

Dissertation
submitted to the
Combined Faculties of Mathematics, Engineering and Natural Sciences
of the Ruperto-Carola-University of Heidelberg, Germany,
for the degree of
Doctor of Natural Sciences

Put forward by
JACOB WILLIAM ISBELL
born in Cedar Rapids, Iowa, USA

Oral examination: 9 November 2022

THE DUSTY HEARTS OF AGN:
UNDERSTANDING THE CENTRAL PARSEC THROUGH
IMAGING

REFEREES: PROF. DR. KLAUS MEISENHEIMER
PROF. DR. CORNELIS P. DULLEMOND

SUPERVISORS: PROF. DR. KLAUS MEISENHEIMER
 DR. JÖRG-UWE POTT

Abstract

Active galactic nuclei (AGN) play a crucial role in the formation and evolution of their host galaxies. Understanding the dust in the vicinity of supermassive black holes is key to understanding how AGN are fed and how they interact with their hosts. This circumnuclear dust traces both radiation-driven outflows and dense molecular gas which feeds the AGN. Circumnuclear dust structures are also thought to be responsible for determining AGN types and are a key feature of AGN unification. In this thesis, I use mid-infrared observations of circumnuclear dust to reveal its morphology and temperature. I use the new mid-infrared interferometer, MATISSE, to observe NGC 1068 and the Circinus Galaxy, two of the nearest AGN. Using these observations, I have produced the very first mid-infrared images of the circumnuclear dust. Analysis of these images in the L -, M -, and N -bands reveals the temperature and structure of the circumnuclear dust at sub-parsec resolution. In both galaxies, the images show a disk-like substructure which obscures emission from the accretion disk. Perpendicular to the disk – in the polar direction – warm dust extends for several parsecs. The temperature of this polar dust remains high ($\gtrsim 200$ K) at large distances from the accretion disk, indicating significant clumpiness of the dust. Additionally, substructures revealed in the images provide direct evidence for clumpy or filamentary polar dust. To complement the high-resolution study of these two AGN, I compiled a low-resolution catalog of 119 AGN. The reported L - and M -band fluxes indicate 44 targets for further high-resolution study with MATISSE. The $L - M$ and $M - N$ colors are compared to various radiative transfer models of the circumnuclear dust and show that models with clumpy polar dust provide the best match. Together the high- and low-resolution studies indicate a disk+wind model of the circumnuclear dust: the disk is a geometrically-thin obscuring structure, which plays a role in the Seyfert 1/2 dichotomy and is associated with the inflowing material that feeds the AGN; the wind (shown via the polar dust) is clumpy, warm, and associated with massive outflows from the AGN to the host galaxy. The combination of low-resolution mid-infrared fluxes/colors with the first sub-parsec images of the circumnuclear dust places strong constraints on both future modeling and AGN unification.

Zusammenfassung

Aktive galaktische Kerne (AGN; Englisch “active galactic nuclei”) spielen eine entscheidende Rolle bei der Entstehung und Entwicklung ihrer Galaxien. Das Verstehen des Staubs in der Umgebung supermassereicher Schwarzer Löcher ist der Schlüssel zum Verständnis, wie AGN gespeist werden und wie sie mit ihren Wirtsgalaxien interagieren. Dieser zirkumnukleare Staub zeigt sowohl strahlungsgetriebene Ausströmungen als auch dichtes molekulares Gas, das die AGN ernährt. Man nimmt an, dass zirkumnukleare Staubstrukturen auch für die Bestimmung der AGN-Typen verantwortlich sind und ein Schlüsselmerkmal für die AGN-Vereinigung darstellen. In dieser Arbeit verwende ich Beobachtungen des zirkumnuklearen Staubs im mittleren Infrarot, um dessen Morphologie und Temperatur zu ermitteln. Ich verwende das neue Mittelinfrarot-Interferometer MATISSE, um NGC 1068 und die Circinus-Galaxie, zwei der nächstgelegenen AGN, zu beobachten. Mit Hilfe dieser Beobachtungen habe ich die allerersten Mittelinfrarot-Bilder des zirkumnuklearen Staubs erstellt. Die Analyse dieser Bilder in den L -, M - und N -Bändern zeigt die Temperatur und die Struktur des zirkumnuklearen Staubs mit einer Auflösung unterhalb eines Parsec. In beiden Galaxien zeigen die Bilder eine scheibenartige Substruktur, die die Emission der Akkretionsscheibe verdeckt. Senkrecht zur Scheibe - in der polaren Richtung - erstreckt sich warmer Staub über mehrere Parsec. Die Temperatur dieses polaren Staubs bleibt auch in großer Entfernung von der Akkretionsscheibe hoch ($\gtrsim 200$ K), was auf eine erhebliche Verklumpung des Staubs hinweist. Darüber hinaus liefern die auf den Bildern erkennbaren Substrukturen direkte Hinweise auf klumpigen oder fadenförmigen polaren Staub. Um die hochauflösende Studie dieser beiden AGN zu ergänzen, habe ich einen niedrig aufgelösten Katalog von 119 AGN zusammengestellt. Die verzeichneten L - und M -Band-Flüsse qualifizieren 44 AGN für weitere hochauflösende Untersuchungen mit MATISSE. Die $L - M$ - und $M - N$ -Farben werden mit verschiedenen Strahlungstransfermodellen des zirkumnuklearen Staubs verglichen und zeigen, dass Modelle mit klumpigem polarem Staub die beste Übereinstimmung liefern. Zusammengefasst deuten die hoch- und niedrigauflösenden Studien auf ein Scheiben- und Windmodell des zirkumnuklearen Staubs hin: Die Scheibe ist eine geometrisch dünne, verschleiernde Struktur, die bei der Seyfert 1/2-Dichotomie eine Rolle spielt und mit einfließendem Material in Verbindung steht; der Wind (dargestellt durch den polaren Staub) ist klumpig, warm und mit massiven Ausströmungen aus dem AGN verbunden. Die Kombination von niedrig aufgelösten Flüssen/Farben im mittleren Infrarot mit den ersten Sub-Parsec-Bildern des zirkumnuklearen Staubs stellt sowohl für die zukünftige Modellierung als auch für die AGN-Vereinigung eine starke Einschränkung dar.

Contents

Abstract	v
Zusammenfassung	vii
1 Introduction	1
1 Active Galactic Nuclei	2
1.1 History of AGN	2
1.2 What is an AGN – Types and Classification	3
1.3 Dissecting AGN using their Spectral Energy Distributions	6
2 Circumnuclear Dust	9
2.1 MIDI Results and Paradigm Shift	9
2.2 Modeling	11
2 LM AGN Atlas – Comparisons of Torus Models with Local AGN	13
1 Introduction	14
2 Sample Selection and Observations	15
3 Data Reduction	17
3.1 Two-Gaussian Fitting	17
3.2 Flux Calibration	18
4 The Flux Catalogs	20
4.1 L- and M-band Nuclear Flux Table	21
4.2 L- and M-band Extended Flux Table	21
5 Comparison to Dust Emission Models	21
5.1 Stellar Contamination	29
5.2 Accretion Disk Spectra	31
5.3 Polar Elongation and the 3-5 Micron Bump	33
6 Estimating VLT Fluxes from <i>WISE</i>	34
6.1 Potential VLTI/MATISSE Targets	36
7 Summary and Conclusions	36
3 Mid-Infrared Interferometry	39
1 Interferometry Basics	39
1.1 Monochromatic Point Source	39
1.2 Resolving Power of the Interferometer	42
1.3 Polychromatic Sources	42
1.4 Coherence Length and Coherence Time	43
1.5 Extended Sources	43
1.6 Closure Phases	44
2 MATISSE Instrument	46
2.1 The Very Large Telescope Interferometer	46
2.2 Instrument Setup	47

2.3	Observing Procedure	48
2.4	Data Reduction and Data Products	50
3	Image Reconstruction	57
3.1	Minimization with Regularization	57
3.2	Regularization Functions	58
3.3	IRBis	61
3.4	Quantifying Image Fidelity	64
4	Thermal imaging of dust hiding the black hole in NGC 1068	65
1	Main Article	66
1.1	NGC 1068, the archetype Seyfert 2 galaxy	67
1.2	New infrared observations and images	67
1.3	Thermal modelling and dust mineralogy	68
1.4	Radio images and the hidden black hole	70
1.5	Spatial distribution of the dust	71
1.6	Conclusions	72
2	Methods	72
2.1	Description of the observations	72
2.2	Data flagging, reduction and calibration	74
2.3	Morphological Analysis of the Data	79
2.4	Radio Data	87
2.5	Spectral Energy Distribution Modeling	88
5	The dusty heart of Circinus	93
1	Introduction	94
2	MATISSE Observations	96
3	<i>N</i> -band Analysis	97
3.1	Data reduction and calibration	97
3.2	Image reconstruction	102
3.3	Measuring the dust temperature distribution	107
3.4	Comparisons to radiative transfer models	110
4	<i>L</i> - and <i>M</i> -band Analysis	112
4.1	MATISSE <i>LM</i> data reduction and calibration	112
4.2	Modeling and Image Reconstruction	113
4.3	Measuring Component Temperatures	119
5	Discussion and Interpretation	121
5.1	Central Point Source	123
5.2	Central disk	125
5.3	Polar extension	126
5.4	Overall morphology	129
6	Conclusions	130
6.1	<i>N</i> -band Results	130
6.2	<i>LM</i> -band Results	131
6.3	Final Remarks	131
6	Conclusions and Outlook	133
1	Conclusions	133
2	Outlook	136
2.1	Future MATISSE Observations and Analysis	136
2.2	The Large Binocular Telescope Interferometer	139
2.3	New and Future Observatories	140

Publications by the Author	141
References	143
Appendices	153
A Selected Epochs and Observing Conditions	153
A.1 Special Cases	153
B Calibration Strategy	153
C MATISSE <i>N</i> -band correlated fluxes	154
D MATISSE <i>N</i> -band closure phases	154
E VISIR-SAM data	157
F Dirty beam for Circinus in the <i>N</i> -band	157
G Imaging with and without ATs	157
H Image error estimates	162
I SKIRT model parameter variation	162
J MATISSE <i>LM</i> -band correlated fluxes	164
K MATISSE <i>LM</i> -band closure phases	164
Acknowledgements	173

Chapter 1

Introduction

At the heart of every galaxy lies a cosmic monster. This monster, the **supermassive black hole** (SMBH) at the center of each galaxy, drives the formation and evolution of its host. **Active galactic nuclei** (AGN) are an incredible phenomenon directly related to the SMBH. Because they are among the most energetic objects in the Universe, we can use the AGN to explore exotic physics, to understand galaxy evolution, and to search for the first galaxies to form in the Universe.

The AGN begins where the SMBH ends; an accretion disk of material (whose ultimate fate is to end up falling onto the SMBH) glows with the light of superheated material. Gas and dust surrounding the accretion disk are key components of the AGN, and are thought to serve as a fuel reservoir for the accretion disk. However, the specifics of these structures are poorly understood. Less understood still is the exact role the AGN plays in the evolution of a galaxy, though its impact is commonly thought to be significant. Yet not all galaxies have AGN (including our own Milky Way), which begs the questions of “how” and “under what circumstances” do AGN form. Or does each galaxy go through active and quiescent phases? These questions have puzzled astronomers and intrigued the public for nearly a century.

This thesis focused on the structure of gas and dust which lies on the outer edge of the accretion disk. With a thorough understanding of this structure, we can describe how the AGN is fed and furthermore how it interacts with its host galaxy – giving crucial insights into the AGN lifecycle and into galaxy evolution. The best tool to understand this structure, and in particular to understand the distribution of the dust within it, is mid-infrared interferometry.

This thesis is organized as follows:

- The remainder of Chapter 1 will briefly describe AGN, the Unified Model, and the importance and role of circumnuclear dust.
- Chapter 2 explores a statistical sample of AGN at sub-arcsecond resolution, measuring the L - and M -band fluxes of many of AGN for the first time. The $L - M$ and $M - N$ colors are used to compare the populations of several types of AGN to state-of-the-art circumnuclear dust models. It is based on [Isbell et al. \(2021\)](#).
- Chapter 3 introduces interferometry and the MATISSE instrument, and then describes the algorithms used to reduce MATISSE data and reconstruct infrared interferometric images.

- Chapter 4 focuses on the prototypical Seyfert 2, NGC 1068, using MATISSE. The first-ever mid-infrared interferometric images of the circumnuclear dust in an AGN are presented and interpreted. Here particular emphasis is put on the methodology, as it was a groundbreaking use of MATISSE in this way. It is based on [Gómez Rosas et al. \(2022, incl. J. Isbell\)](#)
- Chapter 5 explores the nearby AGN, the Circinus Galaxy, presenting and discussing the first MATISSE observations and images of its circumnuclear dust in the L -, M -, and N -bands. It is an expanded version of the work published in [Isbell et al. \(2022\)](#).
- Finally, Chapter 6 contains a summary of the work contained in this thesis and presents an outlook toward the future of studying the circumnuclear dust in AGN.

1 Active Galactic Nuclei

1.1 History of AGN

AGN have been known for less than a century, but they have risen quickly to prominence due to their impact on galaxy evolution and use in studying cosmology. They are also the most energetic objects known to astrophysics, and they serve as both particle accelerators and a means to study supermassive black holes. The following, except where noted, comes primarily from information in [Combes \(2021\)](#) and [Carroll & Ostlie \(2006\)](#).

AGN were first detected in 1908 by Edward Fath ([Fath, 1909](#)), who was measuring the emission lines of what he called “spiral nebulae.” In particular, NGC 1068 showed six bright emission lines that were remarkably different from the known spectra of stars. In 1926, Edwin Hubble claimed that these were extragalactic sources ([Hubble, 1926](#)). These “peculiar galaxies,” however, were first characterized by Carl Seyfert nearly twenty years later ([Seyfert, 1943](#)). Using optical telescopes, he found a population of galaxies which had very bright cores, sometimes brighter than the rest of the galaxy. These galaxies also had very broad (primarily hydrogen) emission lines (up to $\gtrsim 1000 \text{ km s}^{-1}$), setting them apart from a stellar population like in our Milky Way. Based on the width of their emission lines, AGN are now classified as Seyfert 1 or Seyfert 2.

As radio telescopes improved throughout the 1940s and 1950s, maps of the sky were made. The numerous and bright objects found were not like the stars observed in the optical. Many of these radio sources were extended and they were *very* distant. They were associated with other galaxies, and they showed jets and bubble-like structures called lobes. Radio galaxies tended to have strange spectra, with broad emission lines, a bright continuum, and absorption features unlike those found in other radio sources. Moreover, some of these galaxies had very bright cores, which appeared just like stars. The aptly-named “quasi-stellar radio sources” or **quasars** were claimed by Marten Schmidt to be very distant galaxies whose spectra have been redshifted by cosmological distances. Schmidt identified the first (and brightest) quasar, 3C273, as a galaxy at $z = 0.158$ or approximately 2 billion lightyears away ([Schmidt, 1963](#)). This was a radical idea, as this was the most distant object discovered at the time. The vast distances of quasars also has an important implication: they are *extremely* luminous. In addition to quasars, some radio galaxies have strange shapes. It was later found that the radio emission in these galaxies comes from jets leaving the nucleus of the galaxy and extending for kiloparsecs (kpc). It seems then, that the extremely luminous objects at the centers of these galaxies are capable of interacting with their hosts.




Fig. 1.1: Anatomy of an AGN and AGN unification diagram from Beckmann & Shrader (2012). At the center is a supermassive black hole surrounded by extremely hot in-falling material in the accretion disk. Further out, a hypothetical toroidal dusty absorber blocks a direct view toward the accretion disk for observers at certain orientations. The relative inclination of this structure is held responsible for differentiating between AGN types and is a key component of the unified scheme of AGN. However, its shape is unknown and a key question of this thesis.

1.2 What is an AGN – Types and Classification

In the decades since their discovery, AGN have been classified and reclassified, but today there are a variety of AGN sub-types which share common features and differ only in the specifics of their manifestation. In short, an AGN is an accreting supermassive black hole. More specifically, what we see is the emission arising from the superheated material spiraling towards the SMBH in an accretion disk. While every galaxy hosts an SMBH in their center, not all galaxies have AGN. It is generally agreed now that galaxies undergo quiet and active phases throughout their life, and an AGN is the active phase. What triggers these phases is still debated, but gas-rich mergers of galaxies are a possible explanation.

The accretion disk and SMBH together are commonly called the **central engine**. The emission from the central engine is the release of gravitational energy as material in the accretion disk spirals inward. The existence of black holes was shown to be inevitable in general relativity, and since then, SMBHs have been used to explain the AGN phenomenon. It becomes clear why when considering the vast energy which can be released by such an object: for a mass m coming from infinity toward a black hole with mass M , the gravitation potential energy is $E = GMm/R_s$. The Schwarzschild radius is given by $R_s = 2GM/c^2$, and therefore a particle can expend energy up to $1/2mc^2$. For a single hydrogen atom, this is a mere 7.5×10^{-4} erg, but all the material of a star like our Sun would produce 9×10^{53} erg. Even with a typical release of $0.1mc^2$ rather than $1/2mc^2$, an amount of material between $1 - 10 M_\odot$ can cause the accretion disk to shine $1000\times$ brighter than the 200 billion stars in the Milky Way.

A black hole, however, can not accrete material infinitely fast. As more material is accreted, the central engine glows brighter. But as the central engine glows brighter,

infalling material can be pushed away through radiation pressure. The luminosity at which radiation pressure balances out gravity is called the **Eddington Luminosity**. More material coming in beyond this limit will be blown away back into the host galaxy. The Eddington Luminosity for ionized hydrogen is given by $L_{\text{Edd}} = 4\pi GcMm_p/\sigma_T$, where m_p is the proton mass and σ_T is the Thomson cross-section. It is immediately apparent that this luminosity depends on the mass of the SMBH, and that more massive black holes can accrete more material and shine more brightly. Radiation pressure sets a limit on how fast black holes can grow¹, and the presence of bright quasars with $M \sim 10^9 M_\odot$ in the early Universe ($z = 7.64$) raises questions about the initial size and formation of SMBHs². Nonetheless, the Eddington luminosity is useful because it gives a physically motivated rule of thumb for what an AGN is: it must have a so-called Eddington ratio of $L_{\text{bol}}/L_{\text{Edd}} \gtrsim 10^{-5}$. The Milky Way is not an AGN by this criterion.

We now have a list of the necessary components of an AGN, but the varied outward appearances and features of observed AGN has led to a “zoo” of AGN types. The fundamental idea that these myriad objects, which look very different, are all manifestations of the same structure is called the **Unified scheme** of AGN. The Unified Model states that the primary differences between AGN are 1) their orientation relative to us, 2) their SMBH mass, and 3) their stage in the active phase. A schematic of AGN unification from [Beckmann & Shrader \(2012\)](#) is shown in Fig. 1.1. This figure shows a smooth torus, but the true shape of the dusty, absorbing structure is hotly debated and is the focus of this thesis. Of course, interactions with other galaxies and numerous other effects create apparent exceptions to this scheme, but this simple idea has been relatively successful in the last few decades in describing AGN. In the following subsections, the different AGN types will be briefly described.

1.2.1 Seyfert Galaxies

A Seyfert galaxy is characterized by its bright, featureless continuum which appears to come from a compact source, and by its emission lines. In **Seyfert 1** galaxies, the emission lines come from both allowed transitions (e.g., the hydrogen Lyman and Balmer series) and forbidden transitions (e.g., [OIII]). The allowed transition lines are especially broad ($\sim 1000 \text{ km s}^{-1}$), but the forbidden transition lines also exhibit Doppler broadening of $\sim 500 \text{ km s}^{-1}$. Even the “narrow” lines in Seyfert 1s are broad compared to inactive galaxies. In **Seyfert 2** galaxies, both forbidden and allowed lines are observed, but they are much narrower ($\sim 500 \text{ km s}^{-1}$). The broad, featureless continuum is also brighter in Seyfert 1s than in Seyfert 2s, and in Seyfert 1s this continuum tends to outshine the rest of the galaxy.

Seyfert galaxies intermediate to 1 and 2 exist, denoted 1.5, 1.8, etc. These classifications are purely spectral, and the exact classification depends on the relative strengths of the broad and narrow lines and the $\text{H}\alpha$ flux. Notably, the spectra of Seyfert 1.5s have been seen to become Seyfert 2 spectra over the span of several years. In Seyfert 1 and Seyfert 1.5 galaxies, relatively strong X-ray emission is observed which exhibits fast variability (hours to days). Seyfert 2 galaxies, on the other hand, show little X-ray emission. If one assumes that the X-ray emission is there, arising from the central engine, then the “missing flux” indicates an absorbing screen of hydrogen with column density $10^{26} - 10^{28} \text{ m}^{-2}$.

In NGC 1068, the prototypical Seyfert 2, broad emission lines were found in polarized light. The polarization indicates that the broad emission is scattered by dust and

¹It is possible to exceed the Eddington luminosity, a so-called super-Eddington phase, but this phase cannot be long-lived.

²This is the highest redshift quasar at the time of writing ([Wang et al., 2021](#)).

reflected far from the BLR, even though it is otherwise obscured from direct observations. Based on these observations, and later supported by other galaxies and other wavelengths, Antonucci & Miller (1985) proposed the Unified Scheme and the need for an obscuring dust structure based on polarized broad emission found in NGC 1068. Follow-up observations and theoretical work have indicated that this structure is geometrically thick and possibly toroidal; the so-called “dust torus” will be explored further in a later section and throughout this thesis. The study of Seyferts was, and continues to be, important for our understanding of AGN Unification.

1.2.2 Radio Galaxies

Radio astronomy was kick-started in the 1930s by Karl Jansky, who measured the radio emission toward the Milky Way center, in the constellation Sagittarius (Jansky, 1933). His legacy is undoubted (and indeed the flux density unit used throughout this thesis is named after him: 1 Jansky = 1 Jy = 10^{-26} W m⁻² Hz⁻¹), but there was little enthusiasm from the community following his discovery. Nonetheless, not long later, Grote Reber discovered the first discrete radio source other than the sun, Cygnus A (Reber, 1944). Later studies found an optical counterpart to this very large radio emission, an elliptical galaxy at $z = 0.057$. This was the first of many, and radio galaxies have since been split into two classifications similar to Seyferts. Broad line radio galaxies (BLRGs) exhibit broad and narrow emission lines, and narrow line radio galaxies (NLRGs) have spectra dominated by narrow emission lines.

Radio galaxies have a diverse range of shapes, and their radio emission can arise from the galaxy’s core, from a galactic halo, from kiloparsec-scale jets, and from radio lobes far from the galaxy. These lobes are gigantic; they are sometimes ~ 20 kpc in diameter and have been found 750 kpc from the host galaxy. Very often, jets connect the galaxy’s core to the radio lobes. Jets are typically several kpc to 50 kpc long. While jets are often symmetric (two-sided) about the galaxy, in the brightest radio galaxies the jets are one-sided. The reasons for this are still debated. Based on the relative brightness of the core and of the radio lobes, these galaxies are sometimes further classified into Fanaroff-Riley class I (FRI; core-dominated emission) and FR class II (FRII; lobe-dominated emission). Radio emission, especially in the jets and lobes, is dominated by synchrotron emission which indicates strong magnetic fields. Synchrotron emission also has a nearly constant spectral slope, and is therefore relatively easy to separate from the thermal emission of the AGN.

Unlike Seyfert galaxies which are typically spiral galaxies, radio galaxies are typically ellipticals, and indeed none of the strongest radio galaxies are spirals. These elliptical galaxies typically host SMBHs with very high mass ($\gtrsim 10^9 M_\odot$) which drives up the Eddington luminosity and causes the average Eddington ratio to be low. Radio galaxies tend to occur in a specific Eddington ratio range (Ho, 2002; Sikora et al., 2007): if ϵ_{Edd} is too low the nucleus is inactive; if it is too high, radiation pressure and related effects become dominant. Seyfert galaxies exhibit very little to no radio emission. Additionally, Seyfert galaxies are $\sim 100\times$ more abundant in the local Universe than radio galaxies. However, the radio galaxy, Centaurus A (Cen A), is one of the closest AGN and is only 3 Mpc away. A census of AGN indicates that about 10% of AGN are radio-loud, and the other 90% are radio-quiet.

1.2.3 Quasars and Quasi-Stellar Objects

As mentioned above, quasars (or quasi-stellar radio sources) were discovered in 1963 by Maarten Schmidt. These distant objects are extremely luminous, with the core of the galaxy dominating the emission so much that the galaxies appear star-like. Radio-loud

quasars have both core and jet emission, and the jets show a similar range in morphology to the general radio galaxy population. Quasars, however, have much higher bolometric luminosities: $L_{\text{bol}} \sim 5 \times 10^{39}$ W is typical, which is $10^5 \times$ brighter than the Milky Way. Like the general population of AGN, roughly 90% of quasars are radio-quiet. In this case, the objects are technically called “quasi-stellar objects” (QSOs), but the term quasar is often used as a catch-all.

Quasars exhibit a continuous spectrum which spans 15 orders of magnitude in frequency. At the low frequency end, the spectrum falls off abruptly as it is dominated by synchrotron emission. At higher, infrared frequencies, there is a turnover due to dust. Quasars exhibit an ultraviolet (UV) excess, and between 10^{14} and 10^{16} Hz a feature called the “blue bump” is found. This UV emission comes from the extremely hot accretion disk. Like Seyfert 1s and BLRGs, quasars exhibit very broad ($\sim 10^4$ km s $^{-1}$) emission lines (particularly hydrogen lines) along with narrow emission lines. Important for cosmological studies, quasar spectra also exhibit absorption lines from e.g., the Lyman series and metals such as C IV and Mg II. These lines are used to diagnose the interstellar/intergalactic medium between the distant quasar and the Milky Way, and to study the properties of the Universe at different epochs.

1.2.4 Blazars

Blazars are characterized by their high variability and large degree of polarization. For instance, blazars can vary by $\sim 30\%$ in 24 hours. The archetypical blazar, BL Lacertae (BL Lac), exhibited a factor of 15 change over the span of several months. Blazar spectra show a nearly featureless continuum, with very weak lines, both broad and narrow. Because the emission is dominated by synchrotron emission and is highly polarized, it is thought that blazars are almost exactly face-on AGN, and our line of sight is essentially along a radio jet. They are sometimes called Type 0 AGN.

1.2.5 Emission Line Galaxies

The final AGN classification is emission line galaxies. These galaxies have historically been classified as AGN due to their broad emission lines, but mounting evidence shows this is dominated by star forming regions (HII regions) rather than SMBH accretion. Low Ionization Nuclear Emission Regions (or LINERs) are associated with AGN having low nuclear luminosities. In these regions, there are strong emission lines from low-ionization species, such as the forbidden transitions [OI] and [NII]. The spectra of LINERs are similar to those from low luminosity Seyfert 2s. For this reason, it is thought that LINERs could be the low-luminosity limit of AGN. Furthermore, LINER emission is found in the centers of many spiral galaxies, given high enough sensitivity. On the other hand, LINER emission is also present in HII regions and starbursts, so it is unclear if LINER emission is truly an AGN phenomenon. In the local Universe there are also “starburst” galaxies or HII galaxies which are both dominated by star formation rather than SMBH accretion. (Ultra)luminous Infrared Galaxies, or (U)LIRGs are now thought to be the high- z equivalent to local starbursts. Despite not being AGN based on the Eddington ratio criterion, these regions are often included in AGN surveys (e.g., in Ch. 2) due to historical or purely spectral classifications.

1.3 Dissecting AGN using their Spectral Energy Distributions

The spectral energy distributions (SEDs) of AGN give insights into their anatomy. The spectra (νF_ν) of AGN are relatively flat, with significant emission at all wavelengths from the radio to γ -rays. A labeled example is shown in Fig. 1.2. The following describes




Fig. 1.2: Schematic representation of an AGN SED with important features in various wavelength regimes labeled. The black line is the observed flux. Figure by F. Shankar and presented in Monaco (2018).

the phenomena inferred from the SED from high to low energy, which essentially start from the SMBH and progress outward to larger scales.

Starting at short wavelengths, X-rays and γ -rays make up roughly 10% of the bolometric luminosities of AGN. They exhibit variations on the scales of days and weeks, indicating that they arise from very small scales, likely near the SMBH. The X-ray component of AGN has a continuum which exhibits a power-law spectrum with the form $F_\nu \propto \nu^\alpha$ and extends to TeV photon energies (Horan & Weekes, 2004). The typical spectral slope of $\alpha \sim 0.8$ indicates the emission is non-thermal (e.g., Tozzi et al., 2006), and instead arises from inverse Compton scattering in a corona near the accretion disk. There is also a “soft excess” of X-rays below 1 keV (Arnaud et al., 1985) and a “reflection hump” at ~ 20 keV (George & Fabian, 1991).

In the ultraviolet (UV) the SED shows the **blue bump**. This bump, which extends from 10 to 400 nm, is thermal emission consistent with temperatures between 10^4 and 10^5 K. Such high temperatures indicate a powerful origin, and the favored explanation is an accretion disk of material around the SMBH (initially by e.g., Salpeter, 1964; Shields, 1978, and now widely accepted). In the UV and optical bands, broad ($\gtrsim 1000 \text{ km s}^{-1}$) emission lines are observed in Seyfert 1, quasars, and BLRGs. Short temporal variations of these lines indicate that they arise from a compact region called the **broad line region** (BLR) which is ~ 1 light day (≈ 0.0001 pc) across. In all AGN types, narrow ($\lesssim 500 \text{ km s}^{-1}$) emission are observed in these same bands; these lines are found slightly farther from the AGN in the **narrow line region** (NLR).

In the infrared, from 2 to 300 μm , a large peak due to thermal emission is observed. This peak is separated from the blue bump by a dip in the SED at $\sim 1 \mu\text{m}$ and is called the **infrared bump**. The infrared bump is consistent with thermal dust emission at $T \lesssim 1500$ K (Barvainis, 1987). The dust temperature peaks at this value due to sublimation; dust heated further will be broken apart. It is clear from the amount of thermal dust emission that a relatively large ($\gtrsim 1$ pc) structure containing dust resides near the SMBH. This structure is commonly called the “dusty torus” due to

early attempts at explaining AGN Unification with a toroidal structure, but throughout this thesis it will be referred to as **circumnuclear dust**, as our understanding of its morphology has evolved (partially due to this work). Dust emission in AGN extends to the submillimeter.

Finally, AGN can be either radio loud or radio quiet, so here the SEDs vary. When there is radio emission, it is found to be largely non-thermal. Instead it is a power law with $\alpha \approx -0.5$, which is consistent with **synchrotron emission** from the jets. There is a small contribution in the radio from **free-free emission** which is caused by ionized gas at $T \sim 10^4$ K.

Detailed study of the SEDs and emission lines of AGN has revealed a few consistent features which tell us the components of an AGN. They are given from smallest to largest below. Scales are given both in terms of the SMBH mass via the Schwarzschild radius $R_S = 2GM_{\text{BH}}/c^2$ and in physical scales. Starting from the center of the AGN we find

- A supermassive black hole with mass $M_{\text{BH}} = 10^6 - 10^{10} M_{\odot}$. While it has no contribution to the SED, its gravitational pull is necessary for the energy released by the accretion disk. The SMBH extends by definition to $1R_S = 10^{-3}$ pc.
- An accretion disk surrounding the SMBH. It starts at $1R_S$ and extends to $\ll 100R_S = 0.1$ pc. The accretion disk consists of optically thick plasma which is superheated as it falls toward the SMBH to temperatures $10^4 - 10^6$ K. The soft X-rays and UV radiation seen in the SEDs come from this material. The accretion disk ultimately is responsible for all of the AGN emission, as the dust emission is re-radiated light which originated here.
- A broad line region (BLR), consisting of dense clouds which orbit the SMBH at very high velocity ($\sim 10^4$ km s $^{-1}$; causing strong Doppler broadening of the lines). This region is also $\sim 10^4$ K and reverberation mapping shows that it is light days ($\sim 0.001 - 0.1$ pc $\sim 100R_S$) across. The BLR is not observed in NLRG or Seyfert 2 AGN. Spectral classification of AGN depends on the detection or non-detection of the broad emission lines which originate here.
- A structure composed of circumnuclear dust (AKA the dusty torus) on scales 1 - 10 pc ($10^{3-4}R_S$), which surrounds the above components. Its inner radius is set by the sublimation temperature of the dust, approximately 1500 K. The dust is heated by the accretion disk's radiation. Its orientation relative to us is thought to explain the differences between Seyfert 1 and Seyfert 2 AGN. The thermal emission from this structure spans the near- and mid-infrared and causes the so-called infrared bump.
- A narrow line region (NLR) in which gas clouds orbit the SMBH at a distance of 10-100 pc = $10^{4-5}R_S$ and an extended narrow line region (ENLR) out to 1000 pc = 10^6R_S . The (E)NLR extends above and below the circumnuclear dust and is visible in all AGN types. The gas here is photoionized by the accretion disk's radiation. The photoionized gas is observed via forbidden transitions (e.g., [OIII]) and the emission lines have typical widths of ~ 500 km s $^{-1}$.

In the roughly 10% of AGN which are radio bright, there are radio jets which span a vast range in scales, from 0.01 to 10^6 pc, and which emit synchrotron radiation. At the ends of these jets there can be radio lobes hundreds of kiloparsecs across and found far from the host galaxy.

2 Circumnuclear Dust

Due to their vast gravitational pulls and energy outputs, AGN play a key role in galaxy evolution. AGN feedback, via radiation pressure, jets, and winds, can regulated star formation in the host galaxy. State-of-the-art cosmological simulations include this feedback as a key component in order to match observations at $z = 0$ (e.g., [Pillepich et al., 2018](#)). Understanding the dust in the vicinity of supermassive black holes is key to understanding how AGN are fed and how they interact with their hosts because the dust traces dense molecular gas which feeds the AGN. Additionally, large, obscuring dusty structures are thought to be responsible for both funneling material toward the central engine, and for distinguishing between Seyfert 1 and Seyfert 2 AGN. In order to fully understand the accretion process and the life cycle of an AGN, one must understand the parsec-scale dust structures surrounding it.

The previous sections have shown that the circumnuclear dust is a key component of the SED and a necessary part of AGN unification schemes. Circumnuclear dust is, moreover, the primary focus of this thesis work, and so it is described in more detail below. The information below presents a broad overview of the discovery and conceptualization of the circumnuclear dust, and so it is by no means comprehensive. The review articles by [Antonucci \(1993\)](#) and more recently [Hönig \(2019\)](#) are the basis of this text, and the curious reader is encouraged to peruse those articles for more details.

The concept of an obscuring dust structure which hides a Seyfert 1 nucleus and gives an AGN the appearance of a Seyfert 2 was first proposed by [Antonucci & Miller \(1985\)](#). Those authors found, in polarized light, broad emission lines as from a Seyfert 1 in the prototypical Seyfert 2, NGC 1068. In non-polarized light, the broad lines are completely obscured, so the medium must be optically thick. The degree of polarization ($\approx 16\%$) was explained by an optically thick dusty medium which scattered and polarized the light from the BLR. [Antonucci & Miller \(1985\)](#) then speculated that this structure could be present in all Seyfert 2 AGN, and in each one a Seyfert 1 nucleus is simply obscured; in essence, Seyferts 1 and 2 are all the same objects, but it is our viewing angle which determines the classification. The authors proposed a toroidal structure for simplicity, and the idea of a “dusty torus” stuck despite later evidence of a more complex morphology.

The presence of a circumnuclear structure, made up of dusty molecular clouds gained traction in the following years. Nearly all Seyfert 2 and 60% of Seyfert 1 AGN showed thermal emission, easily interpreted as light from the accretion disk being re-emitted by warm dust ([Miley et al., 1985](#)). A theoretical explanation for a geometrically and optically thick torus of molecular clouds was put forth by [Krolik & Begelman \(1988\)](#). The Unified Scheme of AGN, as it was beginning to be called (e.g., [Antonucci, 1993](#); [Urry & Padovani, 1995](#)) seemed to work remarkably well, but the nature of the obscuring structure – the key component of the scheme – remained speculative. In order to measure the morphology of the warm dust directly, at the parsec scale, astronomers turned to infrared interferometry, and in particular the MIDI instrument.

2.1 MIDI Results and Paradigm Shift

Infrared interferometry allowed astronomers to probe the circumnuclear dust emission for the first time, whereas previous studies relied only on absorption to probe the structures. Though only a few dozen AGN are near and bright enough to be studied in this way, studies using these methods turned the concept of a “dust torus” on its head. Returning to NGC 1068 (Seyfert 2, 14.4 Mpc away), studies using speckle interferometry ([Weigelt et al., 2004](#)) and N -band long-baseline interferometry with MIDI ([Jaffe et al., 2004](#); [López-Gonzaga et al., 2014](#)) showed that in addition to an unresolved compo-

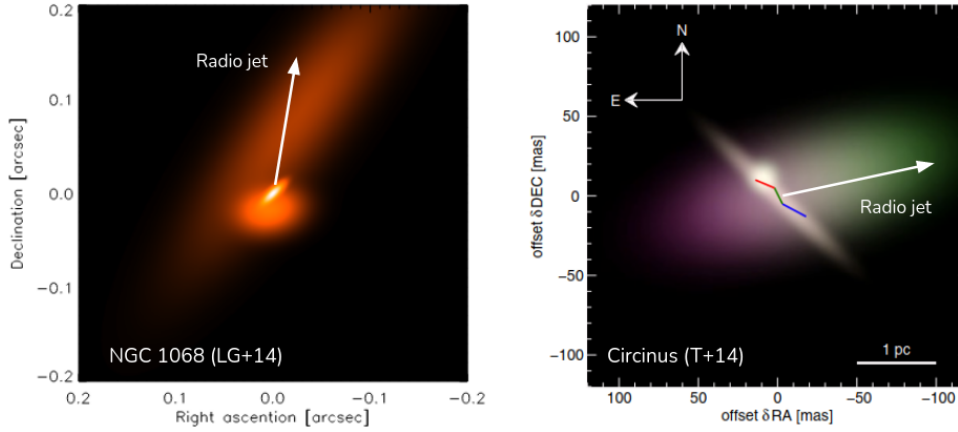


Fig. 1.3: Polar dust in NGC 1068 and Circinus as revealed by MIDI. (*left*) Model fit at $12\ \mu\text{m}$ of the circumnuclear dust in NGC 1068 by López-Gonzaga et al. (2014). (*right*) N -band false-color model fit of the circumnuclear dust Circinus by Tristram et al. (2014). Both figures have been modified to show the radio jet direction (approximately the polar direction).

ment, a significant percentage of the dust emission was 1) resolved at the parsec scale and 2) extended roughly along the jet direction (i.e. the polar direction). Similarly in the Circinus Galaxy (a “true” Seyfert 2, 4.2 Mpc away), studies with MIDI showed an unresolved component accompanied by a disk-like component and warm polar dust extending several pc (Tristram et al., 2007, 2014). The thermal dust emission measured with MIDI from NGC 1068 and from Circinus is shown in Fig. 1.3.

At around this time, the floodgates opened. Hönig et al. (2012); Hönig et al. (2013) found polar extended dust in NGC 424 and NGC 3783, respectively. Using MIDI, Burtscher et al. (2013) showed that in 23 nearby Seyferts 1 and 2 a two-component (resolved + unresolved) structure was the norm, and significant elongations were found in eight sources. López-Gonzaga et al. (2016) similarly showed that in 23 Seyferts, five had polar-elongated dust. At lower resolution, Asmus et al. (2016) found polar-extended dust out to 100 pc from the nucleus in 21 nearby AGN. It was furthermore shown by several authors (López-Gonzaga et al., 2016; Stalevski et al., 2017; Leftley et al., 2018; Stalevski et al., 2019) that the polar emission was preferentially aligned with the *edge* of the ionization cone. It was at during this dynamic decade that authors such as López-Gonzaga & Jaffe (2016) were “challenging the standard model.”

At the time of writing, the dust hasn’t settled³, though a paradigm shift has taken place. Based on the mid-infrared observations with MIDI as well as molecular gas and cool dust observations done in the sub-mm (with e.g., ALMA, see Hönig, 2019, for a list of references), a two component model has emerged. The two-component model includes 1) a geometrically thin but optically thick disk and 2) hollow polar dust cones made up of individual clouds or clumps; see Fig. 1.4 for a schematic. Knowledge about specific morphology of this model, let alone a full physical description, is however limited by several factors. First, the results are based on simple Gaussian models with relatively few interferometric observations from MIDI. Second, the mm-wavelength data are at lower resolution than the MIDI data, and capture only the large-scale structure. Lastly, before GRAVITY and MATISSE, wavelengths shorter than $8\ \mu\text{m}$ were poorly explored at high angular resolution in these AGN. One of the main goals of this thesis is to measure the morphology of the circumnuclear dust structure in a model-independent way using the new instrument MATISSE.

³Pardon the pun.




Fig. 1.4: Schematic of the circumnuclear dust from Hönic (2019).

2.2 Modeling

Before the advent of infrared interferometry, the circumnuclear dust could be observed only indirectly, and so numerous theoretical models of the underlying structure were produced, radiation transfer (RT) models in particular. Krolik & Begelman (1988) proposed a model which was composed of numerous individual and optically thick clouds, and the geometrical thickness of the structure was supported by heating from e.g., supernovae or cloud-cloud collisions. This “clumpy” or filamentary structure was supported observationally in Circinus by Tristram et al. (2007). While a clumpy medium was viewed as necessary from the beginning, computational difficulty caused the first models to use smooth distributions (Pier & Krolik, 1992). Modeling then focused on modifying the dust distribution and composition (e.g., van Bemmél & Dullemond, 2003; Schartmann et al., 2005; Fritz et al., 2006). The first clumpy modeling of the circumnuclear dust was produced by Nenkova et al. (2002), and was soon followed by numerous authors with varied approaches and assumptions (e.g., Dullemond & van Bemmél, 2005; Hönic et al., 2006; Schartmann, 2007; Nenkova et al., 2008a,c).

Following the finding that polar dust is present in many of the nearest AGN, models were once again modified. State of the art RT models of the circumnuclear dust include clumpy dust, an equatorial (disk) component, and polar dust (e.g., Hönic & Kishimoto, 2017; Stalevski et al., 2017). These models vary in that some include diffuse dust between the clumps, but the broad strokes are the same. These models provide a good fit to the SEDs of nearby AGN (González-Martín et al., 2019, this work), and the models of Stalevski et al. (2017, 2019) reproduce the MIDI-observed quantities in Circinus.

Though the RT models can give insights into the morphology of the circumnuclear dust structures, they cannot alone explain how these structures form. Parallel to the RT modeling, much work has been done to hydrodynamically model the inner few parsecs. Hydrodynamical work by Wada (2012); Wada et al. (2016) showed geometrically thin disks and radiation-driven fountains which push a significant amount of the dust into the polar direction. Full radiation-hydrodynamical (RHD) modeling is difficult, but numerous authors have now shown with RHD that radiation-driven outflows are expected (Schartmann et al., 2011; Namekata et al., 2014; Williamson et al., 2019, to name a few), though whether they are polar or radial is unclear (Williamson et al., 2019). The scales and resolution of these models are comparable to what is obtainable with MATISSE, and so comparisons can now be made which inform the modeling toward a true physical

description.

Chapter 2

The Subarcsecond Mid-Infrared View of Local Active Galactic Nuclei. IV. The L- and M-band Imaging Atlas

While later chapters utilize a new instrument to study individual AGN at the sub-parsec scale, much can still be done with archival data to inform us about the nature of circumnuclear dust. This chapter, based on [Isbell et al. \(2021\)](#), uses an archival sample of 119 nearby AGN observed with VLT ISAAC to

1. Identify a sample of AGN which can be studied interferometrically with MATISSE, and
2. Use the mid-infrared colors of a statistical sample of AGN to guide future radiative transfer modeling of the circumnuclear dust.

Both aspects of this work are crucial as we move forward, as only the statistics from a MATISSE large program in combination with comparisons to state of the art circumnuclear dust models can reveal the true nature of AGN unification and the dusty structures surrounding the supermassive black hole. I was the lead author of this manuscript; I wrote all the text and created all the figures. Daniel Asmus was instrumental in the data reduction process because he provided and aided in the use of the pipeline VIPE. Marko Stalevski provided the SKIRTOR+Polar Wind models and aided in their interpretation. All co-authors (L. Burtscher, D. Asmus, J.-U. Pott, P. Couzy, M. Stalevski, V. Gámez Rosas, and K. Meisenheimer) provided helpful feedback on manuscript structure and data interpretation..

Brief Summary

We present the largest currently existing subarcsecond 3-5 μm atlas of 119 local ($z < 0.3$) active galactic nuclei (AGN). This atlas includes AGN of 5 subtypes: 22 are Seyfert 1; 5 are intermediate Seyferts; 46 are Seyfert 2; 26 are LINERs; and 20 are composites/starbursts. Each AGN was observed with VLT ISAAC in the *L*- and/or *M*-bands between 2000 and 2013. We detect at 3σ confidence 92 sources in the *L*-band and 83 sources in the *M*-band. We separate the flux into unresolved nuclear flux and resolved flux through two-Gaussian fitting. We report the nuclear flux, extended flux, apparent size, and position angle of each source, giving 3σ upper-limits for sources which are undetected. Using *WISE* *W1*- and *W2*-band photometry we derive relations predicting the nuclear *L* and

M fluxes for Sy1 and Sy2 AGN based on their *W1-W2* color and *WISE* fluxes. Lastly, we compare the measured mid-infrared colors to those predicted by dusty torus models SKIRTOR, CLUMPY, CAT3D, and CAT3D-WIND, finding best agreement with the latter. We find that models including polar winds best reproduce the 3-5 μm colors, indicating that winds are an important component of dusty torus models. We find that several AGN are bluer than models predict. We discuss several explanations for this and find that it is most plausibly stellar light contamination within the ISAAC *L*-band nuclear fluxes.

1 Introduction

Understanding the dust in the vicinity of central supermassive black holes is instrumental to understanding how active galactic nuclei (AGN) are fed and powered. Large, obscuring dusty structures are held responsible for both funneling material toward the central engine, and for distinguishing between Type 1 and Type 2 AGN. In the Unified Model of AGN (Antonucci, 1993; Urry & Padovani, 1995; Netzer, 2015), a central obscuring torus of dust is oriented such that the broad-line region of the AGN is directly visible (Type 1) or such that its observation is blocked by the torus (Type 2). The dust making up the sublimation ring (the dust closest to the AGN) is $\sim 1500\text{K}$ and is best observed in the near-infrared (NIR). The extended dust of the torus, on the other hand, is most readily observable in the thermal infrared (3-25 μm). Interferometric observations of AGN in the *H*-band (Weigelt et al., 2004), in the *K*-band (Wittkowski et al., 1998; Kishimoto et al., 2011; Gravity Collaboration et al., 2020), and in the *N*-band (e.g., Burtscher et al., 2013; López-Gonzaga et al., 2014; Tristram et al., 2014; Leftley et al., 2019) conclusively show hot ($\gtrsim 100\text{K}$) dust in the vicinity of AGN (0.1pc - 100 pc) and provide strong evidence that the torus is clumpy. Clumpy media have moreover been argued in theory as necessary to prevent the destruction of dust grains by the surrounding hot gas (Krolik & Begelman, 1988). Following Nenkova et al. (2002), a clumpy formalism has been used in many radiative transfer models of tori (e.g., Nenkova et al., 2008b; Schartmann et al., 2008; Hönig & Kishimoto, 2010; Stalevski et al., 2016), reproducing the spectral energy distributions (SEDs) and spectral features of the *N*-band particularly well. The 3-5 μm bump (see e.g., Edelson & Malkan, 1986; Kishimoto et al., 2011; Mor & Netzer, 2012; Hönig et al., 2013), however, has remained difficult to properly model. Recent modeling suggests that this feature can be explained by the inclusion of a wind-driven outflow originating at the sublimation ring and propagating orthogonal to the disk (the disk+wind model; Hönig & Kishimoto, 2017).

A large body of work using spectral energy distribution fits to local AGN (e.g., Ramos Almeida et al., 2009; Alonso-Herrero et al., 2011; Lira et al., 2013; García-González et al., 2016; García-Bernete et al., 2019; Martínez-Paredes et al., 2020) suggests that *L* and *M* observations at high sensitivity and angular resolution are required to study the physical properties of the 3-5 μm radiation bump. In fact, Lira et al. (2013) emphasize that spectral information at 5 μm is necessary to properly constrain their SED fits. This mid-infrared (MIR) bump is expected to originate from dust radiating at intermediate spatial scales: outside of the accretion disk and the hot dust sublimation zone, but still inside of any extended polar dust emission further out. In the near future, those spatial scales will be directly resolved in detail in the *L*- and *M*-bands with the new instrument Very Large Telescope Interferometer (VLTI) MATISSE, which allows for simultaneous *L*-, *M*-, and *N*-band interferometric observations, but which also requires accurate estimates of nuclear target fluxes (Lopez et al., 2014).

A primary goal of this chapter is to anticipate such future interferometric investigations of dusty AGN in the thermal infrared. We build on the SubArcSecond MidInfraRed

Atlas of Local AGN (SASMIRALA; [Asmus et al., 2014](#), hereafter A14), which presented an N - and Q -band imaging atlas of nearby AGN at subarcsecond resolution. In this work, we extend this atlas to the L - and M - bands for 119 nearby ($z < 0.3$) AGN, at a threefold increase in angular resolution compared to the N -band. We derive new spatial flux information at the seeing limit of the excellent Cerro Paranal site, and systematically explore how these fluxes relate to those measured in space with the *WISE* $W1$ and $W2$ bands. We then investigate how our measurements compare to the expectations derived from existing clumpy torus models.

The examination herein of L - and M -band fluxes in local AGN in a statistically relevant sample fulfills two goals: (*i*) direct observational evidence of the fact that LM -flux in excess of the classical hot torus radiation is a typical feature of nearby AGN; and furthermore, (*ii*) the presentation of an atlas and systematic characterization of the spatially resolved radiation properties to aide the sample selection for future, detailed interferometric imaging of that excess radiation to further understand its origin.

This chapter is structured as follows: in §2 we present the sample, discussing its selection and observation. In §3 we discuss the data reduction and present the measured fluxes. In §4 we present and describe the L - and M -band flux catalogs. In §5 we compare the MIR colors of our sample to those predicted from various clumpy torus models. In §6 we compare the L and M fluxes to those measured with *WISE* bands $W1$ and $W2$ respectively. We summarize and conclude the chapter in §7.

2 Sample Selection and Observations

The program from which the majority of sources were observed (ESO ID 290.B-5133(A); PI: Asmus) was a survey of AGN designed to complement the subarcsecond N - and Q -band AGN sample of A14 with 3 and 5 μm images. Out of the original sample of 253 objects, 59 were observed in June and July 2013 with the ISAAC (Infrared Spectrometer and Array Camera; [Moorwood et al., 1999](#)) instrument on the Very Large Telescope (VLT) before it was decommissioned. Two stars were observed as flux calibrators during each night of the observing program: HD 106965 and HD 130163.

To supplement this sample, we searched the ESO archive for all L - and/or M -band ISAAC observations of local ($z < 0.3$) active galaxies contained in A14. We focus on 8m class telescopes in the Southern Hemisphere because we need to resolve as much of the central region as possible to properly separate the AGN itself from its host galaxy. We did not include archival observations taken with VLT NaCo because they include only 9 AGN which are not part of this sample, providing a small statistical gain for a large data reduction overhead. The ISAAC archival programs contain 60 of the individual targets and were proposed with a variety of goals, but they each contain nearby, optically-classified active galaxies observed in at least one of the L - and M -bands. We include 20 AGN observed from the archival programs which were not in the A14 sample. As these are archival data, the selection of calibration stars and the frequency of their observation is inconsistent. Whenever possible, we gather the calibration sources taken with the same instrumental setup on the same night as the AGN. Several sources were repeated in different programs.

Each AGN was observed with the L' -filter ($\lambda_c = 3.78 \mu\text{m}$; hereafter referred to as the L -band), with the narrow M_{nb} -filter ($\lambda_c = 4.66 \mu\text{m}$; hereafter referred to as the M -band), or with both. We list the ISAAC programs included in this work in Table 2.1, with the principle investigator and number of targets observed in each filter.

The final sample includes 119 AGN of four broad classes, which we group based on their optical classifications as in A14:

- Seyfert 1 (Sy1): contains 1, 1.2, 1.5, 1n

Prog. ID	PI	$N_{\text{Obs.,L}}$	$N_{\text{Obs.,M}}$
65.P-0519	Krabbe	15	38
67.B-0332	Marco	20	12
70.B-0393	Lira	40	38
71.B-0379	Lira	30	36
71.B-0404	Brooks	3	4
072.B-0397	Galliano	2	2
074.B-0166	Galliano	14	0
085.B-0639	Asmus	14	9
290.B-5133	Asmus	69	69

Table 2.1: VLT ISAAC observing programs entering into this analysis. The sum of the $N_{\text{obs},X}$ columns can be larger than 119 because several sources were observed in multiple epochs.




Fig. 2.1: Histogram of the AGN optical classifications in this sample. Darkened regions indicate the number of sources detected with $\text{SNR} \geq 2$ with both the L and M filters.

- Intermediate Seyferts (Int. Sy): contains 1.8, 1.9, 1.5/2
- Seyfert 2 (Sy2): contains 1.8/2, 1.9/2, 2
- Low-ionization nuclear emission region galaxies (LINERs): contains L, L:, L/H, S3
- Composites/Starbursts (Cp): contains Cp, Cp:, H.

Optical classifications for each of the AGN come from A14 when available and are listed with the individual sources otherwise. As A14 compiled all optical classifications from the literature, there are multiple classifications for some objects (e.g., Sy 1.5/2). In Fig. 2.1 we show the distribution of AGN classes in the sample.

The final sample is then as follows: 21 are Seyfert 1; 5 are Intermediate Seyferts; 46 are Seyfert 2; 29 are LINERs; and 16 are so-called ‘Cp’ or Composites/Starbursts. Throughout this chapter we color-code these types consistently as blue, purple, red, orange, and green, respectively. While there are 119 total AGN, not all of them were observed in both bands; instead there are 95 L -band observations and 107 M -band observations. The final sample thus includes 87 AGN with measurements in both bands. There are 20 AGN included in this work which were not part of the original A14 N - and Q -band sample. Their optical classifications (from the literature) are given in Table 2.2.

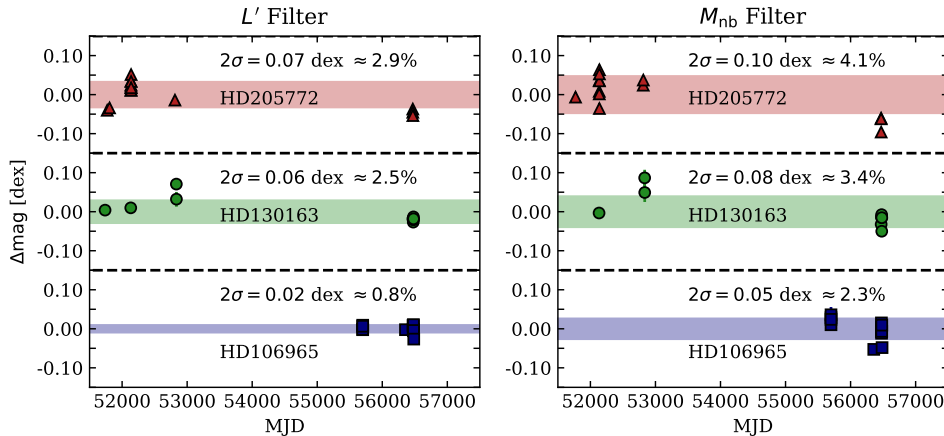


Fig. 2.2: Stability of the 3 most often observed calibrators over ~ 10 years. We show the change in brightness (Δmag) relative to the median in dex (magnitudes can be computed by multiplying these values by -2.5). On the (*left*) we show the L' filter and on the (*right*) we show the M_{nb} filter.

3 Data Reduction

The data were reduced using the custom Python tool, VISIR and ISAAC Pipeline Environment¹ (VIPE; Asmus et al., in prep.). This pipeline applies the following algorithm:

1. Combine the data at each individual chop (and nod) position
2. Combine the individual exposures of each nodding cycle into pairs
3. Combine the nodding pairs into a single exposure, taking the jitter offsets into account
4. Find the brightest source beam in the total combined image
5. Extract the positive and negative beams from the different chop/nod positions by attempting to fit every beam in every nodding pair. If this is not possible, e.g., because the source is too faint, extract and combine beams at calculated chop/nod positions instead.

In comparison to the default pipeline, this method does a better job removing the sky background in especially the M-band, resulting in several novel $\geq 2\sigma$ NIR detections (e.g., of NGC 5278). In the exemplary case of the faintest L -band source we detected, 3C321, the signal-to-noise ratio of the detection increases from $\text{SNR}_{\text{DRS}} = 1.01$ with the default pipeline to $\text{SNR}_{\text{VIPE}} = 6.71$ with VIPE.

3.1 Two-Gaussian Fitting

Our primary interest is the unresolved, nuclear flux capturing the emission of the central engine. While for the nearest AGN, we may detect extended thermal dust emission, for more distant AGN it is likely the nuclear emission also contains significant contribution of stellar light from the host. We aim to separate these two disparate components by fitting two elliptical Gaussians to each reduced image: one to represent the unresolved emission, and one to represent the extended emission. This method is quite commonly used in MIR interferometric data (e.g., Burtcher et al., 2013) to disentangle extended emission and the central engine.

¹<https://github.com/danielasmus/vipe>

As we were primarily interested in the unresolved component of each AGN, we needed to have an estimate of the point spread function (PSF) of each observation. For this, we did an initial round of fitting only the calibrators with single elliptical Gaussians. We found that the PSF size can vary by up to $\sim 10\%$ within an individual night. We can then set limits on the double-Gaussian fit; one component is set to have the major and minor axes ($\pm 10\%$) of the calibrator measured closest in time, while the second, larger component is required to have axes at least 10% larger than the central component. This accomplishes two things: 1) it effectively ignores the small amount of non-Gaussian central flux in the PSF, and 2) it wholly separates the extended and unresolved components, reducing the number of fit degeneracies and prevents the extended component from mistakenly fitting any PSF residuals. To reduce the number of total parameters, we assumed that the Gaussians are concentric.

Both the fitted Gaussians’ parameters and the error estimates are obtained through Markov-Chain Monte Carlo likelihood maximization. We sample the parameter space using the package `emcee` (Foreman-Mackey et al., 2013). The log-probability function to be maximized is given by the typical formulation

$$p(\vec{\theta}, c | \vec{x}, y, \sigma) \propto p(\vec{\theta}) p(y | \vec{x}, \sigma, \vec{\theta}, c). \quad (2.1)$$

with measurements y at positions x , parameters $\vec{\theta}$ and error estimates σ scaled by some constant c . For maximum likelihood estimation, the log likelihood function for an arbitrary model $f(x, \vec{\theta})$ can be represented as

$$\ln p(y | \vec{x}, \sigma, \vec{\theta}, c) = -\frac{1}{2} \sum_n \left[\frac{(y_n - f(x_n, \vec{\theta}))^2}{s_n^2} + \ln(2\pi s_n^2) \right], \quad (2.2)$$

where $s_n^2 = \sigma_n^2 + c^2 f(x_n, \vec{\theta})^2$, and c represents the underestimation of the variance by some fractional amount. We estimate the best-fit value as the median of each marginalized posterior distribution and the 1σ errors from the values at 16th and 84th percentiles.

Finally, we define the nuclear Gaussian flux ($F_{\text{nuc.gauss}}$) as the integrated flux of the PSF-sized, so-called “unresolved” component, and the extended Gaussian flux ($F_{\text{ext.gauss}}$) as the integrated flux of the second, larger component. In the remainder of this paper, AGN “nuclear flux” refers to $F_{\text{nuc.gauss}}$, emphasizing that for sources closer than the median distance of 45.6 Mpc, at the average fitted calibrator size of ≈ 425 mas, this area covers the central ≤ 100 pc region of the AGN.

3.2 Flux Calibration

We flux-calibrate each AGN flux measurement ($F_{\text{nuc.gauss}}$, and $F_{\text{ext.gauss}}$) with the equivalent measurement of the calibration star observed which

1. was observed closest in time to the target, to minimize changes in atmospheric transmission and seeing
2. has *either* both *L*- and *M*-band flux in van der Blik et al. (1996) or has spectral type $\in \{O, B, A, F\}$.

The spectral type selection is explained in detail in Appendix B, but in short we choose stars which have an effective temperature high enough such that the NIR color $L - M \approx 0$. This means that even when the catalog is missing a measurement in one of the two bands, the other can be reliably estimated.

In the majority of cases, calibrators were observed within 6 hr of the target but there are several nights in which no calibration source was observed. For these nights we




Fig. 2.3: Cutouts for 4 representative sources: Cen A (*top left*) has clear detections in both bands; 3C353 (*top right*) is detected in only the *L*-band; 3C317 (*bottom left*) is detected in neither band; and NGC 4594 (*bottom right*) shows extended emission. Panels in grayscale are classified as non-detections, while those in color have $\text{SNR}_{\text{gauss}} \geq 3$ in either Gaussian component and are classified as detections. With each cutout we present 1-D slices across the center of the image in both the *x*- and *y*-directions. Data in these slices are shown in black, and the profiles of the fitted elliptical Gaussians are plotted in red (nuclear), blue (extended), and magenta (sum). All images are presented with log-scaling. The ellipse in the bottom left of each cutout represents the fitted FWHM of the PSF calibrator.

estimate the long-term stability of the transfer function, by examining the flux stability of a few calibrators over many nights. Three calibrators were observed often between 2000 and 2013: HD 130163, HD 106965, and HD 205772. These three calibrators allow us to examine the stability of the measured flux over time. In Fig. 2.2, we show the flux variations of these sources in the *L*- and *M*-bands. From this, we see that in the *L*-band the 2σ flux variation is less than 3% for all sources, and is as small as 0.8% for HD 106965. We also find that the *M*-band 2σ flux variations are slightly larger, but all smaller than 5%.

For the AGN which were calibrated with these “primary” stars, we add the 2σ flux variation directly to the flux uncertainty. For those calibrated with other stars, which were often observed only once, we cannot derive a similar 2σ value. Therefore, we add 3% and 5% for the *L*- and *M*-bands, respectively; values which are slightly larger than the mean standard deviations of the 3 “primary” calibrators. Finally, for those sources which have no calibration star observed in the same night, we use whichever of the three often-observed calibrators was observed closest in time and add the 3σ uncertainty to the flux error estimate. The 3σ uncertainties are roughly 4% and 6% in the *L*- and *M*-bands, respectively. The presented flux errors combine the fitting uncertainties from both the AGN and the calibration star as well as the flux variations of the calibrators.

When sources were observed on more than one epoch, we report only the “best” measurement. We typically select the epoch with smallest seeing. The sample is quite heterogeneous, however, so a simple definition of “best” is not satisfactory. We therefore report each source with its observation date, seeing, and any VIPE reduction flags in the online version of the full article (Isbell et al., 2021)². We opt not to use the mean of the measurements, as sources were often re-observed only because of poor weather or instrumental errors. There are, however, 16 total (13 in *L*, 9 in *M* with some

²Available in machine-readable form at <https://iopscience.iop.org/article/10.3847/1538-4357/abdf3>




Fig. 2.4: L -flux vs M -flux for all sources, with non-detections given as upper limits. The colors are the same as in Fig 2.1. Sources with measured fluxes exhibiting $\text{SNR}_{\text{gauss}} \geq 3$ in both bands are filled.

overlap) sources which were observed multiple times under “good” conditions (i.e., no instrumental errors, no clouds, seeing $< 0.7''$). For these we find the measured fluxes to be quite stable: mean variations of 10.23% are found in the L band total fluxes; and mean variations of 15.63% are found in the M -band total fluxes. These values are comparable to the statistical errors derived from the calibrators and the fitting errors. In none of these sources do we see signs of significant brightening nor dimming over a 10 year span.

Additionally, we fit several targets with special conditions to extract double nuclei (e.g., Arp 220) or to locate the AGN at the center of a much brighter stellar disk (e.g., NGC 7552). We list all of these exceptional sources and the approach used to measure each in Appendix A.

3.2.1 Extended Emission

The ISAAC PSF is slightly non-Gaussian, so not all of the flux is recovered by a Gaussian fit, not even for the spatially unresolved calibrators. We find that on average $87 \pm 6\%$ of the flux measured in a $1''$ circular aperture is recovered in the L -band by such a Gaussian fit. Comparable point-source fitting in the M -band gives a similar value, $88 \pm 5\%$. This post-fit PSF residual does not ultimately affect nuclear flux calibration, however, because both the unresolved AGN flux component and the calibrators experience this in the same way.

However, the small difference between the true PSF and the Gaussian PSF approximation can be ignored when flux calibrating the extended source, since the PSF fine-structure is lost when convolving with a larger source. We therefore calibrate only the Gaussian “unresolved” fluxes with the fitted fluxes of the calibrators, while we calibrate the extended fluxes with the Jy/cts conversion derived from circular aperture ($1''$) measurements of the calibrators.

4 The Flux Catalogs

We define a detection for at least one of the fitted flux components ($F_{\text{nuc.gauss}}$ or $F_{\text{ext.gauss}}$) as having a calibrated $\text{SNR} \geq 3$. For non-detections we report only an upper limit of $3\sigma_{\text{gauss}}$. We report fluxes with $\text{SNR} \geq 2$ only when the accompanying flux component has $\text{SNR} \geq 3$; this most often affects the extended flux components.

4.1 L- and M-band Nuclear Flux Table

We present the measured nuclear L - and M -band fluxes and flux upper-limits of 119 active galaxies in Table 2.2. We include the N -band parent sample fluxes from A14 for reference. Our reported fluxes are the nuclear, unresolved component of our two-Gaussian fit. We present these fluxes separately, because the unresolved flux is most relevant for study of the AGN and interferometric follow-up observations (such as with VLTI/MATISSE).

We detect 92/95 sources in the L -band and 83/119 in the M -band. For all AGN with both L - and M -band measurements, there exist no cases where the AGN was detected in M but not in L . We show cutouts and their fitted Gaussians for a representative sample of sources in Fig. 2.3, and present all of the cutouts in the online version of the article. Fig. 2.4 shows the measured L -band flux versus the M -band flux for all sources and/or their upper limits. Here we do not see any significant difference between the different AGN classes, but rather a tight linear relation between the two bands. We detect 19/20 Sy1, 4/5 Int. Sy, 36/46 Sy2, 5/29 LINERs, and 11/16 Cp in both bands.

4.2 L- and M-band Extended Flux Table

We present the measured *extended* L - and M -band fluxes of the active galaxies in Table 2.3. In the table we also display the physical extent of the emission (from the fitted Gaussian axes), the minor/major axis ratio, and position angle (PA) of the source. We also include N -band extended emission PAs from [Asmus et al. \(2016\)](#) for reference. In the final column we give the unresolved flux to total flux ratio, f_{nuc} .

In total we find significant resolved emission accompanying 73 and 42 AGN in the L' and M bands, respectively. In the L' -band, we find that 15 Sy1, 3 Syi, 32 Sy2, 14 LINERs, and 9 Cps exhibit extended emission. In the M -band 9 Sy1, 1 Syi, 14 Sy2, 6 LINERs, and 11 Cps had extended emission. Of the detected AGN, Sy2 were the most likely to show extended emission, with 32/36 resolved.

4.2.1 A Note on Position Angles

We note that the errors on the fitted PAs are quite large, and we caution the reader that even bright AGN showed large changes in fitted PA despite having consistent fitted flux and fitted FWHM values across epochs. NGC 1068, for example, exhibited PAs of $-28 \pm 38^\circ$ and $19 \pm 9.1^\circ$ in the 2001 and 2004 epochs, respectively. The PSF PA varied by a similar amount. The unstable PSF of ISAAC leads us to suggest that the reader uses these PAs cautiously. We show L , M , and N PAs of the AGN with significant extended emission and PAs reported in [Asmus et al. \(2016\)](#) in Fig. 2.5. We include therein the optical PAs from HyperLEDA ([Makarov et al., 2014](#)) for comparison. In a few objects (3C 273, MCG-6-30-15, Circinus, Fairall 49, MGC-3-34-64, NGC 1068, and NGC 1386) there is relatively good agreement between the L , M , and N PAs. In only NGC 1386 and Superantennae S do the HyperLEDA angles match the L angle. The variation per source, however, is so great that we cannot draw any definite conclusions about the PA relations. AO-aided observations with ERIS or with ELT-METIS will be necessary to make progress in this line of inquiry.

5 Comparison to Dust Emission Models

We compare our detected AGN MIR colors to those predicted by various dusty torus models. The SKIRTOR models ([Stalevski et al., 2016](#)) are heterogeneous, consisting of high-density clumps and low-density interclump media. The CAT3D models ([Hönig &](#)

Table 2
L- and *M*-band AGN Nuclear Flux Catalog

Target Name	R.A. (J2000) (hh:mm:ss)	Decl. (J2000) (dd:mm:ss)	AGN Type	<i>L</i> Flux _{nuc} (mJy)	<i>M</i> Flux _{nuc} (mJy)	<i>N</i> Flux (mJy)
3C 273	12:29:06.70	+02:03:09.00	1	91.67 ± 7.34	77.81 ± 7.09	289.5 ± 51.1
3C 317	15:16:44.50	+07:01:18.00	2 ^a	≤0.41	≤2.35	≤2.7
3C 321	15:31:43.50	+24:04:19.00	2	0.47 ± 0.07	≤4.63	≤50.30
3C 327	16:02:27.40	+01:57:56.00	1	0.92 ± 0.05	1.27 ± 0.20	70.8 ± 21.3
3C 353	17:20:28.20	-00:58:47.00	2 ^a	1.26 ± 0.10	≤4.94	8.0 ± 2.3
3C 403	19:52:15.80	+02:30:24.00	2	5.00 ± 0.19	5.21 ± 0.43	94.2 ± 12.8
3C 424	20:48:12.00	+07:01:17.00	2	≤0.43	≤2.16	≤1.6
Arp 220E	15:34:57.07	+23:30:14.74	L ^b	≤0.29	3.27 ± 0.18	...
Arp 220W	15:34:57.07	+23:30:14.74	L ^b	13.02 ± 1.03	3.80 ± 0.48	...
Cen A	13:25:27.60	-43:01:08.99	2	261.69 ± 19.58	305.82 ± 18.28	1524.2 ± 152.4
CGCG381-051	23:48:43.90	+02:14:42.58	H	≤0.29	≤2.85	...
Circinus	14:13:09.90	-65:20:20.99	2	458.16 ± 39.18	676.41 ± 44.58	8326.9 ± 1049.2
ESO 103-35	18:38:20.30	-65:25:38.99	2	18.06 ± 1.97	40.67 ± 4.73	≤258.70
ESO 138-1	16:51:20.10	-59:14:05.00	2	50.90 ± 3.45	94.10 ± 4.80	759.9 ± 35.7
ESO 141-55	19:21:14.10	-58:40:13.00	1.2	46.85 ± 2.53	53.51 ± 4.11	148.7 ± 35.9
ESO 286-19	20:58:26.80	-42:39:00.01	Cp:	5.23 ± 0.33	16.33 ± 0.61	411.9 ± 41.2
ESO 323-32	12:53:20.30	-41:38:08.00	1.9	3.47 ± 0.12	2.04 ± 0.35	56.7 ± 7.4
ESO 323-77	13:06:26.10	-40:24:53.00	1.2	159.61 ± 9.27	119.07 ± 10.12	346.7 ± 79.4
ESO 506-27	12:38:54.60	-27:18:28.00	2	10.56 ± 0.87	23.32 ± 0.89	179.5 ± 17.9
ESO 511-30	14:19:22.40	-26:38:41.00	1	14.23 ± 0.57	13.65 ± 0.66	52.2 ± 5.2
Fairall 49	18:36:58.30	-59:24:09.00	2	52.25 ± 5.24	62.60 ± 7.74	≤213.20
Fairall 51	18:44:54.00	-62:21:53.00	1.5	36.03 ± 4.05	49.31 ± 5.57	391.4 ± 39.1
IC 3639	12:40:52.90	-36:45:21.00	2	15.39 ± 0.66	25.73 ± 1.39	386.1 ± 38.6
IC 4329A	13:49:19.30	-30:18:34.00	1.2	186.10 ± 7.07	231.37 ± 9.78	1157.7 ± 99.3
IC 4518W	14:57:41.20	-43:07:56.00	2	18.81 ± 1.41	35.44 ± 1.28	199.4 ± 31.7
IC 5063	20:52:02.30	-57:04:08.00	2	35.06 ± 2.40	58.31 ± 5.83	820.6 ± 57.8
IC 5179	22:16:09.04	-36:50:36.48	H ^b	...	≤2.19	...
IRAS 13349+2438	13:37:18.70	+24:23:03.00	1n	133.40 ± 11.12	120.27 ± 12.80	476.4 ± 63.3
IRASF00198-7926	00:21:54.21	-79:10:09.55	2	11.10 ± 1.93	25.74 ± 2.99	...
LEDA 170194	12:39:06.28	-16:10:47.09	2	3.09 ± 0.22	3.09 ± 0.25	42.4 ± 8.2
M87	12:30:49.40	+12:23:28.00	L	6.41 ± 0.24	2.69 ± 0.24	20.8 ± 2.9
MCG+2-4-25	01:20:02.66	+14:21:42.21	H ^b	...	7.34 ± 1.92	...
MCG-0-29-23	11:21:12.44	-02:59:04.40	2 ^a	4.87 ± 0.17	6.00 ± 0.48	...
MCG-3-34-64	13:22:24.50	-16:43:42.00	2 ^a	23.62 ± 1.67	39.81 ± 5.47	530.6 ± 65.4
MCG-6-30-15	13:35:53.70	-34:17:44.00	1.5	65.24 ± 4.07	56.78 ± 6.75	340.8 ± 62.9
Mrk 331	23:51:26.73	+20:35:12.51	Cp ^a	...	≤2.41	...
Mrk 463	13:56:02.97	+18:22:16.84	1.5 ^a	68.53 ± 9.56	35.56 ± 7.95	...
Mrk 509	20:44:09.70	-10:43:25.00	1.5	108.46 ± 4.81	117.10 ± 5.40	256.4 ± 29.0
Mrk 841	15:04:01.20	+10:26:16.00	1.5	14.58 ± 1.48	23.03 ± 1.46	163.3 ± 47.6
Mrk 897	21:07:45.80	+03:52:40.00	Cp	7.42 ± 0.47	≤9.10	8.2 ± 2.7
NGC 63	00:17:45.51	+11:27:02.73	H ^b	...	0.71 ± 0.10	...
NGC 253	00:47:33.10	-25:17:18.00	Cp:	78.76 ± 2.68	121.74 ± 4.86	≤1038.5
NGC 424	01:11:27.60	-38:05:00.00	2	203.46 ± 7.91	225.63 ± 8.14	736.2 ± 222.9
NGC 520	01:24:34.65	+03:47:32.23	Cp ^a	...	≤2.32	...
NGC 660	01:43:02.25	+13:38:48.11	L ^a	...	≤2.09	...
NGC 986	02:33:34.13	-39:02:42.14	H ^c	...	4.72 ± 0.53	...
NGC 1008	02:37:57.11	+02:04:12.59	L/H ^b	≤0.05	≤2.94...	...
NGC 1068	02:42:40.70	-00:00:48.00	2 ^a	1940.88 ± 241.52	2687.53 ± 374.86	≤3567.20
NGC 1097	02:46:19.00	-30:16:30.00	L	5.55 ± 0.24	5.63 ± 0.51	16.8 ± 2.3
NGC 1125	02:51:40.48	-16:39:05.55	2 ^a	3.51 ± 0.23	6.76 ± 0.46	...
NGC 1368	03:33:39.77	-05:05:24.29	2 ^a	0.67 ± 0.06
NGC 1365	03:33:36.40	-36:08:25.00	1.8	196.43 ± 6.59	152.32 ± 6.74	360.7 ± 36.1
NGC 1386	03:36:46.20	-35:59:57.01	2	18.65 ± 0.98	29.31 ± 2.29	299.3 ± 62.3
NGC 1511	03:59:32.73	-67:38:00.00	H ^b	...	2.44 ± 0.17	...
NGC 1566	04:20:00.40	-54:56:16.00	1.5	10.55 ± 0.38	...	≤30.00
NGC 1614	04:33:59.90	-08:34:44.00	Cp:	...	7.13 ± 0.64	≤345.6
NGC 1667	04:48:37.10	-06:19:12.00	2	1.48 ± 0.08	...	≤2.60
NGC 1808	05:07:42.30	-37:30:47.01	Cp:	6.30 ± 0.51	10.46 ± 0.56	328.8 ± 34.2
NGC 3125	10:06:33.34	-29:56:06.31	L/H ^b	≤0.75	≤3.47	...
NGC 3281	10:31:52.10	-34:51:13.00	2	...	116.64 ± 5.26	486.4 ± 50.5
NGC 3660	11:23:32.30	-08:39:31.00	1.8 ^a	2.99 ± 0.13	1.46 ± 0.15	≤15.30
NGC 4038/9	12:01:54.88	-18:53:05.95	1 ^a	...	2.43 ± 0.26	...
NGC 4074	12:04:29.70	+20:18:58.00	2	12.53 ± 0.84	10.04 ± 0.72	68.3 ± 19.2
NGC 4235	12:17:09.90	+07:11:30.00	1.2	14.49 ± 0.97	5.71 ± 1.15	36.0 ± 6.5

Table 2.2: *L*- and *M*-band AGN Nuclear Flux Catalog

Table 2
(Continued)

Target Name	R.A. (J2000) (hh:mm:ss)	Decl. (J2000) (dd:mm:ss)	AGN Type	L Flux _{muc} (mJy)	M Flux _{muc} (mJy)	N Flux (mJy)
NGC 4253	12:18:26.80	+29:48:45.39	1n ^a	23.99 ± 6.38	39.53 ± 2.28	...
NGC 4261	12:19:23.20	+05:49:31.00	L	0.64 ± 0.06	≤5.84	13.2 ± 3.0
NGC 4278	12:20:06.80	+29:16:51.00	L	≤0.57	≤2.66	2.5 ± 0.6
NGC 4303	12:21:54.90	+04:28:25.00	2	...	≤1.89	6.1 ± 0.8
NGC 4303	12:21:54.90	+04:28:25.00	2	1.65 ± 0.13	≤4.07	6.1 ± 0.8
NGC 4374	12:25:03.70	+12:53:13.00	2 ^a	0.32 ± 0.04	≤2.21	≤8.0
NGC 4388	12:25:46.70	+12:39:44.00	2	34.11 ± 1.23	44.11 ± 2.10	187.8 ± 32.8
NGC 4418	12:26:54.60	-00:52:39.00	2	2.03 ± 0.20	6.26 ± 0.61	1426.8 ± 167.5
NGC 4438	12:27:45.60	+13:00:32.00	L/H	4.04 ± 0.26	2.20 ± 0.43	10.3 ± 2.8
NGC 4457	12:28:59.00	+03:34:14.00	L	3.01 ± 0.22	≤4.30	6.2 ± 1.7
NGC 4472	12:29:46.80	+08:00:02.00	2/L	≤0.41	≤2.09	≤8.6
NGC 4501	12:31:59.20	+14:25:13.00	2	1.50 ± 0.16	≤0.31	3.7 ± 0.5
NGC 4507	12:35:36.60	-39:54:33.01	2	59.67 ± 5.96	62.07 ± 12.71	622.8 ± 64.3
NGC 4579	12:37:43.50	+11:49:05.00	L	16.91 ± 0.66	14.99 ± 1.02	74.6 ± 4.5
NGC 4593	12:39:39.40	-05:20:39.00	1	32.31 ± 0.40	36.15 ± 2.54	227.4 ± 37.6
NGC 4594	12:39:59.40	-11:37:23.00	L	2.74 ± 0.32	≤4.42	4.4 ± 1.4
NGC 4746	12:51:55.40	+12:04:59.00	L/H	≤0.42	≤4.66	≤10.9
NGC 4785	12:53:27.30	-48:44:57.01	2	1.11 ± 0.08	≤2.17	≤17.4
NGC 4941	13:04:13.10	-05:33:06.00	2	4.54 ± 0.29	2.01 ± 0.34	76.1 ± 8.7
NGC 4945	13:05:27.50	-49:28:06.00	Cp	≤0.86	≤2.58	22.1 ± 6.9
NGC 5135	13:25:44.10	-29:50:01.00	2	12.63 ± 0.70	19.31 ± 2.75	132.0 ± 25.8
NGC 5252	13:38:16.00	+04:32:33.00	1.9	16.17 ± 1.20	25.19 ± 1.31	68.7 ± 6.9
NGC 5363	13:56:07.20	+05:15:17.00	L	0.98 ± 0.12	≤4.72	≤1.30
NGC 5427	14:03:26.10	-06:01:51.00	2	2.95 ± 0.10	2.09 ± 0.26	≤20.0
NGC 5506	14:13:14.90	-03:12:27.00	2	343.79 ± 12.95	359.44 ± 26.41	870.8 ± 65.6
NGC 5548	14:17:59.50	+25:08:12.00	1.5	3.29 ± 0.89	≤4.99	≤77.50
NGC 5643	14:32:40.70	-44:10:27.99	2	8.56 ± 0.45	14.11 ± 1.55	254.1 ± 68.7
NGC 5728	14:42:23.90	-17:15:11.00	1.9 ^a	2.89 ± 0.22	2.76 ± 0.32	49.1 ± 7.1
NGC 5813	15:01:11.20	+01:42:07.00	L:	≤0.38	≤2.06	≤6.7
NGC 5953	15:34:32.40	+15:11:38.00	Cp	6.01 ± 0.63	≤2.29	≤29.5
NGC 5995	15:48:25.00	-13:45:28.00	1.9	60.09 ± 3.02	54.06 ± 4.99	332.4 ± 46.8
NGC 6000	15:49:49.69	-29:23:14.20	H ^a	...	2.72 ± 0.46	...
NGC 6221	16:52:46.33	-59:13:00.99	Cp	12.53 ± 0.47	11.30 ± 0.62	103.8 ± 21.1
NGC 6240N	16:52:58.92	+02:24:04.78	Cp	16.65 ± 0.56	18.47 ± 0.68	≤7.40
NGC 6300	17:16:59.50	-62:49:14.00	2	31.34 ± 1.73	49.80 ± 2.54	553.6 ± 162.1
NGC 6810	19:43:34.20	-58:39:20.00	Cp	7.10 ± 0.56	4.85 ± 0.51	44.4 ± 13.2
NGC 6814	19:42:40.60	-10:19:25.00	1.5	10.57 ± 0.40	12.04 ± 0.69	95.6 ± 23.5
NGC 6860	20:08:46.90	-61:06:01.00	1.5	35.24 ± 1.50	39.72 ± 2.47	206.1 ± 24.1
NGC 6890	20:18:18.10	-44:48:24.00	1.9 ^a	8.12 ± 0.35	9.45 ± 0.66	116.6 ± 25.6
NGC 7130	21:48:19.50	-34:57:04.00	Cp	6.91 ± 0.30	14.39 ± 0.75	104.5 ± 21.1
NGC 7172	22:02:01.90	-31:52:11.00	2	48.96 ± 1.92	61.82 ± 3.30	185.0 ± 18.6
NGC 7213	22:09:18.16	-47:07:59.05	L/H ^a	...	≤3.39	...
NGC 7314	22:35:46.20	-26:03:02.00	2 ^a	7.94 ± 0.43	18.47 ± 0.83	61.5 ± 11.8
NGC 7479	23:04:56.70	+12:19:22.00	2	12.84 ± 0.58	...	695.1 ± 95.1
NGC 7496	23:09:47.30	-43:25:41.00	Cp	6.59 ± 0.26	6.89 ± 0.64	169.5 ± 10.7
NGC 7552	23:16:10.80	-42:35:04.99	L/H	...	≤2.13	≤63.3
NGC 7582	23:18:23.50	-42:22:14.00	Cp	196.65 ± 7.49	91.64 ± 5.07	443.2 ± 79.3
NGC 7590	23:18:54.80	-42:14:21.00	2:	0.68 ± 0.05	≤1.79	≤10.6
PG 2130+099	21:32:27.80	+10:08:19.00	1.5	17.64 ± 3.15	44.28 ± 1.44	187.8 ± 20.9
PKS 1417-19	14:19:49.70	-19:28:25.00	1.5	6.20 ± 0.22	3.04 ± 0.40	≤10.10
PKS 1814-63	18:19:35.00	-63:45:48.00	2	3.71 ± 0.12	2.09 ± 0.19	27.7 ± 5.3
PKS 1932-46	19:35:56.60	-46:20:41.00	1.9	≤0.33	≤2.02	≤2.0
Superantennae S	19:31:21.47	-72:39:21.21	2	16.72 ± 0.81	26.91 ± 1.52	221.5 ± 62.8
UGC 2369 S	02:54:01.91	+14:58:17.54	Cp ^b	...	1.18 ± 0.11	...
Z 41-20	12:00:57.90	+06:48:23.00	2	2.87 ± 0.13	2.18 ± 0.34	29.3 ± 2.9

Notes. Here Flux_{muc} indicates the flux from the fitted unresolved component. N -band fluxes and most AGN classifications are from A14 except where otherwise noted.

^a Véron-Cetty & Véron (2010).

^b Yuan et al. (2010).

^c Hernán-Caballero & Hatziminaoglou (2011).

^d Ho et al. (1997).

^e Hameed & Devereux (1999).

(This table is available in machine-readable form.)

Table 2.2: continued. The machine readable form can be found at <https://iopscience.iop.org/article/10.3847/1538-4357/abdfd3>

Table 3
L-band and *M*-band AGN Extended Flux Catalog

Target Name	L Flux _{ext} (mJy)	$\Theta_{L, ext}$ (pc)	$PA_{L, ext}$ (deg)	r_L $\equiv \theta/\Theta$	M Flux _{ext} (mJy)	$\Theta_{M, ext}$ (pc)	$PA_{M, ext}$ (deg)	r_M $\equiv \theta/\Theta$	PA_M (deg)	f_{inc} L (M)
3C 273	69.0 ± 28.9	3160.2 ± 419.9	28.8 ± 48.2	0.9	73.6 ± 22.4	3160.2 ± 292.4	65.5 ± 32.2	0.9	114.0 ± 34.0	0.57 (0.51)
3C 317	≤0.4	≤2.1
3C 321	≤0.7	≤3.9	89.0 ± 0.0	...
3C 327	≤0.4	≤2.8	97.0 ± 0.0	0.68 (...)
3C 353	≤0.7	≤4.7	0.66 (...)
3C 403	≤0.7	≤4.0	60.0 ± 0.0	0.88 (0.59)
3C 424	≤0.4	≤2.3
3C 424	≤1.5	≤2.4
Arp 220E	2.2 ± 0.3	7.4 ± 0.2	83.8 ± 1.6	0.6	≤2.4
Arp 220W	168.3 ± 80.0	11.6 ± 1.6	66.2 ± 58.6	0.9	182.1 ± 55.1	11.6 ± 1.1	5.0 ± 66.2	0.9	74.0 ± 42.0	0.85 (0.77)
CGCG381-051	≤0.3	≤2.3
Circinus	≤1.7	591.1 ± 207.2	26.1 ± 3.2	40.4 ± 32.6	0.9	100.0 ± 10.0	1.00 (0.53)
ESO 103-35	22.9 ± 7.2	350.3 ± 29.2	-37.1 ± 42.1	0.9	≤5.3	109.0 ± 12.0	0.44 (0.91)
ESO 138-1	51.2 ± 16.2	180.6 ± 19.6	9.3 ± 31.2	0.9	39.3 ± 17.8	180.6 ± 24.9	-3.7 ± 50.6	0.9	121.0 ± 21.0	0.50 (0.70)
ESO 141-55	37.2 ± 9.1	1062.2 ± 92.8	91.3 ± 62.1	0.9	≤5.9	120.0 ± 47.0	0.56 (0.93)
ESO 286-19	≤0.6	≤4.3	0.90 (0.80)
ESO 323-32	3.3 ± 0.8	539.9 ± 54.6	-6.8 ± 9.0	0.7	≤4.1	76.0 ± 33.0	0.38 (...)
ESO 323-77	≤0.8	114.5 ± 39.3	221.1 ± 24.9	85.9 ± 21.2	0.9	95.0 ± 27.0	1.00 (0.51)
ESO 506-27	9.0 ± 4.0	658.8 ± 83.2	-1.3 ± 71.9	0.9	≤4.6	93.0 ± 8.0	0.54 (0.84)
ESO 511-30	≤0.6	≤4.4	82.0 ± 3.0	0.96 (0.77)
Fairall 49	75.8 ± 21.5	587.7 ± 46.2	-21.4 ± 55.9	0.9	87.7 ± 27.7	587.7 ± 47.0	-6.3 ± 45.1	0.9	42.0 ± 19.0	0.41 (0.41)
Fairall 51	53.5 ± 15.9	399.6 ± 33.1	-26.2 ± 42.0	0.9	≤7.4	180.0 ± 0.0	0.40 (0.90)
IC 3639	11.0 ± 2.6	253.0 ± 24.9	75.4 ± 50.9	0.9	9.0 ± 4.3	253.0 ± 31.2	-9.9 ± 48.8	0.9	105.0 ± 32.0	0.58 (0.74)
IC 4329A	60.4 ± 28.6	573.3 ± 92.0	4.0 ± 66.6	0.8	≤9.1	65.0 ± 9.0	0.76 (0.98)
IC 4518W	10.9 ± 4.5	440.9 ± 54.9	78.4 ± 47.5	0.9	≤4.6	0.63 (0.89)
IC 5063	42.4 ± 5.4	203.3 ± 30.9	14.3 ± 47.0	0.9	≤55.9	108.0 ± 5.0	0.45 (...)
IC 5179	≤2.0
IRAS 13349+2438	≤0.8	110.7 ± 45.1	2641.6 ± 275.1	59.6 ± 30.7	0.9	16.0 ± 19.0	1.00 (0.52)
IRAS F0198-7926	22.9 ± 5.8	4.8 ± 0.4	24.0 ± 9.6	0.8	25.1 ± 11.5	4.8 ± 0.5	15.0 ± 32.1	0.9	...	0.33 (0.50)
LEDA 170194	1.6 ± 0.8	642.1 ± 89.5	2.5 ± 33.2	0.8	≤2.7	44.0 ± 18.0	0.66 (0.54)
M87	3.2 ± 1.4	61.3 ± 8.4	72.1 ± 47.7	0.9	≤4.5	0.56 (...)
MCG+2-4-25	≤2.0	- (0.80)
MCG-0-29-23	13.5 ± 1.4	10.5 ± 0.6	133.6 ± 3.7	0.6	≤2.0	0.27 (0.77)
MCG-3-34-64	33.9 ± 8.9	442.9 ± 33.4	106.4 ± 41.6	0.9	56.0 ± 22.3	442.9 ± 45.0	120.5 ± 29.0	0.9	51.0 ± 8.0	0.41 (0.43)
MCG-6-30-15	38.0 ± 18.4	115.4 ± 18.1	100.0 ± 44.3	0.9	53.8 ± 25.9	115.4 ± 16.2	83.4 ± 13.1	0.8	106.0 ± 13.0	0.63 (0.52)
Mrk 509	≤0.6	≤4.3	107.0 ± 24.0	0.99 (0.97)
Mrk 841	9.9 ± 4.4	598.7 ± 71.2	11.0 ± 53.8	0.9	≤3.8	155.0 ± 31.0	0.60 (0.87)
Mrk 897	3.5 ± 1.6	677.7 ± 94.8	-7.3 ± 52.1	0.9	0.68 (...)
Mrk 331	≤2.1
Mrk 463	101.4 ± 43.8	5.5 ± 0.7	39.0 ± 27.7	0.9	126.4 ± 27.9	5.5 ± 0.4	47.5 ± 12.5	0.9	...	0.41 (0.21)
NGC 6300	≤0.7	≤6.4	112.0 ± 4.0	...
NGC 253	82.9 ± 10.7	42.8 ± 3.7	130.4 ± 3.0	0.4	88.9 ± 27.8	42.8 ± 4.6	134.0 ± 5.0	0.4	...	0.49 (0.58)
NGC 424	≤0.6	81.0 ± 33.4	439.4 ± 72.7	-0.2 ± 58.0	0.8	80.0 ± 0.0	1.00 (0.74)
NGC 520	≤1.9
NGC 660	≤1.8
NGC 986	≤4.0
NGC 1008	≤0.2	≤3.2

Table 2.3: *L*- and *M*-band AGN Extended Flux Catalog

Table 3
(Continued)

Target Name	L Flux _{ext} (mJy)	$\Theta_{L, \text{ext}}$ (pc)	$P_{A, \text{ext}}$ (deg)	r_L $\equiv \theta/\Theta$	M Flux _{ext} (mJy)	$\Theta_{M, \text{ext}}$ (pc)	$P_{A, \text{ext}}$ (deg)	r_M $\equiv \theta/\Theta$	$P_{A, \text{y}}$ (deg)	f_{inc} L (M)
NGC 1068	1818.3 ± 874.9	83.1 ± 9.4	-19.0 ± 38.4	0.9	4204.3 ± 1488.6	83.1 ± 8.0	-17.4 ± 25.2	0.9	175.0 ± 3.0	0.53 (0.40)
NGC 1097	8.3 ± 1.3	133.2 ± 9.1	-0.4 ± 68.5	0.9	≤1.9	0.40 (0.75)
NGC 1125	6.0 ± 1.3	5.3 ± 0.4	130.5 ± 10.0	0.8	≤1.6	0.37 (0.82)
NGC 1368	4.4 ± 0.3	8.1 ± 0.3	104.9 ± 8.7	0.9	0.13 (...)
NGC 1365	48.8 ± 24.0	173.4 ± 40.9	2.7 ± 58.1	0.7	≤1.7	82.0 ± 25.0	0.80 (0.99)
NGC 1386	20.9 ± 5.0	114.2 ± 8.9	-32.6 ± 10.4	0.8	18.9 ± 8.3	114.2 ± 14.1	0.9 ± 52.7	0.9	177.0 ± 20.0	0.47 (0.61)
NGC 1511	≤1.5	-(0.63)
NGC 1566	11.1 ± 1.8	118.2 ± 9.3	1.5 ± 32.8	0.9	157.0 ± 37.0	0.49 (...)
NGC 1614	15.6 ± 2.1	...	72.5 ± 9.6	0.8	...	-(0.31)
NGC 1667	2.0 ± 0.4	597.4 ± 51.5	11.4 ± 3.7	0.5	54.0 ± 0.0	0.42 (...)
NGC 1808	15.7 ± 2.2	72.6 ± 3.6	29.8 ± 5.2	0.7	9.0 ± 2.1	72.6 ± 7.4	42.6 ± 13.9	0.8	...	0.29 (0.54)
NGC 3125	≤0.6	≤3.1
NGC 3281	≤5.0	176.0 ± 16.0	-(0.96)
NGC 3660	1.8 ± 0.3	324.6 ± 23.7	-8.8 ± 56.8	0.9	≤1.4	0.49 (...)
NGC 4038/9	≤1.9
NGC 4074	≤0.7	≤4.4
NGC 4235	11.9 ± 4.0	177.3 ± 18.8	-1.9 ± 61.9	0.9	8.7 ± 2.0	177.3 ± 5.8	74.5 ± 4.5	0.5	189.0 ± 0.0	0.95 (0.71)
NGC 4253	33.1 ± 14.4	6.5 ± 0.7	-18.0 ± 29.9	0.9	≤4.6	66.0 ± 61.0	0.55 (0.39)
NGC 4261	≤0.6	≤5.2	0.41 (0.90)
NGC 4278	≤0.5	≤2.5
NGC 4303	≤2.1
NGC 4303	5.9 ± 0.7	119.5 ± 5.4	-10.4 ± 26.3	0.9	≤4.1	0.22 (...)
NGC 4374	≤0.4	≤2.0
NGC 4388	≤0.4	≤2.6	28.0 ± 34.0	0.99 (0.95)
NGC 4418	1.7 ± 0.8	207.9 ± 27.8	99.9 ± 9.2	0.7	≤4.5	0.55 (0.60)
NGC 4438	12.6 ± 1.4	133.5 ± 6.0	-23.2 ± 3.0	0.6	≤5.1	0.24 (...)
NGC 4457	15.5 ± 1.2	139.6 ± 4.8	112.6 ± 6.7	0.8	≤3.7	0.16 (...)
NGC 4472	≤0.4	≤1.9
NGC 4501	13.7 ± 1.2	144.1 ± 4.9	43.1 ± 5.1	0.8	≤2.4	153.0 ± 0.0	0.10 (...)
NGC 4507	51.5 ± 22.9	259.3 ± 35.5	78.1 ± 45.6	0.9	≤6.4	111.0 ± 22.0	0.54 (0.92)
NGC 4579	17.1 ± 4.3	82.5 ± 9.4	85.8 ± 48.6	0.9	7.2 ± 3.3	82.5 ± 13.6	87.3 ± 42.7	0.8	...	0.50 (0.67)
NGC 4593	22.8 ± 11.1	155.6 ± 22.4	78.5 ± 47.6	0.9	25.6 ± 7.0	155.6 ± 14.9	-10.3 ± 53.0	0.9	103.0 ± 19.0	0.58 (0.58)
NGC 4594	38.9 ± 2.0	120.0 ± 3.0	90.8 ± 0.9	0.4	≤3.8	0.07 (...)
NGC 4746	≤0.4
NGC 4785	4.0 ± 0.5	436.7 ± 23.5	71.9 ± 6.7	0.8	≤2.0	0.22 (...)
NGC 4941	4.2 ± 1.4	115.7 ± 13.3	17.2 ± 32.7	0.9	≤3.4	116.0 ± 11.0	0.52 (...)
NGC 4945	54.9 ± 3.5	42.6 ± 0.9	133.4 ± 1.0	0.4	≤2.3	0.04 (...)
NGC 5135	8.0 ± 3.3	202.1 ± 28.0	80.4 ± 46.0	0.9	≤3.4	66.0 ± 48.0	0.61 (0.85)
NGC 5252	13.3 ± 4.8	344.6 ± 38.6	69.0 ± 50.8	0.9	≤4.8	144.0 ± 24.0	0.55 (0.85)
NGC 5363	8.4 ± 0.7	149.6 ± 5.9	46.1 ± 6.0	0.8	≤4.2	0.10 (...)
NGC 5427	≤0.3	≤2.1	0.90 (...)
NGC 5506	123.8 ± 49.6	165.1 ± 30.0	15.2 ± 49.8	0.8	152.4 ± 47.4	165.1 ± 8.0	46.8 ± 10.8	0.7	60.0 ± 56.0	0.73 (0.70)
NGC 5548	59.0 ± 3.8	410.8 ± 12.3	11.6 ± 5.8	0.9	≤5.0	200.0 ± 17.0	0.05 (...)
NGC 5643	10.9 ± 2.3	91.2 ± 8.4	100.3 ± 49.2	0.9	≤3.6	52.0 ± 17.0	0.44 (0.82)
NGC 5728	3.0 ± 1.1	358.3 ± 41.7	97.7 ± 29.8	0.9	≤4.5	98.0 ± 25.0	0.49 (...)
NGC 5813	≤0.3	≤1.9
NGC 5953	≤0.3	≤2.0	0.96 (...)

Table 2.3: continued.

Table 3
(Continued)

Target Name	$L \text{ Flux}_{\text{ext}}$ (mJy)	$\Theta_{L, \text{ext}}$ (pc)	$\text{PA}_{L, \text{ext}}$ (deg)	r_L $\equiv \theta/\Theta$	$M \text{ Flux}_{\text{ext}}$ (mJy)	$\Theta_{M, \text{ext}}$ (pc)	$\text{PA}_{M, \text{ext}}$ (deg)	r_M $\equiv \theta/\Theta$	PA_N (deg)	f_{inc} $L (M)$
NGC 5995	44.7 ± 13.4	613.3 ± 69.7	-12.2 ± 40.1	0.9	53.7 ± 19.0	613.3 ± 69.4	75.1 ± 39.8	0.9	111.0 ± 39.0	0.57(0.50)
NGC 6000	≤2.1 (...)
NGC 6221	12.8 ± 2.9	103.5 ± 9.0	-12.5 ± 7.5	0.6	≤3.8	0.49 (0.75)
NGC 6240N	12.4 ± 1.5	1510.1 ± 52.4	-12.3 ± 3.3	0.5	≤1.8	0.57 (0.91)
NGC 63	1.2 ± 0.5	...	56.4 ± 12.4	0.6	...	- (0.98)
NGC 6810	28.4 ± 2.5	199.4 ± 7.1	10.7 ± 2.4	0.6	10.0 ± 2.0	199.4 ± 12.3	9.3 ± 3.0	0.5	...	0.20 (0.33)
NGC 6814	5.4 ± 1.5	108.3 ± 13.0	84.2 ± 59.7	0.9	≤3.6	97.0 ± 10.0	0.66 (0.78)
NGC 6860	15.0 ± 5.1	448.9 ± 56.5	86.1 ± 60.4	0.9	≤4.9	53.0 ± 49.0	0.70 (0.91)
NGC 6890	6.9 ± 1.5	164.6 ± 15.8	-16.3 ± 39.3	0.9	≤1.3	102.0 ± 0.0	0.54 (0.88)
NGC 7130	7.7 ± 1.2	406.2 ± 28.3	9.7 ± 49.4	0.9	≤2.8	0.47 (0.89)
NGC 7172	31.5 ± 7.6	215.5 ± 23.0	16.7 ± 50.5	0.9	≤3.4	86.0 ± 10.0	0.61 (0.95)
NGC 7213	≤3.1 (...)
NGC 7314	6.8 ± 1.5	73.4 ± 6.7	91.4 ± 25.7	0.9	≤2.5	87.0 ± 43.0	0.54 (0.89)
NGC 7479	10.0 ± 2.5	168.9 ± 17.5	-3.4 ± 37.0	0.9	82.0 ± 0.0	0.56 (...)
NGC 7496	4.7 ± 1.2	133.4 ± 13.9	16.6 ± 48.1	0.9	≤2.3	0.58 (0.77)
NGC 7552	≤2.3 (...)
NGC 7582	60.2 ± 26.9	246.0 ± 58.0	19.1 ± 42.0	0.7	83.2 ± 22.0	246.0 ± 24.2	100.7 ± 56.8	0.9	...	0.77 (0.52)
NGC 7590	4.0 ± 0.3	226.6 ± 8.2	-29.0 ± 3.7	0.7	≤1.7	0.14 (...)
PG 2130+099	20.7 ± 9.5	1525.4 ± 180.2	67.8 ± 31.8	0.9	≤4.5	158.0 ± 27.0	0.45 (0.92)
PKS 1417-19	≤0.7	≤4.2	98.0 ± 0.0	0.91 (...)
PKS 1814-63	≤0.4	≤2.1	25.0 ± 0.0	0.90(...)
PKS 1932-46	≤0.3	≤2.1 (...)
Superantennae S	9.4 ± 2.8	1553.8 ± 189.1	6.5 ± 32.8	0.8	17.4 ± 4.9	1553.8 ± 132.7	78.1 ± 56.5	0.9	151.0 ± 26.0	0.64 (0.61)
UGC 2369 S	≤1.9 (...)
Z 41-20	≤0.7	≤4.1	172.0 ± 36.0	0.76 (...)

Note. Here Flux_{ext} indicates the flux from the fitted resolved emission, and we only include Θ and PA for those sources with $>2\sigma$ detections. N -band PAs are from Asmus et al. (2016). Here $r_x \equiv \theta_{\text{FWHM},x}/\Theta_{\text{FWHM},x}$. (This table is available in machine-readable form.)

Table 2.3: continued. Available in machine-readable format.

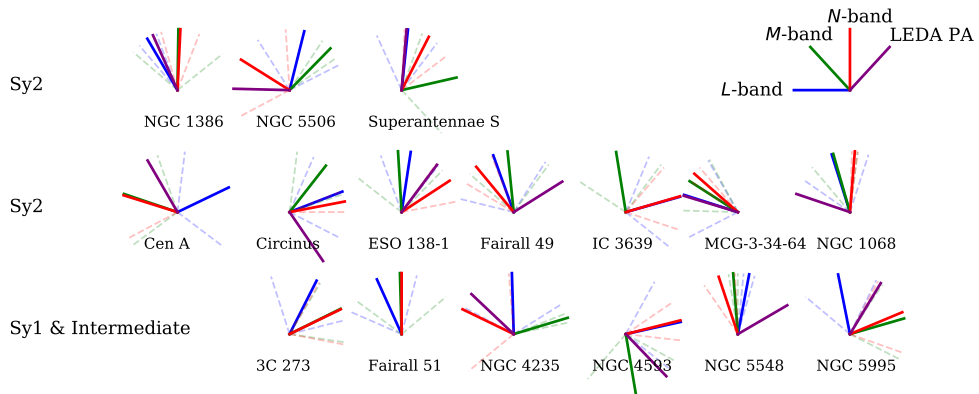


Fig. 2.5: The measured infrared position angles (PAs) of sources with 1) $\text{SNR} \geq 2$ extended emission in the L - and M -bands, and 2) a reported PA in [Asmus et al. \(2016\)](#). We include the optical position angle from HyperLEDA ([Makarov et al., 2014](#)) in purple for reference. The L , M , and N position angles are shown in blue, green, and red, respectively. 1σ errors are given as faint dashed lines in the same color as the reported angle.

[Kishimoto, 2010](#)) and the CLUMPY models ([Nenkova et al., 2008c](#)) resemble classical clumpy tori made up of spherical clouds, while CAT3D-WIND ([Hönig & Kishimoto, 2017](#)) includes an additional hollow cone in the polar region, representing a dusty wind driven by radiation pressure. Numerous previous works have performed SED fitting to test various torus models (e.g., [Ramos Almeida et al., 2009](#); [Alonso-Herrero et al., 2011](#); [Lira et al., 2013](#)), and they have emphasized the importance of the $5\mu\text{m}$ flux. We therefore focus on the NIR, utilizing the new L - and M -band measurements from this work, presenting a spectral slope plot similar to that in e.g., [Hönig & Kishimoto \(2017\)](#).

In addition, we include in comparison a new, unpublished model library consisting of a flared disk and a polar wind³. We discuss the model only briefly here, and only in comparison to SKIRTOR ([Stalevski et al., 2016](#)). The disk component is defined by its angular width and optical depth and is similar to that in SKIRTOR, but with dust distributed smoothly. The polar wind takes the shape of a hollow cone parametrized by half-opening angle and radial extent. The dust composition is the same as in SKIRTOR (53 % silicates and 47 % graphite), and with the same power-law grain size distribution ([Mathis et al., 1977](#)), but with larger grains (between 0.1 and 1 μm), as suggested by flat extinction curves and silicate feature profiles ([Gaskell et al., 2004](#); [Shao et al., 2017](#); [Xie et al., 2017](#)). For the hollow cone, we calculated an additional set of models with only graphite grains: if dust in the wind is driven away from the sublimation zone, it might be expected to be silicate-poor. This is a preliminary model library with limited parameter space: the current number of SEDs is an order of magnitude smaller than its precursor, SKIRTOR. However, it is well-motivated by the model of the resolved MIR images of the Circinus AGN ([Stalevski et al., 2017, 2019](#)). Even though this model library is to be developed further and tested against larger datasets, it is useful for our purpose here, since it represents an extension of the SKIRTOR models, thus allowing us to isolate the effect of polar winds on the flux ratios. While at this stage any results obtained with this model are to be taken with caution, it is indicative that, as it will be shown below, even though it is much more limited in terms of parameters and number of SEDs, it provides a significantly better match to the observed color space than its

³For further information on the SKIRTOR+wind model, please contact M. Stalevski.

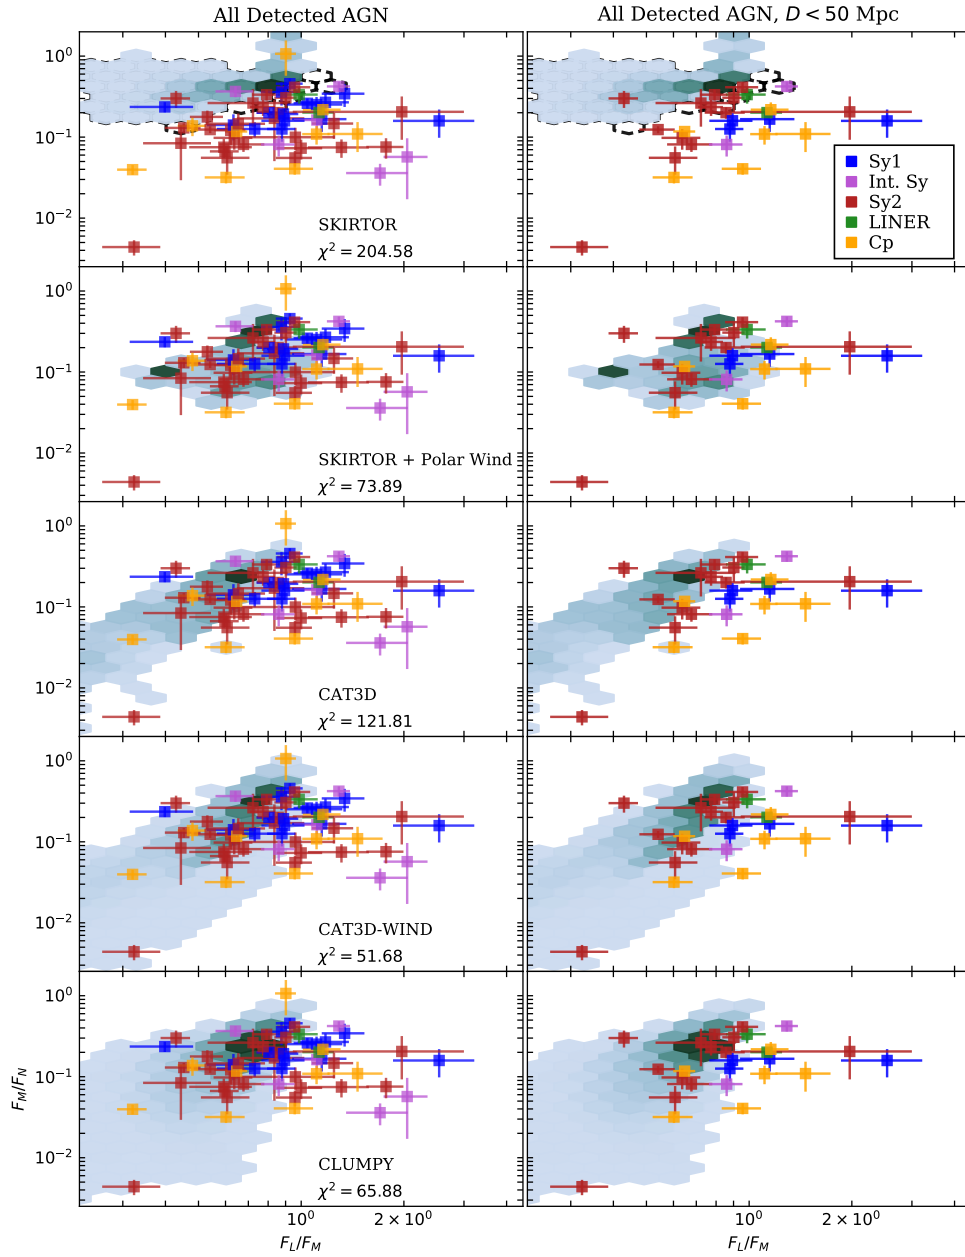


Fig. 2.6: Comparisons of L/M and M/N flux ratios to those predicted by torus models SKIRTOR (Stalevski et al., 2016), SKIRTOR+Polar Wind, CLUMPY (Nenkova et al., 2008c), CAT3D (Hönig & Kishimoto, 2010), and CAT3D-WIND (Hönig & Kishimoto, 2017). Only includes sources with $\text{SNR}_{\text{gauss}} > 3$ in both the L - and M -bands. Colors are the same as in Fig. 2.1. In the (*left*) panels we show all detected sources in our sample. In the (*right*) panels we show only AGN closer than 50 Mpc in our sample. The χ^2 value is measured from each data point to the model SED which provides the best fit. For SKIRTOR we show the effects of using the Kishimoto et al. (2008) accretion disk spectrum (described in §5.2) as a dashed contour.

precursor. Also, unlike any other included model library, the color space covered by this model does not extend significantly beyond the observed colors.

For each model SED produced by these setups, we extract the L -, M -, and N -band fluxes. We then compare the flux ratios F_L/F_M and F_M/F_N to those from the L - and M -band fluxes from this sample and the N -band fluxes from A14, shown in Fig 2.6. We compare the models using the χ^2 value, measured from each galaxy to the SED in each model that it matches most closely. We primarily consider Seyfert types 1, intermediate, and 2 in this analysis, as torus orientation aims to explain their differences in AGN Unification. We include the other AGN (e.g., LINERs and Cp) as they may have dusty tori but their mid-infrared emission might not be dominated by them.

Our observations favor torus+wind models such as CAT3D-WIND and the new SKIRTOR+wind. Specifically, CAT3D-WIND has the lowest χ^2 value, $\chi^2_{\text{CAT3D-WIND}} = 51.68$, which is a very sharp improvement from the base CAT3D $\chi^2_{\text{CAT3D}} = 121.81$. Similarly, the agreement with SKIRTOR is greatly improved with the inclusion of the polar wind: $\chi^2_{\text{SKIRTOR}} = 204.58 \rightarrow \chi^2_{\text{SKIRTOR+wind}} = 73.89$. CLUMPY exhibits a similar agreement, $\chi^2_{\text{CLUMPY}} = 65.88$, notably *without* the inclusion of winds. The χ^2 values serve here primarily as a means for simple comparison, as the various torus models have vastly different parameter spaces and numbers of SEDs. Our results nonetheless agree with González-Martín et al. (2019) who found that CAT3D-WIND best matched their Swift/BAT-selected sample, with CLUMPY performing second-best. Moreover, similar to González-Martín et al. (2019), we find that the AGN colors are confined to a much smaller region of the parameter space than the torus SEDs. With our simple color-color comparisons, however, we cannot distinguish whether these regions have unphysical parameters or whether the discrepancy is driven by selection effects.

We find 13 of our Seyfert sample are “bluer” than predicted by any model F_L/F_M . We note that none of the torus models produces SEDs wherein $F_L/F_M \gtrsim 1$. We focus on color-color comparisons rather than SED fitting as they utilize the photometric data we have measured in this work, but we note that full SED fitting would be required to fully test the validity of the various torus models. Our comparisons are thus only qualitative, but they do give hints as to a preferred “torus” feature: the polar wind.

There are two explanations for this apparent discrepancy which are discussed in the following subsections: stellar contamination to the nuclear flux and incomplete representation of the 3-5 μm bump in the torus models.

On average, we find that Sy1 are bluer than Sy2 in the NIR ($\log(F_{L,\text{Sy1}}/F_{M,\text{Sy1}}) \equiv \log(f_{LM,\text{Sy1}}) = -0.05 \pm 0.18$ vs. $\log(f_{LM,\text{Sy2}}) = -0.16 \pm 0.18$ and that they are also slightly bluer in the MIR ($\log(F_{M,\text{Sy1}}/F_{N,\text{Sy1}}) \equiv \log(f_{MN,\text{Sy1}}) = -0.70 \pm 0.17$ vs. $\log(f_{MN,\text{Sy2}}) = -0.91 \pm 0.37$). In fact, Sy1 primarily populate the blue limit of the the torus SED models in F_L/F_M . The Sy2 sample shows much more diversity in both flux ratios, and there are several outliers which are much bluer than models predict. Only two LINERs were significantly detected in both L and M , so we cannot draw any conclusions about this population. The extremely mixed Cp subsample shows a large amount of scatter, but tends to be extremely “blue” compared to the Seyferts and to the models ($\log(f_{LM,\text{Cp}}) = 0.05 \pm 0.24$). It does not seem as if this diagram can be used to reliably separate different AGN types.

5.1 Stellar Contamination

The physical scales at which we extract the “nuclear” flux in this sample vary greatly, from ~ 10 pc to more than 1.5 kpc. For essentially all sources at $D_{\text{AGN}} \geq 50$ Mpc, corresponding to an unresolved area $\theta_{\text{FWHM}} \gtrsim 100$ pc, we expect a significant contribution from the “red” tail of the stellar emission to our L -band fluxes. This stellar contribution would serve to make our AGN artificially bluer than the torus models. We therefore

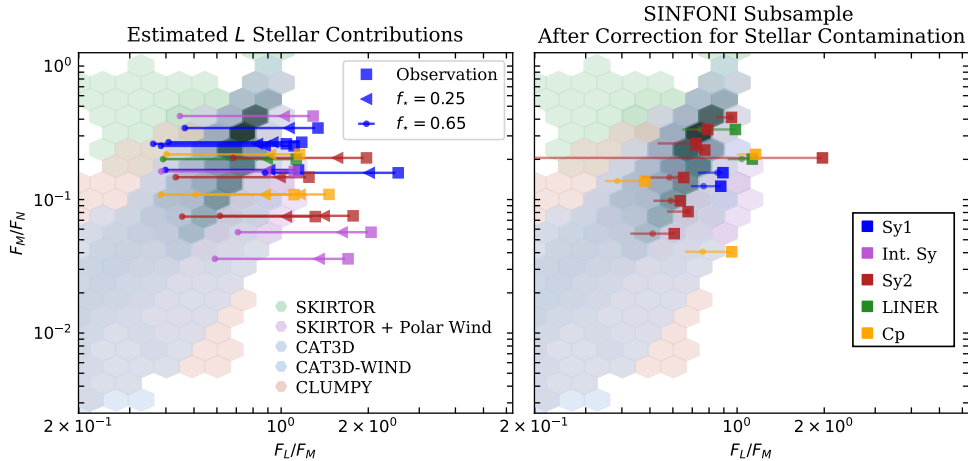


Fig. 2.7: Comparisons of L/M and M/N flux ratios to those predicted by torus models SKIRTOR (green and purple; [Stalevski et al., 2016](#)), CAT3D (gray; [Hönig & Kishimoto, 2010](#)), CLUMPY (orange; [Nenkova et al., 2008c](#)), and CAT3D-WIND (blue; [Hönig & Kishimoto, 2017](#)). Only includes sources with $\text{SNR}_{\text{gauss}} > 2$ in both the L - and M -bands. Colors same as in Fig. 2.1. In the (*left*) panel we show the sources which are bluer than models allow with corrections made for various $f_{*,L}$. In the (*right*) panel we show the 21 AGN which had SINFONI fluxes in [Burtscher et al. \(2015\)](#) which we used to estimate L -band stellar contribution using Eq. 2.3.

perform a simple distance cut, $D \leq 50$ Mpc, in order to focus on the polar dust region. We show this in Fig. 2.6. We find that after this cut, only two Seyfert galaxies are strictly $\geq 1\sigma$ bluer than the models. They are NGC 4235 (Sy1) and NGC 1365 (Sy1.8). This simple test then indicates qualitatively that increased stellar emission on larger scales could bias the AGN with $D > 50$ Mpc to bluer colors.

More quantitatively, in [Alonso-Herrero et al. \(2001\)](#) the stellar contribution to the L -band within a $3''$ diameter aperture was estimated to be up to 25% for a sample of 12 Sy2 AGN (except for NGC 5252 which had 65% stellar contribution). In the M -band it was estimated to be $< 10\%$. This indicates a blueward flux-ratio shift of $\log(F_L/F_M) \lesssim 0.1$. For their Sy1 sample (5 AGN) they assumed that 100% of the L -band flux was non-stellar in origin. We show the effects of stellar contamination in Fig. 2.7 by plotting the expected flux ratios of the AGN in our sample assuming they have 0%, 25%, and 65% stellar L -flux.

With a 25% L -band correction, 7 of the AGN in the full sample are still bluer than the models predict. If we allow for up to 65% of the L -flux to come from a stellar origin (i.e., the maximum percentage found in [Alonso-Herrero et al., 2001](#)), the shift increases to $\log(F_L/F_M) < 0.5$, which would make all the outliers in Fig. 2.6 consistent with the models.

We do not include an estimate of the M -band stellar flux because *i*) [Alonso-Herrero et al. \(2001\)](#) found it to be strictly $< 10\%$; *ii*) [Assef et al. \(2010\)](#) show that the AGN flux is steeply rising between the L - and M -bands while the stellar contribution steeply drops, indicating the M -band stellar contribution is $\lesssim 5\%$; and *iii*) subtracting an M -band contribution would actually serve to make AGN in Figs. 2.6-2.7 agree even less with the models by shifting them toward bluer colors.

5.1.1 SINFONI-Estimated Stellar Fluxes

For a subset of 21 AGN we can estimate the stellar contamination using K -band fluxes and AGN fractions derived by [Burtscher et al. \(2015\)](#). Using SINFONI, these fluxes were measured in the central $1''$ diameter aperture of each galaxy. Their NIR AGN fractions

($f_{\text{AGN,K}}$) were derived from spectral fitting in this central region and checked using the stellar CO equivalent width. We then used the average color of old stellar populations in stellar bulges, $K - L = 0.09$ mag, as measured from the empirical elliptical galaxy template of [Assef et al. \(2010\)](#), to calculate the expected L -band stellar magnitude as

$$L_{\text{mag},\star} = K_{\text{mag,total}} + \text{AGN}_{\text{mag,correction}} - 0.09, \quad (2.3)$$

where $\text{AGN}_{\text{mag,correction}} = -\log_{10}((100 - f_{\text{AGN,K}})/100)$ is a correction for the measured AGN fraction of the galaxy. We converted these magnitudes into ISAAC L fluxes.

We measured the ISAAC L fluxes of these 21 AGN using a central, 1" diameter aperture. We list the 1" aperture fluxes, the Gaussian-fitted nuclear fluxes, the estimated stellar fluxes, and the measured AGN fractions in Table 2.4. We show a comparison in Fig. 2.8 to illustrate the AGN fractions. Note that the typical unresolved nuclear flux in our sample has a FWHM of 0.75", so the stellar fractions could be slightly less in the unresolved nuclear fluxes on nights with average or better seeing. Note also that Circinus was measured in the SINFONI sample with a 0.5" diameter aperture, which we match here; it is marked in Table 2.4 and Fig. 2.8.

We find that a rather large fraction of the AGN (13/21) should have $> 10\%$ of their ISAAC L nuclear flux come from stars rather than the unresolved nucleus. Two of these sources, NGC 4303 and NGC 4261, are shown to have $\approx 70\%$ stellar flux in the L -band. This is in rough agreement with the K -band $f_{\text{AGN,K}} < 10\%$ estimates found by [Burtscher et al. \(2015\)](#). All but two of the AGN in this subsample agree with the findings of [Alonso-Herrero et al. \(2001\)](#), i.e., stellar contributions $\leq 65\%$.

We finally compare these AGN once more with the torus models in Fig. 2.7 (right panel), but now with their K -band estimated stellar flux subtracted. Following this stellar correction, all of the AGN in this subsample are compatible with the torus SED models of CAT3D-WIND, CLUMPY, and SKIRTOR+wind. The stellar corrections required for CLUMPY are larger than for the torus+wind models as shown in Fig. 2.7. To correct all of the AGN nuclear fluxes in this sample, we would need similar SINFONI K -band data and stellar CO equivalent width measurements. For the nearest AGN, however, we expect this correction to be on the order of 1%, as exhibited by Circinus and NGC 1068.

5.2 Accretion Disk Spectra

The assumed spectrum of the accretion disk can play a large role in the final model SED. Traditionally a broken power-law has been assumed, i.e.,

$$\lambda F_{\lambda} \propto \begin{cases} \lambda^{1.2} & 0.001 \leq \lambda \leq 0.01 \mu\text{m} \\ \lambda^0 & 0.01 < \lambda \leq 0.1 \mu\text{m} \\ \lambda^{-0.5} & 0.1 < \lambda \leq 5 \mu\text{m} \\ \lambda^{-3} & 5 < \lambda \leq 50 \mu\text{m} \end{cases} \quad (2.4)$$

(see e.g., [Sanders et al., 1989](#); [Kishimoto et al., 2008](#), and references therein). The spectrum produced by Eq. 2.4 was used in the models CLUMPY ([Nenkova et al., 2008c](#)) and SKIRTOR ([Stalevski et al., 2016](#)). More recently, however, a bluer accretion SED has been adopted which is strongly supported by QSO observations (e.g., [Zheng et al., 1997](#); [Manske et al., 1998](#); [Vanden Berk et al., 2001](#); [Scott et al., 2004](#); [Kishimoto et al., 2008](#)). This SED is relatively blue as it has a shallower power-law falloff $\nu F_{\nu} \propto \nu^{4/3}$

Target	L Flux _{1as} [mJy]	L Flux _{nuc.gauss} [mJy]	Est. L Flux _* [mJy]	f_{AGN}
Circinus	628.0 ± 6.5	459.3 ± 42.5	5.4	0.99
IC 5063	57.0 ± 0.1	57.0 ± 2.2	1.4	0.98
NGC 1068	3111.7 ± 6.2	2012.8 ± 238.8	70.9	0.98
NGC 1097	8.9 ± 0.1	5.6 ± 0.2	3.5	0.61
NGC 1386	30.1 ± 0.2	18.8 ± 1.0	4.6	0.85
NGC 1566	14.7 ± 0.1	10.6 ± 0.4	5.4	0.63
NGC 4261	1.9 ± 0.1	0.6 ± 0.1	1.6	0.16
NGC 4303	4.3 ± 0.1	1.6 ± 0.1	2.8	0.33
NGC 4303	4.3 ± 0.1	3.8 ± 0.2	2.2	0.50
NGC 4388	29.3 ± 0.1	33.9 ± 1.3	5.7	0.80
NGC 4501	9.2 ± 0.3	1.5 ± 0.2	5.1	0.44
NGC 4579	27.5 ± 0.2	16.9 ± 0.6	6.7	0.76
NGC 4593	49.0 ± 0.2	32.3 ± 2.5	4.6	0.91
NGC 5135	18.6 ± 0.2	12.6 ± 0.7	11.5	0.38
NGC 5506	400.6 ± 0.3	343.8 ± 12.5	1.2	1.00
NGC 5643	17.0 ± 0.1	8.5 ± 0.5	8.5	0.50
NGC 6814	13.3 ± 0.1	10.6 ± 0.4	3.0	0.77
NGC 7130	12.3 ± 0.1	6.9 ± 0.3	3.1	0.75
NGC 7172	67.4 ± 0.2	49.1 ± 2.0	1.5	0.98
NGC 7496	9.3 ± 0.1	6.6 ± 0.3	5.4	0.42
NGC 7582	150.8 ± 0.6	86.1 ± 5.3	2.4	0.98

Table 2.4: Measured 1" diameter aperture flux, measured nuclear Gaussian flux, estimated stellar flux using SINFONI *K*-band measurements and Eq. 2.3, and measured AGN fraction for 21 cross-matched AGN.




Fig. 2.8: Comparison of SINFONI-estimated L -band stellar contributions and our measured 1" aperture fluxes from ISAAC. Overplotted are the lines showing what we expect if 100%, 65%, 25%, 10% and 1% of the flux is stellar in origin. Circinus is marked separately because it was measured with an aperture of 0.5" in both ISAAC and SINFONI observations. AGN are color coded as in Fig. 2.1.

between 0.3 and 3 μm :

$$\lambda F_\lambda \propto \begin{cases} \lambda^1 & \lambda < 0.03 \mu\text{m} \\ \lambda^0 & 0.03 < \lambda \leq 0.3 \mu\text{m} \\ \lambda^{-4/3} & 0.3 < \lambda \leq 3 \mu\text{m} \\ \lambda^{-3} & 3 \mu\text{m} < \lambda, \end{cases} \quad (2.5)$$

and it was used in both CAT3D and CAT3D-WIND (Hönig & Kishimoto, 2010, 2017, respectively).

In order to test the role the model accretion disk spectrum plays in the L - and M -bands, we replaced Eq. 2.4 with Eq. 2.5 in the SKIRTOR models using the method outlined in Yang et al. (2020). The results of this replacement are shown in Fig.2.6. The overall effect of this replacement is to shift the models' flux ratios F_L/F_M “blueward”, i.e., to the right in our plots. After this replacement, the SKIRTOR models with the bluer accretion disk are a better match to the data $\chi^2_{\text{SKIRTOR}} = 204.58 \rightarrow \chi^2_{\text{Bluer AD}} = 174.17$, but this shift is not enough to explain observations with $F_L/F_M > 1$. We show the effects of changing the accretion disk spectrum in Fig. 2.6 as a dashed contour in the SKIRTOR panels. The CLUMPY library provides two sets of models: one with only dust emission from the torus, and another which includes the accretion disk. However, the accretion disk emission is not traced in the radiative transfer calculations. Rather, it is up to the user to decide which one to use, e.g. based on the probability to have a clear line of sight to the central source for the given model. While this is not fully consistent treatment, we tested both versions and found the difference to be very small ($\Delta\chi^2_{\text{CLUMPY}} < 1$). We conclude that including the accretion disk flux in CLUMPY models has an insignificant effect on the model predictions. In both cases, only a radically different accretion disk spectrum could explain these data, so stellar contamination is still favored.

5.3 Polar Elongation and the 3-5 Micron Bump

It is also possible that the MIR flux is underestimated in modern torus models. It has been well documented that many QSOs and Sy1 galaxies exhibit what is called the 3-5 μm bump (e.g., Edelson & Malkan, 1986; Kishimoto et al., 2011; Mor & Netzer, 2012; Hönig et al., 2013). This feature was thought to be caused by the presence of inter-clump dust, but recent work by Hönig & Kishimoto (2017) claims that the addition of wind orthogonal to the accretion disk can explain it. In CAT3D-WIND, Hönig & Kishimoto (2017) this feature is apparently caused by a large amount of dust in a “puffed-up” region in the vicinity of the sublimation ring or by winds which remove clouds from the central few pc and cause the SED to be split into the hot dust emission from the center and the cold dust emission from a region elongated in the polar direction. A similar conclusion is valid for the models based on SKIRTOR with polar winds: the additional dust of low optical depth at the base of the wind is responsible for additional 3-5 μm emission and “bluer” colors.

It appears that within our sample, the differences between observations and the models are sufficiently explained by stellar contamination in the L -band. The AGN in our sample agree best with CAT3D-WIND(Hönig & Kishimoto, 2017), indicating that the inclusion of the polar wind describes AGN at this resolution quite well. Moreover, the SKIRTOR (Stalevski et al., 2016) models fit the data much better after the inclusion of a polar wind. On the other hand, we find that CLUMPY (Nenkova et al., 2008c) provides good fits to the data without polar winds. It will be interesting to further investigate the nuclear regions of NGC 1365 and NGC 4235 with MATISSE, testing whether these “too blue” Seyferts are sufficiently explained by stellar contamination.

6 Estimating VLT Fluxes from *WISE*

The Wide-field Infrared Survey Explorer (*WISE*) space telescope (Wright et al., 2010) has observed a large number of AGN in the MIR at high sensitivity but limited spatial resolution. Moreover, the ALLWISE AGN catalog contains 1.4 million AGN selected using 3 *WISE* bands (Secrest et al., 2015). This catalog and others like it present a rich source from which to draw AGN for further study in the MIR. The median point-spread-function in the *W1*-band has a FWHM of 6", making it nearly impossible to spatially distinguish thermal torus emission from star formation in the foreground or even the central few hundred parsecs of even nearby AGN. The large spatial areas probed by *WISE* in each AGN will bias the flux to be comparatively larger than seen from the 8-10m class telescopes used in MIR interferometry. In order to select AGN from a much larger parent sample for MATISSE followup from *WISE*, we study the relationships between *WISE* colors, *WISE* fluxes, and the fraction of nuclear flux observable at the VLT(I).




Fig. 2.9: (top) $L/W1$ flux ratio vs. $W1-W2$. (bottom) $M/W2$ flux ratio vs. $W1-W2$. For each panel only detected ($\text{SNR}_{\text{gauss}} > 2$ in each band) sources are included. Colors same as Fig 1. The black line at $W1-W2=0.8$ indicates the AGN selection cutoff of Stern et al. (2012), which is also the point after which $F_L/F_{W1} \approx 1$ in our Sy1 sample. Blue and red dashed lines are fitted piece-wise functions (Eq. 2.5-2.8) for the Sy1 and Sy2 samples, respectively. gray shaded regions are the 1σ fit errors.

We explore the relationships between the measured ISAAC fluxes and the existing *WISE* fluxes for each source in order to determine a scaling between the space-based MIR measurements and nuclear fluxes measured at the VLT. By providing an estimate for the fluxes observed at the VLT, we can select AGN for follow-up observations directly from the *WISE* catalog. This is not straightforward because the *WISE* telescope has a lower resolution than the VLT and because the filters *W1* and *W2* are both wider than and offset from either L' or M_{nb} . The resolution in the *W1* band is 6.1" and is 6.4" in the *W2* band (Wright et al., 2010), while these ISAAC observations are seeing limited at $\lesssim 1''$. This means that in general we expect $F_{\text{ISAAC}}/F_{\text{WISE}} \leq 1$, as there will be a contribution from non-nuclear flux, especially at shorter wavelengths wherein the stellar contribution is higher. In Fig. 2.9 we show the ratios F_L/F_{W1} and F_M/F_{W2} for each detected ($\text{SNR}_{\text{gauss}} > 2$) AGN in our sample. Note that F_L/F_{W1} and F_M/F_{W2} are occasionally > 1 . This is due to the offset in the filters: L is redder than *W1*, and it therefore measures a larger relative flux from the nuclear dust emission in the same AGN. Nevertheless, for both the Sy1 and the Sy2 subsamples we see a linearly

increasing trend in $F_{\text{ISAAC}}/F_{\text{WISE}}$ for blue $W1-W2$ colors until a ‘‘saturation’’ point at $W1-W2 \approx 0.8$ mag. This point occurs at roughly the same $W1-W2$ color in both subsamples in both bands. To further quantify this, we fit a linear piece-wise function to each subsample of the form

$$f(W1 - W2) = \begin{cases} ap + b, & \text{if } (W1 - W2) \geq p \\ a(W1 - W2) + b, & \text{else} \end{cases} \quad (2.6)$$

where p is the fitted point at which $F_{\text{ISAAC}}/F_{\text{WISE}} \approx 1$. The fitted lines are shown in Fig. 2.9 with 1σ errors shaded in gray. Fitting was done using `curve_fit` from `scipy.optimize`, employing the Levenberg-Marquardt algorithm. Errors are estimated from the diagonal of the covariance matrix, with the implicit assumption that they are uncorrelated. We find that for the Sy1 sample, the L flux matches the $W1$ flux for $W1 - W2 \geq 1.05 \pm 0.03$ mag. For the Sy2 sample this cutoff is $W1 - W2 \geq 0.90 \pm 0.01$ mag, much redder than for the Sy1 sample. For Sy1 and Sy2 galaxies we can estimate the L nuclear flux from the $W1$ flux as

$$\log_{10} F_{L,\text{Sy1}} = \quad (2.7)$$

$$\begin{cases} \log_{10} F_{W1} - 0.05 \pm 0.002, & \text{if } (W1 - W2) \geq 1.05 \pm 0.03 \\ \log_{10} F_{W1} + (1.12 \pm 0.03)(W1 - W2) - (1.23 \pm 0.02), & \text{else} \end{cases}$$

$$\log_{10} F_{L,\text{Sy2}} = \quad (2.8)$$

$$\begin{cases} \log_{10} F_{W1} - 0.23 \pm 0.004, & \text{if } (W1 - W2) \geq 0.90 \pm 0.01 \\ \log_{10} F_{W1} + (1.51 \pm 0.03)(W1 - W2) - (1.61 \pm 0.01), & \text{else.} \end{cases}$$

The WISE $W1-W2$ color at which the ISAAC and *WISE* fluxes match is similar to $W1-W2$ color used to select AGN in WISE. Stern et al. (2012) define an AGN selection color cut of $W1-W2 \geq 0.8$ mag. We plot the Stern et al. (2012) criterion in Fig. 2.9 as a black dashed line for comparison. This cut was motivated by the desire to find high-redshift AGN ($z \gtrsim 3$), so it is unsurprising that local AGN may exhibit much bluer $W1-W2$ colors. Indeed, Mateos et al. (2012) show that $W1-W2$ is redder for high-redshift AGN.

In the M -band we find similar trends, albeit with significantly more scatter. The Sy1 sample shows a match between M and $W2$ when $W1 - W2 \geq 0.78 \pm 0.02$ mag; while the Sy2 sample does the same at the much redder $W1 - W2 \geq 1.19 \pm 0.02$ mag. For Sy1 and Sy2 galaxies we estimate the M -band nuclear flux from the $W2$ flux as

$$\log_{10} F_{M,\text{Sy1}} = \quad (2.9)$$

$$\begin{cases} \log_{10} F_{W2} - 0.23 \pm 0.03, & \text{if } (W1 - W2) \geq 0.78 \pm 0.02 \\ \log_{10} F_{W2} + (1.48 \pm 0.15)(W1 - W2) - (1.38 \pm 0.10), & \text{else} \end{cases}$$

$$\log_{10} F_{M,\text{Sy2}} = \quad (2.10)$$

$$\begin{cases} \log_{10} F_{W2} - 0.16 \pm 0.01, & \text{if } (W1 - W2) \geq 1.19 \pm 0.02 \\ \log_{10} F_{W2} + (0.73 \pm 0.03)(W1 - W2) - (1.03 \pm 0.03), & \text{else.} \end{cases}$$

We find that in both bands, Circinus is ‘‘underluminous’’ relative to the general trend. It is likely here that we are over-resolving emission that *WISE* cannot spatially distinguish, as it is at a distance of only 4 Mpc. Another of the outliers, NGC 4355 (a.k.a. NGC 4418), is a compact obscured nucleus (CON) known to exhibit an SED unusual for Sy2 galaxies (e.g., Costagliola et al., 2011; Ohyama et al., 2019). We also

note that NGC 5953 (the extremely “blue” galaxy above) is underluminous for either the Sy1 or Sy2 trends, but roughly agrees with other Cp AGN in our sample.

In summary, for Sy1 with *WISE* colors $W1-W2 \geq 1.05$ mag the *WISE* *W1* flux is a good proxy for the *L* flux at the VLT. For Sy2 with *WISE* colors $W1-W2 \geq 0.90$ mag the *WISE* *W1* flux is a good proxy for the *L* flux at the VLT. We present functions which use the *W1-W2* color and the *W1* (or *W2*) flux to reliably estimate the *L*-band (or *M*-band) fluxes observable at the VLT with instruments such as VLTI MATISSE, ERIS, and CRIRES.

6.1 Potential VLTI/MATISSE Targets

The short atmospheric coherence times in the *L*, *M*, and *N* bands severely limit the sensitivity of self-tracked interferometric observations in these bands with VLTI/MATISSE. This necessitates bright targets for detailed study. The recommended limits for *L*-band target flux are 75 mJy when using the UTs and 1 Jy when using the ATs. Using these limits, we can identify 13 potential MATISSE targets in our sample, 2 of which are observable with the ATs. Most of these (10/13) have been previously observed either with MIDI or already with MATISSE as part of commissioning and/or guaranteed time observations. Fortunately, we can use the above formulae to estimate *L*-band fluxes from various *WISE* catalogs and identify potential targets. Here a potential target means *L*-band fluxes greater than those listed above and declination $\delta < 20^\circ$.

Using the MIR photometrically-selected *WISE* AGN Candidates Catalogs (Assef et al., 2018), we apply Eq. 2.7 to estimate the *L*-flux. We use the Assef et al. (2018) R90 catalog (~ 4 million AGN), which has 90% reliability, to identify 57 potential targets, 2 of which can be observed with the ATs.

In addition to using photometrically selected catalogs, we cross-matched the Véron-Cetty & Véron (2010) AGN catalog with the ALLWISE catalog to create a sample of optically classified AGN that had been detected using *WISE*. Using then the *W1* and *W2* fluxes, we computed the *L*-band fluxes using Eq. 2.7. In this way, we identify 44 AGN suitable for follow-up, 4 of which can be observed with the ATs. Many of the AGN in the Véron-Cetty & Véron (2010) catalog are not included in color-selected *WISE* AGN catalogs (e.g., Assef et al., 2018; Secretst et al., 2015) because the color cuts employed focus on high-redshift sources (e.g., Mateos et al., 2012). Thus this second approach is helpful if one wishes to focus on nearby AGN.

7 Summary and Conclusions

In this work we present a MIR flux catalog of 119 AGN, extending the work done by Asmus et al. (2014) from the *N*- and *Q*-bands to the *L*- and *M*-bands. This is the largest existing subarcsecond catalog of AGN in these bands. Each AGN was observed using VLT ISAAC in at least one of the *L*- and *M*-bands between 2000 and 2013. We include local ($z < 0.3$) AGN of 5 optical classifications: 21 are Seyfert 1; 5 are Intermediate Seyferts; 46 are Seyfert 2; 29 are LINERs; and 16 are so-called ‘Cp’ or Composites. We report two tables: one with nuclear fluxes in the *L*- and *M*-bands, and one with resolved emission fluxes, sizes, and PAs in both bands. The nuclear and resolved emission were separated by fitting one Gaussian to the PSF and one to the flux on scales larger than the PSF. We detect 98 sources in the *L*-band and 81 sources in the *M*-band. We found resolved *L*-band emission in 73 of the AGN.

We compared the flux ratios F_L/F_M and F_M/F_N to those predicted by several suites of AGN torus models: CAT3D (Hönig & Kishimoto, 2010), CAT3D-WIND (Hönig & Kishimoto, 2017), CLUMPY (Nenkova et al., 2008c), SKIRTOR (Stalevski et al., 2016),

and SKIRTOR models with polar winds. We find that the inclusion of a polar wind component significantly improves the agreement between our measurements and model predictions. CAT3D-WIND provides the best overall match to our measurements. The SKIRTOR (Stalevski et al., 2016) models fit the data much better after the inclusion of a polar wind. Notably, we find that CLUMPY (Nenkova et al., 2008c) provides good fits to the data without polar winds. No set of models, however, produces values of $F_L/F_M \gtrsim 1$, which we measure in 10 Seyfert galaxies. We discuss two possible explanations of this: stellar contamination in the relatively large physical scales probed; and underestimation of L -band accretion disk flux in torus models.

We favor the stellar contamination hypothesis, as several of the AGN in this sample (e.g., NGC 5252) were shown in Alonso-Herrero et al. (2001) to have up to 65% of their L -band emission come from stellar sources rather than the AGN. We also use K -band flux measurements of 21 cross-matched AGN from Burtscher et al. (2015) to estimate the L -flux contribution from stars using the Assef et al. (2010) $K - L$ colors for old stellar populations. After removing this estimated stellar contribution, all of the AGN become consistent with CAT3D-WIND, CLUMPY, and SKIRTOR+wind. We find that the stellar corrections necessary to match CLUMPY are $\approx 10\%$ larger than for the torus+wind models.

We measured the effect of changing the accretion disk spectrum in the SKIRTOR models to a QSO-motivated and relatively bluer spectrum. We find that while it does improve the agreement between these models and the observations, it does not fully reproduce the observed F_L/F_M . Nonetheless, follow-up observations of these NIR “blue” AGN may provide insight into the mechanisms driving the 3-5 μ m bump and the formation of these dusty tori.

We lastly derived relations between the reported WISE $W1$ and $W2$ fluxes and the L and M fluxes observable at the VLT. These relations (Eqs. 2.7-2.10) can be used to estimate the NIR fluxes for sources not included in this survey, using only the WISE $W1$ and $W2$ bands. This is especially useful for potential VLTI/MATISSE targets, allowing one to determine their observability and the possibility of resolving nuclear extended dusty structures.

This MIR AGN atlas holds a significant portion of local AGN of all optical classifications. It represents a statistically relevant sample suitable for AGN unification studies and interferometric follow-up. The need for such interferometric follow-up with VLTI/MATISSE is evinced by the fact that even our high-resolution 8.2 m telescope data has a significant fraction of sources with unresolved MIR dust emission.

Chapter 3

Mid-Infrared Interferometry

The resolving power of a telescope is directly related to the size of its primary mirror: the bigger the better. Light reflected by the circular mirror interferes with itself due to path length differences from various points on the mirror to the image plane, undergoing diffraction. The smallest measurable separation between two observed point sources (e.g., distant stars) in such a system for with circular aperture is described by

$$\sin \theta = 1.22 \frac{\lambda}{D}, \quad (3.1)$$

where θ is the angular resolution, λ is the wavelength of the observation, and D is the diameter of the telescope's primary mirror. This is known as the **Rayleigh criterion**. For the largest existing telescopes, $D = 10$ m, which yields an angular resolution of $\theta = 250$ mas at $10 \mu\text{m}$. Even the under-construction Extremely Large Telescope (ELT), with $D = 30$ m gives only 83 mas resolution. In the nearest AGN (~ 4 Mpc away) the ELT resolution probes a spatial scale of 1.6 pc, leaving crucial details of the circumnuclear dust unresolved. Additionally, only a handful of AGN are this close, and to understand the general AGN population we must look 10s of Mpc away where the entirety of the circumnuclear dust structure is unresolved at the $\gtrsim 4$ pc scale. There are two ways to resolve the resolution problem: build infeasibly large telescopes (with $D \gtrsim 50$ m) or use several telescopes in conjunction in a process called **interferometry**.

1 Interferometry Basics

The following derivations follow [Buscher & Longair \(2015\)](#) as well as lectures I attended at the Very Large Telescope Interferometer Summer School (2018, Lisbon).

1.1 Monochromatic Point Source

Astronomical interferometry is quite similar to Young's double slit experiment. It differs only in that the first screen with slits is replaced with two telescopes at positions \vec{x}_a and \vec{x}_b , and the second screen is replaced with the astronomical instrument. The separation vector between the two telescopes, also called the **baseline** is given by

$$\vec{B} = \vec{x}_b - \vec{x}_a. \quad (3.2)$$

The two telescopes are pointing at a source with position \vec{S} relative to the center of the baseline. The pointing direction of the telescopes is then described by the unit vector \vec{s}

$$\vec{s} = \frac{\vec{S}}{|\vec{S}|}. \quad (3.3)$$

If we assume the emitting source is at a large distance, the electromagnetic field can be described as a plane wave with angular frequency $\omega = 2\pi\nu$. The electric field at each telescope a and b is then given by

$$\vec{E}_a = E_{a,0} e^{i\vec{k}\cdot\vec{x}_a - i\omega t} \quad (3.4)$$

$$\vec{E}_b = E_{b,0} e^{i\vec{k}\cdot\vec{x}_b - i\omega t}$$

$$\vec{E}_b = E_{b,0} e^{i\vec{k}\cdot\vec{x}_a + i\vec{k}\cdot\vec{B} - i\omega t}, \quad (3.5)$$

with the wave vector $\vec{k} = -k\vec{s} = 2\pi/\lambda\vec{s} = -\omega/c\vec{s}$. Without loss of generality, we can omit the common phase element $i\vec{k}\cdot\vec{x}_a$, yielding

$$\vec{E}_a = E_{a,0} e^{-i\omega t} \quad (3.6)$$

$$\vec{E}_b = E_{b,0} e^{-ik\vec{s}\cdot\vec{B}} e^{-i\omega t}, \quad (3.7)$$

which tells us that the relative phase shift between the two signals is related to the difference in the optical path $\vec{s}\cdot\vec{B}$. An additional phase shift is introduced through the distance from each telescope to the detector given by d_a and d_b . The electric field from each telescope at the location of the detector is then

$$\vec{E}_a = E_{a,0} e^{ikd_a} e^{-i\omega t} \quad (3.8)$$

$$\vec{E}_b = E_{b,0} e^{ikd_b} e^{-ik\vec{s}\cdot\vec{B}} e^{-i\omega t}. \quad (3.9)$$

We can now compute the average intensity from each telescope at the location of the detector:

$$\begin{aligned} I_{\text{avg}} &\equiv |E|^2 = \vec{E} \cdot \vec{E}^* \\ &= (\vec{E}_a + \vec{E}_b) \cdot (\vec{E}_a^* + \vec{E}_b^*) \\ &= \vec{E}_{a,0}^2 + \vec{E}_{b,0}^2 + \vec{E}_{a,0}\vec{E}_{b,0} e^{ik(d_a - d_b + \vec{s}\cdot\vec{B})} + \vec{E}_{a,0}\vec{E}_{b,0} e^{-ik(d_a - d_b + \vec{s}\cdot\vec{B})} \\ &= I_a + I_b + 2\sqrt{I_a I_b} \cos k(d_a - d_b + \vec{s}\cdot\vec{B}). \end{aligned} \quad (3.10)$$

Both I_a and I_b can be related to the source intensity I_{src} via $I_x = \eta_x I_{\text{src}}$, in which η_x captures differences in telescope light collecting power, filter efficiency, etc. The **optical path difference** (OPD), written as $\delta \equiv d_a - d_b + \vec{s}\cdot\vec{B}$, encapsulates the internal phase shifts caused by the path from the telescopes to the detector ($d_a - d_b$) as well as the phase shifts caused by the position of the source and the pointing of the telescope ($\vec{s}\cdot\vec{B}$). Keeping this in mind, we can simplify Eq. 3.10 to

$$I_{\text{avg}} = I_{\text{src}}(\eta_a + \eta_b + 2\sqrt{\eta_a \eta_b} \cos \delta k). \quad (3.11)$$

This is the average intensity of the interfered light reaching the detector after traveling through telescopes a and b .

The above derivation was for an image plane or **Fizeau interferometer**, meaning that the light from each telescope is focused directly on the detector and the interference occurs there. This is analogous to Young's double-slit experiment. If the light is combined at some pupil plane before the detector, then the system is called a **Michelson**

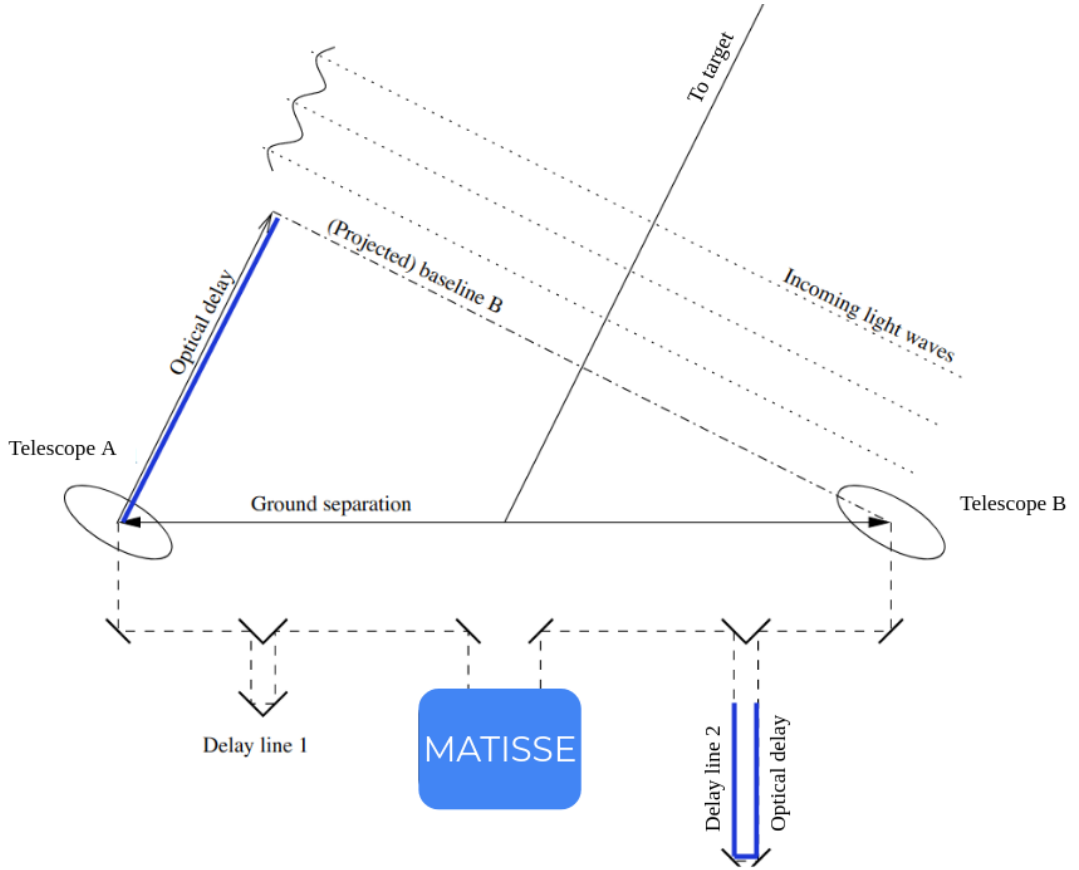


Fig. 3.1: Schematic of an interferometric system consisting of two telescopes separated by a baseline B . The delay lines compensate for the optical path difference introduced by the angle of the incoming light. The light ultimately enters the interferometric instrument, in this case MATISSE. This schematic is a modified version of one by Roy van Boekel (priv. comm.).

interferometer because it was the setup employed in the famous Michelson and Morely experiment which measured the speed of light. The combination of light in the pupil plane is performed with a **beam combiner**, a transmissive mirror which then aims the combined light toward the detector. This is the approach used at the Very Large Telescope Interferometer. The above derivation must be slightly modified due to the reflection occurring in the beam combiner. After passing through the beam combiner, the electric fields have the form

$$\vec{E}_1 = \sqrt{t}\vec{E}_a + \sqrt{r}\vec{E}_b e^{i\pi/2}$$

$$\vec{E}_2 = \sqrt{r}\vec{E}_a e^{i\pi/2} + \sqrt{t}\vec{E}_b, \quad (3.12)$$

$$(3.13)$$

where the factor $e^{i\pi/2}$ comes from the reflection of the light in the (typically half-silvered) mirror, and the factors t and r represent the transmission and reflection coefficients, respectively. This light then reaches the detector with intensities

$$\begin{aligned} I_1 &= tI_a + rI_b + 2\sqrt{trI_aI_b} \sin k\delta \\ I_2 &= rI_a + tI_b - 2\sqrt{trI_aI_b} \sin k\delta. \end{aligned} \quad (3.14)$$

The two intensities are complementary (opposite in phase) due to the reflection, and

can be subtracted to find the total intensity, I :

$$\begin{aligned} I &= I_1 - I_2 = (t - r)I_a + (r - t)I_b + 4\sqrt{trI_aI_b} \sin k\delta \\ &= (t - r)(I_a - I_b) + 4\sqrt{trI_aI_b} \sin k\delta \\ &= I_{\text{src}}[(t - r)(\eta_a - \eta_b) + 4\sqrt{tr\eta_a\eta_b} \sin k\delta]. \end{aligned} \quad (3.15)$$

In an ideal interferometer the telescopes have the same efficiency, and $\eta_a = 1 = \eta_b$. An ideal beam combiner will have $t = 0.5 = r$. Substituting in these values causes the first term, the background intensity $I_{\text{bck}} = I_{\text{src}}(t - r)(\eta_a - \eta_b)$ to vanish, leaving only the interferometric component:

$$I_{\text{int}} = 4I_{\text{src}}\sqrt{tr\eta_a\eta_b} \sin k\delta \quad (3.16)$$

$$I_{\text{int}} = 2I_{\text{src}} \sin k\delta. \quad (3.17)$$

As stated above, the pair of telescopes a and b make up a baseline. A schematic of the interferometric system consisting of two telescopes, the resulting optical path difference, and finally an interferometric instrument is given in Fig. 3.1. The derivation in this section applies to each pair of telescopes in an array, and the number of measured baselines increases rapidly with the number of telescopes. More specifically, for an array of N telescopes, there are $\binom{N}{2} = \frac{N!}{2(N-2)!}$ baselines measured simultaneously. For the MATISSE instrument with $N = 4$ this means there are 6 sets of interference fringes measured per observation.

1.2 Resolving Power of the Interferometer

Consider two sources \vec{S} and \vec{S}' ; these sources can be distinguished when the peak of the fringe pattern of one source is found at or further than the first minimum of the fringe pattern of the second source. That is

$$k(d_a - d_b + \vec{s} \cdot \vec{B}) = k(d_a - d_b + \vec{s}' \cdot \vec{B}) + \pi, \quad (3.18)$$

which simplifies to $(\vec{s} - \vec{s}') \cdot \vec{B} = \lambda/2$. Now, using the definition of the wave number, $k = 2\pi/\lambda$, and by projecting the separation of the two sources onto the physical baseline, $(\vec{s} - \vec{s}') \cdot \vec{B} = \theta_{\text{Rayleigh}} B_{\text{proj}}$, we obtain the Rayleigh criterion for an interferometer

$$\theta_{\text{Rayleigh}} = \frac{\lambda}{2B_{\text{proj}}}. \quad (3.19)$$

The resolving power of the interferometer is directly related to the **projected baseline**, B_{proj} and to the wavelength of the observations. The UT baselines at the VLTI extend to 130.2 m, giving resolution limits of 9.5 mas at 12 μm and 2.8 mas at 3.5 μm , a vast improvement over both existing and planned single-dish telescopes.

1.3 Polychromatic Sources

In the preceding sections, it was assumed that the light entering the interferometer was monochromatic. It is, however, straightforward to consider light which is emitted from astronomical sources with a wide range of wave numbers, k , and which enters the telescope through filters with finite passbands. We can then modify Eq. 3.17 to represent the total intensity over all wave numbers k :

$$I_{\text{int}} = \int 4I_{\text{src}}(k)\sqrt{t(k)r(k)\eta_a(k)\eta_b(k)} \sin(k\delta) dk, \quad (3.20)$$

where the source intensity, filter efficiency, transmission efficiency and reflection efficiency are now allowed to vary with wavelength.

Top-hat filter functions can be seen simple approximations of real filters, including the L , M , and N bands used in MATISSE. It is therefore instructive to consider a source with constant intensity $I_{\text{src}}(k) = I_0$ passing through an ideal beam combiner ($t = 0.5 = r$) and a top-hat filter function centered on k_0 with width Δk and value η_0 inside $k_0 - \Delta k/2 < k < k_0 + \Delta k/2$ and 0 elsewhere. The integration in Eq. 3.20 becomes

$$I_{\text{int}} = 2I_0\eta_0 \int_{k_0 - \Delta k/2}^{k_0 + \Delta k/2} \sin(k\delta) dk \quad (3.21)$$

$$I_{\text{int}} = 2I_0\eta_0\Delta k \sin(k_0\delta) \frac{\sin(\Delta k/2\delta)}{\Delta k/2\delta}$$

$$I_{\text{int}} = 2I_0\eta_0\Delta k \sin(k_0\delta) \text{sinc}(\Delta k/2\delta). \quad (3.22)$$

This result is similar to 3.17 with the addition of the sinc function modulation. This means, moreover, that when $|\Delta k/2\delta| > 1$, this equation becomes small; the implications of this are discussed below in §1.4.

Finally, noting that the sinc function in Eq. 3.20 is the Fourier transform of the filter function, we can more generally express the polychromatic behaviour of the interferometer in terms of an arbitrary filter function $\eta(k)$:

$$I_{\text{int}} = 2I_0\eta_0\Delta k \sin(k_0\delta) \mathcal{F}[\eta(k)](\delta). \quad (3.23)$$

1.4 Coherence Length and Coherence Time

Returning to the example of a top-hat filter function, we see that when $|\Delta k/2\delta| > 1$, Eq. 3.22 becomes small. In practice, this means that fringes are only observed if the optical delay δ is smaller than some coherence length, Λ_{coh} . Let

$$1 = (k_0 - (k_0 - \Delta k/2))\Lambda_{\text{coh}} = \frac{2\pi}{\lambda_0} - \frac{2\pi}{\lambda_0 + \Delta\lambda} \Lambda_{\text{coh}}$$

$$1 = \frac{2\pi\Delta\lambda}{\lambda_0^2 + \lambda_0\Delta\lambda} \Lambda_{\text{coh}} \approx \Lambda_{\text{coh}} \frac{\Delta\lambda}{\lambda_0^2}$$

$$\implies \Lambda_{\text{coh}} \approx \frac{\lambda_0^2}{\Delta\lambda} = R\lambda_0, \quad (3.24)$$

where it is assumed that $\Delta\lambda \ll \lambda$, which is valid for the MATISSE bands, and the definition $R \equiv \Delta\lambda/\lambda$ is used. For the L -band ($3 < \lambda < 4 \mu\text{m}$), the coherence length is $\Lambda_{\text{coh}} \approx 12.25 \mu\text{m}$. Because the coherence length depends on the width of the band, dispersing the light with spectral resolution can vastly increase the coherence length, as $\Lambda_{\text{coh}} = R\lambda_0 - \lambda_0/4R$. For MATISSE, with spectral resolution $R = 30$ in the L -band, this increases the coherence length to $\Lambda_{\text{coh}} \approx 105 \mu\text{m}$. The coherence length is important because the delay lines must keep the beam pathlengths the same within this value, otherwise the observed signal diminishes in strength. In essence then, it sets the accuracy tolerance of the optical path in an instrument. High contrast fringes are observed when the accuracy of the delay lines is much smaller than the coherence length.

1.5 Extended Sources

The sources which are most interesting to observe are not pure point sources, and instead have polychromatic, extended flux. These extended sources can, however, be thought of as a superposition of many point sources which interfere. Similar to the polychromatic case then, one can describe an extended source at a given wavelength as

$$I_{\text{int}} = \int 4I_{\text{src}}(\vec{s}) \sqrt{tr\eta_a\eta_b} \sin(k(d - \vec{s} \cdot \vec{B})) d\vec{s} \quad (3.25)$$

(in the more general, polychromatic case, I_{src} , t , r , η_a , and η_b are all functions of wave number k and the integral becomes a double integral also evaluated over dk). Now, one can use the Euler formula to expand the sin term into exponentials

$$I_{\text{int}} = \frac{2\sqrt{tr\eta_a\eta_b}}{i} \left[e^{ikd} \int I_{\text{src}}(\vec{s}) e^{ik\vec{s}\cdot\vec{B}} d\vec{s} - e^{-ikd} \int I_{\text{src}}(\vec{s}) e^{ik\vec{s}\cdot\vec{B}} d\vec{s} \right]. \quad (3.26)$$

Each integral term is a Fourier transform of the source intensity distribution $I_{\text{src}}(\vec{s})$:

$$\mathcal{I}(\vec{B}_{\text{proj}}/\lambda) = \int I_{\text{src}}(\vec{s}) e^{ik\vec{s}\cdot\vec{B}} d\vec{s} = \int I_{\text{src}}(\vec{s}) e^{2i\pi\vec{s}\cdot\vec{B}/\lambda} d\vec{s} = \mathcal{F}[I_{\text{src}}(\vec{s})](\vec{B}_{\text{proj}}/\lambda). \quad (3.27)$$

It is often useful to consider individually the amplitude and phase of the complex value arising from the Fourier transform. This means that we can write $\mathcal{I}(\vec{B}_{\text{proj}}/\lambda) = \underline{I}(\vec{B}_{\text{proj}}/\lambda) e^{i\phi(\vec{B}_{\text{proj}}/\lambda)}$, where \underline{I} represents the absolute value and ϕ represents the phase of the complex function. They are both functions of **spatial frequency** $\vec{u} = \vec{B}_{\text{proj}}/\lambda$, which is commonly a two-dimensional vector with components $\vec{u} = (u, v) = (B_{\text{proj},x}/\lambda, B_{\text{proj},y}/\lambda)$. These aptly-named uv -**coordinates** are the position at which the Fourier transform of the source intensity is evaluated. Each telescope pair forms a baseline which is projected onto the uv -plane, and more observed baselines means a denser sampling of $\mathcal{I}(\vec{B}_{\text{proj}}/\lambda)$, and crucially leads to a better characterization of the source intensity via modeling or imaging.

The Fourier transform of a real function such as $I_{\text{src}}(\vec{s})$ is Hermitian:

$$I(-\vec{u}) = \underline{I}(-\vec{u}) e^{i\phi(-\vec{u})} = \underline{I}(\vec{u}) e^{i\phi(-\vec{u})} = \mathcal{I}^*(\vec{u}). \quad (3.28)$$

This implies that if the two telescopes making up B_{proj} are swapped, then only the phase of the measured fringe pattern changes sign and the intensity remains unchanged.

The resulting intensity is often called **complex correlated flux**. It is, however, sometimes useful (for example in imaging) to consider a normalized quantity referred to most commonly as the **complex visibility**. The visibility is the correlated flux normalized by the correlated flux on the zero-length baseline (i.e., the total flux of the source),

$$V(u, v) = \frac{\mathcal{I}(u, v)}{\mathcal{I}(0, 0)}. \quad (3.29)$$

The above derivation has led us to the statement,

The visibility of a source is equal to the normalized Fourier transform of the source intensity distribution.

This is the essence of the **van Cittert-Zernike theorem**, in the limit that the incoming light is spatially incoherent (van Cittert, 1934; Zernike, 1938), a valid assumption of most astronomical sources due to their distance. It is therefore straightforward to calculate the visibility on an arbitrary baseline for an arbitrary source distribution. This calculation is a crucial part of forward modeling and imaging, as it is possible to simply change an input intensity distribution and compare it to the observed visibilities. An example of several model images and their corresponding visibilities are shown in Fig. 3.2, illustrating the connection between an extended flux distribution and its observed visibilities.

1.6 Closure Phases

Following the van Cittert-Zernike theorem, the complex visibility of the source is simply the Fourier transform of the source flux distribution on the sky. While the Fourier transform of the source flux distribution is a complex number with amplitude and phase, so

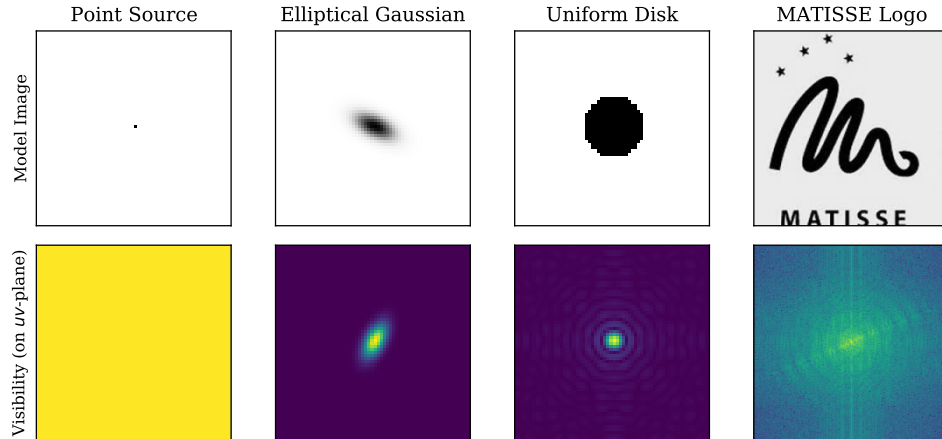


Fig. 3.2: Image and visibility pairs illustrating the van Cittert-Zernike theorem. The various model images are in the (*top*) row, and their corresponding visibilities from the (normalized to 1) Fourier transforms are shown in the (*bottom*) row.

far only the amplitude has been considered. The phase component is crucial for imaging, but it is highly sensitive to e.g., atmospheric turbulence and telescope imperfections. The phase is therefore unstable, and the short atmospheric coherence times and small isoplanatic angles in the mid-infrared make it nearly impossible to calibrated directly. Instead, a quantity called the **closure phase** is introduced which cancels out the effects from the atmosphere and from the individual telescopes.

The closure phase was developed in radio astronomy to compensate for poor phase stability in early radio interferometers (e.g., [Jennison, 1958](#)), but it is still widely used today for self-calibration. Since then it was shown that this effect also applies at shorter wavelengths such as the mid-infrared (first by [Baldwin et al., 1986](#)). The below derivation follows largely from [Monnier \(2007\)](#).

The measured electric field at a telescope a is considered to be

$$\tilde{E}_a^{\text{measured}} = \tilde{G}_a \tilde{E}_a^{\text{source}} = |\tilde{G}_a| e^{i\psi^G} \tilde{E}_a^{\text{source}}, \quad (3.30)$$

where \tilde{G}_a carries all telescope specific modifications to the electric field. The amplitude of \tilde{G}_a represents changes to the intensity from things such as poor adaptive optics corrections, mirror reflectivity, etc. The phase of \tilde{G}_a represents phase shifts from primarily atmospheric turbulence. Recalling Eqs. 3.10 and 3.29 which together show for baseline $a - b$ that $V_{ab} \propto \tilde{E}_a \tilde{E}_b^*$. The interfered light then takes the form

$$\begin{aligned} V_{ab}^{\text{measured}} &= \tilde{G}_a \tilde{G}_b^* \tilde{E}_a^{\text{source}} \tilde{E}_b^{\text{source}*} = |\tilde{G}_a| |\tilde{G}_b| e^{i(\psi_a^G - \psi_b^G)} V_{ab}^{\text{source}} \\ &= |\tilde{G}_a| |\tilde{G}_b| |V_{ab}^{\text{source}}| e^{i(\psi_a^G - \psi_b^G + \phi_{ab}^{\text{source}})}. \end{aligned} \quad (3.31)$$

Here the atmospheric phase shifts for telescope x are written as ψ_x and the source intrinsic phase for the baseline $a - b$ is written as $\phi_{ab}^{\text{source}}$. This equation shows that a phase shift is introduced to the measured visibilities purely due to the individual telescope and atmospheric effects.

Consider three measured baselines $a - b$, $b - c$, and $c - a$ which form a *closure triangle*. Any phase shift introduced by telescope a affects all of the baselines which include a but only those baselines. Moreover, the phase shift introduced from a on baseline $a - b$ is equal in magnitude but opposite in sign from that on $c - a$. There is no phase shift from a on baseline $b - c$. We then consider the sum of the measured phases from each baseline, ϕ_{xy} , that make up the closure triangle. This sum is the definition of the closure

phase ϕ_{abc} :

$$\begin{aligned} \phi_{abc} &= \phi_{ab}^{\text{measured}} + \phi_{bc}^{\text{measured}} + \phi_{ca}^{\text{measured}} & (3.32) \\ &= \phi_{ab}^{\text{source}} + (\psi_a - \psi_b) + \phi_{bc}^{\text{source}}(\psi_b - \psi_c) + \phi_{ca}^{\text{source}} + (\psi_c - \psi_a) \\ &= \phi_{ab}^{\text{source}} + \phi_{bc}^{\text{source}} + \phi_{ca}^{\text{source}} + (\psi_a - \psi_a) + (\psi_b - \psi_b) + (\psi_c - \psi_c) \\ \phi_{abc} &= \phi_{ab}^{\text{source}} + \phi_{bc}^{\text{source}} + \phi_{ca}^{\text{source}} & (3.33) \end{aligned}$$

and the telescope specific effects cancel out. The closure phase is considered a “good quantity” because the observed value is always equal to the intrinsic value from the source.

To measure a single closure phase, the interfered fringes on three baselines (corresponding to three telescopes) must be measured *simultaneously*. More generally, the number of closure phases that can be measured with an interferometric array of N telescopes is $\binom{N}{3}$, but only $\binom{N-1}{2}$ of those are independent. Closure phases were not measurable with MIDI ($N = 2$) because it measured one baseline at a time. MATISSE on the other hand with $N = 4$ telescopes records three independent closure phases and a fourth which is a combination of the other three. The ability to measure closure phases was a key goal of the current generation of mid-infrared interferometers because phase information is crucial for imaging. More specifically, closure phases have several key properties which inform us about the source flux distribution:

- Closure phases are independent of translations of the image, and instead measure internal phase information.
- Closure phases are 0° or 180° for point-symmetric images. Intermediate values measure the *skew* or asymmetry of the source flux.
- Closure phases are only non-zero (or non- 180°) when the target is resolved.

While it is difficult to interpret closure phases directly these properties are used in the image reconstruction or modeling process to inform the final image.

2 MATISSE Instrument

The Multi AperTure mid-Infrared Spectro-Scopic Experiment (MATISSE) at the Very Large Telescope Interferometer is the state of the art instrument in mid-infrared interferometry (Lopez et al., 2022). It was developed jointly by Laboratoire Lagrange (Université Côte d’Azur, Observatoire de la Côte d’Azur); the Max-Planck-Institut für Radioastronomie; Universität Kiel; Universiteit Leiden; NOVA (ASTRON: Netherlands Institute for Radio Astronomy); Universität Wien; and the **Max-Planck-Institut für Astronomie**. MATISSE has been in operation at the Cerro Paranal site in Chile since 2018 and was opened for science operations in April 2019. It combines the light from either the four 8.2 m diameter unit telescopes (UTs) or from four 1.8 m auxiliary telescopes (ATs), measuring the interfered signals in the L , M , and N bands (2.8-4.2 μm , 4.5-5.0 μm , and 8.0-13.0 μm , respectively). The LM and N band fringes are recorded simultaneously on two different detectors. Though the M band is measured, the instrument is optimized for the L and N bands.

2.1 The Very Large Telescope Interferometer

The Very Large Telescope Interferometer (VLTI) is located atop the Cerro Paranal mountain in the Republic of Chile. This site, at $24^\circ 40' \text{ S}$, $70^\circ 25' \text{ W}$, is within the Atacama Desert and is therefore one of the most arid observatories on the planet. The




Fig. 3.3: (*left*) Photo taken from the telescope platform showing an AT (white, in front) and a UT (silver, in back). (*right*) the optical tunnels underneath the telescope platforms which bring the light beams from all the telescopes together.

dry air, coupled with the high altitude of the observing platform (2635 m), makes the site exceptionally transparent for infrared observations where water vapor causes large amounts of absorption. The VLTI is 135 km from the nearest city, Antofagasta; its remoteness means that there is minimal light pollution and minimal atmospheric dust from human activities. The VLTI is owned and operated by the **European Southern Observatory** (ESO), which is a multinational organization for astronomical research made up of 16 member states and two strategic partners (Chile and Australia).

The VLTI is currently home to three interferometric instruments in the near and mid infrared: PIONIER (studying the H band), GRAVITY (studying the K band) and MATISSE (studying the L , M , and N bands). These instruments use either the four 8.2 m UTs, which are fixed in position and have physical baselines¹ ranging from 46 to 130 m, or four 1.8 m ATs, which can move to 30 different stations and have baselines ranging from 8 to 200 m. The ATs, however, are often fixed in four set configurations: small, medium, large, and astrometric. Connecting the telescopes is a series of tunnels which contain optical delay lines and the optics to bring the light to each instrument. Fig. 3.3 shows an AT in comparison to a UT as well as a photo of the optical tunnels running underneath them. The instruments themselves are located in the interferometric laboratory. The UTs and ATs are each served by adaptive optics (AO) systems in order to reduce the effects of atmospheric dispersion (MACAO on the UTs and NAOMI on the ATs) before the light is sent to the interferometric instruments.

2.2 Instrument Setup

The MATISSE instrument is made of several key components: “warm” optical elements at ambient temperature (called the WOP) and two separate cryostats (one each for the LM and the N bands) which contain cooled optics (called the COB) and the detectors. A diagram of the instrument is given in Fig. 3.4, and its components are described

¹The projected baselines can be shorter, and they depend on the source altitude and azimuth.




Fig. 3.4: Schematic of the MATISSE instrument. Red elements are at ambient temperature and blue elements are on the cold optical bench cooled by the cryostat. Figure from [Lopez et al. \(2014\)](#).

below.

The WOP receives the four VLTI beams from the UTs or ATs, and it first routes them to two independent beam commuting devices (BCDs). A BCD can be in either ‘OUT’ mode (meaning it is out of the optical path) or ‘IN’ mode (meaning the beams are commuted). MATISSE uses four BCD configurations: ‘OUT-OUT’, ‘OUT-IN’, ‘IN-OUT’, and ‘IN-IN’. The BCDs are used to change the order of the fringes on the detector which allows one to estimate and remove detector-specific features. The beams are then spectrally separated using dichroics in order to form the LM -band and the N -band beams. Before being passed to the cryostats, the beams are modulated using piezo actuators to ensure that the path lengths for all beams in a band are the same. Finally, the WOP includes “internal” sources used for alignment and spectral calibration, which deliver beams identical to the VLTI beams.

The $L + M$ -band and N -band cryostats are very similar, and they cool the optics to < 40 K. Within the cryostats, the beams pass through spatial filters (e.g., a slit or pinhole) before reaching the beam splitters. The beam splitters are optionally used; when in use they split the beams into interferometric and photometric channels. This is decided by the user for each observation based on the goals of the project. Next, the light reaches the spectral filters and dispersion elements and passes through them to finally reach the detector, where the beams interfere. The AQUARIUS detector used for the N -band operates at 10 K, and the HAWAII-2RG detector used for the L and M bands operates at 40 K.

2.3 Observing Procedure

MATISSE operates in both the LM - and N -bands simultaneously. In each band, it is possible to either send all the photons to the interferometric channel (the so-called “High-

Sens mode”) or to split them into photometric and interferometric channels (“SiPhot mode”). In the N -band, SiPhot mode is not offered because the detector behaves non-linearly and prevents proper Kappa matrix estimation. Photometry is instead offered in the N -band using the interferometric channels; following the interferometric exposures, one shutter at a time is opened to record the photometry of each telescope. Therefore, MATISSE’s standard operation is called a “Hybrid mode,” wherein the N -band uses HighSens and the LM -band uses SiPhot.

A standard observational block (OB) generates 14 exposures in each band. These consist of many shorter detector integration times (DITs): in the N -band 20 ms is used for low spectral resolution, and 75 ms is used for high spectral resolution; in the LM -band the user can select DITs from 75-125 ms. If an external fringe tracker (such as GRA4MAT) is used, the DIT can go up to 10 s.

The standard HighSens OB proceeds as follows:

1. One 30 s sky exposure is recorded for each the OUT-OUT and IN-IN BCD configurations.
2. Four 60 s interferometric exposures are then taken, using each of the four BCD configurations and without chopping.
3. Eight 60 s photometric exposures are taken with chopping.

Here chopping refers to alternating between taking exposures of the target and exposures of an empty patch of sky. Chopping is used to subtract the (large) contribution of the sky to the total flux in the infrared.

In SiPhot mode, the standard procedure is

1. One 30 s sky exposure is recorded for each the OUT-OUT and IN-IN BCD configurations.
2. Four 60 s interferometric + photometric exposures are then taken, using each of the four BCD configurations and without chopping.
3. Eight total 60 s interferometric + photometric exposures are then taken, using each of the four BCD configurations twice and with chopping.

In both cases, approximately 13 minutes of “open shutter” exposure happens during an OB. An OB typically lasts 27 minutes, including pointing the telescopes, starting adaptive optics guiding, and finding the fringes for fringe tracking. The interferometric exposures within an OB (step 2 in both modes above) can be repeated a number N_{cycles} times to increase the exposure time without adding telescope overheads. The recommended observing pattern in a night alternates between observing the calibration star (CAL) and the science target (SCI), and is often referred to CAL-SCI-CAL. Often the calibrators used before and after the science target are different, one used for the LM -bands and one for the N -band.

Finally, during the observations, fringes must be searched for, detected, and automatically tracked. The MATISSE Near-Real Time Software (NRTS) keeps the fringes within a fraction of the coherence length ($\Lambda_{\text{coh}} \approx \lambda^2/\Delta\lambda = R\lambda$, where λ is the observing wavelength and R is the spectral resolution, see §1.4). The NRTS uses the dispersed fringe tracking algorithm of [Koechlin et al. \(1996\)](#), which in essence cleans each interferometric frame using the mean sky, takes the 2D Fourier transform of this, and locates the coherently integrated fringes. The spacing of these fringes is converted from pixels to optical path delay (OPD), and the optics are adjusted to prevent the fringes from moving. From these estimates, a so-called “waterfall diagram” is made, which shows

the fringe peak positions as a function of time. It is extremely useful for diagnosing the quality of the observation, and the first appearance of fringes on this diagram is celebrated heartily each time it happens.

An example of the observing interface is shown in Fig. 3.5, taken during observations of ESO 323-G77. It shows the observing sequence on the left, the waterfall diagram in the middle (in addition to an interface to manually search for fringes), information about the beams to the right of that, and finally it shows the instantaneous detector readout on the far right. It serves as a nice summary of the above paragraphs.

2.4 Data Reduction and Data Products

MATISSE has an official ESO-supported data reduction pipeline. This data reduction software (DRS) produces the calibration maps, reduces the raw interferograms and puts them in OIFITS files, calibrates the data, and can produce reconstructed images. The following algorithmic information is collected from the MATISSE reference paper (Lopez et al., 2022) and the MATISSE Pipeline Manual² (which I assisted in writing).

The DRS can produce and use premade bias, flat-field, non-linearity, and bad-pixel maps to calibrate the raw frames as is typical for single-dish telescopes' charge-coupled device (CCD) calibrations. Additionally, a *shift map* is used to calibrate the spectral distortion for a given spectral resolution and a *kappa matrix* is used to compute the projection between the interferometric and photometric channels, the intensity ratio between interferometric and the photometric channels, and the spatial shift. The CCD calibration and data reduction process are each performed independently for the *LM* and *N* bands, primarily because of the different detectors used in each band (HAWAII-2RG in the *LM* and AQUARIUS in *N*) which have different readout noise properties and sensitivities.

Following the CCD calibrations, a 1D Fourier transform of the cleaned interferometric signal, $I(u)$, is given by

$$I(\vec{u}) = M_B(\vec{u}) \sum_{i=1}^4 n_{Bi} + M(\vec{u}) \sum_{i=1}^4 n_{*i} + \sum_{i=1}^4 \sum_{j=2, j>i}^4 M(\vec{u} - \vec{u}_{ij}) \sqrt{n_{*i} n_{*j}} V_{ij}. \quad (3.34)$$

Here n_{Bi} is the number of thermal background photons in beam i , n_{*i} and n_{*j} are the numbers of photons produced by the target object in beams i and j , $M_B(\vec{u})$ is the low frequency peak from thermal background emission, $M(\vec{u})$ is the low frequency peak from the interferometer itself, $M(\vec{u} - \vec{u}_{ij})$ is the fringe peak of the interferometer at spatial frequency $\vec{u}_{ij} = \vec{B}_{ij}/\lambda$ (\vec{B}_{ij} is the baseline from telescope i to telescope j), and finally V_{ij} is the complex visibility. It is important to note that $M_B(0) = M(0) = 1$ and that the third term represents the coherent fluxes C_{ij} of the six baselines measured by MATISSE. This 1D Fourier transform can be applied at each wavelength to construct the 2D fringe pattern. The DRS removes the low frequency peaks (the first and second terms of Eq. 3.34) using a multi-stage modulation technique. Background-removal is a crucial step because the thermal background is variable and often greatly exceeds the target coherent flux. The stages are chopping, spatial modulation, and temporal modulation; the exact process is outlined in Petrov et al. (2007).

Following the background removal, the interferometric observables can be estimated either coherently or incoherently. Incoherent estimators require no knowledge of the optical path delay (OPD) modulation and are also called *speckle-like*. Incoherent estimators work well for bright sources, but perform poorly for faint sources because the phase shifts due to the OPD cause the signal to be washed out. Coherent estimators, on

²The MATISSE Pipeline Manual can be found at https://www.eso.org/sci/facilities/paranal/instruments/matisse/doc/MATISSE_USERMANUAL.pdf

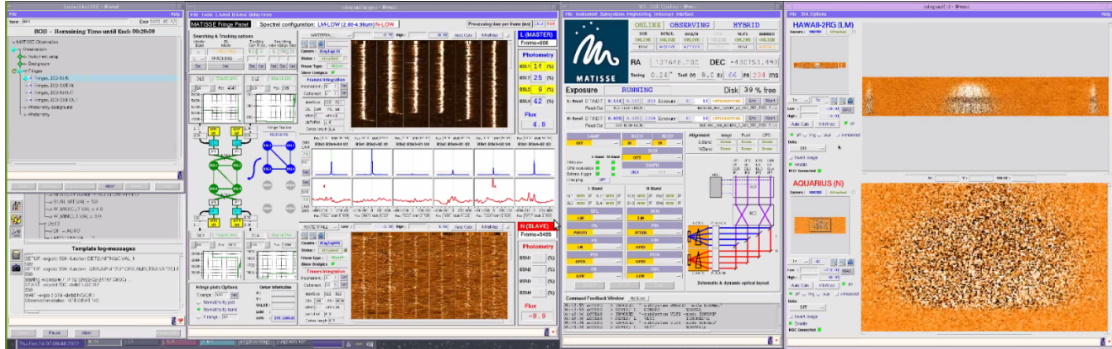


Fig. 3.5: Example observers interface from the observation of ESO323-G77. It shows the observing sequence on the left, the waterfall diagram in the middle (in addition to an interface to manually search for fringes), information about the beams to the right of that, and finally it shows the instantaneous detector readout on the far right.

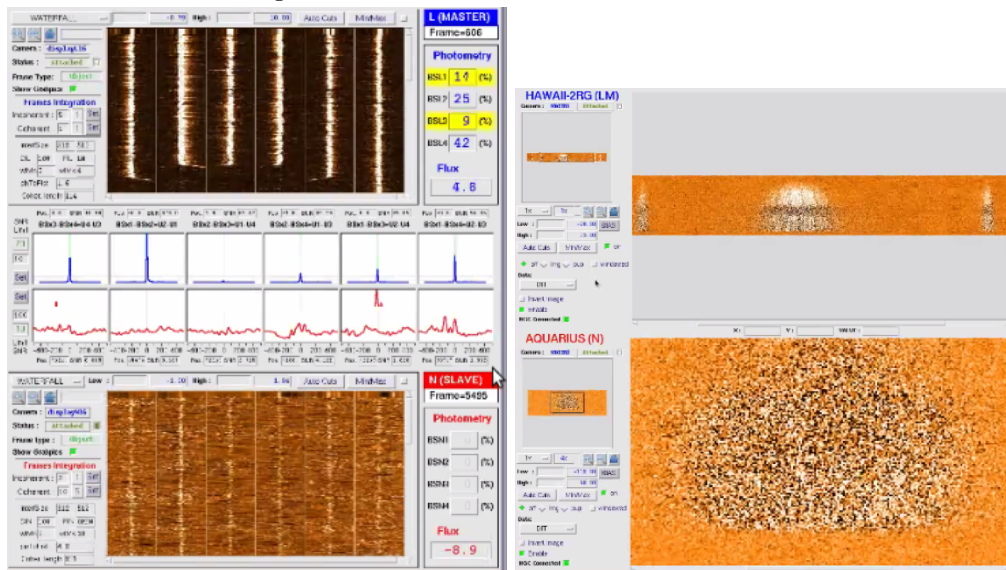


Fig. 3.6: (left) Cutout showing the waterfall diagram. On the top the six L -band fringe tracks are shown, in the middle instantaneous slices of the six fringe tracks are shown for the L -band in blue and the N -band in red, and on the bottom are the six N -band fringe tracks. These plots update every few seconds, filling from the top and cascading downward (hence the name waterfall). This allows the observer to see both the instantaneous status and a history of the fringe tracking in order to judge the quality of the observations. (right) The instantaneous detector readout for both the LM -band (top) and N -band (bottom). The interferometric channel is in the center of each panel (occasionally fringes can be seen by eye, but only in bright sources) and flanking it are the four photometric channels (though they are difficult to see for this faint source).

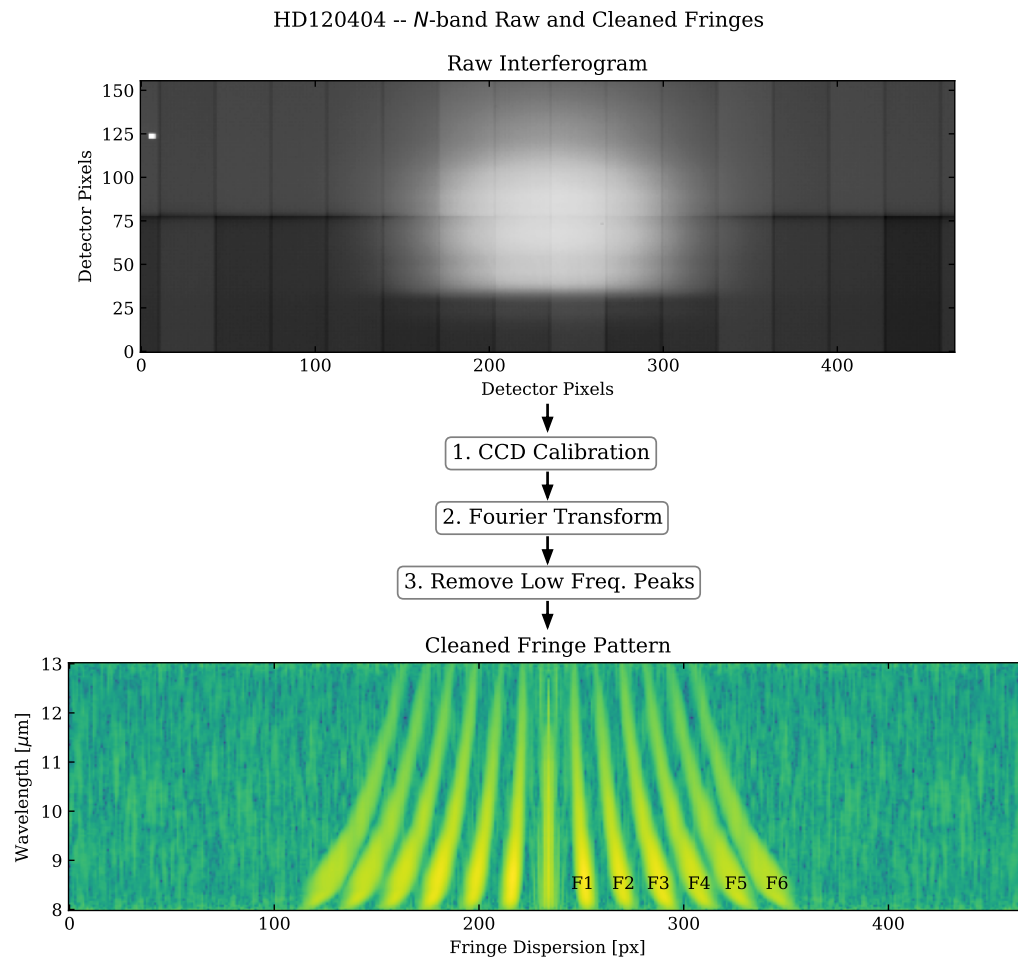


Fig. 3.7: Raw interferogram (*top*), the data processing steps (*middle*), and the cleaned fringes (*bottom*) for HD120404. Note the grid-like pattern on the detector which must be removed. The fringe pattern is mirrored about the center, and so the six fringes F1, F2, F3, F4, F5, and F6 are each shown twice. Both images are in log-scaling.

the other hand, take the OPD modulation into account in order to add the co-phased signal constructively. Throughout an observation, a number of interferometric frames are taken. For both coherent and incoherent methods, the above procedure is repeated for *each frame* independently to build up a series of 2D Fourier transforms (Eq. 3.34 for each wavelength) referred to as the *fringe patterns*. An example of the raw interferogram and its corresponding cleaned fringes for a bright calibrator (HD120404 with $F_{12 \mu\text{m}} = 13 \text{ Jy}$) is shown in Fig. 3.7. The interferometric observables are then extracted from the cleaned fringe pattern using the methods detailed below.

2.4.1 Incoherent Estimators

Incoherent estimators use the time-averaged fringe patterns from the entire observation, without correction for phase shifts due to the OPD.

Squared Correlated Flux: The squared correlated flux³ C_{ij}^2 for a baseline B_{ij} is given by

$$C_{ij}^2(\lambda) = \sum_{\vec{u}} \langle |I(\vec{u}, \lambda, t)|^2 - \beta \rangle_t, \quad (3.35)$$

where $\langle \dots \rangle_t$ is the time average, \vec{u} is the spatial frequency integrated between $(\vec{B}_{ij} - D)/\lambda$ and $(\vec{B}_{ij} + D)/\lambda$, D is the pupil diameter, and β is the bias present in each fringe peak (estimated from the value between fringe peaks). The pupil here refers to the spatial-filter pinhole with diameter $1.5\lambda/d$ in the L band and $2\lambda/d$ in the N band with d representing the telescope diameter.

Squared Visibility: The squared visibility is simply the squared coherent flux divided by the total photometric flux from either the UTs or the ATs used in the observation. Consider a photometric beam $P_i(\lambda, t)$ from each telescope i , then the average spectrum from each beam is given by $P_i(\lambda) = \langle P_i(\lambda, t) \rangle_t$. The squared visibility is computed using this value and the squared coherent flux:

$$V_{ij}^2(\lambda) = \frac{C_{ij}^2(\lambda)}{P_i(\lambda)P_j(\lambda)}. \quad (3.36)$$

Closure Phase: Recalling that the Fourier transform given by Eq. 3.34 is a complex number which has both an amplitude and a phase (the argument of the complex number), one can compute the closure phase for a telescope triplet ijk , ϕ_{ijk} in the following way:

$$\phi_{ijk}(\lambda) = \text{Arg} \left[\sum_{\vec{u}_1, \vec{u}_2} \langle I(\vec{u}_1, \lambda, t) I(\vec{u}_2, \lambda, t) I^*(\vec{u}_1 + \vec{u}_2, \lambda, t) \rangle_t - \gamma \right]. \quad (3.37)$$

The coordinates \vec{u}_1 , \vec{u}_2 , and $\vec{u}_1 + \vec{u}_2$ correspond to the baselines \vec{B}_{ij} , \vec{B}_{jk} , and \vec{B}_{ik} , respectively, integrated over the pupil width as for the squared correlated flux. The parameter γ represents the photon bias in the bispectrum. This photon bias contains both an additive and a multiplicative term due to the combination of the fringes. The estimation of this bias is non-trivial, and the method used in the DRS is given in e.g., [Gordon & Buscher \(2012\)](#). However, in the limit where read-noise is negligible and photon noise dominates, $\gamma \approx |\beta_{ij}|^2 + |\beta_{jk}|^2 + |\beta_{ik}|^2 - 2N$ where β_{ij} is the photon bias used in Eq. 3.35 for a baseline \vec{B}_{ij} and N is the mean number of photons in the interferogram ([Wirnitzer, 1985](#)). For bright sources, the γ term is negligible.

³The terms correlated flux and coherent flux are used interchangeably throughout this work.

2.4.2 Coherent Estimators

Coherent estimators rely on a reliable estimate of the residual atmospheric OPD so that phase shifts can be corrected. The fringe tracker of MATISSE minimizes this during the observation, but nonetheless a phase shift remains. The OPD in each frame is modeled as

$$\phi_{\text{atm}}(\lambda, t) \approx \phi_0(t) + 2\pi \frac{\delta(t)}{\lambda}, \quad (3.38)$$

where $\phi_0(t)$ is the achromatic phase due to atmospheric dispersion and $\delta(t)$ is the residual atmospheric phase shift. The term $\delta(t)/\lambda$ is often called the *group delay*. In the DRS, two terms $\phi_0(t)$ and $\delta(t)$ are estimated using a maximum likelihood fit to each frame based on Schutz et al. (2016).

Using the estimate of $\phi_{\text{atm}}(\lambda)$, the phase in each frame can be corrected by the term $e^{-i\phi_{\text{atm}}(\lambda, t)}$. Now we can estimate the correlated flux, visibility, and differential phase. The closure phase is only computed incoherently in the current implementation of the DRS, and the absolute phase is not useful without a phase calibrator (in the mid-infrared the sky is too variable to use a nearby star as a phase calibrator).

Correlated Flux: The correlated flux is integrated coherently in the form

$$|C_{ij}(\lambda)| = \left| \sum_{\vec{u}} \langle I(\vec{u}, \lambda, t) e^{-i\phi_{\text{atm}}(\lambda, t)} \rangle_t \right|, \quad (3.39)$$

where once again u is the spatial frequency integrated between $(\vec{B}_{ij} - D)/\lambda$ and $(\vec{B}_{ij} + D)/\lambda$ and D is the pupil diameter.

Visibility: Analogous to the incoherent case, the visibility is simply the correlated flux normalized by the photometric flux:

$$V_{ij}(\lambda) = \frac{|C_{ij}(\lambda)|}{\sqrt{\langle P_i(\lambda, t) P_j(\lambda, t) \rangle_t}}. \quad (3.40)$$

Typically $|C_{ij}(\lambda)|$ is used as the numerator, but to compute the *complex visibility* one can instead use the complex value of the right hand side of Eq. 3.39.

Differential Phase: As stated before, the absolute phase can not currently be calibrated due to the high variability of the sky in the infrared. However, one can consider the differential phase, as relative phase shifts with wavelength of the target object are independent of the current phase shift due to the sky or the instrument. Starting from some reference wavelength λ_{ref} (commonly 9.7 μm in the N band), the differential phase is computed as

$$\phi_{ij}(\lambda) = \text{Arg} \left[\sum_{\vec{u}} \langle I(\vec{u}, \lambda, t) e^{-i\phi_{\text{atm}}(\lambda, t)} \rangle_t - \sum_{\vec{u}} \langle I(\vec{u}, \lambda_{\text{ref}}, t) e^{-i\phi_{\text{atm}}(\lambda_{\text{ref}}, t)} \rangle_t \right]. \quad (3.41)$$

Though the differential phase was used extensively in the analysis of MIDI data, the closure phase is more widely used with MATISSE data because it gives some amount of absolute phase information necessary for narrow-band imaging. The differential phase is crucial for chromatic imaging (Sanchez-Bermudez et al., 2018, incl. J. Isbell), but this is yet to be widely employed.

2.4.3 Calibration of Interferometric Observables

Compared to the extraction of the interferometric observables, their calibration is relatively straightforward. In short, each observable is measured on both the science target

and the calibration star, and because we “know” what the observations from the calibrator should be, a correction factor can be calculated.

The (squared) visibilities and correlated fluxes are calibrated in a similar way, but the correlated fluxes require one additional step. First, consider the calibration star. Ideally, this star is completely all resolved at all MATISSE wavelengths. From Eq. 3.29 we can compute that this results in visibilities of exactly 1 on all baselines. Any deviations from 1 are then instrumental and should be divided out. In practice, however, the star may be partially resolved. Its true visibilities are then calculated using its known diameter (from e.g., [Cruzalèbes et al., 2019](#)) and the zeroth-order Bessel function of the first kind $J_0(x)$. This assumes that the star can be approximated as a uniform disk, and for the vast majority of cases this assumption is valid. The calibration of the visibility on each baseline is then

$$V_{\text{sci,calibrated}}^\alpha(\lambda) = \frac{V_{\text{sci,raw}}^\alpha(\lambda)}{V_{\text{cal,raw}}^\alpha(\lambda)} \times [2J_0(\pi dB/\lambda)]^\alpha, \quad (3.42)$$

with d the diameter of the star in radians, λ the observed wavelength(s), B the baseline length in meters, and $\alpha \in [1, 2]$ determines whether visibilities or squared visibilities are calibrated. The calibration of correlated fluxes the same (with $\alpha = 1$), with the added step of multiplying by the known spectrum of the calibrator:

$$F_{\text{sci,calibrated}}(\lambda) = \frac{V_{\text{sci,raw}}(\lambda)}{V_{\text{cal,raw}}(\lambda)} \times 2J_0(\pi dB/\lambda) \times F_{\text{cal,true}}(\lambda) \text{ [Jy]}. \quad (3.43)$$

The closure phase is somewhat simpler. For an unresolved source, the closure phase on any triangle should be exactly 0° . Moreover, even if the star is marginally resolved, it should be point-symmetric, and therefore still have closure phase of exactly 0° . Any deviations from zero in the closure phase of the calibrator are directly subtracted from the closure phase of the target. This is done individually for each BCD configuration because the BCDs introduce sign flips in the phase as the order of the fringes is commuted. The differential phase is calibrated identically.

2.4.4 Error Estimates

For both phases and visibilities (or correlated fluxes), the BCD configurations are used to compute error estimates on the observables. The observations are repeated at least four times throughout an OB, once for each BCD configuration. Following the calibration steps above, four independent measurements of the various observables are available. Taking phase sign flips from the BCD configurations into account, the mean value of each observable is then calculated to be the “final” observable. The 1σ error is estimated statistically using the standard deviation of those four observations. The accuracy of this estimation increases with the number of independent observations, and it is not uncommon for the interferometric part of an OB to be repeated a number N_{cycles} times. Particularly for faint sources, using the additional information from these repeated snapshots yields much more accurate fluxes and error estimates.

2.4.5 Using the DRS

The MATISSE DRS is build using the ESO Common Pipeline Library, and as such can be used from within each of the [esorex](#), [esoreflex](#), or [gasgano](#) environments. Throughout this work, I have used the esorex interface because it offers the greatest amount of control over the data reduction process, allowing one to set any of hundreds

Table 3.1: Suggested non-default LM -band reduction parameters for `mat_raw_estimates`.

Parameter	useOpdMod	coherentAlgo	corrFlux	compensate
Value	FALSE	2	FALSE	pb,rb,nl,if,bp,od
Parameter	cumulBlock	spectralBinning		
Value	FALSE	5		

Table 3.2: Suggested non-default N -band reduction parameters for `mat_raw_estimates`.

Parameter	useOpdMod	coherentAlgo	corrFlux	compensate
Value	TRUE	2	TRUE	pb,rb,nl,if,bp,od
Parameter	coherentIntegTime	cumulBlock	spectralBinning	
Value	< 0.2	FALSE	11	

of parameters. In practice, however, most of these parameters remain unchanged. Furthermore, the DRS was designed as a series of “atomic” recipes that can be wrapped and combined in larger “master” recipes.

The master recipe most often used is called `mat_raw_estimates`, and it performs the CCD calibration as well as the extraction of the interferometric observables. The default is to use the incoherent approach which returns raw squared visibilities, raw closure phases, and an estimate of the OPD in each frame (although it is not applied). Coherent integration can be changed by setting `corrFlux=TRUE` and `coherentAlgo=2`, which returns the correlated flux, the closure phase, the differential phase, and the estimated OPD. There are a number of other parameters which are changed as recommended by the MATISSE consortium, regardless of integration method. The recommended non-default reduction parameters for each band are listed in Tables 3.1 and 3.2.

After the CCD calibration and visibility, closure phase, etc. extraction by `mat_raw_estimates` for both the target and the calibrator, the observables must be calibrated. For the incoherently processed data, this is done using the master recipe `mat_cal_oifits`. The recipe takes as inputs the target raw observables, the calibrator raw observables, and the name of the calibrator (or its diameter). The calibrated visibility is obtained by dividing the target visibility by the calibrator visibility and then multiplying by the expected visibilities for a star with the given diameter. The final closure phase is obtained by subtracting the calibrator closure phase from the target closure phase. No recipe exists at the time of writing for the calibration of correlated fluxes, but a method to do this is given in Ch. 5.

The DRS has been optimized for sources with an L band flux $F_{3.6} > 0.5$ Jy or an N band flux $F_{12} \gtrsim 1$ Jy. Moreover, it uses $3.6 \mu\text{m}$ as the reference wavelength for the estimation of the OPD in the L band. This works well for bright sources or for blue sources (such as young stellar objects). For red objects, like Seyfert 2 galaxies, OPD corrections in the L band are severely impacted by low SNR and the M band should be used. At the time of writing, there is no option to change the OPD reference wavelength. For the data reduction of Circinus, it was necessary to manually extract the fringes using the M band as a reference.

Fortunately, the CCD calibration steps and the production of the clean interferogram could be used directly from the DRS. These steps resulted in a time series of 2D Fourier transforms of the interferogram (called `OBJ_CORR_FLUX_*.fits`). In each frame, the positions of the fringes could be measured and shifted (i.e., phase-corrected) before the correlated flux, differential phase, and closure phase could be extracted using the same methods as described above. I wrote the manual extraction procedure in Python for the express purpose of reducing Circinus L band data. This method is used extensively in

Ch. 5.

3 Image Reconstruction

Interferometric data are difficult to interpret directly, except in the simplest of cases (point source, Gaussian, or uniform disk), and representations of the data as images are necessary. With MIDI, these representations were made through forward modeling using simple components (e.g., several elliptical Gaussians and a point source), but they failed to capture the inherent details of the source. Image reconstructions, on the other hand, are model free representations of the source flux and can capture the details otherwise hidden by prior assumptions.

Image reconstructions were not possible using MIDI data because there were no absolute or closure phases available. Phases are a crucial part of image reconstructions because they carry information about the spatial distribution of the flux (see example figure). The current generation of infrared interferometers at the VLTI, GRAVITY and MATISSE, both provide closure phases and make image reconstructions possible. Image reconstructions are used extensively throughout this thesis, and so the reconstruction algorithms, typical usage, and resulting image fidelity are described in this section.

3.1 Minimization with Regularization

Direct Fourier inversion of infrared interferometric data is both not possible (due to the lack of absolute phases) and not recommended (due to sparse uv -coverage), and so image reconstructions instead rely on regularized minimization to iteratively find the image which best matches the data. This process is similar to forward model fitting, but with each pixel in a grid allowed to vary. One can immediately see that the number of parameters can greatly exceed the amount of data available to constrain them, and so this inversion problem is considered *ill-posed*. In other words, there are many images which provide a “good fit” to the data and the global minimum may either be unphysical or difficult to identify. The use of **regularization** ensures that the image does not become overly complex and remains physically plausible; this is especially important for sparse uv -coverage.

An algorithm for astronomical image reconstruction using the principle of maximum entropy was first presented by [Skilling & Bryan \(1984\)](#). Since then, numerous image reconstruction packages for infrared interferometric data have been developed based on the same principle (BSMEM, etc). More recently, however, algorithms have started using more generalized methods such as Markov chain Monte Carlo sampling of the parameter space (e.g., SQUEEZE) or gradient descent (e.g., IRBis) to find the best image rather than pure maximum entropy. Therefore, while the method of [Skilling & Bryan \(1984\)](#) did not include regularization, the modern approaches can and do include it as a key component.

We will now consider an image, I , sampled at a collection \vec{x} of (u, v) points. The basic principle of image reconstruction is to find the image with the minimum cost Q as defined

$$Q(y, \sigma_y, I(\vec{x})) = \frac{1}{2} \chi^2(y, \sigma_y, I(\vec{x})) + \sum_i^n \mu_i R(I(\vec{x})_i), \quad (3.44)$$

where y and σ_y are the observed quantities and their uncertainties, respectively; $R(I(\vec{x})_i)$ is the regularization function calculated at pixel i ; and μ is the weighting factor (also called the **hyperparameter**) used to scale the relative weighting of the regularization. The value of the hyperparameter plays a large role in the final images, and so care must be

taken to find the best μ for a given source. Different regularization functions emphasize different aspects of the image, and there are six commonly used regularization functions which are described in more detail in the next section along with a method to estimate the best hyperparameter value.

It is not straightforward to find the global minimum, and different image reconstruction software packages approach this in different ways. The most-used methods are Markov chain Monte Carlo sampling of the parameter space and conjugate gradient descent. Different packages also implement slightly modified versions of the χ^2 function (in particular, changing the relative importance of the closure phases and the visibilities), but the principle is the same. The specific implementation of the cost function in IRBis is described below.

3.2 Regularization Functions

Regularization is a crucial component of image reconstruction, as the *a priori* assumptions they provide limit the parameter space of this ill-posed problem. While any function could be used for regularization, there are several which are common to the three most-used infrared reconstruction packages: MiRA, SQUEEZE, and IRBis. The six regularization functions used in IRBis are described below, and a comparative example is shown in Fig. 3.8. While the implementations from IRBis are described, they are nearly identical in other software packages. In several cases, the regularization functions are only subtly different, but it is useful during image reconstruction to compare the results of them all to look for features which are constant and which are dependent on regularization.

The regularization functions described below include a prior term as is used in both IRBis and SQUEEZE. This prior image, P , has the same shape as the reconstructed image, I , and serves as an initial guess. The default value is a uniform image, but random noise or any user-generated image can be used. In general, the regularization functions penalize large deviations from this prior image.

3.2.1 ℓ_2 -Norm

The ℓ_2 -norm (or Euclidian norm) penalizes large pixel values anywhere in the image and results in the minimum amount of flux than can be used to reproduce the observables.

$$R(I, P) = \sum_i^n I_i^2 / P_i \quad (3.45)$$

3.2.2 Maximum Entropy

The maximum entropy regularizer behaves very similarly to the original maximum entropy image reconstruction methods of Skilling & Bryan (1984). The entropy of a set of positive numbers f_1, f_2, \dots, f_n was described by Shannon (1948); Skilling & Bryan (1984) as

$$S(f) = - \sum_i^n p_i \log p_i \quad (3.46)$$

$$p_i = f_i / \sum f.$$

Using this definition, the entropy measures “the number of bits of information needed to localize the position $[i]$ of a single radiated photon” in an image (Skilling & Bryan, 1984).




Fig. 3.8: Comparison of different regularization functions for poor (top) to good (bottom) uv -coverage. The regularization functions emphasize different aspects but converge with increasing uv -coverage. Only five of the six IRBis regularization functions are used here because Quadratic Tikhonov is nearly identical to the ℓ_2 -Norm. PDQ means Pixel-Difference Quadratic.

Because newer reconstruction packages rely on minimization rather than maximization, the negative sign is dropped, and the slightly modified (to include the prior) form

$$R(I, P) = \sum_i^n (I_i \log(I_i/P_i) - I_i + P_i) \quad (3.47)$$

is used.

3.2.3 Pixel-Difference Quadratic

The pixel-difference quadratic regularizer penalizes large jumps in flux between adjacent pixels. It takes the form

$$R(I, P) = \sum_{x,y}^n (|I(x, y) - I(x + dx, y)|^2 + |I(x, y) - I(x, y + dy)|^2) / P(x, y), \quad (3.48)$$

where $I(x, y)$ is the image evaluated at pixel coordinate (x, y) and $dx = dy = 1$.

3.2.4 Edge-Preserving

The edge-preserving regularizer is similar to the pixel-difference quadratic in that it measures the difference between adjacent pixels. It differs, however, in the inclusion of an edge term, ϵ , which is the allowed flux jump at the edge of a structure:

$$R(I, P) = \sum_{x,y}^n \left[(|I(x + dx, y) - I(x, y)|^2 + |I(x, y + dy) - I(x, y)|^2 + \epsilon^2) - \epsilon \right] / P(x, y). \quad (3.49)$$

3.2.5 Smoothness

The smoothness regularizer is similar to the pixel-difference quadratic and in that it penalizes large changes between adjacent pixels, but in this case diagonal pixels are considered. The smoothness regularizer takes the form

$$R(I, P) = \sum_{x,y}^n (|I(x, y) - I(x + dx, y + dy)|^2 / P(x, y)), \quad (3.50)$$

where $I(x, y)$ is the image evaluated at pixel coordinate (x, y) and $dx = dy = 1$.

3.2.6 Quadratic Tikhonov

The quadratic Tikhonov regularizer penalizes differences from the prior image. This regularizer should be used if the prior is well constrained. In the uniform prior case, this regularization function behaves similarly to the ℓ_2 -norm. The quadratic Tikhonov regularization function takes the form

$$R(I, P) = \sum_i^n (|I_i - P_i|^2). \quad (3.51)$$

3.3 IRBis

Image **R**ecreconstruction software using the **B**ispectrum (IRBis; Hofmann et al., 2014, 2016) is the reconstruction package designed specifically for MATISSE. It is included in the MATISSE DRS under the recipe name `mat_cal_imatec` so that it can be accessed and installed in the same framework as the rest of the data reduction steps. IRBis searches for the image, I , with the minimum cost

$$J(I) \equiv Q(I) + \mu R(I), \quad (3.52)$$

where $Q(I)$ is uv -weighted χ^2 , μ is the hyperparameter, and $R(I)$ is the chosen regularization function. IRBis does not use either the squared visibilities or the closure phases, but instead uses *both* of them in a combination called the bispectrum. The bispectrum is a complex-valued function containing the squared visibilities and the closure phases:

$$B(\vec{f}_1, \vec{f}_2) = \sqrt{V^2(\vec{f}_1)V^2(\vec{f}_2)V^2(\vec{f}_1 + \vec{f}_2)} e^{i\phi(\vec{f}_1, \vec{f}_2)}, \quad (3.53)$$

where $\phi(\vec{f}_1, \vec{f}_2)$ is the closure phase from the triangle defined by spatial frequencies \vec{f}_1 , \vec{f}_2 , and $\vec{f}_1 + \vec{f}_2$, with $\vec{f}_1 = \vec{u}_1/\lambda_{\text{obs}} = (u/\lambda_{\text{obs}}, v/\lambda_{\text{obs}})$. IRBis then defines the χ^2 of the bispectrum to be

$$Q(I) = \int_{\vec{f}_1, \vec{f}_2 \in M} \frac{w_d(\vec{f}_1, \vec{f}_2)}{\sigma^2(\vec{f}_1, \vec{f}_2)} |\gamma_0 B_k(\vec{f}_1, \vec{f}_2) - B_{\text{obs}}(\vec{f}_1, \vec{f}_2)|^2 d\vec{f}_1 d\vec{f}_2, \quad (3.54)$$

where M is the set of all observed spatial frequency vectors \vec{f} , $w_d(\vec{f}_1, \vec{f}_2)$ is the uv weight density at the selected spatial frequency, $\sigma(\vec{f}_1, \vec{f}_2)$ is the uncertainty of the bispectrum element at the selected spatial frequencies, and γ_0 is a normalization factor for the bispectrum.

The best image is searched for using the conjugate gradient descent algorithm ASA_CG (Hager & Zhang, 2005). ASA_CG is a bound-constrained, nonlinear, active set optimization algorithm. The variables to be optimized are the pixel intensities, and they are constrained to be larger than 0. The input to the minimization algorithm is the cost function $J(I_k)$ (Eq. 3.52) at iteration k and its first derivative with respect to the pixel intensities

$$\frac{\partial J(I_k)}{\partial I_k} = \frac{\partial Q(I_k)}{\partial I_k} + \mu \frac{\partial R(I_k)}{\partial I_k}. \quad (3.55)$$

Gradient descent algorithms take steps in the direction of the maximum gradient, so that they reach minima efficiently. It is easy to fall into local minima with this procedure, but varying the step size and introducing perturbations (called “kicks”) can help the algorithm to escape local minima and converge on the global solution.

Finally, a reconstruction quality estimator q_{rec} is defined to select the best image reconstructed with the selected parameters (described below):

$$q_{\text{rec}} = \frac{1}{4} (|\chi_{V^2}^2 - 1| + |\rho_{\rho_{V^2}} - 1| + |\chi_{\text{CP}}^2 - 1| + |\rho_{\rho_{\text{CP}}} - 1|). \quad (3.56)$$

The χ^2 terms are calculated in the normal way, but the “residual ratios”, $\rho_{\rho_{V^2}}$ and $\rho_{\rho_{\text{CP}}}$, are defined as the ratio of the sum of all positive residuals to the sum of all negative residuals for the squared visibilities and for the closure phases, respectively. For example, the residual ratio of the squared visibilities is given by

$$\rho_{\rho_{V^2}} = \frac{\int_{\vec{f} \in M_+} [V_k^2(\vec{f}) - V_{\text{obs}}^2(\vec{f})]/\sigma_{\text{obs}}(\vec{f}) d\vec{f}}{\int_{\vec{f} \in M_-} [V_k^2(\vec{f}) - V_{\text{obs}}^2(\vec{f})]/\sigma_{\text{obs}}(\vec{f}) d\vec{f}}, \quad (3.57)$$

where M_+ is the set of spatial frequencies with positive residuals and M_- is the set of spatial frequencies with negative residuals. Each term is inverted if its value is between 0 and 1. The quality estimator q_{rec} defined in this way converges toward 0 for a perfect reconstruction.

During the work carried out in the following chapters, I have developed a similar reconstruction quality parameter which allows one to set the relative importance of the squared visibilities and the closure phases,

$$q = \frac{\alpha}{N_{V^2}} \sum_{i=1}^{N_V} \frac{(V_{\text{obs}}^2 - V_{\text{model},i}^2)^2}{\sigma_{V^2,\text{obs}}^2} + \frac{\beta}{N_\phi} \sum_{j=1}^{N_\phi} \frac{(\phi_{\text{obs}} - \phi_{\text{model},i})^2}{\sigma_{\phi,\text{obs}}^2}, \quad (3.58)$$

where α and β serve as weights for the squared visibilities and the closure phases, respectively. This function is used exactly as the default one, but it allows the user to further suppress noisy data (especially in cases where the uncertainties are underestimated). Moreover, IRBis allows the user to select different cost functions (the $Q(I)$ term in Eq. 3.52): cost function 1 is equivalent to setting $\alpha = \beta = 1$; cost function 2 is equivalent to $\alpha = 0$ and $\beta = 1$; and cost function 3 uses the bispectrum as defined in Eq. 3.54 (the third cost function was introduced in Hofmann et al. (2022)). Using the above-defined Eq. 3.58 one can select the best images using the same criterion as used in the optimization run.

3.3.1 IRBis Parameter Search

IRBis (accessed through `esorex mat_cal_imarec`) has 53 individual parameters the user can change in order to find the optimal image as of DRS version 1.6.0. Fortunately, only a handful of these are changed often and the rest work (for the most part) with the default value.

The following parameters are most often changed by the user:

1. The reconstruction field of view (FOV) in milliarcseconds.
2. The number of pixels (e.g., 128×128 or 256×256 ; optimally a power of 2 for the Fast Fourier Transform). This in combination with the FOV sets the pixel scale of the image.
3. The starting and/or prior image (e.g., a Gaussian, a uniform disk, etc.)
4. The regularization function to be used (1 - 6), with the prior image set by a model image or simple shapes such as Gaussians. Optionally a negative value (-1 to -6) can be used to select the regularization function, but in this case a uniform prior is used.
5. *Either* a starting value for the hyperparameter μ , a step size to change it by, and a number of steps n *or* a length n array of μ values to use.
6. A starting value for the object mask, a step size to change it by, and a number of steps m . The object mask is a binary mask; pixels inside the mask are allowed to vary during optimization and pixels outside are set to a constant.
7. The cost function (1, 2, or 3 as defined in the previous section).

IRBis does a total of $m \times n$ reconstructions for the given settings, with m the number of mask radii and n the number of μ steps. The best image for the $m \times n$ runs is selected by default with Eq. 3.56. To use Eq. 3.58, one must set $m = n = 1$ and calculate the quality of the reconstruction manually.

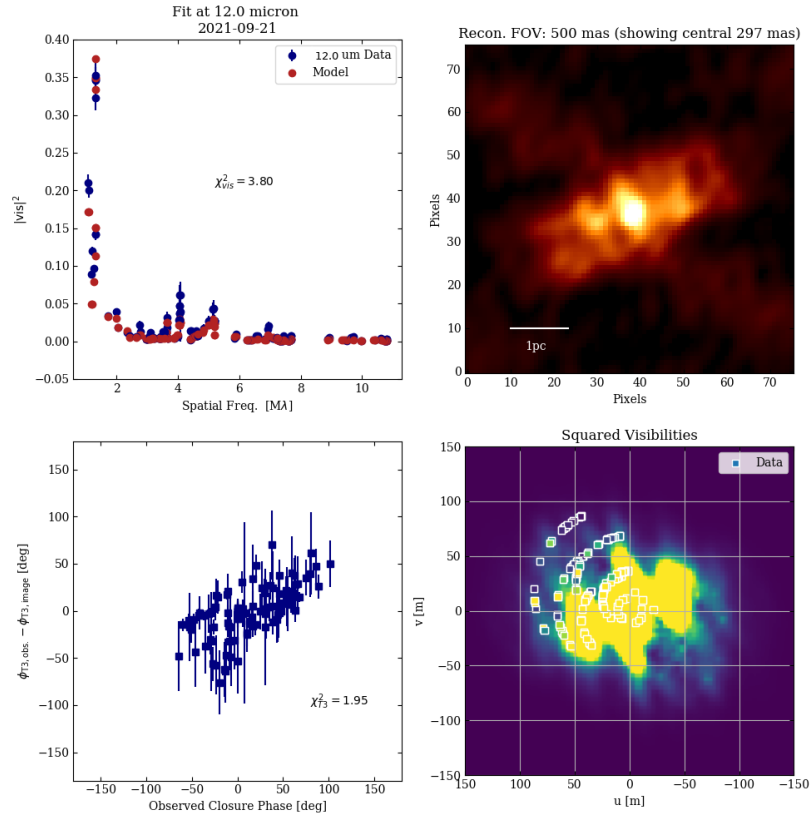


Fig. 3.9: Example of an image reconstruction quality assessment plot. In the (*top left*) are the squared visibilities as a function of baseline length for the model and for the observed data. In the (*top right*) is a cutout of the central 300 mas of the reconstructed image. In the (*bottom left*) are the residuals of the closure phases for the reconstructed image. In the (*bottom right*) are the observed squared visibilities plotted over the model squared visibilities on the filled uv -plane.

Each of these parameters can strongly affect the resulting image reconstruction, particularly in cases of sparse uv -coverage. To find the best image reconstruction given the large parameter space, I have developed a parallelized grid search algorithm which creates a large number of “optimal” images with IRBis and then uses Eq. 3.58 to select the best image. The algorithm consists of three steps:

1. A large number $N = f \times r \times m \times n \times c$ of parameter files are created (here f is the number of fields-of-view to use, r is the number of regularization functions, m is the number of object mask radii, n is the number of hyperparameter values, and c is the number of cost functions).
2. Once these numerous parameter files are made, parallel processes are spawned in batches of up to 32, each calling `esorex mat_cal_imatec` with settings defined by the parameter file.
3. The reconstruction quality (Eq. 3.58) is calculated and a plot comparing the model and data squared visibilities and closure phases is made for each image.

An example of this quality assessment image is shown in Fig. 3.9. The pixel number is not varied in this setup because for parameter exploration 128×128 pixels provides a good balance between details and computation speed (Fast Fourier Transforms scale as $O(N \log N)$ where N is the number of elements). Once the ideal parameters have been identified a 256×256 pixel image can be made.

3.4 Quantifying Image Fidelity

The following two approaches were used throughout the development of the work in this thesis to estimate the errors on the flux in a given pixel. These error estimates are crucial for *a)* identifying real structures in the image and *b)* estimating uncertainties on aperture fluxes and fitted parameters such as temperature. The first method is similar to bootstrapping in that images are reconstructed using a random subset of the uv -coverage. The second method compares the reconstruction of simulated observation of a model to the model itself. These methods are described below.

3.4.1 Delete- d Jackknife

While we can estimate the uncertainty of the correlated flux or closure phase measurements for a single uv -point using repeated snapshots, sub-exposures of an integration, and even adjacent uv -points, it is important to know quantify the impact that individual (possibly erroneous) measurements can have on the reconstructed image. We do this through delete- d jackknife resampling of our uv -coverage (the method is developed in [Shao & Wu, 1989](#)). In each Monte Carlo realization we randomly discard 10% of each the squared visibilities and the closure phases. For Circinus with 150 uv -points, this choice satisfies the criterion for being asymptotically unbiased: $\sqrt{n} < d < n$, where n is the sample size and d is the number of deleted elements. We then perform the image reconstruction at each wavelength using the best parameters from the grid search and save the results. After 100 realizations, we calculate the median and standard deviation in each pixel of each image. The median image at each wavelength is used as the final image. The standard deviation image serves as an error map with which we calculate the SNR at each pixel.

3.4.2 Reconstructing a Model

The previous method reduces the impact of individual uv -points, but it does not fundamentally change the uv -coverage. The specific sampling of the uv -plane, particularly in the case of sparse sampling, can impact the apparent elongation and size of reconstructed structures. To test the impact of this sampling (in other words, to disentangle real features from the so-called “dirty beam” which is the Fourier transform of the uv -sampling), one can do a model reconstruction. Starting from a simple model (e.g., a Gaussian), we take the 2-D Fourier transform of the model image and sample it at the given uv -points and closure triangles. Using these model observables (with or without added noise) we reconstruct an image using the same methods as for the real data. The resulting image is compared with the initial model image to calculate a residual map. From this map, one can either use the mean residual as a pixel error estimate, or use the residual map directly as an estimate of the 1σ errors. This method was applied to the sparsely uv -sampled NGC 1068 observations in Ch. 4.

Chapter 4

Thermal imaging of dust hiding the black hole in NGC 1068

The principles of interferometry and the new instrument, MATISSE – as described in the previous chapter – were used to study the circumnuclear dust in the prototypical Seyfert 2, NGC 1068. It was this AGN which inspired the unified model, and it is a key target for helping astronomers to understand the circumnuclear dust.

This chapter is based on [Gómez Rosas et al. \(2022, incl. J. W. Isbell\)](#). It has been modified to match both the style and language used throughout this thesis. As it was published in *Nature*, the formatting is different than the other papers: key results are presented in the first pages, and detailed methodology follows in a separate section. This format has been maintained, though sections which contain repeated information have been omitted for brevity and clarity. I was second author of this work, and my specific contributions were:

1. Assistance with planning and execution of the observations.
2. Reduction of the *LM*-band data, and the crucial tests to determine which spectral binning to use for the closure phases (§2.2.3).
3. Determination of image orientation through reconstruction of del Sco (§2.2.7).
4. Estimation of the image errors (§2.3.3).
5. Development of the Gaussian model fitting code used in §2.3.5.
6. Scientific interpretation and discussion of the morphology and temperature of the circumnuclear dust.
7. Data presentation, specifically Figs. 4.1, 4.9, and 4.10a.

This work would not have been possible without the contributions of the entire MATISSE collaboration, who tirelessly worked to bring us from first-light to science products. Violeta Gómez Rosas, Walter Jaffe, and I worked very closely on the main text. The full author list and their specific contributions are given on the *Nature* website with the online version of the article¹.

¹<https://www.nature.com/articles/s41586-021-04311-7>




Fig. 4.1: (*a-d*) Images are constructed from the *L*, *M* and *N* bands at median wavelengths of 3.7 μm (*a*), 4.6 μm (*b*), 8.5 μm (*c*), and 12 μm (*d*). The dashed white contours indicate the 3σ noise level and the white ellipses indicate the resolution. (*e*), Red–green–blue composite color image of the 3.7 μm (blue), 4.6 μm (green), and 12 μm (red) images. The ellipses in the composite image illustrate the possible midplane of the inhomogeneous obscuring layer discussed in §1.5. The innermost small ellipses indicate two possible positions of the masked sublimation limit around the position of the supermassive black hole, and the outermost dot-dashed ellipse shows the maximum measured disk extent (10 pc). Scale bar, 1 pc. $\Delta\alpha$ and $\Delta\delta$ are the position offsets from the brightest feature, measured in true mas, in the right ascension and declination directions, respectively.

1 Main Article

In the widely accepted ‘unified model’ solution of the classification puzzle of AGN (Antonucci, 1993), the orientation of a dusty accretion torus around the central black hole determines the appearance of the AGN. In ‘type-1’ systems, the bright nucleus is visible at the center of a face-on torus. In ‘type-2’ systems the thick, nearly edge-on torus hides the central engine. Later studies suggested evolutionary effects (López-Gonzaga et al., 2016) and added dusty clumps and polar winds (Asmus et al., 2016) but left the basic picture intact. However, recent high-resolution images (Gravity Collaboration et al., 2020) of the archetypal type-2 galaxy NGC 1068 (Seyfert, 1943; Antonucci & Miller, 1985), suggested a more radical revision. The images displayed a ring-like emission feature that was proposed to be hot dust surrounding the black hole at the radius where the radiation from the central engine evaporates the dust. That ring is too thin and too far tilted from edge-on to hide the central engine, and *ad hoc* foreground extinction is needed to explain the type-2 classification. These images quickly generated reinterpretations of the dichotomy between types 1 and 2 (Vermot et al., 2021; Prieto et al., 2021). Here we present new multi-band mid-infrared images of NGC 1068 that detail the dust temperature distribution and reaffirm the original model. Combined with radio data (J. Gallimore, priv. comm.), our maps locate the central engine that is below the previously reported ring and obscured by a thick, nearly edge-on disk, as predicted by the unified model. We also identify emission from polar flows and absorbing dust that is mineralogically distinct from that towards the Milky Way center.

From the polarization properties of the Seyfert type-2 galaxy NGC 1068 (Antonucci & Miller, 1985), earlier work (Antonucci, 1993) has established the unified model of

AGN, which explained the type 1-2 dichotomy with a physically and optically thick dust torus that hid the nucleus when viewed nearly edge-on. More recent observations and theory have refined the picture: the dust must be clumpy, its inner rim should be near the sublimation radius (r_{sub}), and radiation pressure can push the dust into an outflowing wind along the torus axis (Hönig, 2019).

1.1 NGC 1068, the archetype Seyfert 2 galaxy

For NGC 1068, r_{sub} is ~ 0.4 and ~ 0.2 parsecs (pc) for silicate grains with radii between 0.005 and 0.25 μm and for large graphite dust particles with radii $> 0.1 \mu\text{m}$, which have sublimation temperature (T_{sub}) of ~ 1200 and ~ 2000 K, respectively, assuming a bolometric luminosity $L_{\text{bol}} = (0.4 - 4.7) \times 10^{45} \text{ erg s}^{-1}$ (see Gravity Collaboration et al., 2020; Barvainis, 1987; Baskin & Laor, 2018, for details). Directly verifying the proposed torus geometry requires large interferometers to image these very small structures at infrared wavelengths (2-20 μm). Early MIDI/VLTI 8-12 μm interferometer data was too limited to permit direct imaging (Jaffe et al., 2004; Raban et al., 2009; López-Gonzaga et al., 2014), but model-fitting showed that complex, asymmetrically structured polar emission actually dominates the infrared luminosity. Recent extensive GRAVITY/VLTI data at 2 μm wavelength allowed a higher resolution image that showed an incomplete ring and other, fainter, structures (Gravity Collaboration et al., 2020). The authors interpreted the ring as a planar-circular feature of 0.24 pc radius seen 20° from edge-on and, after considering other possibilities, advocated identifying this feature as the hot sublimation ring with the central engine (CE) at its center, primarily because its radius agreed with the expected value of r_{sub} for large graphite grains. This geometry breaks with the Unified Model. The ring is too thin and too inclined to cover the CE and the authors invoke an additional component, a circumstantial foreground dust cloud to block the emission from the Broad Line Region (BLR). This break with the established model is based on the GRAVITY temperature estimate of the ring, and its consequent identification with r_{sub} , but in fact the wavelength range of GRAVITY is too restricted to establish accurate temperature values in the range of T_{sub} . The newly commissioned Multi-aperture Mid-Infrared Spectroscopic Experiment at the Very Large Telescope Interferometer (MATISSE/VLTI) interferometer observes in the wavelength range from 3 to 13 μm , which is ideal for mapping the dust temperature distribution. NGC 1068 was therefore a major target for MATISSE observations.

1.2 New infrared observations and images

In Sept. 2018 and Nov. 2019, we observed NGC 1068 with the MATISSE instrument (§2.1; Lopez et al., 2022), which combines very high angular resolution with a broad spectral range. The observations covered three infrared bands: *L* (3–4 μm), *M* (4.5–5 μm) and *N* (8–13 μm). The MATISSE interferometric data is densely enough sampled (Fig. 4.5) for image reconstruction. We utilized three independent image reconstruction systems: IRBis (Hofmann et al., 2014), MIRA (Thiébaud, 2008), and a system based on the CLEAN algorithm (Högbom, 1974). They all converged to similar images. The IRBis images are shown in Fig. 4.1.

In the *L*- and *M*- bands we see two similar structures: (1) a Northern asymmetric incomplete elliptical ring, $25 \times 12 \text{ mas}$ ($1.7 \times 0.9 \text{ pc}$) with Position Angle (PA) $\sim -45^\circ$: the Northern Complex (NC), and (2) to the South, a fainter, less well-defined source: the Southern Extension (SE). The SE is not an interferometric artifact; it persists over wavelength and is independent of reconstruction method. The *N*-band images show a bright central source similar in extent and PA to the NC, and additional diffuse structures at larger scales.

MATISSE data do not provide absolute astrometric positions that would allow direct registration of the different bands, so we have registered the images by two criteria: (1) morphological similarity; (2) consistency of the resulting spectral energy distributions (SEDs). This is equivalent to superimposing the brightest features in each band. This registration is shown in Fig. 4.1. The SEDs extracted at various points after these registrations were smooth and physically plausible (compare Figs. 4.2 and 4.3). This was not the case with significant (>3 mas) relative shifts of the images. Furthermore, we find a good match between NC and the K -band ($2 \mu\text{m}$) GRAVITY image4 when we align the GRAVITY inner ring with the brightest component of NC (Fig. 4.4c). The GRAVITY K -band flux agrees with our SED fit (see Fig. 4.3) for the brightest feature in our images, indicated as ellipse E1 in Fig. 4.2.

1.3 Thermal modelling and dust mineralogy

To understand the thermal state of the dust in NGC 1068, we extracted spectrophotometry from nine areas or apertures based on the recovered geometry (Fig. 4.2). We applied the apertures to the three reconstructed images and to a multi-Gaussian model fit of the data (explained in §2.3.5). We tested multiple image reconstructions and modelling methods and found that the fluxes extracted from each yielded equivalent results within the error bars (§2.3.7 and Fig. 4.12). The Gaussian fits proved more convenient for computing SEDs at many wavelengths, as they provide uncertainty estimates via probability distributions, and are used below. The SEDs for the brightest central aperture E1 and for the high-temperature SE are plotted in Fig. 4.3, the rest in Fig. 4.13.

To add a physical interpretation to these SEDs, we modelled each with emission curves from a blackbody with a cold foreground absorbing screen. Most components required a second blackbody at a distinct temperature, a situation supported by radiation transfer modelling of dusty clumps (Nenkova et al., 2008b). We model each blackbody as:

$$S_\nu(\lambda) = A\eta B_\nu(T)e^{-\kappa(\lambda)N_{\text{ext}}} \quad (4.1)$$

A is the area of each ellipse; B_ν is the Planck function. The fitted parameters are the dust temperature T , the area filling factor η for each blackbody, and the cold foreground dust surface density: N_{ext} (in $\mu\text{g cm}^{-2}$). $\kappa(\lambda)$ is the wavelength-dependent dust opacity per unit surface density. Contributions from silicates, iron, and carbon, as well as the dust grain size distribution determine $\kappa(\lambda)$ in the mid-infrared. The most prominent spectral feature is the strong silicate absorption near $10 \mu\text{m}$. We first attempted fitting the SED of areas E1 and SE with a “standard interstellar” curve for $\kappa(\lambda)$ which closely resembles that measured toward the Galactic center (Barvainis, 1987; Fritz et al., 2011). This was unsuccessful; at $8 \mu\text{m}$ these curves contain a large opacity which is incompatible with the high measured fluxes, see Fig. 4.3. To clarify the dust composition, we tested theoretical opacity curves of common interstellar substances; we could only fit the measurements with dust particles comprised primarily of amorphous olivine ($[\text{Mg}_2, \text{Fe}_2]\text{SiO}_4$ Min et al., 2005) with an admixture of up to 20% by weight of carbonaceous dust and with relatively small dust particles ($0.1\text{--}1.5 \mu\text{m}$). These specifications apply to the foreground dust only; we have assumed the warm dust to emit as an optically thick blackbody. For the absorbing dust, $N_{\text{ext}} = 1000 \mu\text{g cm}^{-2}$ corresponds to optical depths $\tau_{3.4 \mu\text{m}} = 2.0$, $\tau_{9.7 \mu\text{m}} = 2.4$, $\tau_{10.6 \mu\text{m}} = 2.5$, and $\tau_{12.0 \mu\text{m}} = 1.2$

The contribution of carbon is confirmed by the presence of a visibility feature close to $3.4 \mu\text{m}$ on several baselines (Fig. 4.6), but no evidence of aromatic C-H features near $3.3 \mu\text{m}$ are seen. In preliminary high resolution N -band spectra (Fig. 4.8) the shape of the silicate feature at the longest baselines suggests the presence of crystalline grains

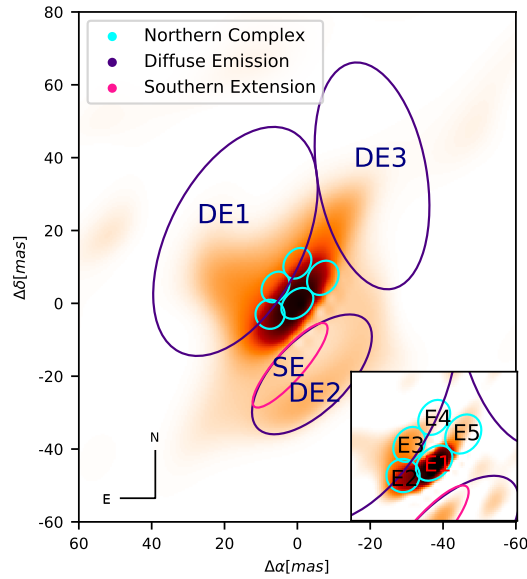


Fig. 4.2: The labelled apertures are superimposed on MATISSE N -band image. Inset, 30×30 mas enlargement of the central region (the northern complex) with labelled small apertures, superimposed on MATISSE L -band image.

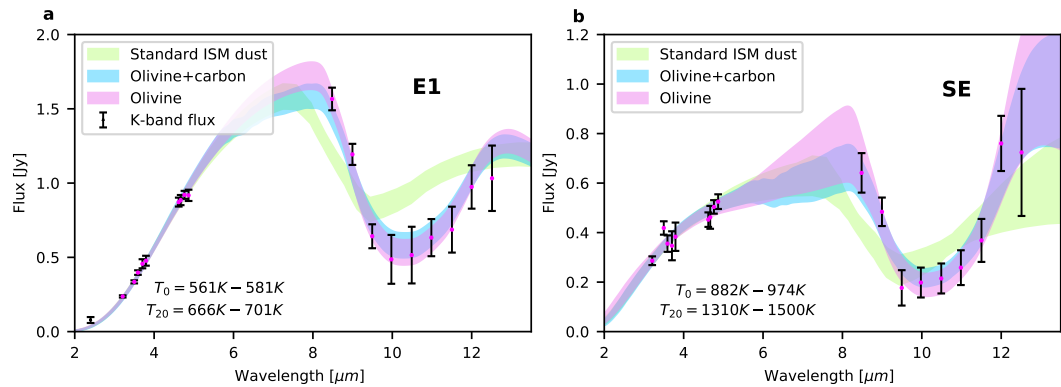


Fig. 4.3: (a) One-temperature fit to the bright E1 component. The black point at $2.2 \mu\text{m}$ gives the GRAVITY K -band flux (GRAVITY Collaboration, private communication) inside E1. The dots in magenta define the median values of the integrated fluxes at the different wavelengths. These and the error estimates are defined in §2.5. (b) Two-temperature blackbody fit to MATISSE SED of the southern extension. The shaded regions confine models inside 1σ , color-coded by different kinds of dust: green, cyan and magenta correspond to standard ISM dust, olivine with 20% carbon and olivine with no carbon, respectively. The colors are semi-transparent to show the regions where the models overlap. T_0 and T_{20} refer to the hot component temperatures of the fits using 0% and 20% of carbon, respectively.

Table 4.1: Fitted blackbody parameters for NGC 1068 in nine apertures. Values for each dust composition (0 or 20% carbon) are given as ranges $\pm 1\sigma$ from the minimum χ^2 solution.

Aperture	Carbon Content [%]	T_{cold} [K]	T_{hot} [K]	N_{ext} [$\mu\text{g cm}^{-2}$]
E1	0	561-581	-	396-547
	20	666-701	-	448-610
E2	0	188-469	703-1179	281-530
	20	206-506	829-1461	386-728
E3	0	217-376	860-1095	67-386
	20	225-371	917-1412	92-489
E4	0	232-389	555-818	281-530
	20	250-412	633-1200	386-621
E5	0	223-285	717-1098	1-100
	20	220-287	712-1275	1-189
DE1	0	229-238	771-865	1-34
	20	228-242	820-870	1-43
DE2	0	139-228	417-475	825-1000
	20	153-292	543-703	1000-1468
DE3	0	148-162	-	1-197
	20	148-162	-	1-241
SE	0	203-354	846-1031	530-853
	20	173-300	1333-1500	464-619

that would imply high temperature grain reprocessing. Further interpretation of these features will be presented in a future paper.

Fig. 4.3 shows the best fits using pure olivine and a mix of olivine with carbon (cyan, magenta). Fits using standard ISM dust opacities (green) do not match the silicate absorption profile. The fraction by weight of carbon is limited to $\leq 20\%$; with the best fit near no carbon. Increased carbon fraction correlates with higher temperatures. We obtain a good fit with a single-black body for E1 and DE3, the other ellipses require two blackbodies. The temperature range for E1 is 560–700 K, while for SE the hot component requires temperatures from 880–1500 K. We present the map and table of temperatures for all the ellipses in Table 4.1.

For all apertures, T_{hot} does not decrease systematically with distance from the inferred black hole position. For dust grains of fixed properties in radiation equilibrium from a central source we would expect $T \propto r^{-\alpha}$, with $0.3 < \alpha < 0.5$ depending on grain size. Although we only measure radii projected on the sky, the lack of a systematic temperature decrease with distance suggests heating by a mechanism other than direct radiation from the central engine.

The extinction varies greatly on small scales, implying it arises close to the sources. It increases rapidly and systematically to the South-West of E3, extending at least 30 mas in projection from E1 to cover SE and DE2 with the highest opacities. The Northern diffuse apertures show almost no extinction.

1.4 Radio images and the hidden black hole

In Fig. 4.4 we compare MATISSE images to an ALMA (Atacama Large Millimeter/-submillimeter Array) continuum image at 1.2 mm wavelength (Impellizzeri et al., 2019) and to a VLBA image at 1.4 cm (§2.4, J. Gallimore, priv. comm.) with the overlaid positions of a series of water masers (Greenhill et al., 1996) which mark warm, dense




Fig. 4.4: (a) MATISSE *N*-band image, and contours of ALMA 256-GHz radio image. The contours give [0.1, 0.2, 0.4, 0.8, 0.95] of the peak (0.85 mJy per beam). (b) MATISSE *L*-band image and contours of VLBA 22-GHz image. The white asterisk marks the probable black hole position. The small dots mark the maser positions and their colors indicate their redshifts relative to the galaxy systemic velocity. Note the smaller field of view. The contour levels are [0.1, 0.2, 0.4, 0.8, 0.95] of the peak (3.0 mJy per beam). (c) MATISSE *L*-band image and contours of GRAVITY *K*-band image. Contours are [0.1, 0.2, 0.5] of the peak.

molecular gas. The morphological similarities between infrared and radio images are striking. The smoothed MATISSE 12 μm image and the ALMA image show X-shaped emission of similar dimensions and orientation which can be registered to < 3 mas accuracy (§2.4). With this cross-identification, the 3.7 μm MATISSE image and the VLBA image can be superimposed (Fig. 4.4b) to the same accuracy. The strong resemblance in shape and orientation of the Northern arc of the VLBA and the *LM*-band NC supports this alignment. This fixes the position of the 22 GHz peak, which we assume marks the position of the black hole, at slightly below the lower edge of component E2, and sets the positions of the masers along the lower edge of NC.

The centimeter and millimeter radio emission arise as bremsstrahlung from very hot ionized gas (greater than approximately 106 K; discussed in §2.4). The association between this gas and the warm dust structures on 1–6-pc scales is puzzling because, if co-spatial, the gas would destroy the dust on short time scales. The association may only be a projection on the plane of the sky, but the morphological similarity strongly suggests connected physical processes. Both sources resolve into complex morphologies that cannot be characterized by a single position angle. The dominant axis of the millimeter source seems to be near position angle -30° , and that of the centimeter source is approximately -60° . The latter is comparable to the that of the nuclear radio structure S1 (Gallimore et al., 2004).

1.5 Spatial distribution of the dust

We estimate the temperature of the brightest emission feature, E1, to be 684 ± 17 K, well below T_{sub} . Together with the position of the VLBA peak and the gradients in temperature and extinction toward the South, this strongly suggests that the black hole and sublimation rim lie below E1 and E2 and are hidden in *LM*-bands by a band of highly optically thick dust, which is responsible for the observed type 2 classification. The opacity weakens to the South to reveal the SE. The proposed position of the black hole and temperatures we recover agree better with the interpretation of the brightest feature of the GRAVITY image (Fig. 4.4c) as cool (720 K) dust (their model 2), rather than a hot disk (Gravity Collaboration et al., 2020). A layer of cold dust with ~ 0.8 covers the whole SW quadrant, including DE2, to a projected distance of ~ 2.5 pc, but does not cover DE1 or DE3 in the North. The overall pattern is consistent with a large

sheet or disk of dust, inclined away from us, with major axis at PA $\sim -45^\circ$, that covers the Southern structures, but exposes DE1 and DE3. The disk opacity is highest in a 0.3 pc region between E1/E2 and SE, covering the central engine, but it flares up to partly cover the NC. On scales of 50-100 pc, the geometry of the Narrow Line Region and other structures implies that its pole is about 10° from the plane of the sky (Das et al., 2006; Poncelet et al., 2008), a similar inclination to the line of sight would imply a radial extent of our obscuring sheet in excess of 10 pc. The extent and orientation of this obscuring sheet at this inclination are indicated in Fig. 4.1. The components DE1 and DE2 probably represent dust blown out of the center; their orientation matches wind structures at larger scales (García-Burillo et al., 2019). The winds may originate within the hot irregular NC and SE structures.

1.6 Conclusions

In conclusion, we find the dusty structure that plays the role of the “historical torus” with an optically thick ring obscuring the central engine at parsec scales and a less optically thick disk extending to at least 10 pc. The distribution of the warm and hot dust reveals two main axes, one along the densest obscuring structure, at position angle -45° , and a second perpendicular axis parallel to the molecular outflows. However, the detailed structures are asymmetric in all bands, suggesting that they are not in full thermal and dynamical equilibrium. This conclusion is supported by the difference in PA between the minor axis of NC ($\sim 45^\circ$) which represents the angular momentum of the dust system at pc scales, and the PA of the radio jets ($\sim 10^\circ$; Evans et al., 1991; Gallimore et al., 1996), which indicates the black hole spin axis, and the PA perpendicular to the ultraviolet polarization directions ($\sim 7^\circ$; Kishimoto, 1999; Antonucci et al., 1994), which may indicate the spin axis of the sublimation ring. The innermost ellipses in Fig. 4.1 show two possibilities for the orientation of this ring: aligned with the ultraviolet polarization (white) and with the NC structures (yellow). The cold dust grain properties appear quite different from the standard Milky Way ISM dust. The grains are sub-micron Mg-rich olivines with some carbon. The corresponding properties of the hot dust seen in emission remain open issues. Future MATISSE observations of additional AGN will show whether NGC 1068 is atypical or prototypical. These high resolution, broadband images from MATISSE open new avenues for investigations of the dust-radiation interactions in AGN, such as detailed radiation transfer modelling of the clouds, studies of the evolution of carbon and silicate grain mineralogy in AGN environments, and more globally, a synthesis of observations with radio, millimeter, and infrared instruments to connect the dust cloud physics on scales of tenths to tens of parsecs.

2 Methods

Here we detail the MATISSE (Multi AperTure mid-Infrared Spectro-Scopic Experiment Lopez et al., 2022) the observations and the data reduction. We compare the results of several different image reconstruction techniques for our data. We describe the methods for extracting spectrophotometric data from several spatial areas of interest, and for fitting multi-component black-body emission models to these spectra, as presented in §1. Finally, we present details of the radio observations and the results of the black-body fits for those apertures not given in §1.

2.1 Description of the observations

NGC 1068 was observed during MATISSE commissioning in September 2018 (under program 0103.C-0143) and again with Guaranteed Time Observation (GTO) in 2019




Fig. 4.5: uv -coverage of MATISSE observations of NGC 1068.

(under 0104.B-0322). We use both datasets in our work, both are available at the ESO archive. The detailed logs of both the science and calibrator observations are publicly available², and the uv -coverage is shown in Fig. 4.5.

The instrument was configured in spectrally resolved Low Resolution “Hybrid” mode. In this mode the LM -band optical system uses a prism dispersing element ($R \simeq 30$) in the Simultaneous Photometry mode (SIPHOT), where a 1-to-3 beam-splitter separates the photometric and interferometric signals which are detected simultaneously. This allows visibilities, the ratio of these signals, to be measured in a way that is largely insensitive to the variations in the atmospheric Strehl ratio, which are important at these wavelengths. The N -band observations use an $R \simeq 30$ prism in High Sensitivity (HISENS) mode, because the high sky backgrounds require separate photometric observations, and the Strehl ratio variations are less important. A detailed description of the available MATISSE observing modes can be found in the MATISSE reference article and at the ESO website³.

To calibrate the instrument and atmosphere, we chose interferometric calibrator stars from a catalogue of mid-IR interferometric standard stars (Cruzalèbes et al., 2019), that were at similar airmass as the target ($< 10^\circ$ from the target), within a factor 2 in either L - or N - flux to the target, and whose angular diameters are small and well known.

The calibrators for each observation are given in the observing log along with the quantity τ_0 which describes the atmospheric coherence time in the V -band in milliseconds; it is a measure of the reliability of the visibilities. It has proved to be a more consistent estimator of data quality than measurements of the atmospheric seeing. Typically, $\tau_0 > 5$ ms indicates high reliability, while $\tau_0 < 3$ ms indicates low reliability (Lopez et al., 2022). Most of the commissioning data included here were taken under intermediate ($\tau_0 = 3 - 5$ msec) to good ($\tau_0 > 5$ msec) conditions; the GTO data were mostly taken under good conditions.

A standard observation contains six visibilities, from the distinct pairings of the four telescopes. The projections of the baseline vectors on the plane of the sky at the source are the uv -coordinates of the observations, and these coordinates specify which components of the Fourier transform of the sky brightness are represented in the visibilities. ED Fig. 4.5 shows the combined uv -coordinates of all our measurements. The maximum extent of the uv -coverage sets the effective resolution of the interferometric synthetic images; incompleteness in the filling of the uv -plane tends to create “aliased” artifacts in these images.

²The observing logs and reduced data can be found at https://github.com/VioletaGamez/NGC1068_MATISSE.

³<https://www.eso.org/sci/facilities/paranal/instruments/matisse.html>

2.2 Data flagging, reduction and calibration

2.2.1 Introduction

In this section we describe data measured in the L -, M -, and N -bands during the September 2018 commissioning runs and the GTO observation of November 2019. Typical of commissioning, not all data were of useful quality, either because of instrumental problems or because of bad weather, and we have excluded data deemed to be less reliable. In some cases, we applied special reduction steps that are not necessary for routine operations. In the following sections we describe:

- Pre-reduction data selection and editing
- Pipeline reduction of raw data
- Post-reduction selection and editing

2.2.2 Data selection and flagging

LM -band chopping status: When estimating non-interferometric photometry, we only include exposures during which the telescope secondary mirrors were chopped at ~ 1 Hz between target and sky to allow accurate removal of sky backgrounds.

Revision of chopping status labels: The MATISSE operating software contains problems that sometimes resulted in mislabelled chopping status of individual frames as T/U/S=Target/Unknown/Sky in the archived data files. In this project we corrected this with specially written software. Unfortunately, this labelling problem has persisted but the current versions of the MATISSE pipeline procedures detect and correct it automatically.

2.2.3 Pipeline Data Reduction

We removed known instrumental signatures from the remaining observations and extracted interferometric measures with the procedures described below. We then selected reliable exposures based on signal/noise (S/N) criteria, and finally calibrated the reduced measures using the calibrator observations. The output measures are the squared visibilities for each baseline, the closure phases for each triplet of baselines, and the photometric (one-telescope) flux, averaged over the set of exposures. Each measure is a function of wavelength.

We reduced the data with the ESO/MATISSE `esorex` Data Reduction Software (DRS) pipeline. The DRS evolved during our analysis to accommodate newly discovered issues, and different versions were applied to different data subsets. For most LM -band data we used version 1.4.1; for N -band data and “binned” LM -band data we used version 1.5.2 which contained new algorithms specific to these modes. As of version 1.5.2 of the DRS, the phase convention of the N -band is flipped with respect to the phase convention of the L - and M -bands. This causes a 180° rotation in the final images (see §2.2.7). The data and images of this paper have been corrected for this.

Here we have reduced data with the incoherent method in contrast to the previous $10\ \mu\text{m}$ MIDI data that was reduced in coherent mode (e.g., López-Gonzaga et al., 2014), because at the current time the incoherent method has been better validated. In the context of the MATISSE DRS, we specified this procedure by using the parameters `corrFlux=FALSE`, `OpdMod=FALSE` and `compensate=[pb,nl,if,rb,bp,od]` for the master recipe `mat_raw_estimates`. The reduced squared visibilities (V^2) and closure phases are shown in Figs. 4.6 and 4.7.

Extracting closure phases requires taking the triple product of the three baselines included in a “triangle” of three telescopes. This non-linear procedure greatly magnifies the effect of noise if the noise per frame and per spectral channel is $\ll 1$. We mitigate this effect by binning the raw data in the spectral direction before extracting the triple products. This binning can however reduce the visibility contrast when the interferometric phase varies rapidly with wavelength. In the low-resolution mode that we used, both detectors are oversampled, and a modest binning increases the S/N while avoiding contrast loss. For *LM*-band we use the default DRS bin size of 5 pixels ($0.25 \mu\text{m}$) for the visibilities and closure phases; in the *N*-band we use 7 pixels ($0.3 \mu\text{m}$) for the visibilities. We use 21 pixels ($1 \mu\text{m}$) for the closure phases in the *N*-band. However, at short *L*-band wavelengths ($\sim 3.2 \mu\text{m}$), the low single exposure S/N yielded poor closure phase accuracy. We separately reduced the data at $3.2 \mu\text{m}$ using a DRS binning parameter (`-spectral_binning`) of 11 pixels ($0.5 \mu\text{m}$) for closure phases, retaining the original reduced data at 5-pixel binning for the visibilities. We find that over the entire band, closure phase uncertainties were decreased from $\sim 10^\circ$ to $\sim 3^\circ$ on as a result of binning.

2.2.4 Data selection

After the pipeline reduction, we selected reliable sub-exposures by evaluating the frame-to-frame stability of both the photometric and interferometric data. Low photometric fluxes generally indicate a failure of the VLTI Adaptive Optics (AO) system to correct the atmospheric seeing. Low interferometric fluxes can be caused by either AO failure or extremely unstable atmospheric conditions. The versions of DRS that we use does not support this frame-by-frame inspection; therefore, we used the Expert Work Station package adapted from the MIDI reduction system (EWS; Jaffe, 2004). Future versions of DRS will include this statistical grading option. The final post-processing of the accepted data sets was performed in DRS.

Using EWS, we flagged a frame as acceptable if it met selection criteria based on its associated photometry and correlated flux.

A photometry frame was rejected if the photometry counts in a frame were less than either:

- 0.5 times the 80th percentile of counts for the single-dish telescope flux
- 6 times the RMS noise for all frames

For interferometric frames, we first estimated the correlated flux as the peak of the Fourier transform of the data as a function of wave number. We reject frames in which this peak is less than either:

- 0.001 times the 85th percentile of all baselines and frames (to eliminate weak “ghosts” caused by internal reflections)
- 0.5 times the 85th percentile of fluxes for this baseline or less than 7 times the rms noise (*LM*-bands) or 9 times the noise (*N*-band). This eliminates frames where the atmosphere was exceptionally unstable or the AO system failed.

Any sub-exposure in which more than 50% of the frames passed both criteria was accepted for DRS processing. In general, either just a few or almost all frames of a given sub-exposure are adversely affected by AO or atmospheric problems. The above criteria gave a clean separation between these two cases.

2.2.5 Calibration of the target measures

For the calibration of the interferometric data of NGC 1068 we used the `esorex` recipe `mat_cal_oifits`, including only the selected sub-exposures of the target and calibrators. The different BCD settings that we used have different instrumental signatures, and are calibrated separately. We then averaged together the calibrated visibilities and closure phases for the various settings. An in-depth description of the BCD calibration and the sign convention can be found in the detailed AMBER data reduction (Millour et al., 2004) and in Ch. 3. The final calibrated squared visibilities and closure phases are plotted in Figs. 4.6 and 4.7.

2.2.6 Photometric Calibration

While MATISSE is an interferometric instrument, in SIPHOT mode, it also captures a single dish spectrum from each individual telescope. For an extended target, such as NGC 1068, this represents the flux through a circular aperture of diameter 0.135". This low-resolution spectrum has 64 spectral channels between 2.9 μm and 5.0 μm (the *L*- and *M*-bands). To connect these spectra to an absolute flux scale, we observed two standard stars with known model-atmosphere spectra from the Cohen catalogue (Cohen et al., 1999). These stars were HD 9362 (del Phe), observed on both observing nights, and HD 188603 (b Sgr), on September 25th only. These stars are too well-resolved to serve as interferometric calibrators. In the *N*-band we followed a similar procedure, using the Cohen calibrator HD 18322 (Eta Eri) as standard. On each night the photometric spectra from the interferometric calibrator stars were cross-calibrated with the Cohen standard stars to bring them to an absolute scale. These spectra were then applied to the instrumental photometric spectra of NGC 1068 to bring them to the same scale. All of the spectra so calibrated were averaged together to give the final target spectrum, and the errors in the process were estimated by variations between the individual spectra. The final photometric spectrum of NGC 1068 is shown in Fig. 4.8, with comparison to broad-band photometric estimates from other instruments.

2.2.7 Determination of Image Orientation

In complicated interferometric systems, there is occasionally at commissioning time an uncertain phase sign ambiguity which translates to a 180° rotation ambiguity in reconstructed images. To validate the proper orientation of each band, we both performed image reconstruction and fit two Gaussian sources to the known binary del Sco. This unequal binary was observed on 19 May 2018 during MATISSE commissioning. Based on previous data (Meillard et al., 2011), we expect the separation at this time to be 178 mas with the fainter star almost directly to the North of the brighter star. In the model fitting, however, we do not enforce this assumption, and instead use uniform priors over the parameter space. Both image reconstructions and Gaussian modelling of this source reproduce the angular separation of the binary. From these reconstructions/models we find that the *N*-band is flipped 180° relative to the *LM*-bands. We correct all images and models in this thesis for this artificial rotation.

2.2.8 Calibrated V^2 and Closure Phase

From the accepted exposures, we obtained 48 independent squared visibility measurements (24 in each band) and 16 (*LM*-band) + 17 (*N*-band) closure phases. Figs. 4.6 and 4.7 plot the calibrated squared visibilities and closure phases as a function of wavelength, along with model values provided by the multi-Gaussian fits, described in §2.3.5, for the same *uv*-positions.

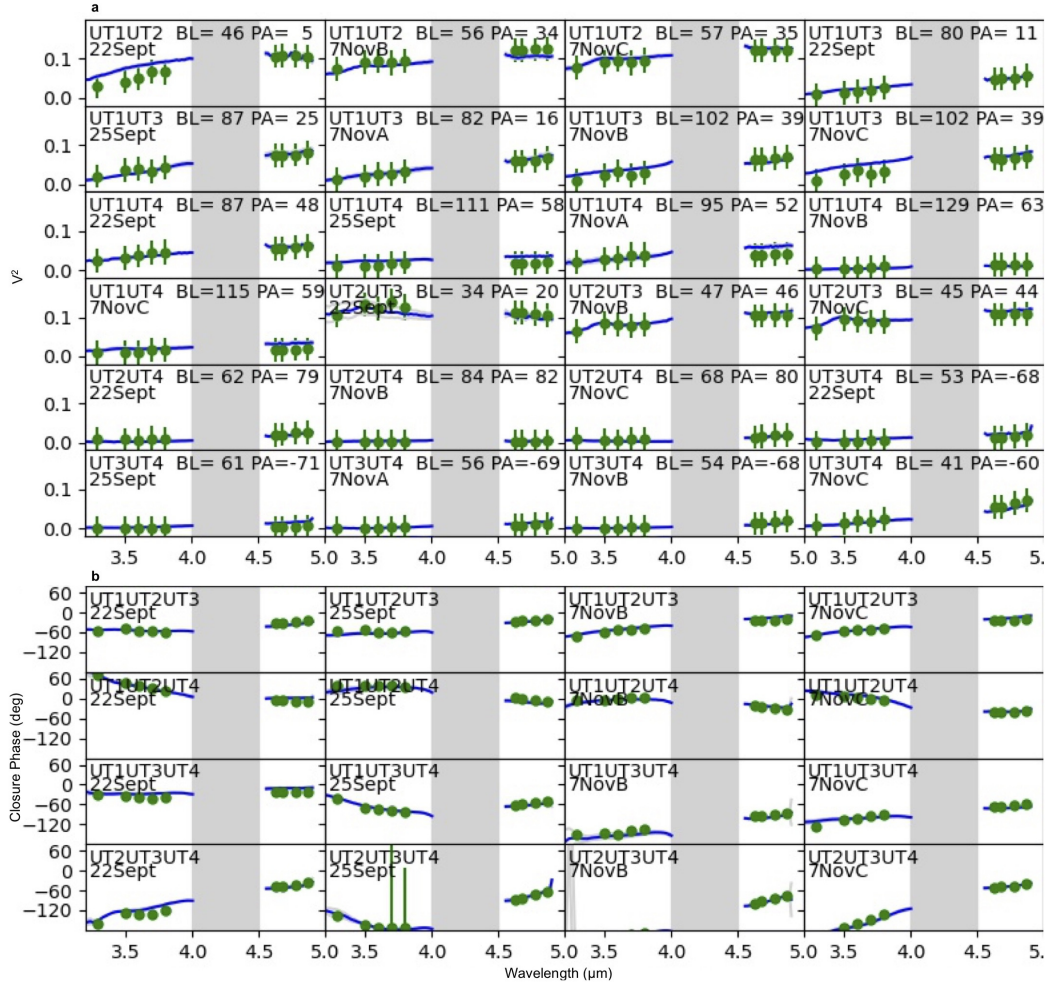


Fig. 4.6: (a) Squared visibilities for NGC 1068 in the L - and M -bands. The blue lines show observed values, averaged over sub-exposures; the thin gray lines show individual sub-exposures in order to illustrate the measurement uncertainties, but are often hidden behind the blue lines. The green points with error bars show values predicted by the multi-Gaussian models. The error bars represent the r.m.s. sum of the measurement errors and the uncertainties of the model parameters. The distance between models and observations shows that a limited number of Gaussians cannot exactly represent the true sky or that we do not have a sufficient uv -coverage and/or resolution. The gray bands mark the atmospheric non-transmission band. The labels indicate the telescope pairs for each baseline, the baseline length (m) and position angle (degrees), and the specific exposure label from the observation log (the observing log and reduced data can be found here: https://github.com/VioletaGamez/NGC1068_MATISSE). (b) Closure phases (degrees) using the same color code as above. The labels indicate the telescope triplets and the specific exposure label from the observing log.

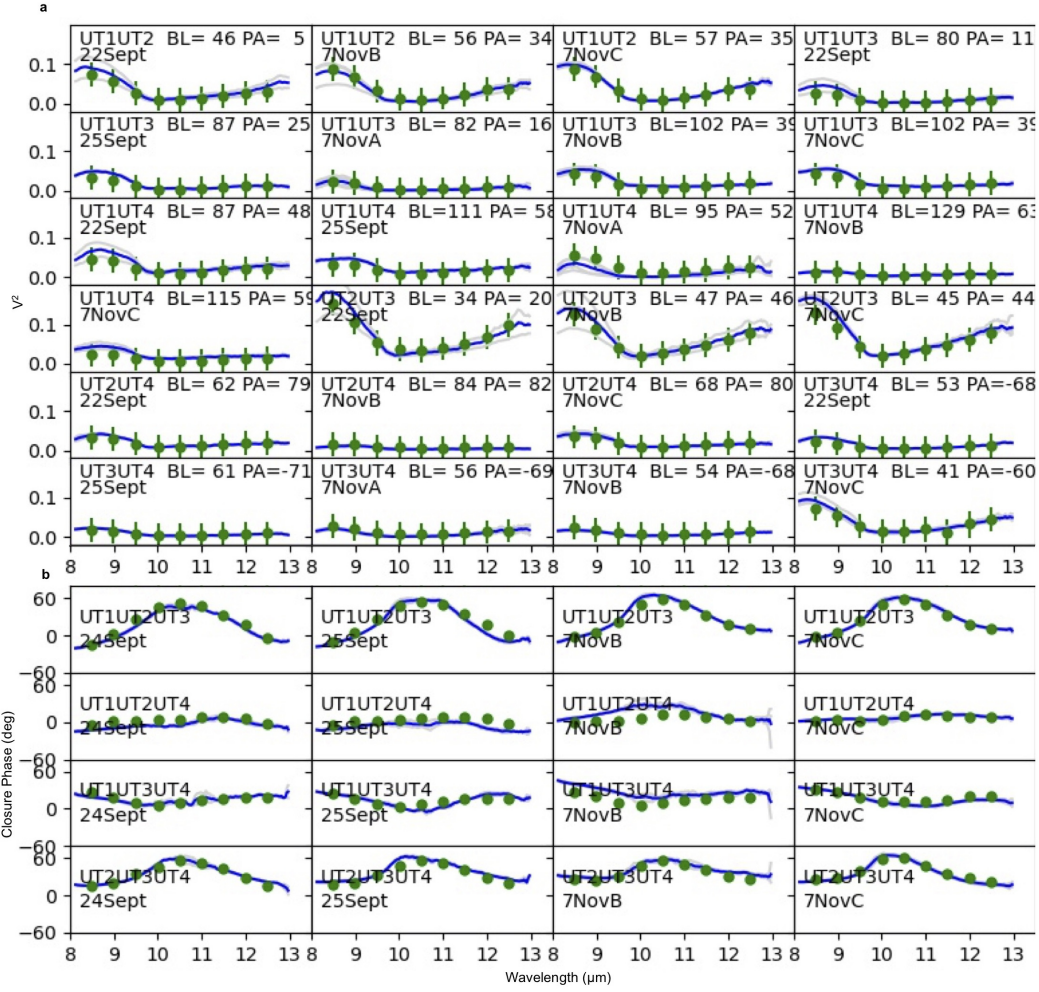


Fig. 4.7: The quantities plotted and the symbols used are the same as Fig. 4.6 but for the N -band.




Fig. 4.8: (a) Average single telescope spectrum of NGC 1068 in LM -bands (black solid line) and N -band (blue solid line). The error bars represent uncertainties estimated from the differences between different dates and calibrators. The yellow stars refer to VLT/ISAAC L' - and M -band single-dish flux estimates from Isbell et al. (2021, see Ch. 2), while the green triangle corresponds to a VLT/NACO M -band flux from Prieto et al. (2010). (b) The silicate absorption feature observed on two baselines at high spectral resolution ($R \sim 300$) during a single MATISSE commissioning snapshot. The long (85 m) baseline shows the broader, double-peaked profile characteristic of crystalline, reprocessed grains (van Boekel, 2004) while the short (~ 30 m) baseline does not. The difference between the curves shows that the crystallinity varies over the source.

2.3 Morphological Analysis of the Data

2.3.1 Introduction

Figs. 4.6 and 4.6 demonstrate that the visibilities vary strongly with uv -position, reaching near-zero values at position angles near 80° . This is a clear sign that we are resolving the object. Additionally, the closure phases display large non-zero values varying strongly with wavelength and with very different behaviours among the four different triplets of telescopes. This indicates that the target is strongly asymmetric.

Historically, optical/infrared long baseline interferometry has yielded too few uv -points to provide the direct image synthesis that is common at radio wavelengths, but recent instruments such as GRAVITY and MATISSE were designed to collect uv -data more efficiently. Experiments with poor or sparse uv -coverage are usually analysed by parameterizing rather limited models to fit the data; from uv -rich data direct images can be examined. Our current dataset is of intermediate richness and invites experimentation into the optimum methods for extracting physical quantities. In this section we present four methods that we applied to obtain detailed morphology and reliable fluxes from the instrumental measures: MIRA, IRBis, a multi-Gaussian model and a point-source model.

2.3.2 Interferometric Image Reconstruction

Aperture-synthesis imaging with optical/infrared arrays, such as the VLTI, is an ill-posed inverse problem. Images are reconstructed by minimizing a cost function that includes both data and regularization terms. The data term contains the measured data (for example, closure phases and visibilities), and the regularization term consists of some *a priori* assumptions about the target, for example, positivity or smoothness. In the last three decades, several image reconstruction algorithms were developed for aperture-synthesis imaging with infrared arrays. The performance of these algorithms was evaluated in several blind tests initiated by the International Astronomical Union Working Group on Optical/IR Interferometry (Lawson et al., 2004a; Cotton et al., 2008a; Baron et al., 2012; Sanchez-Bermudez et al., 2016).

In this paper, image reconstruction of the central thermal dust emission of NGC 1068 was performed with two software packages, MIRA and IRBis, to verify the reconstruction results. The details of the packages and the application to our data are described below.

2.3.3 Image Reconstruction with IRBis

The IRBis software package is part of the MATISSE data reduction pipeline. For the reconstruction process we tested the two built-in minimization engines, the three cost functions, and the six regularization functions (Hofmann et al., 2014).

The current version of the MATISSE image reconstruction package IRBis consists of two optimization routines: a) a large-scale bound-unconstrained nonlinear optimization routine called ASA_CG (engine 1; Hager & Zhang, 2005) and b) the new built-in optimization routine, a limited memory algorithm for bound-constrained optimization called L-BFGS-B (engine 2; Byrd et al., 1995). The version of IRBis described in Hofmann et al. (2014) contains two different data terms in the cost function. The errors of the calibrated data used in the image reconstruction were derived from the scatter of the different data sets recorded during one observation, each lasting about 10 min. All data sets of each observation were used separately. The starting image of each reconstruction attempt was a circular Gaussian, fitted to the measured visibilities of the target. Two priors were tested: a) this fitted Gaussian and b) a constant. From the many reconstructions that were calculated, the optimal was selected using the image




Fig. 4.9: In order to quantify the fidelity of the reconstructions shown in Fig. 4.1, we performed analogous reconstruction on an artificial model. The model consisted of seven Gaussians, similar to our multi-Gaussian model for the dust emission. We simulated visibility and closure phase data for this model for our uv -coverage; we added noise to the simulated data similar to that in the observations. We then performed image reconstruction using IRBis with identical reconstruction parameters to those used in Fig. 4.1. (a) The input 7-Gaussian model. (b) The IRBis reconstructed image. (c) The reconstructed image minus the input model. In all cases the color scale represents the fraction of the peak intensity of the original model. The r.m.s. errors in the residual maps were 2.3% of the peak brightness. This indicated that most of the artifacts present in Fig. 4.1 result from the uv -coverage rather than noise on the observed quantities. In Fig. 4.1 we have drawn white contours at $3\sigma = 6\%$ of the peak. Features brighter than this certainly represent true source emission.

quality measure, q_{rec} . The quality measure q_{rec} consists of the reduced- χ^2 values of the closure phases and visibilities and a so-called residual ratio. Because of the sparse and inhomogeneous uv -coverage of the data, it was helpful for the image reconstruction process to reduce the weight measures at uv -points in dense clusters relative to measures in low density regions of the uv -coverage. This was done with different powers of the inverse uv -density.

LM-band: The field of view and pixel grid used for the reconstructions in the L and M bands was 340×340 mas and 256×256 pixels, respectively. Image reconstruction experiments were performed in the wavelength ranges $3.05\text{--}3.4 \mu\text{m}$, $3.4\text{--}4.0 \mu\text{m}$, and $4.55\text{--}4.91 \mu\text{m}$. The radius of the applied image space mask was increasing between 70 mas and 100 mas during each reconstruction run. Because of the large pixel number, only engine 2 (L-BFGS-B) was used because engine 1 (ASA_CG) is computationally expensive for large numbers of pixels. Fig. 4.1 and Fig. 4.11 show the reconstructions obtained from data recorded with the UTs. The reconstruction parameters and quality are listed in Table 4.2. The reconstructions were not convolved with the interferometric resolution and have spatial resolutions of $\sim \lambda/2B$ (where B is the length of the longest baseline observed), corresponding to 2.6 mas, 3.0 mas, and 3.8 mas for the respective wavelength bins.

N-band: For the N -band reconstruction, the field of view and pixel grid used were 360×360 mas and 128×128 pixels, respectively. The image reconstructions were performed in the wavelength range $8.0\text{--}9.0 \mu\text{m}$, $10.0\text{--}11.0 \mu\text{m}$, and $11.5\text{--}12.5 \mu\text{m}$. The radius of the space mask was increasing between 160 and 184 mas during each reconstruction run. Fig. 4.11 shows the N -band reconstructions obtained from data recorded with the UTs. The reconstruction parameters and quality are listed in Table 4.2. The reconstructions are not convolved with the interferometric resolution and have spatial resolutions of $\sim \lambda/2B$, corresponding to 6.7 mas, 8.3 mas and 9.5 mas, respectively.

Estimates of IRBis image artifacts: To quantify the reliability of the features seen in reconstructions shown in Fig. 4.1, we performed an analogous reconstruction

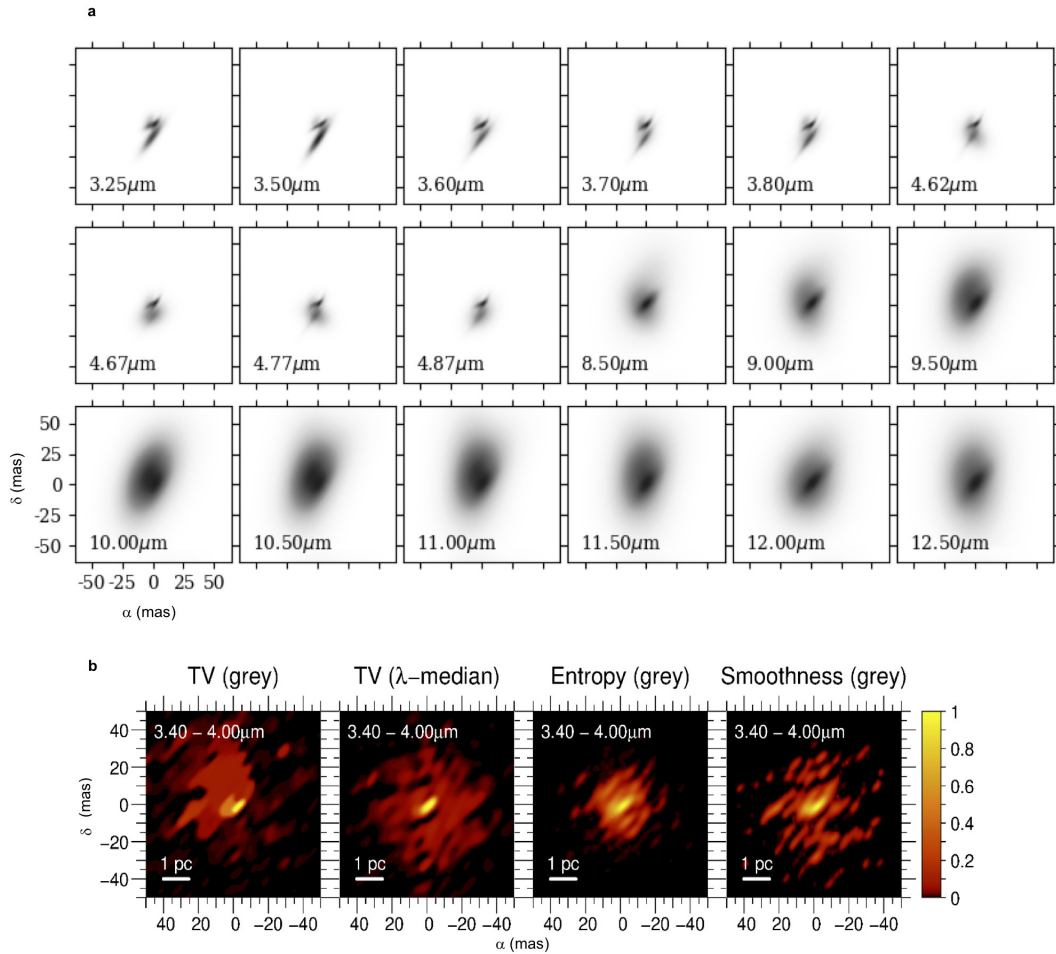


Fig. 4.10: (a) Image representation of models obtained from the Gaussian modelling approach described in §2.3.5. We use square root intensity scale. (b) Image reconstruction with MIRA using different approaches (§2.3.4). From left to right: using total variation (TV) regularizer, on a large bandwidth (‘gray’ reconstruction); using the same regularizer, but independently reconstructing images at each wavelength, and computing a median over the wavelength interval; using maximum entropy regularizer (gray reconstruction); using smoothness regularizer (gray reconstruction).

on an artificial model. The artificial model consisted of seven Gaussians similar to our multi-Gaussian model for the galaxy (see §2.3.5). We created visibility and closure phase data for this model with the same uv -coverage as our galaxy data, and similar noise. We performed a reconstruction with IRBis with identical regularization models and parameters. Images of the input model, the reconstruction and the residuals (reconstruction - input) are shown in Fig. 4.9. This process indicated that most of the artifacts visible in Fig. 4.1 are caused by the limited uv -coverage, rather than by noise, and that the r.m.s. values of these artifacts was $\sim 2\%$ of the peak brightness. In Fig. 4.1a–d we have drawn white contours at $3\sigma = 6\%$ of the peak. Features brighter than this almost certainly represent true sky emission.

2.3.4 Image Reconstruction with MIRA

The MIRA image reconstruction package is developed and distributed by Thiébaud (2008). It works directly on closure phases and squared visibilities, computing χ^2 for each quantity by using an approximation on the noise model that linearizes these quantities in the complex plane. The minimization routine in MIRA, VMLM-B, is a variable metric method with limited memory and bound constraints on the input parameters, which is

basically a variant of the BFGS algorithm.

MIRA does not introduce weighting in the uv -coverage as in IRBis, nor a global parameter-search process. No weighting can produce many artifacts in the reconstruction process when using a very sparse uv -coverage. Therefore, two methods were used to reduce and assess these artifacts in the reconstructed images.

Low-frequency filling (LFF; Millour et al., 2008) was used to extrapolate the visibilities at low frequencies, where the visibilities variations are deterministic (Lachaume, 2003). This method usually strongly reduces the ‘ghost image’ artifacts in the image reconstruction process. We experimented with several LFF sizes, and used at the end a size of 80 mas. We reconstructed images in the same central wavelengths and bandwidth as with IRBis for a direct comparison of the result.

The MIRA images shown in Fig. 4.11 were reconstructed using a total variation regularizer, hyperparameter value of 10,000, 128 pixels, and a pixel size depending on the wavelength: 1 mas for the L -band, and 3 mas for the N -band. The start image was random, and the prior image was a Gaussian with a full-width at half-maximum (FWHM) of 32 mas (LM bands) and 96 mas (N -band).

The second method we experimented with was to reconstruct images at each MATISSE-provided wavelength independently and computing a median image across a defined band, selected for the previous MIRA and IRBis reconstruction. This method is known to vastly reduce, but not completely remove, the typical image reconstruction artifacts that arise from a sparse uv -coverage. This method helped us distinguish the structures we consider as ‘real’ from other structures that clearly come from uv -coverage artifacts. We show a comparison of broadband MIRA reconstructions in the 3.4–4.0 μm band in Fig. 4.10. The second image displays the median of a series of narrow-band images within the total band, the others show images with constructed with different regularizers that attempted to fit all wavelengths in the band simultaneously, taking account of the different effective resolution at each band. The median image better matches the features found in other bands and by other reconstruction techniques.

2.3.5 Multi-Gaussian Model Fitting

We adapted the well established method of Gaussian model fitting directly in the uv -plane to estimate fluxes. This is analogous to the method used in previous MIDI analysis of NGC 1068 (López-Gonzaga et al., 2014).

There are several differences in this work: we fitted various components based on the shape of the main sources found through image reconstruction; we use squared visibilities and closure phases; and we use a more sophisticated Markov chain Monte Carlo (MCMC) Bayesian methodology to estimate parameters and their uncertainties (using `emcee` from Foreman-Mackey et al., 2013). We fitted independent models to a selection of wavelengths across the entire wavelength range provided by MATISSE. Unconstrained fitting with many parameters leads to low residuals, but also to poorly conditioned and degenerate solutions. To avoid this, we imposed several constraints. First, the positions and shapes of the components had to approximate the images generated by the reconstructions. Secondly, we enforced local continuity in wavelength by making simultaneous fits at two (in L -band) or three (in M -band) different wavelengths, with 0.05 μm separations along a large range of wavelengths.

We determined the minimum number of Gaussians in each band which resulted in both closure phase residuals $< 10^\circ$ and squared visibility residuals less than 0.1 and still visually resemble the morphology revealed in the image reconstruction. For the LM bands we needed seven Gaussians to recover the structures; for the N band we only needed three Gaussians. Each Gaussian has six free parameters: its (l, m) position, relative flux, FWHM major and minor axes, and the position angle of its major axis.

The MATISSE visibilities and closure phases are not sensitive to the absolute position of the source complex on the sky or to the sum of the fluxes of the components. To remove these non-measured degrees of freedom from the model, we fix one Gaussian component to the center of the image, and the positions of the center of the other components are specified by their distances and position angles relative to the first. We also define the central brightness of each component relative to the fixed Gaussian.

We calculate the squared visibilities and closure phases from the best-fit model intensity distribution and estimate uncertainties for these quantities from the parameter posterior probability distributions. These are shown in Figs. 4.6 and 4.7. We use a similar process to estimate the uncertainties of the measured spectral energy distributions (SEDs) discussed in §2.5. The list and definitions of fitted parameters are in Table 4.2.

In Fig. 4.10 we show the best-fitting models at all modelled wavelengths. In Fig. 4.11 we show the best-fitting models at four chosen wavelengths for comparison to other methods, 3.6, 4.77, 8.50 and 11.5 μm . We show the best-fitting parameters for the Gaussian model at three wavelengths, one in each band, in Table 4.2. In the L and M bands we can identify two main structures separated primarily along north–south axis; we refer to them as NC (northern complex) and SE (southern extension). In the N -band, at 12 μm , the distribution of the intensities given by our three Gaussians look similar to the two central sources of model 1 of López-Gonzaga et al. (2014). We compare the model squared visibilities and closure phases to the data in Figs. 4.6 and 4.7. The average r.m.s. of the residuals are 8.2, 7.6, and 8.3° in closure phase, and of 0.01, 0.01, and 0.01 in V^2 , in the L , M , and N bands, respectively.

In the Gaussian modelling, we did not require a background component, owing to the already extended nature of the Gaussians. As a check, we performed a similar fit with an additional extended component, but in all cases the resulting residuals were larger. This indicates that the majority of the flux comes from the Gaussians in our models at the different wavelengths. We therefore distribute the single-dish flux among all the components fitted at each wavelength. When determining the photometry, the value and uncertainties are calculated by bootstrapping, with replacement, the last 4,000 samples of the associated MCMC chain.

2.3.6 Point-source Model Fitting

We further model the interferometric data using a point-source modelling method analogous to clean-like image reconstruction (Högbom, 1974). The point-source model allows us to create an independent model image of the data, to be compared to the image reconstruction, while also providing photometry at the same narrow wavelength bins as the Gaussian modelling.

The point-source model: The base-model fitting method is described in previous work (Leftley et al., 2021). We use the method that fits an image consisting of N_p point sources to the squared visibility and closure phase data. Unlike the described model, we also include an extended background component and do not include the scale factor (which is not relevant for such a bright object). We chose the extended background source to be an elongated Gaussian, unfixed in position and amplitude, with a fixed major-axis FWHM of 70 mas, minor-axis FWHM of 50 mas, and major-axis position angle of 0°. The background Gaussian’s fixed parameters were determined from the N -band Gaussian models. Each image is 256×256 pixels, we use 200 walkers and 1,500 iterations.

Verifying the image reconstruction: When determining geometry independently of the image reconstruction with the point-source model, we used the random starting position option for the walkers. We separated the observed data into 0.1 μm wide bins centered at 3.55 μm , 3.65 μm , 3.8 μm , 3.9 μm , 4.65 μm , and 4.8 μm . The data in these

bins are not averaged during fitting but instead used to extend coverage of the uv -plane, that is, assuming that the different wavelengths see the same structure at a different angular resolution.

Initially, we fitted the 4.8 μm bin. The first of the two fits performed when using a random starting position was performed for the N_p range of three to 16 sources. It was found that the preferred value of N_p for the chosen bin width is 9. We used the result as the starting point of the second fit for the 4.8 μm bin as well as the other wavelength bins.

Point-source photometry: The $N_p = 9$ point-source model is a good representation of the observables and sufficient to compare geometry but the discrete nature of the model introduces uncertainty to any fixed-aperture photometry. To solve this, we overfitted the data. Each of the nine point-source starting positions from the geometric modelling was tripled and randomly shifted in position by a normal distribution of 0.2 mas, giving 27 point sources. An $N_p = 27$ fit for each wavelength bin was performed with this starting position. The resulting models are visually similar to the geometric models and a good description of the data. The photometry and its uncertainties at each wavelength are calculated by performing aperture photometry on a bootstrapped, with replacement, selection of the 100,000 samples provided by the last 500 iterations of the 200 walkers.

2.3.7 Comparison of Reconstruction Methods

Morphology: We find excellent agreement in source morphology using the various modelling and reconstruction techniques. The results of the four imaging techniques at different wavelengths are shown in Fig. 4.11. The L -band images have the most substructure as a result of the smaller angular resolution, so we focus our quantitative comparison on this band. In this band, regardless of approach, we find two main features: an incomplete ellipse with a major axis position angle of approximately -45° and a secondary southern source. In the image reconstructions we take a conservative uncertainty on the size measurements equal to 1 pixel (1.325 mas in IRBis; 1 mas in MIRA). The measured outer extent of the broken ellipse in each method are (minor \times major): 8.1×13.8 mas with a position angle of -45° using IRBis; 7.5×13.2 mas with a position angle of -45° using MIRA; 7.6×12.2 mas with a position angle of -40° using Gaussian fits; and 6.5×13.6 mas with a position angle of -45° using point-source reconstruction. We also measure the distance between the center of the broken ring and the center of the secondary source: 15.6 mas with IRBis; 17.13 mas using MIRA; 15.3 mas using Gaussian modelling; and 17.2 mas using point-source reconstruction. We note that the southern extension in the Gaussian modelling is much larger than for the other reconstruction methods. We hypothesize that this is a result of the simplicity of a Gaussian component, and the true southern source may have multiple components as seen in the MIRA, IRBis and point-source reconstructions. Nonetheless, within uncertainties the dimensions of the broken ring and the separation of the sources agree remarkably well, especially considering our sparse uv -coverage. We see a similar level of consistency within the M -band and between the two bands. More qualitatively, we find less agreement in the N -band, especially at longer wavelengths where MIDI revealed considerable extended flux to the north (López-Gonzaga et al., 2014). With baselines exclusively > 30 m and the smaller field of view of the UTs relative to the MIDI AT observations, we expect to resolve-out this extended component, but we cannot rule out that it contributes to the MATISSE N -band flux. In particular, this extended component can be broken up into ‘noise’ peaks in image reconstruction. With this in mind, we find a consistent position angle -45° bright central structure in each N -band reconstruction. It is accompanied in each case by extended flux: to the north in both MIRA and Gaussian modelling; and




Fig. 4.11: From left to right: the MIRA image reconstruction, the IRBis image reconstruction, the over-fitted point source model (convolved with the beam), and Gaussian model for four selected wavelengths. The plot uses a 0.6 power color scaling for visual purposes. Each method reveals similar structures and morphology.

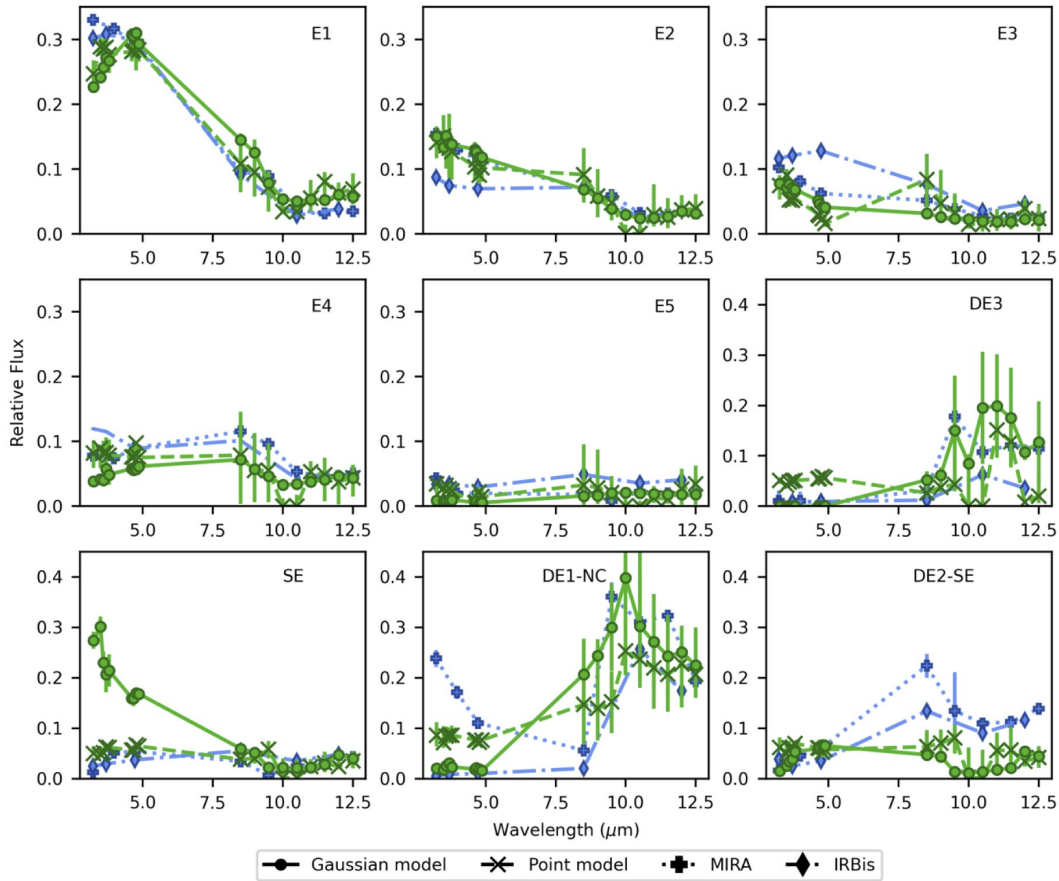


Fig. 4.12: Comparison of extracted photometry using different images/models of NGC 1068. The photometry between methods generally agrees in spectral shape between different methods. The most notable exceptions are SE and E3 which still produce similar temperatures from SED modelling between methods.

roughly evenly north–south in the IRBis reconstruction and point-source reconstruction. We can conclude only that the central bright source is consistent and that it must be accompanied by extended flux. At $8.5\mu\text{m}$ where this extended background is less prominent, we find excellent agreement between the techniques: a central ~ 15 mas (1 pc) component with a position angle of -45° flanked to the north and south by fainter extended flux. In the IRBis images, we see hints of the broken ring continuing from the *LM*-band. As we see consistent morphology across both methods and wavelengths, we report the first images of this inner region of NGC 1068 with high confidence.

Photometry: We compare the photometry from both models and the image reconstructions for the apertures: E1 to E5, DE3, southern extension (SE), DE1 with northern complex contributions removed, and DE2-SE (see Fig. 4.2). The results are shown in Fig. 4.12. The photometry shown is the relative flux, under the assumption that all flux in the image sums to one. Because the apertures do not cover the whole image, all apertures for one image may not sum to one.

In the northern complex (E1-E5), we find good agreement in all ellipses with all methods. There is a slight discrepancy in the IRBis flux in E3, primarily in the *M* band. Overall, we find that the photometries from the independent methods are similar in spectral shape and relative flux. The similar spectral shape will lead to similar intrinsic temperatures, absorption signatures and chemistry. This suggests that the photometry used for the associated SEDs is reliable and mostly insensitive to methodology.

For the southern extension, we find good agreement between the images and point-

source model, but the Gaussian model shows much larger fluxes in the L and M bands. We tested whether this was a background effect by introducing a background component to the Gaussian model that matched the one included in the point-source model. We found that this results in a poorer fit and no important change in the LM -band fluxes. When we use the background-included model fluxes for the SED fitting (assuming no carbon), we find similar temperatures. Furthermore, the inclusion of the background did not affect the extracted photometry of the northern complex. Finally, we also fitted an SED, assuming no carbon, created from the southern extension photometry of the point-source model. We find a hot component temperature of 808–846 K instead of 882–974 K, which is cooler but not in conflict with the results. We hypothesize that the difference between methods for the southern extension is due to a limitation of Gaussian modelling, that is, the southern extension is complex and not well described by a Gaussian. For the diffuse emission (DE) apertures, we find relatively large scatter. Although more scattered, the DE apertures have similar overall spectral shapes between methods.

In conclusion, we find that the aperture photometry is suitably reliable for our scientific undertaking and structurally consistent between the methods presented in this work. Although extracting the fluxes from any method should produce suitable SEDs for the temperature modelling, we choose to use the Gaussian modelling. Using a model instead of a reconstruction allows us to easily extrapolate through the model and data uncertainties using the derived parameter distributions from MCMC modelling as well as allowing narrower wavelength bands for more detailed SEDs. The Gaussian model specifically is the most comparable to previous studies of NGC 1068 with MIDI (Jaffe et al., 2004; Raban et al., 2009; López-Gonzaga et al., 2014). Furthermore, it provides fine control over the brightness distribution, which is useful when determining the importance of revealed structures.

2.4 Radio Data

The nucleus of NGC 1068 was observed using the High Sensitivity Array (HSA), which includes the Very Long Baseline Array (VLBA), the phased Very Large Array (VLA), and the Green Bank Telescope (GBT) (J. Gallimore, priv. comm.). The observations took place on 8-9 February 2020 and 21-22 March 2020. The receivers were tuned to the 22.235080-GHz transition of water vapor, offset to the redshift of NGC 1068. The details of the observations, and particularly the analysis of the detected H_2O masers will be published separately. Data reduction, including bandpass calibration, fringe-finding and aperture synthesis imaging were performed using standard techniques in the Astronomical Image Processing System (AIPS; Wells, 1985; Greisen, 1990, 2003). The final calibration of interferometric phase was based on the brightest (few hundred mJy) maser sources of NGC 1068, and phase solutions were transferred to the nearby astrometric calibrator J0239-02, located 2.7° from NGC 1068. By evaluating the offset sky positions on images of the astrometric calibrator, J. Gallimore and V. Impellizzeri (priv. comm.) were able to determine the absolute positions of the NGC 1068 maser sources to an accuracy of about 0.3 mas.

The final products are a data cube with axes of sky coordinates and radial velocity and a single-channel image of 22-GHz radio continuum. Continuum was detected in line-free channels only on the shortest interferometric baselines. To produce the continuum image, J. Gallimore and V. Impellizzeri (priv. comm.) averaged the line-free channels at the ends of the observed bandwidth and applied a Gaussian taper with 50% weight at 30 million-wavelengths during Fourier inversion. The resulting continuum image was then deconvolved using the multiscale CLEAN algorithm (Cornwell, 2008) and restored with a Gaussian beam of 4.3×3.3 mas, at a position angle of -21° . This image is displayed as

contours in Fig. 4.4b. The peak brightness in this image is 3.0 mJy per beam. The total flux density of the recovered 22-GHz continuum is $S_\nu = 13.8 \pm 0.3$ mJy. The centroid position of the resolved continuum is $\alpha_{J2000} = 02 : 42 : 40.70901$, $\delta_{J2000} = 00 : 00 : 47.9448$.

This bright central radio emission is almost certainly free-free radiation from gas hotter than 10^6 K because of its flat spectrum between 5 and 22 GHz, its brightness temperature of $> 2 \times 10^6$ K at 5 GHz and its inverted spectrum below this frequency (Gallimore et al., 2004). The nature of the diffuse 256 GHz emission (Impellizzeri et al., 2019) seen in Fig. 4.4 a is less clear; it may be a combination of free-free emission and the long-wavelength tail of the infrared dust emission. The total 256-GHz flux is 12.7 ± 0.1 mJy, composed of an unresolved point of 6.6 mJy and a diffuse component of 6.1 mJy (Impellizzeri et al., 2019). The peak brightness in Fig. 4.4c is 0.85 mJy per beam. The estimated total free-free flux, extrapolated from lower frequencies, is 13 ± 1 mJy, similar to the measured total flux (Cotton et al., 2008b). The blackbody fits of the diffuse *N*-band infrared emission can be extrapolated to 256 GHz (1.2 mm) and lead to an estimated flux of 30 mJy, which exceeds the measured flux, but this assumes that the emitting dust is optically thick even at millimeter wavelengths. This requires a large number of very large, that is, millimeter-sized, dust grains, which seems unlikely. Galactic dust particles have emission efficiencies that scale at large wavelengths $\sim \lambda^{-1}$ (Mathis et al., 1977). Thus if the 12 μm emission optical depth is not $\gg 1$, the expected 1.2 mm flux would be about 100 times smaller than the blackbody extrapolation. From these considerations we believe that the diffuse emission also arises most probably from hot gas rather than warm dust.

2.4.1 Radio-infrared Relative Positions

We determined the position of the VLBA 22-GHz peak relative to the 3.7 μm peak by a three-step process. First, we determined the displacement of the ALMA 256-GHz image (Fig. 4.4a) with respect to the MATISSE 12 μm image by cross-correlating the images. The formal error in this process—essentially the width of the cross-correlation function divided by its signal/noise ratio—was 1.6 mas. Secondly, we assume the MATISSE 12 μm peak coincides with the MATISSE 3.7 μm peak. The relative positions can be determined to an accuracy of < 3 mas. Lastly, we can align the VLBA 22-GHz image with the ALMA 256-GHz map to < 1 mas, based on their absolute astrometry. The three steps yield Fig. 4.4b, where the relative positions are accurate to ~ 3 mas.

2.5 Spectral Energy Distribution Modeling

2.5.1 SED Estimates

In §2.3.5, we produced Gaussian models that well reproduced the observed squared visibilities and closure phases. We multiplied the relative fluxes of each aperture (see Fig. 4.12) in each band by the photometric flux densities determined in Methods section ‘Data flagging, reduction and calibration’ (Fig. 4.8) for the same wavelength range. This results in the actual flux densities of the components in these bands. For the uncertainties, we propagated the derived error from the Gaussian model fluxes and the observed photometry. To determine the distribution of dust temperature, we defined nine main areas of interest (shown in Fig. 4.2), five small ellipses over the structure of the northern complex (E1–E5), one medium ellipse over the southern extension (SE), and three large ellipses (DE1–DE3) to encompass the extended emission that predominates at larger wavelengths. The small areas were defined to cover the features seen in the *L*- and *M*-band multi-Gaussian models, with as many ellipses as possible but larger

than the resolution limit in the N -band, so we could integrate their fluxes over the three bands. The large ellipses (DE1 to DE3) were defined to cover the remaining areas with the extended and/or diffuse emission in the N -band multi-Gaussian models, based on the narrow band models, so to recover any spectral changes along the N -band. Because the N -band multi-Gaussian models only used three components, we matched that number in ellipses to not to cause an over-sampling. Because of this, the latter do not perfectly gather all the flux or emitting areas at every wavelength, and also because the structures change substantially along the N band owing to the Si feature. We then summed the contribution from each Gaussian component in each of the defined areas to obtain spectral energy distributions (SEDs). The plotted points with error bars in ED Fig. 3 and Fig. 4.3 represent the estimated SEDs in the nine selected apertures.

2.5.2 SED Fitting

For each aperture we fitted the formula given in §1.3 with one blackbody where this was sufficient, or else with two blackbodies. For each set of measured fluxes, we performed a brute force evaluation of χ^2 over the range of parameters in the formula. We then determined the minimum value χ_{\min}^2 and the parameter envelope $\chi^2 - \chi_{\min}^2 < \sqrt{2n_p}$, where $n_p = 6$ is the number of fitted parameters. For a nonlinear, multi-parameter model, this envelope corresponds in confidence to the 1σ level for a single Gaussian variable. The shaded areas in Fig. 4.13 and Fig. 4.3 indicate the union of all SEDs of the models that met this criterion.

Opacity laws and dust mineralogy: The thermal modelling process requires the definition of the foreground opacity law, $\kappa(\lambda)$. The standard ISM extinction profile with a Mathis–Rumpl–Nordsieck (MRN) size distribution (Mathis et al., 1977), which resembles the extinction measured toward the Galactic center (Fritz et al., 2011), resulted in bad fits to the silicate absorption feature and the fluxes at $8 \mu\text{m}$. (see Fig. 4.3). We then tried fitting κ with a series of theoretical curves modelling amorphous silicate grains of various sizes with various ratios of pyroxene to olivine and of iron to magnesium (Min et al., 2005). Dust containing large amounts of pyroxene shows excessive absorption at $8 \mu\text{m}$. The same is true for particle sizes greater than approximately $5 \mu\text{m}$ and iron-poor silicate mixtures. We conclude that the absorbing material in front of the hot emitting regions contains predominantly Fe/Mg olivine particles of small to moderate size ($0.1\text{--}1.5 \mu\text{m}$). Owing to the $3.4 \mu\text{m}$ carbonaceous features present in the L band, we additionally tested varying amounts of amorphous carbon particles of the same size.

Two-temperature model results: We show in Fig. 4.3 and Fig. 4.13 the results of the two-component blackbody fits to the SEDs of the ellipses, for models using no carbon and models including 20% of carbon by weight. A larger carbon contribution produced considerably poorer fits. As more carbon is added, the temperatures of the hot component rise. We list in Table 4.2 the ranges of the values of the parameters for the models with fits inside 1σ at each wavelength. For the fits we set a maximum T_{hot} of 1500 K, a minimum η_{hot} of 10^{-5} (below that the contribution of the flux is irrelevant), and a minimum N_{ext} of $1 \mu\text{g cm}^{-2}$ (lower values become degenerate with the combination of high temperatures of ~ 1000 K). There is some cross-talk between the parameters N_{ext} and T_{hot} starting from around 900 K –that is, if we use a higher temperature for the hot component then we have to add a larger amount of extinction.

2.5.3 Dust composition

The composition, crystallinity and size distribution of the dust grains contain clues to their origin and thermal history. These in turn contain clues about the accretion

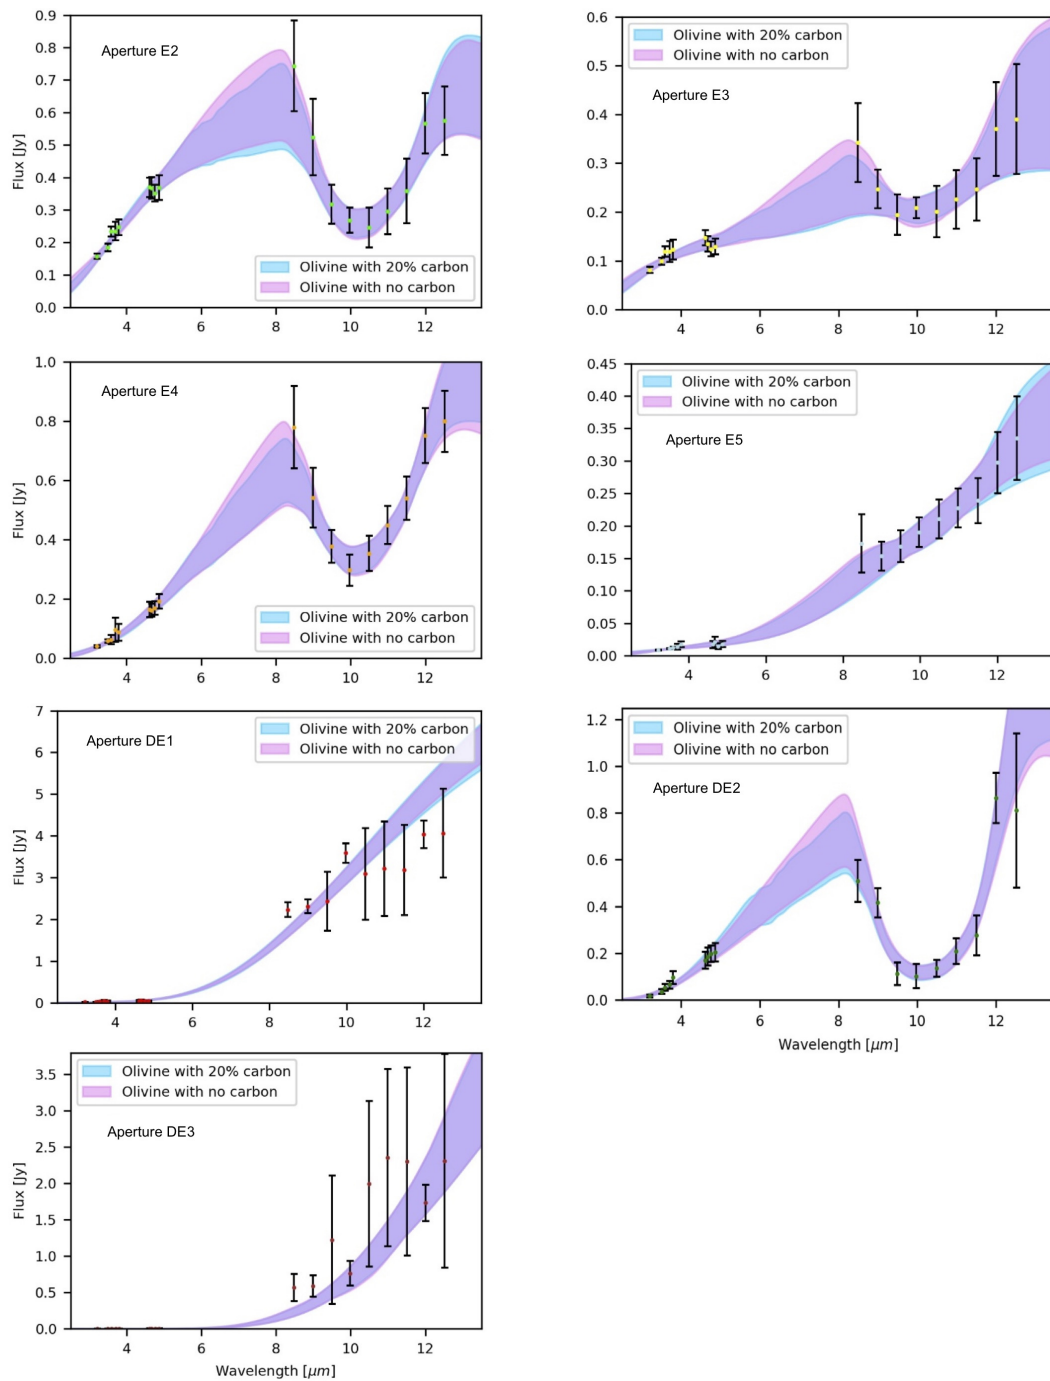


Fig. 4.13: The figures are labelled with the aperture names defined in Fig. 4.2. The shaded areas show all models falling inside 1σ of the photometry, considering both pure amorphous olivine (magenta) and a mix of olivine and 20% amorphous carbon by weight (cyan). The plots for apertures E1 and SE are in Fig. 4.3.

processes for dust and gas. The *N*-band silicate feature, discussed in §1.3 is striking in its weak opacity near $8\ \mu\text{m}$, in contrast to ‘standard interstellar’ curves or the opacity toward the Galactic center. In our Galaxy, low $8\ \mu\text{m}$ opacity characterizes the diffuse ISM, rather than dust in dense molecular clouds where grains are larger and have icy mantels that add opacity at this wavelength (Fritz et al., 2011). Our opacity curve in fact most closely resembles Galactic extinction curves measured in low-density regions far from the center (Zasowski et al., 2009). The comparison to theoretical curves indicates a higher olivine/pyroxene ratio than in the general Galactic diffuse medium, but the possible relation of this ratio to the ISM density or distance to the Galactic center is not well studied. Some authors have attributed the $8\ \mu\text{m}$ peak to SiC grains (Köhler & Li, 2010), and the $3.4\ \mu\text{m}$ C–H feature may hint at some carbon-rich absorbing grains, but the general good match to olivine rich silicates, and the crystalline features seen in higher-resolution spectra (see next section and Fig. 4.8) favor the silicate interpretation.

2.5.4 Thermally-Reprocessed Silicates

During the 24 September 2018 commissioning run, we took a MATISSE snapshot using the *N*-band HIGH resolution grating ($R = 300$). At this resolution some of the baselines showed a flattened, double-peaked, silicate absorption feature (Fig. 4.8). This spectral form indicates that the grains are crystalline rather than amorphous. Analogous to the grains near young stellar objects they have probably been melted and recrystallized in a hot environment (van Boekel, 2004). This feature can also be seen in previously published high-resolution MIDI data (López-Gonzaga et al., 2014).

a)		Wavelength range [μm]	Closure Phase error [deg] ⁽¹⁾	V^2 S/R ⁽²⁾	FOV [mas] ⁽³⁾	N_{pix} [pixels] ⁽⁴⁾	Cost func. ⁽⁵⁾	Regularisation ⁽⁶⁾	Engine number ⁽⁷⁾	uv-weight ⁽⁸⁾					
		3.05 - 3.4	15,1	8,6	340	256	1	+ 2	2	1					
		3.4 - 4.0	4	13,3	340	256	1	- 4	2	0,5					
		4.55 - 4.91	4,5	14,6	340	256	1	- 4	2	1					
		8.0 - 9.0	1,2	13,2	360	128	1	- 4	1	0,5					
		10.0 - 11.0	4,3	8	360	128	1	+ 2	1	1					
		11.5 - 12.5	2,4	11,6	360	128	1	+ 2	1	1					
b)		3.7 μm		4.67 μm		c)		8.5 μm		12 μm					
Gauss #	Param.					Gauss#	Param.								
1	ϑ	15,8	+ 4,0	- 4,0	17,7	+ 3,5	- 4,0	1	ϑ	50,4	+ 4,9	- 4,3	55,9	+ 4,4	- 4,6
	θ	9,8	+ 2,9	- 2,5	12,6	+ 2,1	- 2,0		θ	32,0	+ 5,1	- 5,3	37,0	+ 3,9	- 3,7
	PA	-38,1	+ 102,7	- 38,0	-57,6	+ 114,7	- 25,3		PA	-17,1	+ 14,0	- 15,6	-14,5	+ 12,7	- 11,5
	f_{rel}	3,4	+ 2,5	- 1,6	5,9	+ 1,5	- 1,6		f_{rel}	1,0	+ 0,0	- 0,0	1,0	+ 0,0	- 0,0
	l	-0,6	+ 0,6	- 0,5	-1,0	+ 1,0	- 0,9		l	0,0	+ 0,0	- 0,0	0,0	+ 0,0	- 0,0
	m	0,9	+ 0,1	- 0,1	0,9	+ 0,2	- 0,1	m	0,0	+ 0,0	- 0,0	0,0	+ 0,0	- 0,0	
2	ϑ	8,1	+ 2,0	- 2,3	7,9	+ 1,6	- 2,5	2	ϑ	15,8	+ 1,2	- 1,3	22,8	+ 3,0	- 3,4
	θ	2,7	+ 0,8	- 0,9	3,0	+ 1,1	- 1,5		θ	5,9	+ 0,7	- 0,9	8,0	+ 1,4	- 1,5
	PA	-51,8	+ 151,5	- 7,8	-56,9	+ 154,7	- 13,2		PA	-43,8	+ 3,0	- 2,7	-42,1	+ 7,5	- 5,5
	f_{rel}	4,2	+ 1,5	- 1,1	4,6	+ 1,2	- 1,2		f_{rel}	1,2	+ 0,4	- 0,3	0,4	+ 0,7	- 0,1
	l	-1,5	+ 1,1	- 1,0	-1,3	+ 0,9	- 0,7		l	-3,2	+ 2,6	- 2,4	-2,0	+ 2,6	- 2,5
	m	14,9	+ 0,6	- 0,7	14,7	+ 0,6	- 0,6	m	15,0	+ 1,9	- 1,8	17,1	+ 2,0	- 1,8	
3	ϑ	13,6	+ 1,3	- 1,1	13,9	+ 1,6	- 1,4	3	ϑ	26,7	+ 2,8	- 3,1	34,5	+ 16,0	- 3,4
	θ	5,8	+ 1,1	- 1,1	6,3	+ 1,9	- 2,0		θ	19,7	+ 2,6	- 3,2	32,6	+ 3,1	- 4,1
	PA	-7,2	+ 1,2	- 1,0	-7,0	+ 1,1	- 1,3		PA	50,6	+ 14,0	- 23,3	6,2	+ 47,7	- 59,6
	f_{rel}	1,6	+ 1,5	- 1,1	1,0	+ 1,1	- 0,8		f_{rel}	1,4	+ 0,7	- 0,5	0,7	+ 1,0	- 0,2
	l	-9,0	+ 0,8	- 1,1	-8,9	+ 1,8	- 1,3		l	3,2	+ 2,7	- 2,7	7,3	+ 3,6	- 6,2
	m	15,4	+ 0,8	- 1,0	15,3	+ 1,4	- 1,0	m	21,6	+ 2,2	- 2,4	26,5	+ 2,8	- 13,5	
4	ϑ	8,9	+ 1,0	- 1,0	9,2	+ 1,2	- 1,3	4	ϑ	2,6	+ 0,7	- 0,8	2,2	+ 0,7	- 0,6
	θ	2,6	+ 0,7	- 0,8	2,2	+ 0,7	- 0,6		θ	-32,6	+ 1,1	- 1,1	-33,1	+ 1,5	- 1,4
	PA	-32,6	+ 1,1	- 1,1	-33,1	+ 1,5	- 1,4		PA	3,7	+ 1,5	- 1,3	4,7	+ 1,2	- 1,1
	f_{rel}	3,7	+ 1,5	- 1,3	4,7	+ 1,2	- 1,1		f_{rel}	3,1	+ 0,7	- 0,8	3,5	+ 0,7	- 0,6
	l	3,1	+ 0,7	- 0,8	3,5	+ 0,7	- 0,6		l	16,6	+ 1,0	- 0,8	16,4	+ 1,0	- 0,6
	m	16,6	+ 1,0	- 0,8	16,4	+ 1,0	- 0,6	m	9,7	+ 1,2	- 1,1	10,2	+ 1,4	- 1,2	
5	ϑ	9,7	+ 1,2	- 1,1	10,2	+ 1,4	- 1,2	5	ϑ	5,6	+ 0,9	- 1,1	5,6	+ 1,5	- 1,4
	θ	5,6	+ 0,9	- 1,1	5,6	+ 1,5	- 1,4		θ	-29,9	+ 1,2	- 1,1	-29,9	+ 1,3	- 1,3
	PA	-29,9	+ 1,2	- 1,1	-29,9	+ 1,3	- 1,3		PA	2,3	+ 1,5	- 1,0	1,6	+ 0,8	- 0,6
	f_{rel}	2,3	+ 1,5	- 1,0	1,6	+ 0,8	- 0,6		f_{rel}	-4,0	+ 0,9	- 0,8	-3,6	+ 0,8	- 0,7
	l	-4,0	+ 0,9	- 0,8	-3,6	+ 0,8	- 0,7		l	20,0	+ 0,8	- 0,9	20,4	+ 1,5	- 1,7
	m	20,0	+ 0,8	- 0,9	20,4	+ 1,5	- 1,7	m	9,2	+ 0,9	- 1,0	9,2	+ 1,0	- 0,9	
6	ϑ	9,2	+ 0,9	- 1,0	9,2	+ 1,0	- 0,9	6	ϑ	3,4	+ 1,0	- 1,1	3,1	+ 1,0	- 1,1
	θ	3,4	+ 1,0	- 1,1	3,1	+ 1,0	- 1,1		θ	-20,7	+ 1,1	- 1,4	-20,3	+ 1,6	- 1,7
	PA	-20,7	+ 1,1	- 1,4	-20,3	+ 1,6	- 1,7		PA	1,0	+ 1,3	- 0,7	0,7	+ 0,7	- 0,4
	f_{rel}	1,0	+ 1,3	- 0,7	0,7	+ 0,7	- 0,4		f_{rel}	7,3	+ 1,1	- 1,0	7,8	+ 1,2	- 1,3
	l	7,3	+ 1,1	- 1,0	7,8	+ 1,2	- 1,3		l	25,0	+ 0,9	- 1,0	24,9	+ 1,6	- 1,3
	m	25,0	+ 0,9	- 1,0	24,9	+ 1,6	- 1,3	m	22,5	+ 1,0	- 1,0	22,5	+ 1,6	- 1,2	
7	ϑ	22,5	+ 1,0	- 1,0	22,5	+ 1,6	- 1,2	7	ϑ	3,2	+ 1,0	- 0,9	2,8	+ 1,5	- 1,3
	θ	3,2	+ 1,0	- 0,9	2,8	+ 1,5	- 1,3		θ	-35,4	+ 1,3	- 1,5	-35,5	+ 1,7	- 2,1
	PA	-35,4	+ 1,3	- 1,5	-35,5	+ 1,7	- 2,1		PA	2,5	+ 2,0	- 1,6	0,6	+ 1,9	- 0,5
	f_{rel}	2,5	+ 2,0	- 1,6	0,6	+ 1,9	- 0,5		f_{rel}	-2,0	+ 0,9	- 1,0	-1,8	+ 1,3	- 2,0
	l	-2,0	+ 0,9	- 1,0	-1,8	+ 1,3	- 2,0		l	-2,7	+ 1,3	- 1,3	-3,2	+ 2,8	- 1,8
	m	-2,7	+ 1,3	- 1,3	-3,2	+ 2,8	- 1,8	m							

d) Aperture		Carbon content	T_{cold} [K]		η_{cold}		T_{hot} [K]		η_{hot}		N_{ext} [$\mu\text{g}/\text{cm}^2$]		χ^2			
		%	min	.. max	min	.. max	min	.. max	min	.. max	min	.. max	min	.. max		
E1	0	561	..	581	0,04	..	0,06	N.A.	N.A.	396	..	547	10	..	13	
	20	666	..	701	0,04	..	0,04	N.A.	N.A.	448	..	610	15	..	18	
E2	0	188	..	469	0,04	..	1,00	703	..	1179	0,00050	..	0,00050	281	..	530
	20	206	..	506	0,04	..	0,62	829	..	1461	0,00100	..	0,00660	386	..	728
E3	0	217	..	376	0,04	..	0,50	860	..	1095	0,00050	..	0,00130	67	..	386
	20	225	..	371	0,05	..	0,34	917	..	1412	0,00050	..	0,00130	92	..	489
E4	0	232	..	389	0,06	..	0,43	555	..	818	0,00060	..	0,00950	281	..	530
	20	250	..	412	0,09	..	0,38	633	..	1200	0,00040	..	0,00940	386	..	621
E5	0	223	..	285	0,10	..	0,36	717	..	1098	0,00005	..	0,00040	1	..	100
	20	220	..	287	0,09	..	0,38	712	..	1275	0,00004	..	0,00040	1	..	189
DE1	0	229	..	238	0,10	..	0,13	771	..	865	0,00001	..	0,00002	1	..	34
	20	228	..	242	0,09	..	0,13	820	..	870	0,00001	..	0,00001	1	..	43
DE2	0	139	..	228	0,09	..	1,00	417	..	475	0,00299	..	0,00918	825	..	1000
	20	153	..	292	0,06	..	1,00	543	..	703	0,00175	..	0,00775	1000	..	1468
DE3	0	148	..	162	0,40	..	1,00	N.A.	N.A.	1	..	197	13	..	17	
	20	148	..	162	0,40	..	1,00	N.A.	N.A.	1	..	241	13	..	17	
SE	0	203	..	354	0,03	..	0,51	846	..	1031	0,00060	..	0,00140	530	..	853
	20	173	..	300	0,04	..	1,00	1333	..	1500	0,00070	..	0,00090	464	..	619

Table 4.2: (a) Parameters used in the IRBis image reconstructions. (1) Mean closure phase error in the data (deg); (2) Mean signal-to-noise ratio of the squared visibility in the data; (3) Field of view of the reconstruction in mas; (4) Pixel grid dimensions ($x = y$) used in the reconstruction; (5) Cost function number as defined in the text; (6) Regularization functions numbers: 2 and 4 correspond, respectively, to maximum entropy and total variation, and signs indicate used prior: + means a fitted Gaussian and - means a constant; (7) Optimization engine number as defined in the text; (8) Power of the uv -density weight. (b) L - and M -band best-fit parameters. Θ , FWHM of the major axis in milliarcseconds (mas); θ , FWHM of the minor axis (mas); PA, position angle of the major axis (degrees), f_{rel} , relative integrated flux of the Gaussian component (%); l , position west of the centre for the component (mas); m , position north of the centre (mas). (c) N -band best-fit parameters. Symbols as in b. (d) Extended list of SED parameters, for each area defined in Fig. 4.2, assuming either 0% or 20% carbon content (by weight) in the absorbing screen. T_{hot} and T_{cold} are temperatures and η_{hot} and η_{cold} are the filling factors of the two blackbodies, except for E1 and DE3, which are adequately described by a single cold blackbody. N_{ext} is the extinction and χ^2 approximately describes the fit quality.

Chapter 5

The dusty heart of Circinus

Though NGC 1068 is the brightest nearby AGN, the Circinus Galaxy is much closer. Its proximity allows us to study the circumnuclear dust down to tenths of a parsec scale using MATISSE. We can compare the circumnuclear dust structures found in this Seyfert 2 AGN to those found in NGC 1068 to start building a picture of the key components of Seyfert 2s.

This chapter¹ is based largely on the *N*-band study of Circinus published in [Isbell et al. \(2022\)](#), but has been extended to include new analysis of the *L*- and *M*-bands. I was the lead author of the published manuscript. Joel Sanchez-Bermudez and Karl-Heinz Hofmann provided invaluable advice with regards to image reconstruction; Joel also reduced the VISIR-SAM data and wrote parts of §3.1.3. Marko Stalevski provided the disk+wind radiative transfer models and aided in their interpretation. All text and figures, unless otherwise noted, were produced by me, taking into account the feedback of the co-authors. All co-authors (K. Meisenheimer, J.-U. Pott, M. Stalevski, K. R. W. Tristram, J. Sanchez-Bermudez, K.-H. Hofmann, V. Gámez Rosas, W. Jaffe, L. Burtscher, J. Leftley, R. Petrov, B. Lopez, T. Henning, G. Weigelt, F. Allouche, P. Berio, F. Bettonvil, P. Cruzalebes, C. Dominik, M. Heininger, M. Hogerheijde, S. Lagarde, M. Lehmitz, A. Matter, F. Millour, S. Robbe-Dubois, D. Schertl, R. van Boekel, J. Varga, and J. Woillez) provided helpful feedback on manuscript structure and data interpretation. The new, *LM*-band section was written by me for this thesis.

Brief Summary

We present the first sub-parsec scale images of the circumnuclear dust in the nearby Seyfert 2, the Circinus Galaxy. Previous observations have revealed complex structures and polar dust emission but interpretation was limited to simple models. The new Multi AperTure mid-Infrared Spectro-Scopic Experiment (MATISSE) makes it possible to image these structures for the first time. We observed the Circinus Galaxy with MATISSE at the Very Large Telescope Interferometer (VLTI), producing 150 correlated flux spectra and 100 closure phase spectra in the *LMN*-bands. The novel inclusion of closure

¹This work makes use of ESO Programs 099.B-0484(A), 0104.B-0064(A), 0104.B-0127(A), 106.214U.002, and 105.205M.001. The images in Fig. 5.4 are available in electronic form at the CDS via anonymous ftp to [cdsarc.u-strasbg.fr](ftp://cdsarc.u-strasbg.fr) (130.79.128.5) or via <http://cdsweb.u-strasbg.fr/cgi-bin/qcat?J/A+A/>.

phases makes interferometric imaging possible for the first time. We reconstructed images in the three bands at $\sim 3 - 10$ mas resolution. We fit blackbody functions with dust extinction to several aperture-extracted fluxes from the images to produce a temperature distribution of central dusty structures. We find significant substructure in the circumnuclear dust: central unresolved flux of ~ 0.5 Jy, a thin disk 1.9 pc in diameter oriented along $\sim 45^\circ$, and a $\sim 4 \times 1.5$ pc polar emission extending orthogonal to the disk. The polar emission exhibits patchiness, which we attribute to clumpy dust. Flux enhancements to the east and west of the disk are seen for the first time. We distinguish the temperature profiles of the disk and of the polar emission: the disk shows a steep temperature gradient indicative of denser material; the polar profile is flatter, indicating clumpiness and/or lower dust density. The unresolved flux is fitted with a high temperature, ~ 370 K. The polar dust remains warm (~ 200 K) out to 1.5 pc from the disk. We attribute approximately 60% of the $12 \mu\text{m}$ flux to the polar dust, 10% to the disk, and 6% is unresolved; the remaining flux was resolved out. The recovered morphology and temperature distribution resembles modeling of accretion disks with radiation-driven winds at large scales, but we placed new constraints on the subparsec dust. The spatially resolved subparsec features imaged here place new constraints on the physical modeling of circumnuclear dust in active galaxies; we show strong evidence that the polar emission consists of dust clumps or filaments. The dynamics of the structures and their role in the Unified Model remain to be explored.

1 Introduction

Active galactic nuclei (AGN) are thought to play a crucial role in the formation and evolution of their host galaxies. Moreover, understanding the dust in the vicinity of supermassive black holes is key to understanding how AGN are fed and how they interact with their hosts. The dust traces dense molecular gas which feeds the AGN. Large, obscuring dusty structures are thought to be responsible for both funneling material toward the central engine, and for distinguishing between Seyfert 1 and Seyfert 2 AGN. In the original Unified Model of AGN (Antonucci, 1993; Urry & Padovani, 1995; Netzer, 2015), a central obscuring torus of dust is oriented such that the broad-line region of the AGN is directly visible (Seyfert 1) or such that its observation is blocked by the torus (Seyfert 2; hereafter Sy2). So in order to fully understand the accretion process and the life cycle of an AGN, one must understand the parsec-scale dust structures surrounding it.

The so-called torus is comprised of several key features which vary in temperature from < 200 K to 1500 K and scale from tenths of a parsec to hundreds of parsecs. The inner edge is the radius at which radiation from the accretion disk (AD) causes the dust to sublimate. The sublimation radius is dependent on both the luminosity of the AD and the dust composition, but typically ~ 0.1 pc for dust evaporating at 1500 K, for a $L \sim 1 \times 10^{10} L_\odot$ AGN. Beyond the sublimation zone, it is thought that a dense disk or torus of material is responsible for both hiding the broad line region (BLR) in Sy2 AGN and for feeding the AD. Previous mid-infrared (MIR) interferometric studies revealed that many “tori” have an additional component in the form of a polar extension (see, e.g., Hönicg et al., 2012; Burtscher et al., 2013; López-Gonzaga et al., 2016; Leftley et al., 2018), the Circinus Galaxy’s chief among them (Tristram et al., 2007, 2014). The polar component is thought to be a radiation-driven outflow (e.g., Wada, 2012; Wada et al., 2016), and it can represent a key mechanism of AGN feedback. This is called the fountain model, and it was shown by Schartmann et al. (2014) to reproduce the MIR polar extension and dusty hollow cone in the Circinus Galaxy (hereafter Circinus). A key finding of spectral energy distribution (SED) fits to nearby AGN as well as comparisons

to radiative transfer models is that the dust in the central structures (and particularly in the wind) must be clumpy, allowing dust to reach high temperatures and exhibit “blue” spectra even at large distances from the AD (Krolik & Begelman, 1988; Nenkova et al., 2008a; Hönig & Kishimoto, 2017; Martínez-Paredes et al., 2020; Isbell et al., 2021). The exact nature of these components and how they are connected to each other and to the host galaxy remains an open question. A holistic model of the central dust distribution is shown in Izumi et al. (2018), but only the resolution offered by infrared interferometry can probe the subparsec details of the dust near the active nucleus.

The Multi AperTure mid-Infrared Spectro-Scopic Experiment (MATISSE) is the second-generation MIR interferometer on the Very Large Telescope Interferometer (VLTI) at the European Southern Observatory (ESO) Paranal site (Lopez et al., 2014; Lopez et al., 2022). MATISSE combines the light from four unit telescopes (UTs) or four auxiliary telescopes (ATs) measuring six baselines in the L -, M -, and N -bands simultaneously. MATISSE furthermore introduces closure phases to MIR interferometry. The combination of the phase measurements on any three baselines $\phi_{ijk} \equiv \phi_{ij} + \phi_{jk} - \phi_{ik}$ is called the closure phase; this summation cancels out any atmospheric or baseline-dependent phase errors (Jennison, 1958; Monnier, 2003). Closure phases are crucial for imaging because they probe the spatial distribution of target flux and because they are unaffected by atmospheric turbulence. Recent imaging studies of NGC 1068 with GRAVITY (Gravity Collaboration et al., 2020) and MATISSE (Gómez Rosas et al., 2022) have illustrated the power of this approach in revealing new morphological details and spatially resolved temperature measurements of the circumnuclear dust. Until this work, NGC 1068 was the only AGN to have been imaged with MATISSE.

Circinus is of particular interest as it is one of the closest Sy2 galaxies (at a distance of 4.2 Mpc Freeman et al., 1977; Tully et al., 2009) and the second brightest in the MIR (only fainter than NGC 1068). Circinus is a prototypical Sy2 galaxy, exhibiting narrow emission lines (Oliva et al., 1994; Moorwood et al., 1996) and an obscured broad-line region (BLR; Oliva et al., 1998), as well as bipolar radio lobes (Elmouttie et al., 1998) and an optical ionization cone (Marconi et al., 1994; Maiolino et al., 2000; Wilson et al., 2000; Mingozi et al., 2019). Additionally, Circinus exhibits a Compton-thick nucleus and a reflection component in X-rays (Matt et al., 1996; Smith & Wilson, 2001; Soldi et al., 2005; Yang et al., 2009). Finally, in- and outflows and spiral arms have been observed in CO down to ~ 5 pc scales (Curran et al., 1998; Izumi et al., 2018; Tristram et al., 2022), further indicating the complexity of the central structures.

Circinus was observed extensively with the first generation MIR interferometer, the MID-infrared Interferometric instrument (MIDI; Leinert et al., 2003), in the N -band (e.g., Tristram et al., 2007, 2014, hereafter T14). These observations showed a warm (~ 300 K) dust disk roughly aligned with the water maser emission (Greenhill et al., 2003), but the flux was dominated by large scale ($\gtrsim 100$ mas) emission roughly orthogonal to the disk. The orientation of the large scale emission’s major axis was found to differ significantly from the optical ionization cone central angle ($\text{PA}_{\text{opt.}} = -45^\circ$ vs $\text{PA}_{\text{dust}} = -73^\circ$), and follow-up modeling work by Stalevski et al. (2017, 2019) has indicated that the polar-extended dust emission may come from an edge-brightened outflow cone.

The proximity and declination of Circinus (at around -60°) make it an ideal target for imaging with MATISSE, as it provides high spatial resolution (10 mas = 0.2 pc) and because its nearly circular uv -tracks aid in the production of high fidelity reconstructions. MATISSE provides the first MIR measurements of the closure phase, which sample the (a)symmetry of a source and are crucial for image reconstruction. Previous analysis relied on Gaussian model fitting, which is a smooth, simplified representation of the source emission; but interferometric image reconstruction has the potential to build on

Table 5.1: VLTI/MATISSE observations entering this analysis

TPL Start	Target	N_{cycles}	τ_0 [ms]	Seeing ["]
2020-03-13T04:02:11	Circinus	2	6.4	0.73
2020-03-13T04:56:22	Circinus	1	7.1	0.58
2020-03-14T03:53:00	Circinus	2	7.1	0.65
2020-03-14T04:31:58	Circinus	2	4.9	0.91
2020-03-14T04:51:12	Circinus	4	7.3	0.63
2020-03-14T07:57:12	Circinus	2	6.6	0.54
2020-03-14T08:57:48	Circinus	4	8.0	0.47
2021-02-28T06:32:19	Circinus	2	10.8	0.79
2021-02-28T07:42:00	Circinus	2	8.8	0.81
2021-06-01T03:10:17	Circinus	2	4.7	0.70
2021-06-01T04:29:41	Circinus	2	5.8	0.54
Calibrators				
2020-03-13T04:40:24	HD120404	1	6.0	0.56
2020-03-14T05:59:29	HD120404	1	7.9	0.48
2020-03-14T08:31:10	HD120404	1	7.4	0.55
2021-02-28T06:18:58	HD120913	1	9.0	0.79
2021-02-28T07:07:46	HD120404	1	5.8	1.06
2021-06-01T02:38:46	HD119164	1	5.2	0.78
2021-06-01T03:59:25	HD120404	1	6.2	0.47

Seeing and coherence time (τ_0) values are given from the start of each observing block; N_{cycles} is the number of observed interferometric cycles, consisting each of four 1 min long exposures with changing configurations of the BCD.

these results through model-independent sampling of the source structure. Herein we present the first image reconstructions of the N -band circumnuclear dust in Circinus.

This section is organized as follows. In §2 we present the observations entering this work. In §3 the data reduction and verification, imaging, and temperature fitting of the N -band data is presented and discussed. In §4 this discussion and analysis is extended to the LM -bands. In §5 we analyze the various components of the central dust structure in Circinus and discuss their implications. Finally, we conclude and summarize in §6.

2 MATISSE Observations

The MATISSE observations of Circinus were carried out on 13–14 March 2020, 27 Feb 2021, and 31 May 2021 as part of guaranteed time observations. Data were taken with low spectral resolution in both the LM - (3–5 μm) and N -bands (8–13 μm). The observations were taken using the unit telescope (UT) configuration, with physical baselines ranging from 30 m to 140 m. At 12 μm this corresponds to angular resolutions between 9 and 41 mas with a “primary beam” of 153 mas. Each observation sequence consists of two sky exposures, a number of exposure cycles, N_{cycles} , consisting each of four 1 min interferometric exposures with different configurations of the beam commuting device (BCD) of MATISSE, as well as optional photometric exposures while chopping (for details see [Lopez et al., 2022](#)). Near the end of the night of 14 March 2020, we opted to repeat more exposure cycles to reduce the overhead time of re-acquisition on the target. The exact number of exposure cycles, along with the atmospheric conditions at the start of each observation, are given in Table 1. The observing conditions on 14 March 2020 were excellent, while on 13 March 2020 high-altitude cirrus negatively impacted acquisition, guiding, and adaptive optics in several individual exposures; we note that the final correlated flux error estimates on this night are higher. Observations on 28 Feb and 01 Jun 2021 were unaffected by such issues. We show the combined uv -coverage of




Fig. 5.1: MATISSE uv -coverage from all 25 exposure cycles. Squares denote observations taken in 2020, while circles represent observations from 2021. The mean correlated flux between 11.5 and 12.5 μm is used as the color scale. The discrete color binning is done in 0.2 Jy intervals, based on the measured correlated flux uncertainties.

all the observations in Fig. 5.1.

On each night, we observed the calibration star HD120404 ($F_{12\mu\text{m}} = 13$ Jy) directly before and/or after the Circinus observations. The atmospheric conditions at the start of each calibrator observation are given in Table 1. This star serves as a spectral calibrator, an instrumental phase calibrator, and an instrumental visibility calibrator. It has a MIR spectrum given by van Boekel (2004), and its diameter is given as 2.958 mas in Cruzalèbes et al. (2019). During the Feb. and May 2021 observations, we observed secondary calibrators, HD120913 ($F_{12\mu\text{m}} = 5.7$ Jy) and HD119164 ($F_{12\mu\text{m}} = 1.2$ Jy) in order to perform cross-calibration and closure phases accuracy checks.

3 *N*-band Analysis

This section is organized as follows. In §3.1 the reduction and calibration of the MATISSE observations is presented. In §3.2 we lay out the interferometric image reconstruction process and final image reconstruction parameters. We also compute image errors and assess the morphology of the resulting structure. In §3.3 we measure the temperature distribution of the dust in the central structure via blackbody fitting.

3.1 Data reduction and calibration

3.1.1 MATISSE data reduction

The *N*-band data are reduced using the MATISSE data reduction software² (DRS) version 1.5.1. We used the coherent reduction flags `corrFlux=TRUE` and `coherentAlgo=2` in order to produce correlated fluxes using the coherent integration algorithm as employed in the MIDI Expert Work Station (EWS; Jaffe, 2004) and used in T14. We also use spectral binning 21 px ($= 1 \mu\text{m}$) and the default values for all other parameters.

The correlated flux, $F(u, v, \lambda)$ was then calibrated in the standard way:

$$F_{\text{targ}}^{\text{cal}}(u, v, \lambda) = F_{\text{targ}}^{\text{raw}}(u, v, \lambda) / F_{\star}^{\text{raw}}(u, v, \lambda) \times F_{\star}^{\text{tot}}(u = 0, v = 0, \lambda), \quad (5.1)$$

²<https://www.eso.org/sci/software/pipelines/matisse/>

where F_{\star}^{raw} is the raw flux (in counts) of the calibrator, F_{\star}^{tot} is the catalog flux of the calibrator, and $F_{\text{targ}}^{\text{raw}}$ is the raw flux of the target. This assumes the calibrator is unresolved; for the selected calibrators with diameter < 3 mas this is the case.

Within an observing cycle, individual exposures are taken minutes apart. The standard deviation of these correlated flux measurements is used as an uncertainty estimate, typically 0.2 Jy at all baselines. The uncertainties measured in this way broadly agree with the DRS-estimated values. The squared visibilities are finally calculated as $V^2(u, v, \lambda) = [F_{\text{targ}}^{\text{cal}}(u, v, \lambda)/F_{\text{targ}}^{\text{cal}}(u = 0, v = 0, \lambda)]^2$, where the “zero-baseline” flux is the arithmetic mean of the photometric flux spectra measured by each of the 8.1m UTs.

The photometric flux was initially reduced via incoherent processing in the DRS (using `corrFlux=FALSE`). This mode extracts the photometric flux passing through the $2\lambda/D$ pinhole in each UT (0.61” at $12\ \mu\text{m}$). This is not computed for each observing block, as the N -band photometry cycle adds 10 minutes to each observation, but once per epoch we recorded the photometry. The N -band photometry we obtain from the DRS is a factor ~ 3 larger than expected from the MIDI and VISIR observations in T14, 36 ± 4 Jy vs 12 ± 1 Jy at $12\ \mu\text{m}$. We doubt temporal flux variations in the source, as none of the correlated fluxes at any spatial scale exhibit a similar change since 2008 (see §3.1.4). When using EWS (Jaffe, 2004), which was used previously for the MIDI observations, we extract a photometric flux of 12.4 ± 0.5 Jy at $12\ \mu\text{m}$ for the same set of observations. This indicates that the photometric flux only exhibited a change due to the spatial filter used in each software; EWS employs a narrow Gaussian filter while DRS employs a wider top-hat. To compare consistently to the MIDI data, the EWS total flux value is used throughout this work.

We assume the calibration stars are symmetric and have zero closure phase on all phase triangles – any deviations from zero represent instrumental phase errors. As a first step in closure phase calibration, deviations from zero phase in the calibrator $\delta\phi_{\star,ijk}(\lambda)$ are subtracted from the target phase: $\phi_{ijk,\text{targ}}^{\text{cal}} = \phi_{ijk,\text{targ}}^{\text{raw}} - \delta\phi_{\star,ijk}$. A typical MATISSE observation cycle includes 4 configurations of the BCD which serve to calibrate the closure phase. The varied BCD configurations (called out-out, in-in, in-out, out-in) should be identical save for sign flips on individual closure loops (as $\phi_{ijk} = -\phi_{ikj}$). We then average the star-calibrated closure phases. We first calculate the temporal mean value for each individual BCD configuration, as they are each repeated a number N_{cycles} times. Finally the mean of the four BCD configurations serves as the closure phase value at each wavelength, and the standard deviation is used as an estimate of our closure phase uncertainty (on the order of 15° for Circinus, on the order of 1° for HD120404). We note, however, that all closure loops which include the ~ 130 m baseline, UT1-UT4, have systematically higher uncertainties due to the low signal-to-noise ratio (S/N) correlated flux on this baseline. The uncertainty is on the order of 50° for Circinus, which means that only the closure triangles UT1-UT2-UT3 and UT2-UT3-UT4 provide high-precision phase information.

We have measurements in a total of 25 MATISSE exposure cycles, corresponding to 150 correlated flux measurements and 100 closure phase measurements. A subsample of the calibrated correlated fluxes is shown in Fig. 5.2, the remaining correlated fluxes are shown in Appendix Fig. C.1, and all closure phases are given in Appendix Fig. D.1. We define a position angle in the uv -plane as $\tan\psi = v/u$; we have sampled essentially all ψ between 0 and 110° , although the sampling is not uniform. This becomes especially noticeable on the longer baselines (> 100 m). Two long-baseline regions at $\psi \approx [10, 40]^\circ$ and $\psi \approx [80, 110]^\circ$ are highly sampled, while a more sparse region is present between $\psi = 45^\circ$ and 60° . On the shorter baselines, no such gaps are present.

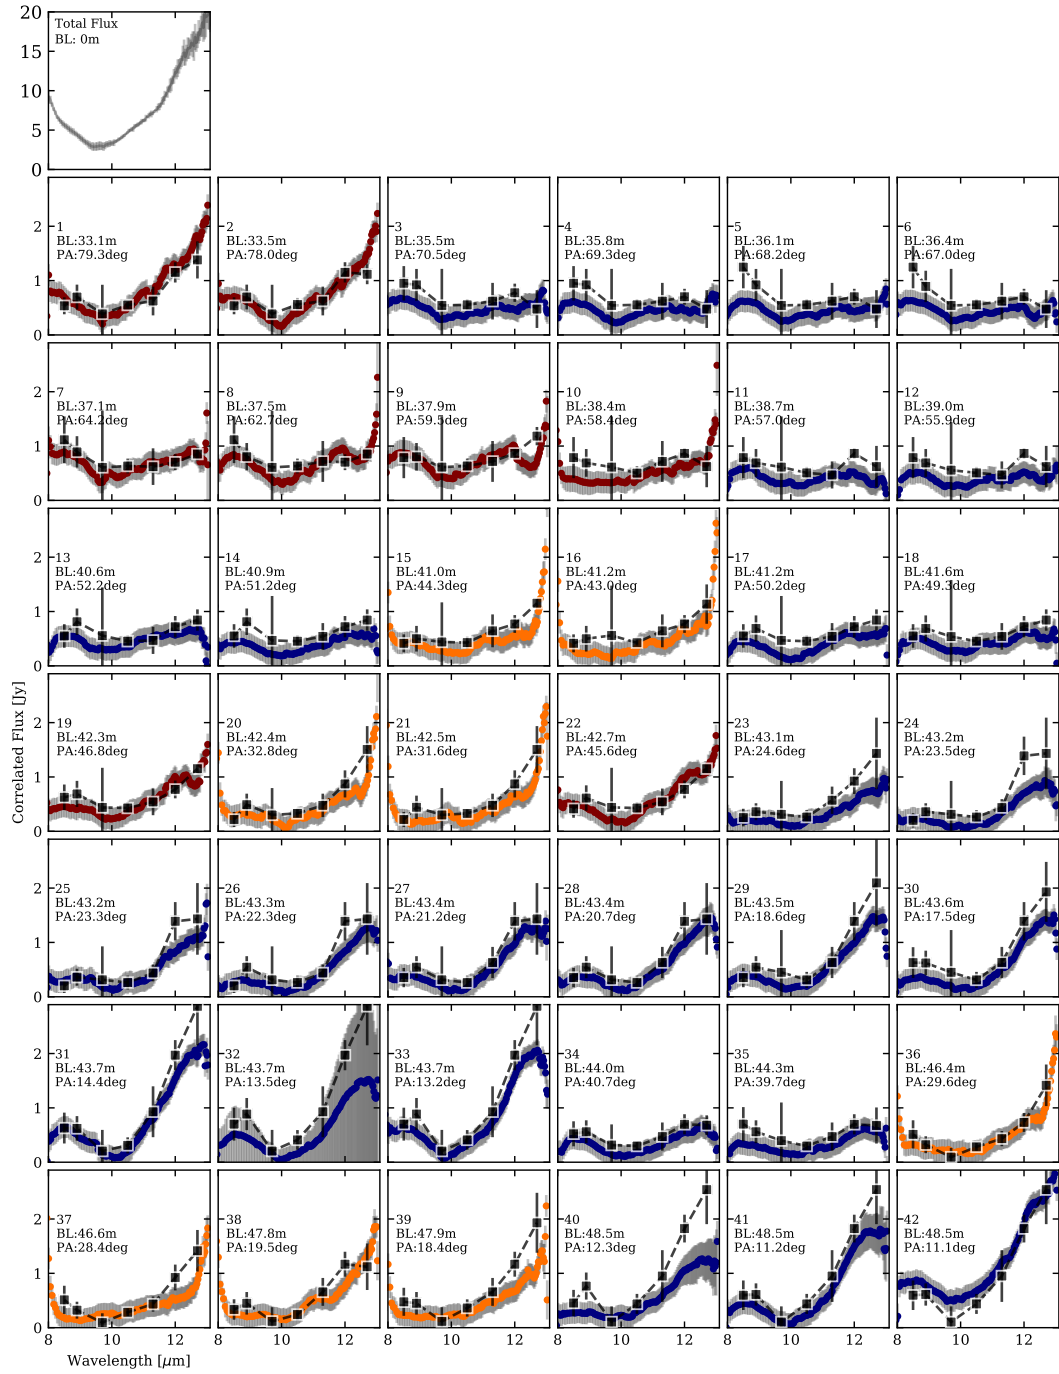


Fig. 5.2: Subsample of *N*-band correlated flux data for Circinus from March 2020 (blue), February 2021 (yellow), and May 2021 (red). The black points are simulated values extracted from the final images, with errors estimated using the 1σ error maps (described in §3.2.2). The total photometric flux is included in the first panel. Presented errors come from both the calibrator flux uncertainty and the statistical variation of the observables within a set of observing cycles. Near 8 and 13 μm one can see flux variations due to the edges of the atmospheric window. The remaining correlated fluxes are shown in Fig. C.1.

3.1.2 MIDI observations

We include short-baseline MIDI observations from T14. These short baselines provide the small spatial frequencies necessary for imaging or modeling of the large-scale structure in Circinus. These data were reduced using the MIDI EWS (Jaffe, 2004). The exact procedure is given in T14. These data contain the correlated flux, the visibility amplitude, and the wavelength differential phase. We calculate the squared visibility as $V^2 = (F_{\text{corr}}/F_{\text{tot}})^2$. Both the MATISSE and MIDI data have been calibrated with the same calibration star, HD120404. The MIDI data do not provide closure phases, so we select only a small number of AT baselines rather than fully incorporating the MIDI uv -coverage. We selected the baselines to have (i) a projected baseline $< 35\text{m}$; and (ii) u, v spacing of at least 8.1 m (the UT-diameter). This leaves us with 18 baselines from the small configuration. In the MATISSE OIFITS format, these 18 baselines correspond to 12 closure phase loops, which we give as $0 \pm 180^\circ$ such that these nonexistent closure phases have no weight on imaging. This assumption is supported by the closure phase measurements of the VLT spectrometer and imager for the mid-infrared (VISIR) sparse-aperture-masking data.

3.1.3 VISIR sparse-aperture-masking data

Circinus was observed with the sparse-aperture-masking (SAM) mode of VISIR. The observations were taken in the N -band ($\lambda_0 = 11.3 \mu\text{m}$; Filter Name = 11_3_SAM) on 02 June 2017 (099.B-0484A). The data consisted of five observing blocks on the science with interwoven observations with the calibrator star HD 125687. Each data set in the sequence SCI-CAL was observed with a DIT=142 milliseconds and NDIT=6 exposures. The data reduction consisted of two parts. The first one uses the ESOREX data reduction pipeline offered by ESO³. It allowed us to correct for (i) the background, (ii) the bad pixels, (iii) to extract the interferograms from the chopping sequence and (iv) to center each frame on a 256×256 pixel grid. Frames with low signal-to-noise or with bad cosmetics were discarded manually from the data. Once the interferograms were cleaned, we extracted the interferometric observables from them.

To obtain the squared visibilities and closure phases from the data, we used the CASSINI-SAMPip⁴ software (see e.g., Sanchez-Bermudez et al., 2020). This algorithm fits the interferogram directly on the image plane, methods with similar performance based on fringe fitting are described by Greenbaum et al. (2015) and Lacour et al. (2011). The code uses a Single Value Decomposition method to obtain the interferometric observables. The algorithm works with monochromatic data and uses a sinc-filter for compensating the wavelength smearing of the broad-band VISIR filter. Each frame in the data was fitted independently. The uncertainties in the observables were obtained by averaging the observables of the six frames in each data set of science and calibrator, respectively. With the seven pin-holes mask available on VISIR, 21 squared visibilities and 35 closure phases were obtained per data set. The minimum baseline produced with the VISIR non-redundant mask has a length of 1.67 meters ($\lambda_0/2B_{\text{min}} = 600 \text{ mas}$) and the maximum one a length of 6.28 meters ($\lambda_0/2B_{\text{max}} = 184 \text{ mas}$), respectively. Figure E.1 shows, as example, one snapshot of the recorded interferogram of the science target and the uv -coverage obtained with our observations. Once the raw observables were extracted, the data were calibrated by dividing the squared visibilities of the target over the ones of the calibrator star; the closure phases were calibrated by subtracting the closure phases of the calibrator from the ones of the target. Figure E.2 shows the calibrated observables versus spatial frequency.

³<https://www.eso.org/sci/software/pipelines/visir/visir-pipe-recipes.html>

⁴<https://github.com/cosmosz5/CASSINI>




Fig. 5.3: Comparison of MIDI and MATISSE correlated flux values on baselines cross-matched within 4 m. Color scale is difference in $\sigma_{F_{\text{corr}}}$ from the MATISSE observations. Only two uv -points are $1\sigma_{F_{\text{corr}}} < \Delta_{F_{\text{corr}}} < 2\sigma_{F_{\text{corr}}}$ discrepant between the MIDI and MATISSE observations spaced more than 10 years apart.

3.1.4 Correlated flux stability

Combining the MIDI and MATISSE datasets taken ≥ 10 years apart depends on the assumption that both the structure and photometry of Circinus are stable in the same period. T14 reported possible flux variation of Circinus between 2008 and 2009. Moreover, there may be instrumental biases which are not properly calibrated. Therefore, we compare the correlated flux values taken using MATISSE in 2020 and 2021 with those at similar u, v coordinates reported in T14. We identify and compare 30 baselines from MIDI and MATISSE which are within 4m in u, v distance of each other; these are shown in Fig 5.3. We find excellent agreement between the two epochs, with $> 90\%$ of baselines consistent within the $1\sigma_{F_{\text{corr}}} \approx 0.2$ Jy calibrated correlated flux errors. Only two baselines are discrepant at $12 \mu\text{m}$ by $> 1\sigma_{F_{\text{corr}}}$, but agree within $2\sigma_{F_{\text{corr}}}$. We find that there are no significant changes in correlated or total flux over the last ≥ 10 years.

3.1.5 Combination of MIDI and MATISSE data

Information at a large range of spatial frequencies is necessary for robust imaging of a source. The MATISSE UT observations have a shortest baseline of ~ 30 m, which causes structures larger than 82.5 mas to be resolved out at $12 \mu\text{m}$. Without MATISSE AT observations, we lack constraints on the large-scale structure. We know, however, that there is large-scale structure out to ≥ 600 mas from MIDI, VISIR-SAM, and VISIR data (T14, this work, and [Asmus et al., 2014](#), respectively). In order to *a*) avoid resolving out structure which is shown to be present, and *b*) constrain the locations of small-scale structures, we perform the image reconstruction using a combination of MIDI and MATISSE data. The practice of including small spatial frequency data via modeling or data supplementation is common in imaging (e.g., [Cotton, 2017](#)). The MIDI data do not include closure phase measurements, so as stated above we set the values to $0 \pm 180^\circ$ during imaging.

We claim that such a data supplementation is valid in the case of Circinus for the following reasons. First, both the MIDI modeling and VISIR-SAM data show that closure phases are small on these scales ($\phi \leq 10^\circ$ in the T14 modeling; $\phi = 0.1 \pm 2.5^\circ$ in the VISIR-SAM data). Secondly, it is safe to combine the squared visibility measurements directly, as we show in §3.1.4 that the fluxes are stable on all scales over the last 17 years. Finally, we can combine the AT and UT data despite their different

inherent spatial filtering because at 30 m baselines, the MIDI AT and MATISSE UT data give consistent correlated flux values, indicating that they probe the same structure. We finally note that the 18 included MIDI AT baselines represent only a small fraction of the imaging data, and serve primarily as a spatial constraint. The results of imaging both with and without the AT data are described further in Appendix G, but in summary the primary small-scale features remain stable in either approach.

3.2 Image reconstruction

The primary advantage of MATISSE over MIDI is that the availability of closure phases makes it possible to reconstruct high-fidelity images. We employ the image reconstruction software, IRBis (*I*mage *R*econstruction software using the *B*ispectrum; Hofmann et al., 2014, 2016), which was designed for MATISSE and is incorporated into the standard data reduction package. IRBis includes six regularization functions, two minimization engines, and myriad fine-tuning parameters such as the pixel scale, hyperparameter, and object mask. For the VISIR-SAM data, a completely independent image reconstruction process was carried out. We kept the image reconstruction for the MATISSE+MIDI data and that for the VISIR-SAM data separate due to dynamical range concerns, the different wavelength ranges, and because the spacial scales they measure are completely independent. We first focus on the image reconstruction of the MATISSE+MIDI data, with MIDI closure phases assumed to be $0 \pm 180^\circ$ (see §3.1.5). The image reconstruction for the VISIR-SAM data will be discussed in §3.2.5.

3.2.1 MATISSE image reconstruction

We select seven wavelength bins in which to produce independent images: $8.5 \pm 0.2 \mu\text{m}$, $8.9 \pm 0.2 \mu\text{m}$, $9.7 \pm 0.2 \mu\text{m}$, $10.5 \pm 0.3 \mu\text{m}$, $11.3 \pm 0.3 \mu\text{m}$, $12.0 \pm 0.2 \mu\text{m}$, and $12.7 \pm 0.2 \mu\text{m}$. Any spectral information within each bin is averaged, producing a series of “gray” images. Each bin is imaged with a range of regularization functions and hyperparameters (hereafter μ ; essentially a scaling on the amount of regularization), with the best selected via a modified χ^2 function:

$$q = \frac{\alpha}{N_{V^2}} \sum_{i=1}^{N_V} \frac{(V_{\text{obs}}^2 - V_{\text{model},i}^2)^2}{\sigma_{V^2,\text{obs}}^2} + \frac{\beta}{N_\phi} \sum_{j=1}^{N_\phi} \frac{(\phi_{\text{obs}} - \phi_{\text{model},i})^2}{\sigma_{\phi,\text{obs}}^2}, \quad (5.2)$$

with α and β serving as weights on either squared visibilities or the closure phases. In IRBis, there are three “cost functions” which vary the relative weighting of the closure phases and squared visibilities during the image reconstruction process. For cost function 1, $\alpha = \beta = 1$; for cost function 2, $\alpha = 0$ and $\beta = 1$. Cost function 3 is more complex, using the χ^2 coming from the sum of the bispectrum phasors and the squared visibilities (Hofmann et al., 2022); in essence replacing the closure phases in the second term of Eq. 5.2 with the bispectrum. We use cost function 1 for the quality assessment of best-fitting images.

In order to produce images, we perform a grid search of the IRBis parameters, varying the field of view (FOV), the pixel number, the object mask, the regularization function, the hyperparameter μ , the cost function, and the reduction engine (ASA-CG or L-BFGS-B, see Hofmann et al. (2016) for more details). We use uniform weighting in the uv -plane, corresponding to `weighting=0` in IRBis. An initial best image is selected in each wavelength bin using Eq. 5.2, and a follow-up round of imaging using the best regularization function and pixel scale is performed. Regularization is a crucial component of ill-posed problems such as image reconstruction where the number of free parameters ($\approx N_{\text{px}}^2$) is much larger than the number of data points. Regularization is

Table 5.2: Final image reconstruction parameters

λ [μm]	Reg. ^{<i>a</i>} Func.	μ^b	FOV ^{<i>c</i>} [mas]	Obj. Mask ^{<i>d</i>} [mas]	Cost ^{<i>e</i>} Func.	χ^2 ^{<i>g</i>} [V^2, ϕ_{T3}]
8.5 ± 0.2	2	0.5	600	120	3	[2.6,3.8]
8.9 ± 0.2	1	0.01	600	160	3	[1.7,3.4]
9.7 ± 0.2	2	0.18	600	120	3	[0.5,0.9]
10.5 ± 0.3	5	0.08	500	120	1	[0.7,0.3]
11.3 ± 0.3	3	0.51	500	120	1	[2.3,1.4]
12.0 ± 0.2	3	0.51	500	120	1	[6.4,1.6]
12.7 ± 0.2	5	0.30	600	140	1	[7.2,1.0]

a: the IRBis regularization function; *b*: the weight on the regularization function (AKA the hyperparameter); *c*: the field of view of the reconstructed image; *d*: the radius of the object mask employed by IRBis in mas; *e*: the cost function used in reconstruction, as described in Eq. 5.2 and in (Hofmann et al., 2022); *g*: the χ^2 terms from the final images entering Eq. 5.2 for the squared visibilities and closure phases, respectively.

the enforcement of an a priori constraint (e.g., smoothness, compactness, edginess, etc.) to prevent overfitting, but the strength of enforcement is set by the hyperparameter. Starting from the initial images, we finely vary the hyperparameter to construct *L*-curves—diagnostic comparisons between the amount of regularization and the residuals of the reconstruction (first applied by Lawson & Hanson, 1995). One identifies the “elbow” of the curve as the image with optimal regularization parameters. This selection is necessary to strike a balance between over-regularization and allowing too many image artifacts to manifest. We give the final parameters for the reconstructions in Table 5.2. We note that different regularization functions in the same wavelength bin often result in very similar morphology, implying that the result is robust and simply not a consequence of regularized noise. Furthermore, the cost function has little effect on the final morphology or image quality and primarily aids convergence. We show the reconstructed images in Fig. 5.4, separating the continuum images from those inside the Si absorption feature. We also show the flux-weighted mean of the continuum images in Fig. 5.5 which represents an *N*-band image. Finally, we show the fit quality of each image in Figs. 5.2 - D.1; we simulate the correlated fluxes and closure phases represented by each image at each *uv*-point and compare to the observed data. We see that overall the images trace the closure phase and correlated flux spectra well, although specific wavelengths at a handful of *uv*-coordinates are discrepant.

3.2.2 Image error analysis

We use the values in Table 5.2, which represent the “best” reconstruction parameters, to estimate the image-plane uncertainties. We do this through delete-*d* jackknife resampling of our *uv*-coverage (the method is developed in Shao & Wu, 1989). In each Monte Carlo realization we randomly discard 10% of each the squared visibilities and the closure phases (i.e., 15 squared visibilities and 10 closure phases). This choice satisfies the criterion for being asymptotically unbiased: $\sqrt{n} < d < n$, where *n* is the sample size and *d* is the number of deleted elements. We then perform the image reconstruction at each wavelength using the parameters given in Table 5.2 and save the results. After 100 realizations, we calculate the median and standard deviation in each pixel of each image. The median image at each wavelength is used as the final image shown in Fig 5.4. The standard deviation image serves as an error map with which we calculate the S/N at each pixel. The error maps and S/N maps are given in Appendix H. We use the median image at each wavelength for our morphological analysis. The patchiness of the extended structure at 12.0 and 12.7 μm is moreover confirmed through measurement of flux within a 14 mas aperture at several points in the polar emission. Taking the image




Fig. 5.4: Compilation of MATISSE images reconstructed independently in each wavelength bin. Pixel scale and field of view are matched in all panels. The FWHM of the beam is shown in the bottom right corner of each panel. The top row consists of images of the continuum emission, and the bottom row holds the images within the Si absorption feature. The bottom row also includes the reconstructed VISIR-SAM image (panel *h*), which has the field of view of the MATISSE images overlaid. All images are scaled to the power of 0.6. Contours are drawn at $5\times$ the mean image error in each wavelength channel (see §3.2.2).

errors into account, the differences between adjacent bright and dark regions (e.g., at $[(51.6, 23.4), (51.6, 46.9)]$ mas and $[(18.8, 37.5), (32.8, 37.5)]$ mas from the image center) are $\geq 2\sigma$.

3.2.3 Morphology

In the final, independently reconstructed images (shown in Fig. 5.4), we find several consistent and key features. We discuss each below and have labeled them in Fig 5.5 for reference.

A central disk-like component. This component is resolved in the NE-SW direction (≈ 1.9 pc across), but unresolved at all wavelengths in the NW-SE direction. Its orientation varies slightly in the different wavelength channels: along $PA_{disk} \approx 45^\circ$ in the 8.5 and 8.9 μm images; and along $PA_{disk} \approx 30^\circ$ in the images red-ward of the Si feature.

Central, unresolved flux. It is ≈ 10 mas ($= 0.2$ pc) to the NE of the center of the disk in the 12 μm image. This is the brightest feature of the image at all wavelengths.

Significant extended emission in the polar direction ($PA \sim 295^\circ$), perpendicular to the maser emission and roughly aligned with the radio jet (see Fig. 5.5 for the orientations). The large-scale emission is more prominent at longer wavelengths. In the 11.3, 12.0, and 12.7 μm images, the extended emission is approximately symmetric about the photo-center, and it is roughly 4×1.5 pc. This emission is notably not smooth, and shows patchiness far above the noise level.

Two bright components, forming a rough line with the photo-center at $PA_{E-W} = -80^\circ$ and superimposed on the polar emission, are observed for the first time. These substructures become more prominent at longer wavelengths, but are nonetheless present in all channels. They each extend to ~ 65 mas ($= 1.2$ pc) from the center and are both roughly 30 mas across at 12 μm .

At each wavelength we measure the flux contributions of the unresolved component,




Fig. 5.5: *N*-band continuum image and component labels. In the (*left*) panel, we show the flux-weighted mean MATISSE image – a proxy for an *N*-band continuum image. Contours are drawn at $5\times$ the flux-weighted mean of the individual image errors. The cyan dashed ellipses represent the FWHM of the Gaussians fitted to MIDI observations of Circinus by T14. In the (*center*) panel we show the same image as a contour map with levels at $[5, 10, 20, 40, 80]\times$ the $5\times$ the flux-weighted mean of the individual image errors. Key morphological features are labeled: the 1.9 pc disk with $i \gtrsim 83^\circ$, the polar emission, and the polar flux enhancement. In the (*right*) panel we show the same image with the Greenhill et al. (2003) masers overplotted. The black dashed line represents the direction of the radio jet (Elmouttie et al., 1998). The cyan lines show the central PA and opening angle of the optical ionization cone (Fischer et al., 2013). Pixel scale and field of view are matched in both panels. All images and contours are scaled to the power of 0.65. The FWHM of the beam is shown in the bottom right corner of each panel. North is up and east is to the left.

Table 5.3: Measured fluxes of circumnuclear dust components

λ [μm]	F_{polar} [Jy]	F_{disk} [Jy]	$F_{\text{unres.}}$ [Jy]	f_{polar} [%]	f_{disk} [%]	$f_{\text{unres.}}$ [%]
8.5 ± 0.2	1.84 ± 0.07	0.46 ± 0.03	0.39 ± 0.01	42.5 ± 1.5	10.5 ± 0.6	9.0 ± 0.2
8.9 ± 0.2	1.18 ± 0.08	0.29 ± 0.04	0.36 ± 0.01	37.4 ± 2.5	9.2 ± 1.1	11.3 ± 0.4
9.7 ± 0.2	1.40 ± 0.16	0.28 ± 0.06	0.28 ± 0.02	48.6 ± 5.6	9.8 ± 2.1	9.6 ± 0.7
10.5 ± 0.3	3.28 ± 0.34	0.60 ± 0.09	0.29 ± 0.02	64.7 ± 6.6	11.8 ± 1.7	5.6 ± 0.4
11.3 ± 0.3	5.42 ± 0.42	1.00 ± 0.11	0.49 ± 0.03	63.9 ± 4.9	11.8 ± 1.3	5.7 ± 0.3
12.0 ± 0.2	7.79 ± 0.58	1.61 ± 0.17	0.77 ± 0.04	64.0 ± 4.8	13.2 ± 1.4	6.3 ± 0.3
12.7 ± 0.2	9.44 ± 0.93	2.04 ± 0.28	0.96 ± 0.07	61.5 ± 6.1	13.3 ± 1.8	6.3 ± 0.4

Aperture fluxes (*left*) and fractions of the total photometric flux (*right*) for the polar emission, disk, and unresolved component in each image reconstruction. Fractional values do not sum to 100% because some of the total flux is resolved out by the ≥ 30 m baselines.

the disk, and the polar emission. These values are the total flux inside elliptical apertures placed at the image center with dimensions $(10 \times 10 \text{ mas})$, $(100 \times 10 \text{ mas})$ with $\text{PA} = 45^\circ$, and $(220 \times 120 \text{ mas})$ with $\text{PA} = -65^\circ$, respectively. The contributions from the disk and unresolved component are subtracted from the largest aperture. Similarly, the contribution from the unresolved component is subtracted from the disk aperture. These values are presented in Table 5.3. The fractional contribution of the unresolved component increases to shorter wavelengths, indicating that it contains relatively hot dust; conversely, the polar emission contribution increases at longer wavelengths because it is cooler.

There are several features which, while containing a large amount of flux, we consider to be artifacts of the image reconstruction process. A first, simple criterion is to take a S/N cut of $\sigma_{\text{image}} \geq 3$, using the errors derived in §3.2.2. This simple cut agrees well with the following, more detailed considerations. If structures increase their size or radial distance from the photo-center linearly with wavelength, it is likely that they are artifacts of the *uv*-coverage. This is complicated, however, as different wavelengths probe different temperatures in the thermal infrared, and thus real structures may become

“larger” at longer wavelengths where cooler dust is observed. An example of an artifact which varies with wavelength is the pair of arc-like emission features ~ 100 mas to the NE and SW of the photocenter in the $9.7 \mu\text{m}$ image. These appear to correspond to the secondary peaks of the dirty beam (Appendix Fig. F.1). Finally, we assume that structures in the continuum should vary smoothly between adjacent imaging bins, and so we only consider those structures which are present in both the 8.5 and $8.9 \mu\text{m}$ images and those structures which are in all of the 11.3 , 12.0 , and $12.7 \mu\text{m}$ images to be real. We take the flux-weighted average of these five continuum images to produce a proxy for an N -band image. This is shown in Fig. 5.5. The continuum-average image emphasizes consistent features of the images while suppressing artifacts.

3.2.4 Effects of uv -coverage on image morphology

In this section we check what effects the attainable uv -coverage could have on our final images. On the two longest baselines, we have no uv -coverage for $\psi > 110^\circ$, and even more notably, we have no uv -measurements on any baseline for $135^\circ < \psi < 180^\circ$. These uv -holes are currently unavoidable due to VLTI delay line shadowing on the UT1-UT4 and UT2-UT3 baselines for Circinus at DEC = $-65 : 20 : 21$. While MATISSE can be used in a two-baseline configuration, we would not be able to measure the closure phases necessary for imaging. Nonetheless, this uv -region has been shown by T14 to be important, as the disk-like structure present in their modeling is primarily constrained by long baselines in this direction. T14 reports MIDI measurements of the UT1-UT3 baseline (~ 90 m) in the uv -region we cannot currently measure. In order to test the effects of including measurements at these ψ , we performed a second round of imaging, incorporating MIDI baselines with $\text{BL} \in [30, 100]$ m, $\psi_{\text{MIDI}} \in [100, 180]^\circ$, and which are separated by at least 4 m in uv -coordinates. These criteria resulted in 18 additional baselines with correlated flux measurements from MIDI. As there were no closure phases for these baselines, we use the same procedure as when including the MIDI AT measurements, setting $\phi_{\text{T3,MIDI}} = 0 \pm 180^\circ$.

Adding these 18 MIDI UT correlated fluxes to the 150 MATISSE measurements and the 18 MIDI AT measurements, we produced independent images at 8.9 and $12.0 \mu\text{m}$. At $8.9 \mu\text{m}$ the resulting image is essentially unchanged, indiscernible by eye from the image shown in Fig. 5.4; an explicit comparison is shown in Fig. G.1. At $12.0 \mu\text{m}$, however, the disk becomes more prominent and changes position angle slightly, while all other features remain constant. The disk-like structure in the initial imaging lies along $\text{PA}_{\text{disk}} \sim 35^\circ$, while after the addition of MIDI UT baselines the same disk-like structure lies along $\sim 40^\circ$. The latter value more closely resembles the $46 \pm 3^\circ$ given in T14. However, given the size of the beam at $12 \mu\text{m}$, 9 mas, the disk orientation could quite easily vary in the image by $\sim 3^\circ$. The overall differences in the image plane are small when we include these baselines, so we proceed in our analysis without the MIDI UT measurements. It is clear from T14, however, that these baselines are important to understand the size and orientation of the disk-like structure in Circinus, and the planned doubling of the VLTI delay lines will make closure phase measurements including these baselines possible.

3.2.5 VISIR-SAM image reconstruction

The VISIR-SAM data allowed us to reconstruct an image of the target’s large scale structure at $11.3 \mu\text{m}$. We reconstruct the VISIR-SAM image and the MATISSE images separately, rather than combining the uv -coverage for two reasons: first, the longest SAM baselines (≈ 7 m) are shorter than the shortest MIDI-AT baselines, meaning there is no overlap in measured visibilities; second, the SAM data exhibit squared visibilities

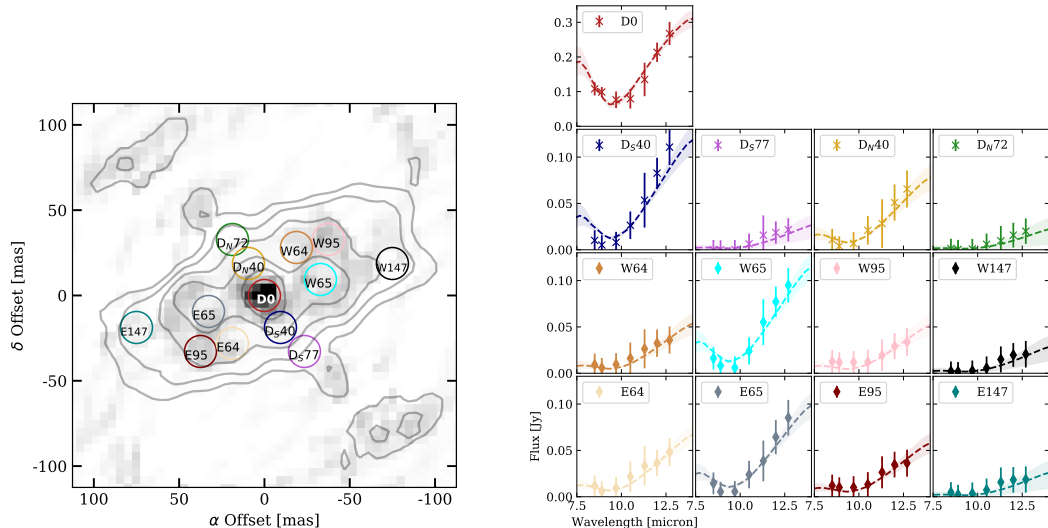


Fig. 5.6: In the (*left*) panel we show the continuum image contours on top of the $8.9 \mu\text{m}$ image. Overplotted on these are the 13 apertures we used for SED extraction. The contours are at [5, 10, 20, 40, 80] times the continuum image error (see Fig. 5.5). All images and contours are scaled to the power of 0.65. In the (*right*) panel we show the extracted mean flux at each wavelength in each aperture as well as the best-fitting blackbody functions to each SED. These fits had uniform priors on T and Gaussian priors on A_V with $\mu = 28.5 \text{ mag}$ and $\sigma = 8.1 \text{ mag}$. The shaded regions represent 1σ uncertainties as estimated from the posterior probability distributions.

> 0.4 , which are much larger than the MATISSE values on longer baselines and would result in dynamical range issues during imaging. The data are nonetheless useful as a way to contextualize the MATISSE images and to measure the true extent of the polar emission. We used BSMEM (Buscher, 1994; Lawson et al., 2004b) to reconstruct the VISIR-SAM image. This code uses a regularized minimization algorithm to recover an image from infrared interferometric data. The regularized optimization engine uses a trust-region gradient-descent method with entropy (i.e., the sum of the logarithm of the pixel values in the image grid) as regularization function. Images were reconstructed using squared visibilities and closure phases simultaneously. The reconstructed image used a pixel scale of 15 mas in a pixel grid of 512×512 pixels. The code converges to a χ^2 close to unity. Figure 5.4 shows our VISIR-SAM image.

3.3 Measuring the dust temperature distribution

The images produced above supply not only morphological information, but also information about the temperature and optical depth of the dust in different regions. In this section, we fit one-temperature blackbody models with extinction to a series of apertures. As shown in Ch. 4, Gaussian modeling, point-source fitting, and image reconstruction all resulted in similar extracted SEDs in NGC 1068; therefore we can with confidence use the extracted apertures from our reconstructed image to undertake a temperature analysis of Circinus. We first convolve the images to the beam of the lowest resolution reconstruction ($12.7 \mu\text{m}$, corresponding to 10.1 mas). The individual images are aligned using cross-correlation before SED extraction, but effectively the photo-centers are simply matched. Then we define 13 apertures (shown in Fig. 5.6) which are 5 px (23.4 mas) in diameter and which do not overlap; their exact locations were chosen by hand to cover key features of the disk and polar emission. This is $> 2\times$ larger than the lowest resolution “beam” in our images, and in this way we do not make claims based on any hyper-resolved features. We extract the mean flux from each aperture in each image and

estimate the flux error from the same apertures on the error maps estimated in §3.2.2. Finally, we add the calibration error of the total photometric flux at each wavelength in quadrature to the extracted flux error.

We fit a single blackbody (BB) curve with absorption to each aperture-extracted spectrum with the form⁵

$$I(\lambda, T, A_v) = BB_\nu(\lambda, T) e^{-A_v/1.09\tau(\lambda)/\tau_v}, \quad (5.3)$$

where $\tau(\lambda)/\tau_v = \kappa(\lambda)/\kappa_v$, and we use the standard interstellar medium $\kappa(\lambda)$ profile from Schartmann et al. (2005) which is based on the standard interstellar medium profile of Mathis et al. (1977). Fitting of T and A_V is done in two iterations using Markov Chain Monte Carlo sampling with the package `emcee` (Foreman-Mackey et al., 2013). Final values in each iteration are the median of the marginalized posterior probability distribution. The 16th and 84th percentiles of the resulting temperature and extinction distributions are used as the 1σ fit uncertainties⁶.

In the first iteration, we use uniform prior probability distributions with $T \in [100, 600]$ K and $A_V \in [0, 100]$ mag. We find that the extinction does not vary significantly across the field. The central aperture, fit with the smallest uncertainties, shows $A_V = 28.5^{+8.5}_{-7.7}$ mag which is $\tau_{9.7} = 2.0^{+0.6}_{-0.5}$ using the mass-extinction profile from Schartmann et al. (2005). The other apertures have nominally higher values, but large uncertainties which make the differences insignificant. Only D_S40, W65, and E65 show differences $> 1\sigma$ from the central value.

In the second iteration, we use again a uniform prior for temperature ($T \in [100, 600]$ K), but a Gaussian prior for A_V with $\mu = 28.5$ mag and $\sigma = 8.1$ mag based on the fit to D0 in the first iteration. The central aperture, D0, serves as a good estimate of the overall extinction because it *a)* has the highest S/N and *b)* has significant flux on both sides of the Si absorption feature. The resulting temperatures are consistent with the unconstrained case but are typically lower. The qualitative behavior of the temperature distribution is unchanged, but the fitted uncertainties are greatly diminished due to the degenerate nature of A_V and T for a fit to the Rayleigh-Jeans tail of the Planck function. In the following discussion, we therefore use the values from the second iteration. We show the best fit parameters for each aperture in each iteration in Table 5.4.

We do not find evidence of an extinction gradient across the field, indicating that there is a relatively uniform screen of foreground absorption. In the first fitting iteration, with A_V allowed to vary, the mean extinction values are similar to the east and west. In the second iteration, we restricted A_V around 28.5 mag, to get better constraints on the temperature. Based on Hubble *K*-band imaging, Wilson et al. (2000) estimated an extinction of $A_V = 28 \pm 7$ toward a compact (< 2 pc) nucleus. Burtscher et al. (2016) measured a value of $A_V = 27.2 \pm 3$ using SINFONI in the *K*-band. Roche et al. (2006) found $2.2 \leq \tau_{9.7} \leq 3.5$ using T-ReCS on Gemini-South. Previous measurements are nearly identical to the fitted value in D0, $28.7^{+8.5}_{-7.7}$ mag, and furthermore consistent with the rest of the field. Uniform absorption, however, is in contrast to the $\Delta\tau = 27$ arcsec⁻¹ gradient across the polar emission measured by T14. This discrepancy is puzzling, but we recognize there are major differences between our approach and that of T14. Specifically, T14 used differential phases and Gaussian modeling due to the lack of closure phase data. Their differential phases were measured on the UT and AT baselines, and thus probe larger-scale material than the MATISSE UT closure phases alone. On the other hand, we use no differential phases and had to assume an unconstrained AT closure phase value of $0 \pm 180^\circ$. However, we note that on the UT baselines (probing $\lesssim 1$ pc scales), the

⁵We do not include a “graybody” emissivity here because the two-parameter model provides robust fits to the spectra.

⁶Valid only because the resulting distributions are approximately Gaussian.

9.7 μm closure phases are well matched by our images *without* an extinction gradient. The phase signals are instead produced by small-scale structure that was smoothed out in the Gaussian modeling approach of T14. The two approaches emphasize different aspects of the data, but differential phases could be included with the closure phases in future work in chromatic image reconstructions. Future closure phase measurements at 9.7 μm are required on shorter baselines (e.g., with MATISSE AT observations) to directly measure the Si absorption across the large-scale component.

We separate the apertures into two rough categories based on their locations. Those oriented NE and SW from the photocenter at $\text{PA} \sim 30^\circ$ are labeled as “disk” apertures, based on the presence of a thin disk-like structure in both our reconstructed images and in the Gaussian modeling of T14. The other points, extending NW and SE from the photocenter are labeled “outflow” apertures, as they lie in the direction of the polar extension. The extracted spectra and the fitted blackbody curves (with uncertainty estimates as shaded regions) are shown in Fig. 5.6. The two-dimensional temperature distribution based on the fits is shown in Fig. 5.7. We find that on average the “disk” apertures show a much steeper temperature falloff with projected distance than the “outflow” apertures.

Table 5.4: Fitted blackbody parameters for each of the 13 image-extracted spectra

Aperture	Dist. [pc]	Temp. [K]	First	Iteration	Second	Iteration
			A_V [mag]	$\tau_{9.7}$	Temp. [K]	A_V [mag]
Disk						
D0	0.00	367^{+30}_{-26}	$28.5^{+8.5}_{-7.7}$	$2.0^{+0.6}_{-0.5}$	–	–
D _S 40	0.40	358^{+68}_{-56}	$66.5^{+22.3}_{-22.8}$	$4.6^{+1.5}_{-1.6}$	281^{+17}_{-17}	$33.3^{+7.0}_{-7.1}$
D _N 40	0.40	297^{+64}_{-50}	$58.2^{+28.9}_{-29.0}$	$4.0^{+2.0}_{-2.0}$	249^{+17}_{-17}	$29.8^{+7.7}_{-7.7}$
D _S 77	0.77	221^{+49}_{-43}	$60.2^{+28.1}_{-36.7}$	$4.1^{+1.9}_{-2.5}$	198^{+19}_{-29}	$29.2^{+8.1}_{-8.2}$
D _N 72	0.72	208^{+51}_{-49}	$58.8^{+29.0}_{-35.7}$	$4.0^{+2.0}_{-2.5}$	191^{+22}_{-41}	$29.2^{+7.9}_{-8.0}$
NW	Polar	Ext.				
W64	0.64	246^{+57}_{-35}	$44.9^{+36.1}_{-28.5}$	$3.1^{+2.5}_{-2.0}$	228^{+16}_{-16}	$28.4^{+7.9}_{-7.6}$
W65	0.65	333^{+64}_{-47}	$58.6^{+23.6}_{-21.1}$	$4.0^{+1.6}_{-1.5}$	277^{+17}_{-17}	$32.0^{+7.4}_{-7.0}$
W95	0.95	244^{+55}_{-34}	$45.5^{+34.4}_{-29.2}$	$3.1^{+2.4}_{-2.0}$	227^{+17}_{-16}	$28.7^{+8.0}_{-8.1}$
W147	1.47	220^{+50}_{-43}	$58.2^{+29.1}_{-35.2}$	$4.0^{+2.0}_{-2.4}$	200^{+19}_{-29}	$29.0^{+7.9}_{-8.1}$
SE	Polar	Ext.				
E64	0.64	261^{+59}_{-38}	$44.4^{+34.2}_{-26.9}$	$3.1^{+2.4}_{-1.8}$	240^{+15}_{-14}	$28.6^{+8.2}_{-7.5}$
E65	0.65	325^{+66}_{-52}	$61.6^{+24.8}_{-25.5}$	$4.2^{+1.7}_{-1.8}$	266^{+17}_{-16}	$31.4^{+7.5}_{-7.3}$
E95	0.95	248^{+54}_{-33}	$42.8^{+34.8}_{-27.0}$	$2.9^{+2.4}_{-1.9}$	231^{+15}_{-11}	$28.1^{+8.1}_{-8.2}$
E147	1.47	214^{+50}_{-44}	$54.9^{+32.3}_{-35.0}$	$3.8^{+2.2}_{-2.4}$	197^{+21}_{-36}	$29.0^{+8.0}_{-8.0}$

$\tau_{9.7}$ is the simple conversion from A_V to the optical depth of the Si feature based on the $\kappa(\lambda)$ curve from Schartmann et al. (2005) and is included only for comparison to previous results, namely T14. Projected distances in parsec are given from the central aperture, D0, with Circinus 4.2 Mpc away (Freeman et al., 1977). We measure the inclination to be $i \gtrsim 83^\circ$, so the correction from projected to physical distance is small. The two rightmost columns are the results of refitting with a Gaussian prior on $A_V = 28.5 \pm 8.1$, based on the initial fit to D0.

3.3.1 Temperature gradient analysis

In current modeling, the dust in the outflow is anisotropically illuminated by a face-on accretion disk. We can use the temperature profile of the outflow to characterize the dust environment.

We begin with a comparison to the simple analytic model of Barvainis (1987):

$$T_{gr}(r) = 1650 \left(\frac{L_{\text{acc}}}{r^2} \frac{\text{pc}^2}{10^{10} L_{\odot}} \right)^{1/5.6} e^{-\tau_{\text{uv}}/5.6} \text{K}, \quad (5.4)$$




Fig. 5.7: Two-dimensional temperature distribution as fitted in each of 13 apertures. Colors at the edge of each circle match those given by the aperture definitions in Fig. 5.6. Temperatures within the polar extension remain high even at large distances from the center when compared to the “disk” apertures. The contours are at [5, 10, 20, 40, 80] times the continuum image error and scaled to the power of 0.65 (see Fig. 5.5).

where L_{acc} is the luminosity of the accretion disk in L_{\odot} , r is the distance from the accretion disk in parsec, and τ_{uv} is the optical depth to the ultraviolet continuum. Here we use $L_{acc} = 6 \times 10^9 L_{\odot}$, which is the lower bound on estimates of the accretion disk luminosity in Circinus ($6 \times 10^9 - 7 \times 10^{10} L_{\odot}$), as inferred from X-ray (Arévalo et al., 2014; Ricci et al., 2015), IR (T14) and optical (Oliva et al., 1999) observations. We also use $A_V = 40$ mag $\rightarrow \tau_{uv} = 7.2$, set roughly by the mean of the first-iteration fitted extinction values in Table 5.4 and converted using the dust extinction curve of Schartmann et al. (2005), but we note that the best-fitting value, $A_V = 28.5^{+8.5}_{-7.7}$ mag, would result in even higher temperatures at a given radial distance. With these assumed values, we compare the radial temperature profile of the optically thin, continuous dust environment described by Barvainis (1987) to the fitted SED temperatures of the “out-flow” in Fig. 5.8. At all radii, the Eq. 5.4 temperatures are larger than the measured Circinus temperatures by a factor of ~ 2 . This is not completely unexpected, as the inefficient re-radiation by the dust grains in the Barvainis (1987) model leads to higher temperatures at large radii; this model should be considered an upper limit on the dust temperatures at a given radius. Moreover, the Barvainis (1987) model does not take the anisotropy of the radiation into account. We also plot the expected temperature profile arising due to the simplest case of radiation equilibrium for perfectly efficient blackbodies as given in Tristram et al. (2007) for comparison.

3.4 Comparisons to radiative transfer models

3.4.1 Clumpy torus models

Modern AGN “torus” modeling takes the clumpiness of the dust, as implied from infrared interferometry, as well as anisotropic illumination from the accretion disk into account. We compare the temperatures at different radial distances in the standard clumpy torus model of Schartmann et al. (2008) to those fit in Circinus. These models consist of a wedge-shaped torus filled with randomly placed, optically-thick spheres of dust. The clump density falls off with radius, r , from the anisotropically illuminating source as $\rho \propto r^{-0.5}$ and the clump size increases as $a \propto r^{1.0}$. These models consist only of a

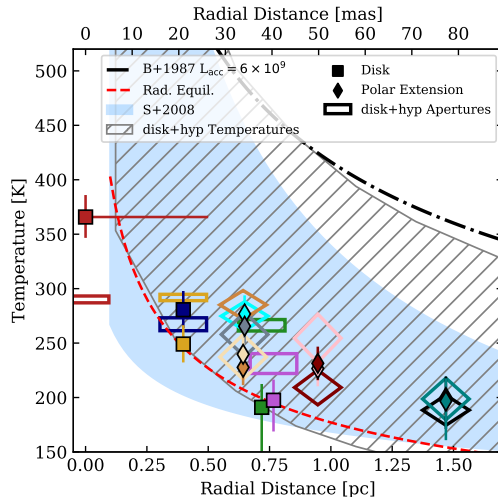


Fig. 5.8: Radial temperature profile of the “torus” in Circinus based on 13 extracted apertures. Measurements along the disk-like structure in the center of image are marked with squares; measurements in the polar extensions are marked with diamonds. Colors correspond to the apertures given in Fig. 5.6. *B+1987* is the radial profile from Barvainis (1987) and is given as the solid black line. The radial profile arising from simple radiation equilibrium is given as the dashed red line. In both analytic profiles the luminosity of the Circinus accretion disk is assumed to be $L_{\text{acc}} = 6 \times 10^9 L_{\odot}$. The shaded blue region labeled *S+2008* represents the range of temperatures of the dust clumps at each radius in the standard clumpy torus model shown in Schartmann et al. (2008), Fig. 3. The hashed region represents the range of temperatures of dust cells in the disk+hyp model. The boxes/diamonds represent fitted blackbody temperatures in our 13 apertures applied to disk+hyp models based on Stalevski et al. (2019); each box (for disk apertures) or diamond (for polar apertures) center is the median temperature T and the height of each represents the range of temperatures fitted to the disk+hyp models.

puffy “disk” with half-opening angle $\theta = 45^\circ$, as they predate the observations of polar dust in Circinus. The clumpy torus models produce a range of dust temperatures as a function of radius which serve as a theoretical bound on the temperature distribution in the central few parsec. The temperatures found by our blackbody fits are clearly within these theoretical bounds of the model, c.f. Fig. 5.8. A similar result was already found by Tristram et al. (2007, 2014).

3.4.2 Disk + wind models

More recently, Stalevski et al. (2017, 2019) undertook radiative transfer modeling of VISIR imaging data, the MIR SED, and MIDI interferometric data of Circinus. Their best-fitting model (presented in Stalevski et al., 2019) consists of a compact, dusty disk and a hollow hyperbolic cone extending in the polar direction (hereafter disk+hyp). In this modeling, a parameter grid for the radiative transfer models was searched such that the overall SED as well as the interferometric observables were well reproduced. This was not a model fit, but rather an exploration of the parameter space. For comparison with the MATISSE data, we start from the best model of Stalevski et al. (2019) and vary its parameters with finer sampling of the parameter space. We significantly expand the explored range of the parameters which define the clumpiness: the number of clumps (i.e. filling factor) and different random realizations of the clumps’ positions (set by the “seeds” for the random number generator). Using the MATISSE *uv*-coverage, we simulate the squared visibilities and closure phases of each model image and compute the χ^2 to the data (the comparisons and resulting model are shown in Fig. I.1). This comparison placed constraints on the system inclination ($i \sim 85^\circ$), the hyperboloid opening angle ($\theta_{\text{OA}} \sim 30^\circ$), the disk Si feature depth ($\tau_{\text{Si,DSK}} \sim 15$), and the outer

radius of the disk ($r_{\text{out}} \sim 3$ pc). The closure phase comparison favored a small number of clumps. We then performed a finer parameter search based on these constraints, focusing on the filling factors of the disk and the hyperboloid. After comparing with the MATISSE data, the parameter values defining the boundaries of the model geometry remain unchanged (the dusty disk outer radius, angular width, optical depth; hyperboloid shell position, width and optical depth). However, our modeling converged on a smaller number of clumps (30% less than in the MIDI model) and found that random positions of the clumps have a significant impact on the quality of the fits. The selected model exhibits a sky covering fraction of 78% due to the dust clumps at $0.55 \mu\text{m}$. We show in Fig. 5.8 the average dust temperature as a function radius; these are indicative temperatures obtained by averaging the local thermal equilibrium temperatures over the dust species and grain sizes. We finally extract fluxes in each of the 13 apertures and fit blackbody temperatures using Eq. 5.10 to the disk+hyp model grid at $\lambda \in [8.53, 8.91, 9.29, 9.70, 10.12, 10.56, 11.02, 11.50, 12.00, 12.52] \mu\text{m}$.

We compare the extracted model spectrum in each aperture to the observed spectra. We quantify this through the χ^2 but do not perform any model fitting. These comparisons are shown in Fig. 1.2. In the polar extension, the model and image extracted fluxes and temperatures agree well. The preferred model of [Stalevski et al. \(2019\)](#) includes the polar dust flux-enhancements E-W of the center. Along the disk, and particularly in the central aperture, D0, we see significant discrepancies. The central aperture temperature is ~ 100 K less in the models than in the observations and the extracted flux is $\lesssim 10\%$ of the observations. These discrepancies may indicate that the model disk is perhaps too dense. The disk apertures D_S77 and D_N77 also show much lower observed temperatures than the model predicts, indicating that the model disk can be further improved. Given that the outer radius, angular width, average edge-on optical depth and inclination of the disk appear to be well constrained, it is likely that the disk is actually inhomogeneous, or perhaps with a gap, thus allowing more warm emission to escape. *LM*-band images at ~ 3 mas resolution are required to further constrain the disk component in modeling of Circinus. The MATISSE observations and imaging of Circinus agree very well with clumpy modeling, but it is beyond the scope of this work to place constraints on the specifics of a clumpy medium.

In the next section, we will utilize the *LM*-band MATISSE data, with ~ 3 mas resolution, to probe the hotter dust at small scales both within the disk and perhaps at the origins of the polar extension. We show that the central aperture can contain dust near the sublimation temperature, and a detailed study of the *LM*-data can give insights into the density, thickness, and perhaps the clumpiness of the disk.

4 *L*- and *M*-band Analysis

The following section contains new, unpublished work which extends the analysis of the *N*-band from the previous section down to the *L*- and *M*-bands using the same observations and similar techniques.

4.1 MATISSE *LM* data reduction and calibration

The *LM*-band data for Circinus and the calibrators listed in §2 were initially reduced using DRS version 1.5.1. We reduced the data both coherently and incoherently. For the coherent reduction we used the flags `corrFlux=TRUE` and `coherentAlgo=2` in order to produce correlated fluxes. For the incoherent reduction we used the flag `corrFlux=FALSE`. In both cases use spectral binning 11 px ($= 0.5 \mu\text{m}$) and the default values for all other parameters. The DRS is not optimized for coherent reduction of

the *LM*-bands, but coherent integration is necessary for faint sources (see e.g., the flux limits given in Lopez et al. (2022)). The above settings (both coherent and incoherent) resulted in strange spectra with 1) no *M*-band flux and 2) sharp emission features at 3.7 μm ; but analysis of the intermediate products (i.e., the cleaned interferogram) found neither of these features. It was found that a bias floor was present in the DRS-reduced data. This is likely due to the fringe search being optimized for 3.6 μm , but the Circinus spectrum is very “red” and the *L*-band flux is very low. Instead, a fringe extraction using the *M*-band was necessary.

The data are then re-reduced using a custom python script (found here). The custom pipeline uses the intermediate products of the DRS, specifically the complex cleaned interferograms (files called OBJ_CORR_FLUX). Using the 4.6 μm flux, the 6 fringes are identified and extracted in each exposure snapshot and each BCD configuration. Additionally, a bias “fringe” per frame is extracted far from the science fringes. The extracted fringes for each baseline and the extracted bispectra for each closure triplet are then bias corrected and temporally averaged (over the exposure cycle). See Ch. 2.4.5 for more details. The resulting correlated fluxes and closure phases are computed for each BCD independently, and the final values are taken as the mean of the 4 BCD configurations. The final errors are the standard deviations of the 4 BCD configurations. In the *L*-band the correlated flux errors are typically 1.8 mJy, and in the *M*-band the correlated flux errors are typically 5-10 mJy. In both bands the closure phase errors are in general quite large, $\gtrsim 90^\circ$ and their use is limited. This process is done for both the calibrators and Circinus, and the resulting observables are calibrated as usual. Reductions using the DRS and the custom pipeline for both a calibration star and for Circinus are shown in Fig. 5.9 for comparison and validation of the approach.

The correlated flux, $F(u, v, \lambda)$ is then calibrated in the standard way:

$$F_{\text{targ}}^{\text{cal}}(u, v, \lambda) = F_{\text{targ}}^{\text{raw}}(u, v, \lambda) / F_{\star}^{\text{raw}}(u, v, \lambda) \times 2J_0(\pi dB/\lambda) \times F_{\star}^{\text{tot}}(\lambda), \quad (5.5)$$

where F_{\star}^{raw} is the raw flux (in counts) of the calibrator, F_{\star}^{tot} is the catalog flux of the calibrator, $F_{\text{targ}}^{\text{raw}}$ is the raw flux of the target, $B = \sqrt{u^2 + v^2}$, and d is the diameter of the calibration star. The squared visibilities are finally calculated as $V^2(u, v, \lambda) = [F_{\text{targ}}^{\text{cal}}(u, v, \lambda) / F_{\text{targ}}^{\text{tot}}(\lambda)]^2$, where the total flux $F_{\text{targ}}^{\text{tot}}$ comes from the shortest baseline correlated fluxes; in this case we use the azimuthal maximum of the 30-35 m correlated fluxes as an estimate. The 30 m total flux and sample of all the correlated fluxes are shown in Fig. 5.10 (the remaining correlated fluxes are shown in Appendix Fig. J.1). We use this 30 m correlated flux rather than the “zero-baseline” correlated flux because squared visibilities can be scaled somewhat arbitrarily; the relative changes give substructure and the absolute changes are indicative of background emission. The 8.2 m total flux from ISAAC observations (458.16 ± 39.18 mJy; Isbell et al., 2021) would result in extremely small squared visibilities ($\sim 3 \times 10^{-4}$) and cause numerical issues with little gain in understanding of the source.

The closure phases are calibrated as in Ch. §3.1, but they instead use the results of the custom pipeline. Due to the low signal-to-noise ratio (S/N) for the visibilities on many baselines, only four closure phase triangles give reasonable values. The others are all dominated by noise, as mentioned above. The closure phases are shown in Appendix Fig. K.1.

4.2 Modeling and Image Reconstruction

Both image reconstruction and modeling fitting rely on minimization of a cost function, q , measuring the flux distribution’s similarity to a weighted combination of the closure phases and squared visibilities, which is given in Eq. 5.2. In the case of image

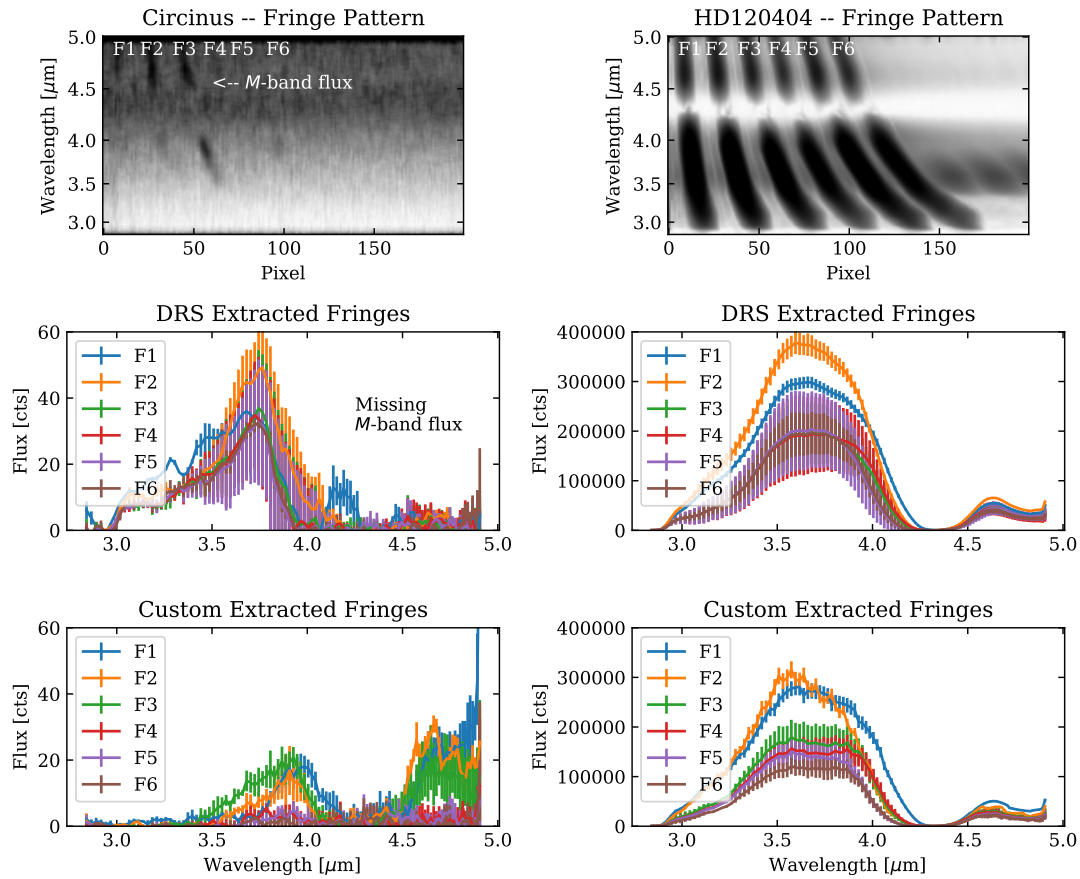


Fig. 5.9: Comparison of DRS-extracted and custom-extracted correlated fluxes for Circinus (*left*) and a calibration star, HD120404 (*right*). From top to bottom, the panels show the cleaned fringe pattern (the Fourier transform of the cleaned interferogram), the six fringes extracted with the DRS, and the same six fringes extracted with our method. In the Circinus fringe pattern, there is significant flux in the *M*-band which is missing in the DRS fringes but it is present in our method.

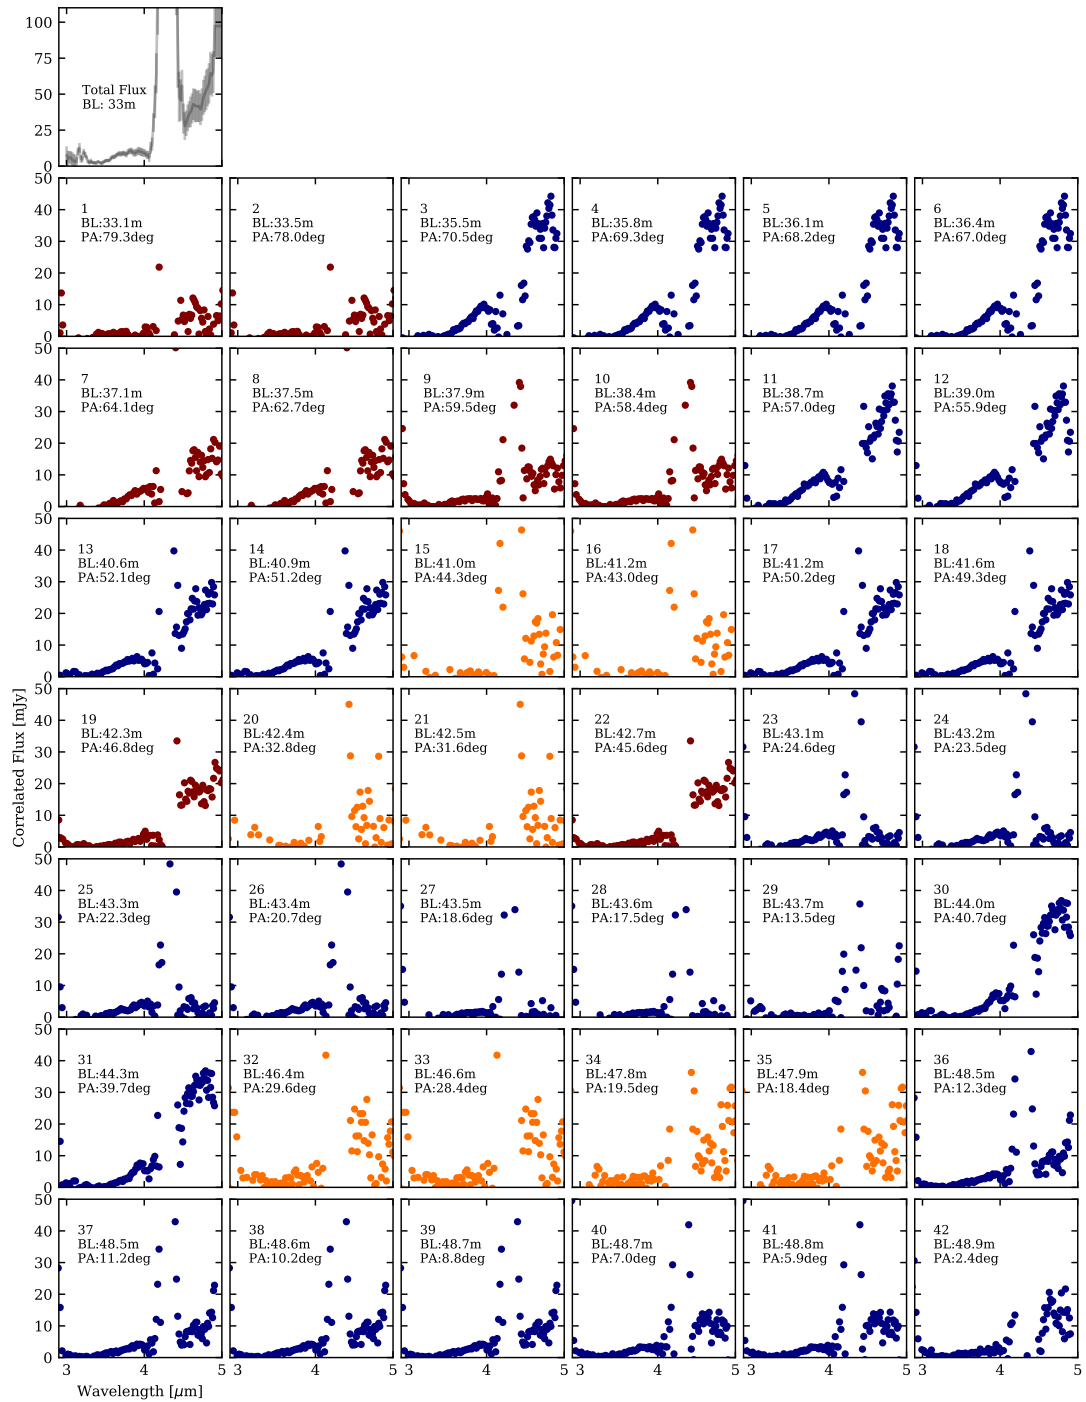


Fig. 5.10: Excerpt sample of calibrated correlated fluxes for Circinus in the *LM*-band. The top left panel shows the 33 m “total flux” used to compute squared visibilities. The colors indicate the observing date: March 2020 (blue), February 2021 (yellow), and May 2021 (red). Continued in Appendix Fig. J.1.

reconstruction, a regularization term is added yielding the cost function for an image, I , sampled at the vector of uv -coordinates \vec{x}

$$J(I(\vec{x})) = q(V_{\text{model}}^2, \phi_{\text{model}}) + \mu R(I(\vec{x})), \quad (5.6)$$

where q represents Eq. 5.2, V_{model}^2 and ϕ_{model} are the squared visibilities and closure phases of the image flux distribution sampled at \vec{x} , μ is the so-called “hyperparameter” which sets the amount of regularization applied and R the regularization function to be applied (the regularization functions used in IRBis are given in Ch. 3 §3.2).

We first attempted model independent image reconstruction, as these are the first interferometric observations of Circinus in the LM -bands, and the source flux distribution is unknown. Following the relatively simple results of the image reconstruction, we performed follow-up modeling of the flux using elongated Gaussian components.

4.2.1 Image Reconstruction

Following the procedure in §3.2 we use IRBis to reconstruct images in the LM -band. We select two wavelength bins in which to produce independent images: $3.7 \pm 0.1 \mu\text{m}$ and $4.7 \pm 0.1 \mu\text{m}$. Any spectral information within each bin is averaged, producing a series of “gray” images. Each bin is imaged with a range of regularization functions and hyperparameters (hereafter μ ; essentially a scaling on the amount of regularization), with the best selected via Eq. 5.2. We perform a grid search of the IRBis parameters, varying the field of view (FOV), the pixel number, the object mask, the regularization function, and the hyperparameter μ . We use non-uniform weighting in the uv -plane, setting `weighting=0.5` in IRBis to de-emphasize the sparsely-sampled and low-S/N points on baselines longer than ~ 60 m.

An initial best image is selected in each wavelength bin using minimization of Eq. 5.2, and a follow-up round of imaging using the best regularization function and pixel scale is performed. We give the final parameters for the reconstructions in Table 5.5. The images are shown in Fig. 5.11; here it is apparent that the sparse uv -coverage and low S/N of the correlated fluxes have resulted in significant image artifacts, particularly in the M -band. Image errors are estimated as in §3.2.2.

Despite the artifacts, the results of 1) a point source in the L -band and 2) a disk-like component in the M -band at $\approx 45^\circ$ position angle are robust above the noise. The point-like source is present in the L -band, but it is not present in the M -band image reconstruction. There is, nonetheless, a peak in flux in the M -band image at the same pixel position as the point source. The disk-like component is roughly 4 mas in width. This is comparable to the resolution at $4.7 \mu\text{m}$ (3.7 mas), so the disk width is only marginally resolved. The disk extends 20.2 mas in FWHM along $\text{PA} \approx 46^\circ$. There are several lines of flux which cross the M -band image along $\text{PA} \sim 40$; these are certainly artifacts due to their regular spacing, symmetry, and flux falloff with radius. Peaks along these lines indicate potential real polar flux, but they are at too low significance to analyze robustly.

4.2.2 Gaussian Modeling

The image reconstructions result in simple components: a point source at $3.7 \mu\text{m}$ and a disk-like structure at $4.7 \mu\text{m}$. These simple structures can be modeled using elongated Gaussian components, which serves to minimize the effect of the dirty beam and image artifacts due to sparse uv -coverage and a small number of closure phases.

The fitted Gaussian components are fixed to the center of the image and only their major axis Θ , their minor-to-major axis ratio $r \equiv \theta/\Theta$, the position angle of the major

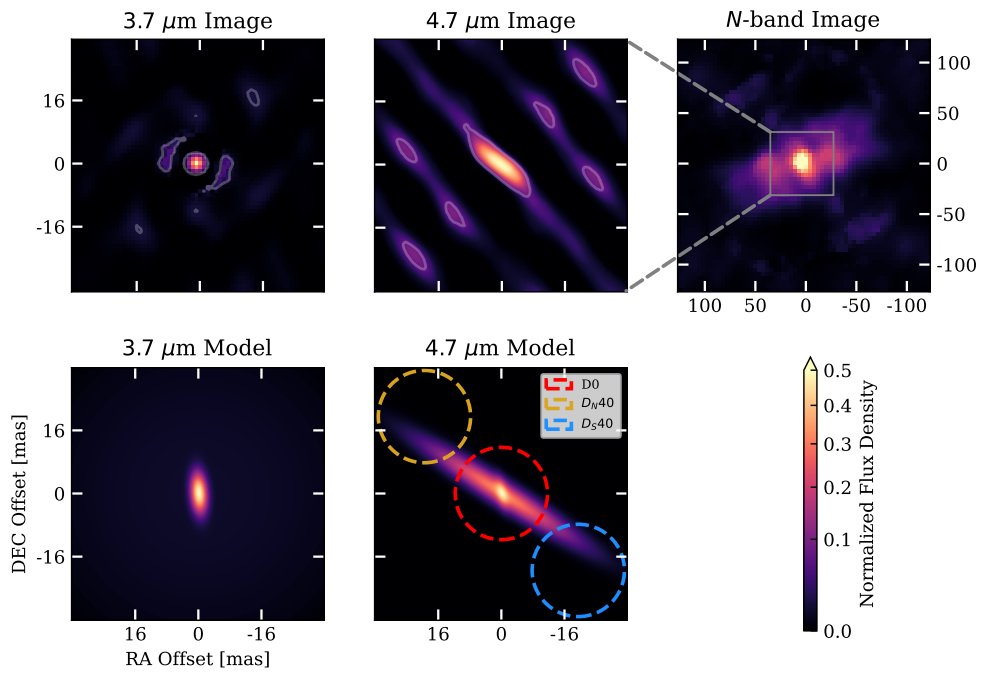


Fig. 5.11: Images and models of the Circinus *LM*-band dust. (*top*) Image reconstructions at 3.7 and 4.7 μm . The contours are at $10 \times \sigma_{\text{Im}}$, estimated from the error maps produced by delete-*d* Jackkifing the *uv*-coverage (see §3.2.2). (*bottom*) Gaussian model fits to the same wavelengths. Also included is the *N*-band continuum image from Fig. 5.5 for reference. The colored circles in the bottom-right panel illustrate the extraction apertures for the measured fluxes.

Table 5.5: Final image reconstruction parameters

λ [μm]	Reg. ^a Func.	μ^b	FOV ^c [mas]	N_{px} -	Obj. Mask ^d [mas]	Cost ^e Func.	χ^2 ^g [V^2, ϕ_{T3}]
3.7 ± 0.1	1	0.01	128	128	120	1	[1.31, 0.4]
4.7 ± 0.1	1	0.01	128	256	50	1	[1.14, 0.4]

a: the IRBis regularization function; *b*: the weight on the regularization function (AKA the hyperparameter); *c*: the field of view of the reconstructed image; *d*: the radius of the object mask employed by IRBis in mas; *e*: the cost function used in reconstruction, as described in Eq. 5.2 and in (Hofmann et al., 2022); *g*: the χ^2 terms from the final images entering Eq. 5.2 for the squared visibilities and closure phases, respectively.

axis ϕ , and their relative flux f are allowed to vary. We fit a number of Gaussian components N_{comp} to each wavelength channel, but in general favor models with fewer components. The relative flux f of one Gaussian is fixed to 1. The fitted Gaussians' parameters and error estimates are obtained through Markov-Chain Monte Carlo likelihood maximization. We sample the parameter space using the package `emcee` (Foreman-Mackey et al., 2013). The log-probability function to be maximized is given by the typical Bayesian formulation

$$p(\vec{\theta}, c | \vec{x}, y, \sigma) \propto p(\vec{\theta}) p(y | \vec{x}, \sigma, \vec{\theta}, c). \quad (5.7)$$

with measurements $y = (V_{\text{model}}^2, \phi_{\text{model}})$ at uv -coordinates \vec{x} , parameters $\vec{\theta}$ and error estimates σ scaled by some constant c . For maximum likelihood estimation, the log likelihood function for the models is written as

$$\ln L(y | \vec{x}, \sigma, \vec{\theta}, c) = -\frac{1}{2} \sum_n \left[q(y) + \ln(2\pi s_n^2) \right], \quad (5.8)$$

where q (Eq. 5.2) is the cost function for the squared visibilities and closure phases produced by the model with parameters $\vec{\theta}$, c represents the underestimation of the variance by some fractional amount, and $s_n^2 = \sigma_n^2 + c^2 f(x_n, \vec{\theta})^2$. We estimate the best-fit value as the median of each marginalized posterior distribution and the 1σ errors from the values at 16th and 84th percentiles. Because the closure phases are very low S/N, we fix all components to the center of the image, and we set $\beta = 0$ in Eq. 5.2 to fit the squared visibilities alone.

We first fit one Gaussian (i.e., $N_{\text{comp}} = 1$) based on the simplicity of the image reconstruction, but Gaussian fits with more components were attempted as well. We favor models with fewer parameters based on the Akaike information criterion (AIC; Akaike, 1981). The AIC for a model with k parameters and maximum likelihood L is

$$a = 2k - 2 \ln(L). \quad (5.9)$$

The model with the minimum value of a is considered the “best” because it is a sufficient representation of the data without overfitting. For each fitted model with Gaussians fixed at the center, we fit up to $k = 4N_{\text{comp}} - 1$ parameters (because the flux f of one component is fixed). The fit results for $3.7 \mu\text{m}$ with $N_{\text{comp}} = 1$ and $4.7 \mu\text{m}$ with $N_{\text{comp}} \in \{1, 2, 3\}$ are given in Table 5.6 and the best fitting models are shown in Fig. 5.11.

In the L -band, a true single-Gaussian model does not perform well. It becomes large in order to produce low visibilities at short baselines, but then the long-baseline visibilities are far too low. An augmented single-Gaussian model was then fit, wherein a second, large component was added. This second component has fixed size (128×128 mas) and orientation, but its flux is allowed to vary. This model (called 1+) gives a marked improvement in AIC over the single-Gaussian model with only one additional parameter. It results in a marginally extended source with FWHM 7.1×2.8 mas. A

Table 5.6: Fitted Gaussian Parameters for the *LM*-bands

λ [μm]	N_{comp} -	Θ [mas]	r -	ϕ [deg]	f -	$\ln L$ -
3.7 ± 0.1	1	$16.3^{+0.8}_{-0.6}$	$0.63^{+0.13}_{-0.05}$	$42.5^{+7.4}_{-8.2}$	1	257
3.7 ± 0.1	1+	$7.1^{+0.2}_{-0.4}$	$0.4^{+0.1}_{-0.1}$	$4.6^{+1.1}_{-0.6}$	1	267
		128	1	0	$0.07^{+0.01}_{-0.01}$	
3.7 ± 0.1	2	$6.5^{+0.5}_{-1.7}$	$0.3^{+0.2}_{-0.1}$	$3.3^{+8.7}_{-1.2}$	1	270
		$18.9^{+1.2}_{-0.6}$	$0.6^{+0.2}_{-0.1}$	$44.0^{+6.4}_{-8.2}$	$0.78^{+0.15}_{-0.32}$	
4.7 ± 0.1	1	$25.7^{+2.4}_{-1.2}$	$0.2^{+0.1}_{-0.1}$	$58.4^{2.5}_{-2.5}$	1	247
4.7 ± 0.1	2	$30.2^{+0.6}_{-1.3}$	$0.14^{+0.10}_{-0.07}$	$57.6^{+2.1}_{-3.1}$	1	262
		$5.4^{+0.6}_{-1.3}$	$0.4^{+0.3}_{-0.2}$	$12.8^{+10.7}_{-4.1}$	$0.6^{+0.3}_{-0.2}$	
4.7 ± 0.1	3	$28.4^{+2.2}_{-1.5}$	$0.13^{+0.11}_{-0.09}$	$52.4^{+3.6}_{-1.9}$	1	263
		$3.7^{+1.2}_{-0.3}$	$0.7^{+0.2}_{-0.3}$	$79.5^{+5.3}_{-4.1}$	$0.8^{+0.1}_{-0.4}$	
		$48.8^{+2.4}_{-1.4}$	$0.2^{+0.4}_{-0.2}$	$84.4^{+2.1}_{-5.4}$	$0.1^{+0.2}_{-0.1}$	

two-component Gaussian model produces a similarly extended source (6.5×2.0 mas) with a much less-extended second component. The two-component model produces the same AIC value as the 1+ model; for the rest of this work, we use the simpler 1+ model.

In the *M*-band, all three modeling results include in a disk-like component with a similar size and orientation; this disk-like component reproduces the image reconstruction’s morphology. In the two- and three-component fits, a point-like source is introduced on top of the disk. In the three-component model, a diffuse, roughly polar extended source is added. While its orientation is suggestive (based on the *N*-band polar emission), this additional component is disfavored by the AIC. If there is a signature of the polar dust in the *M*-band, it is at low significance, and more observations would be necessary to confirm it. The model with two Gaussian components is preferred at $4.7 \mu\text{m}$, and it will be used for the rest of this work. We note, however, that the flux in the central aperture (see the following section) differs by only 1% between the one- and two-component models. The selected model is marginally resolved with a width of 4.2 mas (the $4.7 \mu\text{m}$ resolution is 3.7 mas) and has a major axis with $\text{FWHM} = 30.2^{2.4}_{-1.2}$ mas and $\text{PA} = 57.6^{2.1}_{-3.1}$. The increased size of the disk in the modeling when compared to the images likely compensates for large scale flux which is allowed in the imaging. This scenario is supported by the observation that the PA of the disk component decreases with an increasing number of fitted components.

4.3 Measuring Component Temperatures

In §3.3, circular apertures with diameter 23.4 mas were used to extract the flux from each *N*-band image at a number of locations. We make the assumption here that the *LM* models can be astrometrically matched with the *N*-band images by aligning the photocenters. This was done for NGC 1068 in G3mez Rosas et al. (2022), and the cross-correlation matching done in the *N*-band was in the end equivalent to photocenter matching. Therefore, we use the same aperture diameter and distribution as §3.3.

In the *L*-band though there is a point source within the central aperture, D0, much of the flux is contained in the background component. This is true in both the imaging and in the 1+ model. Accordingly, only 86% and 37% of the total *L*-band flux is contained in the aperture D0 for the model and image, respectively. In the *M*-band, the majority of the flux is found in the central aperture with only a minuscule amount falling in the

Table 5.7: *L*- and *M*-band Extracted Fluxes

Aperture ^{<i>b</i>}	Image	Reconstruction	Gaussian	Modeling ^{<i>a</i>}
	$F_{3.7\mu\text{m}}$ [mJy]	$F_{4.7\mu\text{m}}$ [mJy]	$F_{3.7\mu\text{m}}$ [mJy]	$F_{4.7\mu\text{m}}$ [mJy]
D0	3.0 ± 1.7	13.2 ± 3.5	7.0 ± 1.6	28.5 ± 7.8
D _S 40	$0.4^{+0.6}_{-0.4}$	3.5 ± 1.1	≤ 2.2	3.5 ± 1.1
D _N 40	$0.4^{+0.6}_{-0.4}$	3.8 ± 1.1	≤ 2.2	3.5 ± 1.1

Error estimates contain the contribution from the total flux uncertainty as well as the model uncertainties. Image reconstruction fluxes are typically lower than the Gaussian models because they include background flux and artifacts which both take away flux from the primary components.

a: Fluxes come from the two-component model, but both the one-component and two-component Gaussian models give similar (within 1%) extracted fluxes. *b*: Apertures from the *N*-band analysis that aren't listed can be considered to be upper limits ($\leq 2\sigma_{\text{im}}$), with $F_{3.7\mu\text{m}} \leq 1.2$ mJy and $F_{4.7\mu\text{m}} \leq 2.2$ mJy.

disk apertures D_S40 and D_N40. In the imaging and modeling of neither band is there significantly measured flux in the polar direction. For these regions we present an upper limit from the “sky” background $F_{\text{upperlim.}} \leq 2\sigma_{\text{sky}}$ in the image reconstructions. All apertures from the *N*-band analysis that aren't listed in Table 5.7 are considered to have only upper limits: $F_{3.7\mu\text{m}} \leq 1.2$ mJy and $F_{4.7\mu\text{m}} \leq 2.2$ mJy.

The Gaussian models and image reconstructions give slightly different morphologies, and therefore yield different flux measurements. We consider both sets of results independently. The 3.7 and 4.7 μm fluxes for the apertures from each imaging method are given in Table 5.7. The uncertainties in each case come from both the image/model uncertainty (each is described above) and the uncertainty on the total correlated flux (at 30 m). The total flux uncertainty (~ 5 mJy at 4.7 μm) dominates in both imaging and modeling.

We fit a two-blackbody (BB) curve with absorption to each aperture-extracted spectrum with the form

$$I(\lambda, T, A_v) = \sum_{i=1}^2 \eta_i BB_\nu(\lambda, T_i) e^{-A_{v,i}/1.09\tau(\lambda)/\tau_v}, \quad (5.10)$$

where η is an absolute flux scaling due to the filling factor of the dust in the aperture, $\tau(\lambda)/\tau_v = \kappa(\lambda)/\kappa_v$, and we use the standard interstellar medium $\kappa(\lambda)$ profile from Schartmann et al. (2005) which is based on the standard interstellar medium profile of Mathis et al. (1977).

Fitting of T , η , and A_v is done in two iterations using Markov Chain Monte Carlo sampling with the package `emcee` (Foreman-Mackey et al., 2013), similar to the approach in §4.2.2. Final values in each iteration are the median of the marginalized posterior probability distribution. The 16th and 84th percentiles of the resulting temperature and extinction distributions are used as the 1σ fit uncertainties.

For the first BB component, we use uniform prior probability distributions with $T_1 \in (100, 500]$ K, $\eta_1 = 1$, and $A_{v,1} \in [20, 37]$ mag. These priors are based on the *N*-band fit results, particularly $A_v = 28.5^{+8.5}_{-7.7}$ mag for D0 (see §3.3). The second component is forced to be strictly hotter and smaller than the first component, resulting in the uniform priors $T_2 \in (500, 1500]$ K, $\eta_2 \in (10^{-3}, 0.1]$, and $A_{v,2} \in (0, 700]$ mag. For the central aperture, which should cover the sublimation zone, another set of priors is also used. They come from the assumption that dust is the sublimation temperature is indeed present but can be heavily obscured. As the sublimation radius for dust at 1500

Table 5.8: Fit Results for the Image Reconstructions

Aperture	N_{BB}	T_1 [K]	$A_{V,1}$ [mag]	η_2 -	T_2 [K]	$A_{V,2}$ [mag]
D0	1	343_{-5}^{+6}	$22.7_{-1.4}^{+2.4}$	-	-	-
D0	2	343_{-5}^{+6}	$22.6_{-1.2}^{+2.2}$	$0.05_{-0.03}^{+0.04}$	831_{-306}^{+407}	479_{-179}^{+146}
D0	2*	343_{-5}^{+6}	$22.6_{-1.3}^{+2.4}$	$0.03_{-0.02}^{+0.02}$	1500*	520_{-135}^{+125}
D_N40	1	266_{-16}^{+13}	$32.5_{-5.5}^{+2.7}$	-	-	-
D_N40	2	252_{-16}^{+15}	$29.9_{-5.7}^{+4.5}$	$0.06_{-0.03}^{+0.03}$	1042_{-371}^{+314}	340_{-91}^{+74}
D_S40	1	286_{-11}^{+8}	$33.0_{-4.3}^{+2.3}$	-	-	-
D_S40	2	281_{-13}^{+10}	$32.3_{-5.3}^{+3.0}$	$0.05_{-0.03}^{+0.04}$	840_{-329}^{+434}	403_{-165}^{+184}

K is given in Barvainis (1987) as

$$r_{\text{sub}} = 1.3L_{UV,46}^{0.5}T_{1500}^{-2.8} \text{ pc}, \quad (5.11)$$

and Moorwood et al. (1996) report $L_{UV} = 5 \times 10^9 L_{\odot}$, we estimate $r_{\text{sub}} = 0.05 \text{ pc} = 2.8 \text{ mas}$ at 4.2 Mpc distance. This gives an upper limit on $\eta_2 \leq (2.8/11.7)^2 = 0.06$, where 11.7 mas is the aperture radius. Therefore, we define the priors of the second component: $T_2 = 1500 \text{ K}$, $\eta_2 \leq 0.06$, and $A_{V,2} \in (0, 700] \text{ mag}$. This fit gives a rough estimate on the minimum amount of extinction necessary to hide dust at the sublimation temperature. The recovered temperatures for the one- and two-component blackbody fits to the image reconstruction fluxes and to the Gaussian model fluxes are given in Tables 5.8 and 5.9, respectively. The fitted SEDs for the Gaussian models using both one and two blackbodies are shown in Fig. 5.12. The nuclear K -band flux from Burtscher et al. (2015) for Circinus is shown in Fig. 5.12, and it serves as an upper limit on the near-infrared flux; all extrapolated blackbody flux values are far below this limit.

For the aperture D0, the image reconstructions’ fluxes result in lower fitted temperatures than the Gaussians’. The fitted T_1 values in both the images and the Gaussian models, however, are consistent within the uncertainties to the temperature inferred from the N -band alone ($367_{-26}^{+30} \text{ K}$). Due to image artifacts, we favor the temperatures from the Gaussian model. The fitted “cool” component temperature T_1 has essentially the same value both with and without the additional hot blackbody. Furthermore, in neither the images nor the models is an additional, hot component necessary. Large extinction values ($\sim 450 \text{ mag}$) are preferred, and the hot component makes up only 0.2% of the $3.7 \mu\text{m}$ flux and 1.3% of the $4.7 \mu\text{m}$ flux. In fact, looking at the posterior probably distributions for all apertures, only second component extinction values $\gtrsim 250 \text{ mag}$ are allowed in any of the fits.

For the disk apertures, D_N40 and D_S40 , the results are consistent from both the modeling and image reconstructions. The results, furthermore, are consistent with the N -band fitted temperatures and extinctions. A second component is once again disfavored due to the large fitted extinction values and the fact that the fit is not markedly improved with the addition of this component.

5 Discussion and Interpretation

MIR interferometry of Circinus has revealed several major components of the thermal dust: a disk-like central emission, large-scale polar emission, and a central point source along the disk. Image reconstruction has recovered these features in unprecedented detail and brought forward new substructures. The morphological features are labeled in Fig. 5.6. In the following, we examine each of these features separately. After the

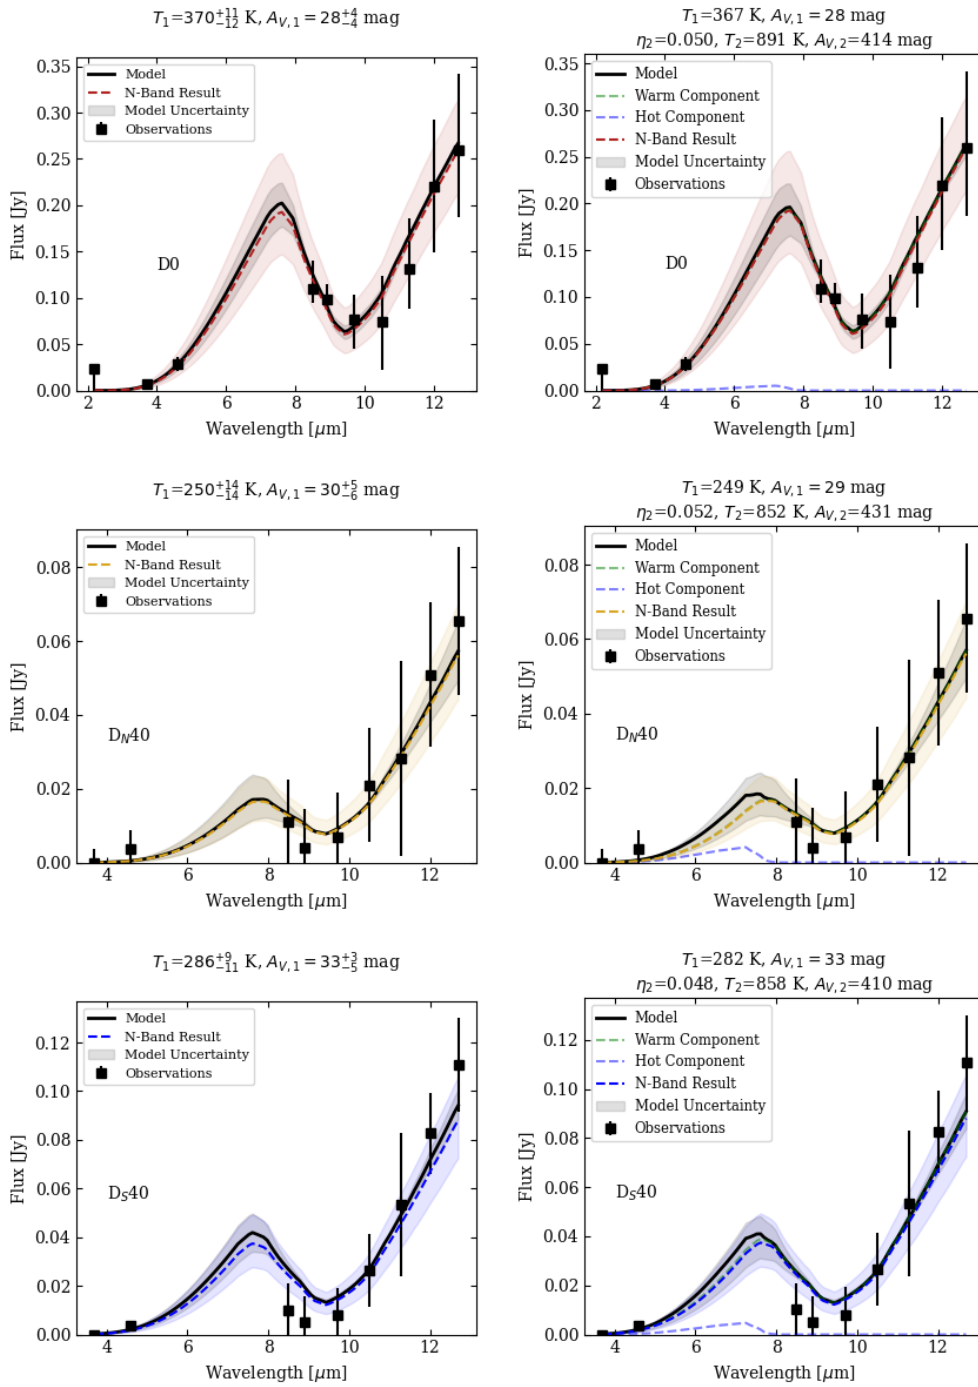


Fig. 5.12: One-blackbody (*left*) and two-blackbody (*right*) fits for the aperture-extracted Circinus *LMN* fluxes. The colors are the same as in Fig. 5.6, with D0 in red (*top row*), DN40 in yellow (*middle row*), and DS40 in red (*bottom row*). There is little discernible improvement in fit quality with the addition of a second component, and this second component is in all cases highly extinguished. The fits using the *N*-band data alone are included for comparison. In aperture D0 the *K*-band measurement from [Burtscher et al. \(2015\)](#) is included as an upper limit for the near-infrared flux.

Table 5.9: Fit Results for the Gaussian Models

Aperture	N_{BB}	T_1 [K]	$A_{V,1}$ [mag]	η_2	T_2 [K]	$A_{V,2}$ [mag] -
D0	1	370^{+11}_{-12}	$28.4^{+3.9}_{-3.7}$	-	-	-
D0	2	367^{+13}_{-15}	$27.8^{+4.2}_{-4.2}$	$0.05^{+0.03}_{-0.03}$	891^{+399}_{-354}	413^{+197}_{-191}
D0	2*	367^{+13}_{-15}	$27.7^{+4.0}_{-4.2}$	$0.03^{+0.02}_{-0.02}$	1500*	459^{+164}_{-178}
D_N40	1	250^{+15}_{-15}	$29.7^{+4.7}_{-5.7}$	-	-	-
D_N40	2	250^{+13}_{-14}	$30.0^{+4.3}_{-5.6}$	$0.05^{+0.03}_{-0.03}$	880^{+416}_{-335}	433^{+183}_{-178}
D_S40	1	301^{+10}_{-13}	$42.7^{+5.6}_{-7.0}$	-	-	-
D_S40	2	281^{+10}_{-13}	$32.1^{+3.1}_{-5.3}$	$0.05^{+0.04}_{-0.03}$	904^{+401}_{-362}	421^{+173}_{-148}

discussion of the individual features, from the smallest scales to the largest, we discuss the overall morphology.

The orientations of the central structures in Circinus are compared to those of the optical ionization cone and of the warped maser emission in the center. The well-studied optical ionization cone has a central axis along $\text{PA}_{\text{opt}} = -52^\circ$ and a projected half-opening angle between 36° and 41° (see, e.g., Marconi et al., 1994; Maiolino et al., 2000; Wilson et al., 2000; Fischer et al., 2013; Mingozzi et al., 2019). The observed ionized emission is not symmetric; it only extends toward the NW with no optical counterpart seen to the south, though a southern counterpart can be seen in the NIR (Prieto et al., 2004). The ionization cone is thought to coincide with an outflow of dense material, driven by radiation pressure and fed by a gaseous nuclear bar (Maiolino et al., 2000; Packham et al., 2005). Notably, the O[III] and $\text{H}\alpha$ emission in the ionization cone is much brighter along its southern edge ($\text{PA} \sim -90^\circ$). The ionization cone is observed out to ~ 40 pc from the nucleus (Wilson et al., 2000).

The warped H_2O maser disk was separated by Greenhill et al. (2003) into 3 components: the blueshifted emission ($0.11 < r \lesssim 0.4$ pc; $\text{PA}_{\text{maser,blue}} = 56 \pm 6^\circ$), the central emission ($0 < r < 0.11 \pm 0.02$ pc; $\text{PA}_{\text{maser,central}} = 29 \pm 3$ deg), and the redshifted emission ($0.11 < r \lesssim 0.4$ pc; $\text{PA}_{\text{maser,red}} = 56 \pm 6^\circ$). The central maser emission, which may trace the orientation of the accretion disk and the dense material around it, is nearly perpendicular to the radio jet axis ($\text{PA}_{\text{jet}} = 115$ and $295 \pm 5^\circ$; Elmouttie et al., 1998), which is not aligned with the central axis of the optical ionization cone. These orientation markers are shown in Fig. 5.5 for comparison to the MATISSE images.

5.1 Central Point Source

We find a central, bright component unresolved at all wavelengths (≤ 2.9 mas at $3.6 \mu\text{m}$, ≤ 6.7 mas at $8.5 \mu\text{m}$, and ≤ 10.1 mas at $12.7 \mu\text{m}$). It is in the same position relative to the image photocenter in each reconstruction, and is therefore likely the same unresolved object present throughout. This point source is consistently found ≈ 10 mas to the NE of the photocenter of the disk. Our central aperture, D0 (Fig. 5.6, Table 5.4), is centered on this point source and the extracted fluxes are brighter than the surrounding features by more than a factor of 2. We find that the fitted blackbody is relatively hot, 367^{+30}_{-26} K. While this source is well-fit by a single blackbody, we note that this is difficult to motivate physically and only serves as an estimate.

In the L -band, both image reconstruction and Gaussian modeling reveal a point-like source in the L -band. It is marginally resolved in the modeling, with $\text{FWHM} = 7.1^{+0.2}_{-0.4}$ mas, but is a true point source in the image reconstructions. Fortunately, both sizes give similar results in the SED analysis. In the M -band, the point source is present in Gaussian modeling but not in the image reconstruction. We identify the LM point

source with the point source found in the N -band based on 1) its position on the disk and 2) extrapolation of the N -band fitted temperature (367 K) to the L -band is fully consistent with the recovered flux.

These results are similar to those of T14, who found a central unresolved component lying along the disk-like component. Their point source was shifted 14 mas to the NE of the disk-center, similar to the 10 mas which we find. T14 measured the temperature of this component to be 317 ± 22 K, which is $\sim 2\sigma$ lower than our measured temperature. The temperature difference is perhaps a result of the overlapping contributions of the three Gaussian components in T14, while we fit an isolated mean temperature at each extraction location. Nonetheless, no directly visible hot ($\gtrsim 900$ K) dust is found by either T14 or this work.

The central aperture is almost certainly probing a column of much cooler dust along the line of sight and it may indeed reach dust at the sublimation temperature. A large range of spatial scales and temperatures are being merged into one aperture because of projection effects. It is thus difficult to draw any strong conclusions about the temperature in this feature without the LM -bands which should be more sensitive to hot dust. The L - and M -band fluxes measured using VLT/ISAAC by [Isbell et al. \(2021\)](#) represent the AGN flux within 630 mas, and are certainly an upper limit on the LM flux within the central aperture. However, if we perform a two-blackbody fit to the ISAAC LM measurements in addition to our central aperture fluxes, we see that a very compact and extinguished 1500 K blackbody in addition to a larger 310 K blackbody fit the data very well. So, it is possible that the central aperture contains dust at the sublimation temperature, but the LM MATISSE data are crucial for probing this region due to their higher spatial resolution and ability to sample hotter emission.

5.1.1 Where is the hot dust?

The L - and M -band structures are fully consistent with the “warm” dust observed in the N -band, so the question arises: where is the hot dust? The typical picture of the circumnuclear dust includes dust at or near the sublimation temperature (~ 1500 K for silicate dust; see [Hönig, 2019](#), for a review). Yet these observations give no trace of dust at this temperature. It is either completely obscured or simply not there.

In the two-blackbody models, high temperatures were always accompanied by very high absorption ($A_V \gtrsim 250$ mag). In several models we made the assumption that the sublimation temperature dust was there at the scale of the sublimation radius. The fitting results indicate that the emitting region must be small (i.e. the sublimation zone is geometrically thin and edge on) and/or the extinction must be very high. The geometrically thin disk found in the M - and N -bands must then be relatively dense, and could act as the obscuring structure in the AGN unification scheme. Using the X-ray column density $N_H = 10^{24}$ cm $^{-2}$ from [Matt et al. \(1999\)](#) we can estimate the amount of extinction to the accretion disk, which serves as an upper limit on the extinction possible to the sublimation zone. The correlation given by [Predehl & Schmitt \(1995\)](#) relates the column density to A_V :

$$N_H(\text{cm}^{-2}) = (1.79 \pm 0.03) \times 10^{21} A_V. \quad (5.12)$$

For Circinus we then obtain $A_V = 5586 \pm 6$ mag, which is more than $10\times$ the extinction necessary to hide the sublimation zone. Of course, this is an upper limit, but we cannot at this time rule out that the sublimation zone is obscured by a dense, geometrically thin disk. Radiation-hydrodynamical models tend to produce optically thick, geometrically thin disks (see e.g., [Wada, 2012](#); [Williamson et al., 2019](#)), and these observations could constrain the scale and density of the disk.

The second possibility is that there is simply no dust at the sublimation temperature in Circinus. This circumstance is speculative but exciting, as it could indicate that Circinus is leaving an active phase and is no longer actively heating the dust. Circinus is not unusual compared to other Seyferts; its Eddington ratio is 0.2 with SMBH mass $10^{6.2}M_{\odot}$ (Greenhill et al., 2003), and its broad line region was observed in polarized light (Oliva et al., 1998). There is therefore no indication that it should be ending its active phase, but nonetheless, there is no trace of the expected sublimation zone. It could also indicate that the dust stops farther out than the sublimation radius, though this would need to be explored further.

5.2 Central disk

We find a thin disk-like structure along $\approx 30^{\circ}$ in the N -band. It is present in all wavelength channels, but most prominent at longer wavelengths. In the continuum image, this disk is almost 1.9 pc in diameter and is unresolved in width. The extent of the disk is set by the 5σ contours at 12 and $12.7 \mu\text{m}$. Considering the dirty-beam (Appendix F), we expect artifacts in the form of secondary lobes at ~ 100 mas from the center along the disk PA, which indeed manifest themselves as low surface-brightness features near the edges of the images. Nonetheless, the central part of the disk in our images has a high flux density and is robustly detected at $S/N \geq 5$.

In the M -band, we find a structure resembling the disk-like component in the N -band images. The flux and extent of the M -band disk (both in imaging and modeling) agree well with a simple extrapolation of the N -band flux and fitted temperatures. While the position angle of the M -band structure differs from the N -band results (46° in the M -band image vs $\approx 30^{\circ}$ in the N -band), this is likely due to the effectively sparse uv -coverage from the low S/N correlated fluxes on intermediate to long baselines. The LMN SED fitting strongly supports the interpretation that the M -band disk is the same ~ 300 K structure observed at a different wavelength. The disk width is marginally resolved ($4.23 \text{ mas} = 0.08 \text{ pc}$). The scale of the disk –and the emerging three-dimensional picture of the structure from the dust at different wavelengths and temperatures– places constraints on the circumnuclear dust structure in future modeling.

Evidence that the dust in this disk is relatively dense comes from the blackbody fits performed on the “disk apertures.” Here we see that the temperature falls quickly as one moves farther from the photo center; indeed, the lowest fitted temperatures in the image occur in the disk at a projected distance of only 0.7 pc from the center. Taken together, the disk apertures (D_0 , D_{S40} , D_{N40} , D_{S77} , D_{N72}) exhibit a much steeper radial temperature gradient than apertures in the polar direction. The temperature profile of the disk is shown in both Fig. 5.8 and in the images themselves, as the disk becomes much less prominent at short wavelengths, indicating that the dust is relatively cool and the emission drops off significantly below $9 \mu\text{m}$. The steep temperature gradient is possibly indicative of a dense environment wherein only the innermost dust has a direct view toward the accretion disk, and the outer clouds are heated only through re-radiation and photon scattering (e.g., Krolik, 2007).

The disk component places constraints on the inclination of the system. Assuming that the disk is both thin and axisymmetric with diameter 1.9 pc, the fact that we do not resolve the width of the disk ($\leq 9.5 \text{ mas} = 0.18 \text{ pc}$ at $12 \mu\text{m}$) indicates an inclination $i \gtrsim 83^{\circ}$. This is in agreement with the best disk+hyp model with $i \sim 85^{\circ}$ matched to the closure phases and squared visibilities. This estimate can be considered a lower limit, as a more-realistic “puffed-up” disk would be thicker. $\gtrsim 83^{\circ}$ is closer to edge-on than on the galactic scale ($\sim 65^{\circ}$ Freeman et al., 1977; Elmouttie et al., 1998), the ALMA CO(3-2) tilted ring estimation ($i \gtrsim 70^{\circ}$ Izumi et al., 2018), and the T14 estimate for the MIR disk ($i > 75^{\circ}$). However, Izumi et al. (2018) note that from 10 pc inward the

inclination seems to increase, eventually reaching $i \sim 90^\circ$ for the warped H₂O maser disk (Greenhill et al., 2003). The relatively dense dust of the MIR disk is consistent with the above and may lie in the same plane as the maser emission. The above assumes that the disk-like emission is the edge of the disk, rather than from reflected light on the top of a disk-like structure (see e.g., T14; Stalevski et al., 2019). We make this assumption because we observe no absorbing band on either side of the disk-emission.

The disk is aligned very well with the inner position angle of the H₂O maser emission ($29 \pm 3^\circ$; Greenhill et al., 2003) as well as with the compact nuclear disk (CND) at 10s of pc found in ALMA CO(3-2) and [CI](1-0) ($32 \pm 1.9^\circ$; Izumi et al., 2018). The entirety of the warped maser disk, moreover, fits within our ≤ 9.5 mas = 0.18 pc thick dust disk. It is for this reason that we place the maser emission in the center of our disk; we do not have absolute astrometry from MATISSE, and so we must base the correspondence on the coincidence of PA and scale. Through Gaussian modeling, T14 also found a thin disk oriented along $46 \pm 3^\circ$ and with a FWHM of 1.1 ± 0.3 pc. The size of the disk in the T14 modeling is similar to what we measure. The T14 disk orientation differs slightly from that of our imaged disk, but they *a*) used differential phases rather than closure phases in their modeling; and *b*) used Gaussian modeling which simplifies the structure and may combine components. In §3.2.4 we found that with the T14 *uv*-coverage, our image disk could be oriented along $\approx 40^\circ$.

This dense disk of dust may play the role of the classical “torus”, obscuring a direct view toward the BLR. However, we find two competing phenomena. First, we see in the central aperture that hot dust is present, and depending on the exact distribution of the dust in the *LM* bands, we may even have a direct line of sight to dust at/near the sublimation temperature. This is, however, somewhat at odds with the steep temperature gradient we see across the disk. The thin disk must somehow be dense enough to shield some or most of the dust from directly seeing the sublimation zone or the central engine, but clumpy or low-density enough that we can see evidence of hot dust at or near the sublimation temperature. Authors such as Kishimoto et al. (2011) and Hönig et al. (2012) hypothesize that a “puffed-up” inner region (a few sublimation radii in size) may act as the classical obscuring torus.

5.3 Polar extension

We find a large-scale structure oriented in the same direction in all *N*-band wavelength channels. This structure is referred to as a “polar extension” because its primary axis lies perpendicular to the AGN orientation and along the radio jet. The polar extension in our imaging is a large ($\sim 4 \times 1.5$ pc) structure made up of warm (> 200 K) dust with major axis along $\approx -60^\circ$. This larger envelope contains significant substructure: most prominently enhanced brightness directly E and W of the disk center. The polar emission exhibits “patchiness” at a significance $\geq 3\sigma$ on scales similar to the beam size, most prominently in the 12.0 and 12.7 μm images. Patchiness in the image could arise from clumpy dust emission, though it is unlikely that we resolve individual clumps at this scale (10 mas = 0.19 pc). Nonetheless, these images provide direct evidence that the polar emission is not a smooth, continuous structure.

We find that the substructures of the polar emission exhibit spatial variation in temperature. At a similar projected distance, apertures E65 and W65 are marginally hotter than W64 and E64 (~ 270 K vs ~ 230 K). Additionally, the dust comprising the polar extended regions remains warm (~ 200 K) out to a projected distance of ~ 1.5 pc from the center of the structure. This behavior is significantly different than the dust temperature gradients along the disk, indicating differences in environment and density. The dust in the polar direction is likely less dense and/or more clumpy, as high temperatures at large distances require a relatively unobscured line of sight to

the accretion disk. As shown in Fig. 5.8, the temperatures in the polar emission are entirely consistent with predictions from radiative transfer modeling of clumpy media (e.g., Schartmann et al., 2008; Stalevski et al., 2019), however only the latter reproduce the interferometric observables. At much lower resolution, the MIR SEDs of nearby AGN have shown that clumpy formalism is necessary to reproduce the relatively “blue” spectra indicative of an abundance of warm dust (e.g., Nenkova et al., 2008a; Stalevski et al., 2016; Hönig & Kishimoto, 2017). At the parsec scale, our results support the predictions of clumpy models.

In the LM -bands we find no strong evidence of polar dust. While the M -band three-Gaussian model fit to the squared visibilities produced an extended component with $PA \sim 90^\circ$, this result is not robust and the AIC indicates it could be the result of overfitting noisy data. The lack of polar dust is not entirely surprising, given the temperatures of the so-called E-W flux enhancements seen in the N -band. In the M -band, the components called E65 and W65 would have only ~ 1.4 mJy. With image flux uncertainties of 1.1 mJy, this flux would not be robustly detected. The rest of the polar dust is even colder, and certainly falls below the detection threshold. Either further observations or improvements in the calibration of the correlated fluxes will be necessary to identify polar dust at these wavelengths. Both the images and modeling of the LM -bands required “resolved-out” flux to be present in order to reproduce the low squared visibilities. This diffuse L -band flux may simply be the “blue” tail of the ~ 250 K polar dust. Assuming that the entire 4×1.5 pc N -band polar dust structure has emits at the average temperature of 250 K, then up to 200 mJy of diffuse flux could be present at $3.7 \mu\text{m}$ on large scales. It would largely be resolved out at these wavelengths (at $3.7 \mu\text{m}$ the primary beam is $46.5 \text{ mas} = 0.9 \text{ pc}$); nonetheless the ≈ 5 mJy extended flux can be explained by the large-scale polar N -band structure alone. The rest of this subsection therefore focuses on the N -band.

5.3.1 E-W flux enhancements

The morphology we recover is in accordance with previous single-dish N -band estimates of the polar dust, and with the MIDI results of T14. The primary position angle of the polar extension was estimated from VLT/VISIR observations to be $-80 \pm 10^\circ$ (Asmus et al., 2016). Similarly, the modeling done by T14 resulted in a 93_{-12}^{+6} mas FWHM (≈ 2 pc) Gaussian component with $T = 304_{-8}^{+62}$ K and with a major axis along -73 ± 8 deg. Both the single-dish PA and that of the large Gaussian component in T14 are directed more closely to E-W orientation than our imaging suggests. This is likely explained by the lack of resolution and the simplicity of the Gaussian modeling; the large structure in our imaging shows significant non-uniformities. Namely, enhancements in flux directly to the E and to the W of the image photocenter. If one considers a flux-weighted mean of the polar emission in our imaging, it would certainly be more similar to the $PA = 75 \pm 8^\circ$ as seen in T14. Indeed, the analysis by Stalevski et al. (2017, 2019) claims that the T14 large component is a simplified representation of an edge-brightened outflow cone, and they use this hypothesis to explain the discrepancy between the orientation of their polar outflow and the true pole of Circinus.

We put forward two possible explanations for the bright E-W substructure of the polar emission. The first is that the accretion disk in Circinus is tilted with respect to the central dust structures. If one considers that the central maser emission traces the orientation of the accretion disk (supported by the agreement with the radio jet position angle, assuming the jets originate in the central region), then one can relatively simply explain the asymmetric illumination of the polar extension. We show in Fig. 5.5 a line tracing the radio jet orientation. This line touches both the E and W flux-enhanced regions of the image. Due to the anisotropic nature of accretion disk emission,

any dust some angle θ away from the “face” of the accretion disk is illuminated by a factor $\propto \cos\theta(2\cos\theta + 1)$ less than the dust which does see the “face” (Netzer, 1987). The features we observe end more abruptly than this function suggests, but this could be due to patchiness or clumpiness of the dust. The idea of an accretion disk tilted with respect to the large-scale structures in Circinus is not new. Greenhill et al. (2003) suggests that the orientation of the accretion disk should only be “weakly coupled via gravity to the surrounding large-scale dynamical structures” because the central engine has a sphere of influence with a radius of only a few pc (Curran et al., 1998). Using VISIR images, MIDI observations, and the SED of Circinus, T14 as well as Stalevski et al. (2017, 2019) hypothesized that a warped or tilted accretion disk (as described by e.g., Petterson, 1977; Nayakshin, 2005) was required to asymmetrically illuminate the polar dust in their modeling. Hydrodynamic modeling of the central structures by Wada (2012) predicts that symmetric radiation-driven outflow cones should form perpendicular to the accretion disk. So while our observations suggest that the illumination of the polar dust is asymmetric –possibly from a tilted accretion disk– we cannot at this time explain why or how such a tilt occurred. The second possibility is that there is simply more material along the E-W direction; indeed Greenhill et al. (2003) speculated that the warped accretion disk could channel material in the nuclear outflow. This hypothesis is in better agreement with the Wada (2012) modeling, as in this case the polar outflows would be symmetric w.r.t. the accretion disk. The higher temperatures of the E-W flux-enhancements with respect to the apertures at the same projected distance argue in favor of the direct-illumination hypothesis. An overdensity of material should exhibit cooler temperatures due to dust self-shielding (as seen in the disk). This is merely a qualitative agreement, and in order to distinguish between these two hypotheses, detailed modeling of the formation of the outflow cones in the presence of a warped accretion disk will be crucial.

5.3.2 Connection to larger scales

It is clear in the MATISSE imaging that the majority of thermal dust emission in the center of Circinus comes from the polar extension, but its full extent is poorly constrained. In our imaging, any structure larger than those probed by the shortest baselines is resolved out; this means for imaging using the MIDI AT baselines we are not sensitive to structures larger than 688 mas at $12\ \mu\text{m}$. This is strictly an upper limit, however, and in the image reconstruction process we *a*) limit the FOV to 600 mas, and *b*) apply an object mask with a radius 160 mas. The object masking heavily suppresses any structure which falls outside of the specified radius. We can, nonetheless, confidently state that there is *N*-band emission out to ~ 1.5 pc from the center to both the NW and SE, and that the emission shows a flux enhancement to the E and W of the image center.

The VISIR-SAM data were fit in the image plane with a Gaussian having FWHM 3.3×2.2 pc and major axis along $PA = 72^\circ$. This is larger than either the MATISSE images or the MIDI modeling (with FWHM = 2 pc), indicating that the MATISSE images do not capture the true extent of the structure. The position angle of the SAM data matches the T14 result, though it is likely also flux-biased toward the South due to the E-W flux enhancements. Continuing to lower resolution, the *N*-band VLT/VISIR images in Asmus et al. (2014, 2016) show that in Circinus roughly 60% of the flux is extended farther than 5.24 pc and at $PA = 100 \pm 10^\circ$. It is clear that the polar structures we see in our images extend continuously outward past 5 pc. Future *N*-band observations with the MATISSE ATs will yield the first closure phase measurements of the $\gtrsim 1$ pc dust, further improving our imaging capabilities beyond the MIDI data.

5.4 Overall morphology

We present the first model-independent images of the circumnuclear dust in Circinus. The recovered combination of a geometrically thin disk and large-scale polar emission supports previous MIR interferometric findings, but newly imaged substructures hint at complexity unmatched in existing modeling. In particular, we find that the disk is simultaneously dense and yet allows emission from hot dust to radiate through; we find that an unresolved component lies 10 mas NE of the photocenter along the disk; and we find significant flux enhancements in the polar emission E and W of the disk-center.

The size of circumnuclear dust structures has been shown to vary with AGN luminosity (e.g., Kishimoto et al., 2011; Burtscher et al., 2013). The scales measured herein of the circumnuclear structures in Circinus –namely a thin disk with diameter 1.9 pc and $\gtrsim 4$ pc polar emission– with $L_{\text{AGN}} = 6 \times 10^9 - 7 \times 10^{10} L_{\odot}$ (Arévalo et al., 2014; Ricci et al., 2015; Tristram et al., 2014; Oliva et al., 1999) place a constraint on the luminosity-dependent scaling of the dust structures in AGN. Leftley et al. (2019) showed that the ratio of extended flux to unresolved flux increased with Eddington ratio (ϵ_{Edd}), claiming that this implied the presence of more dust in a radiation-driven wind for a higher ϵ_{Edd} . Circinus, with $\epsilon_{\text{Edd}} \sim 0.2$ (Greenhill et al., 2003), is dominated by polar dust emission. We measure the flux of the unresolved component to be $F_{\text{pt},12\mu\text{m}} = 0.77 \pm 0.04$ Jy, which is $6.3 \pm 0.3\%$ of the total flux at 12 μm . At 8.9 μm we measure $F_{\text{pt},8.9\mu\text{m}} = 0.39 \pm 0.01$ Jy, which is $9.0 \pm 0.2\%$ of the total flux. The fraction at 12 μm is significantly smaller than previously reported (20% and 10% at 12 μm in Leftley et al., 2019; López-Gonzaga et al., 2016, respectively), but they relied on simple two-Gaussian modeling of MIDI data.

Disk+wind radiative-transfer models (Hönig & Kishimoto, 2017; Stalevski et al., 2019) have recently been invoked to explain the polar emission found in a number of nearby AGN (e.g., Tristram et al., 2007; Burtscher et al., 2013; Tristram et al., 2014; López-Gonzaga et al., 2014; López-Gonzaga et al., 2016; Leftley et al., 2018). Fits to the NIR and MIR SEDs of nearby AGN have shown that disk+wind models provide the best match to the overall SED, reproducing the MIR flux through large-scale emission and NIR flux via reflected light from the accretion disk in the windy outflow (e.g., Martínez-Paredes et al., 2020; Isbell et al., 2021). The disk+wind morphology in the radiative-transfer models is supported by hydrodynamical and radiation-hydro modeling (Wada, 2012; Wada et al., 2016; Williamson et al., 2020; Venanzi et al., 2020), but has had few direct observational constraints. The images presented in this work, with a thin disk (1.9 pc \times ≤ 0.18 pc) and polar emission ($\sim 4 \times 1.5$ pc) perpendicular to it, resemble the disk+wind models only in broad strokes. Modifications to the disk+wind model in Stalevski et al. (2019) explain the E-W flux enhancements in the polar emission via a tilted accretion disk, but the dynamical stability of such a shift in radiation pressure remains untested. Hydrodynamical models produce structures symmetric about the accretion disk (Wada, 2012; Venanzi et al., 2020), so tilting with respect to the dusty structures may play a larger role. Whether this is specific to Circinus or a more general feature remains to be explored.

Only one other AGN has been imaged with MATISSE so far: NGC 1068. Imaging work by Gámez Rosas et al. (2022) has revealed a quite different circumnuclear dust morphology than we recover. In NGC 1068 at 12 μm , they find a disk-like structure ~ 2 pc in diameter with emission extending nearly perpendicular to it, similar to what we see with the disk and E-W polar flux enhancements. However, the NGC 1068 and Circinus morphologies differ significantly at other wavelengths. At 8.5 μm and in the *LM*-bands, the NGC 1068 emission is resolved into a ring-like structure with 720 K dust embedded within. We have shown that hot dust can make up the unresolved flux in Circinus, and the *LM* data can help clarify the situation, as they probe the $\gtrsim 500$ K dust

morphology. Finally, [Jaffe et al. \(2004\)](#); [López-Gonzaga et al. \(2014\)](#), and [Gómez Rosas et al. \(2022\)](#) showed that in NGC 1068, the standard ISM dust we use does not reproduce the observed SEDs. The effects of varying dust composition will be explored in future work. Circinus and NGC 1068 are laboratories in which to study the circumnuclear dust at extremely small physical scales, but the ongoing the MATISSE AGN Program aims toward a statistical understanding of the central dust scaling and relation to the SMBH.

6 Conclusions

In this work we present the first images of the circumnuclear dust in the Seyfert 2 galaxy Circinus. These images were reconstructed with IRBis using 150 correlated fluxes and 100 closure phases in the L -, M -, and N -bands from VLTI/MATISSE. Closure phase measurements of Circinus are reported here, and their novel inclusion in MATISSE observations makes imaging possible for the first time. The above results are largely in agreement with previous observations from MIDI ([Tristram et al., 2007, 2014](#)) and VISIR ([Asmus et al., 2014](#)). But our images, moreover, are model-independent and show new substructure which can be used to further constrain physical modeling of circumnuclear dust in AGN.

6.1 N -band Results

Through analysis of the interferometric observables and the images reconstructed in seven independent wavelength channels we

1. Show that correlated flux measurements on individual baselines have not changed over the last 17 years, implying that the underlying structures remain unchanged from the MIDI observations obtained between 2004 and 2011.
2. Find significant substructure in the circumnuclear dust. The circumnuclear dust can be separated into several components: central, unresolved flux; a thin disk 1.9 pc in diameter; polar emission ($\sim 4 \times 1.5$ pc) extending orthogonal to the disk and exhibiting patchiness; and flux enhancements E and W of the disk embedded within the polar dust.
3. Report that the polar dust makes up $\sim 60\%$ of the total flux, increasing toward longer wavelengths. The unresolved flux makes up $\lesssim 10\%$, increasing toward shorter wavelengths and further hinting at the presence of hot dust.
4. Measure SEDs in 13 apertures across the structures and fit temperature and extinction values to blackbodies in those apertures. We fit hotter dust temperatures ($T = 367_{-26}^{+30}$ K) in the central aperture along with warm dust ($T \gtrsim 200$ K) 1.5 pc from the center, indicating a clumpy circumnuclear medium. We clearly distinguish the radial temperature profiles of the disk and the polar extension: the disk shows a steeper temperature gradient indicating dense material; the polar emission shows a much flatter temperature profile with warm temperatures out to 2 pc from the center.
5. Recover a remarkably symmetric object, in terms of both flux and temperature distributions. We fit $A_V = 28.5_{-7.7}^{+8.5}$ mag, consistent with the galactic-scale value ($A_V = 28 \pm 7$ [Wilson et al., 2000](#)). We find no evidence of an absorption gradient across the field, in contrast to previous results (i.e., [Tristram et al., 2014](#)). Our new results indicate the presence of a foreground dust screen with very little local variation.

6. See that on large scales, the recovered morphology of the N -band dust in Circinus resembles the results of disk+wind modeling (e.g., [Wada et al., 2016](#); [Stalevski et al., 2019](#)), but new questions are raised because the subparsec dust is imaged here for the first time. We find that the temperature distribution is well-reproduced by the clumpy torus models of [Schartmann et al. \(2008\)](#) and [Stalevski et al. \(2019\)](#). The [Schartmann et al. \(2008\)](#) models do not, however, match the imaged morphology. The disk+hyp models better match the structure, but discrepancies are found in the central and disk apertures, indicating modifications to the disk component are necessary in the models. Using a suite of disk+hyp models based on [Stalevski et al. \(2019\)](#), we find that a large range of clump densities and disk filling-factors can match the data within the uncertainties of the images and interferometric observables.
7. Discover inhomogeneities in the polar dust emission: namely significant patchiness on scales of the resolution element; and flux enhancements directly to the E and W of the disk. The here-discovered patchiness is the first direct evidence that the polar dust is not a smooth, continuous structure but is rather clumpy. The E-W flux enhancement raises questions about the relation of the accretion disk to the larger dust structures.

6.2 LM -band Results

Following up the N -band analysis, we present the first-ever L - and M -band interferometric observations of Circinus. These observations allowed us to reconstruct images and fit Gaussian models to the L - and M -bands. Using these images and models we

1. Find a thin disk whose width is marginally resolved ($0.08 \text{ pc} = 4.23 \text{ mas}$). This disk is shown to be the spectral continuation of the the disk imaged in the N -band to shorter wavelengths, as the measured fluxes correspond to the fitted N -band temperatures. In addition to this thin disk, there is point-like source found in the L - and M -bands which was identified with the N -band point source based on the LMN -band SED fit.
2. Show that there is no trace of no hot dust ($T \sim 1500 \text{ K}$) in the circumnuclear dust structure of Circinus. By assuming the dust is there, we find that obscuration of $A_V \gtrsim 250 \text{ mag}$ is necessary to reproduce the measured fluxes. With dust extinction this high, the imaged disk could then play the role of the obscuring “torus” in the unified scheme of AGN.

6.3 Final Remarks

The imaged substructures and temperature distributions presented herein serve as a direct constraint on future physical modeling of the circumnuclear dust. They range from $\sim 0.1 - 5 \text{ pc}$ in scale and are seen in a wide wavelength range ($3\text{-}13 \mu\text{m}$). In the modern disk+wind model (e.g., [Hönig, 2019](#); [Stalevski et al., 2019](#)), the thin disk we image is related to inflowing material, while the polar emission represents a radiation-driven outflow. How these components relate to large-scale ($\gtrsim 10 \text{ pc}$) structures and furthermore to the host galaxy can be tested in both hydrodynamical modeling and future observations, specifically with the MATISSE ATs. It is clear that the classic geometrically-thick torus is not present in our imaging, but the (nearly) rotationally symmetric structures we recover can play much the same role; yet the detailed implications for AGN Unification remain to be explored through modeling and the MATISSE AGN Program.

Chapter 6

Conclusions and Outlook

1 Conclusions

The unified scheme of AGN proposed a dusty obscuring structure, which in Seyfert 2 AGN blocks observers' view to the broad line region, and in Seyfert 1 AGN is tilted to allow a direct sight-line. From this proposal sprung the idea of a geometrically and optically thick toroidal structure called the “dusty torus.” In the last decade, however, there has been a paradigm shift sometimes called “the death of the torus” wherein our idea of the circumnuclear dust has evolved significantly.

Observations with MIDI and VISIR revealed polar extended dust in a number of AGN, Circinus and NGC 1068 in particular (Tristram et al., 2014; López-Gonzaga et al., 2014; López-Gonzaga et al., 2016; Asmus et al., 2016). A thorough understanding of the morphology and temperature of this dust is crucial to justifying and explaining AGN unification. This circumnuclear dust structure, moreover, traces the material which both feeds the AGN and which is ejected via outflows, interacting with the host galaxy. The dust thus informs us about AGN evolution, AGN feedback, and AGN unification, and its study one of the primary science drivers of the new mid-infrared interferometric instrument MATISSE.

The aim of this thesis has been to reveal the nature of the circumnuclear dust through model-independent interferometric imaging using MATISSE. The circumnuclear dust structures range in scale from tens of parsecs to tenths of parsecs, and interferometric observations are crucial because only they allow us to measure the dust at very small angular resolution (3 mas = 0.06 pc in the L -band, 10 mas = 0.2 pc in the N -band for Circinus). Unlike previous mid-infrared interferometers, MATISSE provides closure phases (which inform us about the spatial distribution of the dust) and efficiently samples the uv -plane with its four telescopes (corresponding to six simultaneous baselines). This novel combination permits image reconstruction, a significant step forward in probing the true nature of the circumnuclear dust emission.

In this thesis, a large amount of work has been focused on two nearby AGN: Circinus and NGC 1068. Circinus is of particular interest as it is one of the closest Sy2 galaxies (at a distance of 4.2 Mpc Freeman et al., 1977; Tully et al., 2009) and the second brightest in the MIR (only fainter than NGC 1068). Circinus was observed extensively with the first generation MIR interferometer, MIDI (Tristram et al., 2007, 2014), and was a key driver of the “death of the torus” due to its prominent polar dust emission. NGC 1068

is *the* archetypical Seyfert 2 galaxy, and studies of its nucleus in polarized light inspired the unified scheme of AGN (Antonucci & Miller, 1985). It was also observed extensively with MIDI (Jaffe et al., 2004; López-Gonzaga et al., 2014), and it was shown to have significant polar emission and indications of a dust composition unlike the Milky Way.

In this thesis, I present the first L -, M -, and N -band images of Circinus and NGC 1068. These images, along with spectral energy distribution (SED) fits, have revealed similarities as well as significant differences between these two sources. In order to determine which aspects are general features of circumnuclear dust and the unified scheme, and which are object-specific peculiarities, a larger sample of AGN was necessary. At lower angular resolution, a statistical sample of AGN was studied in the L - and M -bands in order to 1) identify a larger number of AGN which can be studied in detail with MATISSE and 2) statistically determine which radiative transfer model geometries best reproduce the mid-infrared colors of AGN. The high- and low-resolution studies of AGN in the mid-infrared have revealed several common features and place strong constraints on circumnuclear dust modeling.

1. Polar Dust:

In both Circinus and NGC 1068, the presence of polar dust is confirmed with model-independent imaging. It is also found to be much more complex than the MIDI models suggested, with flux enhancements and asymmetries. In particular, both Circinus and NGC 1068 show enhanced emission near one edge of the ionization cone. In Circinus this is hypothesized to be related to a warped or tilted accretion disk (e.g., Stalevski et al., 2019), and the measured temperature enhancements of this region support this hypothesis. In both Circinus and NGC 1068, the polar dust extends $\gtrsim 1$ pc, though in NGC 1068 it is more cone shaped than in Circinus (perhaps due to its lower inclination). Lower resolution analysis of the mid-infrared colors ($L - M$ and $M - N$) of 119 AGN shows a strong preference for models which include polar dust.

2. Clumpiness:

The MATISSE images of Circinus and NGC 1068 at several wavelengths allow one to measure the temperature distribution of the dust. In Circinus, the polar dust is $\gtrsim 200$ K out to 1.5 pc from the central engine. I compare the temperature distribution to both analytical curves and radiative transfer models of clumpy media. I find that clumpy models agreed much better, as direct sight-lines to the central engine can allow dust to have a relatively high temperature far from the accretion disk. Additionally, in Circinus the imaged polar dust shows flux variations several times larger than the image noise, indicating a filamentary or patchy structure consistent with clumpy modeling. Similarly in NGC 1068, the dust is found to be $\gtrsim 200$ K at approximately 1 pc from the proposed SMBH location, indicating clumpiness in this AGN as well.

3. Dense Central Disk:

In neither Circinus nor NGC 1068 is evidence of hot dust found (i.e., with ~ 1500 K, the sublimation temperature of silicate dust). In Circinus, which is nearly edge-on ($i \geq 83^\circ$) a geometrically thin disk is found perpendicular to the polar emission in the N -band images. This disk shows a very rapidly decreasing radial temperature trend which is more consistent with a continuous-medium model than a clumpy model. Using the LM data, I show that very high extinctions ($A_V \gtrsim 250$ mag) are necessary to obscure the expected 1500 K dust. The high extinction and steep temperature gradient are indicative of a dense central disk which could play the role of the classical dusty torus. In NGC 1068, a dark band is found to

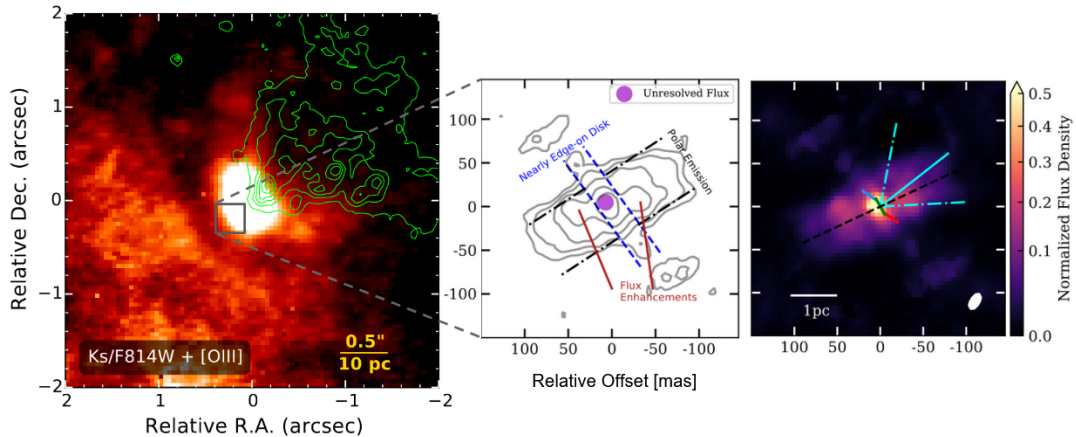


Fig. 6.1: Comparison of MATISSE polar emission to the large-scale photoionization cone. (*left*) Ratio of Ks/F814W filters from Mezcua et al. (2016) with [OIII] contours from Wilson et al. (2000) overplotted in green. (*center and right*) Zoom-in to the reconstructed MATISSE image of the circumnuclear dust in Circinus (explained in Ch. 5). The polar dust in the MATISSE image extends roughly in the same direction as the ionization cone; the opening angle of the ionization cone is shown in cyan in the rightmost panel. The placement of the MATISSE images relative to the optical image is only an estimate.

separate the emission into two roughly symmetric regions. To the north of this band, warm dust (~ 800 K) is found in a ring-like shape associated with the polar outflow. South of this band is highly extinguished and warm dust (800 – 1000 K). Based on the position of the SMBH as inferred from the 1.3 mm and 1.4 cm images, this dark band is associated with a dense, obscuring disk. These two objects thus imply that a dense obscuring disk accompanies the polar outflow in Seyfert 2 AGN and it may be the orientation of this disk which plays a large role in separating Seyferts 1 and 2. In the Circinus images, an unresolved source is found along the disk. This point source is associated with the warmest emission found in this circumnuclear structure (≈ 370 K) and may represent a lucky sight line to warmer dust due to clumpiness – as the current models do not reproduce this feature, this hypothesis and the amount of clumpiness in the dense disk remain to be explored.

While Circinus and NGC 1068 have their differences in terms of appearance, their common features are suggestive and supported by the statistical study of AGN in the L - and M -bands. The images, model fits, and parameter searches conducted throughout this thesis provide strong constraints on future models. They moreover suggest a new geometry to be considered in the unified scheme of AGN, wherein a geometrically thin, optically thick disk obscures the central engine, and the clumpy polar dust traces an interaction channel with the host galaxy via radiation-driven winds.

Fundamentally, the disk+wind structure I describe leads to a preferred direction for the AGN to interact with its host. The intense radiation from the AGN is blocked by the disk component, meaning that radiation-driven feedback (ionization and/or winds caused by radiation pressure) is primarily in the polar direction. The polar dust structures in Circinus and NGC 1068 align well with the large scale (10s of pc) ionization cones in each galaxy, implying a link between the regimes (Fig. 6.1 shows the MATISSE images relative to the photoionization cone to illustrate their alignment). The circumnuclear disk+wind structure would then directly determine which regions of the host galaxy are ionized by the AGN. Furthermore, the intense radiation in this direction drives dusty outflows (e.g., the radiation fountain model Wada et al., 2016), changing the composition of the host galaxy non-uniformly.

Yet the polar direction is not the only channel for host-AGN interaction. In Circinus,

ALMA observations of the dust continuum (Izumi et al., 2018; Tristram et al., 2022) show spiral arms of dusty material extending ~ 100 pc from the center, originating from a disk roughly orthogonal to the ionization cone – with a disk position angle closely matching the disk component imaged with MATISSE. Izumi et al. (2018) find that this disk rotates with Keplerian velocities, but they associate deviations of ~ 40 km s $^{-1}$ to streaming inflows toward the AGN. The inflowing material would then eventually find itself in the MATISSE disk, which likely feeds material toward the accretion disk. The presence of the dense disk (or conversely, the lack of a sight line to the accretion disk) could then indicate a wealth of material with which to power the AGN. This idea adds complexity to the unified scheme of AGN, as the Sy1/Sy2 dichotomy could be affected both by orientation and by the current amount of available fuel for the AGN. Quantifying the relative importance (i.e., size, flux, etc.) of the polar dust in comparison to the disk will give critical insights into the in- and outflow channels of the circumnuclear dust (see §2.2 and Fig. 6.2).

Low resolution studies (e.g., Ch. 2) indicate that Circinus and NGC 1068 are not atypical Seyfert 2 AGN. Their infrared colors are indistinguishable from the larger sample. Their Eddington ratios are likewise typical for AGN in the MIDI sample (Leftley et al., 2019). It is therefore not unreasonable to assume that these AGN are representative, at least in many ways, of the general Seyfert 2 population. A larger sample, however, is nonetheless necessary in order to probe AGN with different luminosities, environments, and evolutionary stages. The catalog work included in this thesis has identified 44 further AGN to study with MATISSE, pushing the sample beyond what was done with MIDI. Clearly, much remains to be done in terms of modeling and a statistical study of AGN at the milliarcsecond scale with MATISSE, but great strides have been made towards understanding the nature of the circumnuclear dust in AGN.

2 Outlook

2.1 Future MATISSE Observations and Analysis

MATISSE is still a new instrument, and the investigation of AGN in the L - and M -bands is just beginning. The L - and M -bands serve as a bridge from the warm large-scale dust emission observed with MIDI in the N -band to the compact, hot dust emission observed with GRAVITY in the K -band. Additionally, the efficient uv -sampling and closure phases of MATISSE will allow us to push beyond the N -band studies completed with MIDI, applying the only recently possible technique of image reconstruction. While NGC 1068 and Circinus have been studied extensively with the UT array, much remains to be done in both these AGN at higher resolution and with the short-baseline AT array. Moreover, these are both Sy 2 AGN, and any exploration of AGN unification needs to include a variety of AGN types.

Several questions remain unanswered from the MIDI era, and several new ones have been posed based on the unprecedented detail achieved in NGC 1068 and Circinus. The polar dust is shown to be clumpy, but can we quantify the amount of clumpiness? In neither NGC 1068 nor Circinus is dust at the sublimation temperature found, but is this hot dust even present? Or is it obscured by something analogous to the dusty torus? What are the dynamics of the dusty structures? What are the chemical properties of the dust, and is it being reprocessed by the central engine? Are NGC 1068 and Circinus typical AGN; if so, what drives the observed differences between them? Can we devise a unified model for not only the Sy 2 population, but for AGN in general? In order to answer these questions, further exploration of Circinus and NGC 1068 as well as infrared interferometric observations and images of further targets are necessary.




Fig. 6.2: Comparison of wind-dominated circumnuclear dust (*left*) to disk-dominated dust (*right*). The models are given at both model resolution (0.04 pc; grayscale) and at the resolution of the LBTI or the MATISSE ATs with ~ 20 m baselines (~ 2 pc; orange) for an AGN similar to Circinus but located 10 Mpc away. Windy models indicate that outflow processes dominate, while disk-like models indicate that inflows play a more important role. Models are from M. Stalevski (priv. comm.) but are based on the models presented in [Stalevski et al. \(2019\)](#).

1. Further exploration of existing data:

In addition to Circinus and NGC 1068, a handful of other AGN have already been observed as part of the MATISSE Guaranteed Time Observations (GTO) program: Centaurus A (Cen A; radio galaxy), 3C273 (a quasar), NGC 3783 (Seyfert 1), and NGC 424 (Seyfert 1). In preliminary analysis, Cen A and NGC 3783 both show non-zero closure phases and low visibilities, indicating that they contain resolved structures. They have both been observed several times, and have sufficient uv -coverage to attempt imaging. These AGN are particularly interesting: Cen A is the nearest AGN and has significant non-thermal emission in the N -band ([Meisenheimer et al., 2007](#)); NGC 3783 will potentially be the first Seyfert 1 imaged with MATISSE, and these images can be immediately linked to GRAVITY K -band images ([Gravity Collaboration et al., 2021](#)). Additionally, MATISSE observations of NGC 1068 were obtained in high N -band spectral resolution ($R \sim 200$), and further analysis of these data will allow us to spatially resolve the dust composition, better understanding dust dynamics in the polar region. Finally, MATISSE observations with the 1.8 m ATs are strongly affected by a correlated flux bias at ~ 5 Jy in the N -band. While AT observations of NGC 1068 and Circinus have been obtained, this piston-induced bias cannot yet be corrected. Efforts by the MATISSE consortium are underway to de-bias the data (Walter Jaffe, James Leftley priv. comm.), and once this method is verified, the existing AT data can be utilized. The AT data have been obtained on short baselines, and will provide the first closure phase measurements of the large-scale structures ($\gtrsim 2$ pc in Circinus) and will allow us to better connect the AGN to its host galaxy.

2. New observations of NGC 1068 and Circinus:

The aforementioned AT observations of Circinus and NGC 1068 are few in number (i.e., one snapshot in Circinus), and should be supplemented with further snapshots. The short-baseline observations are crucial for MATISSE imaging of AGN, as they measure the true extent of the polar dust. The importance of short baselines is shown for Circinus in Ch. 5 §3.2 and for NGC 1068 in Fig. 6.3. Measuring temperatures of the polar dust helps constrain clumpiness and the origin of the material. Full imaging of the large-scale structure is crucial for linking the sub-pc circumnuclear dust structures to the host galaxy. Existing and future single-dish (e.g., the VLT and the James Webb Space Telescope) observations measure the dust emission at 10s of parsec, and the AT observations can bridge the gap down to

the sub-parsec dust structures. In addition to AT observations, using the upgraded VLTI we can explore several new questions. High resolution L -band spectra, obtainable now with GRA4MAT¹, will explore a curious $3.4 \mu\text{m}$ emission feature found on some baselines which is possibly nano-diamonds (Gámez Rosas et al., incl. J. Isbell, in prep.). Also, there is some controversy about the temperature of the dust found in NGC 1068 (Gravity Collaboration et al. (2020) and this work find ~ 1500 K and 800 K, respectively), and this will certainly be explored further with higher fidelity images due to more observations with both GRAVITY and MATISSE. Finally, at this time the longest MATISSE baseline (UT1-UT4) cannot observe Circinus for long after it transits due to delay line shadowing; but the delay lines will be upgraded, allowing us to fill an unexplored part of the uv -plane.

3. An AGN survey with MATISSE:

Circinus and NGC 1068 have many similarities, but they also have many key differences; which aspects are common to a general AGN population and the necessary components of a unified model can only be explored via a statistical sample of AGN. The number of MATISSE-observed AGN is increasing each ESO period, but we are essentially limited to the same ~ 20 AGN that were part of the MIDI Large Program (Burtscher et al., 2013). Upgrades to the VLTI, in particular GRAVITY+ and MATISSE-Wide aim to permit the observations of many more AGN which are fainter and more distant, through the use of nearby stars for fringe tracking. Correcting for atmospheric phase shifts using a bright star’s fringes will allow for longer coherent integration of AGN fringes, resulting in increased sensitivity. While the 20 MIDI-era AGN are a good start for a MATISSE AGN sample, a systematic large scale survey is crucial for the understanding of the general AGN population. For this purpose, in Ch. 2 I identified 44 AGN for study with MATISSE, doubling the MIDI sample. Moreover, GTO observations are restricted to the scheduled night, making them susceptible to weather losses. The MATISSE observing strategy and planning procedure has evolved sufficiently over the last few years to make a large-scale survey possible in service mode, guaranteeing “decent” to “good” atmospheric conditions, a crucial requirement for faint sources. Careful considerations, based on our new understanding of NGC 1068 and Circinus, would be necessary to determine the required uv -spacing to characterize these AGN as efficiently as possible. A MATISSE survey would give sizes and orientations of the circumnuclear dust in the LMN bands, allowing us to look for statistical differences between AGN types and to link these measurements to fundamental parameters such as Eddington ratio (e.g., Leftley et al., 2019, for MIDI). We can also explore whether Seyfert 1s contain the same circumnuclear structures or whether they are less dusty than Seyfert 2s.

4. Comparison to simulations:

Only a small portion of this work was dedicated to comparisons between models and observations, and so further work can and should be done. Firstly, the images of Circinus and NGC 1068 place new constraints on the temperature distribution of the circumnuclear dust. Temperature distributions can place a strong constraint on dust clumpiness, as it allows dust to remain warm at relatively large distances from the central engine. Detailed comparison to models, and even fitting of the parameters in models similar to those shown in Ch. 5 §3.4 can inform the dust distributions of future models. The models of Circinus from Stalevski et al. (2019) were compared to the MIDI observations and to the existing SEDs. The new LM data from both ISAAC (Ch. 2) and from MATISSE (Ch. 5 §4) should be utilized

¹GRAVITY fringe tracking for MATISSE; in science verification at the time of writing.




Fig. 6.3: Demonstration of need for short baselines in NGC 1068. From (*left*) to (*right*): the VISIR image of NGC 1068 from [Asmus et al. \(2016\)](#) which shows extensions far to the north and south of the nucleus; the MIDI model from [López-Gonzaga et al. \(2014\)](#) fitted to combined UT and AT observations; the MATISSE image reconstruction from Ch. 4 made with only UT observations. The large scale flux is resolved out in the rightmost panel because it uses only baselines > 30 m. The green circle represents the resolution of the LBTI with a 22.8 m baseline.

to update the radiative transfer models. There are already indications in the N -band that the disk component of the models should be tweaked to reproduce the fluxes and temperatures found in Circinus (Ch. 5 §3.4). Finally, new work in radiation-hydrodynamical (RHD) models by authors such as [Williamson et al. \(2020\)](#) and [Venanzi et al. \(2020\)](#) should be compared to the MATISSE images for consistency. For example, in Circinus we showed significant flux enhancements within the polar dust which could be explained by an accretion disk tilted with respect to the dust structures. Existing RHD models produce much more uniform polar emission, and the dynamical lifetime of such an offset structure should be explored with new RHD simulations.

2.2 The Large Binocular Telescope Interferometer

New ground is being tread with mid-infrared interferometers other than MATISSE. The Large Binocular Telescope Interferometer (LBTI) instrument NOMIC is a Fizeau imaging interferometer that operates on a 23 m baseline and consists of two 8.4 m diameter mirrors ([Ertel, 2020](#)). Fizeau imaging differs from the Michelson-style interferometer used at the VLTI in that there is no beam combiner and the interference happens directly on the image plane. NOMIC operates in the N -band (8-13 μm) and complements the VLTI well for several reasons. First, it is in the northern hemisphere (Mt. Graham is located at $32^{\circ}42'06''$ N $109^{\circ}52'17''$ W) and thus can observe AGN blocked to MATISSE. Second, the 8.4 m telescopes operate on a shorter baseline than the UTs at the VLTI, meaning that they give interferometric resolution comparable to the ATs but with the sensitivity of the UTs. In AGN, this high sensitivity on (relatively) large spatial scales allows us to study the polar dust and its interaction with the host galaxy.

Recently, I was Principle Investigator of an accepted observing proposal which aims to study NGC 1068 with NOMIC, complementing the UT baselines. NGC 1068 is one of the few AGN accessible to both interferometers. These observations should happen in Autumn 2022, and can test whether NGC 1068 is disk-dominated or wind-dominated (i.e., whether it is primarily being fed or is primarily outflowing; see Figs. 6.2 and 6.3). While the observations are promising, there is no standard reduction of the data from this new instrument, and a procedure to reduce and image the data will have to be

produced. Fortunately, the Fizeau imaging with the LBTI has been proven through imaging of Jupiter’s moon, Io (Conrad et al., 2015; de Kleer et al., 2017). Observations with NOMIC can be a precursor study of AGN for future 30 m class telescopes such as the Extremely Large Telescope (ELT; $d = 39.3$ m), the Thirty Meter Telescope (TMT; $d = 30$ m), and the Giant Magellan Telescope (GMT; $d = 23.4$ m).

2.3 New and Future Observatories

The next decade will be exciting for the study of AGN, with new advancements both in sensitivity and resolution. First, the James Webb Space Telescope (JWST) was launched in December 2021 and has started science observations as of summer 2022. JWST observes in the visible and infrared ($0.6 - 28.3 \mu\text{m}$) with a primary mirror of 6.5 m. While diffraction-limited images from JWST will have lower resolution than existing VLT instruments (e.g., VISIR), the lack of atmosphere means that JWST’s integration times can be long and the thermal background is limited, resulting in much deeper images. Studies with VLT/VISIR showed polar dust to ~ 100 pc in several nearby AGN; JWST will be able to probe cooler, fainter, and more diffuse emission from the polar dust, giving a more holistic picture of dusty, radiation-driven winds in AGN. Additionally, the MIRI instrument can obtain high resolution spectra of the polar emission, giving insights about its composition. Finally, JWST provides sparse aperture masking capabilities, giving interferometric results with ≤ 6.2 m baselines, yielding a direct complement to MATISSE observations.

The next generation of ground based telescopes (GMT, TMT, and ELT) will have lower resolution than interferometric instruments such as MATISSE, but will be complementary to the UT baselines which are > 30 m. The ELT, for example, will be able to directly image the large-scale circumnuclear emission in a very short amount of time, while now we require several snapshots with the ATs or the LBTI to produce models or (in the best case) image reconstructions which have relatively low dynamical range. Most importantly, the synergy of these observatories – the deep observations of JWST, the studies of large-scale emission with the ELT, and the high resolution images from MATISSE and GRAVITY – can be combined to test AGN unification, the origin and extent of polar winds, how the AGN is fed, and the feedback of the AGN on its host. Studies similar to the work by Asmus et al. (2014) and in Ch. 2 can be carried out with the 30 m class telescopes efficiently, separating the nuclear flux (i.e. the circumnuclear dust) from diffuse galactic or stellar emission in a large number of galaxies. The amount and extent of circumnuclear material can then be linked to fundamental parameters of the AGN such as its Eddington ratio or SMBH mass.

I look forward to expanding the MATISSE results with analysis of NGC 3783, Cen A, and a MATISSE AGM Large Program. The joint capabilities of the current and upcoming infrared facilities will give us a holistic view of this crucial dust structure within the next decade.

Publications by the Author

Publications Used in this Thesis

Sanchez-Bermudez, J., Millour, F., Baron, F., et al. 2018, *Experimental Astronomy*, 46, 45

Isbell, J. W., Burtscher, L., Asmus, D., et al. 2021, *ApJ*, 910, 104

Gómez Rosas, V., Isbell, J. W., Jaffe, W., et al. 2022, *Nature*, 602, 403

Lopez, B., Lagarde, S., Petrov, R. G., et al. 2022, *A&A*, 659, A192

Isbell, J. W., Meisenheimer, K., Pott, J.-U., et al. 2022, *A&A*, 663, A35

Other Publications

Fu, H., Hennawi, J. F., Prochaska, J. X., Mutel, R., Casey, C., Cooray, A., Kereš, D., Zhang, Z.-Y., Clements, D., Isbell, J., Lang, C., McGinnis, D., Michałowski, M. J., Mooley, K., Perley, D., Stockton, A., & Thompson, D. 2016, *ApJ*, 832, 52

Fu, H., Isbell, J., Casey, C. M., Cooray, A., Prochaska, J. X., Scoville, N., & Stockton, A. 2017, *ApJ*, 844, 123

Fu, H., Steffen, J. L., Gross, A. C., Dai, Y. S., Isbell, J. W., Lin, L., Wake, D., Xue, R., Bizyaev, D., & Pan, K. 2018, *ApJ*, 856, 93

Xue, R., Fu, H., Isbell, J., Ivison, R. J., Cooray, A., & Oteo, I. 2018, *ApJL*, 864, L11

Isbell, J. W., Xue, R., & Fu, H. 2018, *ApJL*, 869, L37

Brown, S., Bergerud, B., Costa, A., Gaensler, B. M., Isbell, J., LaRocca, D., Norris, R., Purcell, C., Rudnick, L., & Sun, X. 2019, *MNRAS*, 483, 964

Publications by the MATISSE Consortium (incl. J. Isbell)

Varga, J., Hogerheijde, M., van Boekel, R., et al. 2021, *A&A*, 647, A56

Hocdé, V., Nardetto, N., Matter, A., et al. 2021, *A&A*, 651, A92

Kokoulina, E., Matter, A., Lopez, B., et al. 2021, *A&A*, 652, A61

Weigelt, G., Hofmann, K.-H., Schertl, D., et al. 2021, *A&A*, 652, A140

Hofmann, K.-H., Bensburg, A., Schertl, D., et al. 2022, *A&A*, 658, A81

Chiavassa, A., Kravchenko, K., Montargès, M., et al. 2022, *A&A*, 658, A185

Drevon, J., Millour, F., Cruzalèbes, P., et al. (2022), arXiv e-prints, arXiv:2208.10845.

References

- Aalto, S., Wilner, D., Spaans, M., et al. 2009, *Astronomy and Astrophysics*, 493, 481
- Akaike, H. 1981, *Journal of Econometrics*, 16, 3
- Alonso-Herrero, A., Quillen, A. C., Simpson, C., Efstathiou, A., & Ward, M. J. 2001, *The Astronomical Journal*, 121, 1369
- Alonso-Herrero, A., Ramos Almeida, C., Mason, R., et al. 2011, *The Astrophysical Journal*, 736, 82
- Antonucci, R. 1993, *Annual Review of Astronomy and Astrophysics*, 31, 473
- Antonucci, R., Hurt, T., & Miller, J. 1994, *ApJ*, 430, 210
- Antonucci, R. R. J. & Miller, J. S. 1985, *The Astrophysical Journal*, 297, 621
- Arévalo, P., Bauer, F. E., Puccetti, S., et al. 2014, *ApJ*, 791, 81
- Arnaud, K. A., Branduardi-Raymont, G., Culhane, J. L., et al. 1985, *MNRAS*, 217, 105
- Asmus, D., Hönig, S. F., & Gandhi, P. 2016, *The Astrophysical Journal*, 822, 109
- Asmus, D., Hönig, S. F., Gandhi, P., Smette, A., & Duschl, W. J. 2014, *Monthly Notices of the Royal Astronomical Society*, 439, 1648
- Assef, R. J., Kochanek, C. S., Brodwin, M., et al. 2010, *The Astrophysical Journal*, 713, 970
- Assef, R. J., Stern, D., Noirot, G., et al. 2018, *The Astrophysical Journal Supplement Series*, 234, 23
- Baldwin, J. E., Haniff, C. A., Mackay, C. D., & Warner, P. J. 1986, *Nature*, 320, 595
- Baron, F., Cotton, W. D., Lawson, P. R., et al. 2012, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 8445, *Optical and Infrared Interferometry III*, ed. F. Delplancke, J. K. Rajagopal, & F. Malbet, 84451E
- Barvainis, R. 1987, *The Astrophysical Journal*, 320, 537
- Baskin, A. & Laor, A. 2018, *MNRAS*, 474, 1970
- Beckmann, V. & Shrader, C. R. 2012, *Active Galactic Nuclei*
- Burtscher, L., Davies, R. I., Graciá-Carpio, J., et al. 2016, *Astronomy and Astrophysics*, 586, A28
- Burtscher, L., Meisenheimer, K., Tristram, K. R. W., et al. 2013, *Astronomy and Astrophysics*, 558, A149

- Burtscher, L., Orban de Xivry, G., Davies, R. I., et al. 2015, *Astronomy and Astrophysics*, 578, A47
- Buscher, D. F. 1994, in *Very High Angular Resolution Imaging*, ed. J. G. Robertson & W. J. Tango, Vol. 158, 91
- Buscher, D. F. & Longair, F. b. M. 2015, *Practical Optical Interferometry*
- Byrd, R. H., Lu, P., Nosedal, J., & Zhu, C. 1995, *SIAM J. Sci. Comput.*, 16, 1190–1208
- Carroll, B. W. & Ostlie, D. A. 2006, *An introduction to modern astrophysics and cosmology*
- Cohen, M., Walker, R. G., Carter, B., et al. 1999, *The Astronomical Journal*, 117, 1864
- Combes, F. 2021, *Active Galactic Nuclei: Fueling and feedback*
- Conrad, A., de Kleer, K., Leisenring, J., et al. 2015, *AJ*, 149, 175
- Cornwell, T. J. 2008, *IEEE Journal of Selected Topics in Signal Processing*, 2, 793
- Costagliola, F., Aalto, S., Rodriguez, M. I., et al. 2011, *Astronomy and Astrophysics*, 528, A30
- Cotton, W., Monnier, J., Baron, F., et al. 2008a, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7013, *Optical and Infrared Interferometry*, ed. M. Schöller, W. C. Danchi, & F. Delplancke, 70131N
- Cotton, W. D. 2017, *PASP*, 129, 094501
- Cotton, W. D., Jaffe, W., Perrin, G., & Woillez, J. 2008b, *A&A*, 477, 517
- Cruzalèbes, P., Petrov, R. G., Robbe-Dubois, S., et al. 2019, *arXiv e-prints*, arXiv:1910.00542
- Curran, S. J., Johansson, L. E. B., Rydbeck, G., & Booth, R. S. 1998, *A&A*, 338, 863
- Das, V., Crenshaw, D. M., Kraemer, S. B., & Deo, R. P. 2006, *AJ*, 132, 620
- de Kleer, K., Skrutskie, M., Leisenring, J., et al. 2017, *Nature*, 545, 199
- Dullemond, C. P. & van Bemmell, I. M. 2005, *A&A*, 436, 47
- Edelson, R. A. & Malkan, M. A. 1986, *The Astrophysical Journal*, 308, 59
- Elmouttie, M., Haynes, R. F., Jones, K. L., Sadler, E. M., & Ehle, M. 1998, *MNRAS*, 297, 1202
- Ertel, S. 2020, in *Ground-Based Thermal Infrared Astronomy - Past, Present and Future*, 17
- Evans, I. N., Ford, H. C., Kinney, A. L., et al. 1991, *ApJ*, 369, L27
- Fath, E. A. 1909, *Lick Observatory Bulletin*, 149, 71
- Fischer, T. C., Crenshaw, D. M., Kraemer, S. B., & Schmitt, H. R. 2013, *ApJS*, 209, 1
- Forbes, D. A., Norris, R. P., Williger, G. M., & Smith, R. C. 1994, *The Astronomical Journal*, 107, 984

- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *Publications of the Astronomical Society of the Pacific*, 125, 306
- Freeman, K. C., Karlsson, B., Lynga, G., et al. 1977, *A&A*, 55, 445
- Fritz, J., Franceschini, A., & Hatziminaoglou, E. 2006, *Monthly Notices of the Royal Astronomical Society*, 366, 767
- Fritz, T. K., Gillessen, S., Dodds-Eden, K., et al. 2011, *ApJ*, 737, 73
- Gallimore, J. F., Baum, S. A., & O’Dea, C. P. 2004, *The Astrophysical Journal*, 613, 794
- Gallimore, J. F., Baum, S. A., O’Dea, C. P., & Pedlar, A. 1996, *ApJ*, 458, 136
- Gómez Rosas, V., Isbell, J. W., Jaffe, W., et al. 2022, *Nature*, 602, 403
- García-Bernete, I., Ramos Almeida, C., Alonso-Herrero, A., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 486, 4917
- García-Burillo, S., Combes, F., Ramos Almeida, C., et al. 2019, *A&A*, 632, A61
- García-González, J., Alonso-Herrero, A., Hernán-Caballero, A., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 458, 4512
- Gaskell, C. M., Goosmann, R. W., Antonucci, R. R. J., & Whysong, D. H. 2004, *ApJ*, 616, 147
- George, I. M. & Fabian, A. C. 1991, *MNRAS*, 249, 352
- González-Martín, O., Masegosa, J., García-Bernete, I., et al. 2019, *ApJ*, 884, 11
- Gordon, J. A. & Buscher, D. F. 2012, *A&A*, 541, A46
- Gravity Collaboration, Amorim, A., Bauböck, M., et al. 2021, *A&A*, 648, A117
- Gravity Collaboration, Dexter, J., Shangquan, J., et al. 2020, *Astronomy and Astrophysics*, 635, A92
- Greenbaum, A. Z., Pueyo, L., Sivaramakrishnan, A., & Lacour, S. 2015, *ApJ*, 798, 68
- Greenhill, L. J., Booth, R. S., Ellingsen, S. P., et al. 2003, *The Astrophysical Journal*, 590, 162
- Greenhill, L. J., Gwinn, C. R., Antonucci, R., & Barvainis, R. 1996, *ApJ*, 472, L21
- Greisen, E. W. 1990, in *Acquisition, Processing and Archiving of Astronomical Images*, 125–142
- Greisen, E. W. 2003, in *Astrophysics and Space Science Library*, Vol. 285, *Information Handling in Astronomy - Historical Vistas*, ed. A. Heck, 109
- Hager, W. W. & Zhang, H. 2005, *SIAM Journal on Optimization*, 16, 170
- Ho, L. C. 2002, *ApJ*, 564, 120
- Hofmann, K. H., Bensberg, A., Schertl, D., et al. 2022, *A&A*, 658, A81
- Hofmann, K.-H., Heininger, M., Schertl, D., et al. 2016, 9907, 99073H, conference Name: *Optical and Infrared Interferometry and Imaging V* ISBN: 9781510601932

- Hofmann, K.-H., Weigelt, G., & Schertl, D. 2014, *Astronomy and Astrophysics*, 565, A48
- Högbom, J. A. 1974, *Astronomy and Astrophysics Supplement Series*, 15, 417
- Hönig, S. F. 2019, *ApJ*, 884, 171
- Hönig, S. F., Beckert, T., Ohnaka, K., & Weigelt, G. 2006, *A&A*, 452, 459
- Hönig, S. F. & Kishimoto, M. 2010, *Astronomy and Astrophysics*, 523, A27
- Hönig, S. F. & Kishimoto, M. 2017, *The Astrophysical Journal Letters*, 838, L20
- Hönig, S. F., Kishimoto, M., Antonucci, R., et al. 2012, *ApJ*, 755, 149
- Hönig, S. F., Kishimoto, M., Tristram, K. R. W., et al. 2013, *The Astrophysical Journal*, 771, 87
- Horan, D. & Weekes, T. C. 2004, *New A Rev.*, 48, 527
- Hubble, E. P. 1926, *ApJ*, 64, 321
- Impellizzeri, C. M. V., Gallimore, J. F., Baum, S. A., et al. 2019, *ApJ*, 884, L28
- Isbell, J. W., Burtscher, L., Asmus, D., et al. 2021, *ApJ*, 910, 104
- Isbell, J. W., Meisenheimer, K., Pott, J.-U., et al. 2022, *A&A*, 663, A35
- Izumi, T., Wada, K., Fukushige, R., Hamamura, S., & Kohno, K. 2018, *ApJ*, 867, 48
- Jaffe, W., Meisenheimer, K., Röttgering, H., Leinert, C., & Richichi, A. 2004, 222, 37
- Jaffe, W. J. 2004, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 5491, *New Frontiers in Stellar Interferometry*, ed. W. A. Traub, 715
- Jansky, K. G. 1933, *Nature*, 132, 66
- Jennison, R. C. 1958, *MNRAS*, 118, 276
- Kishimoto, M. 1999, *ApJ*, 518, 676
- Kishimoto, M., Antonucci, R., Blaes, O., et al. 2008, *Nature*, 454, 492
- Kishimoto, M., Hönig, S. F., Antonucci, R., et al. 2011, *Astronomy and Astrophysics*, 536, A78
- Koechlin, L., Lawson, P. R., Mourard, D., et al. 1996, *Appl. Opt.*, 35, 3002
- Köhler, M. & Li, A. 2010, *MNRAS*, 406, L6
- Krolik, J. H. 2007, *ApJ*, 661, 52
- Krolik, J. H. & Begelman, M. C. 1988, *The Astrophysical Journal*, 329, 702
- Lachaume, R. 2003, *A&A*, 400, 795
- Lacour, S., Tuthill, P., Amico, P., et al. 2011, *A&A*, 532, A72
- Lawson, C. L. & Hanson, R. J. 1995, *Solving least squares problems (SIAM)*

- Lawson, P. R., Cotton, W. D., Hummel, C. A., et al. 2004a, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 5491, New Frontiers in Stellar Interferometry, ed. W. A. Traub, 886
- Lawson, P. R., Cotton, W. D., Hummel, C. A., et al. 2004b, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 5491, New Frontiers in Stellar Interferometry, ed. W. A. Traub, 886
- Leftley, J. H., Höinig, S. F., Asmus, D., et al. 2019, *The Astrophysical Journal*, 886, 55
- Leftley, J. H., Tristram, K. R. W., Höinig, S. F., et al. 2021, *ApJ*, 912, 96
- Leftley, J. H., Tristram, K. R. W., Höinig, S. F., et al. 2018, *ApJ*, 862, 17
- Leinert, C., Graser, U., Richichi, A., et al. 2003, *The Messenger*, 112, 13
- Lira, P., Videla, L., Wu, Y., et al. 2013, *The Astrophysical Journal*, 764, 159
- Lopez, B., Lagarde, S., Jaffe, W., et al. 2014, *The Messenger*, 157, 5
- Lopez, B., Lagarde, S., Petrov, R. G., et al. 2022, *A&A*, 659, A192
- López-Gonzaga, N., Burtscher, L., Tristram, K. R. W., Meisenheimer, K., & Scharmann, M. 2016, *A&A*, 591, A47
- López-Gonzaga, N. & Jaffe, W. 2016, *A&A*, 591, A128
- López-Gonzaga, N., Jaffe, W., Burtscher, L., Tristram, K. R. W., & Meisenheimer, K. 2014, *Astronomy and Astrophysics*, 565, A71
- Maiolino, R., Alonso-Herrero, A., Anders, S., et al. 2000, *ApJ*, 531, 219
- Makarov, D., Prugniel, P., Terekhova, N., Courtois, H., & Vauglin, I. 2014, *Astronomy and Astrophysics*, 570, A13
- Manske, V., Henning, T., & Men'shchikov, A. B. 1998, *Astronomy and Astrophysics*, 331, 52
- Marconi, A., Moorwood, A. F. M., Origlia, L., & Oliva, E. 1994, *The Messenger*, 78, 20
- Martínez-Paredes, M., González-Martín, O., Esparza-Arredondo, D., et al. 2020, *The Astrophysical Journal*, 890, 152
- Mateos, S., Alonso-Herrero, A., Carrera, F. J., et al. 2012, *Monthly Notices of the Royal Astronomical Society*, 426, 3271
- Mathis, J. S., Rumpl, W., & Nordsieck, K. H. 1977, *ApJ*, 217, 425
- Matt, G., Fiore, F., Perola, G. C., et al. 1996, *MNRAS*, 281, L69
- Matt, G., Guainazzi, M., Maiolino, R., et al. 1999, *A&A*, 341, L39
- Meilland, A., Delaa, O., Stee, P., et al. 2011, *A&A*, 532, A80
- Meisenheimer, K., Tristram, K. R. W., Jaffe, W., et al. 2007, *Astronomy and Astrophysics*, 471, 453
- Mezcua, M., Prieto, M. A., Fernández-Ontiveros, J. A., & Tristram, K. R. W. 2016, *MNRAS*, 457, L94

- Miley, G. K., Neugebauer, G., & Soifer, B. T. 1985, *ApJ*, 293, L11
- Millour, F., Petrov, R., Malbet, F., et al. 2008, in *2007 ESO Instrument Calibration Workshop*, ed. A. Kaufer & F. Kerber, 461
- Millour, F., Tatulli, E., Chelli, A. E., et al. 2004, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 5491, *New Frontiers in Stellar Interferometry*, ed. W. A. Traub, 1222
- Min, M., Hovenier, J. W., & de Koter, A. 2005, *A&A*, 432, 909
- Mingozzi, M., Cresci, G., Venturi, G., et al. 2019, *A&A*, 622, A146
- Monaco, P. 2018, *Radiative Processes Active Galactic Nuclei*
- Monnier, J. D. 2003, in *EAS Publications Series*, Vol. 6, *EAS Publications Series*, ed. G. Perrin & F. Malbet, 213
- Monnier, J. D. 2007, *New A Rev.*, 51, 604
- Moorwood, A., Cuby, J.-G., Ballester, P., et al. 1999, *The Messenger*, 95, 1
- Moorwood, A. F. M., van der Werf, P. P., Kotilainen, J. K., Marconi, A., & Oliva, E. 1996, *Astronomy and Astrophysics*, 308, L1
- Mor, R. & Netzer, H. 2012, *Monthly Notices of the Royal Astronomical Society*, 420, 526
- Namekata, D., Umemura, M., & Hasegawa, K. 2014, *MNRAS*, 443, 2018
- Nayakshin, S. 2005, *MNRAS*, 359, 545
- Nenkova, M., Ivezić, Ž., & Elitzur, M. 2002, *The Astrophysical Journal Letters*, 570, L9
- Nenkova, M., Sirocky, M. M., Ivezić, Ž., & Elitzur, M. 2008a, *The Astrophysical Journal*, 685, 147
- Nenkova, M., Sirocky, M. M., Ivezić, Ž., & Elitzur, M. 2008b, *The Astrophysical Journal*, 685, 147
- Nenkova, M., Sirocky, M. M., Nikutta, R., Ivezić, Ž., & Elitzur, M. 2008c, *The Astrophysical Journal*, 685, 160
- Netzer, H. 1987, *Monthly Notices of the Royal Astronomical Society*, 225, 55
- Netzer, H. 2015, *Annual Review of Astronomy and Astrophysics*, 53, 365
- Ohyama, Y., Sakamoto, K., Aalto, S., & Gallagher, III, J. S. 2019, *The Astrophysical Journal*, 871, 191
- Oliva, E., Marconi, A., Cimatti, A., & di Serego Alighieri, S. 1998, *Astronomy and Astrophysics*, 329, L21
- Oliva, E., Marconi, A., & Moorwood, A. F. M. 1999, *A&A*, 342, 87
- Oliva, E., Salvati, M., Moorwood, A. F. M., & Marconi, A. 1994, *A&A*, 288, 457
- Packham, C., Radomski, J. T., Roche, P. F., et al. 2005, *ApJ*, 618, L17
- Petrov, R. G., Malbet, F., Weigelt, G., et al. 2007, *A&A*, 464, 1

- Petterson, J. A. 1977, *ApJ*, 214, 550
- Pier, E. A. & Krolik, J. H. 1992, *ApJ*, 401, 99
- Pillepich, A., Nelson, D., Hernquist, L., et al. 2018, *MNRAS*, 475, 648
- Poncelet, A., Sol, H., & Perrin, G. 2008, *A&A*, 481, 305
- Predehl, P. & Schmitt, J. H. M. M. 1995, *A&A*, 293, 889
- Prieto, M. A., Meisenheimer, K., Marco, O., et al. 2004, *ApJ*, 614, 135
- Prieto, M. A., Nadolny, J., Fernández-Ontiveros, J. A., & Mezcua, M. 2021, *MNRAS*, 506, 562
- Prieto, M. A., Reunanen, J., Tristram, K. R. W., et al. 2010, *Monthly Notices of the Royal Astronomical Society*, 402, 724
- Raban, D., Jaffe, W., Röttgering, H., Meisenheimer, K., & Tristram, K. R. W. 2009, *Monthly Notices of the Royal Astronomical Society*, 394, 1325
- Ramos Almeida, C., Levenson, N. A., Rodríguez Espinosa, J. M., et al. 2009, *The Astrophysical Journal*, 702, 1127
- Reber, G. 1944, *ApJ*, 100, 279
- Ricci, C., Ueda, Y., Koss, M. J., et al. 2015, *ApJ*, 815, L13
- Roche, P. F., Packham, C., Telesco, C. M., et al. 2006, *MNRAS*, 367, 1689
- Salpeter, E. E. 1964, *ApJ*, 140, 796
- Sanchez-Bermudez, J., Millour, F., Baron, F., et al. 2018, *Experimental Astronomy*, 46, 457
- Sanchez-Bermudez, J., Rosales-Guzmán, A., Morales, H., et al. 2020, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 11446, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 114461O
- Sanchez-Bermudez, J., Thiébaud, E., Hofmann, K. H., et al. 2016, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9907, *Optical and Infrared Interferometry and Imaging V*, ed. F. Malbet, M. J. Creech-Eakman, & P. G. Tuthill, 99071D
- Sanders, D. B., Phinney, E. S., Neugebauer, G., Soifer, B. T., & Matthews, K. 1989, *ApJ*, 347, 29
- Schartmann, M. 2007, Ph.D. Thesis
- Schartmann, M., Krause, M., & Burkert, A. 2011, *MNRAS*, 415, 741
- Schartmann, M., Meisenheimer, K., Camenzind, M., Wolf, S., & Henning, T. 2005, *Astronomy and Astrophysics*, 437, 861
- Schartmann, M., Meisenheimer, K., Camenzind, M., et al. 2008, *Astronomy and Astrophysics*, 482, 67
- Schartmann, M., Wada, K., Prieto, M. A., Burkert, A., & Tristram, K. R. W. 2014, *MNRAS*, 445, 3878

- Schmidt, M. 1963, *Nature*, 197, 1040
- Schutz, A., Thiébaud, É., Soulez, F., et al. 2016, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9907, *Optical and Infrared Interferometry and Imaging V*, ed. F. Malbet, M. J. Creech-Eakman, & P. G. Tuthill, 99073D
- Scott, J. E., Kriss, G. A., Brotherton, M., et al. 2004, *The Astrophysical Journal*, 615, 135
- Secret, N. J., Dudik, R. P., Dorland, B. N., et al. 2015, *The Astrophysical Journal Supplement Series*, 221, 12
- Seyfert, C. K. 1943, *ApJ*, 97, 28
- Shannon, C. E. 1948, *Bell System Technical Journal*, 27, 623
- Shao, J. & Wu, C. F. J. 1989, *The Annals of Statistics*, 17, 1176
- Shao, Z., Jiang, B. W., & Li, A. 2017, *ApJ*, 840, 27
- Shields, G. A. 1978, *Nature*, 272, 706
- Sikora, M., Stawarz, Ł., & Lasota, J.-P. 2007, *ApJ*, 658, 815
- Skilling, J. & Bryan, R. K. 1984, *MNRAS*, 211, 111
- Smith, D. A. & Wilson, A. S. 2001, *ApJ*, 557, 180
- Soifer, B. T., Neugebauer, G., Matthews, K., et al. 1999, *The Astrophysical Journal*, 513, 207
- Soldi, S., Beckmann, V., Bassani, L., et al. 2005, *A&A*, 444, 431
- Stalevski, M., Asmus, D., & Tristram, K. R. W. 2017, *MNRAS*, 472, 3854
- Stalevski, M., Ricci, C., Ueda, Y., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 458, 2288
- Stalevski, M., Tristram, K. R. W., & Asmus, D. 2019, *MNRAS*, 484, 3334
- Stern, D., Assef, R. J., Benford, D. J., et al. 2012, *The Astrophysical Journal*, 753, 30
- Thiébaud, E. 2008, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7013, *Optical and Infrared Interferometry*, ed. M. Schöller, W. C. Danchi, & F. Delplancke, 70131I
- Tozzi, P., Gilli, R., Mainieri, V., et al. 2006, *A&A*, 451, 457
- Tristram, K. R. W., Burtscher, L., Jaffe, W., et al. 2014, *Astronomy and Astrophysics*, 563, A82
- Tristram, K. R. W., Impellizzeri, V., Villard, E., et al. 2022, *A&A*, submitted
- Tristram, K. R. W., Meisenheimer, K., Jaffe, W., et al. 2007, *Astronomy and Astrophysics*, 474, 837
- Tully, R. B., Rizzi, L., Shaya, E. J., et al. 2009, *AJ*, 138, 323

- Urry, C. M. & Padovani, P. 1995, *Publications of the Astronomical Society of the Pacific*, 107, 803
- van Bemmell, I. M. & Dullemond, C. P. 2003, *A&A*, 404, 1
- van Boekel, R. J. H. M. 2004, PhD thesis
- van Cittert, P. H. 1934, *Physica*, 1, 201
- van der Blik, N. S., Manfroid, J., & Bouchet, P. 1996, *Astronomy and Astrophysics Supplement Series*, 119, 547
- Vanden Berk, D. E., Richards, G. T., Bauer, A., et al. 2001, *The Astronomical Journal*, 122, 549
- Venanzi, M., Hönig, S., & Williamson, D. 2020, *The Astrophysical Journal*, 900, 174
- Vermot, P., Clénet, Y., Gratadour, D., et al. 2021, *A&A*, 652, A65
- Véron-Cetty, M.-P. & Véron, P. 2010, *Astronomy and Astrophysics*, 518, A10
- Wada, K. 2012, *The Astrophysical Journal*, 758, 66
- Wada, K., Schartmann, M., & Meijerink, R. 2016, *The Astrophysical Journal Letters*, 828, L19
- Wang, F., Yang, J., Fan, X., et al. 2021, *ApJ*, 907, L1
- Weigelt, G., Wittkowski, M., Balega, Y. Y., et al. 2004, *Astronomy and Astrophysics*, 425, 77
- Wells, D. C. 1985, in *Data Analysis in Astronomy*, ed. V. di Gesu, L. Scarsi, P. Crane, J. H. Friedman, & S. Levialdi, 195
- Williamson, D., Hönig, S., & Venanzi, M. 2019, *The Astrophysical Journal*, 876, 137
- Williamson, D., Hönig, S., & Venanzi, M. 2020, *The Astrophysical Journal*, 897, 26
- Wilson, A. S., Shopbell, P. L., Simpson, C., et al. 2000, *AJ*, 120, 1325
- Wirnitzer, B. 1985, *Journal of the Optical Society of America A*, 2, 14
- Wittkowski, M., Balega, Y., Beckert, T., et al. 1998, *Astronomy and Astrophysics*, 329, L45
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, *The Astronomical Journal*, 140, 1868
- Xie, Y., Li, A., & Hao, L. 2017, *ApJS*, 228, 6
- Yang, G., Boquien, M., Buat, V., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 491, 740
- Yang, Y., Wilson, A. S., Matt, G., Terashima, Y., & Greenhill, L. J. 2009, *ApJ*, 691, 131
- Zasowski, G., Majewski, S. R., Indebetouw, R., et al. 2009, *ApJ*, 707, 510
- Zernike, F. 1938, *Physica*, 5, 785
- Zheng, W., Kriss, G. A., Telfer, R. C., Grimes, J. P., & Davidsen, A. F. 1997, *The Astrophysical Journal*, 475, 469

Appendices

A Selected Epochs and Observing Conditions

The full table of selected observations is given in [Isbell et al. \(2021\)](#), and in particular is given in machine-readable form at <https://iopscience.iop.org/article/10.3847/1538-4357/abdfd3>.

A.1 Special Cases

A.1.1 Arp 220 – A Known Double Nucleus

Special care was taken for Arp 220 (a.k.a. IC4553) which is known to have a double nucleus (e.g., [Soifer et al., 1999](#); [Aalto et al., 2009](#)). This is the only galaxy in which we fit three elliptical Gaussians to the emission, rather than the usual two; one for each nucleus, and one for any background/extended emission. We report the flux of each nucleus separately in Table 2.2.

A.1.2 NGC 7552 – An AGN with a Starburst Ring

This (possibly dormant) AGN is located at the center of a bright starburst ring ([Forbes et al., 1994](#)). To avoid fitting starburst regions rather than the AGN, we limit the fit to within the central 1" of the galaxy.

A.1.3 Faint Sources

There were 13 sources which are marginal detections with our two-Gaussian approach, but which are clear detections "by-eye." To reduce the number of free parameters and increase the flux significance of the results, we fit the following with only one Gaussian. The FWHM of this Gaussian is set equal to that of the PSF calibrator closest in time, as in the two-Gaussian case. The AGN: 3C321, 3C327, 3C424, ESO 323-32, M87, NGC 63, NGC 986, NGC 3660, NGC 4038/9, NGC 5427, PKS 1814-63, UGC 2369 S, Z 41-20.

B Calibration Strategy

As we took many of these sources from archival programs with various setups, we had to define a consistent calibration strategy. While most of the calibrators chosen have L magnitudes in the catalog of [van der Bliik et al. \(1996\)](#), many do not have M magnitudes, and none have M_{nb} . We therefore first make the assumption that $M - M_{\text{nb}} = 0$. Secondly, we investigate the relationships between the $L - M_{\text{nb}}$ color, L magnitude, and spectral type of the calibrator, shown in Fig. B.1. We find that for stars of type O, B, A, and F the color $L - M_{\text{nb}} \approx 0$ with very little scatter ($\sigma_{L-M} \lesssim 0.05$). This results in the calibration strategy below:

```
function CALIBRATE TARGETS
```

```

for each targ observation do
  for each band  $\in [L, M_{nb}]$  do
    1) Find the calibrator (calib) observed closest in time which
      a) was observed within 6hrs of the target
      b) (has both L' and M band catalog values)  $\vee$  (is spectral type  $\in [O, B,$ 
A, F])
    2) Read  $F_{calib,real}$  from band or calculate from  $\neg band$  if necessary
    3) Compute  $F_{targ,cal} = F_{targ,raw}/F_{calib,raw} \times F_{calib,real}$ 
    4) Compute the error from the relative errors on the individual fits
  end for
end for
end function

```

where catalog values refers to the NIR catalog of ESO calibrators from [van der Bliik et al. \(1996\)](#); and $F_{targ,raw}$ and $F_{calib,raw}$ are the fitted integrated fluxes in counts for the target and calibrator, respectively. Finally, we list each target and its calibrator in Table A.1.

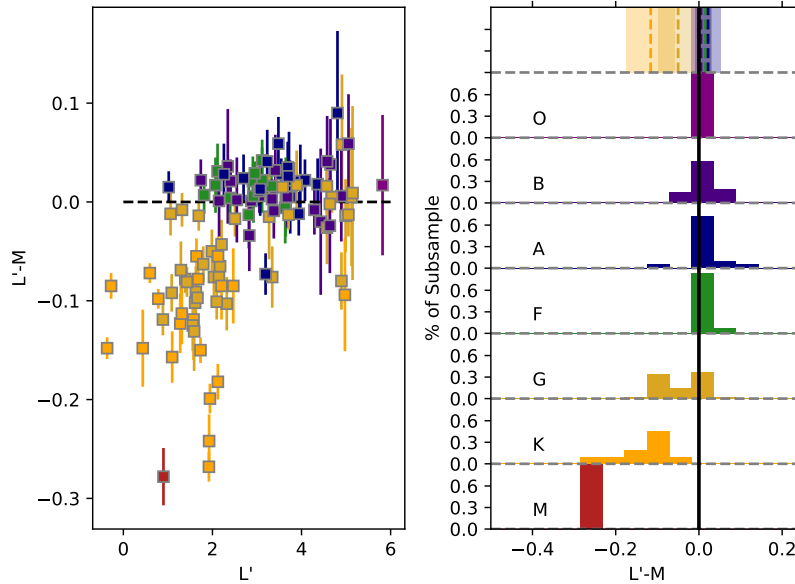


Fig. B.1: The L-M colors for all calibration stars in [van der Bliik et al. \(1996\)](#), separated by spectral type. For stars of type earlier than G, we find that $L-M \approx 0$. Bins are spaced every 0.05 mag.

C MATISSE N -band correlated fluxes

In Figs. 5.2 and C.1 we present the N -band correlated flux for each baseline, reduced and calibrated as described in Ch. 5 §2. The total photometric flux (“the zero-baseline flux”) is included in the first panel of Fig. 5.2.

D MATISSE N -band closure phases

In Fig. D.1 we present the N -band closure phase spectrum for each closure triangle, reduced and calibrated as described in Ch. 5 §2.

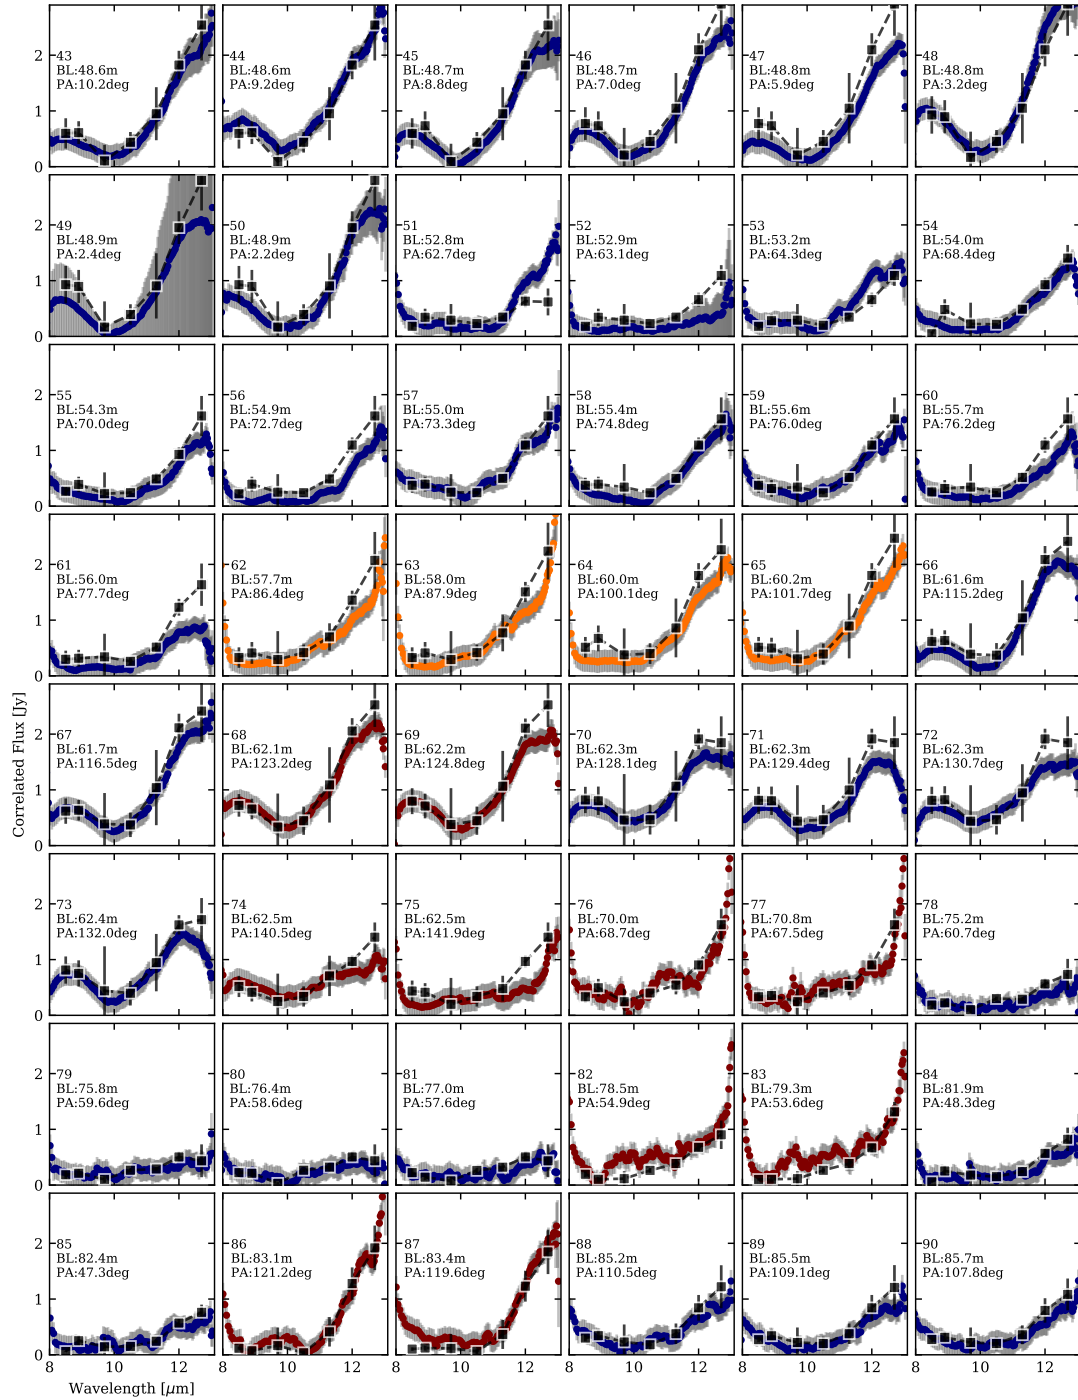


Fig. C.1: Continuation of Fig. 5.2.

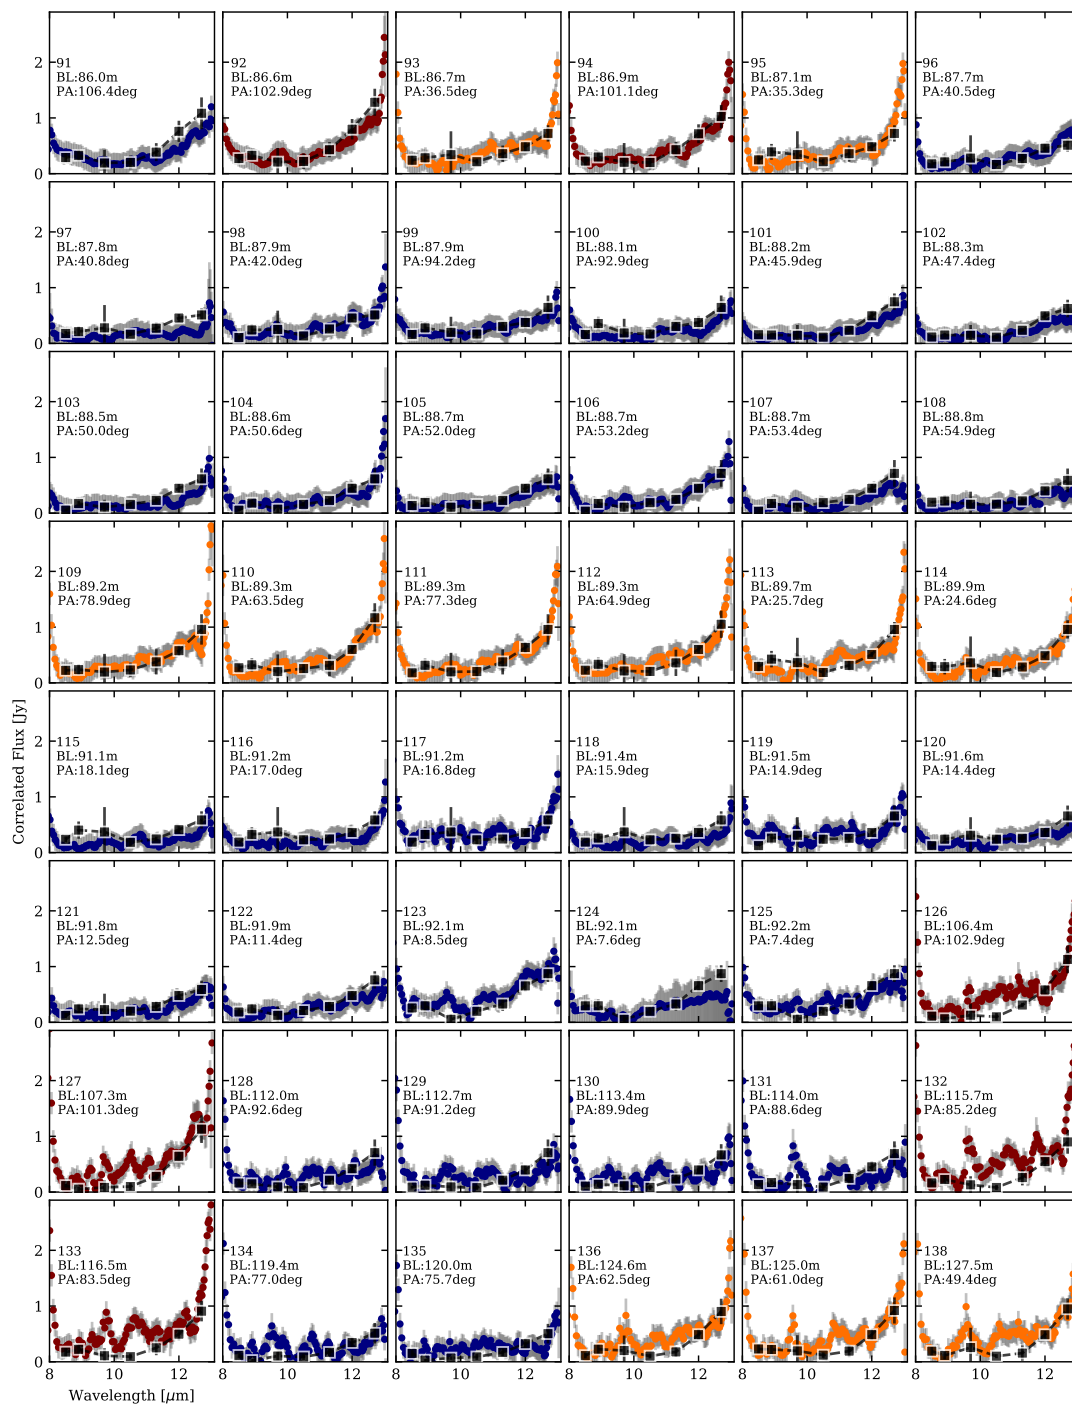


Fig. C.1: continued.

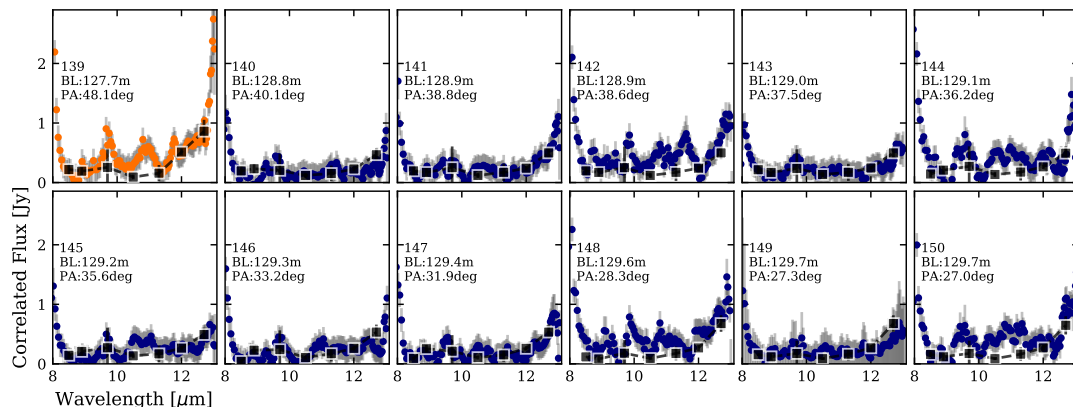


Fig. C.1: continued.

E VISIR-SAM data

In Figs E.1 and E.2 we show the measured VISIR-SAM data from the reduction as described in Ch. 5 §3.1.3.

F Dirty beam for Circinus in the N -band

We estimate the dirty beam in the typical way in order to identify image artifacts. In the uv -plane, we set the squared visibility at each uv -point we observed (± 4.1 m, the UT radius) to 1 and the surrounding points to 0. We set the phase to 0 deg across the uv -plane. We finally take the inverse Fourier transform of this complex array to obtain an estimate of the dirty beam (shown in Fig. F.1).

G Imaging with and without ATs

In Fig. G.1 we show the effects of imaging with and without the MIDI AT baselines. We stress that the MIDI AT baseline inclusion is necessary due to the resolved nature of this AGN, as shown in both MIR interferometric and single-dish observations. The MIDI AT baselines require the synthesis of closure phase triangles in order to match the IRBis formatting. We set the closure phases involving these baselines to $0 \pm 180^\circ$, such that they do not bias the imaging. We justify the inclusion of these baselines through the following arguments: First, the correlated flux values for all 30 MIDI uv -points within 4m of a MATISSE point show $< 2\sigma$ variation over 10 years (Ch. 5 §3.1.4). Second, the AT baselines from MIDI transition continuously to the MATISSE UT baselines around 30m (i.e., variations within the 0.2 Jy correlated flux uncertainties). Finally, VISIR-SAM imaging of Circinus shows $0.1 \pm 2.5^\circ$ closure phases on ≤ 6.3 m baselines (Ch. 5 §3.1.3). This agrees with T14's Gaussian modeling of the MIDI data which gives $\approx 0^\circ$ closure phases for baselines ≤ 30 m.

Nonetheless, it is instructive to see which structures arise as a result of the MATISSE-only imaging. We show the 12 μm UT-only reconstruction in Fig. G.1 alongside the Gaussian model of T14 which used both UTs and ATs from MIDI (without closure phases) and the 12 μm image reconstruction as detailed in Ch. 5 §3.2. We see that the central ≈ 1 pc is nearly identical in the two images, and notably the bright features E-W of the center remain prominent in both setups. The disk-like component is perhaps even more obvious in the UT-only image, given the same color scaling. The largest difference

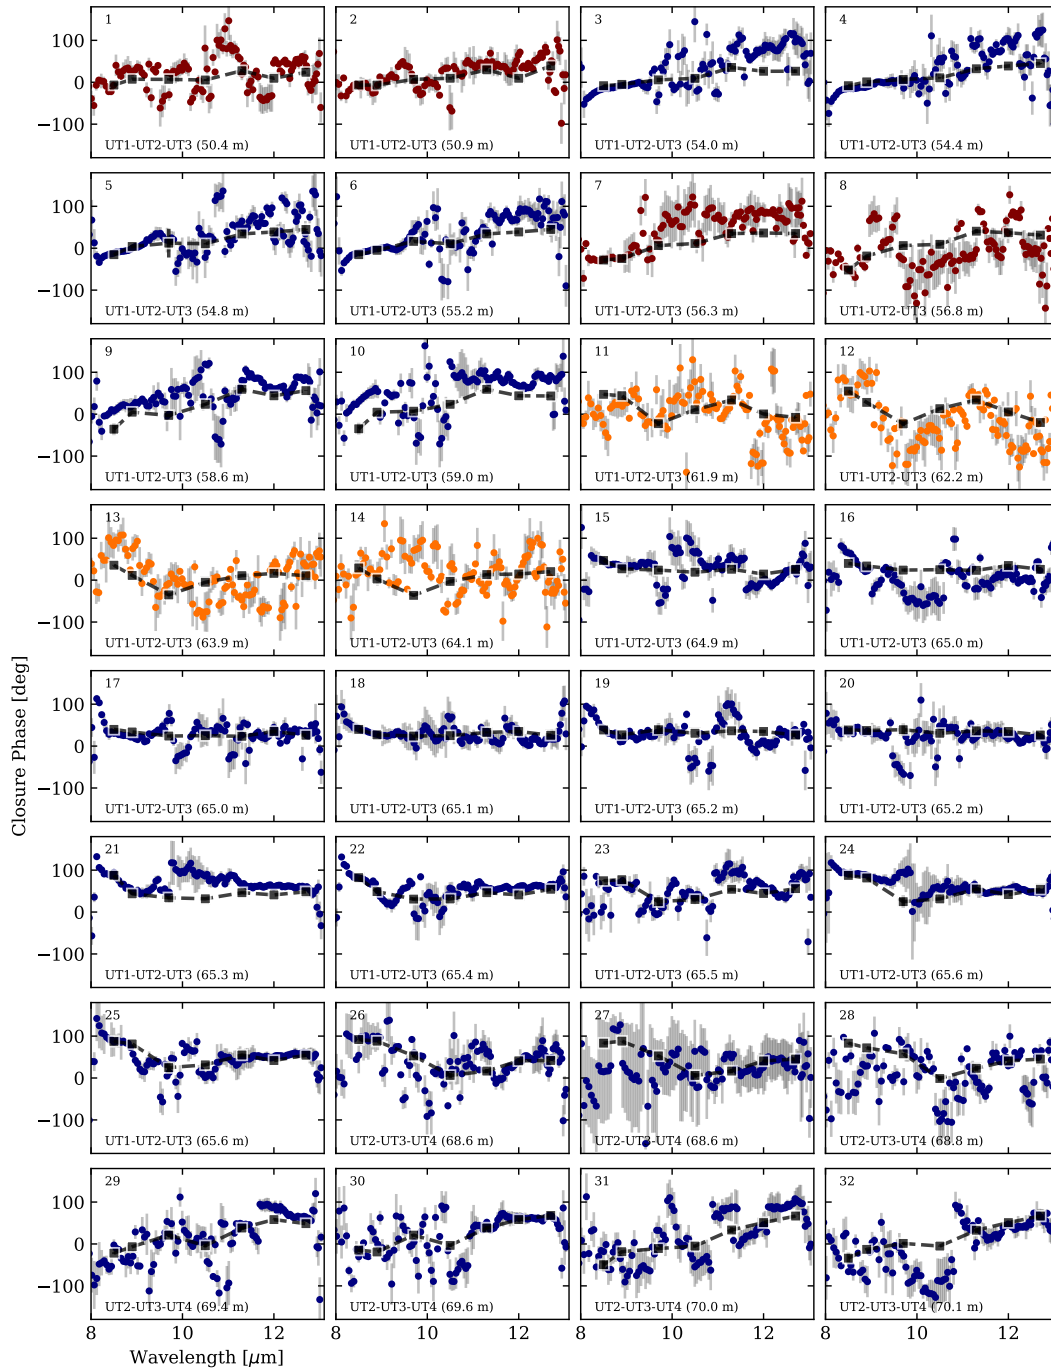


Fig. D.1: N -band closure phase data for Circinus from March 2020 (blue), February 2021 (yellow), and May 2021 (red). Presented errors come from both the calibrator phase uncertainty and the statistical variation of the observables within a set of observing cycles. The black points and the simulated values extracted from the final images, with errors estimated using the 1σ error maps (described in §3.2.2). The panels are sorted by length of the longest projected baseline in the closure triangle.

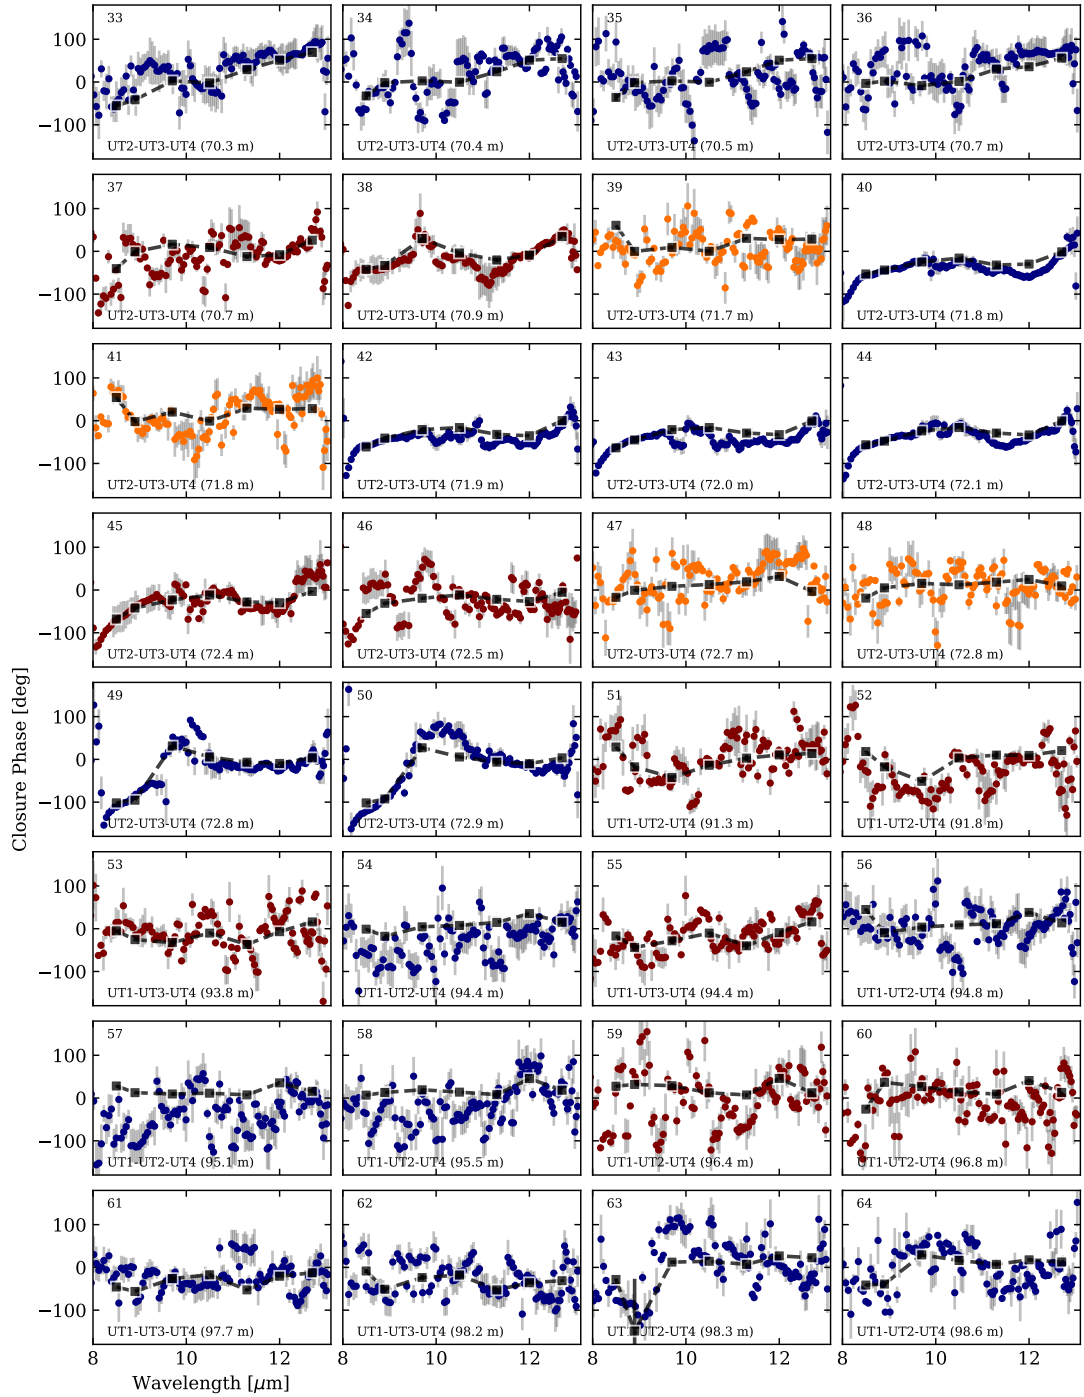


Fig. D.1: continued.

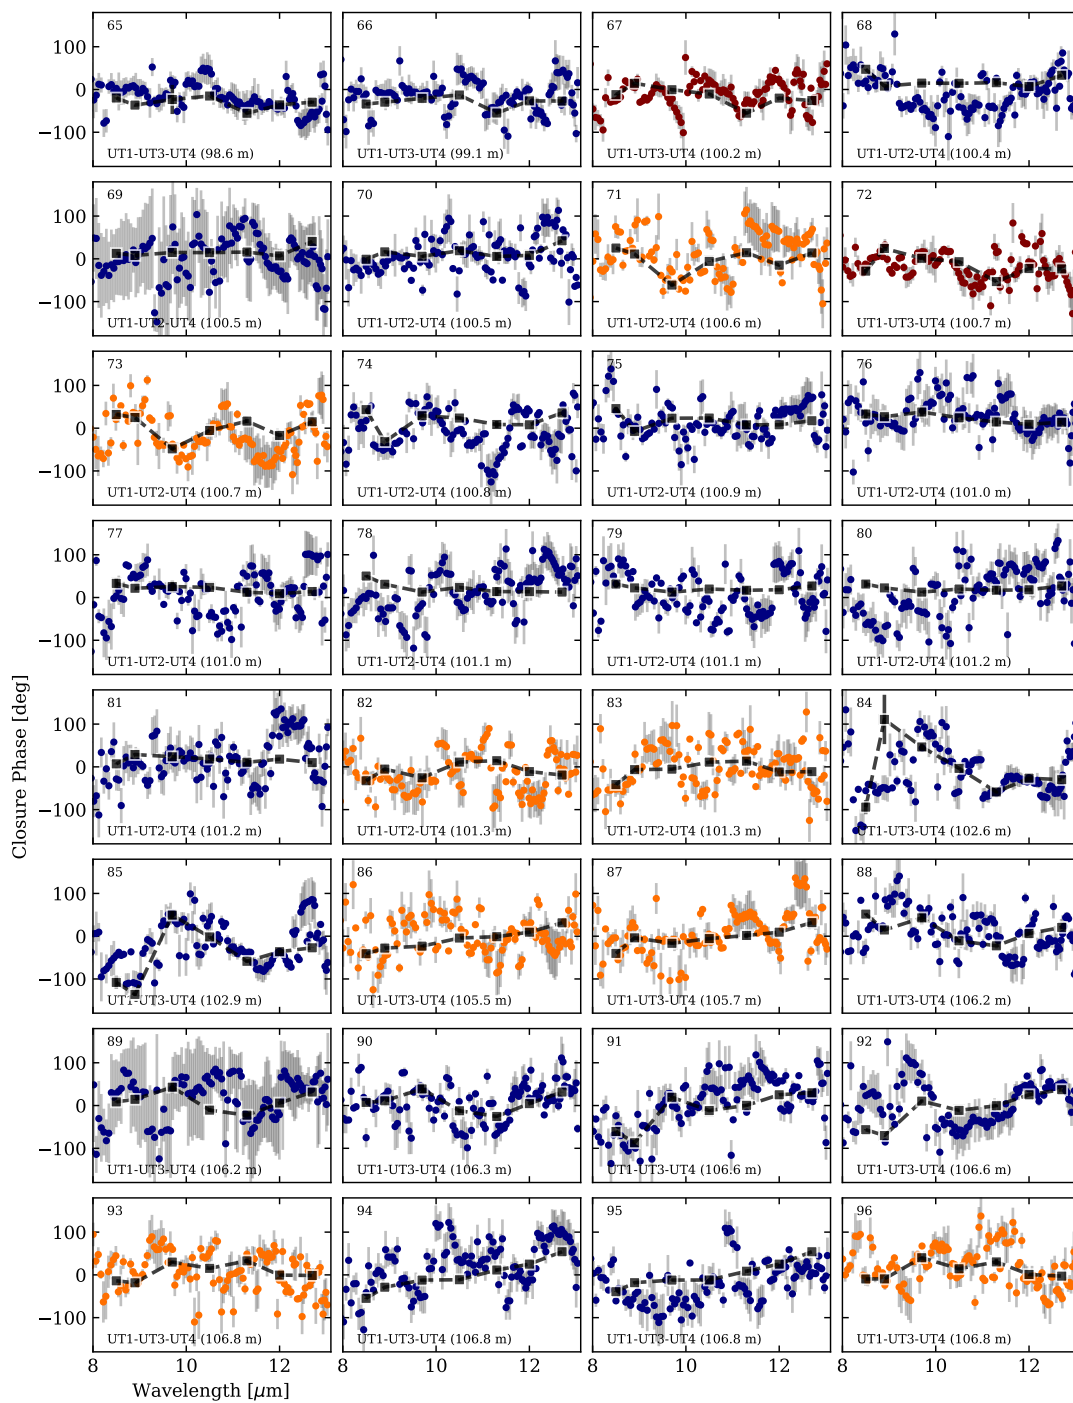


Fig. D.1: continued.




Fig. D.1: continued.

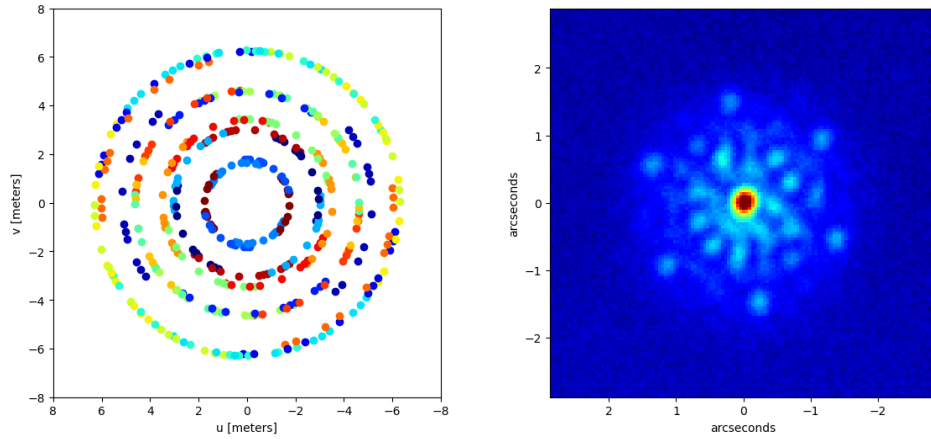


Fig. E.1: uv -coverage obtained with the VISIR-SAM data (*left*), the different colors indicate the 21 different baselines in the data. Snapshot of the Circinus interferogram obtained with the VISIR-SAM data (*right*).




Fig. E.2: VISIR-SAM data used in the diagnosis of large-scale closure phases. In the (*left*) panel we show squared visibilities. The individual observations are in gray, and the mean value over the cycles at a given baseline is in red with 1σ error bars coming from the standard deviation of the cycles. In the (*right*) panel we show the same for closure phase, with mean values in green.




Fig. F.1: Dirty beam estimated for the combined MIDI AT and MATISSE UT uv -coverage of Circinus. On the (*left*) we show the final uv -coverage, and on the (*right*) we show the resulting dirty beam with square-root scaling inside a 600 mas window. The displayed dirty beam is for the N -band but can be trivially scaled to the L -band.




Fig. G.1: Comparison of image reconstruction using different uv -samplings for Circinus. In the (*leftmost panel*) we show the image resulting from MATISSE UT uv -coverage alone. In the (*second panel*) we show the MIDI UT+AT Gaussian model from [Tristram et al. \(2014\)](#). In the (*third panel*) we show the image reconstruction resulting from the combination of MATISSE UT and MIDI AT uv -coverage. In the (*rightmost panel*) we show the results of imaging using the MATISSE UT, the MIDI AT, and the MIDI UT data, with closure phases in the MIDI data set by the T14 Gaussian model. The interior structures (a disk, an unresolved source, and bright E-W flux enhancements) are present in all image reconstructions, implying their fidelity.




Fig. H.1: Final image reconstructions and error estimates of Circinus in the N -band. From (*left*) to (*right*): the median image, error map, and S/N map for $8.9 \mu\text{m}$ as estimated from *delete-d* jackknifing. The white contour is $S/N=3$ and the black contour shows $S/N=5$.

between the images is the lack of large-scale extended flux, but this is expected as the UTs shortest baseline corresponds to $\approx 40 \text{ mas}$ at $12 \mu\text{m}$, and structures larger than this are suppressed.

H Image error estimates

We performed *delete-d* jackknifing ([Shao & Wu, 1989](#)) to estimate the errors present in our final images (see Ch. 5 §3.2.2). We present the final images, the error maps, and the S/N maps in Figs. H.1-H.7. We use the S/N maps to determine which morphological features we trust. We perform an S/N cut of ≥ 3 on the final images to *a*) define where valid apertures can be located, and *b*) determine the extent of large features.

I SKIRT model parameter variation

In Fig I.1 we compare the extracted squared visibilities and closure phases to the observed values for a range of model parameters. The simulated squared visibilities and closure phases use the uv -coverage of the MATISSE UTs. The parameters are the depth of the silicate feature in the disk and in the hyperboloid ($\tau_{9.7}$), the outer radius of the




Fig. H.2: As Fig. H.1, but for $8.9 \mu\text{m}$.




Fig. H.3: As Fig. H.1, but for $9.7 \mu\text{m}$.




Fig. H.4: As Fig. H.1, but for $10.5 \mu\text{m}$.




Fig. H.5: As Fig. H.1, but for $11.3 \mu\text{m}$.




Fig. H.6: As Fig. H.1 but for 12.0 μm .




Fig. H.7: As Fig. H.1, but for 12.7 μm .

disk (Disk R_{out}), the opening angle of the hyperboloid, the relative number of clumps (N_{cl}), and the inclination of the model (i where 90° is edge-on). The comparisons place constraints on the system inclination ($i \approx 85^\circ$), the hyperboloid opening angle ($\theta_{\text{OA}} \approx 30^\circ$), the disk Si feature depth ($\tau_{\text{Si,DSK}} \approx 14$), and the outer radius of the disk ($r_{\text{out}} \approx 3$ pc). The closure phases provide clearer constraints.

In Fig. 1.2 we show a comparison via χ^2 between the model spectra and observed spectra in each of our 13 apertures defined in §3.3. We see that flux in the central apertures is under-represented in the models, indicating that modifications to the disk component (of e.g., clumpiness or thickness) may be necessary.

J MATISSE LM -band correlated fluxes

In Figs. 5.10 and J.1 we present the LM -band correlated flux for each baseline, reduced and calibrated as described in Ch. 5 §4.1. The total photometric flux computed using the ≈ 30 m baselines is included in the first panel of Fig. 5.10.

K MATISSE LM -band closure phases

In Fig. K.1 we present the LM -band closure phase spectrum for each closure triangle, reduced and calibrated as described in Ch. 5 §4.1.

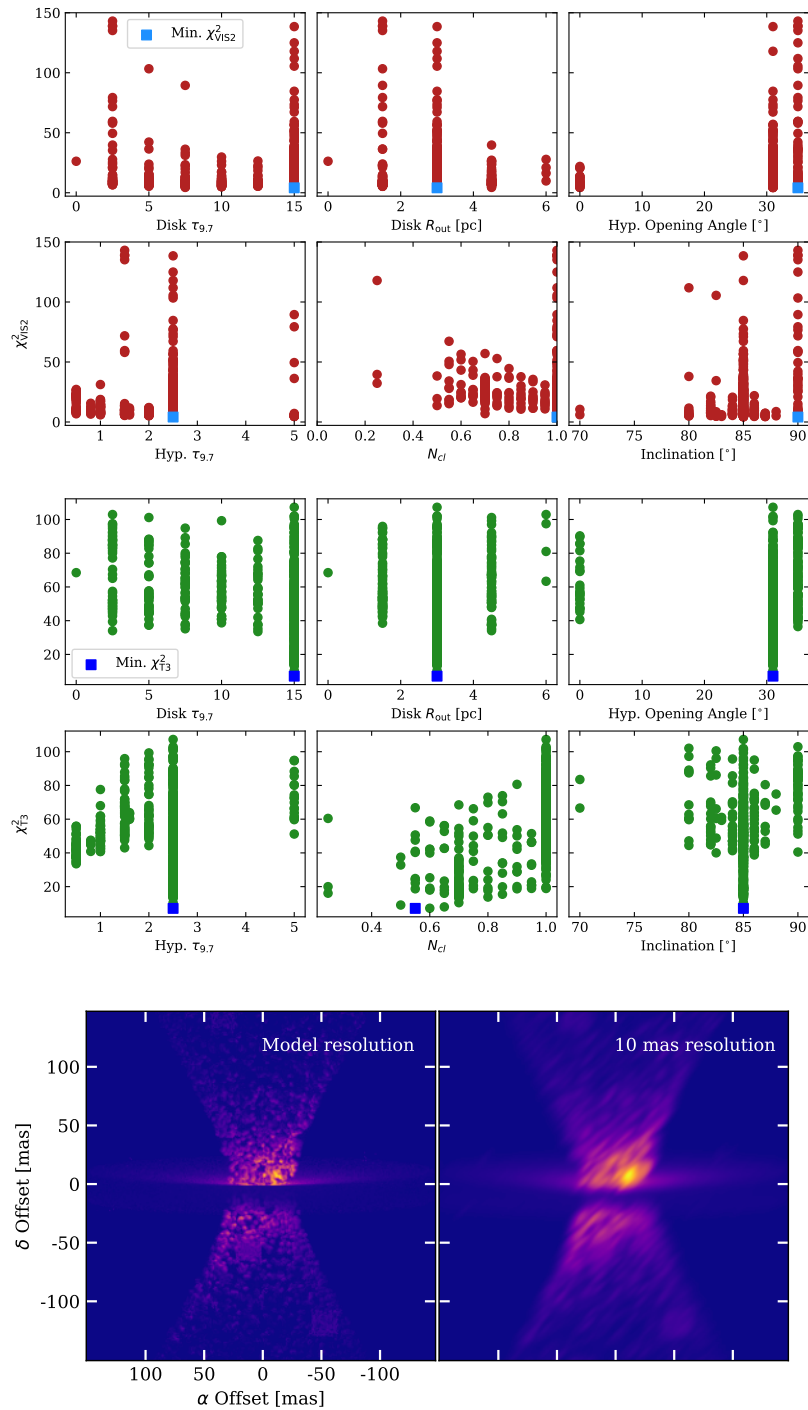


Fig. I.1: Comparisons via χ^2 of simulated observables – squared visibilities and closure phases – to the *N*-band MATISSE data of Circinus for a range of disk+hyp model parameter values. In the (*top*) six panels we show the squared visibility comparisons. In the (*middle*) six panels we show the closure phase comparisons. In the (*bottom*) panels we show the model with parameters favored by the χ^2 comparison at both its native and 10 mas resolution.

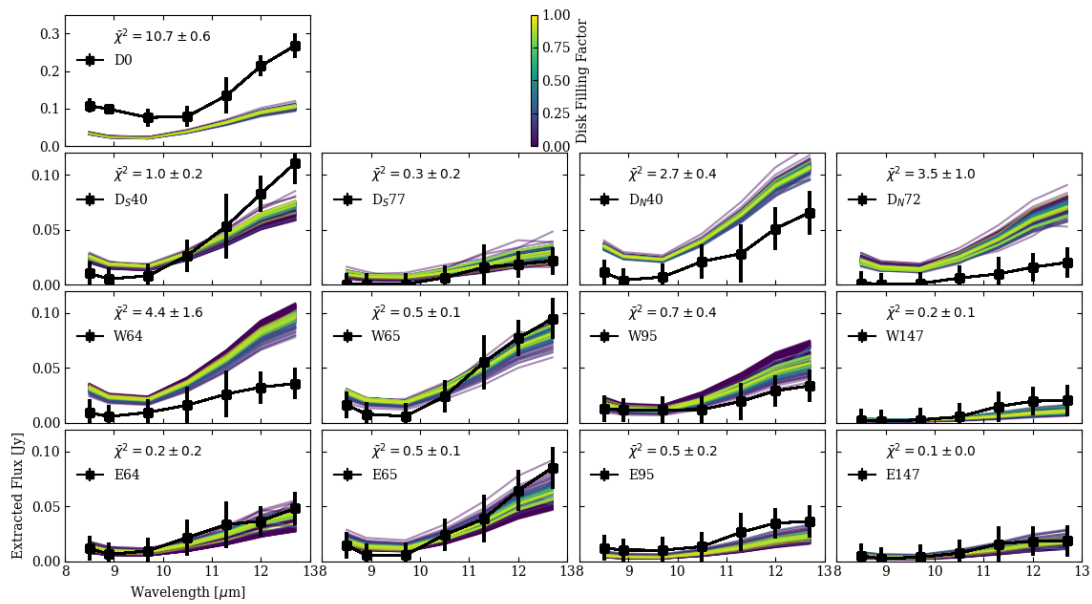


Fig. 1.2: Comparisons of measured aperture-extracted spectra of Circinus in the N -band to those of disk+hyp models with disk filling factor varied. The displayed χ^2 in each panel is the mean χ^2 of all models and the ranges are given by the standard deviation of the model values. At large radii, the models agree well with observations, but the unresolved central flux is under-represented in the models.

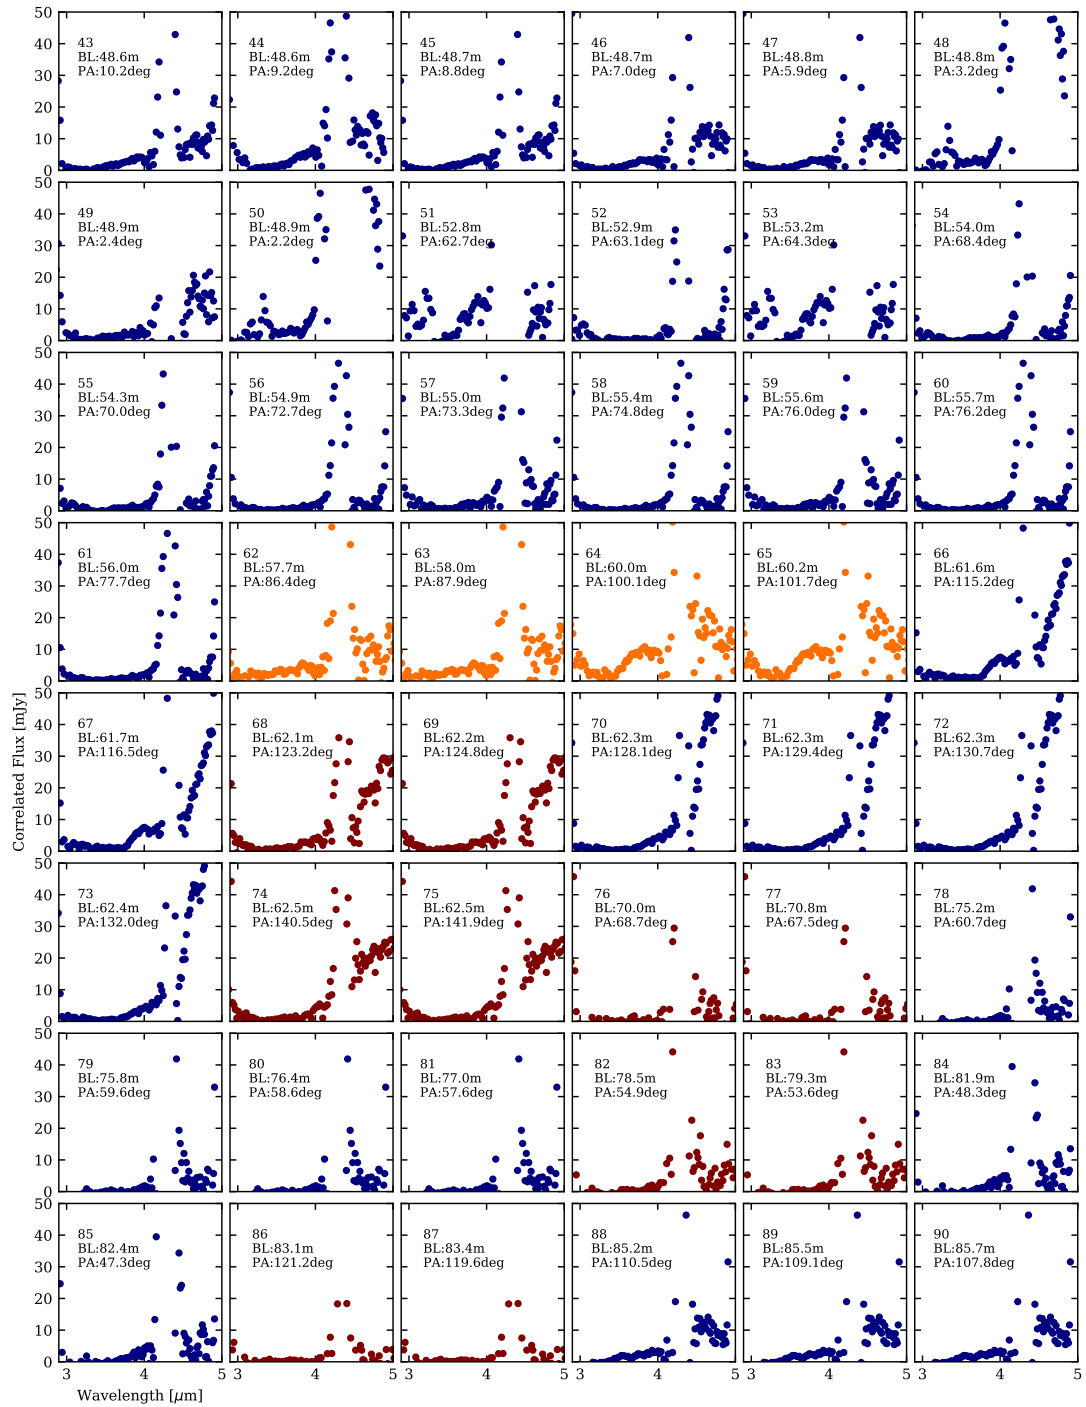


Fig. J.1: Continuation of Fig. 5.10.

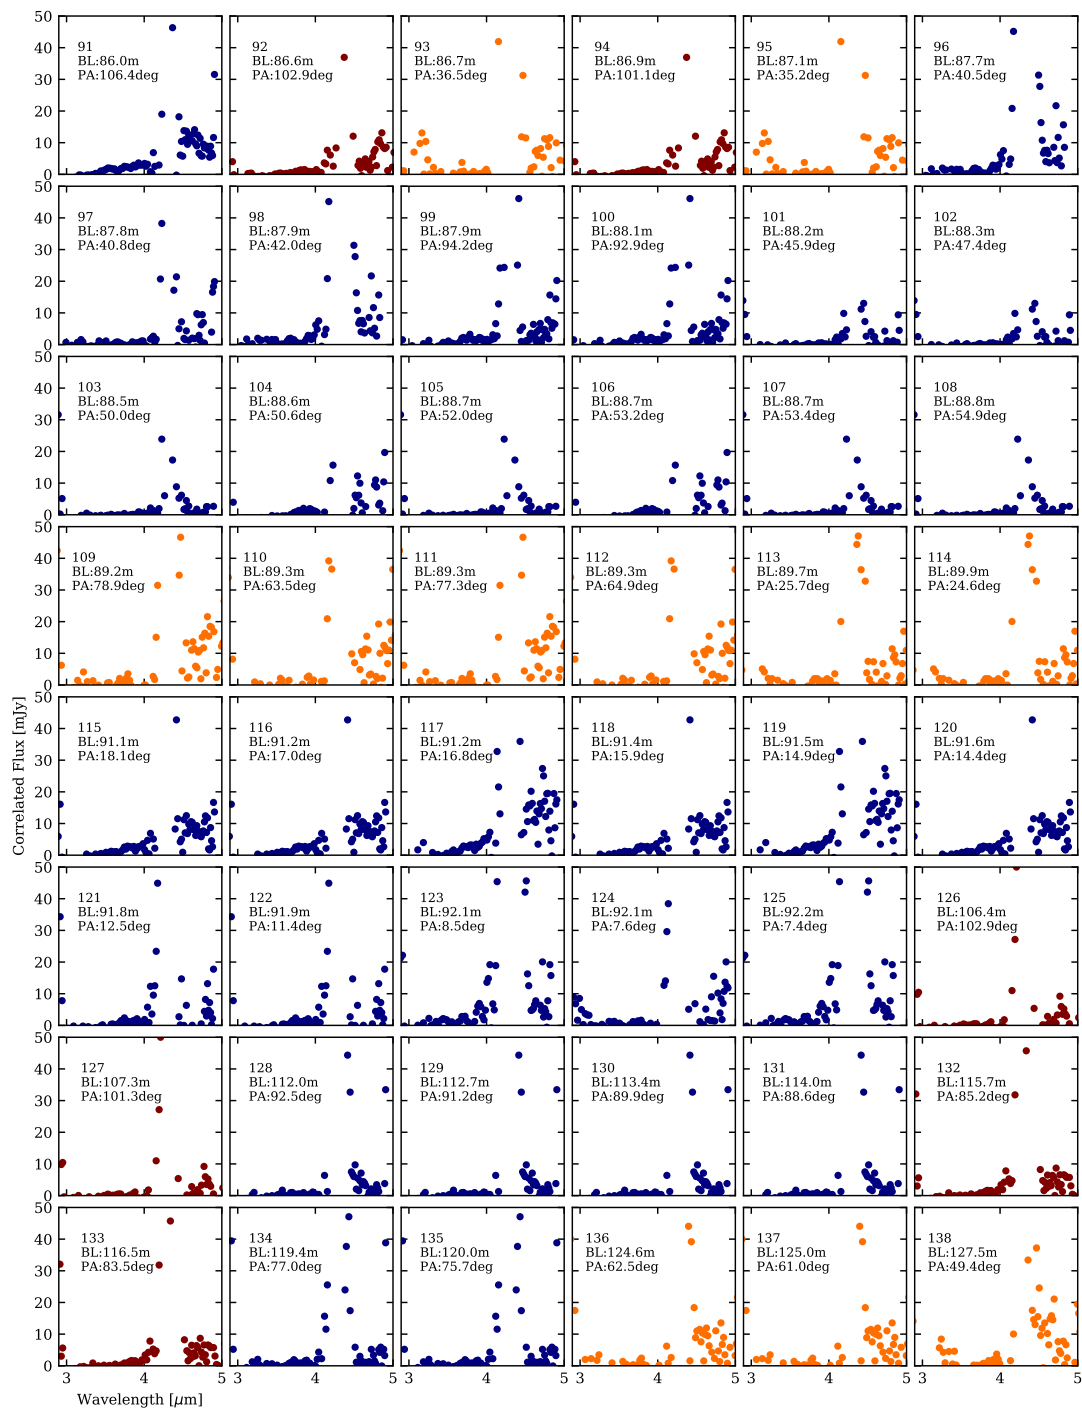


Fig. J.1: continued.




Fig. J.1: continued.

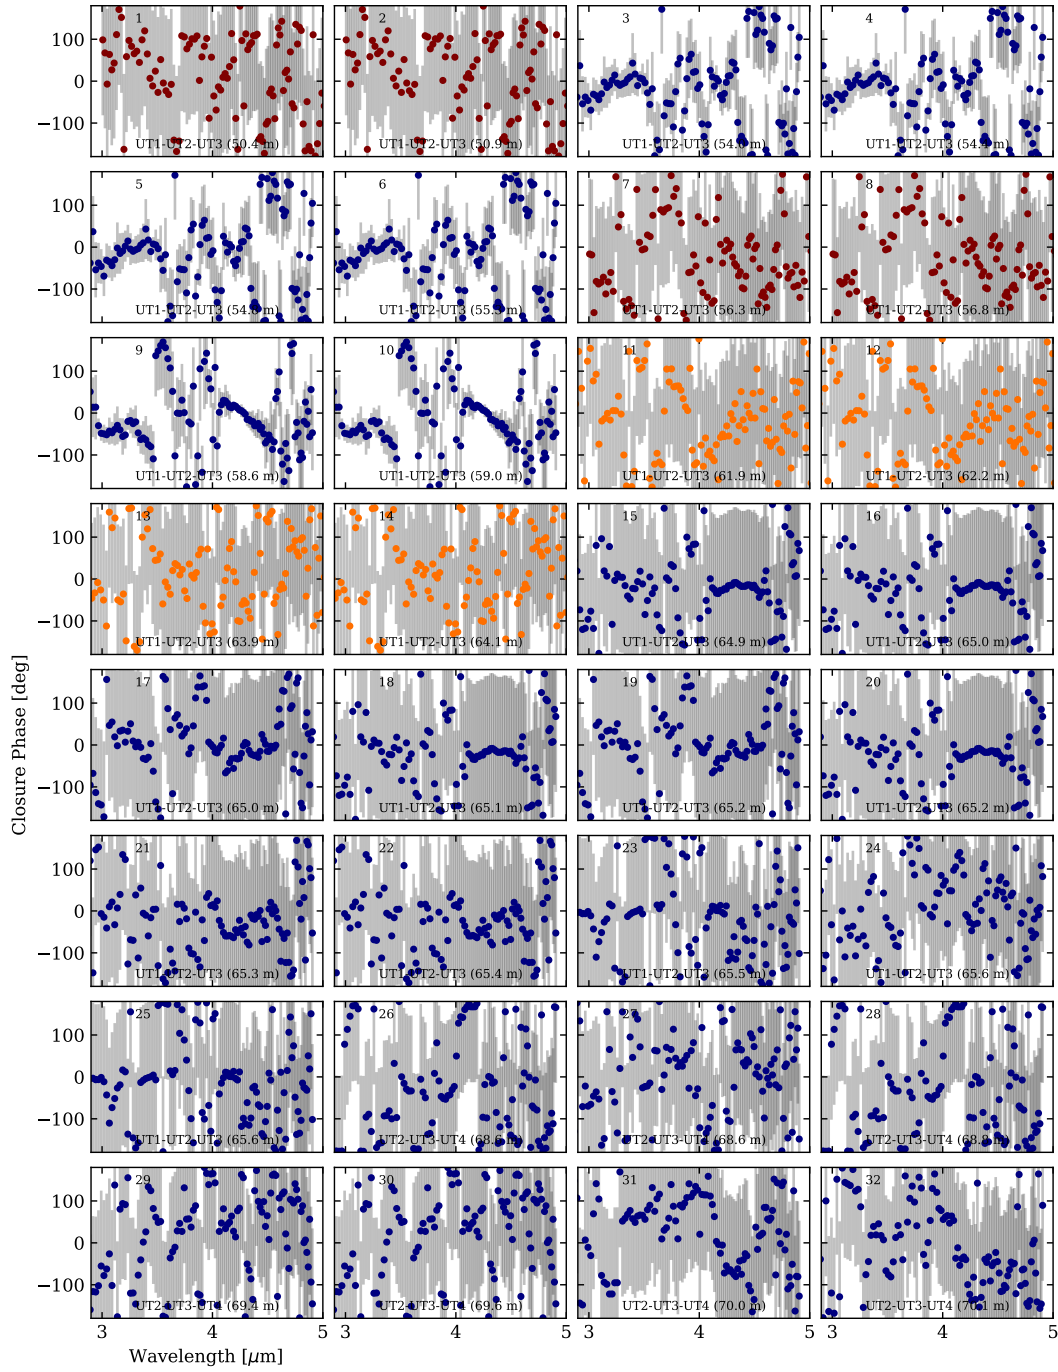


Fig. K.1: *LM*-band closure phase data for *Circinus* from March 2020 (blue), February 2021 (yellow), and May 2021 (red). Presented errors come from both the calibrator phase uncertainty and the statistical variation of the observables within a set of observing cycles. The panels are sorted by length of the longest projected baseline in the closure triangle.

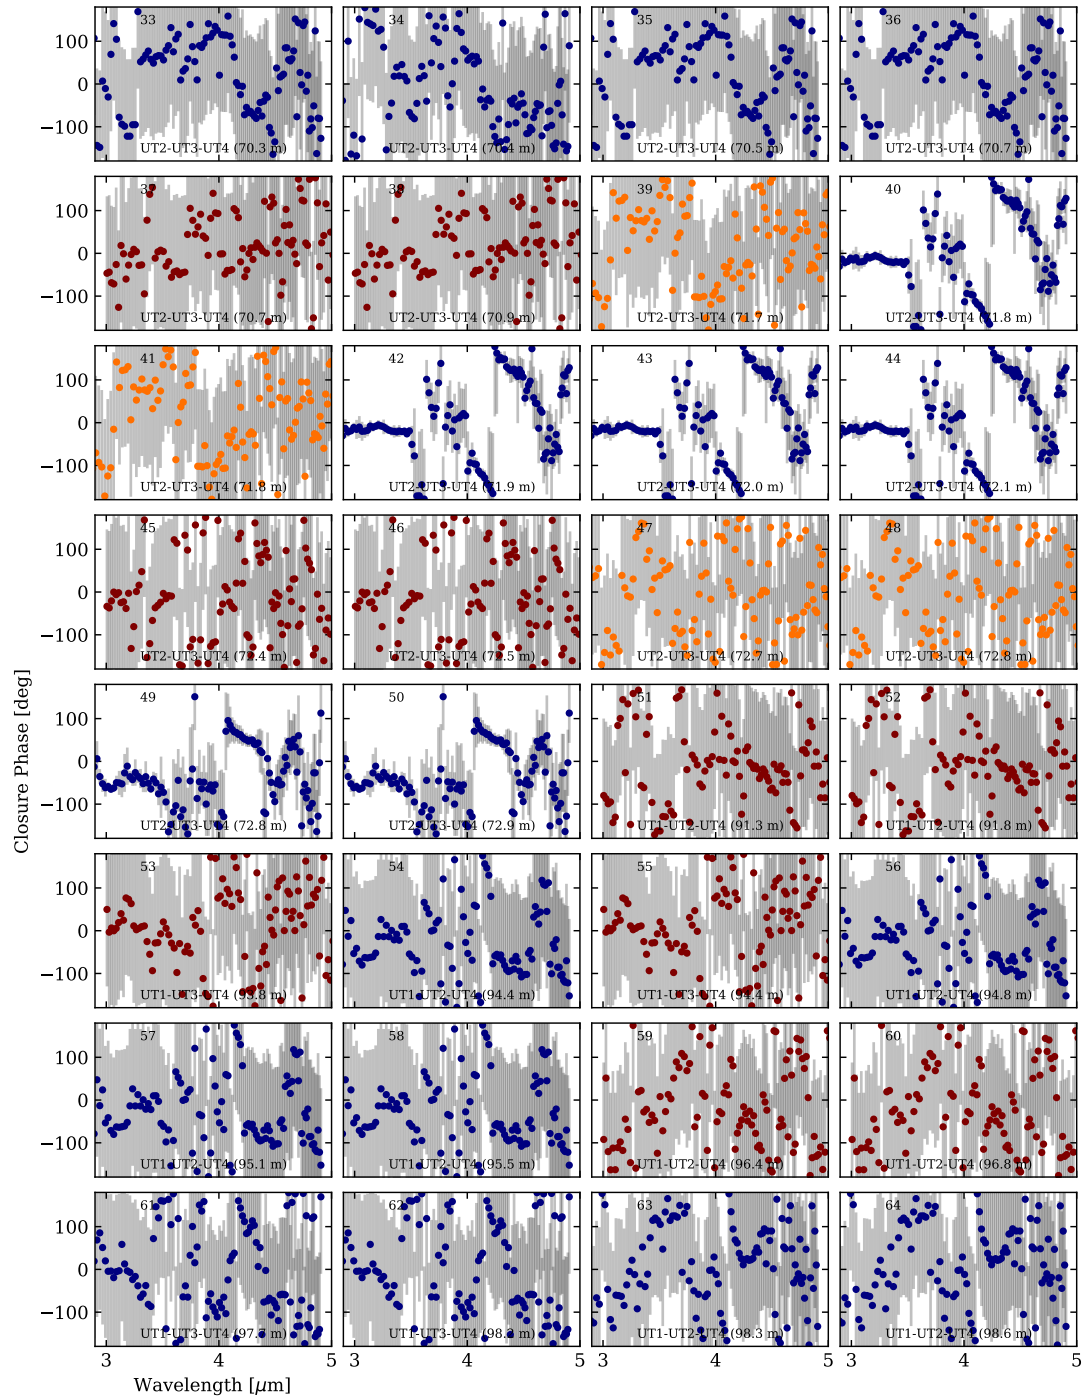


Fig. K.1: continued.

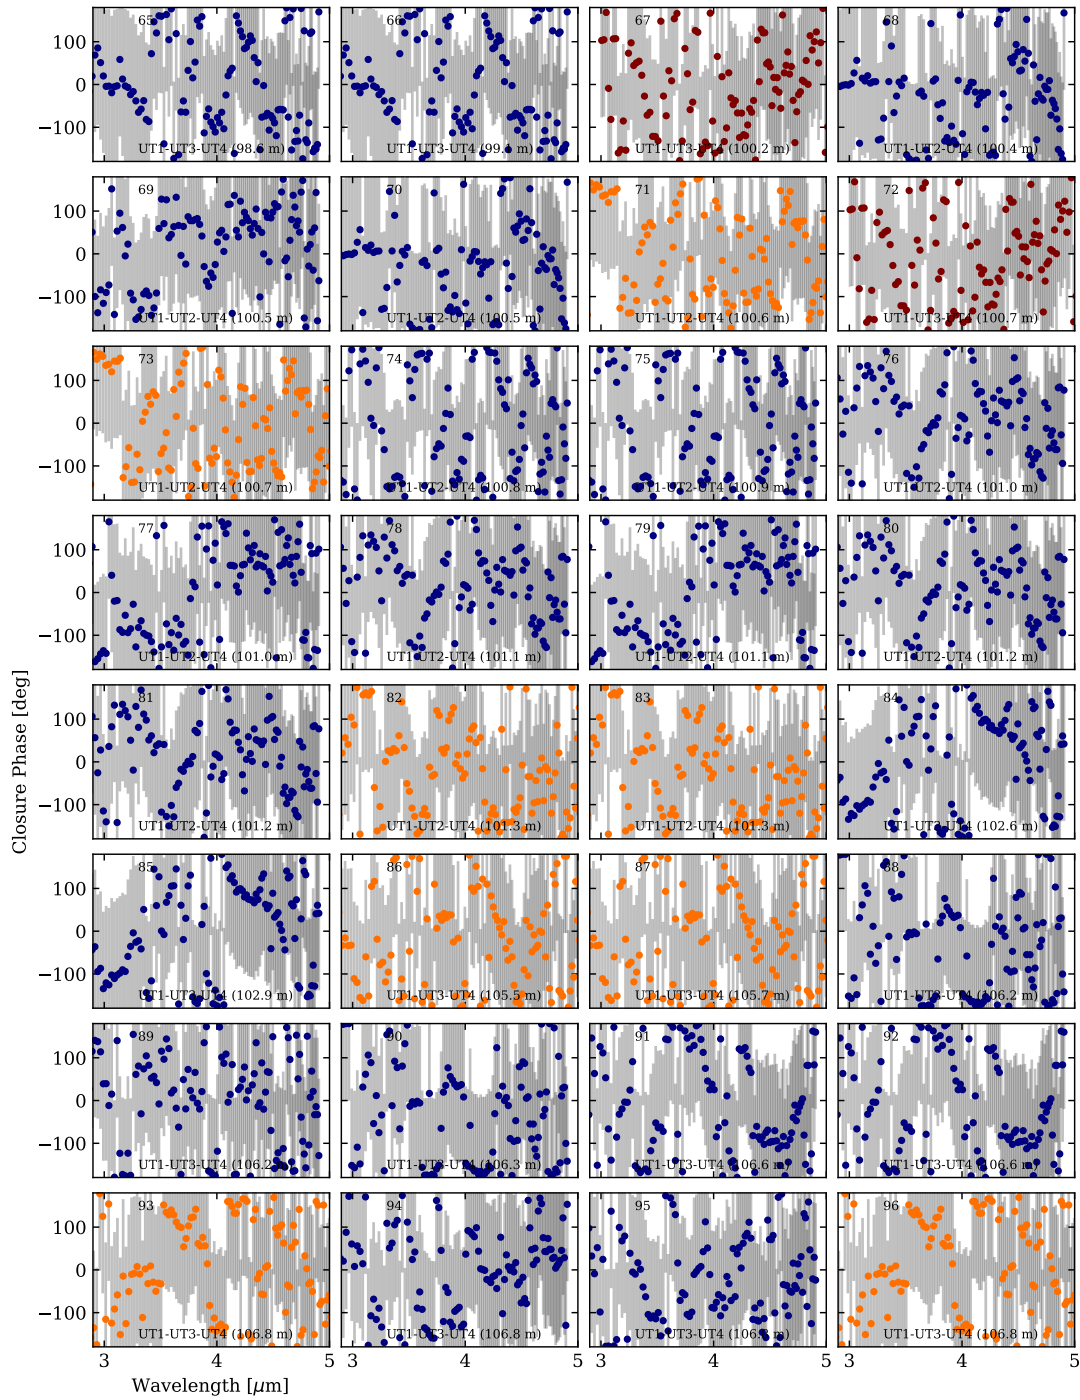


Fig. K.1: continued.




Fig. K.1: continued.

Acknowledgements

No undertaking such as this is done in a vacuum, and there are numerous people I would like to thank for their support, advice, and friendship over the last four years. I couldn't have completed this without you.

I wish to thank the referees, Klaus Meisenheimer and Cornelis Dullemond, for taking the time to read and review this (long) thesis. I also thank the anonymous referees of the papers which entered this thesis, as their comments improved the quality and impact of the results.

My supervisors, Klaus Meiseheimer and Jörg-Uwe Pott, have not only provided me with invaluable guidance, they also eased my transition to working on another continent. Their trust and confidence in me has been a constant source of inspiration and motivation over the last four years. They gave me the room to make mistakes, but crucially also gave me the tools to learn from them. Their flexibility in adjusting to remote work during the pandemic and during illness cannot be overstated, as even in these unusual times we were able to remain in close contact with vivid discussions. So to you both, herzlichen Dank!

I would like to thank my family, whose love fills oceans and whose support spans continents. My parents, Jake and Sherri, and my sister, Katie, have always had my back. My entire life, they have fostered my curiosity and given me every opportunity to try new things to find my passion. I wish also to thank my extended family: Grandma Linda, Uncle Shawn and Aunt Sara, Uncle Chris and Aunt Brooke, and my cousins Logan, Alex, Will, Candace, and Carly. Thank you to "Uncle" Bob and "Aunt" Becky for being like family to me. A special thanks goes to my Grandma Nancy, who always stood behind me and nurtured my exploration of the world around me, even when that meant breaking all her flower pots. I can never thank you all enough.

I would like to thank Domenica Barone for all of her support, inspiring curiosity, and (most importantly) patience during the completion of this work. Completing a thesis is never easy, but she made it feel easy, and she has filled the last years with joy, laughter, and motivation.

I would like to thank the entire MATISSE Collaboration for all of their hard work and assistance over the last four years. I would especially like to thank Walter Jaffe, Romain Petrov, Gerd Weigelt, James Leftley, Violeta Gámez Rosas, and Anthony Meiland for staying up on many, many observing nights both on Paranal and at our desks. To Konrad Tristram and Leonard Burtscher, thank you for sharing your expertise in AGN interferometry with me and for the fruitful collaboration.

Thank you to Christian Fendt, Huong Witte-Nguy, and the IMPRS-HD program. The IMPRS-HD program not only allowed me to work in the stimulating environment of Heidelberg, but it gave me the opportunity to meet many fantastic people and make lifelong friends. It was my pleasure to be an IMPRS representative, as I wanted to give back even a little bit to this wonderful program. I would like to thank the entire IMPRS-HD cohort, but especially Thomas Jackson, Vincent Carpenter, Lizxandra Flores, Francesco Conte, Giancarlo Mattia, Mischa Breuhaus, Marcelo Barraza, Alina

Böcker, Melanie Kaasinen, Martin Schlecker, Neven Tomicic, Gideon Yoffe, Arv Hughes, Alex Dimoff, Eric Rohr, Grigorii Smirnov-Pinchukov, and Irina Smirnova-Pinchukova. All of you have filled my stay in Heidelberg with great times and wonderful memories. Thank you also to my friends back home who never felt far away: Devin Parker, Gavin Parker, Luke Moen, and Bonnie Duong.

Thank you to the sci-an team, Oliver Völkel, Tobias Moldenhauer, and Siddhant Deshmukh. I couldn't imagine a better set of teammates to revolutionize science with. A special thanks goes to Oliver who is my closet friend in Heidelberg (quite literally because we're roommates; and yes I copied your joke). The last four years would not have been the same without you. I'm very excited to see what sci-an brings, and I'm even more excited to find out with you all.

I would like to thank the astronomer-musicians David Melon Fuksman, Gabriele Picchieri, Siddhant Deshmukh, Evert Nasedkin, Riccardo Franceschi, Paul Heeren, Vince Carpenter, and Mateo Mazzani. Together we formed three bands ("The Dead Cellars," "P and the VJs," and "Turbo Normal") played the highly sought-after IMPRS-HD Holiday Party, and filled the Untere Straße with *Chameleon*. You are all rock stars!

Thank you to Jasper Halekas, Robert Mutel, and Hai Fu for helping start my journey into research. I would like to thank Hai especially for all of his mentorship and friendship. He showed me what being a scientist was really like, he prepared me for the grueling PhD applications, and he took me on my first real observing trip. His passion for not only astronomy, but for learning in general, is infectious. I would also like to thank Brent Patterson, who helped me discover a love of science and who guided me towards becoming a researcher.

Thank you to Oliver Völkel, Vincent Carpenter, Melanie Kaasinen, and the doctors of the Uniklinikum Heidelberg for rescuing me when I had a brain bleed. Thank you to Marten Scheuck and Arv Hughes for taking the time to proofread parts of this thesis. Thank you to the German tax payer for funding my research. Thank you to the MPIA for hosting me and for the thriving research atmosphere. And finally, thank you to all of the people who post on Stack Overflow; your questions and answers made it possible for me to write all the code necessary for this work.

I am truly grateful for you all.