MASTERARBEIT | MASTER'S THESIS

Titel | Title

A high precision mass measurement of the central black hole in the massive elliptical galaxy NGC 2513 with ALMA molecular gas dynamics

verfasst von | submitted by

Shelley-Anne Deborah Harrisberg BSc

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of

Master of Science (MSc)

Wien | Vienna, 2025

Studienkennzahl lt. Studienblatt | Degree programme code as it appears on the student record sheet:	UA 066 861
Studienrichtung lt. Studienblatt | Degree programme as it appears on the student record sheet:	Masterstudium Astronomie
Betreut von | Supervisor:	Dipl.-Phys. Dr. Sabine Thater

Acknowledgements

This Master thesis has been a very special time for me and it would be untoward of me not to acknowledge the many people that have made its completion a possibility.

First and foremost, I would like to thank my supervisor, Dr Sabine Thater. Over the course of this project, Sabine's knowledge and expertise of the field of black hole mass measurements via dynamical means, together with her understanding of the scaling relations that are so important in motivating our interest in these objects, has proved invaluable. Moreover, her passion for the field and continuing curiosity in new things to be learned have been an inspiration. As a supervisor, she has provided support over and beyond her duties with regular meetings, constructive criticism and feedback. Her patience and willingness to explain the simplest things enabled me to grow in confidence and never be afraid to ask questions. They opened the doors of learning. There is so much more to be said, but, in short: Thank you Sabine for everything!

I would next like to thank my co-supervisor, Dr Dieu D. Nguyen. Without Dieu, this project could not have taken place. Dieu was the Principle Investigator responsible for initiating and submitting the project to gain ALMA observation time of the molecular gas in NGC 2513. Besides providing this data, Dieu both wrote and kindly provided the code that I initially used to produce the first moment maps of the molecular gas disc and is also the author of the code used in the modelling that was used in this thesis. Dieu's expertise in working with ALMA data and in the dynamical modelling of molecular gas has proved of immeasurable help. I would like to thank him for the many meetings that we had where he assessed the work I had done and provided advice for improving the imaging and quality of the moment maps. Thank you, Dieu, for your guidance, patience and knowledge. It has been an inspiration and a privilege!

Another person without whose help I could not have completed this project is Hai N. Ngo. Hai not only carried out the dynamical modelling that produced our final results, but was instrumental in giving me access to the modelling code and final data products so that I could understand the process and produce my own plots of the results. In addition, he introduced me to the code used to create the final moment maps and PVD, and kindly produced some of the plots included in this thesis. Hai most willingly shared his excellent knowledge of Bayesian forward modelling and, in particular, of the steps taken to do the modelling of the molecular disc in NGC 2513. He was always ready with clear, concise answers to questions that I had and to assist in any way possible. Thank you, Hai, for your patience, readiness and amazing work!

I would like to offer special thanks to Ass. Prof. Dr Alvaro Hacar who provided me with an account on the Institute's "ALMA" computer where I was able to carry out the imaging of the ALMA data sets with the software package CASA. He also prompted me to join ALMA's European Chapter, assisting in the application process, so that I could gain direct access to the data and to updates on ALMA news. Although working in a different field to mine, he provided me with many useful tips on masking, combining the data sets and imaging. I also gained so much valuable knowledge from his seminars on ALMA, and, quite apart, learned from him that molecules are even more exciting than I thought. His knowledge of, and passion for radio astronomy, the ISM and ALMA have been so valuable to me.

A very special contribution to this project comes from Dr Alice Zocchi. Alice stepped in as an additional proof-reader and assisted in a supervisory way as writing got underway. Alice gave this project time and energy far exceeding any expectations. She was ready with bi-weekly meetings, fixing LaTeXproblems, streamlining texts, fixing my log notation, offering advice and offering highly constructive critical feedback. Her in-depth knowledge of MCMC modelling and excellent ability to explain things were of invaluable assistance, including the penny-dropping moment when I finally understood how walkers actually walk! Her endless patience in providing a sympathetic ear assisted so much in overcoming so many obstacles. Thank you, Alice. You are a gem!

A special thanks goes to Prof. Glenn van de Ven, head of the research group The Dynamics of Stellar Systems, and to all the amazing group members. Not only have I been a welcome member of the group since my Bachelor days, but the group has been an important platform for me to present and share the results of my Master project. In particular, I would like to thank Glenn for providing the funding that enabled me to go on my first ever conference, the YAGN meeting in 2023, and, in addition, for me to travel to South Africa to attend the IAU GA 2024 that was held in Cape Town. In both cases, I was privileged to represent our Institute and to have the opportunity to present my results in a public sphere.

A special thanks likewise goes to Prof. Bodo Ziegler, head of the Extragalactic Astrophysics research group. Bodo not provided me with a work-space, but also provided funding for my participation in the AECC 2024 organized by the ÖGAA where, again, I was able to present my work publicly. As part of his group, I was given the opportunity to present ongoing results at group meetings, gaining immensely from his depth of knowledge, experience and insightful criticism. Thank you Bodo!

I am much indebted to Francisco Aros for providing me with the Overleaf File used to type-set and produce this thesis! Thank you, Francisco, you were a life-saver!

Many thanks go to my fellow supermassive black hole Master students Malika Duffek and Anirban Roy Chowdhury with whom I also shared a work-space, the ups and downs of our projects and from whom I received practical tips and much encouragement. Also thanks to Maria Koller, Léah Skuza and Julia Lamprecht for encouragement, tips, company and advice. Special mention also to Iva Simonovic who helped me "pull myself together", and to Johannes Müller and Roland Bergauer, who were there from the start.

Last but not least, I thank my boyfriend and partner Mike Bremer who has always believed in me and who has never stopped encouraging me in my studies or in the hope that my research will one day be translated into Latin: De rationibus gurgitum atrissimorum gigantissimorumque ingentissima mole vel gravitate praeditorum. Whatever, happens, Mike, "Et hätt noch immer jot jejange!"

I dedicate this thesis to the memory of my parents Clifford Simon Harrisberg and Marjorie Selma Harrisberg.

For my father, Clifford ... How we used to seek out a dark patch, looking for ways to escape the light and gaze upon the other lights up in the big black sky.

For my mother, Marjorie ... As you watched the Sun sink to the horizon from your room, you always said to me: It's bigger than all of us.

Kurzfassung

Supermassive schwarze Löcher (SMBHs) befinden sich in den Zentren der meisten Galaxien. Mit Massen im Bereich von $10^5 - 10^{10}\,M_\odot$ ist ihr Gravitationspotenzial beträchtlich, doch dies erstreckt sich nur über den innersten Bereich der Galaxie. Es gibt trotzdem mehrere empirische Korrelationen zwischen SMBHs und verschiedenen Eigenschaften ihrer Muttergalaxien, die weit über den Einflussbereich des SMBHs hinausreichen. Dies lässt auf eine Art Koevolution schließen. Die stärkste dieser Korrelationen ist die Skalierungsbeziehung zwischen der Masse des SMBH und der effektiven stellaren Geschwindigkeitsdispersion $M_{BH} - \sigma_*$. Um diese Korrelationen besser zu verstehen und somit einen umfassenderen Einblick in die Evolution der Galaxien zu erhalten, ist ein großer und robuster Datensatz von SMBH Massenmessungen erforderlich, insbesondere im Bereich niedriger und hoher Massen, wo es noch relativ wenige Messungen gibt.

In dieser Arbeit stelle ich die Ergebnisse der ersten Massenmessung des SMBH in der massereichen elliptischen Galaxie NGC 2513 vor. Die Messung wurde mit der Methode der kalten molekularen Gasdynamik unter Verwendung von hochauflösenden Beobachtungen des Atacama Large Millimetre/sub-millimetre Array (ALMA) durchgeführt. Zunächst modelliere ich die Leuchtmasse der Galaxie, indem ich die Multi-Gauß‘sche Expansionsmethode (MGE) auf ein hochauflösendes Bild der Galaxie des Hubble Space Telescope (HST) anwende. Dann reduziere ich die ALMA-Daten mit dem CASA-Paket (Common Astronomy Software Applications), um einen Datacube zu erhalten, der in die KINematic Molecular Simulation (`KinMS`)-Routinen zur Modellierung der molekularen Gasdynamik in einem Bayes'schen Rahmen eingespeist werden kann. Ich erhalte eine Best-Fit-Messung für die SMBH-Masse von $(1.23^{+0.14}_{-0.23}) \times 10^9\ M_\odot$ und eine entsprechende J-Band Masse-Leuchtkraft-Beziehung M/L von $2.00^{+1.82}_{-1.15}\ M_\odot/L_{\odot,J}$, beide auf dem Konfidenzniveau von 68 %. Die Massenmessung ist konsistent mit der $M_{BH} - \sigma_*$-Relation und liegt gut innerhalb der Streuung. Als erste SMBH-Massenmessung, die an dieser Galaxie durchgeführt wurde, stellt dieses Ergebnis einen wichtigen Beitrag zur Datenbank der SMBH-Massenmessungen im Allgemeinen dar, und darüber hinaus zur wachsenden Anzahl von Messungen, die mit kalter Molekulargasdynamik durchgeführt wurden.

Abstract

Supermassive black holes (SMBHs) have been found to exist in the centres of most galaxies. With masses in the range of $10^5 - 10^{10}\,\mathrm{M}_\odot$ their gravitational potential is significant, yet it extends only over the innermost region of the galaxy. Despite this, there exist several empirical correlations between SMBHs and various properties of their host galaxies that extend well beyond the SMBH's sphere of influence, thereby suggesting some form of co-evolution. The strongest of these is the scaling relation between the SMBH mass and effective stellar velocity $M_{BH} - \sigma_*$. In order to better understand these correlations and thus gain greater insight into galaxy evolution, a large and robust dataset of SMBH mass measurements is required, particularly in the low- and high-mass regime where measurements are still relatively few.

In this thesis, I present the results of the first mass measurement of the SMBH in the massive elliptical galaxy NGC 2513. The measurement is obtained via the method of cold molecular gas dynamics using high-resolution observations from the Atacama Large Millimetre/sub-millimetre Array (ALMA). I first model the luminous mass of the galaxy by applying the multi-Gaussian expansion (MGE) method on a high-resolution Hubble Space Telescope (HST) image of the galaxy. I then image the ALMA data with the Common Astronomy Software Applications (CASA) package to obtain a velocity cube that can be fed into the KINematic Molecular Simulation (`KinMS`) routines for forward modelling in a Bayesian framework of the molecular gas dynamics. I obtain a best-fit measurement for the SMBH mass of $(1.23^{+0.14}_{-0.23}) \times 10^9\ \mathrm{M}_\odot$ and corresponding J-band mass-to-light ratio M/L of $2.00^{+1.82}_{-1.15}\ \mathrm{M}_\odot/\mathrm{L}_{\odot,J}$, both at the 68 % confidence level. The mass measurement is consistent with the $M_{BH} - \sigma_*$ relation, lying well within the scatter. As the first SMBH mass measurement to be conducted on this galaxy, this result presents an important contribution to the database of SMBH mass measurements in general and, moreover, to the growing sample of measurements conducted with cold molecular gas dynamics.

Contents

List of Tables

List of Figures

1. Introduction

Supermassive black holes (SMBHs) are some of the most intriguing objects in the Universe. Starting out as mere theoretical predictions, they have since been found in the centres of most massive galaxies (with masses $\geq 10^{10}$ $M_\odot$; see e.g., Kormendy & Ho 2013), while mounting evidence suggests that they also frequently exist in the centres of lower mass galaxies ($\leq 10^{10}$ $M_\odot$; see e.g., Reines & Volonteri 2015; Nguyen et al. 2018). In both cases, this is regardless of the morphological type of the host galaxy. With masses in the range of $10^5 - 10^{10}$ $M_\odot$ (see e.g., Kormendy & Ho 2013), SMBHs are extremely compact objects with spatial extents that are tiny (on the scale of astronomical units) in comparison to that of their host galaxies (on the scale of kiloparsecs). Moreover, their gravitational influence, while significant, extends only over the innermost region of the galaxy (typically on scales of parsecs). Despite this, several strong empirical correlations have been found to exist between SMBHs and various properties of their host galaxies that suggest some form of co-evolution (see e.g., Magorrian et al. 1998; Ferrarese & Merritt 2000).

This co-evolution most likely involves self-regulating mechanisms such as feedback from Active Galactic Nuclei (AGN) often coupled with mergers (Silk & Rees 1998; Di Matteo et al. 2008; Kormendy & Ho 2013). While theoretical models of galaxy evolution support this, the exact nature of this co-evolution is not fully understood, nor is the formation process of the SMBHs themselves. Furthermore, the correlations have proved to be more complex than originally thought, with evidence that suggests that SMBHs correlate differently with their host galaxy depending on the host's morphology and mass (e.g., McConnell & Ma 2013), as well as the respective components where the SMBH resides (such as disc, bulge or pseudo-bulge, bar, etc.; e.g., Kormendy & Ho 2013; Saglia et al. 2016).

Gaining a more thorough understanding of these nuanced relations between SMBH masses M_{BH} and the properties of their host galaxies should enable a greater insight into galaxy evolution (e.g., Ferrarese & Merritt 2000) and, possibly, the formation of the SMBHs themselves. To achieve this, a large and robust dataset of SMBH mass measurements over the entire range of galaxy morphologies and masses is required. While the existing dataset of SMBH measurements lies in the hundreds (see e.g., Bosch 2016), measurements in the high- and, especially, low-mass ranges of SMBHs are still relatively under-represented. In turn, it is also vital to have accurate and well-calibrated methods of measuring SMBH masses across the spectrum of galaxy morphologies and masses (Davis et al. 2017a). This project aims to add a new measurement to the dataset of SMBH mass measurements in high mass galaxies and, in particular, to add to the dataset of measurements made using cold molecular gas dynamics.

In this Chapter, I introduce the concepts that underlie the method of using spatially resolved dynamics to measure SMBH masses in Section 1.1 and described in-depth the method of cold molecular gas dynamics in Section 1.2. Section 1.3 outlines the various correlations that exist between SMBHs and their host-galaxies, focusing, in particular, on the scaling relation between the SMBH mass and stellar velocity dispersion. In Section 1.4 I introduce and explore the target galaxy of this project, NGC 2513, also deriving several SMBH mass estimates for the putative SMBH of NGC 2513 by using the scaling relation.

1.1. Dynamical methods

1.1.1. General concept

Dynamical modelling of the spatially resolved kinematics of a selected tracer has revealed itself to be one of the most accurate methods of measuring SMBH masses (cf. e.g., Davis et al. 2017a).[1] The method rests on the principle of modelling the kinematics of the tracer within the region where the gravitational potential of the SMBH dominates. This region is commonly referred to as the SMBH's sphere of influence and is characterized by its radius, R_{SOI}. This is most usually defined as

$$R_{SOI} = \frac{GM_{BH}}{\sigma_*^2} \ , \tag{1.1}$$

as derived by Peebles (1972), where G is the gravitational constant, M_{BH} is the mass of the SMBH, and σ_* is the velocity dispersion of the stellar population of the bulge or central region of the galaxy (see Section 1.3 and, e.g., Ferrarese & Ford (2005)). Essentially it can thus be described as the radius within which the circular velocity of the tracer resulting from the potential of the SMBH increases above the surrounding stellar velocity dispersion (Barth et al. 2016b).[2]

The motion of any material at distances further than a few thousand Schwarzschild radii from the SMBH (so that relativistic effects can be ignored), but within the sphere of influence, has been shown to be principally Keplerian (see e.g., Ferrarese & Ford 2005). This means that the orbital velocity of material within the sphere of influence can be modelled by assuming circular velocity. While it is the gravitational potential of the SMBH that is the dominant driving force of this motion, the combined gravitational potential of stars, gas, dust, dark matter and any other material within the region will also contribute. Thus, dynamical modelling of a tracer within this region will determine the total gravitational potential in the vicinity. To extract the contribution to the potential by the SMBH, the contributions by the luminous mass (stars, gas and dust) and dark matter need to be modelled and then subtracted from the total mass. This is illustrated in the following relations

$$\begin{aligned} M_{tot} &= M_{lum} + M_{BH} + M_{DM} \\ \Rightarrow M_{BH} &= M_{tot} - M_{lum} - M_{DM} \ , \end{aligned} \tag{1.2}$$

[1] For a summary of non-dynamical methods, especially as applicable to AGN, see e.g., Gliozzi et al. (2024); Casura et al. (2024) and references therein.

[2] An alternative definition of the R_{SOI} given in Merritt (2004) is the radius at which the enclosed stellar mass M_* is twice that of the SMBH mass: $M_*(r < R_{SOI}) = 2M_{BH}$.

where M_{tot} is the total mass, M_{lum} is the luminous mass and M_{DM} is the dark matter mass (Ferrarese & Ford 2005; Thater 2019). The contribution of dark matter in the innermost region of a galaxy is small enough to usually be ignored in the modelling and is, in fact, often quite difficult to constrain this close to the galactic centre. It is thus most commonly accounted for in the error budget of the SMBH mass. The contribution of the luminous mass, however, is significant and must be accounted for. Using high-resolution optical or infrared images, a model of the two-dimensional surface brightness is created and then deprojected to obtain the three-dimensional intrinsic luminosity density. This in turn is converted to a mass density by multiplication with a mass-to-light ratio M/L that is adopted for the galaxy (see e.g., Ferrarese & Ford 2005, and Chapters 2 and 4 for detailed discussion.).

The influence of the SMBH on the kinematics of any tracer drops off rapidly outside the sphere of influence, and the kinematics of the tracer beyond this region is not driven solely by the SMBH, but by other material, mainly stars. One of the key issues in measuring SMBH masses by dynamical means is thus to resolve the R_{SOI}. Indeed, Kormendy & Ho (2013) report on how SMBH mass measurements can differ quite substantially, depending on whether the R_{SOI} is resolved or not. It has, however, since been shown and argued that resolving the R_{SOI} is not a rigid criterion for reasonably accurate measurements of SMBH masses. For measurements using stellar kinematics, the use of three-dimensional integral-field spectroscopy (IFS) over two-dimensional long-slit spectroscopy has significantly expanded the available kinematic information and allowed successful measurements of SMBH masses below the R_{SOI} resolution limit (e.g., Krajnović et al. 2009; Thater et al. 2017). For measurements made with cold molecular gas kinematics, Davis (2014) has argued that the characteristic enhancement of circular velocity deriving from the SMBH's potential can be measured at radii greater than the R_{SOI} and that this criterion should be modified for SMBH mass measurements with molecular gas.[3] A more recent analysis of the R_{SOI} and alternative metrics that can quantify the physical scales probed by kinematic tracers is given in Zhang et al. (2024).

Nonetheless, high-resolution data remains essential, not only for modelling the kinematics of the tracer, but also for modelling the luminous mass as described above. Thus, the introduction of the Hubble Space Telescope (HST)[4] and of ground-based telescopes, especially those in the 10-m class, that use adaptive optics[5] has been instrumental in improving resolution for observations in the optical and infrared wavelengths (e.g., Kormendy & Ho 2013; Thater 2019). In this respect, the newly launched James Webb Space Telescope (JWST) and forthcoming Extremely Large Telescope (ELT) should contribute

[3] For a full discussion on the issue and alternative resolution criteria for cold molecular gas measurements, see Davis (2014) For further discussion on the definition of the R_{SOI} and its use as a criterion for selecting galaxies that have measurable SMBH masses, see e.g., Merritt (2004); Gültekin et al. (2009); Batcheldor (2010).

[4] The HST, as a space telescope, is free of any distortions originating in the Earth's atmosphere that commonly affect imaging. In terms of resolution, Kormendy & Ho (2013) comment that it offered approximately five times higher spatial resolution than ground-based telescopes of the time.

[5] Adaptive optics provides real-time correction for image distortions caused by atmospheric turbulence using a sophisticated system of deformable mirrors, wavefront sensors and computers (see, for example, `https://elt.eso.org/telescope/adaptiveoptics/`).

significantly (see e.g., Nguyen et al. 2025). Similarly, for tracers that are observed in the radio regimes, the use of radio interferometry and the construction of ever better radio interferometers has been game-changing.

1.1.2. Dynamical tracers

A variety of tracers have been used to successfully measure the kinematic signature associated with the presence of a SMBH at a galaxy's centre. A good fraction of SMBH detections and SMBH mass measurements have been carried out using stellar kinematics (see e.g., Kormendy & Ho 2013; Bosch 2016; Thater 2019). In the Milky Way, this is possible using the resolved orbits of individual stars (Ghez et al. 2008; Gillessen et al. 2009). In external galaxies, where the spatial resolution of individual stars is not possible, measurements are made with integrated stellar kinematics using either long-slit (e.g., Dressler & Richstone 1988; Magorrian et al. 1998) or integral field spectroscopy (e.g., Emsellem et al. 1999; Verolme et al. 2002; Thater et al. 2019).

Besides stars, gas in various forms can be used to trace gravitational potentials. A particularly accurate tracer is water megamaser emission from the circumnuclear accretion discs of some AGNs (e.g., Miyoshi et al. 1995; Greene et al. 2010).[6] While providing extremely precise measurements, this technique is applicable only to SMBHs in AGN that have masing material. Ionized gas, on the other hand, is far more common in the Universe and has been extensively used as a tracer (e.g., Ferrarese et al. 1996; Sarzi et al. 2001), with the majority of SMBH measurements in the Bosch (2016) sample coming from ionized gas (see also Thater 2019). While the dynamical modelling of gas kinematics is less expensive and simpler than that of stellar kinematics, accurate measurements of warm gas, as discussed briefly below, can be very problematic owing to its behaviour as a collisional fluid (Kormendy & Ho 2013). The use of cold molecular gas as a tracer is a relatively recent development and is fast proving to be a viable and useful technique.

1.2. Cold Molecular Gas Dynamics

The employment of cold molecular gas to measure SMBH masses was pioneered by Davis et al. (2013) and a figure of merit for it was presented in Davis (2014). This technique uses, as its tracers, asymmetrical molecular gases such as CO, HCO^+ and HCN. These types of molecules are capable of undergoing rotational transitions that result in emission in the millimetre wavelength regime.[7] It was the substantial improvement in observational capacity in these wavelengths that enabled molecular gas kinematics to become a viable method of SMBH mass measurement. To achieve the spatial resolutions necessary to

[6] Masers are the microwave equivalent of lasers. Here, maser stands for "microwave amplification by stimulated emission of radiation".

[7] For more details on the physics of molecular transitions, see, e.g., https://www.cv.nrao.edu/~sransom/web/Ch7.html and https://vallance.chem.ox.ac.uk/pdfs/MolecularEnergyLevelsNotes.pdf, or see Harrisberg (2020) for a summary specific to the use of molecular gas emission for SMBH mass measurements.

make SMBH mass measurements, observations in these wavelengths need to be carried out by interferometers (see Section 3.1).

The first such interferometer to offer sufficient spatial resolution in these wavelengths was the Combined Array for Research in Millimetre Astronomy (CARMA), which was in service from 2006 to 2015 and which, at a frequency of 230 GHz (equivalent to a wavelength of 1.3 mm[8]) had a spatial resolution of 0.18″ (Davis 2014). Data from CARMA was used in the measurements carried out by Davis et al. (2013) and Onishi et al. (2017). All other measurements to date have been carried out using data from the Atacama Large Millimeter/submillimeter Array (ALMA). Indeed, the ALMA interferometer has been pivotal in extending the potential of using cold molecular gas as a tracer. Davis (2014) cites, in direct comparison to the values given just above for CARMA, a predicted spatial resolution of 0.011″ at 345 GHz (0.87 mm) for ALMA in its full operational capacity.[9] He further argues that the potential exists to measure SMBH masses in over 10^5 local galaxies using ALMA data and, with sufficient sensitivity, that measurements would be possible for SMBHs with masses greater than $\sim 4 \times 10^8\,\mathrm{M_\odot}$ at any redshift.

Besides the tremendous advantages offered by ALMA, the technique itself has several other characteristics in its favour. (i) Probably the most significant advantage of molecular gas as a tracer is that it is dynamically cold (hence *cold* molecular gas[10].). All gas behaves as a collisional fluid and is hence prone to non-gravitational perturbations that can detract from pure Keplerian motion (Kormendy & Ho 2013). Molecular gas, however, while not completely immune to turbulence, is not as susceptible to radiation pressure, shocks or magnetic fields as ionized gas is (see e.g., Davis 2014; Ruffa et al. 2023). In this respect, molecular gas is far more appealing as a tracer and far more likely to exhibit pure Keplerian motion in response to the SMBH potential than ionized gas. In a review of studies of molecular gas kinematics, Ruffa & Davis (2024) find that ETGs that are not Brightest Cluster Galaxies (BCGs)[11] and that do not host AGNs show regularly rotating molecular gas with insignificant or no departure therefrom. (ii) The technique is applicable to all galaxy types across the Hubble sequence, including irregular and dwarf galaxies, as long as molecular gas is present in their centres. In this respect, Young et al. (1995) find a detection rate of 88 % for CO gas in spiral galaxies, while Young et al. (2011) report a detection rate of 22 % for ETGs. Thus, while a substantial number of ETGs would not be suitable, the sample size of potential galaxies is still significant. (iii) The technique is applicable to both active and non-active galaxies. For measurements with stellar dynamics, AGN tend to compromise stellar features in the host-galaxy spectra, which makes accurate modelling difficult. For reasons cited in point (i), ionized gas in the case of AGN is particularly unsuitable.

The cold molecular gas dynamics technique also has some limitations that need to be recognized to avoid bad target selection and to correctly assess possible systematic

[8] In the radio regime, it is more common to give the observing wavelength in terms of frequency.

[9] The angular resolution of an interferometer depends not only on the observing frequency or wavelength, but also on the baseline length and configuration that is used. This and other details of the ALMA array and radio interferometry are discussed in detail in Chapter 3.

[10] Molecular as is also cold in terms of temperature, with typical temperatures of ~ 10 K

[11] The authors define these explicitly as massive ETGs that reside at the centre of rich clusters.

errors. (i) The presence of molecular gas itself does not guarantee suitability of the target galaxy: the gas should have an adequate distribution in the vicinity of the SMBH and have a sufficient surface brightness to permit high-resolution mapping (Davis 2014).[12] (ii) Non-circular motion, turbulence and a high gas velocity dispersion are still factors that might detract from Keplerian motion and must be accounted for in the modelling (see Chapters 4 and 5). (iii) Despite the optimistic outlook for measuring SMBH masses at higher redshift mentioned above, the technique is, for the moment, still limited to galaxies in the local Universe with most measurements lying within 100 Mpc. (iv) A contributing factor to the last point is that, as mentioned above in Section 1.1, luminous mass models depend on high-resolution images in the optical or, preferably, in the near-infrared (see Chapter 2). Thus targets are limited to those with available HST or ground based telescopes with adaptive optics assisted imaging (Davis 2014) and, as mentioned, JWST.

Despite these limitations, measurements with cold molecular gas dynamics have provided a steadily increasing number of reasonably accurate measurements. Zhang et al. (2024) list 29 published molecular gas SMBH mass, while the updated sample of Harrisberg (2020) (see Table A.1 in Appendix A) includes five measurements by Combes et al. (2019), bringing the total number of measurements to 34. The sample of measurements to date includes galaxies across most morphological types, galaxies with both active and non-active nuclei, low- and high-mass galaxies and has provided measurements of SMBH masses ranging from $\sim 10^5$ to $\sim 10^9\,\mathrm{M}_\odot$. Of particular interest are measurements made in the high- and low-mass regimes that would be crucial for exploring the outer regions of the $M_{BH} - \sigma_*$ and other scaling relations. For instance, while the sample of measurements from Bosch (2016) covers SMBH masses in the range of $\sim 10^6$ to $\sim 10^{10}\,\mathrm{M}_\odot$, the majority of measurements in this sample lie between $\sim 10^7$ and $\sim 10^{8.5}\,\mathrm{M}_\odot$, with the very low- and very high-mass regions being relatively poorly populated. Cold molecular gas dynamics has shown promise in contributing exactly to these extreme mass demographics. As part of the mm-Wave Interferometric Survey of Dark Object Masses (WISDOM) project (Onishi et al. 2017, and following papers),[13] Davis et al. (2020) successfully measured the mass of a low-mass SMBH in the dwarf elliptical galaxy NGC 404 at $\sim 5 \times 10^5\,\mathrm{M}_\odot$.[14] On the high-mass end, the cold molecular gas sample to date boasts 10 SMBH mass measurements of mass $> 10^9\,\mathrm{M}_\odot$ (see updated Harrisberg 2020).

Indeed, by exploiting the increasing capabilities of ALMA, the range of SMBH mass should increase substantially and this, in turn, should help better constrain the SMBH-galaxy scaling relations (cf. Davis et al. 2017a).

1.3. SMBH mass correlations with the host galaxy

McConnell & Ma (2013) list several correlations between SMBH masses and properties of

[12] Gas that is not evenly distributed or that has holes in the centre of the gas disc is more problematic, but measurements can still be made (e.g., Onishi et al. 2015; Nguyen et al. 2020).

[13] https://wisdom-project.org/

[14] Davis et al. (2020) refer to the black hole as an intermediate mas black hole, similar to Nguyen et al. (2019) who refer to black holes with $10^3 < M_{BH} < 10^6$ as intermediate mass black holes.

their host galaxies that have been reported in the literature. At large scales, these include a correlation between the SMBH mass M_{BH} and the circular velocity or dynamical mass of the dark matter halo in the outer regions of the galaxy, which, it has been argued (and also disputed), acts as an indication of a correlation or coupling between the SMBH mass and dark matter halo mass at earlier epochs (Ferrarese 2002; Volonteri et al. 2011; Kormendy & Bender 2011). Another large-scale correlation, originally presented by Burkert & Tremaine (2010), has been found between the SMBH mass and the total mass of globular clusters in the galaxy (see also, e.g., Harris & Harris 2011).

At smaller scales, correlations exist in the form of several important scaling relations between the SMBH mass and properties of the host galaxy's more central regions. These include relations between the SMBH mass and photometric properties of the bulge, namely the bulge luminosity L_{Bu} (e.g., Kormendy & Richstone 1995; McLure & Dunlop 2002) and concentration of bulge light (Graham et al. 2001), as well as a correlation between the SMBH mass and stellar mass of the bulge M_{Bu} (e.g. Magorrian et al. 1998).[15]

The strongest of these scaling relations is that between the SMBH mass and stellar velocity dispersion σ_*, which was first reported by Ferrarese & Merritt (2000) and Gebhardt et al. (2000). The stellar velocity dispersion is most commonly taken to be the stellar velocity dispersion as measured within an aperture equivalent to one effective radius R_e of the galaxy or a fraction thereof (see e.g., Sahu et al. 2019; Bosch 2016) and is often referred to as the effective velocity dispersion.[16] An alternative that is often used is the so-called central velocity dispersion, as available from the Hyperleda database (Makarov et al. 2014).[17] This comprises homogenized measurements of velocity dispersion taken within a fixed aperture of 0.595 kpc for a given galaxy.[18]

Further research (Tremaine et al. 2002; Gültekin et al. 2009; Beifiori et al. 2012; McConnell & Ma 2013; Bosch 2016; Saglia et al. 2016) conducted on increasingly larger SMBH samples has confirmed the $M_{BH} - \sigma_*$ relation, but has also revealed that, as with the other scaling relations (and as mentioned above), the slope of the correlation appears to depend on galaxy morphology and mass, as well as the component of the galaxy in which the SMBH sits.[19]

Another factor that may affect the slope is the type of method used to measure the SMBH mass (see Section 1.1) as using different methods can lead to (sometimes significantly) different mass results (see e.g., van den Bosch et al. 2016, and references therein). Figure 1 in Bosch (2016) illustrates the $M_{BH} - \sigma_*$ relation. Here, the data points are differentiated according to the measurement method used so as to determine

[15] In elliptical galaxies, where the whole galaxy is bulge-like, the "bulge" most often refers loosely to the entirety of the galaxy. However, recent studies suggest that the morphology of ellipticals is more complex and that, regarding the $M_{BH} - M_{Bu}$ relation, the SMBH mass is physically more related to the cores of ellipticals (see Zhu et al. 2021, and also references within).

[16] It is also sometimes referred to as the bulge stellar velocity dispersion.

[17] `http://atlas.obs-hp.fr/hyperleda/a001/`

[18] For more in-depth discussion on problems associated with the discrepancies in the way in which stellar velocity dispersion is measured and defined, see e.g., Sahu et al. (2019).

[19] McConnell & Ma (2013) and Saglia et al. (2016) investigate the relation according to the surface brightness profiles of ETGs, considering as separate subsets, core and power-law ETGs. Saglia et al. (2016) also investigate the relations in bulges vs. pseudo-bulges (see also Kormendy & Ho 2013).

any potential biases that may arise due to systematic errors particular to the respective methods or to the fact that each method probes a slightly different region of the galaxy's potential well (Bosch 2016). Indeed, using different methods to measure SMBH masses is useful (and, arguably, essential) in providing yardsticks for comparison and cross-referencing of individual measurements. Nonetheless, having a large enough database of measurements made with a single technique would enable fitting the $M_{BH} - \sigma_*$ and other scaling relations on a more homogeneous basis (see also Thater 2019).

Despite the various sources of discrepancies in slope, Bosch (2016) concludes that the $M_{BH} - \sigma_*$ correlation is universal, spanning all galaxy types and SMBH mass ranges, and that it is independent of any global photometric parameters (see also Gültekin et al. 2009; Beifiori et al. 2012). Indeed, the fit is tight enough to enable an estimate of any prospective SMBH mass using the relevant fit equation and is often used as an indirect method of measuring SMBH mass, especially in AGN where dynamical methods are not possible (see e.g., Gliozzi et al. 2024). Such an estimate is also useful as a prelude to conducting a direct measurement with dynamical methods, as it allows a prediction of the expected mass range and, once it has been conducted, provides a useful comparison to the dynamical measurement. I perform such a mass estimate for the mass of the putative SMBH in NGC 2513, and this is described below in Section 1.4.1.

1.4. NGC 2513

NGC 2513 is a quiescent massive elliptical galaxy with the morphological classification of E3 in the NASA Extragalactic Database (NED).[20] It is part of the Catalog of Nearby Poor Clusters (White et al. 1999) where it is listed as part of a group with 9 members grouped into 2 sub-clusters. It is also one of the massive ellipticals to have been included in the currently ongoing MASSIVE survey (Ma et al. 2014, and following papers).[21] Drawing from the Two Micron All Sky Survey[22] (2MASS, Skrutskie et al. 2006) High Density Contrast group catalogue (Crook et al. 2007), Ma et al. (2014) classify NGC 2513 as a brightest group galaxy of a 4-member galaxy group and upon investigation of the angular momentum and rotation of galaxies in their sample,Veale et al. (2017) find NGC 2513 to be a slow rotator. Figure 1.1 shows the Hierarchical Progressive Surveys (HiPS) 2MASS combined J, H, K band image of NGC 2513.

The distance adopted throughout this thesis for NGC 2513 is 70.8 Mpc and is derived from group-corrected flow velocities (Ma et al. 2014). At this distance, $1''$ is equivalent to $\sim$337.82 pc. Using the extinction corrected total absolute K-band magnitude of -25.52 mag computed by Ma et al. (2014) from magnitudes from the 2MASS Extended Source Catalog (XSC; Jarrett et al. 2000), Veale et al. (2018) estimate the stellar mass with Eq. 2 of Cappellari (2013a) to be $\log_{10}(M_*/\mathrm{M}_\odot) = 11.69$ (see also Section 2.3.10). They calculate an effective radius R_e of $13.9''$ (equivalent to ~ 4.7 kpc at the distance of

[20] https://ned.ipac.caltech.edu/

[21] MASSIVE is a volume-limited, multi-wavelength, integral-field spectroscopic and photometric survey of the most massive ETGs within 108 Mpc. See https://blackhole.berkeley.edu/.

[22] https://irsa.ipac.caltech.edu/Missions/2mass.html

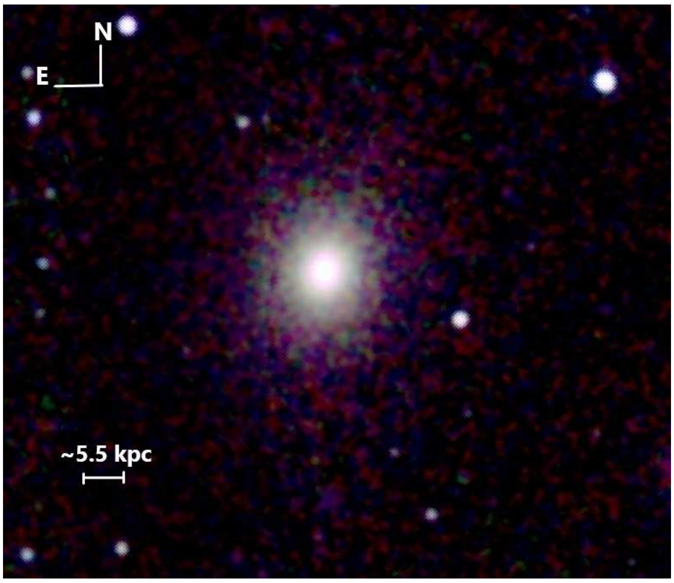

Figure 1.1.: The HiPS 2MASS combined J, H, K band image of NGC 2513, where the J, H, and K bands are centred at 1.23μm, 1.66μm and 2.16μm, respectively. Image available from https://irsa.ipac.caltech.edu/frontpage/.

70.8 Mpc) using Eq. 3 of Ma et al. (2014) and surface brightness profiles from the 2MASS XSC catalogue, and they list an optical-half light radius of 24″ ($\sim$ 8.1 kpc) from the NASA–Sloan Atlas catalogue. Veale et al. (2018) measure the average luminosity-weighted dispersion within 1 R_e as $\sigma_* = 253\,\mathrm{km\ s^{-1}}$, which is the value I assume throughout this thesis.

Using new infrared observations with the HST Wide Field Camera 3 (WFC3; see Section 2.2.1), Goullaud et al. (2018) investigate, among other characteristics, the presence and structure of dust in massive ellipticals and find that NGC 2513 is one of the galaxies that has dust in its centre. Interstellar dust is actually a good indication of the presence of cold molecular gas as it is in this phase of the interstellar medium that the two tend to coexist (see e.g., Williams 2005) and can, to an extent, aid in selecting galaxies that have molecular gas (Davis 2014). Nonetheless, the presence of dust offers its own challenges, especially in modelling the luminous mass of the galaxy (see Chapter 2 and, in particular, Section 2.3.6).

The left panel of Figure 1.2 is a composite J-band image in false colour of NGC 2513 comprising several exposures taken with the HST WFC3 and showing the galaxy's full extent (see Section 2.2 for details on the HST images). The right panel of Figure 1.2 is a zoom-in of a single-exposure HST J-band image of the galaxy taken with the WFC3 showcasing the galaxy's inner regions where the dust lane running through the centre is clearly visible.

Figure 1.3, taken from Goullaud et al. (2018), shows the radial profiles of various isophotal parameters for NGC 2513 as a function of the geometric mean radius $r_{geo} = \sqrt{ab}$,

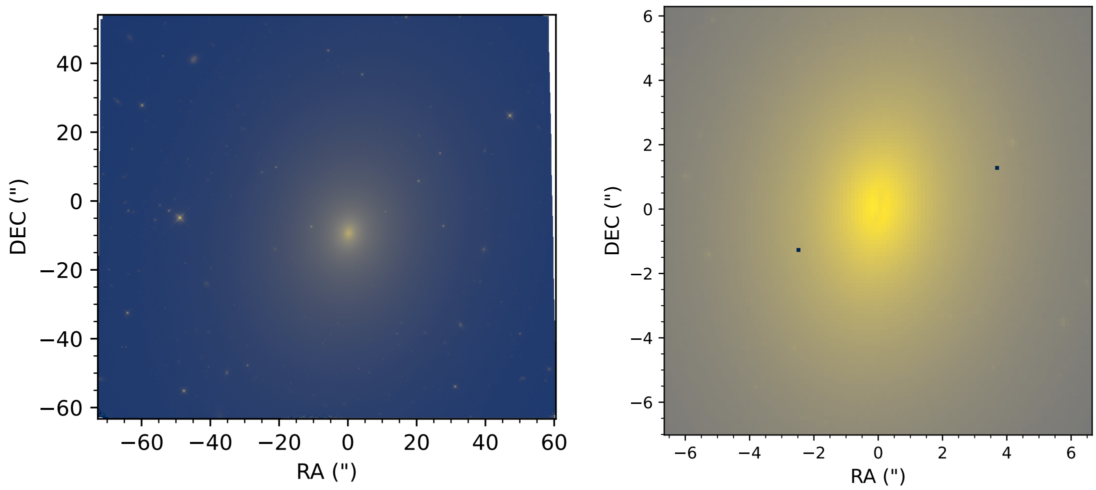

Figure 1.2.: *Left panel*: HST composite J-band image of NGC 2513 in false colours showing the full extent of the galaxy. *Right panel*: Zoom-in of a single exposure HST J-band image of the galaxy's centre (also in false colours) where the dust lane is visible. For the assumed distance of 70.8 Mpc, $1''$ is equivalent to ~ 337.86 pc in both plots. Both images were taken with the HST WFC3.

where a is the semi-major axis and b the semi-minor axis of the galaxy. From top to bottom and from left to right, the panels show the AB magnitude surface brightness profile μ_{AB}, ellipticity ϵ, position angle (PA), higher order deviations from ellipticity A_3, B_3, A_4 and B_4, and isophotal drifts from the galactic centre along both the semi-major and semi-minor axes Δa and Δb, respectively.[23] The horizontal dashed line in the surface brightness profile marks the sky background (see Section 2.3.3). The galaxy appears fairly regular in most respects with a small deviation of the PA towards the centre, which is probably caused by the dust.

1.4.1. Preliminary SMBH mass estimate for NGC 2513 using the $M_{BH} - \sigma_*$ relation

To get an estimate of the SMBH mass for NGC 2513 using the $M_{BH} - \sigma_*$ relation, I adopted the power-law fit as defined in McConnell & Ma (2013):

$$\log_{10}(\,M_{BH}/\mathrm{M}_\odot) = \alpha + \beta\,\log_{10}\,X\ , \tag{1.3}$$

where α is the intercept and β is the slope and $X = \sigma_*/km\ s^{-1}$.[24] I took the best fit values for the parameters in Eq. 1.3 from several sources in the literature for samples of all galaxy types and the respective subsamples of ETGs alone (with the exception

[23] For a definition of the higher order deviations from ellipticity A_3, B_3, A_4 and B_4 and more in-depth discussion of the parameters, see Goullaud et al. (2018).

[24] In McConnell & Ma (2013), the stellar velocity dispersion is normalized with $X = \sigma_*/200\ \mathrm{kms}^{-1}$, which is often the practice.

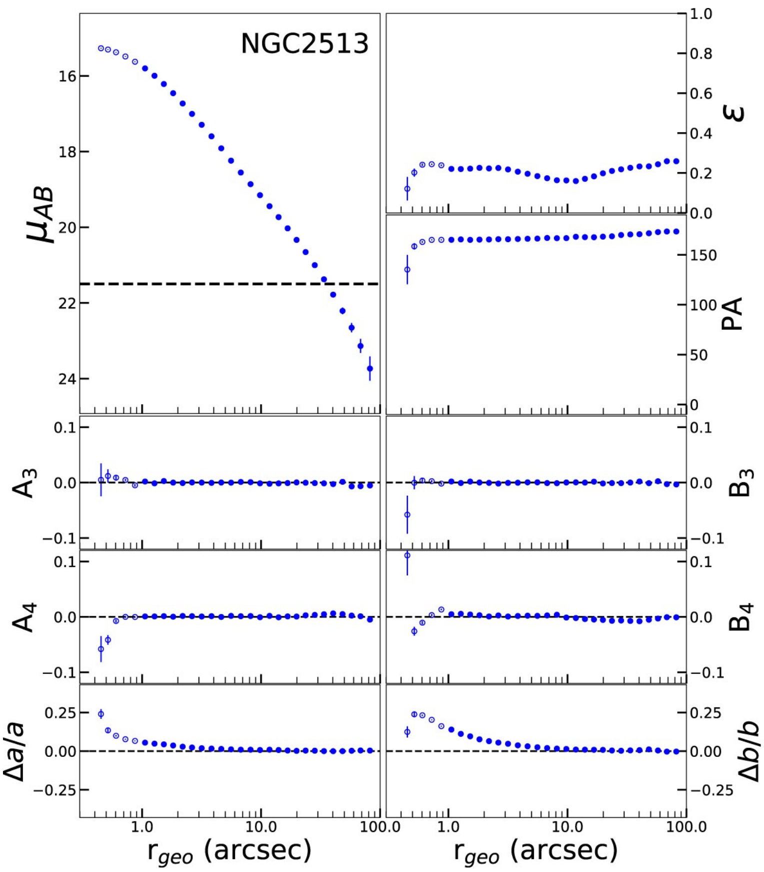

Figure 1.3.: Radial profiles of isophotal parameters of the galaxy NGC 2513 as a function of the geometric mean radius. Figure taken from Goullaud et al. (2018).

of Bosch (2016), who only provides a fit for all galaxy types). In addition, I also used best-fit values for the parameters from both the original and updated all-galaxy samples and their corresponding ETG subsamples from Harrisberg (2020), all of which include only SMBH mass measurements made from cold molecular gas dynamics.[25]

The recovered estimates for the mass of the putative SMBH in NGC 2513 are listed in Table 1.1. They are mostly within the upper range of $10^8\,M_\odot$, the only exception being the estimate from the original ETG subsample from Harrisberg (2020), which is at a low $10^9\,M_\odot$. The difference in values can be attested, for instance, to differences in the respective samples, such as the number of galaxies available, the fitting method used and, for the all-galaxy samples, the ratio of galaxy type.[26] The original fits from Harrisberg (2020) were, for example, based on a still relatively small sample (18 galaxies in

[25] As no stellar velocity dispersion measurement was available for Fairall 49 (cf. Zhang et al. 2024), this galaxy is excluded from the sample that is plotted. In total, the all-galaxy sample comprises 33 galaxies, the ETG subsample, 25.

[26] For a full discussion on the respective samples and the fits see the references in Table 1.1.

Table 1.1.: Mass estimates of the putative SMBH in NGC 2513 obtained from several $M_{BH} - \sigma_*$ best-fits taken from the literature; the respective R_{SOI} values computed from these mass estimates using Eq.1.1.

	1	2	3	4	5	6
M_{BH} All galaxies (10^8 $M_\odot$)	7.87	7.34	7.02	7.81	7.44	7.30
M_{BH} ETGs (10^8 $M_\odot$)	8.33	—	9.84	8.00	10.30	7.75
R_{SOI} All (pc)	52.90	49.32	47.19	52.50	50.00	49.07
R_{SOI} All ($''$)	0.16	0.15	0.14	0.16	0.15	0.15
R_{SOI} ETGs (pc)	56.10	—	66.14	53.80	69.22	52.08
R_{SOI} ETGs ($''$)	0.17	—	0.20	0.16	0.20	0.15

Notes. The mass estimates are based on fit equations taken from the following sources. 1: McConnell & Ma (2013); 2: Bosch (2016); 3: Saglia et al. (2016); 4: Sahu et al. (2019); 5: Harrisberg (2020); 6: Harrisberg (2020) updated. For Saglia et al. (2016), the ETG fit is specifically for core ellipticals. The updated and revised sample of Harrisberg (2020) is given in Table A.1.

total), while all 14 new galaxies added to the revised sample are ETGs.[27] The revised all-galaxy sample is thus heavily biased towards ETGs. When using the best-fit parameters obtained by considering the respective subsamples of ETGs, a larger value is recovered for the estimated SMBH mass. This is congruent with the fact that ETGs tend to be more massive than Late Type Galaxies (LTGs) and would thus tend, according to the $M_{BH} - M_{Bu}$ relation, to have more massive SMBHs.

With these SMBH mass estimates, it is also possible to do a preliminary calculation of the R_{SOI} using Eq. 1.1. The results are included in Table 1.1 in both parsecs and arcseconds. These estimates of the R_{SOI} are useful in determining the angular resolution required to resolve the region where the SMBH mass dominates and for assessing whether the available data has comparable angular resolution (see Section 3.2).

Figure 1.4 shows the $M_{BH} - \sigma_*$ relation for, respectively, the samples of all galaxies from Bosch (2016) and subsample of ETGs from McConnell & Ma (2013), each including measurements made from a mix of dynamical methods. Both samples are overlaid, respectively, with the updated all-galaxy sample and updated ETG subsample of SMBH mass measurements made using cold molecular gas dynamics from Harrisberg (2020). Also shown in the plots are the data points for NGC 2513, where the SMBH mass has been estimated from the $M_{BH} - \sigma_*$ relation as derived from, respectively, the updated all-galaxy sample and updated ETG subsample from Harrisberg (2020).

[27] I include in the ETG sample, all lenticular galaxies; I also include PGC 11179, which is a red nugget and NGC 4786 which is classified as cD peculiar. See Table A.1.

1.4.2. NGC 2513 as the target object

The presence of a molecular gas disc in its centre makes NGC 2513 an obvious target for a SMBH mass measurement with cold molecular gas dynamics. As a massive elliptical with relatively high stellar velocity dispersion, it is also highly likely to boast a SMBH mass in the high-end regime based on both the $M_{BH} - M_{Bu}$ and $M_{BH} - \sigma_*$ relations discussed in Section 1.3. As such, it was selected as a target for observations of the ^{12}CO(2-1) emission line with ALMA based on a proposal by my co-supervisor Dr. Dieu D. Nguyen (see Chapter 3). As discussed in Section 1.2, high-mass SMBHs are of particular interest in view of expanding the scaling relation at the high-mass end. While the cold molecular gas technique has been applied to galaxies more distant than NGC 2513 (see Table 1 in Zhang et al. (2024)), at a distance of 70.8 Mpc, NGC 2513 populates the outer ranges of measurements possible with this method. This therefore provides an excellent opportunity to once again test the molecular gas dynamics method at higher distances and to add to the database of galaxies here. The available data through its inclusion in the various 2MASS surveys and in the MASSIVE survey also make it an attractive candidate. Perhaps most significant, however, is that no SMBH mass measurement has yet been made on this galaxy with any method. Such a measurement will thus no doubt constitute an important addition to the SMBH mass measurement database and to that of measurements made with cold molecular gas.

1.5. Overview of the thesis

The following provides a brief summary of the overall structure and chapter content of this thesis

Chapter 2 discusses the construction of the luminous mass model via the multi-Gaussian expansion (MGE) method. I describe the HST images available for the galaxy and the selection of the image that I used for the MGE as well as the generation of the Point Spread Function (PSF) for the WFC3 instrument and its F110W filter that was used to take the image. I then proceed to explain the MGE method and its application to the image of NGC 2513 in some detail, including a description of the dust-masking I carried out to successful model the image's surface brightness. I then present the final results that will be used in the dynamical modelling to retrieve the three-dimensional intrinsic luminosity density needed to obtain the luminous mass.

Chapter 3 details the procedure of imaging the radio data used for the dynamical modelling. I first introduce the concept of radio interferometry and the necessity of using this technique to attain high-resolution images in the radio regime. I explain the basic concepts and terms of interferometry necessary to fully understand the imaging process. I then briefly describe the ALMA interferometer and its capabilities, also detailing the observations carried out to get the high- and low-resolution data used in the imaging. For the imaging I used the Common Astronomy Software Applications (CASA). After exploring the various tool available in CASA, I explain the steps required to combine the high- and low-resolution data and subsequently image the combined data. I conclude with a presentation of one- and two-dimensional visualisations of the resulting three-dimensional data cube that will be used in the modelling.

Chapter 4 discusses the dynamical modelling of the imaged ALMA data cube with the KINematic Molecular Simulation (`KinMS`) tool via Bayesian analysis with Markov chain Monte Carlo (MCMC) forward modelling procedures. I explain the various steps in setting up the model of the molecular gas disk and total mass of the galaxy. The latter includes the SMBH mass and the results of the MGE from Chapter 2. I then give some background to the key concepts of Bayesian analysis and the MCMC process and link this to the settings chosen for the MCMC forward modelling as carried out here. I conclude by presenting and describing the fixed and free parameters that are included in the model.

Chapter 5 presents and discusses the results of the (MCMC) forward modelling. I examine possible error sources and carry out an appraisal of the recovered SMBH mass in the context of the $M_{BH} - \sigma_*$ scaling relation. Chapter 6 summarizes the project and presents an outlook for further work.

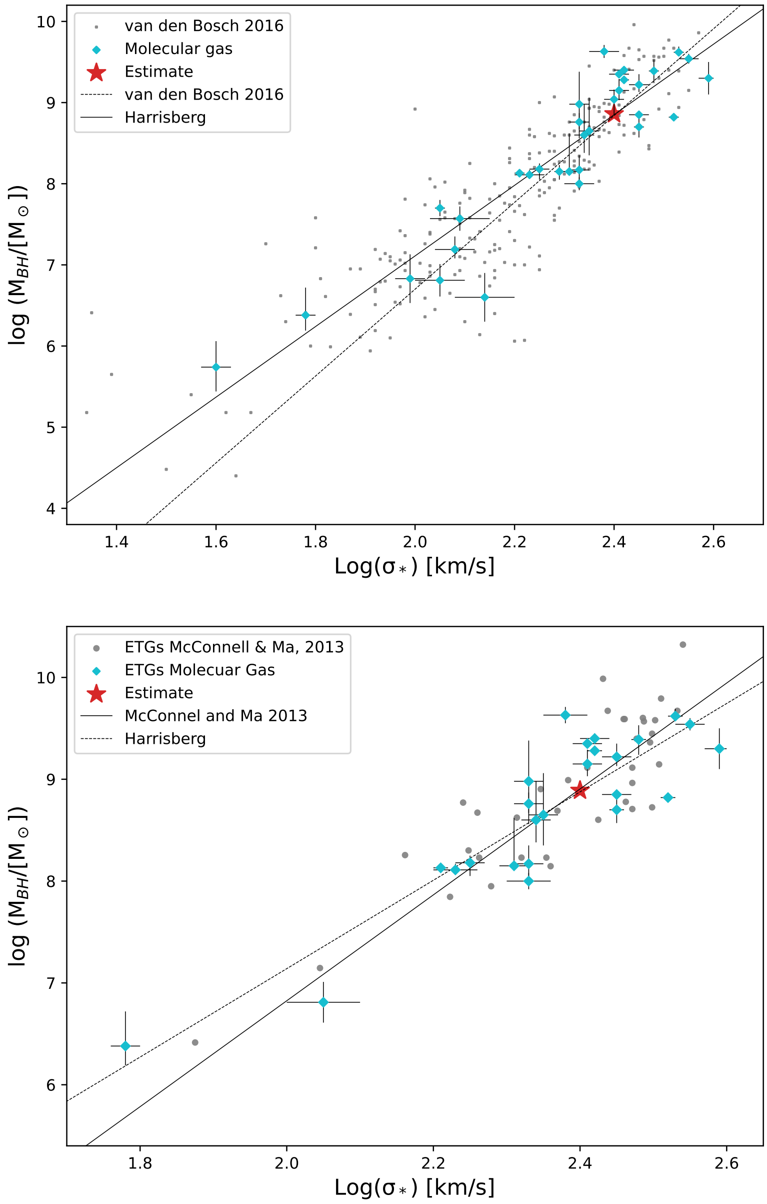

Figure 1.4.: *Top panel*: The $M_{BH} - \sigma_*$ scaling relation for all galaxy types. The grey dots are the sample from Bosch (2016). The cyan diamonds are the updated sample from Harrisberg (2020). The red star is NGC 2513 with an estimated SMBH mass from the $M_{BH} - \sigma_*$ relation. The black solid line is the fit from Bosch (2016): $\log_{10}(M_{BH}/\mathrm{M}_\odot) = 8.32 + 5.35 \log_{10}(\sigma_*/200 \text{ km s}^{-1})$. The black dashed line is the updated fit from Harrisberg (2020): $\log_{10}(M_{BH}/\mathrm{M}_\odot) = -1.59 + 4.35 \log_{10}(\sigma_*/\text{km s}^{-1})$. *Bottom panel*: The same, but for the ETG subsample from McConnell & Ma (2013) (grey dots) and updated ETG subsample from Harrisberg (2020) (cyan diamonds). The black solid line is the fit from McConnell & Ma (2013): $\log_{10}(M_{BH}/\mathrm{M}_\odot) = 8.39 + 5.20 \log_{10}(\sigma_*/200 \text{ km s}^{-1})$. The black dashed line is the updated fit from Harrisberg (2020): $\log_{10}(M_{BH}/\mathrm{M}_\odot) = -1.54 + 4.34 \log_{10}(\sigma_*/200 \text{ km s}^{-1})$.

2. Measurement of the luminous mass

2.1. Introduction

The process of attaining the luminous mass can be broken down into three essential steps: 1) model the projected two-dimensional surface brightness, 2) retrieve the three-dimensional intrinsic luminosity density through deprojection and 3) acquire the total mass by multiplying this result with the mass-to-light ratio M/L as presented in Eq. 2.1:

$$M_{\rm lum} = \frac{M}{L} \cdot L_{\rm obs} \tag{2.1}$$

In this Chapter, I describe the processes required for step 1 in detail. (Steps 2 and 3 are attained in the modelling process described in Chapter 4.) To construct an accurate model of the surface brightness density I applied the multi-Gaussian Expansion (MGE) method developed by Emsellem et al. (1994) and Cappellari (2002) to a high-resolution Hubble Space Telescope (HST) image. The following sections detail the process.

2.2. The HST images

To get accurate measurements of the luminous mass, it is imperative to use good quality, high-resolution images of the galaxy. In this sense, the HST has proved to be invaluable in attaining SMBH mass measurements. This is largely afforded by the fact that the HST has a high angular resolution and, as a space telescope, is not affected by atmospheric considerations, such as seeing, as are ground-based telescopes.

The HST image of NGC 2513 that I used for the MGE modelling was obtained from the Standard HST Archive collection of the ESA Hubble Science Archive.[1] The image derives from a single observation carried out on 2016-10-30 with the HST's Wide Field Camera 3 (WFC3) using the F110W filter, which is a wide-band filter in the near-infrared. The observation comprised five consecutive exposures, each with a total integration time of 499.23 seconds. Four of the images have been combined into a single composite image. I opted to download this composite image and its four constituent images to do an initial investigation of the galaxy and to select an appropriate image for the MGE modelling. Details of these images and the process of the image selection are given below after a brief introduction to the WFC3 camera and the F110W filter.

[1] https://hst.esac.esa.int/ehst/

Table 2.1.: Characteristics of the WFC3 **IR** channel.

Spectral range (nm)	Pixel format	Pixel scale (arcsec)	Field of view (arcsec)	No. of filters
800 - 1700	1014 x 1014	0.135 x 0.121	136 x 123	15

2.2.1. The WFC3 camera and F110W filter

The WFC3[2] was installed in 2009 as a replacement for Hubble's Wide Field Planetary Camera 2 (WFPC2). It has two independent channels, one sensitive at the ultraviolet and visible range (the **UVIS** channel) and one at the near infrared (the **IR** channel). Each channel has its own optical system, grisms, detectors and filters. The images of NGC 2513 were obtained with the IR channel's F110W filter. Details of the **IR** channel are given in Table 2.1[3]See: `https://hst-docs.stsci.edu/wfc3ihb/chapter-2-wfc3-instrument-description/2-1-optical-design-and-detectorsandhttps://hst-docs.stsci.edu/wfc3ihb/chapter-7-ir-imaging-with-wfc3/7-5-ir-spectral-elements`

Of the 15 channels constituting the **IR** channel, the F110W filter is particularly suited for deep imaging. It is a wide filter that spans the ground-based Y and J bands with a central wavelength of 11 696.34 Å. It should be noted that an IR filter is ideal for creating accurate luminous mass models as imaging in the infrared limits the effects of dust extinction. Due to their shorter wavelengths, visible and especially blue light waves are far more susceptible to absorption and scattering by dust grains than are red light waves. Using a bluer filter thus means that a substantial amount of blue light is lost, making stars or galaxies appear redder than they are. The loss of light means an inaccurate modelling of the intrinsic light that will result in a galaxy of lower luminous mass. This in turn would result in a SMBH mass that is too high.

2.2.2. HST image selection

For the purposes of this thesis, I downloaded the composite image and its four member constituents, all taken with the F110W filter, which is a *J*-band filter in the WFC3's IR band. The main characteristics of these five images are listed in Table 2.2.

A visual inspection of the images in the `SAOImageds9`[4] imaging and visualization application revealed that the four individual images appear fairly similar. The composite image appears cleaner with sharper contrast than its constituents, which results from combining and smoothing over the four constituent images. Looking at the headers of each image reveals small differences in the values of such parameters as the number of good pixels and signal-to-noise ratio (given in Table 2.2), as well as the maximum, minimum and average values of the good pixels. In plots of the four images created in

[2]The main reference for this section is the current WFC3 Instrument Handbook, which can be found at: `https://hst-docs.stsci.edu/wfc3ihb`

[3].

[4]`https://sites.google.com/cfa.harvard.edu/saoimageds9`

Table 2.2.: Characteristics of the HST images.

Observation ID	Exposure time (s)	Scale (arcsec)	No. good pixels	Max. S/N
icv424010	1996.94	0.120	–	–
icv424**clq**	499.23	0.121	987759	97.73
icv424**cnq**	499.23	0.121	987757	91.53
icv424**cpq**	499.23	0.121	987758	95.64
icv424**crq**	499.23	0.121	987756	97.79

Notes. The four individual HST images of NGC 2513 (rows 2−5) were combined to form the composite image (row 1). All images were captured with the WFC3's F110W filter in the **IR** band. The scale (Col. 3) refers to the conversion between 1 pixel and its width in arcseconds and was taken directly from each respective image using the region statistics tool of `SAOImageds9` (see Section 2.3.3). The number of good pixels (Col. 4) and maximum signal-to-noise ratio (Col. 5) were taken from the headers of the respective FITS files.

`Python`, however, it is possible to see slight differences in brightness: the images icv424cpq and, especially, icv424crq appear brighter, particularly at larger radii, than the other two images, especially icv424clq. This was confirmed on performing MGE fits on the images, which is discussed below. The four images and composite image are shown in Fig. 2.1.

To inspect each image more methodically, I created a surface-brightness profile of each of the four constituent images and the composite image. I first cropped the images from the original box-size of 1014×1014 pixels ($\sim$122.70$''$ $\times$ 122.70$''$) to have a diameter along the horizontal axis in the range of $\sim$ 170 pixels, which is $\sim$20.57$''$.[5] The cropped image covers a radius of about 10.29$''$ from the galactic centre, which, for comparison, is a little over half the semi-major axis of 17.97$''$ obtained by Quenneville et al. (2023) for the half-light elliptical isophote from K-band imaging with the Canada–France–Hawaii Telescope (CFHT). Cropping the images was done to both avoid some of the bright foreground stars at larger radii and to focus on the more central regions of the galaxy.

I then created the profiles by taking a slit of one pixel in width along the horizontal axis of the cropped images and plotting the brightness in electrons per second[6] against the horizontal axis in arcseconds. The profiles, overlaid on one another, are shown in Fig. 2.2. Here, the offset of the profiles from the centre has been introduced artificially so as to give a better view of the individual profiles.

Each profile displays a slight kink on the right-hand side of the profile just below the peak. This is most probably due to obscuration by dust. The kink is most marked in the profiles of the images icv424cnq, icv424cpq and icv424crq, while that in icv424clq is relatively smooth. The composite image has the least marked kink, which again is most probably due to the smoothing over of the four images in the process of their combination. For purposes of performing an accurate MGE, a smoothed, composite image is not the

[5] The conversion to arcseconds is calculated using the scale of 0.121$''$ to 1 pixel as given in Table 2.2

[6] This is the unit of brightness of the HST images of NGC 2513 as given in the headers.

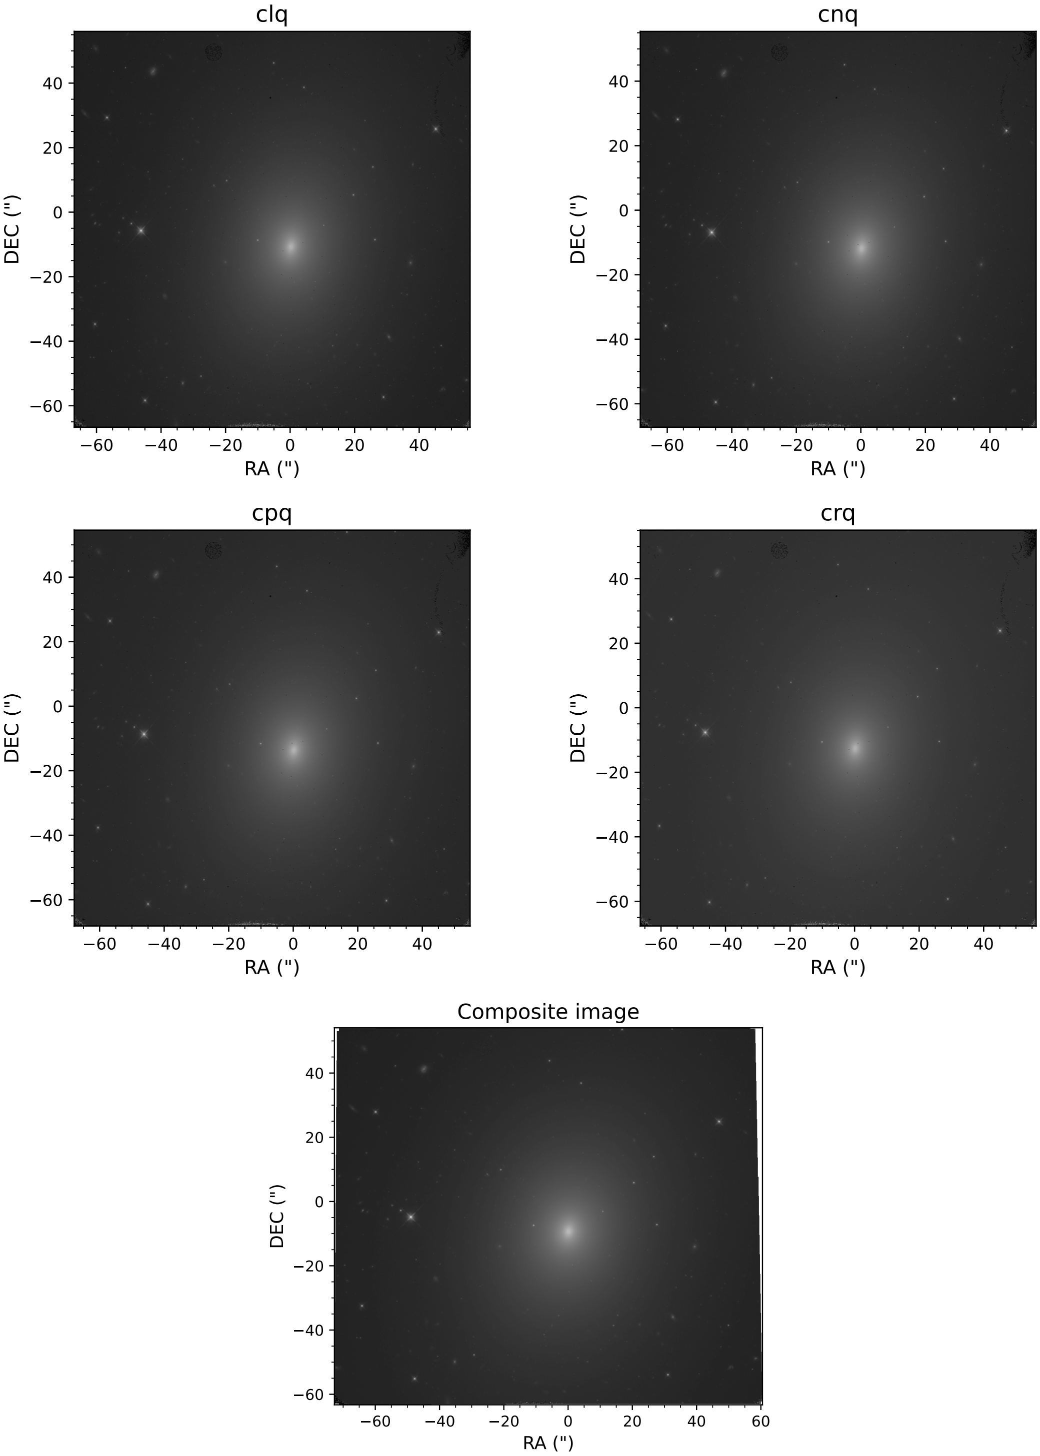

Figure 2.1.: The four individual HST images (shown out to ∼65″) together with the composite image formed by their combination (shown out to ∼70″).
First row, left to right: icv424clq, icv424cnq. *Second row, left to right*: icv424cpq, icv424crq. *Bottom row*: Composite image icv424010.

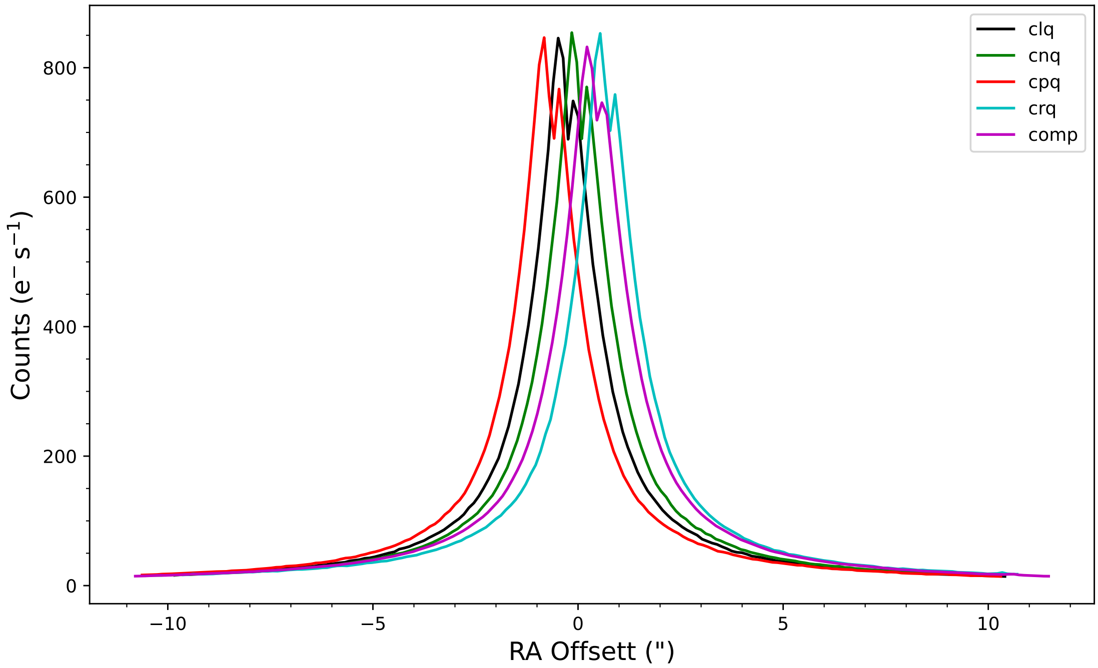

Figure 2.2.: Surface brightness profiles of the four HST images and of the composite image obtained from their combination. Each profile has been slightly offset from the centre for better viewing and comparison.

best choice as information is lost in the smoothing.

Of the four member images, icv424crq shows the highest count-rate and it also has the best S/N ratio (although only by 0.06 – see Table 2.2). However, the kink in icv424clq is relatively smooth and less pronounced than in the other images and could indicate that the inner regions are slightly less obscured by dust or any artifacts.

Another method of assessing the quality of the images is to perform an MGE on the images to investigate how the differences affect the models. I thus performed MGE modelling using the `MgeFit` package[7] by Cappellari (2002) on all four member images. (A full discussion of the MGE method and its application to this galaxy is given below in Section 2.3.)

Figure 2.3 shows the two-dimensional contour plots of the galaxy's isophotes (see 2.3.1) in black and the MGE fit contours in red for all four images. These 2-D plots are the visualisation of the main results from the MGE fit and, although only a preliminary run, give a reasonably good idea as to how well the galaxy's light distribution is mapped and how well the fit works.

The whole-galaxy contour plots for icv424clq and icv424cnq are reasonably similar, although the fit for icv424clq appears to be slightly better than for icv424cnq. The latter also contains contours extending slightly further out than for icv424clq, indicating that the

[7] `http:purl.org/cappellari/software`

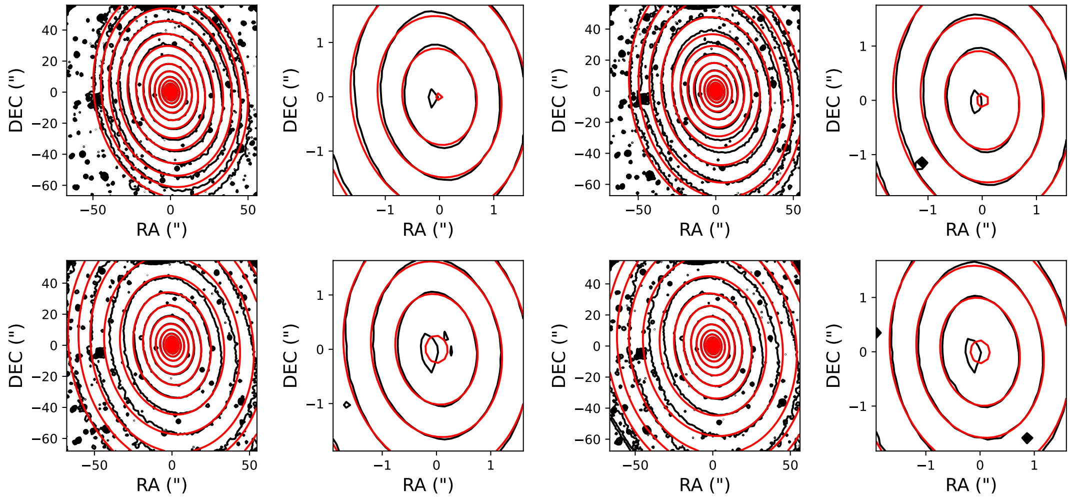

Figure 2.3.: Two-dimensional contour plots of the galaxy isophotes (black curves) and best-fit MGE contours (red curves) for the four HST *J*-band images.
Top row, left to right: icv424clq, icv424cnq; *Bottom row, left to right*: icv424cpq and icv424crq. The left panel in each case shows the contour plot of the whole galaxy; the right panel shows a zoomed-in plot of the galaxy's central regions.

image contains brighter pixels in this region. The zoomed in plot for icv424cnq contains an artifact about 1″ left of the centre, which is absent from icv424clq. The whole-galaxy plots for icv424cpq and icv424crq are more extended than for the other two images. This confirms the findings discussed above that these images, especially icv424crq, have a more extended brightness than the other two images.[8] The zoomed-in plot of icv424crq also contains two artefacts, while that for icv424cpq has some interesting mini-isophote contours both in the centre just adjacent to the central isophote and just outside of the isophote at 1″ left of centre. The offset of the central isophote in all the zoomed-in plots is most probably due to obscuration by the dust lane (see Section. 2.3.6).

The results indicate that the image icv424clq produces the cleanest results with the least artefacts. The galaxy is also not too extended, indicating that the MGE code avoids including pixels that do not belong to the galaxy. I therefore chose to work with this image for the modelling of the luminous mass. Henceforward, all references to the HST *J*-band image refer specifically to the icv424clq image.

2.3. Multi-Gaussian Expansion (MGE)

In this Section, I discuss the process of modelling the projected two-dimensional surface brightness of NGC 2513 on the high-resolution HST image detailed in Section 2.2.2 using the multi-Gaussian Expansion method by Cappellari (2002).

[8]This can also be nicely seen in the plots of the `FIND_GALAXY` and `SECTORS_PHOTOMETRY` produced in the MGE fit (see Section 2.3.3).

2.3.1. The MGE method

The idea behind any given method in modelling a galaxy's luminosity is to identify those points (usually pixels) in the available image that are of equal brightness. The curves that link these points are referred to as isophotes. The isophotes must then be fitted with a suitable method to model the two-dimensional surface brightness distribution.

As most galaxies tend to have isophotes that are close to elliptical in shape, a suitable and common method of modelling a galaxy's isophotes is with ellipses of increasing semi-major axis (Milvang-Jensen & Jørgensen 1999; Cappellari 2002). Any deviations from ellipses can be parameterised by means of Fourier expansions (Milvang-Jensen & Jørgensen 1999). However, as Cappellari (2002) points out, strong deviations cannot easily be photometrically modelled with only a few Fourier terms, which in turn causes problems with dynamical modelling of the fitted isophotes. Cappellari (2002) also notes that this is a disadvantage particularly for multi-component galaxies such as spirals or lenticulars that may contain a bulge, main disc and possibly nuclear disc or bar components.

An alternative approach, and one that can circumvent these limitations, is to model the galaxy's isophotes using sums (i.e., a series expansion) of two-dimensional Gaussian functions (Cappellari 2002). This is now commonly referred to as the multi-Gaussian Expansion (MGE) method. The advantage of using Gaussian functions derives mainly from their mathematical properties. The most relevant of these are that Gaussian functions are their own Fourier transforms and that the convolution of two Gaussian functions is also a Gaussian function. Considering that one of the most elemental principles of galaxy imaging and modelling is the convolution of a galaxy's image with the point spread function (PSF) of the observing telescope (see Section 2.3.3), working with Gaussians simplifies calculations significantly. Indeed, Bendinelli (1991) remarks that the analytical properties of Gaussians allow for the reformulation of complicated problems in a simpler manner and to obtain information from images on structures that are smaller than the data resolution.

The application of multi-Gaussian expansions to actual galaxy images began with Bendinelli (1991), who showed that MGEs can be applied to determine the surface brightness profile of a galaxy image in the simple spherical case.[9] The principle, as outlined by Bendinelli (1991), is that by convolving the spatial intensity of a galaxy with the PSF, both of which are expanded in a sum of Gaussians, one can in turn obtain the observed intensity, again as a sum of Gaussians. The MGE method was first applied to real galaxies by Monnet et al. (1992), who also expanded its application to triaxial galaxies. Emsellem et al. (1994) further generalised the method to allow for non-centred Gaussians, which can better account for components such as offset bars and discs.

The seminal basis for most MGE modelling today and for the MGE modelling in this work is by Cappellari (2002), which, building on the principles of these previous works, introduced an algorithm for automatically fitting MGEs of galaxy images.

[9] Multi-Gaussian expansion had previously been successfully used to determine the PSF of under-sampled CCD frames and to obtain an approximation of the expected PSF of the HST (see Bendinelli 1991, and references therein).

2.3.2. MGE: Formalism and Principles of Application

The following is taken largely from Cappellari (2002) and briefly outlines the mathematical formalism of multi-Gaussian expansion as set out in that paper as well as the basis for the algorithm used in the `MgeFit` code by Cappellari (2002). The paper only considers concentric Gaussians with circular PSF; the generalised case is given in Emsellem et al. (1994).

As mentioned in Section 2.3.1, the idea is to express the two-dimensional projected surface brightness of a galaxy as a sum of Gaussians. Consider a system of coordinates (x', y', z') centred on the galaxy nucleus and with z' pointing towards the observer. For polar coordinates (R', θ') on the plane of the sky (i.e., on the x', y' plane), the MGE projected surface brightness $\Sigma(R', \theta')$ can then be defined as

$$\Sigma(R', \theta') = \sum_{j=1}^{N} \frac{L_j}{2\pi\sigma'^2_j q'_j} \exp\left[-\frac{1}{2\sigma'^2_j}\left(x'^2_j + \frac{y'^2_j}{q'^2_j}\right)\right] , \qquad (2.2)$$

(see Eq. 1 in Cappellari (2002)), with $x'_j = R' \sin(\theta' - \psi_j)$ and $y'_j = R' \cos(\theta' - \psi_j)$. Here N is the number of Gaussian functions with total luminosity L_j, dispersion σ'_j along the major axis, position angle (PA) ψ_j, measured from the y'-axis to the major axis, and observed axial ratio q'_j, with $0 \leq q'_j \leq 1$, where $q' = b/a$, and a and b are the semi-major and the semi-minor axes of the Gaussians, respectively.[10]

The surface brightness of the image must be convolved with the normalised PSF of the instrument for accurate modelling. To facilitate this, the PSF itself can be expressed as a sum of M circular Gaussians:

$$\mathrm{PSF}(R') = \sum_{k=1}^{M} \frac{G_k}{2\pi\sigma^{*2}_k} \exp\left(-\frac{R'^2}{2\sigma^{*2}_k}\right) , \qquad (2.3)$$

(see Eq. 3 in Cappellari (2002)), where G_k is a weighting factor multiplying the normalized Gaussian with $\sum_{k=1}^{M} G_k = 1$, R' is defined as in Eq. 2.2 and σ^*_k is the dispersion of the Gaussians. Taking into account that the luminosity L of any Gaussian with a surface brightness Σ does not change after convolution with the PSF, i.e. $\bar{\Sigma} = \Sigma * \mathrm{PSF}$, the convolved Gaussian can then be expressed as

$$\bar{\Sigma}(R', \theta') = L \sum_{k=1}^{M} \frac{G_k}{2\pi\bar{\sigma}'^2_k \bar{q}'_k} \exp\left[-\frac{1}{2\pi\bar{\sigma}'^2_k}\left(x'^2 + \frac{y'^2}{\bar{q}'^2_k}\right)\right] , \qquad (2.4)$$

(see Eqs. 4 and 5 in Cappellari (2002)), with $\bar{\sigma}'^2_k = \sigma'^2 + \sigma^{*2}_k$ and $\bar{\sigma}'^2_k \bar{q}'^2_k = \sigma'^2 q'^2 + \sigma*^2_k$.

The fitting algorithms are explained in detail in Cappellari (2002). Briefly, the simplest non-linear 1- and 2-D fitting procedure require sampling the photometric profile of the galaxy logarithmically, solving the non-linear least-square minimization for the dispersion

[10] Note that this definition of q' is for the oblate axisymmetric case, which is what is assumed in this discussion and for the modelling of NGC 2513.

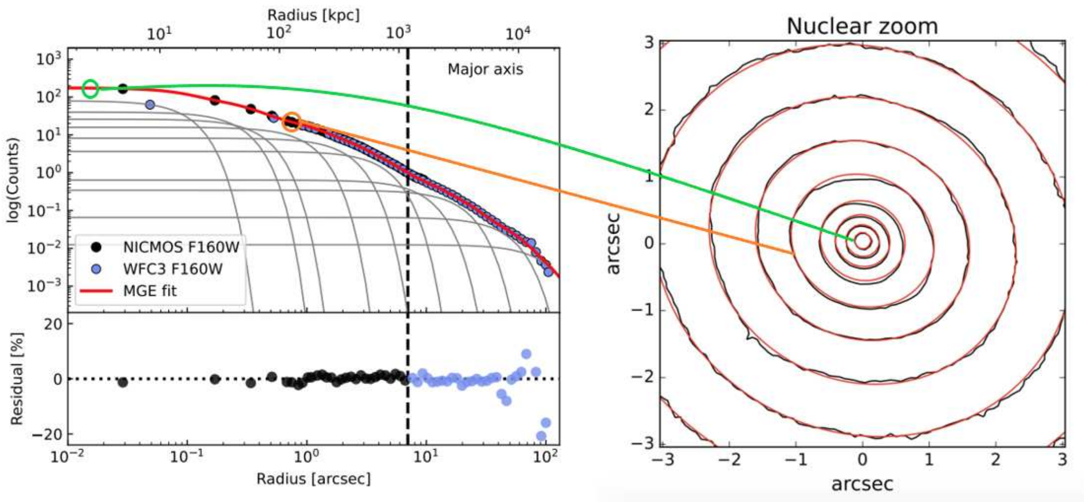

Figure 2.4.: MGE fit of the HST NICMOS image of NGC 6958 by Thater et al. (2021). *Left panel*: 1-D plot of the series expansion of Gaussians (grey curves) generated by the MGE fit and the resulting fit (red curve). The black and blue dots are the data points from the NICMOS F160W filter and WFC3 cameras, respectively. *Right panel*: 2-D contour plots of the galaxy isophotes (black curves) and MGE fit (red curves). The green and orange circles and lines indicate corresponding points on the respective plots.

of the Gaussians σ_j and finding the best-fitting values for the total luminosity L_j of each Gaussian (where σ_j and L_j are from Eq. 2.2). The application of the non-linear 2-D fitting procedure of the `MgeFit` code on NGC 2513 is described in Section 2.3.3, as is the process of convolving the galaxy image with the PSF as defined in Eq.2.3.

The idea of fitting the surface brightness with an expansion of Gaussians is illustrated in Figure 2.4. The left panel shows the series expansion of Gaussians (grey curves) and resulting fit (red curve) to the data (blue and black dots) in one dimension. The right panel shows the corresponding two-dimensional contour plots (right panel) of the isophotes (black curves) and fit (red curves) resulting from the MGE fit performed on the HST NICMOS[11] image of NGC 6958 by Thater et al. (2021). The green and orange circles and lines indicate examples of two corresponding points on the 1-D and 2-D fit.

2.3.3. MGE fit of NGC 2513

I performed the MGE fit using the `MgeFit` python package by Cappellari (2002). This software package provides an automated algorithm designed to produce accurate MGE parameterizations of a galaxy's surface brightness. The fitting procedure consists of four steps, each of which is performed by a corresponding procedure in the package. The

[11]NICMOS (Near Infrared Camera and Multi-Object Spectrometer) was in service on the HST till about 2008, covering wavelengths of 800 to 2500 nm.

procedures, which I describe directly below, are called within the framework of the main fitting function of the code. Before implementing the procedures, various input parameters need to be defined. These are then called as arguments within the various procedures.

Input parameters

The *sky level* refers to the background brightness of the sky. This incorporates any light, both natural and artificial, that does not emanate from the astronomical object or objects under observation. Sources that originate from within or close to the Earth's atmosphere are, for instance, airglow (both natural and artificial) and the auroras, while sources from without the Earth's atmosphere include zodiacal light and diffuse galactic light (Barentine 2022).

In order to acquire an accurate model of the galaxy's surface brightness, the sky level must be subtracted from the galaxy image. Failure to do so, or estimating a sky-level value that is too low, will result in a photon count and subsequent surface brightness density that is too high, and this, in turn, would result in a corresponding luminous mass that is too high. Conversely, a sky-level value that is too high will produce a surface brightness and corresponding luminous mass that is too low.

The effects of inaccurate sky levels are illustrated in Figure 2.5 which shows the 1-D profile plots of the series expansion of Gaussians of two MGE fits for NGC 2513 (left panels) along with the variation of relative error along the profiles for each (right panels). Each MGE was constructed with a sky level value that deviates significantly from the sky level of 2.0 $e^- \, s \, pixel^{-1}$ that was found for NGC 2513 (see below). The top figure shows a sky level of 0.5 $e^- \, s \, pixel^{-1}$, which is too low. In this case, the background sky is far too bright and will most likely be modelled as part of the galaxy. This can be seen in the shape and extent of the Gaussian components of the expansion and in the resulting fit, which tapers off well beyond the extent of the galaxy, covering regions as far out and beyond as $\sim 60''$ (~ 29.96 kpc). The right panel shows a sky level of 5.0 $e^- \, s \, pixel^{-1}$, which is too high. Here, the background sky is virtually non-existent and a good portion of the galaxy's outer regions will not be modelled. In this case, the contours drop off much more steeply and the galaxy's extent covers a mere $10''$ (~ 3.38 kpc), which is lower than the effective radius of $13.9''$ (equivalent to ~ 4.7 kpc, see Section 1.4). These issues of poor fitting are also reflected in the relative error variations, which deviate quite significantly in regions above $\sim 20''$ and $\sim 10''$ for low and high sky levels, respectively. For comparison, the corresponding 1-D contour plots of the best MGE fit shown in Figure 2.14 show a fit that extends out to $\sim 50''$ (~ 16.89 kpc) and the variation in relative error is comparatively minimal.

It should be noted that it is highly likely that the sky background is not constant over the image. Thus, assuming a single value for the sky level is a relatively strong assumption. Nonetheless, it is sufficient for the purposes of the modelling.

To get the sky value for the image of NGC 2513, I used the image viewing tool `SAOImageds9`. Selecting pixels from several parts of the image that are free of any background stars and far from the galaxy, I found a sky-level value of ~ 2.0 $e^- \, s^{-1} \, pixel^{-1}$. Alternatively, it is also possible to create a small region on a suitable part of the image

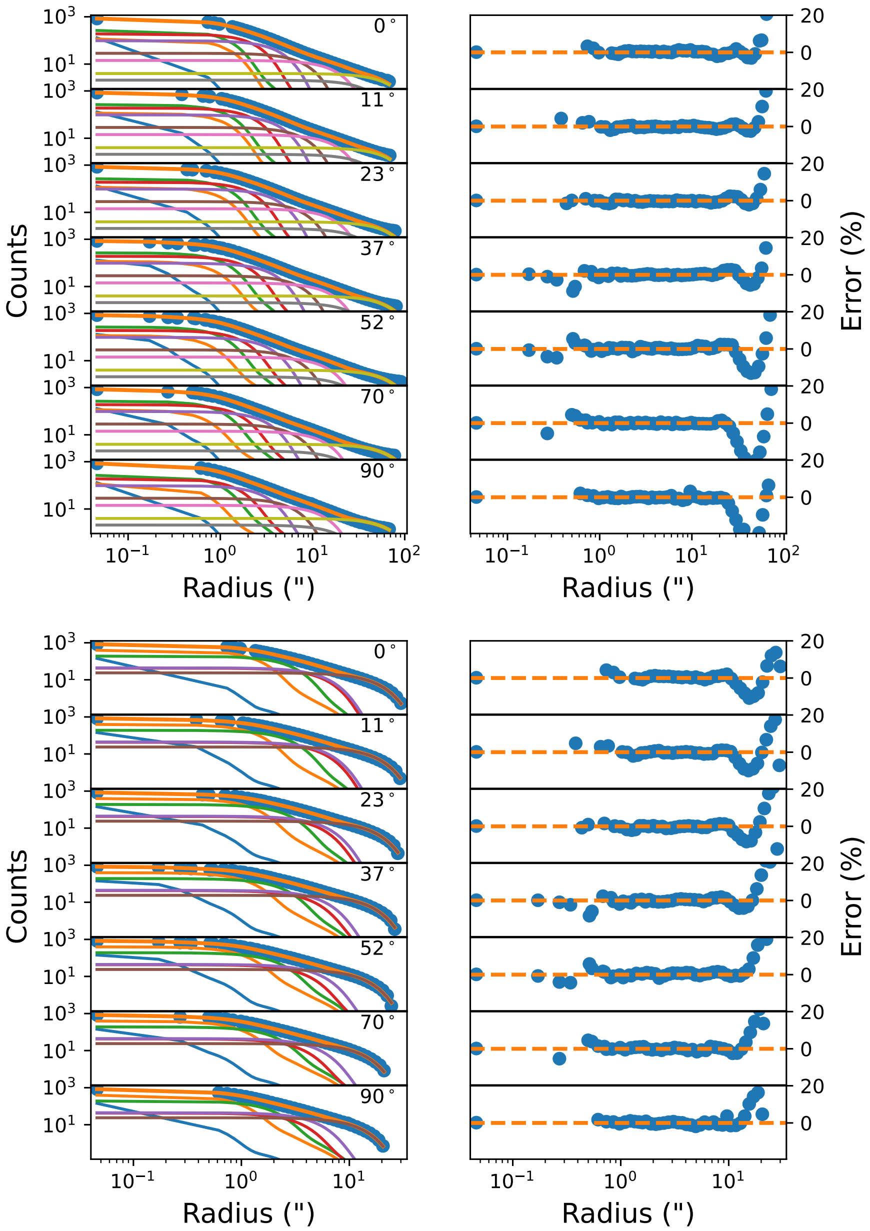

Figure 2.5.: Effect of sky level on the photometry of galaxy NGC 2513. The top left panel shows a 1-D MGE fit modelled from the HST *J*-band image of NGC 2513 with a sky-level of 0.5 $e^{-}\,s\,pixel^{-1}$, which is too low with respect to the estimated value of 2.0 $e^{-}\,s\,pixel^{-1}$. The top right panel shows the variation of the relative error along the profiles. Bottom panels are the same as top ones, but with a sky level of 5.0 $e^{-}\,s\,pixel^{-1}$, which is too high with respect to the estimated value. The coloured curves in the 1-D MGEs are the individual Gaussian components of the expansion, the upper thick orange curve in the 1-D MGEs is the fit and the blue dots are the data. For full details, see also Figure 2.14.

in the viewing tool and take the mean from the region's statistics. This method also produced values around $\sim 2.0\ \mathrm{e^{-}\ s^{-1}\ pixel^{-1}}$. More rigorous approaches are detailed in both Goullaud et al. (2018) and Jensen et al. (2021), both of which include galaxies from the MASSIVE Survey (Ma et al. 2014), among them NGC 2513. The values they find are $2.21 \pm 0.06\,\mathrm{e^{-}\ s^{-1}\ pixel^{-1}}$ and $2.00 \pm 0.06\,\mathrm{e^{-}\ s^{-1}\ pixel^{-1}}$, respectively. Thus, a value of $\sim 2.00\ \mathrm{e^{-}\ s^{-1}\ pixel^{-1}}$ for the sky-level is consistent with the values in the literature.

The *minimum level* sets a minimum count level below which any counts will be ignored. For instance, setting a minimum level of $0.2\,\mathrm{e^{-}\ s^{-1}\ pixel^{-1}}$ as I did for this model, means that any pixel with counts $< 0.2\,\mathrm{e^{-}\ s^{-1}}$ will be ignored in the fitting. This parameter essentially determines the extent of the galaxy. If the level is set too high, then pixels containing less light, but which still belong to the galaxy, will be omitted from the modelling. The opposite is true for a minimum level that is too low. In this sense, the minimum level compensates for the fact that a perfect sky subtraction is virtually impossible. In the hypothetical case of an optimal sky-level value, one would have a perfect sky-subtraction and all the light would emanate only from the galaxy. In this case, the minimum level could be set to zero. To account for the fact that, even with a good estimate of sky level, a certain amount of sky brightness may still be present in certain pixels, a minimum count level reduces the risk of including these pixels in the model.

The *image scale* is the conversion between one pixel and its width in arcseconds for that particular image. The image is composed of pixels and this is the unit used for the actual fitting process. However, for an MGE model to be used in any dynamical modelling, the results must be expressed in physical units rather than the instrumental units of the fit's output. Hence, for the modelling, the unit of pixels must be converted to arcseconds. The image scale is, nonetheless, included as an input in the MGE fitting procedure so that the scale of the resulting contour maps of the fit can be expressed in terms of arcseconds. The scale can be attained directly from the region statistics in `SAOImageds9`. For this image, I get a value of $0.121\,\mathrm{arcsec\,pixel^{-1}}$ (see Table 2.2).

Finally, the *number of Gaussians* determines the maximum number of Gaussian functions allowed in the fit. Cappellari (2002) finds that a value in the range of 10−15 Gaussians produces an accurate MGE fit.[12] For this fit, I chose a total of 12 Gaussians. It should be noted, however, that not all Gaussians are necessarily used in the MGE fit. Negative- and zero-valued Gaussians are discarded.[13] For my final fit, a total of 10 Gaussians were found useful and 2 were discarded.

The values for the input parameters described above are summarized in Table 2.3.

[12]This is the ideal number of Gaussians for the non-linear fitting algorithm described in Cappellari (2002), which is the most economic in terms of speed. The fully linear algorithm given in Cappellari (2002) will ensure finding the lowest χ^2 value, but requires a higher number of Gaussians and might not necessarily provide a significantly lower error.

[13]Negative-valued Gaussians can be included if needed (e.g. to fit boxy bulges), but this must be specified by the user (see documentation of the `Mgefit` python package).

Table 2.3.: Input parameters for the MGE model.

Sky level ($e^- s^{-1}$ pixel^{-1})	Minimum level ($e^- s^{-1}$ pixel^{-1})	Image scale (arcsec pix^{-1})	No. of Gaussians
2.0	0.2	0.121	12

The Point Spread Function (PSF)

The standard PSF modelling package for HST since its launch has been the software package `Tiny Tim` (Krist et al. 2011).[14]. The software consists of two basic programs, `tiny1` and `tiny2`, and of a third program, `tiny3`, specific to the ACS and WFC3 PSFs. In `tiny1`, the user enters parameters such as the instrument to model, the filter, the PSF diameter and the date of observation. The PSF is computed in `tiny2`. As the WFC3 (like the ACS) has significant geometric field distortions, `tiny3` must be used to compute the distorted PSF. Running `Tiny Tim` produces a `FITS` image of the PSF. This is shown in Figure 2.6.

As can be seen in Eq. 2.3, the normalised PSF as expressed by a sum of Gaussians is characterised by three free parameters: the polar coordinate radius R' of the PSF, the standard deviation or dispersion of each Gaussian σ_k^* (corresponding to the width) and weighting factor G_K which must be normalized with $\Sigma_{k=1}^{M} G_k = 1$. The values of G_K are given in photon counts or electrons per second (depending on the telescope) and the dispersion σ_k^* is given in pixels. It is these two parameters that need to be incorporated into the main fitting function of the MGE code for convolution with the MGE surface brightness. For convenience and because this is the term used in the output of the MGE, I refer to G_k in terms of counts for the rest of this discussion.

To obtain these parameters, I modelled the PSF image obtained from `Tiny Tim` (see Figure 2.6) as a sum of circular Gaussians by performing an MGE fit on the PSF image. In performing the modelling, there are a few important considerations concerning the input parameters discussed above. Firstly, as the sky itself has not been imaged in the PSF, the *sky level* must be omitted from the input. Secondly, the *minimum level* can be set to zero, as the objective is to image the entire extent of the PSF. Thirdly, the scale can be omitted as this has no bearing on the results. As mentioned above, the MGE fit of the galaxy's image is modelled in instrumental units, i.e. counts (in electrons per second for the HST) and pixels. The units of the normalized PSF thus also need to be in these instrumental units to be successfully convolved with the galaxy's image. If a scale is used for the visualization, then it should be that of the camera being used, which, in this case, is the WFC3 which has an average absolute pixel scale $0.13''$ pixel^{-1}. Finally, when performing the fit on the PSF image, one needs to ensure that the model

[14]The package is available at: `https://www.stsci.edu/hst/instrumentation/focus-and-pointing/focus/tiny-tim-hst-psf-modeling` It should be noted that the software, however, is no longer supported, nor is it particularly suited for observations with the WFC3. A database of PSFs for this camera is being developed at: `https://www.stsci.edu/hst/instrumentation/wfc3/data-analysis/psf/psf-search`.

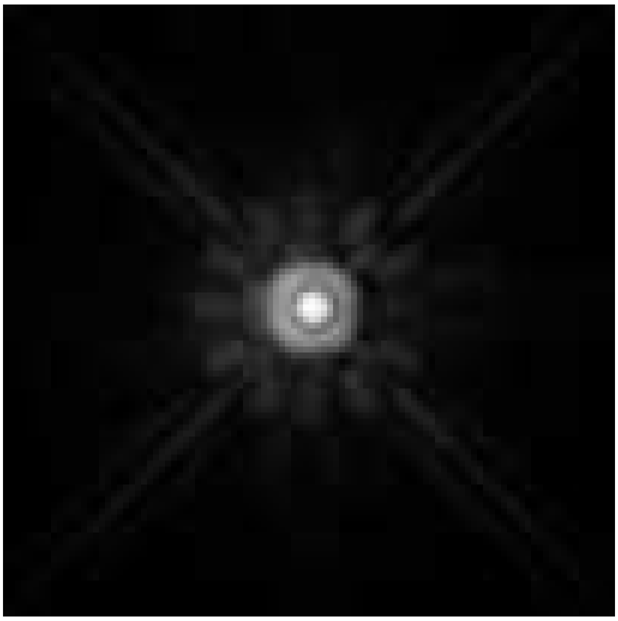

Figure 2.6.: PSF generated in `Tiny Tim` for the HST WFC3 F110W filter observation of NGC 2513 on 2016-10-30.

produces circular Gaussians. This is controlled by the axial ratio q', which determines the ellipticity of the Gaussians (see Section 2.3.2). A value of $q' = 1.0$ describes a circle. Hence by setting the bounds to $0.999 \leq q \leq 1.000$, one can ensure a model of circular Gaussians. After completing these preliminaries, the MGE fit can be executed to obtain the total counts and dispersions of each Gaussian, taking into account that the counts given in the results are not normalized.

The dispersions and non-normalized counts of the best fitting Gaussians are given directly in the results of the MGE fitting. After normalizing the counts, they can be entered together with the results for the dispersion into the MGE fitting function for the galaxy. Both the normalized counts (`normpsf`) and dispersions (`sigmapsf`) of the Gaussians are shown in Table 2.4. For comparison, the dispersions are also given in arcseconds, derived from multiplying the values in pixels with the scale of the WFC3 camera given above. These parameters are included as arguments in the MGE fitting procedure (see Main fitting procedures below).

The disadvantage of modelling PSFs as circular Gaussians via the MGE method is that high-resolution images, such as those from the HST, generally have PSFs that are non-circular, often displaying extended wings and asymmetries (Emsellem et al. 1994). Examples of both wings and asymmetry can clearly be seen in Figure 2.7, which displays the MGE fit of the PSF depicted in Figure 2.6. The black contours depict the PSF itself, while the red contours show the Gaussian fit.

Table 2.4.: Normalised counts and dispersions of the Gaussians reproducing the PSF shown in Figure 2.6 as obtained through the MGE fit.

`normpsf` (e^{-} s^{-1})	`sigmapsf` (pixels)	`sigmapsf` (arcsec)
0.46	0.99	0.13
0.36	3.15	0.41
0.09	10.72	1.39
0.04	20.61	2.68
0.05	35.27	4.59

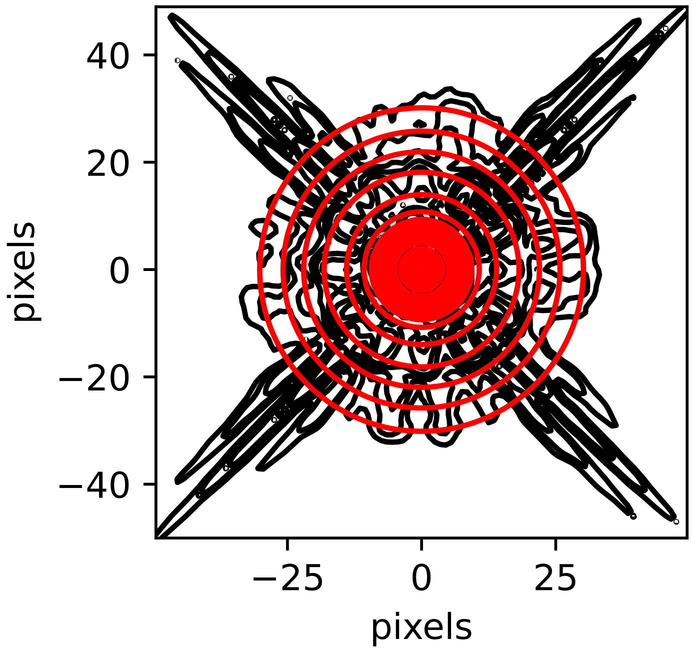

Figure 2.7.: MGE model of the PSF depicted in Figure 2.6 for HST's WFC3 F110W filter for the observation of NGC 2513 on 2016-10-30.

Main fitting procedures

In the following, I describe the procedures contained in the `MgeFit` package and how I use them to generate the MGE fit using the input parameters and PSF discussed and given above. The main references for this section are the README and documentation of the `Mgefit` python package, and Cappellari (2002).

The four procedures in the MGE fitting are as follows: locate the galaxy, obtain the galaxy's photometry, perform the MGE fit and, finally, display the results of the MGE fit as 2-D contour plots.

The first step is to locate the galaxy in the image. This is carried out in the procedure

FIND_GALAXY. Its purpose is to provide a precise quantitative definition of the galaxy's centre, its orientation with respect to the image axes (given by the position angle) and its ellipticity. In effect, it looks for the largest connection of pixels that lie above a certain level of intensity by computing the weighted first and second moments of the intensity distribution of the image. The plot of the FIND_GALAXY for the HST J-Band image of NGC 2513 is shown in Figure 2.8. Also depicted in the plot are the galaxy's major and minor axes (solid and dotted green line, respectively) .

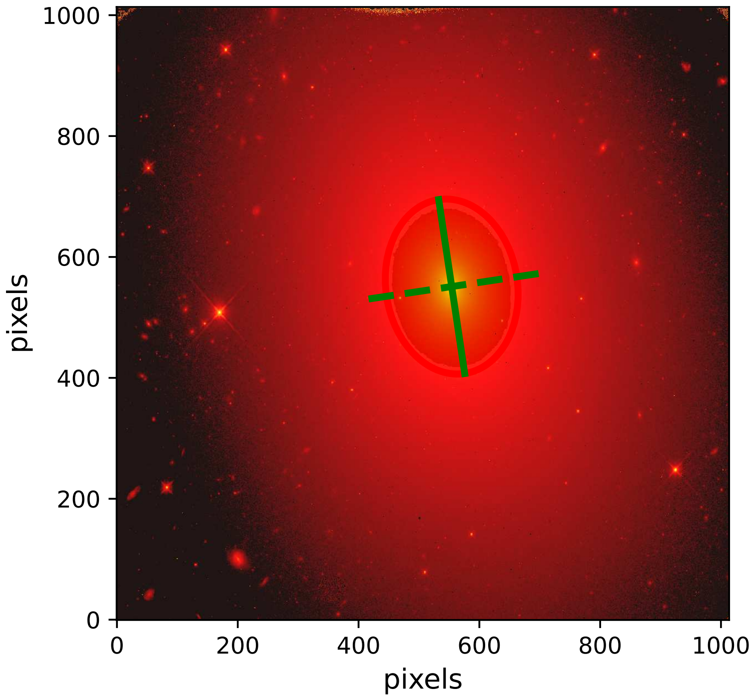

Figure 2.8.: Plot from the FIND_GALAXY procedure from the MGE fit of the HST J-Band image of NGC 2513 based on the input parameters and PSF given above. The major and minor axes of the galaxy (solid and dotted green line, respectively) are also depicted

The FIND_GALAXY procedure results in several output values that are used in the subsequent procedures. These are: the x and y coordinates of the galaxy centre (equivalent to the brightest galaxy pixel), the x and y coordinates of the luminosity weighted galaxy centre, the position angle θ measured clockwise from the image's x-axis, the astronomical position angle measured counter-clockwise from the image's y-axis, the galaxy's ellipticity, and the major axis (the maximum distance of the selected pixels from the luminosity weighted galaxy centre).

The ellipticity is given by the standard definition of $\varepsilon = 1 - b/a = 1 - q'$, where a and b are the galaxy's semi-major and semi-minor axis, respectively. Their ratio q' is the axial ratio or so-called "flattening" of the galaxy first introduced in Eq. 2.2. The values

Table 2.5.: Results from the `FIND_GALAXY` module for the MGE fit on NGC 2513.

Pixels used	40 474
Galaxy centre (x, y), (pixels)	553, 551
Luminosity weighted centre (x, y), (pixels)	554.22, 551.06
Theta (deg)	81.5
Astro PA (deg)	8.5
Ellipticity	0.25
Major axis (pixels)	132.2

obtained for NGC 2513 are given in Table 2.5.

The second procedure is `SECTORS_PHOTOMETRY`. In this procedure, the galaxy is divided into sectors that are linearly spaced in angle between the major (0°) and minor (90°) axes, and logarithmically spaced in radius. Photometry of the galaxy is then measured along the elliptical annuli along these sectors. Among the required input parameters for this procedure are the output from `FIND_GALAXY`, namely the galaxy's ellipticity, the position angle θ and the coordinates in pixels of the galaxy centre. In addition, optional parameters include the minimum count level discussed above under the input parameters and an argument that implements any masks that have been defined for the fit (see Section 2.3.6).

The photometry is given as three output vectors, R_j, θ_j and I_j (radius, angle and intensity), which are then used in the subsequent fitting procedure (see below). R_j is the radius in pixels of the surface brightness measurements, θ_j is the polar angle in degrees of the surface brightness measurements as taken from the galaxy's major axis, and I_j contains the surface brightness measurements themselves, which, for this image, are in $e^- s^{-1}$. In effect, R_j and θ_j are the polar coordinates of points at which the surface brightness counts are to be fitted. For the MGE fit of the NGC 2513 HST J-band image, 19 sectors were used in the fit. A plot showing the sectors produced via the `SECTORS_PHOTOMETRY` procedure is shown in Figure 2.9.

The third procedure, `MGE_FIT_SECTORS`, performs the MGE fit itself. As required input parameters, the procedure takes the results of the photometry derived in the previous step, R_j, θ_j and I_j, a well as the ellipticity from `FIND_GALAXY`. In addition, there are several optional input keywords that I specify. Several of these have been discussed above: the number of Gaussians, the `normpsf` and `sigmapsf` defining the PSF to be convolved with the galaxy image before fitting, and the image scale that defines the number of arcseconds to a pixel. In addition I set the arguments `linear` and `qbounds`. The former defines whether the fit should follow the fully linear or, alternatively, non-linear 2-D algorithms described and discussed in Cappellari (2002) (see also Footnote 12). For the purposes of this fit, I decide on the non-linear algorithm, which should suffice to give an accurate fit.

The keyword `qbounds` can be included to set a lower and upper limit to the observed axial ratio q' of the fitted Gaussians. If omitted, these lower and upper limits are set to a default of $[0.05, 1.0]$. The upper value of $q' = 1.0$ derives from the definition of q', which as a ratio of the semi-minor axis b to the semi-major axis a, cannot be greater

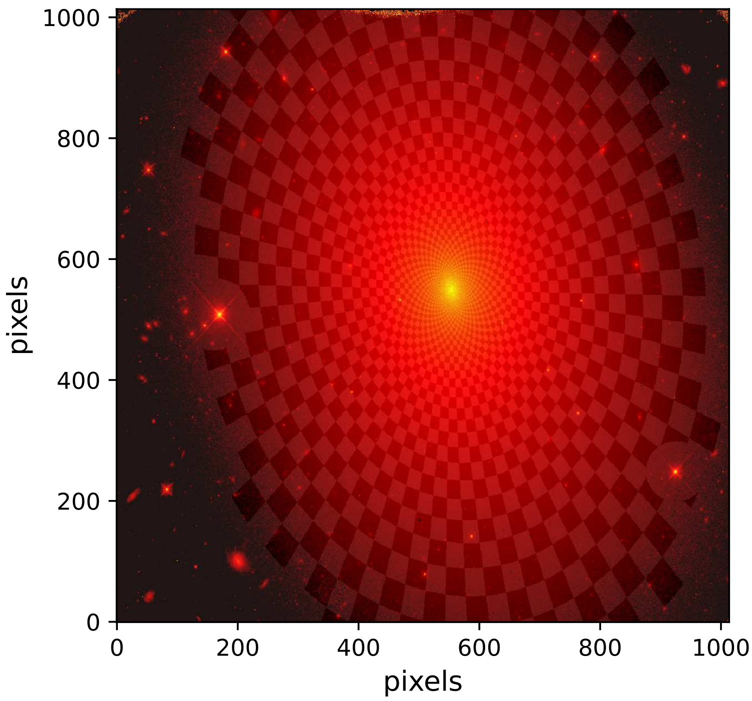

Figure 2.9.: Plot generated by the `SECTORS_PHOTOMETRY` procedure for the MGE fit of the HST J-band image of NGC 2513 showing the sectors along which the photometry is performed.

than 1.0 (see also the discussions for Eq. 2.2 in Section 2.3.2 and `FIND_GALAXY` directly above). The lower bound of $q' = 0.05$ means that no Gaussians flatter than $q' = 0.05$ will be included in the fit. The lower bound of the value q' should not be too flat (i.e., too small) as this can affect the deprojection of the image. This is discussed more fully below under the regularisation of the fit and in Chapter 4.

The main results of the fit are the total counts in $e^- s^{-1}$, the dispersion in pixels σ and the observed axial q' ratio of each best-fitting Gaussian component.[15] Also given are the standard deviation, mean absolute deviation between the fitted MGE and the data, and minimum χ^2 value of the fit.

Included in the `MGE_FIT_SECTORS` results is a series of one-dimensional plots depicting the convolved Gaussian components of the series expansion and resulting MGE best-fitting model to the data along the sectors created to perform the galaxy photometry. Examples of these plots were given in Figure 2.5.[16] The plots created in the MGE fit for NGC 2513 are shown in Figure 2.14, where they are included in the discussion on regularization and the final MGE results in Section 2.3.9.

[15] Here, the total counts and σ are the G_k and σ_k^* discussed in Section 2.3.2 and in the Section on the PSF.

[16] Also, compare the left panel of Figure 2.4 Here, though, the plot is produced by the 1-D version of the fitting algorithm.

The last procedure of the MGE fitting process is `MGE_PRINT_CONTOURS`. It uses the output of `MGE_FIT_SECTORS` to generate two-dimensional contour plots of the galaxy's isophotes (i.e., the data), and MGE best fit to the data (compare the right panel of Figure 2.4). These 2-D contour maps, besides displaying the final results, are particularly useful as a visual comparison of the model with the data and can be used in preliminary runs of the MGE fit to adapt parameters of the fit. The procedure produces both a map of the entire galaxy showing the fit in its extended regions and a zoomed-in map of the central regions so that the goodness of fit can be determined in both areas.

The construction of the plots is based on the the astronomical position angle (PA) as determined in `FIND_GALAXY`, the x and y coordinates of the brightest galaxy pixel from `FIND_GALAXY` (i.e., the galaxy centre) and the results of `MGE_FIT_SECTORS` (the total counts, dispersion σ and observed axial ratio q'). Other parameters that can be set are the image scale that allows the plot to be displayed in physical units, the PSF (`sigmapsf normpsf`) that must be convolved with the Gaussian components, and an argument that allows the implementation of any masks.

2.3.4. Masking

As noted directly above, the MGE procedures provide for the application of masks on the image to be modelled. Masking is a useful tool to either prevent unwanted light sources in the image from being modelled or to exclude areas obscured by dust. For the modelling of NGC 2513, examples of both instances occur. Masks were applied to two bright foreground stars (i.e., stars not belonging to NGC 2513, but to the Milky Way) and to the prominent dust lane in the centre of the galaxy.

2.3.5. Masking bright stars

To obtain an accurate parametrization of the galaxy's surface brightness and hence its intrinsic luminosity and luminous mass, it is important that only light from the galaxy is modelled. In producing an accurate MGE model, it is thus necessary to mask out bright foreground stars as these will adversely affect the mass model. Luminous stars that do not belong to the galaxy can produce significantly higher counts than any produced even by the galaxy's brightest pixels. This leads to bad fitting in those regions and a subsequently inaccurate model of the luminous mass. These stars can be easily masked by applying a boolean image mask, which, in the MGE procedure, has the same dimensions as the image. True values are included in the photometry, while false values will be ignored.

Stars can most often be effectively masked with a circular boolean array. The MGE fit code incorporates a function that defines a circular aperture which is useful for producing such circular masks. The function returns an array where the value of each element is its distance from a specified centre, which should approximate the centre of the bright star. The size of the mask is defined by the circle's radius, which is also specified by the user. The true values of the array are those areas of the image that are *not* masked, and the false values are those areas that *are* masked. For the two stars in the image of NGC 2513, I approximated the x and y coordinates of the centre of each star by creating a region

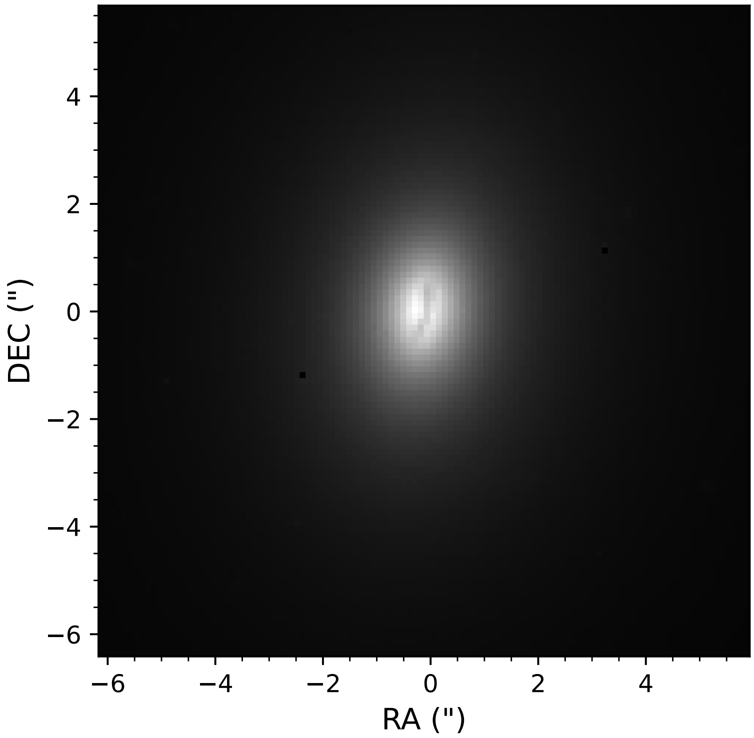

Figure 2.10.: A zoom in of the HST J-band image icv424clq clearly showing the dust running through the central region.

around each star in `SAOImageds9` and reading the centre from the region statistics. From the centre, I chose a radius of 50 pixels. After performing the MGE, I could then adjust the values slightly to improve the fit. The resulting masks are depicted in purple in the left panel of Figure 2.15.

2.3.6. Dust masking

Although elliptical and lenticular galaxies tend to contain far less gas and dust than spiral galaxies, a significant sample of these galaxies do contain dust. Sadler & Gerhard (1985) report an estimated 40 % of nearby ellipticals with dust, and high quality optical imaging has revealed that 50 %–80 % of nearby ellipticals and lenticulars contain dust (see e.g., Goudfrooij et al. 1994a,b; van Dokkum & Franx 1995; Tran et al. 2001, and references therein). The presence of dust often goes hand in hand with the presence of molecular gas (e.g., see Young et al. 2011, and references within.) and is therefore a good indicator of the presence of molecular gas. Thus it is not surprising that NGC 2513 has a significant dust lane that can clearly be seen in the zoom in of the HST J-band image in Figure 2.10.

The presence of dust, however, is problematic for creating an accurate luminous mass model. This is due to the absorption and scattering of light, predominantly in the blue wavelength regime where the wavelength is similar to the typical dust-grain size. This serves to effectively remove shorter wavelengths from the light reaching the observer, making the object appear dimmer (*extinction*) and redder (*interstellar reddening*). As the luminous mass is directly proportional to the amount of light observed, the loss of light in shorter wavelengths must result in a lower luminous mass.

If two images at two different wavelength filters are available, dust extinction can be corrected by examining the colour excess due to the reddening $E(B-V)$ in pixels across the affected region.[17] Because NGC 2513 only has high-resolution images available in one filter (F110W), any dust correction is restricted to this one filter and an alternative method must be used. In the following sections, I introduce the methods explored in this thesis.

2.3.7. Dust masking using the galaxy's radial surface brightness profile

Thater et al. (2017) developed a dusk-masking method that uses the galaxy's characteristic radial surface brightness profile to ascertain which pixels might be affected by dust. This is done by fitting the lower envelope of the surface brightness profile with a suitable function and identifying those pixels that lie below the fit. Assuming that most pixels at a certain radius will have a similar surface brightness, the pixels that lie below the average are those that have suffered extinction due to dust. The following details the method and my approach for the image of NGC 2513.

A galaxy's radial surface brightness profile characterises the surface brightness as a function of radius r, i.e., as a function of the distance from the galactic centre. The surface brightness is computed by averaging over all the pixels in concentric circular or elliptical annuli.

For the radial profile of NGC 2513, I initially tried out profiles constructed on both circular and elliptical annuli. To construct the former, I used the function from the `mge_fit_example` file available in the `MgeFit` package described above in Section 2.3.5. As mentioned, this produces an array in which the value of each element is its distance from the specified centre. For the radial profile constructed using elliptical annuli, I used a function based on the definition of an ellipse. Both profiles are shown in Figure 2.11.

I chose to work first with the radial surface brightness profile constructed with circular annuli, partly because the `MgeFit` example code already incorporates the function mentioned above for the specific purpose of constructing masks for the MGE fit.

I first plotted the surface brightness profile shown in the left panel of Figure 2.11 against a radius going out to $\sim 5''$ from the galaxy's centre and then fitted the lower envelope of the curve with an exponential function. The radial surface brightness profile and fit are shown in Figure 2.12. The pixels falling below the fitted curve are, in principle, those pixels affected by dust obscuration and it is these pixels that need to be masked for the MGE fit.

The mask was constructed by defining a boolean array across the pixels falling below

[17] A full description of this method can be found, for example, in Cappellari (2002), Thater et al. (2017) and Cohn et al. (2021).

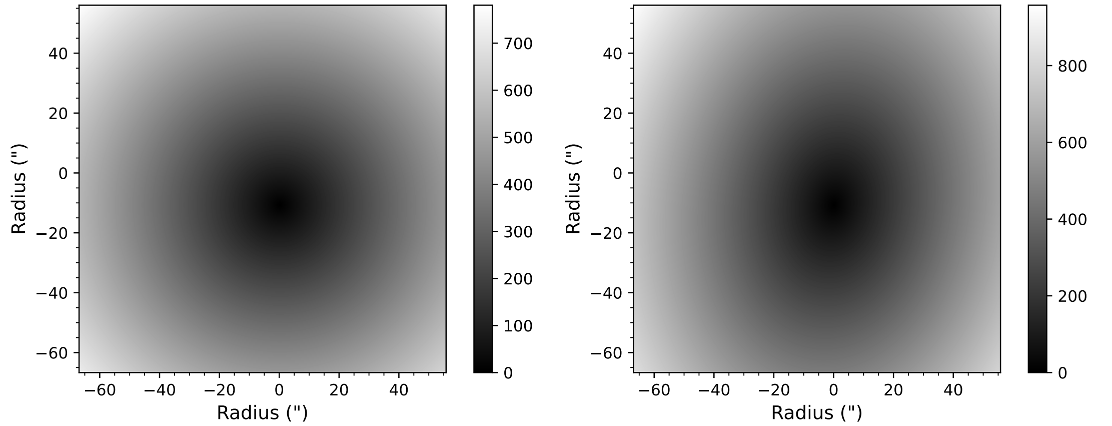

Figure 2.11.: Radial profiles of NGC 2513 using concentric circular annuli (*left panel*) and concentric elliptical annuli (*right panel*). The colourbar indicates distance from the galaxy's centre.

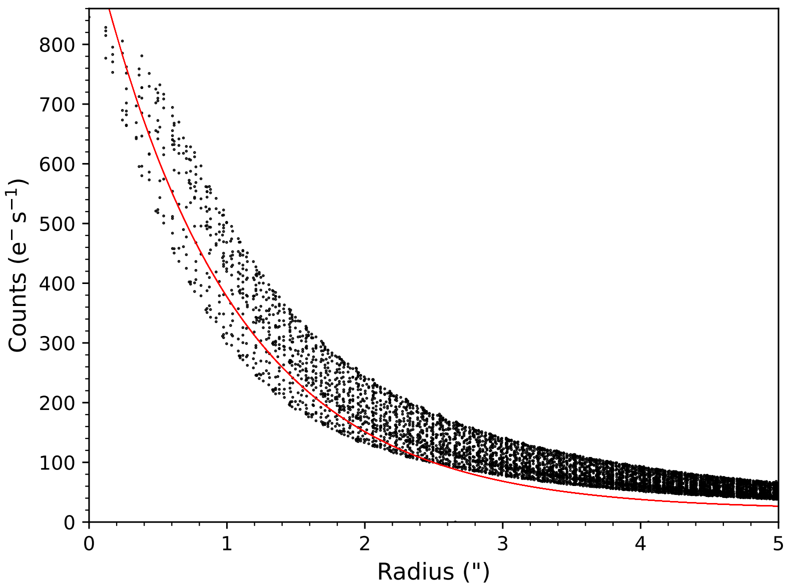

Figure 2.12.: Characteristic surface brightness profile of the first ∼5″ for NGC 2513. Each dot is one pixel in the image, while the red curve is the fit for the surface brightness. Pixels falling below this curve are those pixels affected by dust obscuration that need to be masked for the MGE fit.

the fitted curve. To test whether the mask fitted the shape of the dust lane in the galaxy, I then plotted the mask superimposed on the inner region of the galaxy. This is shown in the left panel of Figure 2.13. The right panel shows the same region of the galaxy, but without the mask for better comparison. (Note that as the mask is constructed in units of pixels, I likewise plot the image offsets in pixels.)

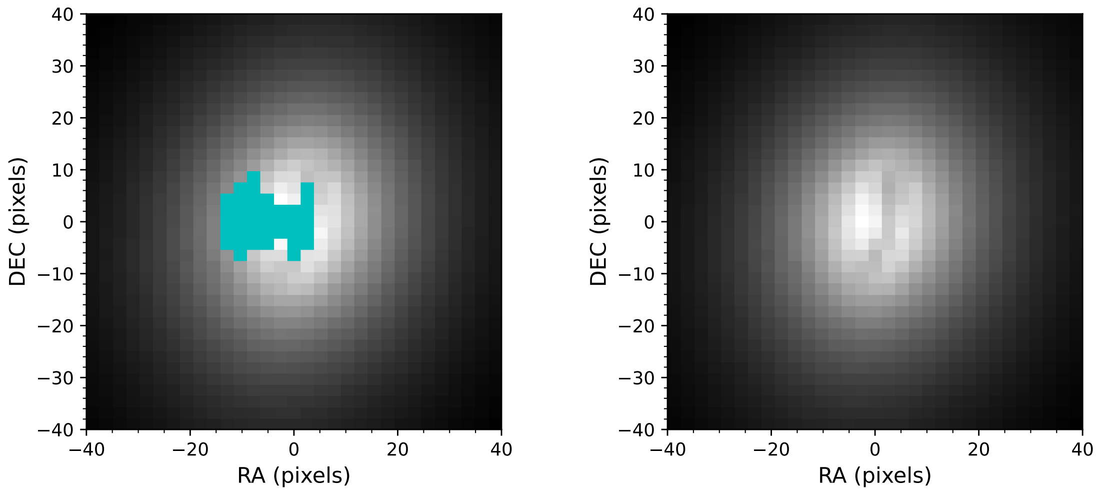

Figure 2.13.: *Left panel*: Zoom-in of the HST J-band image of NGC 2513 overlaid with the mask derived from fitting the radial profile of the galaxy and masking the pixels below the fit. *Right panel*: The same zoom-in of the HST J-band image of NGC 2513 without mask showing the dust lane for clearer comparison with the mask.

The right section of the mask (approximately RA pixels 3−4) does mimic the dust lane fairly well. However, a substantial number of pixels not belonging to the dust lane are also masked in an extended region to the left of the image's centre.

To double-check the procedure, I created a mask using a radial profile constructed from the elliptical annuli shown above in Figure 2.11. The resulting mask had a slightly different shape, but still comprised a large extended area to the left of the galaxy's dust lane.

While this method has been successfully used on the galaxies NGC 4414 in Thater et al. (2017) and NGC 6958 in Thater et al. (2021), it failed to produce an adequate mask for this galaxy. The reasons for this are not, as yet, fully understood, but will be further examined in future work.

2.3.8. Direct masking

An alternative, albeit less rigorous method, is to simply identify the dust-affected pixels by eye. I proceeded by opening the HST image in `SAOImageds9` and enlarging the image so that the individual pixels were clearly visible. Based on the coordinates of the dust-affected pixels, which were easily identifiable, I created a single mask or rectangular shape, again by defining a boolean array. I then tested the mask in a trial MGE fit of the galaxy image with the dust mask included. The results proved to be a worse fit than the MGE created without any dust mask, with a notably higher χ^2 value in the masked version.

On re-evaluating the region affected by dust in the image, it seemed that the region could be better covered by two separate rectangularly shaped masks. I applied these masks and again ran a trial MGE fit to test the mask's effectiveness. In this case, the results were excellent with a good contour fit and lower χ^2 value than in the unmasked MGE fit. The dust mask is depicted in purple in the right panel of Figure 2.15.

2.3.9. Regularisation and final results of the MGE fit

After performing the MGE and adjusting the parameters to create the best possible fit, the MGE fit must be regularised. The regularisation process serves to increase the minimum value of the observed axial ratio q' of the Gaussians. (Recall that the higher q' is, the rounder the Gaussian, and the lower q' is, the more elliptical or "flatter" the Gaussian.) The preference for higher q' values derives from the relation of the deprojected intrinsic axial ratio q (3-D) to the projected value of q' (2-D) and the inclination i of the galaxy. For an oblate axisymmetric case, which is what I model here,[18] this is given by

$$q^2 = \frac{q'^2 - \cos^2(i)}{\sin(i)} \tag{2.5}$$

(see Eq. 10 in Cappellari (2002)). It is clear that for $q'^2 < \cos(i)$, the relation is undefined. Thus, the lowest q' value, i.e. the flattest Gaussian in the MGE fit, dictates the minimum possible inclination for which the MGE can be used in dynamical modelling (Cappellari 2002). The Gaussians themselves, however, are not real physical quantities and thus do not necessarily have physical significance. Thus, any minimum inclination imposed by the flattest Gaussian does not necessarily provide for robust dynamical modelling (Cappellari 2002)).

The process of regularisation aims to counteract this problem. It does so by increasing the minimum axial ratio q' in the MGE fit until the galaxy contour can no longer be accurately reproduced by any number of Gaussians. This maximum value of the flattest q' (q'_{max}) gives a good approximation of the minimum inclination for which the galaxy can be deprojected (Cappellari 2002). The regularisation is implemented with `MGE_FIT_SECTORS_REGULARIZED`, which, as a wrapper for `MGE_FIT_SECTORS`, accepts the same input arguments as the latter.

An interesting instance of q' values that were too low occurred in the initial MGE fit I performed on NGC 2513. A reasonable assumption is that the minimum inclination of the galaxy should be $i \sim 50°$, where $0° \leq i \leq 90°$, with $i = 90°$ corresponding to edge-on and $i = 0°$ to face-on systems. With the default minimum value of $q' = 0.05$ in the sectors fitting procedure, the regularised MGE fit produced a flattest Gaussian with $q' = 0.182$. From the preliminary dynamical modelling with `KinMS`, it became apparent that this minimum q'-value resulted in an inclination too low to be successfully modelled (see Chapter 4.) For $q' = 0.182$, even an inclination value of $i = 90°$ produces an undefined result (note also that for $i = 0°$, the deprojection is degenerate). To avoid this problem, it was decided to set a minimum q-bound value of 0.5, i.e., `qbounds` $= [0.5, 1.0]$. This

[18] For simplicity, one assumes oblate axisymmetry in dynamical modelling of molecular gas with `KinMS`.

resulted in a minimum q'-value of 0.65, which, using Eq. 2.5 as a guide, gave a minimum inclination of 49.5°, which is comparatively close to the assumption of 50° quoted above and which should enable reasonable dynamical modelling.[19]

The final results of the regularised MGE fit for a total number of $N = 10$ Gaussians are given in Table 2.6, along with the final χ^2 value, standard deviation and mean absolute deviation.

Table 2.6.: Final regularised MGE results for the HST/WFC3 F110W light profile of NGC 2513.

Total counts ($e^-\,s^{-1}\,pixel^{-1}$)	σ (pixels)	Observed q
2260.58	0.742835	0.7
46981.6	6.07054	0.65
57534.5	7.46661	0.75
138248	12.7695	0.65
222958	22.6232	0.686922
181567	35.9078	0.740954
339662	63.9312	0.75
489200	119.687	0.701065
442293	188.835	0.748412
803981	303.464	0.65
Final qbounds = $[0.650, 0.750]$		
$\chi^2 = 0.642$		
Standard deviation = 0.026		
Mean absolute deviation = 0.016		

Notes. For each of the 10 Gaussian components, Col. 1 lists the surface brightness, Col. 2 the dispersion (standard deviation) and Col. 3 the axial ratio.

The series of 1-D plots showing the MGE best-fitting model to the HST WFC3/F110W data (referred to above under `MGE_FIT_SECTORS`) are shown in the left panel of Figure 2.14. The plots also show the individual 10 Gaussian components of the expansion used for the fit. Each plot shows the results for one of the 19 sectors that are linearly spaced in angle between the major and minor axes of the galaxy to cover the entire galaxy and along which the photometry was measured. Seven of these sectors are depicted here. The right panel shows the radial variation of the relative error along the profiles.

Contour plots (see `MGE_PRINT_CONTOURS` above) of the final results of the MGE fit are shown in Figure 2.15. Both plots show the galaxy isophotes of the HST WFC3/F110W J-band image in black with the best-fit MGE contours superimposed in red. The left panel shows the whole galaxy out to a radius of $\sim$50″ ($\sim$16.89 kpc), including the bright foreground stars that have been masked in purple. The right panel shows a zoom-in of

[19] A minimum search range value of $i_{min} = 60.5$ was eventually used in the KinMS modelling (see Chapter 4).

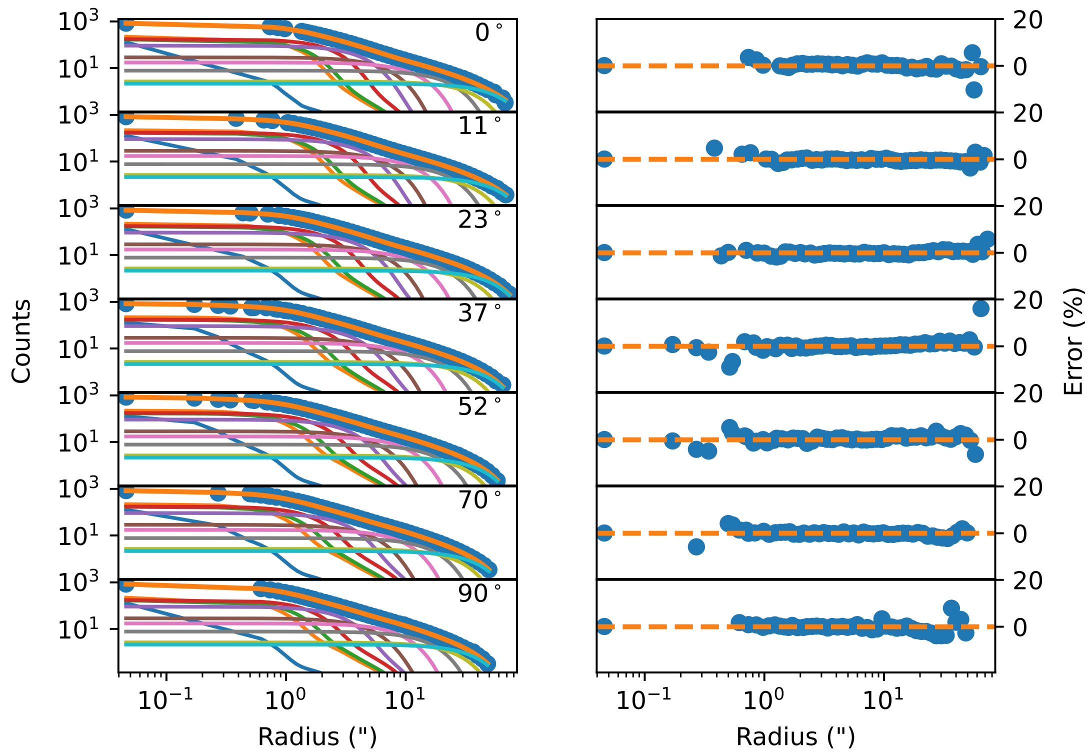

Figure 2.14.: *Left panel*: 1-D plots showing the MGE best-fitting model (orange curve) to the HST WFC3/F110W data (blue dots). The results for seven of the 19 sectors along which the photometry was measured are depicted here. The sectors are linearly spaced in angle between the major (0°) and minor (90°) axes of the galaxy. Also depicted are the $N = 10$ Gaussian components of the expansion (coloured curves) used for the fit. *Right panel*: Variation of the relative error along the profiles. (Cf. Figure 6 in Cappellari (2002))

the central region out to a radius of $\sim 2''$ ($\sim$675.71 pc) with dust mask in purple. The effects of the dust and its masking are visible in the poor alignment of the central red MGE contour to the black contour of the image, which is shifted off-centre to the left. Despite this inconsistency, the overall fit for the central region is excellent. The fit for the extended regions of the galaxy is also reasonably good, with small deviations from the data mainly in the south-east region around the bright star that has been masked here.

2.3.10. Converting to surface brightness in physical units

The results of brightness for each Gaussian and their dispersion (total counts and σ, respectively) that are obtained from the MGE modelling are in instrumental units: the brightness is given as electrons per second, and the dispersion is given in pixels. These units need to be converted to physical units that can be used in dynamical modelling. A standard example for this conversion is given in the README of the `MgeFit` documentation and

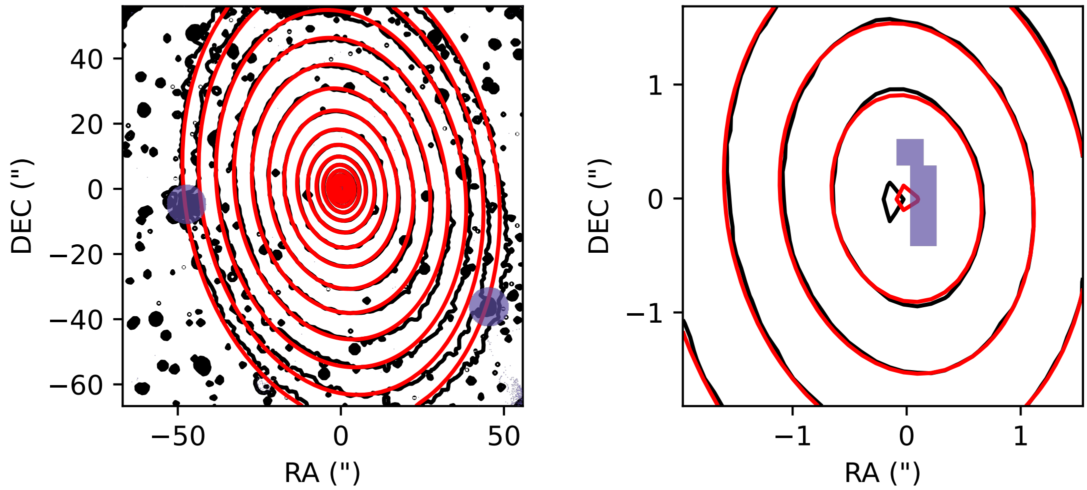

Figure 2.15.: Contour plot of the regularised MGE model (red contours) overlaid on the isophotes of the HST WFC3/F110W image of NGC 2513. *Left panel:* Extended view of the entire galaxy, showing the two bright stars that have been masked (purple). *Right panel*: Zoomed-in view of the inner regions around the galaxy's centre. The central dust-lane has been masked for the fit (purple).

is, except for the first step, based on standard photometric procedures.

The process consists of three steps: 1) convert the total counts of each Gaussian into the corresponding peak surface brightness C_0 in counts pixel^{-1}; 2) convert the instrumental units to surface brightness in mag arcsec^{-2} for the particular filter of the camera being used; 3) convert these surface brightness values into surface brightness I' in $L_\odot\,\mathrm{pc}^{-2}$.

The first step is carried out with the following equation, which is derived using Eq. 1 in Cappellari (2002) and which is Eq. 1 in the README:

$$C_0 = \frac{\mathrm{TotalCounts}}{2\pi\,\mathrm{SigmaPixels}^2 q\mathrm{Obs}} \ . \tag{2.6}$$

Since the HST image used here has the units of $e^-\,s^{-1}$, the brightness of each Gaussian is also given in $e^-\,s^{-1}$ and these must first be converted to photon counts. This is done by multiplying each brightness value by the exposure time (in seconds) to get to units of electrons and then dividing these results by the gain G, which is defined as $G = N_e/N_c$, where N_e is the number of electrons per pixel and N_c is the number of photon counts per pixel.[20]

The conversion to surface brightness is adapted from Eq. 2 in the README[21] for the WFC3/F110W filter, which is a J-band filter (see Section 2.2.2). The surface brightness μ_J in mag arcsec^{-2} is given as:

$$\mu_J = ZP + 0.1 + 5\log_{10}(\mathrm{SCALE}) + 2.5\log_{10}(\mathrm{EXPTIME}) - 2.5\log_{10}C_0 - A_J \ , \tag{2.7}$$

[20]See e.g., `http://spiff.rit.edu/classes/phys445/lectures/gain/gain.html`
[21]This equation is for the I-band WFPC2/F814W filter as adapted from Holtzman et al. (1995).

Table 2.7.: Values for conversion to surface brightness μ_J.

Vega ZP (mag)	AB ZP (mag)	Scale (arcsec pixel^{-1})	Exposure time (s)	A_J (mag)
26.042	26.819	0.121	499.234	0.02

Notes. Vega ZP is the Vega magnitude zero point; similarly, AB ZP is the AB magnitude zero point.

where ZP is the photometric zero point, SCALE is the pixel-to-arcsecond scale, EXPTIME is the exposure time, C_0 is the counts from Eq. 2.6 and A_J is the galaxy extinction in the J-band. The value 0.1 is a correction for infinite aperture as given in the README.

For the calculation, I used both the Vega magnitude zero point and AB magnitude zero point for comparison, taking the values for both from the STScI IR Inverse Sensitivity (Zeropoint) Tables.[22] For the KinMS modelling, however, only those values derived from the Vega magnitude zero point were used (see Chapter 4). The scale is that derived for the image as described above in Section 2.2.2. The exposure time was obtained from the image header and the galaxy extinction A_J for NGC 2513 as imaged by the WFC3/F110W filter is from Schlafly & Finkbeiner (2011) as retrieved from the NASA/IPAC Extragalactic Database (NED).[23] The values used for the calculation can be found in Table 2.7. The values for C_0 and resulting μ_J are given in Table 2.8.

The last step is to convert the surface brightness μ_J, which is expressed in units of mag arcsec^{-2}, to a surface brightness I' in $\mathrm{L}_\odot\,\mathrm{pc}^{-2}$. This is done by means of the following expression, which is Eq. 3 in the README:

$$I' = \left(\frac{64800}{\pi}\right) 10^{0.4(M_{\odot,J}-\mu_J)}\ , \tag{2.8}$$

where $M_{\odot,J}$ is the absolute magnitude of the Sun in the J-band. The absolute solar magnitude values for both the Vega and AB zero point were taken from Willmer (2018).[24] As values for $M_{\odot,J}$ for the WFC3/F110W filter were not available, I used the $M_{\odot,J}$ values for the HST NICMOS 3 camera's F110W filter. A comparison with the values of WFC3's F105W filter showed that the values were very similar. Both sets of values are listed in Table 2.9.

The results derived from Eq. 2.8 are the converted physical values for the surface brightness. The dispersion σ can be trivially converted from pixels to arcseconds by multiplying all values obtained from the MGE with the image scale. The final resulting Gaussian parameters (which, as noted in the README, are distance independent) are given in Table 2.10, and are ready to use for dynamical modelling.

[22] https://www.stsci.edu/hst/instrumentation/wfc3/data-analysis/photometric-calibration/ir-photometric-calibration
[23] https://ned.ipac.caltech.edu/
[24] As retrieved from https://mips.as.arizona.edu/~cnaw/sun.html

Table 2.8.: Results for conversion to surface brightness μ_j.

C_0 (counts pixel^{-1})	μ_J Vega ZP (mag arcsec^{-2})	μ_J AB ZP (mag arcsec^{-2})
186003.626	15.11	15.88
2336.739	16.29	17.07
43732.513	16.68	17.46
41455.485	16.74	17.51
20155.420	17.52	18.30
6040.215	18.83	19.61
3521.634	19.41	20.19
548.169	20.31	21.08
526.731	21.47	22.25
426.877	21.71	22.48

Notes. Columns: (1) C_0 is the peak surface brightness in counts pixel^{-1} for each respective Gaussian of the fit. (2) Surface brightness in mag arcsec^{-2} derived from the Vega magnitude zero point using Eq. 2.7 with the scale, exposure time and extinction given in Table 2.8. (3) The same as for Col. 2, but using the AB magnitude zero point.

Table 2.9.: Values for the solar absolute magnitude.

	$M_{\odot,J}$ Vega ZP (mag)	$M_{\odot,J}$ AB ZP (mag)
NIC3/F110W	3.82	4.52
WFC3/F105W	3.89	4.51

The deprojection to the 3D surface brightness density is carried out within the modelling procedures of`KinMS` (see Chapter 4). As a preliminary check of the soundness of the Gaussian parameter values obtained after conversion, one can, however, estimate the galaxy mass by calculating the galaxy's total luminosity. This can be done by summing over the total luminosity of each Gaussian as given by

$$L_j = 2\,\pi\, {I'}_j {\sigma'}_j^2 q'_j \ , \tag{2.9}$$

which is Eq. 4 in the README. For this calculation, the unit of σ'_j needs to be converted from arcseconds to parsecs. Multiplying the resulting total luminosity by the mass-to-light ratio gives the galaxy mass (see Eq. 2.1). Using Eq. 2.9, I calculated the total luminosity for NGC 2513 for both the Vega and AB zero points, respectively and, using a value for the mass-to-light ratio of $M/L = 2.0$ that was used for preliminary `KinMS` modelling, I obtained respective estimates of the galaxy's total luminous mass. These are given in Table 2.11 where they are compared to the stellar mass for NGC 2513 estimated by Veale et al. (2018) from the total absolute K-band magnitude M_k using Eq. 2 in Cappellari (2013b). The absolute K-band magnitudes themselves were computed with the method

Table 2.10.: Final Gaussian parameters.

j	I'_j Vega ($L_{\odot,J}\,pc^{-2}$)	I'_j AB ($L_{\odot,J}\,pc^{-2}$)	σ'_j (arcsec)	q'_j
1	12996.838	12107.032	0.090	0.70
2	4355.724	4057.517	0.734	0.65
3	3055.770	2846.562	0.903	0.75
4	2896.665	2698.350	1.545	0.65
5	1408.342	1311.922	2.737	0.687
6	422.055	393.159	4.344	0.741
7	246.071	229.224	7.734	0.75
8	108.177	100.771	14.480	0.701
9	36.805	34.285	22.845	0.748
10	29.827	27.786	36.713	0.65

Notes. Columns: (1) Summation index j of the Gaussians. (2) F110W surface brightness of each respective Gaussian using the Vega zero point. (3) As for Col. 2, but using the AB zero point. (4) The dispersion (standard deviation) of each respective Gaussian. (5) Axial ratio of each respective Gaussian.

described in Ma et al. (2014) based on magnitudes from the 2MASS XSC database.

The masses obtained from the Vega and AB zero points in this work are fairly similar, but both are lower than the estimate from Veale et al. (2018). This may be due to the fact that the mass estimate from Veale et al. (2018) is in the K-band, which, centred at 2.20 µm, is less susceptible to dust than the J-band, which is centred at the shorter wavelength of 1.25 µm. As mentioned in Section 2.3.6, a longer wavelength loses less light than a shorter wavelength to dust obscuration, which in turn leads to a higher luminous mass.[25] The difference in the mass estimates from this thesis to that in Veale et al. (2018) are notable, but nonetheless fall into the same order of magnitude and thus seem reasonable.

[25] Another contribution to the higher mass may be the field of view of the images used but this information was not immediately available.

Table 2.11.: Values for total luminosity and mass estimate of NGC 2513.

L_j Vega ($L_{\odot,J}\,pc^{-2}$)	L_j AB ($L_{\odot,J}\,pc^{-2}$)	$\log_{10}(M_*)$ Vega ($M_\odot$)	$\log_{10}(M_*)$ AB ($M_\odot$)	$\log_{10}(M_*)$ MASSIVE ($M_\odot$)
6.35×10^{10}	5.92×10^{10}	11.10	11.07	11.69

Notes. Columns: (1) Total luminosity of NGC 2513 using the Vega zero point. (2) Total luminosity of NGC 2513 using the AB zero point. (3) Mass estimate of NGC 2513 using the luminosity obtained with the Vega zero point and $M/L = 2.0$. (4) Same as for Col. 3, but using the AB zero point. (5) Mass estimate of NGC 2513 by Veale et al. (2018) from the 2MASS XSC total absolute K-band magnitude Ma et al. (2014).

3. Measuring the molecular gas kinematics

The molecular gas kinematics were obtained from the ALMA observations of the carbon monoxide CO gas in the vicinity of the SMBH. In order to do this, the ALMA data was first calibrated and then imaged to obtain the continuum emission of the gas and a *velocity cube*[1] of the line emission of the gas, in this case the ^{12}CO(2-1) line. A velocity cube encapsulates the observed velocity of the molecular gas as a function of its position. Through modelling the cube with the `KinMS` modelling tool, the total potential of the galaxy within the radius in which the gas has been observed can be obtained (see Chapter 4).

Moment maps of the intensity, velocity and velocity dispersion, as well as a position-velocity-diagram (PVD) of the gas can be extracted from the velocity cube. These act as a useful visual aid in assessing not only the quality of the imaged cube, but the nature of the distribution and kinematics of the gas. In particular, the velocity map should indicate whether the motion of the gas is Keplerian and the PVD should indicate the signature increase of velocity near the galaxy centre that is characteristic of gas in orbit around a SMBH. The maps and PVD from the observed data can also be compared with those derived from the model produced with `KinMS`.

In this Chapter, I first introduce the concept of interferometry and its importance in providing improved angular resolution. I then describe and discuss the imaging process, which uses the Common Astronomy Software Applications (`CASA`) package.[2] In conclusion, I present and discuss the moment maps and PVD produced from the imaged velocity cube.

3.1. Radio Interferometry

3.1.1. Motivation: A quest for resolution

Radio interferometry, otherwise known as *aperture synthesis*, is a technique that allows for vast improvement in spatial (i.e., angular) resolution. Its principle is to combine the signals received from the sky by two or more physically separated antennas thereby mimicking a single telescope of much larger diameter. The significance of a large diameter size is apparent from the definition of a telescope's angular resolution, which is diffraction limited. That is, it is subject to the Rayleigh criterion, which stipulates that for any

[1] The velocity cube is sometimes referred to as a data cube.
[2] https://casa.nrao.edu/

circular aperture, two point sources of equal brightness are just resolved when the peak of the diffraction pattern of one lies directly on the first minimum of the other's diffraction pattern. This is equivalent to the minimum angle of separation between the two sources, which for small angles, where $\sin(\theta) \approx \theta$, is given by:

$$\theta \approx \frac{1.22\,\lambda}{D}\ , \tag{3.1}$$

where θ is the minimum angle of separation in radians, λ is the wavelength and D is the aperture's diameter (see Schroeder 2000a).[3] The larger the diameter, the smaller the angular limit of resolution, i.e., the higher the resolution. For radio and millimetre wavelengths, however, even a large diameter gives poor resolution in comparison to optical wavelengths. While current ground-based optical telescopes (the largest of which are in the range of 10 m) usually have an angular resolution of $1''$ or better, the 10-m Effelsberg single-dish radio telescope has an angular resolution of $35''$ at a wavelength of 1.3 cm.

There is a technological limit to the size that a single dish antenna can reach[4] and therefore a limit on the angular resolution. By combining an array of smaller dishes, this limit can be overcome. The distance between any two antennas in an array is referred to as the *baseline*[5] and interferometers can offer baseline lengths in the range of kilometres, thus significantly increasing angular resolution. In the case of ALMA, the maximum baseline is 16 km. For a frequency of 230 GHz (which is approximately equivalent to a wavelength of 1.3 mm), the angular resolution is in the range of 20 mas.[6]

3.1.2. Basic principles of interferometry

The following provides an explanation and discussion of certain key concepts in radio interferometry that are necessary for understanding the data and the imaging process. Unless otherwise specified, the main sources of reference are Cortes et al. (2023), the lecture 'Aperture Synthesis' from the National Radio Observatory (NRAO) Crash Course in Radio Astronomy and Interferometry[7] and lectures from the NRAO's 16th Synthesis Imaging Workshop.[8]

Radio telescopes cannot detect individual photons in the manner that optical telescopes do, as the energies of radio photons are far too low. Instead, the electromagnetic (EM) power from a source in the sky is converted to voltage oscillations (voltage response) that can be processed and measured.[9] The EM power per unit frequency that is received

[3] See also https://pressbooks.bccampus.ca/introductorygeneralphysics2phys1207/chapter/27-6-limits-of-resolution-the-rayleigh-criterion/

[4] The largest current single-dish antenna is the FAST radio telescope in China with a diameter of 500 m and total reflective area of 300 m.

[5] As discussed in Section 3.1.2, the baseline is, strictly speaking, a vector defining both the direction and physical separation of the two antennas.

[6] See https://almascience.eso.org/about-alma/alma-basics.

[7] https://science.nrao.edu/science/meetings/presentation/jdf.webinar.2.pdf

[8] https://science.nrao.edu/science/meetings/2018/16th-synthesis-imaging-workshop/16th-synthesis-imaging-workshop-lectures.

[9] Reference: Lecture series on Radio Astronomy, University of Vienna, Winter Semester 2021-2022 by Hans Oloffson, https://ufind.univie.ac.at/en/course.html?lv=280527&semester=2021W

by an antenna, dP_{rec}, from a distant, unpolarized source is proportional to the source's brightness, i.e., its specific intensity I_ν, and to the effective area A_e of the receiver:

$$dP_{rec} = I_\nu A_e d\nu d\Omega \ , \tag{3.2}$$

where $d\nu$ is a small range (bandwidth) of frequencies and $d\Omega$ is the solid angle from which the power is received. Here, it should also be noted that the preferred unit for power is the *Jansky* (Jy) where $1\,\mathrm{Jy} = 10^{-26}\,\mathrm{W\,m^{-2}\,Hz^{-1}}$.

The *antenna power response*, i.e., its *sensitivity*, is a summation of all EM power brought to the focus of the telescope. To find an expression for this total power as received by an antenna from the two-dimensional sky, we first define θ and φ as orthogonal directional sky coordinates. We can then define the normalized power response of an antenna as $P_n(\theta, \varphi)$ as the power radiated per steradian in the direction (θ, φ) divided by the peak power P_{max} per steradian in that direction:[10]

$$P_n(\theta, \varphi) = \frac{P(\theta, \varphi)}{P_{max}} \ . \tag{3.3}$$

In addition, we define $I_\nu(\theta, \varphi)$ as the directional function of the sky brightness. The *total power* received by an antenna is then given by the integral of the sky brightness distribution and normalized power response over frequency and solid angle for the effective area of that antenna:

$$P_{rec} = A_e \iint_{4\pi} I_\nu(\theta, \varphi) P_N(\theta, \varphi) d\nu d\Omega \ . \tag{3.4}$$

Due to diffraction, the antenna response is dependent on the angle taken from the on-axis pointing direction of the antenna. For on-axis wavefronts where the paths of the EM power are all equal, the power is summed constructively at the focus, while for off-axis wavefronts where the paths are unequal, the power is not always summed constructively at the focus. In particular, for off-axis angles of λ/D radians (where λ is the observing wavelength and D is the antenna's diameter), the path-difference across the diameter results in destructive interference at the focus. Additionally, because the projected diameter along the off-axis direction is less than the true diameter, the power received from that direction (and hence, the antenna's relative sensitivity) is less than for the on-axis case. Figure 3.1 illustrates the two cases.

Several important definitions derive from the diffraction pattern and power response of an antenna. These are illustrated in Figure 3.2, which shows a representation of the normalized one-dimensional antenna power response given for an idealized 12-m parabolic antenna that is uniformly illuminated at 350 GHz.

The dependence of the power response on the angular offset from on-axis wavefronts is clearly visible, with the power response rising and declining according to the constructive and destructive interference of the wavefronts. The central Gaussian-like feature is referred to as the *primary beam* or *antenna beam size*, while the smaller Gaussian features that

[10] See `https://www.mikepeel.net/physics/Radio%20Astronomy/Lecture%20notes.pdf`. Accessed 24.02.25 at 22:29:51.

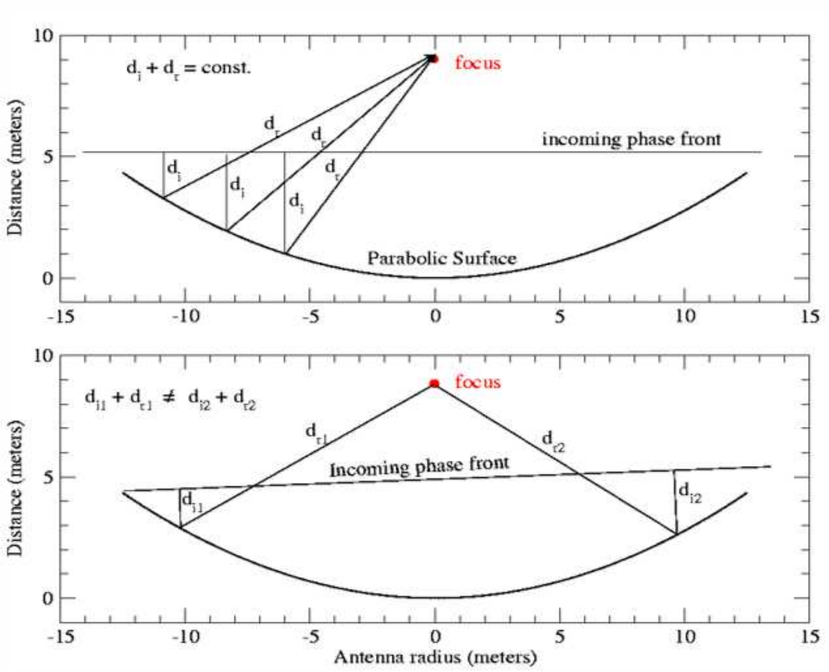

Figure 3.1.: Schematic of an incoming plane-parallel wavefront to an antenna of diameter D which is then reflected to a focus. The x-axis is the antenna radius in metres and the y-axis is the distance from the antenna surface in metres. The *top panel* demonstrates the on-axis case where the paths of the EM power are all equal and power is summed constructively at the focus. The *bottom panel* demonstrates the off-axis case. Here, the path lengths of the EM power differ, which results in less constructive interference at the focus. Figure taken from Cortes et al. (2023).

form at larger angles are referred to as *sidelobes*. Although the first side-lobes have a relative response of only 1.74 % that of the primary beam, incident emission at angles larger than those of the primary beam, if bright enough, can contribute significantly to the received EM power. Part of the imaging process is to reduce the effects of these sidelobes.

The *width* of the primary beam is described by the Half Power Beam Width (HPBW) or, analogously, Full Width at Half Power (FWHP). In one dimension, this is simply equivalent to the Full Width at Half Maximum (FWHM). In radians, this is given as:

$$\mathrm{HPBW}_{\mathrm{PrimaryBeam}} = \frac{1.02\ \lambda}{D} \ , \tag{3.5}$$

where D is the diameter of the receiver. The angular distance between the first nulls (i.e., where the interference is at a maximum and hence completely destructive) is called the First Width Between Nulls (FWBN) and is given, again in radians, as:

$$\mathrm{FWBN}_{\mathrm{PrimaryBeam}} = \frac{2.44\ \lambda}{D} \ . \tag{3.6}$$

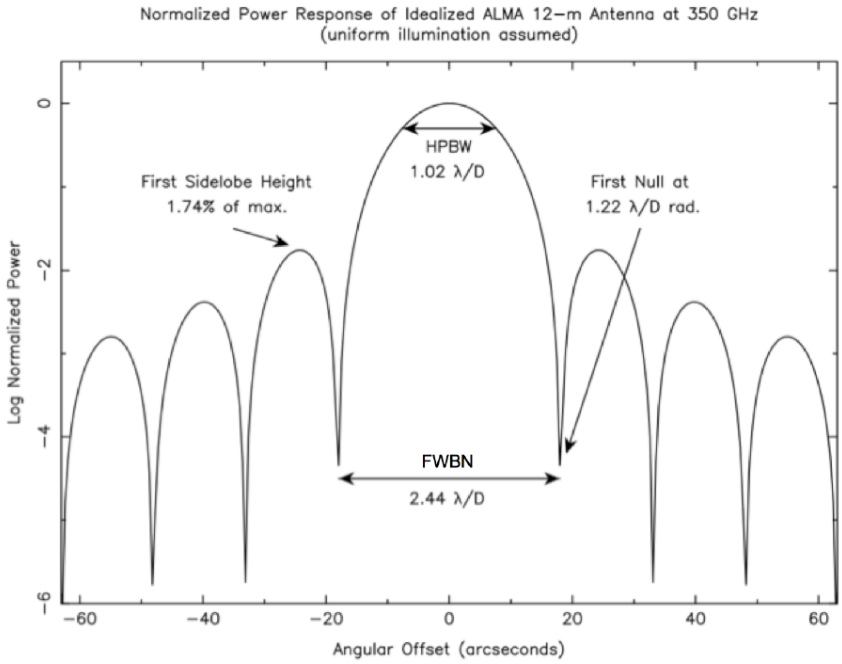

Figure 3.2.: Schematic of the normalized 1-D antenna power response for a 12-m parabolic antenna, uniformly illuminated at 350 GHz. The x-axis is the angular offset in arcseconds, the y-axis is the normalized power in log. Also shown are the Half Power Beam Width (HPBW) of the primary beam, the Full Width at Half Power (FWBN), first null and first sidelobe height. Definitions of these parameters are given in the text. Figure taken from (Cortes et al. 2023).

The angle of the first null, which is the *angular resolution* of the antenna, is given by half this value, which corresponds to the quantity introduced in Eq. 3.1.

Recalling that the EM power induces voltage oscillations, the antenna response can also be expressed in terms of a voltage with $P(\theta) \propto V^2(\theta)$. For an aperture subdivided into N elements, the total voltage is equal to the sum of the voltage contributions from each element $\Delta V_i(t)$:

$$V(t) = \sum_i \Delta V_i(t) \ . \tag{3.7}$$

The average power received is then proportional to the time average of the square of the summed voltage outputs. This, in turn, can be expressed as the sums of identical element pairs (autocorrelation term) and of different element pairs (cross-correlation term):

$$\begin{aligned} \langle P \rangle \propto \left\langle \left[\sum_i V_i \right]^2 \right\rangle &= \sum \sum \langle \Delta V_i \Delta V_k \rangle \\ &= \sum \langle \Delta V_i^2 \rangle + \sum \sum_{i \neq k} \langle \Delta V_i \Delta V_k \rangle \ . \end{aligned} \tag{3.8}$$

The basis of aperture synthesis is that on condition that the contributions from each element i and k arrive in phase at the focus, the elements need not be contiguous. That is, each cross-correlation term $\langle \Delta V_i \Delta V_k \rangle$ (which is the sum of only two of the N elements) can be measured with two physically separated antennas at positions i and k by computing the time averaged product of their output voltages. This is done with a multiplying or correlating receiver, otherwise known as a *correlator*. The condition described here is, however, only valid for emission that is received on-axis. The power response of a pair of antennas also incorporates off-axis emission that results in phase differences of the signals from the two antennas and this must be accounted for.

In order to understand the full power response of a pair of antennas, it is useful to consider an idealized 1-D interferometer consisting of a single pair of antennas as shown in Figure 3.3. The baseline is a vector $\boldsymbol{b}$ that defines both the direction of one antenna towards the other and the distance between them. The distance between the antennas is measured in units of the observing wavelength λ so that, for a distance between antennas L, $\boldsymbol{b} = L/\lambda$. The direction in which the antennas point (towards a particular source on the sky) is given by the unit vector $\boldsymbol{s_0}$, which lies at an angle θ from the meridian. The direction denoted by $\boldsymbol{s_0}$ is indicated with the dashed lines in Figure 3.3 and is also the path along which on-axis wavefronts travel.

The projection of the baseline towards $\boldsymbol{s_0}$ from the perspective of the source is $u = \boldsymbol{b}\cos\theta$, which, with its dependence on $\boldsymbol{b}$, is also measured in units of the observing wavelength. In effect, u is the spatial frequency of the sinusoidal phase delay in the E-W direction. To describe the 1-D position on the sky at a small angle off the axis denoted by α, we define $l = \sin\alpha$, which is the direction cosine with respect to u. That is, l is the angle between the direction vector $\boldsymbol{s_0}$ and u and effectively denotes the E-W direction on the plane of the sky.

An on-axis wavefront (where $\alpha = 0$) that is incident to both antennas must travel an extra distance of $\boldsymbol{b} \cdot \boldsymbol{s_0}$ to arrive at antenna 1. The resulting geometric delay of $\tau_g = \boldsymbol{b} \cdot \boldsymbol{s_0}/c$ (where c is the speed of light) causes a phase difference. This is compensated for by introducing a path delay equivalent to τ_g for antenna 2, thus ensuring that the signals from both antenna arrive at the correlator with the same phase.[11]

Off-axis wavefronts travelling along a path indicated by the solid lines in Figure 3.3 arrive at an angle $\alpha > 0$ from the axis and must travel an extra path $x = u\sin\alpha = ul$ over and above the geometrical delay to reach antenna 1, which cannot be compensated for. Here x is the distance given in the observing wavelength λ. The extra path length results in a phase difference ul that varies sinusoidally as a function of angle. The voltage response of antenna 2 can now be expressed as the product of the voltage response of antenna 1 with this phase factor:

$$V_2 = V_1 e^{2\pi i(ul)} \ . \tag{3.9}$$

To cover the two-dimensional sky-plane and two-dimensional baseline projection onto the sky, we expand the baseline $\boldsymbol{b}$ to a 2-D vector with components $\boldsymbol{b_1}$ and $\boldsymbol{b_2}$ and introduce v orthogonal to u with $u = b_1\cos\theta$ and $v = b_2\cos\varphi$, where φ is the angle from the

[11] See also: https://science.nrao.edu/science/meetings/presentation/jdf.webinar.2.pdf.

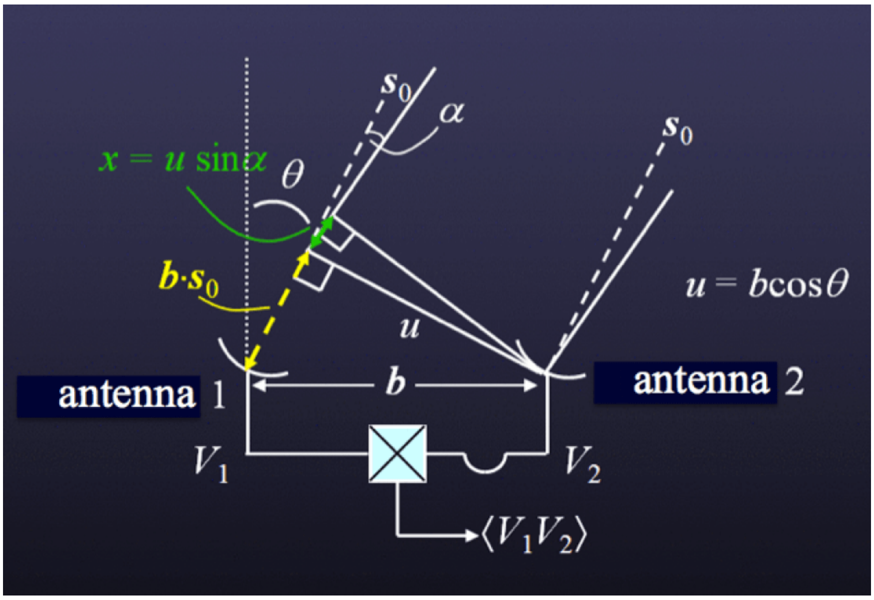

Figure 3.3.: Schematic representation of an ideal 1-D two-antenna interferometer. The baseline is given by the vector $\boldsymbol{b}$. The antennas point to a source on the sky with position given by the unit vector $\boldsymbol{s_0}$. The projected baseline in the direction of $\boldsymbol{s_0}$ from the perspective of the source is given by $u = \boldsymbol{b}\cos\theta$. The extra path traversed by off-axis wavefronts to reach antenna 1 given by $x = u\sin\alpha$. The antennas are connected to a correlator where the voltages from each are multiplied together and then averaged. Figure taken from Cortes et al. (2023).

meridian to the unit vector $\boldsymbol{s_0}$, but orthogonal to θ. Here v is the spatial frequency of the sinusoidal phase delay in the N-S direction. We also introduce β as a direction on the sky orthogonal to α with corresponding direction cosine $m = \sin\beta$ with respect to v, effectively denoting the N-S direction on the plane of the sky. The extra path traversed in this new direction is $y = v\sin\beta = vm$, where y is the distance measured in units of the observing wavelength. The 2-D voltage response of antenna 2 can then be expressed as:

$$V_2 = V_1 e^{2\pi i(ul+vm)} \; , \tag{3.10}$$

where the phase centre (which is also one of the variables in the `tclean` task described in Section3.3.1) is set at the on-axis position $\boldsymbol{s_0}$ with $l = m = 0$.

The quantities u and v span the so-called uv-plane, a plane that lies tangential to the celestial sphere with u and v in the E-W and N-S directions, respectively. The concept of the uv-plane and its relation respective to the lm-image plane is shown in Figure 3.4.

As mentioned above, the correlator multiplies and then time averages the incoming signals from the antenna pair. We can now express this output in terms of the direction

[12] https://science.nrao.edu/facilities/alma/naasc-workshops/nrao-cd-spf1/BasicsInterf.pdf. Accessed 21.11.2024 at 13:12:02.

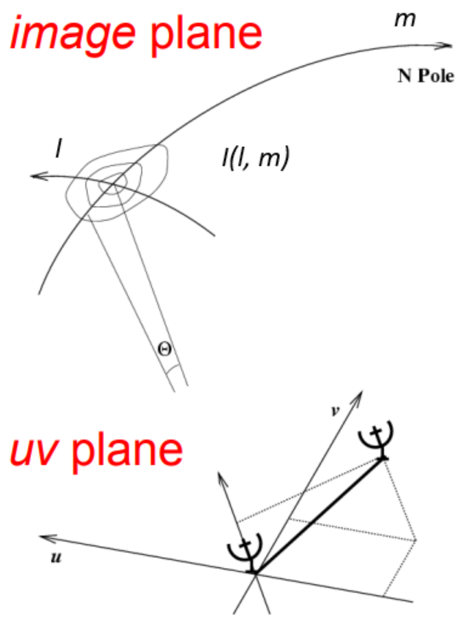

Figure 3.4.: Illustration of the uv-plane (bottom sketch) and image plane (plane on the sky) spanned by l and m (top sketch). The bottom sketch shows two antennas pointing in the direction of a source on the sky and the uv-plane spanned in the E-W direction by u and in the N-S direction, by v. The baseline between the antennas projected onto the uv-plane is indicated by the dotted line joining the two antennas. The top sketch shows the image plane spanned by l and m, which are direction cosines with respect to u and v, respectively. The circular contour lines indicate the source with brightness $I(l, m)$. Figure taken from Introduction to Radio Interferometry.[12]

cosines l and m:

$$\langle V_1 V_2 \rangle = \iint \langle V_1(l,m) V_2(l,m) \rangle \, \mathrm{d}l \, \mathrm{d}m \tag{3.11}$$

and, using the relation given in Eq. 3.10, this can be written as:[13]

$$\langle V_1 V_2 \rangle = \iint \langle V_1(l,m)^2 \rangle e^{2\pi i (ul+vm)} \, \mathrm{d}l \, \mathrm{d}m \ . \tag{3.12}$$

From the relation $V^2 \propto P$ (Eq.3.8) and $P \propto I_\nu$ (Eq. 3.2), the above expression can be written in terms of $I(l, m)$, which is the intensity distribution on the sky:

$$\langle V_1 V_2 \rangle \propto \iint I(l,m) \; e^{2\pi i (ul+vm)} \, \mathrm{d}l \, \mathrm{d}m \ . \tag{3.13}$$

The right-hand side of the above expression is known as the *complex visibility* $\mathcal{V}$ and it is this quantity that the correlator measures. To obtain the full mathematical definition, we recall the definition of P_n given in Eq. 3.3 and note that, for a circular uniformly

[13]For the complete derivation of Eqs. 3.11 and 3.12, see Cortes et al. (2023).

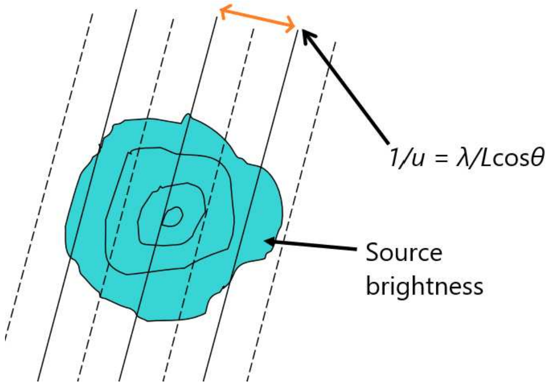

Figure 3.5.: Illustration of the interference or fringe pattern (dashed and dotted lines) produced by a pair of antennas for a source brightness on the sky (blue patch with contours). The spacing between peaks (solid lines) is given by $1/u = 1/\boldsymbol{b}\cos\theta = \lambda/L\cos\theta$. Figure adapted from the presentation on Aperture Synthesis from the Crash Course in Radio Astronomy and Interferometry.[15]

illuminated aperture, P_n is, in fact, the Airy function $\mathcal{A}$.[14] In terms of l and m, we then have $P_N(l, m) = \mathcal{A}(l.m)$. Using Eq. 3.4 to incorporate the normalized response of each individual antenna and by defining the complex visibility as a function of the spatial frequencies u and v, we can then write:

$$\mathcal{V}(u, v) = \iint \mathcal{A}(l, m)\ I(l, m) e^{2\pi i(ul+vm)}\,\mathrm{d}l\,\mathrm{d}m = Ae^{i\varphi}\ . \tag{3.14}$$

Thus, we see that the complex visibility is in fact a complex number with amplitude A, which encapsulates the sky brightness of the source, and phase φ, which holds information on the source's location relative to the phase centre. The complex visibility is in fact also the Fourier transform of the intensity distribution on the sky (sky brightness) in the image plane, $I(l, m)$. Conversely, the intensity distribution on the sky is the inverse Fourier transform of the complex visibility in the uv-plane (also known as the visibility plane). Thus, to recover the sky brightness distribution $I(l, m)$, it is necessary to take the inverse Fourier transform of $\mathcal{V}(u, v)$

$$I(l, m) = \iint \mathcal{A}(l, m)\ \mathcal{V}(u, v) e^{-2\pi i(ul+vm)} du dv \tag{3.15}$$

and then to divide by $\mathcal{A}(l, m)$, which is the last step in the image reduction process (see Section 3.3.1).

As indicated above, the phase delay is sinusoidal and, as such, so is the response of the interferometer. For a pair of antennas at physical distance $\boldsymbol{b} = L/\lambda$, this creates an

[14] For a definition of Airy functions, see e.g., https://personal.math.ubc.ca/~cbm/aands/page_446.htm (accessed 13.02.2025 at 10:45:54) or Schroeder (2000b).

[15] https://science.nrao.edu/science/meetings/presentation/jdf.webinar.2.pdf. Accessed 27.02.2025 at 23:31:17.

interference or fringe pattern with a spacing (in 1-D) of $1/u = 1/\boldsymbol{b}\cos\theta = \lambda/L\cos\theta$. In other words, the Fourier transform of a single visibility measurement is a sinusoid with fringe spacing of $1/u$.[16] This is illustrated in Figure 3.5.

Following from this definition of the fringe spacing, it is clear that the interference of the respective signals from the antenna pair samples the sky brightness only at particular angular scales. These scales are proportional to the wavelength of the sinusoid and are inversely proportional to the projection of the antenna separation on the sky. This means that for the true sky brightness to be recovered, a sufficient coverage of visibilities on the *uv*-plane is necessary. This is aided to an extent by the fact that, as visibilities are samples of a complex-valued function, each sampling gives two visibilities: (u, v) and its complex conjugate $(-u, -v)$. However, certain factors in the array configuration and observation planning can also contribute to good uv-coverage. These include: 1) To sample as many *uv*-points as possible, there should be numerous antennas in the array, each at a different baseline length from the others to avoid redundancy. 2) By extending the total exposure time on the target (also through repeated observations), the projected baselines will change with the apparent motion of the target due to the Earth's rotation (depending on its altitude and azimuth). This will allow the antenna pairs to sample multiple visibilities across the *uv*-plane. 3) By configuring the array to have various baseline lengths, different parts of the *uv*-plane can be sampled.

An example of *uv*-coverage for two configurations of the Submillimeter Array (SMA) is given in Figure 3.6. Examples of the *uv*-coverage for the data in this work are shown in Figure 3.12.

Even if *uv*-coverage is good, the sampling of the sky brightness is nonetheless imperfect, as there will always be certain scales that are not covered in any one particular observation. This results in gaps in the *uv*-coverage, as there is simply no information on the complex visibilities in these regions. In particular, there are limits to both the minimum and maximum scale of an interferometer. The minimum scale is the smallest scale at which detail can still be discerned in the image and is, in essence, the *angular resolution* of the interferometer. It can be approximated as:

$$\theta_{res} \propto \lambda/L_{max} \ , \tag{3.16}$$

which recalls the diffraction limit given in Eq. 3.1. Here λ is the wavelength of observation and L_{max} is the maximum baseline length.

The maximum scale is the largest angular scale at which the interferometer is still sensitive to emission. It is limited by the shortest baseline in the array configuration. The shortest baseline possible is itself limited to a baseline length equal to an antenna diameter, which means that interferometers cannot sample frequencies lower than this length. This results in the so-called *zero-spacing problem*, namely that there is always a lack of visibility coverage in the vicinity of the origin of the *uv*-plane. The *maximum*

[16] See https://science.nrao.edu/science/meetings/presentation/jdf.webinar.2.pdf

[17] https://science.nrao.edu/science/meetings/presentation/jdf.webinar.2.pdf. Accessed 10.07.2025 at 18:56:30.

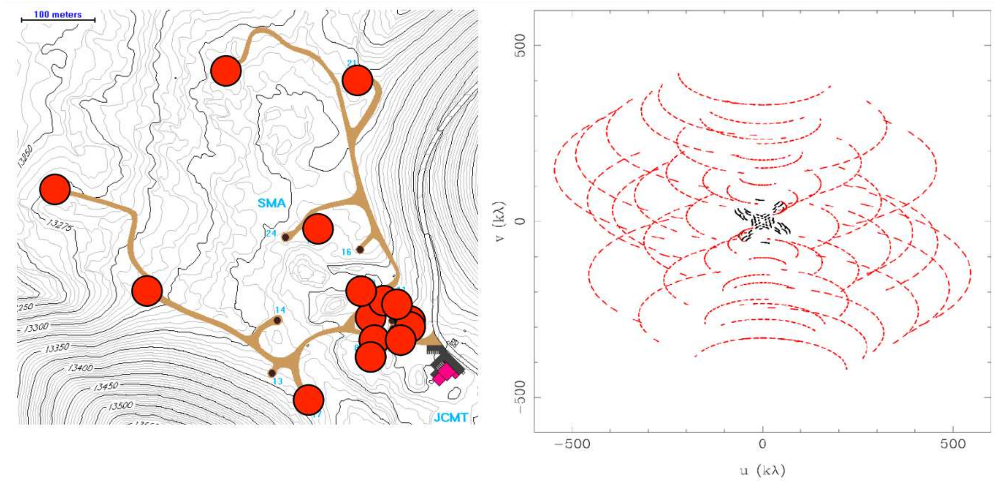

Figure 3.6.: Illustration of two configurations of the SMA. The left panel shows the physical antenna configurations: one spread out with long baselines, the other compact (right bottom corner). The right panel shows the equivalent *uv*-coverage: the red contours are from the long baseline configuration, the black contours are from the compact array configuration. The characteristic curved lines of the contours reflect the Earth's rotation. Figure from Aperture Synthesis, Crash Course in Radio Astronomy and Interferometry.[17]

recoverable scale (MRS) of an interferometer can be approximated as:

$$\theta_{MRS} \propto \lambda / L_{min} \ , \tag{3.17}$$

where λ is again the observing wavelength and L_{min} is the minimum baseline length of the configuration.

The visibilities corresponding to the unobserved scales (i.e., the gaps) are typically assigned a value of $\mathcal{V}(u, v) = 0$. Performing the Fourier transform of the visibilities to get the sky brightness distribution will include these gaps in the Fourier domain and result in an image containing artifacts. This is known as the *dirty image*. Mathematically, the dirty image $I^D(l, m)$ can be expressed as the inverse Fourier transform of the product of the visibilities $\mathcal{V}(u, v)$ with the sampling distribution $B(u, v)$:

$$I^D(l, m) = FT^{-1}\{B(u, v)\mathcal{V}(u, v)\} \ . \tag{3.18}$$

Here $B(u, v)$ characterises the ensemble of discrete (u, v) points that are sampled in the *uv*-plane:

$$B(u, v) = \sum_k (u_k, v_k) \ . \tag{3.19}$$

In turn, the Fourier transform of the sampling distribution is none other than the point spread function (PSF) of the interferometer. This is often referred to as the *synthesized beam* or *dirty beam*:

$$b(l, m) = FT^{-1} B(u, v) \ . \tag{3.20}$$

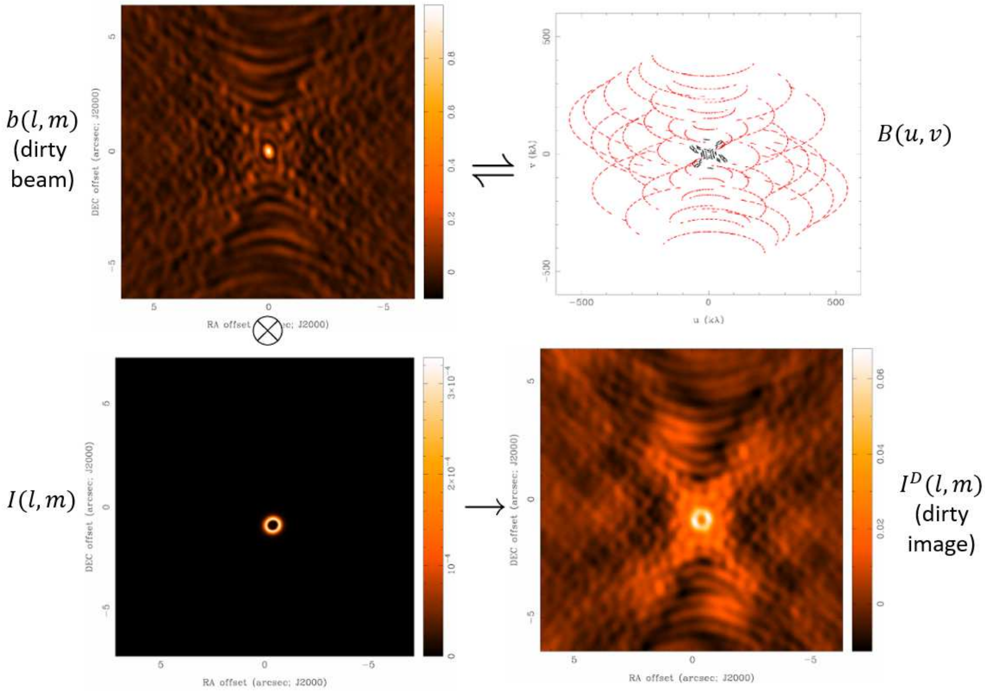

Figure 3.7.: Examples of four key imaging concepts in radio interferometry. *Upper left panel*: the dirty beam $b(l, m)$ (or synthesized beam), which is the equivalent of the PSF. *Upper right panel*: the related sampling distribution $B(u, v)$ of discrete points in the uv-plane (i.e., the uv-coverage). *Bottom left panel*: the true sky distribution $I(l, m)$, which is a model of the emission created in the cleaning process. *Bottom right panel*: the dirty image that results from convolving the true distribution $I(l, m)$ with the dirty beam $b(l, m)$. Figure adapted from (Cortes et al. 2023).[18]

Applying the convolution theorem, which states that the Fourier transform of a convolution of two functions is the product of the Fourier transforms of each of those function, we can use Eq. 3.20 and the definition of $\mathcal{V}(l, m)$ given in Eq. 3.14 to rewrite the expression for the dirty image given in 3.18 as follows:

$$I^D(l, m) = b(l, m) * I(l, m)\mathcal{A}(l, m) \ . \tag{3.21}$$

Thus we see that the dirty image is the convolution of the true sky brightness distribution $I(l, m)$, as modified by the antenna power response $\mathcal{A}(l, m)$, with the PSF $b(l, m)$.

Figure 3.7 shows examples of these concepts. The upper left panel shows the dirty beam $b(l, m)$, while the upper right panel shows the corresponding visibilities sampled in the uv-plane $B(u, v)$, i.e., the uv-coverage for that configuration of baselines. The lower left panel shows a model of the true sky distribution $I(l, m)$, while the lower right panel shows the dirty image $I^D(l, m)$ obtained from the convolution of the true sky brightness $I(l, m)$ with the dirty beam $b(l, m)$.

[18] See also: https://space.asu.cas.cz/~barta/ARC-doc/wilner_synthesis10.pdf

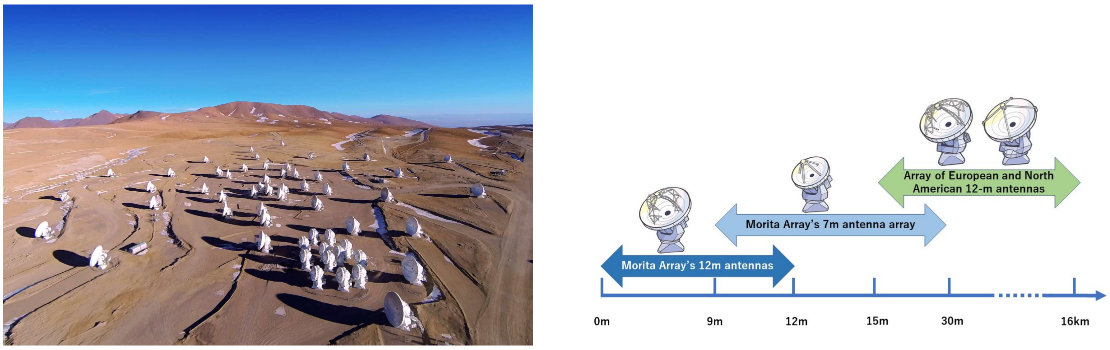

Figure 3.8.: *Left panel*: Aerial view of the ALMA interferometer in its Cycle 7 configuration showing the compact 7-m array in the foreground. *Right panel*: Schematic illustration of the baseline lengths and antenna diameters of the three arrays. Images taken from How ALMA Works Vol. 2.[20]

These concepts are fundamental to the process of imaging, where the main objective is to deconvolve the dirty beam from the dirty image I^D or *clean* the dirty image. This is discussed below in Section 3.3.1.

3.2. ALMA Observations

3.2.1. The ALMA Interferometer

The ALMA interferometer[19] consists of a total of sixty-six antennas. Fifty of these are 12-m antennas that constitute the "12-m array". With its minimum baseline of 15 m and maximum baseline of 16 km, it offers ALMA's more extended baseline configurations that are necessary for high-resolution observations. It is complemented by the Atacama Compact Array (ACA; also known as the Morita Array), which itself has two components: the "7-m array" consisting of twelve closely spaced 7-m antennas and the "Total Power Array" (TP) that consists of four 12-m antennas that function as single-dish telescopes. The minimum and maximum baseline of the 7-m array are 9 m and 30 m, respectively. While these shorter baselines offer the more compact configurations necessary for wide-field imaging of extended structures at larger angular scales, the TP is able to address the zero-spacing problem mentioned above in Section 3.1.2, if necessary (Cortes et al. 2023). Figure 3.8 shows an aerial view of the ALMA interferometer in its Cycle 7 configuration (left panel) and a schematic illustration of the baseline lengths and antenna diameters of the three arrays (right panel).

[19]The main source of reference for this sub-section is Cortes et al. (2023). See also `https://almascience.nrao.edu/about-alma/alma-basics`.
[20]`https://alma-telescope.jp/en/column/almabasics-2e.html`. Accessed 22.11.2024 at 15:06:54.

3.2.2. Observations

The ALMA data for this project were acquired in two sets of observations of the ^{12}CO(2-1) line using the ALMA 12-m array, one with a lower angular resolution to provide adequate sensitivity on larger scales and the other with a much higher angular resolution to provide an accurate probe of the region close to the SMBH.[21] The observations were carried out as part of the project "Probing the Over-massive Black Hole Populations in Compact, High-Dispersion Galaxies" (project code: 2019.1.00036.S) with my co-supervisor, Dr. Dieu D. Nguyen as PI (Principle Investigator). To cover the ^{12}CO(2-1) line's redshifted frequency of 227.01 GHz (with rest frequency $\nu = 230.54$GHz), ALMA band 6 with its frequency range of 211−275 GHz was used for both observations.

The first observation was carried out on 2019-11-25 with a total time on the science target NGC 2513 of 12.09 minutes.[22] The minimum and maximum baselines were 15.10 m and 313.70 m, respectively, providing a maximum recoverable scale (MRS) of 10.90″ and an angular resolution of 1.10″. The resulting beam size was 1.38″ × 1.13″, which corresponds to approximately 466.24 pc × 381.78 pc for the distance of 70.8 Mpc adopted for NGC 2513 in this work. The second observation was carried out on 2021-03-24 to 2021-03-25 with a total time on the science target of 45.33 minutes. In this case, the minimum and maximum baselines were 15.20 m and 1231.40 m, respectively, yielding a correspondingly smaller MRS of 4.30″ and higher angular resolution of 0.31″. The beam size was 0.37″ × 0.33″, corresponding to approximately 125.01 pc × 111.49 pc.

As discussed in Section 1.1, to properly detect the dynamical effect of the SMBH on the tracer and successfully model its mass, the R_{SOI} should ideally be resolved. For the SMBH in NGC 2513, the R_{SOI} was estimated to lie between 0.14″ and 0.20″ (47.19 pc and 69.22 pc; see Table 1.1). For a conservative estimate, I took the median value of the estimates for the ETGs of $R_{SOI} = 0.17''$ (∼57.45 pc). This gives an equivalent diameter for the sphere of influence of 0.34″ (∼114.87 pc), which is just 1″ smaller than the mean beam size of the high-resolution observation of 0.35″ (∼118.25 pc). A more generous estimate using the largest R_{SOI} of 0.20″ (∼67.57 pc) gives an equivalent diameter for the sphere of influence of 0.40″ (∼135.14 pc), which is larger than the mean beam size of the high-resolution observation. Assuming this value, the observational beam's estimated size fulfils the criterion of being no larger than $2 \times R_{SOI}$ (see e.g., Davis 2014; Barth et al. 2016a), which in turn means that the sphere of influence is resolved with the high-resolution observation. The more conservative estimate falls just short, but considering the relaxation of the R_{SOI} criterion discussed in Section 1.1, the angular resolution should still be adequate. On the other hand, the low-resolution observation's MRS of 10.90″ (3.68 kpc) provides the necessary sensitivity out to large enough scales to be able to map the velocity of the gas in these regions.

[21] The references for this sub-section are, unless otherwise stated, the ALMA Science Archive page for this project, which contains all details of the observations (`https://almascience.eso.org/aq/`), as well as the Quality Assurance Report 2 (QA2) and Weblog, both of which are dispatched along with the observational data to the PI and which are also available on the ALMA Science Archive for this project.

[22] Additional time on source was accrued through calibration.

Both observations consisted of four spectral windows (spws), each with a bandwidth of 1875 MHz (equivalent to $\sim$2475 km s^{-1}) and each consisting of 1920 channels, where each channel had a frequency width of 976.56 kHz and equivalent velocity width of approximately 1.2–1.3 km s^{-1}. Three of the spectral windows were targeted at the continuum emission with redshifted central frequencies of 225 GHz, 239.02 GHz and 241.25 GHz, respectively. The remaining spectral window was targeted at the ^{12}CO(2-1) line emission with a redshifted central frequency of 227.01 GHz.

3.2.3. Calibration

The observed data requires calibration of its pointing, amplitude (i.e., flux) and phase, and in the case of spectral line data, how these quantities change with frequency. Changes in phase are often associated with atmospheric effects, especially those due to the presence of water vapour, which can affect the length of the signal path to the antenna.[23] Cortes et al. (2023) remark that calibrators need to be radio sources of known properties such as flux density, frequency spectrum and celestial position. They continue, however, that for calibrations of antenna pointing and bandpass, the main criteria is for the source to have a stable continuum emission. For amplitude and bandpass (i.e. frequency), calibrators should be bright compact sources, while phase calibrators should be moderately bright compact objects. In this respect, quasars are often best suited as calibrators.

In both observations, the quasar J0750+1231 was used to calibrate amplitude, bandpass and atmosphere, while also being used for pointing and water-vapour radiometry (WVR). Phase was calibrated with the quasar J0757+0956, which was also used for WVR. In addition, the high-resolution observation included a *check-source* scan using the quasar J0808+0514. Check-source scans are carried out for observations with angular resolutions $< 0.25''$ (Cortes et al. 2023).

Calibration and initial image reduction are usually carried out at one of the ALMA Regional Centres (ARCs).[24] This initial data processing is most often done using the standard ALMA pipeline, but, if necessary, will be carried out manually by the ALMA staff.[25] Both pipeline and manual processing use the Common Astronomy Software Applications (CASA McMullin et al. 2007), which is the standard software package for all ALMA data processing and which is discussed in detail immediately below in Section 3.3.

For both observations of NGC 2513, calibration and initial imaging were carried out using the ALMA pipeline. The low-resolution data was processed using CASA version 5.6.1-8, while the high-resolution data was processed using CASA version: 6.1.1.15. Both calibrated sets of observational data were obtained upon request from the European ALMA Regional Centre. These data are in the format of the so-called *measurement set* (MS), which is the standard data storage format for ALMA data during analysis.[26]

[23] See also: `https://www.narrabri.atnf.csiro.au/observing/WVRs.html`.

[24] There are three ARCs: the European, North American and East Asian regional centres.

[25] See: `https://almascience.eso.org/processing/science-pipeline`.

[26] Unless otherwise indicated, all references and further information pertaining to the measurement set and ALMA Science Data Model (see next section) can be found under: `https://casadocs.readthedocs.io/en/stable/notebooks/casa-fundamentals.html#MeasurementSet-v2`.

Figure 3.9.: Content of the low-resolution MS of NGC 2513 showing the various directories and tables contained in the MS.

3.3. ALMA data and Imaging in CASA

All information that is recorded during an observation is stored in the archive in the form of an ALMA Science Data Model (ASDM) (Viallefond 2006; Cortes et al. 2023). It is essentially a database arranged in a directory of further directories and tables that contains all the metadata (e.g., headers and observational setup) and binary data (the raw observational data itself) necessary to perform calibration and science. The reduction package CASA uses a storage format known as the aforementioned measurement set, which is similar to the ASDM in that it is also a database containing all relevant interferometric radio data from the observation as well as all corresponding metadata. As with the ASDM, an MS is structured in the form of a directory containing other directories and tables. The visibility data is contained in the parent directory or main table, which itself consists of table files, while other tables containing secondary information can be found in their own corresponding sub-directories. The content of the measurement set of the low-resolution ALMA data for NGC 2513 (with ID `uid___A002_Xea6cf9_X598.ms.split.cal`) is shown in Figure 3.9.

Any actions performed on the data, whether it be viewing the header, calibrating, plotting or imaging the data, are carried out by *tasks* in CASA.[27] Tasks are a Python interface enabling access to CASA's functions. The task commands can be performed directly in the CASA command terminal or written as a Python script, which can then be executed in the CASA terminal. Details of the task as it is carried out in CASA are given in the *logger*. The interaction with CASA through the task interface obviates direct interaction with the tables that constitute a MS, and a detailed knowledge of their structure is not strictly necessary. However, a cursory understanding is useful as some of the arguments required by certain tasks refer to the columns that constitute the tables. The following section discusses the tasks and their most important parameters used for the imaging in this project, which was carried out on the CASA version 6.5.3.28.

3.3.1. Imaging the data in CASA

A summary of the measurement set containing all metadata pertaining to the observation can be obtained with the task `listobs`.[28] The output of `listobs` provides information

[27] See `https://casadocs.readthedocs.io/en/stable/api/casatasks.html`
[28] `https://casadocs.readthedocs.io/en/stable/api/tt/casatasks.information.listobs.html#casatasks.information.listobs`

similar to a FITS header and includes such details as the observer, project ID and length of the observation. It also includes details as to which fields, respectively, are used for the calibrator and science target, the spectral windows and their associated red-shifted and rest frequencies, and the systemic velocity of the science target. An extract of the MS summary provided by the `listobs` task for the low-resolution MS is given in Figure 3.10.

To make it easier to work with the data, it is useful to isolate the science data from the calibrator data. This is done by specifying the field of the science target in the CASA task `split`,[29] which then also saves the split science data in a new MS. As is evident in Figure 3.10, the science target in the low-resolution MS is field 2. I applied the `split` task to both the low- and high-resolution data to include only the science data of each. An extract of the MS summary for the split low-resolution data is given in Figure 3.11, which shows data for only one field, namely that of the science target, NGC 2513. With the smaller amount of data to display, the systemic velocity and rest frequency of the galaxy are also now visible under 'Sources', along with the corresponding spectral window.

Another useful way of examining the data is to plot various parameters of the data with the CASA task `plotms`. Of particular interest is the *uv*-coverage of the visibilities. Figure 3.12 shows the *uv*-coverage for both the low and high-resolution data. Both plots reflect the baseline coverage of the two observations, with minimum and maximum baselines of 15.10 m and 313.70 m for the low-resolution data, respectively, and 15.20 m and 1231.40 m for the high-resolution data, respectively. The elongated curves, especially apparent in the high-resolution *uv*-coverage, reflect the apparent motion of the sky due to the Earth's rotation as described above in Section 3.1.2. The lack of coverage at the centre of the plot due to the zero-spacing problem (also discussed in Section 3.1.2) is easily visible in the low-resolution plot. It is harder to discern in the high-resolution plot due to the much larger spatial scale, but on zooming in, it is still apparent. Both plots show fairly good *uv*-coverage.

For imaging the line emission and creating the velocity cube, it is necessary to first subtract the continuum emission. This is done with the CASA task `uvcontsub`,[30] which facilitates a fitting and subsequent subtraction of the continuum in the *uv*-domain. The idea is to identify the line-free channels across the relevant spectral windows and then to fit the continuum. The fit solution provides a model of the continuum which can then be subtracted from the selected channels in the spectral window. The output is a new measurement set.

I identified the line-free channels by plotting the amplitude against the channels for both sets of data in `plotms` (see Figure 3.13). The excess in emission is apparent in both plots, especially in the low-resolution data. The high-resolution data is noisier with weaker emission in general, but still shows higher emission than average across several channels. Interestingly, there appear to be two peaks in the region of channels 1000 and 1200, respectively.

[29] https://casadocs.readthedocs.io/en/stable/api/tt/casatasks.manipulation.split.html#datacolumn

[30] https://casadocs.readthedocs.io/en/stable/api/tt/casatasks.manipulation.uvcontsub.html#casatasks.manipulation.uvcontsub

```
================================================================================
           MeasurementSet Name:  /data2/shelley/MS_files/uid___A002_Xe407cf_Xb0bf.ms.split.cal      MS Version 2
================================================================================
   Observer: dieunguyen     Project: uid://A001/X13ba/Xb7
Observation: ALMA
Computing scan and subscan properties...
Data records: 668220       Total elapsed time = 1197.02 seconds
   Observed from   25-Nov-2019/07:10:34.0   to   25-Nov-2019/07:30:31.0 (UTC)

   ObservationID = 0         ArrayID = 0
  Date        Timerange (UTC)          Scan  FldId FieldName            nRows     SpwIds   Average Interval(s)    ScanIntent
  25-Nov-2019/07:10:34.0 - 07:15:36.4     3      0 J0750+1231            180600  [25,27,29,31]  [6.05, 6.05, 6.05, 6.05] [CALIBRATE_BANDPASS#ON_SOURCE,CALIBRATE_FLUX#ON_SOURCE,CALIBRATE_WVR#ON_SOURCE]
              07:15:51.8 - 07:16:22.1     4      1 J0757+0956             18060  [25,27,29,31]  [6.05, 6.05, 6.05, 6.05] [CALIBRATE_PHASE#ON_SOURCE,CALIBRATE_WVR#ON_SOURCE]
              07:16:51.8 - 07:24:57.4     6      2 NGC2513               288960  [25,27,29,31]  [6.05, 6.05, 6.05, 6.05] [OBSERVE_TARGET#ON_SOURCE]
              07:25:08.3 - 07:25:38.5     7      1 J0757+0956             18060  [25,27,29,31]  [6.05, 6.05, 6.05, 6.05] [CALIBRATE_PHASE#ON_SOURCE,CALIBRATE_WVR#ON_SOURCE]
              07:25:46.9 - 07:29:48.9     8      2 NGC2513               144480  [25,27,29,31]  [6.05, 6.05, 6.05, 6.05] [OBSERVE_TARGET#ON_SOURCE]
              07:30:00.8 - 07:30:31.0     9      1 J0757+0956             18060  [25,27,29,31]  [6.05, 6.05, 6.05, 6.05] [CALIBRATE_PHASE#ON_SOURCE,CALIBRATE_WVR#ON_SOURCE]
           (nRows = Total number of rows per scan)
Fields: 3
  ID   Code Name                RA               Decl           Epoch   SrcId      nRows
  0    none J0750+1231          07:50:52.045735 +12.31.04.82822 ICRS    0          180600
  1    none J0757+0956          07:57:06.642952 +09.56.34.85228 ICRS    1           54180
  2    none NGC2513             08:02:24.656880 +09.24.48.93120 ICRS    2          433440
Spectral Windows:  (4 unique spectral windows and 1 unique polarization setups)
  SpwID  Name                                 #Chans   Frame   Ch0(MHz)  ChanWid(kHz)  TotBW(kHz) CtrFreq(MHz) BBC Num  Corrs
  25     X1505322256#ALMA_RB_06#BB_1#SW-01#FULL_RES   1920   TOPO  227953.603      -976.562   1875000.0 227016.5914        1  XX  YY
  27     X1505322256#ALMA_RB_06#BB_2#SW-01#FULL_RES   1920   TOPO  225953.603      -976.562   1875000.0 225016.5914        2  XX  YY
  29     X1505322256#ALMA_RB_06#BB_3#SW-01#FULL_RES   1920   TOPO  238121.580       976.562   1875000.0 239058.5914        3  XX  YY
  31     X1505322256#ALMA_RB_06#BB_4#SW-01#FULL_RES   1920   TOPO  240329.236       976.562   1875000.0 241266.2476        4  XX  YY
Sources: 193
  ID   Name                SpwId RestFreq(MHz)  SysVel(km/s)
  0    J0750+1231          0     -              -
  0    J0750+1231          1     -              -
  0    J0750+1231          2     -              -
  0    J0750+1231          3     -              -
  0    J0750+1231          4     -              -
```

Figure 3.10.: Extract of the MS summary provided by the `listobs` task for the low-resolution MS of NGC 2513 including information on: the file name, PI (observer), project ID (uid), dates and time of the observation, fields, calibrator names, science target (NGC 2513), spectral window IDs, and information regarding the spectral windows such as channel width and central frequency.

```
================================================================================
           MeasurementSet Name:  /data2/shelley/MS_files/Xe407_split.ms      MS Version 2
================================================================================
   Observer: dieunguyen     Project: uid://A001/X13ba/Xb7
Observation: ALMA
Computing scan and subscan properties...
Data records: 433440       Total elapsed time = 777.024 seconds
   Observed from   25-Nov-2019/07:16:51.8   to   25-Nov-2019/07:29:48.9 (UTC)

   ObservationID = 0         ArrayID = 0
  Date        Timerange (UTC)          Scan  FldId FieldName             nRows     SpwIds   Average Interval(s)    ScanIntent
  25-Nov-2019/07:16:51.8 - 07:24:57.4     6      0 NGC2513                 288960  [25,27,29,31]  [6.05, 6.05, 6.05, 6.05] [OBSERVE_TARGET#ON_SOURCE]
              07:25:46.9 - 07:29:48.9     8      0 NGC2513                 144480  [25,27,29,31]  [6.05, 6.05, 6.05, 6.05] [OBSERVE_TARGET#ON_SOURCE]
           (nRows = Total number of rows per scan)
Fields: 1
  ID   Code Name                RA               Decl           Epoch   SrcId      nRows
  0    none NGC2513             08:02:24.656880 +09.24.48.93120 ICRS    0          433440
Spectral Windows:  (4 unique spectral windows and 1 unique polarization setups)
  SpwID  Name                                   #Chans   Frame   Ch0(MHz)  ChanWid(kHz)  TotBW(kHz) CtrFreq(MHz) BBC Num  Corrs
  25     X1505322256#ALMA_RB_06#BB_1#SW-01#FULL_RES   1920    TOPO  227953.603      -976.562   1875000.0 227016.5914        1  XX  YY
  27     X1505322256#ALMA_RB_06#BB_2#SW-01#FULL_RES   1920    TOPO  225953.603      -976.562   1875000.0 225016.5914        2  XX  YY
  29     X1505322256#ALMA_RB_06#BB_3#SW-01#FULL_RES   1920    TOPO  238121.580       976.562   1875000.0 239058.5914        3  XX  YY
  31     X1505322256#ALMA_RB_06#BB_4#SW-01#FULL_RES   1920    TOPO  240329.236       976.562   1875000.0 241266.2476        4  XX  YY
Sources: 21
  ID   Name                SpwId RestFreq(MHz)  SysVel(km/s)
  0    NGC2513             4     -              -
  0    NGC2513             25    230538         4628.77818505
  0    NGC2513             26    230538         4628.77818505
  0    NGC2513             27    228500         4628.77818505
  0    NGC2513             28    228500         4628.77818505
  0    NGC2513             29    242750         4628.77818505
  0    NGC2513             30    242750         4628.77818505
  0    NGC2513             31    245000         4628.77818505
```

Figure 3.11.: Extract of the MS summary provided by the `listobs` task for the split low-resolution MS file. The information included is as for Figure 3.10, but only for the field of the science target NGC 2513, which has been reassigned the field ID of 0. Also visible now under sources are the rest frequency and systemic velocity of the science target.

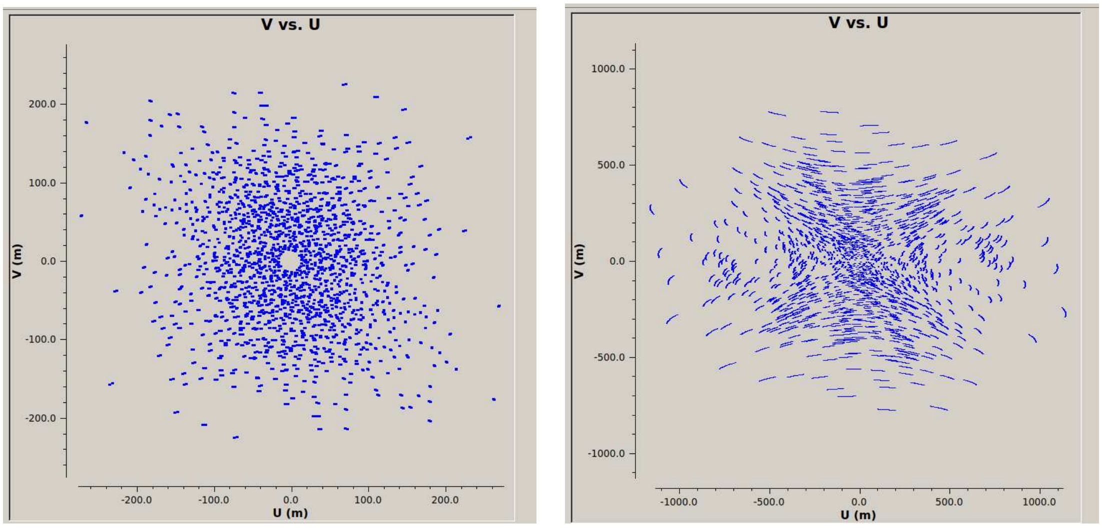

Figure 3.12.: Coverage of the visibilities in the *uv*-plane for the low-resolution (*left panel*) and high-resolution (*right panel*) data. Both plots reflect the different baseline coverage of the two observations. The effect of the zero-spacing problem, where visibilities at a scale less than the minimum baseline are not covered, can clearly be seen in the low-resolution plot. Both plots show reasonably good coverage of the *uv*-plane.

In order to clarify which peak is the line, I made use of the `coloraxis` option offered by `plotms`, which allows one to colour particular parameters such as the field, antenna

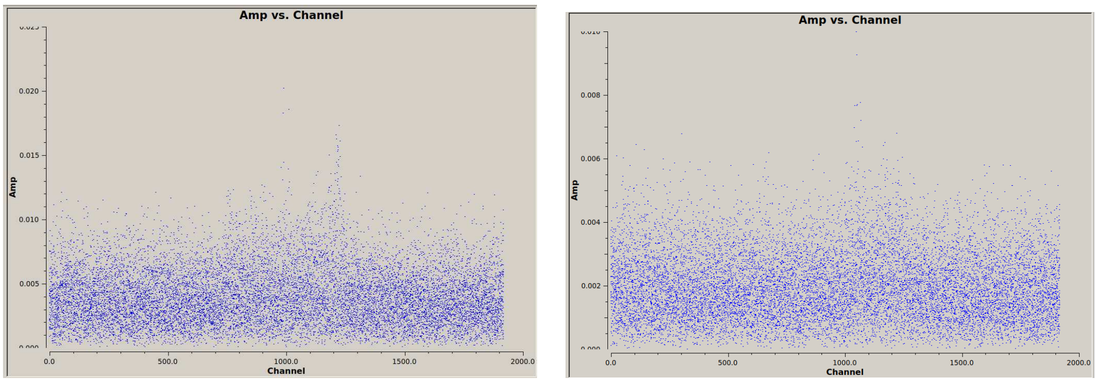

Figure 3.13.: Plots showing the amplitude of the emission against channel for both the low-resolution (*left panel*) and high-resolution (*right panel*) data. Two emission peaks are visible around the channels 1000 and 1200, respectively.

and spectral window of the emission. Choosing to colour the spectral window (spw), I plotted the amplitude against both the spectral windows and frequency for both sets of data. These are shown in, respectively, the top and bottom panels of Figure 3.14. From the MS summary of both sets of data (see e.g., Figure 3.11 for the low-resolution MS), I confirmed that the spectral window containing the red-shifted ^{12}CO(2-1) line is spw 25, which is coloured black in the plots. The plots showing amplitude versus frequency confirm this, with the black-coloured emission lying clearly at the red-shifted frequency of 227.01 GHz of the line .

Interestingly, both sets of plots show that there is a high peak of emission in spw 29 at a frequency centred at 239 GHz. Using the same `coloraxis` option to plot the amplitude against the channels, the plots (shown in Figure 3.15) indicate that the orange-coloured line at 239 GHz seems to dominate, especially in the high-resolution data. This is the peak at channel 1000 identified above in Figure 3.13. The black-coloured peak of the ^{12}CO(2-1), centred around channel 1200, is nonetheless visible in both sets of data.

Based on these plots I identified the line free channels as: 0−1049 and 1351−1919 for the low-resolution data, and 1−1101 and 1299−1919 for the high-resolution data. For the polynomial fit, I chose a fitting order of zero. This fits a constant amplitude continuum, which should suffice for these data and avoid over-fitting.[31]

After subtracting the continuum from each MS file separately, I combined the two sets of continuum-subtracted data into a single visibility file using the CASA task `concat`.[32] As the name suggests, this tool allows the concatenation of several visibility data sets into a single MS. If not stipulated, the task will combine the two MS files with equal weighting. One can, however, adjust the weighting of each MS file depending on the aim of the project. The combination of low- and high-resolution data is always a compromise between sensitivity and higher angular resolution for the reasons described in Section 3.1.2. The

[31] See also: https://casaguides.nrao.edu/index.php/First_Look_at_Line_Imaging_CASA_6.5.4.
[32] https://casadocs.readthedocs.io/en/v6.5.4/api/tt/casatasks.manipulation.concat.html#casatasks.manipulation.concat

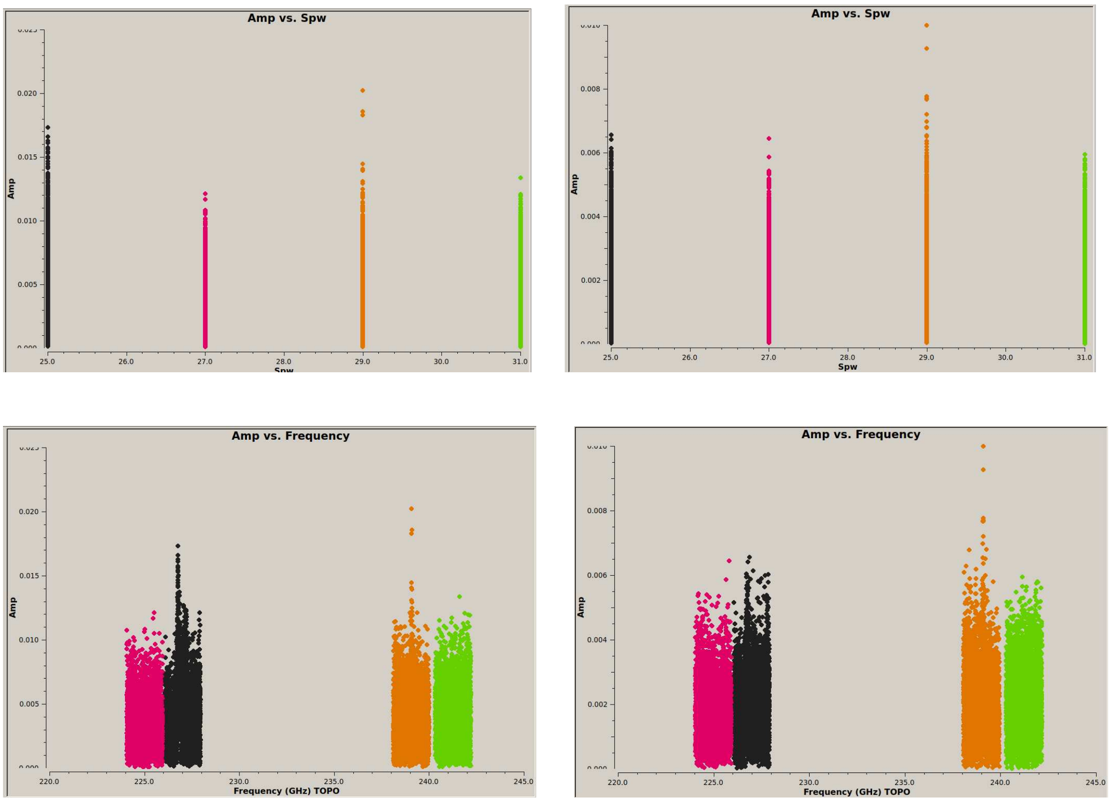

Figure 3.14.: *Top panels*: amplitude versus spectral window (spw) for the low- and high-resolution MS files in the *left* and *right* panels, respectively. The plots are colour-coded for the spectral windows of each MS file. The spectral window containing the line is spw 25. *Bottom panels*: amplitude versus frequency for the low- and high-resolution MS files in the *left* and *right* panels, respectively. The central frequency of the emission in spw 25 (black) corresponds to the red-shifted frequency of 227.01 GHz.

former allows for broader coverage, while the latter provides a better angular resolution of the inner regions around the SMBH.

In this case, it was initially decided to combine the data with a weighting ratio of 70 % to 30 % for the low- and high-resolution MS files, respectively. This weighting would favour sensitivity to the more extended regions around the SMBH and allow for better coverage of the velocity of the gas here. However, on further discussion with my co-supervisor, a more apt choice might have been the exact opposite, with 70 % weighting for the high-resolution data and 30 % for the low-resolution data. Such a weighting would favour better spatially resolving the R_{SOI} of the SMBH, which is the main objective for making an accurate mass measurement. Thus, for further work on this project, I intend to re-image the data with the suggested weighting, taking care, however, that sufficient sensitivity be maintained.

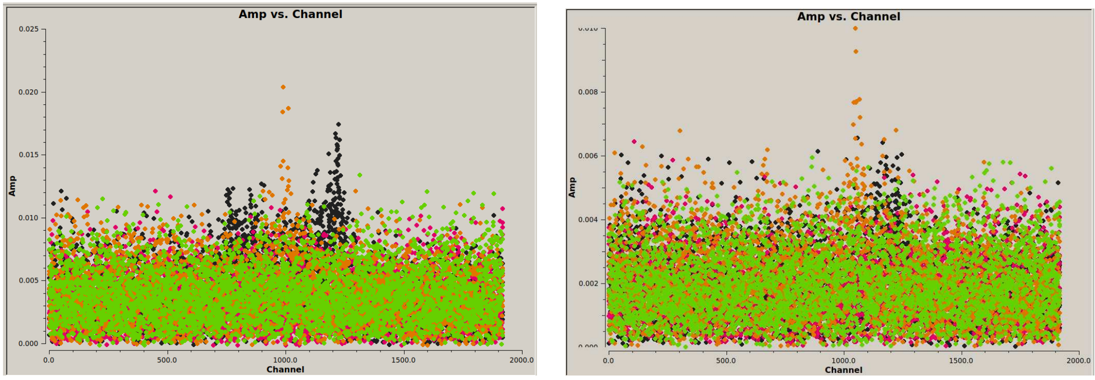

Figure 3.15.: The same as for Figure 3.13, but with spectral windows (spws) colour-coded. The black-coloured spw containing the red-shifted ^{12}CO(2-1) line at 227.01 GHz clearly shows an emission peak centred around channel 1200. Also visible is a peak of the orange-coloured spw 29 centred at 239 GHz just above channel 1000.

I carried out the imaging on the combined data set with the CASA task `tclean`,[33] which is based on the original CLEAN algorithm developed by Högbom (1974). The algorithm is designed to perform a deconvolution of the dirty image from the dirty beam (PSF), creating a model image of the true sky distribution. Very basically, it scans for the brightest pixel (or *clean component*) within a prescribed region of the dirty image and then subtracts a fraction of the dirty beam at that location. The region is usually selected to exclude the main emission of the target object by constructing a mask over it. The subtracted image is the residual image upon which the next iteration of scan and subtraction is performed. This process is continued iteratively until the sidelobes in the image (equivalent to the maximum residual in the region) are reduced to below a threshold equivalent to the noise level. The subtracted fluxes, i.e., the *clean components*, are placed in a model image of the sky. After the stipulated threshold is reached, the final residual image is added to the model of clean components that have in turn been convolved with a *clean beam*. The latter is usually a 2-D Gaussian approximation of the dirty beam's main lobe. The result is the *clean image*, either in the form of a 2-D continuum image or a 3-D data cube.[34] The `tclean` task carries out the image reconstruction by means of an outer loop of *major* cycles and an inner loop of *minor* cycles (CASA Team et al. 2022). The major cycles start out in the data (visibility) domain where an initial model of the visibilities is predicted. The model is subtracted from the data to generate residual visibilities. These are then gridded and inverse-Fourier transformed from the data domain to the image domain, where the minor cycles iteratively perform deconvolution of the residual image with a user-designated clean algorithm.[35] This is illustrated in Figure 3.16.

[33] `https://casadocs.readthedocs.io/en/stable/api/tt/casatasks.imaging.tclean.html`

[34] For further in-depth explanation of the CLEAN algorithm and synthesis imaging in general, see the documentation on `tclean` and, in particular, `https://casadocs.readthedocs.io/en/stable/notebooks/synthesis_imaging.html`

[35] See: `https://www.alma-allegro.nl/wp-content/uploads/2023/10/DRT2023_imaging.pdf`; `https://casadocs.readthedocs.io/en/stable/notebooks/synthesis_imaging.html`

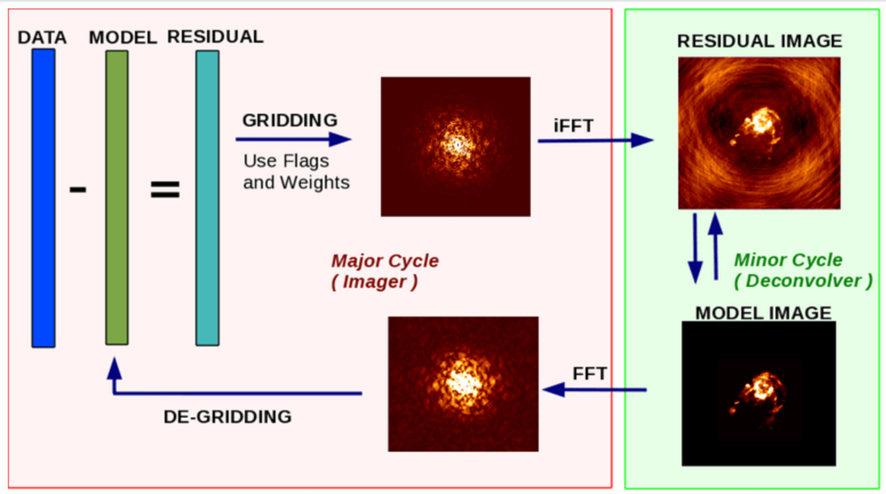

Figure 3.16.: Schematic of the major and minor cleaning cycles in `tclean`. Figure from (CASA Team et al. 2022)

I imaged both the continuum and the $^{12}CO(2\text{-}1)$ line with `tclean`. The continuum imaging is not directly relevant to the focus of this thesis, but will be used for further investigations regarding the distribution and spatial extent of the CO gas. The script for the continuum imaging of the combined data is shown in Appendix A. The script for the $^{12}CO(2\text{-}1)$ imaging with `tclean` is given in Figure 3.17 and, in the following, I discuss the most important input parameters thereof. A summary description of each input parameter is provided in Appendix A.

For the line imaging, I selected the MS to be imaged and specified the spectral windows containing the $^{12}CO(2\text{-}1)$ line, which, in the combined MS, were the two spectral windows spw 25 and spw 57. I also specified the phase centre as given in the `listobs` which, as mentioned in Section 3.1, is necessary for defining the position of the source.

As opposed to the continuum imaging, which produces a 2-D image with spatial coordinate axes, line imaging results in a cube with two spatial coordinate axes (RA and DEC) and a third axis for the velocity of the gas. This is generated in `tclean` by selecting the spectral mode *cube*, which produces an image with spatial dimensions and with multiple channels of a predefined velocity width along the velocity axis. The parameters that define the spectral coordinate system for the cube mode are thus the number of channels, the velocity of the first channel and the width of the channels in terms of velocity.[36] These are called with the inputs `nchan`, `start` and `width`, respectively.

These three parameters depend, amongst others, on the systemic velocity of the system, which can be found in the MS summary called with `listobs`. The aim is to choose a number of channels that will cover all the velocities of the gas from the outer to the inner regions of the gas disc. The systemic velocity for this system given in the MS summary

[36]These parameters can also be given in terms of frequency or as integers.

```
CASA <1>: tclean(vis='combined_ngc2513_1101.ms/',
     ...: imagename='ngc2513_comb37_spw2557_100chn_009arcs_230GHz_nit10000_thresh08_4200kms_050224',
     ...: field='0',
     ...: spw='25,57',
     ...: phasecenter="ICRS 08:02:24.656880 +09.24.48.93120",
     ...: specmode='cube',
     ...: nchan=100,
     ...: start='4200km/s',
     ...: width='10km/s',
     ...: outframe='LSRK',
     ...: restfreq='230.5380GHz',
     ...: cell='0.09arcsec',
     ...: deconvolver='hogbom',
     ...: gridder='standard',
     ...: imsize=[128, 128],
     ...: weighting='briggs',
     ...: robust=0.5,
     ...: restoringbeam='common',
     ...: niter=10000,
     ...: threshold='0.8mJy',
     ...: interactive=True)
```

```
CASA <1>: tclean(vis='combined_ngc2513_1101.ms/',
     ...: imagename='ngc2513_comb37_spw2557_140chn_009arcs_230GHz_nit10000_thresh08_3900kms_090224',
     ...: field='0',
     ...: spw='25,57',
     ...: phasecenter="ICRS 08:02:24.656880 +09.24.48.93120",
     ...: specmode='cube',
     ...: nchan=140,
     ...: start='3900km/s',
     ...: width='10km/s',
     ...: outframe='LSRK',
     ...: restfreq='230.5380GHz',
     ...: cell='0.09arcsec',
     ...: deconvolver='hogbom',
     ...: gridder='standard',
     ...: imsize=[128, 128],
     ...: weighting='briggs',
     ...: robust=0.5,
     ...: restoringbeam='common',
     ...: niter=10000,
     ...: threshold='0.8mJy',
     ...: interactive=True)
```

Figure 3.17.: Scripts for the `tclean` imaging of the ^{12}CO(2-1) line with 100 and 140 channels, respectively.

is 4628.78 km s^{-1}, which is quite quite high.[37] With such a high systemic velocity, a sufficiently large number of channels should be chosen to properly cover the velocity range. Too few channels may result in poor velocity coverage in the outer regions. An example of such lack of coverage is shown in Figure A.2 (see Appendix A), which shows a velocity map for a cube of 50 channels from one of the initial attempts at imaging. Based on this experience, I opted to produce two cubes, one with 100 channels and one with 140 channels, both of which produced reasonable results. The moment maps produced from the 140-channel cube, which was chosen for the modelling, is shown and discussed in Section 3.4.

The start velocity of the first channel should be chosen so that the systemic velocity lies at or within the central channel. For the 100-channel cube, this was 4200 km s^{-1}, and for the 140-channel cube, this was 3900 km s^{-1}. The channel width determines the resolution

[37] This compares quite well with the systemic velocity of 4705 km s^{-1} from SIMBAD, which cites Yu et al. (2022). See `https://simbad.u-strasbg.fr/simbad/sim-basic?Ident=ngc+2513&submit=SIMBAD+search`.

of each channel with regard to line features (i.e., the spectral resolution). A width that is too large will lack resolution, while a width that is too narrow will not capture extended features. Based on comparison to the literature (e.g., Davis et al. 2017a; Kabasares et al. 2022), I chose a channel width of $10\,\mathrm{km\,s^{-1}}$ with respect to the rest frequency of the ^{12}CO(2-1) line given as 230.538 GHz in the MS summary.[38] The parameters **start** and **width** need to be interpreted in a specific reference frame, which is called in **tclean** with **outframe**. The default and most commonly used is the Kinematic Local Standard of Rest or LSRK, which is based on the average velocity of local stars[39] and which automatically corrects for Doppler effects that might arise in longer exposure times.

The cell size, stipulated with the argument **cell**, determines the pixel size in angular units, here arcseconds. A good approximation is to aim for about 3−5 pixels (i.e., cells) across the minor axis of the clean beam.[40] The size of the clean beam is calculated in the **tclean** process itself and for a small dataset, a preliminary quick imaging of the cube can be carried out to get the beam size. However, as most datasets are large and this can be very time-consuming, an alternative is to examine the *uv*-coverage across the longest baseline in wavelengths and obtain an estimate for the number of pixels across the beam with the following relation:

$$\text{cellsize} \approx \frac{206265}{b_l} \cdot \frac{1}{n_{cell}} \,. \tag{3.22}$$

Here b_l is the longest baseline in wavelengths and n_{cell} is the number of cells across the beam. The former can be obtained in **plotms** by setting the x-axis to *UVwave*, which gives the baseline lengths in units of the observing wavelength, as is shown in Figure 3.18. With this method and approximating for 3 pixels across the beam, I obtained an initial value of $\sim 0.08''$, which, after preliminary imaging, I increased to a final value of $0.09''$ to obtain a more suitable resolution.

The gridder serves to resample the weighted visibilities onto a uniform grid in the *uv*-plane. I selected the the **standard** gridder, which is the option most suitable for single pointings, such as the two observations here. This gridder operates in the image plane and corrects for direction-dependent effects.[41] An example of gridding on the *uv*-plane with direction dependent corrections is given in Figure 3.19.

A significant factor affecting the resolution, sidelobe suppression and rms (root-mean-square) noise levels is the weighting of the visibilities (CASA Team et al. 2022). In this process, each visibility is multiplied by a weight before being gridded in the *uv*-plane

---

[38] The rest frequency is also available on the QA2 report or on Splatalogue.

[39] See, e.g., `https://help.almascience.org/kb/articles/what-are-the-frequency-reference-frames-in-casa`

[40] This is the suggestion from a previous tutorial available at the time of imaging, which was for CASA 6.4; a currently available reference for this suggestion is, e.g., `https://casaguides.nrao.edu/index.php/VLA_CASA_Imaging-CASA6.5.2` The updated CASA tutorial for ALMA users suggests aiming for ~ 5 pixels across the beam; see `https://casaguides.nrao.edu/index.php/First_Look_at_Imaging_CASA_6.5.4`.

[41] See e.g., `https://www.alma-allegro.nl/wp-content/uploads/2023/10/DRT2023_imaging.pdf`

[42] `https://casadocs.readthedocs.io/en/stable/notebooks/synthesis_imaging.html#Deconvolution-Algorithms`. Accessed 20.11.2024 at 17:44:00.

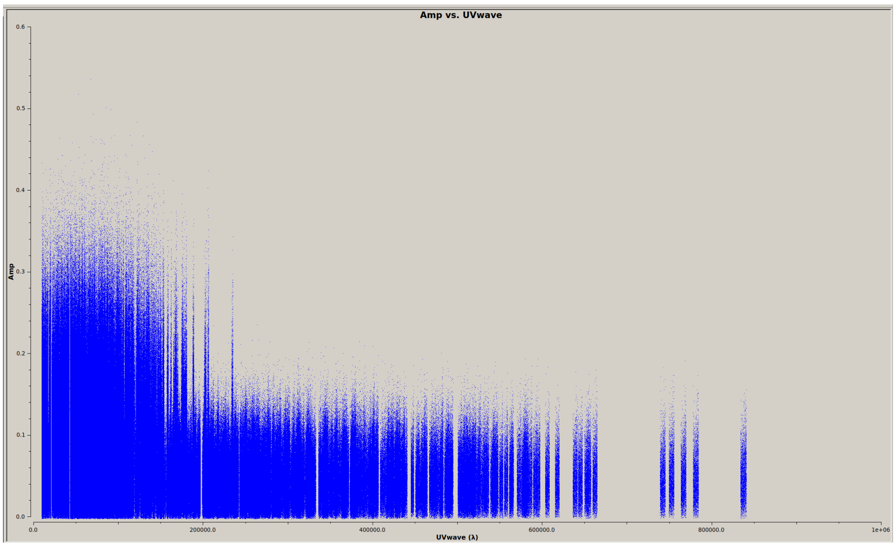

Figure 3.18.: Plot of the amplitude versus the baseline length in units of the observing wavelength for the combined MS. The longest baseline extends out to just below $\sim 840\,000$ wavelengths.

as discussed directly above. The main weighting schemes for interferometric data are: natural, uniform and Briggs (robust). Natural weighting uses the weights of the original data (i.e., the MS) and grids these directly into the *uv*-cells. Natural weighting results in a higher imaging weight for higher *uv* density, producing an image with low noise, but with lower resolution. Uniform weighting first grids the data to cells in the *uv*-plane and then re-weights the *uv*-cells so that they have uniform imaging weights. This produces an image with higher resolution, but at the cost of increased noise. The Briggs weighting scheme offers a compromise between natural and uniform weighting, using a `robust` parameter that takes a value between -2.0 and 2.0 to produce a weighting that is closer to uniform and natural, respectively.[43] A comparison is shown in Figure 3.20. For my imaging, I chose a Briggs weighting scheme with a robust parameter of 0.5.

For the minor cycle deconvolution algorithm, I selected `hogbom`, which is an adapted version of the original Hogbom Clean (Högbom 1974; CASA Team et al. 2022). This deconvolver is suitable for compact structures (i.e., that do not have complex spatial structures) and, because it subtracts the full PSF from the image, is comparatively accurate.[45]

The final stage of cleaning is to 'restore' the model image. This simply means that the model needs to be convolved with a clean beam (as described above in the discussion on

[43]See `https://casadocs.readthedocs.io/en/stable/notebooks/synthesis_imaging.html`; `https://www.alma-allegro.nl/wp-content/uploads/2023/10/DRT2023_imaging.pdf`

[44]`https://casadocs.readthedocs.io/en/stable/notebooks/synthesis_imaging.html`. Accessed 03.03.2025 at 12:57:05.

[45]For a reference to this and for a comparison of the other deconvolver options, see `https://casadocs.readthedocs.io/en/stable/notebooks/synthesis_imaging.html#Deconvolution-Algorithms`.

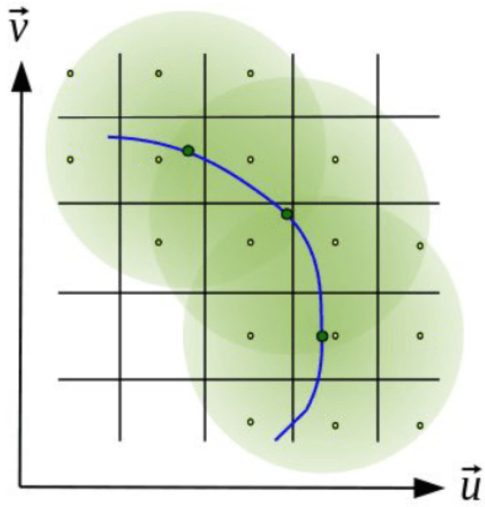

Figure 3.19.: Schematic illustration of gridding on the uv-plane. The visibilities (yellow dots) are weighted and gridded. In the case of the standard gridder, direction dependent corrections are applied, as shown by the green-filled dots. Figure from CASA docs.[42]

the CLEAN algorithm) and the residual image added to the result. The type of beam used is designated by the `restoringbeam` argument. Here I selected the option *common*, which automatically estimates a common beam size or shape appropriate for all image planes.[46]

The `tclean` task can be run in an interactive mode, a non-interactive mode or a combination of both. For basic scripting, this mainly effects three input parameters: the threshold value at which cleaning should stop(`threshold`), the maximum number of iterations or cleaning cycles to be carried out (`niter`), and the mask (`mask`).

In the non-interactive mode, these parameters need to be defined in the `tclean` script prior to running `tclean`. After running the script, both major and minor cycles of the cleaning process run automatically with no further interaction on the part of the user. The advantage of this mode is that it is time-saving and reproducible.

In the interactive mode, which I describe in detail below, the user has manual control over these parameters through a GUI interface, with the advantage that these values can be adjusted during the cleaning process. Moreover, the cleaning process can be stopped at the discretion of the user. To gain experience and a more thorough understanding of the details of cleaning, I decided to use the interactive mode for imaging my data. In the following I discuss the aforementioned three parameters and how they are applied in both modes.

The `threshold` parameter designates a value, usually selected to be several times that of the rms noise in units of Jy. When the maximum residual in the image or selected

[46]Further options for the restoring beam can be found in the `tclean` documentation: `https://casadocs.readthedocs.io/en/stable/api/tt/casatasks.imaging.tclean.html`.

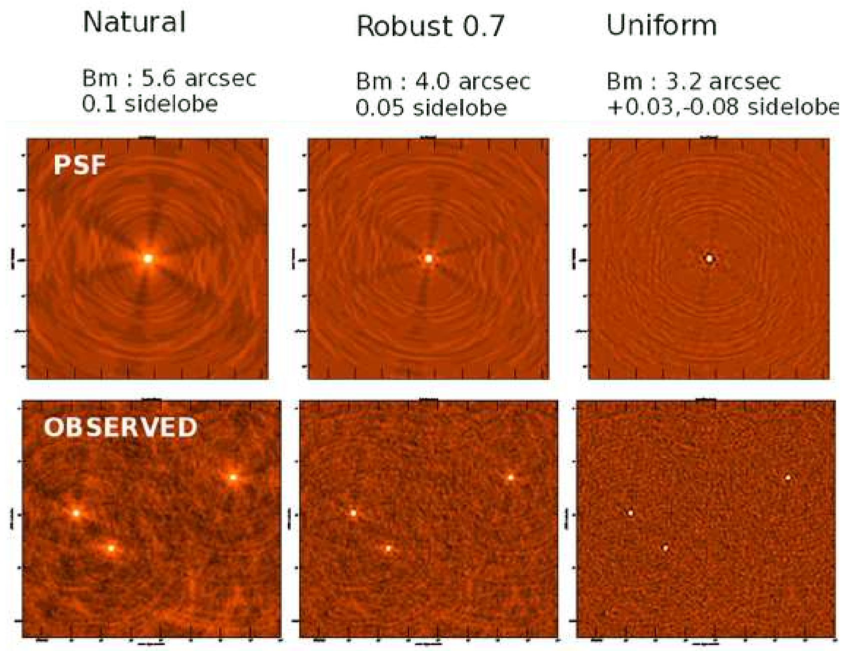

Figure 3.20.: Weighting schemes used in imaging interferometric data. From *left* to *right*: natural, Briggs with robust = 0.7, and uniform. The *upper panels* show the PSF (or dirty beam) for all three schemes, while the *lower panels* show the observed data. Both the beam size (Bm) in arcseconds and sidelobe levels for each weighting are given. Figure adapted from the CASA documentation on imaging.[44]

region outside the mask reaches below the threshold value, cleaning will stop. For non-interactive cleaning, a value in the region of 3 to 5 times greater than the rms noise level ($3\sigma - 5\sigma$) is usually recommended to prevent `tclean` from adding bright noise features to the deconvolved image.[47] For interactive cleaning, a lower value can be selected as a precaution against over-cleaning, although a threshold value is not strictly necessary in this mode.

For my imaging, I decided to select a threshold value based on the rms noise. To estimate the rms noise, I executed a preliminary single run of `tclean` and then opened the cleaned image in the viewer display panel available in CASA, which can be called with the task `imview`.[48] The viewer offers the possibility of selecting a region in the image and attaining statistics on that region, which include, among others, the rms noise. By choosing a region in several of the line-free channels and then averaging the values, I estimated an rms of approximately 0.8 mJy, which is the value I used for the threshold.

In addition to, or as an alternative to the threshold, one can stipulate the maximum number of cleaning cycles with the argument `niter`. This is useful as a backup to the

[47] See e.g., `https://casaguides.nrao.edu/index.php/First_Look_at_Imaging_CASA_6.5.4` and `https://casaguides.nrao.edu/index.php/VLA_CASA_Imaging-CASA6.5.2`

[48] The `imview` task is not currently being maintained and is being replaced by CARTA (Cube Analysis and Rendering Tool for Astronomy).

threshold value that can prevent `tclean` from running too long should the threshold be too low (which would lead to over-cleaning). As such, any random, large number should suffice and for this imaging, I chose a maximum number of 10 000 cleaning cycles. It should be noted that if `niter` is set to zero, no cleaning will take place and a dirty image will be returned.

As mentioned above in the discussion on the CLEAN algorithm, the cleaning process is usually not applied to the entire image, but to a region that excludes the emission of the target. This can be implemented by applying a mask to the target emission area. The mask can be applied either interactively or automatically, depending on the mode selected in which to run `tclean`. In the non-interactive mode, the mask is set in the script with the parameter argument `mask`. A simple mask can be constructed in the form of a box with size stipulated in pixels,[49] while more sophisticated masks can be constructed with CASA's Automasking.[50]

Interactive construction of a mask takes place in the GUI mentioned above. On entering the script in interactive mode, `tclean` initiates the gridder and creates the PSF. Once this is complete, the preliminary cube, consisting of the number of channels stipulated in the script, appears in the GUI viewer panel. Using a shape tool one can manually draw a mask around the emitting region. The viewer gives the option of drawing a mask on each channel separately or of creating a single mask for all channels. As my cube had a high number of channels, I chose the latter option for my imaging, as the former can be time-expensive and does not necessarily improve the result. The viewer also gives the user the choice of running further cleaning iterations automatically (in non-interactive mode), of stopping the cleaning process or of starting a major cycle.

After running a major cycle, the residual emission is displayed in the viewer. The aim is to have the residual emission inside the masked region look similar to the noise outside the region (see Figure A.3). One can either manually end the cleaning task based on this visual comparison or run the cleaning process until the threshold or number of maximum iterations is reached. In my case, I ran two major cleaning cycles.

In both the interactive and non-interactive modes, the cleaning process can be followed on the logger, which details, among others, the running of major and minor cycles, the clean flux recovered thus far, the minimum and maximum residuals and whether the threshold has been reached. This is especially important in the interactive mode, where the user needs to actively watch residual levels to judge over-cleaning and pay attention to the cleaning limits.

After cleaning is finished, several image products are created. These are: the final cleaned image cube, the area that `tclean` cleaned (the clean mask), the model of the clean components used by `tclean`, an image of the primary beam (normalized antenna response or sensitivity map for each antenna), the PSF (synthesized or dirty beam), the residuals and the sum of weights per plane. With the exception of the latter, examples of these products for a single channel at a velocity of 4910 $\mathrm{km\,s^{-1}}$ from the imaging of the 140-channel imaging of the $^{12}CO(2\text{-}1)$ line emission for NGC 2513 are shown in

[49] See `https://casaguides.nrao.edu/index.php/First_Look_at_Imaging_CASA_6.5.4`
[50] See `https://casaguides.nrao.edu/index.php/Automasking_Guide_CASA_6.6.1`

Figure 3.21. It should be noted that this velocity value includes the systemic velocity of $4628.78\,\mathrm{km\,s^{-1}}$, which must be subtracted to get the true velocity of the molecular gas. The clean beam (white ellipse) as calculated in `tclean` of $0.44'' \times 0.37''$,[51] is shown in the cleaned image and PSF.

Although cleaning is essentially complete at this stage, the image cube that has been produced has still not been corrected for the primary beam attenuation of each antenna in the array. As explained in Section 3.1.2, to recover the true intensity distribution of the sky $I(l, m)$, one has to divide by the primary beam (see Eq. 3.15), where the primary beam is the normalized antenna response $P_N(l, m)$ (i.e., the field of view, equivalent to $\mathcal{A}(l, m)$, the Airy function) of each individual antenna in the array. One can either do this directly in `tclean` during the cleaning process by setting the parameter for primary beam correction (`pbcor`) as `True`, or, as I chose to do, by using the task `impbcor`.[52] This task takes the cleaned image and primary beam response products listed above as parameters and produces an output file of the image cube that has been primary beam corrected.

The final product of the image reduction is the primary-beam corrected 3-D data cube (RA, DEC, velocity), which can be fed directly into the `KinMS` modelling software. Figure 3.22 presents a selection of the observed channel maps obtained from imaging the 140-channel cube. Each channel map depicts the molecular gas at a particular velocity, ranging here from $4300\,\mathrm{km\,s^{-1}}$ to $4950\,\mathrm{km\,s^{-1}}$. As indicated above, these values include the systemic velocity of $4628.78\,\mathrm{km\,s^{-1}}$ and this must be subtracted to get the true velocity of the molecular gas. Also depicted in the maps is the synthesized beam (white ellipse) of $0.44'' \times 0.37''$ (equivalent to approximately $149 \times 125\,\mathrm{pc}$).[53]

The following section describes the visualization and preliminary analysis of the data cube.

3.4. Visualisation of the data

To visualize the data and extract information more readily for analysis, one can create 1-D and 2-D projections of the 3-D cube by collapsing one or more of the axes. This is also useful for carrying out a preliminary examination of the imaged cube which, if needs be, can be re-imaged before modelling is carried out. These data products include the integrated spectrum, moment maps and position-velocity diagram (PVD). The following first briefly explains each of these and then discusses how they were produced to analyse the data cube of NGC 2513.

3.4.1. Spectral-line profiles, moment maps and PVDs

Collapsing both positional axes (RA and DEC) and slicing along the velocity axis creates a 1-D spectral-line profile, i.e. integrated spectrum, of the gas where the flux density is

[51] The beam size is available in the cube data header, called with the CASA task `imhead`.

[52] `https://casadocs.readthedocs.io/en/stable/api/tt/casatasks.imaging.impbcor.html`

[53] Recalling that the clean beam is an approximation of the dirty beam's main lobe, the final synthesized beam after cleaning has the same dimensions as the clean beam.

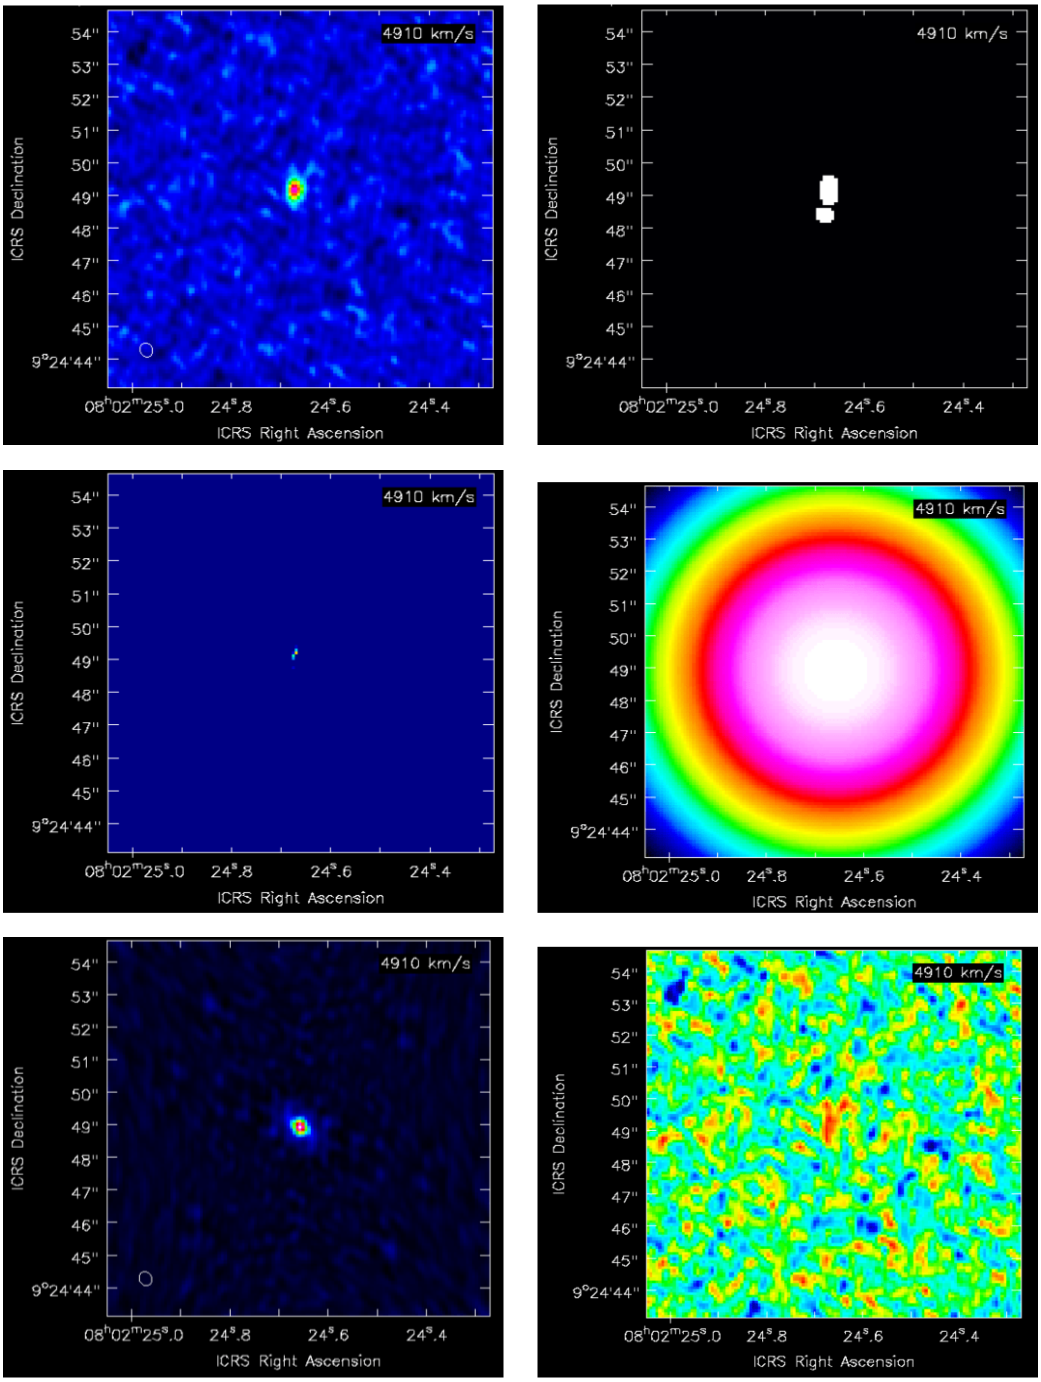

Figure 3.21.: Data products of the 140-channel imaging of the ^{12}CO(2-1) line emission of a single channel at a velocity of 4910 km s^{-1} produced in `tclean` for NGC 2513. Top to bottom and left to right: the cleaned image cube, clean mask, model of the clean components, primary beam, PSF and residuals. The clean beam of $0.44'' \times 0.37''$ is shown as a white ellipse in the cleaned imaged and PSF.

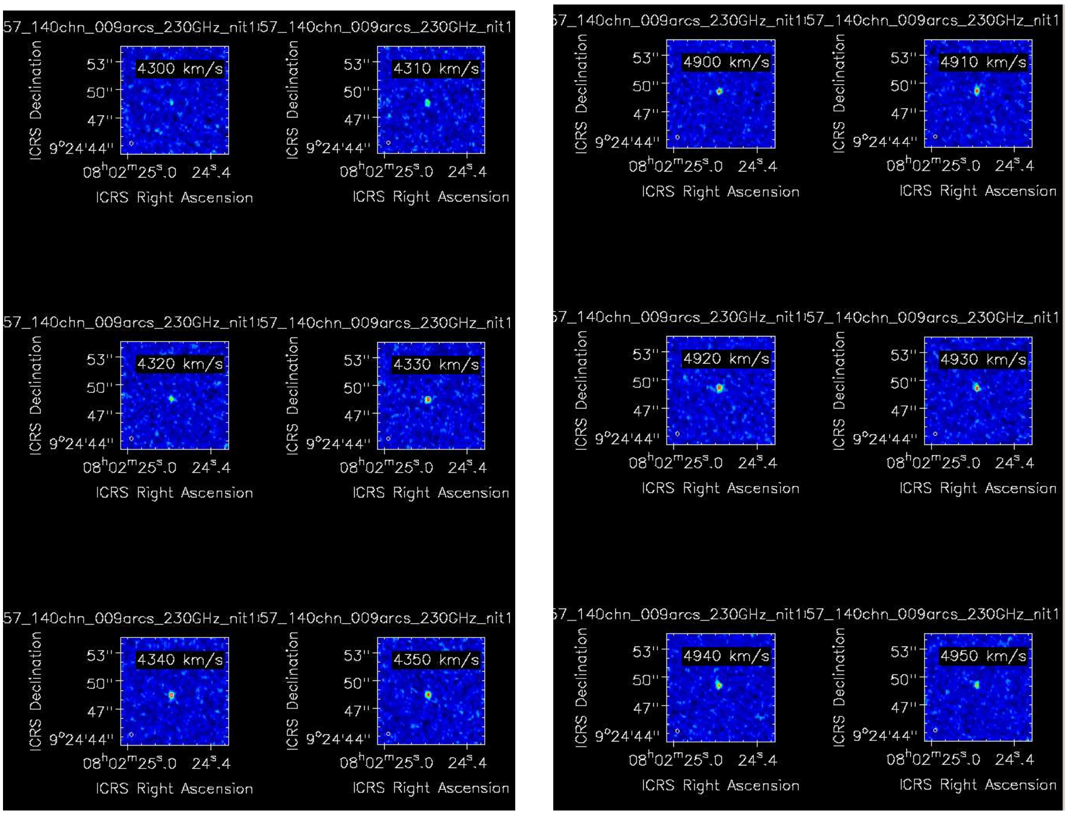

Figure 3.22.: Sample of channel maps produced from the `tclean` imaging of the ^{12}CO(2-1) line emission in NGC 2513 for velocities ranging from $4300\,\mathrm{km\,s^{-1}}$ to $4950\,\mathrm{km\,s^{-1}}$. The synthesized beam of $0.44'' \times 0.37''$ is shown as a white ellipse in the bottom left corner of each map. The maps were produced directly in CASA.

plotted directly against the velocity of the gas. The profile provides a useful visual aid in indicating the type of motion of the gas. In particular, a 'double-horn' profile indicates a rotating disc (see Figure 3.26).

The so-called moment maps yield information in 2-D spatial dimensions of the intensity (zeroth moment), mean velocity (first moment) and velocity dispersion (second moment).[54] The moment maps are produced by computing the following integrals along the velocity axis:[55]

[54]It is also possible to compute higher order moments, but they are not relevant for the analysis here.
[55]See e.g., https://www.alma-allegro.nl/wp-content/uploads/2018/10/Allegro_CASATrainingDay2018_analysisTools.pdf

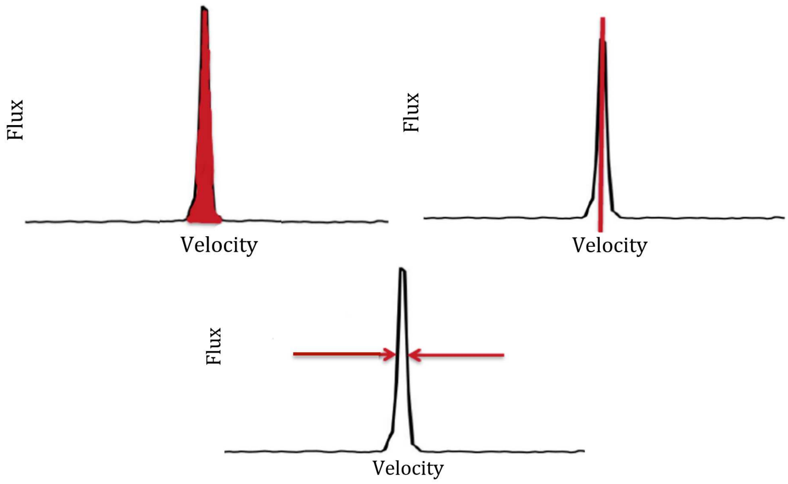

Figure 3.23.: Visual illustration of the concept of the zeroth moment (*top left*), first moment (*top right*) and second moment (*bottom*) for flux plotted against velocity. Figure adapted from lecture on Analysis Tools at the Allegro CASA Training Day 2018.[56]

$$M_0 = \int I_v \, \mathrm{d}v \tag{3.23}$$

$$M_1 = \langle V \rangle = \frac{\int I_v v \, \mathrm{d}v}{\int I_v \, \mathrm{d}v} \tag{3.24}$$

$$M_2 = \sqrt{\langle V^2 \rangle} = \sqrt{\frac{\int I_v (v - \langle V \rangle)^2 \, \mathrm{d}v}{\int I_v \, \mathrm{d}v}} \, , \tag{3.25}$$

where the moments computed are the integrated intensity, intensity-weighted velocity, and intensity-weighted velocity dispersion, respectively. The principle is portrayed visually in Figure 3.23, which shows each moment and its relation to the velocity distribution of the flux.

The position velocity diagram (PVD) is one of the most useful tools in analysing the radial velocity profile of material and in providing a visual aid to the presence of a SMBH. A PVD shows the radial velocity of the material along the line of sight (LOS) as a function of the distance from the centre of rotation. It is constructed by taking a slit (usually a few pixels in width) along the kinematic major axis of the system, i.e., the axis along which the observed rotation is at its highest[57] (see e.g., Kormendy 1982), as shown by the black dotted line in Figure 3.24. The slit contains, at one and the same time, information on the position of the gas relative to the centre of rotation and the radial velocity distribution of the gas at that position anywhere along the slit. Rotating the slit

[56] https://www.alma-allegro.nl/wp-content/uploads/2018/10/Allegro_CASATrainingDay2018_analysisTools.pdf. Accessed 04.12.2024 at 17:28:05.

[57] The kinematic minor axis is the line along which observed the rotation is zero.

around the kinematic major axis gives the position of the gas relative to the system's centre on the x-axis, while the radial velocity distribution lies on the y-axis. The concept is illustrated in Figures 3.24 and 3.25.

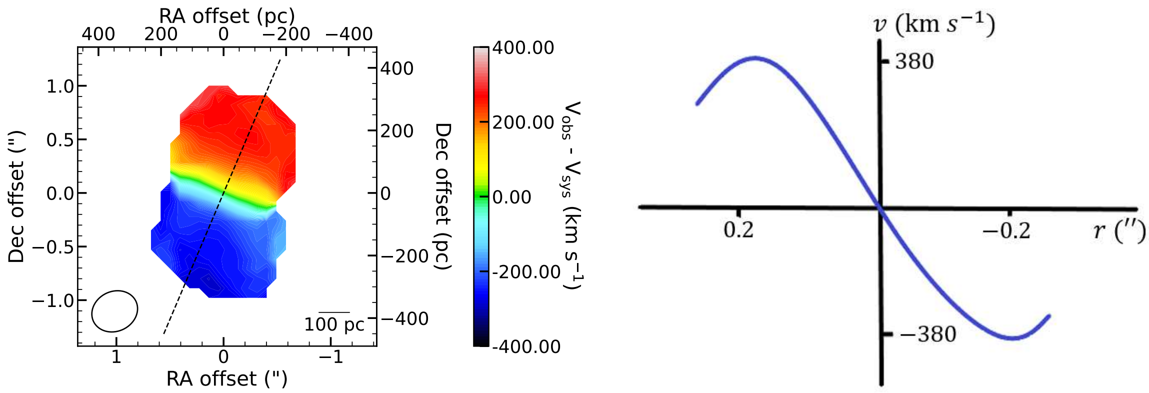

Figure 3.24.: Figure illustrating the concept of PVD construction. The *left panel* shows the first moment (velocity) map of the molecular gas disc in NGC 2513 constructed as part of this thesis. The kinematic major axis along which rotation is at its highest is shown by the dotted line (here equivalent to the position angle of the disc). The *right panel* is a sketch illustrating the corresponding PVD with the position of the gas relative to the centre of the galaxy lying on the x-axis and the velocity of the gas on the y-axis. The peak velocity of $\sim 380\,\mathrm{km\,s^{-1}}$ lies at $\sim 0.2''$.

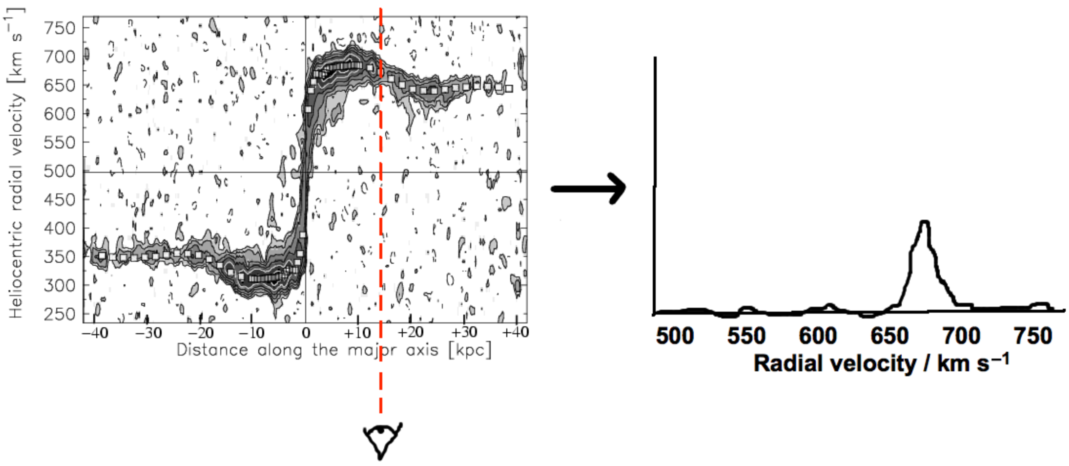

Figure 3.25.: Figure illustrating the concept of a PVD. The *left panel* shows a PVD with the red dashed line indicating a particular line-of-sight (LOS). The corresponding velocity distribution along the LOS is shown in the *right panel*. Figure from Astronomy StackExchange.[58]

[58] https://astronomy.stackexchange.com/questions/13810/create-position-velocity-diagram-from-a-velocity-field. Accessed 02.12.2024 at 11:42:09.

3.4.2. Integrated spectrum, moment maps and PVD of molecular gas in NGC 2513

The integrated spectrum, moment maps and PVD of the ^{12}CO(2-1) gas disc in NGC 2513 were produced using the masked moment technique (e.g., Davis et al. 2017a) with software adapted from the `pymakeplots` package.[59] This technique makes use of applying a three-dimensional mask (along the spatial and velocity axes) to a copy of the clean data cube which has been smoothed in both the spatial and spectral dimensions by convolving the cube with a boxcar function. The effect of the smoothing is to reduce the noise in the image. The mask is created by clipping the smoothed data cube to a multiple of the noise level of the smoothed cube. The mask is then applied to the original, unsmoothed cube in the process of creating the integrated spectrum, moment maps and PVD.

In the case of this data, the mask was produced by selecting all pixels above a flux threshold seven times that of the rms (i.e., 7σ). This is notably higher than the default 3σ, but lower values produced maps that contained an excess of noise. This might be due to under-cleaning the image or incorrect masking in the cleaning process and this will be revisited in later work on the data.

The integrated ^{12}CO(2-1) spectrum of NGC 2513 extracted from the imaged data cube within a $3'' \times 3''$ (approximately 1014 pc $\times$ 1014 pc) region around the galaxy centre is shown in Figure 3.26. The grey dashed line shows the integrated spectrum recovered from collapsing the data cube in the complete image area, i.e., in the image with no mask applied. The solid black line shows the integrated spectrum recovered from integrating within the emission mask produced in the process of creating the maps. The 'double-horn' line profile displayed here is characteristic of a rotating disc.

[59] `https://github.com/TimothyADavis/pymakeplots`

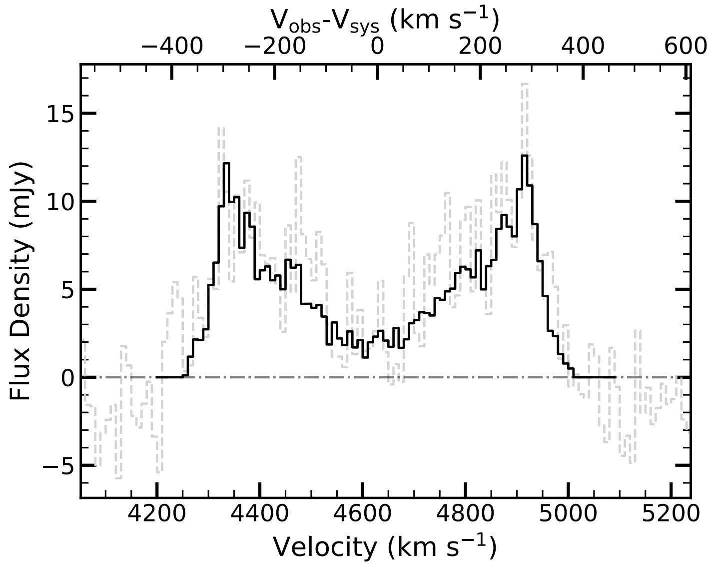

Figure 3.26.: Integrated $^{12}CO(2\text{-}1)$ spectrum extracted from the observed data cube of NGC 2513 within a $3'' \times 3''$ ($\sim$ 1014 pc $\times \sim$ 1014 pc) region around the galaxy centre. The black solid line is the integrated spectrum recovered from the area within the emission mask, the grey dashed line is the integrated spectrum from the entire image. The double-horn profile is typical of a rotating disc.

The moment maps and PVD for the $^{12}CO(2\text{-}1)$ gas disc, obtained from the cleaned data cube, are shown in Figure 3.27. The synthesized beam of $0.44'' \times 0.37''$ (approximately 149 pc $\times$ 125 pc) is indicated by the black ellipse in the lower left corner of the moment maps. The integrated intensity map (top left panel) shows the presence of the molecular gas disc. It is fairly elongated covering an extent of approximately 700 pc $\times$ 400 pc. The outer regions ($\sim$ 100 pc) are fairly diffuse, but there is a clear central concentration of $^{12}CO(2\text{-}1)$ emission. As indicated by the velocity map (top right panel), the gas is regularly rotating with clear Keplerian motion around a kinematic axis as suggested by the integrated spectrum shown in Figure 3.26. The negative-valued, blue coloured part of the map indicates gas travelling towards us, the positive-valued, red coloured part indicates gas moving away from us. The peak velocity values at $\pm$ 400 km s^{-1} are relatively high (compare e.g., Davis et al. 2017a; Kabasares et al. 2024) and certainly of the order to suggest an over-massive object at the centre of the gas disc.

The gas velocity dispersion (bottom left panel) is unusually high at the centre of the gas disc with peak values of $\sim$ 200 km s^{-1} (cf. e.g., Davis et al. 2017b,a; North et al. 2019), although a similarly high gas velocity dispersion can be seen, for example, in the second moment maps of the ETGs NGC 4261 (Boizelle et al. 2021; Ruffa et al. 2023) ($>$ 275 km s^{-1}) and NGC 1332 (Barth et al. 2016a) ($\sim$ 120 km s^{-1}). The map of NGC 2513 is, however, largely flat with constant values of $\sim$ 55 km s^{-1} covering much of the disc. The sharp increase towards the centre is thus unlikely to be intrinsic, but rather due

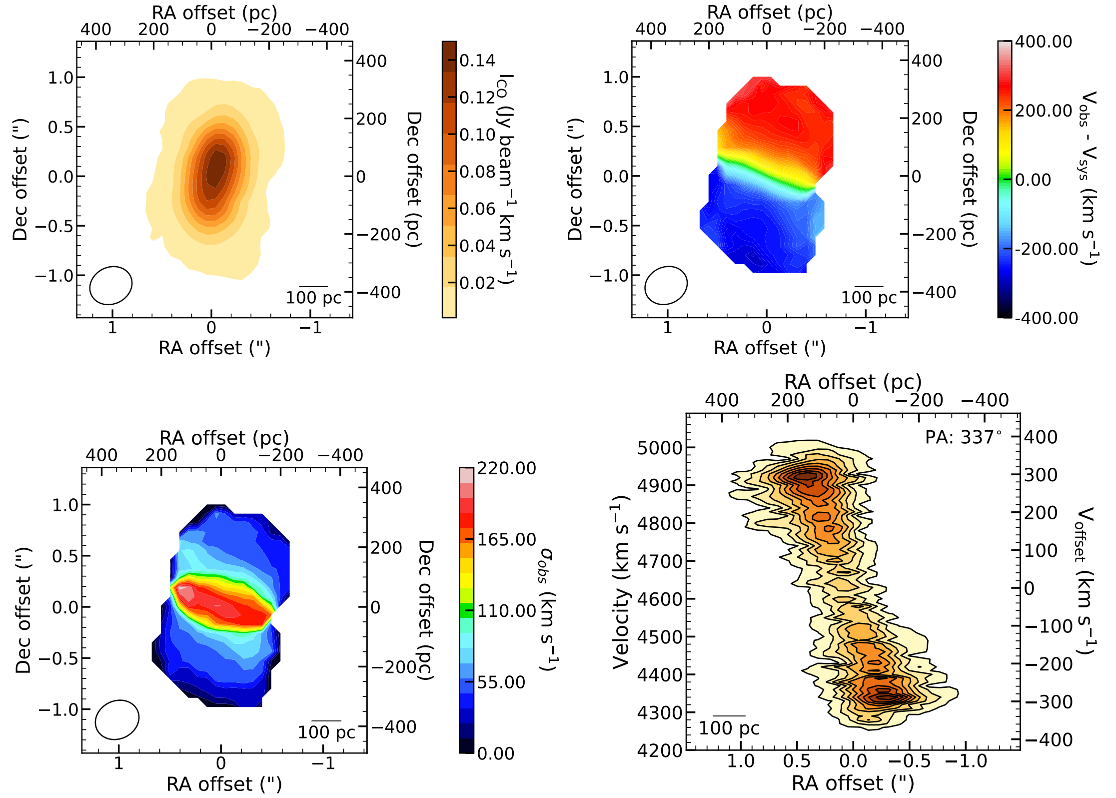

Figure 3.27.: Moment maps and PVD of the ^{12}CO(2-1) emission in NGC 2513 created using the masked moment technique. *Top left*: Integrated intensity map (zeroth moment) of the ^{12}CO(2-1) emission indicating a compact, elliptical shaped disc of emission. *Top right*: Velocity map (first moment) of the ^{12}CO(2-1) emission indicating clear Keplerian rotation of the molecular disc with relatively high peak velocities. *Bottom left*: Velocity dispersion map (second moment) of the ^{12}CO(2-1) emission. *Bottom right*: PVD of the ^{12}CO(2-1) emission extracted along the kinematic major axis. The synthesized beam of $0.44'' \times 0.37''$ is shown as a black ellipse at the bottom left of each moment map.

to the effect of beam smearing that results from velocity gradients across the finite-sized beam coupled to the line-of-sight (LOS) line emission of a disc with inclination $i > 50°$ (cf. Nguyen et al. 2020).[60] Ruffa et al. (2023) also report beam smearing as a likely cause for the high velocity dispersion in NGC 4261, although Boizelle et al. (2021) indicate that turbulence plays a significant role for the gas in this galaxy.

The PVD, (bottom right panel), was constructed along the kinematic major axis of the gas disc, taken along a position angle (PA) of 337° (dashed line shown in the left panel of

[60] For a beam of finite size, the average velocity of the gas as observed in each pixel is weighted by the intensity of the gas emission in that pixel. The effect is more common for discs with inclinations above 50°. A more detailed explanation of beam smearing can be found in Keppler, M. et al. (2019); Boehler, Y. et al. (2021).

Figure 3.24) with a width of 5 pixels. The PA value used here was an initial estimate attained by looping through a series of possible values and is reasonably close to the best-fit PA of 343.28 $^{+2.44}_{-2.36}$ degrees given in the `KinMS` modelling discussed in Chapter 4. The characteristic peaks in the gas velocity within the inner $\sim 0.5''$ (reaching between $350\,\mathrm{km\,s^{-1}}$ and $400\,\mathrm{km\,s^{-1}}$) are again a good indication of a compact massive object at the galaxy's centre. These characteristic peaks in the PVD becomes more evident when the emission is overlaid with models resulting from different SMBH masses produced in `KinMS`, which is discussed in Chapter 4.

4. Dynamical Modelling with KinMS

The dynamical modelling of the data cube was carried out using the KINematic Molecular Simulation (`KinMS`[1]) mm-wave observation simulation tool of Davis et al. (2012), which was designed specifically to model molecular gas.[2] The tool allows for the creation of simulated interferometric data cubes based on input guesses of the gas distribution and kinematics, while taking into account observational effects such as beam smearing, velocity binning, gas velocity dispersion and disc thickness (Davis et al. 2012). The simulated cubes can then be fitted to the observed data. Fitting may be carried out via a simple least-squared minimization routine (see e.g., Davis et al. 2013; Onishi et al. 2015),[3] while for more complex modelling with more parameters, KinMS can be coupled to a suitable Markov Chain Monte Carlo (MCMC) code that allows the identification of the best model parameters through Bayesian analysis (see e.g., Onishi et al. 2017; Davis et al. 2017a). For this work, the best-fitting model parameters were determined using the MCMC method implemented in the `emcee` sampler (Foreman-Mackey et al. 2013).

4.1. The simulated model

The main function of `KinMS` takes, as input, parameters that describe the cube itself, the gas disc and galaxy. Parameters specifying the cube include the cube's dimensions along its three axes (RA, DEC, velocity), cell size and beam size (which are both discussed in Section 3.3.1). Parameters describing the gas disc include its position angle, the gas mass, gas velocity dispersion gas distribution and gas kinematics. Parameters related to the galaxy include its distance and inclination. Parameters may be assigned fixed values or may be included as free parameters in the modelling. The parameters used in this work are listed and discussed below in Section 4.2.5. The parametrization the gas distribution, gas kinematics and mass model for the total mass are discussed directly below in Sections 4.1.1 and 4.1.2.

[1] `https://github.com/TimothyADavis/KinMSpy`

[2] Due to time constraints, the modelling of the imaged data cube for NGC 2513 was carried out by Dieu D. Nguyen and Hai N. Ngo, who also kindly provided details of the `KinMS` code used for the modelling and the final results.

[3] Note that for these two papers, only the PVD is fitted. Later versions of the package allow fitting of the entire data cube also when using the least-squared minimization method, but, because it is computationally more expensive, this method is better suited to modelling with a small number of parameters.

4.1.1. Gas distribution

The gas distribution can be ascertained from the surface brightness profile of the gas. This can be specified analytically, assuming, for instance, a thin exponential disc, Gaussian ring or various combinations and variations thereof.[4] It is, however, also possible to construct a model of the observed gas distribution directly with the `SkySampler`[5] tool described in Smith et al. (2019), which is what was done in this work. This tool obtains the spatial information of the gas distribution independent of the kinematics by integrating along the velocity axis (creating, in effect, a zeroth moment map) of the data cube.

The cube[6] is then sampled with a large number of point sources, where more point sources are generated in locations of brighter molecular gas emission (Davis et al. 2020). The point sources are deprojected into the disc plane at each step of the MCMC sampling process to reproduce the gas distributions for the corresponding position angle and inclination of that iteration in the MCMC framework. Each point source can then be assigned a velocity that is drawn from the kinematic modelling, which is also part of the `KinMS` routine (Davis et al. 2020). To ensure that the gas distribution of the modelled cube matches the observed gas distribution, the integrated flux f is included as a free parameter in the modelling (cf. Nguyen et al. 2021).

4.1.2. Gas kinematics and mass model

As molecular gas is dynamically cold, it is the usual practice when modelling with `KinMS` to assume that the gas follows circular orbits and therefore that its rotation varies only radially (see e.g., Davis et al. 2013; Davis et al. 2017b; Onishi et al. 2017), and Section 1.2). This derives from the fact that cold gas is far less likely to react to dynamical factors such as magnetic fields or ionization pressure and that the dynamics are predominantly governed by the central gravitational potential. However, for a complete profile of the gas' motion and to take observational effects such as beam smearing into consideration, the gas velocity dispersion is included as a spatially constant, but free parameter in the MCMC modelling.

The circular velocity profile of the molecular gas was constructed with the `mge_circular_velocity` procedure of the Jeans Anisotropic Modelling (JAM) package[7] Cappellari (2008). As discussed in Section 1.1, the molecular gas kinematics are governed by the total gravitational potential of the galaxy, which includes the luminous matter (stars, gas and dust) and dark components (dark matter and putative SMBH). The dark matter contribution is sufficiently small in this region of the galaxy to not be included in any modelling. The luminous mass, however, must be accounted for and it is here that the mass model resulting from the MGE parametrization is used as an

[4] See the example tutorial of `KinMS_fitter`: `https://github.com/TimothyADavis/KinMS_fitter`

[5] `https://github.com/Mark-D-Smith/KinMS-skySampler`

[6] Ideally, this should be the cube of the clean component model created in the **`tclean`** process (see Section 3.3.1 and e.g., Figure 3.21) as sampling the convolved data cube directly can lead to over-smearing the data (Smith et al. 2019; Davis et al. 2020). For this work, the convolved data cube was used.

[7] `https://www-astro.physics.ox.ac.uk/~cappellari/software/`

Table 4.1.: Comparison of the MGE described in Chapter 2 and the MGE used for the modelling.

j	$I'_{j,1}$ ($L_{\odot,J}$ pc^{-2})	$\sigma'_{j,1}$ (arcsec)	$q'_{j,1}$	$I'_{j,2}$ ($L_{\odot,J}$ pc^{-2})	$\sigma'_{j,2}$ (arcsec)	$q'_{j,2}$
1	12996.838	0.090	0.70	26492.759	0.38	0.75
2	4355.724	0.734	0.65	7369.552	0.65	0.5
3	3055.770	0.903	0.75	5667.383	1.252	0.5
4	2896.665	1.545	0.65	12182.299	1.252	1.0
5	1408.342	2.737	0.687			
6	422.055	4.344	0.741			
7	246.071	7.734	0.75			
8	108.177	14.480	0.701			
9	36.805	22.845	0.748			
10	29.827	36.713	0.65			

Notes. Columns. (1) Summation index j of the Gaussians. (2) – (4) MGE results from this work as given in Table 2.10 for the Vega zero point; (5) – (7) The same, but for the MGE used in the modelling, which also uses the Vega zero point.

input (see Table 2.6), whereby the surface luminosity density I' is multiplied with the mass-to-light ratio M/L of the galaxy (see Eq. 2.1). In addition to the luminous mass, the putative SMBH mass is also included as an input in the velocity profile function, as is the inclination and distance of the galaxy. While the distance is defined as a fixed parameter in the fitting, the M/L, SMBH mass and inclination are included as free parameters in the MCMC fit.

It should be noted that for the current modelling of the data cube described in this Chapter, a different MGE was used to that presented in Chapter 2, the final results of which were presented in Table 2.10. As discussed in Section 2.3.9, the flattest Gaussian in the MGE fit, which is stipulated by the lowest value of the observed axial ratio q', imposes a minimum inclination for which `KinMS` can successfully model the dynamics of the gas disc. As mentioned in that Section, one of the Gaussians in the initial MGE I produced had a value as low as $q' < 0.2$, which resulted in an inclination too low for successful modelling. Although this was later corrected (as described in Section 2.3.9), the modelling described here was completed with an alternative MGE.[8] This MGE only found four useful Gaussians, thus modelling the inner region of the galaxy, whereas the MGE I produced found a total of 10 useful Gaussians, covering the full extent of the galaxy (see Section 2.3.3). Both MGEs are shown in Table 4.1. Modelling the NGC 2513 data cube in `KinMS` with the MGE that I produced will allow a comparison of the `KinMS` modelling results for both MGEs.

[8] The MGE used in the modelling was produced and kindly provided by Hai N. Ngo.

4.2. Bayesian Analysis

The best-fitting set of model parameters from the data cube were chosen by using a Bayesian analysis technique, a form of statistical analysis based on conditional probability as expressed by Bayes' Theorem.[9] In terms of probability distributions associated with modelling data, the theorem can be expressed as:

$$p(\Theta|\mathbf{x}) = \frac{p(\Theta)\, p(\mathbf{x}|\Theta)}{p(\mathbf{x})} \ , \tag{4.1}$$

where $p(\Theta|\mathbf{x})$ is the so-called *posterior probability distribution* (see e.g., Zocchi et al. 2016). Here, Θ is a vector containing all the parameters used for constructing the model (these are the 'free' parameters referred to above, i.e., the SMBH mass, the M/L ratio, inclination, position angle, etc). The prior knowledge or beliefs regarding each of these parameters are encoded in the *prior distribution* $p(\Theta)$, which is updated after considering the data (or evidence) $\mathbf{x}$. The quantity $p(\mathbf{x}|\Theta)$ is the likelihood function and expresses the probability of the data $\mathbf{x}$ given that Θ is true. The product of the likelihood function and priors are weighted by $p(\mathbf{x})$, which is the marginal probability of the data (or probability of the evidence). The posterior probability $p(\Theta|\mathbf{x})$ therefore quantifies the probability of Θ of being true after considering the data $\mathbf{x}$.[10]

4.2.1. Likelihood Function

The likelihood function can generally be expressed as the product of each likelihood function for each of N points in the observed data:

$$p(\mathbf{x}|\Theta) = \prod_{i=1}^{N} p(x_i|\Theta) \ . \tag{4.2}$$

The goal of the fitting process is to find those parameters that produce a model that most closely resembles the data. More formally, the aim is to find the maximum likelihood estimate, i.e. to find the parameter vector Θ^* that maximizes the likelihood as defined in Eq. 4.2.

In practice, the calculation of the maximum likelihood estimate can be greatly simplified by applying the logarithm to the likelihood function.[11] The use of a logarithm allows the transformation of the *product* over all the data points into the *sum* over the data points, which is much easier to work with.[12] Moreover, as most optimizers work by

[9]Unless otherwise specified, the main reference for this section is Deisenroth et al. (2020).

[10]See also `https://www.stat.cmu.edu/~brian/463-663/week09/Chapter%2003.pdf`

[11]This is allowed as the logarithmic function is a strictly monotonically increasing function and hence leaves the location of the maximum of the likelihood function unchanged. Note that a logarithm of any base is permitted.

[12]The logarithm also provides more numerical stability: the likelihood probabilities can be extremely small and taking the logarithm increases these very small numbers to larger negative values that can be more easily handled by a machine.

minimizing a function, it is conventional to *minimize* the negative log-likelihood rather than to maximize the log-likelihood. The negative log-likelihood $\mathcal{L}$ is then given by:

$$\mathcal{L} = -\log p(\mathbf{x}|\Theta) = -\log \left(\prod_{i=1}^{N} p(x_i|\Theta) \right) = -\sum_{i=1}^{N} \log p(x_i|\Theta) \ . \tag{4.3}$$

Assuming that the errors can be described by Gaussian noise, the likelihood can be expressed in terms of a Gaussian probability distribution. Thus, for a Gaussian distribution $\mathcal{N} = \mathcal{N}(\mu, \sigma)$ with mean μ and standard deviation σ, the Gaussian negative log-likelihood can be written as:[13]

$$\mathcal{L} = -\sum_{i=1}^{N} \log p(x_i|\Theta) = -\sum_{i=1}^{N} \log \mathcal{N}(x_i|\Theta, \sigma) = \frac{1}{2\sigma} \sum_{i=1}^{N} (x_i - \Theta)^2 \ . \tag{4.4}$$

Noting that σ is constant, minimizing $\mathcal{L}$ becomes equivalent to finding the minimum χ^2. For this work, this translates to a comparison of the CO flux in each pixel of the three-dimensional data cube $D_{x,y,v}$ with that in the model cube $M_{x,y,v}$ produced via the `KinMS` modelling function. Here, both cubes are three dimensional with two spatial axes of size $N_x \times N_y$ and one spectral axis in velocity of size N_v (cf. Davis et al. 2017a). The χ^2 is then given by:

$$\chi^2 = \sum_{i=1}^{N_v} \frac{(D_{x,y,i} - M_{x,y,i})^2}{(\mathrm{RMS})^2} \ , \tag{4.5}$$

where the RMS is the noise measured in line-free channels of the data cube and is taken as the measurement uncertainty. The χ^2 must then be normalized by the number of degrees of freedom. Here, this is given by the size of the data cube, i.e., $N_x N_y N_v$. Applying logarithmic rules, this finally gives an expression for the log-likelihood of:

$$\mathcal{L} = -\frac{1}{2}(\chi^2 - N_x N_y N_z) \ , \tag{4.6}$$

which is what is minimized in the fitting procedure (see e.g., Davis et al. 2017a).

4.2.2. Chi-squared uncertainties

When inferring errors with the traditional χ^2 method, the 1σ confidence level (CL) is usually defined by $\Delta\chi^2 = 1$ where $\Delta\chi^2 = \chi^2 - \chi^2_{min}$. For large data sets such as the ALMA data used for this measurement, Van Den Bosch & Van De Ven (2009) noted that, when inferring errors with the traditional χ^2 method, the uncertainties of the χ^2 itself become important and, in particular, statistical errors tend to be underestimated due to the dominance of systematic errors. They thus suggest rescaling the $\Delta\chi^2$ by the standard deviation of χ^2, which, for N number of constraints (i.e., data points) and M number of free parameters, is given as $\sqrt{2(N-M)} = \sqrt{2N}$ for $N \gg M$ so that

[13] For the full derivation of this expression, see Deisenroth et al. (2020).

$\Delta\chi^2 = \sqrt{2N}$. Applying this method to MCMC modelling, Mitzkus et al. (2016) found that an appropriate equivalent rescaling for errors in the MCMC framework is to have $\Delta\chi^2 = (2N)^{1/4}$, which is what was used in this work.

An additional aspect to consider when dealing with ALMA data is that due to convolution of the data with the synthesized beam, pixels in the imaged cube that are closer than 1 FWHM of the beam are strongly correlated and thus cannot provide independent measures of the χ^2 (see e.g., Davis et al. 2017a; Smith et al. 2019; Nguyen et al. 2020). The approach of rescaling the $\Delta\chi^2$ as described here is one method of ensuring that the effect of these correlations is minimal (see e.g., North et al. 2019; Nguyen et al. 2020).[14]

4.2.3. Priors

As mentioned above, the *priors* express prior knowledge concerning the parameters used in the modelling. More explicitly, the priors serve to place reasonable constraints on the ranges of values chosen for the modelling.

Consider, for example, the mass of the SMBH. From the initial estimate based on the $M_{BH} - \sigma_*$ relation (see Section 1.3), I expect a measurement in the region of a relatively high $10^8\,\mathrm{M}_\odot$ or low $10^9\,\mathrm{M}_\odot$. To include the possibility of a lower mass SMBH or no SMBH, while providing for some overreach in the high-mass range, a reasonable range of search values would be, e.g., between $10^2\,\mathrm{M}_\odot$ and $10^{12}\,\mathrm{M}_\odot$. These values, in effect, form the boundaries of the uniform prior distribution for the SMBH mass.[15]

To set reasonable priors, it is often useful to set *flat* priors, where $p(\Theta)$ = constant. This means that all parameter values within the search-range are equally likely to be sampled, i.e., no bias is given to any particular value. In other words, a flat prior provides a first order constraint on the parameter, while not being restrictive. As such, it assumes maximum ignorance within the range that is set and thus offers the most objective sampling. Flat priors were chosen for all the free parameters included in the modelling and are listed in Table 4.2 in Section 4.2.4.

4.2.4. MCMC fitting with the emcee sampler

MCMC is, in essence, a tool for exploring or *sampling* parameter space and finding a distribution that represents a selection of models that describe the data reasonably well. The variance in these models represents how well we are able to constrain the parameters of the models.[16] The parameter space defined by the vector Θ introduced above can be imagined as a grid of values that is explored stochastically by an ensemble of *walkers*. These walkers randomly sample values that fall within a range set by the priors. Each walker starts at a random possible value of Θ. It samples a new random

[14] The issue can also be corrected by introducing the inverse of the covariance matrix that describes the pixel-to-pixel correlations into Eq. 4.5 as in, e.g., Davis et al. (2017a); Smith et al. (2019); Nguyen et al. (2020).

[15] SMBH mass is usually calculated in log-space where a value of zero for no black hole is undefined. The value of $10^2\,\mathrm{M}_\odot$ is a few orders of magnitude below the accepted lower limit for a SMBH mass and as such acts as a proxy for no SMBH.

[16] See e.g., https://prappleizer.github.io/Tutorials/MCMC/MCMC_Tutorial_Solution.html

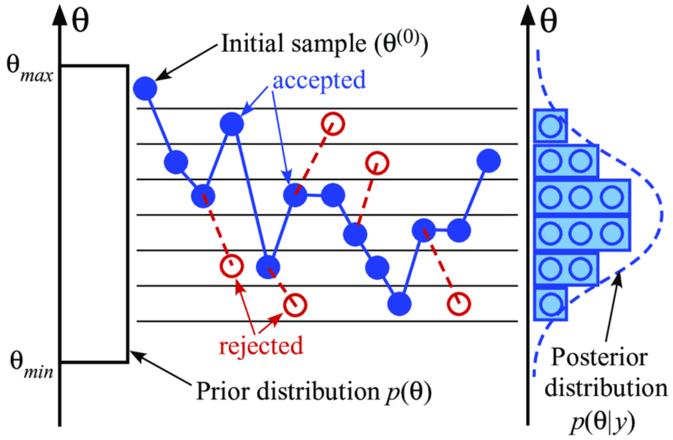

Figure 4.1.: Schematic of an MCMC process of a walker sampling the parameter space defined by a prior distribution $p(\theta)$ within the range of θ_{min} to θ_{max}. The initial sample is $\theta^{(0)}$ and from here the walker steps from one point to the next. Blue filled circles indicate accepted values and each successful step is indicated by a solid blue line. Rejected values are shown by the red circles and their rejected paths by red dashed lines. Figure taken from Tomic et al. (2019).

possible value of Θ by considering the likelihood with regard to that value. That is, a model is generated with that particular Θ value and compared to the given data, usually with a χ^2 minimization as described above. If the χ^2 value of the proposed value lies above a threshold that is set to maximize the sampling process (i.e., the likelihood is good), the walker moves to the new value. If not, the walker considers another value and the process is repeated. This idea is illustrated in Figure 4.1. An important feature is that each future step depends only on the current position of the walker and not on any preceding value. A succession of such steps is know as a *chain*,[17] and the chain is said to converge when the walkers sample a distribution that does not change with any subsequent steps. This stable distribution is in fact the posterior probability distribution defined in Eq. 4.1 and also shown in Figure 4.1.

An MCMC can be run multiple times with a different set of starting values, thus generating multiple chains. Should their outputs be similar, then it is clear that they are sampling the parameter space well, without being stuck in local maxima of the probability distribution.[18]

To get the best-fitting value for each model parameter from the posterior distribution independent of the others, the MCMC *marginalizes* over the the posterior probability surface. This simply involves integrating over all parameters excluding the parameter (or parameters) of interest, which is (are) held fixed. Thus for parameters $\Theta = (\theta_1, \theta_2, ..., \theta_N)$, to find the marginalized probability distribution of parameter θ_1, one integrates over the

[17] See e.g., https://www.publichealth.columbia.edu/research/population-health-methods/markov-chain-monte-carlo

[18] Cf. https://drvalle1.github.io/20_MCMC_convergence.html

remaining parameters:[19]

$$p(\theta_1) = \int p(\theta_1, \theta_2, ..., \theta_N)\ d\theta_2, d\theta_3, ..., d\theta_N\ . \qquad (4.7)$$

The best-fitting value is then taken as the median of the marginalized posterior samples. In effect, the marginalized distribution of a parameter or parameters can be found by projecting the samples from their N-dimensional parameter space onto an n-dimensional histogram. This is usually done by producing a corner plot of the MCMC chain where the distributions are shown as one- and two-dimensional histograms. The former marginalizes over one parameter, the latter over two.[20] The corner plot from the MCMC results for NGC 2513 is shown in Figure 5.1 and discussed in Chapter 5.

As mentioned, this work uses the `emcee` sampler by Foreman-Mackey et al. (2013),[21] which is a Python implementation of the affine invariant MCMC ensemble sampler by Goodman & Weare (2010).[22] To set up a distribution that can be sampled, one needs to define four functions: the model function, the log-likelihood function, a function that defines the priors and the log-probability function that returns the logarithm of the posterior probability distribution. For this work, the model function is the `KinMS` main function described above in Section 4.1. The log-likelihood and priors are discussed above in Sections 4.2.1 and 4.2.3, respectively. The log-probability function calls the log-likelihood and prior functions and, on condition that the parameters and priors are valid, returns the logarithm of the posterior probability distribution.[23]

To sample the resulting distribution, the `emcee` sampler takes the following required arguments: the number of walkers, the number of dimensions in parameter space (i.e., the number of free parameters that should be sampled and modelled) and the log-probability function.

4.2.5. Fitting procedure for NGC 2513

Fitting took place over a 40×40 pixel^2 area and 90 velocity channels of the cube, which was sufficient to cover the main emission features of the gas disc. As mentioned, flat priors with reasonable constraints on their ranges were used for the nine free parameters, which are listed in Table 4.2. Apart from physical considerations, the constraints were chosen based on a comparison of moment maps generated from initial MCMC runs with those generated from the data.

Of these nine free parameters, the RA, DEC and velocity that parametrize the kinematic centre of the gas disc are set as so-called nuisance parameters. These are parameters that are necessary for the modelling, but that are not of primary interest.[24] The spatial

[19] See e.g., https://www.astro.umd.edu/~miller/teaching/astr288a/lecture08.pdf
[20] See https://emcee.readthedocs.io/en/stable/tutorials/line/
[21] https://emcee.readthedocs.io/en/stable/
[22] Affine invariant simply means that the performance of the MCMC is not affected by affine transformations in space.
[23] See https://emcee.readthedocs.io/en/stable/user/sampler/
[24] See e.g., https://www.sciencedirect.com/topics/computer-science/nuisance-parameter.

coordinates (RA and DEC) were constrained to lie within ± 2″ of the ALMA phase centre of the disc, while the velocity offset was allowed to vary within ± 100 km s^{-1} relative to the systemic velocity of 4628.78 km s^{-1}.

The free parameters pertaining to the SMBH and stellar mass, respectively, are the SMBH mass and stellar mass-to-light ratio M/L. As indicated, it is usual to set the SMBH mass in log space. Thus, the prior for the mass of the SMBH M_{BH} was flat in log-space, with the mass being allowed to vary between $\log_{10}(M_{BH}/\mathrm{M}_\odot) = 2$ and 12, a range that encompasses the initial estimates of $\log_{10}(M_{BH}/\mathrm{M}_\odot) \approx 8.5 - 9$. The mass-to-light ratio M/L was assumed to be constant across the fitting area and was allowed to vary between 0 and 50 $\mathrm{M}_\odot/\mathrm{L}_{\odot,J}$. For the purposes of this project, a constant M/L was assumed. Potential problems with this assumption are briefly discussed in Section 5.2.

The remaining four free parameters pertain to the molecular gas disc. These include the disc's inclination, position angle, velocity dispersion and integrated flux. As discussed in Section 2.3.9, the flattest Gaussian in the MGE fit dictates the minimum possible inclination for which the MGE can be successfully used in dynamical modelling. The inclination here was allowed to vary across the full physical range allowed by the MGE with a minimum constraint of 60.5° and maximum constraint of 89.8°. The gas velocity dispersion and position angle were given upper constraints of 100 km s^{-1} and 359°, respectively, while the flux was allowed to vary between 0.1 and 200 Jy km s^{-1}. The upper constraint on the position angle includes values high enough to cover the initial guess discussed in Section 3.4.2. The upper constraint on the velocity dispersion was chosen to account for the relatively high velocity dispersion evident in the central regions of the disc as shown in the second moment map shown and discussed in Section 3.4.2.

Notable fixed parameters were the galaxy distance, systemic velocity and mass of the gas disc and are listed in Table 4.3. The latter, with an upper constraint set at $\sim 3 \times 10^8$ $\mathrm{M}_\odot$ (Davis et al. 2019), was considered too low to have an extensive effect on the total mass and was thus not included as a free parameter to be modelled. It was included as a fixed parameter in the general `KinMS` model (see Section 4.1).

For the MCMC fitting, 900 walkers were used to sample the nine-dimensional parameter space, with the walkers positioned to start in a compact ball around the initial guesses. The initial guesses for these free parameters were selected by doing a preliminary MCMC run with $\sim$ 5000 steps. This initial run also aimed to evaluate the goodness of the fit and to ensure that the MCMC chains converged. For the actual fitting, the best chains were run for a total of 3×10^5 steps with a 50 % burn-in phase[25] to produce a final posterior distribution of the nine free parameters. The posterior probability surfaces were then marginalized over to find the best-fitting values, and the 68 % (1σ) and 99 % (3σ) CLs. These results are presented and discussed in the following Chapter.

[25] A burn-in phase serves to allow the MCMC run to gain independence from the initial values and enter a region where the points sampled are more representative of the distribution being sampled. See e.g., `https://www.johndcook.com/blog/2016/01/25/mcmc-burn-in/`

Table 4.2.: Best fitting model parameters and statistical uncertainties.

Parameter	Search range	Best fit	Error (68% conf.)	Error (99% conf.)
SMBH:				
$\log_{10}(M_{BH}/M_\odot)$	$(2 \rightarrow 12)$	9.09	$-0.08, +0.05$	$-0.48, +0.14$
Luminous mass:				
$M/L(M_\odot)$	$(0 \rightarrow 55)$	2.00	$-1.15, +1.82$	$-3.32, +8.04$
Molecular gas disc:				
Position angle (deg)	$(0 \rightarrow 359)$	343.28	$-2.36, +2.44$	$-7.44, +7.66$
Inclination (deg)	$(60.15 \rightarrow 89.8)$	60.65	$-0.37, +0.79$	$-0.50, +3.19$
Flux (Jy km s^{-1})	$(0.1 \rightarrow 200)$	4.89	$-0.26, +0.27$	$-0.70, +0.80$
Velocity dispersion (km s^{-1})	$(0 \rightarrow 100)$	35.93	$-6.45, +8.21$	$-15.27, +26.64$
Nuisance parameters:				
Centre X offset (arcsec)	$(-2 \rightarrow 2)$	-0.02	$-0.01, +0.02$	$-0.04, +0.04$
Centre Y offset (arcsec)	$(-2 \rightarrow 2)$	-0.08	$-0.02, +0.02$	$-0.05, +0.05$
Centre velocity offset (km s^{-1})	$(-100 \rightarrow 100)$	-22.84	$-4.65, +5.02$	$-16.00, +17.20$

Notes. Columns: (1) Model parameters. (2) Priors for each parameter. Priors are assumed to be flat in linear space for all parameters other than the SMBH mass, which is flat in log space. (3) The best-fit value. (4) The 1σ (68 % CL) uncertainties. (5) The 3σ (99 % CL) uncertainties. The nuisance parameters X and Y are defined relative to the ALMA data phase centre position of ICRS 08:02:24.656880 +09.24.48.93120, while the nuisance parameter of velocity is defined relative to the systemic velocity of 4628.78 km s^{-1}.

Table 4.3.: Fixed parameters assumed for the modelling.

Galaxy distance (Mpc)	Systemic velocity km s^{-1}	Gas disc mass $M_\odot$
70.8	4628.78	3×10^8

References. Distance is from Veale et al. (2018). Systemic velocity is from the ALMA data cube header. Molecular gas mass is from Davis et al. (2019).

5. Results and Discussion

5.1. Results

In this Chapter, I present and discuss the best-fit results obtained from the MCMC fitting. Both the one- and two-dimensional marginalizations of all nine model parameters are depicted in Figure 5.1. The former, which appear along the diagonal of the plot, show the marginalized distribution of each parameter independent of the others. The central dashed line is the median, which is taken as the best-fitting value, while the other two dashed lines indicate the 1σ (68 %) confidence interval. The best-fitting value given by the median together with the 1σ error is given above each histogram. The two-dimensional histograms (contour plots in the remaining panels) show the marginalized two-dimensional distributions. The inner to outer contour lines show the 1σ (68 %), 2σ (95 %) and 3σ (99 %) CLs, respectively. All best-fitting values with their respective 1σ and 3σ uncertainties are shown in Table 4.2.

The contour plots are useful for showing such properties as covariance between parameters, which can in turn indicate degeneracies between these parameters. A strong covariance can be seen, for instance, between the mass-to-light ratio M/L and SMBH mass (although this is accentuated by plotting a linear against a log scale). This is not unusual as the M/L is indicative of the stellar mass contribution to the total mass, and a smaller contribution of the stellar mass must imply a greater contribution from the SMBH mass. Nonetheless, the more dependent the SMBH mass is on the luminous mass model, especially at scales close to or within the R_{SOI}, the more the SMBH mass measurement is susceptible to uncertainties in that luminous mass model. This is discussed more fully in Section 5.2.

The gas velocity dispersion also shows a slight covariance with the flux, the SMBH mass and mass-to-light ratio, respectively, but otherwise, the parameters are largely uncorrelated. The inclination of the gas disc, which is assumed to be the same as that of the galaxy, is relatively poorly constrained and is seen to clip at the lower bounds of the prior. This indicates that, in theory, lower values for the inclination might be possible (see Section 4.2.5), but, as discussed in Section 2.3.9 , deprojection then becomes difficult.

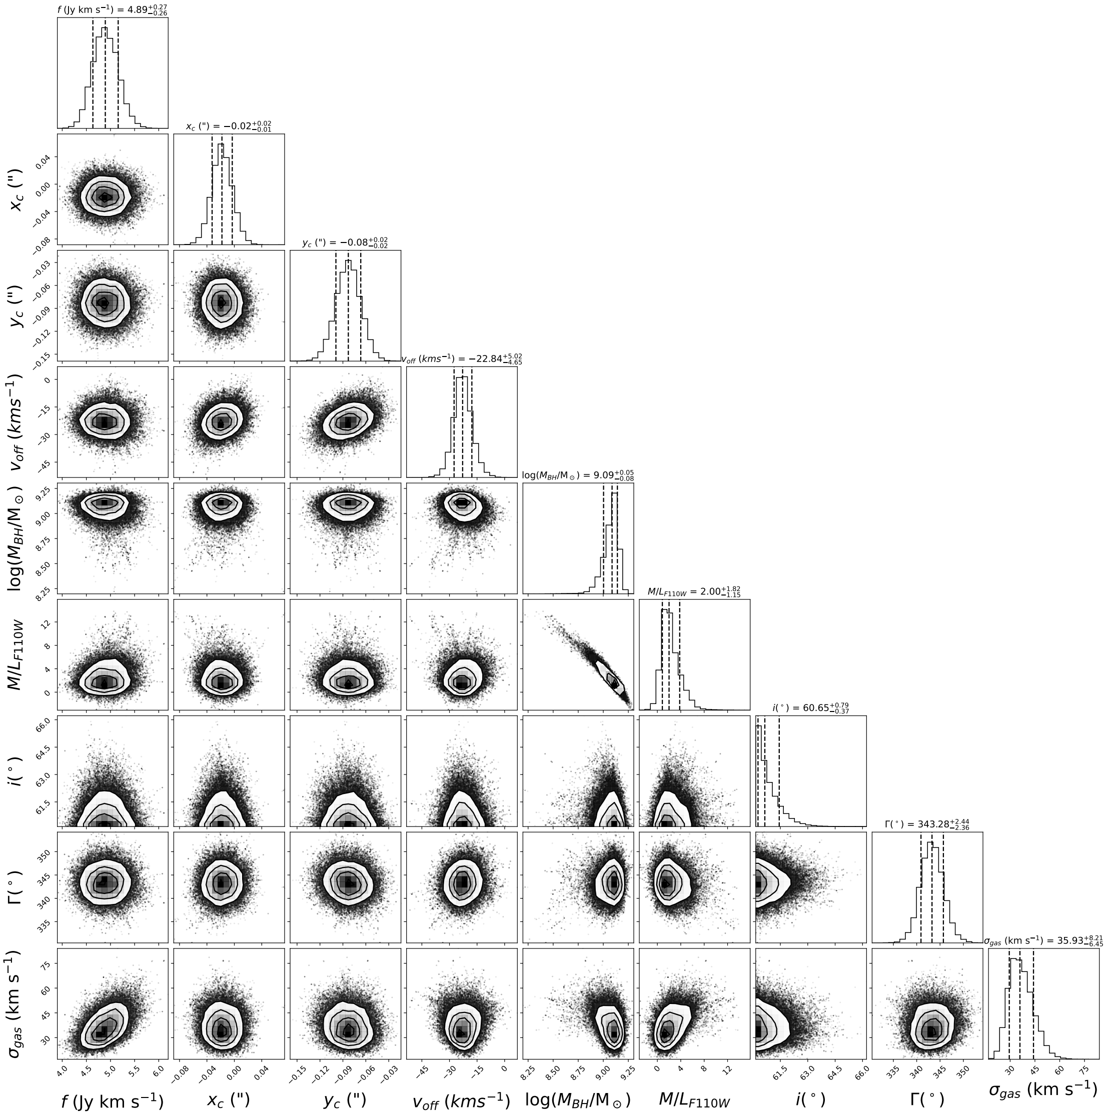

Figure 5.1.: Visualization of the nine-dimensional parameter space explored by the fitting procedure. The top panel of each column shows a one-dimensional histogram of the marginalized posterior distribution of a given parameter. Here, the central dashed line is the median, which is taken as the best-fit value; the outer dashed lines indicate the 1σ (68 %) confidence interval. The two-dimensional contour plots in the panels below show the marginalized two-dimensional distributions of the respective parameters. Here, the inner to outer contour lines show the 1σ (68 %), 2σ (95 %) and 3σ (99 %) CLs, respectively. The chains from which the plot was constructed were kindly provided by Hai N. Ngo.

The results clearly show the presence of a compact dark object with a mass of $(1.23^{+0.14}_{-0.23}) \times 10^9$ M$_\odot$ at the 68 % CL and $(1.23^{+1.36}_{-0.40}) \times 10^9$ M$_\odot$ at the 99 % CL, re-

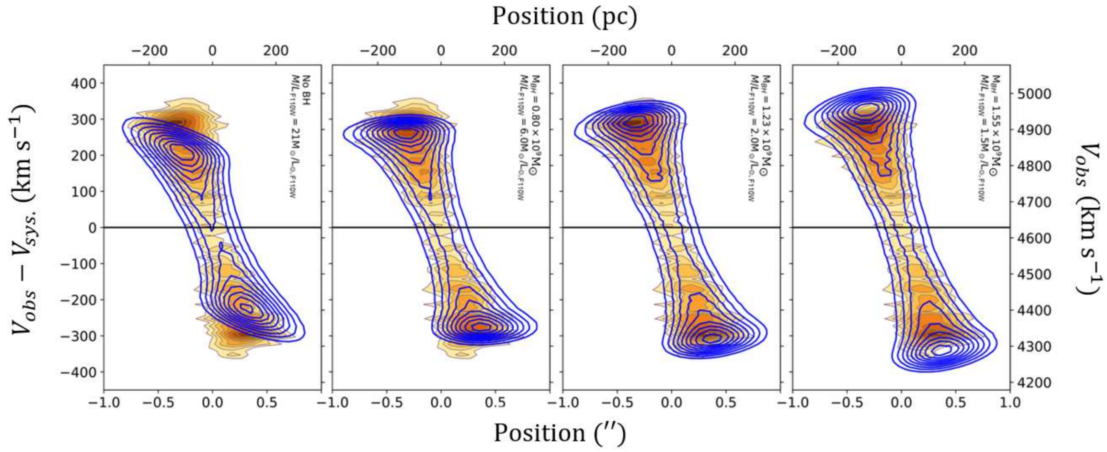

Figure 5.2.: Position velocity diagrams (PVDs) of the observed ^{12}CO(2-1) emission in NGC 2513 (orange with black contours) overlaid with PVDs from the modelled cube (blue contours). The first panel shows a fit with no SMBH; the second panel shows a fit with a SMBH mass that is too low; the third panel shows the best-fit SMBH mass and the fourth panel shows an overly massive SMBH. The legend in each panel gives the SMBH mass and corresponding M/L. Plot kindly provided by Hai N. Ngo.

spectively. Using the definition from Eq. 1.1, this yields a sphere of influence with radius $R_{SOI} = 82.67\,\mathrm{pc}$. The best-fitting J-band mass-to-light ratio, which is assumed to be constant throughout the galaxy, is $2.00^{+1.82}_{-1.15}$ $\mathrm{M}_{\odot}/\mathrm{L}_{\odot,J}$ (68 % CL) and $2.00^{+8.04}_{-3.32}$ $\mathrm{M}_{\odot}/\mathrm{L}_{\odot,J}$ (99 % CL).[1]

Figure 5.2 shows the observed PVD overlaid with the best-fitting model and, for comparison, models with no SMBH, a SMBH with too low a mass and one with too large a mass. The best-fitting model evidently reproduces the kinematics of the molecular gas, while the models with no SMBH and smaller SMBH mass fail to fit the peak increase in velocity towards the centre of the galaxy. On the contrary, the model with an overly large SMBH mass produces velocities in this region that are far too high.

Figure 5.3 shows a comparison between the intensity and velocity maps extracted from the data and those extracted from the simulated cube. Also depicted are the residuals derived by subtracting the model velocity from the data velocity. The model's intensity map reproduces both the flux and extent of the molecular gas reasonably well. The velocity map also reproduces well the one generated from the data. This is underscored by the residuals, which are mostly $\sim 30\,\mathrm{km\,s^{-1}}$. Although several regions reach over $\sim 80\,\mathrm{km\,s^{-1}}$, this is still minimal in comparison to the top velocity of $\sim 400\,\mathrm{km\,s^{-1}}$ reached by the gas.

[1]The lower 3σ error is less than the lower constraint of 0.0 that was set for the M/L prior. This is probably due to a minor error in the MCMC code and will be revisited in future modelling.

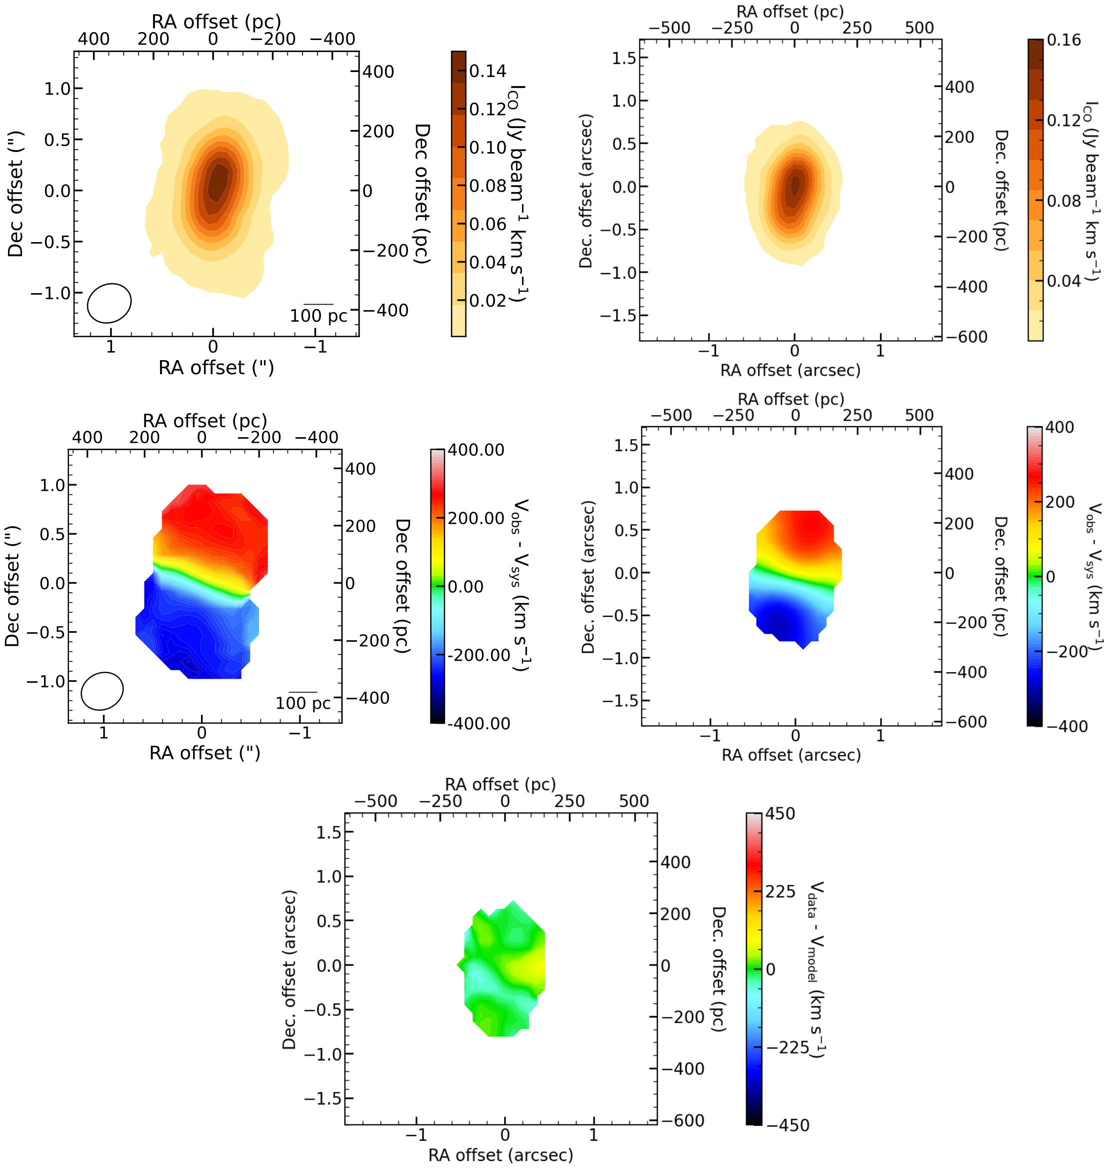

Figure 5.3.: *Top panels*: The zeroth moment (intensity) maps constructed from the data (left) and model (right). *Middle panels*: The first moment (velocity) maps constructed from the data (left) and model (right). *Bottom panel*: Residuals derived by subtracting the model velocity from the data velocity. The beam-size is indicated in the moment maps of the data as an ellipse in the bottom-left corner. The moment maps constructed from the model and plot of the residuals were kindly provided by Hai N. Ngo.

5.2. Main uncertainties

5.2.1. Distance

One of the most significant systematic uncertainties associated with any dynamical SMBH mass measurement is the distance D assumed to the galaxy. The measured SMBH mass M_{BH} scales linearly with the assumed distance (see e.g., Davis et al. 2017b; Smith et al. 2019). This derives from the fact that the modelling assumes circular velocity v_c of the tracer around the SMBH at a distance r with $M_{BH} \propto v_c^2 r$, where $r \propto D$ (Smith et al. 2021).[2] As mentioned in Section 1.4, the distance assumed in this thesis is 70.8 Mpc as taken from Ma et al. (2014) and is derived from using group-corrected flow velocities. Although no associated uncertainties are quoted in Ma et al. (2014), Quenneville et al. (2023) adopt a 10 % uncertainty on distances derived from group-corrected flow velocities for their work. While the distance uncertainty estimate can be included as a second error term on the SMBH mass (see e.g., Davis et al. 2017b), the standard practice, which was adopted in this work, is to exclude it from the systematic uncertainty (see e.g., Smith et al. 2019; North et al. 2019) as the measurement can be corrected to the distance adopted by the proportional scaling (Smith et al. 2021; Nguyen et al. 2021).

5.2.2. Luminous mass model

Another source of uncertainty on the SMBH mass is the luminous mass model produced with the MGE used to estimate the contribution of the luminous mass to the gas kinematics. As mentioned in Chapter 2, an under- or overestimation of the luminous mass will directly affect the SMBH mass. The main sources of error in the mass model can be attributed to the presence of dust and the conversion between light and mass. As discussed in Section 2.3.6, the former can be alleviated by masking the regions affected by dust. Uncertainty in the conversion between light and mass derives mainly from assuming a constant mass-to-light ratio, which, as mentioned in Section 4.2.5, was the assumption in this work. Although sufficient for the purposes of this project, assuming a constant M/L is not always entirely realistic, especially for galaxies containing molecular gas where the star formation rate is significant (Davis et al. 2017a). Davis & McDermid (2016), however, found that M/L gradients in ETGs hosting molecular gas tend to be relatively shallow due to lower star formation rates and a subsequent lower mass fraction of young stars. Examining the star formation rate and looking at the M/L gradient of NGC 2513 more closely would therefore be relevant for future work on this project. Further discussion of the potential problems associated with assuming a constant M/L in dynamical SMBH mass modelling as well as the methodology of radially varying the M/L can be found in e.g., Davis et al. (2017a), North et al. (2019), and Thater et al. (2021). Another related avenue of exploration would be to assess the effect of a spatially varying initial mass function (IMF) on the M/L and, subsequently, on the SMBH mass as in Thater et al. (2023).

[2] For any given angular radius observed, the corresponding physical radius will increase with distance.

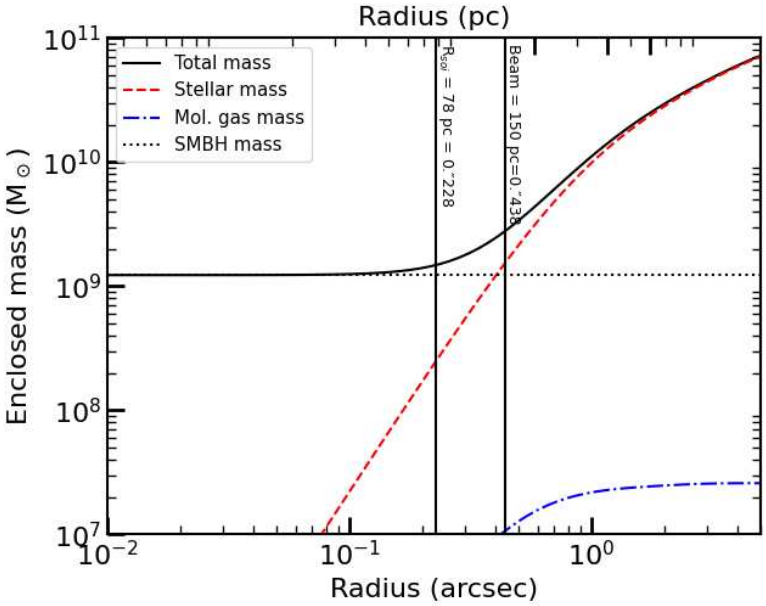

Figure 5.4.: Cumulative mass profile of NGC 2513 showing the total mass (black solid line), SMBH mass (black dotted line), stellar mass (red dashed line), and molecular gas mass (blue dot-dashed line), enclosed as a function of radius. The vertical lines indicate the R_{SOI} and synthesized beam, respectively. Plot kindly provided by Hai N. Ngo.

The effects of the luminous mass model and related uncertainties on the SMBH measurement can be inferred by examining the enclosed stellar mass as a function of radius in relation to its contribution to the total mass. As shown in Figure 5.4, the stellar mass makes only a minimal contribution to the total mass within the R_{SOI}, reaching only about $\sim 10\,\%$ of the total mass at the R_{SOI} itself.[3] In this region, it is clearly the SMBH mass that, at well over $\sim 90\,\%$, dominates the potential. The stellar mass first becomes significant beyond the R_{SOI} at $\sim 0.4''$ ($\sim 135\,\mathrm{pc}$) where it contributes to about $\sim 50\,\%$ of the total mass. It is thus reasonable to assume that the M/L, while correlating with the SMBH mass, has reasonably little effect on the best fit SMBH mass value (see e.g., North et al. 2019).

5.2.3. Circular motion

As mentioned in Section 4.1.2, the modelling assumes the molecular gas to have circular motion. Non-circular motions that result, for instance, from inflows and outflows of gas, streaming or axisymmetric gravitational potential[4] are not taken into account, but could

[3] The R_{SOI} shown in the plot is slightly different from that given above in Section 5.1 as a slightly different velocity dispersion was used to compute it.

[4] Axisymmetry in gravitational potential may result, for instance, from spiral arms (see e.g., Davis et al. 2017a).

naturally affect the modelling. The assumption of predominantly circular motions for the gas does, however, seem reasonable as the first moment map (Figures 3.27 and 5.3) indicates clear Keplerian motion. In addition, the residuals (Figure 5.3) between the data and best-fitting model are fairly flat and do not suggest any strong deviation from circular motion. Any contribution of non-circular motion to the uncertainties is therefore deemed to be minimal.

5.2.4. Gas velocity dispersion

The gas velocity dispersion, which was assumed to be constant for this work, offers another source of uncertainty. Although the centre of the disc shows a high gas velocity dispersion with values reaching above 200 $\mathrm{km\,s^{-1}}$ (see Figure 3.27), the majority of the gas has values of $\sim 55\,\mathrm{km\,s^{-1}}$ and a best-fit value from the modelling of $39.93^{+8.21}_{-6.45}$ km s^{-1} at the 1σ CL. This is still relatively high compared to, for example, Onishi et al. (2017) and North et al. (2019), who report a best-fit gas velocity dispersion of 12.53 ± 0.74 km s^{-1} and 8.32 ± 0.72 km s^{-1} (also at the 1σ CLs) for the ETGs NGC 3665 and NGC 0383, respectively. Nonetheless Ruffa et al. (2023) report a best-value of $61.71^{+5.18}_{-5.11}$ km s^{-1} (at the 1σ CL), which is significantly higher than the best-fit value for NGC 2513, although here, the errors are notably lower. In order to better constrain the gas velocity dispersion of NGC 2513 and take the effects of beam smearing into account, it might be worthwhile to create a model with a gas velocity dispersion that varies as a function of radius as described in Davis et al. (2017a); Nguyen et al. (2020), for future work on this gas disc.

5.2.5. Gas mass

Besides the stellar mass, dark matter and the mass of interstellar material (i.e. gas and dust) also contribute to the total dynamical mass of the galaxy. As mentioned in Section 1.1, the contribution of dark matter in the region being probed by the dynamics of the molecular gas disc is negligible and can be readily accounted for in the error budget of the SMBH mass. While the contribution of gas density can be significant in galaxies with a high gas fraction, ETGs tend to have lower quantities of gas, which are thus dynamically not significant (Davis & McDermid 2016).[5] This is also true for the gas mass in NGC 2513 where Davis et al. (2019) set an upper limit for the H_2 mass at $\sim 3 \times 10^8$ $M_\odot$,[6] which is two orders of magnitude lower than the estimated stellar mass of 6.35×10^{10} $M_\odot$ (see Section 2.3.10) and should thus not contribute significantly to the uncertainties in the mass model. Moreover, the cumulative mass profile in Figure 5.4 shows that the contribution of the molecular gas mass within the inner region of the galaxy considered in

[5] The contribution of interstellar dust to the overall mass is most likely even more insignificant and its contribution is thus ignored. Indeed, Kaviraj et al. (2012), using measurements of molecular gas-to-dust ratios from the literature find an average molecular gas-to-dust ratio of ~ 1000, which is an order of magnitude higher than that for the Milky Way, where dust contributes only about 1 % to the interstellar medium.

[6] Here, Davis et al. (2019) use the standard conversion factor X_{CO} between CO and H_2 and relation given in their Eq. 1.

the dynamical modelling is just under $\sim 10^{7.1}$, which is two orders of magnitude smaller than the best-fit mass of the SMBH.

5.3. The mass of the SMBH in NGC 2513 and the $M - \sigma_*$ scaling relation

This is the first measurement of the SMBH mass in NGC 2513 and, as such, a comparison to any previous mass measurements is not possible. Nonetheless, a comparison with the SMBH mass estimates derived from the $M_{BH} - \sigma_*$ scaling relation (see Section 1.4) can reveal whether the mass is consistent with the estimates and whether if falls within the expected range. Examining its position on the $M_{BH} - \sigma_*$ relation itself can also reveal its consistency with the scaling relation and provide a comparison with other SMBH masses, particularly those in other early-type galaxies.

The best-fit SMBH mass results of $(1.23^{+0.14}_{-0.23}) \times 10^9$ M$_\odot$ (1σ) and $(1.23^{+1.36}_{-0.40}) \times 10^9$ M$_\odot$ (3σ) obtained from the `KinMS` modelling compare relatively well with the estimated masses derived from the $M_{BH} - \sigma_*$ relation in Section 1.4.1. In particular, the estimate of 1.03×10^9 M$_\odot$ obtained from the original ETG fit equation of Harrisberg (2020) falls within the limit of the lower 1σ error of the measurement. The fit equation here, however, was based on a relatively small sample that is not necessarily representative. The measured mass is higher, for instance, than the estimates derived from both the updated all-galaxy sample (7.30×10^8 M$_\odot$) and ETG subsample (7.75×10^8 M$_\odot$) from Harrisberg (2020). On the other hand, the ETG estimate from McConnell & Ma (2013) at 8.33×10^8 M$_\odot$ and especially that of Saglia et al. (2016) at 9.84×10^8 M$_\odot$ fall well within the lower 3σ error of the measurement.

Figure 5.5 shows the $M_{BH} - \sigma_*$ scaling relation for, respectively, the sample of SMBH masses for all galaxy types from Bosch (2016) and subsample of ETGs from McConnell & Ma (2013). In each case they are overlaid, respectively, with the updated all-galaxy sample and updated ETG subsample of SMBH mass measurements made using cold molecular gas dynamics from Harrisberg (2020). Also included in both plots is the data point for the 1σ best-fit SMBH mass measurement for NGC 2513. In both cases, it is consistent with the scaling relation, lying well within the scatter of both the all-galaxy fit of Bosch (2016) and, in particular, of the ETG fit of McConnell & Ma (2013) (black lines). Also shown in these figures are the best-fit $M_{BH} - \sigma_*$ relations from Harrisberg (2020) for the updated all-galaxy sample and ETG subsample, respectively (dashed lines). Here as well, the best-fit mass for NGC 2513 is within the scatter of both lines, especially that for all galaxy types. Recalling, however, that this particular all-galaxy sample is itself heavily populated by ETGs (see Section 1.4.1), it is perhaps not truly representative of all-galaxy types.

In any event, this first SMBH mass measurement for NGC 2513 lies within the expected range for massive elliptical galaxies and sits relatively well on the $M_{BH} - \sigma_*$ relation.

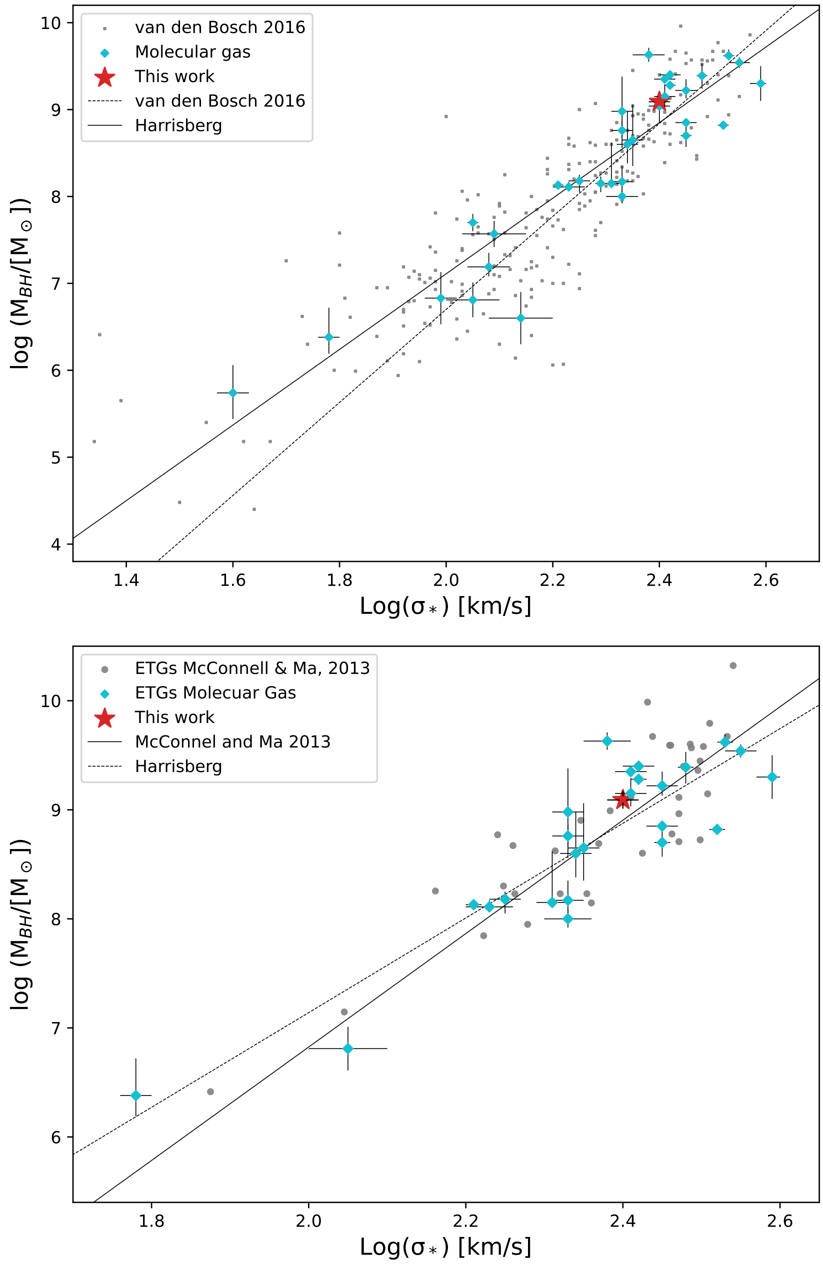

Figure 5.5.: *Top panel*: The same as for the top panel of Figure 1.4, but with the 1σ best-fit SMBH mass for NGC 2513, which is shown with the red star.
Bottom panel: The same as for the bottom panel of Figure 1.4, but with the 1σ best-fit SMBH mass for NGC 2513, which is shown with the red star.

6. Conclusion and Outlook

In this thesis, I presented the first mass measurement to be carried out of the SMBH in the massive elliptical NGC 2513. The measurement was performed via the method of cold molecular gas dynamics using high- and low-resolution ALMA observations of the ^{12}CO(2-1) line. I carried out imaging of the ALMA data to produce a three-dimensional velocity cube with CASA which, together with a luminous mass model created using the MGE method, could be fed into the kinematic modelling software KinMS.

The PVD extracted along the kinematic major axis of the galaxy shows a clear Keplerian rise in rotational velocity within the putative SMBH's gravitational sphere of influence. By modelling the kinematics of the ^{12}CO(2-1) tracer using Bayesian analysis with an MCMC procedure, it was possible to ascertain the presence of a SMBH at the centre of the galaxy and to obtain a best-fitting measurement of its mass. The modelling also provided best-fit values for the stellar M/L and other parameters describing the structure of the molecular gas disc such as inclination, position angle and gas velocity dispersion.

The best-fit SMBH mass is $(1.23^{+0.14}_{-0.23}) \times 10^9$ $M_\odot$ and the best-fit J-band mass-to-light ratio is $2.00^{+1.82}_{-1.15}$ $M_\odot/L_{\odot,J}$, both at the 68 % confidence level. At the 99 % confidence levels, the best-fit SMBH mass and J-band mass-to-light ratio are $(1.23^{+1.36}_{-0.40}) \times 10^9$ $M_\odot$ and $2.00^{+8.04}_{-3.32}$ $M_\odot/L_{\odot,J}$, respectively. With this mass, NGC 2513 lies slightly above the $M_{BH} - \sigma_*$ fits of both Bosch (2016) and Harrisberg (2020) for all galaxy types, and just above the $M_{BH} - \sigma_*$ for ETGs of both McConnell & Ma (2013) and Harrisberg (2020). In each case, however, it lies well within the scatter.

The effects of various sources of error on the modelling were considered. In particular I examined the luminous mass model for which I assumed a constant M/L and which may be affected by dust even after careful dust-masking. Scrutiny of the enclosed stellar mass as a function of radius revealed, however, that the stellar mass should not have too great an effect on the modelling within the R_{SOI}. The velocity map extracted directly from the velocity cube seems to exhibit highly circular rotation, and thus no strong deviations from purely non-circular motion due to turbulence were expected to affect the modelling. However, NGC 2513 does exhibit a relatively high gas velocity dispersion, with a best fit of $35.93^{+8.21}_{-6.45}\,\mathrm{km\,s^{-1}}$. Taken together with the galaxy's inclination, which was modelled at $60.65°^{+0.79}_{-0.37}$, this could be caused by beam smearing, which in turn may result in an over-estimation of the SMBH mass. This will be addressed in future work on the project, which is briefly described below.

Performing the measurement consisted, broadly, of three major parts. The first constituted modelling the surface brightness with the MGE method. This included an understanding of the fitting procedure, creating radial profiles of the galaxy and constructing and applying a dust mask. The resulting MGE clearly provided a good fit and reasonable estimate of the galaxy's total mass. The second was to image the ALMA data into a three-dimensional velocity cube that could be successfully modelled. The resulting maps extracted from the velocity cube indicate a resolved gas disc and features typical of gas in circular motion. The third was to perform dynamical modelling of the velocity cube with `KinMS` using Bayesian analysis to attain best-fit values for the SMBH mass, mass-to-light ratio and other free parameters of the model.

Regarding the measurement of the SMBH mass, there are several aspects that it would be useful to reassess and explore in greater depth not allowed by the scope of this thesis. The first of these is to re-image the ALMA data. In particular (and as mentioned in Section 3.3.1), this would include readjusting the weighting of the combined high- and low-resolution data. Assigning more weighting on the high-resolution data should produce a less noisy data cube with improved spatial resolution and make the modelling more effective. Other aspects in imaging include the implementation of automatic masking to improve accuracy and to prevent possible over-cleaning. Additional avenues also include experimenting with a different weighting scheme such as a different robust parameter or exploring the use of tapering.

Another important goal is to conduct my own modelling of the velocity cube with `KinMS`. In particular, I will use the MGE that I produced (which is described in Chapter 2), assess how this affects the modelling and compare the results with those from the model produced with the MGE described in Chapter 4. Regarding the parameters used in the modelling, there are several avenues to be explored. One is to compute the luminous mass model with a radially varying M/L (as in, e.g., North et al. 2019) and possibly with a spatially varying IMF (as in, e.g., Thater et al. 2023). Another is to have a radially varying velocity dispersion (as e.g. in Davis et al. (2017a)) to address the problem of beam smearing and possible over-estimation of the SMBH mass as described in Section 3.4.2.

Regarding the calculation of the likelihood function and accounting for spatial correlation in the ALMA data, there are several possibilities for constraining the uncertainties. In this thesis, this was done by scaling the rms by a factor of $(2N)^{0.25}$ (see Chapter 4). As mentioned in Section 4.2.2, another option is to produce and then include the full covariance matrix in calculating the likelihood as described in Davis et al. (2017a). Comparison of the errors produced on the estimated parameters with each respective method as in Nguyen et al. (2020) can be useful in assessing whether uncertainties have been underestimated and which method is more suited to this particular data. Concerning the MCMC procedure, an interesting aspect to explore would be, for instance, the use of Gaussian priors rather than flat priors.

The dust masking was initially attempted by constructing a radial profile of the galaxy and then excluding pixels below a fit of the profile. This method, although proven on several other galaxies (e.g., Thater et al. 2017), did not produce good results for NGC 2513. Exploring why this was the case would be useful in terms of understanding and correcting for potential problems associated with this method, especially with the aim of producing an alternative mask for this galaxy and for future work with other galaxies.

There are several other science goals related to, but outside of the SMBH measurement, that could be addressed with the data used in this thesis. The first is to perform a

Sérsic fit to the galaxy's surface brightness to obtain the key parameters for a structural classification of the galaxy. The second is to use the ALMA data to explore the structure of the molecular gas disc more thoroughly. This would include mapping it onto the HST images and examining its spatial extent in relation to the dust disc, which itself could be mapped using thermal dust emission (e.g., see Davis et al. 2017a). A third broader and long-term project would be to measure the stellar velocity dispersion of the galaxy and, in particular, to experiment with varying aperture sizes for the measurement. As mentioned in Section 1.3, discrepancies exist in the way in which stellar velocity dispersion is defined and measured. It would thus be informative to explore this parameter in greater depth.

Lastly, it would be interesting to compare any SMBH mass measurement for NGC 2513 conducted with another dynamical method, for instance with stellar dynamics. The latter might provide an opportunity to try and constrain the dark matter contribution to the potential, as stellar dynamics explores regions further out from the SMBH (see e.g., Thater et al. 2023). In any event, having a second SMBH measurement by another method would provide the possibility of comparing the two measurements and their position on the $M_{BH} - \sigma_*$ and other scaling relations. In this respect, an exploration of the large-scale correlations would also provide some interesting insights.

In summary, the results attained in this thesis provide an important contribution to the database of SMBH mass measurements, not least because this is the first SMBH mass measurement to be conducted on this galaxy. In particular, the mass obtained contributes to the measurements made with molecular gas dynamics and to the population sample of high-mass SMBHs hosted by massive ellipticals. In this respect, the best-fit mass that was obtained fulfilled the expectation of finding a high-mass SMBH typical for such galaxies. Moreover, the successful imaging and modelling of NGC 2513 convincingly tested the method of cold molecular gas dynamics for more distant galaxies lying in the (for the moment) outer ranges possible with this method. The work conducted here also promises interesting opportunities for further research that should in turn contribute to the field of galaxy structure and evolution.

Bibliography

Ann, H. B., Seo, M., & Ha, D. K. 2015, The Astrophysical Journal Supplement Series, 217, 27

Barentine, J. C. 2022, Nature Astronomy, 6, 1120

Barth, A. J., Boizelle, B. D., Darling, J., et al. 2016a, The Astrophysical Journal, 822, L28

Barth, A. J., Darling, J., Baker, A. J., et al. 2016b, The Astrophysical Journal, 823, 51

Barth, A. J., Ho, L. C., & Sargent, W. L. W. 2002, The Astronomical Journal, 124, 2607

Batcheldor, D. 2010, ApJ, 711, L108

Beifiori, A., Courteau, S., Corsini, E. M., & Zhu, Y. 2012, Monthly Notices of the Royal Astronomical Society, 419, 2497

Bendinelli, O. 1991, ApJ, 366, 599

Boehler, Y., Ménard, F., Robert, C. M. T., et al. 2021, A&A, 650, A59

Boizelle, B. D., Barth, A. J., Walsh, J. L., et al. 2019, The Astrophysical Journal, 881, 10

Boizelle, B. D., Walsh, J. L., Barth, A. J., et al. 2021, The Astrophysical Journal, 908, 19

Bosch, R. C. E. v. d. 2016, The Astrophysical Journal, 831, 134

Burkert, A. & Tremaine, S. 2010, The Astrophysical Journal, 720, 516

Cappellari, M. 2002, Monthly Notices of the Royal Astronomical Society, 333, 400

Cappellari, M. 2008, MNRAS, 390, 71

Cappellari, M. 2013a, ApJ, 778, L2

Cappellari, M. 2013b, ApJ, 778, L2

CASA Team, Bean, B., Bhatnagar, S., et al. 2022, PASP, 134, 114501

Casura, S., Ilić, D., Targaczewski, J., Rakić, N., & Liske, J. 2024, MNRAS, 534, 182

Cohn, J. H., Curliss, M., Walsh, J. L., et al. 2023, The Astrophysical Journal, 958, 186

Cohn, J. H., Walsh, J. L., Boizelle, B. D., et al. 2021, ApJ, 919, 77

Combes, F., García-Burillo, S., Audibert, A., et al. 2019, A&A, 623, A79

Cortes, P., Vlahakis, C., Hales, A., et al. 2023, ALMA Cycle 10 Technical Handbook

Crook, A. C., Huchra, J. P., Martimbeau, N., et al. 2007, ApJ, 655, 790

Davis, T. A. 2014, MNRAS, 443, 911

Davis, T. A., Alatalo, K., Bureau, M., et al. 2012, Monthly Notices of the Royal Astronomical Society, 429, 534

Davis, T. A., Bureau, M., Cappellari, M., Sarzi, M., & Blitz, L. 2013, Nature, 494, 328

Davis, T. A., Bureau, M., Onishi, K., et al. 2017a, Monthly Notices of the Royal Astronomical Society, 468, 4675

Davis, T. A., Bureau, M., Onishi, K., et al. 2017b, Monthly Notices of the Royal Astronomical Society, 473, 3818

Davis, T. A., Greene, J. E., Ma, C.-P., et al. 2019, Monthly Notices of the Royal Astronomical Society, 486, 1404

Davis, T. A. & McDermid, R. M. 2016, Monthly Notices of the Royal Astronomical Society, 464, 453

Davis, T. A., Nguyen, D. D., Seth, A. C., et al. 2020, MNRAS, 496, 4061

de Vaucouleurs, G., de Vaucouleurs, A., Corwin, Jr., H. G., et al. 1991, Third Reference Catalogue of Bright Galaxies (Springer, New York, NY (USA))

Deisenroth, M. P., Faisal, A. A., & Ong, C. S. 2020, Mathematics for Machine Learning (Cambridge University Press)

Di Matteo, T., Colberg, J., Springel, V., Hernquist, L., & Sijacki, D. 2008, ApJ, 676, 33

Dominiak, P., Bureau, M., Davis, T. A., et al. 2024, Monthly Notices of the Royal Astronomical Society, 529, 1597

Dominiak, P., Cappellari, M., Bureau, M., et al. 2024, arXiv e-prints, arXiv:2404.11260

Dressler, A. & Richstone, D. O. 1988, ApJ, 324, 701

Emsellem, E., Dejonghe, H., & Bacon, R. 1999, Monthly Notices of the Royal Astronomical Society, 303, 495

Emsellem, E., Monnet, G., & Bacon, R. 1994, A&A, 285, 723

Ferrarese, L. 2002, The Astrophysical Journal, 578, 90

Ferrarese, L. & Ford, H. 2005, Space Sci. Rev., 116, 523

Ferrarese, L., Ford, H. C., & Jaffe, W. 1996, ApJ, 470, 444

Ferrarese, L. & Merritt, D. 2000, The Astrophysical Journal, 539, L9

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, Publications of the Astronomical Society of the Pacific, 125, 306

Gebhardt, K., Bender, R., Bower, G., et al. 2000, The Astrophysical Journal, 539, L13

Ghez, A. M., Salim, S., Weinberg, N. N., et al. 2008, ApJ, 689, 1044

Gillessen, S., Eisenhauer, F., Trippe, S., et al. 2009, The Astrophysical Journal, 692, 1075

Gliozzi, M., Williams, J. K., Akylas, A., et al. 2024, MNRAS, 528, 3417

Goodman, J. & Weare, J. 2010, Communications in Applied Mathematics and Computational Science, 5, 65

Goudfrooij, P., de Jong, T., Hansen, L., & Norgaard-Nielsen, H. U. 1994a, MNRAS, 271, 833

Goudfrooij, P., Hansen, L., Jorgensen, H. E., & Norgaard-Nielsen, H. U. 1994b, A&AS, 105, 341

Goullaud, C. F., Jensen, J. B., Blakeslee, J. P., et al. 2018, ApJ, 856, 11

Graham, A. W., Erwin, P., Caon, N., & Trujillo, I. 2001, The Astrophysical Journal, 563, L11

Greene, J. E., Peng, C. Y., Kim, M., et al. 2010, The Astrophysical Journal, 721, 26

Gültekin, K., Richstone, D. O., Gebhardt, K., et al. 2009, The Astrophysical Journal, 698, 198

Harris, G. L. H. & Harris, W. E. 2011, Monthly Notices of the Royal Astronomical Society, 410, 2347

Harrisberg, S. D. 2020, Bachelor's thesis, University of Vienna, Vienna

Högbom, J. A. 1974, A&AS, 15, 417

Holtzman, J. A., Burrows, C. J., Casertano, S., et al. 1995, PASP, 107, 1065

Jarrett, T. H., Chester, T., Cutri, R., et al. 2000, The Astronomical Journal, 119, 2498

Jensen, J. B., Blakeslee, J. P., Ma, C.-P., et al. 2021, The Astrophysical Journal Supplement Series, 255, 21

Kabasares, K. M., Barth, A. J., Buote, D. A., et al. 2022, ApJ, 934, 162

Kabasares, K. M., Cohn, J. H., Barth, A. J., et al. 2024, The Astrophysical Journal, 966, 132

Kaviraj, S., Ting, Y.-S., Bureau, M., et al. 2012, MNRAS, 423, 49

Keppler, M., Teague, R., Bae, J., et al. 2019, A&A, 625, A118

Kormendy, J. 1982, Saas-Fee Advanced Course, 12, 115

Kormendy, J. & Bender, R. 2011, Nature, 469, 377

Kormendy, J. & Ho, L. C. 2013, ARA&A, 51, 511

Kormendy, J. & Richstone, D. 1995, Annual Review of Astronomy and Astrophysics, 33, 581

Krajnović, D., McDermid, R. M., Cappellari, M., & Davies, R. L. 2009, MNRAS, 399, 1839

Krist, J. E., Hook, R. N., & Stoehr, F. 2011, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 8127, Optical Modeling and Performance Predictions V, ed. M. A. Kahan, 81270J

Lelli, F., Davis, T. A., Bureau, M., et al. 2022, Monthly Notices of the Royal Astronomical Society, 516, 4066

Ma, C.-P., Greene, J. E., McConnell, N., et al. 2014, The Astrophysical Journal, 795, 158

Magorrian, J., Tremaine, S., Richstone, D., et al. 1998, AJ, 115, 2285

Makarov, D., Prugniel, P., Terekhova, N., Courtois, H., & Vauglin, I. 2014, A&A, 570, A13

McConnell, N. J. & Ma, C.-P. 2013, The Astrophysical Journal, 764, 184

McLure, R. J. & Dunlop, J. S. 2002, Monthly Notices of the Royal Astronomical Society, 331, 795

McMullin, J. P., Waters, B., Schiebel, D., Young, W., & Golap, K. 2007, in Astronomical Society of the Pacific Conference Series, Vol. 376, Astronomical Data Analysis Software and Systems XVI, ed. R. A. Shaw, F. Hill, & D. J. Bell, 127

Merritt, D. 2004, in Coevolution of Black Holes and Galaxies, ed. L. C. Ho, 263

Milvang-Jensen, B. & Jørgensen, I. 1999, Baltic Astronomy, 8, 535

Mitzkus, M., Cappellari, M., & Walcher, C. J. 2016, Monthly Notices of the Royal Astronomical Society, 464, 4789

Miyoshi, M., Moran, J., Herrnstein, J., et al. 1995, Nature, 373, 127

Monnet, G., Bacon, R., & Emsellem, E. 1992, A&A, 253, 366

Nagai, H., Onishi, K., Kawakatu, N., et al. 2019, The Astrophysical Journal, 883, 193

Nguyen, D. D., Bureau, M., Thater, S., et al. 2021, Monthly Notices of the Royal Astronomical Society, 509, 2920

Nguyen, D. D., den Brok, M., Seth, A. C., et al. 2020, The Astrophysical Journal, 892, 68

Nguyen, D. D., Ngo, H. N., Le, T. Q. T., et al. 2025, A&A, 698, L9

Nguyen, D. D., Seth, A. C., Neumayer, N., et al. 2019, The Astrophysical Journal, 872, 104

Nguyen, D. D., Seth, A. C., Neumayer, N., et al. 2018, ApJ, 858, 118

North, E. V., Davis, T. A., Bureau, M., et al. 2019, Monthly Notices of the Royal Astronomical Society, 490, 319

Onishi, K., Iguchi, S., Davis, T. A., et al. 2017, Monthly Notices of the Royal Astronomical Society, 468, 4663

Onishi, K., Iguchi, S., Sheth, K., & Kohno, K. 2015, The Astrophysical Journal, 806, 39

Peebles, P. J. E. 1972, ApJ, 178, 371

Quenneville, M. E., Blakeslee, J. P., Ma, C.-P., et al. 2023, Monthly Notices of the Royal Astronomical Society, 527, 249

Reines, A. E. & Volonteri, M. 2015, The Astrophysical Journal, 813, 82

Ruffa, I. & Davis, T. A. 2024, Galaxies, 12

Ruffa, I., Davis, T. A., Cappellari, M., et al. 2023, Monthly Notices of the Royal Astronomical Society, 522, 6170

Ruffa, I., Davis, T. A., Prandoni, I., et al. 2019, Monthly Notices of the Royal Astronomical Society, 489, 3739

Sadler, E. M. & Gerhard, O. E. 1985, MNRAS, 214, 177

Saglia, R. P., Opitsch, M., Erwin, P., et al. 2016, The Astrophysical Journal, 818, 47

Sahu, N., Graham, A. W., & Davis, B. L. 2019, ApJ, 887, 10

Sarzi, M., Rix, H.-W., Shields, J. C., et al. 2001, The Astrophysical Journal, 550, 65

Schlafly, E. F. & Finkbeiner, D. P. 2011, ApJ, 737, 103

Schroeder, D. J. 2000a, in Astronomical Optics (Second Edition), second edition edn., ed. D. J. Schroeder (San Diego: Academic Press), 240–276

Schroeder, D. J. 2000b, in Astronomical Optics (Second Edition), second edition edn., ed. D. J. Schroeder (San Diego: Academic Press), 321–351

Silk, J. & Rees, M. J. 1998, A&A, 331, L1

Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, AJ, 131, 1163

Smith, M. D., Bureau, M., Davis, T. A., et al. 2019, Monthly Notices of the Royal Astronomical Society, 485, 4359

Smith, M. D., Bureau, M., Davis, T. A., et al. 2021, Monthly Notices of the Royal Astronomical Society, 503, 5984

Thater, S. 2019, doctoralthesis, Universität Potsdam

Thater, S. 2019, in ALMA2019: Science Results and Cross-Facility Synergies, 129

Thater, S., Krajnović, D., Bourne, M. A., et al. 2017, A&A, 597, A18

Thater, S., Krajnović, D., Cappellari, M., et al. 2019, A&A, 625, A62

Thater, S., Krajnović, D., Weilbacher, P. M., et al. 2021, Monthly Notices of the Royal Astronomical Society, 509, 5416

Thater, S., Lyubenova, M., Fahrion, K., et al. 2023, A&A, 675, A18

Tomic, S., Beko, M., Camarinha-Matos, L., & Oliveira, L. 2019, Applied Sciences, 10, 272

Tran, H. D., Tsvetanov, Z., Ford, H. C., et al. 2001, AJ, 121, 2928

Tremaine, S., Gebhardt, K., Bender, R., et al. 2002, ApJ, 574, 740

van den Bosch, R. C. E., Greene, J. E., Braatz, J. A., Constantin, A., & Kuo, C.-Y. 2016, The Astrophysical Journal, 819, 11

Van Den Bosch, R. C. E. & Van De Ven, G. 2009, Monthly Notices of the Royal Astronomical Society, 398, 1117

van Dokkum, P. G. & Franx, M. 1995, AJ, 110, 2027

Veale, M., Ma, C.-P., Greene, J. E., et al. 2017, MNRAS, 471, 1428

Veale, M., Ma, C.-P., Greene, J. E., et al. 2018, MNRAS, 473, 5446

Verolme, E. K., Cappellari, M., Copin, Y., et al. 2002, MNRAS, 335, 517

Viallefond, F. 2006, in Astronomical Society of the Pacific Conference Series, Vol. 351, Astronomical Data Analysis Software and Systems XV, ed. C. Gabriel, C. Arviset, D. Ponz, & S. Enrique, 627

Volonteri, M., Natarajan, P., & Gültekin, K. 2011, The Astrophysical Journal, 737, 50

Wenger, M., Ochsenbein, F., Egret, D., et al. 2000, A&AS, 143, 9

White, R. A., Bliton, M., Bhavsar, S. P., et al. 1999, AJ, 118, 2014

Williams, D. A. 2005, Journal of Physics: Conference Series, 6, 1

Willmer, C. N. A. 2018, ApJS, 236, 47

Young, J. S., Xie, S., Tacconi, L., et al. 1995, ApJS, 98, 219

Young, L. M., Bureau, M., Davis, T. A., et al. 2011, MNRAS, 414, 940

Yu, N., Ho, L. C., Wang, J., & Li, H. 2022, The Astrophysical Journal Supplement Series, 261, 21

Zhang, H., Bureau, M., Smith, M. D., et al. 2024, Monthly Notices of the Royal Astronomical Society, 530, 3240

Zhu, P., Ho, L. C., & Gao, H. 2021, The Astrophysical Journal, 907, 6

Zocchi, A., Gieles, M., Hénault-Brunet, V., & Varri, A. L. 2016, Monthly Notices of the Royal Astronomical Society, 462, 696

A. Appendix

The following is the `tclean` script used for 140-channel velocity cube described in Chapter 3:[1]

```
vis: name of input visibility file (MS file)
imagename: name of output images
field: field/s to be imaged
spw: spectral windows and channels to be imaged
phasecenter: phase centre of the image
specmode: defines the spectral mode, e.g. for continuum imaging with one
    channel or line imaging with one or more channels
nchan: number of channels in the output image
start: starting velocity or frequency of the output cube
width: width of each channel in velocity or frequency
outframe: spectral reference frame in which 'start' and 'width' are interpreted
restfreq: rest frequency of the observed line
cell: cell size; determines the pixels in angular units, i.e. pixel to
    arcsecond scale
deconvolver: algorithm to be used in the minor cycle, e.g. Hogbom, Clark
gridder: type of gridder
imsize: size of the resulting output images in pixels
weighting: adjusts the weighting given to each visibility sample e.g., natural,
     uniform, brigss
robust: key parameter for brigss weighting with values between -2.0 (close to
    uniform weighting) and 2.0 (close to natural)
restoringbeam: sets a specific Gaussian restoring beam to construct the final
    restored image from the final model and residuals
niter: maximum number of iterations
threshold: stopping threshold given in mJy with which the peak residual in the
    residual image is compared
interactive (bool = False): set and modify masks and parameters at runtime
```

Figure A.1 shows the `tclean` script for a continuum imaging of the combined low- and high-resolution MS for NGC 2513. For continuum imaging the spectral mode *mfs* is selected. This stands for multi-frequency synthesis and it combines data from all the spectral channels that have been selected into a single image.[2] As no channels are generated, all input with regards to channel width and velocity coverage are omitted.

Figure A.2 shows a velocity (first moment) map for a 50 channel velocity cube of the combined measurement sets that were produced in the initial imaging phase of the project.

[1] Further documentation is given at: `https://casadocs.readthedocs.io/en/stable/api/tt/casatasks.imaging.tclean.html`

[2] See: `https://casaguides.nrao.edu/index.php/First_Look_at_Imaging_CASA_6.6.1`

```
CASA <1>: tclean(vis='ngc2513_combined37_continuum.ms/',
     ...: imagename='ngc2513_combined37_continuum_th08_151124_noninteractive',
     ...: field='0',
     ...: spw='',
     ...: phasecenter="ICRS 08:02:24.656880 +09.24.48.93120",
     ...: specmode='mfs',
     ...: outframe='LSRK',
     ...: restfreq='230.5380GHz',
     ...: cell='0.09arcsec',
     ...: deconvolver='hogbom',
     ...: gridder='standard',
     ...: imsize=[128, 128],
     ...: weighting='briggs',
     ...: robust=0.5,
     ...: restoringbeam='common',
     ...: threshold="0.8mJy",
     ...: niter=10000,
     ...: interactive=False)
```

Figure A.1.: Script for the continuum imaging of the combined MS for NGC 2513.

The lack of velocity coverage in regions above ~ 50 km s^{-1} and below ~ 150 km s^{-1} is apparent.

Figure A.3 shows two stages in the process of imaging the N_2H^+ line in in TW Hydra. The left panel shows the source emission in the centre with mask (white contour). As the cleaning process progresses (indicated by the green arrow), the source emission inside the mask becomes comparable to that of the residuals.[3]

Table A.1 lists the revised and updated sample of SMBH mass measurements made using molecular gas dynamics from Harrisberg (2020). The original sample contained all published measurements up until NGC 404 (Davis et al. 2020) and was assembled primarily to investigate the $M_{BH} - \sigma_*$ relation 1) for all galaxies and 2) for ETGs alone. The use of central velocity dispersion as opposed to effective velocity dispersion was also investigated where possible. The updated and revised sample includes 15 new measurements that have since been published.

[3] See also https://casaguides.nrao.edu/index.php/First_Look_at_Imaging_CASA_6.6.1

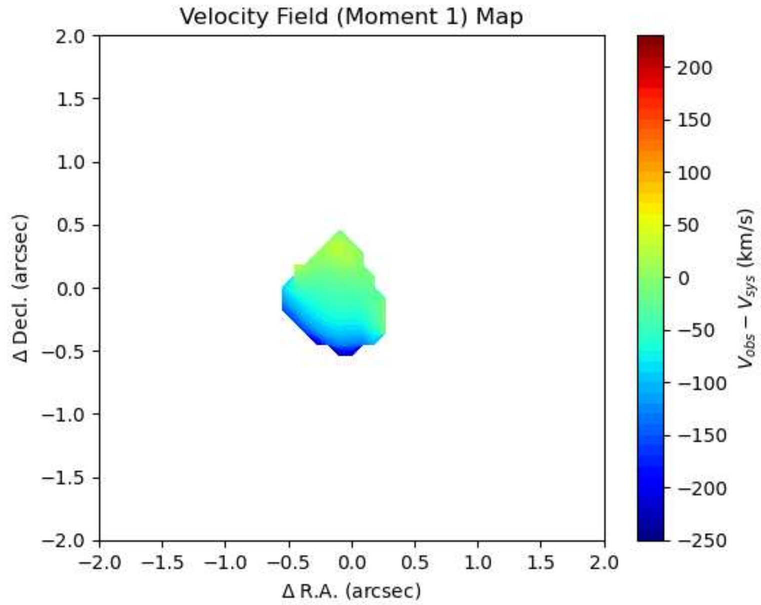

Figure A.2.: Velocity (first moment) map constructed from a 50-channel velocity cube of the ^{12}CO(2-1) emission showing lack of coverage of higher velocities.

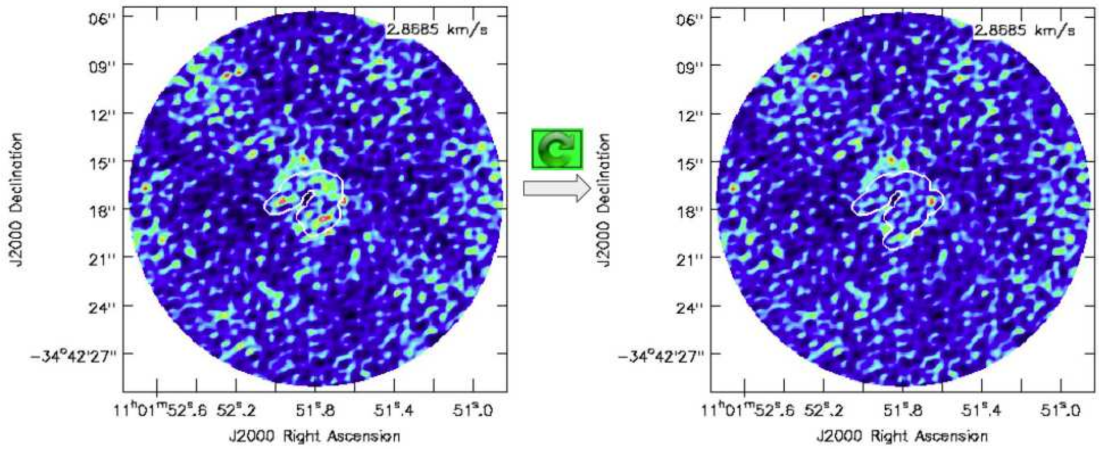

Figure A.3.: Two stages of the cleaning process with `tclean` of the N_2H^+ line in in TW Hydra. The *left panel* shows bright emission within the masked region (white line). The *right panel* shows comparable emission within the mask to the surrounding residuals. The large green arrow between the panels indicates the running of a major cycle of `tclean`. Figure from: ALMA Data Reduction Training Day, 27 Nov. 2023 (`https://www.alma-allegro.nl/wp-content/uploads/2023/10/DRT2023_imaging.pdf`. Accessed 15.07.2025 at 12:11:03.)

Table A.1.: Updated and revised sample of galaxies with SMBH mass measurements from cold molecular gas dynamics from Harrisberg (2020).

ID	Morphology	Type	Activity type	Distance Mpc	$\log_{10}(M_{BH}/M_\odot)$ $M_\odot$	$(\log_{10}(\sigma_*/\mathrm{km\ s^{-1}})/$ $\mathrm{km\ s^{-1}}$	Reference
(1)	(2)	(3)	(4)	(5)	6)	(7)	(8)
NGC 315	FR-I cD	ETG	LINER	70.0 ± 7.0	$9.32^{+0.07}_{-0.03}$	2.53 ± 0.01	Boizelle et al. (2021)
NGC 383	SA0⁻ : [a]	ETG	Radio	66.6 ± 9.9	9.63 ± 0.08	2.38 ± 0.03	North et al. (2019)
NGC 404	dE[b]	–	LINER	3.06 ± 0.37	$5.74^{+0.32}_{-0.30}$	1.60 ± 0.03[e]	Davis et al. (2020)
NGC 524	S0 + (rs)[a]	ETG	–	23.3 ± 2.3	$8.60^{+0.38}_{-0.22}$	2.34 ± 0.02	Smith et al. (2019)
NGC 613	SB(rs)bc[a]	LTG	Seyfert-	17.2 ± 1.7	7.57 ± 0.15	$2.09^{+0.06}_{-0.07}$	Combes et al. (2019)
NGC 0997	E0	ETG	–	90.4 ± 9.0	8.98 ± 0.40	2.33 ± 0.02	Dominiak et al. (2024)
NGC 1097	SB(s)b[a]	LTG	LINER	14.5 ± 1.5	$8.15^{+0.08}_{-0.10}$	2.29 ± 0.01	Onishi et al. (2015)
NGC 1275	pec[a]	–	Seyfert 1.5/LINER	62.5 ± 6.3	$9.04^{+0.12}_{-0.20}$	2.40 ± 0.02	Nagai et al. (2019)
NGC 1326	(R)SB0⁺(r)[a]	ETG	LINER	14.9 ± 1.5	6.81 ± 0.2	2.05 ± 0.05	Combes et al. (2019)
NGC 1332	S0⁻ : (s)[a]	ETG	–	22.3 ± 2.2	8.82 ± 0.04	2.52 ± 0.01[f]	Barth et al. (2016a)
NGC 1365	SB(s)b[a]	LTG	Seyfert I	17.8 ± 1.8	6.60 ± 0.3	2.145 ± 0.1	Combes et al. (2019)
NGC 1380	SA0	ETG	–	17.1 ± 1.7	8.17 ± 0.18	2.33 ± 0.01	Kabasares et al. (2022)
NGC 1566	SABbc[a]	LTG	Seyfert I	7.2 ± 0.7	6.83 ± 0.3	1.99 ± 0.03	Combes et al. (2019)
NGC 1574	S0	ETG	–	19.3 ± 1.9	8.00 ± 0.08	2.33 ± 0.03	Ruffa et al. (2023)
NGC 1672	SB(s)b[a]	LTG	Seyfert	11.4 ± 1.4	7.70 ± 0.1	2.05 ± 0.01	Combes et al. (2019)
NGC 1684	E3	ETG	–	62.8 ± 2.3	$9.15^{+0.14}_{-0.12}$	2.41 ± 0.02	Dominiak et al. (2024)
NGC 3258	E1[a]	ETG	–	31.9 ± 3.9	$9.35^{+0.04}_{-0.04}$	2.41 ± 0.02	Boizelle et al. (2019)
NGC 3504	Sbab[c]	LTG	LINER	32.4 ± 2.1	$7.2^{+0.16}_{-0.11}$	2.08 ± 0.04	Nguyen et al. (2020)
NGC 3557	E3	ETG	LERG	42.5 ± 4.3	8.85 ± 0.02	2.45 ± 0.02	Ruffa et al. (2019)
NGC 3593	SA(s)0/a	ETG	–	7.0 ± 2.0	$6.38^{0.34}_{-0.19}$	1.78 ± 0.02	Nguyen et al. (2021)
NGC 3665	S0gas[c]	ETG	Radio	34.7 ± 3.5	$8.76^{+0.11}_{-0.09}$	2.33 ± 0.02	Onishi et al. (2017)
NGC 4261	FR-I E2-3	ETG	LINER	31.1 ± 2.8	$9.22^{0.13}_{0.09}$	2.45 ± 0.02	Boizelle et al. (2021)
NGC 4429	S0a – S0/Sa[b]	ETG	LINER	16.5 ± 1.6	$8.18^{+0.07}_{-0.13}$	2.25 ± 0.02	Davis et al. (2017b)
NGC 4526	S0[b]	ETG	–	16.4 ± 1.6	$8.65^{+0.41}_{-0.30}$	2.35 ± 0.02	Davis et al. (2013)
NGC 4697	E6[a]	ETG	–	11.4 ± 1.1	$8.11^{+0.06}_{-0.02}$	2.23 ± 0.03	Davis et al. (2017a)
NGC 4751	S0/E6	ETG	–	26.9 ± 2.7	9.54 ± 0.06	2.55 ± 0.02	Dominiak et al. (2024)
NGC 4786	cD pec	ETG	–	62.1 ± 6.2	$8.70^{+0.14}_{-0.13}$	2.45 ± 0.01[g]	Kabasares et al. (2024)
NGC 5193	E pec	ETG	–	45.7 ± 4.6	$8.15^{0.47}0.04$	2.31 ± 0.02[g]	Kabasares et al. (2024)
NGC 6861	E4[d]	ETG	–	27.3 ± 2.7	9.30 ± 0.20	2.59 ± 0.01	Kabasares et al. (2022)
NGC 6958	S0[b]	ETG	–	34.3 ± 3.4	8.13 ± 0.03	2.21 ± 0.01	Thater (2019)
NGC 7052	E	ETG	Radio	69.3 ± 6.9	9.40 ± 0.05	2.42 ± 0.02	Smith et al. (2021)
PGC 11179	red nugget	ETG	–	89.0 ± 8.9	$9.28^{+0.03}_{-0.04}$	2.42 ± 0.01	Cohn et al. (2023)
UGC 2698	E	ETG	–	91 ± 9.1	$9.39^{+0.14}_{-0.15}$	2.48 ± 0.01	Cohn et al. (2021)
Fairall 49	pec	–	Seyfert Ih	86.7 ± 8.7	8.2 ± 0.33	–	Lelli et al. (2022)

Notes: (1) Galaxy ID. (2) Morphology. Values are from the source paper and references therein or NASA/IPAC Extragalactic Database (NED), except for: [a]de Vaucouleurs et al. (1991), [b]Wenger et al. (2000), [c]Ann et al. (2015) and [d]Kormendy & Ho (2013). (3) Hubble type. Here ETG is a broad classification to specify all galaxies I include in the ETG sample for the $M_{BH} - \sigma_*$ fit. (4) Activity type. References are the source paper, NED or Wenger et al. (2000). (5) Distance. Values from the source paper and/or references therein. Where no error was given, I assume 10 % (6) SMBH mass. Errors are 3σ except for NGC 3258, NGC 3557, PGC 11179, UGC 2698 and Fairall 49, which are 1σ. Errors include systematic uncertainties where available. (7) Stellar velocity dispersion. values are from the source paper and references therein except for: [e]Barth et al. (2002), [f]Kormendy & Ho (2013) and [g]Hyperleda Makarov et al. (2014). (8) Reference paper.