ALMA MATER STUDIORUM – UNIVERSITY OF BOLOGNA

School of Science
Department of Physics and Astronomy
Master Degree Programme in Astrophysics and Cosmology

Comparing Shack-Hartmann Wavefront Sensor and Donut Method for Low-Order Aberration Analysis at the 182 cm Copernico Telescope

Presented by:
Leonardo Buoncompagni

Supervisor:
Prof. Gabriele Umbriaco

Co-supervisor:
Dr. Gabriele Rodeghiero

Academic Year 2024/2025
Graduation Date 12/12/2025

Contents

Abstract

This Master's thesis investigates the low-order optical aberrations of the 182cm Copernico Telescope at the Asiago Observatory by comparing two diagnostic techniques: the Shack-Hartmann Wavefront Sensor (SHWFS) and the "Donut" method, a technique that retrieves aberrations from defocused point spread functions. In addition, speckle images were also analysed to study the aberrations induced by the atmosphere. The Donut technique represents a promising approach for measuring non-common path aberrations (NCPA) using scientific images alone, without requiring a dedicated wavefront sensor. This capability is particularly relevant for MORFEO, the Multiconjugate Adaptive Optics module for the Extremely Large Telescope (ELT), where periodic low-order realignment will be necessary during its operations.

A set of observations of bright stars was acquired at multiple elevation angles and at both the Nasmyth and Coudé focal planes of the telescope. These data allowed to characterise how low-order aberrations evolve with elevation as a result of atmospheric turbulence and gravity induced optics misalignments and deformations. Speckle images were analyzed to extract turbulence related parameters, such as the Fried radius, and to verify the WFS measurements by analysing their variations in diameter and ellipticity.

The WFS data were fitted with Zernike polynomials to allow a detailed characterisation of the optical aberrations. The behaviour of aberrations is revealed to be consistent with the combined effect of structural flexures of the telescope at low elevations and the reduced atmospheric path length near the zenith. Speckle analysis confirms this trend, showing a decrease in speckle size with increasing elevation.

The Donut method was then applied to out of focus PSFs, and the retrieved aberration coefficients were compared with those measured by the WFS. When high quality donuts were available, the method reproduced the low-order aberrations with good agreement, typically within the error bars of the reference sensor. Even in cases where the donuts were strongly aberrated, the retrieved coefficients remained within physically ranges. The code also proved to be computationally efficient, producing fits within seconds, which aligns with MORFEO's operational requirements.

Although the current implementation still requires refinement, the results demonstrate that the Donut technique is a reliable and promising tool for estimating low-order aberrations along the entire scientific optical path. With further calibration, the method has the potential to become a fundamental component of MORFEO's alignment strategy and an effective alternative, or complement, to conventional wavefront sensors in adaptive optics systems.

The thesis is organised as follows. Chapter 1 gives an introduction to the scientific context of the work, an overview of the ELT, MORFEO and the main concepts of Adaptive Optics. Chapter 2 presents the theoretical background needed to understand the methods used in the project. Chapter 3 describes the observational setups and data acquisition procedures. Chapter 4 illustrates the numerical implementation of the analysis, with a description of the codes developed for this work. Chapter 5 presents the results of the data analysis. Finally, Chapter 6 contains the conclusions and outlines possible future developments.

Chapter 1

Introduction

1.1 Research topic of the Thesis

This thesis was carried out within the framework of the development and design of the MORFEO instrument. Different observational campaigns were conducted at the 182 cm Copernico Telescope of the Asiago Observatory. The aim was to acquire and compare three types of datasets: speckle imaging, wavefront sensing, and defocused PSFs, called "Donuts", in order to characterise the low-order wavefront aberrations that arise along the telescope optical path.

These measurements were obtained at both the Nasmyth and Coudé focal planes of the telescope. Speckle images and wavefront sensor data were used to study the evolution of optical aberrations with elevation and to quantify the differences between the two foci. This comparison is particularly relevant because the Coudé focus, located several meters below the telescope, includes a long optical train with multiple reflections and refractive elements. Such a configuration provides an excellent testbed for MORFEO, which will be characterised by a long and complex optical path composed of many mirrors.

A challenge for instruments like MORFEO is the presence of non-common-path aberrations (NCPAs). After the dichroic beam splitter, the WFS and the scientific instrument no longer share the same optical path. Even if both experience the aberrations introduced by atmospheric turbulence and the telescope optics, each path contains its own additional optical elements or air gaps that introduce aberrations invisible to the WFS. Since the adaptive optics loop can only correct what the WFS measures, these NCPAs remain uncorrected and appear as residual wavefront errors on the scientific detector, degrading image quality even when the AO system is functioning correctly.

In this context, donut images were acquired to test the Donut method as a low-order wavefront sensing technique, as proposed for the first time by Tokovinin and Heathcote (2006) [1].

The idea is to compare fitted aberrations derived from the donuts with those measured by the Shack-Hartmann WFS in order to evaluate the reliability of the method as an alternative to the WFS. This validation is particularly relevant for MORFEO because the Donut technique uses only the scientific camera and therefore senses aberrations after the entire scientific optical path. As a result, it is sensitive to NCPAs and provides a direct method for measuring and correcting them.

In the end, the Donut technique can constitute a powerful tool for the periodic low-order alignment of MORFEO. Its ability to operate both on-sky and with internal calibration sources makes it particularly suitable for monitoring the misalignments induced by gravity, temperature variations, or the motion of internal components. By allowing the direct measurement of aberrations at the location of the scientific camera, the method addresses one of the central limitations of conventional wavefront sensing and has the potential to significantly improve the optical performance and stability of MORFEO during routine operations (Rodeghiero et al. (2022) [2]).

Figure 1.1: The 182 cm Copernico Telescope at the Asiago Observatory.

1.2 Adaptive Optics in Astronomy

The following introduction explains the physical problem of atmospheric seeing, the mathematical tools used to describe aberrations, how adaptive optics (AO) measures and corrects wavefront errors, the main hardware components, performance and limitations, and the principal components of AO in astrophysics.

1.2.1 Diffraction-limited and seeing-limited telescopes

The resolution limit of an instrument is defined by the Rayleigh criterion, which determines whether two point sources can be distinguished or appear as a single source. Following this criterion, two point sources are resolved if the second point is placed on the first dark ring of the diffraction pattern of the first source, corresponding to an angular separation equal to 1.22 λ/D, where λ is the reference wavelength of the incoming light from the source and D is the aperture of the instrument used to perform the observation.
When observing a star through a telescope, we expect to see an object with the angular diameter of the *Airy disk*, a diffraction pattern characterised by the formula:

$$d_{Airy} = 2.44\frac{\lambda}{D} \quad [rad] \tag{1.1}$$

This condition holds only in the absence of disturbances between the source and the instrument, as in the case of space telescopes. For these instruments we say that their resolution is *diffraction limited.*
However, for ground-based telescopes the situation is much different. Due to the presence of the atmosphere, the observed PSF of a point source (like a star) appears significantly larger. As a consequence, the telescope's resolution is worsened, and the quality of the observation is reduced. This effect is called seeing and is the result of the interaction between light and turbulent atmosphere. In this case the telescope is seeing limited.

1.2.2 Atmospheric effects

Atmosphere has a refractive index that varies in response to changes in its density, pressure, temperature, and composition. Moreover, the atmosphere is turbulent, meaning that air masses of different temperatures and densities are constantly mixing. The amount of turbulence depends on local conditions, such as orography, altitude, and solar irradiation. These effects are most prominent in the troposphere, where the atmosphere is densest. At higher altitudes, the air becomes significantly less dense, and turbulence is negligible.

Atmospheric turbulence is the primary factor responsible for degrading image quality in ground-based astronomical observations. It originates from the dynamics of the atmosphere, particularly within the troposphere, which is the densest and most energy-rich layer of the atmosphere. Solar heating, radiative cooling, and orographic effects generate convective cells, wind shear, and stratified layers with varying temperatures and humidity. These fluid motions lead to spatial and temporal fluctuations in the refractive index of air, $n(\mathbf{r}, t)$, which, in turn, introduce phase distortions on the otherwise planar wavefronts arriving from astronomical sources. The cumulative effect of these perturbations is a loss of image quality, and the fundamental goal of an adaptive optics system is to correct them in order to reconstruct, as accurately as possible, the original wavefront.

The origin of turbulence is strongly linked to the thermal structure of the atmosphere, as described by the adiabatic lapse rate, which expresses how temperature decreases with altitude. In the case of a perfectly dry atmosphere, this profile is given by the dry adiabatic lapse rate, which models the decrease in temperature with increasing height in a perfectly dry atmosphere.

$$\frac{dT}{dz} = -\frac{g}{C_p} = -9.8 \ [K/km] \tag{1.2}$$

Equation 1.2 shows that the temperature decreases by about 9.8 K per kilometer of altitude. A more realistic scenario accounts for the presence of water vapour, which modifies the cooling rate. This is expressed by the wet adiabatic lapse rate:

$$\frac{dT}{dz} = -\frac{g/C_p}{1 + (L/C_p)(dw_s/dT)} \tag{1.3}$$

Where L is the latent heat of condensation of water. Since water vapour stores additional thermal energy, humid air cools more slowly than dry air, as more heat is released during condensation.

Figure 1.2 shows the temperature profile as a function of altitude. The strongest effects of turbulence are confined within the first ~ 20 km, beyond which the atmosphere becomes too rarefied to produce significant effects.

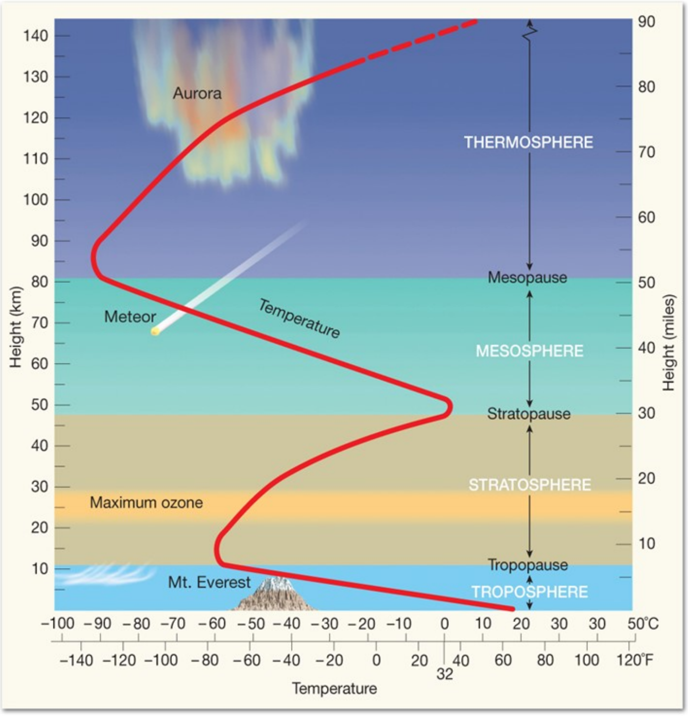

Figure 1.2: Temperature profile of Earth's atmosphere.

Atmospheric turbulence can be described using the Kolmogorov model of turbulence, where the energy contained within the turbulent eddies cascades from large-scale

turbulence down to viscous dissipation at small scales, with a proportionality of $\propto k^{-5/3}$ to the size of the eddies.
These eddies form when a hotter layer is located at a lower altitude than a colder one. In this situation, the hot layer tends to increase its height due to its lower density and mix with the colder one. This mixing is what produces the atmospheric turbulence. Each one of these eddies is characterised by a different refractive index due to the different temperatures and densities. Turbulence can be quantified by the difference in temperature between two points in the atmosphere:

Atmospheric turbulence can be statistically described using the Kolmogorov model of turbulence, in which energy is injected at large scales and cascades down to smaller eddies until it is dissipated by viscosity. The distribution of energy across spatial frequencies follows a power law $\propto k^{-5/3}$, where k is the spatial frequency associated with the eddy size. Turbulent eddies form when a warm layer of air is located under a cooler one: due to buoyancy, the hot air rises and mixes with the colder layer, producing refractive index inhomogeneities. Each eddy is then characterised by different refractive indices, resulting from local variations in temperature and density. The strength of turbulence is ultimately related to the temperature difference between two points in the atmosphere, and can be quantified through the refractive index structure coefficient:

$$Cn^2 = C_T^2[77.6 \times 10^{-6}(1 + 7.52 \times 10^{-3}\lambda^{-2}\frac{P}{T^2}]^2 \tag{1.4}$$

Where

$$D_T(r) =< (T(x+r) - T(x))^2 >= C_T^2 \; r^{2/3} \tag{1.5}$$

Meaning that the larger the temperature gradient between two layers, the stronger the turbulence.

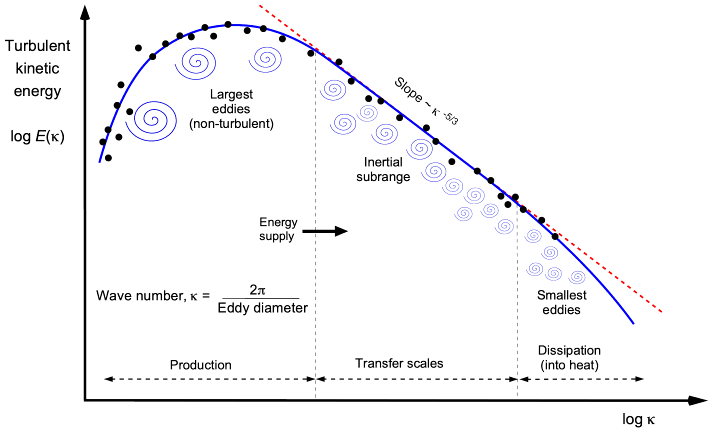

Figure 1.3: Kolmogorov model for turbulence. Credit: Leishman (2023) [3].

Other main parameters that summarise the effect of turbulence in the AO context are:

- **Fried parameter** (r_0) or coherence length. This quantity represents the characteristic size of an atmospheric turbulence cell, or “bubble,” at a given time and location. Within the same Fried radius, the incoming wavefront experiences essentially the same aberrations. The parameter depends on the refractive index structure coefficient C_n^2 (Eq. 1.4) and is defined as:

$$r_0 = [1.67\lambda^{-2}(cos\gamma)^{-1} \int C_n^2(z)dz]^{-3/5} \tag{1.6}$$

where λ is the wavelength used for the observation and γ the zenith angle. The larger the Fried parameter, the wider the region of the sky in which the aberrations remain coherent. Since r_0 scales with wavelength, its value is typically of the order of tens of centimetres in the visible, but can reach about one meter in the near-infrared. This explains why AO correction is often more efficient in the near-infrared, where the larger coherence length allows corrections to be applied over longer timescales and with improved effectiveness.

- **Isoplanatic angle**: it is the angular region within which two sources are contained inside the same turbulent "cell" and are subject to the same wavefront distortions. The isoplanatic angle can also be defined as the on-sky angular separation for which the phase difference introduced by atmospheric turbulence remains below approximately one radian (source: McKechnie (2022) [4]). When observing over a larger field of view, the phase differences beyond this angle become too large, making it impractical for a single correction to maintain good levels of performance across the entire FoV.

  It is described by
  $$\Theta_0 \simeq 0.6\frac{r_0}{h} \tag{1.7}$$
  where h is the altitude of the turbulent layer.

- **Coherence time**: quantifies the temporal stability of the turbulence. It represents the time required for a turbulent cell of size r_0 to drift across the telescope aperture under a wind speed v, and is expressed as
  $$\tau_0 \simeq \frac{r_0}{v} \tag{1.8}$$
  This parameter is particularly important because it defines the temporal scale required for AO corrections. As the turbulent structures are constantly in motion, the aberrations they introduce evolve on timescales set by τ_0. Slower wind speeds correspond to longer coherence times, which in turn allow the AO system to apply corrections over longer intervals, making the compensation process more manageable. It can vary depending on the atmospherical conditions but, for a 1.8 m class telescope, like the Copernico, t_0 is in the order of milliseconds. In general t_0 is larger for larger apertures (to cross a larger aperture, the same r_0 requires more time, assuming the same wind speed, resulting in larger coherence time).

1.2.3 How an AO system works

At its core, an AO system runs a closed-loop feedback cycle:

- Measure the distorted wavefront with a Wavefront Sensor (WFS).

- Compute the correction to the WF shape with a Real-Time Computer (RTC).
- Apply the corrections with the Deformable Mirror (DM) to cancel the measured aberration.

This cycle repeats at a rate determined by the coherence time of the atmosphere (τ_0) and for this reason, larger τ_0 are preferred.

The application of AO is fundamentally a signal processing problem: WFS measurements can be noisy and partial, the reconstructor must mitigate noise, and the DM has a finite stroke range and actuator density. The AO correction behaves like a high-pass filter on WFE modes: modes below a certain cut-off are corrected, those above remain uncorrected and are left as residual noise in the final image. The cut-off depends on the quality of the AO loop, which in turn is determined by the sensitivity of the WFS, the time that the RTC need to compute the corrections and the speed, precision and degrees of freedom that the DM has.

Deformable mirrors

Deformable Mirrors (DMs) represent the main corrective element of an AO system. They are responsible for applying the actual correction to the incoming distorted WF in order to restore image quality. A DM is in general composed by a reflective and flexible membrane that can be deformed repeatedly without damage. The deformation of the mirror is controlled by an array of actuators, which are the components that translate the corrections computed by the RTC into changes in shape of the mirror.

Actuators come in different forms, such as piezoelectric stacks, micro-electro-mechanical systems (MEMS), or magnetic actuators. In all cases, their purpose is to impose localised deformations on the mirror that counteract the distortions of the incoming WF. The idea behind the DM is to superimpose many small deformations across the DM surface to reproduce the conjugate of the aberrated wavefront, flattening it and restoring (as much as possible) the ideal diffraction-limited image quality.

The number of actuators, their spatial location, and their dynamic range (the maximum extent of the stroke that they can impose on the mirror) define the correction capability of the DM. For example, a higher actuator density allows the system to correct for finer-scale aberrations, while actuator stroke determines the maximum amplitude of distortions that can be corrected.

The main kinds of actuators are:

- **Piezoelectric stacked array and bimorphs:** they work with the piezoelectric effect, which allows some particular materials to extend and contract when an electric voltage is applied. The first is composed by an array of small ”pistons” placed under the reflective surface and elongate or contract depending on the voltage applied to them. In the second case, instead, we have a piezoelectric plate that changes shape with the voltage. Both of them can imprint a stroke of a few microns.
- **MEMS**: micro electro-mechanical systems. They are based on electrostatic attraction between the membrane and the electrodes placed under it triggered by the application of a voltage. MEMS are characterised by a very high actuator density, sub-nanometre repeatability and rapid response.

- **Voice-coil (magnetic actuators)**: those appear as an array of electromagnets placed under a reflective membrane with small magnets attached to it. The application of a current activates the electromagnets, which attract or repel the magnets on the membrane, allowing it to change its shape. This system is used in large DM (e.g., ELT M4 and LBT secondary mirrors). These actuators are capable of very large stroke (hundreds of microns) and fast dynamic response, particularly suitable for large telescopes which need faster corrections.

Natural and Laser guide stars

AO systems rely on WF measurement as a first step toward the application of corrections. In order to measure the WFE, the WFS needs a light source to observe and these sources are called guide stars. To effectively perform WF measurement a nearby natural guide star (NGS) is required within the isoplanatic angle if the scientific target (the source that the telescope is observing to retrieve data). In this way both the star and the science target are subject to the same atmospheric turbulence and therefore the corrections valid for the NGS will be valid also for the other target. The probability of finding a guide star depends on two factors: the stellar density in the observed part of the sky (higher in the Galactic bulge, much lower at the poles), and the limiting magnitude of the WFS, which in general struggle when observing faint sources. This probability is called Sky coverage.

Sky coverage can be as low as a few per cent in regions far from the galactic bulge, where bright sources become rarer. For large telescopes the situation improves since their larger apertures collects more photons from fainter stars but, nevertheless, sky coverage represents one of the main limitations for an AO system.

To overcome the problem of sky coverage, laser guide stars (LGSs) are used. LGSs are artificial sources generated by projecting lasers into the atmosphere. A much-used method uses the excitation of the sodium layer at 90–100 km altitude using lasers. This causes a fluorescent emission in the sodium atoms and the light coming from this emission in used as a guide source. Another kind of LGS relies on the Rayleigh scattering of light at the uppermost layers of the troposphere ($\sim 10 - 12 \quad km$).

However, even LGSs has some limitations like the elongation of the source (when projected off-axis, the LGS appears elongated, reducing the accuracy of the measurements), cone effect (differently from natural stars, whose light comes from infinity, LGS sample only a cone of atmosphere above the telescope, not allowing the sensing of the turbulence at higher altitudes) and lack of tip–tilt information (since the up-going and down-going paths of LGS are reciprocal it cannot give information on tip–tilt).

Due to these limitations, AO systems often use both NGSs and LGSs, the first to have a good sampling of the aberrations of the entire FoV of the telescope and the second to obtain information on tip-tilt and integrate the sensing of the entire air column above the telescope.

Strehl Ratio

Strehl Ratio (SR) as a performance metric for an AO system. It is defined as the ratio of the observed peak intensity in the PSF to that of an ideal diffraction-limited system. A

Strehl ratio above 0.8 is considered the threshold for "diffraction-limited" performance in an AO system. Lower Strehl values correspond to broader PSFs, reduced contrast, and worse image quality. SR depends on target brightness, τ_0, RTC speed, DM actuator density and speed so it is a quantity which represents a complete summary of the quality of all the components in an AO system.

AO Modes

Adaptive optics comes in a wide set of modes with different performances and characteristics:

- **Single Conjugate AO (SCAO)**: these systems use one WFS and one DM. They achieve high correction ($0.3 - 0.8$ SR) but only inside a small angle around a guide star (tipycally a FoV of $20''$).
- **Ground-Layer AO (GLAO)**: uses both NGS and LGS to obtain the most precise correction of boundary-layer turbulence (which is the lowest layer of the troposphere and the most turbulent one). GLAO is characterised by a few arcminutes FoV but is capable only to reduce the effect of seeing without reaching high SRs.
- **Laser Tomography AO (LTAO)**: uses one DM, then multiple LGSs, NGSs and WFSs conjugated with different layers of the atmosphere in order to measure the aberrations in the entire air columns above the telescope. It can reach a few tens arcminutes in FoV and SR of between 0.1 and 0.5.
- **Multi conjugate AO (MCAO)**: this the most complicate AO mode. It is composed of 2 or more DMs conjugated to different atmospheric layers and multiple LGSs, NGSs, and WFSs to obtain a large corrected FoV (in the order of arcminutes) and can compensate for turbulence at different altitudes.

Each mode has it limitations in corrected fiel, Strehl ratio, sky coverage, and complexity. As an example, MCAO is overall the best mode because it can reach relatively high SRs over large FoVs, while other modes like SCAO are capable of higher peak performance but the image quality degrades noticeably over large FoVs. So the best mode depends in the end on the specific scientific case that we are interested into.

Uses of AO in Astronomy

AO allows and increase the quality of scientific observations in many areas like:

High-resolution imaging of stellar populations and crowded fields (resolving stars in other galaxies and clusters).

Direct imaging and characterisation of exoplanets requires high contrast in very small separation angles between the host star and the exoplanet.

Precision astrometry, to study the galactic centre dynamics, for example.

Spectroscopy with very high resolution (improves signal-to-noise for faint targets).

1.3 MORFEO and ELT

This Section explains the characteristics and design of the **Multi-Conjugate adaptive Optics Relay For ELT Observations** (MORFEO): the MCAO relay for the

Extremely Large Telescope. MORFEOis an instrument that will be part of the Extremely Large Telescope (ELT), the largest telescope ever built, currently under construction. MORFEO is an MCAO module that will receive the ELT light beam at the Nasmyth platform, perform wavefront reconstruction and correction and deliver a corrected field to subsequent ELT instruments like Multi-AO Imaging Camera for Deep Observations (MICADO), an imager and spectrograph and a future instrument that will occupy the relay's second exit beam. MORFEO is designed to bring the ELT's 39-meter aperture resolution as close as possible to its diffraction limit over an extended FoV (around $1'$), allowing for observations that require very high angular resolution and high sensitivity at near-infrared wavelengths.

The ELT's 39-m primary mirror will have an exceptional diffraction limit and light-collecting capabilities, but atmospheric turbulence will degrade its performance if not corrected with AO. SCAO can reach very high SR on axis, but these corrections rapidly degrade off-axis. Many ELT scientific objectives (e.g. resolved stellar populations, precise astrometry in crowded fields and high-contrast imaging) require correction over a significantly larger field than a SCAO system can provide; for this reason, an MCAO system is required.

MORFEO's MCAO will be constituted by a combination of the ELT adaptive mirror acting as the ground-layer corrector and two DMs conjugated to higher atmospheric layers ($\sim$6.5 km and $\sim$17.5 km). By sensing turbulence with six LGSs and three NGSs it will reconstruct the turbulence distribution above the telescope (performing a tomography) and will apply corrections at the three conjugated layers. The instrument requirement is to achieve a Strehl Ratio of $\geq$ 30% across a large FoV ($\sim$ 1 arcmin) in median conditions with $\sim$ 50% sky coverage (Busoni et al. (2023) [5]) in the K-band.

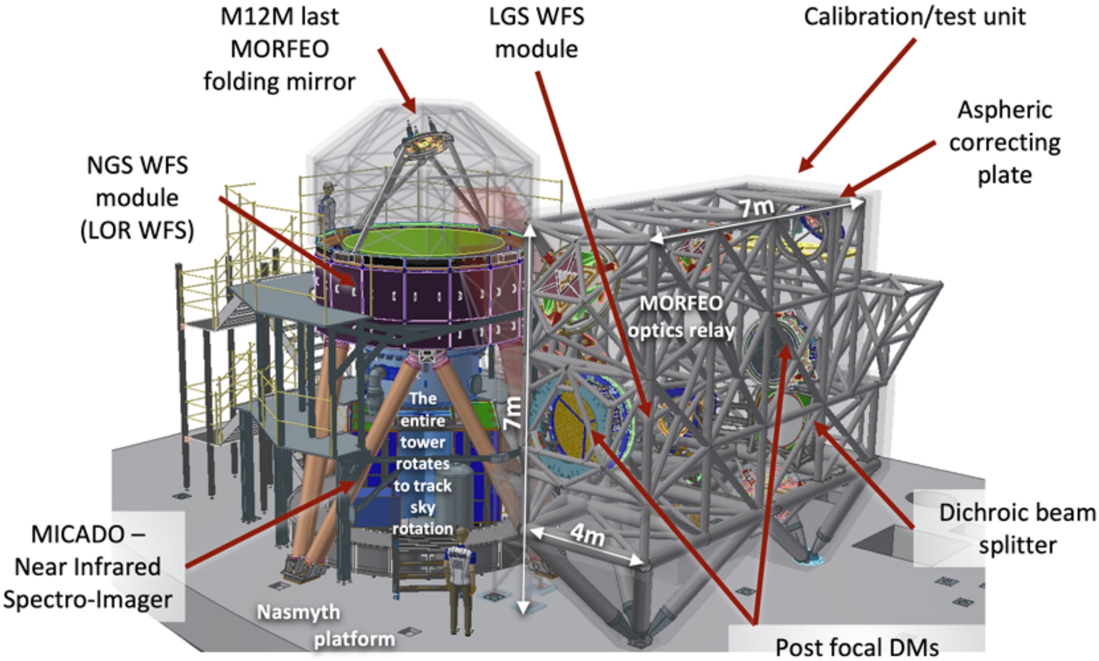

Figure 1.4: Scheme of MORFEO. Credits: Busoni et al. (2023) [5]

MORFEO layout

MORFEO (Fig. 1.5) will be installed on the ELT Nasmyth platform, between the telescope and the scientific instruments. Its main goals are to intercept the ELT beam, perform adaptive correction to it and feed it to MICADO and a future second instrument through a flip mirror. MORFEO will host both the laser guide star and natural guide star wavefront sensor modules, and it will also coordinate the operation of the ELT adaptive mirror, M4.

In the optical path, the incoming beam from the ELT focal plane is folded multiple times through MORFEO via the mirrors M6–M12. Along this path, the beam passes through two deformable mirrors, M9 and M10, which are conjugated to different atmospheric layers ($\sim$6.5 km and $\sim$17.5 km). These, together with the ELT M4 compose the MCAO system. After M10, a dichroic mirror separates the LGS light from the science and NGS channels. The science beam is then directed by a flip mirror (M11) towards MICADO or the second instrument port.

MORFEO's wavefront sensing is based on measurements from six LGS, providing high-order information on the atmospheric turbulence, and three NGS, which sense low-order modes.

The RTC receives the signals from all WFSs, computes the wavefront errors, and reconstructs the turbulence profile through tomographic algorithms that separate high and low order components of the aberrations. With these computations, the RTC finds the shapes to be applied to the three DMs in real time in order to actually perform the adaptive correction.

MORFEO is also quipped with a dedicated Calibration and Test Unit that provides artificial sources to allow also daytime calibration and maintenance, without requiring telescope night-time. (source: Busoni et al. (2023) [5]).

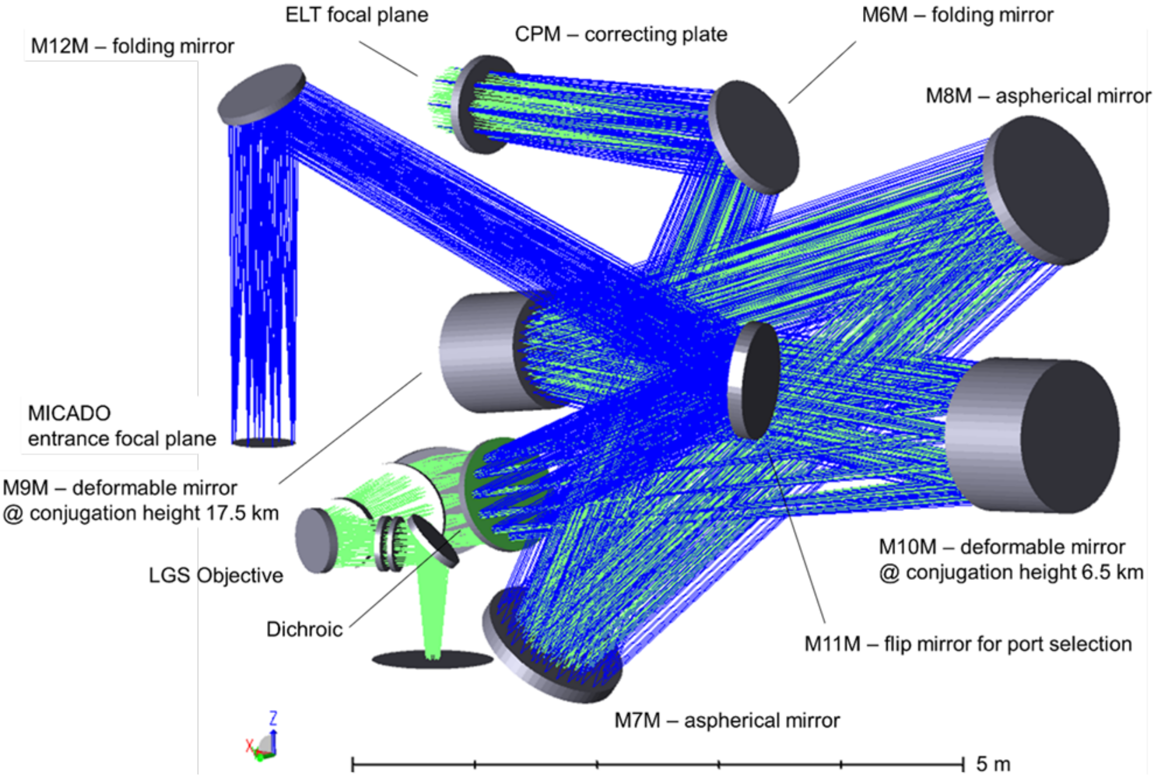

Figure 1.5: Optical design of MORFEO. Credits: Busoni et al. (2023) [5]

1.4 MORFEO and ELT science case

The scientific motivation for the ELT covers almost all major fields of modern astrophysics, from the characterization of exoplanets to the study of the earliest galaxies at the edges of the observable Universe. The ELT, in combination with MORFEO (Fiorentino et al. (2019) [6]), will enable a series of observations that are currently impossible with the existing facilities (ESO (n.d.) [7]).

One of the most anticipated scientific cases is the study of exoplanets: the ELT will be able to detect and characterise planets up to Earth-like masses through radial velocity measurements. In addition, it will perform spectroscopy of planetary atmospheres, providing measurements of chemical composition, temperature and find potential indicators of habitability. Direct imaging of exoplanets and protoplanetary discs will also be possible, allowing the study of planet formation processes. MORFEO's adaptive optics correction will provide near diffraction-limited resolution in the near-infrared, which is fundamental for resolving planetary systems and circumstellar structures.

Another major scientific theme that the ELT will address is the study of resolved stellar populations. Thanks to MORFEO (and MICADO, the imager and spectrograph) the ELT will allow the resolution of individual stars in other galaxies, extending photometric, spectroscopic, and astrometric studies beyond the Local Group. This will make it possible to construct, for example, accurate color-magnitude diagrams, determine luminosity functions, and measure stellar motions in dense clusters. As a result, astronomers will be able to study star formation histories, investigate chemical evolution and the internal kinematics of star clusters and galactic components that are currently inaccessible due to distance or crowding. The ELT is expected to potentially resolve extragalactic globular clusters and stellar structures in the most dense region in galaxy bulges, allowing the study of stellar populations across a wide range of galactic environments.

In the field of galaxy formation and evolution, the ELT will enable studies of both near and high-redshift galaxies with never seen before detail. High resolution imaging and spatially resolved spectroscopy will allow characterization of galactic substructures such as star-forming regions, bulge components, and compact satellite galaxies, even in low mass galaxies at high redshift. MORFEO's multi-conjugate adaptive optics will provide a stable point-spread function over a large field, allowing structural and photometric measurements of faint, distant galaxies with a level of detail never achieved before. These observations will yield stellar masses, ages, star formation rates and chemical enrichment patterns for galaxies up to the epoch of reionization, probing also the galaxy formation and evolution over cosmic time.

Finally, the ELT will contribute also to cosmology and the study of the distant Universe. Observations of objects such as faint, high-redshift galaxies and Type Ia supernovae will provide constraints on galaxy evolution, star formation, and the expansion history of the Universe. ELT's imaging capabilities will allow detection and structural characterization of high-redshift sources, which are fundamental for understanding early galaxy assembly processes, formation of globular clusters and the formation and evolution of galactic substructures.

Many of these scientific goals require almost diffraction-limited imaging and extremely high-resolution spectroscopy over different FoVs. The ELT can only achieve these requirements with the use of an AO system. This is where MORFEO becomes indispensable: by correcting atmospheric turbulence, MORFEO will provide a stable and almost diffraction-limited PSF to MICADO and the future second instrument that will be connected to it. This correction is indispensable to achieve the scientific objectives of the telescope.

Without MORFEO, many of the ELT's most important scientific programmes would be limited or completely impossible to achieve.

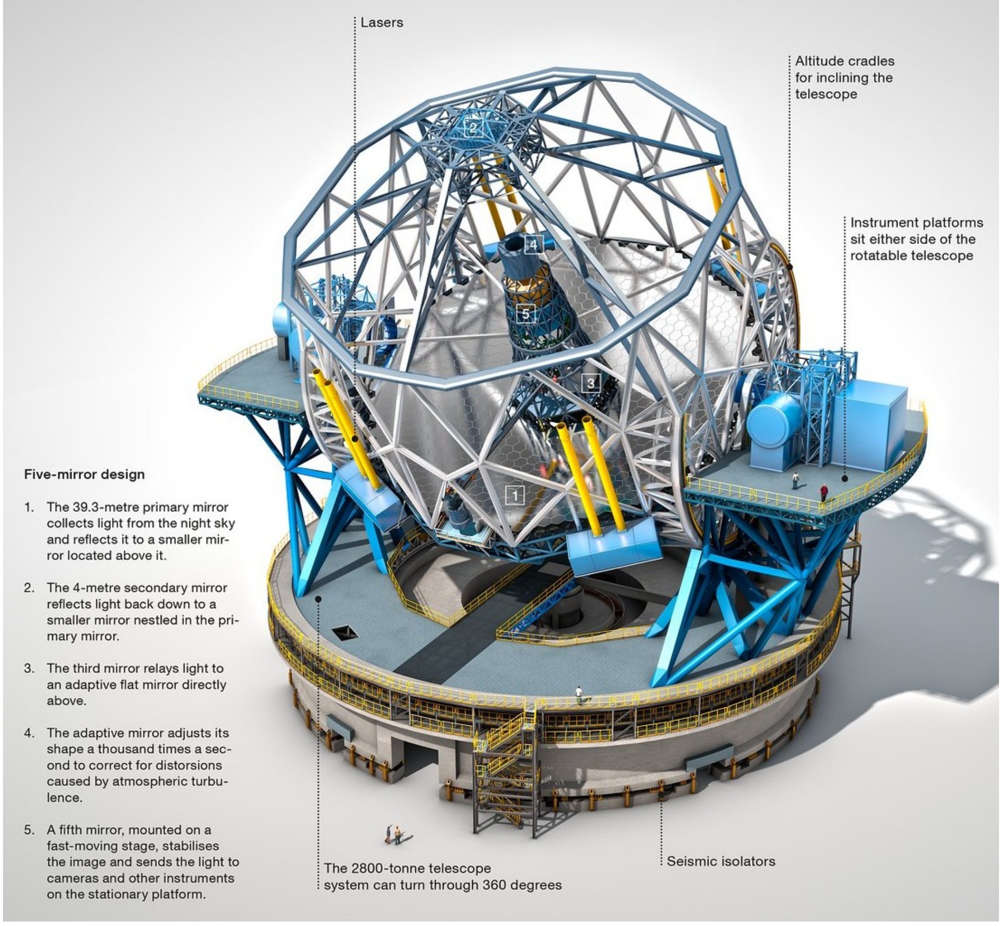

Figure 1.6: Scheme of the ELT with MORFEO and MCADO assembled on the telescope's Nasmyth platform, to scale. (Credits: ESO (n.d.) [7])

Chapter 2

Theoretical Background

2.1 Speckle Imaging

One effective way to measure atmospheric turbulence is through speckle imaging. The idea is to acquire short exposure images of a star (exposure times of the order of a few milliseconds) so that the atmospheric phase distortions are "frozen". In this situation, the observed PSF of a star is no longer a single diffraction-limited Airy disk. Instead, short-exposure images show an interference structure composed of many diffraction-limited "speckles," each with a characteristic angular size of approximately λ/D. Each instantaneous speckle pattern represents a snapshot of the wavefront distortions introduced by the turbulent atmosphere along the telescope's line of sight.

The number and distribution of speckles depend on the atmospheric coherence length r_0, the Fried radius, and the telescope aperture D. Physically, each coherence length of size r_0 across the pupil behaves like a separated sub-aperture which produces one diffraction-limited PSF. As a consequence, the final image is given by the superposition of all these PSFs, one for each turbulent cell contained in the pupil aperture. The smaller the value of r_0 (corresponding to stronger turbulence) and the larger the telescope aperture D, the bigger the number of speckles that appear in the image.

To obtain good speckle images, the exposures must be short because of the atmospheric coherence time τ_0. As previously explained, it corresponds to the time it takes for a turbulent eddy of size r_0 to cross the telescope aperture at a certain wind speed. To "freeze" the turbulence, we need the exposure times of $t_{\text{exp}} \ll \tau_0$. Typical values of τ_0 range from a few to a few tens of milliseconds. This requirement imposes some limits on the applicability of the technique: short exposures reduce the number of photons that reach the sensor, so the observed targets must be very bright to allow the achieving of high enough SNR.

The physical properties of speckles, such as size and ellipticity, are linked to turbulence parameters and can be used to obtain estimates of quantities like r_0, allowing for a rough estimate of low-order optical aberrations. For example, a systematic ellipticity in the observed speckles can be linked to astigmatism introduced by misalignments and deformations in the telescope optics.

In this thesis, speckle analysis is used both to estimate r_0, to compare the magnitude of aberrations between the Nasmyth and Coudé focal planes and to assess the general optical quality at the focal planes by studying the evolution of their diameters with the elevation angle.

Speckle images can also be used to detect closed binary systems that are unresolved in standard long-exposure imaging. By computing the autocorrelation of numerous short-exposure speckle frames (or by averaging the squared modulus of their Fourier transforms to obtain the power spectrum), it is possible to recover the object's spatial frequency information. The resulting pattern in the power spectrum, if the observed system is an unresolved binary, reveals the presence of two distinct sources and also allows for measuring their angular separation even when the binaries have a separation smaller than the diffraction limit of the telescope.
Speckles can also provide estimates of instantaneous atmospheric and instrumental aberrations.

2.2 Wavefront Sensing

2.2.1 Wavefront Basics

In optics, a wavefront is defined as a surface of constant phase of an electromagnetic wave. In an ideal case, such as a plane wave propagating through a vacuum medium, the wavefront is perfectly flat. However, when light passes through a turbulent medium like the Earth's atmosphere, inhomogeneities in the refractive index distort the phase of the incoming wave. Wavefront distortions can also result from defects in the optical system's components. As a consequence, the wavefront becomes irregular, resulting in a degradation of the image quality produced by the optical system, such as a telescope. These distortions are the primary limiting factor for astronomical imaging, whether caused by atmospheric turbulence or the optics themselves. So it is clear that the wavefront and its measurement play a fundamental role within the scope of AO.

2.2.2 Wavefront Sensors

A wavefront sensor records light intensity distributions or other physical quantities (like the displacements of focal spots for a SH-WFS) that are related to the wavefront shape. These measurements are used to reconstruct the phase distortions introduced by the atmosphere or the optical system. Once it is measured, it is possible to understand what the aberrations are that the system is dealing with and quantify the amount of image degradation that we have, with respect to the ideal perfectly plane wavefront, which would return a diffraction-limited resolution. The idea behind wavefront sensing is that WF measurements are used to compute the corrections that an AO system needs to provide (through DMs) to compensate for the aberrations induced both by the atmosphere and the optical system itself, with the scope to increase the quality of the image. In order to perform WFS, we need to measure the wavefront error (WFE): this is the deviation of the wavefront (often measured in μm) from its ideal planar shape. Several types of wavefront sensors exist, each based on different principles, such as:

- **Curvature sensor**: this sensor measures the local curvature of the wavefront by measuring the difference in the light intensity distribution between intra and extra focal images.

- **Pyramid sensor**: it uses a pyramidal prism that splits the incoming light beam into four parts. A four-faceted prism, followed by a relay lens, forms four identical pupil images in the absence of optical aberrations. If aberrations are present instead, the intensity distribution among the pupils will change. The local wavefront gradients can be obtained by measuring the light distribution in the pupils. By doing this, wavefront aberrations can be evaluated from the estimated gradients.

- **Shack-Hartmann wavefront sensor (SH-WFS)**: this kind of sensor (shown in Fig. 2.1), which has been used in this work, is composed by a matrix of lenses that, when placed on the pupil plane of the system, divides the pupil in smaller sub-apertures, each one of them producing a spot on the camera sensor. The displacement δy (as shown in Fig. 2.1b) of the points with respect to their ideal position gives the local derivative of the wavefront. This measurement is then used to retrieve the shape of the incoming WF and subsequently estimate the optical aberrations that characterise the WF. The higher the number of lenses, the higher is the precision in the WF reconstruction. But the more apertures the smaller amount of light comes in each of them, so at a certain point we will be limited by the SNR.

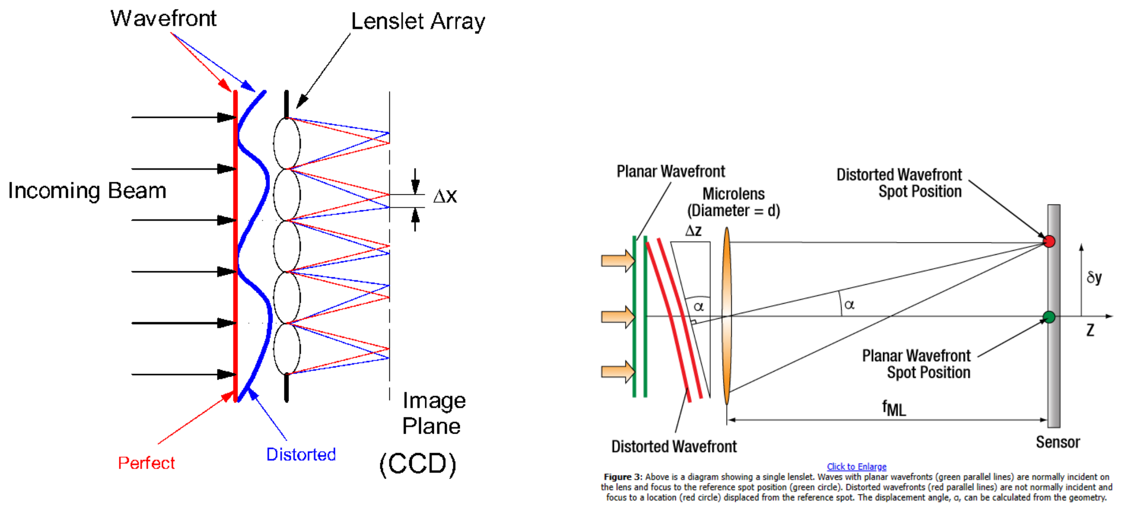

(a) Scheme of a Shack-Hartmann wavefront sensor

(b) Scheme of the spots produced by a SH-WFS

Figure 2.1: Scheme of a Shack-Hartmann wavefront sensor (credits: www.thorlabs.de (n.d.) [8])

Each sensor has its compromises between sensitivity, WF reconstruction capability and photon efficiency; the choice depends on the wavelength, the guide-star brightness and the desired correction order. In general, the purpose of wavefront sensing is to reconstruct the phase distribution across the telescope pupil with sufficient accuracy and speed to allow the AO corrections.

2.2.3 Wavefront Analysis with Zernike Polynomials

Zernike polynomials are a complete set of ortho-normal functions defined inside a unit circle. Orthogonality means that each polynomial describes an independent mode of wavefront aberrations, and that the terms do not overlap with one another. These properties makes them very useful for wavefront representation.

A regular wavefront can be described as the surface of equal phase of a monochromatic light beam. In an ideal case (like a point source at infinity), the wavefront is perfectly plane. But real optical systems introduce aberrations, which can be seen as deviations of the wavefront from its ideal shape. Zernike polynomials are particularly useful in this matter because they allow to represent the WF shape in a mathematical way, by decomposing it in its constituent aberrations.

Each Zernike term corresponds to a specific aberration (like defocus, astigmatism, coma, or spherical aberration). The coefficient relative to each term represents the contribution of that aberration to the total wavefront error (WFE). The orthogonality of the terms, furthermore, ensures that adding more terms does not affect the coefficients of the others. This is an important property of these polynomials which has been widely used in this thesis.

Mathematically, a wavefront $W(r,\theta)$ can be represented using Zernike polynomials:

$$W(r,\theta) = \sum_{n,m} c_n^m Z_n^m(r,\theta) \tag{2.1}$$

where Z is the Zernike polynomial and C is the associated coefficient. The polynomials are defined as follows:

$$\begin{aligned} Z_n^m(r,\theta) &= N_n^m \cdot R_n^m(r) \cdot \cos(m\theta) \quad for \quad m \geq 0 \\ Z_n^m(r,\theta) &= N_n^m \cdot R_n^m(r) \cdot \sin(|m|\theta) \quad for \quad m < 0 \end{aligned} \tag{2.2}$$

Where

$$R_n^m(r) = \sum_{l=0}^{\frac{n-m}{2}} \frac{(-1)^k (n-l)!}{l!(\frac{n+m}{2}-l)!(\frac{n-m}{2}-l)!} r^{n-2l} \tag{2.3}$$

$0 \leq r \leq 1$ while n indicate the radial order and m the azimuthal frequency of the polynomial. Different combinations of n, m represent different contributions to the total WF given by a specific aberration identified by the n, m polynomial.

There are different ways to identify a specific polynomial (and so, the optical aberration): one of the most used is the Noll notation (Noll (1976) [9]) which assigns to each n, m combination a single number. In this way, aberrations are easily identified when using the polynomials to represent the WF.
Fig. 2.2 shows the first Zernike terms with their associated aberration: the lower elements in the "pyramid" give gradually smaller contribution to the total WF.

With the increase in orders, we have finer representations of the WF. Aberrations can be represented by an infinite series in theory, but we usually stop at around the tenth order, depending on the required accuracy. The higher the order of aberration, the smaller its effect on the PSF. The first 13 terms of this polynomial account for about 90 % of the total aberrations, in terms of atmospheric effects, which are characterized

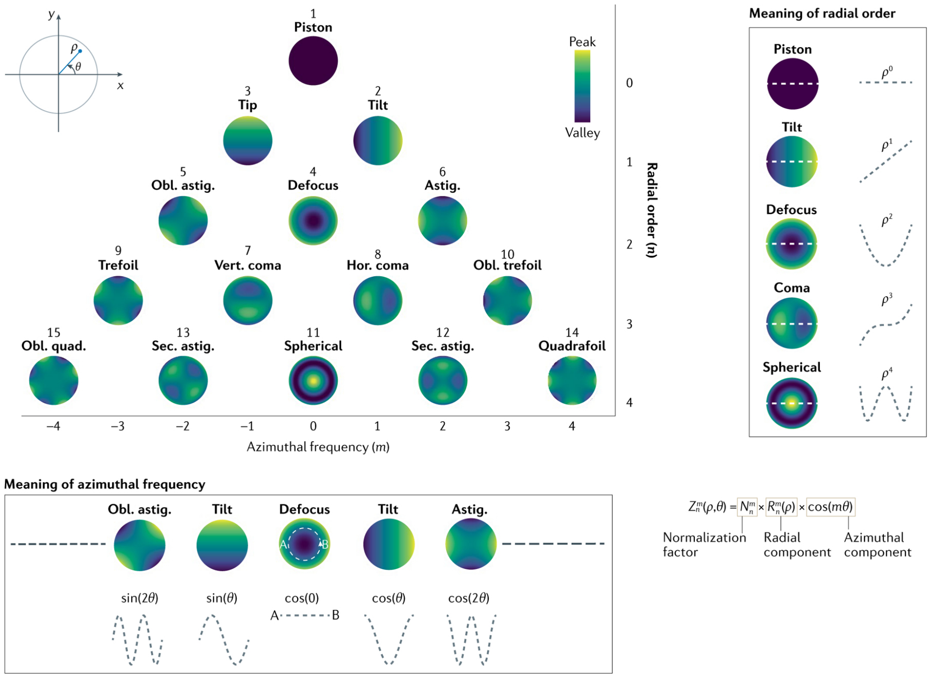

Figure 2.2: The first Zernike polynomials shown with their corresponding optical aberration (Hampson et al. (2021) [10])

by low coherence times and thus change every few milliseconds, depending on the conditions of the observational site.

This method of representing the WF provides two main advantages:

- It allows a description of complex wavefronts using only a few coefficients, since most of the aberration power is contained in low-order modes.
- It provides a physical interpretation of the aberrations, which can be used to drive the correction performed by a DM.

2.3 The "Donut" Method

The "Donut" technique is a method used to estimate low-order wavefront aberrations directly from defocused stellar PSFs. The main advantage of this technique is that it requires only the use of the science imaging path, eliminating the need for a wavefront sensor. As a consequence, with the Donut there is no non-common path and therefore the measured aberrations belong to the true optical path towards the scientific sensor. This characteristic makes the Donut a particularly good measurement in the context of adaptive optics. It provides slow measurements since it requires long exposures and it is not sensitive to high frequency and highly-variable aberrations, so its main uses are on the periodical alignment of the optics. Therefore, this technique can't substitute the measurements of a WFS but solves the non-common path problem for low-order aberrations (mainly defocus and astigmatisms).

The method and its characteristics were presented for the first time by Tokovinin and Heathcote (Tokovinin and Heathcote (2006) [1]).

2.3.1 Physical model and image formation

A fundamental part of the Donut method is the simulation of the image formation. To understand the image formation process, we can start by considering $\boldsymbol{\alpha}$ as the angular coordinates in the image plane and $\mathbf{x}$ as the coordinate in the telescope pupil. If the pupil transmission is $P(\mathbf{x})$ and the wavefront error is $W(\mathbf{x})$ (measured in length units), the complex amplitude in the pupil can be written as:

$$A(\mathbf{x}) = P(\mathbf{x}) \exp\left(i\frac{2\pi}{\lambda}W(\mathbf{x})\right).$$

The corresponding (monochromatic) intensity at a certain angular coordinate $\boldsymbol{\alpha}$ is given by the Fraunhofer integral:

$$I_\lambda(\boldsymbol{\alpha}) = \left|\iint A(\mathbf{x}) \exp\left(-2\pi i \frac{\mathbf{x}\cdot\boldsymbol{\alpha}}{\lambda}\right) d^2\mathbf{x}\right|^2, \tag{2.4}$$

In Tokovinin and Heathcote (2006) [1], the integral (2.4) is numerically implemented by the fast Fourier transform (FFT) performed on a computational grid of size $K \times K$, while the physical linear extent L of the pupil-plane grid must be large enough to contain the telescope pupil: its minimum size is $L \gtrsim D$, where D is the diameter of the telescope pupil. The angular sampling of the FFT is equal to

$$\Delta\alpha \simeq \frac{\lambda}{L},$$

and the total field of view provided by the grid is about $K\Delta\alpha \simeq K\lambda/L$.

The pupil phase map can be represented using Zernike polynomials:

$$W(\mathbf{x}) = \sum_{j=2}^{N_Z} a_j Z_j(\mathbf{x}), \tag{2.5}$$

where the a_j represent the Zernike amplitude terms.

An image based on equation (2.4) contains diffraction-limited features (on a scale λ/D), but in real data, the finer structures are generally affected by atmospheric seeing. This effect is taken into account by convolving the image with a *Gaussian kernel.*

2.3.2 Light Moments

Since the objective is to find the best fit model through an iterative fitting process, some initial guesses are required to avoid local minima and allow a faster convergence towards the solution. Second moments of light provide good first estimates for the low-order terms (mainly defocus and astigmatism). The useful formulas are reported below:

$$\begin{aligned} I_0 &= \sum I_{ij} \\ x_c &= I_0^{-1} \sum x_{ij} I_{ij} \\ y_c &= I_0^{-1} \sum y_{ij} I_{ij} \end{aligned}$$

Here, every combination of indices i, j identifies a single pixel in the image.
The first moments are defined as follows:

$$M_x = I_0^{-1} \sum (x_{ij} - x_c)^2 I_{ij} \tag{2.6}$$

$$M_y = I_0^{-1} \sum (y_{ij} - y_c)^2 I_{ij} \tag{2.7}$$

$$M_{xy} = I_0^{-1} \sum (x_{ij} - x_c)(y_{ij} - y_c) I_{ij} \tag{2.8}$$

$$\tag{2.9}$$

These moments are related to defocus and astigmatism coefficients, reported in angular units, through the relations (Tokovinin and Heathcote (2006) [1]):

$$A_4 = p\sqrt{(M_x + M_y)/2} \tag{2.10}$$

$$A_5 = pM_x(M_x M_y)^{-1/4} \tag{2.11}$$

$$A_6 = 0.5p(M_x - M_y)(M_x M_y)^{-1/4} \tag{2.12}$$

$$\tag{2.13}$$

where p is the plate scale, so the angular dimension of the image pixels.
The physical correspondent values for the aberrations (the WFE in length units) are given by:

$$A_4 = a_4(4\sqrt{3}/D) \tag{2.14}$$

$$A_5 = 0.23\ a_5/D \tag{2.15}$$

$$A_6 = 0.23\ a_6/D \tag{2.16}$$

$$\tag{2.17}$$

where D is the telescope diameter and a_n is the amplitude in length units of the aberrations (defocus, oblique, and horizontal astigmatisms, respectively).

2.3.3 Iterative fit

The Donut fitting procedure described in the paper Tokovinin and Heathcote (2006) [1] is an iterative process where donut images are cyclically generated, compared with the original image untill a best fit model is reached. The fitted model will then give as a result the fitted Zernike coefficents searched. It is based on these main elements:

1. **Initialization:** The algorithm begins with an initial guess for the coefficients, based on simple measurements of the donut image (the second moments of light). A good starting point helps the method converging correctly.

2. **Modelling:** At each step, a synthetic donut image is created from the current parameters. The synthetic image is then compared to the observed image in order to calculate the residuals.

3. **Updating parameters:** The algorithm estimates how sensitive the image is to small changes in each parameter. This information is used to compute a small correction.

4. **Iteration:** The steps above are repeated until the changes in parameters are small or the fit no longer improves significantly. At that point, the cycle stops and returns the fitted Zernike modes.

Chapter 3

Observations and Data Acquisition

3.1 Speckle Acquisition

The analysis of speckles was carried out in two phases. During the first observation night in January 2025, the focus was on studying how variations in exposure time during image acquisition affected the FWHM of the speckles. In the subsequent observations, the objective shifted to investigating how both the FWHM and the ellipticity of the speckles vary with the telescope's elevation angle, and to compare these characteristics between the Nasmyth and Coudé foci.

To obtain the data needed to perform speckle analysis, we prepared a setup composed by a lens with 100 mm focal length, followed by a band-pass red filter, then a 150 mm lens, and finally by a camera (model: PCO Edge 4.2). This system was mounted on an aluminium bar installed at the Nasmyth focus of the telescope. The system provided magnifying power of 3x or 6x, depending on the lenses used.

PCO Edge 4.2 camera

The PCO Edge 4.2 camera is an high-performance scientific camera characterised by very high sensitivity, low noise, and high-speed image acquisition. It features a CMOS sensor with a resolution of 2048×2048 pixels and a pixel size of 6.5 μm, enabling very high-resolution imaging. The camera achieves a quantum efficiency of 82 %, ensuring high sensitivity (in the visible spectrum) and a readout noise of about 0.8 electrons, which allows for a good signal-to-noise ratio also in high-speed imaging. These characteristics are particularly suitable for performing astronomical imaging, where high speeds, resolution and sensitivity are often required (source :excelitas.com (n.d.) [11]).

During the first observation nights, we utilised the setup shown in Fig. 3.2 to perform speckle imaging of the same target star using three different exposure times, thereby allowing us to study the variations that arise due to changes in exposure time. The observed star is Algol (HD 19356), chosen for its position at the time of the observation: it was at the zenith (to minimise the airmass) and near the meridian, to avoid excessive flexures of the telescope that could lead to significant optical misalignment of the optical setup.
We retrieved three sets of images with different rates: 500 Hz, 1 kHz and 10 kHz corresponding to exposure times of 2 ms, 1 ms and 0.1 ms respectively. Each set consists

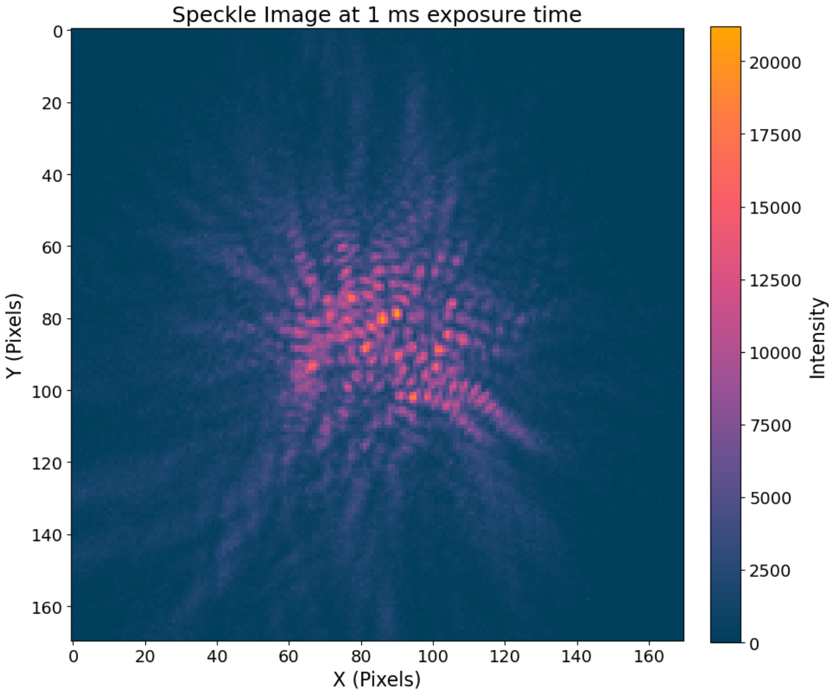

Figure 3.1: Example of speckle image.

of 5000 images, resulting in a total dataset made of three stacks of 5000 images each. Fig. 3.1 shows an example of one of the images obtained.

Then the plate scale was measured. It describes the relationship between angular measurements in the sky (in units of arcseconds) and physical distances on a detector (in units of pixels). To determine the plate scale, we observed the binary star system BD+31424. By measuring the separation between the two stars on the detector (1942 $pixels$) and knowing their apparent angular separation (73 $arcsec$), we computed the plate scale, which was found to be 0.0376 $arcsec/pixel$, corresponding to an equivalent focal length of the telescope of 35501 mm.

For the second part of the analysis, various targets were observed at different elevation angles, to cover the full range of elevations with approximately 10° steps between one target and the other. Observations were always kept as close as possible to the Meridian to minimize mechanical stress on the telescope and on the optical setup mounted on the aluminium support bar.

At the Coudé focus, two distinct setup configurations were employed during separate observation campaigns, one in March 2025 and the other in September 2025. In both cases, the goal was the same as the Nasmyth focus experiment; however, two different setups were used. This was done to allow, during the September 2025 campaign, the simultaneous acquisition of both Donuts and WFE data, while such a thing was not prevented during March 2025. To obtain speckles it was sufficient to move the camera on the system focus and change the lens between the beam splitter and the camera. Fig. 3.3 shows the setup used during March 2025 and Fig. 3.5 the September one.

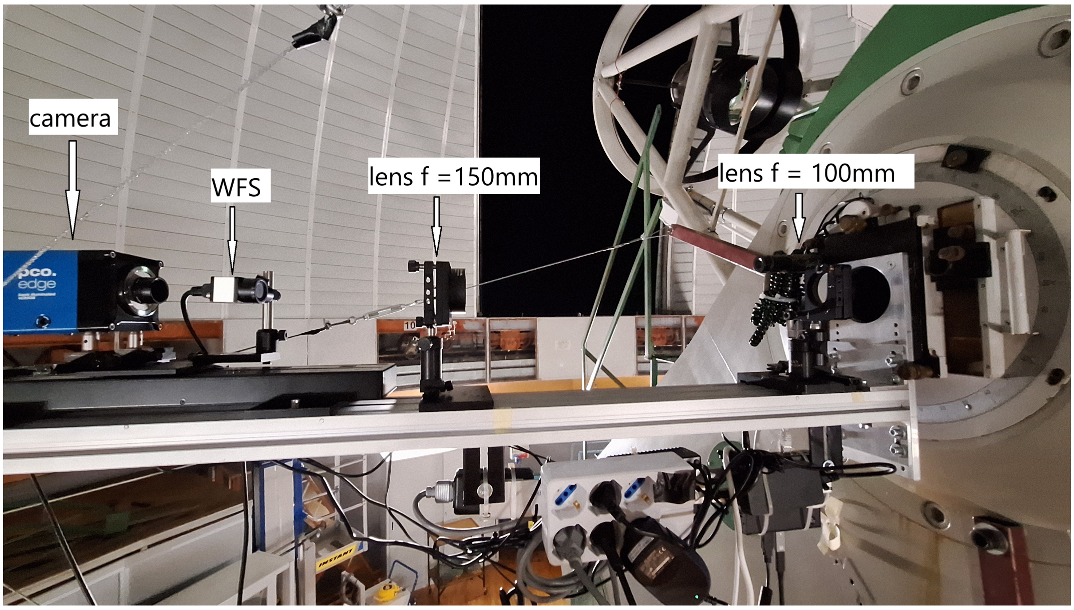

Figure 3.2: Photo of the setup used to perform speckle imaging mounted on the telescope. In this image, the wavefront sensor is also shown in front of the camera. The same setup was used for the other observations performed at the Nasmyth focal plane, with a different lens used to realize different magnifications (3x and 6x).

3.1.1 Telescope's M1 aluminization

It is important to notice that the primary mirror of the telescope underwent aluminization in summer 2025, after the Nasmyth observations conducted in winter 2025 and the Coudé observations performed in September 2025. This process involved removing the mirror from the telescope, re-coating it with a new reflective layer, and reinstalling it. This process is followed by a realignment of the optical system. Such procedures can alter the overall optical quality, both by improving surface reflectivity and by changing the alignment of the optical elements. Consequently, wavefront measurements and speckle characteristics obtained after the aluminization may differ from those recorded earlier, potentially showing reduced aberrations and changing also speckles' characteristics. This effect should be considered when comparing observations taken before and after the aluminization, as it can influence the performance of the telescope and the derived results.

3.2 Wavefront Sensing

For the wavefront sensing data acquisition, we used a Shack-Hartmann sensor, model Thorlabs WFS31-7AR.
The Thorlabs WFS31-7AR is a Shack–Hartmann wavefront sensor designed for use across the visible to near-infrared spectrum (400–900 nm). It employs a microlens array (MLA) with lenslets of 146 µm in diameter, coated with anti-reflection (AR) to increase performances and optimise the amount of light reaching the sensor. The MLA is composed by 73 x 45 lenslets, with a focal length of 5.6 *mm*. The sensor features a CMOS camera with a resolution of 1920 × 1200 pixels. (Source: www.thorlabs.de (n.d.) [8])

WFS calibration

The calibration of the sensor is the first fundamental step that we performed to ensure that the measured spot displacements correspond to true wavefront slopes. The sensor can be calibrated in two ways. The first is the reference-beam calibration, in which the sensor is illuminated with a collimated beam produced by an optical system placed in front of the sensor. Since the incoming wavefront in this setup is ideally flat, any deviation of the measured spot positions from the ideal grid can be attributed to lenslet imperfections or assembly tolerances. These offsets are measured during the calibration phase and later subtracted from all measurements. This method requires dedicated alignment procedures and additional precision optics. The second approach is the fiber-based calibration method, where an optical fiber is placed in front of the sensor acting as a quasi-point source, producing a spherical wavefront with known curvature radius, given by the distance between the sensor and the fiber. The sensor then records the corresponding spot pattern, and determines the reference centroid map in order to perform calibrated measurements. This method revealed to be more practical for our data acquisitions since the fiber calibration requires a simpler optical setup and as a consequence, in this work, we implemented the fiber calibration method.

To obtain the required measurement, the sensor had to be placed on the pupil plane; thus, we prepared different setups at the telescope: one for the Nasmyth and one for the Coudé.

It is worth noting that the observations were conducted over multiple nights, even using different setups. The following description shows the setup used during the observations that took place in March 2025.

The Coudé setup is shown in Fig. 3.3. The setup captures the beam coming from the telescope and splits it using a beamsplitter , which allows some light to enter another camera previously installed on the bench (not useful for this work). Then the light passes through a folding mirror, then a lens ($f = 100mm$), another folding mirror, another lens ($f = 150mm$) and then the WF sensor. The sensor is mounted on a translator to allow precise positioning of the camera. High precision in the position is needed since to get optimal measurements, we need to place the sensor exactly on the pupil. Otherwise, the MLA can't correctly focus the PSFs on the grid (Fig. 2.1b), which are used to estimate the WFE. The setup, designed using Zemax, is shown in Fig. 3.4.

Using these setups, we carried out the observations of our targets. The goal was to observe bright stars (magnitude < 4) located close to the Meridian at the time of observation and distributed across different elevations, in order to study how the telescope elevation affects the alignment status of the system and its WFE.

During the March 2025 observing night at the Coudé focus, data were collected from three targets at four elevation angles: $25°, 40°, 60°$, and $80°$. At the Nasmyth focus, the selected targets were instead observed at $27°, 60°$, and $90°$. For each target, 20 frames were recorded. Each frame corresponds to the average of ten 10 ms exposures, and one frame out of every five is saved. In this way, we have a good representation of the wavefront, ensuring that we sample it over a time interval long enough to allow the study of the time-variability of the wavefront.

Even if the Shack–Hartmann WFS measures the local slope of the wavefront by detecting the δy (see Fig. 2.1b), it also provides a direct computation of the wavefront

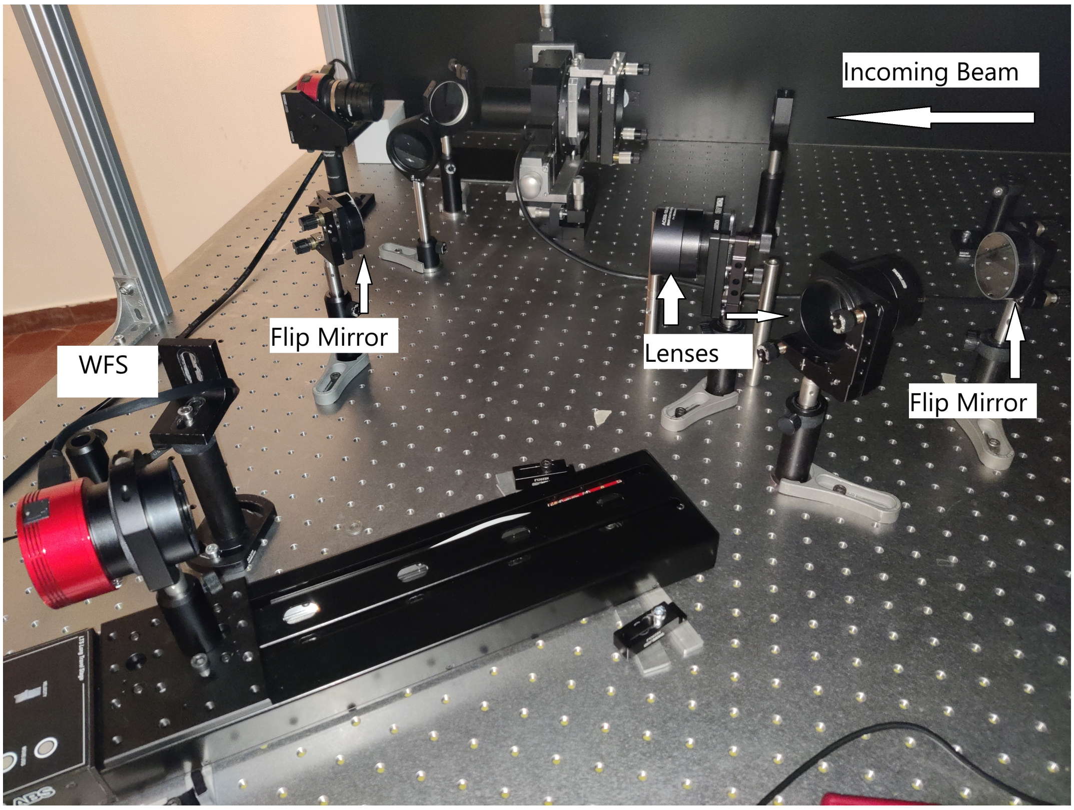

Figure 3.3: Setup used at the Coudé to perform WF measurements

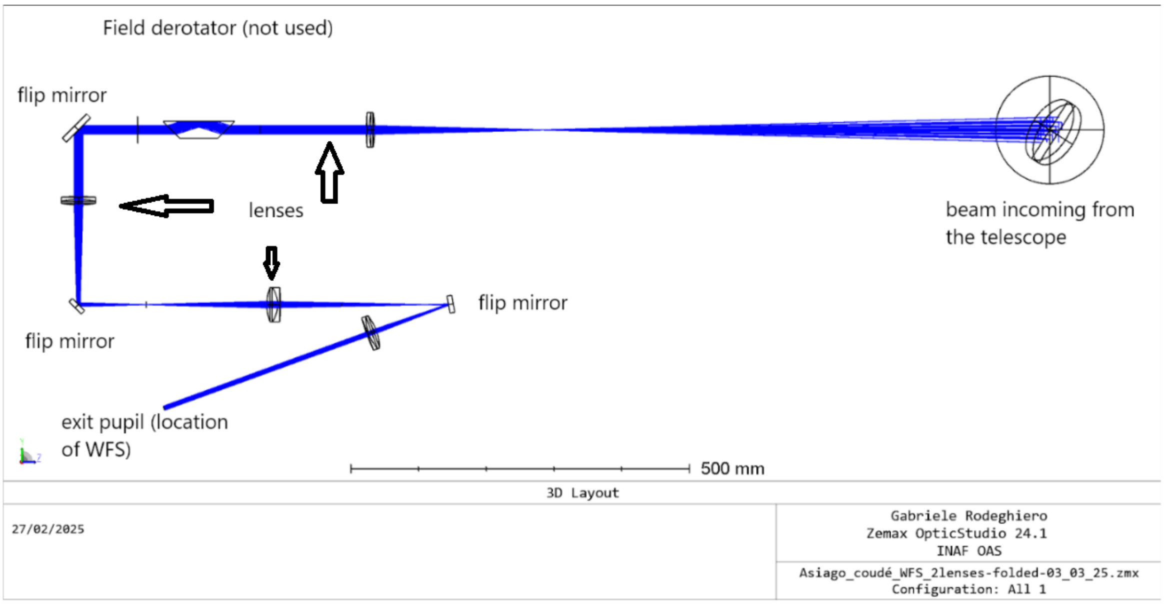

Figure 3.4: Zemax design of the setup used at the Coudè for the SH-WFS during the March observations. (Courtesy of Gabriele Rodeghiero)

error (WFE). The output of each acquisition is a CSV file containing metadata and a 73 × 45 grid with the sampled WFS values. The setup is configured so that the pupil diameter is fully contained within the surfaces of both the microlens array (MLA) and the sensor (10.8 mm × 6.6 mm). This requirement is essential because we needed to sample the entire pupil to obtain reliable measurements of the WFE. For this reason, the WF data obtained during the January 2025 observations had to be discarded, as the pupil diameter was larger than the surface of the detector and could not be sampled correctly.

For the September 2025 campaign, the final observing nights utilized the setup shown in Fig. 3.5 and 3.6. At the Coudé focus an optical bench was arranged to allow the simultaneous acquisition of images and wavefront maps: a cubic beamsplitter divided the incoming beam into two paths, one sent to the camera for Donut images and the other to the WFS. The targets observed in this run were located at elevations of 37°, 45°, 55°, 60°, 70° and 78°. For each target, between 40 and 70 WFS frames were recorded in order to cover the time span of the Donut acquisitions, allowing a direct comparison between the two measurements.

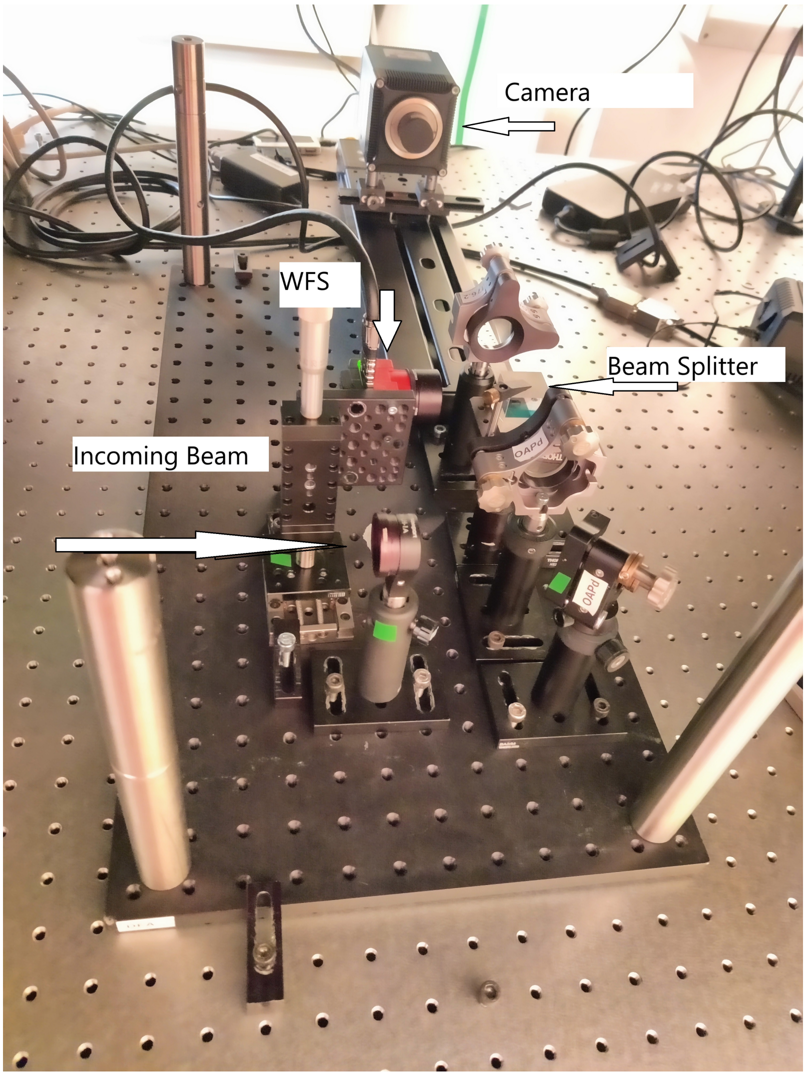

Figure 3.5: Setup used at the Coudé bench for the acquisition of donuts and WF maps during the September 2025 observation campaign.

Field Rotation Analysis

Due to the different telescope pointing, the focal plane undergoes a rotation when the elevation changes. This field rotation affects the measurement of asymmetric aberrations (for example, astigmatism and coma) and, if not taken into account, leads to

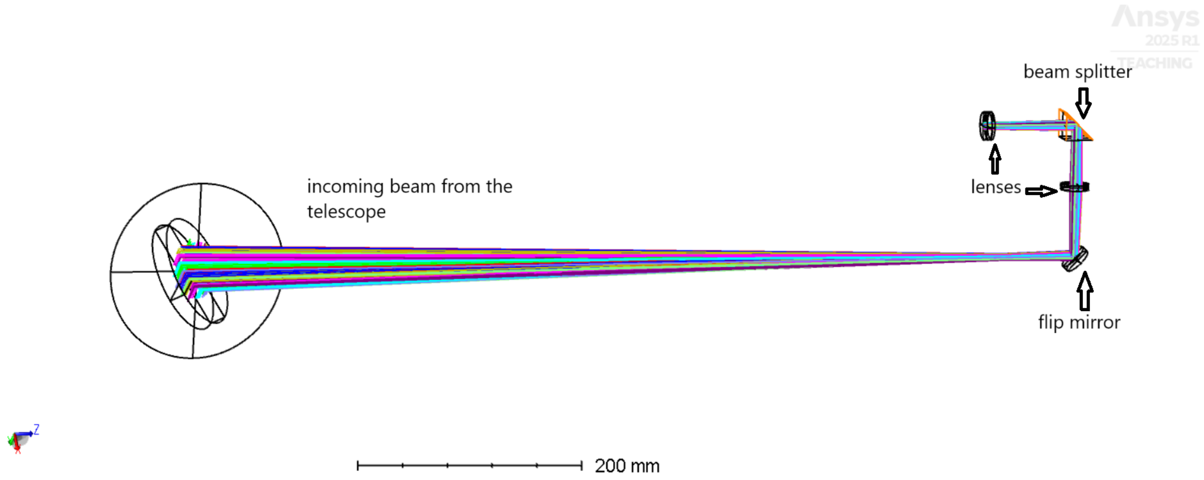

Figure 3.6: Zemax design of the setup in Fig. 3.5 showing the beam splitter and the optical path towards the WFS.

wrong results. The field rotation projects one mode onto another (for example, vertical astigmatism is projected onto oblique astigmatism and vice versa), making it difficult to distinguish between one aberration and the other. To address this issue, I conducted a ray tracing analysis using Zemax to quantify the field rotation as a function of elevation, and then utilised these results to correct the measured aberrations.

The results of this analysis are shown in Fig. 3.7, which reports the field rotation angle as a function of the telescope elevation. The plot shows that the focal plane rotates of approximately 48° between the zenith and the lowest elevation achieved during the observations. Such a large rotation cannot be ignored: if left uncorrected, we have projections on asymmetric modes. As a consequence, the measured Zernike coefficients must be corrected.

Zernike coefficients derotation

To correct for the field rotation we have to derotate the Zernike coefficients bringing all of them to a common "reference frame". Modes with dependences on $\cos(m\theta)$ and $\sin(m\theta)$ project one onto the other when rotated by an angle φ according to the usual 2D rotation by $m\varphi$. The useful modes to correct for this analysis are the astigmatism and coma; we do not correct higher order aberrations since they are not the aim of this study.
In the Noll convention, the two astigmatism modes correspond to $m = 2$ (Z_5 and Z_6), while the two coma terms correspond to $m = 1$ (Z_7 and Z_8). As a consequence, if we have two modes measured in a rotated field (denoted with Z'_j) it is possible to bring them to a common reference frame (so de-rotate the modes, obtaining Z_j) by using the formulas below:

Astigmatism (azimuthal order $m = 2$):

$$Z_5 = Z'_5 \cos(2\varphi) \; - \; Z'_6 \sin(2\varphi), \tag{3.1}$$

$$Z_6 = Z'_5 \sin(2\varphi) \; + \; Z'_6 \cos(2\varphi). \tag{3.2}$$

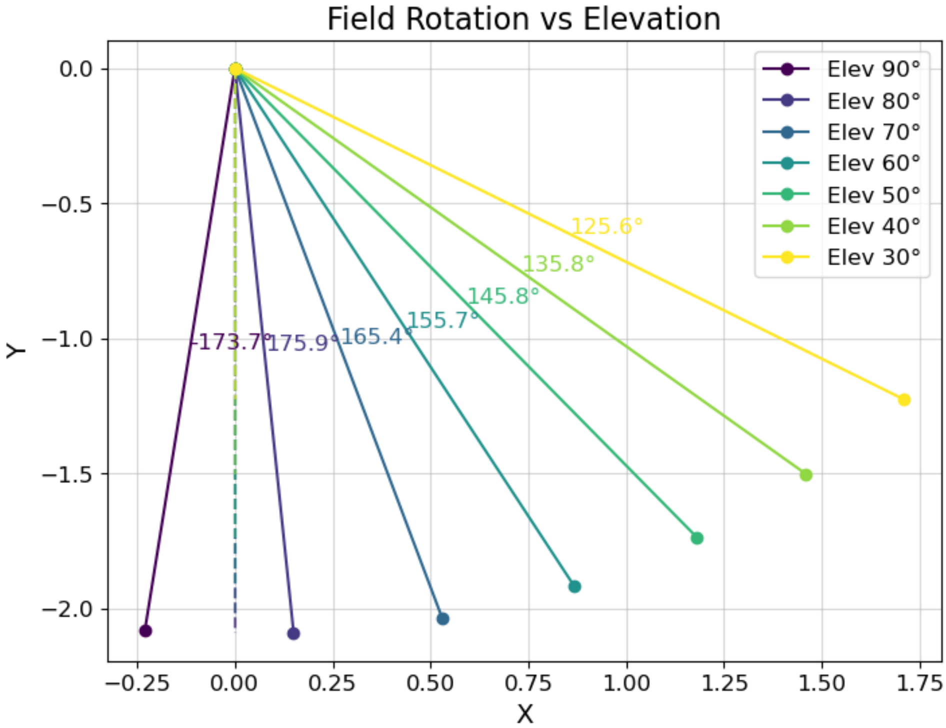

Figure 3.7: Field rotation as a function of the elevation angle. Notice that the images reaching the Coudé focal plane are flipped upside-down and then rotated.

Coma (azimuthal order $m = 1$):

$$Z_7 = Z_7' \cos(\varphi) \; - \; Z_8' \sin(\varphi), \tag{3.3}$$

$$Z_8 = Z_7' \sin(\varphi) \; + \; Z_8' \cos(\varphi). \tag{3.4}$$

In these formulas φ is the focal plane rotation angle (in radians) measured from the reference orientation. These formulas will be used later in the code at the end of the analysis of all the frames, where the final fitted Zernike coefficients are de-rotated to evaluate the evolution of the aberrations as a function of the telescope elevation angle.

3.3 Donut acquisition

Donut images were acquired using the PCO Edge camera (3.1) at both the Nasmyth and Coudé foci. For the Nasmyth focus, we started from the setup shown in Fig. 3.2. Then, all optical elements were removed, and the camera was positioned near the telescope flange. To capture Donut images, the camera must be located a few millimeters away from the focal plane and, since the Nasmyth focus was found to be inside the telescope structure, the camera was positioned as close as possible to the flange.

During the observations at the Nasmyth focal plane, six targets were observed near the Meridian and at various elevation angles (90°, 80°, 47°, 32°, and 27°) to cover a wide range of telescope configurations. For each target, multiple frames were recorded (from 100 to 3000) with exposure times varying from a few tens to hundreds of milliseconds, depending on the target brightness.

However, as Fig. 3.8 shows, the Nasmyth donuts are too out of focus, with even the spider visible and also show filamentary features, like caustics, which are not ideal for the Donut technique. This happened because the distance between the camera and the focal plane was still too large due to the presence of the telescope flange, which blocked the camera from getting closer to the Nasmyth focal plane.

For the Coudé focal plane, we used the setup in Fig. 3.5. In this configuration the camera was mounted on a linear stage, to allow the control of the defocusing.

At the Coudé bench, six targets were also observed at elevations of 37°, 45°, 55°, 60°, 70°, and 78°. These were the same targets used for the WFS data shown in Fig. 5.27, since this setup was prepared to allow simultaneous WF measurements and image acquisition. In this case, 3000 images per target were collected with an exposure time of 80 ms.

In this case, the images were acquired correctly, as the amount of defocus could be freely set without any mechanical constraint, unlike the Nasmyth focus. The spider arms were not visible, and no caustics appeared, resulting in smoother donuts, more suitable for the fitting process.

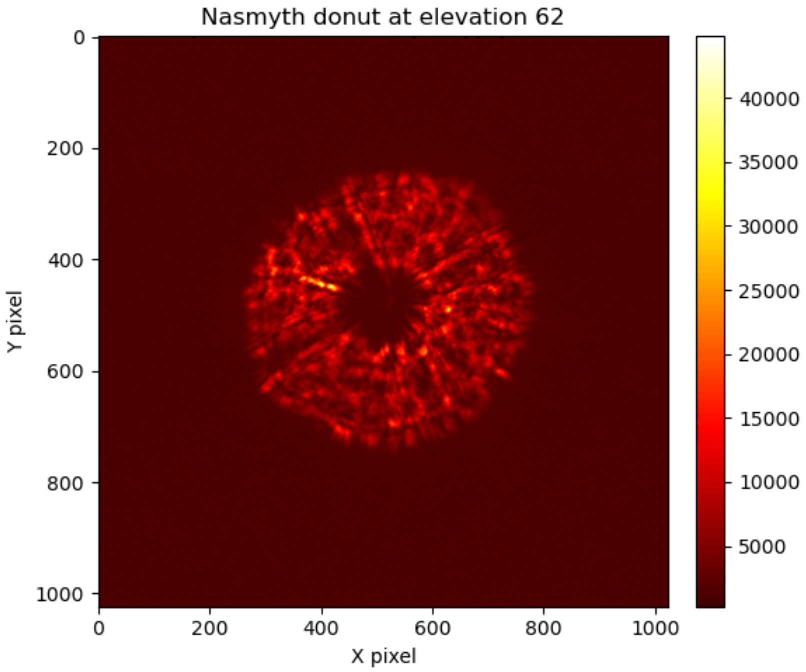

(a) Single donut frame

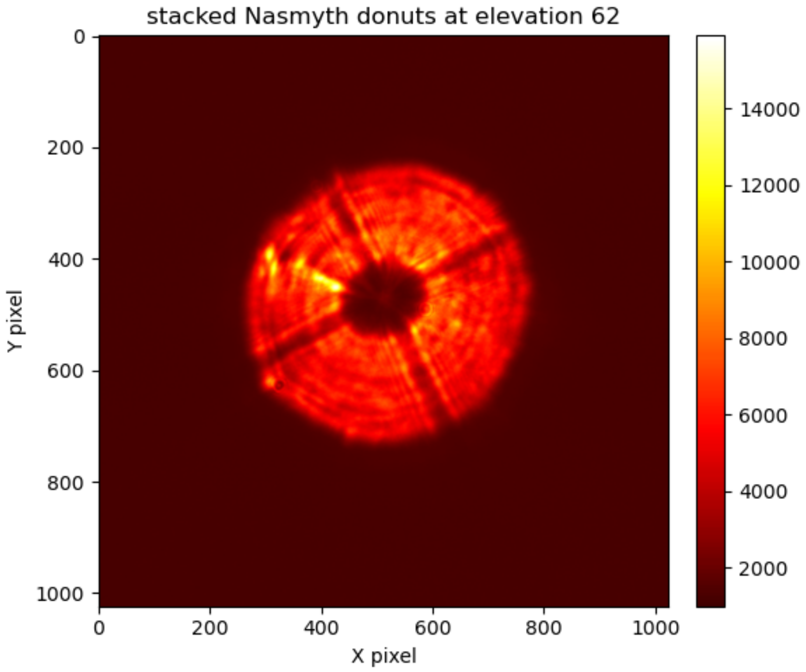

(b) Stack of 500 donuts

Figure 3.8: Example of a single and stacked donut frames obtained at the Nasmyth focal plane. Notice the presence of the spider and the filamentary structures in the single frame.

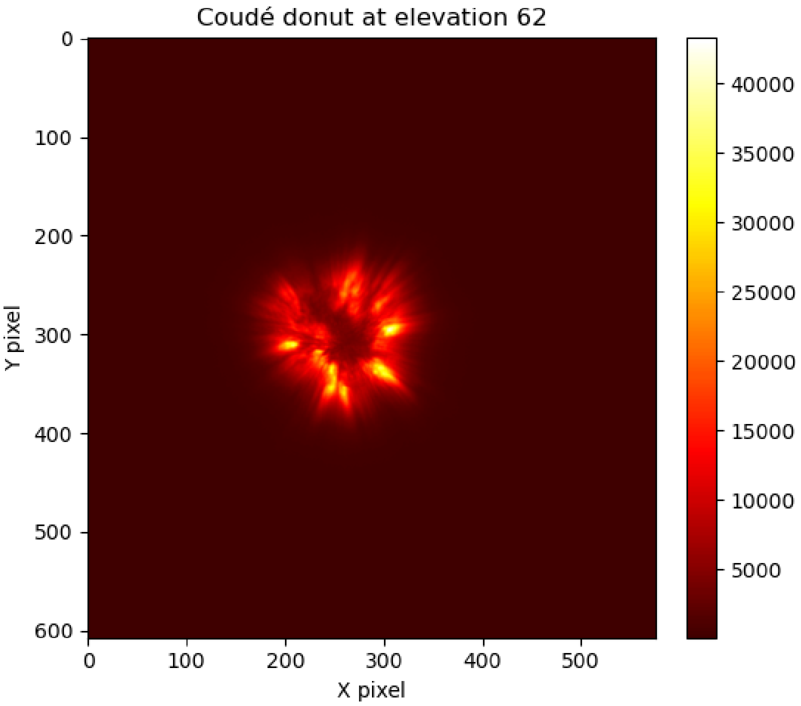

(a) Single donut frame

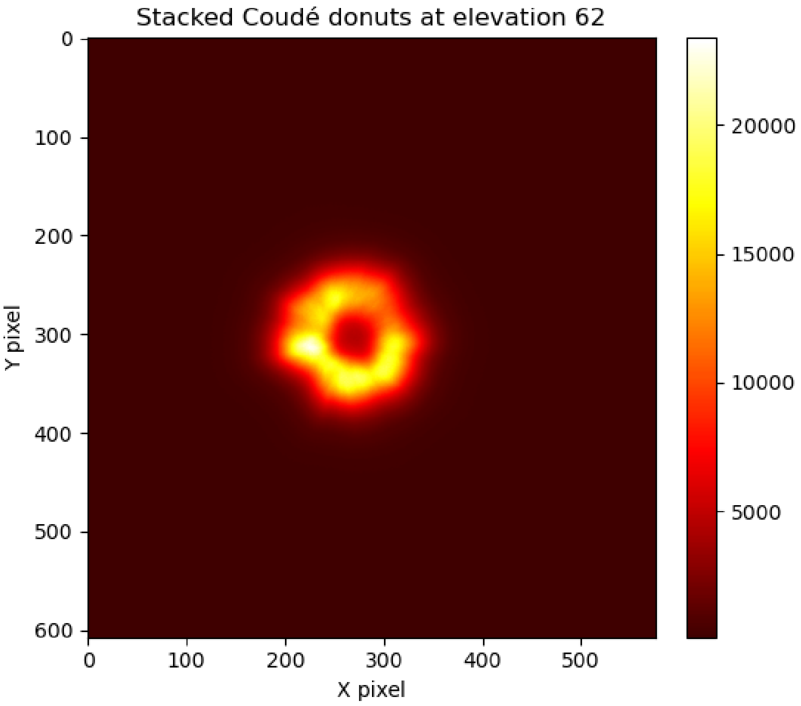

(b) Stack of 500 donuts

Figure 3.9: Example of a single and stacked donut frames obtained at the Coudé optical bench.

Chapter 4

Numerical Implementations

4.1 Speckle Analysis Code

In order to analyse speckle data, I wrote a code which is divided in three main steps:

- Data reading and cleaning
- Speckle search
- Speckle fitting
- PSF Photometry and fitting

The first three steps are repeated for each image in the data cube. Photometry, on the other hand, is performed on the stacked image obtained by combining all the images of a data cube.

This code is available on GitHub at https://github.com/Leonardo-5/Speckle-Analysis-for-AO.

Data reading and cleaning

In the first section, the code opens the file containing images related to one target, selects a specific image, and estimates the background noise, which is given by the median value of the counts in the image. This value, in units of counts/pixel, is then subtracted from every pixel of the image.
In this way we have a clean image that will be used in the next steps. The background is typically estimated by acquiring flat-field images. We didn't obtain such data during the observations because flat fields are required for high-precision photometry, which was not the primary aim of these observations.

Speckle search

The next part of the code involves identifying the speckles through a two-step process. First, the code determines the coordinates of pixels with counts exceeding a specified background threshold. Then, for each pixel, it scans within an Airy disk radius (determined by the telescope aperture and the reference wavelength of the observations) to find any other flagged pixels. If such a pixel is found, the one with the highest count is selected and marked as a speckle.

Next, a final check ensures that no adjacent pixels remain classified as speckles from the previous steps. If two adjacent flagged pixels are detected, only the one with the highest count is considered a speckle.
This process reduces the number of potential speckles from thousands of pixels per image to approximately 200, depending on the threshold.

An example of this process is in Fig. 4.1

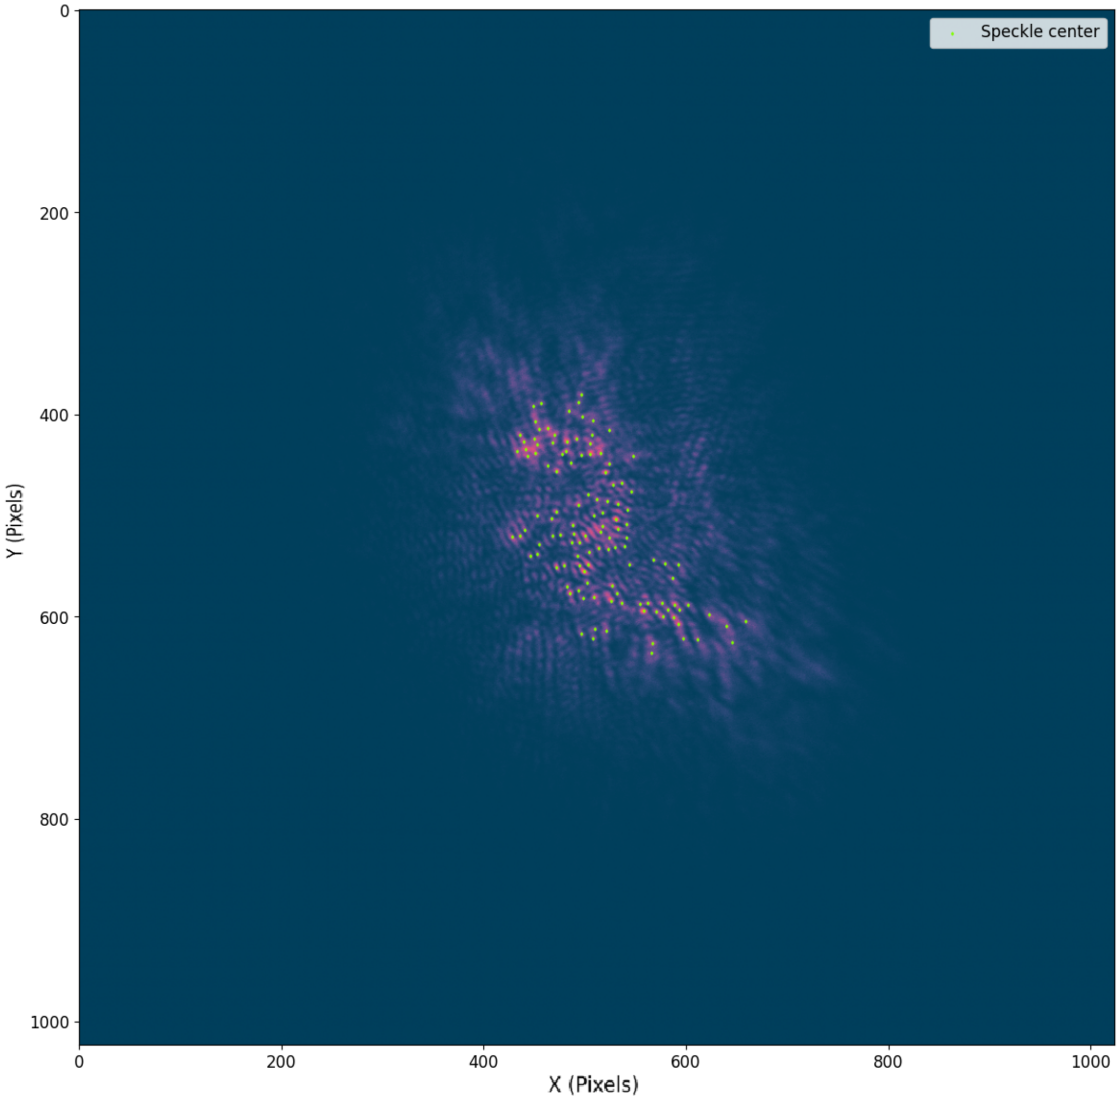

Figure 4.1

Speckle Fitting

Once the speckles are identified, the pipeline proceeds by fitting the speckles with a Gaussian curve in order to obtain the FWHM of all the speckles in the image.
This fit is performed using a 2D Gaussian curve. As an initial estimate for the fitting function, the speckle coordinates are used as the initial guess for the peak position, while the pixel counts give the constraint on the maximum value that the Gaussian curve can reach.
Once the best fit curve is found, the code extracts the FWHM value (both in x and

y axes) and the coordinates of the peak, which will be considered the actual centre of the speckle.

In Fig. 4.2, an example of the Gaussian fit is shown: from these Gaussians the code retrieves the FWHM and the center of each speckle.
This process is repeated for all the images in the stack, after which the code gives as output the FWHM and the positions of all the speckles.
Fig. 4.3 shows a flowchart of the code: the yellow boxes represent functions, while the blue ones are the arrays that store the results.

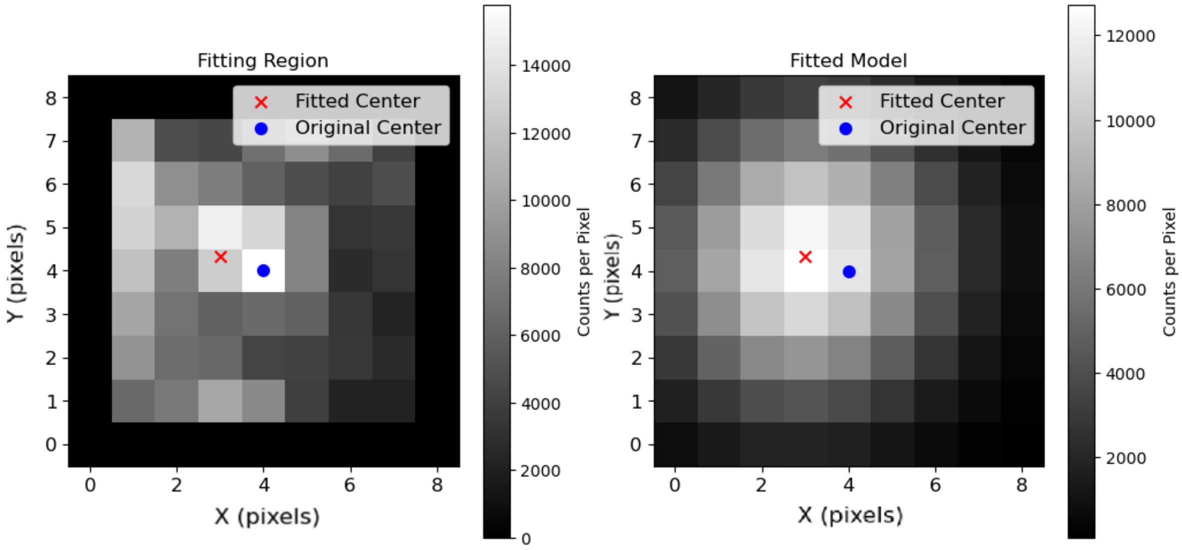

Figure 4.2: Example of Gaussian fit of a single speckle. Notice the photon counts on the left of the fitting region: those counts are part of another speckle, and this represents a problem for the quality of the fitting Gaussian since the code will try to fit those pixels even though they do not belong to the speckle that is being fitted.

Photometry

The final step of the pipeline involves measuring the PSF of the star. This is done in order to estimate the seeing, which is assumed to be equal to the FWHM of the PSF. The single images are stacked one on top of the other to get an image equivalent to a long-exposure image of the star (as seen in Fig. 4.4). The final image contains the average intensity of each pixel. Then, the composed image is smoothed to attenuate the granularity caused by the pixels.

The measurement of the point spread function (PSF) width was conducted using both aperture photometry and PSF fitting.

Aperture photometry is a technique used to measure the total flux of a source by summing the counts within a defined aperture radius centered on the source. The aperture is determined as the radius at which the total counts no longer increase, indicating that the entire PSF has been encompassed. This method is verified using the curve of growth, which plots total counts as a function of photometric radius.
However, the measured counts within the aperture include both the source signal and the local background contribution. To account for this, a background annulus is used

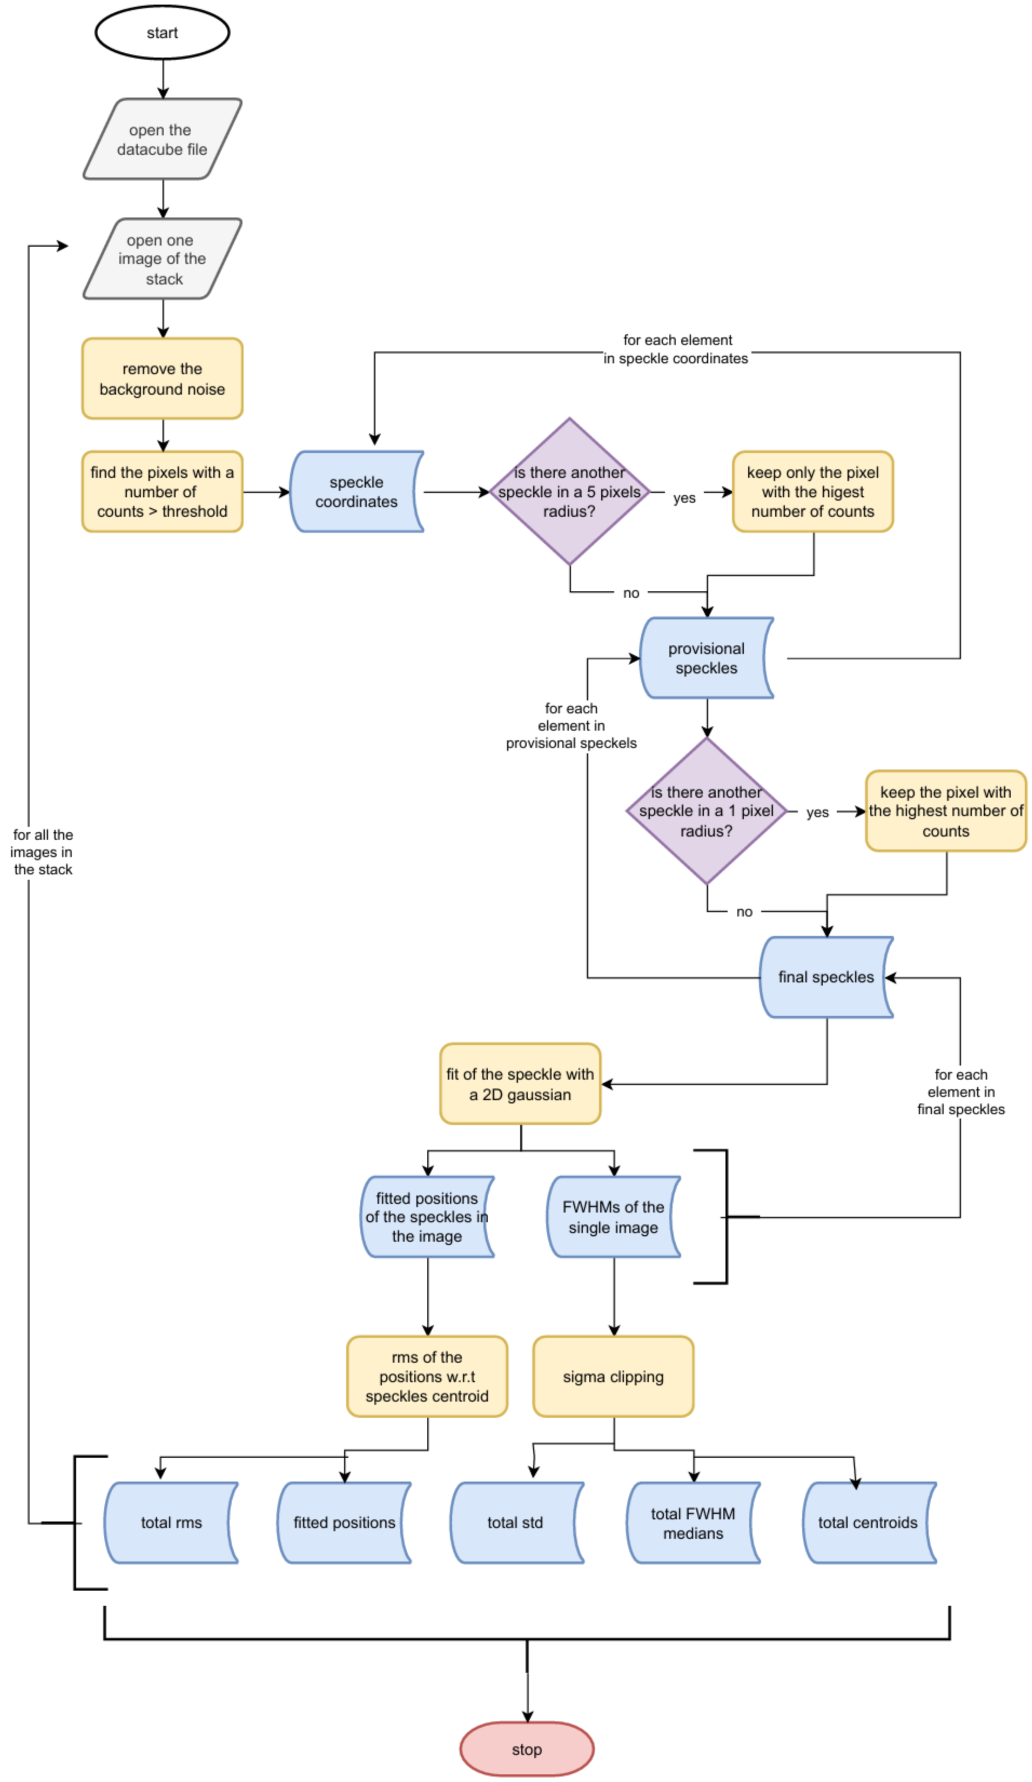

Figure 4.3: Flowchart of the speckle code.

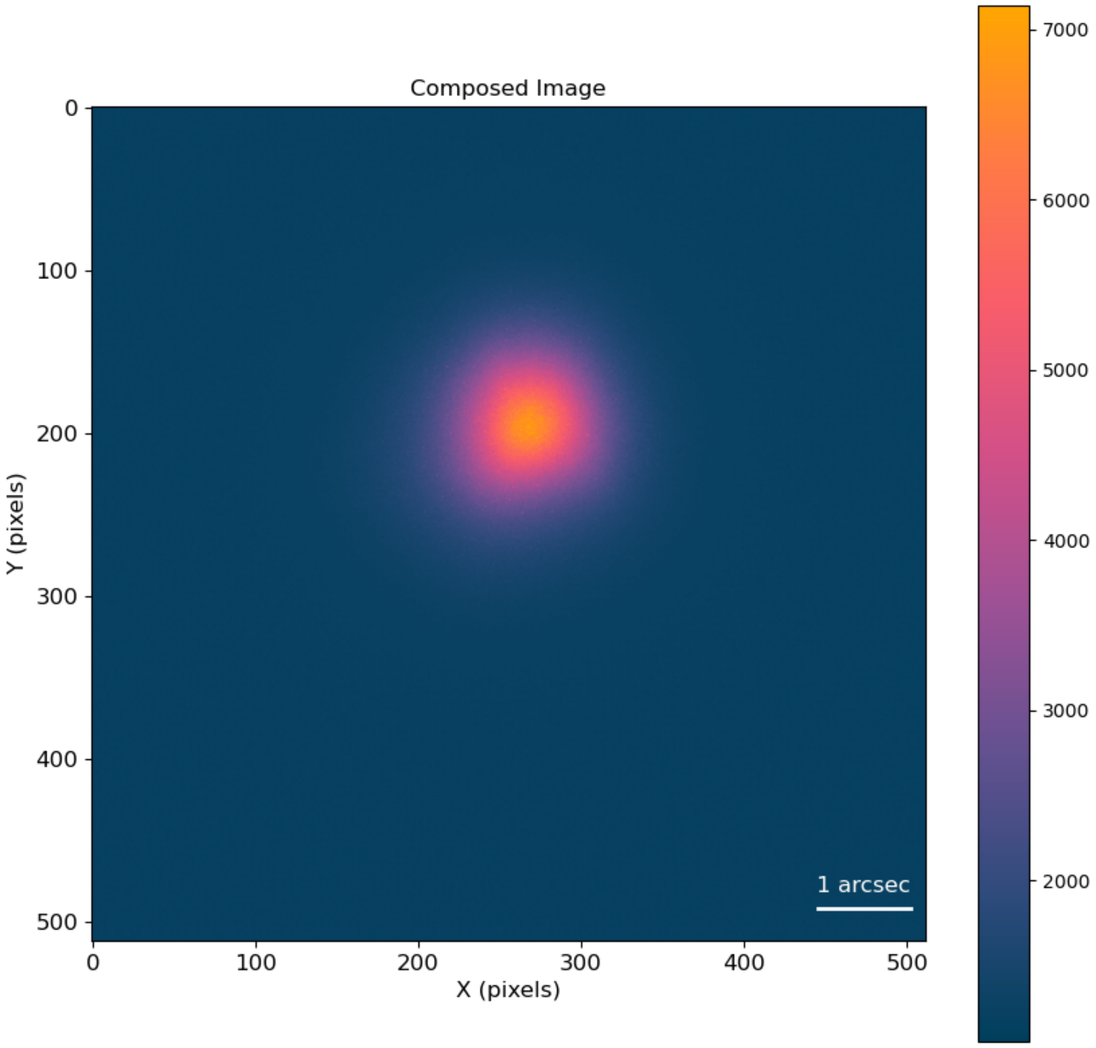

Figure 4.4: Composed image from averaging the counts/pixel of 5000 single 1 ms exposures.

to estimate the background counts per pixel. The annulus is defined far from the PSF to ensure that there is no contamination from the source, allowing for the estimation of the actual background. This background contribution is then subtracted from the total counts within the aperture to obtain the true flux of the source.
A good estimate of the star's PSF size corresponds to the radius that encircles 80% of the total flux. However, in this case, in order to maintain coherence with the results obtained from the PSF fitting, we assume that the width is represented by the radius that contains 50% of the counts.

The PSF fitting is carried out by identifying the curve that best represents the 3D shape of the PSF. In this project, the fitting is performed using a Gaussian function. To enhance accuracy, certain constraints are applied, similar to the speckle fitting process: the peak position is fixed at the center of the star, and the peak value is restricted to the maximum counts within the PSF.

4.2 Wavefront Sensing Code

The WF analysis pipeline, in short, is built in such a way that:

1. It takes as input the frames from the targets produced by the SH-WFS (CSV files),
2. It finds the best-fit set of Annular Zernike polynomials and, more importantly, the coefficients that reproduce the measured WF. Annular polynomials are conceptually the same as the regular ones, but they take into account the obscuration caused by the shadow that the secondary mirror of the telescope produces on the primary,
3. It separates low-order (tip, tilt) from high-order terms (the actual optical aberrations), computes residuals and other metrics (like RMS and peak-to-valley),
4. for each target it gives as an output three wavefront cubes which contain: the total fitted WF, the high-order only WF and the low-order one (in FITS files). The pipeline also produces a table of metrics (in CSV files) to allow fast access to the numerical results.

The code is written in Python and utilises standard libraries such as `Numpy, Scipy, Astropy`, as well as less-common ones like `Poppy`, which proved to be particularly useful since it contains pre-written expressions for the Zernike polynomials.

This code is available on GitHub at https://github.com/Leonardo-5/Donut-Zernike.

4.2.1 Data Reading and Cleaning

As a first step the pipeline takes in input the folder containing all the targets to be studied. Then, every CSV file found in the folder is processed independently, with a different iteration per file. For each file the code uses a processing function that returns a dictionary containing the results. Each CSV is parsed to extract the measured wavefront map (saved as a 2-D array), the sampling coordinates (x, y) and part of metadata (like the estimate of the pupil center given automatically by the SH sensor). Then, from the measured WFE map, the code estimates the pupil centre (x_c, y_c) and the pupil diameter D. These values are used to check for the presence of NaN values inside the pupil region. If such values are present, the frame is rejected, so those files will not be processed. This is necessary to avoid errors in the Zernike polynomials fitting process since the polynomials do not allow ”empty” regions inside the pupil.

Due to malfunctions in the SH-WFS data acquisition system, some targets had only 5 or 6 useful frames. This will be a problem later, since with a smaller sample, we get less reliable estimates for the aberrations affecting the WF.

The code then applies a resampling to the data, increasing the number of data points in the WF without altering its original shape. This resampling is performed with an interpolation function.

A mask is then applied to the data to account for the obscuration due to M2. The pupil coming into the sensor shows the presence of M2 obscuration (see Fig. 4.5). However, SH-WFS raw data show a wavefront map without gaps in it (Fig. 4.6).

This happens because the sensor tries to fill the empty regions by interpolation. For this reason, the code applies a mask where data should not be present to ensure that we are using only the actual measurements.

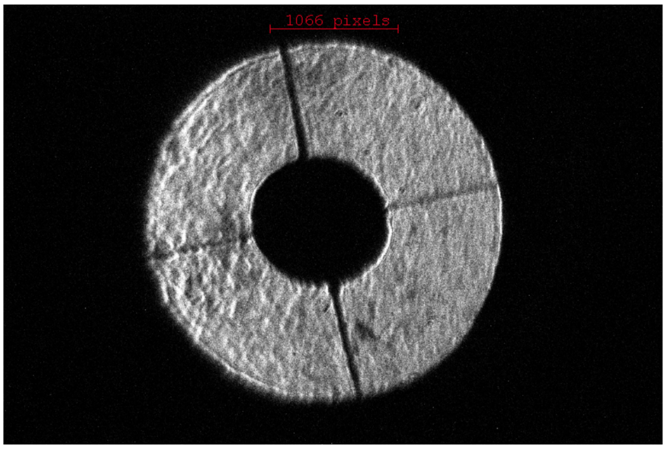

Figure 4.5: Image of the telescope pupil observed from the Coudé with M2 and its support spiders clearly visible.

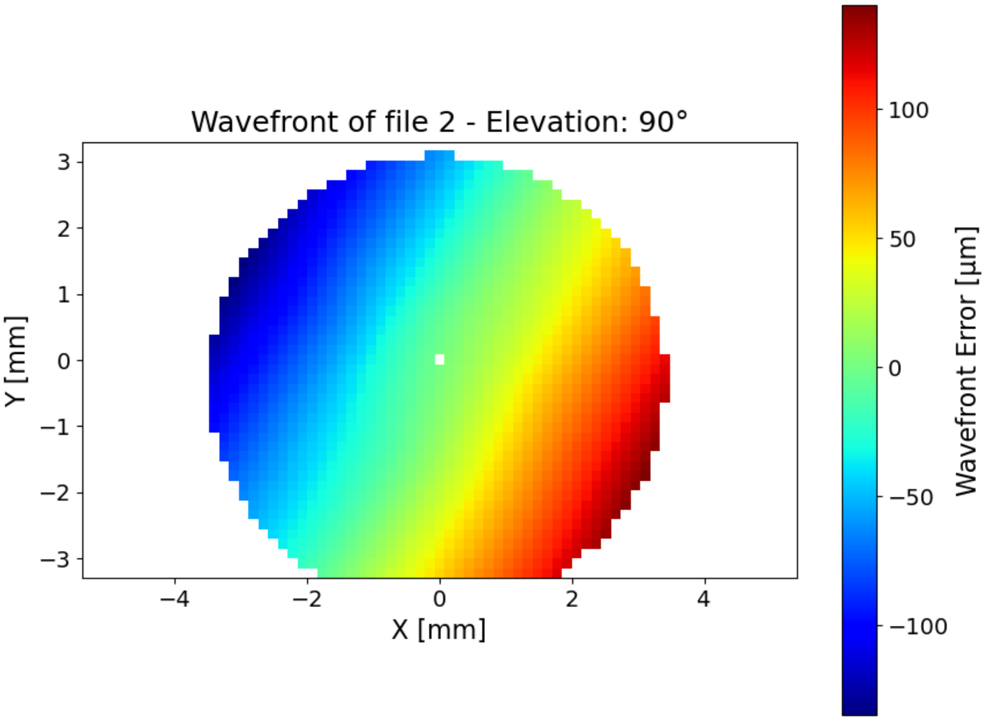

Figure 4.6: One of the wavefront maps measured by the SH showing the absence of the central obscuration.

Cartesian coordinates are then normalized to map the pupil onto the unit disk where Zernike polynomials will be computed. Normalized coordinates are computed as:

$$\begin{aligned} x_n &= \frac{x - x_c}{D/2} \\ y_n &= \frac{y - y_c}{D/2} \end{aligned} \tag{4.1}$$

and then converted to polar coordinates:

$$\begin{aligned} R &= \sqrt{x_n^2 + y_n^2} \\ \Theta &= arctg(y_n, x_n) \end{aligned} \tag{4.2}$$

Since the telescope pupil is centrally obscured, the pipeline also applies the mask to this normalized pupil; values outside the annulus are set to NaN

4.2.2 Wavefront Reconstruction

An annular Zernike basis up to a chosen number of modes N (where $N = 150$ in this case) is generated on the normalized polar grid.
The choice of N must be balanced with the sensor sampling: fitting too many modes than are supported by the data can lead to noise amplification and overfitting.
To build the basis, the function `Zernike1` from `Poppy` is used. This function enables the easy construction of a Zernike basis comprising N elements, where the polynomials are defined on an annulus rather than a full disk.

The Zernike coefficients are obtained by solving a linear least-squares fit between the measured wavefront and the Zernike basis. `Numpy lstsq` function is used. After this process, we have the best fit values for the Zernike polynomials, so we have an estimate for the aberrations.

The pipeline then computes the fitted wavefront using the equation 2.1, leading to a wavefront represented using Zernike polynomials up to an arbitrary number of terms.

A low-order wavefront (WF_{low}) is obtained by summing only the first k fitted modes ($k = 3$ in this case, to exclude the piston and tip/tilt). The high-order wavefront (WF_{high}) is then obtained by the difference between the fitted wavefront (WF_{fit}) and the low-order one, thanks to the orthogonality of the polynomials, which allows the addition or removal of some modes from the total wavefront without changing the others.
This decomposition permits a separate analysis of residual, higher-spatial-frequency aberrations after removing the dominant low-order terms.

In the last computational part of the code, the residuals between the measured WF and WF_{fit} are computed. Residual maps are computed (measured$-WF_{fit}$, measured$-WF_{low}$ and measured$-WF_{low} - WF_{high}$). Then, two other metrics are saved for each frame: the root-mean-square (RMS) of the high-order residual inside the annulus, defined as:

$$RMS = \sqrt{\frac{1}{N_{mask}} \sum_i W_{res,i}^2} \tag{4.3}$$

where N_{mask} is the number of pixels in the residuals map (W_{res}), and the peak-to-valley (PV), defined as the difference between the maximum and minimum residual within the mask. These values, together with the coefficient array (which contains all the fitted Zernike coefficients), provide diagnostics of fit quality.

4.2.3 Results Saving

For each processed frame, the pipeline records a result dictionary containing filename, coefficients, measured wavefront, reconstructed total/low/high components and residual maps of the WF, and the other metrics. After all frames of a target are processed, we have:

- Lists of 2-D arrays (measured WF, fitted WF, low/high, residuals) are stacked into 3-D cubes and written to FITS files.
- A CSV table is produced with one row per frame containing RMS, PV and each Zernike coefficient. Column names are build in a way to include a label, the Noll index and the name of the corresponding optical aberrations.

For each target, the set of results and the pipeline parameters (e.g. mode number N, excluded low modes k, obscuration ratio, resampling ratio) are saved alongside the results to allow reproducibility. The quality of the outputs is presented in the Results chapter (chapter 5).

4.3 The Donut Code

In this section, the numerical pipeline developed to analyse donuts is described. It is based on the code developed by Tokovinin and Heathcote (2006) [1] but implemented in `Python` instead of `IDL`. The code is organised into separate blocks that perform image preprocessing, fitting of models, estimation of Zernike coefficients, and comparison with WFS measurements, respectively.

The code is available on GitHub at: Donut-Code (n.d.) [12]

4.3.1 Image Preprocessing

The code begins by loading the directory containing the image files. For each file, the code performs the following steps:

- **Image loading**: the images are read and converted into arrays.
- **Frame stacking**: since the single frames are obtained with an exposure time of a few milliseconds (as explained in Section 2.3), the code averages the initial frames to reduce noise, increase the signal-to-noise ratio (SNR), and average out the inhomogeneities that characterise the individual frames. In this way, the code obtains an image equivalent to one obtained with a few-second exposure time. An example of a single and stacked image can be seen in Fig. 3.9.
- **Cropping**: the stacked image is cropped down to a square region around the donut to minimise the size of the computing grids.
- **Rebinning**: the cropped image is rebinned to a selected size (e.g. 32×32 pixels) to cancel the smaller features in the donuts that cannot be fitted by the code and can induce errors in the fitting process.

- **Masking**: a circular mask is applied around the rebinned donut to completely remove the background. This prevents background pixels with non-zero counts from affecting the fitting process, as the code would otherwise attempt to model those pixels even though they are not part of the donut, leading to incorrect results.

4.3.2 Donut Modelling and Fitting

For each preprocessed donut image, the code:

- initialises a set of Zernike coefficients describing the optical aberrations up to a chosen order;
- estimates initial guesses for the lowest order aberrations using the formulas shown in 2.14 by applying them to the stacked images;
- computes a model of the donut given the telescope diameter D, central obstruction ϵ, wavelength λ, plate scale p and the initial guesses for the Zernike coefficients;
- performs a non-linear minimisation of the difference between the observed donut and the model using a `numpy` least-squares method. The fitting cycle stops when the N^{th} donut model returns sufficiently low residuals when compared to the observed one or when the code reaches the maximum allowed number of cycles, to handle the case where the code couldn't find a fitting solution.
- estimates the covariance matrix of the fitted parameters in order to compute the uncertainties of the fitted Zernike coefficients.

Numerical Grids

The code uses three different spatial grids to perform its computations. The first grid is the detector grid, which represents the pixel sampling of the observed donuts. The images are cropped around the donut and then rebinned to a size of $N \times N$ (the code uses 32×32). This fixed grid ensures that all fits are performed on the same sampling, independently of the original resolution of the images.

The second grid is the pupil grid: a high-resolution square array of size $N_{\text{pupil}} \times N_{\text{pupil}}$ (set to 64×64 as default in the code) representing the telescope entrance pupil. This is where the wavefront phase is computed and the Zernike polynomials are evaluated. A finer grid is required for this scope because the Zernike modes are more reliably determined if sampled over fine grids.

For the third one, the code defines an image plane grid (so the FFT) of the same size as the pupil grid. The pupil field is Fourier transformed on this grid to generate the model donut image. This means in short, that the image that we see is the FFT of the wavefront phase distribution on the pupil plane. After the FFT, the synthetic donut image is finally or resampled back onto the detector grid so that residuals can be computed pixel by pixel over the observed image (which was rebinned to remove the smaller scale details, which are difficult to reproduce for this code).

4.3.3 Output

The output of the code then is a set of best-fit Zernike coefficients and their associated uncertainties, along with other elements explained in this section.

Error on the best-fit Zernikes

The errors in the fit of the Zernike coefficients are estimated using the Jacobian of the residuals from the non linear least squares fit. With $\mathbf{r}$ we denote the vector of the residuals between the observed and best fit donut image, and $\mathbf{J} = \partial\mathbf{r}/\partial\mathbf{z}$ the Jacobian with respect to the Zernike coefficients array $\mathbf{z}$. The mean squared error is defined as $\mathrm{MSE} = \langle \mathbf{r}^2 \rangle$, and so the covariance matrix of the coefficients is written as

$$Cov(\mathbf{z}) = (\mathbf{J}^T\mathbf{J})^{-1}\,\mathrm{MSE} \tag{4.4}$$

The standard errors for each coefficient are then computed as $\sigma = \sqrt{\mathrm{Cov}}$ and scaled by $\lambda/(2\pi)$ to convert them to length units.

Field Rotation Correction

Since the donut acquisitions are performed at different elevations angles, the pupil rotates with respect to the camera. As a consequence, some Zernike modes (astigmatism and coma) must be transformed into a common reference frame before they can be compared.

To handle this, the code performs a derotation in the same way as for the WFS data (see section 3.2). In this way, it is possible to compare directly the fitted Zernikes and the measured ones coming from the WFS.

Reference WFS Data Reading

If a directory containing WFS data is provided, the code searches for the file whose elevation is closest to the current donut image. If the code finds such file, it extracts the mean measured Zernike coefficients and their standard deviations.

These values are used as reference measurement of the telescope aberrations, to allow validation of the donut fitting method. The WFS values are used as a reference because they are measured quantities and, as a consequence, are more reliable than the fitted ones.

For each processed donut image, the code produces a `CSV` file as output that contains the following elements:

- original and derotated Zernike coefficients;
- uncertainties of all fitted parameters;
- model residuals and reduced chi-square;
- WFS values and their errors (if available).

Plots

As a last step, a function creates a set of four plots for each image stack:

1. **Observed donut image**: in the first plot, the original donut is shown.
2. **Best-fit donut model**: the model corresponding to the best fitted Zernikes is shown next to the observed image for visual comparison.
3. **Fitted Zernike coefficients**: the low-order coefficients are plotted with their associated uncertainties.
4. **Comparison with WFS**: if WFS data are available, the donut and WFS Zernikes are plotted with error bars, and with their difference, to allow a direct comparison between the results.

The script then repeats the entire analysis over all image files contained in the selected directory; it checks for missing data, and prints unexpected exceptions (both in the dataset and in the numerical results) without interrupting the run. Once all files have been processed, the code prints the output path where all CSV files and plots have been saved.

In summary, this pipeline provides an almost completely automated framework that processes donut images, fits donut models, estimates low-order Zernike aberrations, corrects for field rotation and compares the results with WFS measurements.

Chapter 5

Results

In this chapter, the results obtained from the three datasets will be discussed.

5.1 Speckle Results

5.1.1 Speckle evolution with exposure time

As discussed in section 4, the code starts by opening one image of the stack, that appears as shown in Figure 5.1.

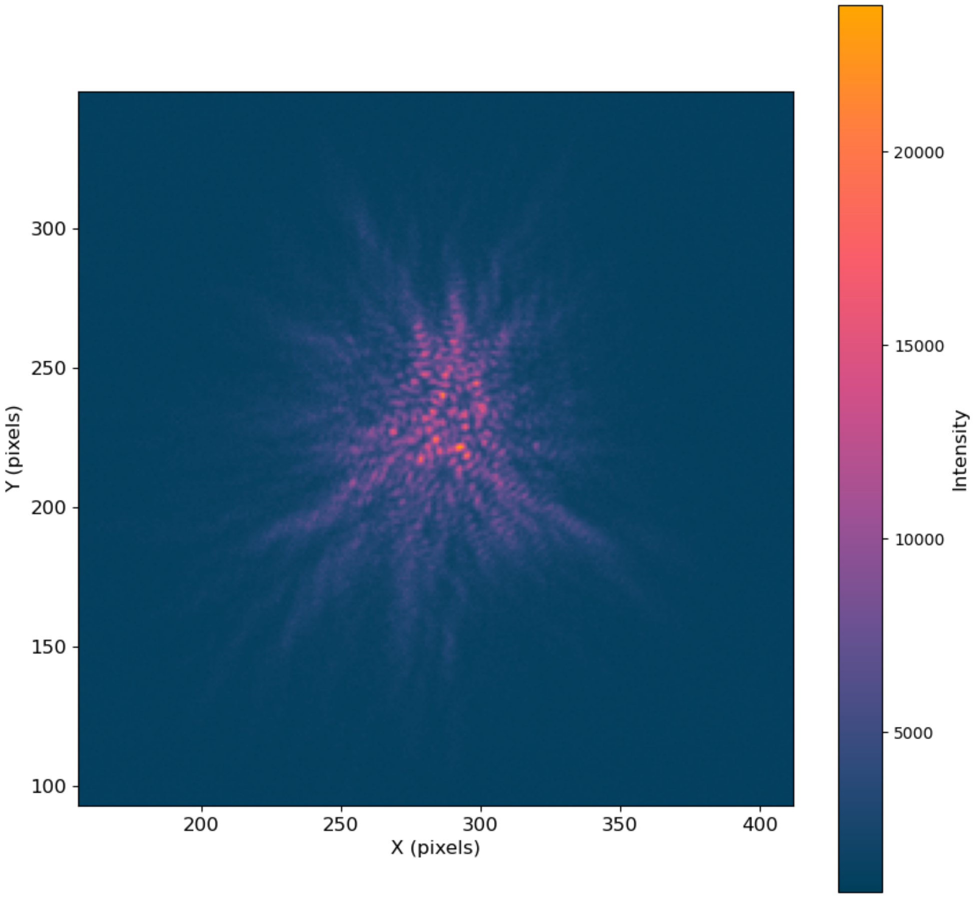

Figure 5.1: Single speckle image taken with a 1 ms exposure time

Then, depending on the counts threshold value, the position of the speckles are found: in the case of Figure 5.2 the threshold is set to 7000 counts per pixel.
Next, each flagged pixel is used as initial guess of the speckle center for the Gaussian fit. The resulting FWHMs are stored and used to produce the histogram in Figure 5.3. We can see how the FWHM of the speckles (the ones found in the same image as Figure 5.1) are distributed. Some of them, however, are not circular as they should be. This problem can be solved by implementing an extra check on the shape of the speckles, and reject the most elliptical ones. Highly elliptical speckles are discarded because they are usually not true diffraction-limited speckles: such shapes can arise from speckle merging, low SNR regions, or can be residuals produced by strong aberrations, especially towards the outer regions of the PSF where the speckle contrast decreases. Including these features would lead to incorrect estimates of the speckle size and as a consequence of the parameters derived from them.

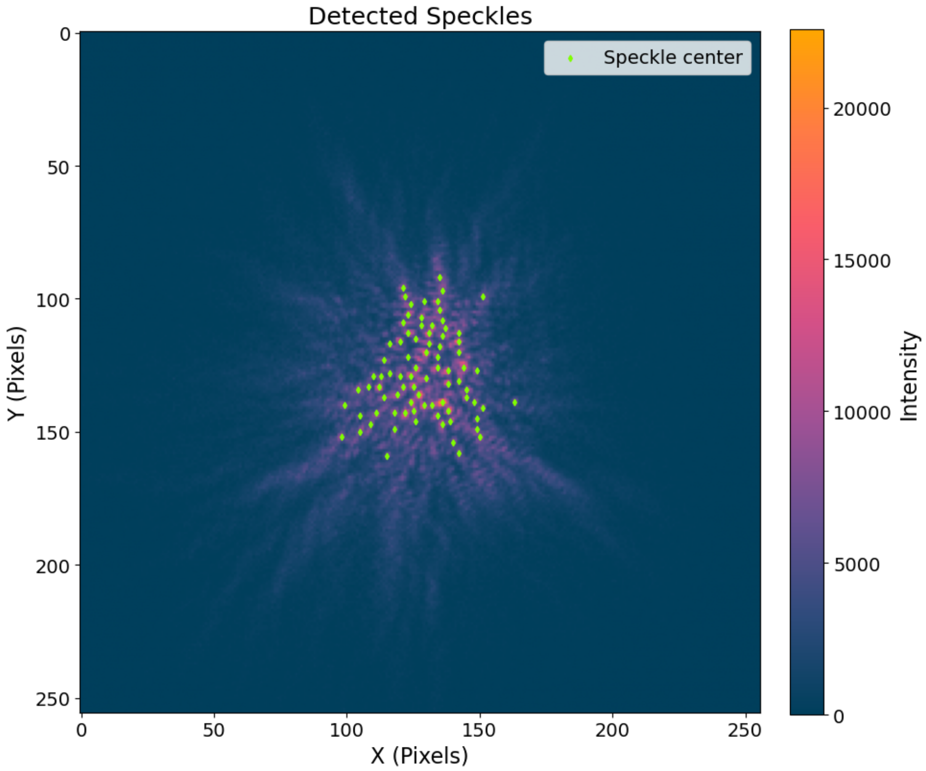

Figure 5.2: Single image with the a marker on the speckles detected by the code using the search procedure described in Section 4.

In this case, we have that the median of the distribution lies ~ 0.173 *arcsec* above the diffraction limit. The spread in the results is likely due to errors in the fitting process, which may have been caused by the crowding of certain areas of the image or by the erroneous recognition of some pixels as speckles that are actually indicative of other effects. Crowding, in particular, represents a problem because if two or more speckles are too close to each other, their external areas tend to overlap, thereby altering the (ideal) Gaussian profile that they should exhibit. This can result in incorrect estimations of their FWHM.
The median value of this distribution is then stored together with the corresponding value of every other image in the stack. These data are finally used to produce the histogram in Figure 5.4b.

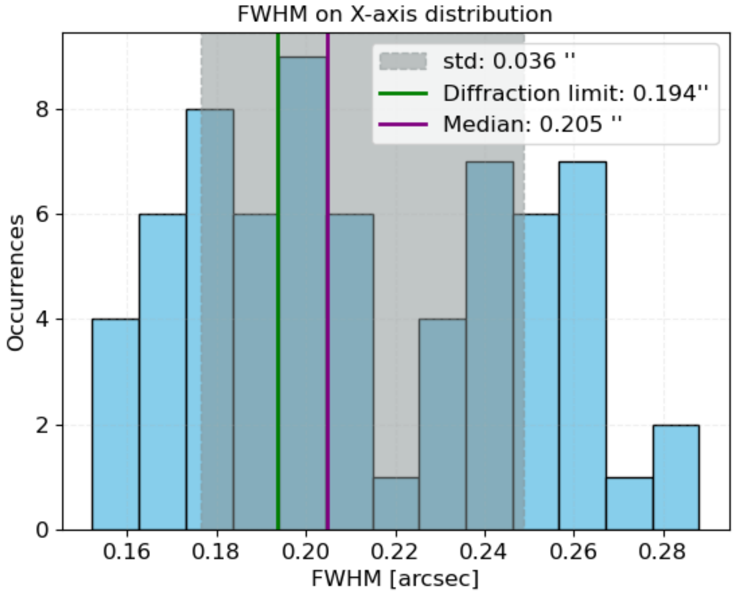

Figure 5.3: distribution of the FWHM of a single image. The green line shows the diffraction limit, the purple one the median of the distribution and the grey band shows the width of the STD of the speckles' FWHM distribution.

The procedure is then repeated until we have the median value of the FWHMs of each image. These data are used to produce a histogram that shows the distribution of the median FWHM of the three stacks:

Figure 5.4a shows the distribution of the medians of the stack obtained with a 2 ms exposure. The median is 214 mas, so 20 mas is larger than the diffraction limit. Figure 5.4b shows instead the data relative to 1 ms of exposure. In this case, the median is 212 mas, 18 mas larger than the diffraction limit. As can be noted, the number of images for which the median is equal to the diffraction limit is larger than can be seen in the 500 Hz case. Finally, the images with 0.1 ms exposure show the best results (Figure 5.5): here the median is just 3 mas larger than the diffraction limit. This means that speckles are not superimposed one on the other; they appear to be almost completely 'frozen' and thus, this exposure time is similar to the coherence time of the turbulence in the atmosphere.
This is expected because a shorter exposure time should give a snapshot of the turbulence and, as a consequence, speckles are expected to approach FWHM values similar to the diffraction limit of the telescope, as long as the telescope is aligned correctly. This trend is also confirmed by the plot in Figure 5.6: the dots represent the median values of FWHM (as shown in Figure 5.4a, 5.4b and 5.5) while the line shows the fitting function between the two variables. Following this relation, we expect to reach a median FWHM equal to the diffraction limit for 19.52 kHz, or ~ 0.05 ms exposure time.

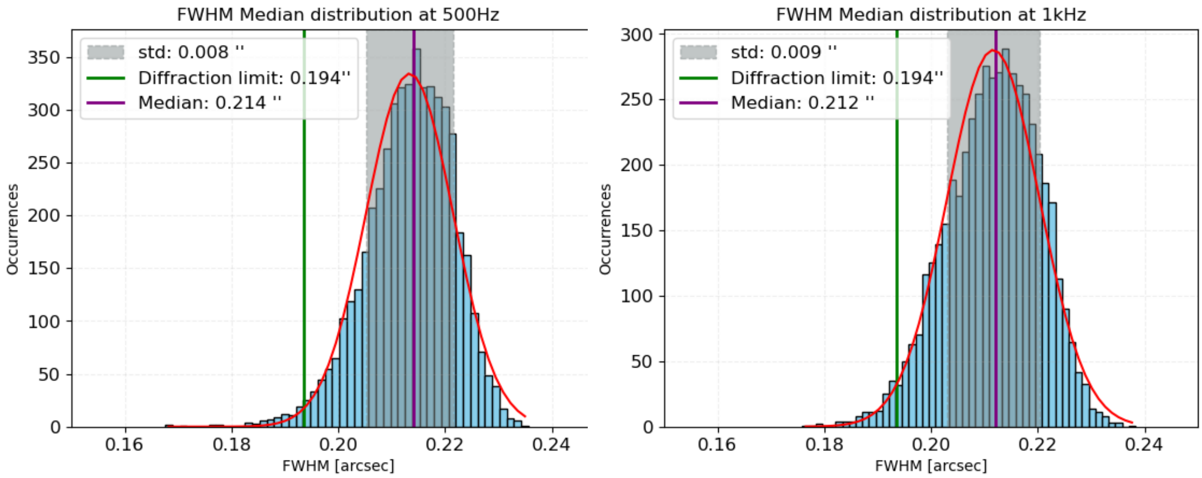

(a) Distribution of the median FWHMs of the Nasmyth speckles at 500 Hz.

(b) Distribution of the median FWHMs of the Nasmyth speckles at 1kHz.

Figure 5.4

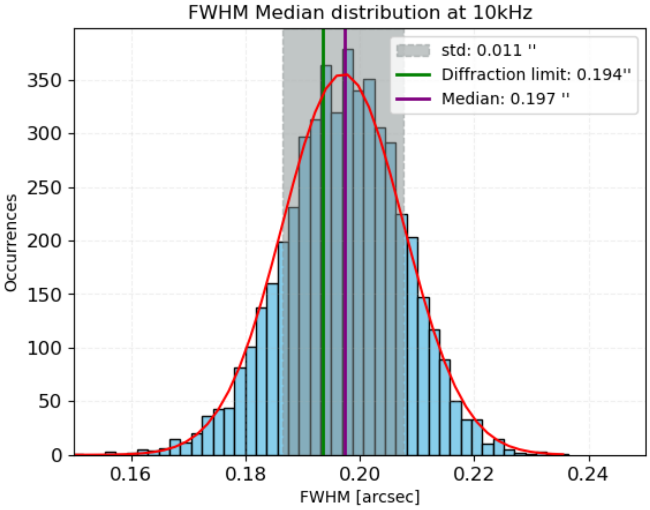

Figure 5.5: Distribution of the median FWHMs of the Nasmyth speckles at 10kHz.

Centroids and RMS

The fitting process also provides information about the positions of the speckles. These data are used to determine the barycenter of the PSF and analyse the distribution of speckles relative to it. Figure 5.7 displays the same image as Figure 5.2, this time including the barycenter. Additionally, one circle represents the root mean square (RMS) of the distances of the speckles from it, while another denotes the distance of the farthest speckle identified by the code.

The RMS provides information on the dispersion of the speckles around their barycenter. Figure 5.8a and 5.8b illustrate this with two histograms showing the RMS distribution across all images in the stack at 1 kHz. The mean value of the RMS provides an indication of the image jitter during the total observation time. The expected behaviour is to find lower values at shorter exposure times due to the reduced total exposure duration. This is confirmed by the plot in Figure 5.9. Each pair of dots

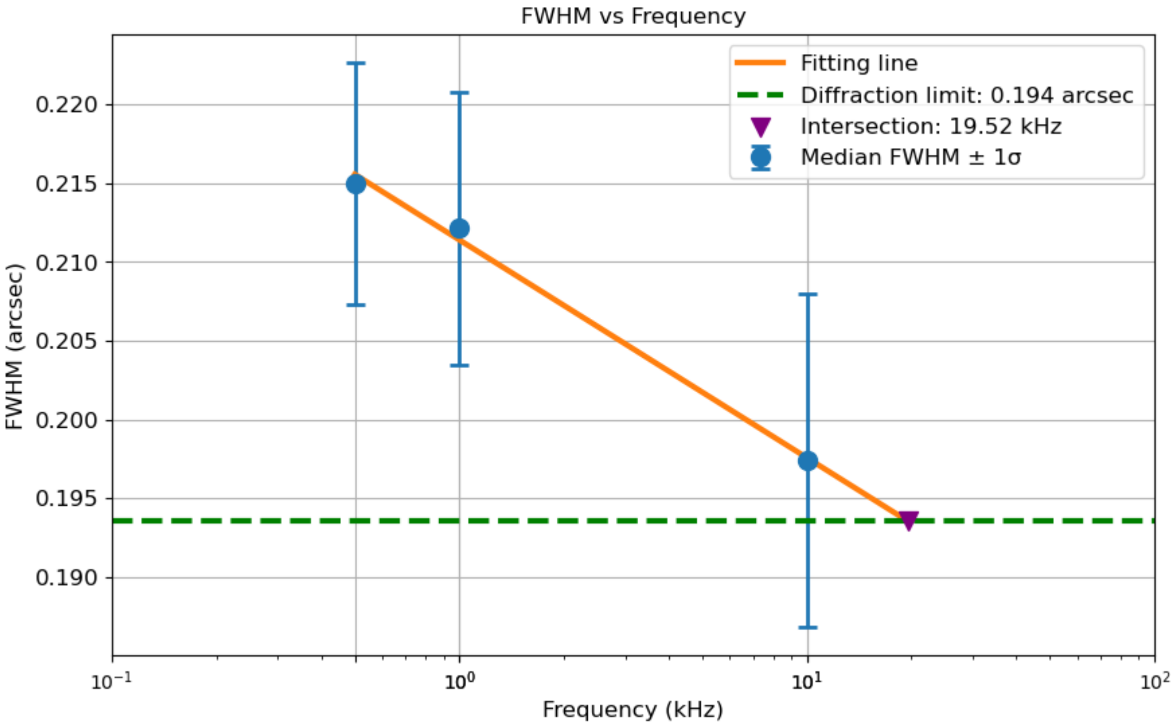

Figure 5.6: FWHM as a function of the frequency of the speckles images. The intersection between the fitting line and the diffraction limit shows the frequency (hence the exposure time) required to obtain speckles with sizes equal to the diffraction limit.

represents the mean value of the RMS along the X and Y axes for the three different image stacks.

By averaging the X and Y positions of the barycenters from individual images, the total centroid of the PSF was determined. Then, the total RMS of the barycenters was calculated. After this, the images were stacked to generate a PSF corresponding to a total exposure time of $5000 \cdot 0.002 = 10\ s$.

PSF fit and aperture photometry

Finding the PSF width is fundamental to obtaining the Fried parameter. Starting from the stacked image, the FWHM of the PSF has been computed using both Gaussian fitting and aperture photometry.

In Figure 5.10, we have the original image in the left panel and the fitted PSF in the right panel. In this case, the FWHM was found to be 77.40 pixels in the X direction and 80.48 pixels in the Y direction. Since the plate scale was found to be $0.0376 \quad arcsec/pixel$, the two FWHM translate to 2.92 and 3.03 $arcsec$ ($2.92 \times 10^3\ mas$ and $3.03 \times 10^3\ mas$) respectively.

The fit is acceptable since the residuals between observed and fitted PSF are in a good agreement. In particular, they show that the maximum discrepancy between real and fitted counts is ~ 600, which corresponds to the 10 % of the PSF's maximum counts, concentrated mainly in the wings of the PSF profile.

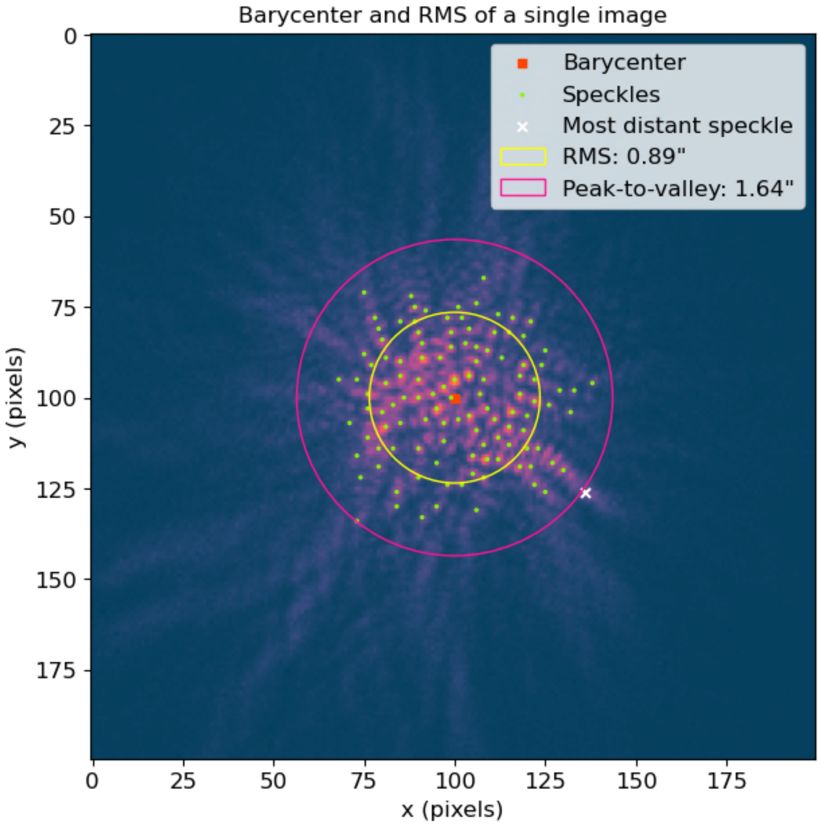

Figure 5.7: Barycenter and RMS of the speckle positions of a single image. The yellow ring shows the root mean square of the position of the speckles with respect to the barycenter. The red one shows the distance of the most distant speckle from the center

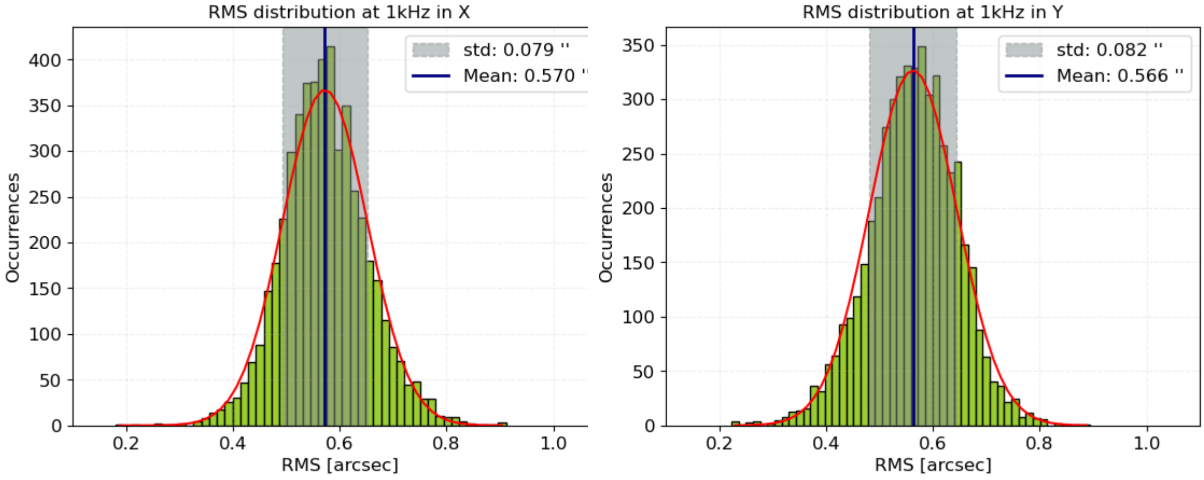

(a) distribution of PSF speckle jitter in X

(b) distribution of PSF speckle jitter in Y

Figure 5.8

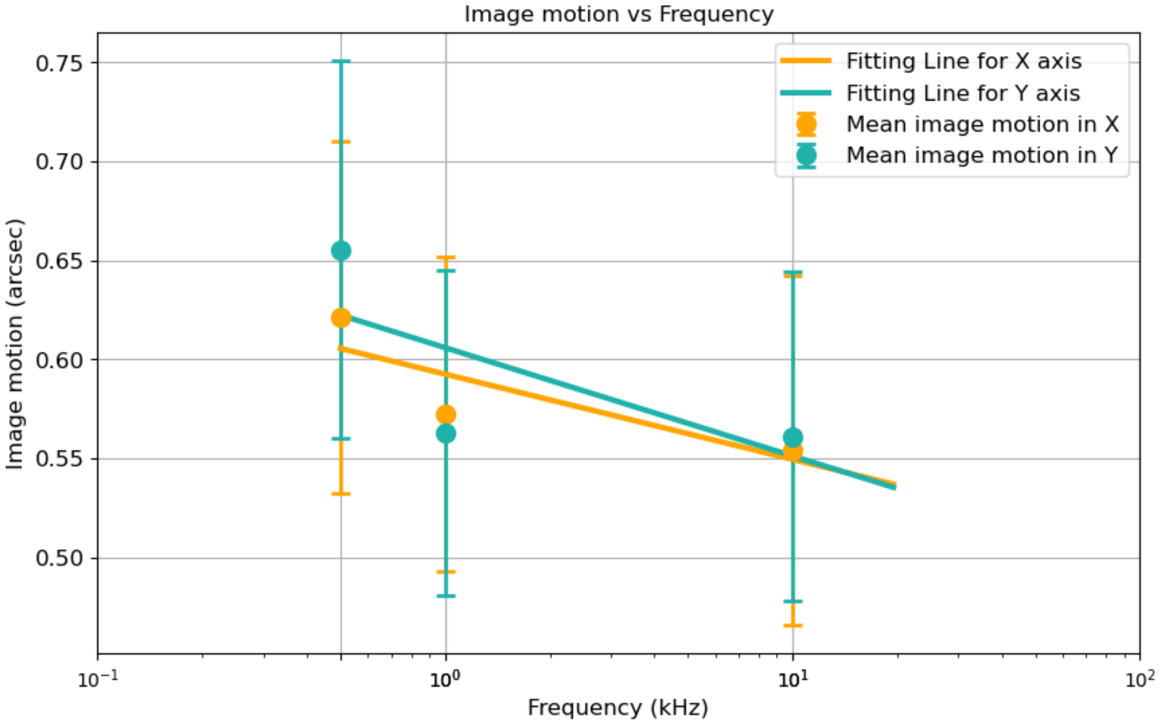

Figure 5.9: Image motion as a function of the frequency. The trend shows a slight decrease with the increase of frequency, both in X and Y axes.

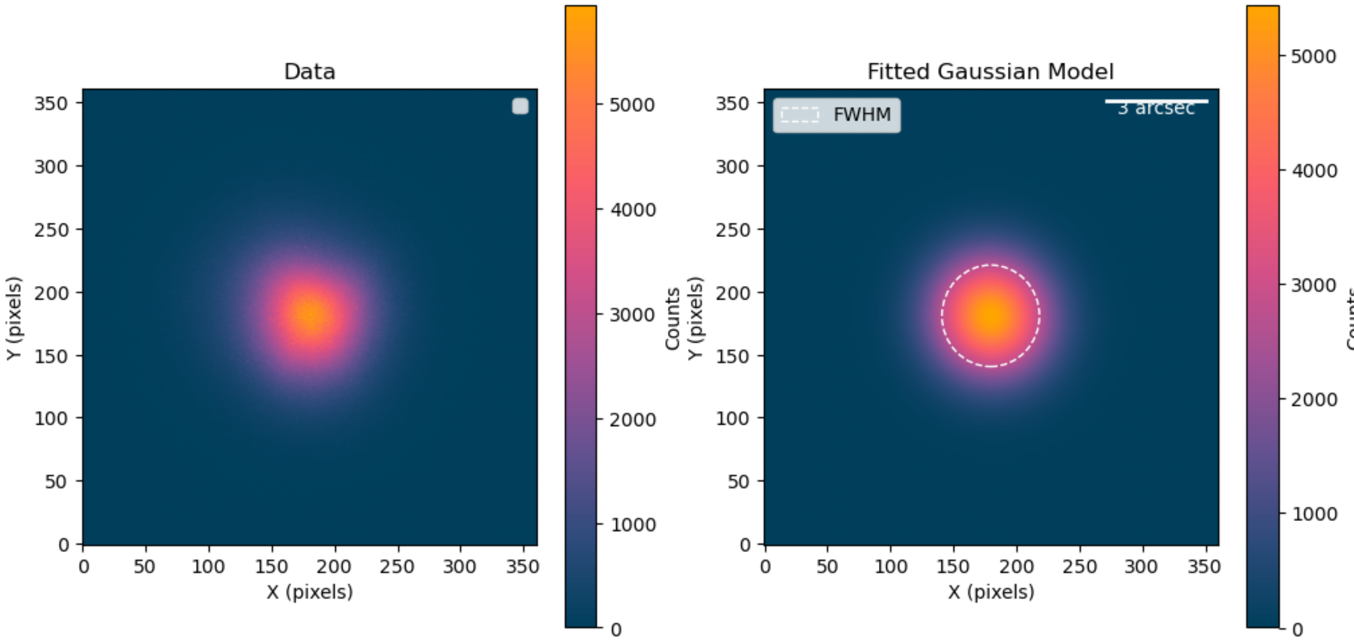

Figure 5.10: Comparison between the observed PSF (left panel) and the fitted one (right panel). The white ring represents the FWHM.

For aperture photometry instead, the results are visible in Figure 5.13. The curve of growth shows that the radius that encloses the 50% of the total flux of the PSF ($r50$) is 43 pixels or 1.62 $arcsec$. Following this result, the diameter of the PSF is 3.24 $arcsec$. The PSF circumference is highlighted in Figure 5.14, along with the back-

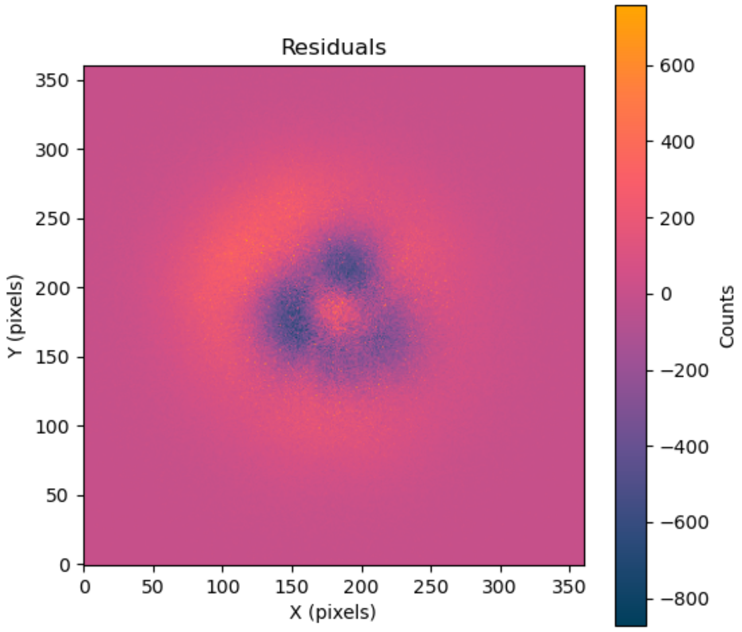

Figure 5.11: 2D map of residuals computed as the difference between the counts per pixel of the original image and the Gaussian fit ones. The largest residuals are located in the wings of the PSF since the true shape is not exactly symmetrical.

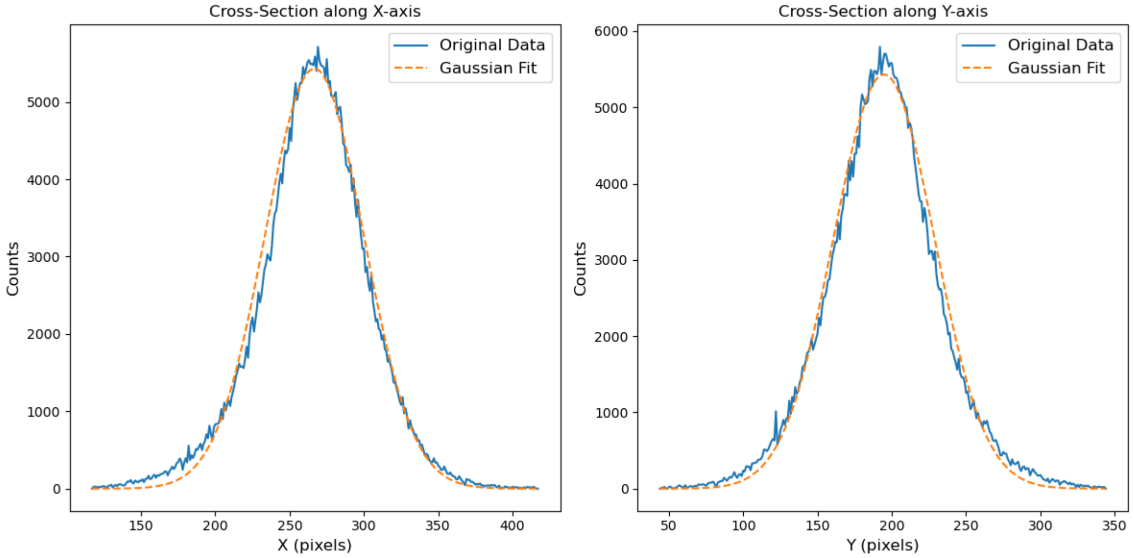

Figure 5.12: PSF and Gaussian profiles along X and Y axes. The small discrepancy between real and fitted PSF confirm the validity of the fit.

ground annulus (used to estimate the background noise) and the maximum aperture reached in the curve of growth (shown in blue). Nevertheless, r_{80} radius (defined as the radius that encloses 80% of the total flux) is typically used to estimate the PSF. In this work, however, r_{50} was chosen instead to maintain consistency between the Gaussian fit FWHM and the corresponding measurement obtained through aperture photometry.

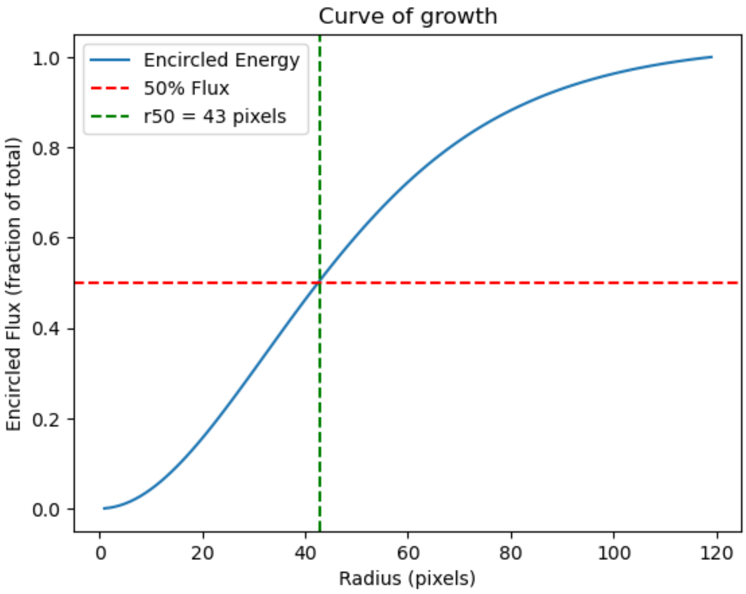

Figure 5.13: Curve of growth showing the radius that encircles 50% of the total flux of the PSF.

Once the FWHM is known, it is possible to compute the Fried parameter through the formula:

$$r_0 = \frac{0.98 \cdot \lambda}{FWHM} \qquad [m] \tag{5.1}$$

Where $\lambda = 700 \pm 90\ nm$ is the central wavelength and the width of the filter used during the observations, and the FWHM is reported in radians in order to obtain a dimensionless number.

The entire process (aperture photometry, PSF fit and r_0 computation) is repeated for the three image stacks. In this way it is possible to compare the results and check for variations that arise with changes in the exposure time.

After this, using the software AstroImageJ (credits: Collins et al. (2017) [13]), it was possible to obtain another estimate for the full width at half maximum (FWHM) of the point spread function (PSF) for each exposure time. This was done to obtain an external confirmation of the results obtained thus far.

A report of these results is shown in Table 5.1.

	500 Hz	**1 kHz**	**10 kHz**
r50 [asec]	3.46	3.23	3.45
FWHM (fit) [asec]	3.16	2.91	2.84
FWHM (ImageJ) [asec]	2.91	2.66	2.71
r0 (fit) [m]	0.044	0.048	0.051
r0 (photometry) [m]	0.041	0.043	0.040
r0 (ImageJ) [m]	0.048	0.052	0.052

Table 5.1: Table with the results derived from PSF fit, aperture photometry and AstroImageJ. Even though there are slight differences, the r_0 is consistent in every case, having a 1 cm range.

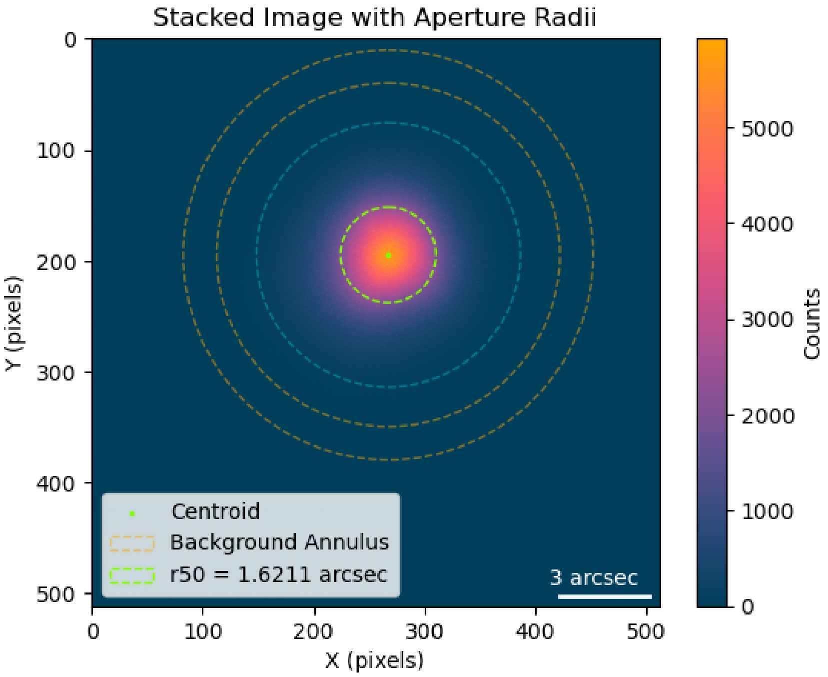

Figure 5.14: PSF with aperture photometry. The green circle is the representation of $r50$, the blue one is the maximum aperture used to produce the curve of growth, and the yellow circles show the region used to estimate the background.

5.1.2 Speckle evolution with elevation

In this section, the evolution of speckle size with changes in elevation angle will be discussed for both focal planes, and the results will be compared to highlight the differences in optical quality between the Nasmyth and Coudé configurations.

Speckles at Nasmyth

The data used for this analysis were obtained during the March 2025 observational campaign. The dataset consists of three targets observed at elevations of 30°, 60°, and 78°. The elevations were chosen to match as much as possible the elevations of the WFS dataset, since the setup did not allow simultaneous data acquisition. For each target, 1000 images were recorded with an exposure time of 10 ms, allowing the study of speckle evolution over a long time span.

The data were analyzed using the pipeline described in Section 4.1. The code produced the FWHM and ellipticity distributions of all fitted speckles across all frames for each targe; this allowed the study of the evolution of the speckle characteristics over the elevation angle instead of an exposure time-wise analysis, as discussed previously.

The distribution of the FWHM of all fitted speckles as a function of elevation is shown in Fig. 5.15.

In this plot, the FWHM of the speckles show the expected trend: speckles observed at higher elevations are smaller than those at lower angles. This happens because speckles are affected by the amount of atmospheric turbulence along the line of sight. At lower elevations, the light passes through a thicker and more turbulent air column, producing more aberrations and, as a consequence, larger speckles. Along with the airmass another cause of the increase speckles' sizes is the gravity-induced misalignment

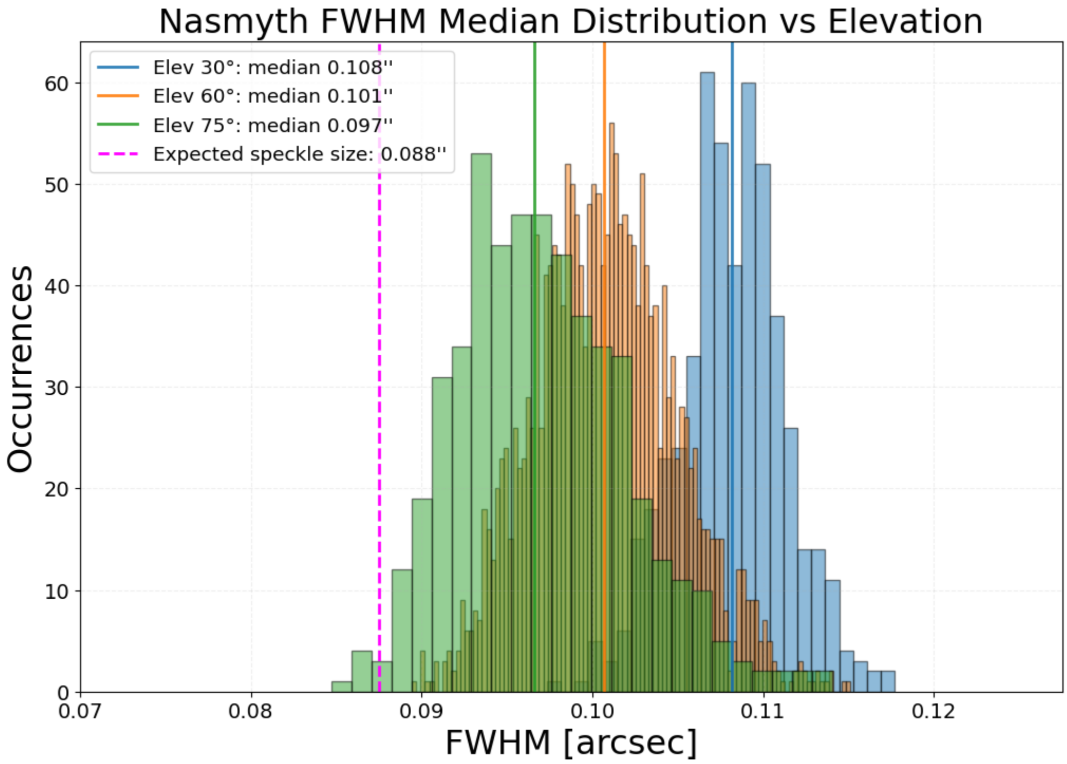

Figure 5.15: Speckle FWHM distribution as a function of the elevation angle at the Nasmyth focal plane

of the telescope optics which gets stronger with the decrease of the elevation angle. The pink dotted line in the plot represents the expected speckle size, which, under ideal conditions, corresponds to the size of the Airy disk. The fact that the observed speckles are larger than the expected size is a sign of the presence of optical aberrations along the optical path, which, as a result, increases the speckle size, particularly for smaller elevations.

Regarding the ellipticity of the speckles, results are displayed in Fig. 5.16. Speckle ellipticity is linked to optical aberrations introduced by the telescope's optics, particularly astigmatism. In an ideal, unaberrated system, speckles should be circular since they behave as Airy disks. However, if astigmatism is introduced, the original circular shape becomes distorted, and takes on a more elliptical form with the increase of astigmatism. Studying this parameter can provide initial information on the system's optical aberrations before using any WFS. Results show that ellipticity reaches a maximum for the target at 60°, suggesting that at this elevation there may be some misalignment or deformation in the mirrors that, in this particular configuration, combine to increase the total astigmatism of the system. This interpretation is consistent with the WFS data. Fig. 5.28 shows a similar trend in the measured astigmatism, consistent with the observed ellipticity of the speckles.

Speckles at Coudé

Speckle images at the Coudé optical bench were obtained during the last observing session in September 2025, using the setup shown in Fig. **??**. Four targets were observed at elevations of 80°, 70°, 50°, and 35°, producing 1000 images per target with an

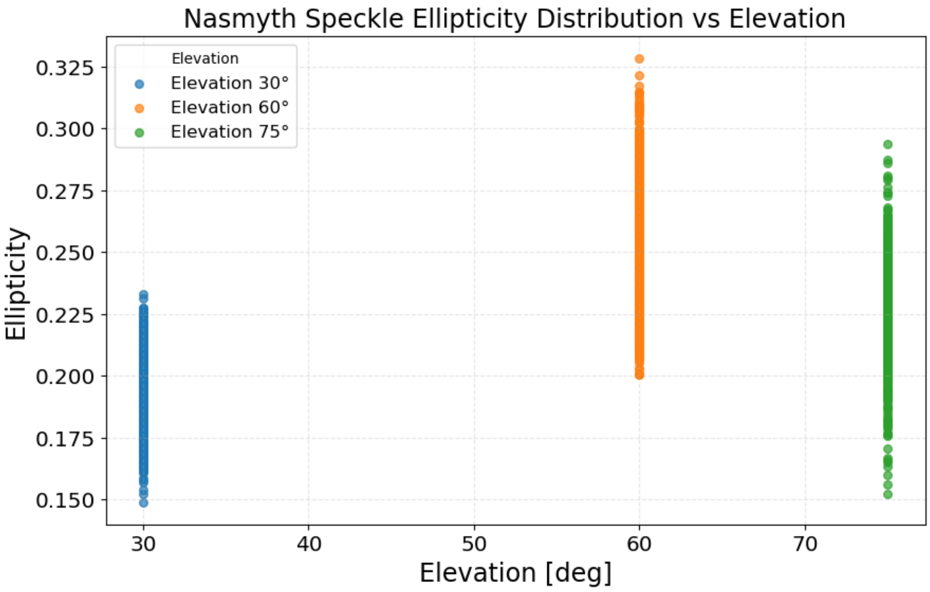

Figure 5.16: Ellipticity distribution of the speckles observed at the Nasmyth focal plane

exposure time of 1 ms.

The evolution of the FWHM distribution is presented in Fig. 5.17. Once again, the trend follows expectations: at higher elevations, we see smaller speckles. The only exception is the target at 35°, which shows smaller speckles than the one at 50°. This anomaly can be attributed to problems in the data.

Due to the long optical path from the telescope to the Coudé bench, a higher amount of optical aberrations is expected compared to the Nasmyth focus, resulting in more degraded images. The consequences of this are shown in Fig. 5.18, where the images appear blurred due to both atmospheric and optical effects. These images are also characterised by low photon counts due to light losses in the long optical path. These factors make the speckles in this case nearly indistinguishable, causing fitting errors and unreliable results for this particular target. There is still coherence in this phenomena: the most problematic case corresponds to the lowest elevation angle, consistent with the effects discussed previously.

Theoretically, by decreasing the exposure time enough, this problem can be solved. Eventually, speckles will begin to be resolved, but then the SNR becomes too low to perform any kind of analysis, so it is a problem of balance between SNR and exposure times low enough to have single speckles resolved. And since the Coudé has large light losses in the optical path, it was not possible to lower the exposure time more than 1 ms.

Regarding the ellipticity (Fig. 5.19, the trend remains on lower values compared to the Nasmyth results; the comparison will be discussed later in more detail. These values can be attributed to the degraded image quality and the short exposure time used, which resulted in lower photon counts per frame. Speckles at the Coudé focus generally appear less resolved than those at the Nasmyth and are less prominent against the background due to the reduced signal level. As a result, the fitting algorithm has greater difficulty in estimating the exact speckle shapes. It is also possible that the larger optical aberrations increase the overall speckle size without introducing significant deformations, smoothing the contribution to ellipticity. In the end, the speckles at the Coudé focus appear larger but rounder compared to those observed at the Nasmyth

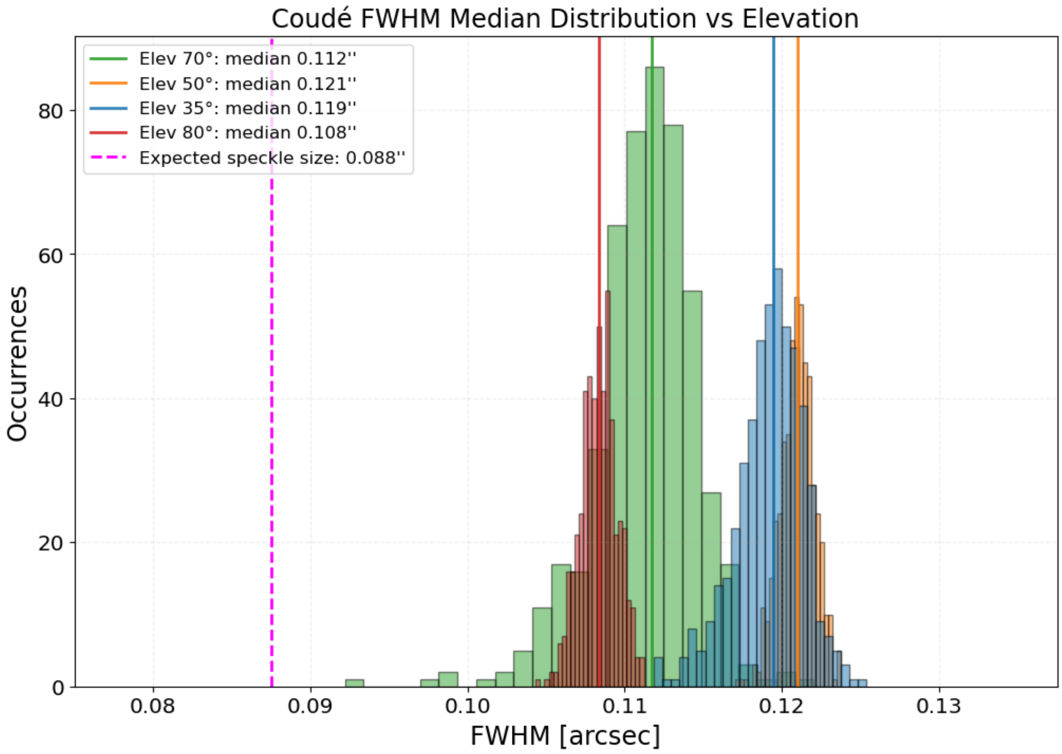

Figure 5.17: Speckle FWHM distribution as a function of the elevation angle at the Coudé optical bench.

focus.

5.1.3 Nasmyth and Coudé comparison

Figures 5.20 and 5.21 show the evolution of the previously discussed quantities in a manner that enables comparison.

The FWHM distributions show that at the Nasmyth focus the speckles are smaller than those observed at the Coudé. This result is expected, as smaller aberrations and, therefore, higher image quality should characterise the Nasmyth focus. In both cases, the speckle size decreases with increasing elevation, except for the target at 50° at the Coudé focus, for the reasons discussed earlier. The bars in the plot represent the standard deviations of the FWHM distributions.

The ellipticity results show different trends. At the Nasmyth focus, a peak is observed at 60°, consistent with the peak in astigmatism measured by the WFS. At the Coudé instead, the ellipticity peaks at 70°, but with overall smaller values (so we have rounder speckles in this case), likely due to the lower photon counts and blurrier images as discussed previously. In addition to these effects, this improvement may also be related to the recent aluminization of the primary mirror which was during the summer 2025, so before the September observations (the observations were we gathered the data that is being discussed). The procedure for the aluminization of the mirror requires that it is removed, re-coated and installed back into the telescope, which then requires a realignment of the optics. The realignment could have reduced some alignment-related. As a consequence, it is possible for the speckles produced under these conditions to be less elliptical.

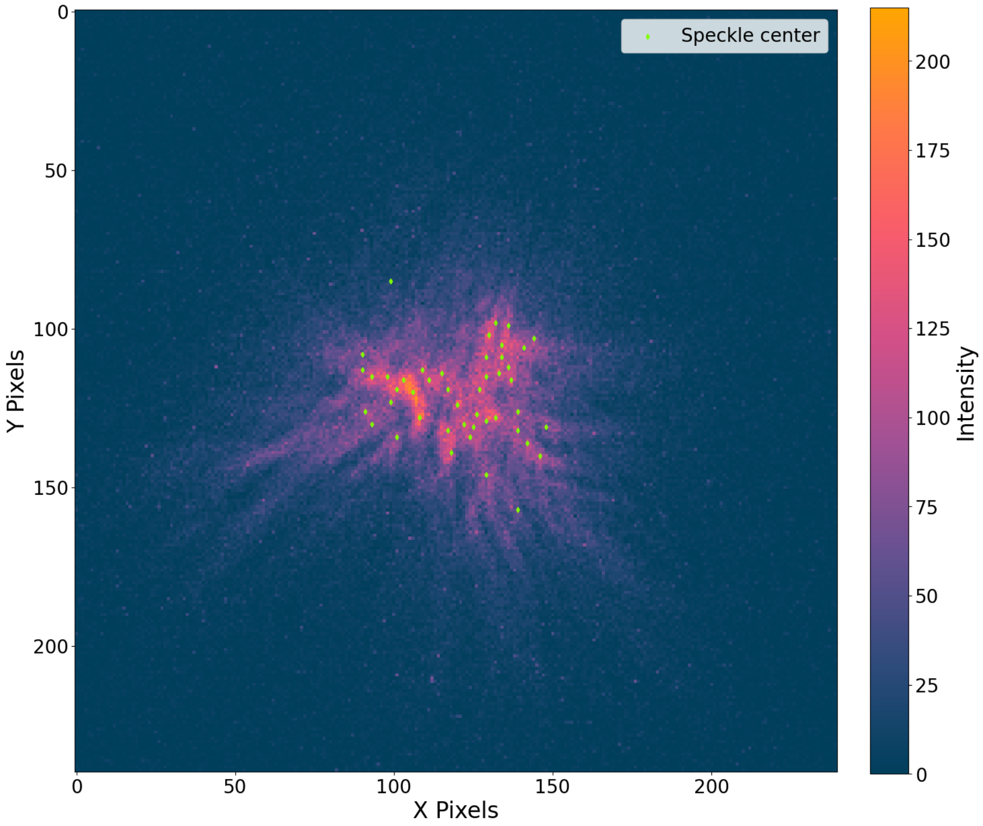

Figure 5.18: Speckle frame at elevation 35° showing the large amount of blur. Single speckles are difficult to distinguish, and this leads to errors in the fitting code and unreliable results. Speckles centers are the locations of speckles as detected by the fitting code.

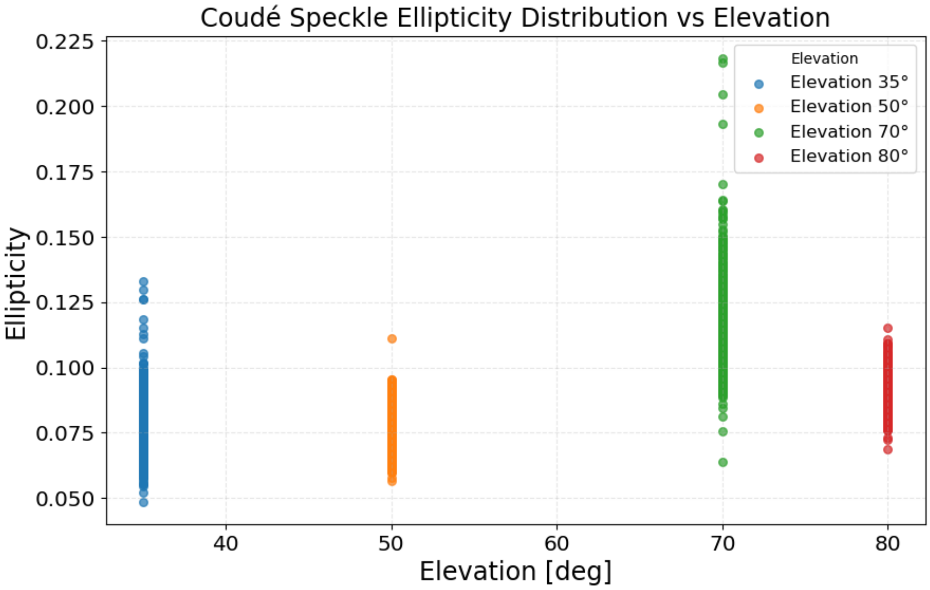

Figure 5.19: Ellipticity distribution of the speckles observed at the Coudé optical bench.

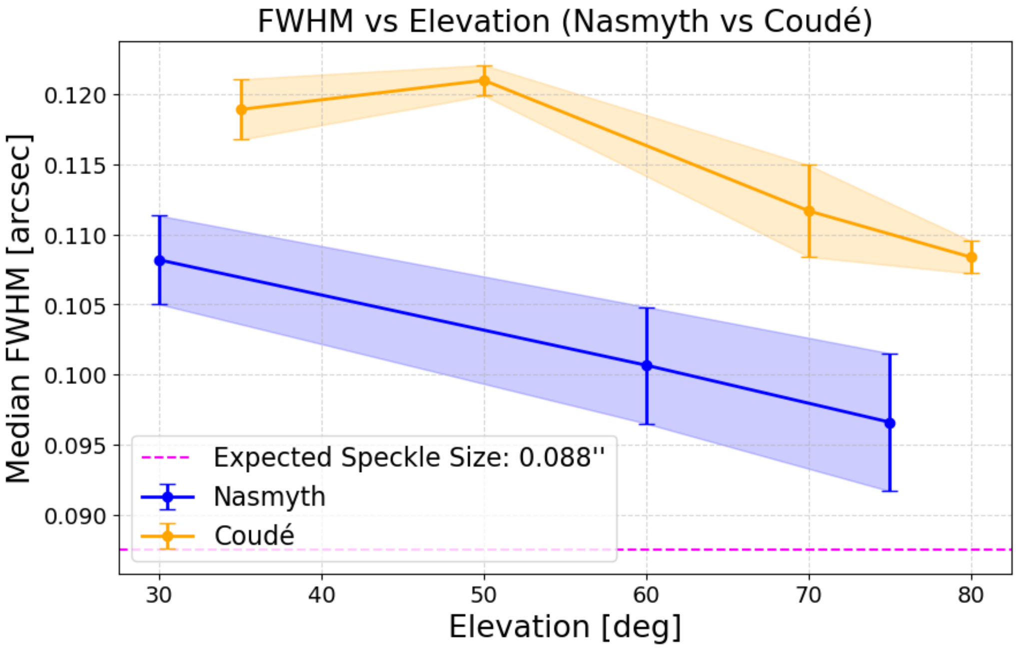

Figure 5.20: Comparison of speckles FWHM evolution as a function of the elevation.

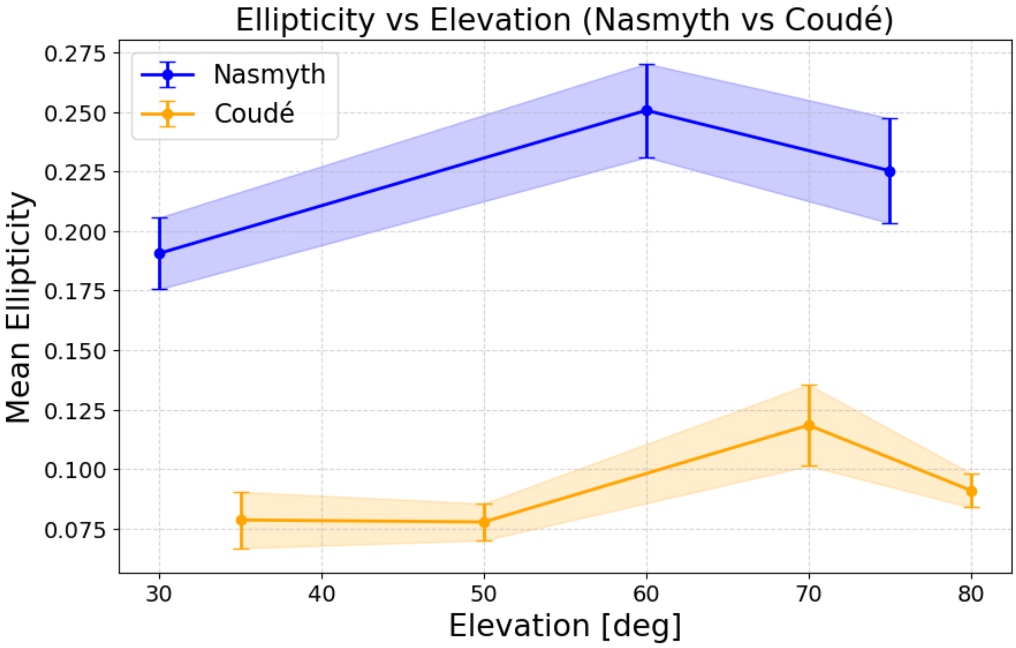

Figure 5.21: Comparison of speckles' ellipticity as a function of the elevation.

5.2 Wavefront Sensing Results

This section presents the results obtained by fitting an annular Zernike expansion to the Shack–Hartmann measurements described in Chapter 4.

As discussed in Section 4.2, the dataset is composed of many frames per target; therefore, we have various WF maps (with the exception of the discarded ones) to fit for each target.
The code initially reads the raw data, applies the mask, and resamples the WF. Fig. 5.22 shows the comparison between the initial data and the resampled data with the applied mask. The resampled one has a slightly smaller radius: this happens because, in this way, we ensure that the pupil is perfectly circular when we resample. Without doing this, we end up having pupils with jagged borders due to the fact that the original pupil, since is represented by a limited amount of pixels, is not completely circular (as can be seen in the bottom-right border of the pupil in Fig. 5.22a).
In Fig. 5.22, one can also notice that the WF maps are slightly out of the bounds of the detector. this does not represent a problem, since the non-sampled part of the pupil is minimal compared to the total surface; thus, we can use these data without losing reliability in the results.

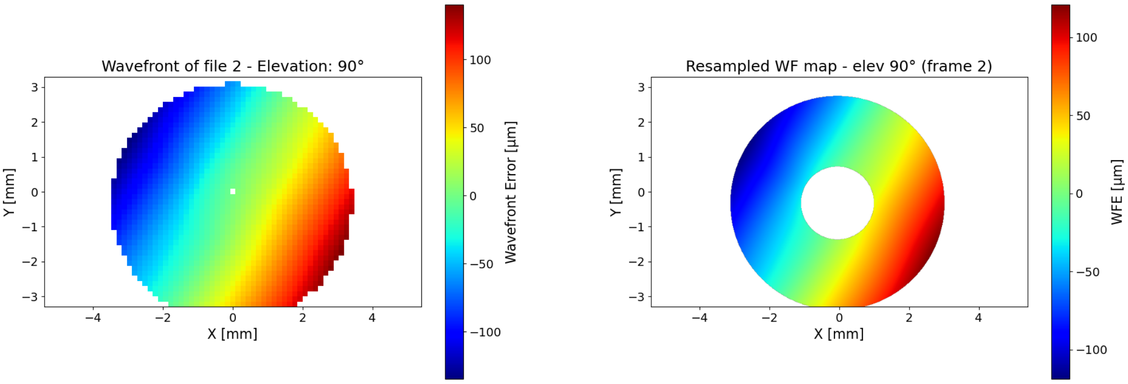

(a) Original WFE map **(b)** Resampled and masked WFE map

Figure 5.22: Comparison between the original and resampled WFE maps measured at the Coudé for a target at the zenith.

In the next step, the code calculates the best-fit Zernike polynomial for all frames corresponding to each target, both for the Coudé and Nasmyth focus. The outcome of this fitting procedure is presented in the following figures (Fig. 5.24 and Fig. 5.24). In each case, the plots on the left display the WFE maps reconstructed from the higher-order Zernike terms. This means that the piston, tip, and tilt components have been removed so that only the contribution of optical aberrations to the total wavefront error (WFE) is highlighted (the WF_{high} seen in 4.2.2).
The plots on the right instead display the residuals of the WFE; thus, these plots encompass all the features that the fitting process could not reproduce.

The results are acceptable since the residual WFE shows values between $-1.5\ \mu m$ and $1.5\ \mu m$ in the worst cases, while the best results show residuals between $-0.5\ \mu m$ and $0.5\ \mu m$.

Only one frame per target is reported here. The residuals are consistently on the order of microns, indicating that the fit is always functioning correctly. However, there

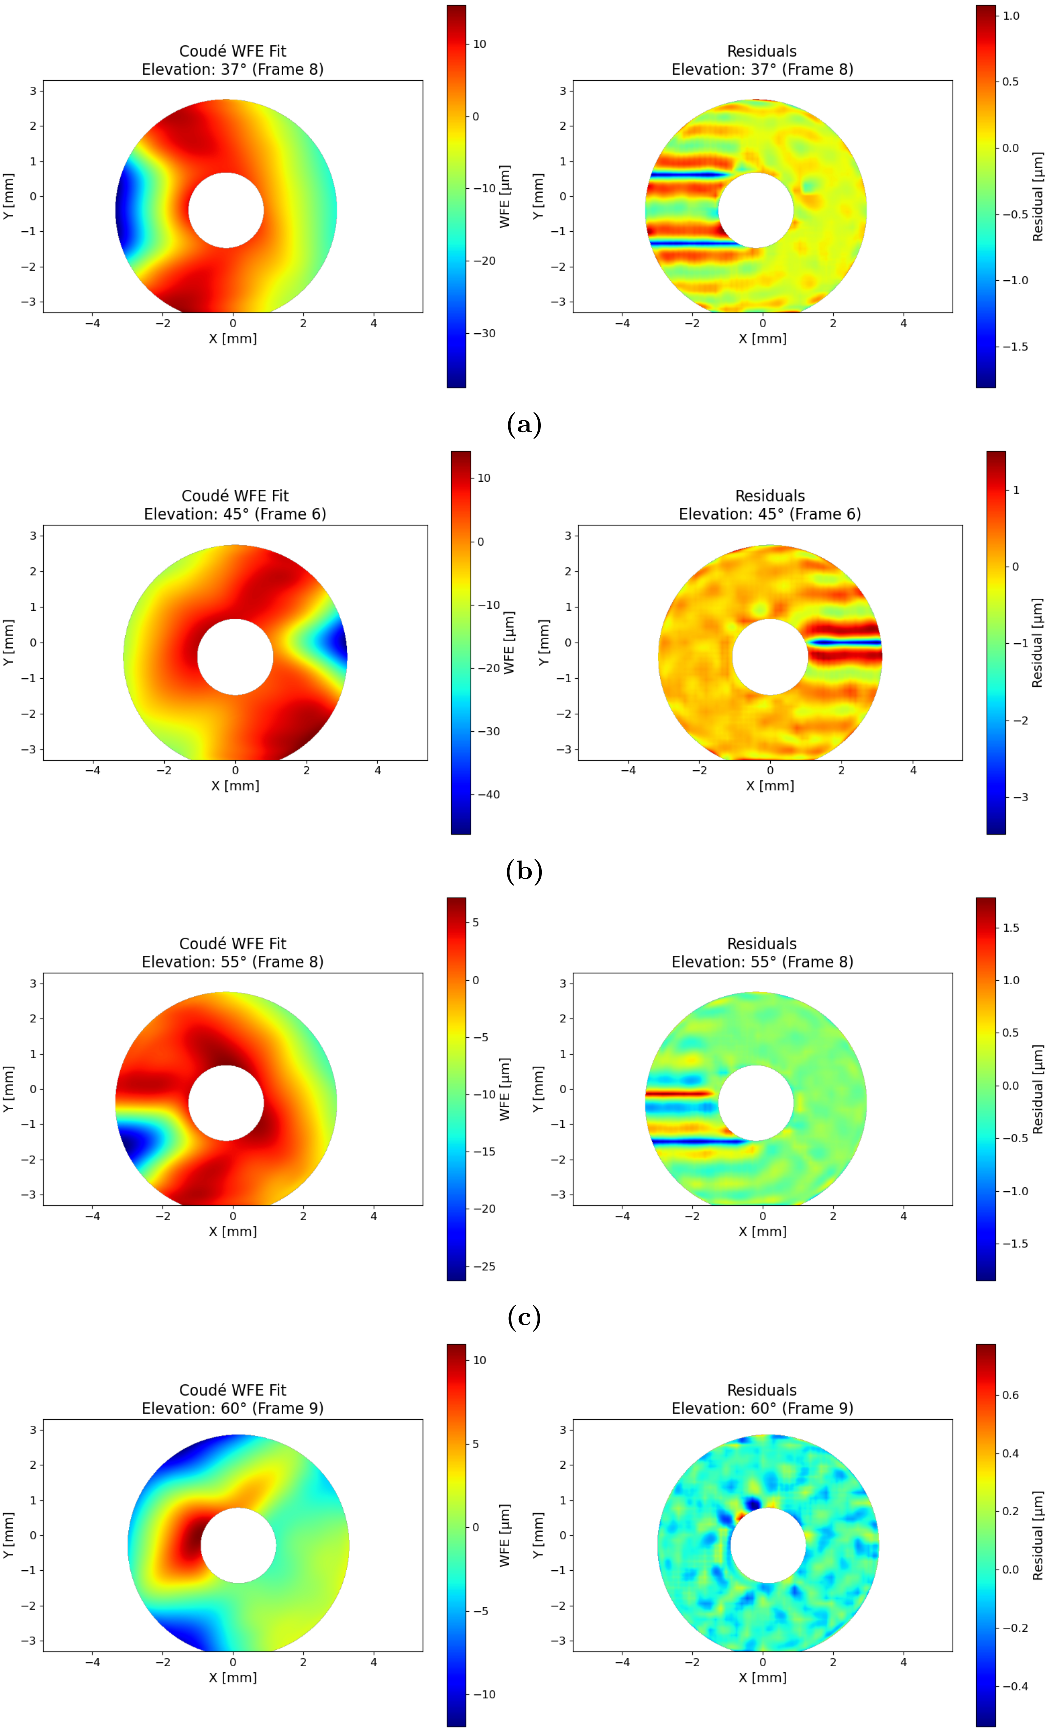
Coudé WFE Fit
Elevation: 37° (Frame 8)
Residuals
Elevation: 37° (Frame 8)
Coudé WFE Fit
Elevation: 45° (Frame 6)
Residuals
Elevation: 45° (Frame 6)
Coudé WFE Fit
Elevation: 55° (Frame 8)
Residuals
Elevation: 55° (Frame 8)
Coudé WFE Fit
Elevation: 60° (Frame 9)
Residuals
Elevation: 60° (Frame 9)
X [mm]
Y [mm]
WFE [µm]
Residual [µm]
−4
−2
0
2
4
−3
−2
−1
0
1
2
3
10
0
−10
−20
−30
−40
5
−5
−15
−25
1.5
1.0
0.5
0.0
−0.5
−1.0
−1.5
1
−1
−2
−3
0.6
0.4
0.2
−0.2
−0.4

(a)

(b)

(c)

(d)

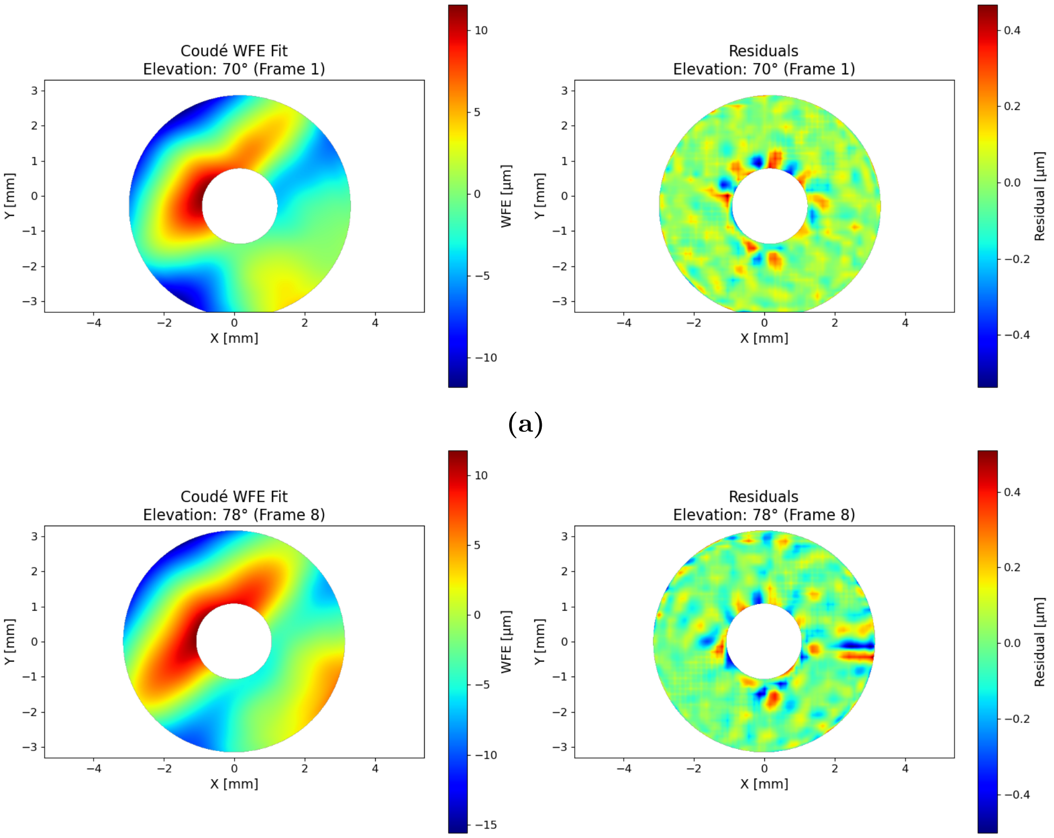

(a)

(b)

Figure 5.24: WFE maps and relative residuals map for the Coudé targets at the different elevation angles.

are discrepancies in the value of the residuals.

Fig. 5.23b shows residuals that are an order of magnitude larger than those in Fig. 5.24a. Those residuals also appear to be organised in filamentary structures rather than merely noise, as seen in the plots with smaller residuals.

These structures could be associated with the Low Wind EffectMilli et al. (2018) [14] observed in other telescopes.

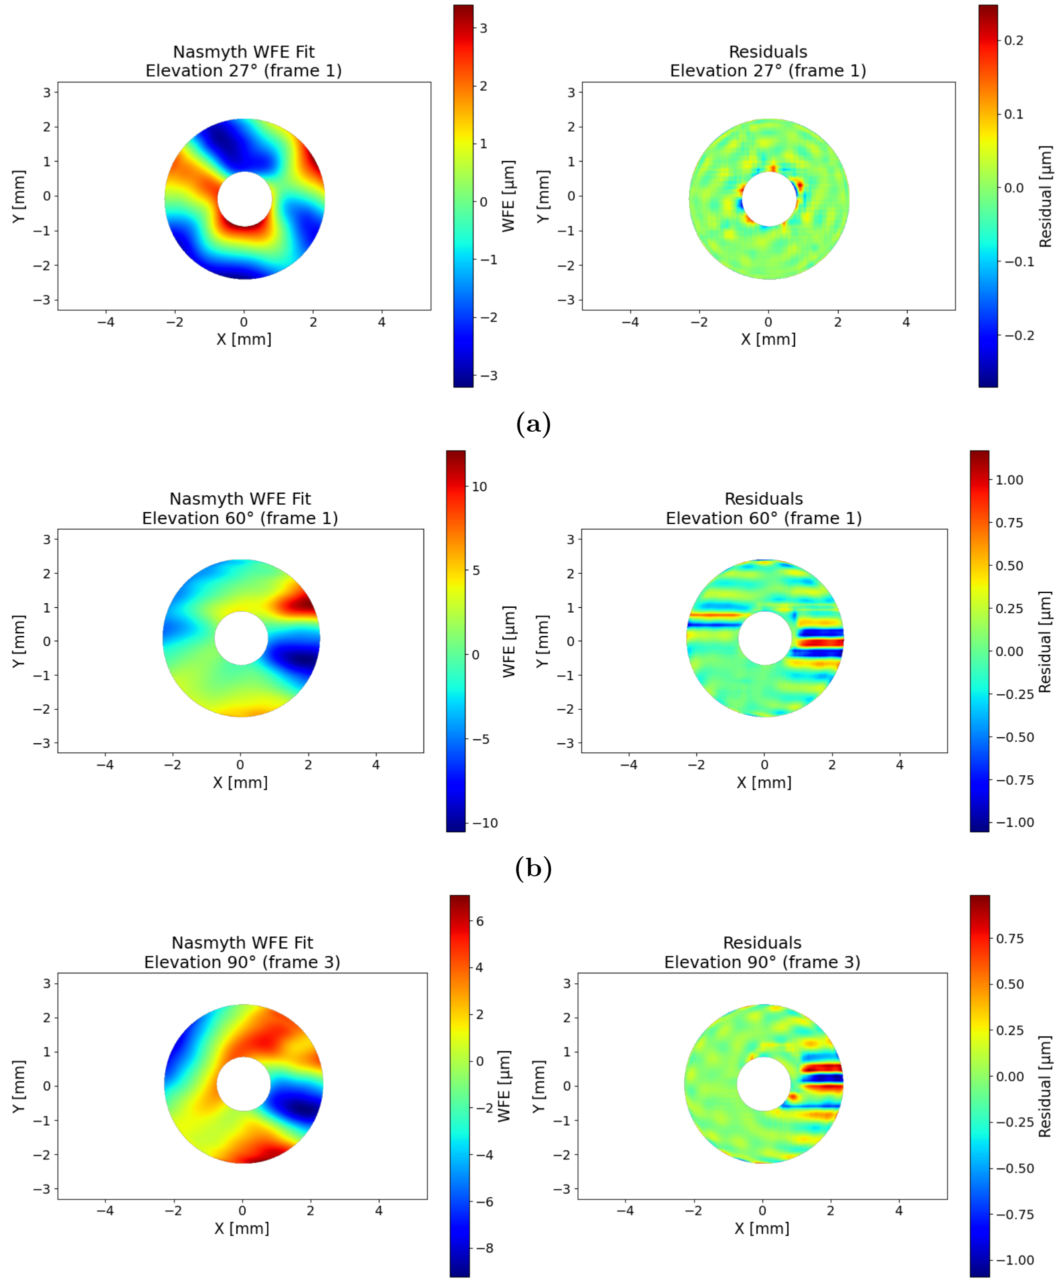

(a)

(b)

(c)

Figure 5.25: WFE maps and relative residuals map for the Nasmyth targets at the 3 elevation angles.

5.2.1 Low Wind Effect

As reported by Milli et al. (2018) [14], the Low–Wind Effect (LWE) is a pupil-scale optical phase disturbance that occurs under very low wind velocity conditions as a consequence of radiative cooling of the telescope spiders. When exposed to the outside environment, the spider loses heat by radiation, cools below the ambient air temperature, and generates thin layers of denser, colder air in their immediate surroundings. These layers create localized refractive index inhomogeneities that distort the incoming wavefront. Under low wind conditions, the cooled air pockets are not efficiently mixed away and remain trapped around the spiders, producing a persistent differential piston

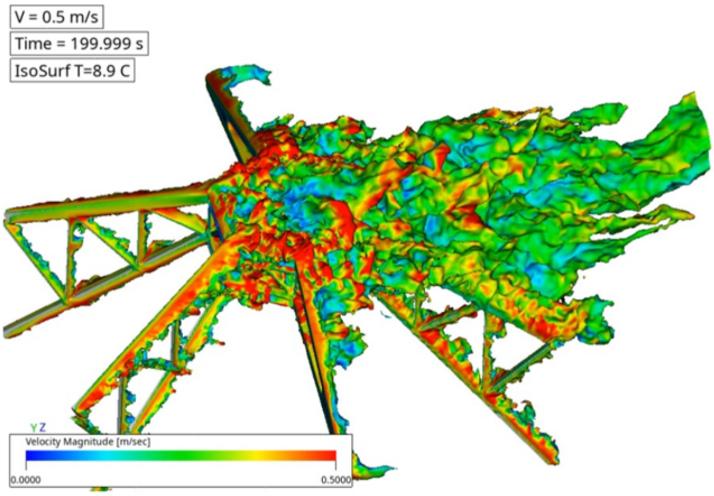

Figure 5.26: Simulation of the air flowing along the spiders in low wind conditions of ELT's M2 support structure. The motion and mixing of air at different temperatures produce turbulence, just like in the layers of the atmosphere. Source: Martins et al. (2022) [15]

and a small tilt between the pupil quadrants delimited by the arms of the spiders. The phenomenon depends on the thermal emissivity of the spiders, and the wind regime, becoming significant for wind speeds below 3 m/s and strong below 2 m/s (Martins et al. (2022) [15]). Numerical simulations show that the flow downstream of the spiders is characterised by thermally stratified wakes and eddies, which produce local optical path difference (OPD) fluctuations that evolve on timescales of a few seconds.

The resulting wavefront is characterised by discontinuities with sharp WFE steps across the spider shadows: each quadrant (if the telescope has 4 spider arms) behaves as an almost independent optical element with its own tip/tilt. Empirical measurements and simulations report phase jumps between adjacent quadrants of up to about 0.5 µm, with typical per-quadrant piston or tilt amplitudes in the range 0.1 µm–0.4 µm. The total phase RMS for moderate events spans roughly 0.02 µm to 0.5 µm. Computational simulations indicate that, at extremely low wind speeds (below 1 m/s), localized OPD "steps" can reach 1.0 µm or more, with peaks approaching 2.0 µm near the spider beams.

These numerical and experimental results are consistent with the residuals observed in the fitted wavefront maps, where the horizontal features present in some of the frames exhibit OPD variations of comparable magnitude, confirming that the observed structures can be attributed to the Low Wind Effect. The fact that the horizontal features in the wavefront maps appear consistently, independently of the telescope's elevation or the focal plane considered, indicates that their origin can be an external effect influencing the telescope.

Field de-rotation

After the fitting process, the numerical values of the aberrations (astigmatisms and comas) are de-rotated to account for the field rotation explained in 3.2. The code extracts the elevation angle of the target, computes the corresponding rotation angle and applies the correction following the formulas 3.1, 3.2, 3.3 and 3.4.

5.2.2 Aberrations as a function of the elevation angle

The WF fitting gives as a result an estimate of the main optical aberrations. For each target, the code produces one estimate per frame, which makes it possible to perform a statistical analysis of how these aberrations vary with the telescope elevation. The results of this analysis are presented in Fig. 5.27 and Fig. 5.28.

In these plots, each point corresponds to the mean value of the Zernike coefficient associated to a given aberration, while the error bars indicate the RMS variation of that coefficient across the considered frames.

At the Coudé focus (Fig. 5.27), results show that some aberrations, in particular defocus and astigmatisms, decrease in amplitude as the elevation increases. This behaviour is consistent with the presence of gravity-induced flexures and turbulence, which are expected to become more relevant at lower elevations. The exception is the coma, which appears to remain nearly constant across all elevations. The stability of the coma indicates that the telescope structure is mechanically stable and that no significant misalignment of the mirrors occurs during observations at different elevations. The Nasmyth focus instead (Fig. 5.28), shows a different behaviour: the main feature here is the increase of all the aberrations for the target at 60° elevation. It is possible that when the telescope is in this position, the flexures in the mirrors as well as in the telescope structure enhance the total aberration, while in other elevations they compensate for each other, resulting in less aberration content.

At the Nasmyth focus (Fig. 5.28), the behaviour of the aberrations is different. The main feature is the increase in the amplitude of all aberrations at an elevation of 60°. A possible explanation is that at this specific pointing angle, structural flexures of the telescope and deformations of the mirrors combine in such a way as to increase the total aberration. At other elevations, these effects may partially compensate for each other, resulting in smaller aberrations.

Another factor that may contribute to this behaviour is the reduced quality of the dataset at the Nasmyth focus. A significant number of frames had to be discarded due to the presence of NaN values (empty pixels within the pupil), which led data unusable. As a result, the analysis relied on a smaller sample with respect to the Coudé focus and more sensitive to statistical fluctuations. There is a possibility therefore that the increase in aberrations is not an actual physical effect, but a consequence of the limited dataset. An example of these discarded WF maps is shown in Fig. 5.29.

Fig. 5.30 shows this outcome. Coudé data behaved exactly as expected, with a decrease in the P-V as a function of the elevation angle. Nasmyth data instead shows a flatter distribution of the P-V, but it is still consistently lower than the Coudé ones. Again, the target at 60° elevation is the one with the worst behaviour, probably due to the same reasons explained before.

The pipeline also provides the P–V values of the high-order wavefront maps (WF_{high}) for each frame. This value is useful for quantifying the overall amplitude of the optical aberrations affecting the system. Larger P–V values are expected at lower elevations, where atmospheric turbulence and gravity-induced mechanical deformations are more significant. The Coudé focus is also expected to show consistently higher P–V values compared to Nasmyth, due to its longer optical path, which introduces more aberra-

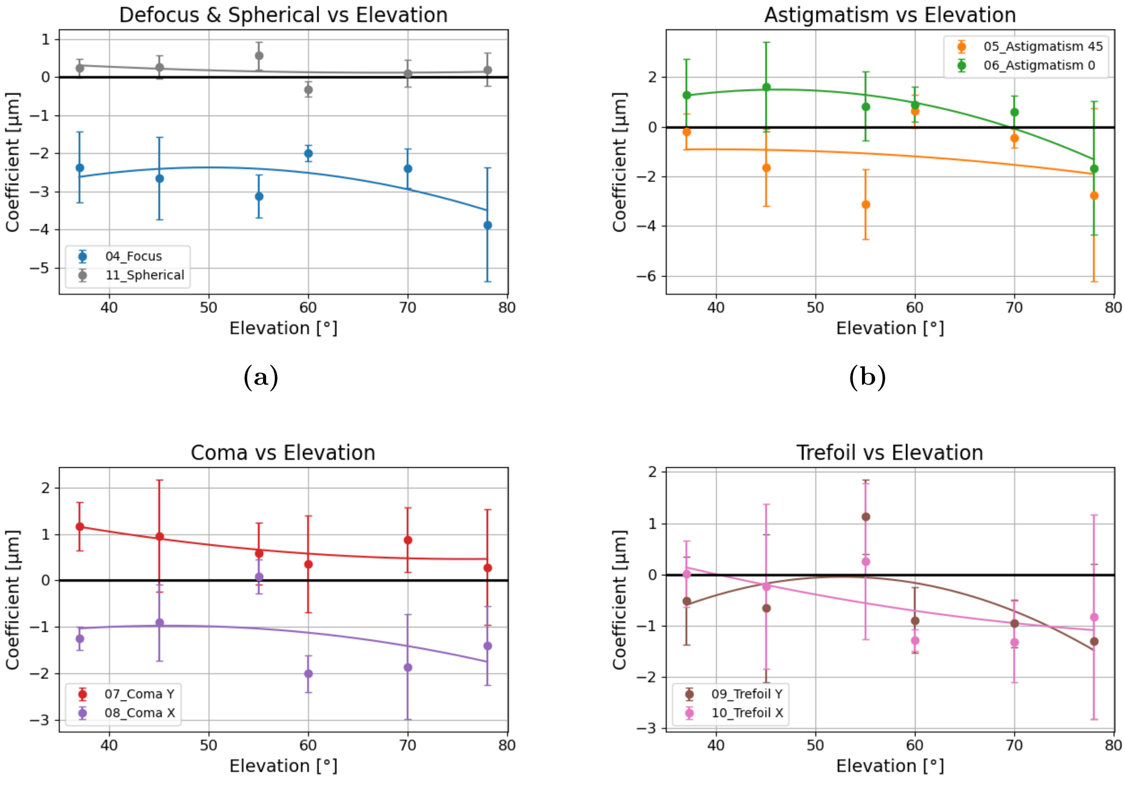

(a) **(b)** **(c)** **(d)**

Figure 5.27: Plots of aberration amplitude as a function of the elevation angle for the Coudé observations.

tions.

Figure 5.30 confirms these expectations. Coudé data exhibit a decrease of P–V with increasing elevation, confirming the predicted behaviour. The exception is for the data at the highest elevation, where the mean P-V is higher due to a larger spread in the values along the dataset. The Nasmyth data, on the other hand, show a flatter trend in P–V, but the values remain systematically lower than those at Coudé. Once again, the target at 60° elevation shows the largest deviation, likely for the same reasons discussed previously.

In Figure 5.31 the RMS values of the fitted wavefront maps with respect to the measured ones are shown. In this case, the trends for the two focal planes are similar to those observed in the P–V trends. The main difference is that, in this plot, the RMS of the residuals at the two focal planes is nearly identical, indicating that even though the Coudé focus exhibits stronger aberrations (i.e. larger P–V values in the wavefront maps), the residual RMS remains low, demonstrating the good quality of the fit. The larger scatter observed in the RMS residuals is associated with features that are more difficult to model in the wavefront error, such as residual aberrations and the effects of the LWE, which is highly dependent on local atmospheric conditions and therefore varies over time.

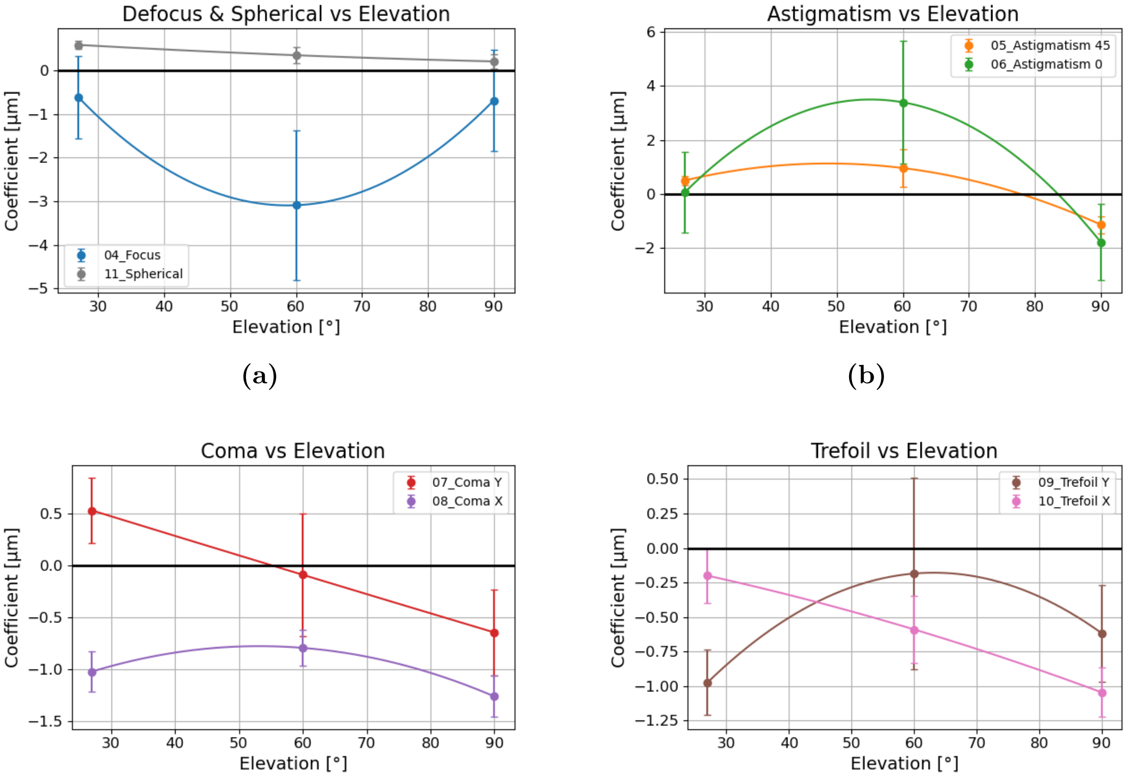

(a) (b)

(c) (d)

Figure 5.28: Plots of aberration amplitude as a function of the elevation angle for the Nasmyth observations.

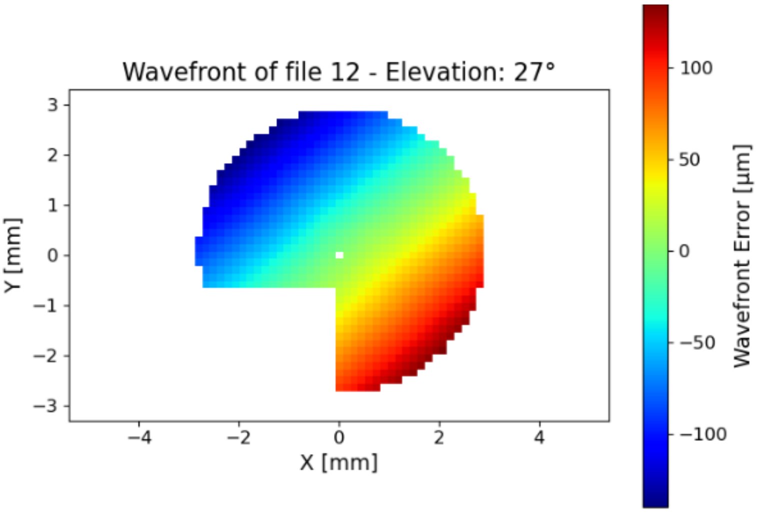

Figure 5.29: One of the pupils with some data acquisition error. The presence of the empty region makes the frame unusable.

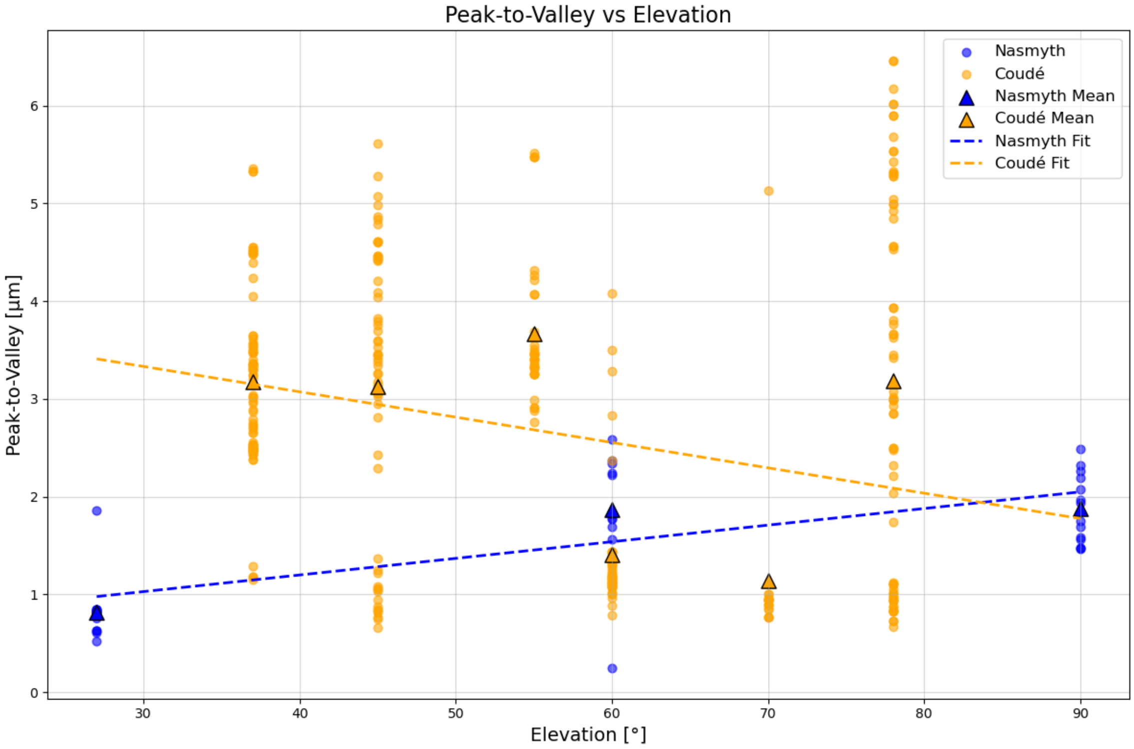

Figure 5.30: Comparison between the mean P-V values of the high-order WF maps at Coudé and Nasmyth. Yellow and blue triangles show the mean values of the P-V values relative to each target at Coudé and Nasmyth respectively. The circles instead represent the P-V of the single frames and the dashed lines show the trend in the P-V evolution with the elevation.

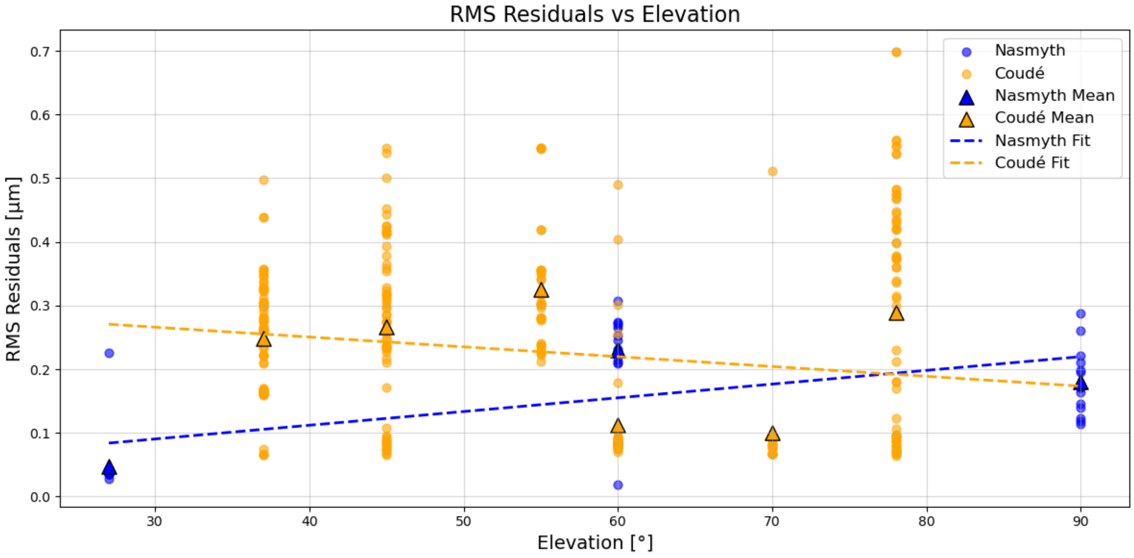

Figure 5.31: Comparison between the mean RMS values of the fitted WF maps residuals at Coudé vs Nasmyth. Yellow and blue triangles show the mean RMS relative to each target at Coudé and Nasmyth respectively. The circles represent the single frames and the dashed lines show the trend as a function of the elevation.

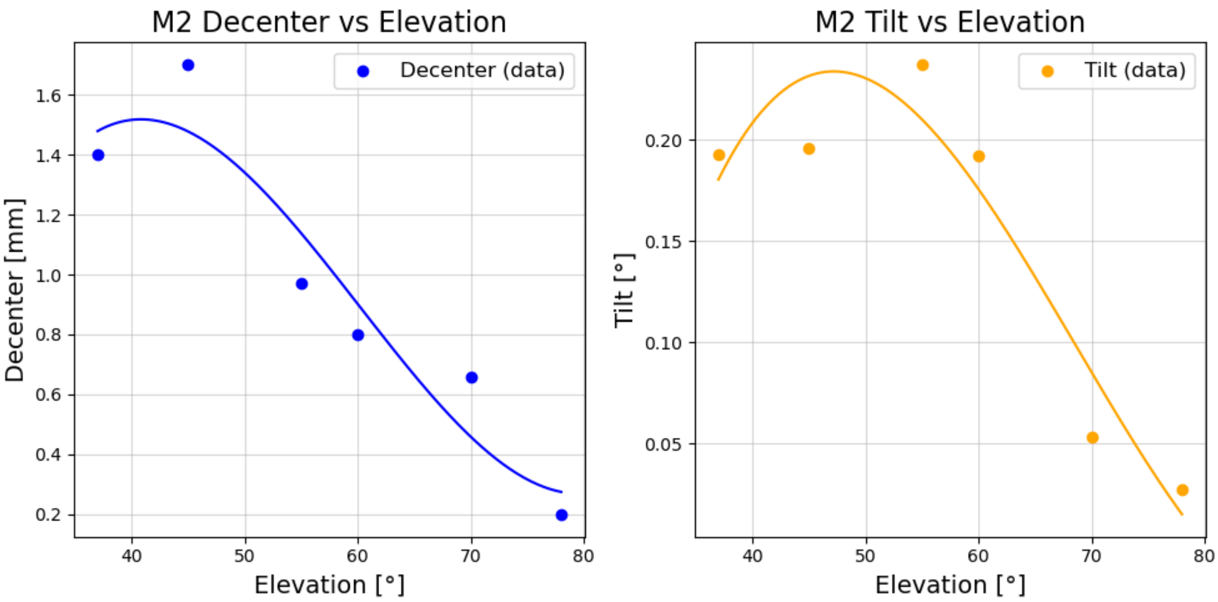

Figure 5.32: Estimated M2 decenter and tilt as a function of elevation, derived from Zemax simulations reproducing the observed aberrations.

5.2.3 Telescope-induced aberrations

To obtain a better physical interpretation of the measured aberrations, I investigated their possible instrumental origin by performing simulations in Zemax using the telescope optical design. The aim was to reproduce the measured aberrations (in particular astigmatism) by introducing plausible misalignments in the telescope structure. I focused on M2 since its location on the structure makes it particularly susceptible to tilt and decenter when the telescope points at low elevations. In the Zemax model I inserted the measured aberration values in the merit function and set the tilt and decenter of M2 as free variables. By minimising the merit function the optimizer found the displacements of M2 that reproduce the observed aberration pattern, effectively simulating how structural misalignments at different pointings translate into optical errors.

The results are shown in Fig. 5.32. As expected, the M2 decenter and tilt increase at the lower elevations, where the structure is subject to the largest mechanical strain. The maximum decenter is of the order of 1.5 mm and the maximum tilt about 0.3°, while the minimum values across the elevation range are approximately 0.2 mm and 0.03°, respectively.

Again, it appears that there is a peak in the tilt at intermediate elevations, as verified by the WFS data. It is possible that, at those particular angles, some misalignments and optical deformations combine in such a way that the net effect is a peak in the aberrations and a minimum in the optical quality of the system.

These simulations are “first order”: they do not include deformations of M1 or the entire structural simulation of the telescope, which would require a finite element analysis. Despite the model’s simplifications, the Zemax results show that M2 misalignment and tilt are a likely dominant cause of the observed astigmatism.

5.3 Donut Results

In this section, the results produced by the donut code are presented. As explained in Section 3.3, the donuts obtained at the Nasmyth focal plane could not be used: their excessive defocus, combined with the presence of the telescope's spider in the images, made them unsuitable for donut fitting. As a consequence, only the Coudé donuts were analyzed. After running the pipeline described in Section 4.3 on the Coudé images, the code returned several plots showing the preprocessed image, the best-fit donut model, and a plot comparing the fitted aberrations with those measured by the WFS.
The plots 5.33, 5.34 and 5.35, show the fitting results for each target observed at the Coudé focus.

In each plot, the top left panel displays the image used as input for the fitting procedure, while the top-right panel shows the best-fit donut produced by the code.

The bottom panel presents a comparison between the observed and fitted quantities. The Y axis contains the different aberrations. The red markers represent the fitted amplitudes of the aberrations (Defocus, Astigmatism, and Coma), along with their associated errors. The blue markers indicate the mean amplitudes of the same aberrations as measured by the WFS (as shown in Section 5.2.2), also shown with their corresponding error bars. On the right side of the bottom panel, the **discrepancy** values (reported in Δ μm) show the difference between the fitted and observed aberration amplitudes. The values are compared in absolute terms because the presence of the beam splitter in the optical system (Fig. 3.6) causes the donut images to be mirrored with respect to the pupil entering the WFS. As a result, the signs of the aberrations are expected to change, and the comparison is therefore made using their absolute values. The RMS $|\Delta|$ shows instead the root mean square of the deltas between measured and fitted values, giving an overall image of the consistency of the method.

The quality of the output is quite variable and strongly depends on the quality of the input image. Figure 5.33b provides an example: in this case, the donut is characterised by irregular edges, and under these conditions the code encounters significant difficulties in finding a good fitting solution. This is also evident in the fitted model, where the streaks visible in the background are produced by the algorithm while attempting to reproduce the features of the original image. As a consequence, the fitted values are expected to be unreliable.

The donut shown in Figure 5.34a illustrates a different situation. Here, fewer sharp structures are present in the original image, allowing the code to generate a model that better represents the donut. In this case, the fitted values of astigmatism fall within the error bars of the reference WFS measurements. The defocus, however, is not well reproduced and is underestimated with respect to the WFS. This occurs when the central region of the donut is either too bright (and so not clearly defined), or not sufficiently round. A non-circular central region results from telescope tracking issues during image acquisition, where image drift during exposure causes misalignment between consecutive frames, leading to an elliptical donut core. This deformation has a negative impact on the defocus estimate.

Figure 5.33b is another example of this behavior: in that case the central region is well defined and no drift is present, which explains why the reconstructed defocus matches the measured value, even though the other aberrations are not accurately reproduced. The same situation occurs in Figures 5.35a and 5.35b, where the donut cores are round (even if not particularly dark as for 5.35b) and the resulting defocus values agree with the WFS measurements.

The donut in Fig. 5.34b shows the best agreement with the measurements: in this case, all aberrations except defocus fall within the measurement uncertainties. The defocus is again poorly fitted, likely due to the central region of the donut not being clearly defined and almost indistinguishable.

In general, the code is able to reproduce low-order aberrations with good accuracy when the input donuts have the right characteristics: a clean shape, a well-defined central region and no strong artifacts. When the donuts are strongly aberrated, affected by drift, or have a not well-defined central region, the code struggles. In these cases, the code tries to reproduce image features that are not related to wavefront aberrations, resulting in inaccurate values.

A second point concerns the reference used for the comparison. In this work, the WFS measurements were treated as the “true” values. However, as discussed in Section 5.2.2, the WFS measurements also show some limitations. Sources of uncertainty such as errors during frame acquisition and variations in illumination conditions introduce fluctuations in the reconstructed wavefronts, leading to highly variable error bars from one target to another. As a consequence, part of the discrepancy between the donut fitted and the WFS measured aberrations may originate not only from the fitting procedure itself but also from the reference data provided by the WFS. In addition, the large error bars observed in some WFS datasets are partially explained by the manner in which the measurements were obtained. The frames were acquired manually using the sensor software, which did not support automated data capture, requiring each frame to be saved individually, impeding the association of one frame to a specific donut image. Moreover, each stored WFS frame corresponded to the average of several short-exposure images (as described in Section 3.2). This implies that the available data were already averaged, making it possible that some of the measured aberrations were actually caused by atmospheric distortions (low-order, slowly-changing aberrations) rather than only by the telescope misalignment. Altogether, these combined effects contributed to the large uncertainties associated with the WFS measurements.

A more reliable validation would require dedicated laboratory tests. In a controlled setup, known aberrations could be introduced into the system, and donut images could be acquired with a priori defined aberration amplitudes. These images would serve as an absolute reference, allowing a precise evaluation of the code’s accuracy and its limitations. Moreover, for the algorithm to converge reliably to the true aberration values, it requires input donuts lower amount of aberrations than those obtained in our measurements. Excessively aberrated donuts are characterised by many difficult to reproduce features and introduce degeneracies in the fitting procedure, making it more difficult for the code to find the correct solution. Laboratory-generated donuts with controlled, moderate aberrations would therefore provide a much more suitable dataset for validating and calibrating the code.

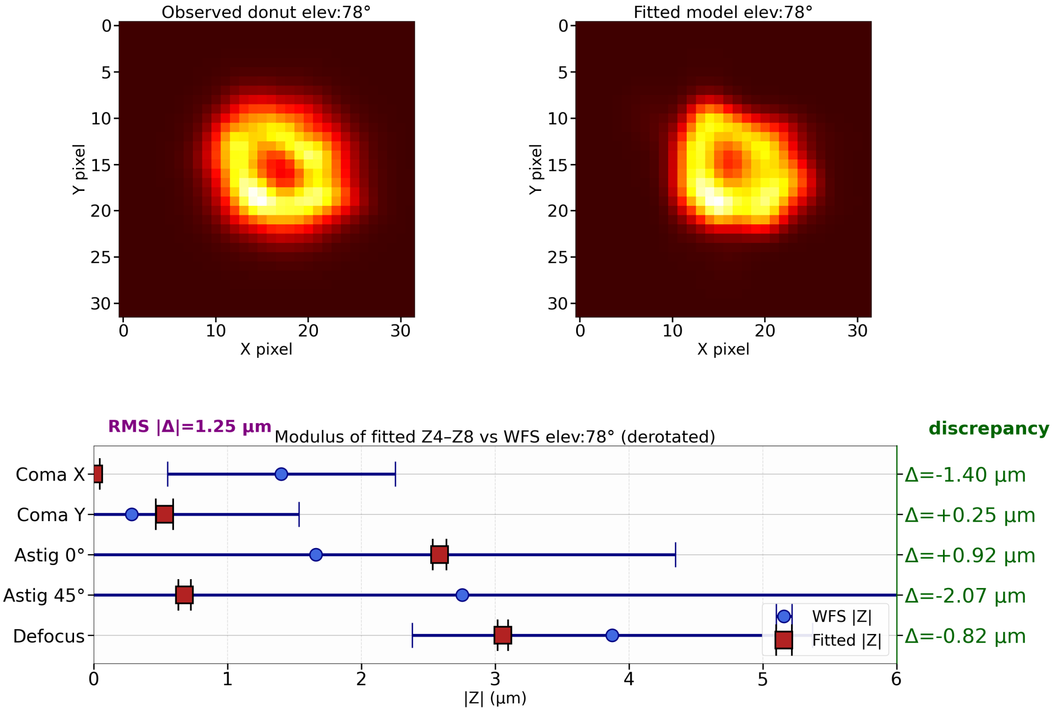

(a) Donut fit results at elevation 78°

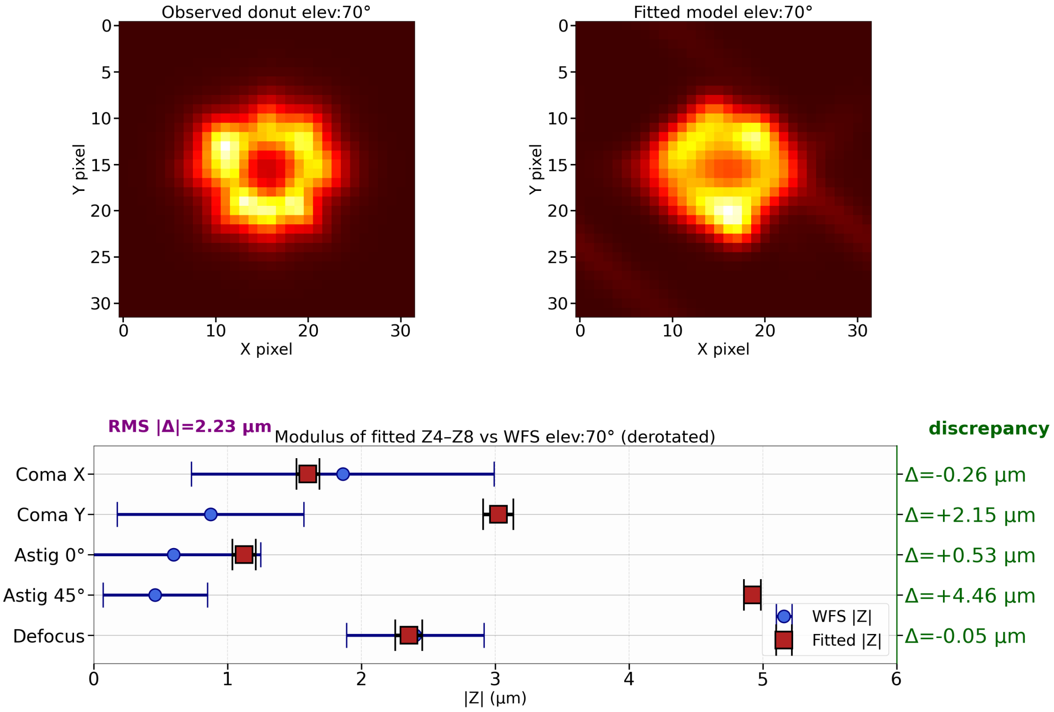

(b) Donut fit results at elevation 70°

Figure 5.33: Fitted Coudé donuts. Notice how figure 5.33b, despite its irregular edges gives a well-reproduced defocus. That happens because the central region of this donut is particularly well defined and in contrast with respect to the remaining structure of the donut.

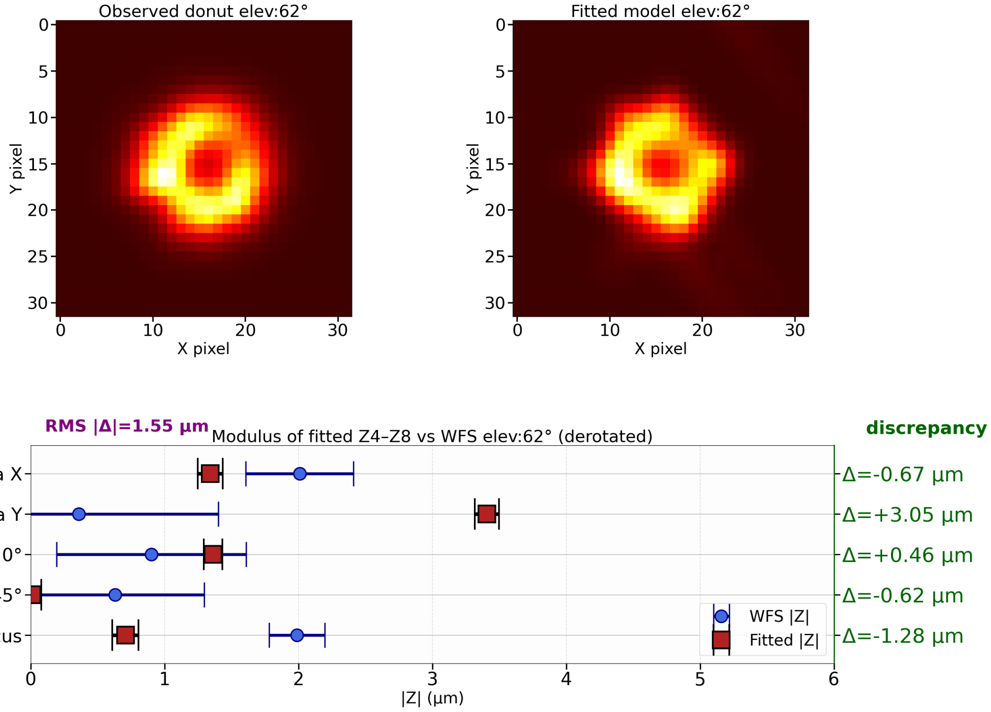

(a) Donut fit results at elevation 62°

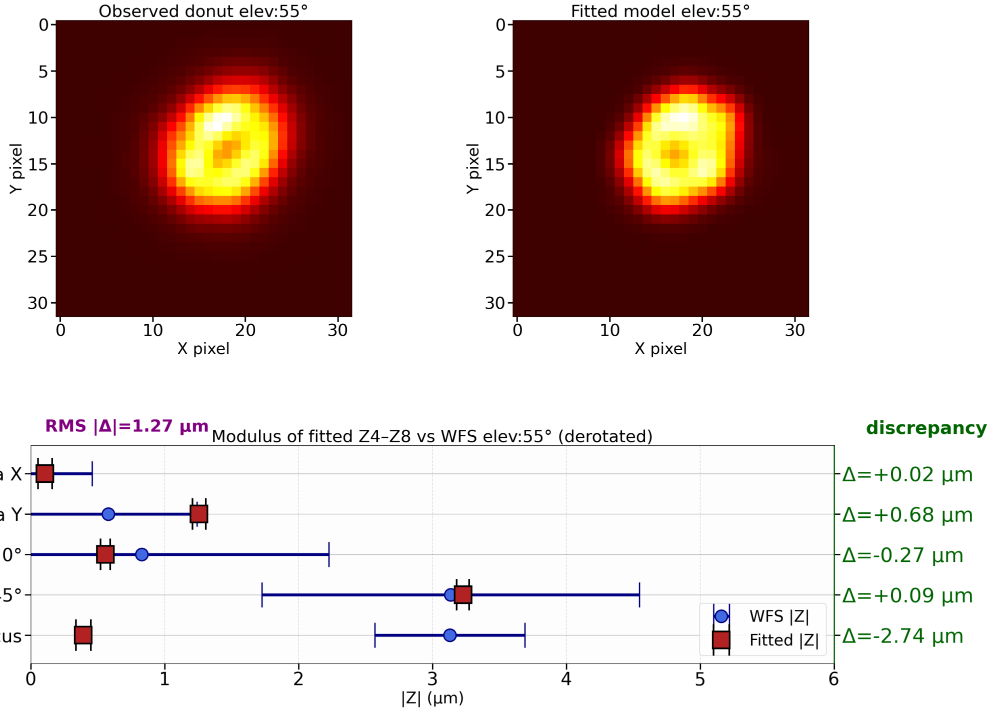

(b) Donut fit results at elevation 55°

Figure 5.34: Fitted Coudé donuts. Figure 5.34b exhibits fitted aberrations well in accordance with the WFS, with the exception of the defocus. This happens because the donut has a regular structure but a not well defined core, which results in a erroneous value of the fitted defocus.

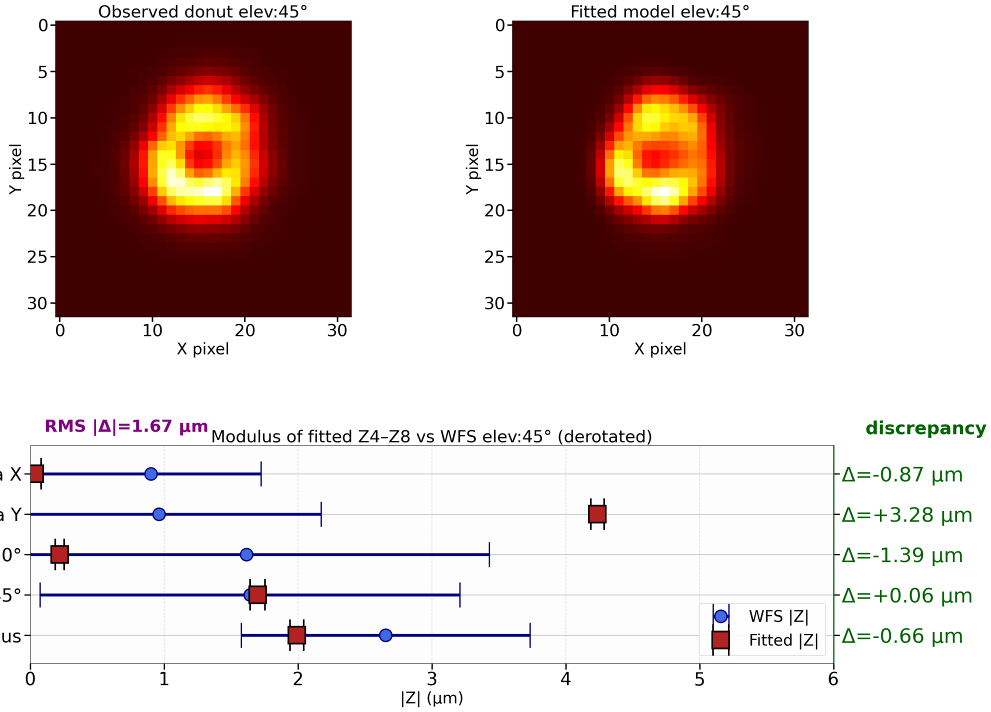

(a) Donut fit results at elevation 45°

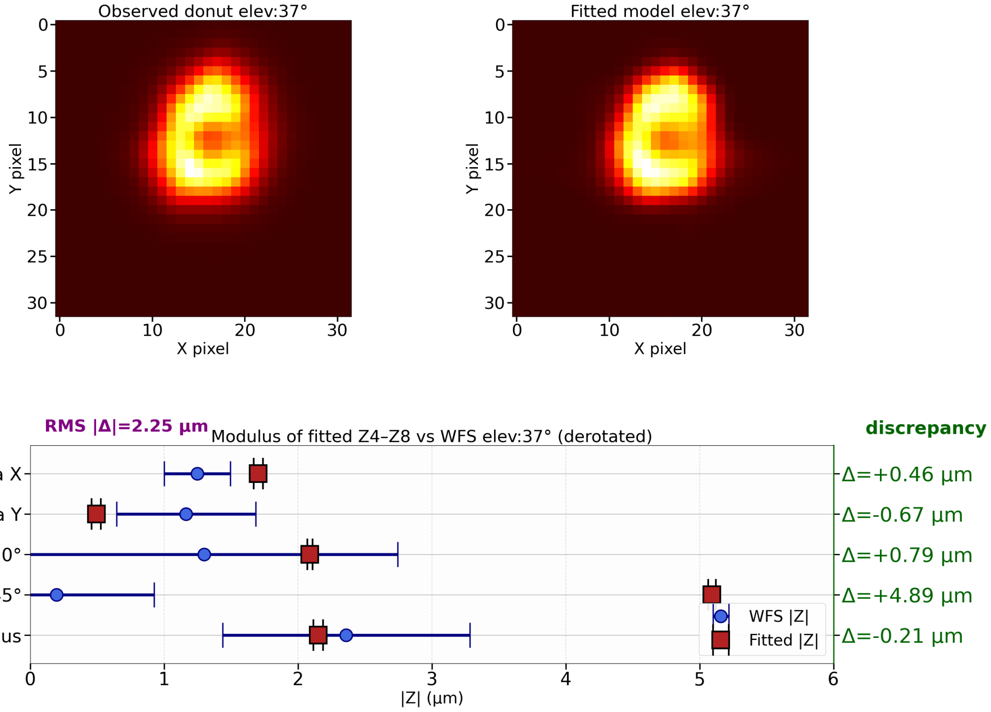

(b) Donut fit results at elevation 37°

Figure 5.35: Fitted Coudé donuts. here 5.35b shows values of coma and defocus

Chapter 6

Conclusions and future developments

The aim of this work was to characterise the aberrations of the Copernico telescope, study their evolution with elevation using speckle images and WFS data and use the WFS as a reference to assess the performance of the donut method for retrieving low-order aberrations directly from the scientific camera. This approach is motivated by the need to measure non-common path aberrations without requiring a dedicated WFS. The donut technique will be used in the future during the periodic realignment of MORFEO. Since the system contains many moving parts and is subject to temperature variations, small misalignments are expected to occur over time. As a consequence, regular low-order realignment will be necessary, and the donut method offers a way to measure aberrations along the entire scientific optical path, not just the WFS path, providing more reliable information and improving the effectiveness of the alignment procedure.

All the data were acquired using setups and sensors prepared and calibrated by us, which provided full flexibility in the choice of targets and observing strategy. This also introduced some issues with the acquired data (e.g., the blurred Coudé speckles in Fig. 5.18 and the unusable Nasmyth donuts); however, overall, the data acquisition was successful.

The characterization of the telescope aberrations from speckle images worked as expected. Speckles exhibited the predicted evolution with elevation, as their size decreased with increasing elevation angle and shorter exposure times. The ellipticity analysis showed a less intuitive behavior (Fig. 5.21), but the same trend appeared in the WFS results and can be explained by the increase in aberration amplitudes at the Coudé focus. In addition, the generally rounder speckles observed at the Coudé focus can be attributed also to the recent aluminization of the primary mirror and the subsequent re-alignment of the telescope, which improved the general optical quality prior to the Coudé observations (but after the Nasmyth ones), which took place in September 2025, a few weeks after the aluminization.

The analysis of the WFS data allowed for a more detailed study of the low-order aberrations. Wavefront maps were successfully reconstructed and fitted (Fig. 5.25). An unexpected result emerged from the residual maps: the presence of horizontal structures that proved to be consistent with the Low Wind Effect, whose characteristics match those reported by Milli et al. (2018) [14] and Martins et al. (2022) [15]. Beyond this, the WFS analysis showed that the Coudé focal plane is more aberrated than the Nasmyth (Fig. 5.27 and 5.28) as expected due to the longer optical path and the additional optics. Moreover, the aberrations at the Nasmyth display a noticeable increase around 60° elevation, likely caused by a combination of mirror misalignments and structural deformations. This behaviour is also reflected in the evolution of speckle ellipticity.

To complete the analysis, a Zemax simulation was carried out to evaluate how the telescope optics respond when the measured aberrations are inserted into the system (Fig. 5.32). The simulation shows that at low elevations, the decenter and tilt of M2 increase due to stronger mechanical strain on the telescope structure. This is consistent with the observed behaviour of the aberrations and supports the interpretation that mechanical misalignments of the optics are a fundamental source of aberrations. The simulation is only a partial representation, as a full analysis would require a FEM model capable of reproducing the structural deformation of the telescope, which Zemax cannot perform. Nevertheless, the overall trend of the misalignments is consistent with what is expected for a structure of this type.

Lastly, the donut code produced results with variability that depends on the quality of the input donuts, as discussed in Section 5.3.

Under favorable conditions, the code successfully reproduces the low-order aberrations with values consistent with the WFS measurements and within their error bars. "Favorable conditions" can be defined as donuts that are regular in shape, not excessively aberrated, and not affected by strong drift or pronounced irregular features. Such conditions were not met for the majority of the donuts obtained during the observations, many of which showed significant distortions and irregular features. These characteristics complicate the reconstruction process and complicate the calibration of the code, due to the numerous uncertainties introduced from multiple sources (the WFS itself and the nature of the on-sky data).

Even when the donuts are strongly aberrated or affected by misalignments, the retrieved coefficients remain within physically plausible ranges for a telescope of this type. In many cases, the discrepancies between the fitted and reference values stay below one micron in modulus, indicating that the algorithm is converging toward the correct solution rather than diverging or producing unphysical results.

A further positive aspect is the computational performance: the code generates model donuts in a matter of seconds, matching the requirements of MORFEO, where alignments must be performed quickly and repeatedly; thus, computational speed is fundamental for this code.

These results show that, although the code in its current state still requires refinement, the donut technique is intrinsically robust. Part of the remaining discrepancies can be attributed to the limitations of the WFS data used for validation: the frames were acquired manually, each saved individually using the sensor software, and the stored images were themselves averages of several short-exposure frames. As a consequence, the resulting wavefronts likely include not only telescope misalignment but also residual low-order atmospheric aberrations, introducing additional uncertainties in the reference data. Future developments should therefore include dedicated laboratory tests, where known aberrations can be imposed and the corresponding donuts acquired in a controlled environment. Such a setup would provide an absolute reference, eliminate the limitations of the on-sky WFS measurements, and allow a precise calibration and optimization of the reconstruction code.

In conclusion, despite the current limitations of the code, the donut technique has proven to be a reliable method for retrieving low-order aberrations directly from scientific images. With further testing and improvements to the code, its performance can only increase, eventually providing MORFEO with a reliable tool for routine alignment during its operations.

Acknowledgements

I would like to thank the people who supported me during this thesis: without their help this work would not have been possible.

First, I want to thank the staff of the Cima Ekar Observatory: L. Tomasella, P. Ochner, A. Frigo, L. Lessio, L. Traverso and M. Mosele, for their help in organizing the observing nights and for providing technical support whenever needed.

I would also like to thank the researchers at INAF Padova: D. Greggio, A. Ballone and E. Carolo as well as the PhD students from Padova O. Rebrysh, R. Van Gaalen, L. Cortese and the PhD students from Bologna A. Taranto and L. Rosignoli. Their support during the observing nights at the telescope, in setting up the optical equipment and in acquiring the data was essential and made the whole process much easier.

Finally, I want to express my deep gratitude to my supervisors, Gabriele Umbriaco and Gabriele Rodeghiero, who have supported me from the very beginning of this project with patience and by teaching me many, many things. I really enjoyed the moments we shared (even the challenging ones) both at the telescope and in the office. What they have taught me will be essential for my future career, and they have been very important to me, both as mentors and as people.

Figure 6.1: The Cima Ekar Observatory in Asiago.

Bibliography

.

1. Tokovinin, A. and S. Heathcote (2006). “Donut: measuring optical aberrations from a single extra-focal image”. In: *Publications of the Astronomical Society of the Pacific* 118.846, pp. 1165–1175. DOI: `10.1086/506972`. URL: `http://arxiv.org/abs/astro-ph/0606388`.

2. Rodeghiero, G. et al. (2022). “The MAORY/MORFEO fine optical alignment and recollimation strategies: preliminary simulations from ‘out of focus’ PSF images”. In: vol. 12185, 121854Z. DOI: `10.1117/12.2627503`. URL: `https://ui.adsabs.harvard.edu/abs/2022SPIE12185E..4ZR`.

3. Leishman, J. Gordon (2023). “Turbulent Flows”. In: URL: `https://eaglepubs.erau.edu/introductiontoaerospaceflightvehicles/chapter/turbulent-flows/`.

4. McKechnie, T. Stewart (2022). “Atmospheric Isoplanatic Angle”. In: *General Theory of Light Propagation and Imaging Through the Atmosphere*. Ed. by T. Stewart McKechnie. Springer International Publishing, pp. 551–587. DOI: `10.1007/978-3-030-98828-9_17`. URL: `https://doi.org/10.1007/978-3-030-98828-9_17`.

5. Busoni, Lorenzo et al. (2023). “MORFEO enters final design phase”. In: 17 pages. DOI: `10.13009/AO4ELT7-2023-046`. URL: `http://arxiv.org/abs/2310.09005`.

6. Fiorentino, G. et al. (2019). *MAORY science cases white book*. DOI: `10.48550/arXiv.1712.04222`. URL: `http://arxiv.org/abs/1712.04222`.

7. ESO (n.d.). *ESO - ELT Science*. URL: `https://www.eso.org/sci/facilities/eelt/science/`.

8. www.thorlabs.de (n.d.). *Shack-Hartmann Wavefront Sensors*. URL: `https://www.thorlabs.de`.

9. Noll, Robert J. (1976). “Zernike polynomials and atmospheric turbulence*”. In: *Journal of the Optical Society of America* 66.3, p. 207. DOI: `10.1364/JOSA.66.000207`. URL: `https://opg.optica.org/abstract.cfm?URI=josa-66-3-207`.

10. Hampson, Karen M. et al. (2021). “Adaptive optics for high-resolution imaging”. In: *Nature Reviews. Methods Primers* 1, p. 68. DOI: `10.1038/s43586-021-00066-7`.

11. excelitas.com (n.d.). *pco.edge 4.2 USB sCMOS Camera — Excelitas*. URL: `https://www.excelitas.com/product/pcoedge-42-usb-scmos-camera`.

12. Donut-Code (n.d.). *Leonardo-5/Donut-Zernike*. URL: `https://github.com/Leonardo-5/Donut-Zernike`.

13. Collins, Karen A. et al. (2017). “AstroImageJ: Image Processing and Photometric Extraction for Ultra-Precise Astronomical Light Curves (Expanded Edition)”. In: *The Astronomical Journal* 153.2, p. 77. DOI: `10.3847/1538-3881/153/2/77`. URL: `http://arxiv.org/abs/1701.04817`.

14. Milli, Julien et al. (2018). “Low Wind Effect on VLT/SPHERE : impact, mitigation strategy, and results”. In: *Adaptive Optics Systems VI*, p. 83. DOI: `10.1117/12.2311499`. URL: `http://arxiv.org/abs/1806.05370`.

15. Martins, Diogo et al. (2022). “Transient wavefront error from cooled air downwind of telescope spiders”. In: *Modeling, Systems Engineering, and Project Management for Astronomy X*. Vol. 12187. SPIE, pp. 492–505. DOI: `10.1117/12.2627683`. URL: `https://www.spiedigitallibrary.org/conference-proceedings-of-spie/12187/1218712/Transient-wavefront-error-from-cooled-air-downwind-of-telescope-spiders/10.1117/12.2627683.full`.