information from frames at different times, the simultaneous imaging of all epochs yields results with a dynamic range of $\sim 8.6 \times 10^5$ and a resolution of ~ 80 micro-arcseconds (μas), representing an improvement of 5-6$\times$ in resolution and more than 200$\times$ in dynamic range compared to previous studies.

In the top two rows of Figure 4.2 we show a representative sample of 12 images out of the 116 video frames reconstructed, while the complete set of frames can be found in Appendix B in Figures 4.11 and 4.12, and in the video available at `https://github.com/mariannafoschi/phdthesis/blob/b847abdb9984f636adf3675cc85df8047ac30544/chapter3-video1-kine3c345.gif`. The images present two distinct regions, a brighter inner jet, propagating westwards, and a more extended plume of diffuse emission bending towards the north, for a total extension of $\sim$13 mas. For three selected frames, a zoom of the compact inner jet is shown in the lower rows of Figure 4.2, with polarization ticks in the left panels, and the plasma velocity vector field (optical flow) in the right panels. The recovered video displays a highly resolved jet structure which propagates from a bright, compact region at the leftmost end. From this location, corresponding to the jet core, plasma is continuously expelled downstream with occasional and noticeable ejection of discrete bright features.

Within the inner few milli-arcseconds, the plasma seems to move ballistically along the path determined by the core's orientation at ejection (e.g. left features in panels b and c of Figure 4.2 and jet launching direction in Figure 4.5), then, at a distance of $\sim$2-3 mas from the core, ejected discrete features turn inward towards the jet axis (e.g. right feature in panel b of Figure 4.2 and trajectories in top right panel of Figure 4.3), before vanishing into the plume (e.g. right feature in panel c of Figure 4.2). These proper motions have been identified as well in previous works through Gaussian model fitting at 15 GHz and other frequencies e.g. Jorstad et al. 2017; Lister et al. 2021; Röder et al. 2024. The "bouncing" of initially ballistic components against the jet edges suggests that the jet is surrounded by a high-pressure interstellar medium (ISM) that constrains and collimates the jet's inner section. Overall the complex motion and dynamics that we observe suggest the presence of a plasma that is pushing through a denser medium, where the jet fluid evolves from a turbulent state in the inner section to a more steady regime in the outer section.

We investigate the dynamics and structure of the relativistic jet shown in the reconstructed video by: computing the jet velocity field with the optical flow method, analyzing the instantaneous and time-average linear polarization maps, and tracking the jet core orientation angle. Thanks to these three analysis approaches, we can draw conclusions on the nature of bright traveling components, the magnitude of the apparent flow speed in the jet, the morphology of magnetic fields in the jet, the extent and regularity of the jet axis precession, and the interactions between the jet and the ISM.

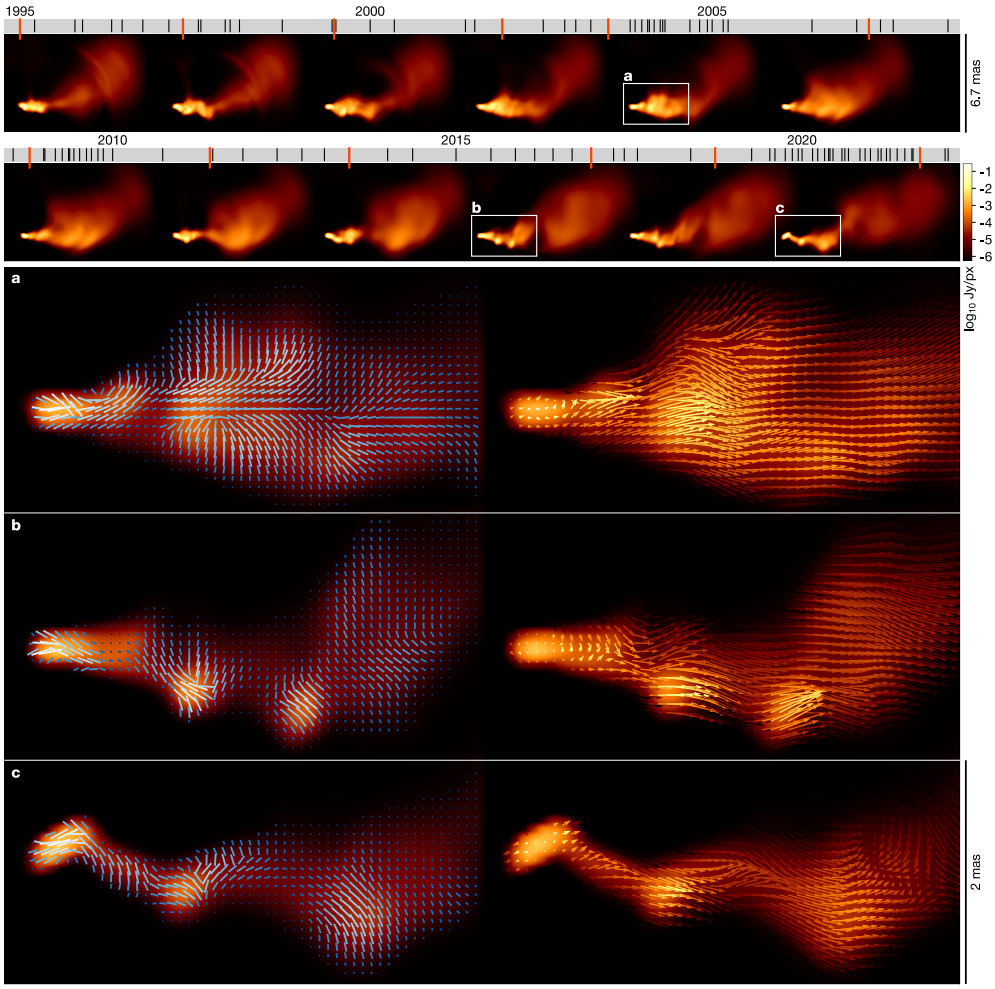

Figure 4.2: **Time-resolved relativistic jet flow in 3C 345.** (Top rows) Sample frames from the `kine` video reconstruction of 116 epochs from the MOJAVE program at 15 GHz. (Bottom rows) Zoom-in of selected frames, showing the linear polarization (left) and the optical flow (right) of the inner jet. Linear polarization is displayed as ticks whose length and color are proportional to the logarithm of the recovered polarization field, with the longest vectors corresponding to a linear polarization of 5.3×10^{-3} Jy/px. Optical flow is displayed as vectors proportional to the recovered instantaneous velocity field, with the longest vectors corresponding to a velocity of 9.0 c. The vectors' color matches the color of the underlying total intensity image. The dates corresponding to the selected frames are marked in red on the timelines.

4.3.1 Jet velocity field

In contrast to the standard kinematic analysis of multi-epoch VLBI data, which relies on fitting Gaussian components to broad features traveling along the jet, we are able to track the instantaneous, local velocity field of the jet's plasma through optical flow estimation. This is enabled by the higher angular resolution of our method's reconstructions, thanks to which we can obtain a pixel-by-pixel velocity vector field, defined over space and time. In the right panels of Figure 4.2, the velocity field is represented as vectors proportional to the pixel velocity. The cadence of the 3C 345 observations is irregular in time, and the separation between adjacent epochs can range from 6 days to more than 1 year. In contrast with other imaging methods, `kine` reconstructs a continuous flux density distribution, allowing for high-quality interpolation between observed frames. For the analysis of the fluid jet dynamics we used the optimized neural network to interpolate frames at times uniformly distributed between the first and last epoch, so that the video reconstruction is not subject to errors derived from irregular sampling or frames widely separated in time.

In the top left panel of Figure 4.3, we present the average apparent speed in the jet (a), with the corresponding standard deviation (b), indicating the amount of variability in the flow speed. We computed the average speed by masking portions of the image with total intensity below 5 times the root mean squared (rms) error of the image. In Figure 4.10 in Appendix B we show the stacking of the masks, which indicates the regions where the estimate of the average is more reliable. By integrating the OF field from a given initial position, one obtains the trajectories followed by volume elements of the jet plasma, as shown in top right panel of Figure 4.3 for trajectories starting from May 2002 (other starting dates show similar behaviors).

Lister et al. (2021) tracked the motion of Gaussian components fitted to MOJAVE observations. Integrating the OF velocity field from the center of the first detection of a component yields trajectories that coincide with those from model fitting as shown in the middle row of Figure 4.3. The trajectories of bright features tracked with OF and the trajectories of those same features obtained from model fitting are in good agreement within the inner $\sim$2 mas from the core, where the Gaussian components are well localized. Further downstream from the core, where model fitting struggles to constrain localized components in favor of diffuse ones, the two methods are in agreement on the overall trajectory, though small discrepancies are present. Given the agreement between the OF motion field and the lower resolution trajectories from prior model fitting (Lister et al. 2021), we trust the optical flow to track the apparent speed of the plasma fluid correctly. However, the OF method is not limited to tracking simple bright components like in model fitting. Indeed, thanks to the high resolution achieved by the `kine` reconstruction and to the locality of the OF vector field, the OF is able to capture not only the velocity of traveling components, but most importantly the velocity of the bulk jet flow in the absence of a passing component.

The highest value of the mean apparent speed reaches a $\beta_{\rm app} = 12 \pm 0.2$ c and are located between 1 and 3 mas from the core in the southern part of the jet, corresponding to the area towards which the majority of the components are ejected. The average speed in other areas of the jet within the first 5 mas is $[9-11] \pm 0.2$ c, while further out in the diffuse emission region it is $[5-8] \pm 0.2$ c, indicating a deceleration in the plasma flow. The speed standard deviation has no significant location dependence and is contained between $[3-6]$ c, indicating that the plasma flow is turbulent and

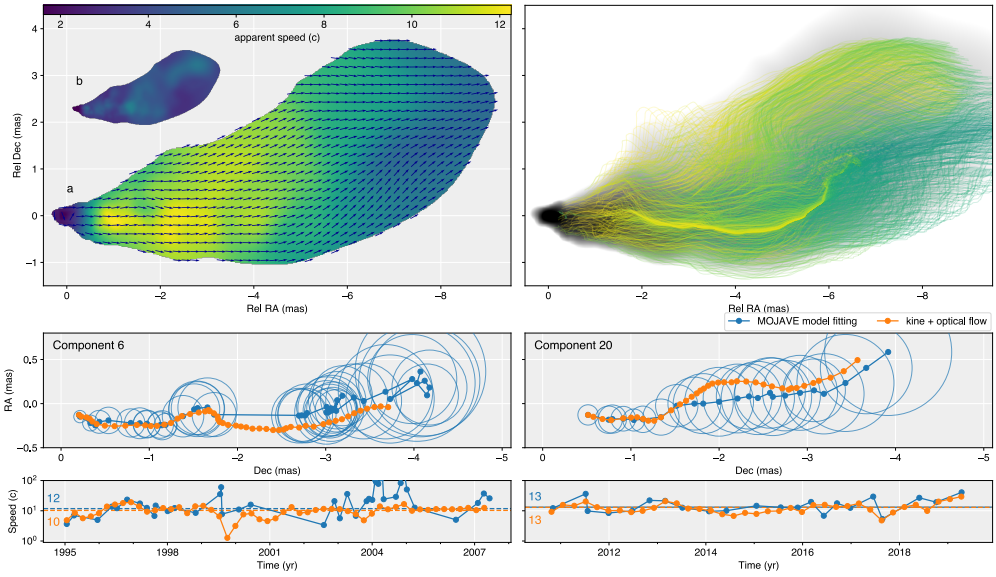

Figure 4.3: **Optical flow velocity field in the jet plasma.** (Top left) Time average (a) and standard deviation (b) of the apparent velocity, with magnitude displayed by the color map and direction by the arrows. (Top right) Trajectories obtained integrating the optical flow starting from May 2002. The trajectories represent the path followed by test particles moving according to the velocity vector field. The color of each trajectory indicates the average speed of the trajectory, with the same scale as panel a. (Middle) Comparison of component trajectories from traditional model fitting (blue) (Lister et al. 2021), and optical flow (orange), for two example components. The blue circles indicate the size of the model-fitted components. (Bottom) Component speed as a function of time. The average speed for each method and component is shown as a dotted line. The velocities reported by (Lister et al. 2021) are $\beta_{\mathrm{app,C6}} = 11.79 \pm 0.19$ c and $\beta_{\mathrm{app,C20}} = 13.16 \pm 0.17$ c respectively.

subject to relevant changes in apparent speed.

We investigate the nature of the bright components visible in the total intensity video by comparing their instantaneous speed, obtained from the OF field along the component trajectory, against the average jet speed. The lower panels in Figure 4.3 show the speed of two example components (C6 and C20) as a function of time. The time averaged component speeds are $\beta_{\mathrm{app,C6}} = 10.3 \pm 0.2$ c and $\beta_{\mathrm{app,C20}} = 13.0 \pm 0.3$ c, similar to the previously measured values of $\beta_{\mathrm{app,C6}} = 11.79 \pm 0.19$ c and $\beta_{\mathrm{app,C20}} = 13.16 \pm 0.17$ c (Lister et al. 2021). While these speeds are indeed similar, for component C6, model fitting and the optical flow yield a difference of 1 c. We attribute this to the intrinsic differences in how these values are obtained (polynomial fitting of model fitted Gaussians in the visibility domain versus feature tracking in the image domain). For the other components, the OF velocity also yields similar values, that match the velocities obtained with model fitting. The magnitude of the components' velocities is comparable to the mean bulk plasma flow speed, as the top left panel in Figure 4.3 indicates. This suggests that the bright components are not in fact shocks, but more likely localized over-densities in the plasma flow or regions with increased emissivity. We further test this hypothesis in the next section by analyzing their polarimetric

properties.

Other information that we can obtain from the optical flow field is that, around 2 mas from the core, trajectories converge towards the jet axis in a fast flowing path (evident in yellow in the top right panel of Figure 4.3), matching the motion of various components visible in the video (see e.g. panel b in Figure 4.2). We also see that the optical flow trajectories bend northwards between 5 and 8 mas, suggesting the presence of an external wind in the ISM that blows from south to north and deviates the initial westward direction of the jet.

OF aims to accurately capture the apparent motion of features in any region of the jet, matching the motion detected by human perception. However, the algorithm relies on image gradients and therefore can only track the motion of visible features, so, in the limit of a perfect laminar flow, the optical flow would report a null velocity, as no feature is observed moving, while, in the presence of a turbulent flow, the reported velocity is expected to be accurate. This explains why the velocity appears to be extremely low in the core and its proximity (right panels in Figure 4.2 and top left panel in Figure 4.3). We consider the OF to provide a reliable estimate of the jet velocity starting from $\sim$0.5 mas from the core. The non-null velocity detected in the core region is caused by residual frame-to-frame misalignment. We consider the time average of the average velocity reported for pixels contained in the core, as the uncertainty on the individual instantaneous OF velocity measurements, since we believe this to be the limiting source of uncertainty. The uncertainties that we report are of the same order of those obtained by component model fitting (Lister et al. 2021), showing that resolution is not the factor restricting the precision of the measurements, which is instead limited by the irregular evolution and motion of the jet.

4.3.2 Polarization and magnetic field

`kine` reconstructions of the data provide a video of the Stokes parameters $\mathcal{I}$, $\mathcal{Q}$, and $\mathcal{U}$ (we neglect Stokes $\mathcal{V}$ in this work), that enables us to observe the evolution of the linear polarization $P = \mathcal{Q} + i\mathcal{U}$ and the fractional liner polarization $m_\ell := |P|/\mathcal{I}$ in relation to structural changes in the jet flow. In Figure 4.2 we show the polarization field in three example frames, with the length of ticks proportional to $\log(P)$.

A transverse-parallel-transverse structure with respect to the jet axis is present in various epochs and seems to be the underlying structure in the EVPA distribution. An example of that is displayed in panel a of Figure 4.2 where the EVPAs are aligned with the direction of the jet along the jet spine, while they are perpendicular to the jet in proximity of the jet borders. Occasionally, in coincidence with the injection of some bright features, the pattern is disrupted by the traveling feature, as shown in panels b and c of the same figure. We present the time-averaged polarization field in panel a of Figure 4.4. In panel b we show the time-averaged fractional polarization m_ℓ, with associated standard deviation in panel e and mask stacking map in Figure 4.10 in Appendix B. In panels c and d of Figure 4.4 we highlight the changes of the EVPA orientation and the fractional polarization across transversal sections of the jet. From 1995 to 2000, most of the 3C 345 observations lack polarimetric data. Therefore, when time averaging, we considered only frames corresponding to observations from June 2001 or later, for which polarimetric data are always present.

We observe that the EVPA distribution recovered in panel a of Figure 4.2, which indicates a magnetic field aligned with the jet at the edges and perpendicular at the

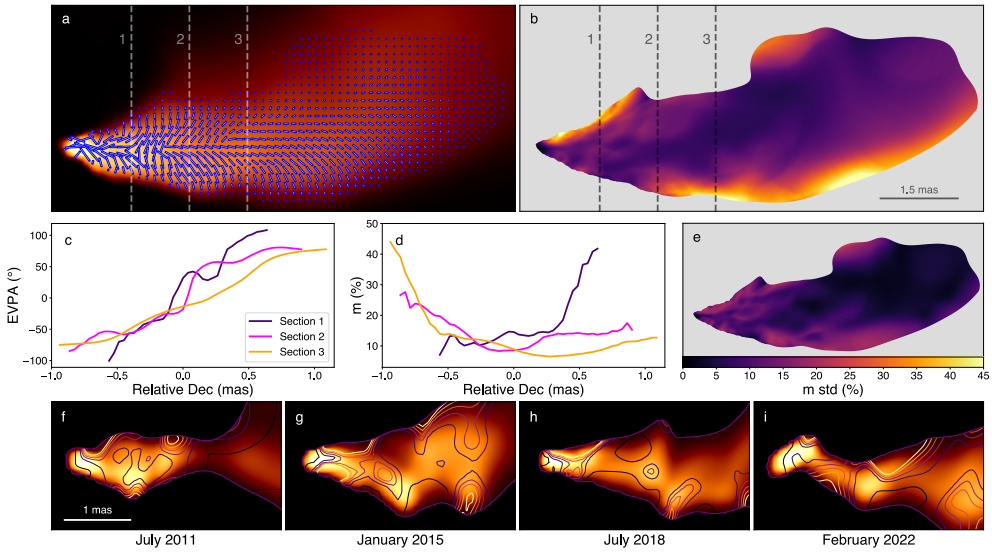

Figure 4.4: **Polarization field structure in the 3C 345 jet.** (a) Time average polarization, shown as ticks proportional to the logarithm of the linear polarization intensity and directed like the electric vector position angle. The background image shows the time average total intensity in color scale. (b) Time average fractional polarization, in the same scale as panel e. (c) EVPA values along the sections indicated in panel a. (d) Fractional polarization along the sections indicated in panel b. (e) Standard deviation of the fractional polarization. (f-g-h-i) Instantaneous frames with total intensity displayed as color map and fractional polarization as contour lines. Contours are shown for $m = 1, 10, 30, 50$, with darker contours corresponding to lower m.

center, persists also in the time average polarization map, with a sharp transition in the inner 2 mas (sections 1 and 2 in panel c of Figure 4.4) and a smoother transition further from the core (section 3). We also observe an increase of the fractional polarization at the edges of the jet (panel d, Figure 4.4), especially the southern one, in agreement with what reported, in lower resolution, by Pushkarev et al. (2023).

The fractional polarization maps are another tool to probe the nature of the traveling bright features and discriminate whether they are shocks. The frames displayed in panels f-i of Figure 4.4 show four of these features, with the corresponding maps of the instantaneous fractional polarization reported below. If the bright components were caused by the propagation of shocks, we would expect that the component's fractional polarization m_ℓ would be higher compared to other areas of the jet (Marscher and Gear 1985). However, even though the motion of a bright feature through the jet causes a local, slight disruption of the underlying polarization structure, we observe that, indeed, there is no correlation between the instantaneous maps of m_ℓ and the position of the features in the total intensity frames. Additionally, since the majority of components are ejected towards the southern side of the jet, we would expect that the fractional polarization of the average map should be higher in that area, but that is not the case.

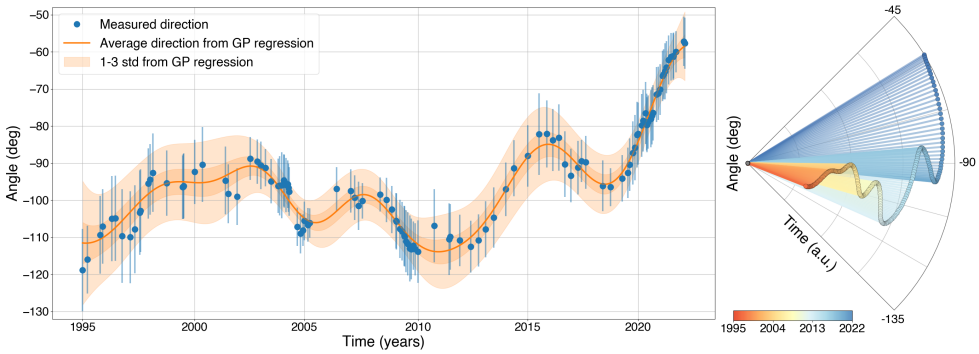

Figure 4.5: **Jet launching direction in 3C 345.** The convention used for the angle is 0 pointing west, with positive values in the north-west quadrant. (Left panel) Blue points indicate the angle computed from the frame corresponding to each observed epoch, with error bars indicating one sigma. The orange line indicates the mean of a Gaussian Process (GP) regression to the angle values, with the orange bands indicating one and three sigmas. (Right panel) The average of the GP regression is plotted in angular coordinates, for a more intuitive representation of the change in the launching direction. Red colors indicate the oldest epochs, while blue colors indicate the most recent epochs.

4.3.3 Core orientation and jet precession

For each frame fo the video, we computed the jet launching direction angle by fitting an elliptical Gaussian to the core and identifying the angle with the orientation of the Gaussian's major axis. The direction is shown as a function of time in Figure 4.5, where blue points indicate the measured direction, while orange points show the average and standard deviation of a Gaussian Process regression to the data. The projected jet launching direction has an overall variation of 60°, with intermediate smaller oscillations, with the jet mostly ejecting material towards -120° to -90° east of north. The variation of the direction angle, however, is irregular and we don't observe any periodicity in the ejection direction.

4.4 Conclusions

In this work we present a comprehensive analysis of the parsec-scale kinematics of the relativistic jet in 3C 345, obtained from a full polarimetric video reconstruction of the source. We observe an inner jet with complex, turbulent motion, characterized by the ejection of bright components from the core.

The orientation of the core varies irregularly spanning a projected angle of 60°, affecting the initial direction of the components, which travel ballistically for the first 2-3 mas before deviating towards the center of the jet and diffusing into a northwards bent plume. Caproni and Abraham (2004) and later Lobanov and Roland (2005), employed a binary black hole (BBH) model (Begelman et al. 1980) to explain the observed optical and radio variability of 3C 345. Caproni and Abraham (2004) observe a quasi-ballistic motion of superluminal features within 1 mas from the core, which they interpret as the result of a non-coplanar secondary black hole inducing torques in the inner parts of the disk. Lobanov and Roland (2005), instead, propose an equal-mass

BBH system paired with a two-fluid model for the jet of the primary BH and fit the combined model to the optical and radio light curve and the trajectory of an ejected superluminal component. They obtain a characteristic rotational period of the accretion disk of 240 years, and an orbiting period for the secondary black hole of 480 years, which induces a precession of the accretion disk axis with a 2570 years period. They attribute the yearly variability timescale observed in the light curve and component ejection process as the result of thermal and mechanical instabilities and accretion rate variations. What we observed instead is that the short-term variability of 3C 345 is directly linked to the change in the orientation of the jet core, which, nevertheless, we find to be irregular rather than periodic or quasi-periodic. The irregular jet precession was also observed in blazar 3C 84 (Foschi et al. 2025a) and could be attributed to instabilities in the accretion disk surrounding the central black hole or a misalignment between the angular momentum of the accretion disk and the black hole spin (Pringle 1997; Dunn et al. 2006; Liska et al. 2018; Liska et al. 2019; Liska et al. 2021), which may cause a warping of the disk and a stochastic variation of the jet's direction.

The polarization pattern displays a transverse-parallel-transverse structure with respect to the jet axis, that is occasionally affected by the passing of bright features. This could be a clear indicator of the presence of a long-lasting toroidal magnetic field threaded to the jet in 3C 345, as previous studies have pointed out (Lyutikov et al. 2005; Fuentes et al. 2018; Fuentes et al. 2021), and which could be associated to a large-scale helical magnetic field with a high magnetic pitch angle, that is dominated by the toroidal component (Gabuzda et al. 2014). In Pushkarev et al. (2023), a very similar, but lower resolution, EVPA pattern is reported for 3C 345 from the stacking of 51 MOJAVE epochs. The authors interpret this pattern as the result of either a shock-induced division of the field between the jet spine and the edges, or a helical magnetic field in which the superposition of differently polarized regions lead to a spine-edge stratification. As discussed in the next paragraph, we discard the shock-ordering case in favor of a helical magnetic field. A future analysis with combined multi-frequency and multi-epoch data will enable the possible detection of rotation measure gradients, thus confirming the existence of a helical magnetic field in 3C 345 (Gabuzda et al. 2015).

We computed the apparent local instantaneous velocity of the jet plasma using optical flow. The average flow speed in the jet changes from 10-12 c within the first 3 mas from the core, down to 5-8 c beyond 6 mas from the core. This shows that the plasma undergoes a deceleration as it travels further from the jet core.

We integrated the velocity field to obtain the trajectories followed by volume elements of the plasma. The OF trajectories that begin close to the core initially diverge from the jet axis, following an apparent ballistic motion, then converge towards it at a distance of 2-3 mas from the core. Other trajectories show that the plasma fluid changes direction following the apparent bend in the jet at 5-8 mas from the core. This suggests the presence of an ISM wind which deviates the jet's directions and steers the plasma flowing along the jet. Another evidence of the interaction between the jet and the surrounding ISM is the increased fractional polarization in proximity of the jet borders, especially the southern one. This may be due to the shearing of the interface layer between the jet and the ISM, which would increase the ordering of the magnetic field on the interface and as a consequence increase the fractional polarization (Meenakshi et al. 2023).

Optical flow trajectories that follow discrete components are consistent with those obtained in previous model-fitting studies, supporting the validity of the method. However, we find no evidence that the bright components coincide with shocked regions in the plasma flow, as previously proposed (Blandford and Königl 1979; Marscher and Gear 1985). Instead, our images show that the instantaneous speeds of the components (10-13 c) are comparable to the time-averaged speeds of the bulk flow in the same regions (11–12 c). In the presence of strong shocks, one would expect a significant difference between the bulk velocity of the shocked material and the pattern velocity of the shock. Nevertheless, it is worth noting that similar values for the shock pattern speed and the plasma flow speed can also arise if the protons are not highly relativistic (Marscher 2006). The fact that the peaks in the fractional polarization maps do not correlate with the positions of the bright features further support a scenario in which no strong shocks are present in the jet. A traveling shock compresses the plasma in the shocked region, amplifying the component of the magnetic field parallel to the shock front due to the increased density. In the case of a turbulent magnetic field or an ordered field dominated by a toroidal component, compression enhances magnetic field ordering and increases the fractional polarization within the shocked region, whereas a predominantly poloidal magnetic field would instead lead to a decrease in fractional polarization (e.g. Marscher and Gear 1985; Hughes et al. 1985). Since the polarization pattern in our observations points to a predominantly toroidal magnetic field, we would expect to see localized increases in fractional polarization at the positions of the bright components if they were shocks. However, this is not observed. Taking all this into consideration, we interpret these bright, compact features as regions of increased pressure caused by plasma turbulence, with enhanced emissivity due to Doppler boosting when they are ejected toward the south. This could not be concluded from earlier image reconstructions of 3C 345 (e.g. Lister et al. 2021) because model fitting is only able to capture the speed of moving components, which would coincide with the pattern speed of a possible shock, but it's not able to recover the underlying flow speed. Additionally, the low resolution of polarization maps, prevented a proper comparison between the fractional polarization of a component and that of the other portions of the jet.

4.5 Discussion

Our higher resolution analysis was made possible because of two properties of our imaging method. The first is that, by imaging all observations simultaneously, `kine` is able to propagate information across frames, improving resolution and dynamic range beyond the super-resolution achievable by frame-by-frame imaging (Figure 4.7, Appendix B). The second is that, thanks to its neural representation, `kine` produces a smooth and continuous model of the flux density distribution, that can be sampled at any required time coordinate, which enables pixel-by-pixel motion analysis using tools such as the optical flow. The application of `kine` to dynamic imaging of multi-epoch observations can provide significant advances in the study of jet kinematics. From highly resolved time-continuous videos, it is possible to recover the instantaneous local velocity field of the jet plasma, instead of being limited to only tracking model-fitted components. An analysis of the multi-epoch observations of other blazar sources with our proposed imaging and analysis method may lead to a revaluation of the nature of assumed traveling shocks, or, instead, to a precise measurement of the pattern and

flow speed associated with the shocks.

Imaging of multi-epoch blazar observations is not the only possible dynamic application of the `kine` method. In fact, the algorithm was originally developed to image single observations of a rapidly varying source, like in the case of horizon-scale observations of the Sgr A* black hole by the Event Horizon Telescope. In this case, the variability timescale of the source is significantly shorter than the full observation time, which violates the aperture synthesis fundamental assumption of source invariance, and the instantaneous coverage is too sparse to constrain individual frames, so dynamic imaging is the only possible approach to spatially resolve the source's variability. An upcoming publication will soon present the validation of the method for extremely sparse observations of the supermassive black hole Sgr A* by the EHT.

Compared to other forward-modeling static imaging algorithms, `kine` achieves super-resolution equivalent to recent Regularized Maximum Likelihood (RML) methods (e.g. `eht-imaging`, Chael et al. 2018), and Bayesian methods (e.g. `Comrade`, Tiede 2022, `Resolve`, Junklewitz et al. 2016), improving over the traditional CLEAN method (Högbom 1974). In contrast to existent dynamic imaging algorithms, like `StarWarps` (Bouman et al. 2018), `eht-imaging` (Johnson et al. 2017), `ngMEM` (Mus and Martí-Vidal 2024), `kine` enforces correlations through implicit regularization and is therefore free from explicit prior assumptions about the morphology or appearance of the source. This ensures that the reconstructions are robust, and at the same time makes the algorithm applicable to any kind of source. Additionally, `kine` doesn't require a careful setting of the regularizer, nor computationally expensive exploration of the hyperparameter space. `kine`'s advantage over dynamic Bayesian methods, like `Resolve` (Arras et al. 2022), is the significantly shorter runtime and the ability to scale to large datasets.

In general, `kine` can be applied to any VLBI dataset, to single or multiple observations and it can perform both static and dynamic imaging, in full polarization. Another advantage of `kine` is that it requires minimal imaging knowledge from the user and, thanks to implicit regularization, it minimizes the bias imprinted by possibly incorrect prior assumptions about the image morphology. Because of this, in addition to its potential for imaging dynamic sources, like blazars, black holes, or X-ray binaries, `kine` is also a promising method for various VLBI tasks which may require automation. More broadly, the method has potential for applications to imaging problems also in other fields of science. For example, in medical diagnosis, magnetic resonance imaging (MRI) needs to solve the same formal problem as VLBI imaging, and `kine` could prove useful for the imaging of moving organs, like the heart or lungs. The deep learning framework employed by `kine`, can be naturally extended to multi-frequency imaging, by treating frequency as one extra dimension of the input coordinates. Future work will be aimed at implementing multi-frequency imaging paired with the estimation of the spectral index map.

4.6 Appendix A. Synthetic data and validation

We validated the algorithm on a synthetic dataset prepared to resemble the morphology, coverage and noise of the real 3C 345 data. The ground truth video consists of a geometric jet model where the jet's emission direction is precessing along a circular orbit and the emitted components move ballistically, expanding in size and decreasing in brightness as they move away from the core. The ground truth model is shown in

the first column of Figure 4.6. For each observation epoch we simulated a synthetic dataset from the jet model, using the corresponding real u-v coverage and applying the same amount of thermal noise corruption. We did not need to introduce complex gain corruption, since the closure quantities used in the imaging process are invariant under gain errors.

We imaged the synthetic data with the same parameters that were later used to image real data. In both cases, we first imaged the total intensity video using only closure phases and logarithmic closure amplitudes, so that the reconstruction would be completely free from any hypothetical residual calibration error. Then we self-calibrated the data to the resulting video, and we imaged the linear polarization video using polarimetric complex visibilities of the self-calibrated data. The reconstructions of synthetic data are displayed in the second row of Figure 4.6 and show that `kine` is able to recover the model with high fidelity. This is reflected by the normalized cross correlation between the ground truth frames and the reconstructed ones, which has an average value of 0.988 for both the total intensity and the linear polarization images, with minimum values of 0.975. The reconstructions achieve a dynamic range of $\sim 10^5 - 10^6$ for both the total intensity and the polarization images.

4.6.1 Static vs dynamic comparison

In the dynamic imaging of 3C 345, the instantaneous coverage corresponding to one frame is the coverage of one full observation. This means that the instantaneous coverage is more than sufficient for snapshot imaging of the individual epochs. For a set of multiepoch observations, `kine` is able to provide both dynamic and snapshot static reconstructions in super-resolution. However, there are some advantages in obtaining a dynamic imaging video rather than a sequence of independent static images.

One reason is that the video provides a continuous representation of the movement and evolution of the source. A sequence of independent images, even if aligned correctly, always presents discontinuities from frame to frame, which cannot be overcome by image interpolation. Having a continuous reconstruction is important for any image domain analysis, especially e.g. for computing the optical flow. Furthermore, the continuous formulation of `kine`'s framework allows for the sampling of the reconstructed video at any time value, providing a motion-preserving interpolation between observed frames.

Another reason is that dynamic reconstructions achieve better dynamic range and eliminate the minor artifacts present in some of the snapshot reconstructions, by leveraging information from neighboring frames. Imaging artifacts don't correspond to any real structure in the source image and, since the u-v coverage and noise change from epoch to epoch, they are not consistent from frame to frame. Because of this, they are not recovered in a time-regularized video, which requires frame-to-frame consistency. The dynamic range of an interferometric image depends on the signal-to-noise ratio (SNR) and the u-v coverage. For some epochs of the 3C 345 observations, SNR and coverage are of poorer quality than the majority of epochs, which reduces the dynamic range of static reconstructions from those epochs. This problem is solved with dynamic imaging because the missing information in one epoch can be recovered by the propagation of information from neighboring epochs. Figure 4.7 shows a comparison between traditional static imaging with CLEAN, snapshot static imaging with `kine` and dynamic imaging with `kine`. Both `kine` reconstructions show consistency

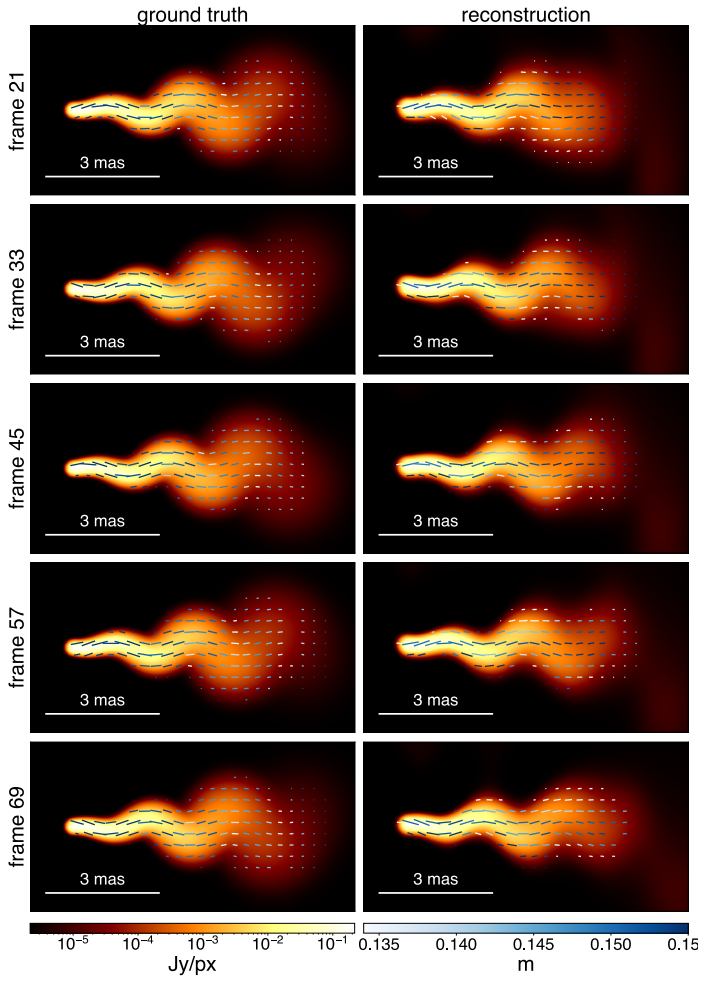

Figure 4.6: **`kine` reconstruction of synthetic data** Validation of the `kine` algorithm on synthetic data mimicking VLBA observations of 3C 345 at 15 GHz. The figure compares a geometric jet model (left column) to the `kine` reconstructions (right column) of synthetic data generated from it. A few sample frames are shown as example. The ticks indicate the direction of the EVPA and their length is proportional to the logarithm of the linear polarization.

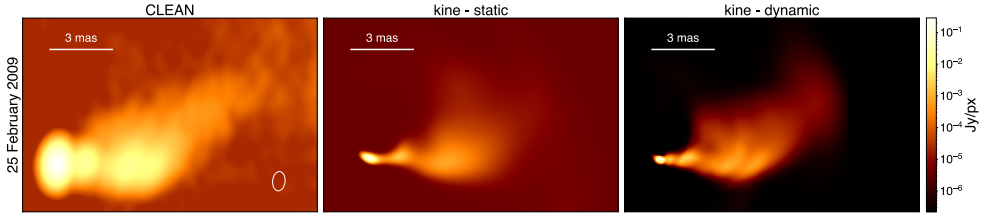

Figure 4.7: **Image comparison among CLEAN, `kine` static, `kine` dynamic.** Comparison between representative reconstructions of VLBA observations of 3C 345 from CLEAN (left column), `kine` static (center column), and `kine` dynamic (right column). The clean image is restored with the nominal beam of (870, 615) μas. The images are rescaled so that the value of the brightest pixel is the same for all images and they are clipped on the lower end to the corresponding noise level. In this way, when plotting the images with the same color scale, the difference in dynamic range is made clear by the background color of the images.

with the CLEAN image, though they are able to resolve significantly more structure in the jet. However, the `kine` dynamic imaging reconstruction achieves better dynamic range than both the CLEAN image and the `kine` static snapshot.

We quantify the dynamic range improvement, by computing, for all epochs, the dynamic range of CLEAN images, `kine` static images and `kine` dynamic frames. The dynamic range is defined as the ratio between the maximum value in the image and the image noise level. We defined the noise level as the average image value in a portion of the image not containing any real emission. The results are displayed in the left panel of Figure 4.8. The average dynamic range is 3.7×10^3 for CLEAN, 3.9×10^4 for `kine` static and 8.6×10^5 for `kine` dynamic, resulting in an improvement of two and a half orders of magnitude for `kine` dynamic over the traditional CLEAN imaging method. Additionally, the dynamic range of the two static snapshot methods undergoes wide oscillations, depending on the quality (SNR and (u, v)-coverage) of each dataset, while the dynamic range in `kine` dynamic varies less. This means that `kine` dynamic is effective in improving the reconstruction quality of a poor dataset by leveraging the information from prior and subsequent epochs.

To compare the effective resolution of the three imaging methods we computed the 2D spatial power spectrum (PS) of each frame. The power spectrum measures the strength of the image signal at different spatial scales and drops in correspondence of the smallest spatial scales resolved. The right panel in Figure 4.8 shows the PS of individual frames (thin lines) and the average PS (thick lines). To determine the resolution of the `kine` images, we blurred the images with Gaussian kernels of different FWHM and we computed the χ^2 between the PS of the blurred images and the PS of the corresponding CLEAN model convolved with a 800 μas circular Gaussian. The FWMH which resulted in the best χ^2 fit of the PS, was then subtracted in quadrature from the reference resolution of 800 μas, to obtain the effective resolution of the `kine` reconstructions. We finally average the effective resolution across all epochs. From this procedure we obtain a resolution of 127 μas for `kine` static and a resolution of 81 μas for `kine` dynamic, representing a improvement factor of 3.7 for `kine` static and 5.8 for `kine` dynamic with respect to CLEAN, whose average nominal resolution of 475 μas is set by the diffraction limit.

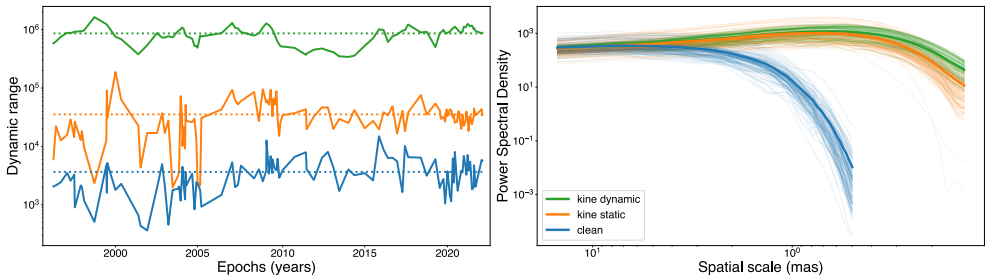

Figure 4.8: **Dynamic range and resolution comparison.** (Left panel) Dynamic range of the `kine` dynamic, `kine` static, and CLEAN reconstructions of 3C 345 observations. Continuous lines indicate the dynamic range of each individual frame, while the dotted lines indicate the average value. (Right panel) Power spectrum of the `kine` dynamic, `kine` static, and CLEAN reconstructions of 3C 345 observations. Thin lines indicate the power spectrum of each individual frame, while the thick lines indicate the average value.

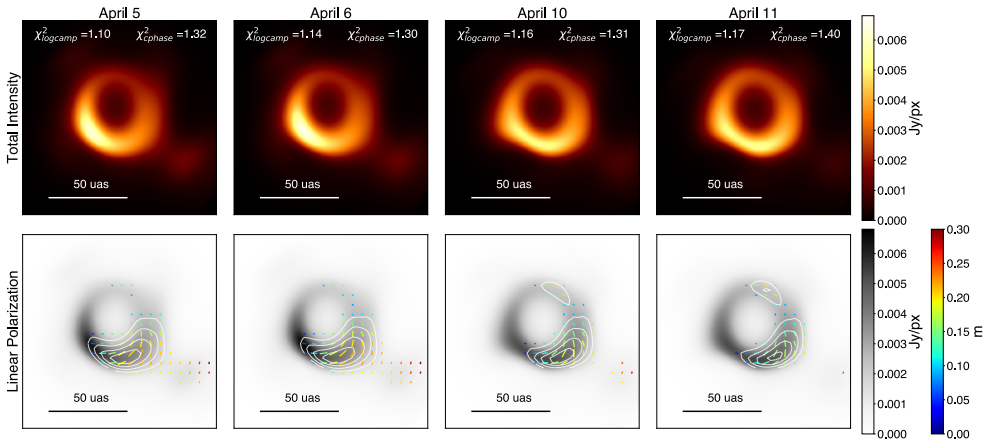

Figure 4.9: **Multi-day dynamic imaging reconstruction of M 87*,** obtained from EHT observations at 230 GHz, conducted in April 2017. The figure shows the four frames from the `kine` video that correspond to the observations days. The top panel shows the total intensity images, while the bottom panel shows the polarimetric images.

We note that the amount of dynamic range and resolution improvements obtained by `kine` are dependent on the quality and quantity of the observations, so they should not be intended as absolute values valid for any kind of dataset.

4.6.2 Testing on M 87* multi-day observations

We also tested `kine` on real, multi-day observations of M 87* at 230 GHz from the EHT campaign in 2017. Specifically, we used observations from April 5th, 6th, 10th, and 11th of 2017 (EHTC 2019b). We chose this dataset to test `kine` on a sparse (u, v)-coverage and because it has already been imaged by the EHT Collaboration with different static imaging methods (EHTC 2019d), by Carilli and Thyagarajan (2022) with CLEAN, and by Arras et al. (2022) using dynamic imaging with the `Resolve` pipeline, making it a suitable benchmark dataset to test `kine` reconstructions.

Dynamic imaging results from `kine` are presented in Figure 4.9, where the frames corresponding to the four observation days are shown. The reconstructions fit the data correctly, as demonstrated by the normalized χ^2s of the closure quantities used in the imaging process. The images present a clear ring-like shape, with a brighter spot in the south-west of the ring and a twisted polarization pattern in the south-east of the ring. A faint extended emission is also visible in the south-east of the ring for all observed days, consistent with the location of the M 87 jet observed at lower frequencies (see e.g. Walker et al. 2018; Kim et al. 2018). To quantify the consistency of our results with previous imaging of the dataset, we extracted the ring features from the frames at each observed day. Our results are presented in Table 4.1 and show agreement with values from the EHT Collaboration.

4.7 Appendix B. Extended data

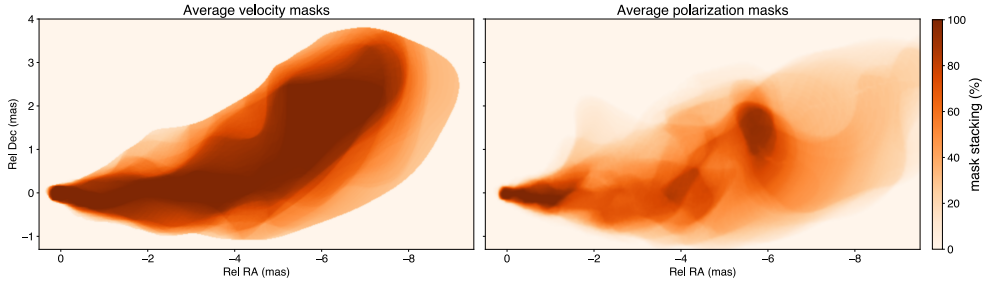

Figure 4.10: **Staking of the instantaneous threshold masks** applied in the computation of the average velocity (left) and average fractional linear polarization (right), the darkest color indicates 100% of epochs used in the average. These maps provide a measure of the confidence in the estimation of average speed/polarization values in different areas of the jet.

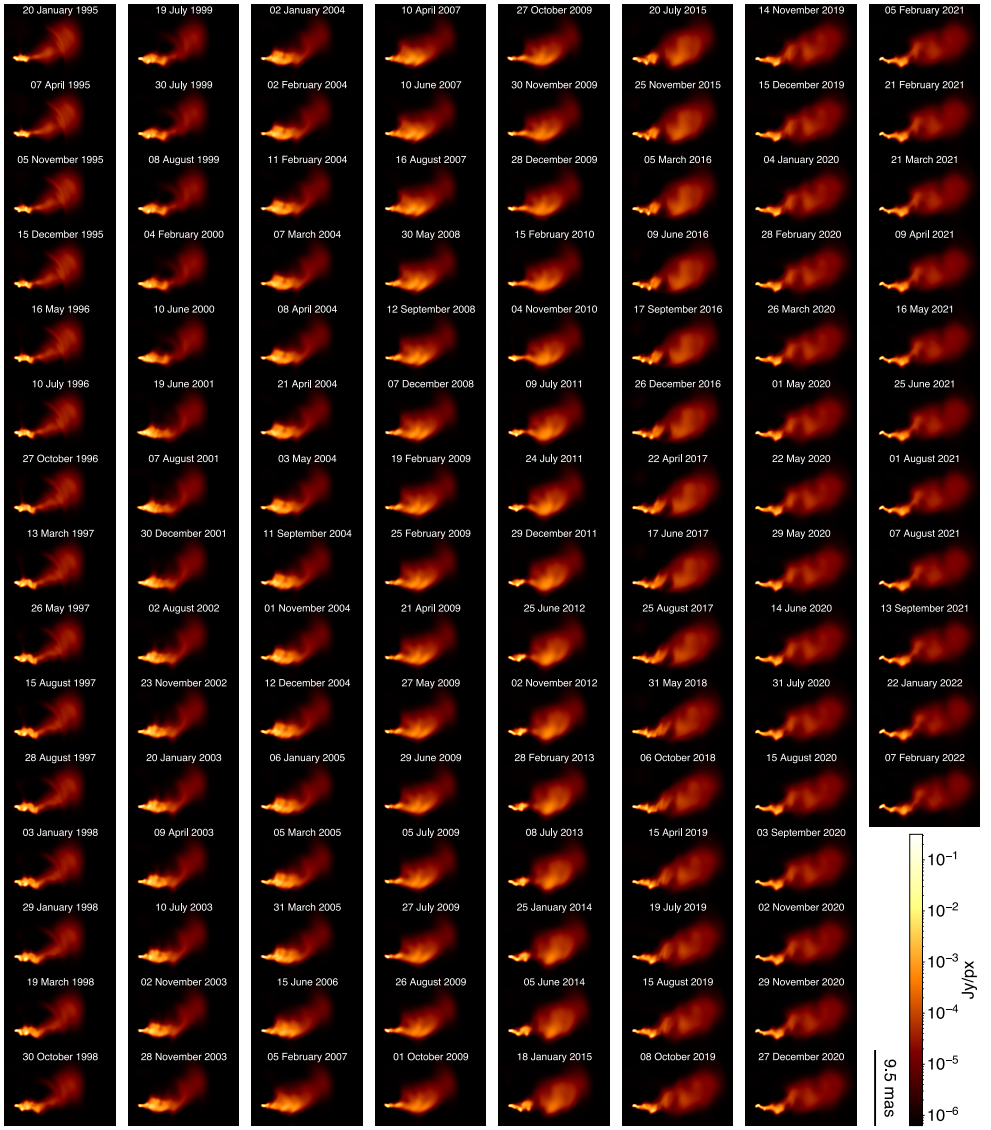

Figure 4.11: **`kine` reconstruction of 3C 345, total intensity.** All 116 total intensity frames of the dynamic imaging reconstruction of blazar 3C 345, from MOJAVE observations with the VLBA at 15 GHz, by the `kine` method.

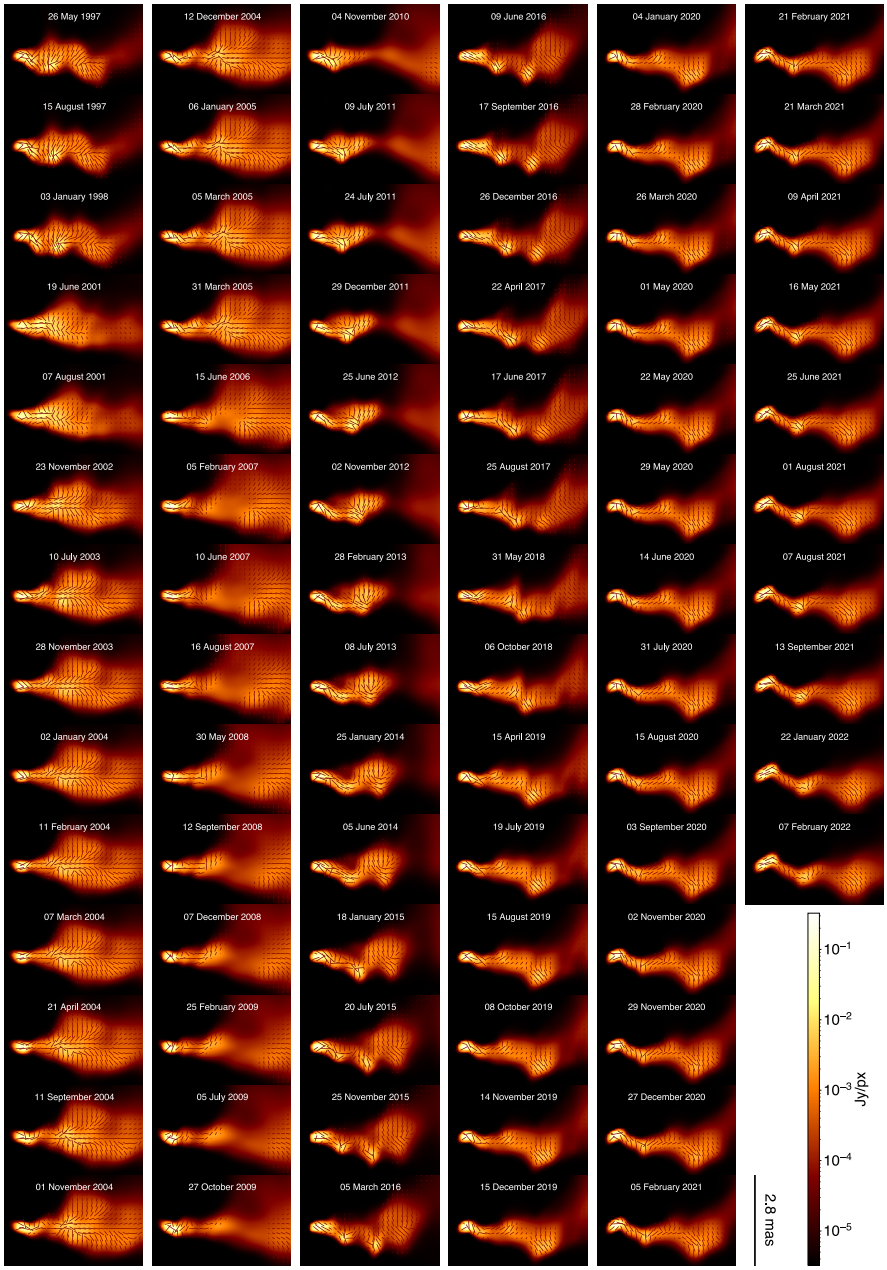

Figure 4.12: **`kine` reconstruction of 3C 345, polarization.** All 116 polarization frames of the dynamic imaging reconstruction of blazar 3C 345, from MOJAVE observations with the VLBA at 15 GHz, by the `kine` method.

	d (μas)	w (μas)	η (°)	A	f_c
DIFMAP					
April 5	37.2 ± 2.4	28.2 ± 2.9	163.8 ± 6.5	0.21 ± 0.03	0.5
April 6	40.1 ± 7.4	28.6 ± 3.0	162.1 ± 9.7	0.24 ± 0.08	0.4
April 10	40.2 ± 1.7	27.5 ± 3.1	175.8 ± 9.8	0.20 ± 0.04	0.4
April 11	40.7 ± 2.6	29.0 ± 3.0	173.3 ± 4.8	0.23 ± 0.04	0.5
eht-imaging					
April 5	39.3 ± 1.6	16.2 ± 2.0	148.3 ± 4.8	0.25 ± 0.02	0.08
April 6	39.6 ± 1.8	16.2 ± 1.7	151.1 ± 8.6	0.25 ± 0.02	0.06
April 10	40.7 ± 1.6	15.7 ± 2.0	171.2 ± 6.9	0.23 ± 0.03	0.04
April 11	41.0 ± 1.4	15.5 ± 1.8	168.0 ± 6.9	0.20 ± 0.02	0.04
SMILI					
April 5	40.5 ± 1.9	16.1 ± 2.1	154.2 ± 7.1	0.27 ± 0.03	7×10^{-5}
April 6	40.9 ± 2.4	16.1 ± 2.1	151.7 ± 8.2	0.25 ± 0.02	2×10^{-4}
April 10	42.0 ± 1.8	15.7 ± 2.4	170.6 ± 5.5	0.21 ± 0.03	4×10^{-6}
April 11	42.3 ± 1.6	15.6 ± 2.2	167.6 ± 2.8	0.22 ± 0.03	6×10^{-6}
kine					
April 5	41.2 ± 4.0	18.2 ± 3.8	161.0 ± 21.2	0.22 ± 0.04	0.3
April 6	41.2 ± 3.6	18.0 ± 3.8	157.2 ± 18.2	0.21 ± 0.05	0.3
April 10	42.3 ± 2.7	17.3 ± 2.6	172.8 ± 29.5	0.17 ± 0.03	0.2
April 11	43.3 ± 2.9	17.4 ± 2.7	169.1 ± 33.2	0.16 ± 0.02	0.2

Table 4.1: **Parameters of the M 87* black hole shadow.** Diameter d, width w, orientation angle η, asymmetry A and floor-to-ring contrast ratio f_c of the M 87* black hole shadow. `kine` results are compared with results from static imaging with `DIFMAP`, `eht-imaging`, and `SMILI` from the EHT imaging of M87 (EHTC 2019d).

5

Dynamic imaging of Sgr A*

Adapted from

Validation of horizon-scale Sagittarius A* video reconstruction with kine

A. Fuentes[†], M. Foschi[†], et al.

In preparation, (2025)

[†]These authors contributed equally to this work

Abstract

The hot plasma orbiting Sagittarius A* (Sgr A*), the supermassive black hole at the Galactic Center, exhibits a variability time-scale of minutes. The Event Horizon Telescope is the only facility capable of resolving the horizon-scale structure near Sgr A*, but this variability introduces unprecedented challenges for image reconstruction. In this work, we extend the capabilities of our novel dynamic imaging algorithm, `kine`, to tackle the specific challenges of EHT data, and we apply it to a large suite of synthetic data generated after EHT observations of Sgr A* on 2017 April 11. The aim is to assess `kine`'s ability to reconstruct the ground-truth dynamics of these models with the extremely sparse coverage of EHT observations in 2017. We test `kine` against a variety of static and dynamic geometric models and general-relativistic magneto-hydrodynamic black hole simulations, which are specifically designed and selected to test different forms of source morphology and intrinsic variability. The simulated datasets were generated with the expected thermal noise and gain corruption, as well as the corruptions introduced by interstellar scattering. We prove that `kine` can discriminate and recover static models without introducing any motion, while recovering with high precision the dynamics of both coherent and incoherent models, both tangential and radial motion, with orbital speeds ranging from 20 to 180 min, and different relative flux densities.

In the case of GRMHD models, which span different inclination angles, spins, magnetization, and electron temperature, we show that `kine` is capable of recovering the correct morphology, position angle, pattern speed, and linear polarization orientation of all models resembling Sgr A* properties. For GRMHD models disfavored by Sgr A* observations, `kine` is not always able to provide a good reconstruction, but is still able to robustly recover the direction of the pattern speed and the orientation of the linear polarization. Remarkably, `kine` can recover also more subtle features like spiral arms.

5.1 Introduction

The Sgr A* black hole at the center of the Milky Way is one of the two supermassive black holes that can be resolved at horizon scales by Very-Long Baseline Interferometry observations with the Event Horizon Telescope (EHT, EHTC 2022a). Imaging Sgr A* is more challenging compared to other sources like M 87*, because of the interstellar scattering produced by foreground material in the galactic plane and because of the intrinsic intra-day variability of the source. EHT observations are extremely sparse and rely heavily on the aperture synthesis technique, which exploits the rotation of the Earth to sample additional spacial frequencies, under the assumption that the source remains unchanged during the observations. This is problematic in the case of Sgr A*, whose emission varies on timescales of 5 to 30 minutes (Wielgus et al. 2022), that are significantly shorter than the few hours required for a full observation. A possible way to address variability and recover its spatial distribution is with dynamic imaging.

Dynamic imaging aims at recovering a video of the source, in which frames are constrained to the corresponding instantaneous observations and correlations are enforced across frames from different times to make up for the sparsity of the instantaneous coverage. Reconstructing the flux density distribution not only in space but also in time adds an additional degree of freedom, compared to static imaging, while the data constraints remain the same. This makes the VLBI imaging problem even more ill-posed than in static imaging, so the imaging algorithms need to rely more on space-time regularization, i. e. enforcing space-time correlations in the output video. Different dynamic imaging methods implement regularization in different ways. Regularized Maximum Entropy (RML) methods, like `eht-imaging` (Johnson et al. 2017), `DoG-HiT` (Müller and Lobanov 2023), or `ngMEM` (Mus and Martí-Vidal 2024), recover a video by optimizing the values of its pixels so to minimize a loss function constraining the video to the data. In this case, regularization is provided by explicit regularizer terms added in the loss function. Bayesian methods, like `StarWarps` (Bouman et al. 2018), or `Resolve` (Arras et al. 2022), assume a parametrized distribution of the video's pixel values and solve for the parameters that yield the most likely video given the data constraint. In this case, regularization is introduced by assuming a correlated probability distribution of the pixels.

Recently, as described in Chapter 4, we proposed a new approach to dynamic imaging using neural field representations. We developed `kine`, an imaging algorithm, which employs a neural network to model the flux density distribution as a continuous parametrized function, which is optimized so to minimize a loss function containing the flux density video to the data (Foschi et al. 2025b). The neural network employed in the algorithm is affected by spectral bias, meaning that it tends to fit for the lowest frequencies first and thus it is not able to produce high resolution outputs. This property provides implicit space-time regularization of the video, without the need of adding explicit, possibly misleading regularizer terms in the loss function. We have proven the validity of the method in the case of multi-epoch observations of evolving blazar sources with the Very Long Baseline Array (Foschi et al. 2025b), and demonstrated the resolution and dynamic range improvement provided by `kine` over other imaging algorithms. However, before applying the method to Sgr A* observations with the EHT, we need to demonstrate that `kine` is able to recover a reliable video from data with the same (u, v)-coverage, signal-to-noise ratio (SNR), gain corruption,

and scattering corruption of real EHT data of Sgr A*. The purpose of the present work is to determine whether the coverage and noise properties of EHT observations of Sgr A* on April 11$^{\text{th}}$ 2017 are good enough to allow reconstruction of a dynamic imaging video with the `kine` pipeline.

Therefore, we used `kine` to image synthetic data produced from a variety of models, with different morphologies, motion, and amount of variability. In particular we considered static geometric models, geometric models with simple dynamics, and ray-traced videos of more realistic general-relativistic magneto-hydrodynamics simulations of accreting black holes. The data are simulated with the same properties of the EHT observations of Sgr A* on April 11$^{\text{th}}$ 2017, since that is the dataset that we aim to image. Specifically, these observations hold the largest time window of the 2017 EHT campaign over which the instantaneous (u, v)-coverage is maximized. In addition, this window starts only $\sim$2 hours after an X-ray flare, making it potentially interesting for reconstructing source dynamics. More details on the observed data and the scientific motivation driving this analysis may be found in EHTC (2025). A second goal of this work is to test the extent of the range of possible videos that can be reconstructed by `kine` from data with this coverage. Specifically we test the limits of the speed of moving objects that can be recovered and the flux ratio between distinguishable features in the video.

In section 5.2, we provide a brief description of the general `kine` imaging algorithm, followed by an explanation of the algorithm variations and the specific pipeline that were developed to image EHT observations of Sgr A*. In section 5.3 we show the application of the pipeline to imaging synthetic data, proving the method's ability to recover the correct ground truth over a variety of sources, and exploring the limits of its capabilities. Finally, in section 5.4, we discuss the method's overall performance and draw conclusions on the robustness of the recovered videos.

5.2 Methods

5.2.1 The `kine` imaging algorithm

`kine` is a novel algorithm for VLBI video reconstruction which models the full polarimetric brightness distribution of a source trough a coordinate-based neural network (Foschi et al. 2025b). The algorithm draws inspiration from Neural Radiance Field networks (NeRFs, Mildenhall et al. 2021), and models the Stokes parameters $(\mathcal{I}, \mathcal{Q}, \mathcal{U}, \mathcal{V})$ of the recorded signal as a continuous function over space and time, parametrized through a Multi-Layer Perceptron (MLP, Popescu et al. 2009):

$$(\mathcal{I}, m_\ell, \chi, m_c)(x, y, t) = \text{MLP}_{\mathcal{W}}(x, y, t)\,, \tag{5.1}$$

where (x, y, t) are the real-valued domain coordinates and $\mathcal{W}$ are the parameters of the network. `kine` recovers the degree of linear polarization m_ℓ, the electric vector position angle χ, and the degree of circular polarization m_c, and then applies a trivial transformation to obtain $\mathcal{Q}, \mathcal{U}$, and $\mathcal{V}$. Similarly to other forward-modeling algorithms, the optimization of the network's parameters is performed through the minimization of a χ^2 loss between the observed data and the recovered video, evaluated over the points of the Fourier space sampled by the interferometer at any given time. The optimization step is repeated iteratively until convergence of the loss function. The

network parameters are optimized only with respect to the data to be imaged, making the algorithm completely unsupervised.

`kine` models the temporal evolution of the source by simultaneously fitting the frames in the video to all available observations. In the algorithm, regularization in space and time is not imposed via explicit prior constrains but rather learned from the data itself by leveraging the spectral bias of coordinate-based MLPs (Rahaman et al. 2019). It has been proven that MLPs output first a low frequency image or video and gradually introduce higher frequencies as the optimization continues. However, convergence is reached before overfitting, preventing the appearance of high-frequency artifacts. A complete description of the method can be found in Foschi et al. (2025b).

5.2.2 New additions to kine

The rapid variability of the horizon-scale plasma accreting onto Sgr A*, in conjunction with the extremely sparse sampling of the EHT over 1 minute intervals, poses specific challenges to video reconstruction. Additionally, visibility phases at 1 mm are quickly corrupted by fluctuations of the atmosphere, making the exploitation of complex visibilities as a data product impossible, even assuming a perfect visibility amplitude calibration. Closure phases, on the contrary, are resilient to changes in the atmosphere (Jennison 1958; Rogers et al. 1974; Twiss et al. 1960), at the cost of providing less data constraints than visibility phases and no information about the source absolute position.

This, in combination with the fact that `kine` does not impose any regularization on how or where the brightness distribution should be localized in the field of view, makes the recovered emission drift slightly from frame to frame. While this effect does not interact with the capacity of `kine` to reconstruct the correct dynamics, it requires video post-processing to remove the drifting, which is not always trivial and can induce interpretation errors. We overcome this issue by assuming that the brightness distribution of an evolving source can be decoupled into a persistent and an variable component. As a consequence, the persistent, or static, component effectively anchors the variable component to the static emission and the drifting disappears. Both the persistent and the variable component can be as small as necessary, resulting in either a purely dynamic or purely static emission. While this approach can be extended to all Stokes parameters, in practice we only need to consider this decoupling for total intensity. Once a good reconstruction in Stokes $\mathcal{I}$ is achieved, visibility phases can be recovered through self-calibration and Stokes $\mathcal{Q}$, $\mathcal{U}$, and $\mathcal{V}$ can be then reconstructed using complex visibilities. We implement this in `kine` by modeling the static and dynamic components through separated MLPs (Figure 5.1), whose output is added to represent the complete predicted video. That is:

$$\mathcal{I}(x, y, t) = \mathrm{MLP}_{d,\mathcal{W}}(x, y) + \mathrm{MLP}_{s,\mathcal{W}}(x, y, t) \, . \tag{5.2}$$

The optimization of the two networks' parameters is carried out in parallel, as the χ^2 loss is computed on the video resulting from the addition of the two networks. We found that, as a byproduct of the parallel but separate modeling of the static and dynamic component, `kine` can recover the time-dependent true structure from synthetic data with significantly higher accuracy.

In order to effectively decompose the video into a static and dynamic component, the total flux assigned to each components must be constrained. Otherwise,

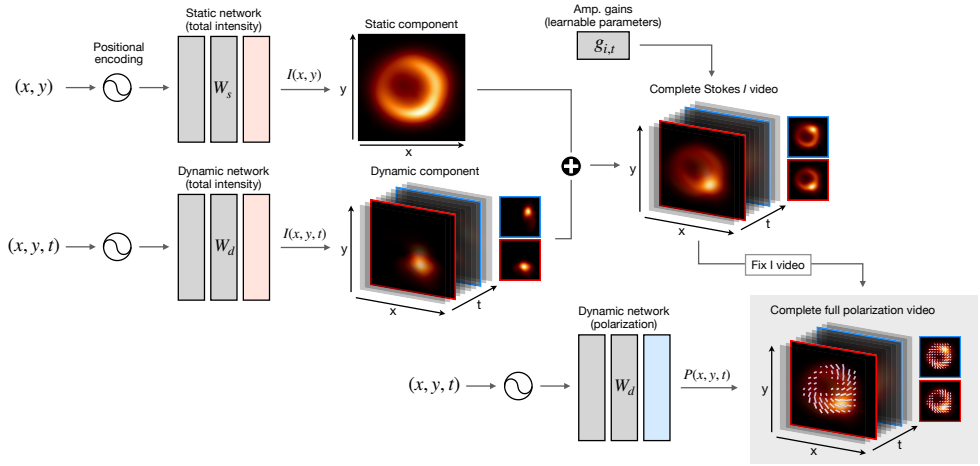

Figure 5.1: **Architecture of the `kine` network.** Diagram of the specialized `kine` network architecture developed for the imaging of Sgr A* observations by the EHT.

the dynamic component could absorb the totality of the emission, leaving the static component to zero and failing the purpose of the decomposition. The lower limit of the dynamic component's total flux is set by the difference between the minimum and maximum value of the source's light curve. Nonetheless, the dynamic component may include persistent dynamic emission with constant flux density. We therefore incorporate an extra term in the loss function which attempts to find the dynamic-to-static flux density ratio by penalizing stationary features in the dynamic component. Specifically, the resulting loss term is expressed as:

$$\mathcal{L} = \sum_{D} \chi^2(\hat{D}, D) + \sum_{x,y} \max_{t} \mathrm{MLP}_{d,\mathcal{W}}(x, y, t) \ , \tag{5.3}$$

where the first term accounts for the goodness-of-fit between the data products $\hat{D}$ computed from the video reconstruction and the observed ones D, while the second encourages persistent emission to be assigned to the static component, ensuring the dynamic one only includes variable emission. We introduce an initial imaging round, employing the above loss function to find the optimal dynamic-to-static flux ratio. In the following imaging rounds the total fluxes of the two components are fixed according to the obtained ratio and the flux-ratio term is removed from the loss function. In practice, we enforce the output of both the static and dynamic networks to have a total flux density of 1 Jy. Then, before addition, we rescale their flux density with the appropriate values S_{static} and S_{dynamic}. In contrast to just enforcing each network's output to have the corresponding flux density, this approach has the benefit of numerical stability, maintaining the network's output always in a similar range of values. The final loss function then takes the form:

$$\mathcal{L} = \mathcal{L}_{\mathrm{data}} + \mathcal{L}_{\mathrm{flux}} \ , \tag{5.4}$$

where

$$\mathcal{L}_{\text{data}} = \sum_{D} \chi^2\left(\hat{D}, D\right) \ , \tag{5.5}$$

$$\mathcal{L}_{\text{flux}} = \left(\sum_{x,y} \text{MLP}_s(x,y) - 1\right)^2 + \sum_{t}\left(\sum_{x,y} \text{MLP}_d(x,y,t) - 1\right)^2 \ , \tag{5.6}$$

while the total intensity output is obtained by:

$$\mathcal{I}(x,y,t) = S_{\text{static}} \cdot \text{MLP}_{s,\mathcal{W}}(x,y) + S_{\text{dynamic}} \cdot \text{MLP}_{d,\mathcal{W}}(x,y,t) \ . \tag{5.7}$$

The static and dynamic MLPs consist of 4 and 6 fully connected layers with 256 neurons each, respectively. We choose a different number of layers because reconstructing a single image (static) versus a set of them (dynamic) requires higher network complexity. Both networks have a one-dimensional final layer where the output is passed through a sigmoid function to enforce positivity in the output $\mathcal{I}(x,y,t)$. For the hidden layers, we use the Gaussian Error Linear Unit function (GELU, Hendrycks and Gimpel 2016) and two "sharper" variants, which are given by the function:

$$f(x) = 0.5x\left(1 + \tanh\left[\sqrt{2/\pi}\left(ax + 0.044715x^3\right)\right]\right) \ , \tag{5.8}$$

where $a = 1$ for the standard GELU, and $a = \{3, 6\}$ for the sharper variants. These two functions are compared in Figure 5.2 and they are used to introduce increasing levels of complexity depending on their sharpness. We further discuss their use in the description of the pipeline. We choose the Adamax optimizer for backpropagation, a variant of the popular Adam optimizer (Kingma and Ba 2014) which provides better results with high learning rates, which in our case are set to 10^{-2} for initialization and to 10^{-3} and 10^{-4} for the optimization of the dynamic and static network, respectively. We use a higher learning rate for the dynamic component so that it converges faster, preventing the static network from interpreting intrinsic variability as high-frequency artifacts in the static component. In this work, the network's weights are pre-optimized to output a given brightness distribution, in contrast to a typical random He uniform (He et al. 2015) weight initialization. This initialization intends to alleviate some of the large uncertainty associated with the imaging problem addressed here, but remains as agnostic as possible with respect to the true source morphology. In a similar fashion to RML methods (e.g. EHTC 2022c), we initialized the network's weights to a uniform disk of 80 μas diameter blurred with a circular Gaussian of 20 μas full width half maximum (FWHM). In later stages of the imaging process, we switch to a more informed initialization, like the median image reconstructed in the previous step. We further discuss this choice in the pipeline description. During initialization we use a higher learning rate because the absence of the undersampled Fourier operator makes convergence faster.

In addition to unknown phase gains, EHT observations of Sgr A* also contain small residual amplitude gains. However, the sparsity of the instantaneous 2017 EHT array makes the information contained in closure phases and (log) closure amplitudes alone insufficient to constraint dynamics on minute time-scales. We therefore incorporate (log) visibility amplitudes as an additional data constraint, while fitting for amplitude

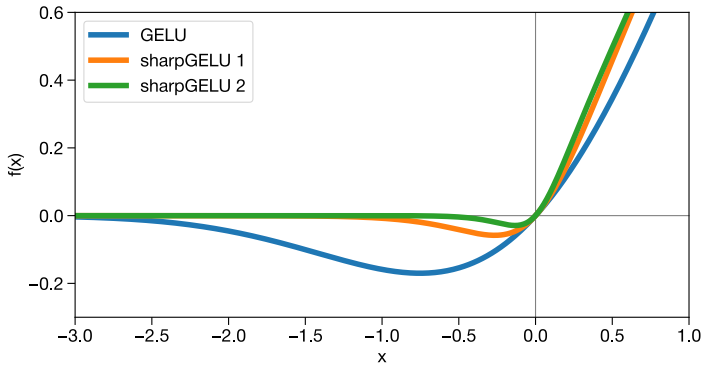

Figure 5.2: **`kine` activation functions.** used at different imaging steps. These are `GELU`, or Gaussian Error Linear Unit (blue), `sharpGELU 1` (orange), and `sharpGELU 2` (green), two custom variants that compress the range of negative values while preserving its differentiability at any point.

gains simultaneously with the video reconstruction process. The joint optimization of the output video and amplitude gain corrections enables the full exploitation of visibility amplitudes, without introducing artifacts derived from calibration errors. Visibility amplitude gains are station-dependent real-valued multiplicative factors $|g_i|$ which relate measured visibility amplitudes $|V_{i,j}|$ and true visibility amplitudes $|\mathcal{V}_{i,j}|$ as:

$$|V_{i,j}| = |g_i|^{-1}|g_j|^{-1}|\mathcal{V}_{i,j}| \,, \tag{5.9}$$

for a baseline formed by telescopes i and j. In `kine`, we model time-dependent amplitude gains through a set of $n_s \times n_t$ learnable parameters, initialized to 1, where n_s is the number of observing sites and n_t the number of time segments (or frames) in the video. This set of parameters includes the amplitude gain corrections $|g_{i,t}|$ for every telescope i at any time instant t, and is applied to the measured data before computing the χ^2 loss term against the visibility amplitudes derived from the predicted video. We constraint the range of possible amplitude gains values between 0.85 and 1.5, which amply accounts for the expected deviations in the 2017 EHT data (EHTC 2019c; EHTC 2022b; Issaoun et al. 2022). To account for the initial states of the network, in which the output video is still far from the optimal solution, we schedule the learning rate of the gain parameter exponentially from 5×10^{-5} to 1×10^{-3} between the first and last iteration of the joint optimization. We discuss `kine`'s effectiveness in recovering amplitude gains in Appendix B (section 5.6).

Contrary to the dataset imaged in Foschi et al. (2025b), where the vast amount of data provided enough information to easily recover the multi-epoch dynamics of the target source, we find that in the case of extremely sparse instantaneous coverage, such as the one we discuss in this work, positional encoding (Vaswani et al. 2017) of the input time coordinate allows for more expressivity of the dynamic network (Rahaman et al. 2019; Mildenhall et al. 2021) and therefore a notable improvement on the recovered dynamics. Effectively, the input coordinates are mapped to a higher dimensionality vector, according to the $\mathbb{R} \longrightarrow \mathbb{R}^{2L+1}$ transformation:

$$\mathbf{x} \longrightarrow (\mathbf{x}, \sin \mathbf{x}, \cos \mathbf{x}, ..., \sin 2^{L-1}\mathbf{x}, \cos 2^{L-1}\mathbf{x}) \,, \tag{5.10}$$

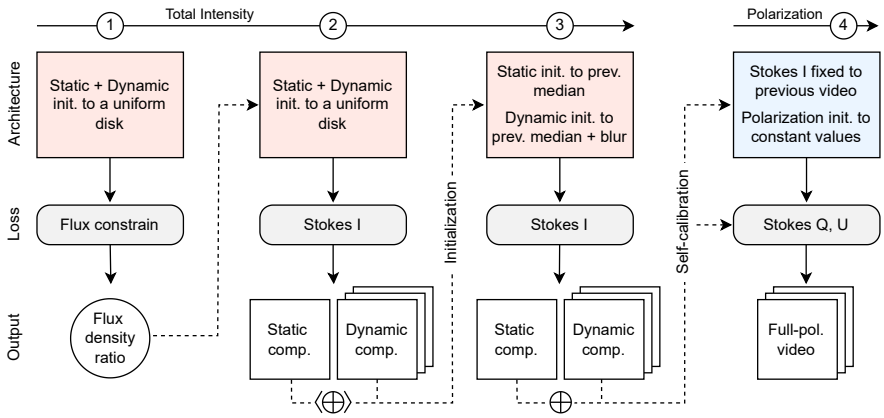

Figure 5.3: **The `kine` imaging pipeline for EHT Sgr A* data.** Reconstructing the variable Stokes I emission consists of three steps: finding the appropriate flux density of the static and dynamic component, refining the network's initialization, and performing the final imaging. Polarization is then recovered using the final Stokes I video for data self-calibration and as an anchor for Stokes $\mathcal{Q}$ and $\mathcal{U}$.

where L is the positional encoding degree. In this work we set $L = 2$ for spatial coordinates and $L = 7$ for the temporal coordinate. These values provide the best results given the space and time resolution of the measured data.

5.2.3 Sgr A* pipeline

The imaging procedure designed for the video reconstruction of synthetic and real Sgr A* EHT data begins with three steps of Stokes $\mathcal{I}$ imaging, aimed at finding the appropriate flux density of the static and dynamic components, refining the networks initialization and obtaining the final video. These are followed by a self-calibration step and one last imaging step for Stokes $\mathcal{Q}$ and $\mathcal{U}$. In this work we do not attempt to recover Stokes $\mathcal{V}$. Prior to imaging, we time-average the data into 60 s intervals within the time window with optimal coverage, that is, the time window during which the maximum number of antennas are observing simultaneously. This ~ 3 hr segment occurs between $\sim$10.85 UT and $\sim$14.05 UT (see Farah et al. 2022; EHTC 2022c, for further details). This results in 99 data segments, or frames. Additionally, we add a 1% extra error budget to account for non-closing errors. As mentioned above, we fit for log-amplitudes (with a simultaneous amplitude gain fitting), closure phases, and log-closure amplitudes (see Thompson et al. 2017, for a complete introduction to VLBI imaging and data products). Visibility amplitudes are rescaled so that $\max_{i,j} |V_{i,j}| = 1$ Jy and they are restored to their original value once the final video solution is achieved. The number of iterations performed in each step was chosen so that convergence was achieved for all the models considered. In the following we describe each of the steps, which are also summarized in Figure 5.3.

1. We perform 12000 optimization iteration with a coarse coordinate grid which outputs a 16×16 pixel video reconstruction. The static and dynamic networks are both initialized to a 80 μas diameter disk contained in a 160 μas field of view (FOV). We optimize the network by performing gradient descent on the loss func-

tion defined in Eq. 5.3, which properly separates the flux density of the static and dynamic component. For instance, the geometric model `mring+hsCW0.60` has a dynamic flux density of 0.6 Jy out of the total 2.7 Jy, that is, 22 % of the total flux density. With this approach, `kine` outputs a dynamic component with 27 % of the total flux density, while being completely agnostic to the ground-truth value. In this step, we discard the final video reconstruction and save the found flux ratio for the next steps.

2. We initialize again the static and dynamic networks to a disk, enclosed in a 160 μas FOV. The input coordinate grid is sampled at 32×32 spatial locations. After convergence is reached, we upscale the final video reconstruction by sampling the network at twice the number of spatial locations, obtaining a video of dimensions $99 \times 64 \times 64$. In this and the remaining steps, the networks are optimized according to the loss function defined in Eq. 5.4, that is, the static and dynamic outputs are enforced to have a total flux density of 1 Jy and then rescaled with the time-dependent flux ratio. We then estimate the intrinsic closure phase variability in the data by computing the $\mathcal{Q}$-metric (Roelofs et al. 2023) on the baseline triangle ALMA-SMT-LMT and we consider the following scenarios. For $\mathcal{Q}$-metric values smaller than 0.4, that is, the upper limit obtained for purely static models (see Fig. 12 in EHTC 2022b), we consider the data to have zero or residual variability. In this case, the static flux density $S_{\rm static}$ is set as the value recovered at the previous step. Otherwise, if the $\mathcal{Q}$-metric value exceeds 0.4, we set the static flux density as the minimum value between the static flux density found before and 90% of the light curve minimum:

$$S_{\rm static} = \begin{cases} S_{\rm static}, & \text{if } \mathcal{Q} < 0.4 \\ \min\left\{S_{\rm static}\,,\, 0.9 \cdot \min S_{\rm total}\right\}, & \text{otherwise} \end{cases} . \tag{5.11}$$

   The networks are optimized for 10000 iterations, reaching χ^2 loss values $\lesssim 1$ for all data products. In general, we achieve a satisfactory video reconstruction already in this step, especially for geometric models.

3. In this step we refine the networks initialization in order to capture more subtle details, which especially benefits complex datasets like GRMHD models. The static network is initialized to the median frame from the video reconstructed in the previous step. The dynamic network is also initialized to the median frame, but blurred with a Gaussian kernel of 30 μas FWHM. We increase the FOV to 200 μas in order to allow for more extended emission. The input coordinate grid is now sampled at 64×64 spatial locations, as well as the final network state, obtaining a final video reconstruction of dimensions $99 \times 64 \times 64$. We optimize the networks for 5000 iterations.

4. We self-calibrate the data in phase and amplitude to the video reconstructed in step 4 and use the self-calibrated data in the final imaging round to recover the linear polarization vector field. In this step, we fix Stokes I using the final

video reconstruction from step 4 and fit for the degree of linear polarization m_l and the electric vector position angle χ. These quantities are then transformed to obtain Stokes $\mathcal{Q}$ and $\mathcal{U}$ (Foschi et al. 2025b). In the case of polarization, the network consists of a single MLP with identical configuration as the total intensity dynamic network, except for the number of output dimensions, which is 2 for polarization, and the number of layers, which is 4 for the polarization network. The loss function takes the form of a χ^2 function between the measured complex visibilities and those recovered for Stokes $\mathcal{Q}$ and $\mathcal{U}$. The network is initialized to a flat video with constant m_l and χ values and dimensions $99 \times 128 \times 128 \times 3$, where the last dimension is the polarization channel ($\mathcal{I}$, $\mathcal{Q}$, and $\mathcal{U}$). We found that a positional encoding of degree $L = 4$ in time ($L = 0$ for space coordinates) provides the best results. Other than this, the network hyperparameters remain unchanged.

As proven in the next section, these steps ensures optimal results in total intensity for the wide variety of synthetic data considered. At the same time, this approach allows us to remain as agnostic as possible with respect to the ground-truth properties and resilient to residual amplitude gains in the data, which we can successfully disentangle from intrinsic variability.

5.3 Results

5.3.1 Imaging synthetic data

Following the suite of validation tests outlined in Dahale et al. (2025), we tested `kine`'s imaging ability on synthetic data generated from various ground truth models, specifically from static geometric models, dynamic geometric models, and GRMHD simulation models. Static models were chosen to test that time variability was not spuriously introduced by the imaging algorithms. Dynamic geometric models served as simple motion recovery tests of easy interpretation, while GRMHD simulations represented the most realistic test case. For all models, we generated synthetic data with the EHT coverage during the optimal time window on April 11, 2017. Realistic thermal noise, gain errors, and scattering corruption were added to the synthetic data (details of the data generation are explained in Dahale et al. 2025 with a summary of the ground truth models properties in Table 1). All the results presented are *on-sky*, meaning that we are not correcting scattering effects (*descattering*) by inflating the noise budget to account for refractive scattering or by de-blurring the reconstructed videos to account for diffractive scattering.

Static geometric models

The first validation test consisted in imaging static geometric models of various morphologies. The goal of the test was to make sure that `kine` is able to reconstruct the correct total intensity morphology and polarization pattern of the models without introducing any significant motion. In particular, we want to ensure that the reconstructions:

1. fit the data correctly,

2. do not introduce additional dynamics when none are present,

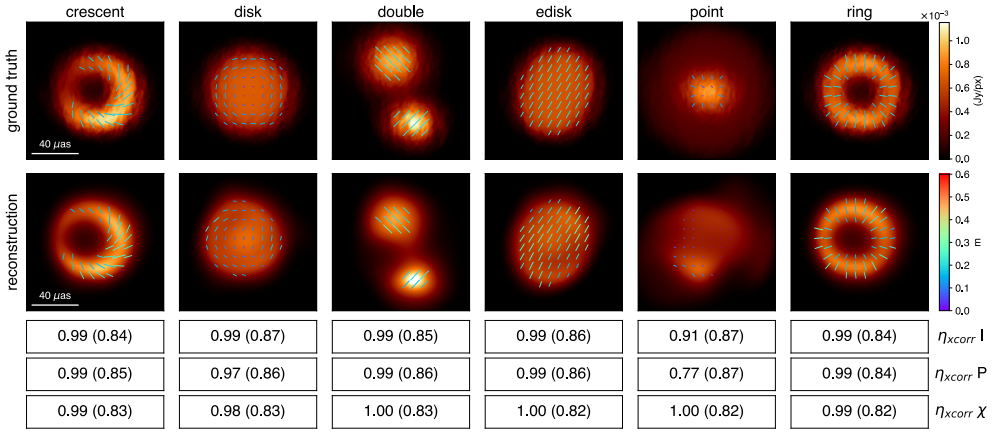

Figure 5.4: **Static models reconstructions.** (Top row) Ground truth static geometric models. (Middle row) `kine` reconstructions. (Bottom rows) Cross correlation between the ground truth and the reconstruction, in total intensity (I), polarization intensity (P), and EVPA (χ). The values in parenthesis are the minimum cross correlation thresholds which are defined in Dahale et al. (2025). The figure displays one frame of the full video reconstructions.

3. recover the correct morphology for Stokes I, Q, U.

We considered six different models, which are shown in the top row of Figure 5.4. The corresponding reconstructions by `kine` are displayed in the second row of Figure 5.4. The figure shows one representative frame of the full videos. The goodness of fit is evaluated with the time-averaged χ^2 of closure quantities in total intensity and polarization.

As shown in rows 1-6 of Table 5.3 in Appendix A (section 5.5), `kine` provides $\chi^2 \sim 1$ for all static models, showing that the reconstructions fit the data appropriately. The reconstructions do not present any additional motion, so we can conclude that indeed the `kine` pipeline does not introduce any motion when there is not any in the ground truth. We test the fidelity of image reconstruction by computing the cross correlation between the ground truth models and the reconstructed videos (lower rows of Figure 5.4). For all models except the point source, the reconstructions achieve cross correlation values of 0.97-1.00, significantly higher than the required thresholds (reported in parenthesis in Figure 5.4, see Dahale et al. 2025 for the threshold definition). For the point source, the total intensity cross correlation is still higher than the threshold, though slightly lower compared to other models. Indeed, from Figure 5.4, we see that the reconstruction presents a small peak with brighter emission and a halo of more diffuse emission, though the peak is misplaced compared to the ground truth. We attribute the slightly worse performance of `kine` on the point source model to the low flux density of the diffuse emission, which is at the threshold of the dynamic range that we can expect to recover with the sparse (u, v)-coverage of EHT observations in 2017 (EHTC 2022b). For polarization, the point source EVPA cross correlation is excellent, while the linear polarization one is slightly below the required threshold. We attribute this to the suboptimal total intensity reconstruction, which is needed for the reconstruction of the polarization image.

We can conclude that the image fidelity is excellent for all models in both total intensity and polarization, except for the absolute value of the point source's linear polarization, which is below the threshold. We remark that the point source reconstruction does not resemble any other model morphology, meaning that the `kine` pipeline is able to fully distinguish different morphologies and reconstruct them without adding any additional variability.

Dynamic geometric models

The second validation test required imaging simple variable geometric models, the majority of which consisting of an m-ring model (Johnson et al. 2020; Roelofs et al. 2023) plus a moving Gaussian hot spot. The goal of the test was to make sure that `kine` is able to recover qualitatively and quantitatively the simple motion present in the models. In particular, we want to ensure that the reconstructions:

1. fit the data correctly,
2. recover coherent orbital motion, with the correct direction and speed,
3. recover non-orbital motion, with the correct direction and speed,
4. recover incoherent motion without bias towards coherent motion,
5. recover dynamics in linear polarization.

We considered seven different dynamic models, which are shown in the top row of Figure 5.5, while the corresponding `kine` reconstructions are displayed in the second row and the dynamic component of the reconstructions is highlighted in the third row. The figure shows one frame of the full videos which are linked in the caption. To test the recovery of orbital motion, we used two m-ring models, with a clockwise (CW) and a counterclockwise (CCW) orbiting hot spot with a period of 80 min. To test the recovery of non-orbital motion, we used an m-ring model with a hot spot crossing the ring and an m-ring model with an off-centered orbiting hot spot. To test bias towards coherent motion, we used an m-ring model with a randomly appearing and disappearing bright spot. And finally, to test the recovery of polarimetric dynamics we employed an m-ring model with an orbiting polarized hot spot and an m-ring model with a spiral polarimetric pattern with varying EVPA pitch angle. From the video in Figure 5.5 we see that the reconstructions provide a high visual resemblance to all ground truth model, both in the full video and in its dynamic component. In the following we quantify the reconstruction fidelity, with particular focus on physically significant observables such as the ring position angle and the speed of moving components.

The goodness of fit of the reconstructions is evaluated with the time-averaged χ^2 of closure quantities in total intensity and polarization. As shown in rows 7-13 of Table 5.3 in Appendix A (section 5.5), `kine` provides $\chi^2 \sim 1$ for all models, showing that reconstructions fit the data appropriately. The quality of the reconstructions is evaluated via various metrics (full details and definitions in Dahale et al. 2025). We measure image fidelity with the normalized cross correlation between ground truth frames and reconstruction frames. The cross correlation is computed separately on the median frame ("static component") and on the residual between the median and the full video ("dynamic component"). The minimum cross correlation thresholds

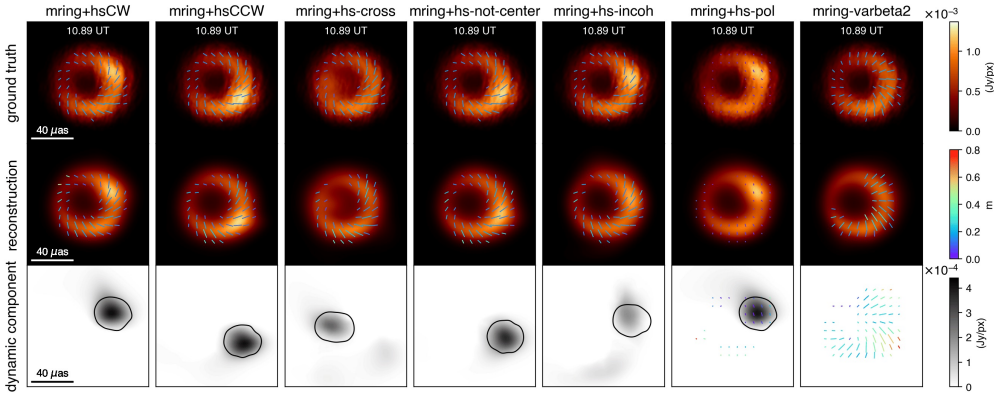

Figure 5.5: **Dynamic geometric models reconstructions.** (Top row) Ground truth dynamic geometric models. (Middle row) `kine` reconstructions. (Bottom row) Dynamic component of the `kine` reconstructions. The figure displays one frame of the full video reconstructions available at `https://github.com/mariannafoschi/phdthesis/blob/main/chapter4-video1-geometric.gif`. A selection of frames is shown for each model individually in the Appendix of the thesis.

are modulated as a function of the instantaneous (u, v)-coverage and visibility SNR. We assess the ability to recover motion by comparing the position of the hot spot, which is located by fitting a crescent + Gaussian model to each frame of the video. To evaluate the fidelity of the linear polarization we focus on the phase $\angle\beta_2$ of the $m = 2$ mode of a β_m decomposition of the polarization field (Palumbo et al. 2020), which is sensitive to the dominant orientation of swirl in the EVPA pattern. Finally, we evaluate how well we recover the speed of orbiting components by comparing the pattern speed (Conroy et al. 2023) of the ground truth and the reconstruction. We present the metrics results for all geometric models in Figure 5.6, which displays the cross correlation of the dynamic component in the left column, the position angle, the linear coordinate, or the $\angle\beta_2$ parameter in the central column, and the pattern speed autocorrelation plots in the right column.

Regarding the recovery of orbital motion, we consider the first two models with clockwise and counterclockwise orbiting models. For both models the dynamic cross correlation is above the required threshold in the vast majority of frames (100% and 85%) and the hot spot position is recovered correctly almost for the totality of frames (96% and 97%). The pattern speed is also recovered correctly, in direction and magnitude, within one sigma uncertainty (details in Dahale et al. 2025). In both cases the quality of the reconstruction is excellent, meaning that the pipeline is fully able to recover and distinguish clockwise motion from counterclockwise motion. Regarding the recovery of non orbital motion, we consider the model with the crossing hot spot and that with the off-centered orbiting hot spot. Also in this case, the dynamic cross correlation is mostly above the threshold (80% and 100%) and the hot spot position is correct for the majority of frames (85% and 100%). The recovered pattern speed of the off-centered hot spot is within one sigma from the true value, while the crossing hot spot pattern speed is two sigma apart from the 0 speed value of the ground truth. This slight deviation from a null pattern speed is caused by a minor fidelity loss in

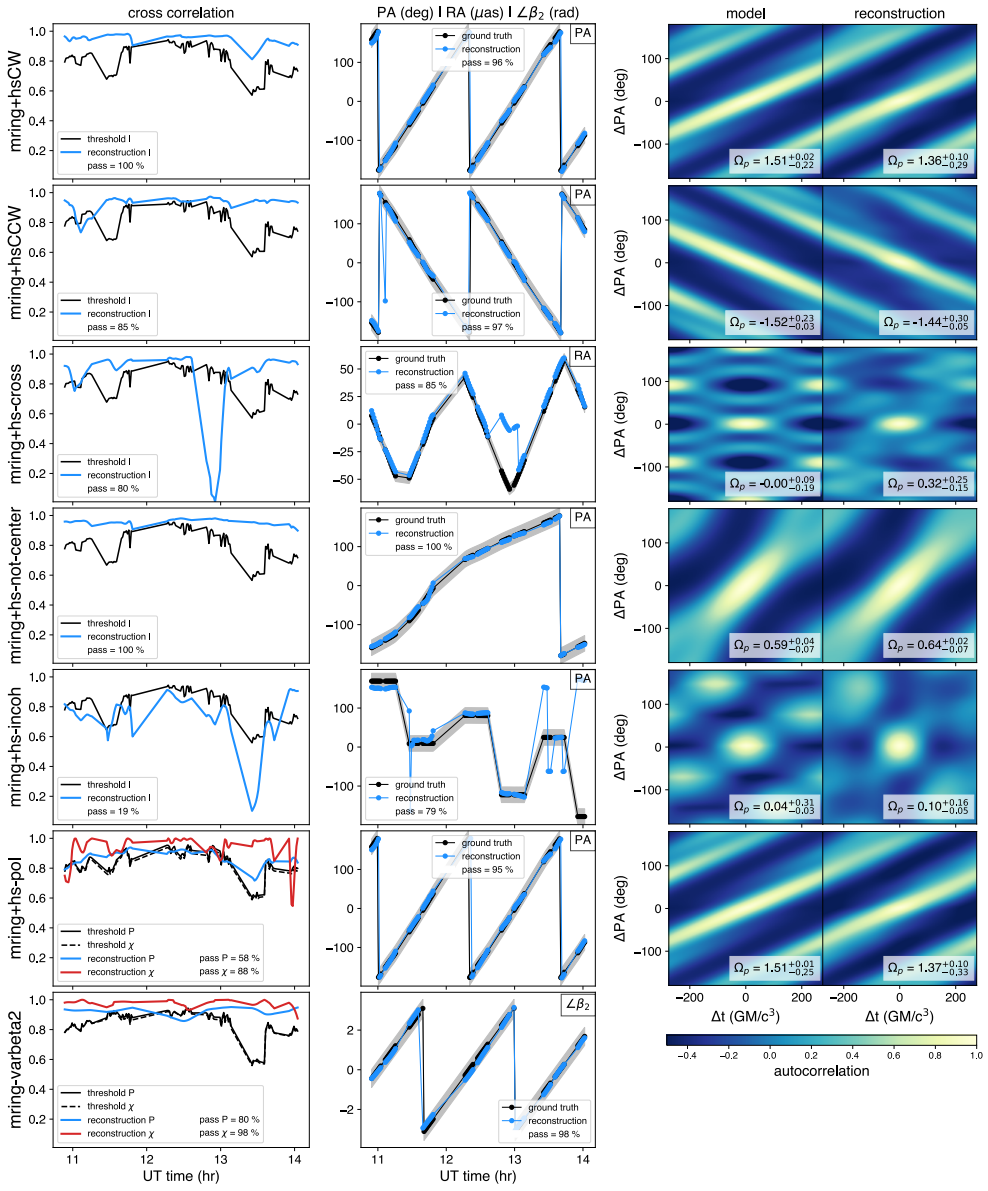

Figure 5.6: **Dynamic geometric models evaluation metrics.** (Left column) Total intensity (I) and polarimetric (P, χ) cross correlation of the dynamic component. (Center column) Position angle (PA) or right ascension (RA) coordinate of the moving Gaussian hot spot, or $\angle\beta_2$ parameter of the polarization vector field. (Right column) Pattern speed autocorrelation plots, with the model one on the left and the reconstruction one on the right. The pattern speed is reported in deg/GMc^{-3}.

the video around 12.80 UT, which is limited to half a scan and does not affect the correct recovery of the motion in other frames. Therefore we conclude that the `kine` pipeline is able to recover with good precision the correct direction and magnitude of both orbital and non orbital motion.

Regarding the recovery of incoherent motion, the dynamic cross correlation of the incoherent hot spot model is slightly below the threshold for many frames (19% passing percentage). However, the position of the hot spot is recovered correctly for the majority of frames (79%) and the recovered pattern speed is 0, in agreement with the ground truth. From this, we can conclude that the pipeline does not hallucinate coherent motion in the presence of incoherent variability, so there is no bias towards either coherent or incoherent motion. The pipeline is also able to track the position of incoherent brightness changes correctly, even though the total flux attributed to local brightness increases might be slightly underestimated.

The recovery of polarization variability is tested for a polarized orbiting hot spot and for a polarization field with changing orientation. In this case, we considered the cross correlation of the linear polarization image $P = \sqrt{Q^2 + U^2}$ and the EVPA image $\chi = \frac{1}{2}\arctan(U/Q)$, instead of the total intensity I image. To quantify the fidelity of the recovered variable features, we considered the position angle for the first model and the variation of the $\angle\beta_2$ parameter for the second one. For the polarized hot spot, the EVPA and polarization intensity cross correlations are higher than the threshold on the majority of cases (58% and 88%). The lower cross correlation in P is likely due to a very low polarization signal outside the hot spot region, combined with a suboptimal recovery of the polarization intensity distribution. However, we note that the orientation of the EVPA pattern is well recovered both inside and outside the hot spot. Regarding the polarization model with changing orientation, the P and EVPA cross correlations are above the thresholds for 80% and 98% of the frames. The position angle and the $\angle\beta_2$ parameter are also recovered correctly for almost all frames (95% and 98%). From these tests we can conclude that `kine` is fully able to recover both local and global changes in the linear polarization field intensity and orientation.

GRMHD simulation models

The last validation test consists in imaging synthetic data obtained from ray-traced videos of GRMHD simulations from the standard `PATOKA` GRMHD library (Wong et al. 2022; Dhruv et al. 2025). These models should present higher complexity and variability than the geometric models and hence provide a more difficult and realistic test case. This final test aims at proving that the `kine` pipeline can provide good reconstructions of complex realistic black hole videos. Given the limited resolution of the EHT array in 2017, we do not expect to recover the fine details of the simulations, but rather we aim to reach high enough fidelity to extract relevant physical information from the reconstruction. Specifically, we want to ensure that the reconstructions:

1. Fit the data correctly,
2. Recover the correct morphology and orientation for Stokes I,
3. Recover the correct morphology, orientation, and dynamics for Stokes Q, U,
4. Recover coherent orbital motion with the correct direction,
5. Recover coherent orbital motion with the correct speed.

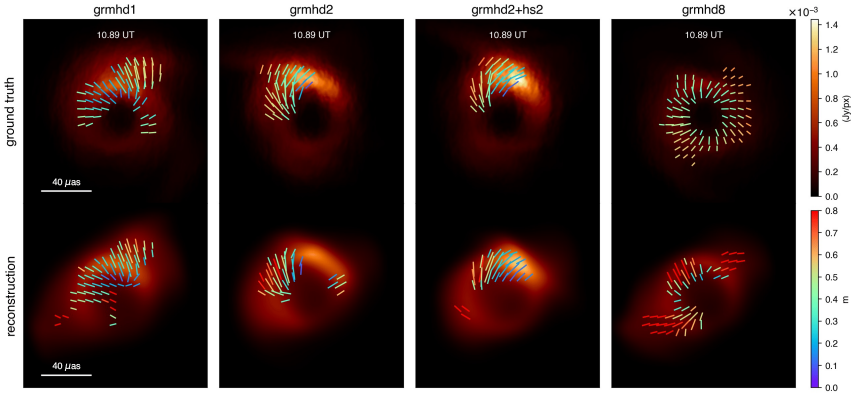

Figure 5.7: **GRMHD models reconstructions.** (Top row) Ground truth GRMHD models. (Bottom row) Corresponding `kine` reconstructions. The figure displays one frame of the full video reconstructions available at `https://github.com/mariannafoschi/phdthesis/blob/main/chapter4-video2-grmhd.gif`. A selection of frames is shown for each model individually in the Appendix of the thesis.

Out of the 360 different simulations in the GRMHD library, we selected a sample of 8 models, designed to span all unique simulation parameter values at least once, with an emphasis on models favored by the data. The parameters characterizing each simulation are the magnetic flux mode (MAD/SANE), the dimensionless black hole spin, the viewing inclination, and the electron temperature parameter. The parameter combinations corresponding to the selected models are listed in Dahale et al. (2025), together with more details about the simulations. We selected 3 GRMHD models (`grmhd1`, `grmhd2`, `grmhd8`) for the main validation tests. Being all MAD models, with prograde spin, high electron temperature, and relatively face-on inclination, the selected models all fall in the region of the simulation parameter space that is favored by Sgr A* observational constraints (EHTC 2022e; EHTC 2024). This choice was made to ensure that the pipeline was able to correctly recover models that behave similarly to what is expected of Sgr A*. To these models we added a variation of `grmhd2` which included an additional geometric orbiting hot spot. Ideally, the imaging pipeline should also be flexible enough to recover morphologies and behaviors that are unlikely for Sgr A*. This is considered as an extra test that is not required to fulfill the validation. Its results are presented in section 5.3.2, and include testing on the `grmhd3`, `grmhd4`, `grmhd5`, `grmhd6`, and `grmhd7` simulations, which include combinations of retrograde, SANE, edge-on, low electron temperature, and zero-spin models.

The ground truth models are shown in the top row of Figure 5.7, with the corresponding `kine` reconstructions in the second row. The figure shows one frame of the full videos which are linked in the caption. The quality of the reconstruction fit to the data is assessed by the χ^2s of closure quantities, which are reported in rows 14-17 of Table 5.3 in Appendix A (section 5.5). The total intensity and polarimetric cross correlation, the position angle, the $\angle\beta_2$ parameter, and the pattern speed metrics are shown in Figure 5.8.

From the video in Figure 5.7, we see that all reconstructions recover variable asymmetric rings, similar to the ground truth models. This is reflected by the high values

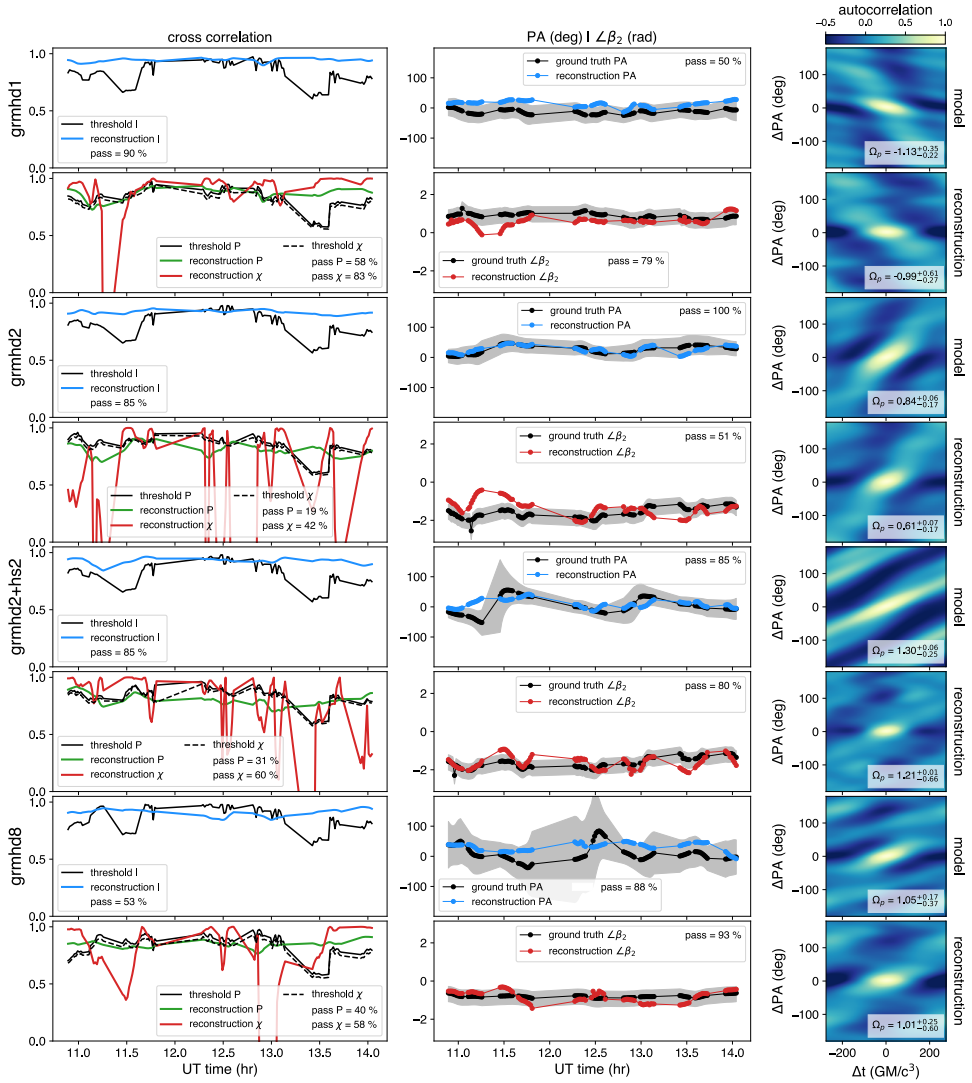

Figure 5.8: **GRMHD models evaluation metrics.** (Left column) Total intensity cross correlation (I, in blue) and polarimetric cross correlation (P in green, χ in red) between reconstruction and ground truth. (Center column) Position angle (PA) of the ring and $\angle\beta_2$ parameter of the polarization vector field. (Right column) Pattern speed autocorrelation plots, with the model one on top and the reconstruction one on the bottom. The pattern speed is reported in deg/GMc^{-3}.

of the total intensity cross correlation, which is above the threshold for the majority of frames for all models (90%, 85%, 85%, 53%). In `grmhd1`, the position angle is recovered within 1 sigma of the true one for half of the frames and within 2 sigmas for all frames. For the other models, the position angle is recovered within 1 sigma for the great majority of frames (100%, 85%, and 88%). The position angle errorbars vary from model to model depending on the model's asymmetry. In fact, for models with low asymmetry, the peak position angle is widespread and therefore less constrained, while in the case of higher asymmetry the PA is more localized. We observe that the position angle in `grmhd1` is recovered with a slight positive bias. We attribute this to the widespread position angle of the model, combined with the presence of an extended arm. However we do not consider this to be statistically relevant because we don't observe the same bias in other models and because the recovered PA is still within the errorbars. Based on these results we can state that the pipeline is able to recover the correct morphology and orientation of the GRMHD models under consideration. The recovered pattern speed is in agreement with the true one within 1 sigma for all the considered GRMHD models, meaning that the pipeline is able to recover both the direction and the magnitude of the velocity of features and patterns orbiting around the black hole. In particular, considering `grmhd2` and `grmhd2+hs2`, we see that the reconstructions successfully recover different pattern speeds showing that coherent orbital motion can be distinguished from the filamentary orbiting patterns typical of GRMHD simulations.

Regarding polarimetric reconstruction, for all models except `grmhd2`, the EVPA cross correlation is above the threshold for the majority of frames, with occasional drops. The cross correlation of the linear polarization magnitude, instead, oscillates around the threshold, and for most models it is often slightly lower than the threshold. This results in a percentage of passing frames for P and χ of (58%, 83%) for `grmhd1`, (19%, 42%) for `grmhd2`, (31%, 60%) for `grmhd2+hs2`, (40%, 58%) for `grmhd8`. The phase of the β_2 parameter, which indicates the radial pitch angle of the EVPA, is well recovered for all models within the errorbars, for the majority of frames (79%, 51%, 80%, and 93%). From this we can conclude that the `kine` pipeline is able to recover correctly the orientation of the polarization vector field, but the amount of recovered linear polarization is at times incorrect.

5.3.2 Extra tests

In addition to the validation tests presented until now, we applied `kine` to other sets of geometric and GRMHD models, to further test the range of motion and morphologies that can be reliably reconstructed. Contrary to the previous ones, passing these tests is not considered a requirement for the validation of the pipeline, since they are aimed at exploring the limitations of the imaging method.

Hot spot speed and flux ratio tests

Here we present `kine` reconstructions of five m-ring plus hot spot models. The first three aim to test `kine`'s ability to recover an orbiting hot spot under different flux ratios between the ring and the hot spot. Specifically we tested models with ring/hot-spot total flux ratios of 17.0, 3.5, and 1.25, in addition to the 8.0 ratio of the models in the validation test. The last two models aim to test the maximum speed of moving features that the pipeline is able to recover. The models consist of an m-ring plus a hot spot orbiting with periods of 40 min and 20 min, respectively.

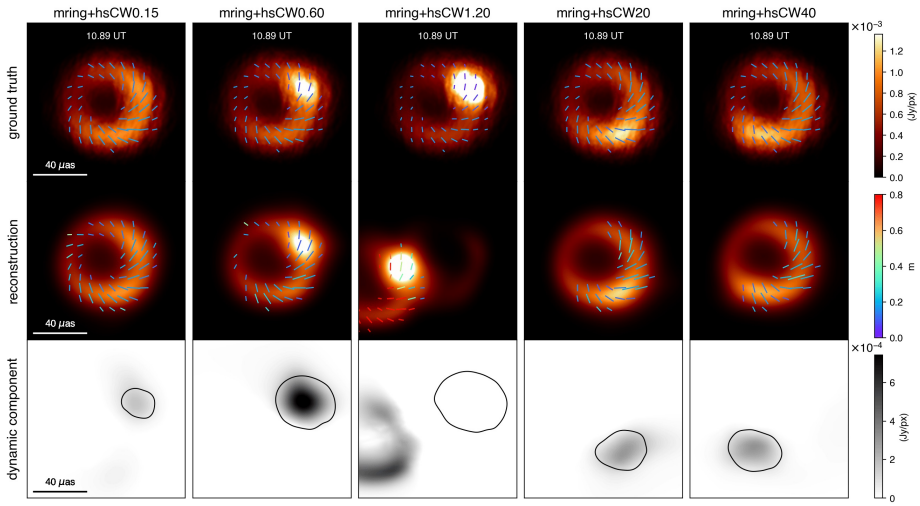

Figure 5.9: **Dynamic geometric models reconstructions, speed and flux tests.** (Top row) Ground truth models. (Middle row) `kine` reconstructions. (Bottom row) Dynamic component of the `kine` reconstructions. The figure displays one frame of the full video reconstructions available at https://github.com/mariannafoschi/phdthesis/blob/main/chapter4-video3-geometricextra.gif. A selection of frames is shown for each model individually in the Appendix of the thesis.

The ground truth models are shown in the top row of Figure 5.9, with the corresponding `kine` reconstructions in the second row and the dynamic component of the reconstructions in the third row. The figure shows one frame of the full videos which are linked in the caption. The quality of the reconstruction fit to the data is assessed by the χ^2s of closure quantities, which are reported in rows 18-22 of Table 5.3 in Appendix A (section 5.5). The cross correlation, the position angle and the pattern speed metrics are shown in Figure 5.10.

Regarding the flux ratio test, the pipeline is able to correctly recover the motion of the hot spot for models with flux ratios of 17.0 or 3.5, as demonstrated by a total intensity dynamic cross correlation above the threshold for 78% and 98% of the frames and the position angle correctly recovered in 87% and 93% of the frames. For both models, the pattern speed of the reconstruction also matches the true value within 1 sigma. However, for the flux ratio of 1.25 `kine` is not able to recover the correct morphology nor dynamics of the model. We believe that this is due to the very low flux density of the ring, which is just slightly above the detection limit set by the expected dynamic range. Indeed `kine` correctly recovers the bright hot spot, as proven by the high total intensity cross correlation (94% frames above threshold), but seems to be "blind" to the ring, allocating the remaining flux incoherently in the frame. This motivation would also explain why the bright spot remains fairly static in the reconstruction, instead of orbiting. Indeed, imaging the data with only closure quantities does not allow to constrain the absolute position of the source and, without the ring acting as an anchoring feature, the bright spot alone is imaged in a fixed position.

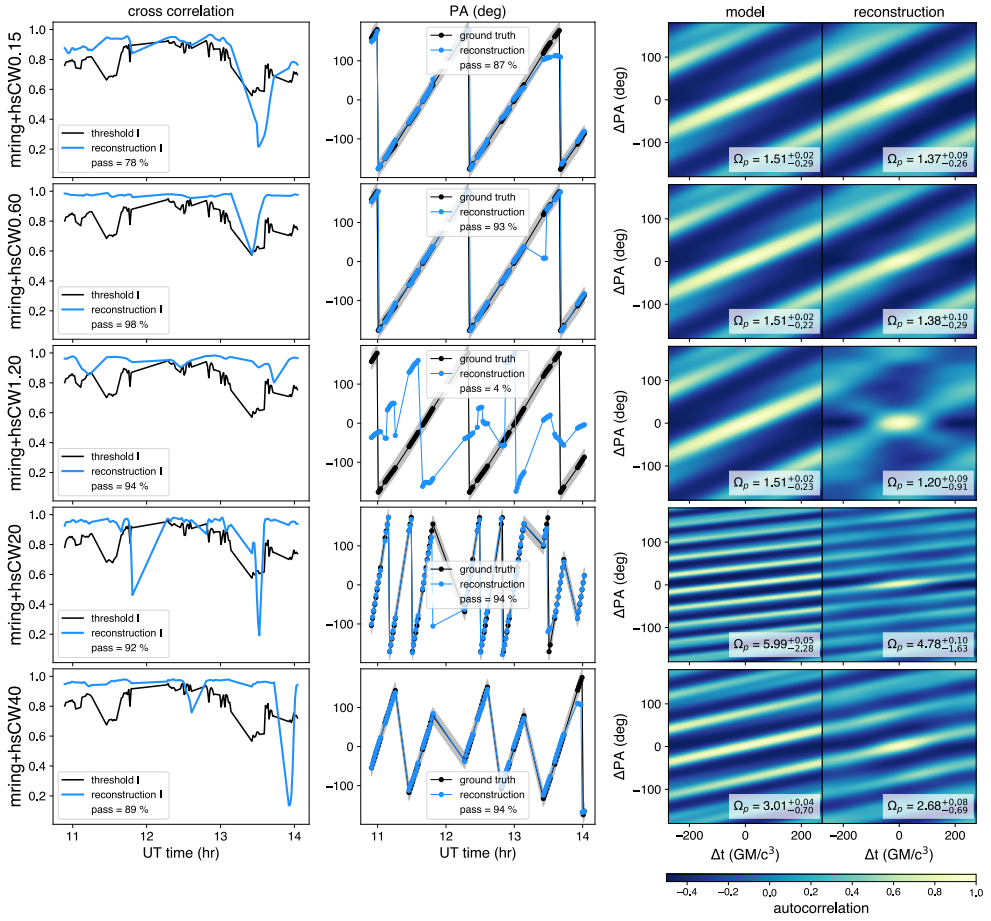

Figure 5.10: **Speed and flux tests with extra geometric models.** (Left column) Total intensity cross correlation of the dynamic component. (Center column) Position angle (PA) of the moving Gaussian hot spot. (Right column) Pattern speed autocorrelation plots, with the model one on the left and the reconstruction one on the right. The pattern speed is reported in deg/GMc^{-3}.

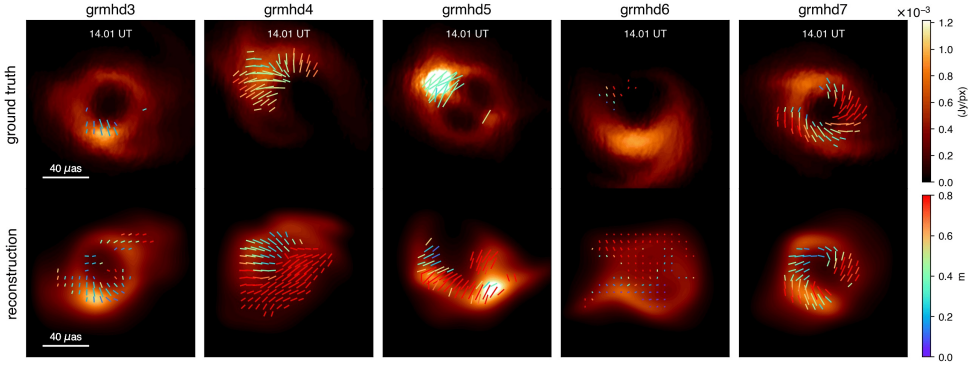

Figure 5.11: **Extra GRMHD models reconstructions.** (Top row) Ground truth GRMHD models. (Bottom row) Corresponding `kine` reconstructions. The figure displays one frame of the full video reconstructions available at `https://github.com/mariannafoschi/phdthesis/blob/main/chapter4-video4-grmhdextra.gif`. A selection of frames is shown for each model individually in the Appendix of the thesis.

Regarding the speed test, `kine` proves to be particularly effective in recovering features moving at high speeds. Indeed, for both models, the total intensity cross correlation is above the threshold in almost all frames (92% and 89%) and the orbiting Gaussian is correctly located in the (quasi) totality of frames (94% and 94%). The pattern speed is also estimated correctly within errorbars. Therefore we can conclude that `kine` can recover features moving with orbiting period down to at least 20 min, meaning that the maximum velocity that can be recovered is not limited by the pipeline, but rather by the frequency of observation sampling and the lengths of inter-scan gaps. In fact, the 2017 EHT observations of Sgr A* consist of multiple ∼10 min scans spaced with gaps ranging from ∼3 min to ∼30 min to observe the calibrator sources.

Extra GRMHD tests

In this section we present `kine` reconstructions of five extra GRMHD simulation videos. In particular, here we want to determine the pipeline's ability to recover GRMHD simulations that depart from what is expected from Sgr A*. Except for `grmhd7` model, the selected simulations consists of parameter combinations that are disfavored by observational constrains of Sgr A* (EHTC 2022e; EHTC 2024). As further detailed in Dahale et al. (2025), the selected simulations include models: `grmhd3`, which is retrograde, `grmhd4`, that is SANE, retrograde, nearly edge-on, with low electron temperature, `grmhd5`, which is fully edge-on, `grmhd6`, which is SANE with zero spin and `grmhd7`, which instead has parameters compatible with Sgr A* observations.

The ground truth models are shown in the top row of Figure 5.11, with the corresponding `kine` reconstructions in the second row. The figure shows one frame of the full videos which are linked in the caption. The quality of the reconstructions fit to the data is assessed by the χ^2s of closure quantities, which are reported in rows 23-27 of Table 5.3 in Appendix A (section 5.5). The total intensity and polarimetric cross correlation, the position angle, $\angle\beta_2$ parameter, and the pattern speed metrics are shown in Figure 5.12.

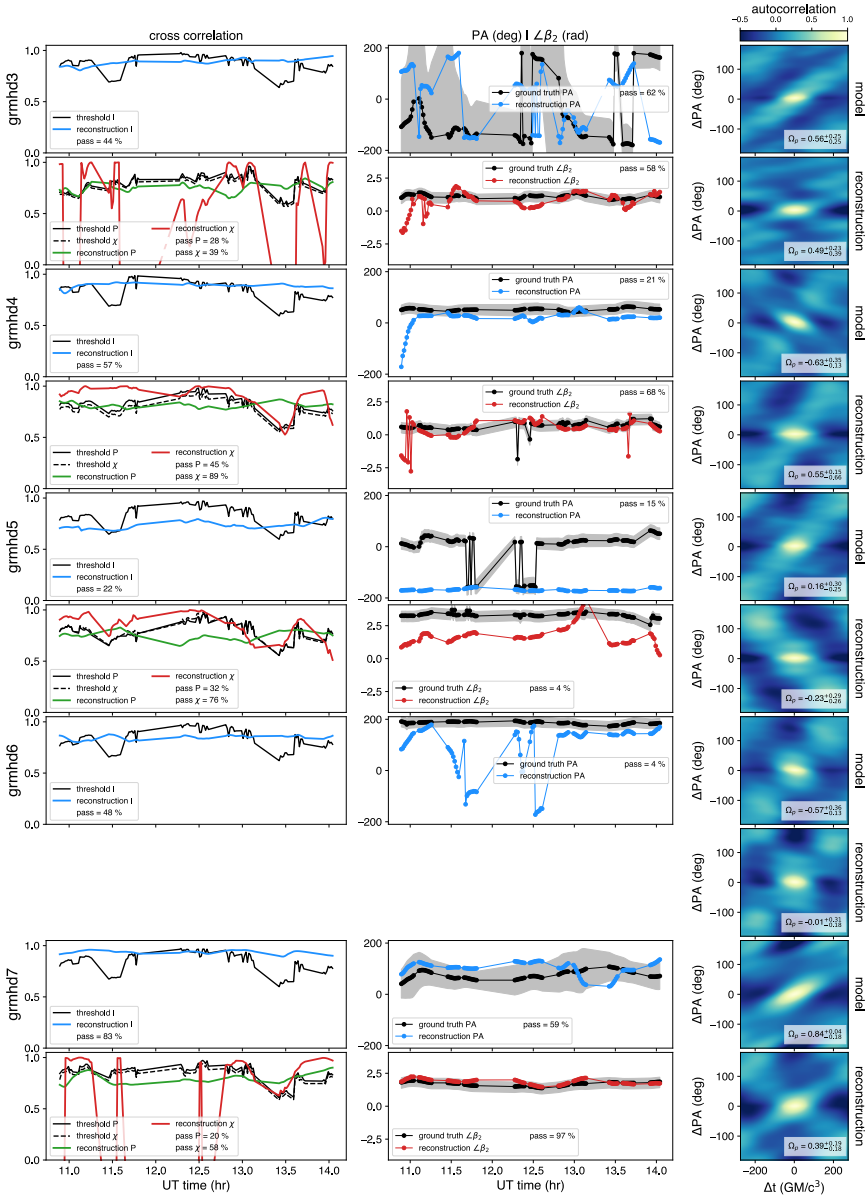

Figure 5.12: **Extra GRMHD models evaluation metrics.** (Left column) Total intensity cross correlation (I, in blue) and polarimetric cross correlation (P in green, χ in red) between reconstruction and ground truth. (Center column) Position angle (PA) of the ring and $\angle\beta_2$ parameter of the polarization vector field. (Right column) Pattern speed autocorrelation plots, with the model one on top and the reconstruction one on the bottom. The pattern speed is reported in deg/GMc^{-3}. For `grmhd6` the polarization metrics are not reported because the polarization signal in the model is lower than the noise level.

From a visual inspection of the video in Figure 5.11, it is clear that reconstructing these models is more difficult compared to almost face-on MAD models with prograde spin and a northern position angle. We examine the reconstruction of each model individually. Regarding `grmhd3`, the reconstruction recovers correctly a ring with low asymmetry, but the total intensity cross correlation is above the threshold in slightly less than half of the frames. The polarization signal is low, and this make its recovery difficult (polarimetric cross correlation above the threshold for 28% and 39% of frames), though global features like $\angle\beta_2$ are well recovered within the errorbars for the majority of frames (58%). In this model, since the asymmetry is minimal, the position angle is difficult to identify clearly even in the ground truth. This inflates the PA error bars and makes the recovered angle compatible with the true one on the majority of frames. The pattern speed is recovered with the right direction and with magnitude within 1 sigma from the true value. We conclude that in the case of models with low ring asymmetry, `kine` is able to correctly measure the direction and speed of orbital motion, as well as the orientation of the polarization field, but it is not able to localize the peak brightness position angle. For the `grmhd4` and `grmhd5` models, the pipeline was not able to recover the basic features of the underlying ground truth. We believe that in the case of `grmh4` this was due to the significant amount of flux extending outside the ring, paired with a high asymmetry, that resulted in portions of the ring barely visible above the nose level. However, `kine` is able to attribute the majority of the flux to the right position, even if not with the correct morphology. In the case of `grmhd5`, instead, we attribute the failure to the high luminosity contrast between the brightest spot in the upper left quadrangle and the ring and diagonal features, which in comparison appear dim. We believe that that our pipeline struggles to detect the ring, and, as a consequence, anchor the other elements in the video. However, for `grmhd5`, `kine` is able to recover a pattern speed compatible with the ground truth null value. We also observe that for `grmhd4` the orientation of the linear polarization is well recovered, as indicated by the high EVPA cross correlation (89% of frames above the threshold) and the good estimate of $\angle\beta_2$ (68% of frames are compatible within 1 sigma). We attribute this success to the high polarization signal present in the model, which resulted in a high polarimetric SNR. In the case of `grmhd6`, the pipeline is able to recover a variable ring structure. However, similarly to `grmhd4`, the high ring asymmetry makes it hard for `kine` to recover the right position angle, resulting in low passing percentage of the total intensity cross correlation and the position angle. In this model, the linearly polarized emission is too low to recover a meaningful signal so the polarimetric metrics are not taken into consideration.

Finally, for `grmhd7`, which has simulation parameters favored by Sgr A* observations, the position angle and the $\angle\beta_2$ parameter are recovered correctly for the majority of frames (59% and 97%) as well as the total intensity structure and EVPA orientation (cross correlations above the threshold for 83% and 58% of the frames), while the magnitude of the linear polarization is correct only for 20% of the frames. The estimated value of the pattern speed is at the border of 1 sigma uncertainty with respect to the true value. Taking into consideration these metrics, we consider that the reconstruction of the `grmhd7` model is successful. Overall, considering all GRMHD models, the direction of the pattern speed and the phase of the β_2 parameter are recovered correctly in all cases except one, making them the most robustly recovered parameters across the GRMHD models.

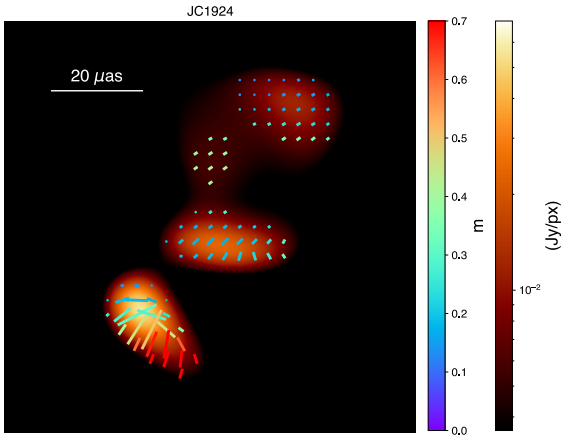

Figure 5.13: **`kine` reconstruction of J1924-2914.** One representative frame from the video reconstruction.

Imaging calibrator sources

As a further test of the ability of `kine` to reconstruct also static sources with a jet-like structure, without introducing extra motion, we applied the pipeline described in this work to real EHT observations of calibrator J1924–2914 on April 11 2017, which has similar coverage of Sgr A* on that day. Testing on real observations is important since it cannot be excluded that real data may contain unknown issues or properties not represented in synthetic data. The J1924–2914 radio source does not present intra-day variability and has been already imaged in both total intensity and polarization by Issaoun et al. (2022), so it represent a good validation dataset. The `kine` video reconstruction of the source, of which we show one frame in Figure 5.13, does not present significant variability. The morphology of the inner jet consists of three distinct components of decreasing brightness, aligned in the southeast-northwest direction. Regarding linear polarization, we observe an EVPA structure mainly oriented along the jet direction, with a slight fan-like structure in the brightest component, and a fractional polarization up to 0.7. This matches the reconstructions of April 11 data provided by multiple pipelines in Issaoun et al. (2022). The correct imaging of J1924–2914 is an additional proof of `kine`'s ability to distinguish correctly between static and dynamic sources, and between ring-like and non-ring-like brightness distributions.

5.4 Conclusions

In this work we extended and tested the `kine` dynamic imaging algorithm, with a specific pipeline developed for 2017 EHT observations of Sgr A*, on a wide range of synthetic data with real coverage and realistic noise conditions. The synthetic data were generated from an extensive suite of ground truth videos of static geometric models, dynamic geometric models, and ray-traced GRMHD simulations. The overall performance of `kine` is remarkably good, with a few isolated failures. We were able to pinpoint the reasons for the few incorrect reconstructions, and since they do not

apply to the characteristics of real Sgr A* data, we believe that the performance of the imaging algorithm on real data should not be hindered. In this section we provide a review of the validation tests and discuss the problematic cases. Table 5.1 displays a summary of all the metrics used to evaluate the fidelity of the reconstructions, stating for each model if the tests proved successful. In Table 5.2 we summarize the answers to the questions asked in the pipeline validation process.

Regarding the reconstruction of the static models, we recovered the correct morphology without additional motion for all models. The polarization field structure is also well recovered for all models except for the absolute value of the linear polarization of the point source. We believe that this is due to the low value of $|P|$ in the ground truth model, combined with a total intensity reconstruction that, while passing the cross correlation test, is suboptimal compared to the other models. Overall, we consider that, for static models, `kine` is able to recover the total intensity and polarimetric ground truth model without hallucinating motion in the video.

From the tests on the geometric dynamic models, we see that `kine` recovers both orbital and non-orbital motion with the correct direction and magnitude for features moving with orbital periods from 180 min up to 20 min. In the model with incoherent motion, `kine` recovers the correct position of the bright spot, but the dynamic total intensity cross correlation is lower than the threshold for the majority of frames. We conclude that the method can track incoherent motion, but not recover the full flux associated with it. Nevertheless, we see that the recovered pattern speed is compatible with zero, as expected from incoherent motion, meaning that the pipeline can discriminate between incoherent and coherent motion without introducing any bias in favor of the latter. The dynamic geometric models with varying polarization are also recovered correctly by `kine`, as proven by the high polarimetric cross correlation and the reconstructed $\angle\beta_2$ parameter, meaning that the pipeline can recover a simple varying polarization field. From the flux-ratio test, we found that `kine` can reconstruct the correct model when the ratio between the flux of the ring and that of the hot spot is greater than 1.25. This is because, below this value, the flux of the ring is too dim to be constrained effectively and therefore it cannot act as an anchor to the moving Gaussian.

From the pipeline's results over the more complicated GRMHD models, we see that, in the case of MAD models with prograde spin, high electron temperature, and relatively face-on inclination, as is expected for Sgr A* from EHT observations (EHTC 2022e; EHTC 2024), the pipeline recovers correctly the ground truth videos, as shown by high cross correlation values and the correct tracking of the position angle. The pattern speed is always recovered with the correct direction and magnitude within the errorbars. The orientation of the polarization vector field is also recovered correctly as shown by the good recovery of the $\angle\beta_2$ parameter and the high cross correlation for the EVPA. However the magnitude of the linear polarization is not recovered sufficiently well. Regarding polarimetric reconstructions of GRMHDs we consider `kine` to partially passing the validation tests, meaning that the pipeline can provide reliable reconstructions for the orientation of the linear polarization but not for its magnitude. Overall this shows that even in the case of a more complex and fast-varying ground truth, `kine` can recover the important physical features of the model, such as the total intensity structure, the orientation of the linear polarization, and the direction and magnitude of the orbital motion. Regarding the tests on extra

Models		χ^2_I	χ^2_P	$\eta_{\text{xcorr I}}$	$\eta_{\text{xcorr }\chi}$	$\eta_{\text{xcorr }P}$	σ_{vis}	PA	$\angle\beta_2$	Ω_p
static	`crescent`	✓	✓	✓	✓	✓	✓	-	-	-
	`disk`	✓	✓	✓	✓	✓	✓	-	-	-
	`double`	✓	✓	✓	✓	✓	✓	-	-	-
	`edisk`	✓	✓	✓	✓	✓	✓	-	-	-
	`point`	✓	✓	✓	✓	✗	✓	-	-	-
	`ring`	✓	✓	✓	✓	✓	✓	-	-	-
geometric	`mring+hsCW`	✓	✓	✓	-	-	-	✓	-	✓
	`mring+hsCCW`	✓	✓	✓	-	-	-	✓	-	✓
	`mring+hs-cross`	✓	✓	✓	-	-	-	✓	-	✗*
	`mring+hs-incoh`	✓	✓	✗	-	-	-	✓	-	✓
	`mring+hs-not-center`	✓	✓	✓	-	-	-	✓	-	✓
	`mring+hs-pol`	✓	✓	-	✓	✓	-	✓	-	✓
	`mring-varbeta2`	✓	✓	-	✓	✓	-	-	✓	-
GRMHD	`grmhd1`	✓	✓	✓	✓	✓	-	✓	✓	✓
	`grmhd2`	✓	✗	✓	✗	✗	-	✓	✓	✓
	`grmhd2+hs2`	✓	✗	✓	✓	✗	-	✓	✓	✓
	`grmhd8`	✓	✓	✓	✓	✗	-	✓	✓	✓
speed limit	`mring+hsCW20`	✓	✗	✓	-	-	-	✓	-	✓
	`mring+hsCW40`	✓	✓	✓	-	-	-	✓	-	✓
flux-ratio limit	`mring+hsCW0.15`	✓	✓	✓	-	-	-	✓	-	✓
	`mring+hsCW0.60`	✓	✓	✓	-	-	-	✓	-	✓
	`mring+hsCW1.20`	✓	✓	✓	-	-	-	✗	-	✓
extra GRMHD	`grmhd3`	✓	✓	✗	✗	✗	-	✓	✓	✓
	`grmhd4`	✓	✗	✓	✓	✗	-	✗	✓	✗*
	`grmhd5`	✓	✗	✗	✓	✗	-	✗	✗	✓
	`grmhd6`	✓	✓	✗	✓	✓	-	✗	-	✗*
	`grmhd7`	✓	✓	✓	✓	✗	-	✓	✓	✗*

Table 5.1: **Reconstruction metrics pass/fail.** Summary of the tests passed and failed by the `kine` reconstructions of the ground truth videos. The asterisks indicates pattern speed estimates that are incompatible with the true value within 1 sigma, but are compatible within 2 sigmas.

GRMHD simulations, `kine` sometimes struggles to recover the correct morphology and dynamics of the ground truth model. In particular, the imaging pipeline fails in the cases of edge-on or almost edge-on models, which are characterized by high asymmetry and the presence of a localized spot significantly brighter than the rest of the emission. However, a scenario in which Sgr A* is edge-on is highly disfavored by EHT and non-EHT observational constrains (EHTC 2022e; EHTC 2024), so we believe that the worse performance of the pipeline on these models will not impact the robustness of the reconstruction of real Sgr A* data. Nevertheless, the direction of the pattern speed and the phase of the β_2 parameter were recovered correctly for all GRMHD models except one, meaning that the estimate of these quantities is robust under any GRMHD scenario. Finally we tested the pipeline on real observations of the radio source J1924-2914 by the EHT on April 11$^{\text{th}}$ 2017, recovering a static polarized image and amplitude gains in agreement with previous results by the EHT (Issaoun et al. 2022).

The quality of the video reconstructions and the extensiveness of the validation tests presented in this work represent a significant improvement in comparison to the first dynamic imaging attempts presented in EHTC (2022c). Most of the validation in EHTC (2022c) revolved around the impact of data pre-processing, specially that concerning subtraction of interstellar scattering, and prior assumptions required by the `StarWarps` pipeline, the only imaging algorithm employed. With respect to data descattering, all results presented here are on-sky data, so all the tests associated to descattering are unnecessary. At the same time, we have proven that with an unified pipeline, the hyperparameters associated to the `kine` pipeline are able to reproduce with high fidelity significantly different synthetic data, with fast, slow, and no variability at all. In EHTC (2022c), we considered GRMHD simulations to be successfully recovered if the overall position angle of the video reconstruction matched within some uncertainty that of the ground-truth model. This resulted in two out of three test cases passed. Here, not only we recover in most cases the position angle of a more extensive suit of GRMHD models, but also reconstruct high-fidelity videos with subtle features such as spiral arms, highly-consistent linear polarization structures, and, in most cases, the correct magnitude and orientation of the ground-truth pattern speed.

The reason why `kine` performs better than other imaging methods, lies in the flexibility of the pipeline's priors and the absence of bias towards specific morphologies or amounts of variability. Thanks to the neural field representation at the core of the imaging algorithm (Foschi et al. 2025b), `kine` enforces space-time regularization in the output video not through explicit, arbitrary regularizer terms, but thanks to the implicit preference of the network for fitting smoother functions (a property known as *spectral bias*, Rahaman et al. 2019). This characteristic makes the pipeline highly expressible and at the same time less biased. Spectral bias is especially important because it allows `kine` to not require a hyperparameter encoding the amount of temporal regularization, since, if the training of the reconstruction is stopped before overfitting, spectral bias, combined with positional encoding, will ensures that the reconstruction has the correct amount of space and time regularization. Additionally, we designed the pipeline for EHT observations of Sgr A* to be as agnostic as possible. This was achieved by building a unified pipeline, with the same set of hyperparameters, to be applied to all synthetic data, by performing the first rounds of imaging with closure quantities only, so to be independent from calibration errors, and by using a simple

Validation tests	static models	geometric models	GRMHD models
Fit the data correctly.	✓	✓	✓
Do not introduce additional dynamics when none are present.	✓	-	-
Recover the correct morphology and orientation for Stokes I	✓	✓	✓
Recover the correct morphology and orientation for Stokes Q, U	✓	✓	✓✗
Recover coherent orbital motion with the correct direction	-	✓	✓
Recover coherent orbital motion with the correct speed	-	✓	✓
Recover non-orbital motion, with the correct direction and speed	-	✓	-
Recover incoherent motion without bias towards coherent motion	-	✓	-
Recover the correct dynamics for Stokes Q, U	-	✓	✓✗

Table 5.2: **Pipeline validation tests.** Summary of the validation tests requirements, satisfied by each set of synthetic data models. The cells in yellow refer to a pass for the orientation of the linear polarization field and a fail for its magnitude.

disk image as initialization, a shape from which the pipeline was easily able to deviate during optimization.

Because of all these considerations, we judge the pipeline to have passed the validation tests, meaning that it can be considered reliable for the reconstruction of the first horizon-scale video of a super-massive black hole, using EHT observations of Sgr A* on April 11 2017. The video reconstructed by `kine` from real Sgr A* data is presented in EHTC (2025), together with a brief analysis of the features and motion displayed in the video.

5.5 Appendix A. Fit quality

We discuss here the quality of fit of the `kine` reconstructions to the synthetic data. For every model, we computed the time-averaged χ^2 of individual data products D.

$$\chi^2_D = \frac{1}{N_t} \sum_{j=1}^{N_t} \frac{1}{k_D N_{D,j}} \sum_{i}^{N_{D,j}} \frac{\left(D_{ij} - \hat{D}_{ij}\right)^2}{\sigma^2_{D,ij}}, \tag{5.12}$$

where the index j runs over all the observed times, from 1 to the total number of observed times N_t, the index i runs over all data corresponding to the observed time t_j, from 1 to the total number of data $N_{D,j}$, and k_D is a normalization factor that takes into account the data product's degrees of freedom. Since the imaging procedure involves amplitude gain fitting and phase self-calibration steps, we chose to consider only the χ^2 of closure quantities, specifically log closure amplitudes, closure phases, and $\breve{m}$, which are not affected by gain errors and calibration. Table 5.3 presents χ^2s for all models. The table shows that all reconstructions provide a good fit to the data, since almost all χ^2s are of the order of 1. The only exceptions are the polarization χ^2s of the mring+hsCW1.20 model, which we fail to recover, and GRMHD models 2, 4, and 6, the latter of which has no significant polarization signal. Nevertheless, the χ^2 values not of order 1 are still lower than 4.

5.6 Appendix B. Gain fitting

We demonstrate our ability to recover visibility amplitude gains disentangled from intrinsic variability in Figure 5.14. In the left panel, we show the ground-truth gains applied to the synthetic data presented in the next sections and those recovered by `kine`. Specifically, these corresponds to the `mring+hsCW` model (see section 5.3.1). In the right panel, we show the amplitude gains derived from the images that different pipelines reconstructed for the jet J1924-2914, presented in Issaoun et al. (2022), and the gains recovered by `kine` during the simultaneous image and gain fitting process. The final image from `kine` can be seen in Figure 5.13. In both cases, for the majority of scans, the recovered amplitude gains are in agreement with the ground-truth gains in the case of Sgr A* synthetic data, and with other established pipelines in the case of J1924-2914. We observe that on a few scans the gain estimates provided by `kine` don't match the ground truth ones. On average, gains recovered by `kine` differ a factor of $(1 \pm 1)\%$ for ALMA, APEX, LMT, and SMA; a factor of $(2 \pm 1)\%$ for SMT and SPT; and a factor of $(3 \pm 2)\%$ for JCMT. However, even when the fitted gains don't match the true ones on all the scans, fitting gains simultaneously in the imaging process makes the reconstructed video closer to the ground truth. Therefore we incorporate

	model	χ^2_{logcamp}	χ^2_{cphase}	χ^2_{mbreve}
1	`crescent`	1.16	1.40	1.07
2	`disk`	0.70	1.32	1.01
3	`double`	0.76	1.41	1.06
4	`edisk`	0.66	1.40	1.98
5	`point`	1.12	1.43	1.60
6	`ring`	1.22	1.31	1.10
7	`mring+hsCW`	1.09	1.10	1.31
8	`mring+hsCCW`	0.81	1.11	1.27
9	`mring+hs-cross`	0.95	1.02	1.38
10	`mring+hs-incoh`	1.27	1.43	1.36
11	`mring+hs-not-center`	1.06	1.00	1.32
12	`mring+hs-pol`	1.22	1.18	1.19
13	`mring-varbeta2`	1.18	1.33	1.17
14	`grmhd1`	0.65	0.92	1.69
15	`grmhd2`	0.78	0.85	3.61
16	`grmhd2+hs2`	0.64	0.72	3.23
17	`grmhd8`	0.72	0.90	1.41
18	`mring+hsCW20`	1.08	0.96	2.15
19	`mring+hsCW40`	1.12	0.98	1.56
20	`mring+hsCW0.15`	1.04	1.08	1.12
21	`mring+hsCW0.60`	0.80	1.21	1.47
22	`mring+hsCW1.20`	0.65	0.83	1.57
23	`grmhd3`	0.82	0.96	1.54
24	`grmhd4`	0.61	1.11	2.34
25	`grmhd5`	0.78	0.81	2.49
26	`grmhd6`	0.64	0.84	1.07
27	`grmhd7`	1.06	0.88	1.57

Table 5.3: **Goodness of fit of `kine` reconstructions.** χ^2 values between reconstructions and synthetic data, for all models and all closure data products.

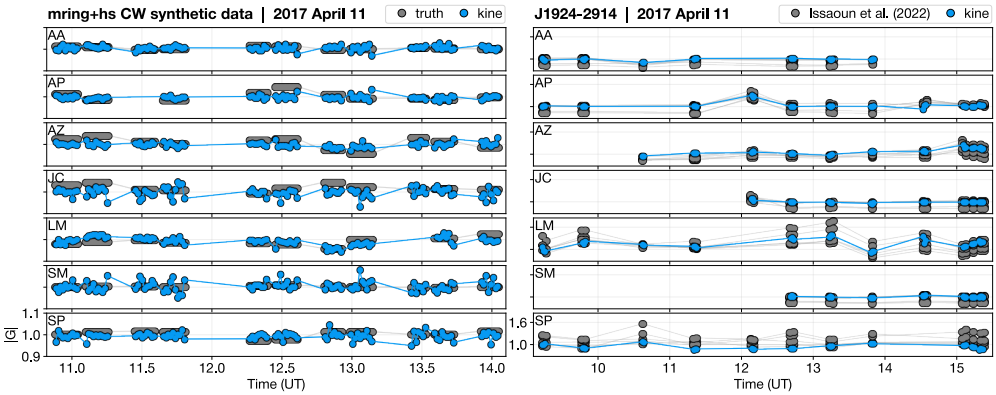

Figure 5.14: **Visibility amplitude gains recovered by `kine`.** Left panel shows the ground-truth gains applied for Sgr A* synthetic data (gray points) and those simultaneously fitted by `kine` while performing a video reconstruction (blue points). Right panel shows the gains recovered by the different imaging algorithms employed in Issaoun et al. (2022) after self-calibration and those simultaneously fitted by `kine`.

gain fitting in the pipeline, but rely on self-calibration to estimate the gains applied in the imaging of the polarimetric visibilities.

Conclusions

In this thesis, we developed, validated, and applied novel imaging and analysis methods to study the dynamical processes of black hole accretion and the launching of relativistic jets from supermassive black holes. First, we focused on imaging multiple observations of the parsec-scale jet in 3C 84, using snapshot reconstructions generated with `eht-imaging`. Thanks to the super-resolution capabilities of this method, we were able to track the time variation of several kinematic properties of the jet. However, discontinuities between frames limited the potential for more detailed temporal analysis. Second, we introduced a new imaging method for VLBI observations based on a neural field representation, capable of reconstructing a continuous video of a source from temporally separated observations. We applied this method to multiple epochs of the relativistic jet in 3C 345. The analysis of the resulting full-polarimetric video provided information about the jet's magnetic field structure, interactions with ISM, the core's irregular precession, and the plasma flow velocity. The developed method was then successfully validated on the extremely sparse coverage of the 2017 EHT array, demonstrating its ability to reconstruct a robust, horizon-scale video of Sgr A*. Here we will detail the results and conclusions drawn from each work.

Evolution of a relativistic jet. We reconstructed images of 3C 84 from 121 epochs of 43 GHz VLBA observations spanning approximately 10 years, using the RML imaging method `eht-imaging`. The images reveal the evolution of the limb-brightened jet that expands southward, ultimately tripling its initial length. We measured the expansion speed of the jet front over time and identified three distinct velocity regimes. During the first two years of observations, the jet expanded at a constant apparent speed of 0.29 ± 0.01 c, while undergoing a morphological transition from an FR I to an FR II structure. Subsequently, the newly formed jet lobe began to inflate, causing a deceleration in the expansion, which proceeded at an apparent speed of 0.228 ± 0.004 c over the next four years. This phase ended with a second morphological transition, this time from FR II back to FR I. Afterward, the jets resumed its expansion at a higher apparent speed of 0.61 ± 0.01 c. Throughout this period, the jet core exhibited rapid, irregular changes in orientation, which in turn affected the jet's direction, resulting in a winding, non-linear structure. Within 0.03 pc from the core, the jet's direction oscillated irregularly, spanning a total projected angle of 80°, including a rapid 60° shift over just two years. Altogether, our results indicate that the jet is propagating through an irregular ISM, characterized by clumps of material with varying densities that affect the jet's speed direction and morphology. Additionally, certain portions of the jet appear darkened which may be explained by the presence of absorbing gas in the foreground. The double morphological transition observed in the jet suggests that the presence of radio lobes at the tip of the jet may be a temporary phase in jet evolution, caused by a higher density of the ambient medium. The historic evolution of 3C 84 supports this interpretation, as remnants of past radio lobes have been observed at multiple scales from the core, in both the jet and the counterjet. The

irregular precession of the jet is more difficult to interpret. Potential causes include instabilities in the accretion disk or a misalignment between the disk's axis and the black hole's spin. Such a misalignment could induce warping in the disk, leading to a stochastic variation in the jet's orientation. The quantitative kinematic results and the qualitative details of the jet evolution presented in this work were made possible by combining a super-resolving imaging technique with the extensive amount of repeated observations provided by blazar monitoring programs. These results represent a change in how jet features can be tracked, shifting from simple component model fitting to specific ad-hoc analysis tailored to the jet morphology.

The `kine` imaging method. We approached the study of variable radio source by developing a dynamic imaging method capable of recovering a video of the source from VLBI observations taken at different times. The method, named `kine`, relies on a coordinate-based neural network to provide a parametric representation of the brightness distribution, which is iteratively optimized to match the observed data. Temporal and spatial regularization of the video is implicitly enforced by the network's spectral bias, which naturally favors smooth outputs by fitting low-frequency components before high-frequency ones. `kine` can recover full-polarization videos from any VLBI data product, without requiring explicit morphological priors, with no dependence on the initialization video, and with minimal tuning of the hyperparameters of the algorithm. We used `kine` to image 116 epochs of 15 GHz VLBA observations of 3C 345 spanning a period of 27 years. This resulted in a full-polarimetric time-continuous video with a resolution 5-6 times higher and a dynamic range two orders of magnitude greater than standard imaging with CLEAN. Compared to "snapshot" imaging with a super-resolution method (like that applied in the imaging of 3C 84), the `kine` video of 3C 345 achieves a factor of $\sim$2 better resolution and more than one order of magnitude improvement in dynamic range, with the additional advantage of providing a model of the source that can be sampled continuously in time. The reconstructed video shows a jet characterized by complex dynamics, with turbulent motion in the inner jet and occasional ejection of bright features from the core. We observe that also in 3C 345 the core changes orientation irregularly, spanning a projected angle of 60°. This result challenges previous interpretations suggesting a regular precession of the jet induced by a binary black hole system. The linear polarization structure shows a transverse-parallel-transverse pattern with respect to the jet axis, indicating the presence of a large-scale helical magnetic field dominated by its toroidal component. We observe an increase in fractional polarization at the jet edges, which we interpret as a result of the shearing of the jet boundary layer from friction with the ISM. To study the jet dynamics, we applied an optical flow technique to measure the local instantaneous velocity field of the plasma. The average flow speed ranges from 10-12 c in the inner jet to 5-8 c in the outer jet, implying a gradual deceleration. By integrating the velocity field, we traced the trajectories of bright components and measured their speeds, finding them to be of the same order of the surrounding average flow velocities. This indicates that the observed bright features are not traveling shocks, as previously proposed, but rather regions of increased emissivity that become Doppler boosted when emitted in the southern direction. If these features were strong shocks, we would expect a discontinuity between the the shock's pattern speed and that of the bulk velocity. We would also expect the shocked region to have a higher fractional polarization caused by the ordering of the magnetic field induced by the shock compression, but neither this is

observed in our video. The measurement of both the component speed and the bulk flow velocity was made possible thanks to the high resolution and dynamic range of the reconstructed video, combined with its time continuity, which were enabled by our novel imaging method. The ability to recover pixel-by-pixel velocity maps represents a major advancement in kinematic and dynamic studies of relativistic jets, especially if applied to decades of observations from monitoring programs.

Dynamic imaging of Sgr A*. We further extended the architecture of the `kine` imaging method to address the challenges posed by the extremely sparse instantaneous coverage of the EHT, paired with the high intra-observation variability of Sgr A*. The new architecture assumes that the video can be decomposed into a static and a dynamic component, which are modeled by separate networks. We tested the pipeline on synthetic data derived from a variety of models, including static geometric models, dynamic geometric models, and ray-traced GRMHD simulations. The synthetic datasets were simulated with realistic thermal noise, complex gain corruption, and interstellar scattering effects, following the EHT (u, v)-coverage on April 11, 2017. To evaluate the reconstruction quality, we employed multiple fidelity metrics, including: data products χ^2s values, image cross correlation, visibility variance, positional tracking of the moving components and the brightness peak, the integrated EVPA pitch angle ($\angle\beta_2$), and the pattern speed of orbiting features. In tests with static models, `kine` successfully reconstructed the source morphology in both total intensity and polarization without introducing spurious variability. The reconstructions of dynamic geometric models demonstrate that the pipeline can accurately recover orbital and non orbital motion with the correct direction and speed, track changes in the linear polarization, and distinguish between coherent and incoherent motion without systematic bias. While the maximum recoverable speed is constrained by the temporal sampling of the observations rather than by `kine` itself, the method's performance is limited by the flux ratio between the static and dynamic components. If the static component is too faint, it cannot effectively anchor the dynamic component. In tests using GRMHD simulations favored by observational constraints on Sgr A*, such as face-on, MAD, and prograde models, `kine` reliably reconstructs the total intensity morphology and the orientation of the linear polarization structure. It also correctly identifies the brightness position angle and recovers the integrated polarization field pitch angle. Furthermore, the reconstructed pattern speed of orbiting features are consistent in both direction and magnitude with the ground truth. However the magnitude of the linear polarization is often not recovered correctly. For GRMHD models that are disfavored by Sgr A* observations, particularly edge-on models with significant asymmetry and faint rings, `kine` sometimes fails to recover the ground truth with such a sparse (u, v)-coverage. Nonetheless, even in these cases, the position angle and polarization orientation are generally reconstructed accurately. Overall, `kine` demonstrates excellent performance, which is especially remarkable given the sparse instantaneous coverage of the EHT array in 2017. We found only a few isolated failure cases, which we were able to understand. Since the conditions that led to incorrect reconstructions are not present in real observations of Sgr A* we believe that the performance of the method on real data should not be hindered. Therefore this study demonstrates that the `kine` imaging method, with the specialized pipeline for EHT data, is able to recover a robust video reconstruction of Sgr A* from April 11, 2017 observations.

Impact and future developments. The major contribution of this work is the development and validation of `kine`, a new dynamic imaging method for VLBI observations. The algorithm can recover videos either from multiple repeated observations of a slowly varying source or from a single observation of a source varying on time scales shorter than the observation time. In the former case, `kine` offers a significant improvement in resolution and dynamic range over traditional methods, thanks to its ability to share information across multiple time steps. This represents an extremely promising tool that will mark a shift in dynamical studies of jets, especially now that the blazar monitoring programs initiated in the 1990s and 2000s have accumulated decades of nearly monthly observations for dozens of sources. Moreover, the continuity of the reconstructed video enables the recovery of the instantaneous projected velocity field of the emitting plasma, from which one can measure the velocity of different elements in the jet and study the details of the turbulent plasma flow. In the second case, `kine` has demonstrated its ability to reconstruct minute-by-minute videos from the extremely sparse observations of the Sgr A* black hole with the EHT. In the near future, this method will be used to image the first robust video of Sgr A*, from real EHT observations conducted on April 11, 2017. From the video, we will be able to measure the direction and speed of the material orbiting the black hole, as well as constrain its spin and the parameters of the accretion models. In the coming years `kine` can also be used in the imaging of multi-day observations of M 87*, which have been proposed for the 2026 EHT observing campaign, as well as to image other single-day observations of Sgr A* after 2017. Future developments to the algorithm will include extending `kine` to multi-frequency VLBI observations, enabling the simultaneous recovery of the spectral index map along with the video. Additionally, the neural field architecture of the method can be reformulated within a Bayesian approach, allowing for estimation of the posterior distribution of the reconstructed videos. These two improvements will make `kine`, which is already a versatile and powerful imaging method, into a truly comprehensive pipeline for a wide range of VLBI imaging tasks.

Conclusiones

En esta tesis hemos desarrollado, validado y aplicado nuevos métodos de análisis e imagen para estudiar los procesos dinámicos de acreción en agujeros negros y el lanzamiento de chorros relativistas desde agujeros negros supermasivos. En primer lugar, nos centramos en la reconstrucción de imágenes individuales a partir de múltiples observaciones del chorro a escalas del pársec en 3C 84, empleando `eht-imaging`. Gracias a las capacidades de superresolución de este método, pudimos seguir la variación temporal de varias propiedades cinemáticas del chorro. Sin embargo, las discontinuidades entre fotogramas limitaron las posibilidades de realizar análisis temporales más detallados. En segundo lugar, introdujimos un nuevo algoritmo de reconstrucción de imágenes para observaciones VLBI basado en una representación de campo neuronal, capaz de reconstruir un vídeo contínuo de una fuente a partir de observaciones separadas temporalmente. Aplicamos este método a múltiples épocas del chorro relativista en 3C 345. El análisis del vídeo resultante en polarización completa proporcionó información sobre la estructura del campo magnético del chorro, las interacciones con el ISM, la precesión irregular del núcleo y la velocidad del flujo de plasma. A continuación, el método desarrollado se validó con éxito en el cubrimiento extremadamente escaso del la red EHT en 2017, demostrando su capacidad para reconstruir un vídeo robusto a escala del horizonte de sucesos de Sgr A*. Aquí detallaremos los resultados y las conclusiones extraídas de cada trabajo.

Evolución de un chorro relativista. Hemos reconstruido imágenes de 3C 84 a partir de 121 épocas de observaciones con el VLBA a 43 GHz que abarcan aproximadamente 10 años, utilizando un método RML. Las imágenes revelan la evolución de un chorro cuyos bordes brillan por encima del eje y que se expande hacia el sur, triplicando su longitud inicial. Hemos medido la velocidad de expansión del frente del chorro a lo largo del tiempo e identificamos tres regímenes de velocidad distintos. Durante los dos primeros años de observaciones, el chorro se expandió a una velocidad aparente constante de 0.29 ± 0.01 c, mientras experimentaba una transición morfológica de una estructura FR I a una FR II. Posteriormente, el lóbulo del chorro recién formado comenzó a inflarse, provocando una desaceleración de la expansión, que prosiguió a una velocidad aparente de 0.228 ± 0.004 c durante los cuatro años siguientes. Esta fase terminó con una segunda transición morfológica, esta vez de FR II de nuevo a FR I. Después, el chorro reanudó su expansión a una mayor velocidad aparente de 0.61 ± 0.01 c. A lo largo de este periodo, el núcleo del chorro mostró cambios rápidos e irregulares de orientación, que a su vez afectaron a la dirección del chorro, dando lugar a una estructura sinuosa y no lineal. A 0.03 pc del núcleo, la dirección del chorro osciló irregularmente, abarcando un ángulo total proyectado sobre el plano del cielo de 80°, incluyendo un rápido cambio de 60° en sólo dos años. En conjunto, nuestros resultados indican que el chorro se propaga a través de un ISM irregular, caracterizado por aglomeraciones de material con densidades variables que afectan a la dirección de la velocidad del chorro y a su morfología. Además, algunas

partes del chorro aparecen oscurecidas, lo que podría explicarse por la presencia de gas absorbente en primer plano. La doble transición morfológica observada en el chorro sugiere que la presencia de lóbulos de radio en el extremo del chorro puede ser una fase temporal en la evolución del chorro, causada por una mayor densidad del medio. La evolución histórica de 3C 84 apoya esta interpretación, ya que se han observado restos de lóbulos de radio antiguos a múltiples escalas desde el núcleo, tanto en el chorro principal como en el eyectado en dirección opuesta. La precesión irregular del chorro es más difícil de interpretar. Entre las posibles causas se incluyen inestabilidades en el disco de acreción o una desalineación entre el eje del disco y el eje de rotación del agujero negro. Tal desalineación podría inducir deformaciones en el disco, dando lugar a una variación estocástica en la orientación del chorro. Los resultados cinemáticos cuantitativos y los detalles cualitativos de la evolución del chorro presentados en este trabajo han sido posibles gracias a la combinación de una técnica de imagen de superresolución con la gran cantidad de observaciones repetidas que proporcionan los programas de monitorizado de blazares. Estos resultados representan un cambio en la forma de rastrear las características del chorro, pasando de un simple ajuste de componentes Gaussianas a un análisis específico ad hoc, adaptado a la morfología del chorro.

El algoritmo de reconstrucción de vídeo `kine`. Abordamos el estudio de las radio fuentes variables desarrollando un método de imagen dinámica capaz de recuperar un vídeo de la fuente a partir de observaciones VLBI obtenidas en diferentes momentos. El método, denominado `kine`, se basa en una red neuronal cuya entrada es un sistema de coordenadas que permite proporcionar una representación paramétrica de la distribución de brillo y que se optimiza iterativamente para ajustarse a los datos observados. La regularización temporal y espacial del vídeo está implícitamente reforzada por el sesgo espectral de la red, que favorece de forma natural los resultados suaves ajustando las componentes de baja frecuencia antes que los de alta frecuencia. `kine` puede recuperar vídeos en polarización completa a partir de cualquier producto de datos VLBI, sin necesidad de información a priori sobre la morfología de la fuente, y con un ajuste mínimo de los hiperparámetros del algoritmo. Utilizamos `kine` para obtener imágenes de 116 épocas de observaciones con el VLBA a 15 GHz de 3C 345 que abarcan un periodo de 27 años. Esto dió como resultado un vídeo contínuo en el tiempo en polarimetría completa con una resolución 5-6 veces superior y un rango dinámico dos órdenes de magnitud mayor que las imágenes estándar con CLEAN. En comparación con las imágenes "instantáneas" obtenidas con un método de superresolución (como el aplicado en las imágenes de 3C 84), el vídeo de `kine` de 3C 345 consigue una resolución dos veces mejor y más de un orden de magnitud de mejora en el rango dinámico, con la ventaja adicional de proporcionar un modelo de la fuente que puede ser muestreado continuamente en el tiempo. El vídeo reconstruido muestra un chorro caracterizado por una dinámica compleja, con movimientos turbulentos en el interior del chorro y expulsión ocasional de elementos brillantes del núcleo. También observamos que el núcleo de 3C 345 cambia de orientación de forma irregular, abarcando un ángulo proyectado sobre el plano del cielo de 60°. Este resultado desafía las interpretaciones previas que sugerían una precesión regular del chorro inducida por un sistema binario de agujeros negros. La estructura de polarización lineal muestra un patrón transversal-paralelo-transversal con respecto al eje del chorro, lo que indica la presencia de un campo magnético helicoidal a gran escala dominado por su com-

ponente toroidal. Observamos un aumento del grado de polarización en los bordes del chorro, que interpretamos como resultado del cizallamiento de la capa exterior del chorro por fricción con el ISM. Para estudiar la dinámica del chorro, aplicamos una técnica de flujo óptico para medir el campo de velocidad instantánea local del plasma. La velocidad media del flujo oscila entre 10-12 c en la parte interna del chorro y 5-8 c en la parte externa del chorro, lo que implica una desaceleración gradual. Integrando el campo de velocidades, trazamos las trayectorias de las componentes brillantes y medimos sus velocidades, encontrando que son del mismo orden que las velocidades medias del flujo circundante. Esto indica que las componentes brillantes observadas no son ondas de choque, como se había propuesto anteriormente, sino más bien regiones de mayor emisividad que se potencian por efecto Doppler cuando se emiten en dirección sur. Si estas componentes fueran choques fuertes, esperaríamos que hubiese una discontinuidad entre la velocidad "pattern" del choque y la velocidad del plasma. También esperaríamos que la región de choque tuviera una un mayor grado de polarización causado por el ordenamiento del campo magnético inducido por la compresión del choque, pero esto tampoco se observa en nuestro vídeo. La medición tanto de la velocidad de las componentes como de la velocidad del flujo fue posible gracias a la alta resolución y al rango dinámico del vídeo reconstruido, combinados con su continuidad temporal, que fueron posibles gracias a nuestro novedoso algoritmo de reconstrucción de imágenes. La capacidad de recuperar mapas de velocidad píxel a píxel representa un gran avance en los estudios cinemáticos y dinámicos de chorros relativistas, especialmente aplicado a décadas de observaciones de programas de monitorizado.

Imágenes dinámicas de Sgr A*. Hemos ampliado aún más la arquitectura del algoritmo `kine` para hacer frente a los desafíos planteados por el escaso cubrimiento instantáneo del EHT, unido a la gran variabilidad intraobservación de Sgr A*. La nueva arquitectura asume que el vídeo puede descomponerse en una componente estática y dinámica, que se modelan mediante redes separadas. Hemos probado el algoritmo en datos sintéticos derivados de una variedad de modelos, incluyendo modelos geométricos estáticos, modelos geométricos dinámicos y simulaciones GRMHD. Los conjuntos de datos sintéticos se simularon con ruido térmico realista, ganancias en amplitud y en fase y efectos de scattering interestelar, siguiendo el cubrimiento (u, v) del EHT del 11 de abril de 2017. Para evaluar la calidad de la reconstrucción, empleamos múltiples métricas de fidelidad, incluyendo: valores χ^2s de los productos de datos, correlación cruzada de imágenes, varianza de las visibilidades, seguimiento posicional de las componentes móviles y el pico de brillo, la distribución integrada de los EVPAs ($\angle\beta_2$) y la velocidad de ciertas características orbitales. En las pruebas con modelos estáticos, `kine` reconstruyó con éxito la morfología de la fuente tanto en intensidad total como en polarización sin introducir variabilidad artificial. Las reconstrucciones de modelos geométricos dinámicos demuestran que el algoritmo puede recuperar con precisión el movimiento orbital y no orbital con la dirección y velocidad correctas, rastrear cambios en la polarización lineal y distinguir entre movimiento coherente e incoherente sin un sesgo sistemático. Aunque la velocidad máxima recuperable está limitada por el muestreo temporal de las observaciones y no por `kine`, el rendimiento del método está limitado por la relación de flujo entre las componentes estática y dinámica. Si la componente estática es demasiado débil, no puede anclar eficazmente la componente dinámica. En las pruebas realizadas con simulaciones GRMHD favorecidas por las ob-

servaciones de Sgr A*, como los modelos face-on, MAD y prograde, `kine` reconstruye con fiabilidad la morfología en intensidad total y la orientación de la estructura de polarización lineal. También identifica correctamente el ángulo de posición del brillo y recupera el ángulo de inclinación del campo de polarización integrado. Además, las velocidades reconstruidas de las características orbitales son consistentes tanto en dirección como en magnitud con el modelo original. Sin embargo, la magnitud de la polarización lineal a menudo no se recupera correctamente. Para los modelos GRMHD que son desfavorecidos por las observaciones de Sgr A*, particularmente los modelos con vista de canto que presentan una asimetría significativa y anillos tenues, `kine` a veces no logra recuperar la morfología original, con un cubrimiento (u, v) tan escaso. No obstante, incluso en estos casos, el ángulo de posición y la orientación de la polarización generalmente se reconstruyen con precisión. En general, `kine` demuestra un rendimiento excelente, lo que es especialmente destacable dado el escaso cubrimiento instantáneo de la red de antennas del EHT en 2017. Solo encontramos unos pocos casos de fallo aislados, que pudimos comprender. Dado que las condiciones que llevaron a reconstrucciones incorrectas no están presentes en las observaciones reales de Sgr A*, creemos que el rendimiento del método con datos reales no debería verse afectado. Por lo tanto, este estudio demuestra que el algoritmo `kine`, con el proceso de imagen ideado para datos del EHT, es capaz de reconstruir de forma robusta un video de Sgr A* a partir de las observaciones del 11 de abril de 2017.

Impacto y desarrollo futuros. La principal contribución de este trabajo es el desarrollo y la validación de `kine`, un nuevo método de imagen dinámica para observaciones VLBI. El algoritmo puede recuperar videos tanto a partir de múltiples observaciones de una fuente que varía lentamente, como a partir de una única observación de una fuente que varía en escalas de tiempo más rápidas que el tiempo de observación. En el primer caso, `kine` ofrece una mejora significativa en la resolución y el rango dinámico en comparación con los métodos tradicionales, gracias a su capacidad para compartir información entre múltiples instantes de tiempo. Esto representa una herramienta extremadamente prometedora que marcará un cambio en los estudios dinámicos de jets, especialmente ahora que los programas de monitorizado de blazares iniciados en los años 1990 y 2000 han acumulado décadas de observaciones casi mensuales para docenas de fuentes. Además, la continuidad del vídeo reconstruido permite recuperar el campo de velocidad instantáneo del plasma emisor proyectado sobre el plano del cielo, a partir del cual se puede medir la velocidad de diferentes elementos del chorro y estudiar los detalles del flujo de plasma turbulento. En el segundo caso, `kine` ha demostrado su capacidad para reconstruir videos minuto a minuto a partir de las observaciones extremadamente escasas del agujero negro Sgr A* con el EHT. En un futuro próximo, este método se utilizará para obtener el primer vídeo de Sgr A*, a partir de observaciones reales del EHT realizadas el 11 de abril de 2017. A partir del vídeo, podremos medir la dirección y la velocidad del material que orbita el agujero negro, así como restringir su spin y los parámetros de los modelos de acrecimiento. En los próximos años, `kine` también podrá utilizarse para la reconstrucción de un vídeo a partir de observaciones de varios días de M 87*, que han sido propuestas para la campaña de observación del EHT en 2026, así como para obtener imágenes de otras observaciones de un solo día de Sgr A* posteriores a 2017. El desarrollo futuro del algoritmo incluirá la extensión de `kine` a observaciones VLBI multifrecuencia, lo que permitirá la recuperación simultánea del mapa del índice espectral junto con el video.

Además, la arquitectura de campos neuronales del método podrá reformularse dentro de un marco bayesiano, permitiendo así la estimación de la distribución posterior de los videos reconstruidos. Estas dos mejoras convertirán a `kine`, que ya es un método de imagen potente y versátil, en una herramienta verdaderamente integral para una amplia gama de tareas de imagen VLBI.

Acronyms

AGN: **A**ctive **G**alactic **N**uclei

ALMA: **A**tacama **L**arge **M**illimeter **A**rray

BBH: **B**inary **B**lack **H**ole

BP: **B**landford-**P**ayne

BZ: **B**landford-**Z**najek

CNN: **C**onvolutional **N**eural **N**etwork

DL: **D**eep **L**earning

EHT: **E**vent **H**orizon **T**elescope

EHTC: **E**vent **H**orizon **T**elescope **C**ollaboration

EVPA: **E**lectric **V**ector **P**osition **A**ngle

FNN: **F**eed-forward **N**eural **N**etwork

FOV: **F**ield **O**f **V**iew

FR: **F**anaroff–**R**iley

FWHM: **F**ull **W**idth at **H**alf **M**aximum

GP: **G**aussian **P**rocess

GR: **G**eneral **R**elativity

GELU: **G**aussian **E**rror **L**inear **U**nit

GRMHD: **G**eneral **R**elativistic **M**agneto-**H**ydro-**D**ynamic

HS: **H**orn **S**chunck

ICM: **I**ntra**C**luster **M**edium

ISM: **I**nter**S**tellar **M**edium

LR: **L**earning **R**ate

MAD: **M**agnetically **A**rrested **D**isk

ML: **M**achine **L**earning

MLP: **M**ulti-**L**ayer **P**erceptron

MLP: **M**agnetic **R**esonance **I**maging

NeRF: **N**eural **R**adiance **F**ield

NN: **N**eural **N**etwork

OF: **O**ptical **F**low

PA: **P**osition **A**ngle

PS: **P**ower **S**pectrum

RML: **R**egularized **M**aximum **E**ntropy

SKA: **S**quare **K**ilometer **A**rray

SNR: **S**ignal-to-**N**oise **R**atio

SMBH: **S**uper**M**assive **B**lack **H**ole

SR: **S**pecial **R**elativity

VLA: **V**ery **L**arge **A**rray

VLBA: **V**ery **L**ong **B**aseline **A**rray

VLBI: **V**ery **L**ong **B**aseline **I**nterferometry

Acknowledgements

Funding. The project that gave rise to these results received the support of a fellowship from "la Caixa" Foundation (ID 100010434). The fellowship code is LCF/BQ/DI22/11940027. The author of this thesis also acknowledges financial support from the Spanish Ministerio de Economía y Competitividad (grant numberPID2022-140888NB-C21) and the Severo Ochoa grant CEX2021-001131-S funded by MCIN/AEI/10.13039/501100011033.

Data. This thesis made use of data from the MOJAVE database that is maintained by the MOJAVE team (Lister et al. 2018). This thesis made use of VLBA data from the VLBA-BU Blazar Monitoring Program (BEAM-ME and VLBA-BU-BLAZAR; http://www.bu.edu/blazars/BEAM-ME.html), funded by NASA through the Fermi Guest Investigator grants, the latest is 80NSSC23K1508.

Professional acknowledgements. I would like to thank Marie Lou Gendron Marsolais for useful suggestions regarding the interpretation of the precession in 3C 84. I also thank Ben Prather and Vedant Dhruv for providing the GRMHD simulations that were useful in developing the `kine` method. I finally thank José M. Martí and Manel Perucho for the discussion and suggestions regarding the interpretation of the 3C 345 video reconstruction and Svetlana Jorstad and Alan Marscher for the useful comments and suggestions about the signatures of weak and strong shocks in jets.

Personal acknowledgements. On a more personal level, I would like to thank many people who have directly and indirectly contributed to the successful completion of this thesis and my doctoral studies.

I would like to start from my supervisor José, who in many occasions has gone out of his way to provide me the best opportunities to showcase my work and to successfully progress in my career. I also thank him for believing in my potential from the very beginning of the PhD.

I want to thank the postdocs in my group for their scientific and personal support, in particular Ilje, whose patient guidance during the first year of my PhD was extremely valuable.

I would also like to thank Antonio for his fundamental support, as a colleague and as a friend. Without him, `kine` would have taken much longer to develop or would have not achieved the same quality. His perfectionism and maximalist tendencies sometimes clashed with my more pragmatic approach and preference for compromise. But I truly believe that the combination of our different working styles (one grounded in theoretical understanding and first principles, the other based in persistent and tenacious empirical testing) contributed to the successful development of the `kine` algorithm. In these years, we spent countless hours working side by side, losing our mind on the code and the tests. Time and time again Antonio listened with empathy to my frustrations, lifted me up in moments of hopelessness and tears, and weathered the pressure and nonsense with me. If I survived the stress of the PhD unscathed is

because I could count on him. He's been the best colleague I could have wished for and I'll miss working with him once I leave the IAA.

I also want to thank Rohan, with whom I've shared my PhD from the beginning to the end. His pro-activeness and his dedication to work motivated me out of my laziness and shyness in multiple occasions. I'm grateful for all the fun moments we had in these years and for having someone who would understand my exasperation. Sharing my office and our group meetings with such a bright fellow student greatly contributed to preserving my sanity.

I want to thank Katie and Aviad for their precious mentorship through the development of `kine` and beyond. I thank them for hosting me at Caltech and introducing me to the computational vision field. I'm grateful for the confidence they placed in my capabilities as an early career researcher. I also owe to Katie the choice of my PhD because it was her TedTalk in 2016 which motivated me to pursue research in machine learning methods for imaging black holes.

I also want to thank the other coauthors of the publications following from this thesis, which I have not mentioned yet. Specifically I would like to thank Brandon for his talented co-developing of `kine`, Svetlana and Alan for sharing their expertise on relativistic jets, and Nick for his significant contributions to `kine`'s validation.

I'm grateful that here is Granada I've found an amazing group of friends among the PhD students at the IAA. I look back with fondness on the dinners, the parties, the hikes, the festivals, and the casas rurales that we shared together. But most of all I will remember the quiet conversations at a bar, at home after dinner, or in front of our cars, talking about love, life, politics, and our futures. I found myself surrounded by brilliant and caring friends whose commitment to activism, society, and the environment has been a source of inspiration and motivation.

A special thanks goes to my roommate Belén, which welcomed me in her home and with whom I've built a solid friendship. During the hardest moments of my PhD, knowing that I would come home to have dinner with her really lifted my spirits. I'll miss our late night talks on the sofa or in the terrace.

I want to thank Marie Lou, Miguel, Roberto, Alberto, Oier, Clara, and Kelley for their friendship and the moments we shared hiking, camping, and climbing. They made the Granada outdoors a truly wonderful place and it's gonna be hard to find a better climbing crew than them. In particular, I thank Roberto for often challenging my views with grace and acuity and for making me widen my perspectives.

I'm also grateful for the friendship of Marta and Ixaka, who joined me with their cheerfulness in our small van life adventures, Borja, who first reached out to me to make me join his office and group of friends, Julio, who brings his sharp and brilliant point of view to any conversation, and Ana, who's a talented artist and one of the coolest persons I know.

I would like to thank my parents Damiano and Silvia, who have been always encouraging and supportive of my PhD studies, or, for that matter, any other study I pursued in my life. The intelligence or education I may have is probably more their merit than mine. I also thank my siblings Matilde and Tommaso, who despite the distance have constantly shown me their love and affection. Especially thanks to Matilde for being the one who most often initiates our calls.

Finally I thank my fiancé Gabriele, whose encouragement, emotional comfort, constant care, and absolute faith in my abilities have been fundamental during my PhD.

Through these years he endured countless flights from Italy to Granada, withstood my too frequent late nights and weekends spent working, and helped me through my worst moments of crisis. He has been the best friend I could have by my side and I'm forever grateful for his character and love.

Bibliography

Abdo, A. A. et al. (2009). "Fermi Discovery of Gamma-ray Emission from NGC 1275". In: *ApJ* 699.1, pp. 31–39. DOI: 10.1088/0004-637X/699/1/31.

Akiyama, K. et al. (2017). "Imaging the Schwarzschild-radius-scale Structure of M87 with the Event Horizon Telescope Using Sparse Modeling". In: *ApJ* 838.1, p. 1. DOI: 10.3847/1538-4357/aa6305.

Antonucci, R. (1993). "Unified models for active galactic nuclei and quasars". In: *Annual Rev. Astron. Astrophys.* 31, pp. 473–521. DOI: 10.1146/annurev.aa.31.090193.002353.

Arras, P. et al. (2022). "Variable structures in M87* from space, time and frequency resolved interferometry". In: *Nat. Astr.* 6, pp. 259–269. DOI: 10.1038/s41550-021-01548-0.

Asada, K. et al. (2006). "The Expanding Radio Lobe of 3C 84 Revealed by VSOP Observations". In: *PASJ* 58, pp. 261–270. DOI: 10.1093/pasj/58.2.261.

Baade, W. and R. Minkowski (1954). "Identification of the Radio Sources in Cassiopeia, Cygnus A, and Puppis A". In: *ApJ* 119, p. 206. DOI: 10.1086/145812.

Balick, B. and R. L. Brown (1974). "Intense sub-arcsecond structure in the galactic center." In: *ApJ* 194, pp. 265–270. DOI: 10.1086/153242.

Beckmann, V. and C. R. Shrader (2012). *Active Galactic Nuclei*. Wiley. ISBN: 9783527666829. DOI: 10.1002/9783527666829.

Begelman, M. C. et al. (1980). "Massive black hole binaries in active galactic nuclei". In: *Nature* 287.5780, pp. 307–309. DOI: 10.1038/287307a0.

Biretta, J. A. et al. (1986). "The Evolution of the Compact Radio Source in 3C 345. I. VLBI Observations". In: *ApJ* 308, p. 93. DOI: 10.1086/164481.

Blandford, R. D. and A. Königl (1979). "Relativistic jets as compact radio sources". In: *ApJ* 232, pp. 34–48. DOI: 10.1086/157262.

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

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

Bolton, J. G. et al. (1949). "Positions of Three Discrete Sources of Galactic Radio-Frequency Radiation". In: *Nature* 164.4159, pp. 101–102. DOI: 10.1038/164101b0.

Bouman, K. L. et al. (2018). "Reconstructing Video of Time-Varying Sources From Radio Interferometric Measurements". In: *IEEE Transactions on Computational Imaging* 4.4, pp. 512–527. DOI: 10.1109/TCI.2018.2838452.

Bradbury, J. et al. (2018). *JAX: composable transformations of Python+NumPy programs*. Version 0.3.13. `http://github.com/jax-ml/jax`.

Broderick, A. E. et al. (2020). "THEMIS: A Parameter Estimation Framework for the Event Horizon Telescope". In:

ApJ 897.2, p. 139. DOI: 10.3847/1538-4357/ab91a4.

Bronzwaer, T. and H. Falcke (2021). "The Nature of Black Hole Shadows". In: *ApJ* 920.2, p. 155. DOI: 10.3847/1538-4357/ac1738.

Caproni, A. and Z. Abraham (2004). "Precession in the Inner Jet of 3C 345". In: *ApJ* 602.2, pp. 625–634. DOI: 10.1086/381195.

Carilli, C. L. and N. Thyagarajan (2022). "Hybrid Mapping of the Black Hole Shadow in M87". In: *ApJ* 924.2, p. 125. DOI: 10.3847/1538-4357/ac3cba.

Carroll, S. M. (2019). *Spacetime and Geometry: An Introduction to General Relativity*. Cambridge University Press. DOI: 10.1017/9781108770385.

Chael, A. A. et al. (2018). "Interferometric Imaging Directly with Closure Phases and Closure Amplitudes". In: *ApJ* 857.1, p. 23. DOI: 10.3847/1538-4357/aab6a8.

Cittert, P. H. van (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.

Cohen, M. H. et al. (1977). "Radio sources with superluminal velocities". In: *Nature* 268.5619, pp. 405–409. DOI: 10.1038/268405a0.

Colaninno, R. C. and A. Vourlidas (2006). "Analysis of the Velocity Field of CMEs Using Optical Flow Methods". In: *ApJ* 652.2, pp. 1747–1754. DOI: 10.1086/507943.

Conroy, N. S. et al. (2023). "Rotation in Event Horizon Telescope Movies". In: *ApJ* 951.1, p. 46. DOI: 10.3847/1538-4357/acd2c8.

Cybenko, G. (1989). "Approximation by superpositions of a sigmoidal function". In: *Mathematics of Control, Signals, and Systems* 2.4, pp. 303–314. DOI: 10.1007/BF02551274.

Dahale, R. et al. (2025). "Validation and evaluation of horizon-scale Sagittarius A* video reconstructions". In: *In preparation.*

Dent, W. A. (1966). "Variation in the Radio Emission from the Seyfert Galaxy NGC 1275". In: *ApJ* 144, p. 843. DOI: 10.1086/148674.

Dewdney, P. E. et al. (2009). "The Square Kilometre Array". In: *IEEE Proceedings* 97.8, pp. 1482–1496. DOI: 10.1109/JPROC.2009.2021005.

Dhruv, V. et al. (2025). "A Survey of General Relativistic Magnetohydrodynamic Models for Black Hole Accretion Systems". In: *ApJ S.* 277.1, p. 16. DOI: 10.3847/1538-4365/adaea6.

Dunn, R. J. H. et al. (2006). "Precession of the super-massive black hole in NGC 1275 (3C 84)?" In: *MNRAS* 366.3, pp. 758–766. DOI: 10.1111/j.1365-2966.2005.09928.x.

Event Horizon Telescope Collaboration (2019a). "First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole". In: *ApJ L.* 875.1, p. L1. DOI: 10.3847/2041-8213/ab0ec7.

— (2019b). "First M87 Event Horizon Telescope Results. II. Array and Instrumentation". In: *ApJ L.* 875.1, p. L2. DOI: 10.3847/2041-8213/ab0c96.

— (2019c). "First M87 Event Horizon Telescope Results. III. Data Processing and Calibration". In: *ApJ L.* 875.1, p. L3. DOI: 10.3847/2041-8213/ab0c57.

— (2019d). "First M87 Event Horizon Telescope Results. IV. Imaging the Central Supermassive Black Hole". In: *ApJ L.* 875.1, p. L4. DOI: 10.3847/2041-8213/ab0e85.

— (2019e). "First M87 Event Horizon Telescope Results. V. Physical Ori-

gin of the Asymmetric Ring". In: *ApJ L.* 875.1, p. L5. DOI: 10.3847/2041-8213/ab0f43.

— (2019f). "First M87 Event Horizon Telescope Results. VI. The Shadow and Mass of the Central Black Hole". In: *ApJ L.* 875.1, p. L6. DOI: 10.3847/2041-8213/ab1141.

— (2022a). "First Sagittarius A* Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole in the Center of the Milky Way". In: *ApJ L.* 930.2, p. L12. DOI: 10.3847/2041-8213/ac6674.

— (2022b). "First Sagittarius A* Event Horizon Telescope Results. II. EHT and Multiwavelength Observations, Data Processing, and Calibration". In: *ApJ L.* 930.2, p. L13. DOI: 10.3847/2041-8213/ac6675.

— (2022c). "First Sagittarius A* Event Horizon Telescope Results. III. Imaging of the Galactic Center Supermassive Black Hole". In: *ApJ L.* 930.2, p. L14. DOI: 10.3847/2041-8213/ac6429.

— (2022d). "First Sagittarius A* Event Horizon Telescope Results. IV. Variability, Morphology, and Black Hole Mass". In: *ApJ L.* 930.2, p. L15. DOI: 10.3847/2041-8213/ac6736.

— (2022e). "First Sagittarius A* Event Horizon Telescope Results. V. Testing Astrophysical Models of the Galactic Center Black Hole". In: *ApJ L.* 930.2, p. L16. DOI: 10.3847/2041-8213/ac6672.

— (2024). "First Sagittarius A* Event Horizon Telescope Results. VIII. Physical Interpretation of the Polarized Ring". In: *ApJ L.* 964.2, p. L26. DOI: 10.3847/2041-8213/ad2df1.

— (2025). "Dynamic evolution of the emission ring in Sgr A*". In: *In preparation.*

Fabian, A. C. et al. (2003). "A deep Chandra observation of the Perseus cluster: shocks and ripples". In: *MNRAS* 344.3, pp. L43–L47. DOI: 10.1046/j.1365-8711.2003.06902.x.

Falcke, H. et al. (2000). "Viewing the Shadow of the Black Hole at the Galactic Center". In: *ApJ L.* 528.1, pp. L13–L16. DOI: 10.1086/312423.

Fanaroff, B. L. and J. M. Riley (1974). "The morphology of extragalactic radio sources of high and low luminosity". In: *MNRAS* 167, 31P–36P. DOI: 10.1093/mnras/167.1.31P.

Farah, J. et al. (2022). "Selective Dynamical Imaging of Interferometric Data". In: *ApJ L.* 930.2, p. L18. DOI: 10.3847/2041-8213/ac6615.

Feng, B. T. et al. (2024). "Event-horizon-scale Imaging of M87* under Different Assumptions via Deep Generative Image Priors". In: *ApJ* 975.2, p. 201. DOI: 10.3847/1538-4357/ad737f.

Forman, W. et al. (1972). "Observations of the Extended X-Ray Sources in the Perseus and Coma Clusters from UHURU". In: *ApJ* 178, pp. 309–316. DOI: 10.1086/151791.

Foschi, M. et al. (2025a). "Evolution, speed, and precession of the parsec-scale jet in the 3C 84 radio galaxy". In: *A&A* 696, A17. DOI: 10.1051/0004-6361/202453406.

Foschi, M. et al. (2025b). "Turbulent flow in the relativistic jet of 3C 345 from neural video reconstructions of interferometric data". In: *Under review.*

Fuentes, A. et al. (2021). "Magnetized relativistic jets and helical magnetic fields. II. Radiation". In: *A&A* 650, A61. DOI: 10.1051/0004-6361/202140659.

Fuentes, A. (2022). "High-Resolution Imaging of Relativistic Jets and Supermassive Black Holes". Ph.D. dissertation. Granada, Spain: University of Granada. https://digibug.ugr.es/handle/10481/79637.

Fuentes, A. et al. (2018). "Total and Linearly Polarized Synchrotron Emission from Overpressured Magnetized Relativistic Jets". In: *ApJ* 860.2, p. 121. DOI: 10.3847/1538-4357/aac091.

Fuentes, A. et al. (2023). "Filamentary structures as the origin of blazar jet radio variability". In: *Nat. Astr.* 7, pp. 1359–1367. DOI: 10.1038/s41550-023-02105-7.

Fujita, Y. and H. Nagai (2017). "Discovery of a new subparsec counterjet in NGC 1275: the inclination angle and the environment". In: *MNRAS* 465.1, pp. L94–L98. DOI: 10.1093/mnrasl/slw217.

Gabuzda, D. C. et al. (2014). "Are spine-sheath polarization structures in the jets of active galactic nuclei associated with helical magnetic fields?" In: *MNRAS* 444.1, pp. 172–184. DOI: 10.1093/mnras/stu1381.

Gabuzda, D. C. et al. (2015). "Evidence for a toroidal magnetic-field component in 5C 4.114 on kiloparsec scales". In: *A&A* 583, A96. DOI: 10.1051/0004-6361/201527185.

Gammie, C. F. et al. (2012). *HARM: A Numerical Scheme for General Relativistic Magnetohydrodynamics*. Astrophysics Source Code Library, record ascl:1209.005. https://ui.adsabs.harvard.edu/abs/2012ascl.soft09005G.

Ghez, A. M. et al. (2008). "Measuring Distance and Properties of the Milky Way's Central Supermassive Black Hole with Stellar Orbits". In: *ApJ* 689.2, pp. 1044–1062. DOI: 10.1086/592738.

Ghisellini, G. (2013). *Radiative Processes in High Energy Astrophysics*. Vol. 873. Springer. DOI: 10.1007/978-3-319-00612-3.

Gibson, J. J. (1950). *The Perception of the Visual World*. Boston: Houghton Mifflin. https://ia600805.us.archive.org/12/items/perceptionofvisu00jame/perceptionofvisu00jame.pdf.

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

Giovannini, G. et al. (2018). "A wide and collimated radio jet in 3C84 on the scale of a few hundred gravitational radii". In: *Nat. Astr.* 2, pp. 472–477. DOI: 10.1038/s41550-018-0431-2.

Gómez, J.-L. et al. (1999). "86, 43, and 22 GHz VLBI Observations of 3C 120". In: *ApJ L.* 521.1, pp. L29–L32. DOI: 10.1086/312170.

Gómez, J.-L. et al. (2000). "Flashing Superluminal Components in the Jet of the Radio Galaxy 3C120". In: *Science* 289.5488, pp. 2317–2320. DOI: 10.1126/science.289.5488.2317.

Gravity Collaboration (2023). "Polarimetry and astrometry of NIR flares as event horizon scale, dynamical probes for the mass of Sgr A*". In: *A&A* 677, p. L10. DOI: 10.1051/0004-6361/202347416.

Hazard, C. et al. (1963). "Investigation of the Radio Source 3C 273 By The Method of Lunar Occultations". In: *Nature* 197.4872, pp. 1037–1039. DOI: 10.1038/1971037a0.

He, K. et al. (2015). "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification". In: *Proceedings of the IEEE international conference on computer vision*, pp. 1026–1034. https://ieeexplore.ieee.org/document/7410480.

— (2016). "Deep Residual Learning for Image Recognition". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778. https://ieeexplore.ieee.org/document/7780459.

Hendrycks, D. and K. Gimpel (2016). "Bridging Nonlinearities and Stochastic Regularizers with Gaussian Error Linear Units". In: *CoRR* abs/1606.08415. DOI: 10.48550/arXiv.1606.08415.

Hiura, K. et al. (2018). "VERA monitoring of the radio jet 3C 84 in the period of 2007-2013: Detection of nonlinear motion". In: *PASJ* 70.5, p. 83. DOI: 10.1093/pasj/psy078.

Hodgson, J. A. et al. (2021). "A Detailed Kinematic Study of 3C 84 and Its Connection to γ-Rays". In: *ApJ* 914.1, p. 43. DOI: 10.3847/1538-4357/abf6dd.

Högbom, J. A. (1974). "Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines". In: *A&A S.* 15, p. 417. https://ui.adsabs.harvard.edu/abs/1974A&AS...15..417H.

Horn, B. K. and B. G. Schunck (1981). "Determining Optical Flow". In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* 281, pp. 319–331. https://ui.adsabs.harvard.edu/abs/1981SPIE..281..319H.

Hornik, K. (1991). "Approximation capabilities of multilayer feedforward networks". In: *Neural Networks* 4.2, pp. 251–257. ISSN: 0893-6080. DOI: 10.1016/0893-6080(91)90009-T.

Hughes, P. A. et al. (1985). "Polarized radio outbursts in BL Lacertae. II. The flux and polarization of a piston-driven shock." In: *ApJ* 298, pp. 301–315. DOI: 10.1086/163611.

Ioffe, S. (2015). "Batch normalization: Accelerating deep network training by reducing internal covariate shift". In: *Proceedings of the 32nd International Conference on Machine Learning (ICML)*. https://dl.acm.org/doi/10.5555/3045118.3045167.

Issaoun, S. et al. (2022). "Resolving the Inner Parsec of the Blazar J1924-2914 with the Event Horizon Telescope". In: *ApJ* 934.2, p. 145. DOI: 10.3847/1538-4357/ac7a40.

Jacot, A. et al. (2018). "Neural Tangent Kernel: Convergence and Generalization in Neural Networks". In: *Advances in Neural Information Processing Systems (NIPS)* 31, pp. 8571–8580. https://dl.acm.org/doi/abs/10.5555/3327757.3327948.

Jansky, K. G. (1933). "Radio Waves from Outside the Solar System". In: *Nature* 132.3323, p. 66. DOI: 10.1038/132066a0.

Janssen, M. et al. (2021). "Event Horizon Telescope observations of the jet launching and collimation in Centaurus A". In: *Nat. Astr.* 5, pp. 1017–1028. DOI: 10.1038/s41550-021-01417-w.

Jennison, R. C. (1958). "A phase sensitive interferometer technique for the measurement of the Fourier transforms of spatial brightness distributions of small angular extent". In: *MNRAS* 118, p. 276. DOI: 10.1093/mnras/118.3.276.

Jennison, R. C. and M. K. Das Gupta (1953). "Fine Structure of the Extra-terrestrial Radio Source Cygnus I". In: *Nature* 172.4387, pp. 996–997. DOI: 10.1038/172996a0.

Johnson, M. D. et al. (2017). "Dynamical Imaging with Interferometry". In: *ApJ* 850.2, p. 172. DOI: 10.3847/1538-4357/aa97dd.

Johnson, M. 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.

Jorstad, S. and A. Marscher (2016). "The VLBA-BU-BLAZAR Multi-Wavelength Monitoring Program". In: *Galaxies* 4.4, p. 47. DOI: 10.3390/galaxies4040047.

Jorstad, S. G. et al. (2005). "Polarimetric Observations of 15 Active Galactic Nuclei at High Frequencies: Jet Kinematics from Bimonthly Monitoring with the Very Long Baseline Array". In: *AJ* 130.4, pp. 1418–1465. DOI: 10.1086/444593.

Jorstad, S. G. et al. (2017). "Kinematics of Parsec-scale Jets of Gamma-Ray Blazars at 43 GHz within the VLBA-BU-BLAZAR Program". In: *ApJ* 846.2, p. 98. DOI: 10.3847/1538-4357/aa8407.

Junklewitz, H. et al. (2016). "RESOLVE: A new algorithm for aperture synthesis imaging of extended emission in radio astronomy". In: *A&A* 586, A76. DOI: 10.1051/0004-6361/201323094.

Kam, M. et al. (2024). "Evolution of the Termination Region of the Parsec-scale Jet of 3C 84 Over the Past 20 yr". In: *ApJ* 970.2, p. 176. DOI: 10.3847/1538-4357/ad51dc.

Kardashev, N. S. and Khartov (2013). "RadioAstron - A telescope with a size of 300 000 km: Main parameters and first observational results". In: *Astronomy Reports* 57.3, pp. 153–194. DOI: 10.1134/S1063772913030025.

Katz, J. I. (1997). "A Precessing Disk in OJ 287?" In: *ApJ* 478.2, pp. 527–529. DOI: 10.1086/303811.

Kim, J. Y. et al. (2018). "The limb-brightened jet of M87 down to the 7 Schwarzschild radii scale". In: *A&A* 616, A188. DOI: 10.1051/0004-6361/201832921.

Kingma, D. P. and J. Ba (2014). "Adam: A Method for Stochastic Optimization". In: *arXiv e-prints*, arXiv:1412.6980. DOI: 10.48550/arXiv.1412.6980.

Kino, M. et al. (2018). "Evidence of Jet-Clump Interaction: A Flip of the Radio Jet Head of 3C 84". In: *ApJ* 864.2, p. 118. DOI: 10.3847/1538-4357/aad6e3.

Kino, M. et al. (2021). "Morphological transition of the compact radio lobe in 3C84 via the strong jet-cloud collision". In: *arXiv e-prints*. DOI: 10.48550/arXiv.2110.11082.

Koller, D. and N. Friedman (2009). *Probabilistic Graphical Models: Principles and Techniques*. Adaptive Computation and Machine Learning. The MIT Press. ISBN: 9780262013192. https://mitpress.mit.edu/9780262013192/probabilistic-graphical-models/.

Kraft, R. P. et al. (2002). "Chandra Observations of the X-Ray Jet in Centaurus A". In: *ApJ* 569.1, pp. 54–71. DOI: 10.1086/339062.

Lai, S. et al. (2025). "Deep learning VLBI image reconstruction with closure invariants". In: *MNRAS* 536.1, pp. 446–461. DOI: 10.1093/mnras/stae2607.

Lee, J. et al. (2018). "Deep Neural Networks as Gaussian Processes". In: *International Conference on Learning Representations*. https://openreview.net/forum?id=B1EA-M-0Z.

Levis, A. et al. (2024). "Orbital polarimetric tomography of a flare near the Sagittarius A* supermassive black hole". In: *Nat. Astr.* 8, pp. 765–773. DOI: 10.1038/s41550-024-02238-3.

LIGO Scientific Collaboration and Virgo Collaboration (2016). "Observation of Gravitational Waves from a Binary Black Hole Merger". In: *Phys. Rev. Lett.* 116.6, p. 061102. DOI: 10.1103/PhysRevLett.116.061102.

Liska, M. et al. (2018). "Formation of precessing jets by tilted black hole discs in 3D general relativistic MHD simulations". In: *MNRAS* 474.1, pp. L81–L85. DOI: 10.1093/mnrasl/slx174.

Liska, M. et al. (2019). "Bardeen-Petterson alignment, jets, and magnetic truncation in GRMHD simula-

tions of tilted thin accretion discs". In: *MNRAS* 487.1, pp. 550–561. DOI: 10.1093/mnras/stz834.

Liska, M. et al. (2021). "Disc tearing and Bardeen-Petterson alignment in GRMHD simulations of highly tilted thin accretion discs". In: *MNRAS* 507.1, pp. 983–990. DOI: 10.1093/mnras/staa099.

Lister, M. L. et al. (2016). "MOJAVE: XIII. Parsec-scale AGN Jet Kinematics Analysis Based on 19 years of VLBA Observations at 15 GHz". In: *AJ* 152.1, p. 12. DOI: 10.3847/0004-6256/152/1/12.

Lister, M. L. et al. (2018). "MOJAVE. XV. VLBA 15 GHz Total Intensity and Polarization Maps of 437 Parsec-scale AGN Jets from 1996 to 2017". In: *ApJ S.* 234.1, p. 12. DOI: 10.3847/1538-4365/aa9c44.

Lister, M. L. et al. (2019). "MOJAVE. XVII. Jet Kinematics and Parent Population Properties of Relativistically Beamed Radio-loud Blazars". In: *ApJ* 874.1, p. 43. DOI: 10.3847/1538-4357/ab08ee.

Lister, M. L. et al. (2021). "Monitoring Of Jets in Active Galactic Nuclei with VLBA Experiments. XVIII. Kinematics and Inner Jet Evolution of Bright Radio-loud Active Galaxies". In: *ApJ* 923.1, p. 30. DOI: 10.3847/1538-4357/ac230f.

Lobanov, A. P. and J. Roland (2005). "A supermassive binary black hole in the quasar 3C 345". In: *A&A* 431.3, pp. 831–846. DOI: 10.1051/0004-6361:20041831.

Lyutikov, M. et al. (2005). "Polarization and structure of relativistic parsec-scale AGN jets". In: *MNRAS* 360.3, pp. 869–891. DOI: 10.1111/j.1365-2966.2005.08954.x.

Marscher, A. P. and W. K. Gear (1985). "Models for high-frequency radio outbursts in extragalactic sources, with application to the early 1983 millimeter-to-infrared flare of 3C 273." In: *ApJ* 298, pp. 114–127. DOI: 10.1086/163592.

Marscher, A. P. (2006). "Relativistic Jets in Active Galactic Nuclei". In: *Relativistic Jets: The Common Physics of AGN, Microquasars and Gamma-Ray Bursts*. Vol. 856. American Institute of Physics Conference Series, pp. 1–22. https://cds.cern.ch/record/896595.

Matthews, T. A. and A. R. Sandage (1963). "Optical Identification of 3C 48, 3C 196, and 3C 286 with Stellar Objects". In: *ApJ* 138, p. 30. DOI: 10.1086/147615.

McCready, L. L. et al. (1947). "Solar Radiation at Radio Frequencies and Its Relation to Sunspots". In: *Proceedings of the Royal Society of London Series A* 190.1022, pp. 357–375. DOI: 10.1098/rspa.1947.0081.

Medeiros, L. et al. (2023). "Principal-component Interferometric Modeling (PRIMO), an Algorithm for EHT Data. I. Reconstructing Images from Simulated EHT Observations". In: *ApJ* 943.2, p. 144. DOI: 10.3847/1538-4357/acaa9a.

Meenakshi, M. et al. (2023). "A polarization study of jets interacting with turbulent magnetic fields". In: *MNRAS* 526.4, pp. 5418–5440. DOI: 10.1093/mnras/stad3092.

Meinhardt-Llopis, E. et al. (2013). "Horn-Schunck Optical Flow with a Multi-Scale Strategy". In: *Image Processing On Line* 3, pp. 151–172. DOI: 10.5201/ipol.2013.20.

Mertens, F. and A. Lobanov (2015). "Wavelet-based decomposition and analysis of structural patterns in astronomical images". In: *A&A* 574, A67. DOI: 10.1051/0004-6361/201424566.

Mildenhall, B. et al. (2021). "NeRF: representing scenes as neural radiance

fields for view synthesis". In: *Commun. ACM* 65.1, pp. 99–106. ISSN: 0001-0782. DOI: 10.1145/3503250.

Müller, H. and A. P. Lobanov (2023). "Dynamic and Polarimetric VLBI imaging with a multiscalar approach". In: *A&A* 673, A151. DOI: 10.1051/0004-6361/202245393.

Mus, A. and I. Martí-Vidal (2024). "New-generation maximum entropy method: a Lagrangian-based algorithm for dynamic reconstruction of interferometric data". In: *MNRAS* 528.4, pp. 5537–5557. DOI: 10.1093/mnras/stae234.

Nagai, H. et al. (2014). "Limb-brightened Jet of 3C 84 Revealed by the 43 GHz Very-Long-Baseline-Array Observation". In: *ApJ* 785.1, p. 53. DOI: 10.1088/0004-637X/785/1/53.

Nagai, H. et al. (2017). "Enhanced Polarized Emission from the One-parsec-scale Hotspot of 3C 84 as a Result of the Interaction with the Clumpy Ambient Medium". In: *ApJ* 849.1, p. 52. DOI: 10.3847/1538-4357/aa8e43.

Nagai, H. et al. (2010). "VLBI Monitoring of 3C 84 (NGC 1275) in Early Phase of the 2005 Outburst". In: *PASJ* 62, p. L11. DOI: 10.1093/pasj/62.2.L11.

Napier, P. et al. (1994). "The Very Long Baseline Array". In: *Proceedings of the IEEE* 82.5, pp. 658–672. DOI: 10.1109/5.284733.

Narayan, R. et al. (2003). "Magnetically Arrested Disk: an Energetically Efficient Accretion Flow". In: *PASJ* 55, pp. L69–L72. DOI: 10.1093/pasj/55.6.L69.

Nesterov, N. S. et al. (1995). "Radio and optical evolution of the Seyfert galaxy NGC 1275". In: *A&A* 296, p. 628. https://ui.adsabs.harvard.edu/abs/1995A&A...296..628N.

O' Riordan, M. et al. (2018). "Observational Signatures of Mass-loading in Jets Launched by Rotating Black Holes". In: *ApJ* 853.1, p. 44. DOI: 10.3847/1538-4357/aaa0c4.

Oei, M. S. S. L. et al. (2024). "Black hole jets on the scale of the cosmic web". In: *Nature* 633.8030, pp. 537–541. DOI: 10.1038/s41586-024-07879-y.

Palumbo, D. C. M. et al. (2020). "Discriminating Accretion States via Rotational Symmetry in Simulated Polarimetric Images of M87". In: *ApJ* 894.2, p. 156. DOI: 10.3847/1538-4357/ab86ac.

Paraschos, G. F. et al. (2022). "Jet kinematics in the transversely stratified jet of 3C 84. A two-decade overview". In: *A&A* 665, A1. DOI: 10.1051/0004-6361/202243343.

Park, J. et al. (2024a). "Discovery of Limb Brightening in the Parsec-scale Jet of NGC 315 through Global Very Long Baseline Interferometry Observations and Its Implications for Jet Models". In: *ApJ L.* 973.2, p. L45. DOI: 10.3847/2041-8213/ad7137.

Park, J. et al. (2024b). "Observational evidence to support a dense ambient medium shaping the jet in 3C 84". In: *A&A* 685, A115. DOI: 10.1051/0004-6361/202347562.

Pauliny-Toth, I. I. K. and K. I. Kellermann (1966). "Variations in the Radio-Frequency Spectra of 3C 84, 3C 273, 3C 279, and Other Radio Sources". In: *ApJ* 146, p. 634. DOI: 10.1086/148941.

Pedlar, A. et al. (1990). "The radio structure of NGC 1275". In: *MNRAS* 246, p. 477. https://ui.adsabs.harvard.edu/abs/1990MNRAS.246..477P.

Penrose, R. (1965). "Gravitational Collapse and Space-Time Singularities". In: *Phys. Rev. Lett.* 14.3, pp. 57–59. DOI: 10.1103/PhysRevLett.14.57.

Perley, R. A. et al. (1984). "The jet and filaments in Cygnus A". In: *ApJ L.* 285, pp. L35–L38. DOI: 10.1086/184360.

Popescu, M.-C. et al. (2009). "Multilayer perceptron and neural networks". In: *WSEAS Trans. Cir. and Sys.* 8.7, pp. 579–588. ISSN: 1109-2734. https://dl.acm.org/doi/abs/10.5555/1639537.1639542.

Porth, O. et al. (2017). "The black hole accretion code". In: *Computational Astrophysics and Cosmology* 4.1, p. 1. DOI: 10.1186/s40668-017-0020-2.

Pötzl, F. M. et al. (2021). "Probing the innermost regions of AGN jets and their magnetic fields with RadioAstron. IV. The quasar 3C 345 at 18 cm: Magnetic field structure and brightness temperature". In: *A&A* 648, A82. DOI: 10.1051/0004-6361/202039493.

Pringle, J. E. (1997). "Self-induced warping of accretion discs: non-linear evolution and application to AGN". In: *MNRAS* 292.1, pp. 136–147. DOI: 10.1093/mnras/292.1.136.

Punsly, B. et al. (2021). "Observing the Time Evolution of the Multicomponent Nucleus of 3C 84". In: *ApJ* 911.1, p. 19. DOI: 10.3847/1538-4357/abe69f.

Pushkarev, A. B. et al. (2023). "MOJAVE - XX. Persistent linear polarization structure in parsec-scale AGN jets". In: *MNRAS* 520.4, pp. 6053–6069. DOI: 10.1093/mnras/stad525.

Rahaman, N. et al. (2019). "On the Spectral Bias of Neural Networks". In: *Proceedings of the 36th International Conference on Machine Learning*. Proceedings of Machine Learning Research 97. Ed. by K. Chaudhuri and R. Salakhutdinov, pp. 5301–5310. https://proceedings.mlr.press/v97/rahaman19a.html.

Raiteri, C. M. et al. (2017). "Blazar spectral variability as explained by a twisted inhomogeneous jet". In: *Nature* 552.7685, pp. 374–377. DOI: 10.1038/nature24623.

Reber, G. (1940). "Notes: Cosmic Static". In: *ApJ* 91, pp. 621–624. DOI: 10.1086/144197.

Röder, J. et al. (2024). "Up around the bend: A multiwavelength view of the quasar 3C 345". In: *A&A* 684, A211. DOI: 10.1051/0004-6361/202349038.

Roelofs, F. et al. (2023). "Polarimetric Geometric Modeling for mm-VLBI Observations of Black Holes". In: *ApJ L.* 957.2, p. L21. DOI: 10.3847/2041-8213/acff6f.

Rogers, A. E. E. et al. (1974). "The structure of radio sources 3C 273B and 3C 84 deduced from the "closure" phases and visibility amplitudes observed with three-element interferometers". In: *ApJ* 193, pp. 293–301. DOI: 10.1086/153162.

Rosenblatt, F. (1958). "The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain". In: *Psychological Review* 65.6, pp. 386–408. DOI: 10.1037/h0042519.

Rybicki, G. B. and A. P. Lightman (1979). *Radiative processes in astrophysics*. Harvard-Smithsonian Center for Astrophysics. https://ui.adsabs.harvard.edu/abs/1979rpa..book....R.

Ryle, M. (1962). "The New Cambridge Radio Telescope". In: *Nature* 194.4828, pp. 517–518. DOI: 10.1038/194517a0.

Ryle, M. et al. (1965). "High-Resolution Observations of the Radio Sources in Cygnus and Cassiopeia". In: *Nature* 205.4978, pp. 1259–1262. DOI: 10.1038/2051259a0.

Sadowski, A. et al. (2014). "Numerical simulations of super-critical black hole accretion flows in general relativity". In: *MNRAS* 439.1, pp. 503–520. DOI: 10.1093/mnras/stt2479.

Savolainen, T. et al. (2023). "RadioAstron discovery of a mini-cocoon around the

restarted parsec-scale jet in 3C 84". In: *A&A* 676, A114. DOI: 10.1051/0004-6361/202142594.

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

Seyfert, C. K. (1943). "Nuclear Emission in Spiral Nebulae". In: *ApJ* 97, p. 28. DOI: 10.1086/144488.

Stawarz, Ł. and M. Ostrowski (2002). "Radiation from the Relativistic Jet: A Role of the Shear Boundary Layer". In: *ApJ* 578.2, pp. 763–774. DOI: 10.1086/342649.

Steffen, W. et al. (1995). "A helical model for the compact jet in 3C345". In: *A&A* 302, p. 335. DOI: 10.48550/arXiv.astro-ph/9505075.

Strauss, M. A. et al. (1992). "A Redshift Survey of IRAS Galaxies. VII. The Infrared and Redshift Data for the 1.936 Jansky Sample". In: *ApJ S.* 83, p. 29. DOI: 10.1086/191730.

Sun, H. and K. L. Bouman (2021). "Deep Probabilistic Imaging: Uncertainty Quantification and Multi-modal Solution Characterization for Computational Imaging". In: *AAAI Conference on Artificial Intelligence (AAAI)*. DOI: 10.48550/arXiv.2010.14462.

Suzuki, K. et al. (2012). "Exploring the Central Sub-parsec Region of the γ-Ray Bright Radio Galaxy 3C 84 with VLBA at 43 GHz in the Period of 2002-2008". In: *ApJ* 746.2, p. 140. DOI: 10.1088/0004-637X/746/2/140.

Tancik, M. et al. (2020). "Fourier features let networks learn high frequency functions in low dimensional domains". In: *Advances in neural information processing systems* 33, pp. 7537–7547. https://dl.acm.org/doi/abs/10.5555/3495724.3496356.

Tavecchio, F. and G. Ghisellini (2014). "On the spine-layer scenario for the very high-energy emission of NGC 1275". In: *MNRAS* 443.2, pp. 1224–1230. DOI: 10.1093/mnras/stu1196.

Thompson, A. R. et al. (1980). "The Very Large Array". In: *ApJ S.* 44, pp. 151–167. DOI: 10.1086/190688.

Thompson, A. R. et al. (2017). *Interferometry and Synthesis in Radio Astronomy, 3rd Edition*. Springer Cham. DOI: 10.1007/978-3-319-44431-4.

Tiede, P. (2022). "Comrade: Composable Modeling of Radio Emission". In: *The Journal of Open Source Software* 7.76, p. 4457. DOI: 10.21105/joss.04457.

Twiss, R. Q. et al. (1960). "Brightness distribution over some strong radio sources at 1427 Mc/s". In: *The Observatory* 80, pp. 153–159. https://ui.adsabs.harvard.edu/abs/1960Obs...80..153T.

Unwin, S. C. et al. (1983). "Superluminal motion in the quasar 3C 345". In: *ApJ* 271, pp. 536–550. DOI: 10.1086/161220.

Urry, C. M. and P. Padovani (1995). "Unified Schemes for Radio-Loud Active Galactic Nuclei". In: *PASP* 107, p. 803. DOI: 10.1086/133630.

Vaswani, A. et al. (2017). "Attention Is All You Need". In: *Advances in Neural Information Processing Systems* 30. https://dl.acm.org/doi/10.5555/3295222.3295349.

Vermeulen, R. C. and M. H. Cohen (1994). "Superluminal Motion Statistics and Cosmology". In: *ApJ* 430, p. 467. DOI: 10.1086/174424.

Walker, R. C. et al. (2000). "VLBA Absorption Imaging of Ionized Gas Associated with the Accretion Disk in NGC 1275". In: *ApJ* 530.1, pp. 233–244. DOI: 10.1086/308372.

Walker, R. C. et al. (2018). "The Structure and Dynamics of the Subparsec Jet in M87 Based on 50 VLBA Observations over 17 Years at 43 GHz". In:

ApJ 855.2, p. 128. DOI: 10.3847/1538-4357/aaafcc.

Weaver, Z. R. et al. (2022). "Kinematics of Parsec-scale Jets of Gamma-Ray Blazars at 43 GHz during 10 yr of the VLBA-BU-BLAZAR Program". In: *ApJ S.* 260.1, p. 12. DOI: 10.3847/1538-4365/ac589c.

Weiß, A. et al. (2008). "LABOCA observations of nearby, active galaxies". In: *A&A* 490.1, pp. 77–86. DOI: 10.1051/0004-6361:200809909.

Wielebinski, R. et al. (2011). "The Effelsberg 100-m Radio Telescope: Construction and Forty Years of Radio Astronomy". In: *Journal of Astronomical History and Heritage* 14.1, pp. 3–21. https://ui.adsabs.harvard.edu/abs/2011JAHH...14....3W.

Wielgus, M. et al. (2022). "Millimeter Light Curves of Sagittarius A* Observed during the 2017 Event Horizon Telescope Campaign". In: *ApJ L.* 930.2, p. L19. DOI: 10.3847/2041-8213/ac6428.

Wilkinson, P. N. (1995). "Very-Long-Baseline Radio Interferometry Surveys of the Compact Structure in Active Galactic Nuclei". In: *Proceedings of the National Academy of Science* 92.25, pp. 11342–11347. DOI: 10.1073/pnas.92.25.11342.

Wong, G. N. et al. (2022). "PATOKA: Simulating Electromagnetic Observables of Black Hole Accretion". In: *ApJ S.* 259.2, p. 64. DOI: 10.3847/1538-4365/ac582e.

Wootten, A. and A. R. Thompson (2009). "The Atacama Large Millimeter/Submillimeter Array". In: *IEEE Proceedings* 97.8, pp. 1463–1471. DOI: 10.1109/JPROC.2009.2020572.

Xie, Y. et al. (2022). "Neural fields in visual computing and beyond". In: *Computer Graphics Forum* 41.2, pp. 641–676. DOI: 10.1111/cgf.14505.

Zensus, J. A. (1997). "Parsec-Scale Jets in Extragalactic Radio Sources". In: *Annual Rev. Astron. Astrophys.* 35, pp. 607–636. DOI: 10.1146/annurev.astro.35.1.607.

Zernike, F. (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.

Zhao, B. et al. (2024). "Single View Refractive Index Tomography with Neural Fields". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 25358–25367. DOI: 10.48550/arXiv.2309.04437.

Zhong, E. D. et al. (2021). "CryoDRGN: reconstruction of heterogeneous cryo-EM structures using neural networks". In: *Nat. Meth.* 18.2, pp. 176–185. DOI: 10.1038/s41592-020-01049-4.

Appendix

Selected frames from the `kine` video reconstructions presented in Chapter 5. The top row shows the ground truth model, the bottom one the reconstruction.

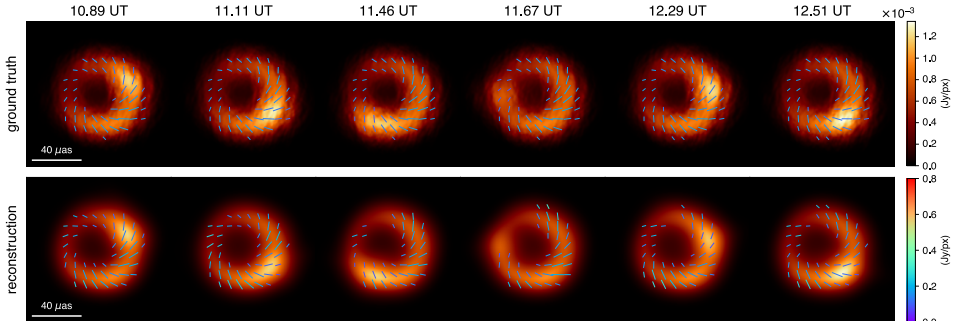

Figure 5.15: **`mring+hsCW` reconstruction with `kine`.**

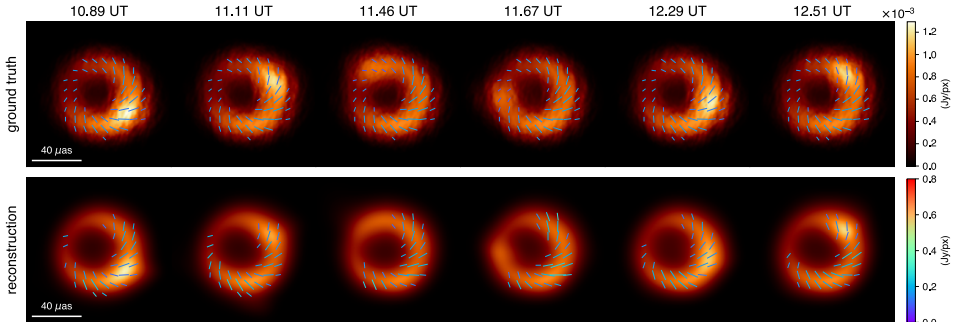

Figure 5.16: **`mring+hsCCW` reconstruction with `kine`.**

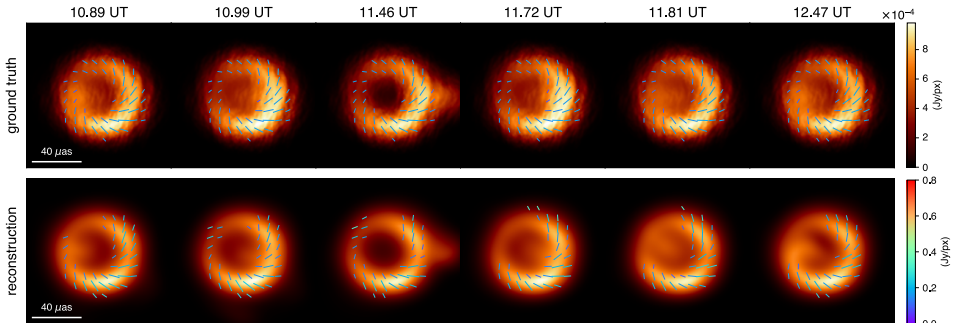

Figure 5.17: **mring+hs-cross reconstruction with kine.**

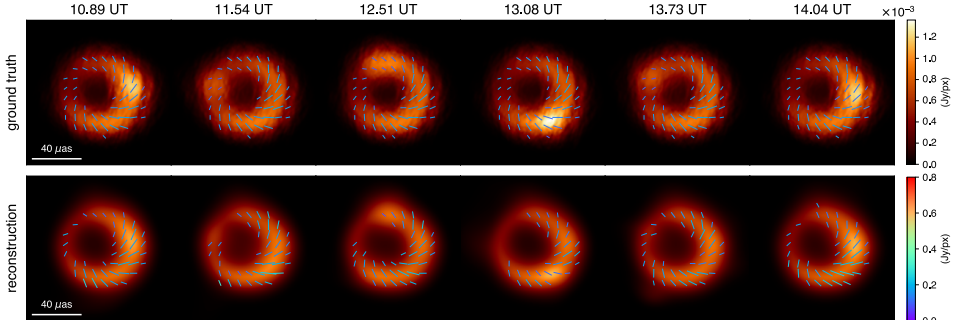

Figure 5.18: **mring+hs-incoh reconstruction with kine.**

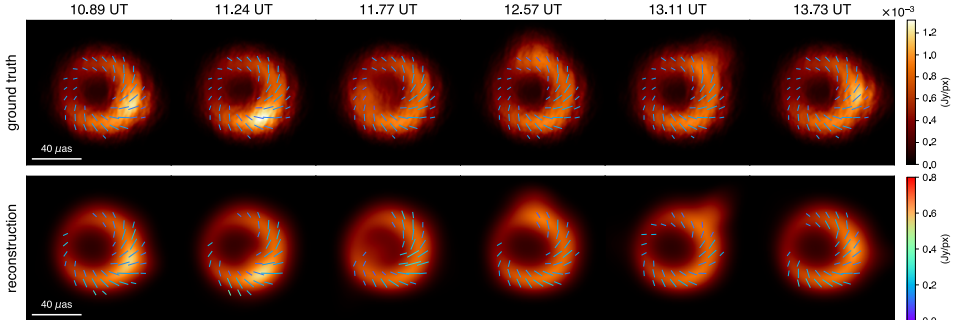

Figure 5.19: **mring+hs-not-center reconstruction with kine.**

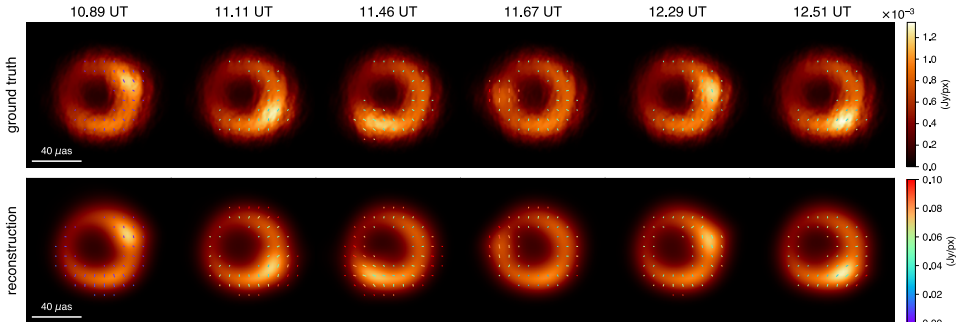

Figure 5.20: **mring+hs-pol reconstruction with `kine`.**

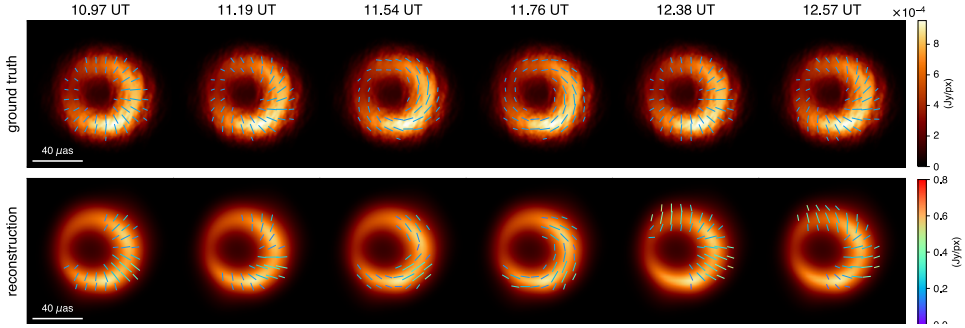

Figure 5.21: **mring-varbeta2 reconstruction with `kine`.**

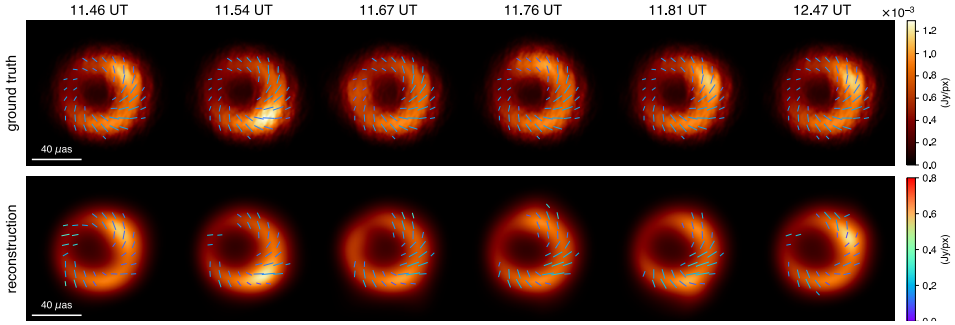

Figure 5.22: **mring+hsCW20 reconstruction with `kine`.**

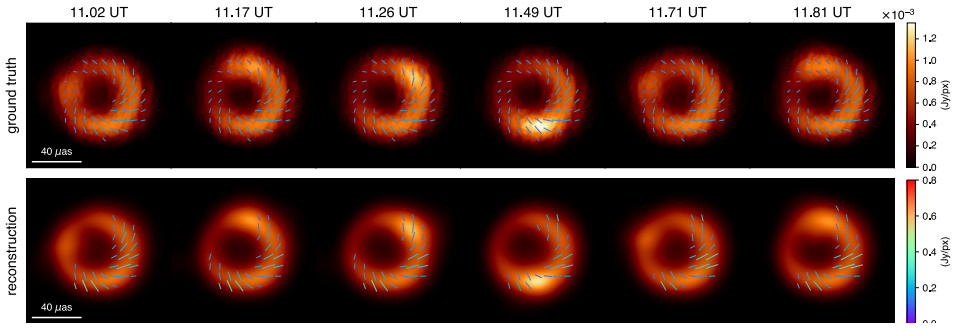

Figure 5.23: **mring+hsCW40 reconstruction with kine.**

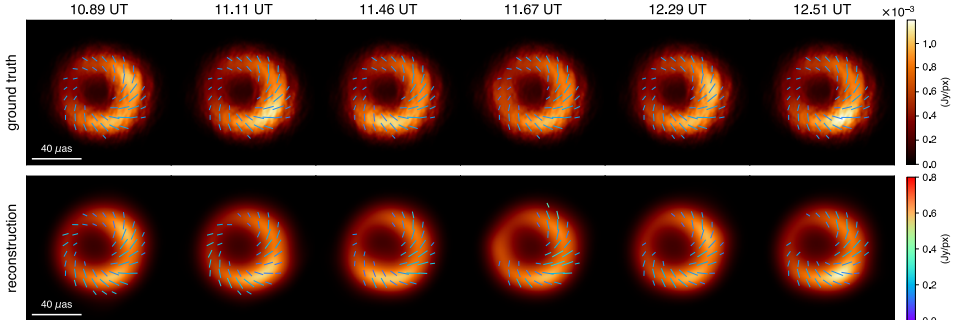

Figure 5.24: **mring+hsCW0.15 reconstruction with kine.**

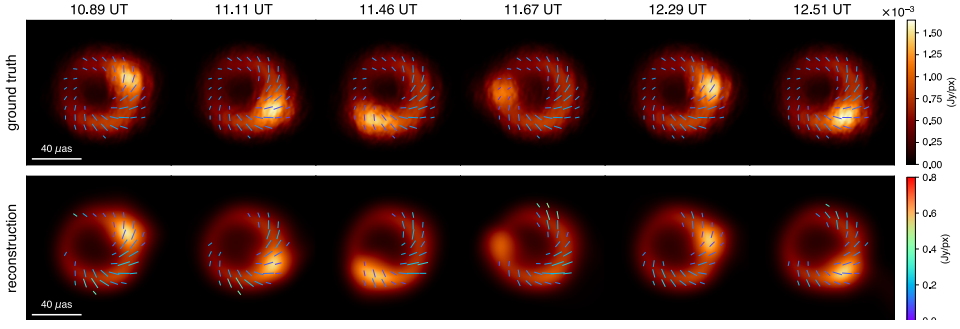

Figure 5.25: **mring+hsCW0.60 reconstruction with kine.**

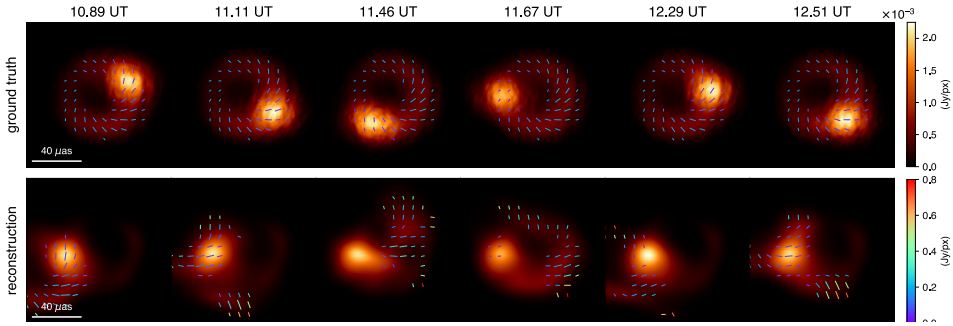

Figure 5.26: **mring+hsCW1.20 reconstruction with kine.**

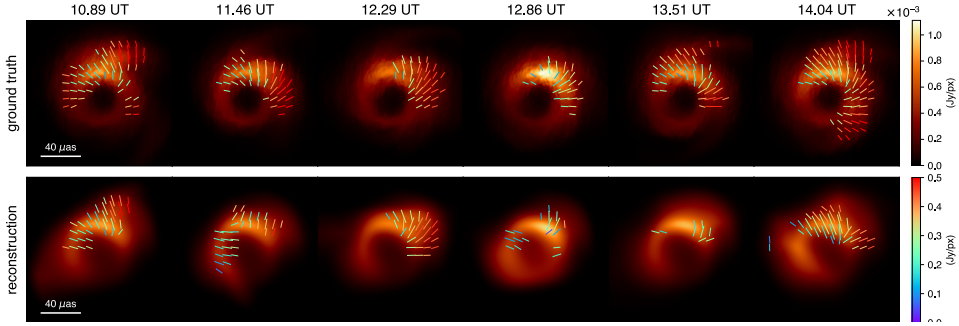

Figure 5.27: **grmhd1 reconstruction with kine.**

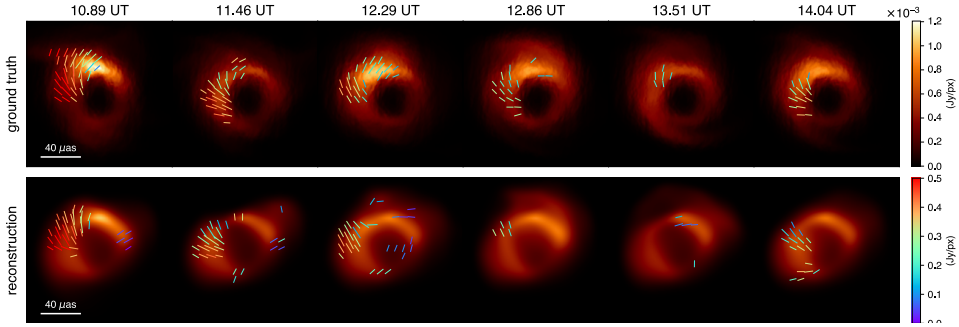

Figure 5.28: **grmhd2 reconstruction with kine.**

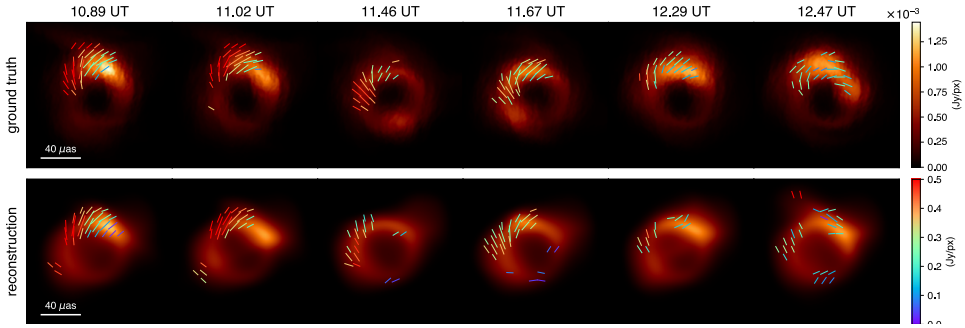

Figure 5.29: **grmhd2+hs2 reconstruction with `kine`.**

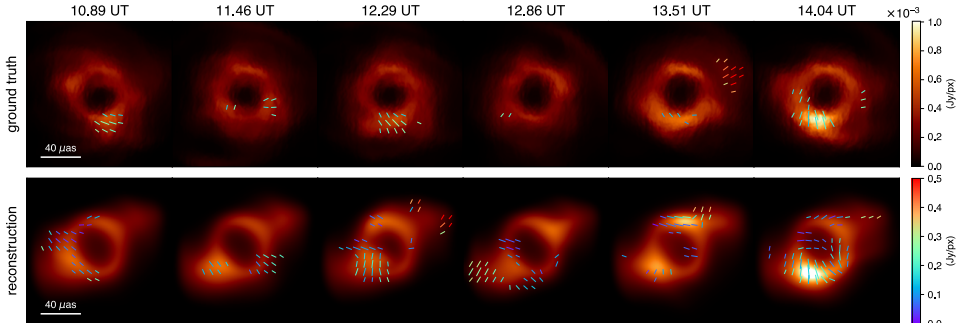

Figure 5.30: **grmhd3 reconstruction with `kine`.**

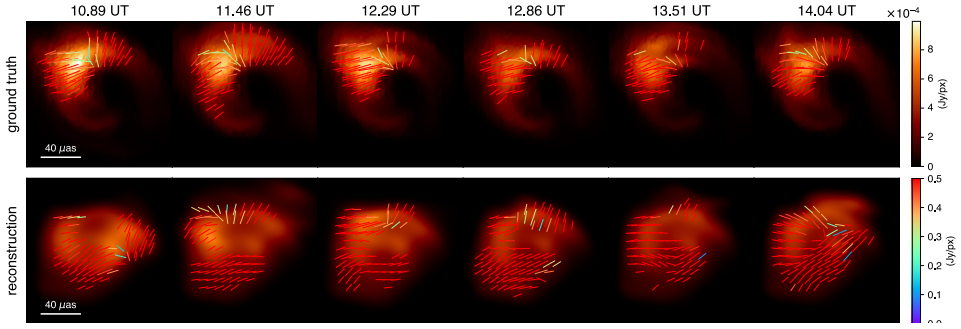

Figure 5.31: **grmhd4 reconstruction with `kine`.**

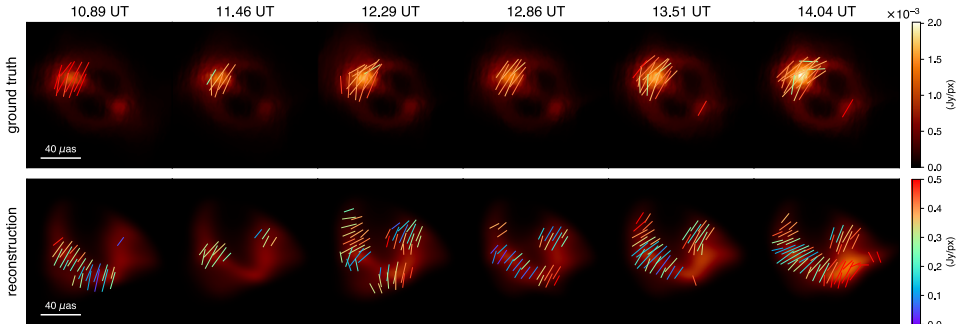

Figure 5.32: **grmhd5 reconstruction with kine.**

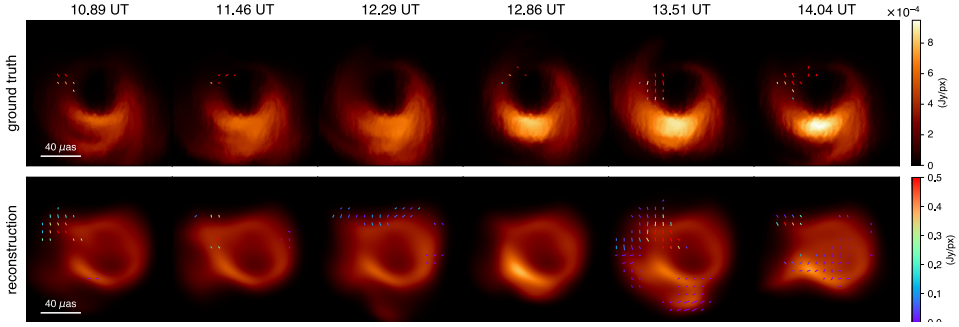

Figure 5.33: **grmhd6 reconstruction with kine.**

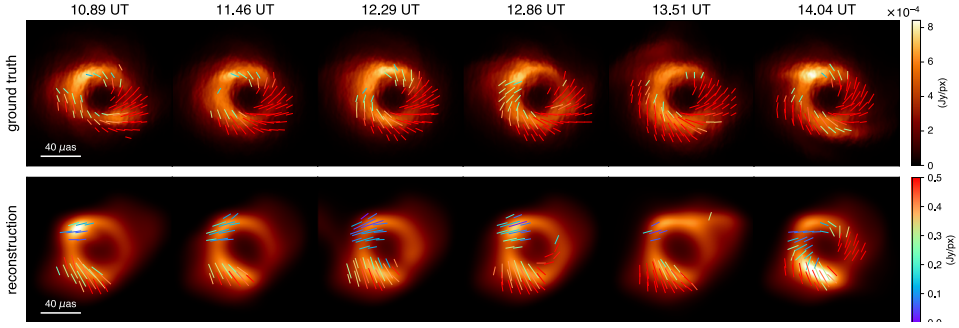

Figure 5.34: **grmhd7 reconstruction with kine.**

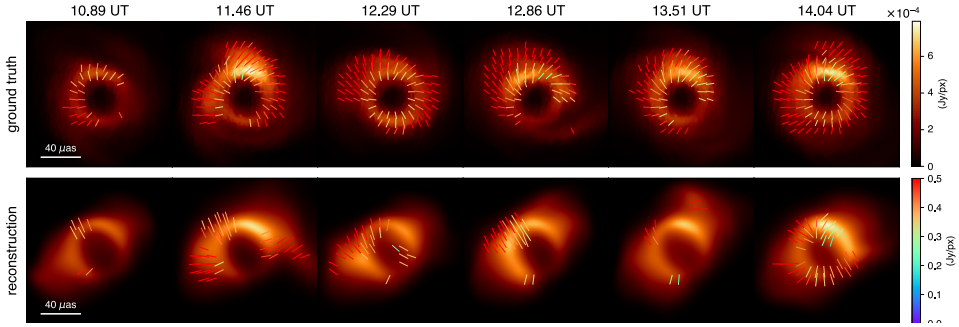

Figure 5.35: **grmhd8 reconstruction with kine.**

UNIVERSIDAD
DE GRANADA

INSTITUTO DE
ASTROFÍSICA DE
ANDALUCÍA

EXCELENCIA
SEVERO
OCHOA

CSIC

UNIVERSIDAD DE GRANADA
IAA INSTITUTO DE ASTROFÍSICA DE ANDALUCÍA
EXCELENCIA SEVERO OCHOA
CSIC