

Measuring and calibrating
Non-Common Path Aberrations in
Adaptive Optics assisted
image-slicer based spectrographs



Álvaro Menduiña Fernández
Mansfield College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2021

Acknowledgements

“In truth, O judges, while I wish to be adorned with every virtue, yet there is nothing which I can esteem more highly than the being and appearing grateful. For this one virtue is not only the greatest, but is also the parent of all the other virtues.”

– Marcus Tullius Cicero, *Pro Plancio*

When I was at school, if someone had told me (*Thunder and lightning. Enter three Witches*)-style that I was going to get a DPhil from the University of Oxford I would have been beyond excited. Now that I look back, I’m *grateful* they forgot to mention the part about the Pandemic, as this would have spoiled the surprise.

Anyway, in what is almost the anniversary of Tiger King (yes, I now only communicate in lockdown euphemisms), I am sitting here writing a few paragraphs to express my sincere gratitude to those who have made this possible, and by **this** I mean my DPhil, not Tiger King.

First and foremost, I want to thank my family (Mamá, Papá, Clara and everyone else) because I owe them everything. They have taught me the value of effort and made me the person that I am today. In these unprecedented circumstances, we have been apart more than what I would have wished. I would like to dedicate all this hard work to Antonio and Ildara, who I am sure would be proud of what I have achieved.

I am very grateful to Catherine, whose love has been a wonderful support over these years, and has made my working days so much happier. Thank you for all the special moments: from a simple coffee break, a trip to the Rad Cam, or a well-deserved Gino’s carbonara at the end of a long day. Thank you for listening to my incessant complaints, for encouraging me to keep going, for making me laugh when I was feeling down, and for making this journey so much easier for me.

I am also grateful to everyone at Oxford who has made me feel welcomed and supported me over the last three and a half years. Thank you Vanessa for helping me so diligently and making my day to day so much better. Thank you John for always being there to help and answer my questions. Thank you Andrea for being such a good friend and helping me feel more positive. And of course, thank you, Matthias and Niranjana for offering me this fantastic opportunity, and for your guidance, supervision and mentorship; it has been a privilege working with you.

In addition, I would like to extend my thanks to Sam and Mike for introducing me to machine learning, sharing their skills and knowledge with me and answering my naive questions.

This would not be a fair Acknowledgements if I did not take a moment to thank the one who always came to my rescue in times of need. Coffee, because I honestly doubt I could have done this without you.

Lastly, I would like to thank Bailey for reminding me every day that as long as there is food in your bowl (or on your plate) and a place to take a nap, the rest is secondary.

Álvaro Menduina Fernández

March, 2021 [1 A.C.]

Abstract

HARMONI is the first-light, visible and near-IR integral field spectrograph for ESO's Extremely Large Telescope (ELT), which promises to revolutionise observational astrophysics with its high spatial resolution, offering quasi-diffraction-limited performance from a ground-based observatory. The performance of Adaptive Optics assisted instruments such as HARMONI is hindered by the presence of differential aberrations between the AO wavefront sensing and science optical paths, known as Non-Common Path Aberrations (NCPA). These aberrations, which primarily affect the spatial point spread function (PSF) in HARMONI, cannot be sensed directly due to the absence of phase information in focal plane intensity measurements, so special techniques are needed to calibrate and pre-compensate them.

This is particularly problematic in high-contrast observations aimed at detecting and characterising exoplanets, as NCPA give rise to quasi-static speckles that resemble planetary signals and severely limit the contrast.

Although extensive research has been devoted to developing NCPA calibration techniques for AO systems and coronagraphic instruments, these have not been demonstrated in the context of integral field spectrographs like HARMONI that contain an image slicer. This optical component, made out of a stack of thin mirrors, slices a 2D field of view into slitlets and re-arranges them into a long pseudo-slit that can be dispersed by a spectrograph. The effects of an image slicer on light propagation could have a negative impact on the behaviour of state-of-the-art techniques for NCPA calibration, and require modifications to account for these effects to ensure they could be applied to instruments such as HARMONI.

This DPhil thesis makes several novel contributions to this field of research. First of all, I have investigated the impact of an image slicer on light propagation and characterised its effects on the PSF following two

approaches: a theoretical one, using a mathematical framework cross-validated with simulations based on optical models of the HARMONI sub-systems; and an experimental one, for which I have designed and conducted an experiment with the engineering model of the SWIFT instrument image slicer.

Secondly, given that image slicers can affect the relationship between phase aberrations and focal plane intensity, and the unproven behaviour of traditional techniques, I decided to explore alternative approaches that could circumvent this drawback. I have developed a novel technique for NCPA calibration that relies on machine learning methods to estimate the aberrations directly from PSF images. I have demonstrated that this technique can discriminate between features from image slicers and features caused by aberrations in the system. I have thoroughly characterised its response to multiple sources of uncertainty and shown how it can be successfully applied to HARMONI.

Lastly, I have developed the capability to create and analyse End-to-End Monte Carlo models of HARMONI that aggregate the different sub-systems to form a realistic representation of how the as-built instrument would behave. This has been used to support our calibration efforts in multiple ways: to characterise the expected level of NCPA in the system and its major contributors, to define requirements for NCPA calibration accuracy based on the intrinsic limitation of field-dependent aberrations, and to train the machine learning calibration models with realistic wavefront maps that reproduce the typical NCPA that we will encounter in HARMONI.

All of these efforts shed light on the question of how NCPA can be calibrated in instruments that contain image slicers, and will pave the way for its future application in HARMONI. The only missing piece of the puzzle is the thorough characterisation of how image slicers impact the behaviour of traditional methods like Phase Diversity and how they could be adapted to account for such effects. Due to time constraints, this could not be fully explored and will be the subject of future research.

Contents

1	Introduction	1
1.1	HARMONI Instrument	1
1.2	Specifications and design review	2
1.2.1	Pre-Optics (IPO) sub-system	3
1.2.2	Integral Field Unit (IFU) sub-system	4
1.2.3	Spectrograph (ISP) sub-system	5
1.2.4	Science Detectors sub-system (ISDS)	7
2	Non Common Path Aberrations	8
2.1	Origins of NCPA	8
2.1.1	Phase vs. Amplitude aberrations	10
2.2	Why care about NCPA?	12
2.3	Modelling NCPA	14
2.4	Research Goals	15
2.4.1	Problem definition	15
2.4.2	Research question	16
2.4.3	This thesis' contribution	17
2.5	Judging Success	19
2.6	Deriving requirements	21
2.6.1	Requirements from Ensquared Energy	21
2.6.2	Requirements from Field Dependent Aberrations	31
2.7	State-of-the-art techniques	32
2.7.1	Focal Plane Sharpening	32
2.7.2	Phase Diversity	33
2.7.3	ZELDA	37
2.7.4	Electric Field Conjugation	37
2.7.5	Self Coherent Camera	38
2.7.6	Coronagraphic Low Order Wavefront Sensor	41

2.7.7	Lyot-based Low Order Wavefront Sensor	41
2.7.8	Differential Optical Transfer Function	42
2.7.9	Discussion	42
3	Image Slicers	44
3.1	Integral field spectroscopy	44
3.2	Calibrating NCPA in HARMONI	45
3.3	Light propagation through Image Slicers	48
3.3.1	Mathematical framework	48
3.3.2	Model validation	53
3.3.3	Impact of coarse sampling on slicer effects	59
3.3.4	Effect of aberrations	62
3.3.5	Final remarks	65
4	Experimental Work - SWIFT	67
4.1	Motivation	67
4.2	Experiment design	68
4.3	Data Reduction and Results	74
4.3.1	Alignment validation	74
4.3.2	Preliminary results	79
4.3.3	Simulations Comparison	84
4.3.4	Stitching the slices	88
4.4	Limitations and future work	91
4.5	Conclusions	92
5	Machine Learning for NCPA calibration	93
5.1	Background and Motivation	93
5.2	Methodology	96
5.2.1	Neural Networks	97
5.2.2	Convolutional Neural Networks	102
5.2.3	Training data	104
5.3	Proof of concept - Image Slicer Effects	106
5.4	Results	109
5.4.1	Architecture	111
5.4.1.1	Readout noise. Hyper-parameter optimisation.	113
5.4.1.2	Field of View	118
5.4.2	Sampling effects. Calibrating coarse scales	120

5.4.2.1	Influence of sampling	120
5.4.2.2	Calibrating at 20 mas	123
5.4.3	Alternative wavefront definitions	127
5.4.3.1	Zernike polynomials vs. Actuator commands	128
5.4.3.2	Pupil mapping errors	133
5.4.3.3	Zernike misalignment	137
5.4.4	Robust calibration - Impact of real effects	141
5.4.4.1	Flat field calibration uncertainties	141
5.4.4.2	Pixel sampling errors	144
5.4.4.3	Anamorphic magnification errors	150
5.4.4.4	Aggregating all effects - Robust calibration model	158
5.4.5	Multiwavelength calibration	163
5.4.5.1	Motivation	163
5.4.5.2	Data format	163
5.4.5.3	Results	164
5.4.6	Interpreting the predictions	168
5.4.6.1	Shapley values	168
5.4.6.2	Impact of noise	175
5.4.6.3	Visualising the layers	178
5.4.7	Impact of the diversity phase	180
5.4.7.1	Defocus strength	180
5.4.7.2	Arbitrary diversities	182
5.5	Going beyond the training data	183
5.5.1	The generalisation problem	183
5.5.2	The challenge of training with telescope data	184
5.6	Advantages and disadvantages	193
6	End-to-End Model	197
6.1	Introduction	197
6.2	Deriving calibration requirements	198
6.2.1	Expected NCPA	203
6.2.2	Intrinsic limitation: Field-Dependent Aberrations	205
6.2.3	Feeding back the correction	212
6.2.4	Chromatic effects in field-dependent aberrations.	215
6.3	Realistic NCPA calibration	218

7	Conclusions and Future Work	223
7.1	Future Work	223
7.1.1	Impact of image slicer effects on Phase Diversity	223
7.1.2	Extending the End-to-End Model work	224
7.2	Conclusions	224
	Bibliography	228

List of Figures

1.1	HARMONI Integral Field Spectrograph (IFS) block definition diagram	3
1.2	HARMONI Field Splitter Unit	4
1.3	HARMONI Image Slicer Module	5
1.4	HARMONI Spectrograph (ISP) sub-system	6
2.1	Origin of Non Common Path Aberrations	9
2.2	Phase vs. Amplitude aberrations	11
2.3	Discovery images of the exoplanet GJ 504 b	13
2.4	Discovery images of the exoplanet 51 Eri b	13
2.5	Venn diagram of the major components of this DPhil thesis	19
2.6	Ensquared Energy (EE) calculation methodology	23
2.7	EE as a function of RMS WFE	25
2.8	Impact of cosine aberrations on the PSF	26
2.9	EE as a function of RMS WFE for the 60×30 mas scale	27
2.10	EE as a function of RMS WFE for the 60×30 mas scale, rescaling	29
2.11	Evolution of EE with RMS wavefront error, all scales in HARMONI.	30
2.12	Principle of the self-coherent camera (SCC)	39
3.1	Impact of aberrations on Strehl ratio (astigmatism and trefoil)	46
3.2	Impact of aberrations on Strehl ratio (astigmatism and coma)	47
3.3	Optical model for the image slicer within the SWIFT instrument	49
3.4	Sketch of an Integral Field Unit sub-system	49
3.5	Approximation of the effect of the pupil mirror aperture	52
3.6	Impact of pupil mirror aperture on the exit slit PSF fringes	53
3.7	Impact of wavelength and pupil mirror aperture	54
3.8	Zemax POP vs Python comparison. 2 PSF zeros, central slice	55
3.9	Zemax POP vs Python comparison. 3 PSF zeros, central slice	55
3.10	Zemax POP vs Python comparison. 2 PSF zeros, adjacent slice	57
3.11	Zemax POP vs Python comparison. 3 PSF zeros, adjacent slice	57

3.12	Zemax POP vs Python comparison. 2 PSF zeros, complete PSF . . .	58
3.13	Zemax POP vs Python comparison. 3 PSF zeros, complete PSF . . .	58
3.14	Image slicer effects, PSF cross-section	60
3.15	Image slicer effects, PSF cross-section, binned	61
3.16	Impact of defocus, nominal vs defocused PSF	63
3.17	Impact of the image slicer, defocused PSF (0.20λ PV)	64
3.18	Impact of the image slicer, defocused PSF (0.50λ PV)	64
3.19	Image slicer effects, PSF cross-section, binned, defocus 0.20λ PV . . .	66
4.1	Schematic diagram of the SWIFT experiment and its sub-systems . .	69
4.2	Pre-optics sub-system	70
4.3	SWIFT Image Slicer and back-optics sub-systems	71
4.4	SWIFT Image Slicer pupil test.	75
4.5	Pre-Optics PSF Test	76
4.6	Pre-Optics PSF Test, Airy pattern fit	76
4.7	Pre-Optics Magnification Test	77
4.8	Complete System Magnification Test	78
4.9	Calibrated detector frame	80
4.10	Impact of back-optics sub-system focus distance	82
4.11	Pre-optics fibre y scan	83
4.12	Experimental PSF slices at the detector	85
4.13	Extension of Fig. 4.12.	86
4.14	Zemax POP simulations of the SWIFT experiment	87
4.15	Stitching the slices to expand the field of view	89
4.16	Stitched sections of the SWIFT PSF	90
5.1	Network diagram for a two-layer neural network	98
5.2	A classic convolutional neural network architecture	103
5.3	Output of the first two convolutional layers of a calibration model . .	104
5.4	Training example of a neural network model with image slicer effects	107
5.5	Comparison between predictions and ground truth of aberrations . .	108
5.6	Impact of model architecture on the calibration performance	112
5.7	ROC curve cumulative fraction	116
5.8	Strehl curve cumulative fraction	117
5.9	Impact of field of view	119
5.10	Impact of PSF sampling	121
5.11	Calibrating coarse spaxel scales. Distribution of Strehl ratios	122

5.12	Calibrating coarse scales. The 20×20 mas case)	124
5.13	Alternative wavefront definitions	128
5.14	Zernike vs Actuator wavefront comparison	129
5.15	Zernike vs Actuator calibration comparison	131
5.16	Zernike vs Actuator results comparison	132
5.17	Zernike vs Actuator density distribution	133
5.18	Impact of actuator pupil mapping errors	135
5.19	Impact of actuator pupil mapping error on Strehl	136
5.20	Impact of actuator pupil mapping error on Strehl, damped correction	137
5.21	Effects of reference Zernike system misalignment	138
5.22	Effects of misalignment angle on residual RMS wavefront error	139
5.23	Evolution of Strehl ratio, Zernike misalignment angle	140
5.24	Impact of flat field uncertainties on residual coefficients	142
5.25	Impact of flat field uncertainties on Strehl ratio	143
5.26	Impact of Nyquist sampling error on residual RMS	145
5.27	Residual RMS vs Nyquist sampling error	146
5.28	Data augmentation for robust training	148
5.29	Impact of anamorphic magnification errors on the PSF	151
5.30	Impact of anamorphic magnification errors on the PSF, averaged	151
5.31	Zernike coefficients resembling ellipticity variations	153
5.32	Zernike polynomials resembling ellipticity variations	153
5.33	Impact of anamorphic magnification errors vs Zernike polynomials	154
5.34	Impact of anamorphic magnification errors on residual Zernike coefficients, ellipticity $a/b = 1.10$	155
5.35	Impact of anamorphic magnification errors on residual Zernike coefficients, ellipticity $a/b = 0.90$	155
5.36	Robust model, cumulative distribution of Strehl ratio	159
5.37	Robust model, impact of uncertainties on final Strehl	160
5.38	Robust model, cumulative distribution of Strehl, readout noise	161
5.39	Robust model, impact of uncertainties on the final Strehl, readout noise	162
5.40	Robust model, testing outside the training range	162
5.41	Example datacube for the multiwavelength calibration experiment	164
5.42	Example of Shapley values	169
5.43	Shapley maps for an actuator coefficient	171
5.44	Shapley values vs the differential PSF features for actuator #1	173
5.45	Shapley values vs the differential PSF features for actuator #4	173

5.46	Relative position of the actuators on the pupil plane	174
5.47	Example of Shapley values in the presence of readout noise	176
5.48	Comparison of Shapley values with and without readout noise	177
5.49	Output of the two convolutional layers of the calibration model	179
5.50	Changes in convolutional layer outputs after poking an actuator	179
5.51	Impact of defocus diversity strength on residual RMS after calibration	181
5.52	Expected NCPA from defocus and astigmatism	186
5.53	Comparison between test coefficients and predictions	189
5.54	Evolution of wavefront error throughout the calibration	190
5.55	Simulating the training at the telescope	192
6.1	Residual RMS wavefront vs. number of Zernike polynomials corrected	199
6.2	RMS WFE contribution of each Zernike polynomial, E2E Monte Carlo	202
6.3	Evolution of residual RMS WFE with number of Zernike polynomials corrected, E2E Monte Carlo simulation	202
6.4	Comparison between NOAO and SCAO modes, RMS wavefront contribution for each Zernike polynomial, E2E Monte Carlo simulation	204
6.5	Comparison between NOAO and SCAO modes, histograms of RMS wavefront error, E2E Monte Carlo simulation	204
6.6	Field-dependent aberrations. Variation in aberration coefficient (Z_4 <i>defocus</i> and Z_6 <i>astigmatism</i>) across slices and IFU paths	206
6.7	Field-dependent aberrations. Variation in aberration coefficient (Z_8 <i>coma</i> and Z_{10} <i>trefoil</i>) across slices and IFU paths	207
6.8	Average range of field-dependent aberrations	209
6.9	Range of correctable RMS NPCA as a function of Zernike polynomials	210
6.10	Residual error vs initial NCPA and number of Zernike polynomials	211
6.11	Impact of field-dependent aberrations on the achievable accuracy	213
6.12	Distribution of field-dependent aberrations	214
6.13	Chromatic field-dependent aberrations. Variation in aberration coefficient (Z_4 <i>defocus</i> and Z_6 <i>astigmatism</i>) across slices and IFU paths	216
6.14	Chromatic field-dependent aberrations. Variation in aberration coefficient (Z_8 <i>coma</i> and Z_{10} <i>trefoil</i>) across slices and IFU paths	217
6.15	Methodology for the E2E HARMONI Monte Carlo wavefront analysis.	218
6.16	Performance of ML model trained on E2E Monte Carlo wavefronts	220
6.17	Impact of readout noise contamination SNR	221

1

Introduction

Come then, and let us pass a leisure hour in story-telling, and our story shall be the education of our heroes.

– Plato, *The Republic*

1.1 HARMONI Instrument

HARMONI is a visible and near-infrared (0.47 to 2.45 μm) integral-field spectrograph (IFS) that will provide the core spectroscopic capability for ESO's Extremely Large Telescope (ELT), offering a range of resolving powers $R = \lambda/\Delta\lambda$ from ~ 3000 to $\sim 17,000$. It will provide simultaneous spectra of $\sim 31,000$ spaxels, arranged in 4:3 aspect ratio contiguous field. HARMONI is designed to become the ELT's workhorse instrument and is expected to revolutionise visible and near-infrared astronomy, capitalising on the large collecting area of the ELT and providing unprecedented levels of spatial resolution and sensitivity.

With regards to **spatial** information, HARMONI offers four spaxel scales of 4×4 , 10×10 , 20×20 and 60×30 milli-arcseconds per spaxel. With around $152 \times 204 = 31008$ spaxels available, this provides an instantaneous field of view covering $0.61''\times 0.82''$, $1.52''\times 2.04''$, $3.04''\times 4.08''$, $6.12''\times 9.12''$ depending on the scale.

In terms of **spectral** characteristics, HARMONI implements a total of 11 spectral bands covering the visible and near-infrared range (0.47 - 2.45 μm) with varying resolving power R , see Table 1.1.

HARMONI is designed to provide a variety of Adaptive Optics (AO) modes, including: Single Conjugate Adaptive Optics (SCAO), Laser Tomography Adaptive Optics (LTAO), no AO at all (NOAO), as well as a High Contrast mode (HCAO) which incorporates SCAO and an apodiser.

Spectral Band	Wavelength Range μm	Resolving Power
'V+R+I'	0.47 - 0.80	≥ 3000
'I+z+J'	0.83 - 1.35	≥ 3000
'H+K'	1.45 - 2.40	≥ 3000
'I+z'	0.83 - 1.05	≥ 7000
'J'	1.05 - 1.32	≥ 7000
'H'	1.45 - 1.80	≥ 7000
'K'	1.97 - 2.40	≥ 7000
'z'	0.828 - 0.902	≥ 17000
'H high'	1.538 - 1.678	≥ 17000
'K short'	2.017 - 2.201	≥ 17000
'K long'	2.199 - 2.400	≥ 17000

Table 1.1: HARMONI Spectral bands and resolving power $R = \lambda/\Delta\lambda$. Data from: HRM-00190 Integral Field Spectrograph Requirements Specification.

When using AO, the diffraction pattern from the telescope is Nyquist-sampled in the spatial direction at the 4×4 mas spaxel scale of HARMONI at around $1.45 \mu\text{m}$. For the remaining scales (10×10 , 20×20 and 60×30 mas) the spatial direction is not Nyquist-sampled.

1.2 Specifications and design review

It is beyond the scope of this thesis to provide a detailed account of whole instrument design and its sub-systems. Nonetheless, for the purpose of understanding this thesis (the calibration of Non-Common Path Aberrations, including their origin but also the available avenues for correction), it is important to have some awareness of the characteristics of the HARMONI instrument and how it operates.

Thus, in this section, I will provide a brief overview of the design of the different sub-systems that are relevant for this thesis. HARMONI is mounted on the Nasmyth platform of ESO's ELT and it is made out of several sub-systems:

- **LGSS: Laser Guide Star Sensors.**
- **NGSS: Natural Guide Star Sensors**, including SCAO, LTAO-NGS and the High Contrast Module (HCM).
- **CARS: Calibration and Relay System**, divided into Calibration Unit, Focal Plane Relay System (FPRS), and LGS Dichroic.

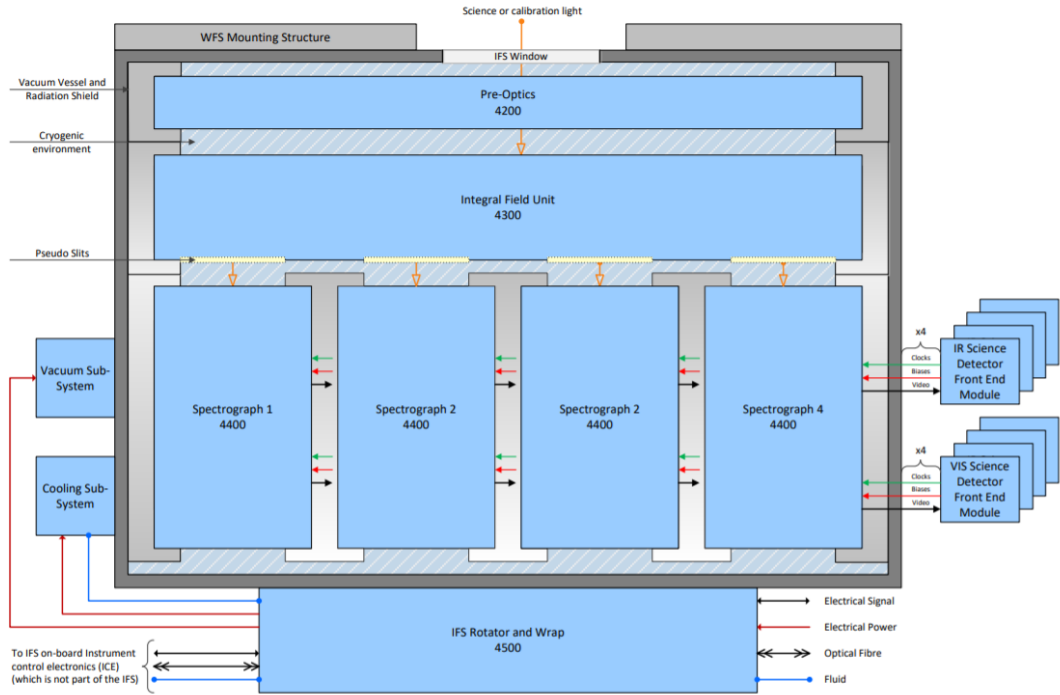


Figure 1.1: HARMONI Integral Field Spectrograph (IFS) block definition diagram showing the different sub-systems and how they interact with each other.

- **IFS: Integral Field Spectrograph**, the heart of HARMONI, divided into Pre-Optics (IPO), Integral Field Unit (IFU), Spectrograph (ISP), and Science Detector modules; all housed inside a cryostat.

1.2.1 Pre-Optics (IPO) sub-system

The role of the FPRS is to relay the light coming from the telescope into the IFS. In order to minimise thermal background, the bulk of the instrument is housed within a cryostat. Starting with the first IFS sub-system (see Fig. 1.1 for reference), the role of the pre-optics (IPO) is to relay the input focal plane and provide four exchangeable spaxel scales. The rectangular field of view coming from the IPO is then reformatted by the IFU sub-system to form four separate slits. Each slit is then fed to one of four identical spectrographs, which collimate the light, disperse it using diffraction gratings and focus the light onto the science detectors. All spectrographs contain a camera designed to operate in the near-infrared regime (0.8 - 2.45 μm), and two spectrographs will also carry a visible camera.




Figure 1.2: Field Splitter Unit that divides the input field into 8 sub-fields arranged in pairs (AB, CD, EF, GH). Each sub-field pair is then sent to one of the 4 Image Slicer Modules (ISM) to be sliced and re-arranged into 2 pseudo-slits.

1.2.2 Integral Field Unit (IFU) sub-system

The Integral Field Unit (IFU) sub-system is responsible for dividing the field and reshaping it to provide 4 output slits arranged in a square at the input to the 4 spectrographs. The input is a rectangular field of view delivered to the IFU by the Pre-Optics (IPO) sub-system. The input field will be divided into eight sub-fields by the Splitting and Relay Module (SRM), which contains the Field Splitter Unit (see Fig. 1.2). After magnification, pairs of sub-fields are delivered to 4 Image Slicer Modules (ISM). Each ISM will slice the field into 2×38 (76) slitlets, as well as demagnifying by $\times 0.13$. The slitlets are arranged along an output slit in a staggered pattern of two rows, ready to be fed to the spectrographs. As shown in Fig. 1.3, the ISM is divided into two assemblies (A and B), each containing one image slicer (38 slices) and a set of pupil and slit mirrors.

The slices of the Image Slicer Modules are 1 mm wide and 50.8 mm long. Due to the anamorphic magnification of the Pre-Optics (IPO), the PSF at the image slicer is elliptical with a 2:1 ratio, elongated in the *across-slice* direction. This is only the case for the fine scales (4×4 , 10×10 and 20×20); the PSF is circular at the slicer for the 60×30 scale, hence the different spaxel sizes in x and y . The reason behind the anamorphic magnification is to ensure the spaxels on sky are square instead of rectangular. The exception of the non-anamorphic 60×30 scale is a design choice that minimises the size of the optics of the spectrograph cameras.

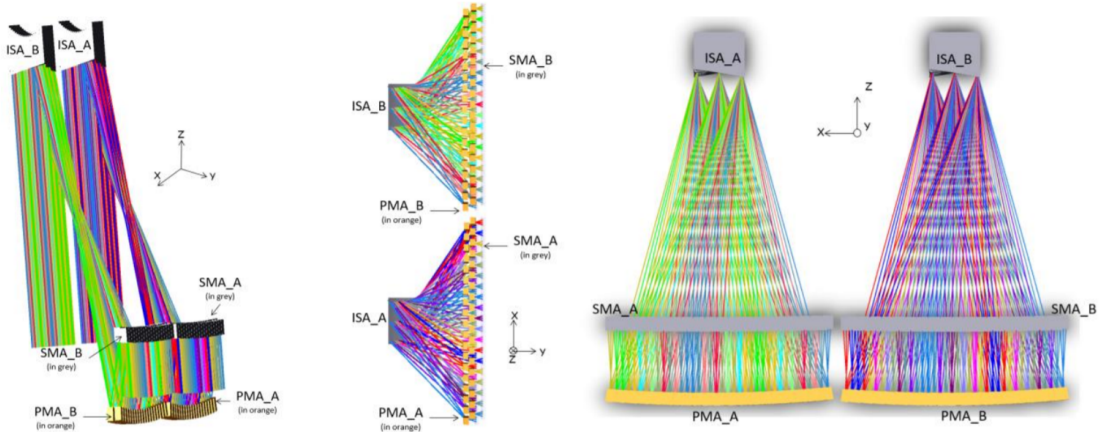


Figure 1.3: The Image Slicer Module consists of 3 optical units called the Image Slicer Unit (ISU), Pupil Mirror Unit (PMU), and the Slit Mirror Unit (SMU). Each unit consists of two assemblies (ISA-A and ISA-B; PMA-A and PMA-B; SMA-A and SMA-B). These assembly designs are slightly different.

1.2.3 Spectrograph (ISP) sub-system

The spectrograph sub-system is responsible for dispersing the light from the slicer with the required spectral resolving power and imaging the spectra onto a detector. Each IFU sub-system creates a single exit slit (~ 541.80 mm) feeding a spectrograph containing an infrared channel (and some also a visible channel). The focal plane of each camera contains a mosaic of two $4K \times 4K$ detectors with $15 \mu\text{m}$ pixels, leading to $\sim 8K$ pixels along the length of the slit. There are four spectrograph sub-systems.

The spectrograph is split into several modules, as shown in Fig. 1.4. The image slicer (IFU) creates a single entrance slit that feeds the light towards the Collimator module. The beam is then redirected towards either the VIS (in blue) or NIR arms (in red), each consisting of 3 modules: a grating module, a camera module and a science detector module. The fold mirrors also have the function of redirecting the light at the correct Littrow angle for the three families of IR spectral configurations (low, medium and high resolution).

The architecture of HARMONI is crucial to this thesis, which revolves around the calibration of Non-Common Path Aberrations as I will introduce in the following Chapter. Despite the fact that we cannot correct for NCPA after the splitting and slicing, as the light follows independent optical paths, the images that we will use to estimate NCPA will inevitably contain the *effects* of such splitting and slicing. For instance, variations in wavefront between different slicer paths (known as field-dependent aberrations) create a limit on the accuracy of the calibration.

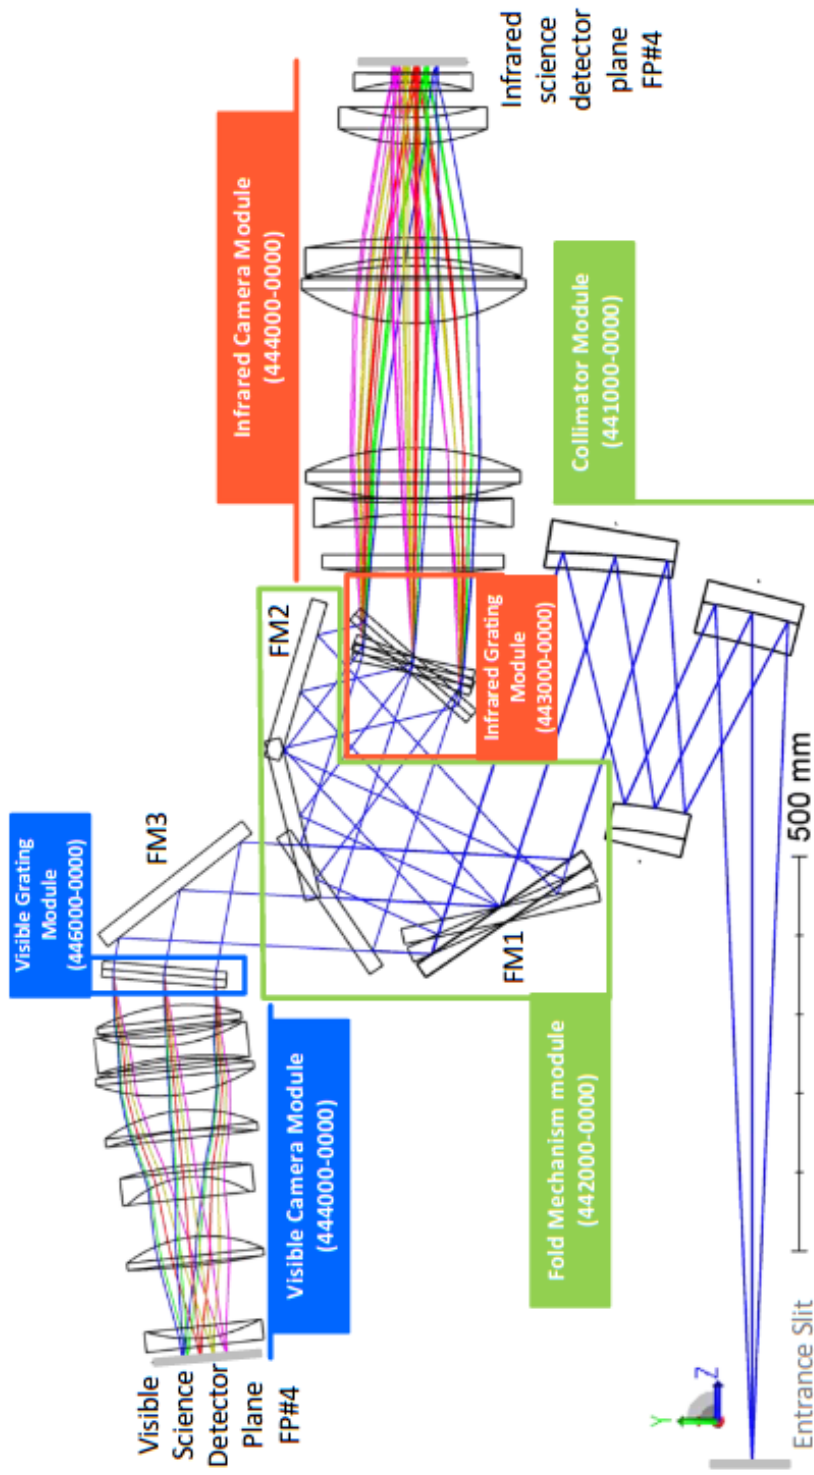


Figure 1.4: HARMONI Spectrograph (ISP) sub-system. The design is fully reflective up to the gratings. The fold mirrors can either direct the light towards the near-infrared disperser and camera, or let it pass towards the visible disperser and camera. The three overlapping gratings represent the three families of infrared dispersers (low, medium or high resolution), mounted on a grating wheel.

Table 1.2: Measured performance of 3 HAWAII 4RG-15 detectors delivered to ESO. QE: quantum efficiency, CDS: correlated double sampling

Parameter [Units]	H4RG-15 Detector		
	#1	#2	#3
Cut-off wavelength [μm]	2.50	2.52	2.51
Median QE 800 nm [%]	87	79	91
Median QE 1230 nm [%]	88	89	85
Median QE 2000 nm [%]	88	91	98
Median noise (CDS) [e^- rms]	11.3	12.4	17.0
Median dark noise [e^- /s/pixel]	0.009	0.012	0.009
Full well capacity [k e^-]	75	90	109
Inter-pixel capacitance [%]	1.1	2.0	2.2

1.2.4 Science Detectors sub-system (ISDS)

Another important part of HARMONI is the detector sub-system, as it determines how data we need for NCPA calibration is recorded and potential effects or contamination that they may contain. Here, a brief description of the system is given.

The HARMONI detector sub-system uses 4 NIR mosaics (one for each spectrograph), each consisting of 1×2 detectors. The detectors are HAWAII 4RG-15 with 4096×4096 pixels ($15 \mu\text{m}$ in size), cooled to an operational temperature of 40 K.

ESO has already received 3 science detectors from Teledyne, allowing for detailed characterisation of their performance, see the summary in Table 1.2. Some of these parameters are relevant for the modelling of noise effects on the PSF and will thus have an impact on NCPA calibration. For example, the full well capacity indicates the number of electrons at which the pixels get saturated, which links directly to the exposure time and the achievable SNR (as will be discussed in detail later on). The changes in median quantum efficiency are indicative of the so-called flat field effect, i.e. variations in sensitivity across pixels, which will be considered in the PSF modelling. The inter-pixel capacitance represents the cross-talk between neighbouring pixels, i.e. how intensity within one pixel can influence the neighbouring intensity (not considered in this thesis). Another important effect is readout noise (10 - 15 e^- per read, typical value), created within the detector electronics and affecting each read. HARMONI will read the detectors using a Sample Up The Ramp (SUTR) method, allowing for multiple non-destructive reads to reduce noise.

In conclusion, the recorded PSF is affected by the behaviour of science detectors, including effects such as readout noise, full well limits, pixel cross-talk, and flat field. For a successful NCPA calibration, these effects must be accounted for.

2

Non Common Path Aberrations

“Begin at the beginning,” the King said, gravely, “and go on till you come to an end; then stop.”

– Lewis Carroll, *Alice in Wonderland*

2.1 Origins of NCPA

The role of Adaptive Optics (AO) systems is to measure and compensate in almost real time wavefront aberrations induced by atmospheric turbulence (Tyson, 2015). In order to achieve this, AO systems typically split the light from the telescope into two separate channels: a wavefront sensing channel that is fed to the AO system for the purpose of estimating the aberrations, and the science channel that follows its own optical path through the instrument up to the detector plane to produce the observations that will be used for scientific analysis.

As a consequence of this, the light that reaches the detector follows a partially different optical path to the light used by the AO system to estimate and correct turbulence-induced aberrations, see Fig. 2.1. Any aberration introduced along this “non-common” path cannot be sensed by the AO system and thus will remain uncorrected, degrading the quality of the observations. It is important to understand that this source of error is independent of the performance of the AO system; it would be there even if the AO system were perfect, simply because it has no way of sensing it.

The aberrations that occur because of this difference in the optical paths between the AO light and the science light are known as **Non Common Path Aberrations** (NCPA). These aberrations can arise inside an optical system from multiple sources, with varying nature when it comes to time scales. They are normally referred to as




Figure 2.1: Origin of Non-Common Path Aberrations. The aberrated incoming wavefront is estimated by the Adaptive Optics subsystem in near real-time using a wavefront sensor. A correction is applied at the deformable to compensate the wavefront. However, an AO-assisted instrument cannot directly sense aberrations that originate after the beamsplitter that separates the wavefront sensing and science channels.

quasi-static aberrations, because they evolve on timescales that can range from several minutes to hours, so depending on the observation time they are seen as almost static.

With regards to the static contribution, NCPA can be caused by manufacturing errors and surface imperfections on optical surfaces along the light path. These can be partially tackled by imposing strict requirements and tight tolerances during the design phase but some residual aberrations are inevitable. During the assembly, integration and testing phase the different optomechanical components are put together. Perfect alignment is impossible, and any alignment errors will create aberrations.

Moreover, the alignment state of the system may vary with time due to variations in thermal and mechanical stresses. Differences in coefficients of thermal expansion between optical surfaces and their mechanical mounts can lead to induced stresses when there are thermal fluctuations inside the system. Such stresses can deform optical surfaces and induce aberrations. In addition, during operation, the IFS can rotate with respect to the Focal Plane Relay System (FPRS) causing NCPA to vary.

Fast varying aberrations can be caused by mechanical vibrations at the telescope, or mechanisms inside the instrument. HARMONI for instance, has to alternate between 4 different spaxel scales and 11 gratings. However, we cannot correct for aberrations in these timescales.

Some of the aberrations described above can be corrected by the Shack-Hartmann 10×10 truth sensors housed in the Natural Guide Star (NGS) arm of HARMONI.

In Chapter 6 we will look into the characteristics of the NCPA in HARMONI, using End-to-End optical models of the instrument. We will demonstrate, that the NCPA are dominated by low order aberrations like defocus, astigmatism and coma. In fact, the dichroic beamsplitter for the SCAO mode greatly contributes with astigmatism. In addition, we will investigate how the strength of aberration coefficients vary across the field of view, as this imposes a limit on the achievable post-calibration accuracy. Finally, we will also assess the impact of chromatic effects on NCPA.

2.1.1 Phase vs. Amplitude aberrations

In principle, NCPA can take the form of both amplitude and phase aberrations. *Amplitude* aberrations $\delta\Pi$ can be variations in the transmission of optical elements or other pupil irregularities; whereas *phase* aberrations $\delta\Phi$ represent variations in optical path, i.e. wavefront errors. Although they both affect the PSF as they modify the complex pupil function $P(\mathbf{u})$:

$$P(\mathbf{u}) = (\Pi + \delta\Pi) \exp\left(2\pi i \frac{\delta\Phi}{\lambda}\right) \quad (2.1)$$

we are only interested in correcting phase aberrations $\delta\Phi$ for several reasons. First of all, from a practical point of view, we can only correct phase aberrations, by applying a correction at the entrance pupil of the instrument using M4 of the ELT, as done for the AO correction. Secondly, it is possible to show that the impact that amplitude aberrations $\delta\Pi$ have on the PSF is comparatively small to the impact of phase aberrations. To demonstrate this, we ran the following experiment.

We simulated PSF images with random aberrations. On the one hand, we simulated images with only amplitude aberrations $\delta\Pi$ and a flat wavefront, such that:

$$P_{\Pi} = (\Pi + \delta\Pi) \exp(2\pi i \cdot 0) \quad (2.2)$$

On the other hand, we simulated images with only phase aberrations and a constant pupil amplitude:

$$P_{\Phi} = \Pi \exp\left(2\pi i \frac{\delta\Phi}{\lambda}\right) \quad (2.3)$$

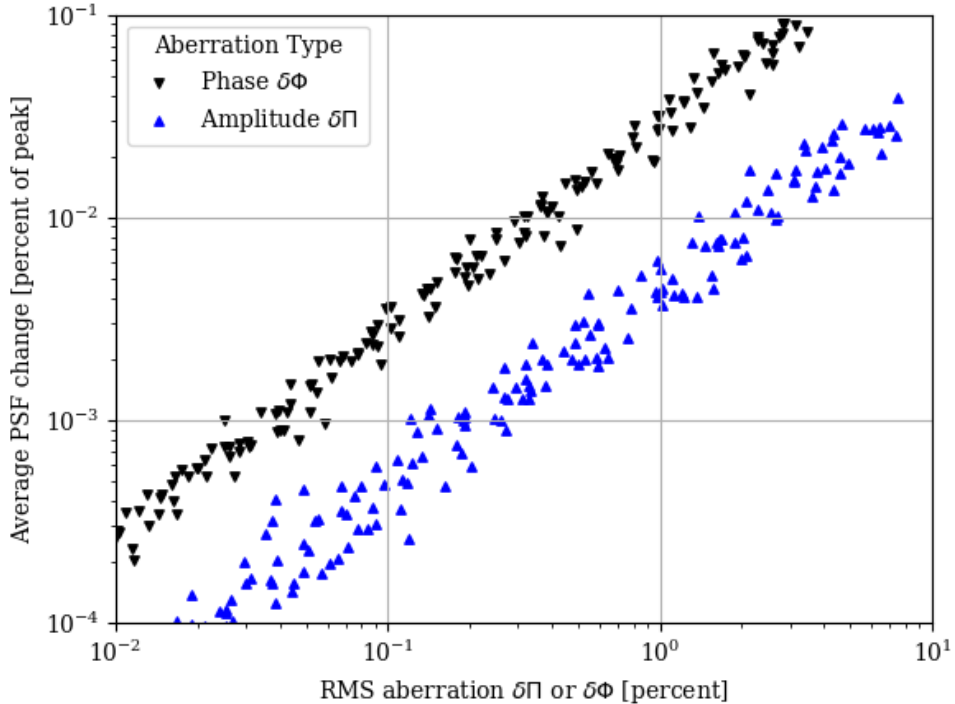


Figure 2.2: Phase vs. Amplitude aberrations. Average change in PSF intensity η_x as a function of RMS aberration content for *phase* $\delta\Phi$ and *amplitude* $\delta\Pi$ aberrations. The impact of amplitude aberrations is comparatively small and is usually disregarded.

For each case, we calculated the average change in PSF intensity η compared to the nominal case:

$$\eta_x = \frac{1}{N^2} \sum_i^N \sum_j^N |\text{PSF}_{ij}(P_x) - \text{PSF}_{ij}(x=0)| \quad (2.4)$$

where x can be either Φ for phase aberrations or Π for amplitude aberrations, and N represents the number of pixels in the focal plane images.

Analysing η allows us to assess how amplitude aberrations affect the PSF compared to phase aberrations. The results are shown in Fig. 2.2 as a function of the RMS aberration. In the case of $\delta\Pi$, this represents the RMS change in pupil amplitude (in percentage). For the case of $\delta\Phi$ this represents the RMS wavefront error in units of percentage of a wave.

The results suggest that the change in PSF intensity is almost an order of magnitude smaller when the aberrations are in the form of amplitude changes compared to phase aberrations. This supports our previous statement that amplitude aberrations (which cannot be corrected directly) are less significant than phase aberrations, which we can correct for once they are estimated through NCPA calibration.

2.2 Why care about NCPA?

If left uncorrected, NCPA degrade the quality of scientific observations as they lower the optical performance of the instrument. Thus, mitigating NCPA is beneficial to all science cases. But NCPA are usually small in magnitude so in the past, they sometimes were overlooked, incorporated as part of the wavefront error budget, in favour of other dominant effects that were driving the performance. However, as new instruments and AO systems kept pushing their performance limits, and the science cases became ever more demanding, NCPA started to gain relevance and popularity.

In particular, in the relatively novel era of exoplanet research, one of the most sought after discoveries is the detection and characterisation of Earth-like analogues, orbiting solar-type stars. In the context of high-contrast imaging of exoplanets, residual quasi-static speckles have been identified as the main limitation to the achievable contrast (Sivaramakrishnan et al., 2008; Soummer et al., 2007). These speckles are mainly caused by residual NCPA, which lead to imperfect light suppression in coronagraphic systems. Due to their quasi-static nature these speckles do not average out in long exposures, in contrast to atmospheric aberrations. Moreover, they resemble exoplanet signals, which makes unambiguous detection even more challenging.

An example of this issue is evident in the discovery of the exoplanet GJ 504 b by Kuzuhara et al. (2013), shown in Fig. 2.3 where the field around the host star is crowded by speckles that resemble exoplanet signals. A similar situation occurs for the discovery of 51 Eri b, as reported in Macintosh et al. (2015), see Fig. 2.4. In both cases, the planets are Jovian and are located at significant distance from their host stars (~ 43.5 AU for GJ 504 b and ~ 13 AU for 51 Eri b), which is not surprising given the selection bias of the direct imaging method towards young, massive planets orbiting at large distances.

Unfortunately, in the quest for Earth-like planets the region closest to the host star is the most relevant; orbital distances of 1 AU will span $0.1''$ for host stars at 10 parsecs or less, but only around 400 stars are located within that distance of the Sun. In order to probe these regions, currently crowded with speckles, coronagraphic instruments need to operate close to their maximum capabilities, which imposes stringent requirements on the level of residual aberrations.

HARMONI is not specially designed to be a planet-finder imager, but rather, its integral-field spectroscopy capabilities are more tailored to the task of spectrally characterising exoplanet atmospheres. For this purpose, techniques have been proposed to detect the absorption of molecular species in exoplanet spectra in AO-assisted




Figure 2.3: Discovery images of the exoplanet GJ 504 b. Panel **a** shows the intensity after suppressing flux from the central star. Panel **b** shows the associated signal-to-noise ratio. Source: [Kuzuhara et al. \(2013\)](#)




Figure 2.4: Discovery images of the exoplanet 51 Eri b, showing a companion planet after PSF subtraction. Source: [Macintosh et al. \(2015\)](#)

integral-field spectrographs like HARMONI. For instance, [Hoeijmakers et al. \(2018\)](#) propose the technique known as **molecular mapping**, which cross-correlates the integral-field spectra with molecular template spectra to search for the spectral signatures of spatially resolved exoplanets that are embedded in the photon noise of the dominating star light.

Since the speckle pattern is chromatic ([Racine et al., 1999](#)), although speckles remain fixed in space for a given filter, they will reside in different locations in other passbands ([Hinkley et al., 2007](#)). In addition, because they are quasi-static in nature with a lifetime ranging from several minutes to hours ([Hinkley et al., 2007](#)), they could come and go across different spectral regions creating spurious spectral features, which could potentially contaminate the molecular mapping and thus the need to correct NCPA to the best of our abilities. Moreover, even though the effects from NCPA can be palliated *a posteriori* during the post-processing of high-contrast data, this does not remove photon noise, so it is advisable to tackle it at the source prior to the actual observations.

The importance of mitigating NCPA is also highlighted in [Carlotti et al. \(2018\)](#) for the case of the High Contrast Module (HCM) in HARMONI, where the presence of quasi-static aberrations represents the main limitation to the achievable contrast.

All of this shows that the estimation and correcting of Non-Common Path Aberrations is one of the most pressing problems for the upcoming generation of astronomical instruments and represents a major technical challenge for the detection and characterisation of exoplanets.

2.3 Modelling NCPA

In order to correct NCPA, it is first necessary to define a way to model and measure the aberrations. For this purpose, throughout this thesis we make use of Zernike polynomials ([Noll, 1976](#)), a basis of global polynomials defined in a unit circle. One of the main advantages of Zernike polynomials is that they follow the types of aberrations that are typically observed in optical systems (astigmatism, defocus, coma, trefoil) so they constitute an effective way of modelling wavefront maps.

There are multiple notations when it comes to defining Zernike polynomials, which vary in terms of the ordering, the index attached to each aberration and the normalisation. Here, we have used the "Zernike Fringe Coefficients" (following Zemax Optics Studio), which starts assigning Z_1 for piston, Z_4 for defocus, Z_5 and Z_6 for astigmatism, etc; and normalises each aberration to peak-to-valley ± 1 . It should be noted

that in specific cases, when dealing with RMS wavefront error, we have modified the normalisation to follow the "Zernike Standard Coefficients" notation (which normalises each aberration to RMS of 1) so that the coefficients directly represented the contribution to the RMS wavefront.

2.4 Research Goals

In this Section, we will explore the main goals for this project as well as the contributions that this thesis has made to this field of research.

2.4.1 Problem definition

As explained above, HARMONI is an integral field spectrograph assisted by adaptive optics. We have seen that this type of instrument can be affected by NCPA, as the adaptive optics sub-system cannot sense the aberrations downstream of the beam-splitter that separates the light used for AO wavefront sensing and the light used for scientific purposes.

I have also explained why NCPA are detrimental to the overall performance of the instrument and, in particular, how they can limit the detectability of exoplanets by creating speckle noise in high-contrast imaging. Thus, it is clear that a strategy for mitigating the impact of NCPA is needed in HARMONI.

But, as we will see in the literature study (Section 2.7), a variety of techniques to measure NCPA already exists and have been applied successfully in other context. Then **why is the case of HARMONI any different?**

The answer to this question lies in the characteristics of the instrument. HARMONI is an integral field spectrograph, capable of obtaining a spectrum at every pixel position for a 2D patch of the sky in a single exposure. This works by reformatting the 2D input field into a long pseudo-slit that can be fed through a spectrograph. In order to achieve this, HARMONI uses an Integral Field Unit (IFU), which contains an optical component known as an **image slicer**: a stack of tilted thin mirrors that slice the field and send light from each slice in different directions, followed by a set of pupil mirrors that re-image each slice to form a pseudo-slit composed of individual slitlets.

It is the presence of an **image slicer** that calls into question the applicability of state-of-the-art techniques for NCPA calibration to the particular case of HARMONI. Without going into technical details yet, as this will be thoroughly explained in Chapter 3, some of those techniques rely on the use of a defocus to break phase

degeneracy, and the response of image slicers to such defocus is not easy to model / compute. Because of their *slicing* function, there is a risk that, in defocused conditions, the light would fall onto neighbouring slices and follow unintentional optical paths. Moreover, significant diffraction takes place at the slicer, which could further distort the PSF. But without detailed characterisation of the impact of image slicers on light propagation, there is no guarantee that conventional techniques could work.

It is within this context that the work of this DPhil project is framed. The calibration of Non Common Path Aberrations in instruments containing image slicers has not been demonstrated before.

2.4.2 Research question

Now that the necessary background has been presented, we can describe the main research objective for this DPhil Thesis:

The research objective is to develop techniques for measuring and calibrating Non-Common Path Aberrations in adaptive optics assisted spectrographs containing Image Slicers.

which can be rephrased into the following research question:

Is it possible to successfully measure and calibrate Non-Common Path Aberrations in the context of adaptive optics assisted spectrographs that contain Image Slicers?

In order to properly address this research question, it is important to first tackle in a consistent and systematic manner a series of related sub-questions that provide valuable insight into the whole problem:

1. **How does the presence of an Image Slicer impact the PSF?** At the beginning of this project, because of unsuccessful past attempts at performing NCPA calibration on an instrument containing an image slicer (SWIFT, [Tecza et al. \(2006, 2008\)](#)), the general hypothesis was that traditional techniques could not work because of the effects that the presence of an image slicer has on the PSF. Thus, the first thing to do is to thoroughly characterise how image slicers influence the PSF.

2. **What alternative techniques could be applied?** If image slicers significantly alter the PSF, it is possible that state-of-the-art techniques for NCPA calibration may not work properly. In that case, two possibilities arise:
 - (a) *Modify conventional techniques to account for slicer effects.*
 - (b) *Develop alternative techniques that are not affected by these effects.*

3. **What constitutes a "successful" calibration?** Once a viable approach to NCPA calibrations is identified, it is necessary to define a way to judge the success of the calibration. This entails defining requirements with regards to the quality of the predictions that can be used to test whether the objective has been reached.

2.4.3 This thesis' contribution

Here, I present the contribution that this DPhil thesis has made to each of the sub-questions needed to answer the *Research Question*:

1. **Image Slicer effects.** The impact of image slicers on the PSF has been thoroughly characterised from two perspectives. First of all, from a *theoretical* standpoint, I have investigated the effect that image slicers have on light propagation in Chapter 3. For this purpose, I have developed a mathematical model that captures the most relevant characteristics of an IFU and used it to simulate the impact of several parameters. In addition, these predictions have been validated using independent simulations with optical models of the HARMONI IFU sub-system in Zemax Optics Studio, by means of Physical Optics Propagation (POP), a technique that accurately captures diffraction effects at the slicer.

Secondly, from an *experimental* perspective, I designed, constructed and carried out an experiment with the engineering model of the SWIFT Image Slicer, to further validate these predictions, see Chapter 4. Both approaches show that image slicers can have a significant effect on the PSF, mainly in the form of fringes parallel to the orientation of the image slicer.

2. **Alternative NCPA calibration techniques.** In light of the results regarding the impact of image slicers on the PSF, the question of *alternative* techniques gained importance.

- (a) *Modifying conventional techniques to account for slicer effects.* Unfortunately, due to time constraints, exploring the possibility of modifying state-of-the-art techniques such as Phase Diversity to include the effects arising from image slicers was not possible.
- (b) *Developing alternative techniques.* This represents the major contribution of this DPhil thesis. In Chapter 5, I introduce a novel approach to NCPA calibration that relies on Machine Learning techniques. It takes advantage of the fact that some of these methods excel at identifying patterns in image data to circumvent the drawback of image slicer effects. The main finding of this study is that machine learning techniques can be successfully applied to estimate NCPA directly from PSF data even when the images contain image slicer features.

This work is further supported by End-to-End model efforts, see Chapter 6. As part of this DPhil thesis, I have developed the most realistic representation of HARMONI yet, the optical End-to-End model that combines the optical models of all sub-systems to simulate the behaviour of the complete instrument. Using Monte Carlo analysis, it is possible to obtain realistic wavefront maps for HARMONI and use them to train the machine learning models.

3. **Requirements definition.** The End-to-End model was also used to estimate how accurate the NCPA predictions have to be, i.e. to define a requirement on the accuracy of the calibration. Field-dependent aberrations, the phenomenon of aberrations varying in magnitude across the field of view imposes a limit on how well aberrations can be corrected, since one can only apply a constant correction at the pupil plane that affects all field points. I used Monte Carlo analysis of the HARMONI End-to-End models to characterise the impact of field dependent aberrations.

The work of this DPhil thesis extends across several disciplines and topics, but as Fig. 2.5 shows, they are all interconnected and together contribute to answering a much larger question: **how can Non-Common Path Aberrations be measured and corrected in the presence of image slicers?**

Undoubtedly, much remains to be done before we can use HARMONI on the ESO’s ELT at Cerro Armazones to deliver astounding observations, but I hope that this thesis is a humble step in the right direction.

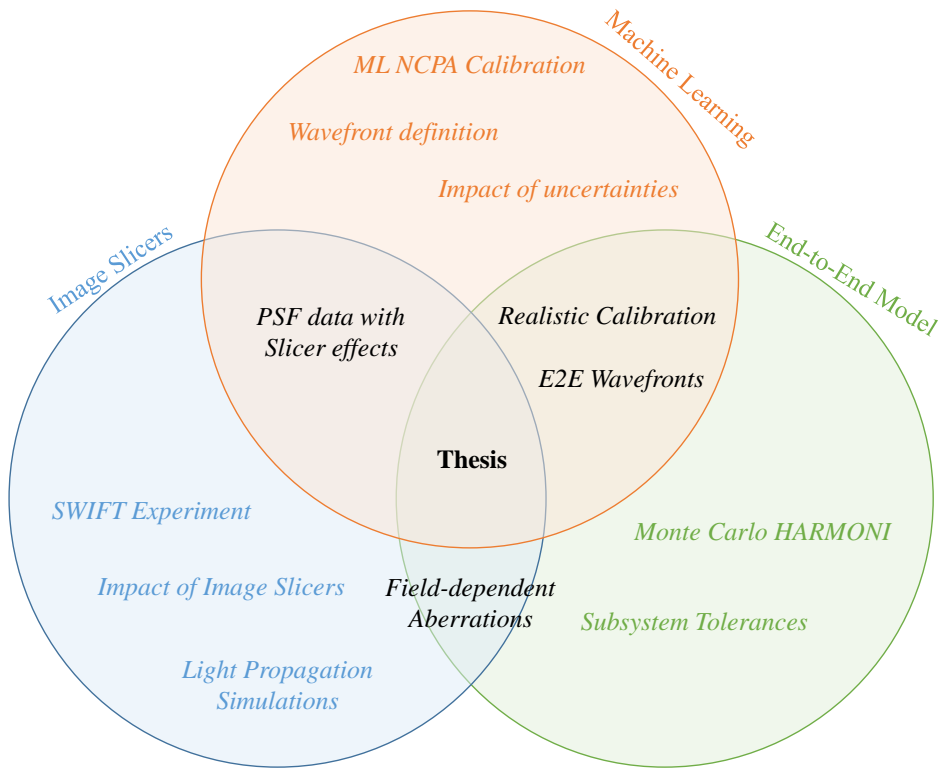


Figure 2.5: Venn diagram of the major components of this DPhil thesis.

2.5 Judging Success

Even though it is easy to state that the general goal of this thesis is to develop a framework of calibration techniques to tackle the problem of Non-Common Path Aberrations (NCPA), one needs to define a way of judging whether this has been achieved and to quantify to what extent these efforts have, or have not been successful. In broad terms, this implies being able to answer the following questions:

- **How much NCPA can we expect to have initially?** In other words, a reasonable estimate for the aberration level of the system, and other relevant characteristics, prior to calibration.
- **How well do we need to correct NCPA?** This entails deriving a set of performance requirements, in the form of quantitative goals that clearly state the level of NCPA that is tolerable after correction, as well as the conditions in which these requirements apply.

These two items combined define both the starting point and the goal for the calibration. A successful framework for NCPA calibration is one that, given the

expected level of NCPA in the system, can deliver a correction that fulfils the defined requirements and does this in a robust and repeatable way. If these conditions are met, then we have succeeded. Although this is the most desirable outcome, it might not always be the case that we can claim the calibration framework is successful. In such circumstances, however, it is important to understand what the main limitations are, and try to answer questions such as: why is the requirement unachievable? What is missing that could be improved in the future?

It should also be noted that an accurate measurement of NCPA is always useful, even if a direct correction is not possible. For instance, normally, variations across the field of view or with wavelength cannot be corrected, but accurately measuring the extent of these effects could be used later in post-processing to correct the observations.

With regards to the first question of **how much NCPA we can expect**, this stems from whatever knowledge is available on the characteristics and behaviour of the instrument. I will address this in Section 6.2.1. Using Monte Carlo instances of the End-to-End HARMONI model, which aggregate the different sub-systems and account for manufacturing and alignment tolerances and other sources of error, we can calculate the expected RMS wavefront error for the as-built HARMONI instrument. This also allows us to investigate the distribution of RMS across different aberrations, to see if the wavefront is dominated by low order Zernike polynomials and to calculate how many Zernike polynomials we need to correct to reach a certain residual error.

What we will see in Section 6.2.1 is that, for the 4×4 mas scale, in H band and SCAO mode, the RMS wavefront error ranges between 15 and 60 nm with an average about 30 nm. In particular, the wavefront is dominated by low order Zernike aberrations, including the astigmatism due to the presence of the dichroic beamsplitter that separates the light for the AO wavefront sensing channel and the science channel.

Regarding the second question of **how well we need to correct NCPA**, this is tackled in Section 2.6. I will take field-dependent aberrations (the fact that the aberration coefficients vary across slices and between IFU paths) as the absolute floor of the calibration. The rationale for this will be provided in Section 6.2.2 using the Monte Carlo End-to-End HARMONI model as well. This is estimated to be around 15 nm for the 4×4 mas scale, in H band.

In conclusion, for the NCPA calibration to be successful we have to develop a technique capable of sensing aberrations in the order of 30 nm (to a maximum of around 60 nm) and correcting them to a residual of around 15 nm.

2.6 Deriving requirements

In the previous section, we explained that in order to assess whether the calibration of NCPA has been successful, we need to have some form of *requirement* that defines the level of acceptable NCPA. In principle, such requirement should arise naturally as a result of a flow-down from the top level science goals of the instrument to its technical specifications. The truth is that most of the time, translating the sometimes broad and general science goals into clear and concise statements regarding NCPA accuracy can be a tortuous path.

In the case of high-contrast imaging, this is relatively straightforward. Given the dramatic effect that quasi-static aberrations have in limiting contrast, the requirement for NCPA correction is usually **”as good as physically possible”**, with typical numbers being around 15-25 nm RMS to ensure raw contrast levels of around 10^{-5} , see [Carlotti et al. \(2018\)](#) for the specific case of the High-Contrast Mode (HCM) in HARMONI.

For less demanding science cases, the answer is not so clear. In any case, in the following sections, we will try to derive a requirement for NCPA calibration from two perspectives: ensquared energy considerations, and field-dependent aberrations.

2.6.1 Requirements from Ensquared Energy

HARMONI is a workhorse instrument designed to deliver high quality observations over a wide range of science cases. This means that the parameter space that HARMONI has to cover is quite vast as the specific instrument requirements vary across science cases, some demanding diffraction-limited PSF sampling, others requiring large spectral coverage or high-resolution spectroscopy, while others lean towards coarse spaxels to be able to observe faint targets.

As a result of this, HARMONI offers 4 different spaxel scales: 4×4 , 10×10 , 20×20 , and 30×60 milli-arcseconds (mas). Only at the 4 mas scale is the HARMONI PSF Nyquist-sampled (at around $1.5\ \mu\text{m}$ wavelength). For the rest of spaxel scales, the sampling is coarser than what the Nyquist theorem requires for the diffraction limited core of the PSF. This wide range of spaxel scales and associated PSF samplings, combined with the myriad of science cases HARMONI can tackle, imply that the NCPA calibration requirements must vary across spaxel scales and observations.

At fine scales (4 mas) where the effect of aberrations is resolved as the PSF is close to Nyquist-sampled, Strehl ratio (closely related to RMS wavefront error) remains the key performance metric. However, at large scales (20 mas or above) for which the

central pixel covers most of the PSF, Strehl ratio loses its relevance as an observer would be more interested in sensitivity than spatial resolution. In such conditions where the PSF is coarsely sampled, NCPA calibration can become more challenging but also less relevant: there is not much point attempting to correct aberrations you cannot even resolve. With such coarse pixels, the effect of aberrations is not entirely resolved, as a large fraction of the PSF is integrated into a single pixel. This makes the ensquared energy of the central pixel a more meaningful metric, and NCPA calibration at such scales should be aimed at ensuring that aberrations do not degrade the integrated intensity too much.

A priori, we are most interested in the highest resolution scale (4 mas spaxel) where the PSF is properly sampled, and there is a direct link between RMS wavefront aberrations and Strehl ratio. However, it should be noted that the HARMONI LTAO Science Requirements document (HRM-00254), which analyses various science cases proposed by the Project Science Team (ESO committee) in the document ESO-191883 and derives requirements for H-LTAO, suggests the PST cases are actually dominated by ensquared energy cases, not Strehl. This means that for a wide range of science cases, NCPA calibration at coarse spaxel scales is relevant.

Consequently, we decided to investigate the effect of aberrations on the ensquared energy to better understand how we can derive NCPA calibration requirements at coarse scales. The way we compute the ensquared energy is as follows: assuming a circular pupil with a central obscuration of $\sim 30\%$ (ESO's ELT), we compute the PSF for varying wavefront error maps $\text{PSF}(\Phi)$ at $1.5 \mu\text{m}$, with a sampling of 0.5 mas per pixel. Then, we aggregate the energy inside a box of 2×2 spaxels for each spaxel scale (4, 10, 20 milli-arcseconds), and relate that energy $E(\Phi)$ that depends on the wavefront map Φ to the nominal ensquared energy E_0 of a PSF with no aberrations. This allows us to track the change in ensquared energy as a function of the wavefront characteristics, such as RMS and spatial frequency content.

$$\delta E = \frac{E(\Phi)}{E_0} \quad (2.5)$$

Fig. 2.6 shows an example PSF with a random Zernike wavefront map, sampled with 0.5 milliarcseconds per pixel. The white box covering 40×40 milliarcseconds represents what a 2×2 spaxel box of 20 mas spaxels would cover. Integrating the intensity within the box allows us to compute the change in ensquared energy due to aberrations at the 20 mas spaxel scale. The reason a 2×2 box is used is because at fine scales, pointing errors that change the location of the image with respect to the




Figure 2.6: The Enquared Energy calculation represents the fraction of total PSF energy contained within a given box size. In the case of HARMONI, the requirement specifies a 2×2 spaxel box. Within this box the intensity is integrated and compared to the nominal case of a PSF without aberrations. *Left*: a PSF with random aberrations sampled with 2 milli-arcsecond (mas) pixels and a box covering 2×2 20 mas spaxels. *Right*: detail of the intensity inside the 2×2 box that will be integrated into a total energy value.

pixel boundaries can greatly affect the result of the ensquared energy calculation in a 1×1 pixel box.

For this analysis, we investigated the behaviour of the ensquared energy at the 4, 10 and 20 mas spaxel scales using Zernike polynomials. For each scale, we generated multiple PSF images with random wavefront maps of varying RMS. The results are shown in Fig 2.7. In order to highlight the influence of the spatial frequency of the aberrations, we ran the analysis in a row-by-row basis of the Zernike polynomials. In other words, each subplot in Fig 2.7 uses wavefront maps containing Zernike polynomials of a given radial order ρ^n ; for example: the top left panel for ρ^2 accounts for astigmatism and defocus (Z_3, Z_4, Z_5 polynomials), while the top right panel for ρ^3 includes coma and trefoil ($Z_6 - Z_9$ polynomials).

The first thing to notice is that at the finest spaxel scale (4 mas) the ensquared energy does not depend on the radial order (a proxy of the spatial frequency of the aberrations). In addition, this scale is the most sensitive to the intensity of the wavefront errors, an RMS of 120 nm already degrades the ensquared energy inside the 2×2 box by 20%. For low order aberrations, the 20 mas spaxel scale retains most of the ensquared energy, 150 nm RMS only degrade it by around 2.5%. However, as higher spatial frequency aberrations are introduced, the loss in ensquared energy becomes more significant.

The main implication of this finding is that at the coarse scales, the level of NCPA calibration needed will depend on the initial characteristics of the wavefront errors Φ . If the NCPA are expected to contain mostly low order aberrations, then a higher RMS can be tolerated before it makes a significant impact on the ensquared energy.

In principle, the stronger impact of high frequency aberrations could be problematic since at such coarse scales it is difficult to measure these aberrations due to poor PSF sampling. In any case, End-to-End model simulations of the HARMONI instrument indicate that the wavefront errors are mostly dominated by low-order aberrations, and that the strength of the aberrations quickly decays with radial order, see Section 6.2.2.

These results raised the question of why the ensquared energy at the finest scale (4 mas) does not depend on the spatial frequency of the aberrations while the coarser scales seem to be sensitive to this parameter. In order to elucidate this we ran the following experiment. We studied a 1D PSF with aberrations in the form of a pure cosine of varying spatial frequency (in cycles / diameter). We compared the residual difference between the aberrated PSF and the nominal one, as shown in Fig. 2.8 for varying spatial frequencies. The cosine aberration has two opposing effects: first of all,

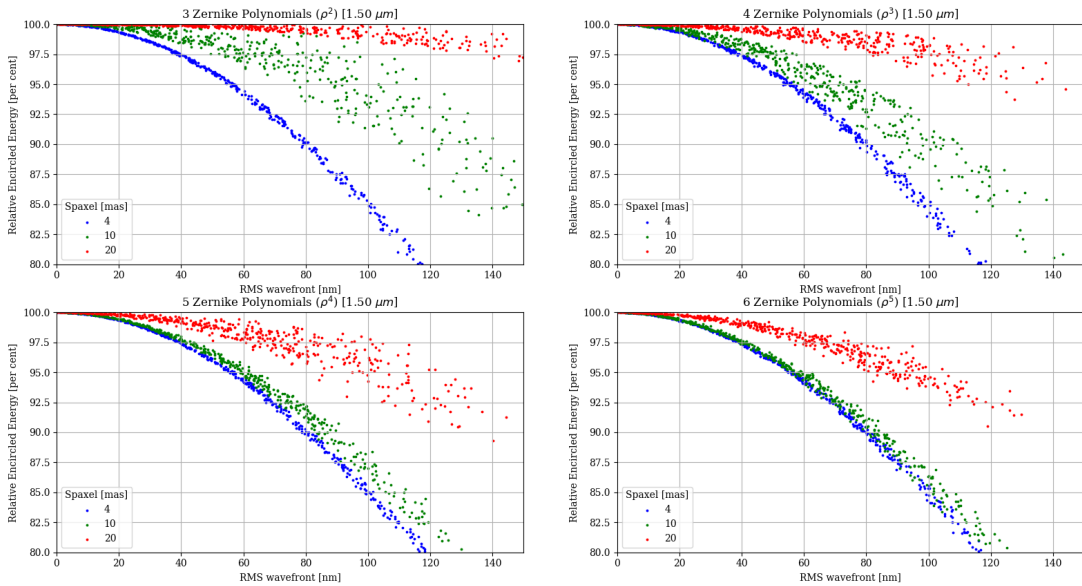


Figure 2.7: Ensquared Energy as a function of RMS wavefront error for the 4, 10 and 20 milliarcsecond spaxel scales (2×2 box). Each panel corresponds to a given Zernike radial order ρ^n used in the wavefront definition. Each dot represents a case of a PSF with random aberrations. The different profiles indicate how the strength of aberrations degrade the ensquared energy for each spaxel scale in HARMONI.

it suppresses the central peak independently of the spatial frequency; and secondly, it excites peaks of slightly lower intensity at both sides of the PSF, that move further out with higher frequencies.

This implies that at coarse scales and low spatial frequencies, the 2×2 spaxel box is large enough to cover both the central peak suppression and the symmetric peak excitation which, when integrated over the whole box, leaves the ensquared energy almost unchanged. As the spatial frequency of the aberrations increases, the peaks move outside the box and only the central peak suppression remains, significantly degrading the ensquared energy.

At the 4 mas spaxel scale, the 2×2 box barely covers the core of the PSF (Nyquist sampling at 1.45 μm at that scale for HARMONI). Consequently, the excited peaks always lie outside the box and the ensquared energy is independent of the spatial frequency of the aberrations.

So far we have analysed 3 out of the 4 available spaxel scales in HARMONI. We decided to treat the 60×30 mas spaxel scale separately due to both its significantly larger size and its rectangular shape. Fig. 2.9 shows the evolution of ensquared energy with RMS wavefront within a 2×2 box for that spaxel scale. The first thing to notice is that this spaxel scale can take a large RMS wavefront error before the ensquared

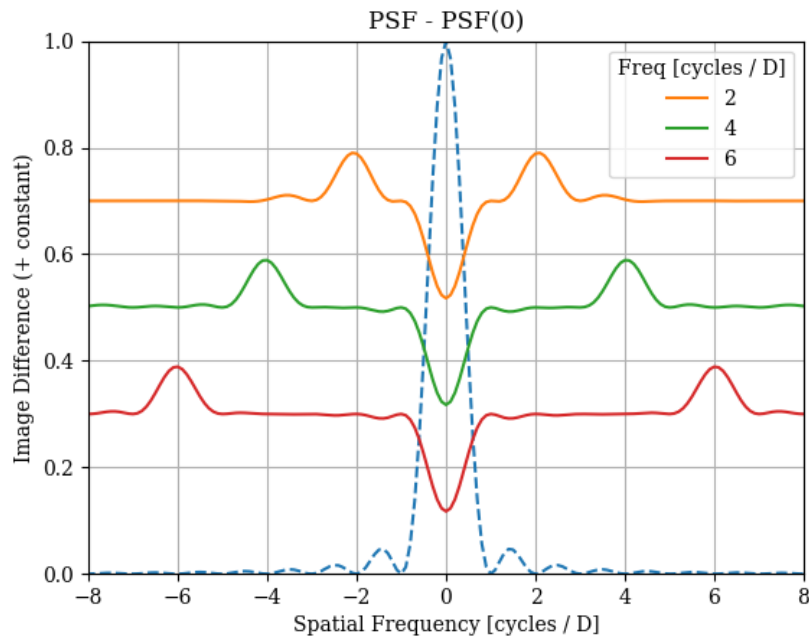


Figure 2.8: Impact of cosine aberrations on the PSF as a function of spatial frequencies. Dotted line shows nominal PSF for reference, the coloured lines indicate the difference between the aberrated PSF and the nominal one. In all cases, the aberration suppresses the PSF core and excites two peaks at either side of the core whose position varies with spatial frequency. These two effects nearly compensate each other when the intensity is integrated over a sufficiently large box. But at the 4 mas scale, the 2×2 box barely covers the core of the PSF, always leaving the 2 excited peaks outside, which explains why there is no influence of the spatial frequency of the aberrations on the ensquared energy at this scale.

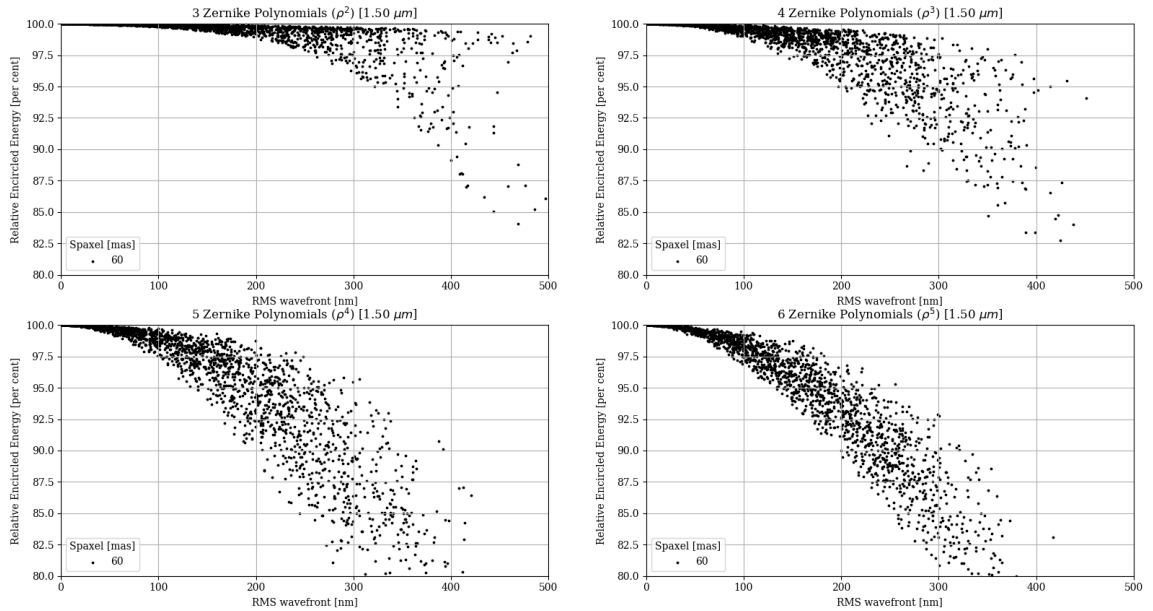


Figure 2.9: Ensquared Energy as a function of RMS wavefront error for the 60×30 milliarcsecond spaxel scales (2×2 box). Each panel corresponds to a given Zernike radial order ρ^n used in the wavefront definition. Each dot represents a case of a PSF with random aberrations.

energy starts to degrade significantly. More than 300 nm RMS of astigmatism / defocus are needed to reduce the ensquared energy by 5%. In the worst case scenario, at ρ^5 radial order 200 nm RMS lead to a reduction of 10% in ensquared energy.

The large spread in ensquared energy for this scale could be explained by the asymmetry in the spaxel box shape, which covers twice the area in one direction. If the wavefront map happens to contain aberrations that predominantly displace intensity along the longest side of the box, the loss in ensquared energy is partially mitigated. As for this analysis we randomly generate the wavefront maps, it could be that some cases are preferentially oriented compared to others.

Another interesting thing to analyse is how the distribution of the aberration coefficients affects the evolution of ensquared energy. So far we have generated wavefront maps for which the Zernike coefficients are uniformly distributed at a given radial order. This is not entirely realistic as usually the strength of the aberrations decays with radial order, as we will demonstrate in Chapter 6. For this new analysis we decided to generate wavefront maps with Zernike polynomials up to order ρ^{14} (approximately 100 polynomials), meaning the wavefronts now include aberrations over a wide range of spatial frequencies. To model variations in spectral distribution we investigated three cases for which we rescale the uniformly distributed coefficients so

Table 2.1: Levels of acceptable NCPA RMS [nm] for each spaxel scale in HARMONI, for several ensquared energy thresholds.

$E(\Phi)/E_0$	Spaxel Scale [mas]			
	4×4	10×10	20×20	60×30
97.5	20	30	60	120
95	35	45	85	150
90	50	65	130	200

that their intensity decays with radial order ρ^n . In other words, for any given random wavefront map we rescale the Zernike coefficients by a decaying factor r , enhancing the low order aberrations and damping the higher order ones.

$$\Phi = \sum_i^N a_i r^{(i/N)} Z_i \quad (2.6)$$

The results for this analysis are shown in Fig. 2.10 for the 60×30 mas scale. The black dots for $r = 1.0$ represent the nominal case where the Zernike coefficients are uniformly distributed. The red and blue dots represent cases where the aberration coefficients are rescaled down such that the higher spatial frequencies are 10 and 100 times weaker than the low spatial frequencies, respectively. There is a slight difference in behaviour for $r = 0.1$, the decay in ensquared energy being less pronounced than for the $r = 1.0$ case: at 200 nm for $r = 1.0$ the ensquared energy ranges around 85-87.5%, while at $r = 0.1$ it ranges around 90-92.5%. Further dampening of high spatial frequencies at $r = 0.01$ and above leads to almost no change in behaviour.

All of this allows us to derive requirements for the NCPA calibration. For instance, if the performance goal is to not degrade the ensquared energy within a 2×2 spaxel box at the 60×30 mas spaxel scale by more than 10%, Fig. 2.10 suggests that around 200 nm is the maximum RMS residual error we can tolerate (for realistic wavefronts where the strength of aberrations decays with spatial frequency).

This analysis was repeated across all spaxel scales for a final comparison. Fig. 2.11 shows the evolution of ensquared energy with RMS for realistic wavefront maps Φ containing 55 Zernike polynomials, with decays ratios of $r = 0.1, 0.01$. The solid lines indicate the 5%-ile, dashed lines indicate 95%-ile. Table 2.1 summarises these findings by specifying the level of RMS that each spaxel case can accept for several ensquared energy thresholds, using the 5%-ile values.




Figure 2.10: Ensquared Energy as a function of RMS WFE for the 60×30 milliarc-second spaxel scales (2×2 box), for several rescaling factors r representing the decay of coefficient strength with aberration order. $r = 1.0$ represents wavefront with equal proportion of low and high order aberrations (which is not entirely realistic). The other two cases $r < 1$ correspond to wavefronts with decaying aberration strength (realistic), showing a more shallow decay of ensquared energy with RMS WFE.

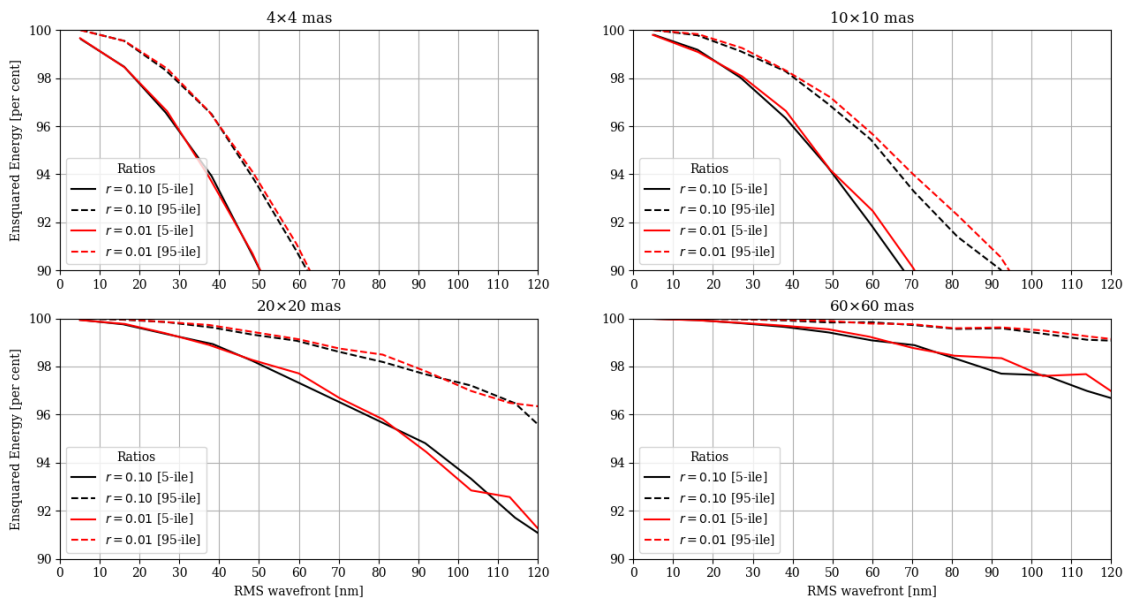


Figure 2.11: Evolution of Ensquared Energy with RMS wavefront error, all scales in HARMONI. Two cases are shown (*red* and *black*) for each scale representing the possible decay factors r of aberration coefficients in the wavefront, showing no influence. Both the 5%-ile (solid) and 95%-ile (dashed) boundaries are shown to indicate the possible range of variation; these values are calculated across many instances of random wavefront maps.

In conclusion, in this section we have shown that the behaviour of ensquared energy within a 2×2 spaxel box widely depends on the particular scale we are investigating (from the finest 4 mas to the coarse 60×30 mas), and on the spatial frequency of the wavefront. With a simple analysis we have demonstrated the reasons behind this variation in behaviour with respect to spatial frequency. Using realistic wavefront maps Φ for which the aberration coefficients decay with radial order, we have calculated the distribution of ensquared energy as a function of RMS wavefront error. These results can be used to derive requirements on how well the NCPA need to be calibrated. Ultimately, these metrics can be used to test whether the algorithms we develop are successful.

2.6.2 Requirements from Field Dependent Aberrations

The pre-compensation of NCPA is limited to a single correction at the entrance pupil plane, irrespective of field variations or chromatic effects. But it is important to remember that the HARMONI field of view is split among four IFU paths (each with its own spectrograph) and that within each path, light follows different optical paths for every slice of the image slicers. This can cause the aberrations to vary depending on the field of view position but, since only one correction is possible, this could potentially lead to residual field-dependent aberrations.

It is important to characterise the range of variability of these field-dependent aberrations because it imposes a limit on the quality of the correction that one can demand from the calibration algorithms. Estimating the NCPA with greater accuracy than the intrinsic range of the field-dependent aberrations is pointless, because the corrected wavefront will be dominated by those aberrations. It should be noted that this effect is rather small and will not limit the performance of the HARMONI instrument, but it does represent a floor on the accuracy of the NCPA correction and thus, it must be taken into account.

In contrast to ensquared energy calculations, the effect of field-dependent aberrations is not straightforward to evaluate as it depends entirely on the characteristics of the instrument. The details on this calculation are reserved for a later chapter (see Section 6.2.2); in the meantime, it will suffice to say that we made use of End-to-End (E2E) Monte Carlo instances of the HARMONI instrument to estimate the effect of field-dependent aberrations. These models are the most realistic representation of a HARMONI as-built system available, and include alignment tolerances and surface irregularities.

The results of that analysis show that field-dependent aberrations represent approximately 15 nm in the finest 4×4 mas scale of HARMONI. Achieving an NCPA correction that is accurate up to that threshold would add in quadrature to a total of approximately 20 nm in RMS wavefront error. Based on the findings from the previous section, we can also conclude that if we reach the 15-20 nm requirement, the ensquared energy would only degrade by a maximum of 2.5%.

2.7 State-of-the-art techniques

Wavefront sensing has always been a key feature of AO systems, given that their main goal is to estimate the phase aberrations caused by atmospheric turbulence and correct them. In AO systems this wavefront estimation is typically done through a dedicated wavefront sensor. However, due to the nature of non-common path aberrations (NCPA), directly using a wavefront sensor is not feasible, as this defeats the purpose of the instrument itself, which is to deliver scientific observations, rather than measure wavefront errors.

Consequently, for the purpose of NCPA calibration special techniques are needed to estimate the wavefront directly with the science detector. This is usually known as focal plane wavefront sensing (FPWFS) since it is performed on the images taken at the focal plane. FPWFS can be very challenging as the relationship between measured focal plane image intensity and pupil plane complex amplitude is both nonlinear and degenerate (Guyon, 2018).

In recent years, stringent error budgets and the increase in design complexity, with a variety of non common optical elements between the wavefront sensor and camera arm, have created the need to perform NCPA correction even within the AO systems themselves (Lamb et al., 2014). Moreover, due to the existing expertise in the topic of wavefront sensing within the AO community, some of the most successful wavefront sensing techniques for NCPA calibration have been developed in that context.

In this section, we provide a review of wavefront sensing techniques that have been developed over the past few decades.

2.7.1 Focal Plane Sharpening

One of the simplest techniques, Focal Plane Sharpening (FPS) uses an optimisation algorithm to maximise the peak of the PSF by varying the shape of the deformable mirror (DM) over a set of aberration modes. The optimiser is typically a relatively simple and robust algorithm such as the simplex method. Conceptually simple and

easy to implement, FPS has the disadvantage of requiring a large number of images or iterations to converge to an optimum solution, especially if one needs to correct for a wide range of aberrations or modes. In addition, if the starting point is a highly aberrated system, non-linearity and the existence of secondary peaks induced by the aberrations can hinder convergence (Lamb et al., 2014).

In general, FPS is never the algorithm of choice for wavefront sensing in complex systems, given its poor performance and its lack of robustness, although it can be used to find a starting point to initialise other algorithms. A comparison between FPS and other more sophisticated techniques like Phase Diversity can be found in Lamb et al. (2017).

2.7.2 Phase Diversity

First introduced by (Fienup, 1982; Gonsalves, 1982; Paxman et al., 1992, 1994), Phase Diversity (PD) was proposed to avoid the degeneracies of phase estimation from focal plane images, by comparing at least two images taken in and out of focus. The nominal image corresponds to the system with some unknown aberrations that we want to estimate, whereas the out of focus image contains those same aberrations plus some known aberration, the *phase diversity*. The additional information coming from the defocused image ensures that the sign of symmetric aberrations are properly calculated.

PD has been used extensively to compensate for NCPA within AO systems; its first application was demonstrated on NAOS, the first AO system of the VLT, used in conjunction with the CONICA near infrared camera (Hartung et al., 2003).

The principle behind PD is relatively simple. Given a numerical model of the image formation for the instrument as a function of the aberrations, one can use numerical optimisation to minimise the quadratic distance between the observed images (focused and defocused) and the model images, and thus find the set of aberrations coefficients that best match the underlying wavefront. Several flavours of PD exist, with slight modifications to their cost function, regularization terms, etc; but the essence of it boils down to the following.

The focal plane image $i(\mathbf{r})$ can be defined as a convolution between the object $o(\mathbf{r})$ and the PSF of the optical system $h(\mathbf{r})$, in the presence of some noise $n(\mathbf{r})$:

$$i(\mathbf{r}) = o(\mathbf{r}) * h(\mathbf{r}) + n(\mathbf{r}) \quad (2.7)$$

where the PSF is defined as the modulus squared of the Fourier transform of the pupil function, which in turn depends on the pupil aperture $\Pi(\mathbf{u})$ and the wavefront aberrations $\phi(\mathbf{u})$.

$$h(\mathbf{r}) = |\mathcal{F}(\Pi(\mathbf{u}) \exp [j2\pi\phi(\mathbf{u})])|^2 \quad (2.8)$$

The *diversity* image can be defined in an equivalent manner

$$i_f(\mathbf{r}) = o(\mathbf{r}) * h_f(\mathbf{r}) + n_f(\mathbf{r}) \quad (2.9)$$

modifying the PSF of the system to account for the diversity in phase $\varphi(\mathbf{u})$:

$$h_f(\mathbf{r}) = |\mathcal{F}(P(\mathbf{u}) \exp [j2\pi\phi(\mathbf{u}) + \varphi(\mathbf{u})])|^2 \quad (2.10)$$

If information on the noise statistics is available, it can be incorporated into the formulation such that PD becomes a Maximum Likelihood Estimation (MLE) problem where the solution is estimated as having maximum probability, given the data (observed images) and the noise. The criterion to optimise is given by the squares of the residual differences between the model images and the observed images:

$$J(\phi) = \frac{1}{2} \left\| \frac{i_{foc} - i(\mathbf{r})}{\sigma_n} \right\|^2 + \frac{1}{2} \left\| \frac{i_{div} - i_f(\mathbf{r})}{\sigma_n^f} \right\|^2 + \mathcal{R}(\phi) \quad (2.11)$$

the term $\mathcal{R}(\phi)$ represents a possible regularisation term, which can take into account our prior beliefs regarding the unknown phase. When this prior is incorporated into the estimation in a Bayesian setting, the problem takes the form of a Maximum A Posteriori (MAP) estimation, similar to the original MLE except for the regularisation.

It should be noted that the choice of defocus as the *diversity* is arbitrary, and any aberration can be used provided that it can be easily introduced in the system.

Using information theory [Dolne and Schall \(2005\)](#) investigates, for various diversity functions, the Cramer-Rao bound on the estimated coefficients, which represents a lower bound on the variance of unbiased estimators of deterministic parameters. After studying the sensitivity of the estimated parameters for each of the diversity polynomials for point and extended sources, [Dolne and Schall \(2005\)](#) concludes that:

- The performance for extended sources is worse than for a point source; because the convolution with the extended source reduces the sensitivity on the aberrations

- The phase-diversity algorithm is efficient for most of the diversity polynomials considered, which included defocus, astigmatism, coma, spherical and other high order aberrations.

In addition, they point out that although defocus seems to perform better the majority of times, other polynomials may be more appropriate in certain cases; therefore one must examine the expected aberrations of the optical system and determine the appropriate diversity phase function that minimises the residual wavefront error. One example of this is [Sauvage et al. \(2012\)](#) where they show that a mixed diversity phase composed of astigmatism and focus can improve the accuracy compared to traditional defocus diversity, by enhancing the convexity of the error metric in a larger area around the minimum. Similarly, the LIFT sensor (Linearized Focal-plane Technique), proposed by [Meimon et al. \(2010\)](#), uses a known phase offset in the form of astigmatism. The LIFT sensor was developed to measure low order aberrations (tip / tilt and defocus) from faint Natural Guide Stars (NGS) in conjunction with Laser Tomographic and Multi-Conjugate Adaptive Optics systems, showing low noise sensitivity compared to a Shack-Hartmann sensor ([Plantet et al., 2013](#)). The choice of astigmatism as phase offset over a more traditional defocus is justified as it improves the shape of the minimisation criterion of the phase estimation, leading to a more robust focus estimation, ([Plantet et al., 2013](#)).

Most PD approaches operate with a nominal and out of focus image, and values between 0.5λ to 1λ peak-to-valley of defocus strength are commonly used [Zhang et al. \(2017\)](#). A recent study by [Zhang et al. \(2017\)](#) explores an alternative approach known as bisymmetric defocuses phase diversity (PD-BD) that uses images taken at two symmetrical defocused planes, instead of the nominal + defocus case. Results seem to indicate that this approach has the potential for better wavefront sensing performance. [Zhang et al. \(2017\)](#) suggests that the reason behind this improvement could be the shape of the error metric, which in the case of PD-BD exhibits steeper gradients around the global minimum, a desirable feature for numerical optimisation.

[Sauvage et al. \(2012\)](#) present an extension of PD to coronagraphic imaging, called coronagraphic focal plane wavefront estimator for exoplanet imaging (COFFEE) to calibrate the aberrations upstream of the focal plane mask, which are known to cause intensity residuals in the images. In the coronagraphic case, it is important to distinguish between upstream and downstream aberrations. For instance, tip/tilt errors upstream affect the centering of the star on the coronagraph and thus have a major impact on contrast performance, whereas tip/tilt errors downstream mostly affect the

positioning of the PSF on the detector plane. As [Sauvage et al. \(2012\)](#) points out, upstream aberrations have a bigger impact on coronagraphic imaging, but the existence of uncorrected downstream aberrations can distort the estimation of upstream aberrations and thus they should not be neglected.

Follow-up studies by [Paul et al. \(2013\)](#) provide a thorough characterisation of COFFEE’s performance by means of numerical simulations, as well as experimental validation of its ability to compensate for the aberrations upstream of a coronagraph.

With the advent of extreme adaptive optics (XAO), the increase in number of actuators imposes severe computational challenges on wavefront sensing techniques. In an attempt to mitigate this issue, [Keller et al. \(2012\)](#) presents a modified PD approach with reduced computational complexity that can handle around 10,000 to 100,000 actuators. This technique, called Fast & Furious (FF) relies on several assumptions: point source, symmetric pupil aperture and monochromatic light. In addition, it exploits mathematical simplifications applicable in the weak aberration regime and uses a specific diversity that reduces computational effort. Although it was initially developed for real-time correction of aberrations in AO systems, this technique is also applicable to NCPA calibration, as its first experimental validation showed ([Korkiakoski et al., 2012](#)).

To overcome some of the limitations of FF, [Korkiakoski et al. \(2014\)](#) propose an extension to the method with a Gerchberg–Saxton (GS)-style error reduction to determine the pupil amplitudes, eliminating the assumption of the pupil amplitudes being even and allowing it to deal with arbitrary pupil shapes.

Typically, the accuracy of PD is closely linked to the optimisation algorithm selected and the characteristics of the physical model ([Zhang et al., 2017](#)). With regards to the algorithm, some gradient-based optimisation algorithms risk getting trapped in local minima, with global optimisation alternatives offering more robustness. For instance, [Zhang et al. \(2016\)](#) propose a two-step hybrid global optimisation algorithm for PD that combines evolutionary particle swarm optimisation to find a rough estimate of the global minimum, with a standard Broyden-Fletcher-Goldfarb-Shanno algorithm. [Qi et al. \(2018\)](#) further improves this approach by introducing a criterion to detect premature convergence followed by a redistribution mechanism to ensure the uniform distribution and randomness of the particles in the search space.

The other factor that normally limits accuracy is the physical model, which is inherently prone to errors, parameter inaccuracies and imperfect knowledge of the system in general. [Blanc et al. \(2003\)](#) investigates the influence of these errors on the wavefront estimation with numerical and experimental data, with particular attention

to uncertainties in the defocus diversity, pupil shape and misalignment, pixel scale errors, etc. A similar study by [Jolissaint et al. \(2012\)](#) confirms that the wavefront accuracy can be very sensitive to errors in the knowledge of the system’s optical parameters, constructs an error budget, and estimates the accuracy margins for the different parameters. Consequently, applying PD in a real system requires thorough calibration to construct a reliable image formation model, as the work by [Gratadour et al. \(2013\)](#) demonstrates for the CANARY Multi-Object Adaptive Optics (MOAO) system. In their work, they show how critical it is to properly model the pupil function, including the true position of spiders and obscurations, as well as variations in illumination, if one is to reach good NCPA correction on-sky.

2.7.3 ZELDA

Another technique worth mentioning is ZELDA (Zernike sensor for Extremely accurate measurements of Low-level Differential Aberrations), which uses a Zernike phase contrast method to calibrate NCPA seen by the coronagraph in exoplanet direct imagers such as VLT/SPHERE ([N’Diaye et al., 2013](#)). Its working principle is rather simple: the coronagraph mask is replaced by a circular phase mask that introduces a $\pi/2$ delay inside the central region of diameter λ/D of the point source image. The diffraction off this zone creates a reference wave in the exit pupil, giving rise to an interference pattern in the exit pupil that represents, in a quasi-linear fashion, the phase map of the wave incident on the coronagraph plane ([Dohlen et al., 2013](#)).

As long as the differential aberrations are sufficiently small in the exoplanet imager, a linear or quadratic relation exists between the wavefront and the modulated pupil intensity, which enables the reconstruction of these aberrations to nanometric accuracy ([N’Diaye et al., 2016](#)).

The ZELDA sensor and its associated algorithm have been tested on SPHERE ([N’Diaye et al., 2018](#)), including on-sky validation ([Vigan et al., 2019](#)). Given its success, it will also be implemented in HARMONI for the high-contrast mode (HCAO) to monitor and correct NCPA ([Carlotti et al., 2018](#)).

2.7.4 Electric Field Conjugation

Electric Field Conjugation (EFC), sometimes known as speckle nulling, is a method for removing quasi-static speckles in a predefined area of the detector in a coronagraphic image using a deformable mirror ([Ruffio, 2014](#)). This method uses a typical diversity methodology (similar to Phase Diversity), and it is based on the idea of

pairwise probing, each probe consists of a positive and negative change to the DM. Under the assumption of small aberrations, the relationship between the difference of a pair of images for a given probe becomes linear on the electric field. This means that one can use this to obtain an estimate of the complex electric field, and then identify the actuator commands that will minimise the speckle intensity in a given area of the detector plane.

EFC cannot correct simultaneously for phase and amplitude errors over the whole image; instead, the correction is typically limited to a small area at one side of the PSF centre, known as the dark hole. This means that EFC is more effective for follow-up studies than for exoplanet discovery, since the area of enhanced contrast (the dark hole) is of limited size and thus the position of the planet needs to be known in advance to properly define the dark hole. EFC has been shown to increase the raw contrast ratio for SPHERE by one order of magnitude (Ruffio, 2014). Laboratory results for pairwise probing and EFC are presented in (Potier et al., 2019).

2.7.5 Self Coherent Camera

Another technique developed for wavefront sensing in coronagraphic systems is the Self-Coherent Camera (SCC). Without introducing any additional NCPA, the SCC uses spatial modulation in the science image to encode the stellar speckles that are directly linked to the wavefront aberrations (Galicher et al., 2010). In order to achieve this, the SCC uses a pinhole next to the Lyot stop of a coronagraphic system to create a reference beam that interferes with the image channel and spatially modulates the stellar speckles in the science image recorded by the detector (Galicher et al., 2019), see Fig. 2.12 for reference.

The SCC has successfully been tested on-sky at the Hale 200 inch telescope at Palomar. One of the main advantages of SCC is its observing efficiency, as each image can be used for both science and NCPA calibration and provides real-time compensation.

In long-exposure coronagraphic images we typically see three features: a smooth halo coming from averaging out post-AO residual speckles that vary rapidly, static speckles that remain constant for all images in the sequence, and quasi-static speckles with a lifetime longer than the exposure time, but shorter than the total observing time. Singh et al. (2019) demonstrates the use of the SCC to correct those quasi-static speckles that vary on timescales of a few images and thus cannot be corrected during post-processing.

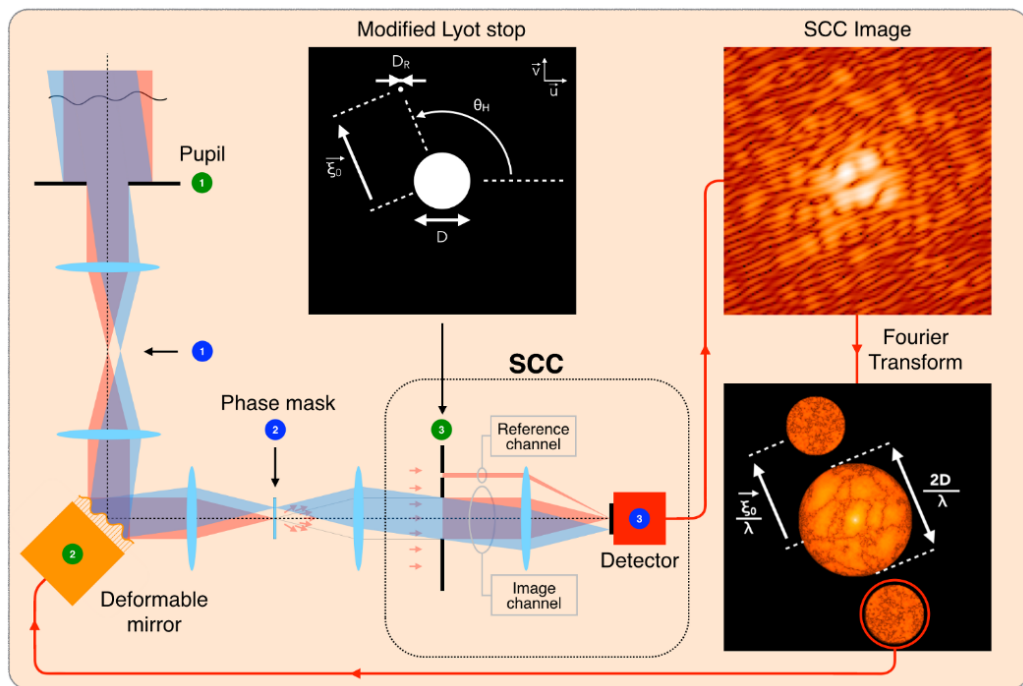


Figure 2.12: Principle of the self-coherent camera (SCC). Originally published in Galicher et al. (2019). The *red* beam represents the stellar light that hits the deformable mirror and goes through a focal plane coronagraphic mask and Lyot stop. In the presence of aberrations, part of the starlight is not entirely suppressed and induces speckles at the detector plane. Using part of the starlight rejected by the focal plane mask (the reference channel), one can spatially modulate the speckle pattern (see top right). The Fourier transform of the SCC image allows direct measurement of the speckle electric field. The *blue* beam represents the planet light, which is not affected by the coronagraph and is not fringed.

Conventional SCC techniques such as the ones from [Galicher et al. \(2019\)](#) and [Singh et al. \(2019\)](#) require relatively long exposures, because of the limitation on the amount of light passing through the pinhole in the Lyot plane. [Gerard et al. \(2018\)](#) presents a redesign of the SCC focal plane mask (FPM) to get enough light through the pinhole to allow measuring and correcting for speckles on all time scales, including millisecond-lifetime atmospheric speckles. This new framework, known as Fast Atmospheric SCC Technique (FAST) records images every few milliseconds and subtracts atmospheric and static speckles without reducing the exoplanet throughput. In principle, this technique removes the intrinsic limitation on raw contrast due to atmospheric and static speckles, allowing for continuous improvement simply by integrating longer ([Gerard et al., 2018](#)). Preliminary laboratory tests for FAST are presented in [Gerard et al. \(2019\)](#) with promising deep contrast gains of more than two orders of magnitude over existing instruments.

A fair and meaningful comparison of the myriad of wavefront sensing and control techniques available is challenging, since most references only deal with one particular approach and the results are applicable only to very different regimes, instrument design constraints, or laboratory conditions. However, some efforts have been made to test multiple techniques under comparable conditions. [Potier et al. \(2020\)](#) compares SCC to a combination of Pairwise probing (PW) and Electric Field Conjugation (EFC) on the same laboratory testbed. In contrast with SCC, which relies on spatial modulation of the speckle intensity, PW uses temporal modulation to retrieve the electric field by combining images recorded after introducing known aberrations (called probes), each probe phase needing a pair of images. The results in [Potier et al. \(2020\)](#) found no significant difference in performance, both techniques being capable of estimating the aberrations to subnanometer accuracy.

One major disadvantage of the SCC technique is that the reference pinhole must be placed at a minimum distance of at least 1.5 times the pupil diameter, which can be difficult to accommodate in existing instruments given the size of the optics, and reduces the sensitivity because of the small amount of light reaching that area. Recently, [Bos \(2021\)](#) proposed a modification of SCC that uses a polariser at the pinhole followed by a polarising beamsplitter downstream of the Lyot stop that relaxes that requirement, allowing the pinhole to be placed closer to the pupil. This Polarisation-Encoded SCC (PESCC) alternative can be used for wavefront sensing in the context of coronagraphic observations.

2.7.6 Coronagraphic Low Order Wavefront Sensor

Another approach to measuring low-order aberrations in coronagraphic systems was proposed by [Guyon et al. \(2009\)](#) based on the idea of reimaging part of the light blocked by the focal plane mask. It uses a special focal plane mask with a central opaque circle, surrounded by a reflective ring that sends part of the light towards an off-axis parabola, which then creates a defocused image. The defocus is important because it allows unambiguous estimation of aberrations (just like PD). CLOWFS relies on the fact that for sufficiently small aberrations, the recorded image is a linear function of the aberration modes.

CLOWFS requires prior calibration to measure the modal responses of the sensor, by deliberately introducing aberrations in the system and then using such responses to decompose the reference image into a linear sum of aberrations. One of the main advantages of this method is that plenty of starlight is available so that the measurements are very sensitive. Nevertheless, this technique is very hardware dependent as it is limited to a narrow domain of Phase-Induced Amplitude Apodisation coronagraphs (PIAA), and it is not applicable to other coronagraphic systems that rely on different techniques for starlight suppression.

2.7.7 Lyot-based Low Order Wavefront Sensor

To circumvent the major drawback of CLOWFS, [Singh et al. \(2014\)](#) proposes an alternative solution that works with diffraction-based Phase Mask Coronagraphs (PMC). This technique known as Lyot-based Low Order Wavefront Sensor (LLOWFS) uses a reflective Lyot stop to reimage the unused starlight after the focal plane mask, in order to sense tip-tilt and defocus aberrations. Similar to CLOWFS, LLOWFS needs to be pre-calibrated by applying a phase map with a controlled amount of tip-tilt and other modes in ideal conditions of no aberration before deploying it in realistic conditions with aberrations.

One of the advantages of LLOWFS is that it is highly photon-efficient as it uses the diffracted starlight that would otherwise be wasted. However, it comes with some of the limitations that CLOWFS already had as it is limited to pointing errors and low order aberrations, arguably because these represent the major source of starlight leakage in coronagraphic systems.

2.7.8 Differential Optical Transfer Function

Another technique called differential optical transfer function (dOTF) uses diversity in the pupil plane, typically in the form of a small obscuration at the edge of the pupil, to generate two images. The difference in optical transfer function between these two images creates a region of two overlapping lobes where the complex pupil field can be probed for wavefront sensing (Codona, 2013). The achievable resolution of this method is limited by the size and shape of the pupil obstruction since this determines the convolution kernel that blurs the original complex amplitudes.

Some of the advantages of this technique include: it is non-iterative (a single step provides the wavefront estimation), it only relies on a small modification of the pupil and it makes no assumption on the strength of the aberrations.

One of the main limitations of dOTF is obtaining a high SNR, since two almost identical images are subtracted, and those images are the focal plane images of a point-source for which most of the light is localized in the core of the PSF (Korkiakoski et al., 2013). In order to overcome this challenge, Korkiakoski et al. (2013) suggests applying a strong defocus to spread the light evenly across the detector, plus a combination of images at varying integration times to enhance dynamic range and averaging across multiple images to further reduce noise. These efforts to reduce noise directly increase the measurement burden, which can be several minutes per dOTF measurement, long enough for the stability of the system to be questioned.

Another point to consider is the feasibility of introducing such pupil diversity. In Codona and Doble (2015) dOTF is proposed as a supplementary or backup approach to segment phasing and NCPA calibration in JWST, using as pupil diversity either a small displacement of a single segment or a small misalignment of the filter wheel. In principle, relying on available hardware is a great advantage of dOTF but it also limits its applicability in the general case, as it is not clear whether suitable hardware would be present in all system designs. Moreover, using available hardware imposes constraints on the achievable resolution, and general performance of dOTF (as it has been pointed out already), meaning that one could be forced to operate in non-optimum conditions. In the case of HARMONI, there is a pupil stop wheel inside the calibration unit that could be used for this purpose.

2.7.9 Discussion

These techniques are not necessarily exclusive. For instance dOTF Codona (2013) could be used to obtain an estimate of the wavefront to serve as starting guess for

other iterative techniques like Phase Diversity. Another example is [Lamb et al. \(2017\)](#) who suggest that the number of iterations and images needed for FPS to converge could be greatly reduced if using a starting point from Phase Diversity; and that Phase Diversity would benefit from an starting point estimated via FPS due to its model independence.

Another interesting example of synergy is shown in [Korkiakoski et al. \(2013\)](#) and [Korkiakoski et al. \(2014\)](#): they use a spatial light modulator (SLM) as the wavefront correction device in the context of the Fast & Furious technique. In this case, it is imperative to calibrate the SLM phase and transmission response to control signals, and use that information to map the location of each SLM pixel to its physical location in the pupil plane. This calibration is done using the dOTF technique.

The latest developments in NCPA tend to be focused in the regime of high-contrast imaging. This is not surprising given the fact that NCPA have been identified as one of the main limiting factors in achievable contrast and the increasing popularity of the quest for detection and characterisation of planets in the habitable zone. Current efforts such as the ZELDA sensor and the Self Coherent Camera are being dedicated to bridging this technical challenge to enable full exploration of a completely new parameter space of high contrast imaging at small angular separation.

However, as we will elaborate later, there has been barely any research regarding the calibration of NCPA in systems that contain **image slicers** such as HARMONI. It is within this context that this thesis makes its novel contribution.

3

Image Slicers

”Every revolutionary idea seems to evoke three stages of reaction. They may be summed up by the phrases: (1) It’s completely impossible. (2) It’s possible, but it’s not worth doing. (3) I said it was a good idea all along.”

– Arthur C. Clarke

3.1 Integral field spectroscopy

Integral field spectroscopy (IFS) allows us to obtain a spectrum at every pixel position of a 2D image in the sky. The principle behind IFS is to reformat the 2D input field so that it can be fed through a conventional spectrograph and thus register the spectrum of every pixel in a single exposure of the detector. This is done by means of an Integral Field Unit (IFU), which in the case of HARMONI contains: an **image slicer**, a stack of tilted thin mirrors that slices the field and sends light from each slice in different directions, followed by a set of pupil mirrors that re-image each slice to form a pseudo-slit composed of individual slitlets. [Eisenhauer and Raab \(2015\)](#) gives an overview of the concepts of visible and IR imaging spectroscopy, focusing on integral field spectroscopy, and the advantages of IFUs over other approaches.

This simultaneous measurement of spatial and spectral information that IFS provides is extremely advantageous and has revolutionised astrophysics over the past few decades. As an example, [Cappellari \(2016\)](#) presents an extensive review on the progress of our knowledge of the structure of early-type galaxies and galaxies in general, thanks to recent advances in IFS. In addition, some of the main science goals of ESO’s ELT will only be achievable thanks to IFS, including the study of resolved stellar populations of elliptical galaxies, which is expected to shed light on the formation mechanisms and evolution of galaxies ([ESO, 2011](#)).

3.2 Calibrating NCPA in HARMONI

Given the importance and popularity of IFS in astrophysics, and the challenging nature of the science cases ESO’s ELT will pursue, successful calibration of NCPA within HARMONI will be essential to ensuring that the instrument delivers its best performance. However, to the best of the author’s knowledge, NCPA calibration in the context of an instrument containing an IFU has never been demonstrated.

From our review of state-of-the-art techniques in Section 2.7 we can identify two major trends with regards to wavefront sensing. First of all, a great deal of effort has been devoted to NCPA calibration in the context of **AO systems**, specially by means of PD whose popularity is quite evident (Lamb et al., 2017; Sauvage et al., 2012, 2007).

Secondly, wavefront sensing and correction in the context of **coronagraphic systems** has gained a lot of attention over the last decade, primarily fuelled by the interest in exoplanet direct imaging and the daunting technical challenges it entails Subedi et al. (2015), Wilby et al. (2017), to cite some examples.

Plenty of studies have explored the propagation of aberrations through a coronagraphic system, developed an understanding of speckle formation, and proposed specific techniques for NCPA calibration in that context (Galicher et al., 2010; Sauvage et al., 2012; Sivaramakrishnan et al., 2008; Soummer et al., 2007).

In contrast, NCPA calibration in IFUs remains an unexplored area. Then, **how is HARMONI going to calibrate Non Common Path Aberrations?** The AO Calibration Unit Design Analysis document (HRM-00583) submitted to the Critical Design Review of the HARMONI Calibration Module states that:

*”it has been determined from past experience that classical methods like phase diversity are not efficient in the case of image slicing spectrographs. It is therefore necessary to proceed by **iterative image optimization using a deformable mirror.**”*

In other words, in the absence of a better alternative, the baseline approach to NCPA calibration for IFUs relies on some form of Focal Plane Sharpening. This method loops over each aberration mode, trying different DM corrections and finding the values that maximise some metric of performance (Strehl ratio, FWHM of a Gaussian fit of the PSF, etc). This is extremely inefficient and time-consuming, as it requires looping over a large number of aberrations trying multiple DM corrections for each case to see how the metric varies.




Figure 3.1: Impact of aberrations on Strehl ratio. *Left*: $\text{PSF}(\Phi)$ with random Zernike aberrations. *Right*: PSF Strehl ratio for the same random wavefront plus astigmatism and trefoil $\text{PSF}(\Phi + aZ_i + bZ_j)$, as a function of aberration intensity. Wavelength: $1.5 \mu\text{m}$, 4 mas spaxels.

In addition, even though the Zernike polynomials are orthogonal in pupil space, they are not in image space. In fact, due to the non-linearity of image formation, metrics such as Strehl ratio can exhibit a landscape with multiple local maxima for some combinations of aberrations. To illustrate this, we ran the following experiment: we generated PSF images with random wavefronts in the form of Zernike polynomials. Then, we calculated how the Strehl ratio varies when we modify the intensity of a given pair of Zernike aberrations Z_i vs Z_j . For example, Fig. 3.1 shows the PSF for a given random wavefront (without astigmatism and trefoil), and the evolution of Strehl ratio with varying intensities of that pair of aberrations. Fig. 3.2 shows an equivalent example with astigmatism and coma. Here the Strehl ratio is calculated as the maximum value of the image pixels, which is not as robust as fitting a Gaussian profile to the PSF and evaluating its peak, but it is faster to compute.

What we found is that spurious peaks in Strehl ratio can be observed with combinations of aberrations that are far from zero. This complicates the task of FPS, as we can encounter peaks in Strehl that correspond to undesired aberrated states. Consequently, when running FPS, it is recommended to repeat the loops over the aberrations a couple of times, to ensure proper convergence.

This calls for more efficient alternative approaches to NCPA calibration that take into account the specificity of image slicers. As stated in the Section 2.4.2, the objective of this DPhil project is to develop **techniques for measuring and calibration NCPA in adaptive optics assisted spectrographs containing image slicers**.




Figure 3.2: *Left*: PSF(Φ) with random Zernike aberrations. *Right*: PSF Strehl ratio for the same random wavefront plus astigmatism and coma PSF($\Phi + aZ_i + bZ_j$), as a function of aberration intensity. Wavelength: $1.5 \mu\text{m}$, 4 mas spaxels.

Therefore, before deciding which technique to use, we must dedicate some time to characterise the impact of image slicers on light propagation and image formation, because this could have important consequences on the applicability of some techniques. When it comes to image slicers effects, two main concerns were identified at the beginning of this project:

1. **Beam footprint clipping.** An important consideration regarding the applicability of traditional techniques for NCPA calibration in the context of image slicers spectrographs is the uncertain interaction between the defocus needed for Phase Diversity and the image slicer. Displacing the slicer out of focus gradually increases the size of the geometric beam footprint. Given the small size of the slicer stack mirrors, there could be a point where the light starts to fall onto neighbouring slices and follow unintentional paths.

Considering that longitudinal defocus Δ , i.e. the displacement due to the defocus in the wavefront is given by $\Delta = 8PF^2$, where P is the peak-to-valley wavefront and F is the system focal ratio f/D ; displacing the focus by Δ increases the size of the geometric footprint to r , where $NA = \sin \alpha = 1/(2F)$ and $\sin \alpha \sim \alpha = r/\Delta$. Combining these expressions we arrive at $D = 2r = 8PF$ so that the size of the geometric beam in out of focus conditions is proportional to the peak-to-valley defocus P and the focal ratio. Given that the size of a HARMONI slicer mirror is 1 mm, we can calculate the maximum peak-to-valley $P = D/8F$ above which the size of the geometric beam footprint is larger than one slice, which corresponds to $P \simeq \pm 0.10\lambda$ for $F = 1360$.

Consequently, for moderate values of defocus above $\pm 0.10\lambda$ commonly used in phase diversity, the beam footprint will cover more than one slice. Depending on the characteristics of the image slicer, this could potentially cause light loss or other shadowing effects that alter the PSF.

2. **Diffraction effects.** The other concern is that due to the size of the slicing mirrors, significant diffraction will take place at the image slicer. How this affects the PSF will be thoroughly discussed in the following section.

3.3 Light propagation through Image Slicers

The first step towards developing a calibration technique that can be deployed in image slicer spectrographs is understanding how the presence of the IFU affects the propagation of light, and ultimately modifies the PSF.

The IFU is a complex optical sub-system that contains multiple optical elements, and modelling the propagation of light through it can be challenging. Nevertheless, here we will use a simplified model that captures most of the functions the IFU is expected to perform.

3.3.1 Mathematical framework

We will assume the light propagation can be modelled with the following 4 planes: input pupil plane, slicer plane, pupil mirror plane, and exit slit plane as shown in Fig. 3.4. The apertures for each plane have been sketched as well.

Starting at the input pupil with an arbitrary aperture Π_p and wavefront $\Phi(\mathbf{u})$, a series of optical elements focus the light onto the image slicer. The image slicer is composed of individual slices, each reflecting a section of the image towards a pupil mirror. The aperture for each slice k can be modelled as an elongated rectangular aperture Π_s^k that masks the electric field at the slicer plane. After the image slicer, each pupil mirror receives the reflected / diffracted light coming from each slice and reimages it back to the exit slit plane to form a set of slitlets. It should be noted that each pupil mirror has a finite aperture Π_m^k that can potentially crop part of the light coming from the slicer.

We can mathematically model the propagation of some unknown wavefront $\Phi(\mathbf{u})$ through the optical system, starting with the complex pupil function $P(\mathbf{u})$:

$$P(\mathbf{u}) = \Pi_p \exp i\Phi(\mathbf{u}) \quad (3.1)$$




Figure 3.3: Optical model for the image slicer within the SWIFT instrument. The image slicer splits the field of view and sends each slice to a pupil mirror / lenslet pair that image it into an exit slit

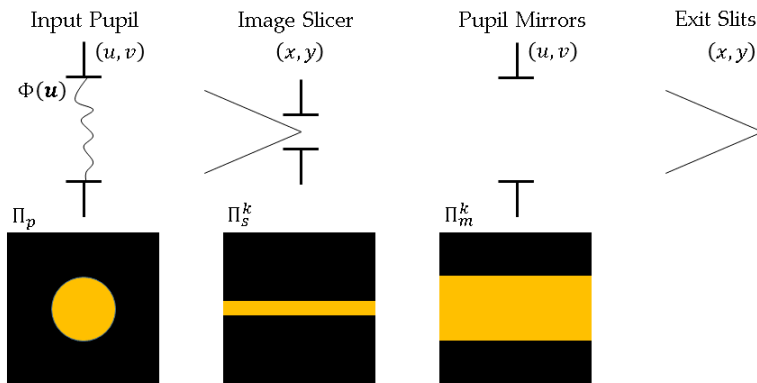


Figure 3.4: Sketch of an Integral Field Unit sub-system with its relevant surfaces and apertures. From left to right: input pupil plane, image slicer plane (made out of adjacent slices with thin apertures), pupil mirror plane (one mirror for each slice of the image slicer) and exit slits.

where Π_p represents the pupil aperture, $\Phi(\mathbf{u})$ is the wavefront at the pupil plane. The electric field at the slicer plane is given by:

$$E_s^k = \mathcal{F}[\Pi_p \exp i\Phi(\mathbf{u})](\mathbf{x}) \quad (3.2)$$

where E_s^k represent the electric field at the slicer plane s for an arbitrary slice k and $\mathcal{F}[\cdot]$ is the Fourier transform. If the slicer plane was the final focal plane, the PSF would be given by:

$$\text{PSF} = \Pi_s^k \|\mathcal{F}[\Pi_p \exp i\Phi(\mathbf{u})]\|^2 \quad (3.3)$$

with Π_s^k representing the masking effect of the slice. However, after the slicer, the electric field is propagated up to the pupil mirror plane m by means of:

$$E_m^k = \mathcal{F}^{-1}[\Pi_s^k \cdot \mathcal{F}[\Pi_p \exp i\Phi(\mathbf{u})]] \quad (3.4)$$

where E_m^k is the electric field at the k -th pupil mirror plane. Using the Convolution Theorem $f * g = \mathcal{F}^{-1}[\mathcal{F}(f) \cdot \mathcal{F}(g)]$ we can write:

$$E_m^k = \mathcal{F}^{-1}[\Pi_s^k] * P(\mathbf{u}) \quad (3.5)$$

so that the electric field at the pupil mirror is given by the convolution between the original pupil function $P(\mathbf{u})$ and the inverse Fourier transform of the slicer mask Π_s^k .

The pupil mirror has its own finite aperture, which adds masking / cropping before propagation to the final exit slit focal plane f :

$$E_f^k = \mathcal{F}[\Pi_m^k \cdot (\mathcal{F}^{-1}[\Pi_s^k] * P(\mathbf{u}))] \quad (3.6)$$

Once again using the Convolution theorem $\mathcal{F}[f \cdot g] = \mathcal{F}(f) * \mathcal{F}(g)$ we can write:

$$E_f^k = \mathcal{F}[\Pi_m^k] * [\Pi_s^k \cdot \mathcal{F}(P(\mathbf{u}))] \quad (3.7)$$

which shows that the final electric field is influenced by two terms: the masking due to the slicer apertures Π_s^k , and the effect of having a finite aperture on the pupil mirrors. In the absence of pupil mirror apertures, the electric field E_f^k would simply be masked by the slicer aperture. However, we need to convolve the electric field with the Fourier transform of the pupil mirror aperture $\mathcal{F}[\Pi_m^k]$. Starting with the definition of the Fourier transform for an arbitrary function $f(u, v)$:

$$\mathcal{F}[f(u, v)] = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(u, v) e^{-iux - ivy} dudv \quad (3.8)$$

Assuming the aperture is just a rectangular aperture along the direction perpendicular to the slices (the one dominated by diffraction) extending up to an arbitrary value $|\bar{v}| = \pi$, such that $f(u, v) = 1, \forall (u, v) \in \{|\bar{v}| \leq \pi\}$ we can write:

$$\mathcal{F}[f(u, v)] = \int_{-\infty}^{\infty} \int_{-\pi}^{\pi} 1 \cdot e^{-iux-ivy} dudv \quad (3.9)$$

it can be shown that this is zero everywhere except for the case of $u = 0$ for which we have:

$$\mathcal{F}[f(u, v)] = \int_{-\infty}^{\infty} \int_{-\pi}^{\pi} e^{-ivy} dudv = C \int_{-\pi}^{\pi} e^{-ivy} dv = \frac{-1}{iy} [e^{-ivy}]_{-\pi}^{\pi} \quad (3.10)$$

which can be simplified to the cardinal sine $\text{sinc}(\cdot)$:

$$\mathcal{F}[f(u, v)] = \frac{-1}{iy} [e^{-ivy}]_{-\pi}^{\pi} = 2\pi \frac{e^{iy\pi} - e^{-iy\pi}}{2iy\pi} = 2\pi \frac{\sin \pi y}{\pi y} = 2\pi \text{sinc}(y) \quad (3.11)$$

This means that the term $\mathcal{F}[\Pi_m^k](x, y)$ corresponds to an empty 2D array except for the line $x = 0$ where it varies as a cardinal sine. Consequently the 2D convolution $\mathcal{F}[\Pi_m^k] * (\Pi_s^k \mathcal{F}[P])$ only takes into account the y direction. To illustrate the effect of the pupil aperture on the final electric field at the exit slit, let us consider the following example. Assuming the electric field at the slicer plane could be represented by a top hat function (the effect of the slicer aperture, but disregarding the diffraction effects that create the PSF), we can calculate the 1D convolution of cardinal sine function of varying spatial frequency k , which represent different pupil mirror aperture sizes:

$$E_f \sim \text{sinc}(ky) * H(y) \quad (3.12)$$

The results of this approximation are shown in Fig. 3.5. The top row represents $\text{sinc}(ky)$ cardinal sines of varying spatial frequency k , while the lower row represents the top hat function $H(y)$ and the result of their mutual convolution E_f . One should note that the pupil mirror finite aperture introduces fringes on the top hat function; in fact, the number of peaks is linked to the spatial frequency k . For a real case where $H(y)$ is substituted by $\Pi_s^k \mathcal{F}[P(\mathbf{u})]$ the slicer-masked focal plane electric field, an analogous result holds: the finite aperture size of the pupil mirrors effectively crops the diffracted light after the slicer, with the main consequence of introducing fringes in the final PSF that vary in number with the relative size of the pupil mirror aperture and the diffracted PSF.

The mathematical framework developed above allows us to simulate the PSF after propagating through an IFU by looping over each slice / pupil mirror pair and




Figure 3.5: Approximation of the effect of the pupil mirror aperture. *Top row*: cardinal sine with varying spatial frequencies, representing the Fourier transform of the pupil mirror aperture. *Lower row*: Top hat function and its convolution with the cardinal sine. The number of peaks is proportional to the spatial frequency of the cardinal sine.

stacking the results into an exit slit PSF. To further illustrate the influence of model parameters such as the pupil mirror aperture size, or the wavelength, we ran the following analysis.

The upper row in Fig. 3.6 shows the image, given by the norm square of the electric field, at the pupil mirror associated with the central slice, for varying aperture sizes. We have chosen the aperture size such that the pupil mirror fits an integer number of zeros of the pupil image (in this case 2 and 3 zeros, respectively). The lower row shows the associated exit slit PSF. As already proven, the number of fringes on the PSF is related to the relative size of the pupil mirror aperture and the diffracted light coming from the image slicer.

Due to the growth of the PSF with wavelength, longer wavelengths are more affected by diffraction at the image slicer. In addition, for a given pupil mirror aperture, proportionally fewer PSF zeros can fit within that size for longer wavelengths. This suggests that the image slicer effects such as the fringes are wavelength dependent. To further illustrate this effect we computed the exit slit PSF for the central slice as a function of both wavelength and pupil mirror aperture size (in number of zeros of the $1.5 \mu\text{m}$ PSF). The results are shown in grid form in Fig. 3.7. Starting with the

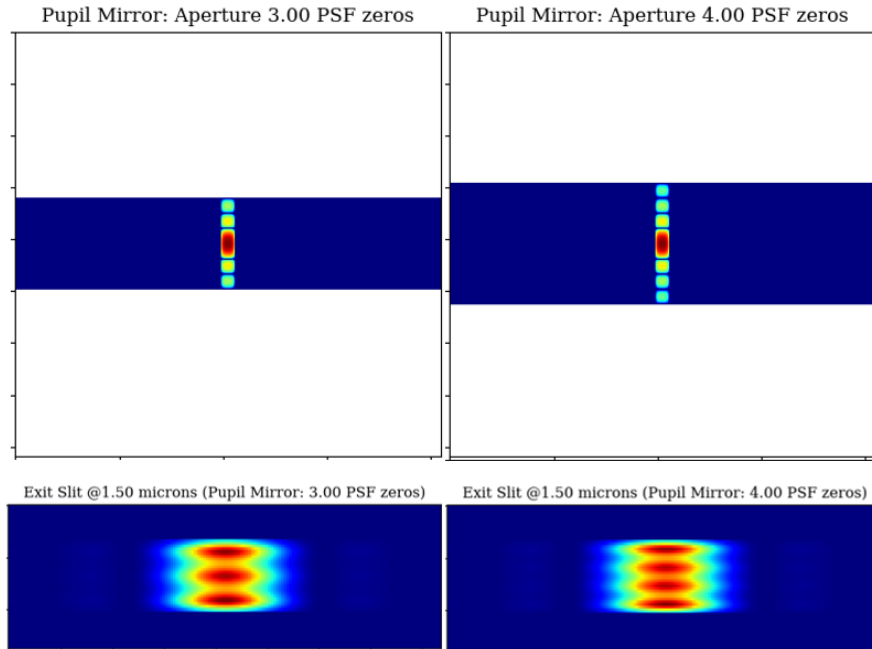


Figure 3.6: Impact of pupil mirror aperture on the exit slit PSF fringes. The number of PSF zeros at each side of the post-slicer diffracted pattern that fit within the pupil mirror aperture dictate the spatial frequency of the fringes.

bottom left corner, the pupil mirror aperture size accommodates 2 zeros of the $1.5 \mu\text{m}$ diffracted PSF. Consequently, 2 fringes are visible in the exit slit PSF at $1.5 \mu\text{m}$, but due to wavelength scaling, only 1 fringe is visible at $3.0 \mu\text{m}$. If we increase the aperture size so that 4 zeros of the $1.5 \mu\text{m}$ PSF fit inside the pupil mirror aperture (see right column), 4 fringes are visible in that exit slit for that wavelength, and 2 fringes are visible at $3.0 \mu\text{m}$.

It should be noted that this comparison between 1.5 and $3.0 \mu\text{m}$ is only qualitative on the number of PSF fringes. The situation is not entirely equivalent because the size of the slices remains constant. Therefore, at longer wavelengths, the relative size of the slicer aperture decreases leading to stronger diffraction. This why it is important to oversize the pupil mirrors to capture the diffracted light after the slicer.

3.3.2 Model validation

The mathematical model we have presented in this section constitutes a simplified version of an IFU that attempts to capture most of its effects, but it comes with its limitations. Therefore, we decided to validate these simulations using an independent tool. For this purpose we used Zemax Optics Studio in its Physical Optics Propagation (POP) mode. Using diffraction calculations, POP propagates a wavefront




Figure 3.7: Exit slit PSF (central slice) as a function of wavelength and pupil mirror aperture size (in terms of number of zeros for the $1.5 \mu\text{m}$ PSF).

through an optical system, accounting for the coherent nature of light. A beam can be defined at the entrance of the system, and its associated complex electric field can be traced, surface by surface, with POP automatically choosing the best propagation algorithm for each diffraction condition. Thanks to this, POP is capable of capturing the diffraction effects that take place at the image slicer. In addition, as Zemax Optics Studio can deal with complicated designs of optical systems, we decided to use the HARMONI optical models to make our simulations as realistic as possible.

To validate the mathematical framework, we replicated the analysis with the HARMONI optical model in Zemax POP, varying the size of the pupil mirror aperture. The results in Fig. 3.8 and Fig. 3.9 show the exit slit PSF (central slice) at $1.5 \mu\text{m}$ for two cases of pupil mirror aperture given by the number of PSF zeros that fit at each side of the core, 2 and 3 respectively. There is good agreement between the Python simulation and the Zemax POP, both showing the PSF fringes that our theoretical calculations predicted. This suggests that the mathematical model we have developed for the light propagation within systems with image slicers captures the most relevant effects coming from the slicer itself and the pupil mirrors finite apertures.

So far we have presented results for the central slice where most of the PSF light is concentrated, but extending the analysis to other slices is straightforward. Recalling the formula for the electric field at the exit focal plane for the k -th slice: $E_f^k = \mathcal{F}[\Pi_m^k] * [\Pi_s^k \cdot \mathcal{F}(P(\mathbf{u}))]$, the slice mask Π_s^k simply acts as a masking for the chosen slice, and $\mathcal{F}[\Pi_m^k]$ the Fourier transform of the pupil mirror aperture will have the same characteristics for any given slice, as the apertures are all the same. This means that the situation for an arbitrary slice is equivalent to that of the central slice: the convolution with the pupil aperture term will induce fringes in the exit slit PSF, and the number of fringes is constant for all slices.

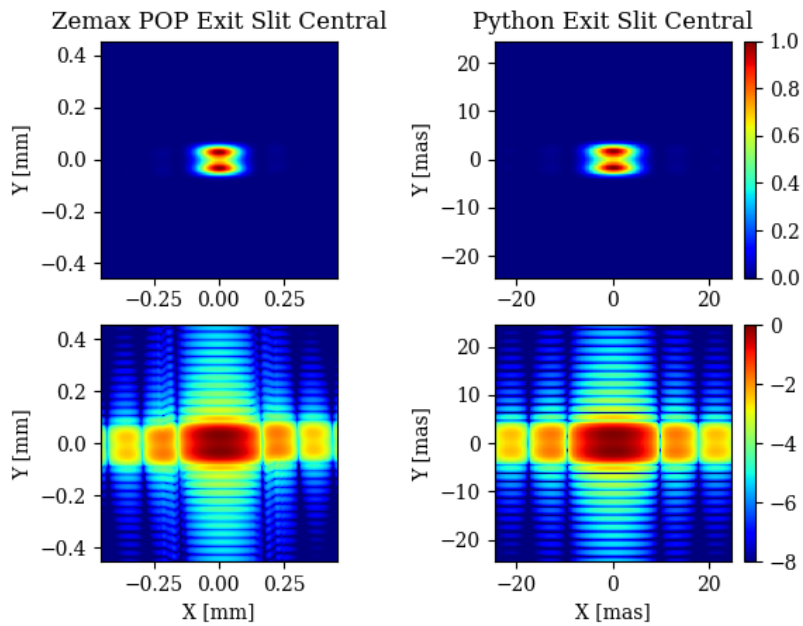


Figure 3.8: Model validation. Comparison between Zemax POP and Python simulations for the $1.5 \mu\text{m}$ exit slit PSF (central slice). Pupil mirror aperture fits 2 PSF zeros at each side of the core. *Upper row*: linear scale, *lower row*: logarithmic scale




Figure 3.9: Model validation. Comparison between Zemax POP and Python simulations for the $1.5 \mu\text{m}$ exit slit PSF (central slice). Pupil mirror aperture fits 3 PSF zeros at each side of the core. *Upper row*: linear scale, *lower row*: logarithmic scale

To demonstrate this, we repeated the previous analysis for the first adjacent slice. The PSF remains centred at the image slicer, but this time we compute the intensity on the neighbouring slice. The results for both Zemax POP and the Python model simulations are presented in Fig. 3.10 and Fig. 3.11, showing good agreement in the number of PSF fringes and overall structure.

In HARMONI, the slitlets associated with each slice are staggered to construct a pseudo-slit that is fed to the spectrograph; therefore each follows an independent optical path through the instrument. However, here we can make use of the simulated slits to reconstruct the complete PSF by stitching them back together. This gives us some insight into the effect of the image slicer on light propagation. Fig. 3.12 and Fig. 3.13 show the complete PSF after recombining the central 9 slices at $1.5 \mu\text{m}$ both with Zemax POP and Python. The structure induced by the image slicer is clearly visible. It is easy to identify the slice boundaries both in linear and logarithmic scale by the sharp decrease in intensity. Moreover, as predicted by the theoretical model of the light propagation, each slice of the PSF presents a number of fringes proportional to the relative size of the PSF and the pupil mirrors finite apertures.

In conclusion, in this section we have used simulations from Zemax POP to validate the theoretical model of the light propagation through image slicers we had developed. We used the HARMONI optical model to ensure our predictions are realistic and applicable. The comparison between Python simulations and the independent Zemax models showed good agreement suggesting that our theoretical model manages to capture the most relevant effects. We have also demonstrated that the pupil mirrors are key in the final PSF structure. Their finite aperture crops part of the strongly diffracted light after the image slicer, inducing fringes in the final PSF at the exit slit, and that the number of fringes is dependent on the aperture size.

This is the first step towards showing that direct application of Phase Diversity in systems containing image slicers could be challenging. Our literature review on NCPA calibration techniques demonstrated that the accuracy of PD is highly dependent on the quality of the optical model. The effects shown here, in particular the PSF fringes, could potentially be mistaken for spurious aberrations if the optical model used in the PD algorithm does not account for image slicer effects. It should be noted that this is true independently of how the image slicer responds to defocus, as all these calculations were done for the nominal configuration.

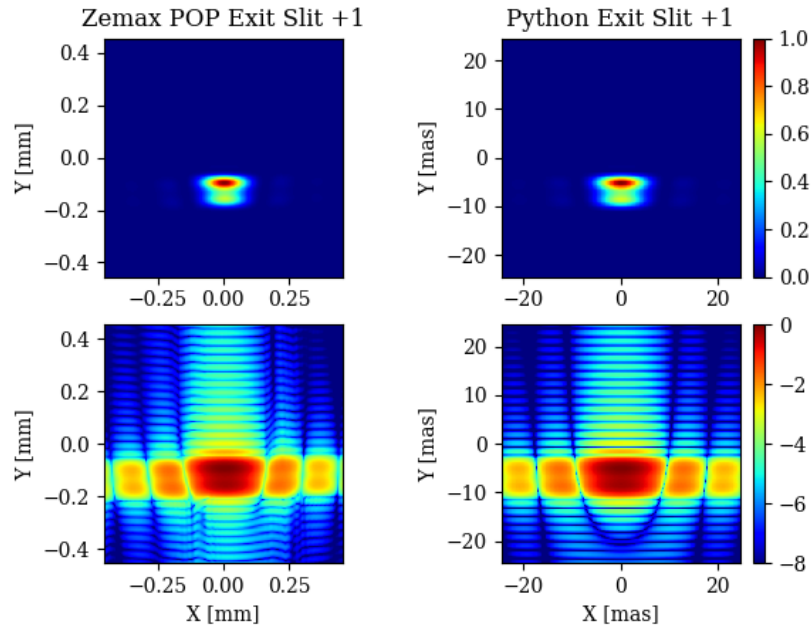


Figure 3.10: Model validation. Comparison between Zemax POP and Python simulations for the $1.5 \mu\text{m}$ exit slit PSF (first adjacent slice). Pupil mirror aperture fits 2 PSF zeros at each side of the core. *Upper row*: linear scale, *lower row*: logarithmic scale




Figure 3.11: Model validation. Comparison between Zemax POP and Python simulations for the $1.5 \mu\text{m}$ exit slit PSF (first adjacent slice). Pupil mirror aperture fits 3 PSF zeros at each side of the core. *Upper row*: linear scale, *lower row*: logarithmic scale




Figure 3.12: Model validation. Comparison between Zemax POP and Python simulations for the 1.5 μm reconstructed PSF at the exit slit. Pupil mirror aperture fits 2 PSF zeros at each side of the core. *Upper row*: linear scale, *lower row*: logarithmic scale

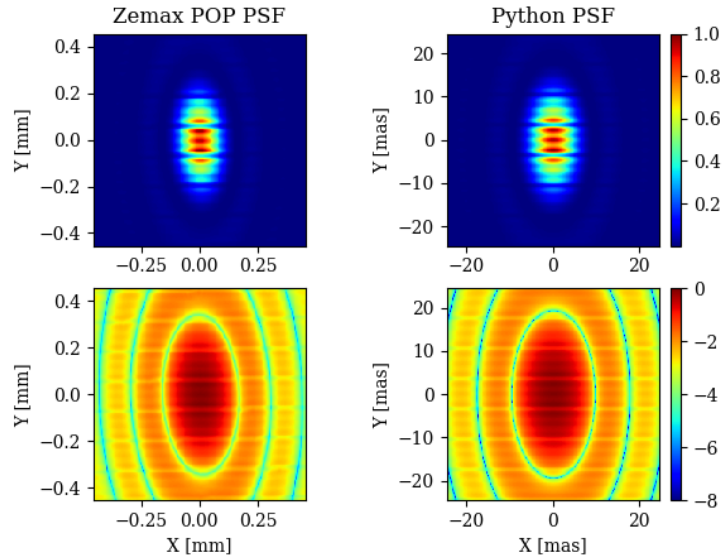


Figure 3.13: Model validation. Comparison between Zemax POP and Python simulations for the 1.5 μm reconstructed PSF at the exit slit. Pupil mirror aperture fits 3 PSF zeros at each side of the core. *Upper row*: linear scale, *lower row*: logarithmic scale

3.3.3 Impact of coarse sampling on slicer effects

Before we move on to the effects of defocus, we will take a closer look at the fringes shown in the previous section. Using the complete PSF calculated with Zemax POP for the nominal HARMONI IFU design, we analysed its cross-section along the axis perpendicular to the slices (y direction), see Fig. 3.14. We have included the PSF profile without slicer effects for comparison (calculated with POP after removing the image slicer aperture).

The effects of the image slicer are clearly visible with marked fringes and substantial light suppression at the boundaries between slices. We have also calculated the change in integrated intensity Δ for each slice due to the presence of the image slicer, i.e. the total intensity within one slice with slicer effects compared to the nominal no slicer case, which ranges around 2-5% relative loss. Repeating that calculation for the complete PSF profile yields a total light loss of around 2.5% of the nominal value.

Although the effects of image slicers are important, it should be noted that, in the case of HARMONI, each slice is sampled with only two spaxels (in the perpendicular direction), which leads to a relatively coarse sampling of the PSF. Because of this, the fringes will not be resolved, as the intensity is integrated into coarse pixels. In order to assess how this affects the PSF, we binned the cross-section profile from Fig. 3.14 with a sampling of two pixels per slice, see Fig. 3.15. The effects from the image slicer are not directly visible, given the coarse sampling, but the total intensity within each pixel is affected, mostly because of the light losses at the slice boundaries.

Despite the fact that the relatively coarse HARMONI sampling of two spaxels per slice partially mitigates the impact of image slicers by not resolving the induced fringes, the PSF is ultimately modified, with a significant loss of intensity at the core compared to a no-slicer PSF. Using the Marechal approximation $S = \exp(- (2\pi\sigma_\Phi/\lambda)^2)$, a suppression on Strehl ratio of 2.5% is equivalent to an RMS wavefront σ_Φ of around 0.025λ , or 37.5 nm for a reference wavelength of 1.5 μm , which represents a non-negligible contribution. This highlights the importance of accounting for image slicer effects in NCPA calibration techniques, to reduce the risk of the algorithms mistaking these features for spurious aberrations.

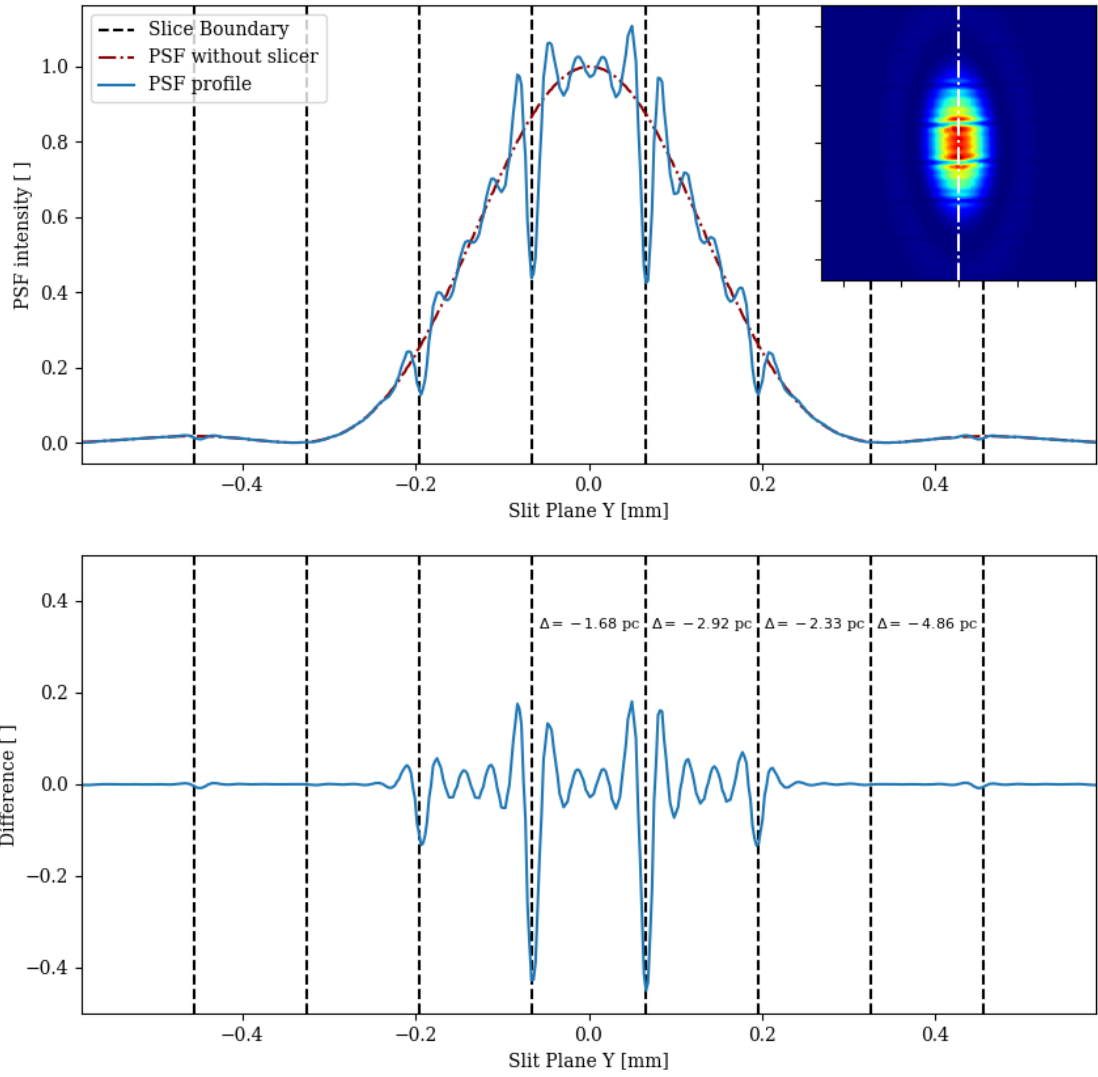


Figure 3.14: Image slicer effects. Zemax POP PSF cross-section perpendicular to the slices. A PSF profile calculated without the effect of the image slicer is shown for reference. Reference wavelength of $1.5 \mu\text{m}$. The Δ values represent the change in integrated intensity for each slice due to the slicer effect.

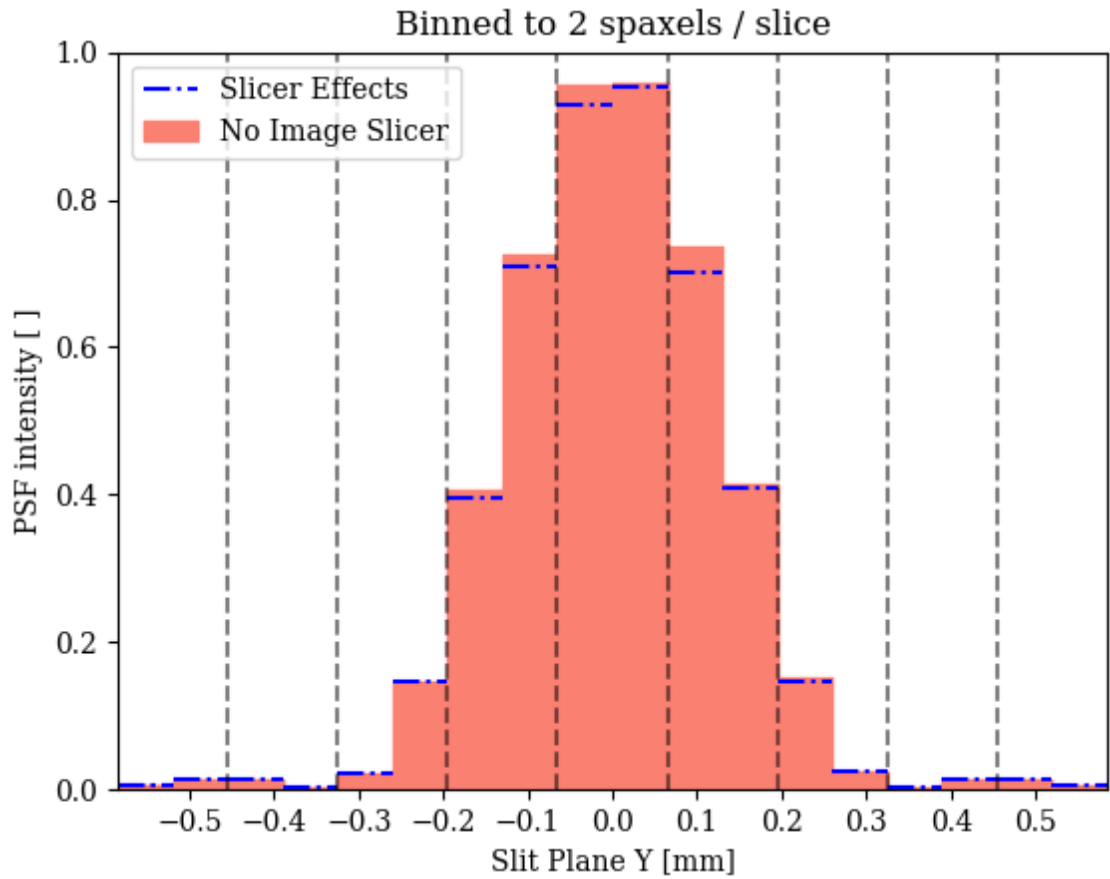


Figure 3.15: Image slicer effects. Zemax POP PSF cross-section perpendicular to the slices, binned to match the HARMONI sampling of two spaxels per slice. Although the fringes are no longer resolved at this sampling, the light losses due to the image slicer are integrated into the pixels, leading to a PSF profile that deviates from the no-slicer case. Nominal PSF, no aberrations, reference wavelength of $1.5 \mu\text{m}$.

3.3.4 Effect of aberrations

The previous section dealt with the behaviour of the light propagation in the presence of image slicers without aberrations. But we have already mentioned that calibration techniques like Phase Diversity require some form of diversity in the wavefront (typically a defocus), and the effect of such aberrations in the light propagation is thought to be disruptive enough to significantly affect the PSF.

For this reason, in this section we will explore the response of the image slicer to phase aberrations, with an emphasis on defocus. Let us begin by recalling the formula for the electric field at the exit slit for the k th slice, as derived with our theoretical model of the light propagation:

$$E_f^k = \mathcal{F}[\Pi_m^k] * [\Pi_s^k \cdot \mathcal{F}(P(\mathbf{u}))] \quad (3.13)$$

Without loss of generality, we can represent the complex pupil function $P(\mathbf{u})$ as a series expansion on the wavefront

$$P(\mathbf{u}) = \Pi_p e^{i\Phi} = \Pi_p [1 + i\Phi + \mathcal{O}(\Phi^2)] \quad (3.14)$$

For simplicity, we will assume the wavefront is sufficiently small that we can truncate the Taylor expansion at order $\mathcal{O}(\Phi)$. Using the Taylor expansion we can write the electric field at the exit slit as:

$$E_f^k = \mathcal{F}[\Pi_m^k] * [\Pi_s^k \cdot (\mathcal{F}[\Pi_p] + i\mathcal{F}[\Pi_p\Phi])] \quad (3.15)$$

Using the distributive property of the convolution we can write

$$E_f^k = \mathcal{F}[\Pi_m^k] * [\Pi_s^k \cdot (\mathcal{F}[\Pi_p])] + i\mathcal{F}[\Pi_m^k] * [\Pi_s^k \cdot (\mathcal{F}[\Pi_p\Phi])] \quad (3.16)$$

where the first term corresponds to the original case of no aberrations ($\Phi = 0$) that we studied in the previous section, and the extra term accounts for the effect of the wavefront aberrations

$$E_f^k = E_f^k(0) + i\mathcal{F}[\Pi_m^k] * [\Pi_s^k \cdot (\mathcal{F}[\Pi_p\Phi])] \quad (3.17)$$

This shows that the aberrations introduce a change in the electric field at the exit slit that is given by the convolution between the Fourier transform of the pupil mirror aperture, and the slicer-masked Fourier transform of the wavefront aberrations Φ ; an approximation that is accurate to order $\mathcal{O}(\Phi^2)$. This term is very similar to the original $\mathcal{F}[\Pi_m^k] * [\Pi_s^k \cdot \mathcal{F}(P(\mathbf{u}))]$ and thus it inherits its characteristics. It will




Figure 3.16: The effect of defocus. Comparison between Nominal PSF and defocus (0.4λ PV).

be spatially modulated by the pupil mirror aperture in the same way, which means that the number of PSF fringes will remain constant in the case of a system with aberrations.

In order to validate this, we used the HARMONI optical model from Zemax. As shown in Fig. 3.16, introducing a defocus (0.4λ PV) did not change the number of PSF fringes. As predicted by the theoretical model, the difference between the images is spatially modulated by the pupil mirror aperture. This can be seen in the region of the central slice, where the change in intensity due to defocus is stronger at the three central fringes. For instance, in the central slice, the three fringes lose more intensity than their surroundings; and in the areas along the first PSF ring the fringes are equivalently enhanced.

To further illustrate the effect of aberrations we decided to compare the PSF with and without the image slicer for several values of defocus intensity. To model the case without the image slicer we simply oversized the aperture values of the slicer, and pupil mirrors; thus removing any diffraction effects. Fig. 3.17 and Fig. 3.18 compare the PSF with and without the image slicer for defocus values of 0.20λ and 0.50λ PV respectively. The presence of the image slicer significantly alters the PSF structure, as we have pointed out throughout this section, specially around the central slice where changes in intensity due to fringes can be as high as $\pm 40\%$. As the strength of defocus grows, so does the extent of the intensity changes as they start affecting neighbouring slices.




Figure 3.17: The effect of the image slicer. Comparison of two defocused images (0.20λ PV), with image slicer (*left*), without image slicer (*centre*) and difference (*right*).




Figure 3.18: The effect of the image slicer. Comparison of two defocused images (0.50λ PV), with image slicer (*left*), without image slicer (*centre*) and difference (*right*).

Finally, we repeated the PSF cross-section analysis and binning for the case of 0.20λ PV defocus, the results are shown in Fig. 3.19. In absolute terms, the image slicer effects are approximately the same as the nominal case, with an excitation of peak fringes around +20% of the PSF peak, followed by a suppression of around -40% at the central slice boundaries.

3.3.5 Final remarks

At the beginning of this chapter, we mentioned a concern regarding the possibility of clipping the geometric beam footprint in out of focus conditions, given that, for common values of defocus, the beam footprint would cover more than one slice. However, the results from Zemax POP simulations show no evidence for major disruptions of the PSF such as significant light losses or shadowing. The defocused images broadly resemble the no-slicer cases, except for the obvious presence of fringes.

Nevertheless, this does not necessarily mean that conventional techniques like Phase Diversity are directly applicable. All the analyses point towards a strong impact of the image slicer on the structure of the PSF; examples of this being the clearly visible slicer boundaries and the fringes. This highlights the importance of accurately modelling the light propagation in order to avoid image slicer effects being mistaken for spurious aberrations.

Unfortunately, no HARMONI hardware is available at the moment to independently validate the predictions presented in this chapter. But this does not mean nothing can be done. In order to further strengthen our case, the following chapter will present the findings of an experiment done using the engineering model of the SWIFT image slicer sub-system.

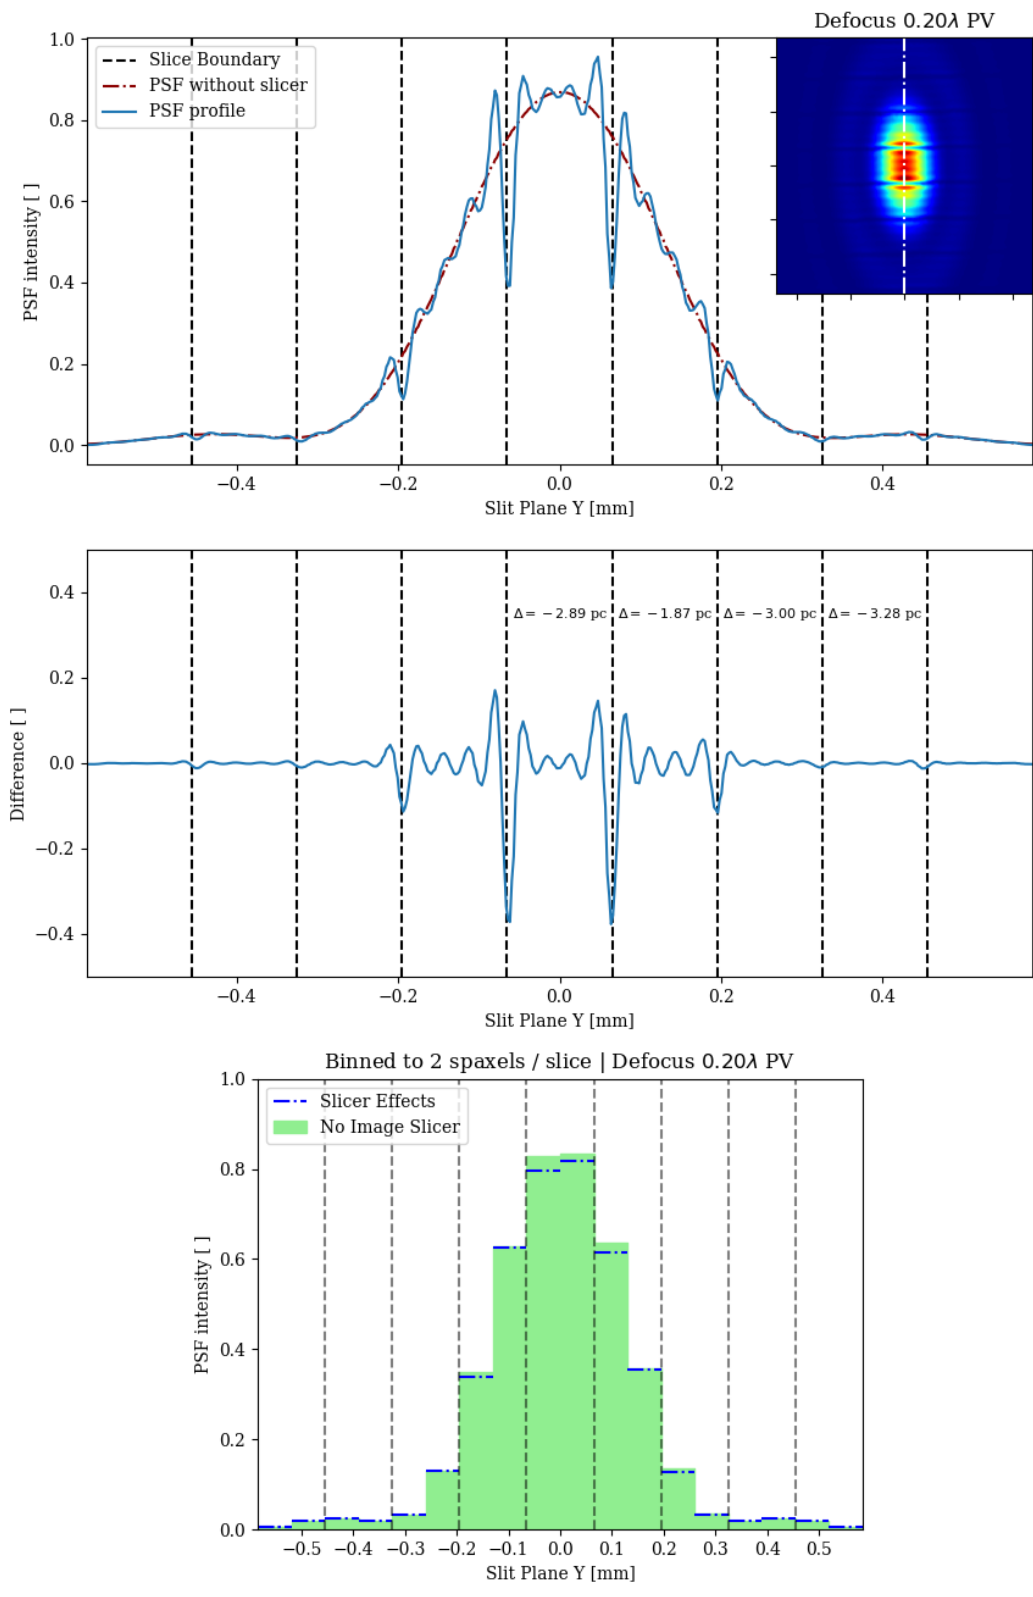


Figure 3.19: Image slicer effects, PSF cross-section, binned, defocus 0.20λ PV.

4

Experimental Work - SWIFT

”To succeed, planning alone is insufficient. One must improvise as well”

– Isaac Asimov, *Foundation*

4.1 Motivation

The main topic of this DPhil thesis is the calibration of non-common path aberrations within the HARMONI instrument project. This type of calibration is relevant during the commissioning phase and also for the daily operation of the instrument, when one wants to ensure it reaches its maximum capabilities.

Unfortunately, when dealing with such large scale projects as the next generation of ELTs and their instruments, it is not uncommon to spend years, if not decades, in the design and development phases. This means that access to the relevant hardware is not always guaranteed, which complicates the experimental validation of techniques and algorithms developed for calibration. This is why most of the work of this thesis relies on simulations and models. For instance, we have made use of the optical models for the different sub-systems as well as the End-to-End model of HARMONI, which although are still under review, represent the most complete and realistic models of the instrument to this date.

Given the characteristics of HARMONI, one of the key concerns regarding NCPA calibration has been the potential impact of image slicers on the PSF and the detrimental effects that this could have on state-of-the-art techniques like phase diversity. Through our modelling work we have shown that, under certain conditions, image slicers can imprint fringes on the PSF, see Chapter 3 and the results from [Menduiña-Fernández et al. \(2020\)](#). These findings are supported by a mathematical model of light propagation through image slicers and independent Zemax POP simulations of

the HARMONI IFU sub-system optical model. This helps build confidence about the validity of these predictions, but lacks the support of observational evidence.

In order to strengthen the case, we decided to complement this theoretical analysis with some experimental work. This is the motivation behind the work we will hereby refer to as the **SWIFT experiment**.

Due to the unavailability of any HARMONI instrument hardware at the time of this DPhil project, we had to consider an alternative. The closest thing we had to an IFS instrument was an engineering model of the image slicer for the SWIFT instrument, an IFS designed to operate at the 200-inch Hale Telescope at the Palomar Observatory (Tecza et al., 2006). The SWIFT image slicer uses two sets of flat mirrors for the slicing and re-arranging of the input field, followed by a mosaic of identical small lenses to de-magnify the input field and to create a short exit slit (Tecza et al., 2006, 2008). Although this design is slightly different than that of HARMONI, as it uses a mosaic of lenslets instead of relying on a stack of mirrors to construct a long pseudo-slit out of individual slitlets, SWIFT is still useful for our purposes of investigating the behaviour of image slicers.

With regards to the goals of the experiment, they can be summarised as follows:

- **Goal 1: Characterise the impact of image slicers on image formation.** Investigate how the presence of an image slicer can distort the PSF, as this could have an impact on the behaviour of calibration algorithms that rely on optical system models.
- **Goal 2: Validate the simulations from Zemax POP and Python.** Both approaches predict the presence of PSF fringes and are in good agreement, but further validation with experimental data is desirable.

4.2 Experiment design

The SWIFT image slicer cannot operate by itself, so a series of optical sub-systems had to be designed around it for the purpose of this experiment. These include *light source*, *pre-optics*, and *back-optics*, its main functions are summarised in Table 4.1, and a diagram of the experiment and the different sub-systems in Fig. 4.1. The final pre-optics sub-system is shown in Fig. 4.2, followed by the SWIFT sub-system containing the image slicer stack in Fig. 4.3.




Figure 4.1: Schematic diagram of the SWIFT experiment and its sub-systems. The *light source* sub-system uses a NIR LED plus a narrowband filter and an optical fibre to inject monochromatic light into the system. The *pre-optics* sub-system uses a series of lenses and a pinhole to deliver a beam of the appropriate f-number to the SWIFT slicer and form a pupil at the correct location. Due to mechanical constraints, the *back-optics* sub-system re-images the exit slits onto the detector focal plane, allowing us to capture 4 slices with a single exposure.




Figure 4.2: Pre-optics sub-system is connected to a y -linear stage to allow scanning, and an xy stage for fine alignment. The fibre, front lens and pinhole are housed inside a cage system; the pinhole acts as pupil stop and is held inside an xy stage for fine alignment. A larger cage system connects the front cage to the SWIFT field lens. The complete sub-system is connected to y stage that allows for vertical alignment between the pre-optics and the SWIFT image slicer stack. The SWIFT field lens mount was custom-designed and manufactured specifically for this experiment.

Sub-system	Optical Functions
<i>Light Source</i>	Provide a (quasi)-monochromatic source to the PO Deliver enough photons at the entrance of the PO
<i>Pre-Optics</i>	Deliver a beam of appropriate F# ratio to the Slicer Ensure a pupil is formed 162 mm behind the Slicer Focus the light at the Slicer Allow adjusting the source in y to scan the PSF
<i>SWIFT Slicer</i>	Slice the input field from the PO and form exit slits
<i>Back-Optics</i>	Re-image exit slits to accommodate the detector back-focal distance Provide an appropriate sampling with small enough pixels Capture 4 slices simultaneously on a single exposure Minimise additional WFE that could distort the slicer effects

Table 4.1: Optical functions for each of the sub-systems comprising the SWIFT experiment. Except for the SWIFT Image Slicer, all other sub-systems had to be designed and constructed for this experiment.




Figure 4.3: SWIFT Image Slicer and back-optics sub-systems. The stack is incomplete and only contains groups of 4 consecutive slices. The different slices of the Image Slicer split the incoming field of view and send each slit to a pair pupil mirror / lenslet to create a pseudo-slit pattern.

With regards to the **light source** sub-system, its main requirements are to provide an approximately monochromatic light in the NIR regime, while acting like a point source and delivering enough flux into the system to minimise the exposure time.

We decided to use a NIR LED source (Thorlabs LED850LN), centred at 850 nm with a FWHM of 55 nm and a half viewing angle of 3.5 degrees. To maximise photon flux we aimed for the highest power available; with a nominal 100 mW of light for a current of around 500 mA at 2.0 V. The LED source is powered by a custom-made PCB with a resistor and micro USB connection.

In order to fulfil the requirement of monochromatic light, we use a narrowband filter (Thorlabs FL05850-10) centred at 850 nm, 10 nm FWHM and 70% transmission. As part of an experimental validation we measured the transmission of the LED and the LED + filter from 700 to 1000 nm, to confirm that the components complied with the specifications and that the light source sub-system delivered the desired 850 nm emission, monochromatic with a FWHM of approximately 10 nm.

The light coming from the LED + filter is focused with a 1:1 system using a biconvex lens (Thorlabs LB1092-B-ML) of 15 mm focal length, and then coupled to a multimode fibre (Thorlabs M67L02) of 25 μm core and 0.10 NA. Considering that the size of the fibre core is 25 μm and that the magnification of the preoptics sub-system is around $4\times$, the source size at the image slicer is approximately 0.1 mm. The width of each slice is 0.47 mm and the FWHM of the PSF at the image slicer spans 4 slices. This implies that the FWHM is around 20 times the size of the light source, which suggests the fibre core is adequately small. Using smaller cores would hinder the coupling with the LED source and significantly reduce the number of photons reaching the detector.

For the **pre-optics** sub-system, the overall goal is to receive the light coming out of the fibre at the exit of the light source sub-system and re-image it onto the image slicer with the appropriate F# (matching the SWIFT design), ensuring that an exit pupil is formed 162 mm behind the slicer focal plane, while keeping the optical aberrations as small as possible. This is achieved with a design that uses two lenses and a pinhole (acting as pupil stop for the system), as shown in Figs. 4.1 and 4.2.

The first optical element is an achromatic lens (Edmund Optics #45-803) of 50 mm focal length, located approximately 1 focal length away from the light source sub-system, its main role is to collimate the light coming from the fibre. This lens is followed by a 150 μm diameter pinhole (Thorlabs P150D) that serves as the pupil stop to accommodate the F#. Ideally one should choose the biggest possible pinhole to minimise the percentage of rejected light and thus improve signal-to-noise ratio,

reduce exposure time and simplify alignment. However, this leads to a trade-off between the focal length and distance of the last lens, as a larger pinhole requires longer focal lengths to fulfil the $F\#$ requirement. We found that 150 μm diameter was a reasonable compromise. After the pinhole, a spare of the SWIFT field lens focuses the light onto the image slicer.

The design of the pre-optics did not come without challenges. To cite one example, due to the unconventional diameter of the SWIFT Field Lens (76 mm) a custom-made lens mount had to be designed and manufactured.

Finally, the rationale behind the design of a **back-optics** sub-system was that the camera available for this experiment had a large back-focal length (34.8 mm). Due to mechanical constraints coming from the camera mounting plate and the image slicer mechanical housing, it was not possible to match the detector plane to the exit slits focal plane of the slicer. We were forced to design a sub-system that could re-image the slits onto the detector plane. The back optics design consists of two matched achromatic pairs (Thorlabs MAP10100100-B), 1:1 $f_1 = 100$ mm, $f_2 = 100$ mm. This sub-system receives the 4 exit slits coming from the slicer (telecentric object space) and re-images them onto the camera, with sufficient distance after the last lens, to accommodate for the long back focal distance of the detector. In addition the backoptics sub-system has a slight magnification of around 1.13 which helps with the sampling by increasing the size of the PSF at the detector plane. The design of this sub-system was also optimised to minimise the additional RMS wavefront as much as possible and thus not alter the PSF coming from the image slicer.

The SWIFT Image Slicer has a nominal magnification of 0.24. Given that the slices are 0.47 mm wide, the apparent size of the slits is 0.1128 mm. With a further magnification of 1.13 for the back-optics sub-system, the slices should cover around 0.1275 mm at the detector plane. The camera used for the experiment has 13.5 μm pixels, which means that the the slices will be sampled with at least 9 pixels. This should provide enough resolution to detect effects such as fringes.

4.3 Data Reduction and Results

4.3.1 Alignment validation

[A] *Pre-Optics Pupil Test.*


In order to validate the alignment of the different sub-systems, we ran a series of diagnostic tests. First of all, a *pupil test* was performed to locate the optimal position for the image slicer. This was done using an alignment pinhole of 1 mm diameter. The pre-optics sub-system is supposed to create a pupil 162 mm behind the focal plane of the image slicer. To verify this, the image slicer was moved out of its nominal position, being substituted by the whole back-optics sub-system, allowing us to record images around the theoretical location of the pupil.

At this point, the back-optics was displaced along the optical axis in steps of 1 mm. At each position, the fibre at the entrance of the pre-optics performed a y scan. Combining multiple exposures of this input field scan allows us to locate the optimum pupil position: outside the pupil plane, the different images will move up and down in y as the fibre changes in field position, creating a blurred image when multiple exposures are added together. Fig. 4.4 shows the results of this *pupil test* with clear blurring of the images at both ends. The optimal pupil position with a well-defined pupil image was found to be 6-7 mm behind the initial position. Fig. 4.4 also compares the image of the pupil at the optimal position to the theoretical value given by the Zemax optical model of the pre-optics sub-system, with good agreement.


[B] *Pre-Optics PSF Test.*

Once the optimal pupil position was confirmed, the back-optics sub-system was moved forward 162 mm to the optimal location of the image slicer. To validate the alignment of the pre-optics with respect to the image slicer, we performed a *PSF test*. Using a pinhole of 100 μm diameter, we recorded the PSF at this focal plane as shown in Fig. 4.5. The experimental PSF was fitted to a theoretical Airy pattern with good agreement, see Fig. 4.6.

This confirms that the experimental PSF is close to a perfect Airy pattern (no visible effects of optical aberrations from the pre-optics) and that the pre-optics is properly aligned with the SWIFT image slicer.



(a) Finding the optimal pupil location behind the SWIFT slicer by displacing the back-optics along the optical axis, while recording the images of a fibre y scan at the entrance of the pre-optics.



(b) Experimental image of the pupil at the optimal position, compared to the design value as prescribed in the Zemax optical model of the pre-optics.

Figure 4.4: SWIFT Image Slicer pupil test.

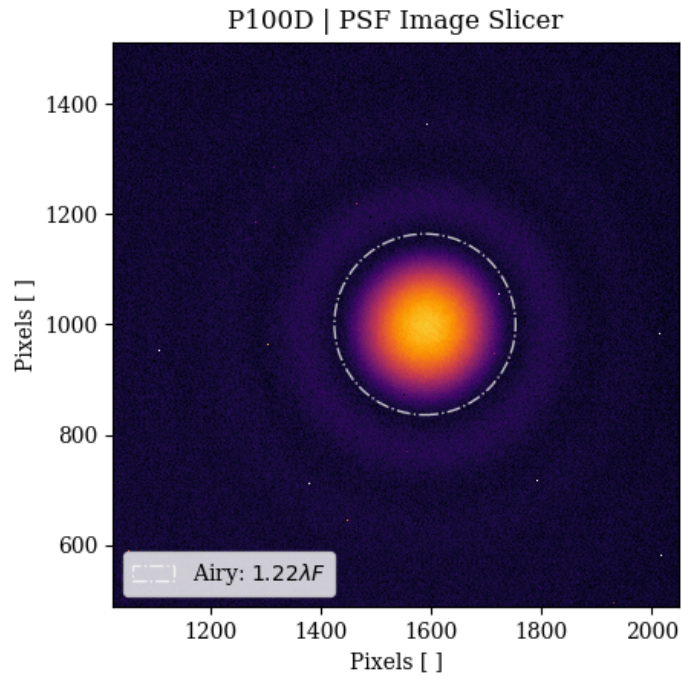


Figure 4.5: Pre-Optics PSF Test. Experimental PSF 162 mm in front of the optimal pupil location. Theoretical location of the first Airy zero is shown.

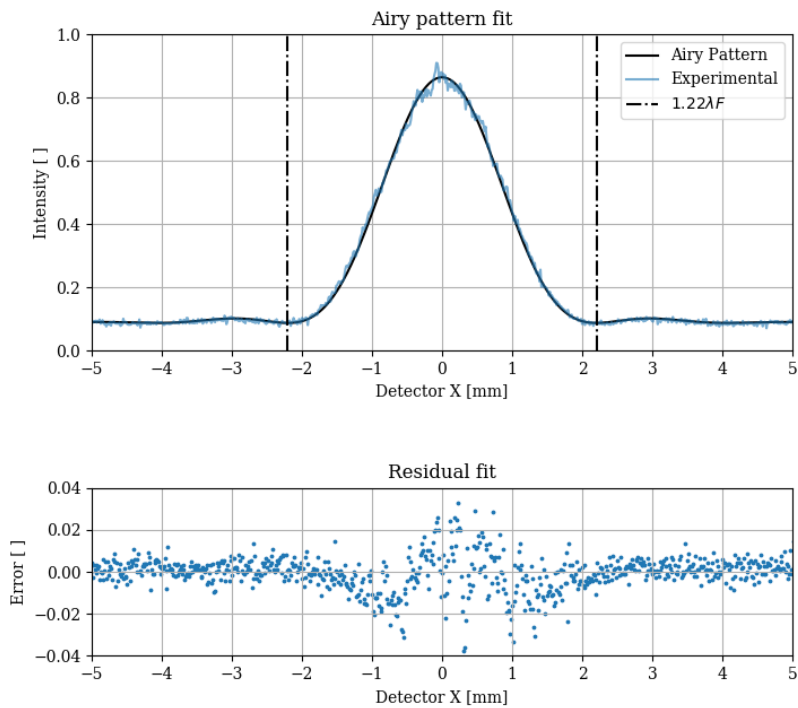


Figure 4.6: Pre-Optics PSF Test, Airy pattern fit. $\lambda = 0.85 \mu\text{m}$, $F = 2135$.

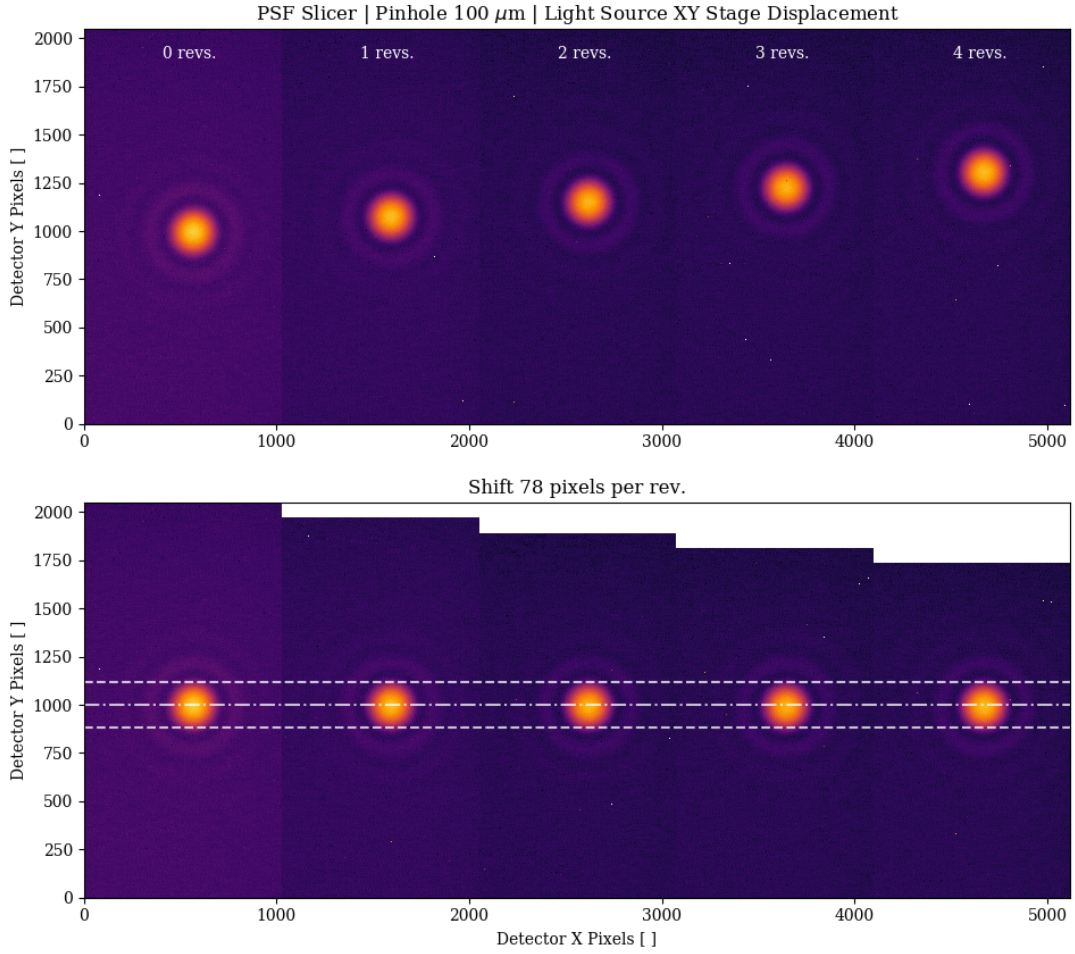


Figure 4.7: Pre-Optics Magnification Test. Displacement of the PSF at the image slicer focal plane as a function of fibre movement (turns of the y stage screw), and number of pixels per turn needed to compensate the shift.

[C] *Pre-Optics Magnification Test.*

Since the engineering model of the SWIFT image slicer has only 4 consecutive slices in the mirror stack, the field of view we can capture with a single exposure is limited to 4 slices, or approximately 1 FWHM PSF at $0.85 \mu\text{m}$. One possible strategy to increase the field of view and better characterise the PSF is to scan the source at the entrance of the Pre-Optics to displace the PSF up and down with respect to the image slicer plane, and stitch together multiple exposures, each covering a different region.

In order to do this, we need to accurately know the magnification of the pre-optics sub-system to calculate the δy_0 displacement of the fibre needed to move the PSF by the correct δy at the slicer ($4 \times$ the slice width, in this case). For this reason, we decided to determine the magnification of the pre-optics sub-system experimentally.

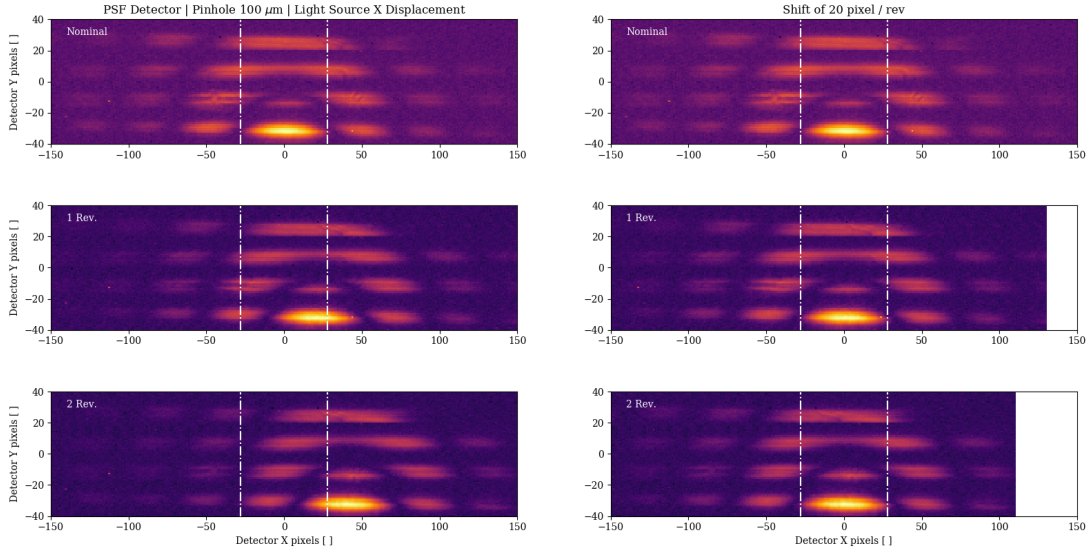


Figure 4.8: Complete System Magnification Test. Displacement of the PSF at the detector as a function of fibre movement (turns of the x stage screw), and number of pixels per turn needed to compensate the shift.

With the detector at the optimum position for the image slicer (recall the previous PSF test), we scanned the fibre at the entrance of the preoptics using the screw of its XY stage. We recorded the displacement of the PSF and calculated the number of detector pixels needed to compensate the shift, see Fig. 4.7.

The XY stage has 100 TPI (threads per inch) screws, meaning one revolution displaces the fibre by 0.254 mm (δy_0). We found that this shifted the PSF by 78 ± 1 detector pixels, corresponding to 1.053 mm (δy), or a magnification of 4.15 ± 0.05 . This value is in good agreement with the nominal design for the pre-optics sub-system, which has a magnification of 4.18.

[D] Complete System Magnification Test.

The previous tests confirmed that the pre-optics sub-system was properly aligned. At this point, a final test was done to confirm the magnification of the complete system to ensure the back-optics sub-system was also correctly aligned.

For the complete system, scanning the fibre in y is no longer an option, since the slicer mirrors are fixed, but scanning the fibre in x displaces the PSF along the slice, which, in turn, displaces it on the detector plane. Thus, we scanned the fibre in x and recorded the displacement of image of the fibre at the detector (more details on the post-processing of the detector frames in Section 4.3.2).

Fig. 4.8 shows the results of this experiment, with four slices roughly stitched together to form an approximate PSF at the detector. For each turn of the XY stage

Sub-system	Nominal	Experimental	Method
<i>Pre-Optics</i>	4.18	4.15 ± 0.05	PO <i>y</i> fibre scan
<i>SWIFT Slicer</i>	0.24	-	-
<i>Back-Optics</i>	1.13	1.13 ± 0.10	Tolerance analysis
<i>Complete System</i>	1.13	1.06 ± 0.11	PO <i>x</i> fibre scan

Table 4.2: Results of the magnification test. Nominal magnification for each sub-system and experimental values.

screw, we found that the PSF shifted by 20 ± 1 pixels, i.e. a displacement of 0.27 mm (δx) for every 0.254 mm (δx_0), with a magnification of 1.06.

Table 4.2 summarises the results of the magnification tests. The nominal magnification for the complete system is 1.13 (combining the pre-optics, SWIFT slicer and back-optics optical models). For the fibre scan experiments, we have assumed uncertainties of ± 1 detector pixels. Since we did not test the back-optics sub-system separately, we ran a Monte Carlo tolerance analysis on the optical model to get an estimate of the experimental magnification.

The experimental value of 1.06 ± 0.11 for the complete system is consistent with the nominal design, suggesting all the sub-systems are correctly aligned. This concludes the alignment validation tests, ensuring the optical system was ready for the experiment.

4.3.2 Preliminary results

The back-optics sub-system allows us to record all four exit slits simultaneously with each exposure. An example of a detector frame is shown in Fig. 4.9, representing a 5-second exposure. The frames are calibrated using master bias and dark current frames to mitigate the effects of detector systematics but, since the slices cover a very limited area of the detector, we do not correct for bad pixels. In order to reduce readout noise contamination and increase the SNR, multiple frames are averaged together.

With regards to sampling, the detector pixels are $13.5 \mu\text{m}$ in size. The apparent size of the slices of the image slicer at the exit of the SWIFT sub-system is around $100 \mu\text{m}$. As the back-optics sub-system was designed with a paraxial magnification of 1.13 (nominal value), this means the slices would be sampled with around 8-9 detector pixels. Each thin mirror of the SWIFT image slicer has a physical size of 0.47 mm.




Figure 4.9: Calibrated detector frame (corrected for bias and dark current) 5-second exposure showing the four exit slits (logarithmic scale). Numbers indicate the ordering of slices on the SWIFT image slicer mirror stack. The experimental setup allows us to capture the 4 PSF slices in a single exposure, but requires post-processing to re-align them and create a composite PSF image. The detail for slice #3 already hints at the presence of fringes.

Moreover, considering a nominal wavelength of 850 nm, 100 μm pinhole diameter and $F\#$ of 2135, we have seen that the FWHM of the PSF should be approximately 1.866 mm at the slicer focal plane. Consequently, each slice covers 1/4 of the FWHM, so that we have access to approximately 1 FWHM in a single exposure (after stitching the slices together). Obviously this limit only applies to the y direction as the slicer is composed of long thin mirrors, so we will also have access to the first couple of PSF rings in the x direction, provided we achieve sufficient SNR.

The preliminary data from a single detector frame (Fig. 4.9) already suggests the existence of slicer effects similar to the ones predicted in Section 3.3.1, showing what appears to be horizontal fringes. To investigate this further, we ran the following test. We displaced the back-optics sub-system with respect to the SWIFT slicer along the optical axis to see if changing the focus position has any impact on the features.

Fig. 4.10 shows the results for several displacements; the four slices have been roughly stitched together without trying to produce a complete PSF. As we sweep along the optical axis in focus, different slices start to show signs of image slicer features. For example, the top slice covering the first Airy ring looks smooth and featureless in the nominal position, but exhibits what appears to be two horizontal fringes at +2 mm. In contrast, the bottom slice covering part of the core of the PSF shows two sharp fringes at +1 mm, which start to blur away at +2 mm.

This is due to the fact that the slices are far apart on the detector plane: adjacent slices are around 900 pixels or 12.15 mm apart, and the two opposite slices are around 1500 pixels or 20.25 mm apart along the diagonal. Therefore, different slices come in and out of focus as the detector plane moves. This was partially mitigated by adjusting the tilt of the detector. What remains constant across these images is the unequivocal presence of fringes, which appear to have a spatial frequency of 2-3 cycles.

In order to investigate whether the presence of these fringes changes with field of view, we ran a y scan of the fibre at the entrance of the pre-optics. The results are shown in Fig. 4.11, in increments Δy of half a turn of the XY stage screw (0.127 mm displacement).

An interesting way of further confirming the experimental magnification of the pre-optics sub-system is implicitly shown in Fig. 4.11. The top two panels ($-7\delta y_0$ and $-2\delta y_0$) show the two edges of the first Airy ring (bottom and top slices respectively). Since the four slices cover 1 FWHM, this corresponds to a movement Δy of the PSF of an Airy ring minus 1 FWHM, at the slicer focal plane of:

$$\Delta y = (2 \cdot 1.22 - 1.028)\lambda F = 2.562 \text{ mm} \quad (4.1)$$




Figure 4.10: Impact of back-optics sub-system focus distance. Roughly stitched slices at different positions along the optical axis, showing clear signs of image slicer features (logarithmic scale). The large distance between slices at the detector plane make it hard to compensate for differential defocus and to adjust the focal distance to a point where the fringes are visible in all slices.




Figure 4.11: Pre-optics fibre y scan. Each panel represent the roughly stitched slices of the PSF after moving the fibre at the entrance of the pre-optics by $\Delta y = 1/2$ turn of the XY stage, or 0.127 mm. Logarithmic scale. Displacing the fibre in y allows us to sample different regions of the PSF field of view, and potentially create a composite image that goes beyond the 4-slice limit.

In this case, δy_0 corresponds to half a turn of the screw, or 0.127 mm, the magnification is $m = \Delta y / \Delta y_0 = 2.562 / (5 \cdot 0.127) = 4.03$.

In order to better characterise the effects of the image slicer on the PSF we looked at the cross-section of the slices, see Fig. 4.12 and 4.13. The physical size of the slices at the detector plane is shown as a grey-filled region to allow a direct comparison of the recorded intensity and the theoretical boundary. The first thing to note is that, instead of looking smooth, the slices of the PSF exhibit a 'dip' in intensity at the centre of the slice, corresponding to the space between the two fringes. Given the relatively coarse sampling with only ~ 9 pixels within the slice boundary, it is hard to infer the presence of higher spatial frequency signals.

4.3.3 Simulations Comparison

We ran a series of POP simulations using the Zemax optical model of the experiment, including all sub-systems. These simulations account for the coherent nature of light, and use diffraction calculations to propagate the wavefront through the different surfaces of an optical system. Since significant diffraction takes place at the image slicer, POP is needed to accurately capture these effects.

For these simulations, we used the same parameters as in the real experiment. We sampled the incoming beam with 1024×1024 arrays to ensure an accurate result during the propagation, but the final arrays were down-sampled to 256×256 with $13.5 \mu\text{m}$ pixels to match the sampling at the detector. To make the results as realistic as possible, we added readout noise contamination after POP, matching a SNR of 2000, determined experimentally from the exposure data.

The source at the entrance of the pre-optics sub-system was displaced up and down in steps of 1 slice width; the results are shown in Fig. 4.14. The first thing to note is that the POP simulations do predict fringes on the PSF, just like it was observed in the actual experiment. As one scans the PSF, the cross-sections reveal that the intensity is not smoothly distributed but shows signs of fringes / fluctuations. Compared to the experimental data (see Figs. 4.12- 4.13), the POP simulations hint at slightly higher spatial frequencies, but given the limited sampling of ~ 9 pixels per slice, this could be due to the inherent challenges of real experiment conditions compared to numerical analysis. In any case, in both cases, at least two clear fringes are visible at both ends of the slice.

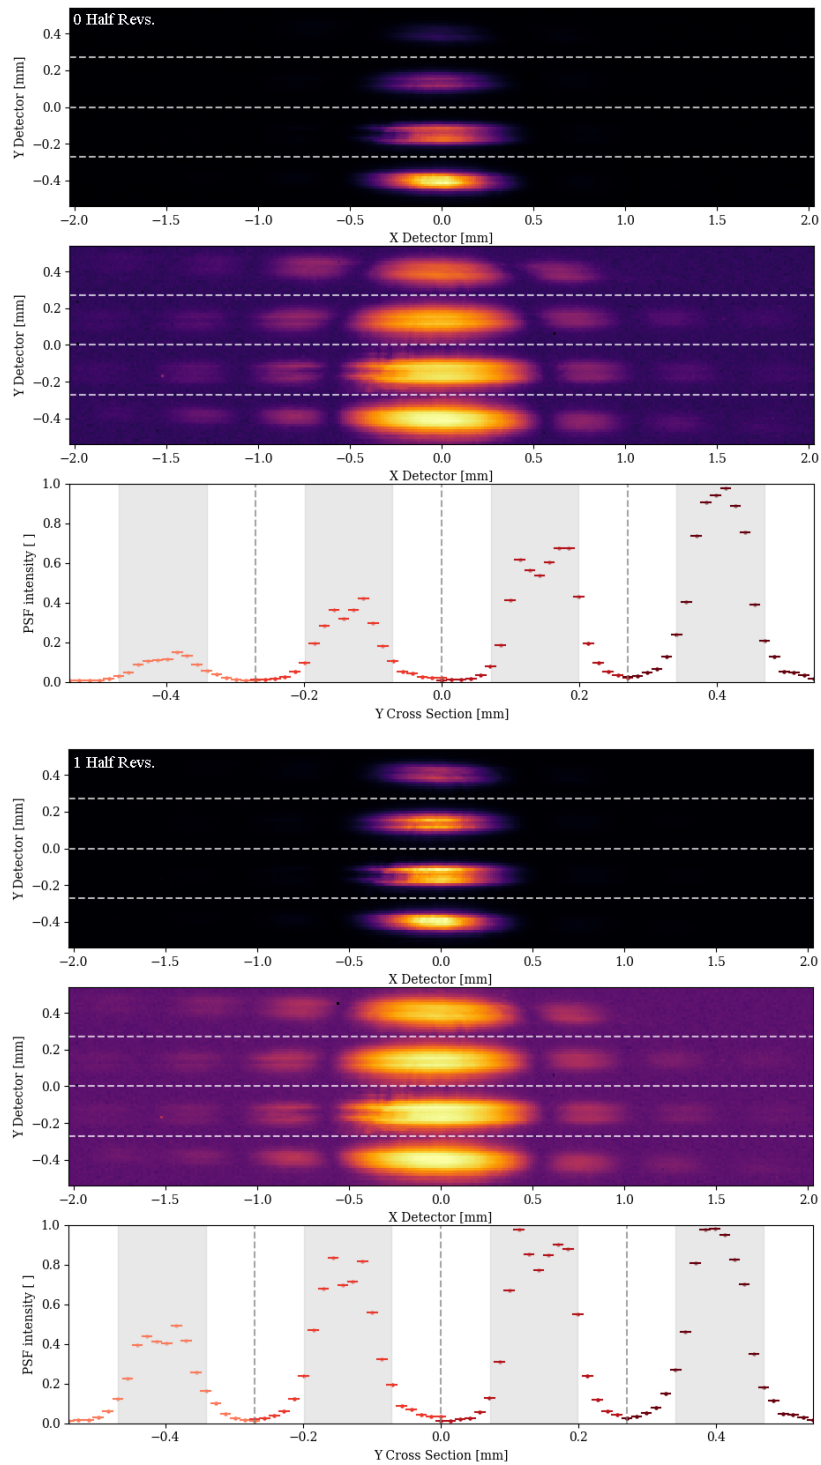


Figure 4.12: Experimental PSF slices at the detector (100 μm pinhole, 0.85 μm wavelength), and cross-sections through the centre of the PSF. Shaded areas represent the physical size of slices at the detector. Source is displaced by half turns of XY stage screw ($\sim 1.12 \times \text{slice width} / \text{half turn}$).

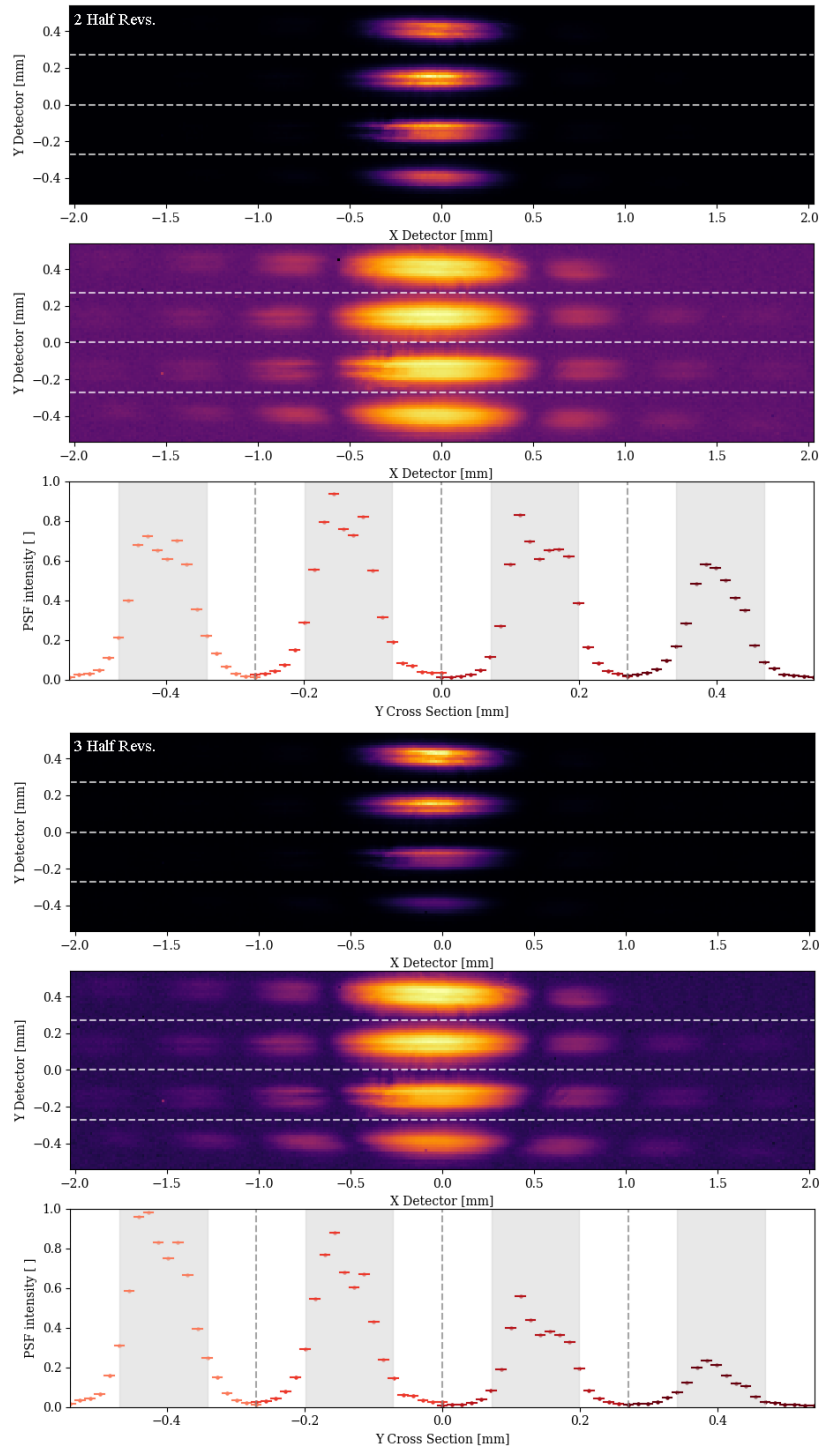


Figure 4.13: Extension of Fig. 4.12.

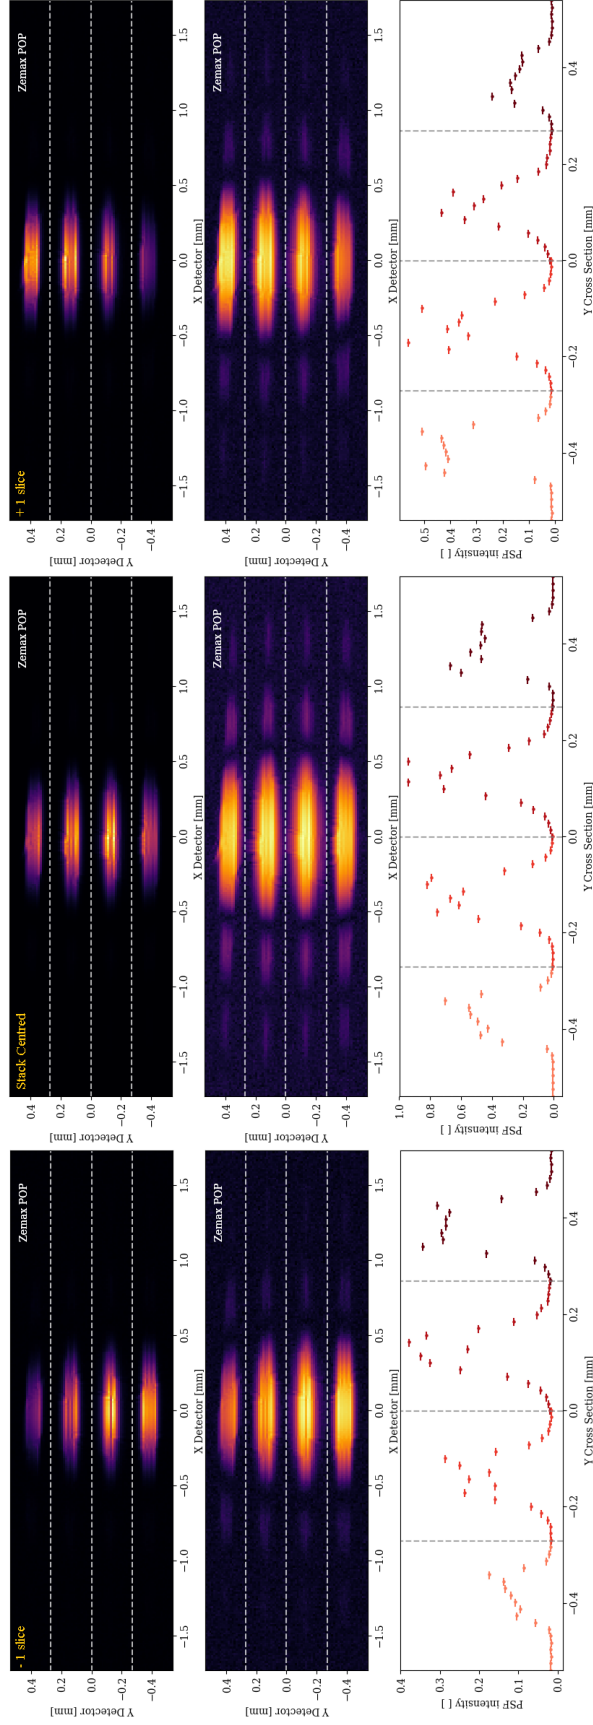


Figure 4.14: Zemax POP simulations of the SWIFT experiment (100 μm pinhole, 0.85 μm wavelength). *Middle*: PSF slices at the detector plane, for a source centred at the image slicer stack, and cross-sections through the centre of the PSF core. *Left & right*: sourced moved by ± 1 slice width. Readout noise with SNR 2500 was added to match the experimental data

4.3.4 Stitching the slices

The engineering model of the SWIFT Image Slicer is incomplete and only contains 4 adjacent slices at the bottom of the mirror stack. This imposes an intrinsic limitation on the field of view that can be captured with a single exposure, which is around 1 FWHM of the PSF at $0.85 \mu\text{m}$ with a $100 \mu\text{m}$ pinhole, and it forces us to scan the light source in y in order to cover a wider area. In addition, a single exposure captures four slices, but these are separated by a significant distance on the detector plane, recall Fig. 4.9.

Consequently, some form of post-processing is needed to re-arrange and combine data from multiple exposures to construct a complete PSF; a process that we will call *stitching*. This can be challenging for multiple reasons. First of all, the lack of stability of the light source between exposures might create the need for flux normalisation. Secondly, in order for the stitching to work, we need to move the source by the correct amount to displace the PSF at the slicer by exactly 1 slicer stack. This might be challenging given uncertainties in the magnification of the pre-optics or general instabilities and movements of the set-up between exposures.

In this section, we will present some preliminary results for this stitching process. The way we approached this analysis can be summarised as follows. Every exposure k produces a set of four slices at the detector plane S^k :

$$S^k = [s_1(x_1^k, y_1^k), \dots, s_4(x_4^k, y_4^k)] \quad (4.2)$$

where $s_i(x_i^k, y_i^k)$ represents the image of the i -th slice for that exposure k , centred at position (x_i^k, y_i^k) on the detector. Fig. 4.9 is an example of such S^k . The first step towards obtaining a PSF is aligning each exposure. For that purpose, we create a model \mathcal{P}^k given by:

$$\mathcal{P}^k(\Delta x_i^k, \Delta y_i^k) = [s_1(x_1^k + \Delta x_1^k, y_1^k + \Delta y_1^k), \dots, s_4(x_4^k + \Delta x_4^k, y_4^k + \Delta y_4^k)] \quad (4.3)$$

which represents the four slices s_i re-arranged according to some arbitrary shifts $(\Delta x_i^k, \Delta y_i^k)$ to form a section of the PSF. For every exposure k , the optimum shifts are found numerically by comparing the model $\mathcal{P}^k(\Delta x_i^k, \Delta y_i^k)$ to a reference PSF with no slicer effects. This way, the four slices for that exposure representing a section of the field of view are located with respect to the PSF. If we repeat this process for several exposures k_1, k_2, \dots, k_n , each corresponding to a shift of the PSF at the slicer

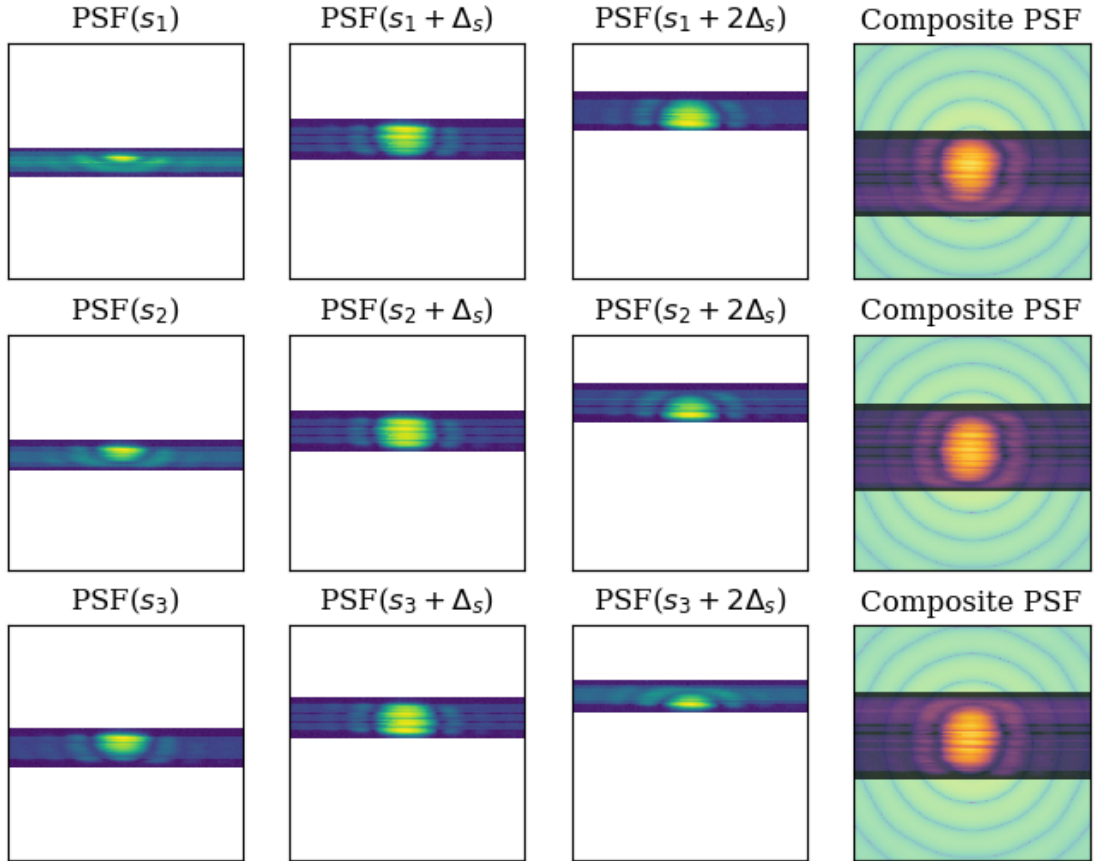


Figure 4.15: Stitching the slices. Each row shows the process of stitching $\mathcal{P}^k(\Delta x_i^k, \Delta y_i^k)$ models for several cases of k , where each exposure is shifted by the width of the slicer stack Δ_s . A panel $\mathcal{P}^k(\Delta x_i^k, \Delta y_i^k)$ represents the result of rearranging and combining the data from four slices at the detector for a given exposure k . Since data is available at intervals smaller than Δ_s , we can combine multiple groups to form several example PSF images with slightly different field of view. Images are in logarithmic scale. A reference PSF is shown to illustrate the quality of the fit.

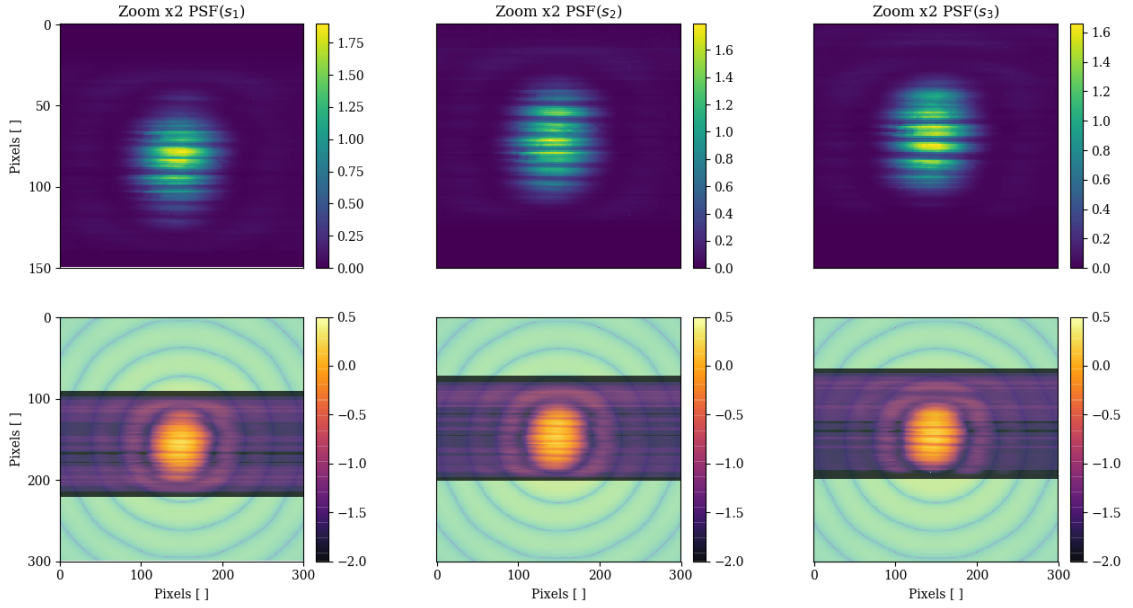


Figure 4.16: Stitched PSF models covering $4\Delta_s$ for a total of 16 slices. Each column represents a shift in the field of view of $\delta \simeq 1$ slice width. Residual errors from the fitting and stitching process are visible, but it shows a preliminary version of the composite PSF image.

plane Δ_s equal to the width of a slicer stack, we can then combine those sections $\mathcal{P}^k(\Delta x_i^k, \Delta y_i^k)$ into a composite image that covers $4n$ slices.

This process is summarised in Fig. 4.15. The first row shows the results of combining 3 consecutive sections of the PSF: $\mathcal{P}^1(\Delta x_i^1, \Delta y_i^1)$, $\mathcal{P}^2(\Delta x_i^2, \Delta y_i^2)$, $\mathcal{P}^3(\Delta x_i^3, \Delta y_i^3)$. The reason we have 3 separate rows is because we have access to exposures taken at y scans with increments that $\delta \simeq 1$ slice width instead of a complete slicer stack width Δ_s , so we can combine them in multiple ways to obtain partial PSF images covering slightly different fields of views.

The results for multiple cases of this stitching process covering $4\Delta_s$ or 16 slices are shown in Fig. 4.16, the only difference being a shift in field of view of $\delta \simeq 1$ slice width. Undoubtedly, this approach is far from perfect as some misalignment is clearly visible and other more sophisticated methods could be explored in the future. Nevertheless, it offers a valuable starting point and provides coverage of the PSF for a total of 16 slices in the perpendicular direction, which spans the core of the PSF plus the first Airy ring. One feature that is worth noting is the apparent presence of gaps between slices. This could be caused by errors in the stitching process but, if we recall Section 3.3.2, these gaps also resemble the gaps identified in Figs. 3.12 and 3.13 where the Zemax POP simulations of multiple slices were stitched together.

4.4 Limitations and future work

Field of view limitations

I have already mentioned that one of the main limitations of this experiment is the restricted field of view that one can access in a single exposure. It is important to remember that this caveat cannot be solved with a full mirror stack without any missing slices. Because of the intrinsic size of the lenslets, the physical separation between the different slices at the exit slit focal plane is already large and covers most of the active area of the detector, see Fig. 4.9. The only feasible solution is to scan the source at the entrance of the pre-optics and stitch together multiple exposure. However, this does come with some caveats:

- *Unstable light source.* Stitching together multiple exposures requires correcting and normalising the flux. This might introduce errors if the light source is not stable.
- *Magnification uncertainties.* In order for the stitching to work, we need to move the source by the correct amount to displace the PSF at the slicer by exactly 1 slicer stack. This might be challenging given uncertainties in the magnification of the pre-optics or general instabilities and movements of the set-up between exposures.
- *Defining the slice masks.* Identifying the boundaries of each slice at the detector plane and defining the appropriate masks to accurately stitch them together is a challenging task, and might require some form of flat-field calibration.
- *Stitching uncertainties.* In the previous section I demonstrated how the process of stitching together multiple exposures is not straightforward and can leave residual errors if the shifts needed to compound all the PSF sections are not accurate enough.

Pixel sampling

Image slicer effects appear to imprint fringes on the PSF, which can be of relatively high spatial frequency. Given the detector pixel size ($13.5 \mu\text{m}$) and the limited size of the slices (~ 9 pixels), frequencies above 3-4 cycles per slice would be hard to sense robustly.

No NCPA control

Another important limitation of this experiment is that we have no control over the NCPA. The system is static and we lack the means to modify the wavefront upstream

to investigate the link between aberrations and image slicer effects. Therefore, this set up cannot be used to generate realistic examples with varying aberrations with which to train the neural network models for NCPA calibration.

4.5 Conclusions

In this chapter, I have provided experimental validation of the effects of image slicers on light propagation. Due to the lack of HARMONI hardware, I have used the SWIFT image slicer engineering model as a substitute. A set of auxiliary sub-systems were designed and constructed around the image slicer to produce a working system that could be used for testing. After aligning the experiment's optical bench, I recorded exposures of the post-slicer PSF covering 4 slices simultaneously. Scanning the light source at the entrance of the pre-optics sub-system allowed me to cover a wider field of view, reaching beyond the core and first Airy ring.

The analysis of the different exposure data has revealed the presence of fringes in the PSF that are consistent with the predictions of the mathematical framework developed in Chapter 3. These features were further confirmed by comparing the experimental data to Zemax POP simulations of the optical model for the experiment.

The results presented in Chapter 3 and here (Chapter 4) address the first of the research sub-questions posed in Section 2.4.2: **"how does the presence of an Image Slicer impact the PSF?"** These findings emphasise the importance of understanding the effects of image slicers and considering them in the context of NCPA calibration. A successful calibration algorithm must take into account how the presence of an image slicer affects light propagation and ultimately modifies the PSF.

In the following Chapter, I will explore the second sub-question of **"what alternative techniques could be applied?"** that naturally follows acknowledging that image slicer effects appear to have a significant effect on the PSF.

5

Machine Learning for NCPA calibration

“People worry that computers will get too smart and take over the world, but the real problem is that they’re too stupid and they’ve already taken over the world.”

– Pedro Domingos, *The Master Algorithm*

5.1 Background and Motivation

In recent years, machine learning techniques have been applied successfully over a wide range of tasks in astrophysics. To cite some relevant examples: morphological classification of galaxies [Banerji et al. \(2009\)](#), identification of faint tidal features [Walmsley et al. \(2018\)](#), detection of bars in galaxies [Abraham et al. \(2018\)](#), identification of galaxy mergers [Pearson et al. \(2019\)](#), deblending of overlapping galaxies [Reiman and Göhre \(2019\)](#), classification of transients [D’Isanto et al. \(2016\)](#), computing the mass of planetary envelopes [Alibert and Venturini \(2019\)](#), transient candidate vetting [Gieseke et al. \(2017\)](#), or detection of exoplanets in high-contrast imaging [Gomez Gonzalez et al. \(2018\)](#) among others.

In part, this paradigm shift has been fuelled by a rapid growth in data size and complexity that calls for more automated and data-driven approaches, leaving behind traditional model fitting techniques that are very time consuming. A powerful example of this is the work by [Hezaveh et al. \(2017\)](#) on the analysis of strong gravitational lenses. They argue that the conventional maximum likelihood fitting of observations is a computationally expensive process that requires sophisticated codes, and expert knowledge of the physical processes and methods involved, taking up to a few weeks

to analyse a single lens. With the advent of large scale surveys like LSST thousands of new strong lenses are expected to be discovered, calling for techniques that can handle the sheer volume of data that the new era of astronomy will produce.

In their work, [Hezaveh et al. \(2017\)](#) tackle this challenge by training a convolutional neural network model to infer the parameters of the lensing model directly from observations of strongly lensed systems. Despite being trained on simulated images, they demonstrate the performance on real data by analysing Hubble Space Telescope (HST) images of SL2S survey lenses. The main advantages of this data-driven approach are that it requires very little pre-processing compared to conventional model fitting techniques that include hundreds of fine-tuned parameters, and that it can handle thousands of strong lenses in a fast and automated way.

Unfortunately, and despite its widespread success in the aforementioned branches of astrophysics, machine learning methods have seen limited application in other domains like instrumentation. Some notable exceptions include the work by [Bonse \(2019\)](#) on atmosphere prediction for extreme adaptive optics (XAO). The temporal bandwidth error, the fact that AO systems are always slightly lagging behind the measured evolution of the atmospheric turbulence leads to imperfect corrections that can degrade imaging performance. In his work, [Bonse \(2019\)](#) demonstrates the use of Convolutional Long Short Term Memory Networks (ConvLSTMs) to predict the atmosphere based on a recent set of wavefront measurements in order to pre-compensate the wavefront before the measurement. Using real datasets from the SPHERE XAO system he shows that ConvLSTMs are able to make good predictions even under harsh conditions, outperforming classical algorithms. His results show that the improvement is so significant that the delay of the AO system with prediction could be up to three times longer until the error becomes identical with the classical correction.

A similar example of a successful application of machine learning to the domain of image-based wavefront sensing is the work by [Paine and Fienup \(2018\)](#). When the initial wavefront error is large, gradient-based techniques like Phase Diversity risk stagnating in local minima with inconsistent PSF results, if the initial guess for the aberration coefficients is not close enough to the true value. [Paine and Fienup \(2018\)](#) show that this becomes statistically more likely as both the RMS WFE and the number of aberration coefficients increase. To tackle this problem, they train a convolutional neural network to predict Zernike coefficients from simulated PSF images using a model of the JWST pupil. Then, they use those estimates to initialise the non-linear optimisation algorithms. Their work shows that, compared to the classical

approach of using a random guess for the wavefront map, using the predictions of the CNN leads to lower RMS WFE and facilitates the convergence of the algorithm.

Using an analogous approach, [Andersen et al. \(2019\)](#) investigate the possibility of using convolutional neural networks for wavefront sensing from a guide star image in astronomical telescopes. They train Google’s Inception v3 network to estimate the Zernike aberration coefficients from millions of simulated in-focus and out-of-focus PSF images with atmospheric phase screens. In a more recent paper, [Andersen et al. \(2020\)](#) show how the neural network works with sequential PSF-pairs that are temporally correlated suggesting closed-loop operation with a deformable mirror could be feasible. Their promising results suggest that, in the future, atmospheric aberrations could be sensed directly from guide star images without the need for an AO wavefront sensor.

The incoming generation of ELTs is revealing many new technical challenges due to both the demanding nature of their science goals and the unique characteristics of these novel systems, including: stringent requirements on raw contrast, or the fact that XAO systems operate in a regime of higher spatial resolution and faster loop speeds (>1000 Hz) than classical AO systems. Moreover, the need for enhanced angular resolution means high dimensional measurements are needed for turbulence prediction making the task more challenging. In some cases, traditional algorithms and techniques struggle to make the leap towards this new parameter space. Examples like [Bonse \(2019\)](#), [Paine and Fienup \(2018\)](#) and [Andersen et al. \(2019\)](#) show how techniques apparently foreign to these domains such as machine learning could potentially help bridge the performance gap and ensure the success of the next generation of telescopes.

One of the reasons behind the impressive success of machine learning techniques in astrophysics research is their ability to identify and recognise complex features in the available datasets, and leverage that information to make predictions or draw conclusions about the underlying physics.

Here we argue that the case of NCPA calibration is not much different from other tasks where these techniques have been applied successfully, and that they offer a potential to outperform traditional approaches to calibration.

In general terms, we can regard the task of NCPA calibration as a regression problem. Defining the complex pupil function $P(u, v)$ for an optical system as:

$$P(u, v) = \Pi(u, v) \exp(2\pi i \cdot \Phi(u, v)) \quad (5.1)$$

where $\Pi(u, v)$ is the arbitrary pupil aperture, i is the imaginary unit, and $\Phi(u, v)$ represents the wavefront in radians, typically defined as the optical path difference divided by the reference wavelength λ . The wavefront $\Phi(u, v)$ can be expressed as:

$$\Phi(u, v) = \sum_k^N a_k \phi_k(u, v) \quad (5.2)$$

a series expansion on an arbitrary basis of aberration functions $\phi_k(u, v)$, a typical choice being Zernike polynomials as they match the optical aberrations usually observed in realistic optical systems: astigmatism, defocus, coma, spherical aberration, etc. But as we will demonstrate in Section 5.4.3 other basis functions can be used without loss of generality.

At the image plane, the PSF for such a system would be directly related to the Fourier transform of the pupil function:

$$\text{PSF}(x, y) = \|\mathcal{F}(P)\|^2 \quad (5.3)$$

in the case of a more complex system like an IFS, we have already shown in Chapter 3 that this can be expressed as the convolution of several effects, but the essence is that the image at the focal plane is related to the pupil function by a form of non-linear mapping, which encodes the effects of aberrations within the optical system.

The regression problem in NCPA calibration boils down to estimating the values of the aberration coefficients a_k that describe the wavefront directly from focal plane images. This task is well suited for supervised learning algorithms, which fit a mapping from feature space (in this case PSF pixels) to some target values (the aberration coefficients) by training a neural network model based on labelled examples.

5.2 Methodology

The essence of the machine learning approach we propose for NCPA calibration is to train a data-driven network to directly infer the aberrations from PSF images. In traditional model fitting approaches such as Phase Diversity (see Section 2.7), one has to define a mathematical model that describes the underlying physical process and fit the observations to such model.

In contrast, machine learning proposes that the data itself is the best descriptor of the process and uses neural networks as a proxy for models. This has the advantage of removing the need for *a priori* knowledge of the systems behaviour, something which

can be hard to capture with mathematical models if the system is complex. On the other hand, extra care needs to be taken to ensure the data properly represents the way the system behaves, as missing dependencies would bias the predictions when applied to the real system.

Later we will see that in practical terms, it is hard to remove the *model* from this equation, because in the absence of hardware from HARMONI, the data must come from some form of model or simulation of the system. But since the machine learning approach relies solely on *data*, nothing stops us from using a whole set of models with uncertain parameters to generate such data and thus impart some robustness to the training, an advantage that traditional techniques lack.

In the following sections, I will thoroughly describe the methodology, including the types of neural networks we use and how these are trained.

5.2.1 Neural Networks

There has been a great deal of *hype* surrounding neural networks, making them seem magical and mysterious, but in reality, they are just non-linear statistical models (Friedman et al., 2001). The term ‘neural network’ has its origins in attempts to find mathematical representations of information processing in biological systems, but from the perspective of practical applications of pattern recognition, however, biological realism would impose entirely unnecessary constraints and is no longer followed (Bishop, 2006).

The best way to understand how neural networks work is by looking at their functional form. The basic neural network model is composed of a series of *layers*, each layer comprising a series of *neurons* that perform the following operations. First, a linear combination of the input variables $\mathbf{x} = [x_1, \dots, x_i, \dots, x_N]$ in the form:

$$a_j^1 = \sum_i w_{ji}^{(1)} x_i + w_{j0}^{(1)} \quad (5.4)$$

where the parameters $w_{ji}^{(1)}$ are known as the *weights* and $w_{j0}^{(1)}$ are the *biases*. The superscript (1) indicates that we are dealing with the 1st layer. The results of this linear combination a_j^1 are known as the *activation* of the layer. It is at this point that the non-linearity of the neural network is realised, in the form of a non-linear *activation function* $h(\cdot)$:

$$z_j^1 = h(a_j^1) \quad (5.5)$$




Figure 5.1: Network diagram for a two-layer neural network from Eq. 5.7. The input, hidden, and output variables are represented by nodes, and the weight parameters are represented by links between the nodes, in which the bias parameters are denoted by links coming from additional input and hidden variables x_0 and z_0 . Arrows denote the direction of information flow through the network during forward propagation. Originally published in [Bishop \(2006\)](#).

Generally, $h(\cdot)$ takes the form of a sigmoidal function such as the logistic sigmoid or hyperbolic tangent. The output of the first layer z_j^1 , serves as the input for the next layer:

$$a_j^2 = \sum_i w_{ji}^{(2)} z_i + w_{j0}^{(2)} \quad (5.6)$$

The original input \mathbf{x} (the image data) is propagated through the different layers (the architecture) until an output is produced.

$$y(\mathbf{x}, \mathbf{w}) = h \left(\sum_i w_{ji}^{(2)} h \left(\sum_i w_{ji}^{(1)} x_i + w_{j0}^{(1)} \right) + w_{j0}^{(2)} \right) \quad (5.7)$$

A network diagram representing this neural network architecture is shown Fig. 5.1. The nodes represent the different components of the input vector, layer variables and output vectors; every set of vertical nodes is a network *layer*. Lines connecting the nodes represent the *weights* w_{ji} .

As a result of this, the neural network model is simply a non-linear function from a set of input variables \mathbf{x} to a set of output variables \mathbf{y} controlled by a vector \mathbf{w}

of adjustable parameters (Bishop, 2006). In other words, the neural network model \mathbf{N} represents a non-linear mapping from input tensor \mathbf{x} to final layer output \mathbf{y} (i.e. prediction):

$$\mathbf{N}(\mathbf{x}|\mathbf{w}) \longrightarrow \mathbf{y} \quad (5.8)$$

In the specific case of NCPA calibration, the input tensor corresponds to a PSF image (or a datacube of multiple PSF images).

$$\mathbf{x} = [\text{PSF}(\Phi), \text{PSF}(\Phi + \epsilon Z_f), \dots] \quad (5.9)$$

where $\text{PSF}(\Phi + \epsilon Z_f)$ represents a defocused PSF needed to break the phase degeneracy. Assuming the wavefront Φ can be described as a series expansion on some generic basis functions ϕ_i :

$$\Phi = \sum_i a_i \phi_i \quad (5.10)$$

the output of the neural network model is the coefficients $\mathbf{a} = [a_1, \dots, a_i, \dots]$ that define the wavefront

$$\mathbf{y} = \mathbf{a} \quad (5.11)$$

so that the network model performs the task of *phase retrieval*:

$$\mathbf{N} \left(\text{PSF} \left(\sum a_i \phi_i \right), \dots | \mathbf{w} \right) \longrightarrow a_i \quad (5.12)$$

based on image data. This is a mapping from *feature* space, in the form of PSF images, to *aberration* space. The neural network model "looks" at a series of PSF images and predicts the aberration content.

In principle, there is no theoretical restriction on the type of output data \mathbf{y} and one could directly retrieve Φ if needed for practical reasons. But we have chosen to retrieve the coefficients a_i for the following reason. If the basis functions ϕ_i are global functions like the Zernike polynomials, the complete wavefront Φ can be described to arbitrary accuracy by a handful of coefficients. This simplifies the training of the neural network. In contrast, predicting Φ directly requires the network to output a pupil map $\Phi(p_x, p_y)$ of pixel values, which: (1) has a limit resolution fixed by the specific network architecture depending on the output layer shape, (2) the number of outputs is significantly larger than when one retrieves a_i if one is to ensure the wavefront is represented with sufficient resolution.

Now that we have clarified the basic architecture of a neural network model, as well as the nature of the data specific to the problem of NCPA calibration, we can move on to the key question of *how does the neural network actually work?* We have mentioned that the model performs a mapping $\mathbf{N}(\mathbf{x}|\mathbf{w}) \rightarrow \mathbf{y}$ that depends on a series of parameters \mathbf{w} , but how does one choose these parameters?

Selecting the network parameters \mathbf{w} is the process known as *training* the neural network model, and it is a critical step towards obtaining a model that produces meaningful predictions. Training implies adjusting the weights $w_{ji}^{(k)}$ and biases $w_{j0}^{(k)}$ of each k -th layer, by means of a numerical optimisation process that attempts to minimise a metric of performance, usually known as the loss function $\mathcal{L}(\mathbf{w})$.

In the case of regression problems, this loss function typically corresponds to a Mean Squared Error calculated over a set of *training examples* $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]$ and associated *labels* $\mathbf{T} = [t_1, \dots, t_n]$:

$$\mathcal{L}(\mathbf{w}) = \frac{1}{N} \sum_n^N \|y(\mathbf{x}_n, \mathbf{w}) - t_n\|^2 \quad (5.13)$$

where $y(\mathbf{x}_n, \mathbf{w})$ represents the predictions of the network for a given example \mathbf{x}_n and choice of weights \mathbf{w} , and t_n represents the ground truth associated to that example. In our case, $\mathbf{x}_n = [\text{PSF}(\Phi_n), \dots]$ and $t_n = \mathbf{a}_n$.

Typically, the loss function $\mathcal{L}(\mathbf{w})$ is optimised using some form of local optimisation algorithm like gradient descent, where the parameters of the network are updated at each iteration τ , using information on the gradient of the loss function:

$$\mathbf{w}_{\tau+1} = \mathbf{w}_\tau - \alpha \frac{\partial \mathcal{L}(\mathbf{w})}{\partial \mathbf{w}_\tau} \quad (5.14)$$

where α is a hyperparameter usually known as the *learning rate*. The main idea behind training neural networks is to compare the predictions for a given set of parameters $y(\mathbf{x}_n, \mathbf{w})$ to the actual values t_n and adjust those parameters \mathbf{w} using information on how the changes in network parameters modify the predictions (the gradient); repeating this process until the network delivers accurate predictions, i.e., acceptably low $\mathcal{L}(\mathbf{w})$.

The information of **how changes in the network parameters \mathbf{w} modify the predictive error** is encoded in $\frac{\partial \mathcal{L}(\mathbf{w})}{\partial \mathbf{w}_\tau}$, and it is typically calculated using a procedure called *error back-propagation*. Here, we will briefly describe how it works.

The loss function $\mathcal{L}(\mathbf{w})$ can be interpreted as a sum over all training examples of individual error functions $E_n(\mathbf{w})$:

$$\mathcal{L}(\mathbf{w}) = \sum_n^N E_n(\mathbf{w}) \quad (5.15)$$

where:

$$E_n(\mathbf{w}) = \sum_k (y_{nk} - t_{nk})^2 \quad (5.16)$$

is the sum across the k -th coefficients $t_{nk} = a_k$ representing the wavefront Φ_n . Calculating the derivative of this error function with respect to the network weights is all one needs to calculate the gradient of the loss function. Using the chain rule for partial derivatives:

$$\frac{\partial E_n}{\partial w_{ji}} = \frac{\partial E_n}{\partial a_j} \frac{\partial a_j}{\partial w_{ji}} \quad (5.17)$$

introducing the notation δ_j , commonly known as the *errors*:

$$\frac{\partial E_n}{\partial a_j} \equiv \delta_j \quad (5.18)$$

and recalling that the activating for the j -th network is given by $a_j = \sum_i w_{ji} z_i$, we can write:

$$\frac{\partial E_n}{\partial w_{ji}} = \delta_j z_i \quad (5.19)$$

which tells us that the error derivative depends on δ , the value at the *output* end of the layer, and z_i the value at the *input* of the layer. The process of obtaining z_i for all layers is often known as **forward propagation**, as it involves starting with the input \mathbf{x} and successively calculating the activations of each layer across the complete network. Calculating the δ_j terms requires a bit more work. Expanding the formula to:

$$\delta_j \equiv \frac{\partial E_n}{\partial a_j} = \sum_l \frac{\partial E_n}{\partial a_l} \frac{\partial a_l}{\partial a_j} \quad (5.20)$$

where l correspond to the all units that the j -th unit is connected to, further down the network layers. In other words, we can express the changes in the error function due to changes in the activation of an intermediate layer j in terms of changes due

to deeper layers l times the changes in those layers as a consequence of varying the activation upstream. Since $a_l = \sum_l w_{lj}h(a_j)$, we can write:

$$\delta_j = \sum_l \delta_l w_{lj} h'(a_j) = h'(a_j) \sum_l w_{lj} \delta_l \quad (5.21)$$

which tells us that the error term δ_j for a previous layer can be calculated from the error terms in the following layer δ_l . Consequently, one can start from the output layer k where $\delta_k = y_{nk} - t_{nk}$ and **back-propagate** the errors through the network up until the entrance. It is this combination of **forward propagation** and **back propagation**, which is nothing more than a practical application of the chain rule, what allows us to calculate the error function gradient $\frac{\partial E_n}{\partial w_{ji}}$.

5.2.2 Convolutional Neural Networks

So far we have described the architecture of a feed-forward network that takes an input vector \mathbf{x} and produces a prediction \mathbf{y} . However, in many applications including NCPA calibration, the nature of the data is not entirely captured by this approach. By working on 1D arrays, it ignores a key property of images, which is that nearby pixels are more strongly correlated than more distant pixels (Bishop, 2006). Very valuable information can be extracted from local features that depend on small areas of the images.

Therefore, it would be desirable for the neural network model itself to be inherently invariant to certain transformations of the input (such as translation), and be receptive to local information. This is the basis for *convolutional neural networks*. Compared to the previous case, not much changes in terms of methodology and architecture, except for the fact that *weights* and *biases* are replaced by *convolutional filters*:

$$a_j^1 = w_{ji}^{(1)} * x_i + w_{j0}^{(1)} \quad (5.22)$$

composed of multiple processing layers to learn representations of data with multiple levels of abstraction. The ultimate goal remains intact: making some predictions given some input data. This is done by transforming the representation at one level (starting with the raw input) into a representation at a higher, slightly more abstract level (Lecun et al., 2015). The convolutional filters and non-linear activation functions present at each layer of the model perform the task of constructing these features of increasing abstraction, and then combine them to arrive at a prediction.




Figure 5.2: A classic convolutional neural network architecture: LeNet-5 for digit recognition. Originally published in [LeCun et al. \(1998\)](#). The network takes as input images of hand-written digits. The different convolutional layers apply a series of filters to the image, to construct a set of feature maps. A set of fully connected layers take these abstract features and output a vector which contains the probability of belonging to each digit class.

A classic example of a convolutional neural network architecture is LeNet-5 for digit recognition [LeCun et al. \(1998\)](#), see Fig. 5.2. The network receives as input an image with a hand-written digit. The first convolutional layers apply a series of filters to that input followed by non-linear activation functions, to construct some feature maps. Then, a series of fully-connected layers operate on those feature maps to produce a vector of predictions that contains the assigned probabilities of belonging to a given digit class (0-9).

To illustrate this with a real example, Fig. 5.3 shows the output of the first two convolutional layers from one of the network models we use to predict aberrations. Each sub-panel is the result of applying a given convolutional filter to the complete stack of outputs from the previous layer. After the first layer we can see some masking has been applied to areas of the PSF image, as well as some transformations. These are the features that the model constructs out of the raw input. The second layer creates increasingly abstract features that no longer resemble the PSF, but still encode the "relevant" data, as understood by the network.

The key aspect is that these features are not designed by humans: they are "learned" from data by exposing the model to a large training set containing pairs of PSF images $\text{PSF}(\Phi_n)$ and their associated wavefront description \mathbf{a}_n and using that information to adjust the \mathbf{w} weights of the convolutional filters that give rise to such features, according to the back-propagation algorithm.




Figure 5.3: Output of the first two convolutional layers of a calibration model. Image on the left is the input PSF with random aberrations. The first layer has 16 filters, the second layer has 8 filters. Each sub-panel is the result of applying a given convolutional filter to the complete stack of outputs from the previous layer

5.2.3 Training data

In the previous section, we mentioned that in order to adjust the weights \mathbf{w} to obtain meaningful predictions one has to use some *training data*. This approach is commonly known as **supervised** learning, because the learning task is done based on some labelled data for which one knows the ground truth. In our particular case, this implies that we need to have a dataset of PSF images with varying aberrations Φ_n and their associated aberration coefficients \mathbf{a}_n to train the network to recognise how the presence of these aberrations modifies the PSF.

The training data must contain, at the very least, an image of the PSF with aberrations $\text{PSF}(\Phi_n)$. But, in reality this is not sufficient to estimate Φ_n because of the existence of degeneracies. Some aberrations suffer from sign-ambiguity, meaning that the PSF for a positive aberration $+Z_i$ is identical to a PSF with the same aberration but opposite sign $-Z_i$:

$$\text{PSF}(+Z_i) = \text{PSF}(-Z_i) \quad (5.23)$$

This makes it impossible to estimate the sign of such aberrations from a single focal

plane intensity measurement. In order to break this degeneracy, techniques like Phase Diversity introduce additional information in the form of a PSF with a *diversity* phase ψ (typically a defocus):

$$\text{PSF}(\psi + Z_i) \neq \text{PSF}(\psi - Z_i) \quad (5.24)$$

Here, the same approach is followed, such that each training example \mathcal{T}_n becomes a datacube with at least two channels:

$$\mathcal{T}_n = [\text{PSF}(\Phi_n), \text{PSF}(\Phi_n + \psi), \dots] \quad (5.25)$$

I say *at least* because in Section 5.4.5 I explore the possibility of using the integral field spectroscopy data readily available in HARMONI to improve the accuracy of the calibration. Given the flexibility of convolutional neural networks, one can easily incorporate additional data, in this case PSF images at multiple wavelengths, an advantage compared to traditional techniques that rely on monochromatic data.

Finally, we have to ask ourselves where the training can be obtained. For reasons that will become apparent soon, ideally one would like to train the model on real PSF images from HARMONI at the telescope, i.e. as close to the real scenario in which the calibration network will be deployed. As this is obviously not possible at the moment, one has to rely on simulations of the HARMONI system. This can be PSF images simulated with the mathematical framework presented in Section 3.3.1 or Zemax POP simulations of IFU as shown in Section 3.3.2 if we need to incorporate image slicer effects. In addition, the wavefront maps Φ_n used for training, can be created with the End-to-End model of the HARMONI instrument to ensure they are as realistic as possible, as I will demonstrate in Section 6.3. But PSF images with slicer effects are very time-consuming to generate so in some sections of this chapter where the impact of slicer effects is not the main concern, or where I try to investigate other effects (for instance, how to use integral field spectroscopy data), a simpler Fourier approach will be used to generate the examples.

To clarify, in the following sections we will use training images that tend to share common characteristics. The PSF images are simulated using a pupil model that mimics the ESO's ELT with a central obscuration at $0.284R$, unless otherwise specified. Monochromatic data will be used at a reference wavelength of $1.5 \mu\text{m}$, except in Section 5.4.5 where data at different wavelengths is used or when specified. The HARMONI 4×4 mas spaxel scale will be frequently used throughout this chapter.

5.3 Proof of concept - Image Slicer Effects

As introduced in Section 2.4.2, one of the research questions of this DPhil thesis is whether or not one can *develop alternative techniques for NCPA calibration that take into account image slicer effects*.

Throughout Chapters 3 and 4 we have demonstrated that the presence of an image slicer can have a significant impact on the PSF, mainly in the form of fringes parallel to the orientation of the slicing mirrors. At the beginning of this Chapter, we have hypothesised that machine learning techniques could be a good alternative to overcome the potential limitations of having image slicer features in the PSF images, given their success in other applications that require identifying patterns or extracting features in data.

In this section, we will demonstrate the feasibility of such an approach by means of a quick proof of concept. We will train a neural network to estimate the aberrations coefficients from PSF images that contain image slicer effects.

The data for this experiment corresponds to simulated PSF images created with the mathematical framework of light propagation through image slicers discussed in Section 3.3.1. As a quick summary, E_f^k the electric field at the exit slit focal plane for the k -th slice of the image slicer is given by:

$$E_f^k = \mathcal{F}[\Pi_m^k] * [\Pi_s^k \cdot \mathcal{F}(P(\mathbf{u}))] \quad (5.26)$$

consequently, the section of the PSF corresponding to the k -th slice for a given wavefront error Φ_n is:

$$\text{PSF}^k(\Phi_n) = ||E_f^k||^2 \quad (5.27)$$

where:

$$P(\mathbf{u}) = \Pi \exp(2\pi i \Phi_n) \quad (5.28)$$

therefore, the PSF images used to train the network will capture the effects from the presence of an image slicer. Each training example contains two PSF images: the nominal image with Φ_n , and a defocused image with $\Phi_n + \epsilon Z_f$ that contains an additional Zernike defocus common to all examples, to break the phase degeneracy. Fig. 5.4 shows one of the training examples

For this experiment, the PSF images were sampled with ~ 0.44 mas pixels to ensure the slicer features are properly resolved, with a total of 100 pixels per image. We




Figure 5.4: Training example containing two channels: *nominal* $\text{PSF}(\Phi_n)$ and *defocused* $\text{PSF}(\Phi_n + \epsilon Z_f)$ used to train a neural network model to predict aberrations directly from PSF data. The PSF images contain image slicer effects to demonstrate that the model can discriminate between intrinsic slicer features and those due to the presence of aberrations.

assume a wavelength of $1.5 \mu\text{m}$, a diversity defocus of around $1/3\lambda$ PV. The wavefront maps Φ_n contain the first 15 Zernike polynomials (excluding piston and tilts). The architecture of the neural network model used here is comprised of 3 convolutional layers, followed by a fully-connected layer whose output corresponds to the aberration coefficients \mathbf{a}_n . The model is trained to minimise the mean squared error between the predictions $\tilde{\mathbf{a}}_n$ and the true coefficients \mathbf{a}_n used to generate the wavefront Φ_n . The training set contained 1,000 samples with varying levels of aberrations. After training, the model is tested on a new dataset containing 200 random samples.

Fig. 5.5 shows a comparison between the *predictions* and *ground truth* for each Zernike aberration across all samples in the test set. What we observe is a strong correlation between the actual aberrations and what the model predicts, which suggests the model has been successfully trained despite the presence of image slicer effects and can discriminate between the different Zernike aberrations. This shows that when confronted with PSF images containing varying level of aberrations, the model is capable of identifying which pixel features correspond to each Zernike polynomial and accurately estimate the underlying coefficient.

To quickly illustrate the performance of this neural network model, we compared the norm of the aberration coefficients before $\|\mathbf{a}_n\|$ and after predicting and correcting the aberrations $\|\mathbf{a}_n - \tilde{\mathbf{a}}_n\|$ in a single iteration: $\|\mathbf{a}_n\| = 214 \pm 42 \text{ nm}$, compared to

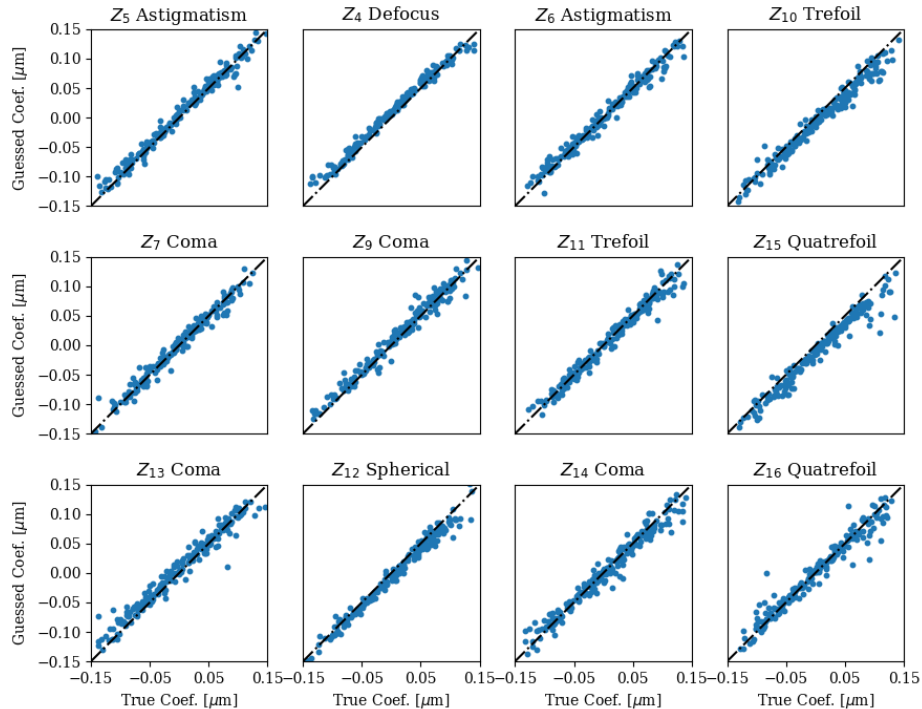


Figure 5.5: Comparison between *predictions* and *ground truth* for each Zernike aberration across all samples in the test set. A strong correlation indicates the model is recognising the different aberrations in the images and estimating the strength.

$\|\mathbf{a}_n - \tilde{\mathbf{a}}_n\| = 40 \pm 20$ nm. In terms of RMS wavefront error, $\sigma(\Phi_n) = 85 \pm 20$ nm and $\sigma(\Phi_n - \tilde{\Phi}_n) = 15 \pm 7$ nm, which shows that the predictions of the model improve the wavefront error.

No real effects such as readout noise, flat field calibration uncertainties and other sources of errors were included in this analysis because we wanted to show the essence of the methodology and prove its feasibility in its simplest form.

As a final note, this calibration approach can be used iteratively, by exposing the neural network model to a PSF image, applying a correction based on the network’s prediction, and feeding the updated PSF back. In some circumstances, iterating might be needed to reach a certain accuracy requirement. But this implies that the network must respond correctly to PSF images that contain a relatively wide range of aberrations: from poorly corrected cases, to almost perfect PSF images. In order to guarantee such performance, the training datasets that we use already span such aberration range, i.e. we linearly rescale the coefficients to expose the network to examples that range from well-corrected to highly-aberrated PSF images.

5.4 Results

The proof of concept from the previous section presented the general idea behind the machine learning approach to NCPA calibration and demonstrated its feasibility. But that opens a whole range of questions about the specific application of this technique in the context of HARMONI.

In this section, we will try to address some of these questions in a consistent manner, providing results and evidence that support our conclusions. These can be broadly classified into two categories: those pertaining to physical processes of the optical system (*physical*) and those that relate to the way the machine learning models operate (*algorithmic*).

The *physical* questions addressed include:

1. **Is it possible to calibrate NCPA even for the coarse spaxel scales?** HARMONI offers a range of spaxel scales (4×4 , 10×10 , 20×20 and 60×30). For the coarser scales above 4×4 , the sampling of the PSF is worse than Nyquist. Since the machine learning approach relies on PSF *features* to identify the presence of aberrations, there is a question of whether this can be applied to such coarse scales where such features are not properly resolved. In Section 5.4.2 we will first explore how the sampling affects the performance of the calibration algorithm, and then propose a way of calibrating the coarser scales that are severely affected by these effects.
2. **Are other wavefront basis functions more suitable than Zernike polynomials?** In Section 5.4.3 we explore the suitability of alternative wavefront definitions such as a basis of actuator commands, and compare that to Zernike polynomials. In addition, we investigate the impact of uncertainties in the actuator pupil mapping and Zernike misalignments on the performance of the calibration.
3. **Can we exploit the IFS data across multiple wavelengths?** Our neural network models can be easily extended to incorporate integral-field spectroscopy data across a whole bandpass; a unique trait compared to traditional techniques that rely on monochromatic data. In Section 5.4.5, we demonstrate how IFS data can improve the accuracy of the calibration.
4. **What is the optimum diversity phase?** The question of which aberration to use as diversity and its strength is addressed in Section 5.4.7. We investigate

the effect that the strength of defocus has on the performance of the calibration models.

The *algorithmic* questions addressed include:

1. **What is the optimum neural network architecture?** There is no *a priori* prescription on how to design a neural network model for a specific task, except for some rules of thumb and prior experience. In Section 5.4.1 we will explore this question and compare the calibration performance for several architecture parameters.
2. **What is the impact of uncertainties and real effects?** The characteristics of the PSF data can have an impact on the behaviour of the calibration model. Effects such as readout noise contamination and the field of view are addressed in Section 5.4.1. In Section 5.4.4, this analysis is extended to pixel scale variations, flat field calibration uncertainties and anamorphic magnifications.
3. **Can we make the neural network robust against model uncertainties?** In Section 5.4.4, we also demonstrate that the calibration neural networks can be made robust against these real effects by means of data augmentation, incorporating to the training set examples coming from multiple instances of optical models with uncertain parameters and other sources of error. Exposing the network to examples with uncertainty enables the network to successfully estimate aberrations for PSF images over a wide range of possible parameters, instead of relying on nominal values.
4. **Is it possible to interpret or understand the predictions of a neural network model?** One common criticism of machine learning techniques is their *black-box* behaviour and how that makes it difficult to interpret the predictions. In Section 5.4.6 we attempt to shed some light on this question, using a technique called *Shapley values* to highlight what features drive the predictions of the calibration models.

5.4.1 Architecture

Unfortunately, when it comes to the design of a neural network model, there is no prescribed method for finding the optimum architecture for a given problem. It is true that one can make educated guesses regarding which *type* of architecture to use, usually informed by the nature of the task and the data involved such as: using **recurrent neural networks** for tasks that exhibit temporal dynamic behaviour and require taking historical data into consideration (e.g. predicting atmospheric turbulence, [Bonse \(2019\)](#)), or using **autoencoders** for tasks that need a compressed representation of the original input. But all of these types of neural networks come with a series *hyper-parameters* such as the number of layers (depth of the network), the type of layers (convolutional, fully-connected), kernel size of convolutional filters, number of filters, the type of activation function, and many others; all of which ultimately determine the *architecture* and have to be defined before training.

For this particular task no rigorous method exists that can provide an *a priori* idea of what architecture parameters will work best for a specific research problem. Apart from some practical recommendations [Bengio \(2012\)](#), researchers are on their own and must resort to trial and error, as well as existing know-how to find the architecture with the smallest generalisation error; i.e. the one that performs best on unknown data.

As a starting point, many follow the steps of historically successful architectures that perform well in similar tasks to the one they are studying, and adjust some parameters to adapt it to their needs (see Fig. 5.2 for the popular LeNet-5 architecture). Others quite literally stand on the shoulders of giants, repurposing Google’s Inception v3 (an image classification network 48 layers deep with 24 million trainable parameters that can classify images into 1000 object categories) and applying it to wavefront sensing in astronomy ([Andersen et al., 2019](#)). In our case, we will use our own custom models using Keras ([Chollet and al., 2015](#)) with significantly fewer layers, that can be trained quickly (a few minutes) on a PC with a standard GPU.

In this section, we explore this issue of architecture design and investigate the effect of different hyper-parameters. For this purpose, we generated a training set of 10,000 PSF datacubes with random aberrations in the form of 52 Zernike polynomials. The PSF images are sampled with 4 mas pixels at a wavelength of 1.5 μm , covering a field of view of 64×64 pixels. We trained multiple calibration models varying the architecture hyper-parameters, and tested their performance on a test set of 1,000 PSF datacubes.

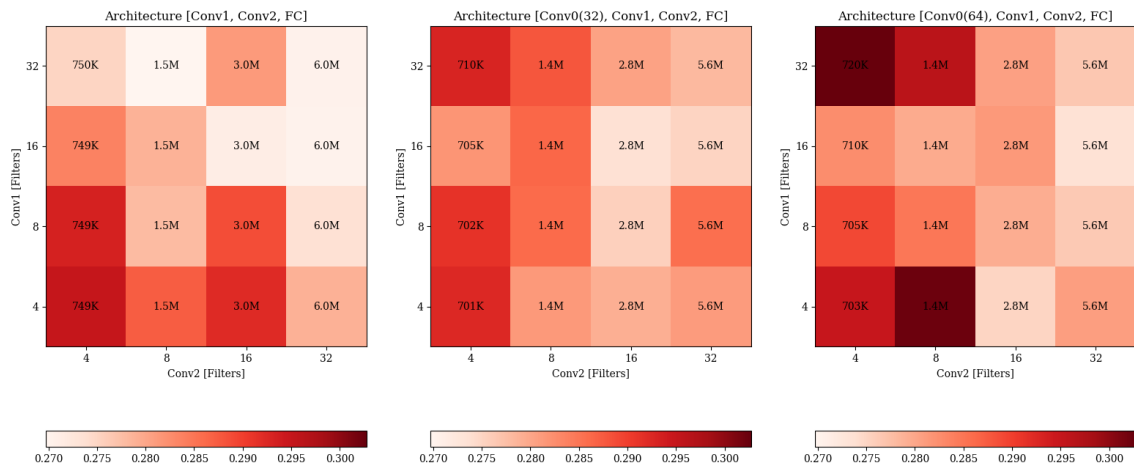


Figure 5.6: Impact of model architecture. Norm of the residual coefficients (*red* shade) for several model architectures, as function of the number of filters in two convolutional layers. The number in each box represents the number of trainable parameters of the network, a proxy for its complexity. Although more filters per layer improves performance, we see a degradation when adding extra layers (deeper models), as shown in the *middle* and *right* panels (models with 3 Conv. layers).

Convolution Layers. Depth and number of filters

Fig. 5.6 shows the average norm of the residual coefficients after calibration for each model (a smaller norm indicates better performance); the numbers in each box represent the number of trainable parameters for each model. The left panel represents models with 2 convolutional layers for varying numbers of filters. To investigate the impact of network depth, we also trained models with 3 convolutional layers, the first layer containing a fixed value of either 32 or 64 filters, see centre and right panels.

The first thing to note is that increasing the depth from 2 to 3 convolutional layers does not lead to an improvement in performance, suggesting shallower models are preferred. In addition, the results indicate that increasing the number of filters is beneficial for models with 2 layers. One interesting finding is that for the models with 3 convolutional layers, increasing the number of filters of the first layer (which barely modifies the total number of trainable parameters) worsens the performance.

Dense Layers. Depth and activation

So far the architectures we have presented contained only one dense layer after the convolutional layers. Here, we will investigate the impact of the number of dense layers and their activation function. We trained models with two convolutional layers (8 and 4 filters respectively), varying the number of dense layers from 1 to 4 as well as their activation function. It should be noted that the last dense layer does not

Activation	Dense Layers			
	1	2	3	4
None	0.287	0.300	0.318	0.328
ReLu	0.287	0.295	0.313	0.344
Sigmoid	0.287	0.322	0.338	0.374
Tanh	0.287	0.296	0.314	0.318

Table 5.1: Norm of the residual coefficients as a function of the number of dense layers and their activation function. The choice of activation function seems to have little influence in the result. However, adding more than 1 dense layer (deeper models) degrades the performance. This goes in line with our previous finding that fewer convolutional layers lead to improved performance.

have an activation. The results from Table 5.1 summarising the norm of the residual coefficients suggest that adding dense layers actually degrades the performance, no matter what activation function is used.

As a summary, in contrast with the widespread believe that models with many convolutional layers are necessary to achieve high performance, we do not see any advantage from using very *deep* models. In fact, the results suggest that the regression task of estimating PSF aberrations can be handled by relatively lightweight models with a few convolutional layers. A possible explanation for this disparity is the fact that deep learning models being used for complex tasks in fields like medicine have to deal with high resolution, multi-modal data (e.g. X-ray, magnetic resonance imaging, optical coherent tomography, echo-cardiograms, blood test data, family history), and identify incredibly abstract features and patterns in order to deliver a prediction, see [Topol \(2019\)](#) and references therein. Fortunately, the HARMONI PSF data is more homogeneous in terms of features and has significantly lower resolution and complexity than an X-ray scan, which could explain why lightweight models are competitive.

5.4.1.1 Readout noise. Hyper-parameter optimisation.

So far we have covered the performance of models trained and tested on *perfect* simulated PSF images, i.e. not accounting for noise effects. The results from the previous section clearly point towards shallow architectures with two convolutional layers followed by a single dense layer. In order to test whether this trend is also true for more realistic datasets, we decided to include readout noise effects in our simulations.

A more detailed account of different real effects is given in later sections, but here it will suffice to say that it is possible to make the calibration models relatively robust to noise contamination with *data augmentation*, a technique commonly used in machine learning to increase the number of available training examples by creating modified copies of existing data. In the context of readout noise, this entails creating multiple copies of the same PSF datacubes at different levels of signal-to-noise ratio (SNR). Readout noise is modelled here as additive Gaussian noise.

For this purpose, we created a dataset of 20,000 PSF datacubes, using the same simulation parameters as in the previous section, except for a field of view (image size) of 32×32 pixels. Results on the impact of *field of view* detailed in Section 5.4.1.2 will show that, in the presence of significant readout noise contamination, the size of the PSF images plays little role. However, reducing the image size can greatly reduce the total number of trainable parameters for a fixed number of convolutional layers by reducing the output size prior to the last dense layer, which speeds up the training.

In terms of data augmentation and the impact of readout noise, we created 3 copies of each PSF example in the 20,000-image training set, with random instances of contamination in the form of additive Gaussian noise at varying σ levels. In other words, during training the model is exposed to 3 copies of the same PSF with varying levels of contamination. Before the results are presented, it is important to clarify the definition of the signal-to-noise ratio (SNR) used here and its possible implications.

In a general case, the signal would be contaminated by different sources of noise such as the shot noise coming from both the object ($N_{obj} = \sqrt{F}$ proportional to the square root of the measured flux) and the sky background ($N_{sky} = \sqrt{N_{pix}b}$ related to the sky background b and the area of pixels covered by the object N_{pix}), as well as the readout noise associated with the detector ($N_{RON} = \sqrt{N_{pix}m\sigma^2}$ which adds in quadrature across the pixels and number of reads m , and is characterised by σ the number of photo-electrons per read). Adding these independent contributions in quadrature, we arrive at a total noise contribution N :

$$N = \sqrt{N_{obj}^2 + N_{sky}^2 + N_{RON}^2} = \sqrt{F + N_{pix}b + N_{pix}m\sigma^2} \quad (5.29)$$

in the context of internal calibration, we can discard the sky background contribution, which leads to an SNR of:

$$\text{SNR} = \frac{S}{N} = \frac{F}{\sqrt{F + N_{pix}m\sigma^2}} \quad (5.30)$$

Hyper-parameter	Search Space	#1 Model	#2 Model
Number of Conv. Layers	[2, 3, 4]	3	4
Conv. Layer Kernel Size	[3, 4, 5]	[3, 5, 4]	[3, 5, 4, 3]
Conv. Layer Filters	[8, 16, 32, 64]	[64, 16, 64]	[8, 32, 8, 32]
Conv. Layer Activation	[Relu, Tanh]	[Relu \times 3]	[Relu \times 3, Tanh]
Learning rate	$[10, 5, 1] \times 10^{-4}$	$5 \cdot 10^{-4}$	$5 \cdot 10^{-4}$

Table 5.2: Architecture hyper-parameter search space and best model results using the Keras-Tuner. The search space indicates the range of possible model architectures. The other two columns indicate the parameters of the models with top performance. For example, the best model had 3 Conv. layers with kernel sizes of 3, 5, and 4 pixels respectively, with 64, 16 and 64 filters per layer, all with ReLU activations.

Here, we have only considered readout noise, in the form of additive Gaussian noise and defined the SNR as the inverse of the RON σ , i.e. $\text{SNR}^* = 1/\sigma = F/(F\sigma)$, where the superscript $*$ indicates that the SNR definition used is unconventional. In other words, the readout in the images considered throughout this thesis is modelled as a Gaussian contamination added to PSF (which has been normalised to peak 1.0), and the SNR cited refers to the relative intensity of that noise.

It should be noted that object photon shot noise has not been considered, i.e. we are assuming a readout noise-dominated case and therefore the validity of the results is limited to such domain. In addition, the custom SNR^* definition allows for arbitrarily large values as it relates solely to the ratio of readout noise electrons to PSF peak. But, when including photon noise (as per Eq. 5.30), such values as we will see (500-750 SNR^*), are not possible due to the full well capacity of the detectors. For instance, given optimistic values for a HAWAII-2RG (well capacity of 80,000 to 100,000, and readout noise of 15-30 e^-), the SNR would be limited to $\simeq 10^5/10^{2.5} = 316$. In conclusion, when looking the values of SNR cited in the thesis, the reader should regard this a relative measure of readout noise contamination on the PSF, rather than the traditional SNR that considers object shot noise.

In the previous section we studied the impact of some architecture parameters, like the number of convolutional and dense layers. For this experiment we take a more rigorous approach, using the Keras-Tuner for hyper-parameter optimisation, with the HYPERBAND methodology (Li et al., 2018). This tool explores different hyper-parameter configurations, allocating more resources to promising configurations while eliminating poor ones, thus sampling the search space more efficiently.

The hyper-parameter search space is summarised in Table 5.2. We allow the tuner to explore calibration models with varying numbers of convolutional layers from 2 up




Figure 5.7: ROC curve showing the cumulative fraction of PSF cases with norm of residual coefficients smaller than a certain percentage of the original value, as a function of iterations. The Area Under the Curve (AUC) is a measure of performance, the ideal being 1.0 for a step function, i.e. all PSF cases with no residual aberrations.

to 4. For each layer, the kernel size, number of filters and layer activation function are allowed to vary; as well as the learning rate for the Adam optimiser.

The tuner ranks the models by their performance on a test set of 1,000 PSF datacubes, the top two are presented in Table 5.2 (their performance was essentially the same). Interestingly, when readout noise is present, slightly deeper models seem to be favoured: 3-4 layers, compared to a single layer in the noiseless case (Table 5.1).

When it comes to evaluating performance of machine learning models, a popular approach is to look at the metric known as receiving operating characteristic (ROC) curve. In the context of machine learning classifiers this curve compares the true positive rate (fraction of images identified correctly) to false positive rate (fraction of negated images incorrectly classified as positive) for different classification thresholds.

For the regression problem of estimating the Zernike coefficients this ROC metric is not directly applicable, but we can define something equivalent. At each iteration we can look at the norm of the residual coefficients (after applying the predicted correction):

$$\bar{r}_k = \bar{a}_k - \mathcal{M}(\mathcal{D}_k) \quad (5.31)$$

where \mathcal{D}_k represents the k -th PSF datacube, $\mathcal{M}(\mathcal{D}_k)$ are the predicted coefficients, and \bar{a}_k are the true coefficients. If we calculate the ratio of the norm of the residual coefficients at the i -th iteration to the norm of the true coefficients:

$$\eta_k^i = \frac{\|\bar{r}_k^i\|}{\|\bar{a}_k\|} \quad (5.32)$$




Figure 5.8: Strehl curve showing the cumulative fraction of PSF cases with norm of residual coefficients smaller than a certain percentage of the original value, as a function of iterations. Higher SNR leads to better performance, AUC closer to 1.

we can calculate the fraction of PSF cases for which the norm of the residual coefficients is below a certain threshold. This is shown in Fig. 5.7 for the top performing model for each SNR value of readout noise in the test set. After the first iteration, approximately 60% of the cases had a norm of the residual Zernike coefficients smaller than 60-70% of the original value, depending on the SNR. As the iterations progress, the residual norm decreases continuously.

In a similar fashion to that of a classifier ROC, the area under the curve (AUC) represents an aggregate measure of the model’s performance. A perfect model ($AUC = 1$) would have all cases with norm of 0% of the original one, with an ROC in the form of a step function.

We can also look at the Strehl ratio after each iteration to get a better idea of the performance, see Fig. 5.8. For a SNR of 750, almost all 1,000 cases have a Strehl ratio greater than 90% after 5 iterations. The number is slightly lower for SNR 500, but still all cases are above 85%. The model starts to underperform at SNR 250, but around 80% of the cases remain above 90% Strehl ratio.

What this shows, first of all, is that a high SNR is always desirable and will lead to an improvement in calibration performance. However, these results also indicate that it is possible to make the model quite robust to contamination from readout noise and still retain much of its predictive performance. We have also shown that the Keras-Tuner allows us to efficiently search for the optimum architecture given the characteristics of the PSF data, and find a set of hyper-parameters that deliver the best performance. This opens the possibility of adjusting our calibration models to realistic training data, once the instrument is delivered.

5.4.1.2 Field of View

One of the factors that could potentially affect the performance of calibration models is the *field of view*, or the number of pixels in the PSF arrays. According to Fourier theory, aberrations with higher spatial frequency modify the PSF further away from its core on the focal plane. Thus, one can hypothesise that a wider field of view might be needed to estimate higher order aberrations or, in other words, that a finite field of view might impose a limit on how many aberrations can be estimated accurately.

In order to investigate this, we ran the following experiment. We generated a training set of 10,000 PSF datacubes with random Zernike aberrations (55 polynomials were included). The reference wavelength is $1.5 \mu\text{m}$, and the PSF images were sampled with 4 mas pixels, covering a total of 128 pixels.

Then, we created four separate calibration models with the same network architecture (two CNN layers followed by a dense layer). The only difference between the models was the input shape of the first layer, which covered 128, 64, 32, and 16 pixels. Each model was trained on the same training set of PSF images, but the images were cropped accordingly to match the input shape, effectively varying the field of view from a maximum of 128 to a minimum of 16 pixels.

After training, we tested the performance of each model on a test set of 1,000 PSF datacubes, evaluating the standard deviation of the residual coefficients for each Zernike polynomial. This provides a metric for how accurate the predictions of each calibration model had been. We did this for multiple cases of SNR readout noise in the PSF images, to investigate whether the level noise had any impact on the field of view dependence.

The results are shown in Fig. 5.9. We have omitted the 128 pixels calibration model because its performance was as bad as the 16 pixels. This could be due to the fact that by keeping the model architecture constant (same number of layers), a wider field of view significantly increases the total number of trainable parameters, with the 128 pixels model having over 3 million parameters. This increases the training difficulty for a fixed number of examples (only 10,000) and could explain why a model with such wide field of view performs badly.

Leaving this aside, the performance for the 64, 32, and 16 pixels calibration models is particularly interesting. The first thing to note is that for the cleanest case (SNR 10^4), we see a clear improvement by increasing the field of view from 16 to 64 pixels. In addition, this improvement is stronger for higher order Zernike polynomials; almost no gain is seen for low order aberrations. This seems to support the hypothesis that a wider field of view helps predict high order Zernikes.

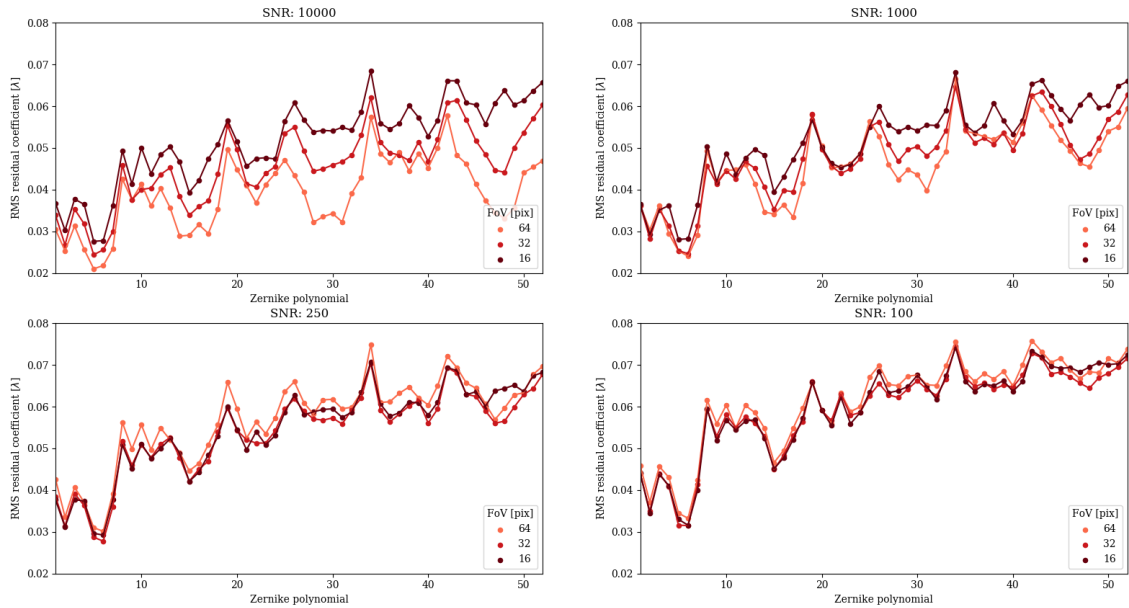


Figure 5.9: Impact of field of view (number of pixels in the PSF images). RMS residual coefficients for each Zernike polynomial, as a function of SNR readout noise. In the limit case of no noise (*top left*), extending the FoV improves the performance, specially for high order Zernikes, as the model has access to high spatial frequency features away from the PSF core.

However, this field of view dependence quickly disappears as we increase the levels of readout noise. In fact, for SNR 250 and SNR 100, the calibration model with 64 pixels has a slightly worse performance than the models with a narrower field of view. This change in behaviour is most likely caused by the calibration models relying more heavily on features close to the PSF core, which have a higher intrinsic intensity compared to the outer regions, when the PSF images are noisier (something that we explore in greater detail in Section 5.4.6.2).

Therefore, when readout noise is significant, a wider field of view does not bring any particular benefit to the calibration. Outer regions of the PSF contaminated with noise provide no meaningful information to the calibration models; if anything, they increase the number of trainable parameters making the training more challenging.

Lastly, we should point out that the residual error seems to increase with the Zernike radial order. This is not necessarily a problem because for these simulations we have used wavefront maps for which all Zernike coefficients were uniformly distributed (i.e. higher order Zernikes were just as strong as low order Zernikes). In reality, the wavefront maps will follow a power spectral density which decays with spatial frequency, with a higher prevalence of low order Zernike polynomials.

In conclusion, in this section we have investigated the impact of field of view on the performance of calibration models. The original hypothesis that increasing the field of view will lead to a better performance as it aids the calibration of higher order aberrations only applies to situations with substantial SNR. Contamination from readout noise quickly suppresses the benefits of extending the field of view to the point that a calibration model with 16 pixels might be preferable to one with 64 pixels, if only for its easy of training due to significantly fewer trainable parameters.

5.4.2 Sampling effects. Calibrating coarse scales

The calibration models we use to estimate NCPA are mainly based on the link between aberrations and PSF features. The whole purpose of the training of a calibration model is to find an optimum mapping from *feature* space, the images of the PSF, to *target* space, the aberration coefficients.

Therefore, one can expect the performance of these models to depend on how well such features are sampled. For instance, if the pixels are too coarse to resolve the core of the PSF, changes in intensity due to aberrations will be integrated into a single pixel, losing valuable information.

5.4.2.1 Influence of sampling

HARMONI offers a range of spaxel scales: 4×4 , 10×10 , 20×20 , and 60×30 milli-arcseconds. Beyond the 4×4 scale, the PSF is not Nyquist-sampled, which raises the question of whether machine learning calibration of NCPA is even feasible under such conditions. Thus, in this section, we will investigate the impact of sampling on the predictions of calibration models, with emphasis on the coarse scales of HARMONI. For this purpose, we ran the following experiment.

We began by creating a training set of 10,000 PSF datacubes with NCPA in the form of Zernike polynomials (33 terms were considered for this simulation). The reference wavelength was $1.5\ \mu\text{m}$, and the images were sampled at 2.0 mas, covering a field of view of 100 pixels. Then, we resampled the PSF images to generate 3 equivalent training sets, each sampled at one of the spaxel scales of HARMONI: 4, 10 and 20 mas respectively. In other words, we created 3 separate versions of the same original training set, but at different PSF sampling values. It should be noted that for this experiment, we have omitted any image slicer effects for two reasons: a matter of speed when generating the training sets (image slicer calculations are significantly more demanding than a simple Fourier PSF), and because at coarse scales the fringes caused by the image slicer will not be resolved at all.

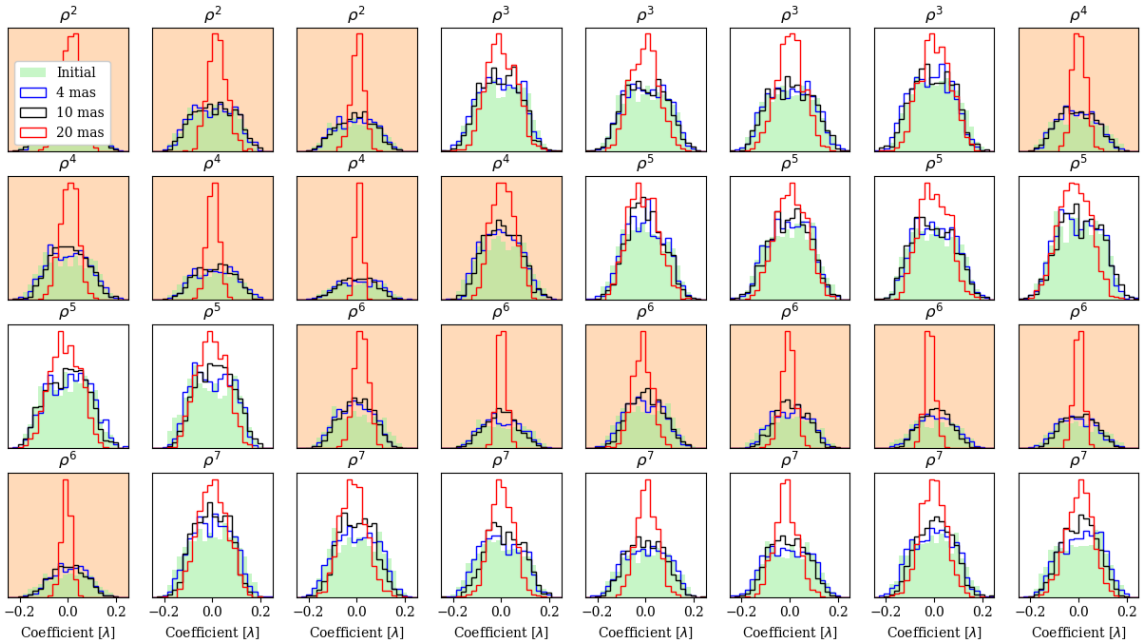


Figure 5.10: Impact of PSF sampling. Histograms of predicted coefficients for 32 Zernike polynomials (1,000 test samples), as function of PSF sampling (4 mas, 10 mas, 20 mas). The shaded cases correspond to Zernike polynomials that are degenerate with respect to the sign. The 20 mas scale shows poor performance, as the predictions are clustered around 0, an indication that the model struggles to estimate the aberrations.

For each training we defined a calibration model with the same architecture parameters: number of layers, filters per layer, output shape. The only difference was the input shape of the PSF datacubes, which varies depending on the sampling to ensure a constant field of view. We trained the three models (4, 10, and 20 mas) on their respective training images, while the *targets* (Zernike coefficients) remained the same for all models, since the original PSF datacubes were identical.

After training, we used each model to predict the Zernike coefficients for a test set containing 1,000 PSF datacubes. The results are presented in the form of histograms for each Zernike polynomial in Fig. 5.10. The green histogram represents the true coefficients that the models should retrieve. The blue, black, and red histograms represent the actual predictions for the calibration models trained on PSF datacubes at 4, 10, and 20 mas respectively. The shaded cases correspond to Zernike polynomials which are degenerate with respect to the sign, i.e. those for which the nominal PSF is identical no matter the sign of the aberration $\text{PSF}(+Z_i) = \text{PSF}(-Z_i)$. This correspond to those polynomials whose radial term $Z_\rho \sim \sum_n \rho^{2n}$ has even exponents.

The first thing to note is that for the 4 and 10 mas models, the predictions follow

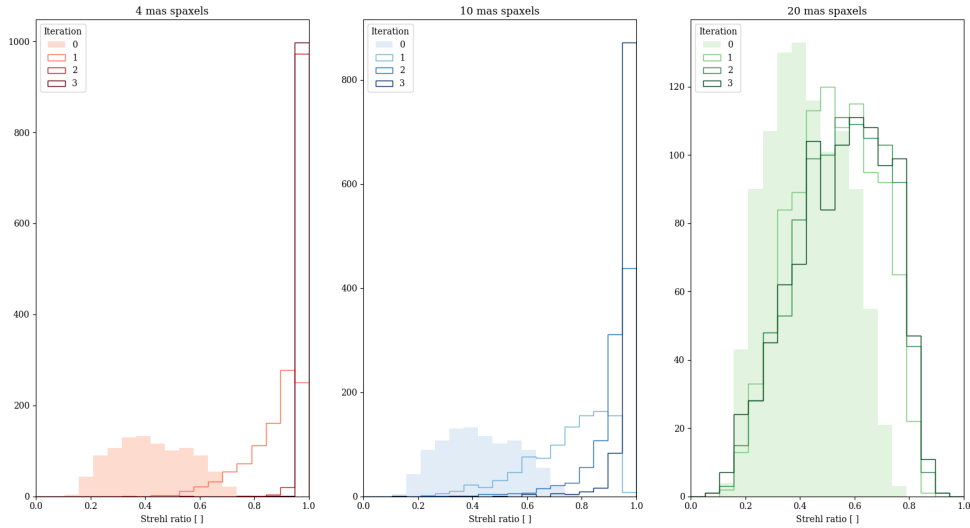


Figure 5.11: Distribution of Strehl ratios in the test set as a function of iterations for each case of spaxel scale: 4, 10 and 20 milli-arcseconds. Because of the difficulty of estimating the coefficients for degenerate Zernike polynomials in the 20 mas scale, the Strehl barely improves.

a similar distribution to that of the true coefficients. This suggests that these models provide reasonable estimates across all Zernike polynomials. In contrast, for the 20 mas calibration model, the situation is very different. While for the non-degenerate Zernike polynomials ρ^{2n+1} the predictions tend to follow the true histogram, in the case of the degenerate polynomials ρ^{2n} the predictions are concentrated at a narrow band around zero, which is what the calibration model would output if the decision process were unclear. These are the cases for which the neural network model relies heavily on information from the defocused channel, the one in charge of breaking the degeneracy.

This overabundance of values close to zero, and the narrow spread of the predictions indicate that the 20 mas calibration model struggles to identify some of the aberrations in the PSF datacubes, and heavily underestimates the coefficients. As a consequence of this, when a correction is applied, it will not lead to a significant improvement in Strehl, compared to the other two models. To test this, we ran the calibration for a total of 3 iterations, repeatedly applying wavefront corrections as provided by each calibration model.

The results in Fig. 5.11 show that for the 4 and 10 mas models the performance is effectively perfect, except for a few cases with Strehl around 80% in the 10 mas case. Unfortunately, for the 20 mas calibration model, the predictions only lead to a modest improvement, with some cases retaining very low Strehl and barely any

cases above 80%. It should be noted that the Strehl ratio is calculated here prior to down-sampling the PSF images to the specific spaxel scale; i.e. the spaxel size will not affect the result. The Strehl ratio is calculated on the 2.0 mas PSF images (prior to down-sampling) as the ratio of PSF peak intensity relative to the unaberrated case.

Interestingly, for the 20 mas scale we see a dependence on the initial Strehl ratio. For PSF images with relatively high Strehl (around 60%), the calibration model manages to deliver predictions that help improve the Strehl over a few iterations. In contrast, when the initial Strehl is around 20-30%, the predictions improve the Strehl ratio very marginally, or not all.

5.4.2.2 Calibrating at 20 mas

In the previous section we have shown that the standard calibration model struggles to improve the Strehl ratio when the PSF is sampled with coarse 20 mas spaxels, driven by its poor performance at estimating the degenerate Zernike polynomials ρ^{2n} .

Therefore, an alternative approach is needed to calibrate NCPA at coarse scales. Since the main issue is the low resolution at which the PSF is sampled, one potential solution would be to combine multiple exposures, after shifting the PSF by some fraction of a pixel, in a similar fashion to the algorithm Variable-Pixel Linear Reconstruction (also known as "Drizzle") see [Fruchter and Hook \(2002\)](#).

However, the CNN models used for calibration are particularly tailored to multi-dimensional datacubes, so we tried to include the shifted PSF images without any data pre-processing. Instead of using a training set containing 2 channels (nominal and defocused) of undersampled PSF images, we generated an extended training set \mathcal{D} that concatenates PSF images shifted by a fraction of a spaxel $\bar{f}_i p$.

$$\mathcal{D} = [\text{PSF}(\Phi, \bar{x}), \text{PSF}(\Phi + \psi, \bar{x}), \dots, \text{PSF}(\Phi, \bar{x} + \bar{f}_i p), \text{PSF}(\Phi + \psi, \bar{x} + \bar{f}_i p)] \quad (5.33)$$

In other words, the channels of \mathcal{D} contain different instances of the same PSF images (both nominal and defocus), shifted by some known fractions of spaxels in the x and y directions. The different channels remain *undersampled* at 20 mas, but the features will slightly change due to the PSF falling on different locations relative to the pixel boundaries.

The idea behind this approach is that the additional information coming from the extra channels would allow the model to disentangle the effect of the aberrations, without the need of enhancing the resolution. During training, the model will attempt to find the optimum set of filters that combine the different channels to produce a




Figure 5.12: Calibrating coarse scales (20 mas). Evolution of Strehl ratio after calibration versus the fraction of a spaxel by which the PSF images were shifted. *Blue* indicates additional channels shifted by f in the positive direction were used, *green* indicates both directions. Wavefronts contained 33 Zernike polynomials.

prediction. In a way, this could be understood as letting the model *learn* its own version of "Drizzle".

The first step to demonstrate the feasibility of this approach was to run the analysis with datasets consisting of PSF images shifted by a fixed f fraction of a 20 mas spaxel in both x, y direction. Fig. 5.12 shows the average Strehl $\mu \pm 1\sigma$ after calibration (3 iterations) as a function of f the fraction of a spaxel by which the images were shifted (33 Zernike polynomials were considered). For reference, we included the performance for the nominal model based on 2 channels of nominal and defocus PSF images at 20 mas (in *red*). The *blue* data correspond to models trained on channels shifted only in the positive direction x, y , whereas the *green* data represent models trained on channels shifted in both \pm directions x, y (i.e. with additional channels).

The first thing to note is that this shifted multi-channel approach leads to a significant gain in performance, compared to the nominal case. This gain appears to slightly depend on f the fraction of a pixel by which the channels were shifted, the optimum being around 3/10 and 6/10. Secondly, including shifts in both positive and negative x, y directions led to a further improvement, with Strehl ratios of as high as 90%.

Shift f	N_z Zernike polynomials				
	18	25	33	42	55
<i>None</i>	67.8 ± 20.1	55.9 ± 17.5	51.8 ± 16.6	49.6 ± 16.1	49.4 ± 14.9
$[\pm 1]$	97.5 ± 4.8	86.6 ± 12.3	82.1 ± 14.3	74.6 ± 13.4	72.0 ± 14.4
$[\pm 2, \pm 1]$	98.5 ± 1.5	89.1 ± 8.3	90.4 ± 8.7	80.4 ± 10.5	77.2 ± 11.7
$[\pm 3, \pm 2, \dots]$	98.5 ± 0.5	96.4 ± 4.2	91.8 ± 8.0	80.5 ± 10.0	76.7 ± 10.6
$[\pm 4, \pm 3, \dots]$	98.9 ± 0.2	95.8 ± 3.7	94.3 ± 6.2	83.4 ± 8.4	80.5 ± 10.1
$[\pm 5, \pm 4, \dots]$	99.4 ± 0.1	97.2 ± 2.4	95.1 ± 5.3	84.7 ± 9.1	82.1 ± 9.7

Table 5.3: Strehl after calibration as a function of the shifted PSF channels included in the training, versus the number of Zernike polynomials in the wavefront maps.. Nominal performance for the 20 mas scale is marked as *None*. Adding shifted channels leads to improved performance, but the returns quickly diminish, suggesting there is an optimum choice, given that compounding extra images increases the calibration effort.

In light of these results, we decided to investigate the *impact of the number of channels* on the final Strehl. Including examples of PSF images shifted by multiple fractions of a pixel might be beneficial. For this purpose, we compared the performance of models trained on PSF images shifted by f in $[\pm 1]$, $[\pm 2, \pm 1]$, $[\pm 3, \pm 2, \pm 1]$, etc. In addition, we calculated this as a function of the number of Zernike polynomials used to model the wavefront maps, to highlight any potential dependence on the radial order.

The results are shown in Table 5.3. Focusing first on the nominal calibration model, we see a consistent deterioration of the final Strehl ratio from $\sim 68\%$ to $\sim 50\%$ as the wavefront maps go from containing 18 to 55 Zernike polynomials. This is because as we include more radial orders, the total number of degenerate polynomials ρ^{2n} for which the network will produce a bad estimate increases, leaving a larger residual contribution to the wavefront error. It is also possible that the number of polynomials itself plays a role, since the model needs to estimate more coefficients from effectively the same information: the pixels of the PSF images.

A similar effect is seen across all cases of calibration models trained on shifted datacubes. The best performance is achieved when the wavefront contain fewer Zernike polynomials. In any case, including additional channels of shifted PSF images leads to a significant improvement compared to the nominal model, which mitigates the dependence on N_z . For instance, while the nominal model struggles to deliver Strehl ratios above 50% for 55 Zernike polynomials, the shifted models achieve 72-82% Strehl.

As far as the impact of the number of channels is concerned, this too depends on N_z . If the wavefront maps primarily contain low order Zernike polynomials ($N_z = 18$),

including channels shifted by $\pm 1/10$ of a 20 mas spaxel will be sufficient to achieve a perfect correction. As the performance increases with the number of channels, if we expect the wavefront maps to contain higher order aberrations, additional channels can be included to compensate that: $[\pm 2, \pm 1]$ already delivers 90% Strehl for $N_z = 33$.

In general, adding extra shifted channels leads to diminishing returns in Strehl. On the one hand, the improvement is marginal, specially when the Strehl is already high. On the other hand, additional channels require shifting the PSF and taking additional exposures, which makes the calibration in a real instrument more time consuming.

As a summary, we have investigated the impact of PSF sampling on the performance of calibration models. For the 4 and 10 mas scales, the nominal calibration models perform remarkably well. Unfortunately, the poor resolution of the 20 mas scale severely hinders the calibration. We have developed a solution to this issue by training a calibration model on datacubes containing multiple PSF images shifted by a fraction of the spaxel. In contrast to other traditional methods for undersampled images, we have not included any pre-processing of the PSF images. Instead, during training the calibration model finds the optimum set of filters to be applied to the shifted PSF images in order to deliver an estimate of the aberrations. We have shown that this alternative approach effectively solves the issues caused by undersampling the PSF, allowing the calibration of NCPA for the 20 mas scale in HARMONI with competitive performance to that of better sampled scales.

Another factor that could potentially help when calibrating coarser scales is the strength of the diversity defocus. A larger defocus would spread the PSF intensity along a wider area and thus, more pixels, possibly mitigating the effect of the coarse sampling.

5.4.3 Alternative wavefront definitions

The family of Zernike polynomials is well known for having some disadvantages. For high radial orders, the wavefront close to the edge of the pupil can diverge quickly, especially when using explicit expressions to evaluate the polynomials that lead to round-off errors (Forbes, 2010). In addition, after estimating the aberrations, one needs to map the Zernike wavefront to whatever basis functions the deformable mirror uses; this inevitably introduces an additional fitting error.

Another thing to note is that, in the context of machine learning for calibration, there is no reason to think that the Zernike family constitutes an optimum choice. It is hard to judge *a priori* what form of wavefront description leads to an optimum performance of the calibration models.

The neural network models work by associating PSF features $\text{PSF}(\Phi_n)$ to a set of coefficients a_i that describe the wavefront for that PSF $\Phi_n = \sum_i a_i \phi_i$. This choice ϕ_i is arbitrary and, in principle, there might exist a choice of basis functions $\Phi_n = \sum_i b_i \varphi_i$ for which the link between pixel features and coefficients is easier to disentangle. A good candidate for this would be a Fourier basis, as it could constitute a 1:1 link between PSF pixels (focal plane spatial frequencies) and the basis coefficients.

Taking all into consideration, we decided to investigate whether we could use alternative wavefront definitions that avoided (or at least mitigated) some of the caveats of Zernike polynomials, and at the same time characterise the behaviour of machine learning calibration models when using such alternative basis.

One alternative we decided to investigate was a wavefront model that mimics the behaviour of the deformable mirror that will be used to pre-compensate NCPA, and operates directly on actuator commands, hereby the **Actuator Model**. For this model, the wavefront Φ can be expressed as:

$$\Phi(x, y) = \sum_i^{N_{act}} \alpha_i \phi(x, y | x_i^c, y_i^c) = H^{act} \cdot \bar{\alpha} \quad (5.34)$$

where α_i represents the command for actuator i , and $\phi(x, y | x_i^c, y_i^c)$ is the influence function for actuator i , representing the wavefront departure at a given point (x, y) when the actuator centred at (x_i^c, y_i^c) is poked. The influence function $\phi(\cdot)$ depends on the architecture of the deformable mirror, but a simple approximation can be a Gaussian function representing the response of the surface of the deformable mirror when one actuator is activated. We can also write the wavefront calculation as the dot product of H^{act} a matrix containing the influence functions for all actuators and a vector of actuator commands $\bar{\alpha}$.




Figure 5.13: Alternative wavefront definitions. Random wavefront for the **Actuator Model**. The actuator centres are marked as black dots. The poking of the actuators creates a localised wavefront departure, compared to the Zernike basis which is *global*: a change in Zernike coefficient affects the complete wavefront.

5.4.3.1 Zernike polynomials vs. Actuator commands

In order to investigate which model is better suited for machine learning calibration, we ran the following experiment. Two identical machine learning calibration models were trained on the same set of PSF images; one being trained to estimate the coefficients of the Zernike basis, the other trained to estimate the coefficients of the basis of actuator commands.

For that purpose, we define a **Zernike Model** with varying number polynomials N_{zern} as a reference for the wavefront description. Then, we define an **Actuator Model** with a total of $N_{act} = 106$ actuators (see Fig.5.13). All wavefronts are calculated using the **Zernike Model**, and are then fitted using Least-Squares to the basis of actuator commands. Therefore, for an arbitrary wavefront Φ , we have a set of Zernike coefficients \mathbf{z} that describe such wavefront, and the equivalent $\bar{\alpha}$ actuator commands, that are accurate up to some small fitting error ϵ_{fit} .

$$\Phi = H^{zern} \cdot \mathbf{z} = H^{act} \cdot \bar{\alpha} + \epsilon_{fit} \quad (5.35)$$

Several examples of such LS fit are shown in Fig.5.14. It should be noted that we slightly oversized the **Zernike Model** normalisation radius to mitigate the divergence

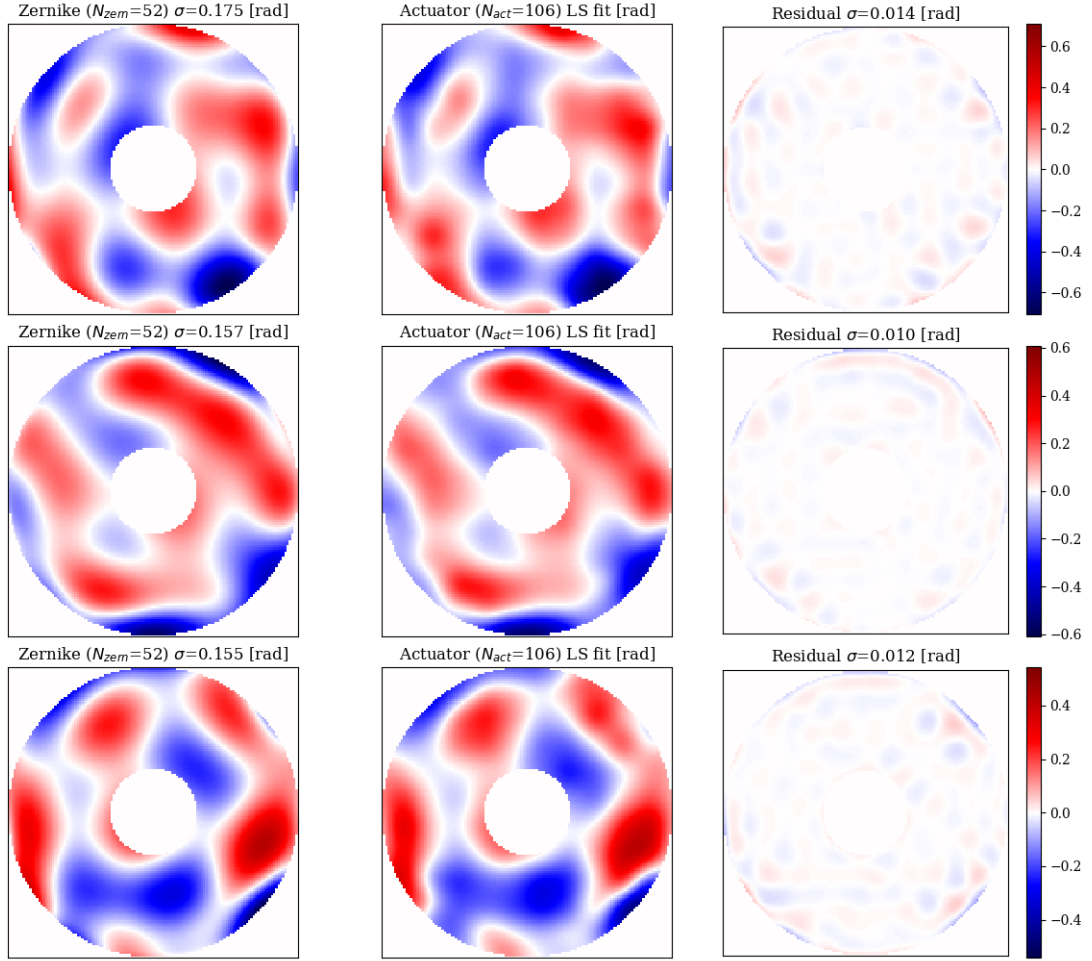


Figure 5.14: Zernike vs Actuator wavefront comparison. Random wavefront generated with the **Zernike Model** (*left*). The Least Squares fit on the **Actuator Model** for that wavefront (*middle*). Residual fit wavefront (*right*). The similarity indicates that both bases can be used to model the same wavefronts.

at the pupil edge, since that is the main source of fitting error for the **Actuator Model**. The LS fit is quite accurate, the RMS of the residual wavefront is typically around 6-8% of the original value, leaving out only some high spatial frequency terms $\epsilon\Phi$. This suggests that both the Zernike coefficients z_i and their associated LS actuator commands $\bar{\alpha}_i$ accurately represent the wavefronts in the PSF images.

$$\Phi = \sum_i z_i Z_i = \sum_i \bar{\alpha}_i \phi_i + \epsilon\Phi \quad (5.36)$$

For this experiment, we generated 10,000 training examples with their associated coefficients z_i and $\bar{\alpha}_i$, both in **Zernike Model** and **Actuator Model**. We use both in-focus and defocus PSF images at $1.5 \mu\text{m}$, sampled with 4 mas spaxels, covering

32×32 pixels. The procedure is summarised below, see Fig.5.15 for reference:

- We begin by generating thousands of random wavefronts $\Phi = \sum_i z_i Z_i$ using the **Zernike Model** for a given number of polynomials N_{zern} .
- Then, we fit the wavefronts to actuator commands $\bar{\alpha}$ of the **Actuator Model**, $\Phi = \sum_i \bar{\alpha}_i \phi_i + \epsilon\Phi$. These are the commands that more closely represent each given wavefront.
- Using the wavefronts, we generate datacubes of in-focus and defocused PSF images. These will be used to train both calibration models.
- We train two (almost) identical calibration models. One is trained to estimate \mathbf{z} Zernike coefficients, the other estimates $\bar{\alpha}$ actuator commands. The overall architecture and training of these CNN models is the same, except for the output shape of the last fully connected layer, which must be equal to N_{zern} and N_{act} respectively. Both models are exposed to the same PSF datacubes, the only difference being the *targets* or labels: either Zernike coefficients or actuator commands.
- Finally, we compare the residual wavefront error for each calibration model after removing the predicted aberrations.

We have seen that both basis are equivalent, at least to first order, when it comes to representing the wavefront Φ . Therefore, the key question here is **whether training the model to estimate the actuator commands $\bar{\alpha}$ leads to a better performance (lower wavefront error) than directly estimating the Zernike coefficients \mathbf{z}** . The CNN calibration models have no preference towards any particular basis and, in fact, they do not rely on a physical interpretation of a given basis for training (such as the Zernike coefficients representing optical aberrations) but rather, they simply create an ad-hoc mapping from feature space (the PSF images) and a target space that can potentially be an entirely arbitrary set of coefficients.

Therefore, it is hard to argue in favour of any wavefront basis prior to actually training the calibration models. With this agnostic approach, we ran the experiment across multiple cases of Zernike polynomials N_{zern} , while keeping the parameters of the Actuator model constant at $N_{act} = 106$. The results are summarised in Table 5.4.

First of all, the LS fit error represents the percentage of RMS wavefront after subtracting the Actuator wavefront from the reference Zernike wavefront, averaged across all 10,000 examples in the training set. In other words, for a model with

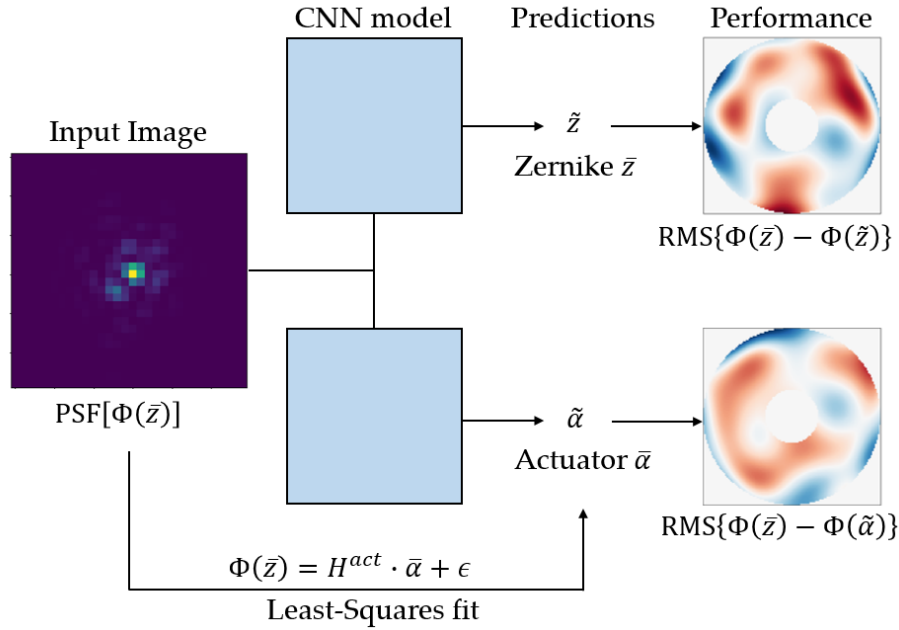


Figure 5.15: **Zernike Model** vs. **Actuator Model**. Two identical calibration models are trained on the same PSF images. One model is trained to estimate the Zernike coefficients used to generate the wavefront. The other model is trained to estimate the actuator commands representing the best fit of that same wavefront.

$N_{zern} = 33$ Zernike polynomials, if we subtract the wavefront given by the LS actuator commands, the residual RMS is around 7.1 % of the original.

With regards to performance, the Zernike and Actuator calibration models exhibit almost identical residual RMS, as shown in the last two columns of Table 5.4. An example is presented in Fig. 5.16. On the left column we can see the in-focus PSF and its associated wavefront; followed in the centre by the two predicted wavefront from the Zernike and Actuator calibration models; and the residual wavefronts on the right column. Qualitatively speaking, both calibration models produce an estimate of the aberrations that resembles the true wavefront; a good starting point for an iterative correction.

Another way of demonstrating the similarity in performance for the two calibration model is presented in Fig. 5.17, which shows the density distribution of initial RMS and residual RMS for the 1,000 PSF images in the test dataset for both calibration models.

These results seem to suggest that there is no substantial difference between the basis of actuator commands and the Zernike polynomials when it comes to estimating aberrations using a machine learning approach. We have demonstrated that it is

N_{zern}	LS fit error [%]	Initial RMS	Zernike RMS	Actuator RMS
33	7.1 ± 1.5	0.104 ± 0.023	0.020 ± 0.011	0.020 ± 0.010
42	7.3 ± 1.3	0.112 ± 0.025	0.033 ± 0.015	0.032 ± 0.014
52	7.8 ± 1.4	0.120 ± 0.026	0.042 ± 0.020	0.040 ± 0.018
63	9.6 ± 1.7	0.128 ± 0.027	0.054 ± 0.023	0.053 ± 0.022
75	12.1 ± 1.9	0.133 ± 0.029	0.065 ± 0.027	0.064 ± 0.026
88	15.0 ± 2.3	0.141 ± 0.030	0.078 ± 0.031	0.077 ± 0.030
102	17.1 ± 2.3	0.146 ± 0.030	0.086 ± 0.033	0.088 ± 0.032

Table 5.4: Comparison between the Zernike and the Actuator calibration models for varying number of Zernike coefficients. The Actuator model has a constant number of actuators ($N_{act} = 106$). Values represent $\mu \pm \sigma$ for a sample of 1,000 test PSF images. RMS units are in [radians].

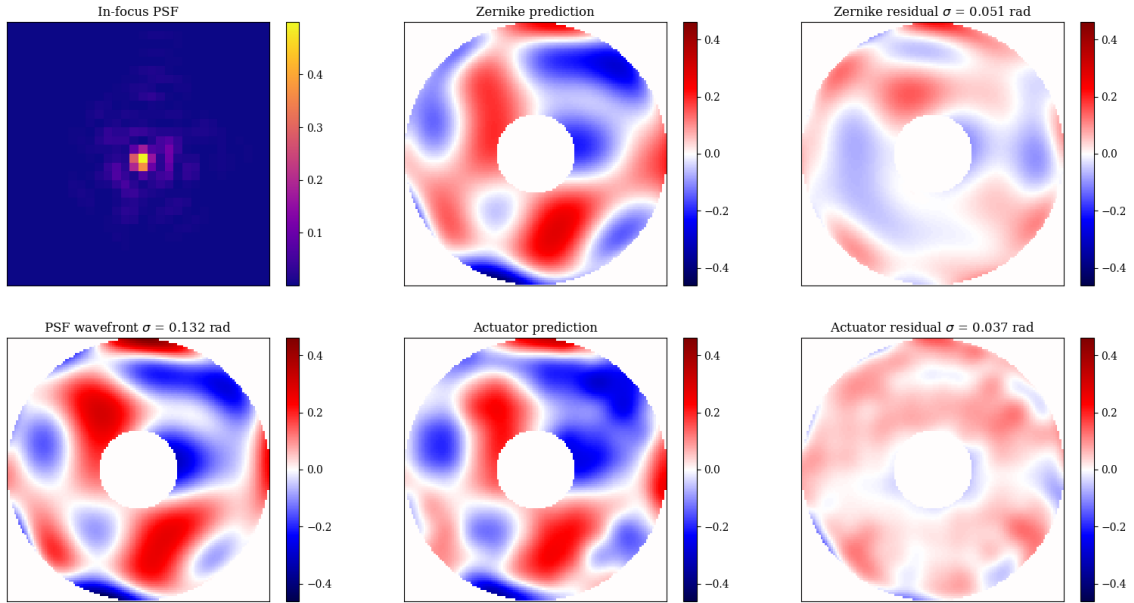


Figure 5.16: Comparison of Zernike ($N_{zern} = 52$) and Actuator calibration models ($N_{act} = 106$). *Left*: In-focus PSF and its associated wavefront. *Centre*: wavefronts predicted by the Zernike and Actuator calibration models for that PSF image. *Right*: residual wavefronts after removing the predictions.




Figure 5.17: Density distribution of RMS *before* and *after* calibration for the Zernike and Actuator calibration models ($N_{zern} = 42$, $N_{act} = 106$). The overlap for the two bases indicate that the performance is identical, no preference for any basis.

possible to train a calibration model to estimate the actuator commands necessary to directly compensate the aberrations with a deformable mirror. This represents an intrinsic advantage compared to the the basis of the Zernike polynomials, which after estimating the coefficients needs to perform a mapping from Zernike space to the influence functions of the deformable mirror, potentially adding a fitting error.

In addition, the Zernike polynomials tend to diverge at the pupil edge; whereas the Actuator model always produces smooth slow-varying wavefront maps. All this seems to indicate that for, practical reasons, the actuator commands would be a more robust basis than the Zernike polynomials.

5.4.3.2 Pupil mapping errors

One potential drawback of the basis of actuator commands could be the existence of a mismatch between the theoretical location of the actuator centres in the pupil plane and their true position, i.e. pupil mapping errors. Variations inside the HARMONI instrument could lead to changes in the actual positioning of the actuators on the pupil plane, which may vary by a significant fraction of the actuator spacing. Consequently, when calibrating NCPA, the wavefront corrections that we are actually applying may

be different from the theoretical one, leading to some residual error.

In order to investigate the potential impact of these effects on the quality of the calibration we ran the following experiment. First of all, we quantified the effect that a misalignment (δ_x, δ_y) of the actuator grid with respect to pupil plane has on both the wavefront error and the PSF. We shifted the grid of actuators by 10% of the actuator spacing Δ , and subtracted the two wavefronts: the one for the nominal grid Φ and the one with a shifted grid Ψ . The nominal wavefront Φ is given by:

$$\Phi = \sum_i \alpha_i \phi_i(x, y | x_i^c, y_i^c) \quad (5.37)$$

where $\phi_i(x, y | x_i^c, y_i^c)$ are the actuator influence functions. If the true location of the actuators on the pupil plane is misaligned by some unknown displacement (δ_x, δ_y) , the correction that we actually apply is:

$$\Psi = - \sum_i \alpha_i \phi_i' = - \sum_i \alpha_i \phi_i(x, y | x_i^c + \delta_x, y_i^c + \delta_y) \quad (5.38)$$

And thus there will be some residual error, even if the predictions α_i are perfect. The results are shown in Fig. 5.18, the *left* column shows Φ and its associated PSF(Φ). The *central* column shows the actual correction Ψ for a shift (δ_x, δ_y) with $||\delta|| = 0.1\Delta$ and $\arctan(\delta_y/\delta_x) = 129$ deg. The differences in wavefront correction $\Phi + \Psi$ due to this misalignment are around 1 order of magnitude smaller than the actual wavefront Φ . This translates into variations in PSF intensity approximately 3 orders of magnitude smaller than the core.

This seems to suggest that the effect on calibration performance should be relatively small. To properly quantify this, we trained a calibration model on 10,000 PSF datacubes with random aberrations, using an PSF model based on actuator commands (106 actuators). The grid of actuators used to generate the wavefronts was perfectly aligned with the pupil plane.

After training, to simulate the effect of pupil mapping errors, we calibrate the aberrations by applying corrections based on a misaligned actuator grid. We begin by showing the calibration model the PSF images with aberrations $\Phi(\mathbf{x})$, where \mathbf{x} is the nominal actuator grid.

$$\text{PSF}(\Phi(\mathbf{x})) \longrightarrow \bar{\alpha} \quad (5.39)$$

The calibration model produces a prediction $\bar{\alpha}$ for the actuator commands to correct the wavefront. However, due to the misalignment $\bar{\delta} = (\delta_x, \delta_y)$, the actual

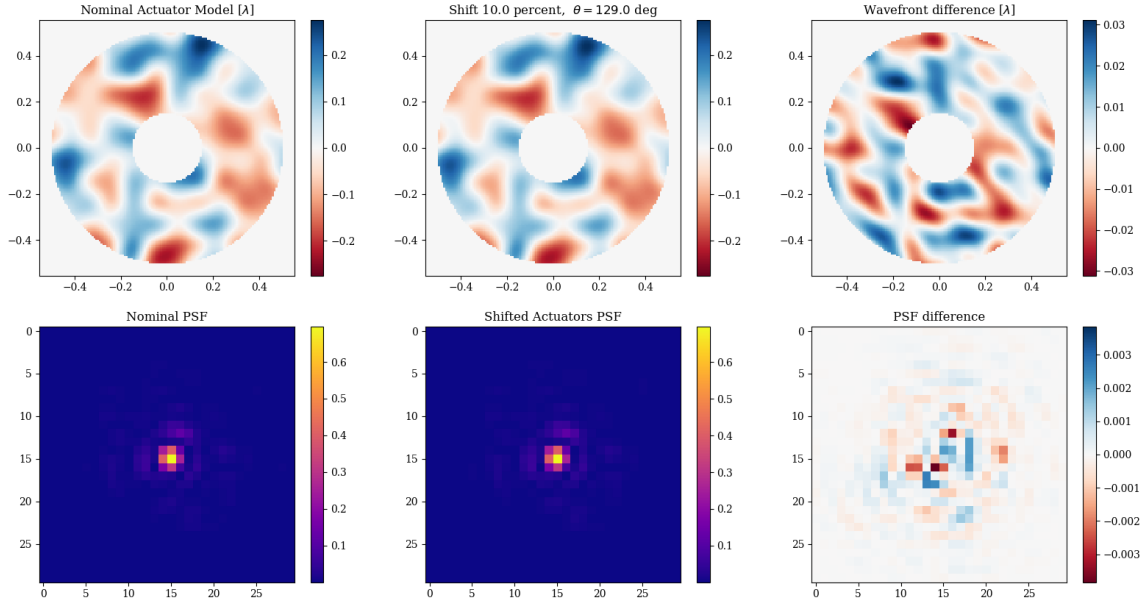


Figure 5.18: Impact of actuator pupil mapping errors. *Left*: nominal case, random wavefront. *Middle*: mapping error between the wavefront definition and the actual position of the actuator grid (shift 10% of actuator spacing). *Right*: residual wavefront error due to misaligned correction and induced PSF changes.

correction that we apply is slightly different from the theoretical one, the residual wavefront being:

$$\varphi = \Phi(\mathbf{x}) - \sum_i \bar{\alpha}_i \phi_i(x, y | \mathbf{x} + \bar{\delta}) \quad (5.40)$$

Consequently, when updating the PSF, it will inherit some residual error φ . We ran the calibration for a total of 4 iterations. The resulting Strehl ratios are shown in Fig. 5.19 as a function of $\|\delta\|$ the misalignment magnitude as a fraction of the actuator spacing Δ . The Strehl ratios represent the average across 1,000 test PSF images. The misalignment angle was randomly chosen for each sample.

The nominal model delivers an almost perfect correction with Strehl well above 95%. As we increase the misalignment, residual errors start to pile up and the effective Strehl ratio starts to degrade. However, the performance loss is relatively flat, with 35% misalignment still retaining 90% Strehl after 4 iterations. Interestingly, for a misalignment of 50% the actuator spacing, the Strehl decays starts to get worse after several iterations.

Despite having a minor / moderate effect on the final Strehl ratio, any loss in performance is undesirable so we decided to try some mitigation strategies. Considering the main effect of misalignment errors is introducing errors in the wavefront




Figure 5.19: Average Strehl ratio (1,000 test images) as a function of actuator grid misalignment $\epsilon = \delta/\Delta$ used to apply wavefront corrections. When the grid is misaligned by more than 1/4 of the actuator spacing, the performance quickly starts to degrade, dominated by residual errors in the correction.

corrections, we tried a strategy of *damped corrections* by which, rather than applying the correction as predicted by the calibration model, we apply a rescaled version, damped by a factor λ between (0, 1).

$$\varphi = \Phi(\mathbf{x}) - \lambda \sum_i \bar{\alpha}_i \phi_i(x, y | \mathbf{x} + \bar{\delta}) \quad (5.41)$$

Fig. 5.20 shows the evolution of Strehl for the *damped correction* strategy with a damping factor $\lambda = 0.5$. The effect of this approach for small misalignment δ is almost negligible, but it becomes beneficial when the misalignment is large, $\delta = 0.5$ delivers over 80% Strehl ratio compared to 70% when $\lambda = 1.0$.

In conclusion, in this section we investigated whether **pupil mapping errors in the form of actuator grid misalignment could put at risk the benefits of an actuator based approach**, compared to the basis of Zernike polynomials. We have shown that for mapping errors below 0.25–0.30 the actuator spacing, the loss in performance is almost negligible, which suggests that determining the true actuator positions at the pupil plane to within 1/4 of the spacing should be sufficient.

If for whatever reason this is not achievable, we can expect significant degradation in Strehl, with an ideal maximum of 0.70 at an uncertainty of 1/2 the spacing. A




Figure 5.20: Average Strehl ratio (1,000 test images) as a function of actuator grid misalignment used to apply wavefront corrections. Damped corrections strategy $\lambda = 0.5$. The damped corrections help mitigate the loss of performance due to misalignment and allow for larger errors in the grid mapping.

possible mitigation strategy is to apply a damped correction, rather than the complete prediction, which helps achieve Strehl ratios above 80% for an uncertainty of 1/2 of the actuator spacing.

5.4.3.3 Zernike misalignment

In the previous section, we looked at the effect that a misalignment of the actuator grid had on the performance of a calibration algorithm. In this section, we will analyse a similar effect for the basis of Zernike polynomials: a misalignment in the form of a rotation or clocking error between the theoretical wavefront basis and the actual positioning of the deformable mirror correction.

Fig. 5.21 shows the residual wavefront after subtracting misaligned Zernike polynomials for an angle θ of 15 degrees. Interestingly, the residual wavefront is given by the same polynomial but rotated and rescaled, where both the rotation angle and the peak-to-valley depend on the misalignment and the Zernike order. One can also calculate the ratio $r(\theta)$ between the residual RMS due to misalignment θ and the




Figure 5.21: Effects of reference system misalignment. *Upper row*: Zernike polynomials Z_i . *Middle row*: Reference system rotated by 15 degrees $Z_i(\theta)$. *Lower row*: Residual wavefront $Z_i - Z_i(\theta)$, between nominal and misaligned Zernike polynomials.

initial RMS, as a function of the angle.

$$r(\theta) = \frac{\sigma[\Phi - \Phi(\theta)]}{\sigma[\Phi(\theta)]} \quad (5.42)$$

Fig. 5.22 shows that ratio $r(\theta)$ for several Zernike polynomials. The number of peaks depends on the azimuth order m of the polynomial: coma ($m = 1$), astigmatism ($m = 2$), trefoil ($m = 3$) and quatrefoil ($m = 4$). The lower panel in Fig. 5.22 shows the residual RMS ratio for wavefront maps with random aberrations; the differences in number of peaks and zeros between the Zernike polynomials saturates the ratio at 1.25 - 1.50. For small misalignment angles θ , the ratio between residual RMS and original RMS becomes linear; and after around 20 degrees, the corrections become useless as the final RMS is in the same order as the original.

At this point, we decided to test the impact of the angular misalignment θ on the performance of the calibration algorithms, in a similar fashion as for the actuator case. A neural network model was trained to estimate the aberration coefficients of an ideal Zernike basis without any misalignment. At the time of testing, however, the corrections are applied using a basis that is rotated by θ , leading to some residual error. The results for the evolution of Strehl ratio with iterations as function of the




Figure 5.22: *Upper panel*: Ratio between residual RMS due to misalignment and initial RMS, as a function of the misalignment angle, for several Zernike polynomials. The number of peaks depends on the symmetry characteristic of each Zernike. *Lower panel*: Same ratio for wavefront maps containing random Zernike polynomials (52 terms considered). The behaviour is linear for a small misalignment but above a certain angle, the residual RMS is around 125-150% the original RMS.




Figure 5.23: Evolution of average Strehl ratio in the test set with iterations as a function of Zernike misalignment angle θ . *Left*: nominal correction. *Right*: damped correction where $\lambda = 0.75$ of the predictions are applied at each iteration. The 5-th and 95-th percentile range is also shown at every iteration for $\theta = 10$ degrees.

misalignment angle θ are shown in Fig. 5.23, for both a nominal correction and a damped correction of $\lambda = 0.75$ of the predictions.

For misalignments above 5 degrees, the performance quickly starts to degrade. Introducing a damped corrections helps mitigate the negative effect of the misalignment, deliver final Strehl ratios between 0.77 and 0.95 even for 10 degrees of misalignment, which is probably a very conservative estimate on the level of uncertainty.

In conclusion, misalignments due to uncertainties in the actual position of the basis functions used to correct the wavefront has a relatively small impact on the residual calibration error, both in the case of actuator commands and Zernike polynomials.

5.4.4 Robust calibration - Impact of real effects

The data used to train the neural network models comes from simulations of the PSF. Given the difficulties that machine learning techniques tend to have when generalising outside the training dataset, one must take care to ensure that the training data is as realistic as possible and accurately represents the test data. This issue is already present in traditional techniques for NCPA calibration like Phase Diversity, where inaccuracies or uncertainties in the modelling of the optical system can negatively affect the quality of the predictions, see [Blanc et al. \(2003\)](#); [Jolissaint et al. \(2012\)](#).

We have already explored this issue in previous sections of this chapter. For instance, Section 5.4.1 investigated the effect of readout noise contamination, while Section 5.4.3 looked at the impact of pupil mapping errors in the actuator basis and misalignments of the Zernike basis. In this section, we will take this analysis one step further and look at effects such as: flat field calibration uncertainties, pixel sampling errors, and anamorphic magnification errors.

In addition, we will demonstrate that calibration neural networks can be made robust against these effects by means of data augmentation. By exposing the networks to examples generated with multiple instances of optical models with uncertain parameters, the networks can be trained to incorporate this uncertainty into the mapping between features and aberrations.

5.4.4.1 Flat field calibration uncertainties

The behaviour of a detector is never perfect; variations in quantum efficiency and other factors can lead to changes in pixel-to-pixel sensitivity. Significant efforts are always devoted to calibrating and compensating this and other detector effects such as dark current, but some residual uncertainty in pixel sensitivity will always remain.

This residual flat field error could potentially hinder the performance of calibration models by rescaling the PSF features in an unpredictable way, causing the models to mistake changes in pixel sensitivity for spurious aberrations. In order to assess whether this could pose a threat to the calibration and provide a lower bound for the quality of the flat field correction, we ran the following experiment.

First of all, we began by training a calibration model with a dataset of 20,000 PSF images with random aberrations in the form of Zernike polynomials (33 terms were considered), sampled at 4 mas at 1.5 μm reference wavelength. The field of view covered 32×32 pixels. After training the model, we tested its performance on a new dataset containing 10,000 PSF images, this time with flat field errors $\text{PSF}_{ff}(\Phi_n)$. We




Figure 5.24: Impact of flat field uncertainties on norm of residual coefficients. *Left*: distribution of norms of residual coefficients after correction, nominal model (no flat field errors). *Right*: Changes in residual norm due to flat field errors and Gaussian fit. The effect of flat field errors appears small.

model the flat field as a uniform distribution of pixel sensitivities centred at 1.0 with range $\pm\Delta$, which rescales the PSF features, $p_{xy} \in \mathcal{U}(1 - \Delta, 1 + \Delta)$.

$$\text{PSF}_{ff}(\Phi_n) = p_{xy} \cdot \text{PSF}(\Phi_n) \quad (5.43)$$

Each PSF datacube in the test was rescaled by a random instance of flat field (both nominal and defocus channels use the same instance). We compare the predictions of the calibration model to those of the *clean* test set to see the effect of the flat field rescaling, the results are shown in Fig. 5.24. The left panel (*blue*) represents the nominal performance of the calibration model, in the form of the distribution of norm of residual coefficients after correction. In other words, in the absence of any flat field effects the norm of the residual coefficients follows that distribution across the 10,000 PSF images in the test set.

The three red panels indicate changes in that residual norm when we include flat field errors for several cases of Δ , the flat field uncertainty. Since the PSF is quite coarsely sampled and the intensity is concentrated in a small fraction of pixels, in order to mitigate potential outliers, each of the 10,000 PSF datacubes is exposed to 5 random instances of flat field. To put it concisely, the panels in Fig. 5.24 show the distribution of *variations* in residual norm after calibration for test PSF images containing flat field errors compared to the *clean* case.

Positive variations in norm indicate that, due to the presence of flat field uncertainties, the calibration model delivered a prediction that was worse than what it delivered when exposed to the clean PSF. Surprisingly, the changes in residual norm are not always positive (worse performance), but appear to be normally distributed




Figure 5.25: Impact of flat field uncertainties on Strehl ratio. *Left*: Strehl ratio before and after correction for the nominal model (no flat field errors). *Right*: variations in Strehl ratio after correction due to flat field effects as a function of initial Strehl. A strong dependence on Strehl is seen, with low Strehl ratio cases being disproportionately affected by flat field uncertainties.

around zero, suggesting that some examples have better performance when distorted by flat field uncertainties.

In any case, these variations are small in comparison to the nominal residual norm, which suggests that the calibration model predictions are barely affected. For instance, with a value of Δ of ± 5 per cent (which could be considered quite a poor flat field calibration), the 2σ range is $5.6 \cdot 10^{-2}$. For an average residual norm of around 0.2λ , that represents less than 3% change. As some cases actually have *negative* variations in residual norm, the 2σ range suggests that $\sim 97.5\%$ of cases had an increase in residual norm due to flat field uncertainties smaller than 3% with respect to the nominal value.

A caveat of these results is that they represent the aggregate across all test PSF images. This is a problem because within the test set, the Strehl ratio varies greatly from 15% to around 75%. The initial Strehl ratio could be a relevant factor when it comes to sensitivity to flat field inaccuracies because it changes the distribution of intensity across the detector: highly aberrated PSFs spread the intensity over many pixels, whereas a PSF with high Strehl ratio has a narrow core concentrated over few pixels.

To test this hypothesis, we ordered the PSF images in the test set according to the initial Strehl ratio, see Fig. 5.25. The left panel compares the initial Strehl ratio to the Strehl ratio after a single correction for the nominal model (no flat field uncertainties). The right panels represent *variations* over that final Strehl ratio when we include flat field effects. For example, for the case of Δ of 5% flat field uncertainty, almost all cases have a Strehl after correction within $\pm 2\%$ of the nominal value, most

being below $\pm 1\%$. Interestingly, we see a strong dependence on the initial Strehl ratio, with lower Strehl ratios having a higher degradation in performance, which corroborates our hypothesis. When the Strehl ratio is above 50% we see very little difference due to flat field uncertainties.

Overall, the results seem to indicate that flat field calibration inaccuracies have a very limited impact on the behaviour of calibration models. Knowing the true pixel sensitivities to within a few percent should be sufficient to ensure an accurate NCPA prediction. In addition, access to IFS data in HARMONI offers a potential mitigation strategy: while the flat field errors affect all areas of the detector, the PSF is recorded at different areas for different wavelengths. This allows us to average out flat field errors by incorporating multiple wavelength channels as we will demonstrate in Section 5.4.5.

5.4.4.2 Pixel sampling errors

Typically, the machine learning calibration models are trained with example PSFs at a fixed spaxel scale, corresponding to the nominal scales for HARMONI (usually 4 mas). Nevertheless, when dealing with real PSF images from the telescope, uncertainties in optical power, alignment tolerances and data reduction errors can cause slight variations in the actual sampling.

If this is not accounted for, it can potentially affect the performance of the calibration model, as the induced PSF scaling can be mistaken by spurious aberrations. To test the impact of unexpected variations in sampling on the performance of the calibration, we ran the following analysis.

We trained a calibration model (actuator based) with 10,000 example PSF images, with random aberrations, at a wavelength of $1.5 \mu\text{m}$ sampled with 4 mas spaxels. After training, we generated an additional dataset with PSF images at a different spaxel scale $4(1 + \epsilon)$, where ϵ represents relative deviations from the nominal scale.

This new dataset was used to test the performance of the model, and compare it to the nominal case. In both cases, we ran 3 iterations of the calibration, updating the PSF images using the predicted aberrations in order to have information on the final wavefront error. The results are shown in Fig. 5.26 for a spaxel error of -5% (3.8 mas). Blue represents the nominal model, trained and tested on 4 mas PSF images. Red represents the same model tested on 3.8 mas PSF images. Darker colour indicates later iterations.

It can be seen that, when we test the model on PSF images that do not match the training set because of sampling errors, the performance of the calibration model




Figure 5.26: Evolution of RMS wavefront with 3 iterations. *Blue*: nominal model trained and tested on PSFs at 4.0 mas. *Red*: nominal model tested on PSFs with a -5% scale error (3.8 mas). The scale error degrades the performance of the model, which struggles to reach 50 nm RMS after 3 iterations, compared to the nominal case.

is degraded. After several iterations, the nominal model converges to less than 50 nm residual RMS wavefront. However, when testing on the new PSF images, the model struggles to bring the RMS down and converges to approximately 62.5 nm.

To properly quantify this effect, we repeated the analysis for multiple values of the sampling error ϵ and compared the RMS wavefront at the end of the calibration. The results are shown in Fig. 5.27.

One important thing to notice is that the performance degradation is not symmetric with respect to the spaxel error. Going towards coarser scales has a substantially larger impact on the predictions than going towards finer scales. We suspect that this is because at finer scales, even though the PSF scaling is off, the features are better resolved, which probably helps the model identify the aberrations. In contrast, when both the scaling is off and the sampling is poorer than the nominal case, the calibration model seems to struggle more.

This means that the nominal model can accommodate some sampling errors as long as they are towards finer scales. For example, a 2.5% error (3.8 mas) leads to almost no degradation in performance.

One of the reasons behind the performance degradation from sampling errors is that machine learning techniques tend to struggle to generalise outside their training




Figure 5.27: RMS WFE post-calibration (3 iterations) as a function of sampling error. For errors above 2.5% the performance severely degrades, as the model struggles with PSF images that are dissimilar to those in the training. The response is not symmetric as negative errors mean finer sampling (< 4 mas), which mitigates the problem.

sets. In other words, they perform badly when tested on examples that differ from the ones used to train the model. A potential solution to this problem is to use data augmentation techniques to include additional examples in the training, which inform the models on the range of variability that one can expect in realistic conditions.

For example, in the case of sampling errors, we can generate multiple copies of the same PSF image, with slight variations on spaxel scale $p_0 + \epsilon$, but the same associated label (the aberration coefficients a_n). That way, we force the model to realise that a particular wavefront can take multiple forms in PSF feature space and that it should aim at estimating the aberrations in a robust way.

$$\begin{aligned}
 \mathcal{T}_1 &= \{[\text{PSF}(\Phi_n, |p_0), \dots], a_n\} \\
 \mathcal{T}_2 &= \{[\text{PSF}(\Phi_n, |p_0 + \epsilon), \dots], a_n\} \\
 &\dots \\
 \mathcal{T}_N &= \{[\text{PSF}(\Phi_n, |p_0 + N\epsilon), \dots], a_n\}
 \end{aligned} \tag{5.44}$$

In order to test this approach, we trained a **robust** calibration model using a training set that includes copies of PSF images with the same wavefront, but varying scale errors $[0.0, \pm 2.5, \pm 5.0, \pm 7.5, \pm 10.0]$ % at the nominal 4 mas scale. To keep the

training consistent, we used the same number of examples in aberration space as for the nominal model: 10,000 PSF images but copied 9 times, one for each scale error values, for a total of 90,000 images.

To test the performance of this robust calibration model we used three separate datasets, of 9 x 1000 images each:

- **Training Points:** these datasets includes random PSF examples with scale errors equal to the discrete values used to train the calibration model.
- **Within Range:** the model was trained on examples with fixed values of scale error ranging $\pm 10\%$. Even if the model performs well for those discrete values, it not entirely clear how the model will perform for PSF images with arbitrary scale errors within that range. To investigate this, we generated datasets with scale errors inside the $\pm 10\%$ range, but at different discrete values to those used for training the model.
- **Outside Range:** to investigate how this robust model responds to images outside its training set, we created other datasets with errors larger than $\pm 10\%$.

The results for all three types of test datasets are shown in Fig. 5.28. The *orange* dots indicate the RMS residual after calibration when the robust model is tested on PSF images with the same scale error values as the training (**Training Points** datasets). The *green* squares correspond to the **Within Range** datasets and the *red* diamonds correspond to the **Outside Range** datasets.

The first thing to notice is that when we use data augmentation and show examples of how scale variations affect the PSF, the model becomes robust against such effects. As a result of this, when tested on examples in the range $\pm 10\%$ at the nominal 4 mas scale, we observe a flat performance response across scales (*orange* dots).

Secondly, when we test the model against images within that $\pm 10\%$ range but at different scale errors to those used for training the model, we find no variations in performance (*green* squares). This suggests that, in a sense, the model can interpolate between the scale error points it has been trained on. This greatly simplifies the training because it allows us to generate PSF images for discrete values of scale errors and expect the model to perform well across the continuous range.

Thirdly, as one moves away from the nominal training range, the model stops generalising properly, and the performance starts to deteriorate (*red* diamonds). In any case, the RMS wavefront after calibration is still lower than the model trained at single spaxel scale, so the model retains some robustness. For instance, the robust




Figure 5.28: Data augmentation for robust training. The *blue* dots show the original behaviour (Fig. 5.27). The remaining data points show the performance for a model trained on images with varying levels of pixel scale errors, up to $\pm 10\%$. The *orange* points indicate the cases used for training. The *green* points show test cases within the range, but at points not used for training; this demonstrates that the model can interpolate and does not just "memorise" the training data. The *red* points represent test cases beyond the $\pm 10\%$ range, for which the performance starts to degrade again. This demonstrates that we can introduce robustness to the model and retain competitive performance across a wide range of pixel sampling errors.

model can tolerate around -20% scale error, whereas the nominal scale model has the same RMS wavefront (>100 nm) for -7.5% scale error.

In conclusion, the effect of uncertainties in spaxel scale can be successfully mitigated by including representative examples in the training sets. This forces the calibration model to become robust against scale variations. As a result of this, we can have a model that reliably estimates the aberrations across a range of $\pm 10\%$ errors (possibly higher if we adapt the training). Moreover, this model can cope with even higher errors with little performance degradation compared to the nominal model so that it can be applied slightly outside its trained range. All of this comes at the cost of having to generate additional examples, which can quickly become a burden, especially if one has to train for robustness against multiple effects (scale errors, photon noise, flat field calibration uncertainties).

5.4.4.3 Anamorphic magnification errors

The PSF at the image slicer of HARMONI has an anamorphic magnification of 2:1 as a nominal value. However, the as-built instrument might have errors in this ratio that would cause the reconstructed PSF to exhibit some spurious ellipticity. Rather than appearing entirely circular, the reconstructed PSF would be slightly elongated.

This might have a detrimental effect on the performance of the calibration algorithm, since we assume a perfectly circular pupil when simulating PSF images for the training sets. For this reason, we decided to characterise the impact of anamorphic magnification errors in the performance of a ML calibration model.

First, we began by characterising *how errors in anamorphic magnification affect the PSF*. For this purpose, we generated multiple PSF images with random Zernike aberrations (reference wavelength 1.5 μm , sampled with 1 mas pixels). Then we generated the equivalent images for a PSF model with 10% anamorphic magnification error. In order to model this, we modified the pupil mask and Zernike model matrix so that they were elliptical with a ratio $a/b = 1.10$, where a represents the semi-major axis (or R_x), and b represents the semi-minor axis (or R_y). This causes the PSF to be elongated 10% along the Y axis, without modifying the aberration content.

We compared the two versions of the PSF: the one with a perfectly circular pupil mask, and the one with ellipticity. The results are shown in Fig. 5.29; we can see that the ellipticity in the reconstructed PSF due to an error in anamorphic magnification of 10% modifies the intensity by approximately 10% of the peak value.

The results from Fig. 5.29 only show the impact over a limited number of examples of random aberrations. In order to investigate whether there is some commonality in the PSF changes across any possible wavefront, we repeated the analysis over 100 random aberrations, and calculated the average difference in PSF features.

Fig. 5.30 shows the average difference in PSF features for several cases of ellipticity. Interestingly, when we average across a wide range of wavefronts, we start to see an underlying pattern: the anamorphic magnification errors significantly modify the PSF along the y direction, mostly affecting two lobes at both sides of the PSF core, as well as the ring structure.

Undoubtedly, such spurious variations will hinder the performance of a calibration algorithm, but the pattern we have seen in Fig. 5.30 raises an interesting question. Do these spurious features affect the predictions for each Zernike aberration equally? Given that there is a specific structure to these PSF features, it is possible that the calibration model might struggle to estimate a particular subset of aberrations that cause changes in the PSF resembling these features.




Figure 5.29: Changes in the PSF due to anamorphic magnification errors. PSFs with random aberrations are compared against their versions with ellipticity $a/b = 1.10$.




Figure 5.30: Changes in the PSF due to anamorphic magnification errors, averaged across 100 random aberrations.

In other words, if there is a subset of Zernike polynomials that affect the PSF in a similar fashion as ellipticity variations, the calibration model could potentially mistake anamorphic magnification errors for spurious aberrations, leading to a bias in the predictions.

In order to investigate this possibility, we ran the following experiment. We performed a non-linear least squares fit on the PSF variations due to ellipticity. The model is given by:

$$\Delta\text{PSF}(\Phi_i = 0, \varepsilon) = \Delta\text{PSF}(\psi(\bar{x}), \varepsilon = 0) + \eta \quad (5.45)$$

where the *observations* $\Delta\text{PSF}(\Phi_i = 0, \varepsilon)$ represent the changes in the nominal PSF intensity (no aberrations $\Phi = 0$) due to ellipticity ε . The non-linear model represents the changes in the PSF introduced by a wavefront $\psi(\bar{x})$ defined by some unknown Zernike coefficients \bar{x} , when there is no ellipticity $\varepsilon = 0$.

We ran this non-linear least squares fit for several values of the ellipticity ratio a/b , and found that a very specific subset of Zernike polynomials generate a wavefront that mimics the PSF changes from anamorphic magnification errors, see Fig. 5.31. This set of Zernike polynomials is shown in Fig. 5.32, and it contains defocus, astigmatism, spherical aberration and some higher order aberrations.

A comparison between the PSF difference caused by this artificial wavefront and the ones induced by an anamorphic magnification error $a/b = 1.05$ is presented in Fig. 5.33. In general, the two sets of features are similar, specially around the ring structure of the PSF. However, in terms of magnitude, the artificial wavefront introduces changes that are around one order of magnitude smaller; plus it invariably reduces the peak intensity.

This suggests that the artificial aberration model represents the changes in PSF features only partially, and that other effects are not being captured. In addition, it is likely that, as the ellipticity a/b increases, the errors will be larger. In any case, the non-linear least squares fit points at a very specific set of Zernike aberrations, which led us to believe that they are likely candidates for calibration errors.

So far we have only covered *how errors in anamorphic magnification affect the PSF*, but we still have to investigate *how these PSF changes affect the performance of calibration models*. In order to explore this, we ran the following experiment.

We trained a CNN calibration model on a dataset of 10,000 PSF datacubes, containing random wavefronts defined on 33 Zernike polynomials (this excludes piston and tilts). The PSF model uses a perfectly circular aperture (with a central obscuration), the PSF images are sampled at 4 mas pixels, with a field of view of 16 pixels,




Figure 5.31: Non-linear least squares fit of the Zernike coefficients inducing PSF changes that resemble those caused by ellipticity variations. For all cases of ellipticity ratio a/b , a subset of Zernike polynomials is identified.




Figure 5.32: Zernike polynomials from the non-linear least squares fit, which could potentially cause a similar effect to ellipticity errors.




Figure 5.33: Comparison between the PSF changes introduced by artificial aberrations and those induced by anamorphic magnification errors. *Lower left*: the artificial wavefront identified by LS as a candidate to mimic the ellipticity changes. *Right column*: although similar, the artificial wavefront does not exactly match ellipticity variations, especially around the core of the PSF.

reference wavelength $1.5 \mu\text{m}$. In other words, this is just a conventional calibration model trained on simulated PSF images assuming no anamorphic magnification errors.

Then, we generated multiple test sets of 1,000 PSF datacubes, using PSF models with some anamorphic magnification errors, i.e. whose pupil and aberration models were elliptical $a/b \neq 1$. We tested the performance of the calibration model for both the *nominal* case and the datasets with errors, running the calibration for 3 consecutive iterations. This exposes the model to PSF images it has been trained on (the ones without anamorphic errors) and PSF images with ellipticity. In order to investigate whether the presence of anamorphic magnification errors affects the predictions of each Zernike polynomial differently, we recorded the residual coefficients separately for each polynomial.

Histograms of residual coefficients are shown in Fig. 5.34 and Fig. 5.35 for a/b ratios of 1.10 and 0.90 respectively. Filled histograms represent the performance for the *nominal* PSF model (circular pupil), black line histograms represent the performance when we test on PSF images that have anamorphic magnification errors. Red histograms highlight Zernike polynomials that the non-linear least squares fit identified

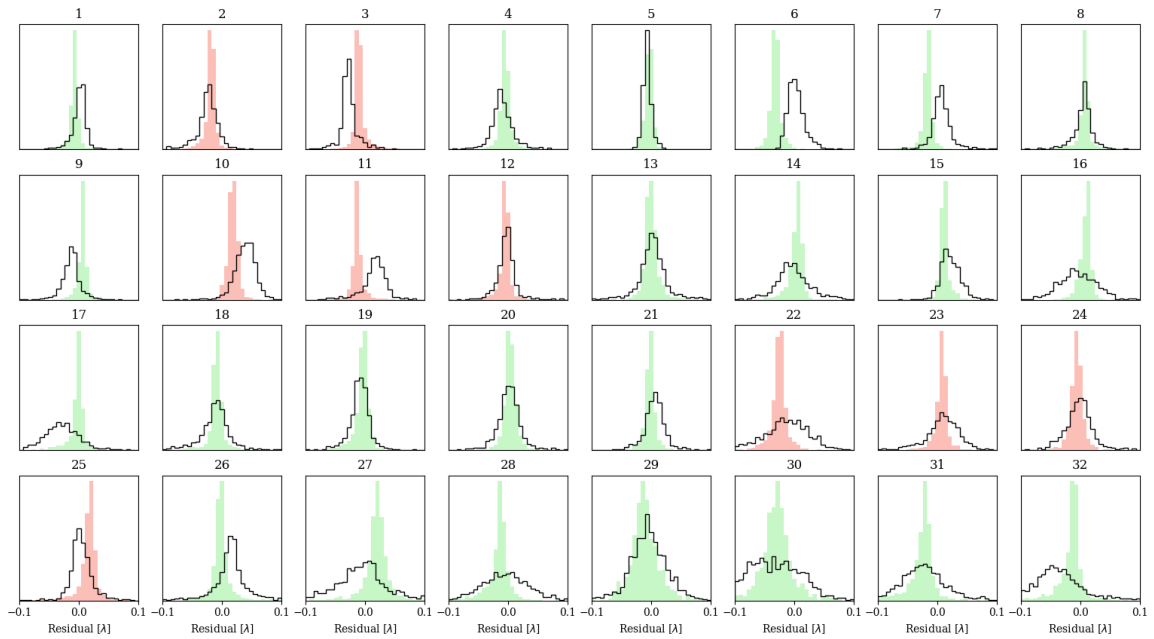


Figure 5.34: Histograms of Zernike residual coefficients after calibration (3 iterations, 1000 samples). Filled values indicate the nominal PSF model (red marks the aberrations identified via non-linear least squares as mimicking anamorphic magnification errors), black line indicates the PSF model with an ellipticity $a/b = 1.10$.

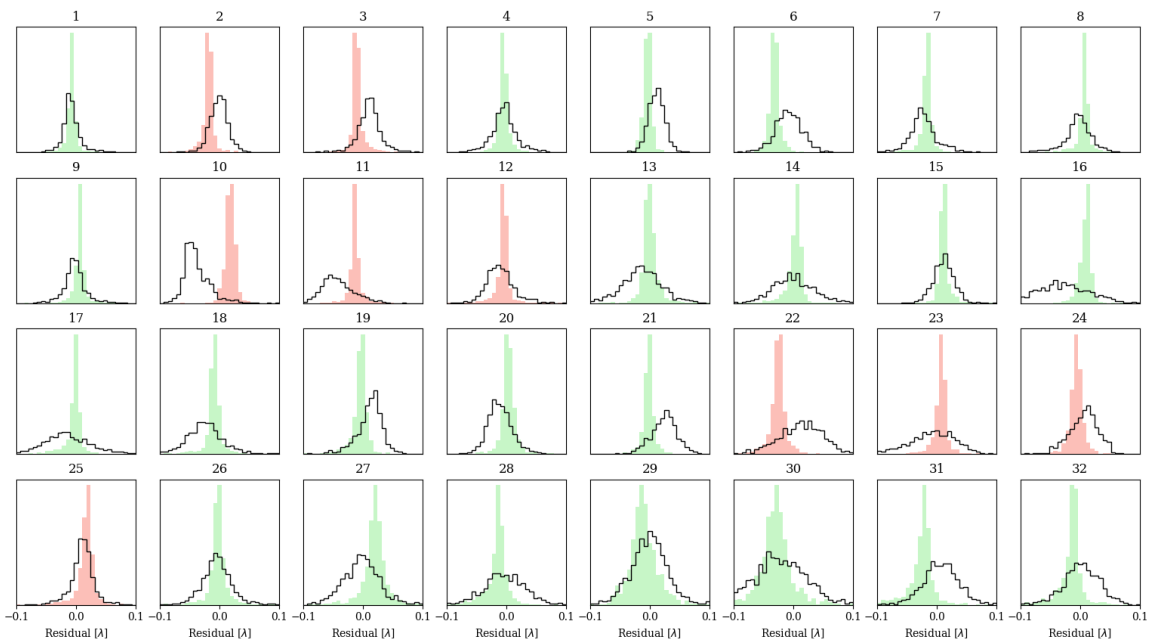


Figure 5.35: Histograms of Zernike residual coefficients after calibration (3 iterations, 1000 samples). Filled values indicate the nominal PSF model (red marks the aberrations identified via non-linear least squares as mimicking anamorphic magnification errors), black line indicates the PSF model with an ellipticity $a/b = 0.90$.

as causing changes in PSF features resembling anamorphic magnification errors.

The first thing to note is that with regards to the aberrations highlighted in red, there is a significant degradation in performance for most of them, leading to either a strong bias in the predicted coefficients, or a greater uncertainty and spread when compared to the nominal PSF model. This appears to support our initial hypothesis that the Zernike polynomials that produce PSF changes resembling anamorphic magnification errors are prone to errors in the predictions.

However, the degradation is not limited to these polynomials only; with most Zernikes showing a worse performance in the presence of anamorphic magnification errors. A surprising exception is polynomial #6 (vertical coma) which exhibits a negative bias for the nominal PSF model, but a better performance for both a/b of 1.10 and 0.90.

In order to quantify the effect that anamorphic magnification errors have in the overall performance, i.e. the combined contributions of all the Zernike polynomials, we calculated the average Strehl ratio across 1,000 PSF datacubes, as function of the iterations. The results for different values of the a/b ratio are shown in Table 5.5. For the nominal PSF model, the calibration achieves an almost perfect Strehl of 97% after 4 iterations. For all other cases, the presence of anamorphic magnification errors leads to some performance loss. However, this is almost negligible, with an a/b ratio of 12.5% retaining a Strehl above 90%.

An interesting thing to note is the asymmetry in the performance. While an $a/b = 1.15$ has 85% Strehl ratio, $a/b = 0.85$ has less than 70% Strehl ratio. One possible explanation for this is the fact that $a/b > 1$ corresponds to an elliptical PSF elongated along the y direction, compared to the nominal PSF, while $a/b < 1$ indicates an elliptical PSF shortened in the y direction. Consequently, $a/b > 1$ leads to a finer sampling in the y direction, which might help resolve the features, whereas $a/b < 1$ creates and artificially coarser sampling.

Finally, we investigated the possibility of mitigating the effect of anamorphic magnification errors by making the training robust. Rather than training the calibration model on perfectly circular PSF images, we created a mixed dataset using 5 PSF models with varying ellipticity $a/b = [0.90, 0.95, 1.00, 1.05, 1.10]$. Each model contributes 2,000 PSF datacubes with random aberrations, for a total of 10,000 training examples (the same as for the nominal model).

After training the robust model, we tested the performance with 9 test datasets with ellipticity ratios a/b between 0.85 and 1.15, slightly outside the training range. Table 5.6 shows the average Strehl ratios across the 1,000 PSF datacubes for each

N_{iter}	Ellipticity Ratio a/b								
	0.85	0.8875	0.925	0.9625	1	1.0375	1.075	1.125	1.15
0	0.431	0.430	0.427	0.439	0.427	0.426	0.424	0.431	0.429
1	0.602	0.659	0.715	0.748	0.764	0.753	0.731	0.707	0.674
2	0.670	0.776	0.855	0.900	0.915	0.904	0.870	0.832	0.785
3	0.686	0.822	0.908	0.946	0.961	0.955	0.927	0.882	0.832
4	0.691	0.835	0.921	0.958	0.974	0.972	0.953	0.908	0.849

Table 5.5: Average Strehl ratios across 1,000 PSF datacubes, as a function of the iterations, for several values of ellipticity ratio a/b . The Strehl ratio remains relatively high even for large ellipticity changes (7.5%).

N_{iter}	Ellipticity Ratio a/b								
	0.85	0.8875	0.925	0.9625	1	1.0375	1.075	1.125	1.15
0	0.432	0.427	0.424	0.426	0.430	0.429	0.425	0.426	0.427
1	0.698	0.720	0.738	0.748	0.755	0.753	0.741	0.731	0.713
2	0.824	0.863	0.885	0.899	0.901	0.898	0.889	0.875	0.850
3	0.883	0.923	0.943	0.953	0.953	0.950	0.943	0.931	0.901
4	0.913	0.953	0.967	0.973	0.972	0.965	0.965	0.954	0.935

Table 5.6: Robust training. Average Strehl ratios across 1,000 PSF datacubes, as a function of the iterations, for several values of ellipticity ratio a/b . The calibration model was trained on PSF datacubes covering a/b ratios between 0.90 and 1.10.

case of a/b . In all cases within the training range, the final Strehl is above 95%. The performance starts to degrade slightly outside the $0.90 < a/b < 1.10$ range, but the Strehl ratios remain above 90% even for a/b of 0.85 and 1.15.

In conclusion, the presence of anamorphic magnification errors induces ellipticity on the PSF after reconstruction. The changes in PSF features due to these errors affect the performance of calibration models trained on perfectly circular PSF examples. However, the loss in performance is moderate, with ellipticities of $\sim 10\%$ maintain a Strehl ratio around 90% after 4 iterations. In any case, these effects can be mitigated with a robust training of the calibration models, including examples of PSF images with varying ellipticity.

5.4.4.4 Aggregating all effects - Robust calibration model

Throughout the previous sections, we have looked at different effects independently of each other, in order to isolate their contribution to the performance degradation of calibration models. In all cases, we have seen that a robust training approach was feasible as a means of mitigating performance degradation. Including training examples with realistic effects such as variations in spaxel scale, anamorphic magnification errors, or readout noise allowed the calibration models to deliver accurate predictions of the aberrations beyond the simplistic scenario of ideal PSF images.

This brought us one step closer to being able to deal with real HARMONI PSF images, but it has major drawback: so far these effects have been treated independently. In this section, we will go one step further by combining multiple real effects to train a robust calibration model on a dataset of more realistic PSF images.

For this experiment, we will incorporate both *modelling* errors that affect the PSF and *detector* systematics. These include:

- **Pixel sampling errors:** variations in spaxel scale around the nominal value. The as-built HARMONI instrument will never have exactly 4.00 mas spaxels. We will include random instances of spaxel scale, representing changes in the PSF sampling.
- **Diversity errors:** variations in applied defocus diversity around a nominal value. We assume there will be some uncertainty in the diversity phase that the as-built instrument actually applies.
- **Anamorphic magnification errors:** residual magnification errors in the as-built instrument can cause spurious ellipticity in the reconstructed PSF images.
- **Flat field uncertainty:** variations in pixel sensitivity not properly calibrated will rescale the PSF features in an unpredictable way. We assume the flat field calibration for the as-built instrument will have some residual uncertainty.
- **Readout noise:** the PSF images are contaminated with readout noise at varying SNR levels.




Figure 5.36: Cumulative distribution of Strehl ratio as a function of iterations, robust model, 52 Zernike polynomials. SNR readout noise ranging from 100 to 500. This includes errors in pixel scale, diversity defocus, ellipticity, and flat field uncertainties.

Every example PSF datacube represents a random sample drawn from this realistic *effects* space. For the purpose of generating realistic PSF images, each effect is modelled as a uniform distribution centred at the nominal value, with a range of $\pm 5\%$, except for the readout noise, which varies from 100 to 500 SNR.

$$\begin{aligned} \mathcal{T}_n &= \{[\text{PSF}_{nom}, \text{PSF}_{foc}], a_n\} \\ \text{PSF}_{nom} &= \epsilon_{ff} \cdot \text{PSF}(\Phi_n | p_0 + \epsilon_p, \epsilon_{a/b}) + \epsilon_{ro} \\ \text{PSF}_{foc} &= \epsilon_{ff} \cdot \text{PSF}(\Phi_n + (f + \epsilon_f)Z_f | p_0 + \epsilon_p, \epsilon_{a/b}) + \epsilon_{ro} \end{aligned} \quad (5.46)$$

In other words, to generate each example PSF datacube we create a PSF model with a random value of spaxel scale ($p_0 + \epsilon_p$), diversity defocus ($f + \epsilon_f$), ellipticity $\epsilon_{a/b}$, flat field uncertainty ϵ_{ff} , and readout noise SNR level ϵ_{ro} . For this experiment, we created a training set of 50,000 PSF datacubes, with random aberrations in the form of 52 Zernike polynomials. The test set comprised 5,000 PSF datacubes, with equivalent parameters to those of the training set.

We found the optimum calibration model using hyper-parameter tuning as described in Section 5.4.1. After training, we ran the calibration for 5 iterations on the test dataset, the results in terms of cumulative Strehl ratio are shown in Fig. 5.36. All samples have final Strehl ratios above 0.70, 80% above 0.85 and 60% above 0.90.

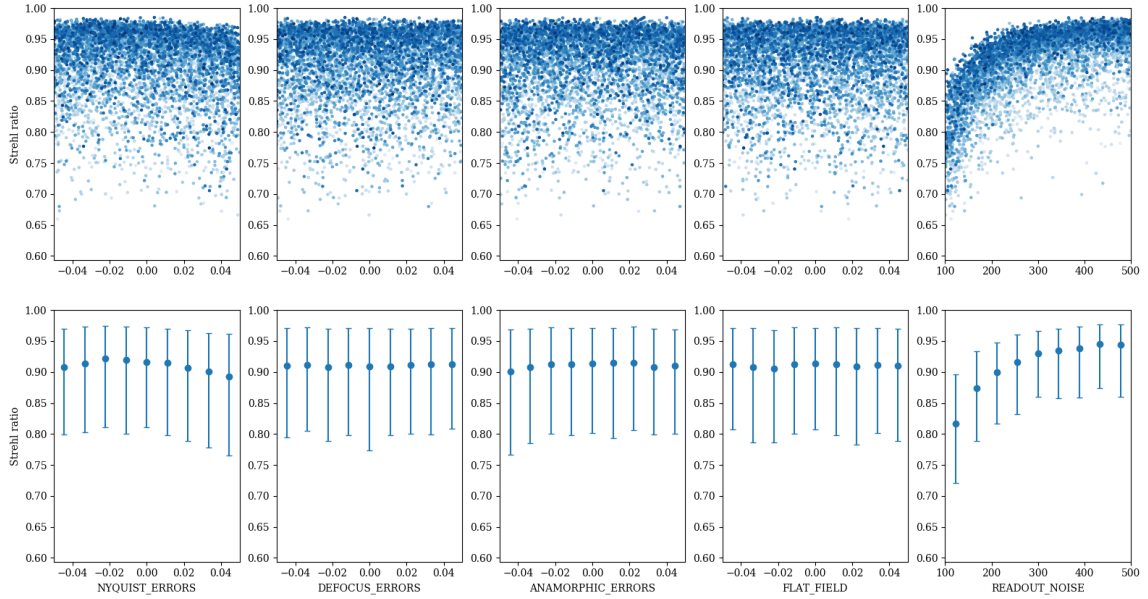


Figure 5.37: Final Strehl ratio after calibration (5 iterations) as a function of the different effects in the robust model, 52 Zernike polynomials. *Upper row*: 5,000 samples of the test set, coloured by initial Strehl ratio (darker shade equals higher Strehl). *Lower row*: binned data, error bars represent 5th and 95th percentile. SNR readout noise ranging from 100 to 500.

In order to assess whether any modelling effect played a significant role in limiting the performance, we decided to plot the final Strehl ratios as a function of each effect, see Fig. 5.37. Each column shows the final Strehl ratio after 5 iterations, the samples are coloured according to the initial Strehl ratio (darker shade indicating a higher Strehl); the lower panels represent binned data with error bars marking the 5th and 95th percentile.

The readout noise SNR level appears to be a critical factor in performance, with values below 200 SNR quickly losing Strehl ratio. For other effects such as diversity defocus or flat field uncertainties the calibration model response is almost flat, a good sign of a robust model. With regards to Nyquist sampling errors the performance is slightly skewed, with positive errors (spaxel scales coarser than 4.00 mas) losing Strehl faster than negative errors.

Since low values of readout noise SNR had such a big impact on final Strehl ratio compared to other effects, we decided to redefine the SNR range and repeat the training; this time ranging from 250 to 500 SNR. In addition, to investigate the effect of the number of Zernike polynomials the model has to estimate, this time we used wavefront maps containing 33 polynomials. The results in terms of final Strehl ratio




Figure 5.38: Cumulative distribution of Strehl ratio as a function of iterations, robust model, 33 Zernike polynomials. SNR readout noise ranging from 250 to 500. Increasing the SNR from 100 to 250 compared to Fig. 5.36 greatly improves the performance.

after calibration are shown in Fig. 5.38. A significant improvement in performance is observed, with all cases above 0.90 Strehl ratio, 80% above 0.95.

As far as each effect is concerned, Fig. 5.39 shows that readout noise still remains the major factor in determining post-calibration Strehl ratio. The behaviour of the calibration model is robust for other effects. In order to test how far this robustness extends, we created a new test dataset for which realistic effects varying between $\pm 10\%$, well outside the $\pm 5\%$ range defined for the training set. We decided to extend the readout noise SNR only towards higher values, because lowering it below 250 would cause a degradation that would eclipse any losses coming from the other effects.

The results are shown in Fig. 5.40, the green shaded region represents the range of training parameters. The first thing to note is that both defocus errors and flat field uncertainties had little impact on the final Strehl ratio even for values outside the training range. This could mean two things: either the changes in the PSF due to these effects are too small to significantly affect the predictions, or (more likely) the calibration model manages to become robust to these perturbations and can generalise outside the initial range.

Interestingly, no degradation in performance is seen at readout noise SNR levels higher than the training range, suggesting the calibration model can deal with images that have lower noise contamination than the ones used in training.

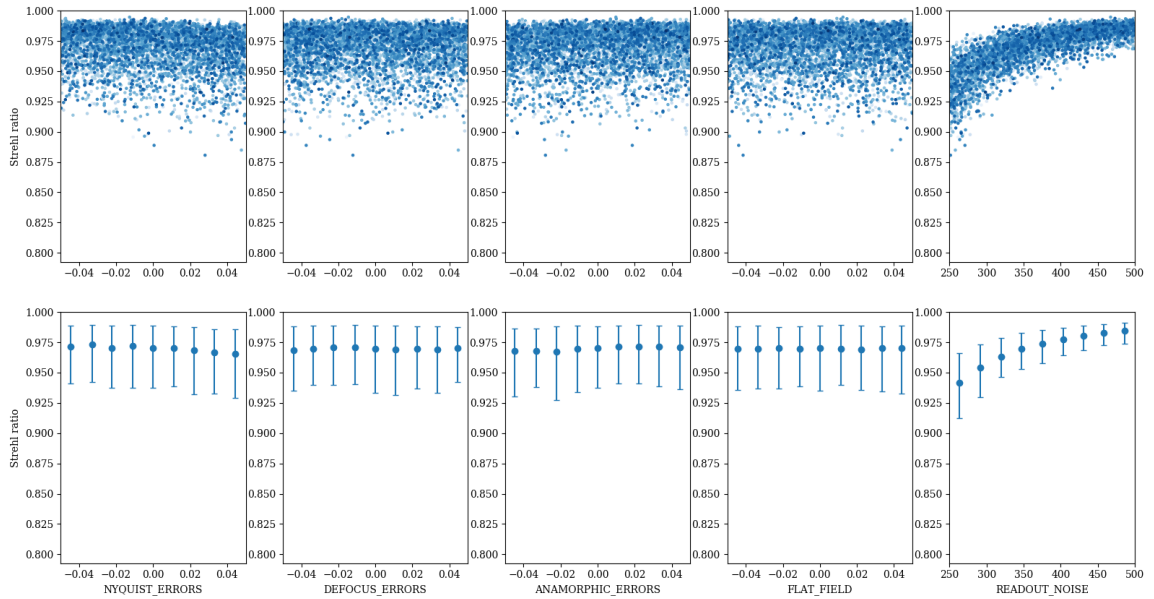


Figure 5.39: Final Strehl ratio after calibration (5 iterations) as a function of the different effects in the robust model, 33 Zernike polynomials. *Upper row*: 5,000 samples of the test set, coloured by initial Strehl ratio (darker shade equals higher Strehl). *Lower row*: binned data, error bars represent 5th and 95th percentile. SNR readout noise ranging from 250 to 500.

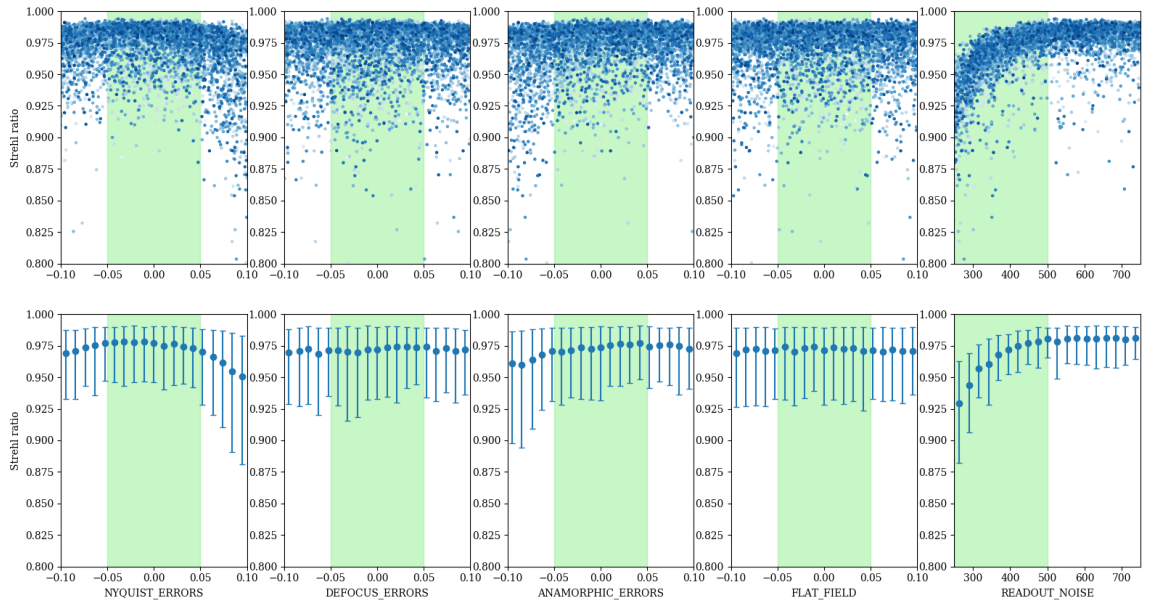


Figure 5.40: Final Strehl ratio after calibration (5 iterations) as a function of the different effects in the robust model, 33 Zernike polynomials. *Upper row*: 5,000 samples of the test set, coloured by initial Strehl ratio (darker shade equals higher Strehl). *Lower row*: binned data, error bars represent 5th and 95th percentile. SNR readout noise ranging from 250 to 500. Green area marks the parameter range used to train the robust model.

The only two effects that exhibit a significant loss in performance when tested outside the training range are spaxel sampling errors and anamorphic magnification errors. For the former, we already showed in Section 5.4.4.2 that positive deviations towards coarser scales (> 4 mas) increased the residual calibration errors faster than negative deviations that improve the PSF sampling. The same behaviour is seen in Fig. 5.40, although the model appears to be relatively robust when exposed to negative deviations. For the latter, a similar pattern holds. Section 5.4.4.3 showed that negative variations in ellipticity (which compress the PSF in the y direction reducing the effective sampling) are more detrimental than positive variations that elongate the PSF. This is also seen in Fig. 5.40.

5.4.5 Multiwavelength calibration

5.4.5.1 Motivation

So far, all the analyses presented here have been *monochromatic*. The PSF datacubes used to train the calibration models always included two channels: the in-focus PSF(Φ) and defocused PSF($\Phi + \epsilon Z_f$) images, at a fixed wavelength λ . This is typically the case for other state-of-the-art techniques for NCPA calibration like PD. Since only one correction can be applied with the DM across the pupil, these techniques normally estimate the aberrations based on PSF images taken at a reference wavelength (Sauvage et al., 2012, 2007).

However, in the case of the HARMONI instrument (an integral field spectrograph), we have privileged access to rich IFS data across whole spectral bands. Relying on PSF information at single wavelength to perform NCPA calibration would be wasteful, specially considering that machine learning techniques thrive on the availability and richness of data to provide accurate predictions. Moreover, convolutional neural network models are ideal for working with multidimensional data, so their application to IFS data should be straightforward.

In light of this, we decided to investigate the possibility of incorporating spectral information into our calibration models.

5.4.5.2 Data format

For this experiment we will operate on simulations of the PSF at multiple wavelengths, both in and out of focus. For instance, for a given NCPA wavefront, we will have a




Figure 5.41: Example datacube $\mathcal{D}(\Phi)$ for the multiwavelength calibration experiment. The top row shows the nominal PSF(Φ, λ_i) as a function of wavelength λ_i . The bottom row shows the defocused channels PSF($\Phi + \psi, \lambda_i$). All these channels are fed to the CNN model as a complete datacube of size $[N_{pix}, N_{pix}, 2N_\lambda]$.

datacube $\mathcal{D}(\Phi)$ containing $2N_\lambda$ channels:

$$\mathcal{D}(\Phi) = [\text{PSF}(\Phi, \lambda_1), \text{PSF}(\Phi + \psi, \lambda_1), \dots, \text{PSF}(\Phi, \lambda_N), \text{PSF}(\Phi + \psi, \lambda_N)] \quad (5.47)$$

for which $\text{PSF}(\Phi + \psi, \lambda_i)$ represents the PSF at a wavelength λ_i for an NCPA map Φ with a diversity defocus ψ . For simplicity, we keep the NCPA aberrations Φ constant (in physical units, not spectral) across all wavelengths, i.e. we do not consider chromatic effects.

For the purpose of this experiment, we will train calibration models on example datasets containing PSF datacubes $\mathcal{D}(\Phi)$ for variable spectral ranges $\lambda_1 - \lambda_N$, and variable number of wavelength channels N_λ . An example datacube is shown in Fig. 5.41 for 5 wavelength channels covering 1.50 - 2.0 μm .

5.4.5.3 Results

To illustrate how multiwavelength information can be exploited, we ran the following simulation. We created a training set of 10,000 PSF datacubes containing random aberrations, with 5 wavelength channels in the range 1.5-2.0 μm . Readout noise is included in the PSF images. Then, we trained a calibration model, varying the number of channels. For example, we begin by training the model using only the first wavelength channel (1.5 μm both in-focus and defocus images), and evaluate the performance in terms of RMS wavefront error (we run 3 iterations, updating the PSF images every time).

N_λ	SNR 500	SNR 250	SNR 125
1	61±20	81±20	110±27
2	52±13	76±22	100±23
3	53±12	71±17	92±21
4	49±11	71±16	89±20
5	48±10	66±15	85±19

Table 5.7: Evolution of RMS after calibration [nm] with number of wavelength channels, for several cases of SNR readout noise. Reference wavelength is 1.5 μm , spectral range 1.5-2.0 μm . Adding extra channels within the same spectral range improves the accuracy of the predictions.

Then, we train the model again, but this time including two wavelength channels (1.50 and 1.625 μm). The model architecture and parameters remain the same, the only difference being that the training set contains additional channels. The results for all cases are shown in Table 5.7. For reference, initial aberrations ranged from 150 to 350 nm.

At this point, one might raise the question of whether the performance improvement of using all 5 channels comes from the fact that the PSF for the last channel (2.0 μm) is effectively better sampled, rather than from an intrinsic benefit of using all available spectral data. To demonstrate that this is not the case, and that there is an advantage in combining all the channels, we trained the model again on a single wavelength channel at 2.0 μm . The results for all cases of readout noise are summarised in Table 5.8. The first thing to note is that, in general, training the model just on the last wavelength channel (2.0 μm) leads to a worse performance than when trained on the first channel (1.5 μm), except for SNR 125 where they are equivalent.

We hypothesise that this could be due to two factors. First of all, the field of view is kept fixed for all channels. At longer wavelengths, with the scaling of the PSF, a smaller fraction of it will fit within the field of view, possibly losing valuable information about the features. Secondly, for a fixed phase diversity in physical units (defocus in this case), at longer wavelengths this diversity represents a smaller fraction of the operating wavelength, which means that the defocused channel is more similar to the in-focus channel and the features become less informative.

The next thing to note is that in all cases, training on the complete set of wavelength channels always led to an improved performance, no matter the levels of readout noise. Although the gains are modest at high SNR, they are significant at SNR 125. This is probably because when the PSF is substantially corrupted by readout noise, the addition of multiple wavelength channels helps mitigate the impact of noise.

$N_\lambda(\lambda)$	SNR 500	SNR 250	SNR 125
1 (1.5 μm)	61 \pm 20	81 \pm 20	110 \pm 27
1 (2.0 μm)	72 \pm 17	94 \pm 21	111 \pm 22
5 (-)	48 \pm 10	66 \pm 15	85 \pm 19

Table 5.8: RMS after calibration [nm] when trained on a single channel, at two wavelengths (1.5 or 2.0 μm), and all 5 wavelengths (1.5 - 2.0 μm). This shows that the performance improvement from using multiple wavelength channels is not due to the fact that at 2.0 μm the PSF is better sampled. Using 5 channels (1.5 - 2.0 μm) outperforms the calibration at the longest wavelength with a single channel (2.0 μm).

$\Delta\lambda[\mu\text{m}]$	SNR 500	SNR 250	SNR 125
1.50 - 1.75	46 \pm 10	61 \pm 14	88 \pm 22
1.50 - 2.00	48 \pm 10	66 \pm 15	85 \pm 19
1.50 - 2.25	51 \pm 9	66 \pm 12	90 \pm 19

Table 5.9: RMS after calibration [nm] when trained on $N_\lambda = 5$, with varying spectral ranges. The spectral range $\Delta\lambda$ seems to have no influence in the performance.

Another interesting thing to look into is whether the spectral range plays any role in the performance gains. Initially we chose the arbitrary range 1.5-2.0 μm . We repeated the analysis, changing the upper wavelength (1.75, 2.0 and 2.25) μm , while keeping the reference wavelength of 1.5 μm and the total number of channels $N_{\text{lambda}} = 5$ constant. The results shown in Table 5.9 suggest that the spectral range has little impact on the performance.

To further investigate the impact of the number of channels, we repeated the analysis on a larger dataset ranging from 1.50 to 1.75 μm with a total of $N_\lambda = 15$ wavelength channels. The results are summarised in Table 5.10. In all cases of readout noise, training the model on the maximum number of channels ($N_{\text{lambda}} = 15$) led to an improvement in RMS after calibration. Interestingly, the improvement is more significant when the readout noise increases: while for a SNR 500 the gain in RMS

N_λ	SNR 500	SNR 250	SNR 125
1	55 \pm 17	74 \pm 20	100 \pm 27
5	46 \pm 18	55 \pm 16	76 \pm 21
10	41 \pm 13	52 \pm 14	66 \pm 18
15	38 \pm 9	50 \pm 11	61 \pm 15

Table 5.10: RMS after calibration [nm] as a function of the number of wavelength channels N_λ , for the range 1.5 - 1.75 μm . The number of channels appears to be the main performance driver, although it shows clearly diminishing returns after 5-10 channels.

was on average $1 \text{ nm} / \lambda$ channel, the gain was approximately 1.6 and $2.6 \text{ nm} / \lambda$ channel, for SNR 250 and 125 respectively.

It is important to point out that for all analyses presented in this section, we have trained the neural network models on PSF images with significant levels of NCPA. The Strehl ratio in the training and test sets ranges from 15 to 75%. Although this demonstrates that the calibration can be done over a wide range of aberration intensities, it does not reflect the situation that we will encounter in HARMONI. As I show in Chapter 6, in the 4×4 scale, NCPA ranges from 30-60 nm RMS, which corresponds to a well-corrected regime with initial Strehl ratios above 90-95%. In order to further demonstrate how this calibration approach could be applied in a more realistic setting, I repeated the analysis, but with a training and test set containing example PSF datacubes with Strehl ratios above 85%. A total of $N_\lambda = 5$ wavelength channels were used covering the range 1.5-2.0 μm , with a SNR of 500.

The results of the calibration for this setting can be summarised as follows: the test set of 500 samples started with an RMS wavefront error of $60.0 \pm 12.4 \text{ nm}$. After the first iteration, the residual RMS wavefront error was $19.5 \pm 3.6 \text{ nm}$, and further improved to $15.4 \pm 1.9 \text{ nm}$ with a second iteration. This shows that with this approach, we can reach an accuracy in the calibration comparable to the intrinsic limit from field-dependent aberrations of 15 nm, which was set as the requirement for the calibration.

In conclusion, we have demonstrated that the ML calibration models are capable of incorporating spectral information that is readily available in IFS data from HARMONI. This is the first demonstration of this approach, compared to more traditional techniques like Phase Diversity, which normally rely on monochromatic data (although are not, in principle, limited to this). Another important point is that no pre-processing of the IFS data is needed, such as a spectral scaling of the wavelength channels and averaging. The neural network models can readily receive a complete IFS datacube and the training can be adapted to accommodate varying number of channels. This approach leads to an improvement in overall calibration performance, the main driver being the total number of wavelength channels N_λ made available to the network. We have also shown that the relative benefit of this approach is greater for cases with higher levels of readout noise, which suggests it could be a good way of counteracting the performance losses due to contamination.

5.4.6 Interpreting the predictions

Arguably, the main criticism of deep neural network models is the lack of explainability and interpretability of their results. They are often regarded as black-box models whose predictions bear little relation to the underlying physical processes. Despite their competitive performance, the challenges of explaining what drove the models to make a particular prediction cast a shadow of doubt on their reliability, which can motivate researchers to adopt simpler models whose results are easier to interpret.

Considerable effort has been devoted over recent years to develop tools to inspect the behaviour of these models, in an attempt to clarify the logic behind their decision-making and interpret the predictions, see [Arya et al. \(2020\)](#). In this section, we will use of these tools to try to understand how the neural network models for NCPA calibration make predictions.

5.4.6.1 Shapley values

Shapley values is a game theory method for the distribution of payout to players participating in a cooperative game, depending on their contribution to the total payout ([Shapley, 1953](#)). Recently, [Lundberg and Lee \(2017\)](#) proposed a unified framework for interpreting predictions based on Shapley values called SHAP (SHapley Additive exPlanations) that assigns each feature an importance value for a particular prediction.

But what is the link between game theory and NCPA calibrations? The features (or pixels) in the PSF image can be regarded as players in a game whose payout is the difference between the prediction and the average prediction for the dataset. The interpretation of the Shapley value for a feature is that: for the given set of features values (a PSF image), the contribution of such feature to the difference between the prediction and the average prediction is its Shapley value. Therefore, the Shapley value for each pixel p_{xy} in a $\text{PSF}(\Phi_n)$ represents a measure of its *importance* when calculating the aberration coefficients a_i^n .

Shapley values can be positive or negative. A positive value indicates that the model saw the presence of that feature as a sign for a_i^n ; whereas a negative value indicates that the model interprets that feature as a sign of the absence of a_i^n or that a_i^n is negative.

Based on this idea, we decided to test whether the Shapley values could help us interpret the predictions of the calibration models and understand what features of the image are driving the decisions. For that purpose, we trained a calibration model

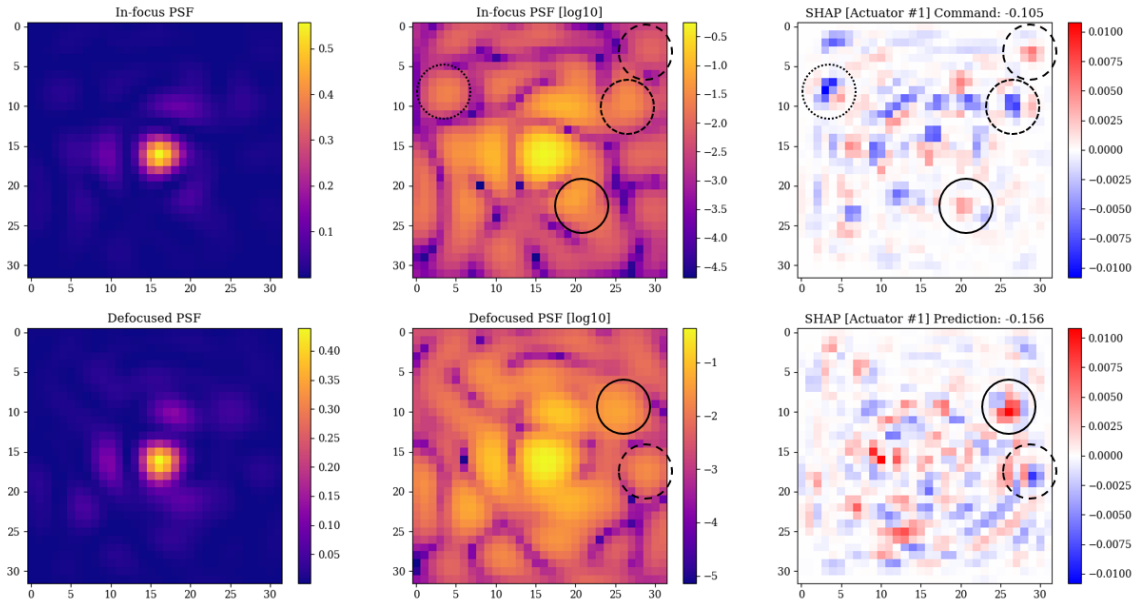


Figure 5.42: Example of Shapley values. *Left*: in-focus and defocus PSF images. *Middle*: logarithmic scale. *Right*: Shapley values for the two PSF images associated with the prediction for the actuator #1 command. Shapley values are a measure of feature importance, highlight which PSF pixels drove each aberration prediction.

based on the **actuator** commands wavefront representation ($N_{act} \sim 50$), using 10,000 PSF images ($1.5 \mu\text{m}$ wavelength).

After training, we use a Python toolbox to calculate the Shapley values for a subset of 500 PSF images from the test set. Instead of using 4.0 mas pixels (HARMONI spaxel scale), we doubled the resolution to 2 mas pixels to better sample the PSF features and facilitate the interpretation of the Shapley values results.

Let us begin with an example. Fig. 5.42 shows the Shapley values for a PSF datacube in the test set. The left column represents the two PSF channels (in-focus and defocused) that the calibration model will use to predict the actuator commands. The central column shows the PSF images in logarithmic scale to highlight the aberrated Airy rings. The right column shows the Shapley values associated with a particular actuator command, for both PSF images. The true value of the command and the prediction are also noted.

The Shapley values tell us which pixel features from the images have influenced the prediction of the calibration model for this actuator command. This serves as a measure of *feature importance*, and quantifies how the model has reached its decision.

In this case, the predicted value is negative, which means that it is the negative Shapley (*blue*) features that have driven such prediction. There are also some features

with positive Shapley (*red*), probably from other actuators contributing to the wavefront. These can be regarded as features that are "uncharacteristic" of this actuator command.

The sum of the Shapley value over all pixel features is equal to the difference between the prediction and the average prediction across the whole dataset, i.e. the marginal contribution of the features. Since the aberration coefficients have zero mean, the summed Shapley value directly equals the predicted actuator command.

If we look at Fig. 5.42 we can learn some things about the behaviour of the calibration model. The first thing to note is that the Shapley values appear to be equally relevant in both PSF channels, which suggests that the model is using both in-focus and defocus data to decide the value of the coefficients.

In general, the calibration model makes use of a significant number of pixels. Rather than relying on specific features, the predictions are driven by the combination of Shapley values across wide areas of the PSF images. Interestingly, the wings of the PSF seem to be quite important in determining the predictions. Despite having relatively low intensity, the Shapley values for these features are high, specially when compared to the core of the PSF.

The Shapley values for one example do not tell us much about the features that are relevant to the model. Thus, we decided to investigate how the Shapley maps for a given actuator change across multiple PSF images. Fig. 5.43 shows the Shapley maps of actuator #1 for 20 different PSF images (in-focus PSF channel only). We decided to split the Shapley values according to the sign of the true actuator command. If the actuator command was positive, we only show the positive Shapley values (and viceversa), to highlight which features contributed positively towards the prediction.

Interestingly, the Shapley maps share some similarities: the same features are present across multiple images. This is important because each of those PSF images carries a random wavefront error composed not only of actuator #1, but also contributions from other ~ 50 actuators; and yet the calibration model tends to identify the same features across very different PSF images in order to predict the actuator command.

But why are these features important to the prediction of the actuator command? We hypothesised that the set of features that the Shapley maps highlight could be related to the differential features introduced in the PSF when a particular actuator is activated. In other words, that the features the calibration model considers relevant are those affected by changes in actuator commands.




Figure 5.43: Shapley maps for actuator #1 for multiple PSF images (in-focus channel only). We only show the Shapley values that match the sign of the true actuator command (above each figure). *Red* indicates positive Shapley, *blue* indicates negative Shapley. Darker shade marks the features that were relevant to the predictions. There appears to be some commonality on the features the model looks for when predicting the actuator coefficient.

The problem with this is that the effect that poking an actuator has on the PSF features depends on the underlying wavefront error, which varies from image to image. To show this, let us consider the complex pupil function with an unknown phase Φ after we slightly poke an actuator by ϵ , with h being its influence function. Introducing a Taylor expansion, under the assumption that ϵ is small:

$$P = \Pi \exp [i(\Phi + \epsilon h)] = \Pi \exp [i\Phi] \exp [i\epsilon h] \simeq \Pi \exp [i\Phi](1 + i\epsilon h) \quad (5.48)$$

we can then calculate the Fourier transform of the pupil function using the linearity property:

$$\mathcal{F}(P) = \mathcal{F}(\Pi \exp [i\Phi](1 + i\epsilon h)) = \mathcal{F}(\Pi \exp [i\Phi]) + \mathcal{F}(i\epsilon h \Pi \exp [i\Phi]) \quad (5.49)$$

To simplify the notation, we will use \mathcal{F}_Φ to denote the contribution from the underlying wavefront Φ , and \mathcal{F}_h to denote the additional contribution from poking one actuator:

$$\mathcal{F}(P) = \mathcal{F}_\Phi + \mathcal{F}_h \quad (5.50)$$

We can write the PSF as:

$$\text{PSF} = \mathcal{F}(P) \cdot \mathcal{F}^*(P) = \mathcal{F}_\Phi \mathcal{F}_\Phi^* + \mathcal{F}_\Phi \mathcal{F}_h^* + \mathcal{F}_\Phi^* \mathcal{F}_h + \mathcal{F}_h \mathcal{F}_h^* \quad (5.51)$$

If we subtract the nominal PSF (before we activate the actuator), which is $\mathcal{F}_\Phi \mathcal{F}_\Phi^*$, and ignore the second order terms $\mathcal{O}(\epsilon^2)$, the difference between PSF images caused by poking an actuator is given by:

$$\delta\text{PSF} \simeq 2\Re(\mathcal{F}_\Phi \mathcal{F}_h^*) \quad (5.52)$$

which inevitably depends on Φ . But, if we average this across many random instances of Φ , we can get an approximation of the differential impact that the actuator has on the PSF image.

To test our hypothesis, we ran the following analysis. First of all, we averaged the Shapley maps across the test PSF images for a given actuator, according to the sign of the command. In other words, we combined the positive and negative Shapley maps as shown in Fig. 5.43 into two average Shapley maps that show which PSF features were preferentially used by the calibration model to estimate both positive and negative commands.




Figure 5.44: Shapley values vs the differential PSF features for actuator #1. *Left*: average Shapley value calculated across PSF images that had a positive command. *Middle*: average Shapley value calculated across PSF images that had a negative command. *Right*: Average differential PSF features; the change that poking actuator #1 has on the PSF (logarithmic scale).




Figure 5.45: Shapley values vs the differential PSF features for actuator #4. *Left*: average Shapley value calculated across PSF images that had a positive command. *Middle*: average Shapley value calculated across PSF images that had a negative command. *Right*: Average differential PSF features; the change that poking actuator #4 has on the PSF (logarithmic scale).




Figure 5.46: Relative position of the actuators on the pupil plane. Actuators #1 and #4 used in Figs. 5.44 and 5.45 are highlighted here.

Then, we estimated the average differential PSF features with 100 pairs of PSF images. For each pair, we select a random wavefront Φ , compute the nominal PSF(Φ) and the PSF after poking the actuator PSF($\Phi + \epsilon h$), and take the difference δ PSF.

Fig. 5.44 shows a comparison between the two metrics for actuator #1, for both in-focus and defocused PSF channels. The differential PSF features are shown in logarithmic scale, since the calibration model relies quite heavily on the PSF wings. The first thing to note is that there seems to be a correlation between the differential features and the Shapley values. The latter exhibit marked fringes in both channels that are aligned with the differential PSF features. This can be confirmed by looking at the equivalent plot for a different actuator, see Fig. 5.45. For reference, the relative position of the actuators on the pupil plane is shown in Fig. 5.46.

Although it is hard to spot, for both actuators, the Shapley maps in the in-focus channel show peaks that approximately match the hexagonal pattern of the first ring of differential features. In contrast, the core of the PSF does not seem to be as relevant in terms of Shapley values as one might expect from the differential feature map. This could be due to the fact that all actuators suppress the PSF peak in approximately the same manner, making this feature less informative than the outer wings.

These results appear to support our initial hypothesis that the features highlighted by Shapely values as being relevant to the predictions are related (at least partially) to the differential features that the actuator introduces in the PSF images. This is reassuring for multiple reasons. First and foremost, it is indicative of a link between the calibration model predictions and an underlying physical process. The predictions are not arbitrary, but can be explained as a informed decision based on the link between PSF features and aberration coefficients.

Moreover, we have seen that the calibration model relies on approximately the same features to predict a particular actuator command, even when exposed to PSF images with random wavefront errors from other actuators. This is a sign of robustness that suggests the model is capable of disentangling the contributions of each actuator.

As a last note, the lack of a perfect match between the differential features and Shapley values could be explained from two perspectives. On the one hand, the calibration model is an imperfect model. It is not guaranteed that during the training stage, the calibration model will "learn" all the relevant features. As its performance is far from ideal, it follows that the Shapley values can include incorrect links between pixel features and predictions.

On the other hand, Shapley maps inform us about the features that have driven the model to make a particular prediction, but they do not directly explain *why* these features are relevant. That is a human task. We hypothesised that the differential features would be a useful way of guessing the aberration coefficients from the PSF images, and it certainly seems like the model takes this into account; but the truth is that, neural networks can learn representations that are not necessarily human-interpretable. It could be that, on top of the differential features, the calibration model is using a series of filters and complicated operations on both PSF channels that, despite being useful to predict aberrations, cannot be boiled down to a simple concept.

5.4.6.2 Impact of noise

We have seen that the wings of the PSF play a big role in the predictions of this calibration model. The Shapley maps look almost logarithmic, highlighting areas of low intensity away from the PSF core. This begs the question of what would happen if we added enough noise to the PSF images to suppress the outer wings of the PSF.

To investigate this, we added readout noise with $\text{SNR} \simeq 250$ to both the training and test datasets and repeated the training, using the same model architecture and parameters. After that, we recomputed the Shapley values and compared them to

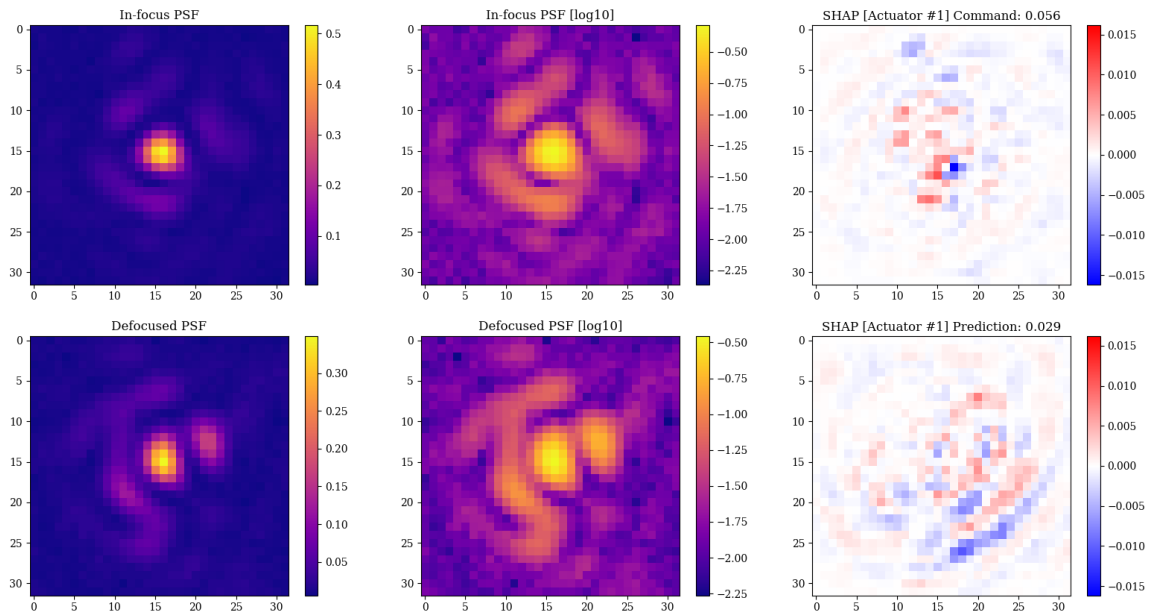


Figure 5.47: Example of Shapley values in the presence of readout noise (SNR = 250). When the images are corrupted with noise, the Shapley values suggest that the model relies less on features outside the PSF core compared to the clean case.

the previous case without noise. Fig 5.47 shows an example of the Shapley values for actuator #1, equivalent to Fig 5.42. The presence of readout noise severely contaminates the areas outside the core and first ring of the PSF, making it almost impossible to distinguish any feature. This certainly affects the Shapley maps, which seem to highlight significantly fewer features in those contaminated areas.

Interestingly, the elongated *red-blue-red* feature at the bottom right corner of the defocused channel is still present, just like in Fig. 5.42, and it seems to significantly contribute to the predicted value.

In order to properly quantify the effect that noise has had in the behaviour of the calibration model, we ran the following analysis. We calculated the average Shapley values for each actuator, as explained in the previous section (see Fig. 5.44 and Fig. 5.45), and then averaged them out across all actuators. This last step provides us with a measure of *feature importance*, by highlighting the areas of the PSF that have contributed most to all possible predictions. We did this for both in-focus and defocused channels, and for both calibration models: with and without noise.

The results are shown in Fig. 5.48. For the original calibration model without noise, the Shapley values follow the structure of the PSF, with particular emphasis on the outer regions. As expected, the defocused channel shows a more uniform distribution, as the defocus spreads the pixel intensity across the image.




Figure 5.48: Comparison of Shapley values for the calibration model trained without noise (*upper row*) and with noise (*lower row*). The values (normalised) represent the positive Shapley averaged across all actuator predictions. In the noisy case, the model appears to be less reliant on features far outside the PSF core, since these are severely contaminated and no longer contain useful information to estimate the aberrations.

In contrast, the calibration model with noise shows a significant dampening of the Shapley values in the outer regions of the PSF, most likely due to the fact that the noise renders much of those features useless to the predictions.

5.4.6.3 Visualising the layers

Another way to try to disentangle the decision process the calibration model applies to decide on a particular prediction is to look at the outputs of the convolutional layers that make up the model. In this particular case the model architecture was composed of two convolutional layers that operate on the datacubes comprised of two channels (in-focus and defocused PSF images), followed by a fully-connected layer that produces as output the predicted actuator commands associated with the PSF.

Each convolutional layer contains a series of 3x3 convolutional filters that are applied to the output of the previous layer, followed by a non-linear activation function. By looking at the output of each layer's activation function, we can probe what the model is doing to the PSF images at each step.

To illustrate this, we show an example with a random PSF image in Fig. 5.49. The calibration model (no noise) contains 16 convolutional filters in the first layer, and 8 in the second layer. Each image in the boxes represents the output of the convolutional layer for a given filter applied to the datacube of in-focus and defocus PSF images. Darker shades of *red* indicate higher pixel intensity, white areas represent values of exactly 0.0, i.e. masked regions.

We can see that for the first convolutional layer, the outputs are relatively similar to the original input. Interestingly, the model applies multiple masks to areas where the PSF exhibits marked features. As we move to the second convolutional layer, the features become more abstract and are harder to interpret.

A priori, the convolutional layer outputs do not tell us much because they represent an intermediate step in the construction of a meaningful prediction. But if we look at them in *differential* terms, by comparing the outputs before and after poking one actuator, we can investigate how the convolutional layers encode the information. For instance, Fig. 5.50 shows the difference between the convolutional layer output for a nominal PSF and the output for a PSF after poking actuator #1. These represent the sensitivity of the calibration model to changes in that particular command. Surprisingly, some of the outputs resemble the logarithmic plot for the PSF differential features (see Fig. 5.44, which seems to suggest that the role of some of the convolutional filters is to enhance these outer wing features, and use them to predict the coefficient.




Figure 5.49: Output of the two convolutional layers of the calibration model (no noise) for a random PSF image. The feature maps get increasingly abstract through the layers.




Figure 5.50: Changes in the first convolutional layer output after poking actuator #1. *Red* indicates positive changes, *blue* indicates negative changes. This shows what features are being activated due to a differential poke.

5.4.7 Impact of the diversity phase

The question of which *diversity* to use has always been a subject of interest in research on Phase Diversity as it has a wide range of implications. From a purely practical perspective, some aberrations might be easier to implement given the characteristics of the instrument. For instance, defocus is a versatile choice as it can be introduced by: displacing the detector along the optical axis, translating a calibration source, or actuating a deformable mirror.

But it can also have implications in the wavefront estimation itself. For example, not all diversity choices remove sign ambiguities. Interestingly, in the case of coronagraphic PD it has been shown empirically that a combination of astigmatism and defocus can increase the area on which the optimisation criterion is convex (Savauge et al., 2012), which facilitates the convergence of local optimisation methods.

From the perspective of a neural network model, there is no *a priori* criterion for selecting the diversity, except for the rule of thumb that the nominal and defocus channels should not be too similar, to allow the model to discriminate any sign ambiguity. In light of this, we decided to approach this from an empirical perspective, and investigate the impact that the diversity has in the performance of the calibration.

5.4.7.1 Defocus strength

The first thing we analysed was how the **strength** of a defocus diversity affects the calibration. This is important because the instrument design could impose some constraints on how much defocus we can introduce. For instance, if defocus is introduced by displacing an optical element, there could be a limit on the displacement due to mechanical constraints.

For this purpose, we ran the following experiment. Using a wavefront model based on actuators commands (50 actuators across the pupil), we began by fitting the Zernike defocus Z_f to the basis of actuator commands $\bar{\alpha}_i$. This gives us the wavefront that is closest to the Zernike defocus, in the least-squares sense.

$$Z_f \simeq \sum_i \bar{\alpha}_i \phi(x, y | x_i^c, y_i^c) \quad (5.53)$$

Then, we trained the same neural network model based on PSF datacubes for which the second channel is the actuator defocus multiple times, varying the intensity f_i of the defocus.

$$\mathcal{M}_i([\text{PSF}(\Phi_n), \text{PSF}(\Phi_n + f_i Z_f)]) \rightarrow \mathbf{a}_n \quad (5.54)$$




Figure 5.51: Impact of defocus diversity strength f on residual RMS after calibration. The RMS error at every iteration is shown, the blue bands indicate the $Q_5 - Q_{95}$ range after calibration.

where:

$$\text{PSF} = \|\mathcal{F}[\Pi \exp(2\pi i\varphi)]\|^2 \quad (5.55)$$

In these simulations, the PSF datacubes also contain readout noise, as this might have an impact on how the model performs. Fig. 5.51 shows the evolution of RMS error at every iteration as a function of the defocus strength f , for two cases of readout noise contamination. The initial RMS in the test PSF datacubes ranged between 100 and 225 nm with an average of 150 nm.

On the one hand, for low values of defocus f , the in-focus and defocus channel images are very similar, and the neural network model struggles to estimate the degenerate aberrations, leading to a relatively poor performance. On the other hand, as we increase the intensity of the defocus, we start to suppress the peak of the defocused PSF and the images become more prone to noise contamination, so the final RMS starts to grow. There seems to be an optimum defocus of around $0.20 - 0.25\lambda$ (reference wavelength of $1.5 \mu\text{m}$). An interesting thing to note is that the levels of readout noise can have a noticeable impact on the performance. For 0.25λ of defocus and SNR 500, the final RMS is around 15 nm, which is the same order of magnitude as the requirement from field-dependent aberrations (the absolute floor for the calibration); however, for SNR 250, the final RMS degrades to around 25 nm.

5.4.7.2 Arbitrary diversities

Despite its popularity as a diversity, there is nothing to suggest (*a priori*) that defocus should be the optimum choice when it comes to the performance of a calibration algorithm. [Sauvage et al. \(2012\)](#) cite the choice diversity phase as crucial and show that, in some cases, a combination of multiple aberrations (defocus + astigmatism) improves the landscape of the minimisation criterion used in Phase Diversity, exhibiting a clear global minimum compared to a series of local minima for a pure defocus. Other studies have shown that the PD techniques can operate successfully over a range of diversity aberrations ([Dolne and Schall, 2005](#)).

For that reason, we decided to explore the possibility of using arbitrary wavefronts as phase diversity. Using the same actuator-based wavefront model as before, we repeated the previous experiment but substituting the defocus for random instances of actuator coefficients. In other words, every time we train a calibration model we use a random wavefront as diversity, and record the performance. Given that the wavefront maps used were entirely random in nature, we see no apparent difference in performance compared to the classic defocus. This seems to suggest that the calibration model can deal with any diversity map, and opens up the possibility of using whatever diversity phase is easier to deploy within the instrument, which may vary depending on the available hardware.

5.5 Going beyond the training data

5.5.1 The generalisation problem

As described in Section 5.2 where we deal with the methodology of the machine learning approach, neural network models are trained and tested using datasets of sometimes very different nature. This can be regarded as a fundamental problem with the structure of the learning task itself. The neural network model is trained to maximise its predictive accuracy on *training* data for which one knows the ground truth, but it will then be deployed and make predictions on new **test** data that it has never seen before, with the ultimate goal of maximising accuracy on that test data. This is somewhat inescapable because one does **not** know what the ground truth is for the test data (otherwise the learning would be pointless) so one has to compromise and train the model using the only dataset for which ground truth exists, the *training* set, while at the same time do everything one can to guarantee that the model will be able to **generalise** outside that dataset.

We can approach this last point of generalisation issues from two different angles: one practical, another one more theoretical. The first one is to use an adequate training strategy, such as including as many representative examples as possible, trying to avoid class imbalance, using regularisation techniques or dropout (Srivastava et al., 2014) to prevent overfitting, among others.

The second one is to try to tackle the root the cause of the problem. The main reason generalisation might fail is because the training data and test data do not come from the same distribution. If they do not share anything in common, then training the model to maximise the predictive accuracy on the training set will not help us perform well on the test set. Consequently, any effort devoted to ensuring the training data matches as close as possible the data on which we are going to test the performance should mitigate the risk of poor generalisation.

In other words, making the training as realistic as possible. In the case of HARMONI, the obvious solution would be to generate the training set on the telescope itself. Using a deformable mirror and a calibration source, one could potentially create a large dataset of PSF images with varying wavefront maps Φ_n . However, until this becomes possible, there is no alternative but to rely on *simulated* data. But this does not mean that we cannot do anything to make the training realistic.

For example, in Chapter 6 I will demonstrate how the End-to-End model of HARMONI can be used to obtain wavefront maps Φ_n that accurately represent the expected range of NCPA and their distribution, and use them to train the neural net-

work. In addition, throughout this chapter I have investigated the impact of many real effects that are likely to affect the PSF images of HARMONI such as readout noise, pixel scale errors, flat field calibration uncertainties, among others. Incorporating these effects into the training set brings the example PSF images closer to the test set and makes the models more robust; and any effort devoted to making the training set more realistic will help mitigate the risk of poor generalisation.

5.5.2 The challenge of training with telescope data

As already mentioned, the ideal scenario is to train the neural network models with PSF images generated *in situ* with HARMONI installed at ESO's ELT. However, even when this becomes a reality, we foresee some difficulties.

Before one can start generating examples with varying levels of NCPA, it is important to first estimate and pre-compensate the NCPA in the as-built HARMONI instrument. The instrument will have some unknown aberrations φ_0 to begin with. If we start creating training examples by adding some known wavefront Φ_n , the training examples will be:

$$\begin{aligned}\mathcal{T}_n &= \{[\text{PSF}(\varphi_0 + \Phi_n), \dots], a_n\} \\ \mathcal{T}_m &= \{[\text{PSF}(\varphi_0 + \Phi_m), \dots], a_m\}\end{aligned}\tag{5.56}$$

where a_n describes the aberration content of the known wavefront $\Phi_n = \sum_j a_j \phi_j$, but inevitably ignore the φ_0 contribution. Since the training of the neural network is based on examples with labels (known as *supervised learning*) one cannot expect the network to learn something it has not been shown. In other words, when exposed to a training example $[\text{PSF}(\varphi_0 + \Phi_n), \dots]$ and the label a_n describing Φ_n , a perfectly trained model would retrieve \tilde{a}_n . In such scenario, the correction would leave a residual of order φ_0 , the original HARMONI NCPA. This creates a vicious cycle: to correct the unknown φ_0 we need a trained calibration network, but how can we train such network if φ_0 is present in all examples and we have no information about it?

One option to mitigate this problem is to estimate φ_0 with a neural network model trained entirely on simulated data. After correcting φ_0 , the residual error could be small enough to be disregarded (i.e. lower than the requirement from field-dependent aberrations). This might be sufficient, but there is no guarantee on how accurate that estimate will be until this is tested in situ.

Another option, is to rethink the training strategy. So far, I have only explored the *mean squared error* as a loss function or criterion to train the network, as commonly

used in regression problems. But one can also think about regression in terms of maximum likelihood estimation.

Instead of producing a single prediction \tilde{y} , the neural network model can be regarded as fitting a conditional distribution $p(y | x)$ over all possible values of y that are compatible with a given example x . Assuming a functional form of a Gaussian:

$$p(y | x) = \mathcal{N}(\tilde{y}(x, w) | \mu, \sigma^2) \quad (5.57)$$

the goal of the training is to adjust the weights w to find estimates that maximise the likelihood. In other words, to find a model that produces estimates \tilde{y} that are most likely given the data x . In practice this is achieved by minimising the negative log-likelihood during training.

When the labels are perfectly known, as it has been the case until now because we knew what aberrations we were adding, this maximum likelihood estimation directly gives rise to the mean squared error metric:

$$\frac{1}{N} \sum_i^N \|y_i - \tilde{y}_i(x_i, w)\|^2 \quad (5.58)$$

with $x_i = [\text{PSF}(\Phi_n), \dots]$, $y_i = a_n$ the known labels, and \tilde{y}_i the predicted labels.

However, as mentioned earlier, this inevitably causes issues when training at the telescope because the labels represent the added Φ_n but we have no way of probing the unknown φ_0 present in the example images $x_i = [\text{PSF}(\varphi_0 + \Phi_n), \dots]$. Instead, it is possible to give the neural network model a log-likelihood as loss function to minimise that incorporates information about how likely the predictions are.

Recalling Eq. 5.57, we know that when creating examples $\text{PSF}(\varphi_0 + \Phi_n)$, the aberrations will be biased by the presence of φ_0 in all training images. Therefore, if we had information about the aberrations in φ_0 we could incorporate this $\mu = \mu(\varphi_0)$ and $\sigma = \sigma(\varphi_0)$ into the likelihood to inform the neural network that, although we are providing labels $y = a_n$ corresponding to Φ_n , the actual aberration state is likely to be biased according to $\mathcal{N}(\tilde{y}(x, w) | \mu(\varphi_0), \sigma^2(\varphi_0))$. But where does that information come from if we do not know φ_0 in the first place?

As we will see in Section 6.2.1, NCPA in HARMONI is dominated by low order aberrations, mostly defocus and astigmatism. Even though we do not know the exact φ_0 present in the as-built instrument, we can know its *characteristics*. Using the End-to-End (E2E) Monte Carlo HARMONI optical models (details in Chapter 6), we can characterise the expected aberration content in φ_0 . Fig. 5.52 shows the distribution of defocus and astigmatism (Zernike Standard coefficients) in the wavefronts across

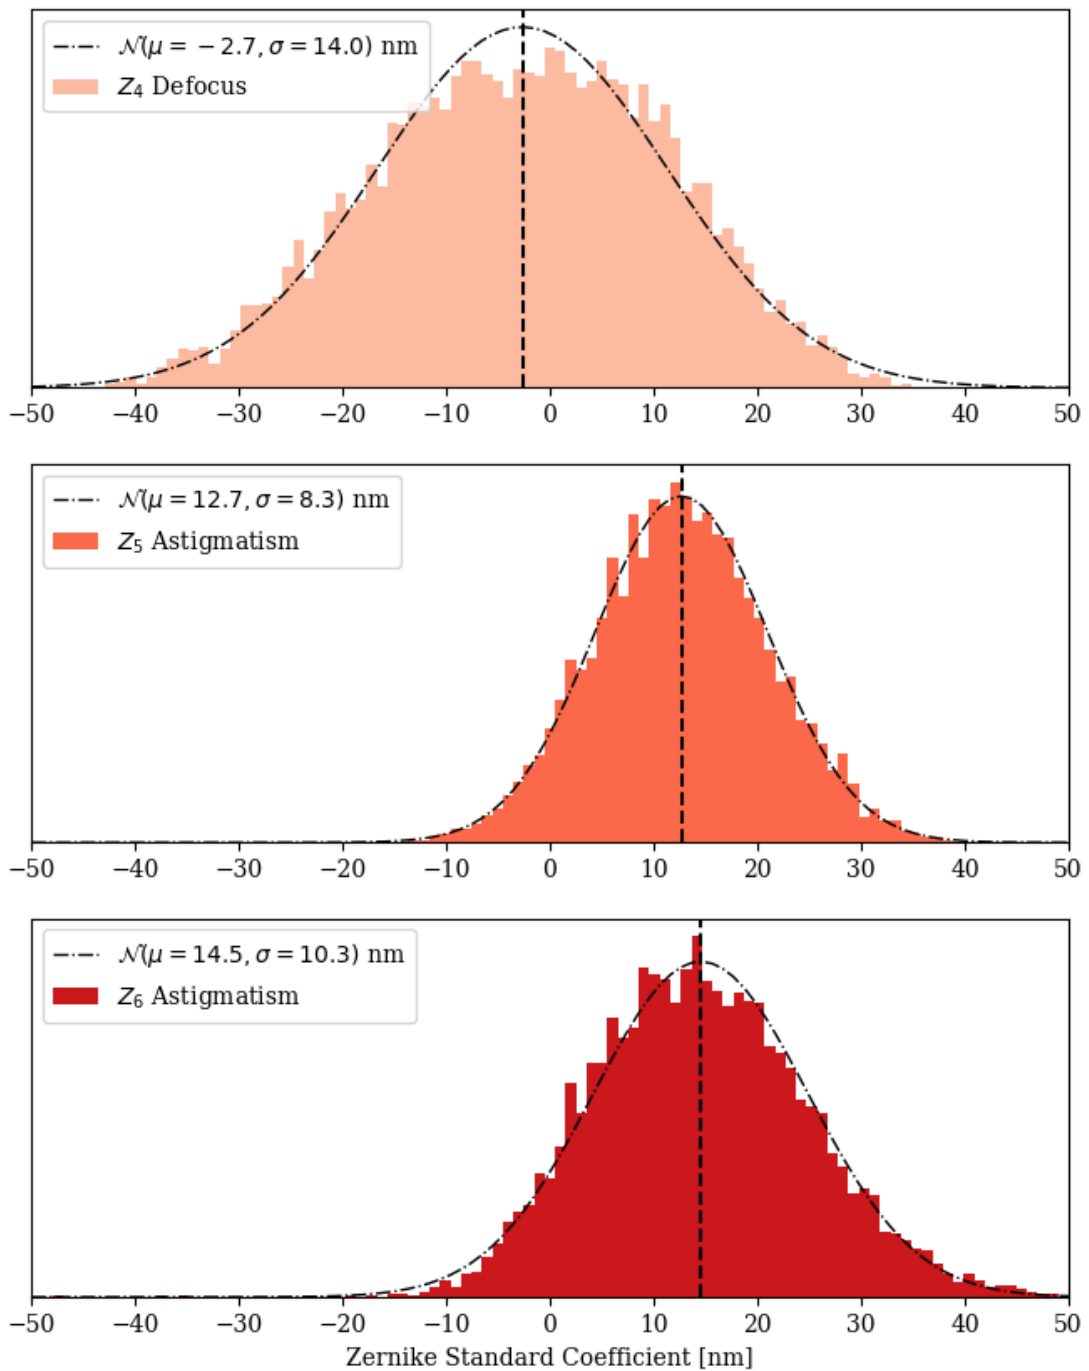


Figure 5.52: Expected NCPA from defocus and astigmatism. Distribution of Zernike Standard coefficients in the detector plan wavefront maps, calculated using the End-to-End (E2E) Monte Carlo HARMONI models, 4×4 scale, SCAO mode, H band (central wavelength $1.625 \mu\text{m}$). The histograms contain data across 30 Monte Carlo instances, all 4 IFU paths with 76 slices each, covering the complete field of view.

30 Monte Carlo instances of the as-built HARMONI instruments in 4×4 scale, SCAO mode, H band.

The results show the average Zernike coefficient and the spreads of the distributions. We can see that the as-built HARMONI instrument is likely to contain around 12-15 nm of Z_5 and Z_6 astigmatism. Fitting the data to a Gaussian distribution allows us to get a probabilistic measure of the unknown φ_0 . Beyond astigmatism and defocus, other aberrations are present in φ_0 but with significantly lower contributions.

With this information, we can modify the training of the network to incorporate this into the likelihood, such that the conditional log-likelihood can be written as:

$$\sum_i^N \log p(y_i | x_i; w) = -N \log(\Sigma) - \frac{N}{2} \log(2\pi) - \sum_{i,j}^{N,n} \frac{[(a_i^j(\Phi_i) + \mu_j(\varphi_0)) - \tilde{y}_i^j]^2}{2\sigma_j^2(\varphi_0)} \quad (5.59)$$

where N represents the number of examples in the training set and the subscript i indicates summation across those examples. The term Σ is related to the trace of the covariance matrix of the aberrations and alongside $N/2 \log(2\pi)$ are constant and do not influence the optimisation of the log-likelihood during training. The only relevant term is the double summation, where n represents the number of aberrations considered in the wavefront (i.e. the total number of Zernike polynomials) and the subscript j indicates the different aberrations.

The term $[(a_i^j(\Phi_i) + \mu_j(\varphi_0)) - \tilde{y}_i^j]^2$ resembles a mean squared error metric, but accounts for the fact that, for some aberrations $a^j(\Phi)$, there will be a bias μ_j due to the presence of φ_0 . In this case, this bias will come in the form of astigmatism (see Fig. 5.52). In addition, the terms σ_j^2 capture how uncertain we are about each aberration. For aberrations beyond astigmatism, we can be quite certain that the aberration content is dominated by the known a^j from the added Φ , so σ_j^2 will be set to low value to indicate the intensity of those aberrations is accurately known. However, for defocus and astigmatism, we have to incorporate the uncertainty from the E2E model, because φ_0 could contain those aberrations, so σ_j^2 will come from Fig. 5.52. In other words, the σ_j^2 can be grouped into a covariance matrix representing the uncertainty for each aberration, which has two types of entries: a constant low value for aberrations beyond astigmatism, and the specific uncertainties derived from the E2E model for defocus and the two types of astigmatism.

In summary, in the presence of an unknown φ_0 present in all examples created at the telescope $x = [\text{PSF}(\varphi_0 + \Phi_n), \dots]$, training the neural network with a mean squared error metric that attempts to match the labels $y = a_n$ associated with Φ_n

and the predictions \tilde{y} on a one-to-one basis would never allow the estimation and correction of φ_0 . I argue that this issue could be mitigated by approaching the training of the neural network as a maximum likelihood estimation that incorporates information about φ_0 so that the network, instead of attempting to match the labels $y = a_n$, provides an estimate \tilde{y} that is "most likely" given that we know that: (1) a_n is present and (2) that there are some unknown aberrations present due to φ_0 , whose statistics we know from the E2E model.

To demonstrate this, I ran the following experiment. The idea behind this was to mimic how the training would take place at the telescope. First of all, I sampled φ_0 from a multivariate Gaussian following the statistic from Fig. 5.52. This wavefront, which contains defocus and astigmatism represents a likely realisation of the NCPA present in the as-built HARMONI instrument. Then, I created example PSF images each with a random wavefront $\Phi_n = \sum_j a_j \phi_j$, but all sharing the underlying φ_0 :

$$\begin{aligned} \mathcal{T}_1 &= \{[\text{PSF}(\varphi_0 + \Phi_1), \dots], a_1\} \\ &\dots \\ \mathcal{T}_n &= \{[\text{PSF}(\varphi_0 + \Phi_n), \dots], a_n\} \end{aligned} \tag{5.60}$$

The labels associated with each example \mathcal{T}_n are the coefficients a_n that describe the known Φ_n , but do not contain any information about φ_0 because it is an unknown wavefront. For Φ_n I used the first 36 Zernike polynomials, whereas φ_0 only contained the 3 dominant aberrations from the E2E simulations (defocus and astigmatism). The datasets used for training and testing contained 5,000 and 500 PSF images respectively, sampled at 4×4 mas, at $1.5 \mu\text{m}$, with a defocus channel at 0.3λ PV. The images also contain readout noise contamination at SNR 500.

As explained above, during training, the neural network attempts to minimise the log-likelihood. For defocus and astigmatism, this likelihood incorporates the bias μ_j as well as σ_j^2 the spread associated with the unknown aberrations in φ_0 from the E2E simulations. For other aberrations, $\mu = 0$ and σ is set arbitrarily at 0.001λ to indicate that we are confident that the only contribution is the added aberrations Φ_n .

To illustrate how the neural network behaves after training, Fig. 5.53 shows a comparison between the true coefficients $a_n(\Phi_n)$ and the predicted coefficients \tilde{y} . The dashed black line indicates the case of a network identically matching labels to predictions. We can see that this is approximately the case for all aberrations except for the ones in φ_0 , which is the desired behaviour and an indication of successful training. In contrast, for astigmatism Z_5, Z_6 the predictions are biased towards the

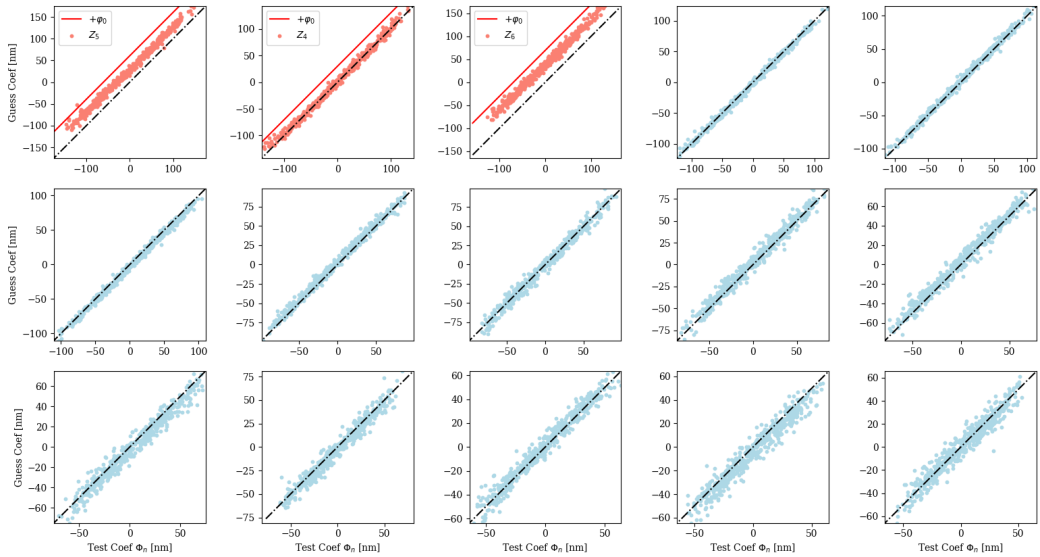


Figure 5.53: Comparison between test coefficients $a_n(\Phi_n)$ and predictions, Zernike PV coefficients. The aberrations marked *red* are defocus and astigmatism present in φ_0 but unknown to the network during training. The predictions for those aberration do not exactly follow the labels a_n because the training based on maximum likelihood estimation incorporates the statistics on φ_0 as given by the E2E Monte Carlo simulations of the expected NCPA.

red line indicating the actual value of the aberrations $a_n(\Phi_n) + a(\varphi_0)$. This shows the effect of the maximum likelihood training: the neural network predicts not just the known aberrations but over-predicts them to account for the likely presence of φ_0 .

But how does this behaviour affect the calibration? I ran the calibration for two iterations on all samples of the test set containing $\varphi_0 + \Phi_n$. The results are shown in Fig. 5.54. The top row shows the φ_0 wavefront present in all examples, representing the unknown as-built NCPA. The second row shows the complete wavefront $\varphi_0 + \Phi_n$. Interestingly, after a couple of iterations the RMS wavefront error is lower than the unknown contribution $\sigma(\varphi_0)$, suggesting that the estimates of the network allows us to correct beyond that threshold.

One important thing to note is that in order to correct for the underlying φ_0 after training, one should not just show $\text{PSF}(\varphi_0)$ to the neural network model, because this type of data was not part of the training. The network has been exposed to images that contain both $\varphi_0 + \Phi_n$. So the ideal approach is to show the network a PSF containing the unknown φ_0 and some known phase Φ representative of the training, and then start correcting.

The results shown above correspond to a single random realisation of φ_0 , but in reality, the actual NCPA in HARMONI could take any form within the distribution




Figure 5.54: Evolution of wavefront error throughout the calibration. The top row shows the unknown wavefront φ_0 present in all examples. Each column corresponds to a sample of the test set. The second row shows the starting wavefront maps $\varphi_0 + \Phi_n$. The final RMS is lower than the underlying error from φ_0 .

derived from the E2E simulations. In order to confirm that this approach to calibration works, I repeated the analysis 100 times, each with a different random instance of φ_0^k drawn from the multivariate Gaussian defined before. Each of these trials represents a potential scenario at the telescope, with some unknown φ_0^k . For every trial, I trained the neural network model, and tested the performance on 500 random PSF images with $\varphi_0^k + \Phi_n$. Then I compared the residual error after calibration (across all 500 random images) to the initial error from φ_0^k .

The results are shown in Fig. 5.55, each datapoint compares the initial unknown φ_0^j RMS to the final RMS after calibration (3 iterations). The rectangles represent the $\Delta_{90} = Q_{95} - Q_5$ range calculated across the 500 samples in each test set. The first thing to note is that in the majority of cases, the residual error after calibration is lower than the unknown φ_0 , as already suggested above, giving us confidence that this method is likely to work at the telescope. The only cases where the calibration error is worse than φ_0 tend to be those for which φ_0 was already small (between 10 and 20 nm). Since this is already close to the requirement from field-dependent aberrations (15 nm), these cases are not particularly problematic. In addition, given that φ_0 is small, this maximum likelihood estimation approach might not be necessary for these cases, and we could go ahead and generate training examples without worrying about the unknown residual.

In conclusion, this experiment should serve as a proof of concept of what could be done to tackle the issue of not knowing the underlying NCPA φ_0 when generating training examples with HARMONI at the telescope. I have demonstrated that it is possible to incorporate the uncertainty on the aberrations into the training with a log-likelihood function that the neural network optimises during training. The uncertainties on the aberrations were taken from the End-to-End model of the HARMONI instrument, which provide the expected range of NCPA from defocus and astigmatism.




Figure 5.55: Simulating the training at the telescope. Comparison between the initial RMS φ_0^j and final calibration residual $(\varphi_0^k + \Phi_n) - \tilde{y}$, for a total of 100 trial runs. φ_0^j is sampled from the E2E model distribution of expected NCPA. The rectangles indicate the $\Delta_{90} = Q_{95} - Q_5$ range calculated across 500 test samples for each trial. Each trial run represents an instance where a neural network model is trained on PSF images that contain some unknown φ_0^j using the method described above, which incorporates uncertainty on the aberrations.

5.6 Advantages and disadvantages

In this section we will explore some of the advantages and disadvantages associated with the machine learning approach to NCPA calibration. Let us start with the bad news: despite its promising performance, machine learning techniques must be used with caution as they:

- **Are sensitive to differences between training data and real data.** The performance of supervised learning algorithms strongly depends on the noise characteristics of the objects in the sample, and as such, an algorithm that was trained on a dataset with particular noise characteristics will fail to generalise to a similar dataset with different noise characteristics (Baron, 2019). Failing to generalise outside the training set is a real danger when the data used to train the models does not accurately capture the characteristics of the data on which they will be tested. Therefore, when training on simulated data because of the lack of access to real data, one should take great care to include as many effects as practically possible to ensure the simulated accurately depict the behaviour of the system. In this thesis, we have tried to mitigate this risk with efforts such as: the use of the End-to-End Monte Carlo HARMONI models to include wavefront maps in the training that are representative of the NCPA one can expect in the as-built system (Section 6.3), or the inclusion of multiple sources of uncertainty like readout noise, flat field calibration errors, diversity defocus errors, variations in anamorphic magnification, and spaxel scale errors (Section 5.4.4).
- **Require large amounts of training data.** One of the major differences between the machine learning calibration approach and conventional techniques like Phase Diversity is that it has to undergo a training stage during which the neural network model is exposed to example PSF images. We have just talked about the *quality* of said data, but machine learning models heavily rely on the *quantity* of data as well, usually requiring thousands of examples to reach competitive performance. When working with simulated data this is not an issue; for instance, the End-to-End Monte Carlo HARMONI models can readily provide thousands of representative wavefront maps. But in the pursuit of *quality*, in an ideal scenario, one would like to train on data generated at the telescope to ensure it is as representative as possible. However, it is not entirely clear how feasible this is, given that generating examples at the telescope would

be substantially more time-consuming. We can think of at least two approaches to tackle this issue.

First of all, having access to at least some real data could help update our knowledge of the system and thus refine the simulations to account for unexpected behaviour. A good example of this is the discovery of the Low Wind Effect (LWE) in VLT/SPHERE, in the form of sharp discontinuities in phase aberrations located around the spiders that appeared under certain low wind conditions due to temperature difference between the spiders and the incoming air flow (Sauvage et al., 2016). Once this issue became apparent, Phase Diversity was successfully modified to account for these high spatial frequency aberrations that the Zernike basis could not account for. A similar approach could be applied to the simulations that we use to generate training data if some unforeseen effect is identified once HARMONI is operational.

Secondly, even when the distributions on training and test set do not perfectly match, it is possible to mitigate its negative impact. The technique of Kernel Mean Matching (KMM) proposes a re-weighting of the available training data so that its distribution more closely matches that of the test data. Proposed by Gretton et al. (2009) it follows the intuitive idea of assigning more weight to examples that are most similar to the real data, and less weight to those which are not as representative or rarely occur. This could be a way of improving the training on simulations by emphasising examples that more closely resemble telescope data, without having to rely on vast amounts of data that cannot be procured in a reasonable time.

Having said this, the machine learning approach to NCPA calibration also offers significant advantages that make it a competitive alternative to traditional techniques:

- **Instantaneous correction.** Once trained, the neural network models deliver a prediction for a test PSF in an almost instantaneous manner, without the need for an iterative local optimisation and free from convergence issues.
- **Continuous improvement.** Deep learning algorithms require very little engineering by hand, so they can easily take advantage of increases in the amount of available computation and data (Lecun et al., 2015). In fact, the whole field of data mining revolves around the question of how to use the huge volumes of historical data that are routinely captured to discover general regularities

(Mitchell, 1999). This offers an opportunity to continuously improve the performance of our models with time by leveraging the data that will be collected at the telescope, such as historical trends on NCPA distribution or additional training examples. In a sense, the performance could "get better with age" as more data becomes available.

An interesting example of this is the work by Walmsley et al. (2020) using *active learning* to intelligently select the most informative galaxies for labelling by volunteers to be then used to train CNN models that classify galaxies on Galaxy Zoo. The idea behind active learning is to select which examples to label in order to train the best model, given that labelling data can be expensive and time-consuming. Walmsley et al. (2020) show that this approach leads to the same predictive performance but with significantly fewer examples. A similar strategy could be applied to the NCPA calibration models: as not all aberrations are equally hard to predict (sign-degenerate Zernike polynomials are harder than aberrations like coma), it could be possible to select which examples the model would benefit the most for learning and thus ease the demands on the amount of real data that we should have to generate at the telescope for training.

- **Exploiting data availability.** Traditional techniques for NCPA calibration like Phase Diversity are usually *monochromatic*; they rely on images taken at a fixed reference wavelength. However, as we have demonstrated in Section 5.4.5, neural network models are well suited for multidimensional data and it is possible to exploit the richness of data that integral field spectroscopy offers across a complete bandpass to improve the accuracy of the calibration. This flexibility to accommodate different data formats and capitalise on all available data is a significant advantage for these techniques.
- **Data-driven and model-agnostic.** The accuracy of PD has been shown to degrade with modelling errors, or uncertainties on the true values of the optical parameters of the system. For instance, Blanc et al. (2003) show that uncertainties in the pixel scale cause errors on the coefficients of radially symmetric aberrations; and that uncertainties in the applied diversity defocus translate into errors in the estimated defocus and other radially symmetric aberrations (Blanc et al., 2003; Jolissaint et al., 2012).

In contrast, as shown in Section 5.4.4, instead of relying on a single optical model whose parameters must be accurately calibrated, the machine learning

calibration models are data driven: they can be trained on PSF images from random instances of optical models with uncertain parameters. This mitigates the risk of the algorithms not working when deployed on the as-built instrument due to mismatches between the numerical models and the actual system.

This capacity to be trained in a model-agnostic way and effectively incorporate uncertainty into the training is a useful trait of neural network models. A good example of this is my work on the End-to-End model, see Chapter 6. Instead of using a single optical model to generate training examples, by relying on multiple Monte Carlo instances of the HARMONI instrument one ensures the calibration network is exposed to a wide range of possible NCPA scenarios.

As a concluding remark, I would like to point out that the goal of this comparison is not to perform a trade-off or to argue that the machine learning approach is superior to any other traditional NCPA calibration technique; especially given that we have not the time to thoroughly characterise the behaviour of phase diversity when confronted with HARMONI PSF images containing image slicer effects.

In fact, the idea is to highlight that, due to the radically different nature of machine learning techniques (data-driven) and phase diversity (model-based), each has its own strengths and limitations but in different aspects. No technique is without flaws and having this complementary behaviour could be the starting point for synergies. Given the challenges of NCPA calibration, having as many options as possible to tackle this problem could be essential.

For example, one of the main limitations of phase diversity is the difficulty of re-centring the in-focus and defocused images with respect to each other. Accurately identifying the location of the optical axis is crucial and can be done through a correlation procedure to align the two images to within 2π radians of tip-tilt (Blanc et al., 2003). Although this is an effect that we have not considered in this thesis, it is not too dissimilar from uncertainties on defocus strength or pixel scale studied here, and data augmentation could potentially be used to make the machine learning models robust against it.

6

End-to-End Model

“All models are wrong, but some are useful”

– George E. P. Box

6.1 Introduction

The End-to-End (E2E) optics model is the most complete and realistic representation of the HARMONI instrument up to now. In addition to analysing the nominal optical design, we have recently developed the capabilities needed to analyse Monte Carlo instances. The different sub-systems provide 100 Monte Carlo files representing possible realisations of their optical systems (including manufacturing and alignment tolerances and surface irregularity), and these are assembled together to construct a Monte Carlo instance of what the as-built HARMONI instrument could look like.

Aside from its usefulness for validating the performance of the as-built instrument, these Monte Carlo files can be a powerful tool for our NCPA calibration purposes. Since they represent realistic examples of possible HARMONI configurations they can be used to extract reasonable wavefront maps. This serves two main purposes.

First of all, it provides us with useful information regarding the distribution of aberrations in the system, allowing us to estimate the expected range of NCPA, as well as the impact of field-dependent aberrations, an intrinsic limitation to the quality of the corrections.

Secondly, it is an opportunity to re-purpose those maps for a more realistic training and evaluation of machine learning calibration algorithms. Instead of training the models on wavefront maps with uniformly distributed aberration coefficients, we can tailor the training to the specific range of aberrations that the neural network models are likely to encounter once HARMONI is deployed. This last point is important

because the performance of machine learning methods trained on simulated data strongly depends on how closely such data follows the real conditions and effects that we will see on the instrument.

6.2 Deriving calibration requirements

In Section 2.5 we discussed that in order to assess whether the NCPA calibration has been successful, it is important to define a *requirement* or a quantitative goal that states the level of acceptable NCPA, post correction. In the following sections, we will make use of the E2E model to derive such requirement based on the intrinsic limit of field-dependent aberrations and to investigate the impact of chromatic effects.

We begin by considering the question of *how many Zernike polynomials should we correct to achieve a certain degree of accuracy?* For an arbitrary wavefront map Φ , the RMS can be expressed as:

$$\sigma(\Phi) = \sqrt{\sigma_0^2 + \dots + \sigma_N^2} = \sqrt{a_0^2 \sigma^2(Z_0) + \dots + a_N^2 \sigma^2(Z_N)} \quad (6.1)$$

assuming identical aberration coefficients $a_i = a, \forall i$, and for a Standard Zernike basis where $\sigma_i^2(Z_i) = 1, \forall i$, we can write:

$$\sigma(\Phi) = a \sqrt{\sum_i \sigma^2(Z_i)} = a\sqrt{N} \quad (6.2)$$

If we apply a correction φ that covers all Zernike polynomials up to an arbitrary order n , the residual RMS is given by:

$$\sigma(\Phi - \varphi) = a\sqrt{0 \cdot \sigma^2(Z_0) + 0 \cdot \sigma^2(Z_1) + \dots + \sigma^2(Z_{n-1}) + \dots + \sigma^2(Z_N)} \quad (6.3)$$

$$\sigma(\Phi - \varphi(n)) = a\sqrt{(N - n)} \quad (6.4)$$

if we calculate the ratio of residual RMS $\sigma(\Phi - \varphi(n))$ to the initial uncorrected wavefront $\sigma(\Phi)$:

$$r(n) = \frac{\sigma(\Phi - \varphi(n))}{\sigma(\Phi)} = \sqrt{\frac{N - n}{N}} \quad (6.5)$$

we see that the residual error depends on the number of Zernike polynomials in the wavefront and the cutoff point n of our correction. It is true that N is, in principle, arbitrary and could go up to infinity for a real wavefront, but the contribution to the




Figure 6.1: Evolution of $r(n)$ the ratio of RMS after correction for up to n polynomials $\sigma(\Phi - \varphi(n))$ and original uncorrected RMS $\sigma(\Phi)$, for several cases of α , the decay factor in the distribution of Zernike coefficients. A sample of 50 randomly generated wavefront maps was used in each case.

total RMS of those high order aberrations quickly decays with order will not affect the result, as long as we choose a value of N that is sufficiently large to accurately represent the wavefront Φ .

In the derivation of this formula, we have assumed that all Zernike aberrations share the same coefficient $a_i = a$, which obviously represents an oversimplification. In order to study the effect of the distribution of coefficients, we calculated $r(n)$ experimentally across a sample of 50 random wavefront maps. The results, see *top panel* in Fig. 6.1, show how the residual RMS decays as we correct more and more Zernike polynomials, assuming the coefficients are uniformly distributed. The results closely follow the theoretical prediction for $r(n)$ from Eq. 6.5; in fact, using a Taylor

series expansion on $\sqrt{1-x} \simeq 1 - x/2$, we can see the RMS should decay linearly as:

$$r(n) \simeq 1 - \frac{n}{2N} \quad (6.6)$$

as long as n is small, or when we correct only a few polynomials. This $r(n)$ trend suggests that one has to correct for a significant number of polynomials in order to reduce the RMS.

However, this assumes the aberration coefficients are uniformly distributed, which is not entirely realistic because, as we will see later in this Chapter, the coefficients decay with radial order in the HARMONI wavefronts. A more realistic distribution of Zernike coefficients could be:

$$a_i = a\alpha^{i/N} \quad (6.7)$$

where the strength of the Zernike coefficient decays from value a for low order polynomials, to $\alpha \cdot a$ for the highest orders. The results in terms of $r(n)$ evolution as a function of number of polynomials corrected for this type of decaying distribution is shown in the lower two panels of Fig. 6.1.

This completely reverses the trend, suggesting that the RMS can be greatly improved by correcting low order aberrations. But how realistic is the α decay assumption on the distribution of RMS across Zernike polynomials? In order to investigate this, we looked at the wavefront maps for the E2E Monte Carlo models; in particular, for every wavefront we calculated the RMS σ_i associated to each Zernike aberration:

$$\sigma_i = a_i \cdot \sigma(Z_i) \quad (6.8)$$

where a_i is the Zernike coefficient and $\sigma(Z_i)$ is the RMS of a wavefront contain the i -th Zernike polynomial. This was done for the 4×4 scale in SCAO mode, which includes the contribution from the dichroic beamsplitter. Fig. 6.2 shows the range of σ_i across all aberrations. As expected, the largest contribution to the RMS comes in the form of low-order aberrations like astigmatism, defocus and coma. After that, the RMS steadily declines with Zernike order. For reference, we have added the profiles for exponential decays with $\alpha = 0.1$ and 0.01 , which capture the general trend of the RMS contribution decreasing. As the E2E Monte Carlo models represent realistic HARMONI wavefronts, we can safely conclude that our previous assumptions on the RMS decay were justified. It is worth pointing out that a significant portion of the variability in the coefficients comes from variations across the field of view within each MC instance, rather than the spread in performance between MC instances.

Fig. 6.1 was based on random wavefront maps. Now that we have access to realistic E2E Monte Carlo wavefronts, we can repeat the same analysis to obtain an accurate prediction of how the residual RMS would evolve if we applied a correction that included the first n Zernike polynomials. The results from Fig. 6.3 suggest that correcting for defocus Z_4 and astigmatism Z_5 and Z_6 already removes on average 50% of the RMS content. Adding a coma correction Z_7 and Z_8 brings the RMS further down to 25% of the initial value (0.15 – 0.4 interval). Further improvement requires correcting for significantly more polynomials, as each contributes less and less to the RMS. Around 20 polynomials are needed to reduce the RMS to 1/10 of the initial value.

This analysis was done using 30 E2E Monte Carlo instances, each containing 4 IFU paths with 76 slices each. For speed and simplicity, we have only considered field points at the centre of each slice, for a total of 304 wavefronts per MC instance. However, these field points cover both sides of the detector focal plane and densely populate the field of view, so they represent a realistic sampling.




Figure 6.2: Contribution to the RMS wavefront for each Zernike polynomial in a sample of 30 E2E files ($30 \times 4 \times 76 = 9120$ samples), 4×4 scale, H band (central wavelength $1.625 \mu\text{m}$), SCAO mode (including dichroic). Exponential decay profiles with $\alpha = 0.1$ and 0.01 are shown for reference. *Left*: linear scale. *Right*: logarithmic scale.




Figure 6.3: Evolution of $r(n)$ the ratio of RMS after correction for up to n polynomials $\sigma(\Phi - \varphi(n))$ and original uncorrected RMS $\sigma(\Phi)$ in E2E Monte Carlo wavefronts (30 E2E files, 4×4 scale, H band (central wavelength $1.625 \mu\text{m}$), SCAO mode including dichroic). *Left*: linear scale. *Right*: logarithmic scale.

6.2.1 Expected NCPA

In Section 2.5 we mentioned that in order to assess whether the calibration techniques will work, first, we need to know **how much NCPA we can expect to have initially**. Secondly, we need to know **how well we need to correct NCPA**, which we will address in the next section by studying the effect of field-dependent aberrations. Knowing this allows us to define a sort of *dynamic range* for the aberrations that the calibration models should cover.

In order to address the first point, we will use the E2E Monte Carlo instances of the HARMONI instrument. Fig. 6.2 already showed the distribution of RMS across the different Zernike polynomials in a sample of E2E MC instances for the SCAO mode. Fig. 6.4 shows a comparison between the NOAO and SCAO modes. The NOAO mode represents the essence of the HARMONI instrument and its sub-systems. We can see from the plot that the main contribution to the RMS comes from low order Zernike aberrations. The primary difference between the NOAO and SCAO mode is the inclusion of a dichroic beamsplitter that separates the light for the wavefront sensing channel and the science channel, which primarily adds astigmatism Z_5 and Z_6 . Its contribution is significant, as it increases Z_5 from around 5 nm to 13 nm (average value) and from 13 nm to 23 nm (95%-ile), and Z_6 from 7 nm to 15 nm (average value) and from 17 nm to 32 nm (95%-ile). It is important to note that NOAO and SCAO are not directly interchangeable: NOAO will be used only for coarse scales (60×30) while SCAO will be used for finer scales (4×4).

Fig. 6.5 shows the histograms of RMS wavefront error for both NOAO and SCAO modes. This covers all IFU paths, and all slices of the image slicer (central point) so it represents a wide coverage of the field of view. For the SCAO mode that includes the dichroic beamsplitter, the RMS is 30 nm on average, but it can reach as high as 55-60 nm in some cases. Compared to the NOAO mode, the SCAO has an additional 5 nm of RMS. These represent the distributions of NCPA that we can expect to encounter for the HARMONI as-built instrument, and answers the first question, at least for the 4×4 scale in H band. This analysis could in principle be extended to other scales and spectral bands, but as it is very computationally intensive we have limited it to this specific configuration. One thing that we double-checked was whether the sample of 30 MC instances was representative of HARMONI: we calculated the distribution of coefficients like Fig. 6.5 but for a single MC model comprised of the *median* instances of every sub-system. The coefficient distribution for that median case was consistent with the distribution for 30 MC shown in Fig. 6.5, which indicates the chosen instances form a representative sample.




Figure 6.4: Comparison between NOAO and SCAO (including dichroic) for the RMS wavefront for each Zernike polynomial in a sample of 30 E2E files ($30 \times 4 \times 76 = 9120$ samples), 4×4 scale, H band (central wavelength $1.625 \mu\text{m}$). The main difference is an astigmatic contribution Z_5 and Z_6 from the dichroic beamsplitter.

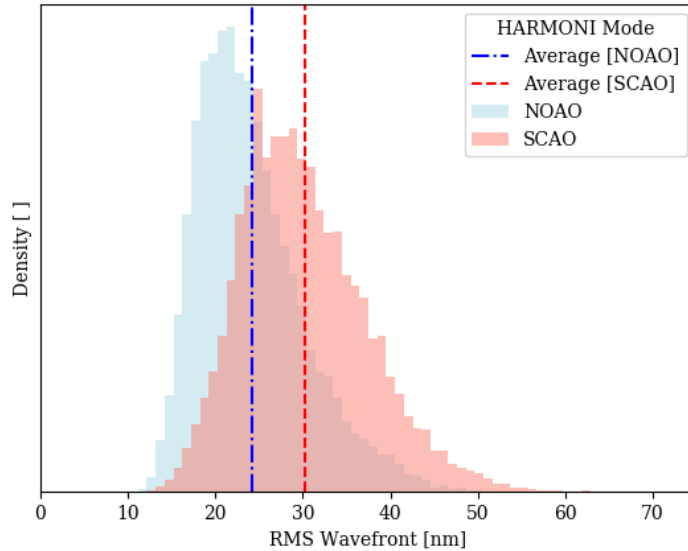


Figure 6.5: Histograms of wavefront error for both NOAO and SCAO (including dichroic) modes. 4×4 scale, H band (central wavelength $1.625 \mu\text{m}$), 30 E2E Monte Carlo instances. The data includes all 4 IFU paths and all 76 slices (field point at the centre of a slice) for every MC instance, i.e. $30 \times 4 \times 76 = 9120$ samples.

6.2.2 Intrinsic limitation: Field-Dependent Aberrations

Due to the differences in optical path after the image slicer in HARMONI, the different field points are exposed to a range of aberrations. In other words, the wavefront error that one field point sees might be slightly different to the wavefront error for another region of the field of view. As we can only apply a single correction for NCPA at the pupil plane, this *field dependency* has important implications on the calibration. In fact, it imposes an upper bound on the accuracy that we should demand from the calibration algorithms, as applying a correction that is more accurate than the intrinsic range of field-dependent aberrations yields no improvement.

In order to derive a requirement for our calibration techniques, we decided to characterise the expected range of field-dependent aberrations in HARMONI. For this purpose, we used the End-to-End (E2E) Monte Carlo models of the HARMONI instrument, which include tolerances for the different sub-systems, and surface irregularities for the spectrograph ISP sub-system.

For each MC model (4×4 scale, H band), we calculated the wavefront at the detector plane with field points located at the centre of the slice, for each slice of the image slicer, and every IFU path, i.e. a sample of $4 \times 76 = 304$ wavefront maps per MC instance. Each of these wavefront maps was decomposed into Zernike polynomials, so that we could track the variation of aberration coefficients across the field of view.

The results for defocus, astigmatism, coma and trefoil are shown for an example MC file in Fig. 6.6 and Fig. 6.7. The first point to note is that, as mentioned earlier, the aberration coefficients are far from being constant across the field. In fact, for some Zernike polynomials there is a bias in the average coefficient between IFU paths. Even if we assume a different correction will be applied to each IFU + ISP section of the field of view, there is still a significant variation in the strength of each aberration between slices. The bands in the plots indicate Δ_{95} the 95% range for the aberration coefficients across all slices. We can see that this variation is in the order of tens of nm RMS for low order Zernike polynomials. This also shows why there is an intrinsic limitation on how well we can correct NCPA. If we were to apply a correction of +20 nm defocus to the IFU-AB, some slices would indeed see their defocus well-corrected, while others would still have a ± 20 nm residual.

Another interesting thing to note is the variation between the sub-field pairs A-B, C-D, E-F, G-H as a clear break around slice #38. This is due to the fact that each sub-field has its own image slicer stack of 38 slices so that the optical paths are very different even within the same IFU sub-system.

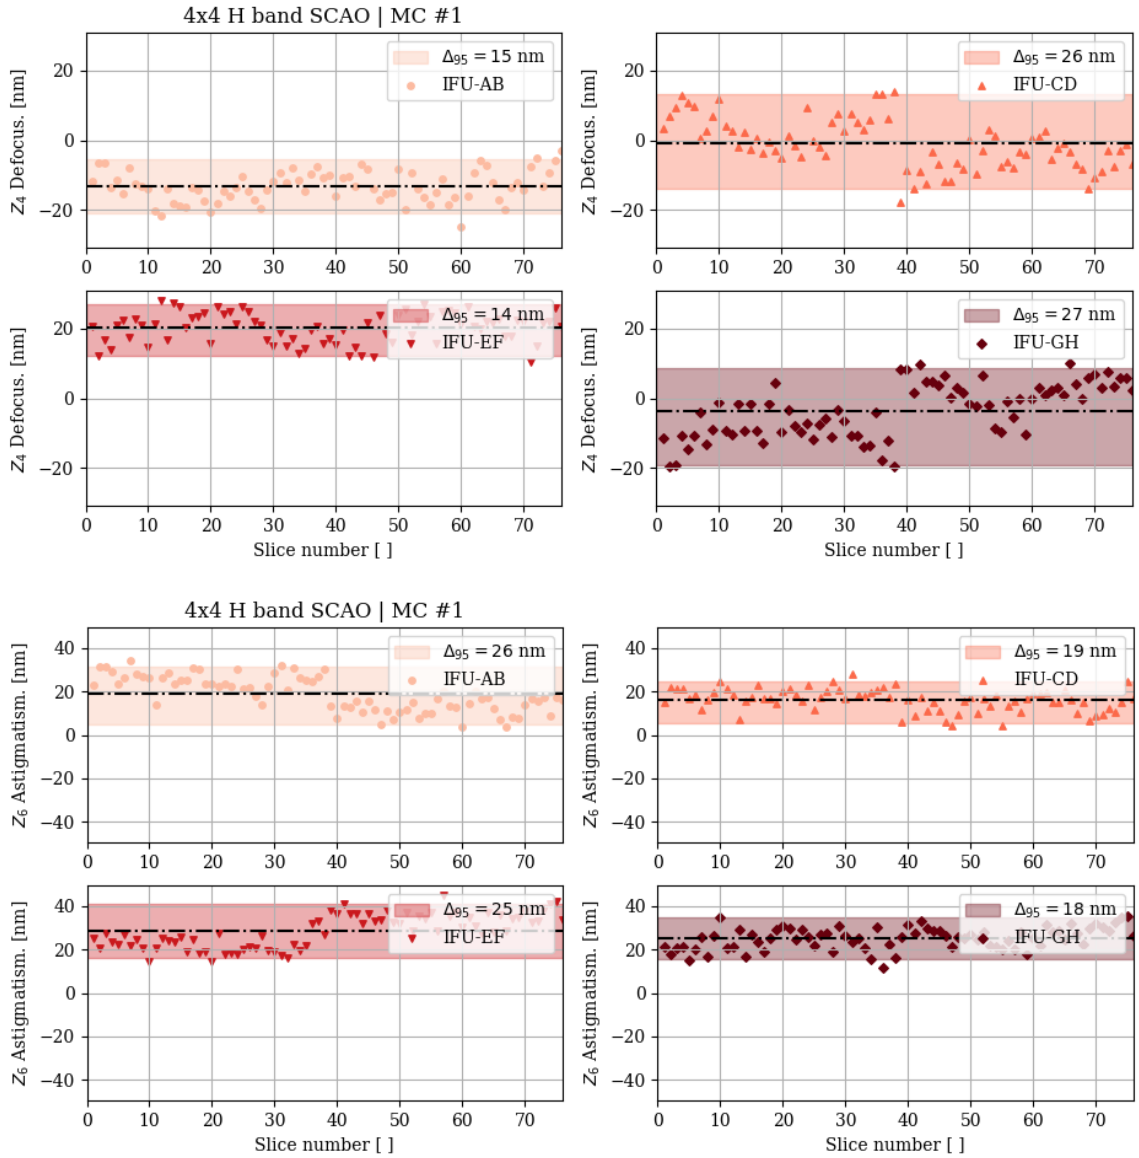


Figure 6.6: Field-dependent aberrations. Variation in aberration coefficient (Z_4 defocus and Z_6 astigmatism) across slices and IFU paths in an End-to-End Monte Carlo instance of the HARMONI instrument, 4×4 scale, H band (central wavelength $1.625 \mu\text{m}$), SCAO mode including dichroic.

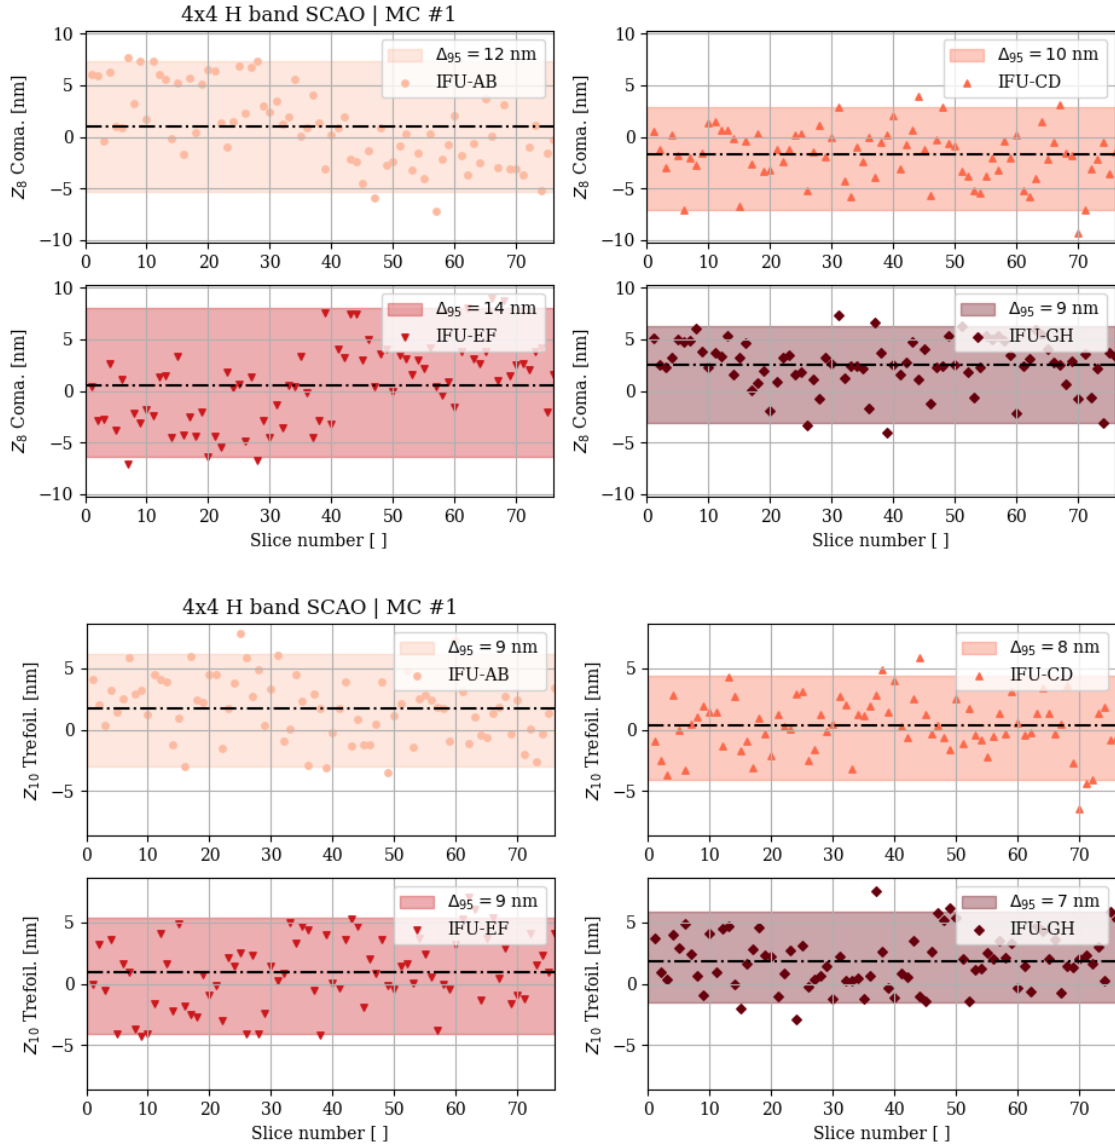


Figure 6.7: Field-dependent aberrations. Variation in aberration coefficient (Z_8 coma and Z_{10} trefoil) across slices and IFU paths in an End-to-End Monte Carlo instance of the HARMONI instrument, 4×4 scale, H band (central wavelength 1.625 μm), SCAO mode including dichroic.

Table 6.1: Average range of field-dependent aberrations $\Delta_{95}(Z_i)$ for each Zernike polynomial. Calculated across a sample of 30 E2E files, 4×4 scale, H band (central wavelength $1.625\ \mu\text{m}$). The central column assumes the FDA is calculated separately for each IFU. The right column calculates the FDA jointly across all IFU paths (one correction for the complete field of view).

Aberration	FDA Range [nm]	FDA Range all IFUs [nm]
Defocus Z_4	22.1	42.1
Astigmatism Z_5, Z_6	15.2, 20.6	22.0, 28.8
Coma Z_7, Z_8	6.8, 11.4	8.4, 12.6
Trefoil Z_9, Z_{10}	8.4, 9.0	8.9, 9.6
Spherical Z_{11}	5.8	5.9
2nd Coma Z_{12}, Z_{13}	6.5, 5.2	6.8, 5.4
Quatrefoil Z_{14}, Z_{15}	6.0, 5.5	6.3, 5.8

We can use the E2E Monte Carlo data to estimate the contribution of field-dependent aberrations to the RMS wavefront error. Using a sample of 30 files, we calculated the average range for each Zernike polynomial (from defocus to trefoil), $\bar{\Delta} = [\Delta(Z_4), \dots, \Delta(Z_{10})]$, see Table 6.1. The central column assumes the range is calculated separately for every IFU, i.e. one correction per field of view section. The right column assumes a single correction is applied to the complete field of view and thus calculates the range across all 4 IFU paths. These estimates are higher because they incorporate the effect of biases between IFU path, as seen in the case of defocus.

Fig. 6.8 shows the same result as Table 6.1, but for each IFU path instead of a global average. This allows us to investigate whether a particular IFU path has higher FDA. Although there are some differences (CD and GH paths have more defocus, while AB and EF have more astigmatism), the total FDA remains approximately constant.

For an arbitrary wavefront, the RMS due to FDA from low order polynomials can be expressed in quadrature:

$$\sigma_{FDA} = \sqrt{\sigma_1^2 + \dots + \sigma_N^2} = \sqrt{\sum a_i^2 \Delta^2(Z_i)} \quad (6.9)$$

where $a_i \in (0, 1)$ are random coefficients representing how much each Zernike is contributing within the FDA range $\Delta(Z_i)$. If we randomly sample many instances of a_i we can estimate a likely value for σ_{FDA} . For example, based on the values from Table 6.1, and 500 trials of a_i assuming a Gaussian distribution leads to $\sigma_{FDA} = 9.7 \pm 3.0$ nm. We have used the fact that $\mu \pm 2\sigma \in \Delta_{95}$ for the calculations.

This implies that even for a perfect NCPA calibration, the 4×4 scale would have around 10 nm of intrinsic (mostly low-order) field-dependent aberrations that cannot




Figure 6.8: Average range of field-dependent aberrations $\Delta_{95}(Z_i)$ for each Zernike polynomial and IFU path. Despite some variations in Z_4 defocus and Z_6 astigmatism, all IFU paths appear to have similar field-dependent aberrations.

be corrected. This assumes we apply one correction per IFU. If we apply a single correction across all IFU channels, the average FDA range $\bar{\Delta} = [\Delta(Z_4), \dots, \Delta(Z_{10})]$ will change slightly for those Zernike polynomials that have bias between IFU, changing the result slightly to $\sigma_{FDA} = 13.8 \pm 5.3$ nm. For reference, the average RMS wavefront in the E2E files is around 25 nm (NOAO) and 30 nm (SCAO), so we are already close to the constraint from field-dependent aberrations.

It should be noted that the field-dependent aberrations are approximately the same in NOAO and SCAO mode, as the contribution from the dichroic affects all field points and the FDA represent *variations* in aberration coefficients.

Based on this threshold from field-dependent aberrations and the results on the evolution of residual RMS from the previous section, we can produce an estimate of how much NCPA we can handle. Recalling the ratio $r(n)$:

$$r(n) = \frac{\sigma(\Phi - \varphi(n))}{\sigma(\Phi)} \quad (6.10)$$

knowing that the correction need not be better than the intrinsic σ_{FDA} , we can infer the acceptable RMS NCPA as:

$$\sigma_{NCPA} = \frac{\sigma_{FDA}}{r(n)} \quad (6.11)$$

which indicates the level of NCPA for which a correction of the first n Zernike polynomials would render a residual indistinguishable from the FDA contribution. If we




Figure 6.9: Inferred range of correctable (acceptable) RMS NCPA as a function of n the number of Zernike polynomials included in the correction. *Left*: ratio of how much NCPA can be corrected relative to the requirement [$\times\sigma_{min}$]. *Right*: correctable RMS [nm] for several values of σ_{min} . Calculated from the $r(n)$ profile in Fig. 6.3 based on E2E MC wavefront data, 4×4 scale, H band (central wavelength $1.625\ \mu\text{m}$), SCAO.

are considering contributions from multiple sources we can readily substitute this for $\sigma_{NCPA} = \sigma_{min}/r(n)$, where σ_{min} represents the requirement for residual NCPA errors.

Using the $r(n)$ profile from the E2E Monte Carlo analysis (recall Fig. 6.3), we show the correctable (acceptable) RMS NCPA in Fig. 6.9. The left panel shows the ratio of how much NCPA is possible to correct relative to the residual requirement σ_{min} , as a function of how many Zernike polynomials are included in the correction. For instance, a correction up to Z_{16} ensures we can always handle at least $5\times\sigma_{min}$ with a likely value of around $9\times\sigma_{min}$. So if we want to bring the residual error down to the level of field-dependent aberrations of $\sigma_{FDA} = 15$ nm, this implies we can confidently correct NCPA of at least $\sigma_{NCPA} = 75$ nm by considering up to Z_{16} .

The right panel in Fig. 6.9 shows the RMS directly, rather than a ratio, for two values of the residual requirement σ_{min} . The conservative values indicate the *minimum* NCPA we can handle (95% confidence). The difference between the conservative and average values arises from the fact that the wavefront maps in the E2E files exhibit a range of RMS distribution across Zernikes. Thus, in some cases the low-order aberration content might be smaller in proportion to the high-order aberrations, meaning that either additional Zernike polynomials are needed to reduce the RMS by the same

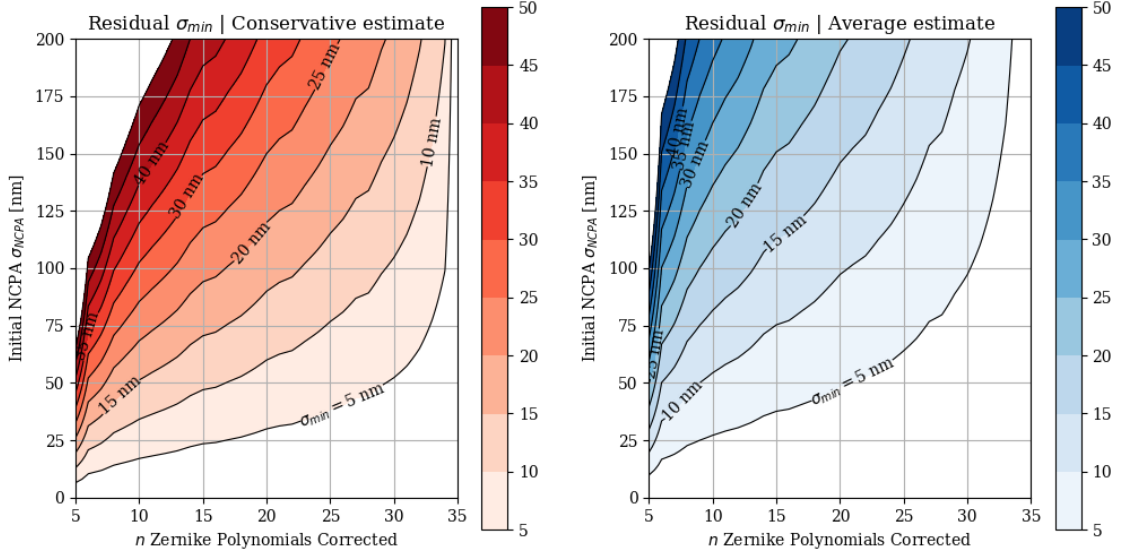


Figure 6.10: Residual RMS error σ_{min} contours as a function of the initial NCPA σ_{NCPA} and the n number of Zernike polynomials included in the correction. Left panel shows a conservative estimate, right panel shows the average estimate. Based on E2E Monte Carlo wavefront data, 4×4 scale, H band (central wavelength $1.625\ \mu\text{m}$), SCAO mode.

amount (larger n), or that we can deal with lower NCPA if we keep n fixed. The conservative estimates are taken from the 5%-ile of the RMS distribution in the E2E files.

This chart can also be read as a prescription on how many Zernikes n to correct in order to reach a certain residual σ_{min} for a given NCPA. For instance, if we know the typical NCPA is $\sigma_{NCPA} = 100\ \text{nm}$, correcting up to Z_{15} would guarantee the final residual will be around $\sigma_{min} = 25\ \text{nm}$. If we want to bring that down to $\sigma_{min} = 15\ \text{nm}$, we have to extend the correction to Z_{25} .

Fig. 6.10 extends that analysis to an arbitrary case, showing the contours of final residual after NCPA calibration σ_{min} as a function of both the initial NCPA σ_{NCPA} and n the number of Zernike polynomials included in the correction. The left panel shows conservative estimates based on limit cases 95%-ile of $r(n)$, while the right panel shows the average case. By following the contour line for a given σ_{min} we can calculate the number of Zernike polynomials needed to reach a certain performance threshold (this could be the constraint from field-dependent aberrations).

6.2.3 Feeding back the correction

To illustrate how one cannot go beyond the limit imposed by field-dependent aberrations we ran the following analysis. Using 30 E2E Monte Carlo instances, we investigated the effect of applying a correction at the entrance pupil. To explain the methodology, let us consider the following sequence:

$$[\psi_u^0 = 0] \longrightarrow \mathcal{H}[\psi_u^0] \longrightarrow \Phi_x^0 = \varphi_x^0(n) + \sum_{i=n+1} a_i Z_i \quad (6.12)$$

where ψ_u^0 represents the initial wavefront at the entrance pupil \mathbf{u} , $\mathcal{H}[\cdot]$ represents the propagation through the HARMONI instrument, Φ_x^0 is the initial wavefront at the detector plane \mathbf{x} . This wavefront Φ_x^0 can be decomposed into a low-order wavefront $\varphi_x^0(n)$ containing n Zernike polynomials plus some higher-order aberrations $Z_i, i > n$.

If we apply a correction $-\varphi_x^0(n)$ at the entrance pupil:

$$[\psi_u = -\varphi_x^0] \longrightarrow \mathcal{H}[\psi_u] \longrightarrow \Phi_x \quad (6.13)$$

the final wavefront error at the detector plane \mathbf{x} will be Φ_x . What we are comparing in this analysis is the initial RMS $\sigma(\Phi_x^0)$ and the final RMS $\sigma(\Phi_x)$ after applying a partial correction at the pupil plane \mathbf{u} based on the original information at the detector plane \mathbf{x} , $\varphi_x^0(n)$ that accounts for the first n Zernike aberrations.

Since only a single correction $-\varphi_x^0$ common to all field points can be applied but $\Phi_x^0(h_u, h_v)$ varies across the field \mathbf{h} , the correction will be sub-optimal for many field points due to the existence of these field-dependent aberrations, and this will be reflected in the final RMS $\sigma(\Phi_x)$.

To illustrate this, we present the results of the analysis for 30 E2E MC instances in Fig. 6.11, for the case of 4×4 , SCAO mode, H band. Each series (red shade) represents an MC instance of HARMONI (4 IFU paths \times 76 slices) where we compare the original RMS wavefront at the detector plane, to the final RMS value, after correcting for the first 10 Zernike polynomials. The correction $-\varphi_x^0$ was arbitrarily selected based on the wavefront map for the first slice and applied uniformly to all slices. The results indicate that most of the field points have a final RMS that is lower than its original value ($\sim 95\%$ cases). Interestingly, in some cases the correction $-\varphi_x^0$ leads to a higher RMS, which can be explained by the existence of field-dependent aberrations (for some field points, the correction adds wavefront error).

Based on the data for all 30 MC instances and field points, we can evaluate the contribution of field-dependent aberrations. Fig. 6.12 shows the density distribution




Figure 6.11: Impact of field-dependent aberrations. Comparison between initial RMS wavefront $\sigma(\Phi_x^0)$ and final RMS wavefront $\sigma(\Phi_x)$ at the detector plane \mathbf{x} , after applying a partial correction $-\varphi_x^0$ at the entrance pupil plane \mathbf{u} . The correction accounts for the first $n = 10$ Zernike polynomials. Each data series (*red* shade) corresponds to one E2E Monte Carlo instance (4 IFU \times 76 slices), 4×4 , SCAO mode, H band. For some field points / slices, the correction $-\varphi_x^0$ (uniform across the field) leads to a worse RMS.




Figure 6.12: Distribution of field-dependent aberrations. Histogram of final RMS wavefront $\sigma(\Phi_x)$ at the detector plane, after applying a correction at the entrance pupil that accounts for the first $n = 10$ Zernike polynomials. Data from 30 E2E MC instances, 4×4 scale, SCAO mode, H band.

of RMS wavefront error after correction, as well as its cumulative distribution. After that correction with the lowest 10 Zernike polynomials, what we observe is mostly the RMS wavefront error due to those intrinsic variations in aberration coefficients across the field of view.

These results from the E2E MC simulations show that the residual RMS from field-dependent aberrations is approximately 16.4 ± 5.6 nm, which is in good agreement with the theoretical prediction presented in Section 6.2.2. One can expect better performance in the areas of the field of view selected as reference to calculate the correction, but this might come at the expense of other regions of the field of view for which the aberration coefficients are different. In conclusion, here we have used the E2E Monte Carlo models to demonstrate how field-dependent aberrations represent an intrinsic limitation to the accuracy of the calibration by comparing the RMS wavefront error before and after applying a correction that is common to all field points. In addition, we have used these data to validate the original predictions for the impact of field-dependent aberrations, which appears to be around 15 nm on average.

6.2.4 Chromatic effects in field-dependent aberrations.

So far we have assumed that we are dealing with a monochromatic case; the results above were calculated for the central wavelength of the HARMONI H band (1.625 μm). However, in the same way that different areas of the field of view had different wavefront errors, different wavelengths can have varying levels of wavefront error. This is due to both wavelength-dependent performance of sub-systems with transmission optics and the fact that after the grating disperses the light, each wavelength follows a separate optical path up to the detector plane.

Therefore, wavefront maps depend not only on field of view (slices) but also on the specific wavelength. Again, as only one correction can be applied to the pupil for all wavelengths within a spectral band, there could be a chromatic variation in the strength of aberrations that imposes a limit on how well we can calibrate them. To investigate whether this chromatic variation is significant, we made use of the E2E MC models. We calculated the wavefront maps across the field of view, for $N_\lambda = 23$ wavelengths in the H band, 4×4 scale, SCAO mode. The analysis is limited to a single MC instance, given that it takes 23 times longer to run it over the whole band.

Fig. 6.13 and Fig. 6.14 show how the coefficients in the wavefront maps vary with wavelength, for defocus, astigmatism, coma and trefoil. The data points show the evolution of each aberration coefficient with wavelength for a particular slice of the image slicer. The coloured band indicates the *spectral* range Δ_{95}^λ for that slice. For comparison, we also show the *spatial* range Δ_{95}^s for each wavelength (grey band), which indicates how much the coefficient changes across the complete field of view; equivalent to the results we showed in Fig. 6.6 and Fig. 6.7, but extended to all wavelengths in the spectral band.

The spectral variation is small compared to the spatial changes Δ_{95}^s , around 5-10 nm. Even when the coefficient varies across wavelengths for a given slice, it tends to do so within the range of variation Δ_{95}^s for monochromatic field-dependent aberrations. The only exception is the defocus term, which increases steadily with wavelength, but unfortunately we cannot correct for this.

In light of this, we conclude that the chromatic effect represents a small contribution to field-dependent aberrations for all aberrations except defocus. Changes in aberration coefficients across the field of view are more substantial than variations with wavelength in the H band. This analysis was repeated for the IZJ band, the widest bandpass and the one with the shortest wavelength in HARMONI, to see if that influenced the chromatic effects, but no significant changes were observed.

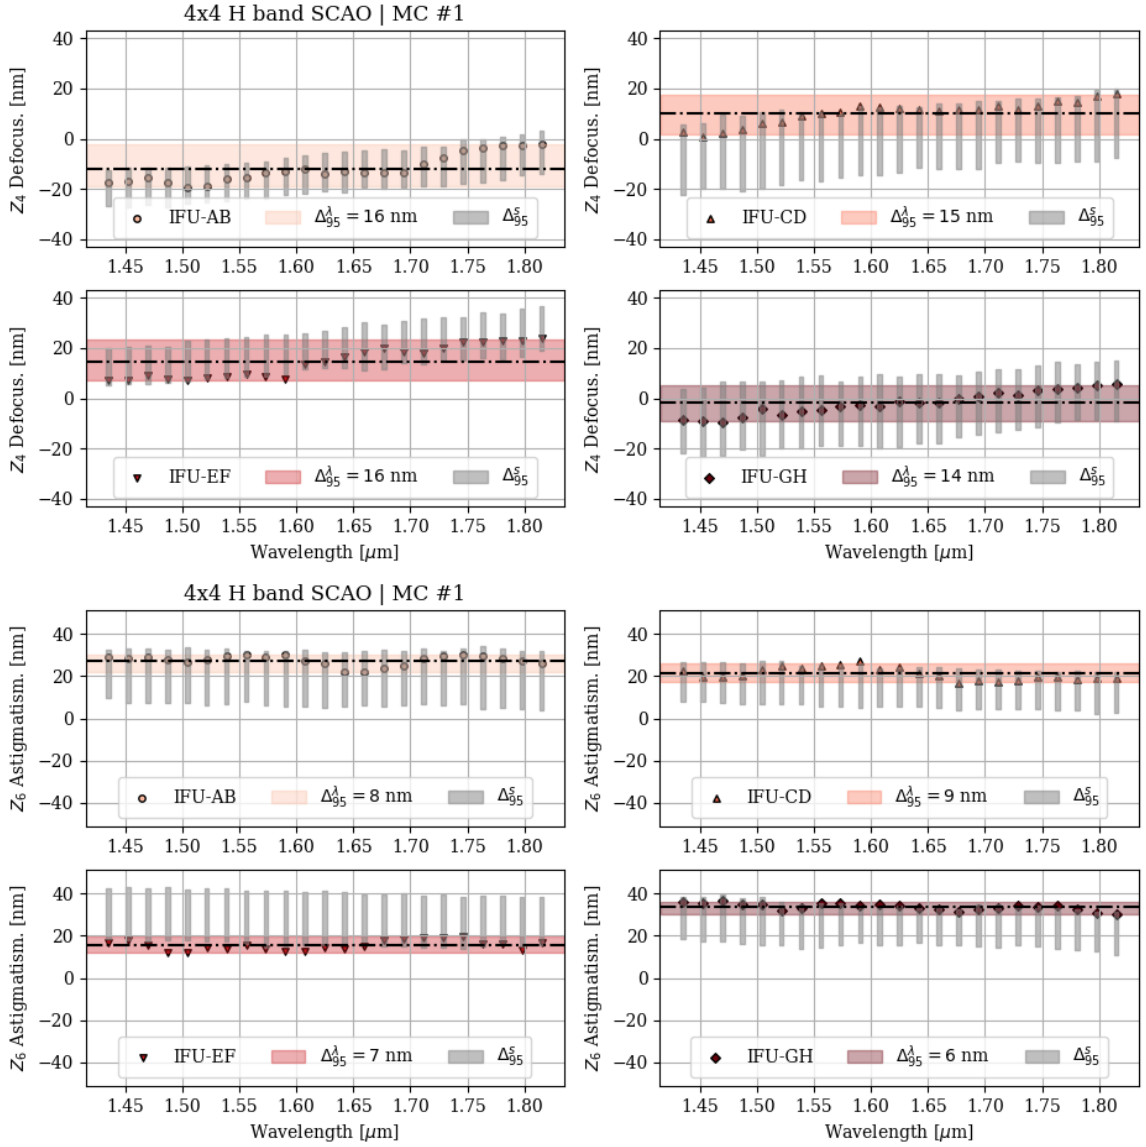


Figure 6.13: Chromatic field-dependent aberrations. Variation in aberration coefficient (Z_4 defocus and Z_6 astigmatism) with wavelength. The coloured band represents the *spectral* range Δ_{95}^{λ} for a single slice across the complete spectral band. The grey bands represent the *spatial* range Δ_{95}^s at a fixed wavelength. Calculated with an E2E Monte Carlo instance of HARMONI, 4×4 scale, H band, SCAO.

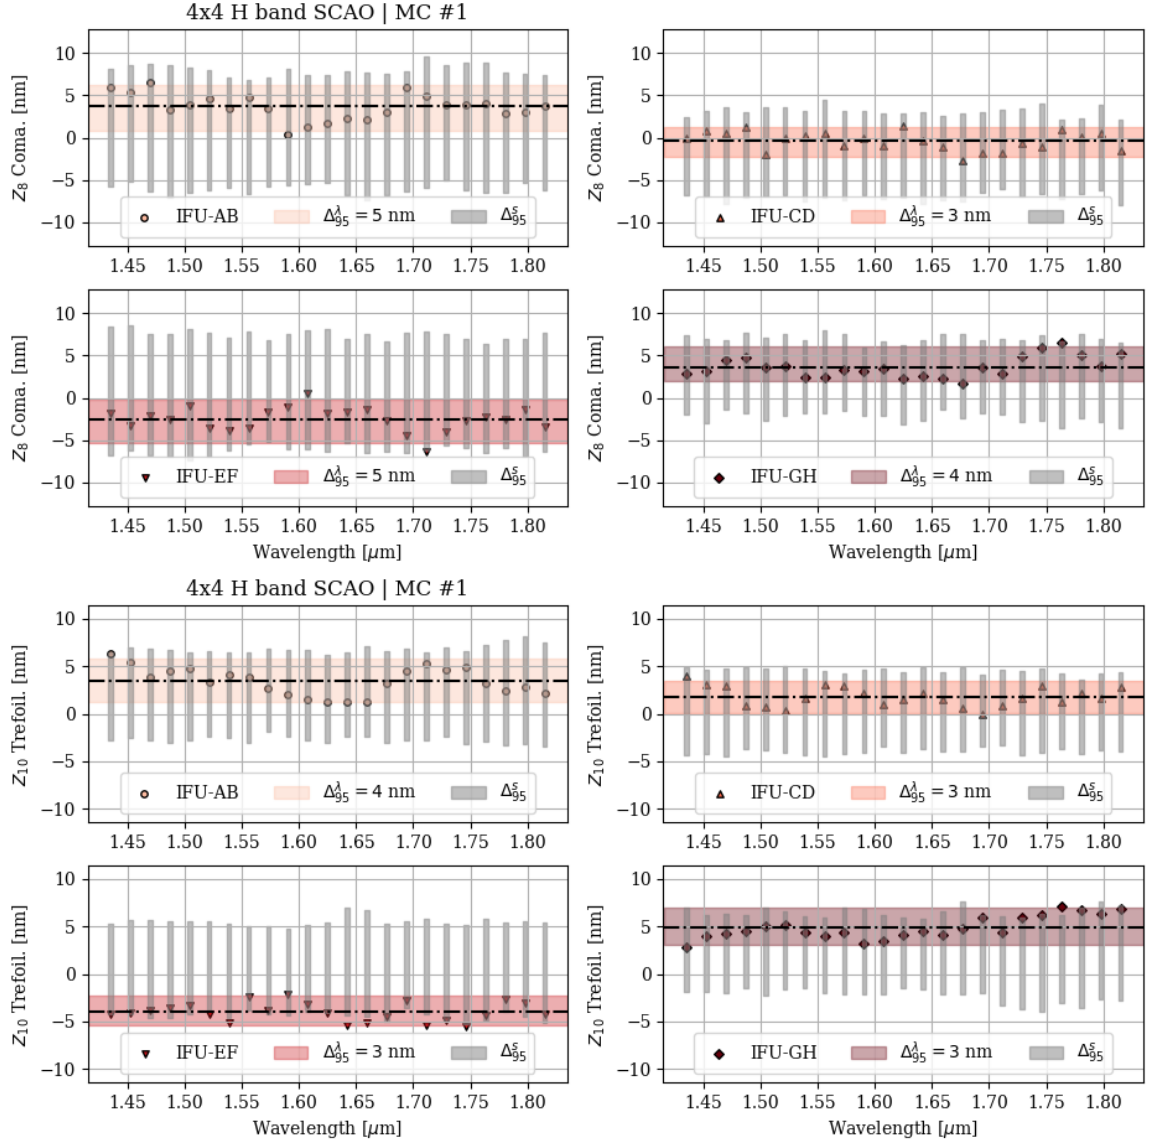


Figure 6.14: Chromatic field-dependent aberrations. Variation in aberration coefficient (Z_8 coma and Z_{10} trefoil) with wavelength. The coloured band represents the *spectral* range Δ_{95}^λ for a single slice across the complete spectral band. The grey bands represent the *spatial* range Δ_{95}^s at a fixed wavelength. Calculated with an E2E Monte Carlo instance of HARMONI, 4×4 scale, H band, SCAO.

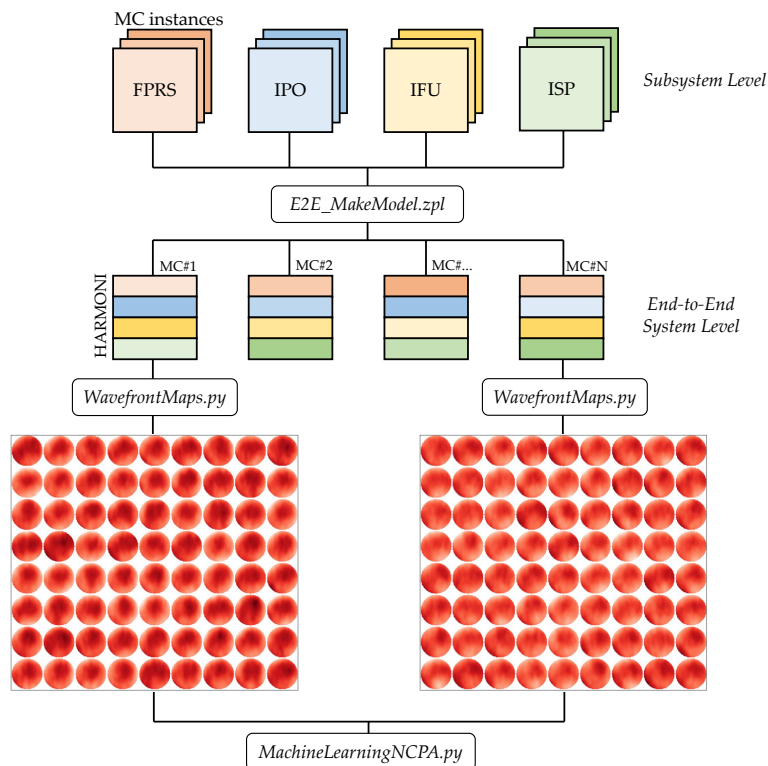


Figure 6.15: Schematic representation of the E2E Monte Carlo wavefront analysis. Each sub-system of HARMONI provides a series of MC instances for their optical models. These are selected and assembled together to construct a MC instance of the as-built HARMONI instrument. Each instance is analysed with the Python E2E toolbox to retrieve the wavefront error at the detector plane for each slice of the image slicer.

6.3 Realistic NCPA calibration

In order to generate realistic wavefront maps we follow these steps, see Fig. 6.15 for reference. Each sub-system of HARMONI (FPRS, IPO, IFU, ISP) delivers a set of Monte Carlo instances of their optical models. These files are selected and assembled to construct a MC instance of the complete HARMONI instrument. Since the E2E files only cover a single IFU path and HARMONI has four different paths, four distinct MC instances are needed to represent a complete HARMONI instance. These share the FPRS and IPO files, but have different IFU models (AB, CD, EF, and GH), as well as different MC instances of ISP sub-systems to be as realistic as possible.

Once these HARMONI MC instances are built, each one is analysed using the Python E2E toolbox that I developed to calculate the wavefront map at the detector

plane. This is done for each slice of the image slicer, for a total of 76 slices. To summarise, every MC instance of HARMONI produces 76×4 samples of wavefront error, representing a realistic distribution of aberrations across the field of view.

Now that we have the capability to generate realistic NCPA wavefronts with the Monte Carlo instances of the HARMONI instrument, we can demonstrate the use of machine learning algorithms for calibration in more representative conditions. For the purpose of this experiment we used 30 HARMONI Monte Carlo models (4×4 scale, H band, SCAO mode), and their associated wavefront maps. This dataset was split into two sets: a *training* set of 5,000 samples and a *test* set of around 750 samples. In each case, the wavefront maps were used to generate simulated HARMONI PSF images that include the level and distribution of NCPA one can expect of the as-built instrument. The PSF images were sampled at of 4×4 milli-arcseconds, assuming a wavelength of $1.50 \mu\text{m}$, to ensure sufficient resolution.

To train the calibration models, one also needs the aberration coefficients associated with each PSF. Since the HARMONI wavefronts contain high spatial frequencies coming from surface irregularities, we decided to limit the description to 78 Zernike coefficients. In other words, every HARMONI wavefront Φ_z receives a *label* containing the 78 Zernike coefficients that best describe that wavefront. This approximate (least-squares) description Φ_0 captures most of the wavefront structure, and is used as ground truth to train the calibration models. Extending the description to capture all the surface irregularity would simply complicate the task of estimating the coefficients without really adding much in terms of error correction.

Once the PSF datacubes are created based on the HARMONI wavefronts Φ_z , the calibration model was trained to estimate the aberration coefficients. For the diversity channel, we use $1/3 \lambda$ defocus. The performance of the algorithm was tested using the *test* set PSF images; the results are shown in Fig. 6.16 for four example wavefront maps. The images used here contained readout noise with SNR 500.

The first thing to note is that the Zernike fit Φ_0 used as ground truth to train the network does capture most of the structure of the HARMONI MC wavefront Φ_z , only missing the high spatial frequency surface irregularities. Secondly, with regards to performance, the calibration model estimates the coefficients with great accuracy, recovering most of the aberration content. Starting from wavefront errors of around 30 nm, typical of this SCAO configuration, the residual after correction (1 iteration only) is smaller than 10 nm, which is below the field-dependent aberration threshold. The remaining wavefront contains a few artefacts close to the edge of the




Figure 6.16: Performance of ML model trained on E2E HARMONI Monte Carlo wavefronts. *Left*: detector plane wavefront maps $[\lambda]$ extracted from the HARMONI MC files. *Left-centre*: Zernike fit to 78 polynomials used as ground truth to represent the wavefront. *Right-centre*: estimated wavefront from the calibration model. *Right*: residual error containing mostly surface irregularity. S indicates the Strehl ratio of the PSF before and after correction. PSF images contaminated with readout noise of SNR 500.




Figure 6.17: Evolution of residual wavefront error after calibration with the SNR from readout noise contamination. Model trained on E2E HARMONI Monte Carlo wavefronts, 4×4 scale, H band, SCAO mode.

pupil, which are caused by the differences in masking between the HARMONI MC wavefronts (Zemax) and the Zernike fits (Python).

Once again, I must highlight the importance of striving for sufficient SNR. I repeated the analysis with the same datasets, varying the levels of readout noise contamination. Fig. 6.17 shows the distribution of RMS wavefront error in the test set (90% range $\Delta_{90}(\sigma_0)$ and average) and the evolution of that metric after calibration as a function of SNR. The limit of field-dependent aberrations is also shown for reference. Readout noise contamination clearly plays a role in limiting the performance of the neural network model. SNR values above 200 seem to be sufficient to reach the requirement, although it is clear that the higher the SNR, the better the accuracy of the calibration.

Before we conclude this Chapter, it is important to point out some of the limitations of this experiment. The Monte Carlo E2E analysis is computationally intensive and time-consuming. Therefore, we have limited this study to a single spectral band (H band), in a monochromatic case, and using field points located at the centre of the slice only. The monochromatic assumption is not particularly relevant, since only a single NCPA correction can be applied ignoring any potential chromatic effects and we have seen that chromatic effects are negligible. With regards to the centre of the

slice, because we are using all the slices, this covers the whole field of view and area of the detector; moreover the fact that we have explored multiple Monte Carlo instances already introduces a level of variability in the wavefronts that could mimic changes along the slice. Therefore, there is no reason to believe that this approach would not work for any other HARMONI configuration.

In conclusion, we have explored the possibility of using the E2E HARMONI Monte Carlo models to characterise NCPA and train the calibration networks using realistic wavefront maps. First of all, we created a Python script capable of extracting the wavefront maps at the detector plane across multiple Monte Carlo instances of the HARMONI instrument. These wavefront maps are representative of the aberrations one can expect in the as-built instrument and thus, serve as a way of characterising NCPA as well as estimating its variability.

Moreover, these realistic wavefronts can be re-purposed to generate HARMONI PSF images that can be used to train the calibration models. To demonstrate this, we randomly selected 5,000 of these wavefronts and trained a simple CNN model to estimate the aberration coefficients. After testing the performance on another set of ~ 750 images, we can conclude that our calibration models can easily deal with realistic HARMONI wavefront maps, paving the way for the application of these techniques in the future HARMONI instrument.

7

Conclusions and Future Work

“Now I’m a scientific expert; that means I know nothing about absolutely everything.”

– Arthur C. Clarke, *2001: A Space Odyssey*

7.1 Future Work

In this section, we provide a brief overview of how the work presented in this thesis could be improved upon and expanded in the future.

7.1.1 Impact of image slicer effects on Phase Diversity

In this thesis, I have explored the question of NCPA calibration in the context of image slicers from multiple viewpoints. First of all, I have investigated the impact of image slicers on the PSF from a theoretical perspective (Chapter 3) and from an experimental perspective (Chapter 4). These efforts have been focused on understanding how slicers affect light propagation, and how this can be modelled. Then, I have worked on developing novel calibration techniques that could circumvent the drawbacks from slicer effects on the PSF images (Chapter 5).

But there remains an open question of how traditional techniques like Phase Diversity would respond to image slicer effects. Studies have pointed out the importance of ensuring that the optical model used in PD accurately reflects the behaviour of the real system, as uncertainties in the parameters can lead to errors in the estimated aberrations [Blanc et al. \(2003\)](#); [Jolissaint et al. \(2012\)](#). In light of this, it is natural to think that the presence of image slicer features on the PSF should have some impact on the behaviour of the PD algorithm.

In the future, it would be interesting to investigate whether this is the case and quantify its extent in order to assess the applicability of PD as an NCPA calibration technique in the context of image slicers. The tools for this research are already available: the mathematical framework I presented in Section 3.3.1 allows us to simulate PSF images with varying image slicer characteristics (this is also supported by Zemax models of the HARMONI IFU), and our colleagues at the Laboratoire d’Astrophysique de Marseille (LAM) have kindly shared a PD code with us.

7.1.2 Extending the End-to-End Model work

Given the computational demands of the End-to-End Monte Carlo calculations, we limited the extent of the analyses on the wavefront maps and the effects of field-dependent aberrations. Only the 4×4 scale was analysed as it is the most relevant for calibration, and I selected the H and IZJ bands out of the 11 available spectral bands in HARMONI.

In the near future, it would be interesting to extend these analyses to the remaining spaxel scales (10×10 , 20×20 and 60×30) and the rest of the spectral bands to obtain a complete characterisation of the NCPA distribution across all instrument configurations and assess the impact of field-dependent aberrations in each case.

7.2 Conclusions

The calibration of Non-Common Path Aberrations has become crucial in recent years, as these aberrations have been identified as the main limitation to high-contrast observations of exoplanets and circumstellar disks. With the gain in popularity of the field of exoplanet research, and the sought-after goal of detecting and characterising an Earth analogue around a solar-type star almost within reach, this technical challenge moved to the forefront of research in instrumentation for astrophysics. The coming generation of extremely large telescopes (ESO’s ELT, the Thirty Meter Telescope, and the Giant Magellan Telescope) are promising to bridge this gap and revolutionise astrophysics for years to come, and the successful calibration of NCPA is an essential element of this endeavour.

The work presented in this thesis demonstrates how NCPA can be measured in instruments containing image slicers and, in particular, in the case of HARMONI at ESO’s ELT. My work has been focused on several sub-questions that are relevant to this problem.

First of all, I have identified a gap in the state-of-the-art literature regarding the calibration of NCPA on instruments such as HARMONI, an integral-field spectrograph containing image slicers. Extensive research has been devoted to developing techniques either in the context of Adaptive Optics sub-systems, or coronagraphic instruments for high-contrast imaging. But instruments like HARMONI make use of an optical component known as an *image slicer* made out of a stack of thin mirrors that takes a 2D field of view and slices it into a long pseudo-slit that can be input to a spectrograph. The effects that an image slicer has on light propagation could have a negative impact on the behaviour of state-of-the-art techniques for NCPA calibration, especially when operating in defocused conditions as required by Phase Diversity. At the start of this project, no clear strategy for NCPA calibration was defined for HARMONI, apart from a cumbersome Focal Plane Sharpening approach iterating over different aberration modes.

The first step towards defining a strategy for NCPA calibration for an instrument that contains image slicers is understanding how exactly this optical component affects the behaviour of the system. Therefore, my initial efforts were focused on characterising the effect of image slicers on light propagation, see Chapter 3. From a theoretical perspective, I first developed a mathematical framework that captures the most relevant characteristics of HARMONI's Integral Field Unit (IFU) and allows me to simulate an image slicer PSF. This framework predicts that the image slicer modifies the PSF introducing a series of horizontal fringes (parallel to the slice boundary), whose spatial frequency depends on the size of the post-slicer diffraction pattern relative to the pupil mirror aperture. This highlights the need to properly over-size the pupil mirrors after the image slicer to capture the diffracted light and mitigate this effect. These effects are partially mitigated due to the coarse HARMONI sampling, with each slice being sampled by only two pixels. Because of this, the fringes are not resolved, but the overall light loss affects the total intensity, specially at the central slice boundaries.

I validated these predictions using Physical Optics Propagation (POP) on the Zemax optical models of the HARMONI IFU, comparing the impact of parameters such as wavelength and aperture size. The good agreement between the two independent approaches provides confidence that these slicer effects are indeed real, but it suffers from the drawback of both methods being based on simulations. In light of this, I decided to further validate this with experimental data. For this purpose, I designed, constructed and conducted an experiment using the engineering model of the SWIFT image slicer (given the absence of HARMONI hardware at this stage). The results

from this experiment presented in Chapter 4 confirm that the PSF exhibits fringes that are consistent with image slicer effects predicted by simulations.

At this point, it is clear that image slicers can have a significant impact on light propagation and ultimately modify the PSF. Although it is not entirely clear and will be a topic of further research in the future, these effects threaten to render traditional techniques like Phase Diversity unusable in the context of HARMONI, unless modified to account for the presence of image slicers. Having taken this uncertainty into account, I decided to explore potential alternatives that could circumvent this drawback. In Chapter 5, I presented a novel technique for NCPA calibration that relies on the use of machine learning methods to estimate the aberrations directly from PSF images.

The first thing to prove was that machine learning techniques could discriminate between features from image slicers and features caused by aberrations in the optical system. I showed that this is indeed true, and that neural network models can effectively "learn" an *ad hoc* mapping from PSF features to the aberration coefficients of the underlying wavefront errors that caused such features.

Having shown this, I extensively characterised the response of this machine learning calibration algorithm to a whole range of parameters, scenarios and sources of uncertainty including: field of view, readout noise contamination, flat field calibration errors, pixel scale errors, diversity defocus errors, and anamorphic magnification errors. I showed that one of the advantages of this machine learning approach compared to traditional calibration strategies is its data-driven nature that allows us to incorporate some level of robustness into the training by including data from multiple instances of optical models with uncertain parameters.

In addition to this, I explored how NCPA calibration could be performed at the different spaxel scales in HARMONI, and demonstrated that for the 20×20 scale where the PSF is far from being Nyquist-sampled a modified strategy would still allow calibration. I also proved that it is possible to exploit the integral-field spectroscopy data readily available in HARMONI and incorporate information across a complete bandpass to enhance the accuracy of the calibration.

Lastly, I dedicated a substantial amount of time to developing the capability to create and analyse End-to-End Monte Carlo models of HARMONI that aggregate the different sub-systems to form a realistic representation of how the as-built instrument would perform. I used this to support the calibration efforts in multiple ways. First of all, to characterise the expected level of NCPA in the system and its

major contributors, showing that the wavefront error is dominated by low-order aberrations in the form of defocus, astigmatism and coma. Secondly, to assess the impact of field-dependent aberrations, the intrinsic limit on the accuracy of the calibration, and use that as a requirement against which I could benchmark the performance of the calibration algorithm. And finally, to train the neural network models with realistic wavefront maps that reproduce the typical NCPA that we will encounter in HARMONI.

In conclusion, measuring Non-Common Path Aberrations within HARMONI using machine learning techniques is feasible as they can deal with the presence of image slicer effects. Although work remains to be done before this becomes a reality, throughout this thesis I have shown that this approach is flexible and offers a series of competitive advantages compared to conventional techniques.

Bibliography

Sheelu Abraham, Arun Aniyan, Ajit K. Kembhavi, et al. Detection of Bars in Galaxies using a Deep Convolutional Neural Network. *Monthly Notices of the Royal Astronomical Society*, 477(1):894–903, nov 2018.

Yann Alibert and Julia Venturini. Using Deep Neural Networks to compute the mass of forming planets. *A&A*, 626(A21), mar 2019.

Torben Andersen, Mette Owner-Petersen, and Anita Enmark. Neural networks for image-based wavefront sensing for astronomy. *Optics Letters*, 44(18):4618, sep 2019.

Torben Andersen, Mette Owner-Petersen, and Anita Enmark. Image-based wavefront sensing for astronomy using neural networks. *Journal of Astronomical Telescopes, Instruments, and Systems*, 6(03):1, 2020.

Vijay Arya, Rachel K E Bellamy, Pin-Yu Chen, et al. AI Explainability 360: An Extensible Toolkit for Understanding Data and Machine Learning Models. *Journal of Machine Learning Research*, 21:1–6, 2020.

Manda Banerji, Ofer Lahav, Chris J. Lintott, et al. Galaxy Zoo: Reproducing Galaxy Morphologies Via Machine Learning. *Monthly Notices of the Royal Astronomical Society*, 406(1):342–353, aug 2009.

Dalya Baron. Machine Learning in Astronomy: a practical overview. *arXiv preprint arXiv:1904.07248*, apr 2019.

Yoshua Bengio. Practical Recommendations for Gradient-Based Training of Deep Architectures. In Grégoire Montavon, Geneviève B. Orr, and Müller Klaus-Robert, editors, *Neural Networks: Tricks of the Trade: Second Edition*, pages 437–478. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.

Christopher M Bishop. *Pattern recognition and machine learning*. Springer, 2006.

- A. Blanc, T. Fusco, M. Hartung, et al. Calibration of NAOS and CONICA static aberrations. Application of the phase diversity technique. *A&A*, 399(1):373–383, 2003.
- Markus Bonse. Master Thesis: Machine Learning based Atmosphere Prediction for Extreme Adaptive Optics. *ETH Zurich*, 2019.
- Steven P. Bos. The polarization-encoded self-coherent camera. *A&A*, 646(A177), jan 2021.
- Michele Cappellari. Structure and Kinematics of Early-Type Galaxies from Integral Field Spectroscopy. *Annual Review of Astronomy and Astrophysics*, 54(1):597–665, sep 2016.
- Alexis Carlotti, François B. Hénault, Kjetil Dohlen, et al. System analysis and expected performance of a high-contrast module for HARMONI. In *Astronomical Telescopes + Instrumentation*, page 352. SPIE-Intl Soc Optical Eng, jul 2018.
- François Chollet and Et al. Keras. <https://keras.io>, 2015.
- Johan L. Codona. Differential optical transfer function wavefront sensing. *Optical Engineering*, 52(9):097105, sep 2013.
- Johan L. Codona and Nathan Doble. James Webb Space Telescope segment phasing using differential optical transfer functions. *Journal of Astronomical Telescopes, Instruments, and Systems*, 1(2):029001, feb 2015.
- A. D’Isanto, S. Cavuoti, M. Brescia, et al. An analysis of feature relevance in the classification of astronomical transients with machine learning methods. *Monthly Notices of the Royal Astronomical Society*, 457(3):3119–3132, apr 2016.
- Kjetil Dohlen, Fabrice Madec, Mamadou N’Diaye, et al. Lab demonstration of the Zernike phase mask near-coronagraph quasi static aberrations sensor, ZELDA. In *3rd AO4ELT Conference - Adaptive Optics for Extremely Large Telescopes*. INAF - Osservatorio Astrofisico di Arcetri Largo Enrico Fermi, 2013.
- Jean J Dolne and Harold B Schall. Cramer-Rao bound and phase-diversity blind deconvolution performance versus diversity polynomials. *Applied Optics*, 44(29): 6220–6227, 2005.

- Frank Eisenhauer and Walfried Raab. Visible/Infrared Imaging Spectroscopy and Energy-Resolving Detectors. *Annual Review of Astronomy and Astrophysics*, 53(1):155–197, aug 2015.
- ESO. *E-ELT Design Reference Mission*. European Southern Observatory, E-TRE-ESO-080-0717 Issue 2, 2011.
- J. R. Fienup. Phase retrieval algorithms: a comparison. *Appl. Opt.*, 21(15):2758–2769, Aug 1982.
- G. W. Forbes. Robust and fast computation for the polynomials of optics. *Opt. Express*, 18(13):13851–13862, Jun 2010.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics New York, 2001.
- A S Fruchter and R N Hook. Drizzle: A Method for the Linear Reconstruction of Undersampled Images. *Publications of the Astronomical Society of the Pacific*, 114:144–152, 2002.
- R. Galicher, P. Baudoz, G. Rousset, et al. Self-coherent camera as a focal plane wavefront sensor: Simulations. *A&A*, 509(1), 2010.
- Raphael Galicher, Pierre Baudoz, Jacques-Robert Delorme, et al. Minimization of non common path aberrations at the Palomar telescope using a self-coherent camera. *A&A*, 631:A143, sep 2019.
- Benjamin L. Gerard, Christian Marois, and Raphaël Galicher. Fast Coherent Differential Imaging on Ground-Based Telescopes using the Self-Coherent Camera. *The Astronomical Journal*, 156(3), jun 2018.
- Benjamin L. Gerard, Christian Marois, Raphaël Galicher, et al. The Fast Atmospheric Self-Coherent Camera Technique: Laboratory Results and Future Directions. *Proceedings of Adaptive Optics for Extremely Large Telescopes VI*, oct 2019.
- Fabian Gieseke, Steven Bloemen, Cas van den Bogaard, et al. Convolutional neural networks for transient candidate vetting in large-scale surveys. *Monthly Notices of the Royal Astronomical Society*, 472(3), aug 2017.
- C. A. Gomez Gonzalez, O. Absil, and M. Van Droogenbroeck. Supervised detection of exoplanets in high-contrast imaging sequences. *A&A*, 613:A71, may 2018.

- Robert A Gonsalves. Phase retrieval and diversity in adaptive optics. *Optical Engineering*, 21(5):829–832, 1982.
- Damien Gratadour, Éric Gendron, and Gérard Rousset. Practical issues with phase diversity for NCPA compensation resolved on the CANARY demonstrator. In *3rd AO4ELT Conference - Adaptive Optics for Extremely Large Telescopes*. INAF - Osservatorio Astrofisico di Arcetri Largo Enrico Fermi, 2013.
- Arthur Gretton, Alex Smola, Jiayuan Huang, et al. Covariate Shift by Kernel Mean Matching. *Dataset shift in machine learning*, 3(4), 2009.
- Olivier Guyon. Extreme Adaptive Optics. *Annual Review of Astronomy and Astrophysics*, 56(1):315–355, 2018.
- Olivier Guyon, Taro Matsuo, and J Roger P Angel. Coronagraphic low-order wavefront sensor: Principle and application to a phase-induced amplitude coronagraph. *Astrophysical Journal*, 693(1):75–84, mar 2009.
- M. Hartung, A. Blanc, T. Fusco, et al. Calibration of naos and conica static aberrations - experimental results. *A&A*, 399(1):385–394, 2003.
- Yashar D. Hezaveh, Laurence Perreault Levasseur, and Philip J. Marshall. Fast automated analysis of strong gravitational lenses with convolutional neural networks. *Nature*, 548(7669):555–557, aug 2017.
- Sasha Hinkley, Ben R Oppenheimer, Rémi Soummer, et al. Temporal evolution of coronagraphic dynamic range and constraints on companions to Vega. *The Astrophysical Journal*, 654:633–640, 2007.
- H. J. Hoeijmakers, H. Schwarz, I. A.G. Snellen, et al. Medium-resolution integral-field spectroscopy for high-contrast exoplanet imaging: Molecule maps of the β Pictoris system with SINFONI. *A&A*, 617, sep 2018.
- Laurent Jolissaint, Laurent M. Mugnier, Chris Neyman, et al. Retrieving the telescope and instrument static wavefront aberration with a phase diversity procedure using on-sky adaptive optics corrected images. In *Adaptive Optics Systems III*, volume 8447, pages 844716–844716–12. SPIE, sep 2012.
- Christoph U. Keller, Visa Korkiakoski, Niek Doelman, et al. Extremely fast focal-plane wavefront sensing for extreme adaptive optics. *Adaptive Optics Systems III*, 8447:844721–844721–10, jul 2012.

- Visa Korkiakoski, Christoph U. Keller, Niek Doelman, et al. Experimental validation of optimization concepts for focal-plane image processing with adaptive optics. In *Adaptive Optics Systems III*, volume 8447, pages 84475Z–84475Z–11. SPIE, sep 2012.
- Visa Korkiakoski, Niek Doelman, Johanan Codona, et al. Calibrating a high-resolution wavefront corrector with a static focal-plane camera. *Applied Optics*, 52(31):7554–7563, oct 2013.
- Visa Korkiakoski, Christoph U. Keller, Niek Doelman, et al. Focal-plane wavefront sensing with high-order adaptive optics systems. *SPIE - Adaptive Optics Systems IV*, jul 2014.
- M. Kuzuhara, M. Tamura, T. Kudo, et al. Direct Imaging of a Cold Jovian Exoplanet in Orbit around the Sun-like Star GJ 504. *The Astrophysical Journal*, 774(1), jul 2013.
- Masen Lamb, David R. Andersen, Jean-Pierre Véran, et al. Non-common path aberration corrections for current and future AO systems. In *Adaptive Optics Systems IV*, volume 9148, page 914857. SPIE, jul 2014.
- Masen Lamb, Carlos Correia, Jean-Francois Sauvage, et al. Quantifying telescope phase discontinuities external to AO-systems by use of Phase Diversity and Focal Plane Sharpening. *Journal of Astronomical Telescopes, Instruments, and Systems*, 3(3):039001, jul 2017.
- Yann LeCun, Leon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998.
- Yann Lecun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553): 436–444, may 2015.
- Lisha Li, Kevin Jamieson, Giulia DeSalvo, et al. Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization. *Journal of Machine Learning Research*, 18:1–52, mar 2018.
- Scott Lundberg and Su-In Lee. A Unified Approach to Interpreting Model Predictions. In *NIPS’17: Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 4768–4777, may 2017.

- B. Macintosh, J. R. Graham, T. Barman, et al. Discovery and spectroscopy of the young jovian planet 51 Eri b with the Gemini Planet Imager. *Science*, 350(6256): 64–67, oct 2015.
- Serge Meimon, Thierry Fusco, and Laurent M Mugnier. LIFT: a focal-plane wavefront sensor for real-time low-order sensing on faint sources. *Optics Letters*, 35(18), 2010.
- Álvaro Menduiña-Fernández, Matthias Tecza, and Niranjana Thatte. HARMONI - first light spectroscopy for the ELT: novel techniques for the calibration of non-common path aberrations. In *Ground-based and Airborne Instrumentation for Astronomy VIII*, page 349. SPIE-Intl Soc Optical Eng, dec 2020.
- Tom M Mitchell. Machine learning and data mining. *Communications of the ACM*, 42(11):30–36, 1999.
- M. N’Diaye, K. Dohlen, T. Fusco, and B. Paul. Calibration of quasi-static aberrations in exoplanet direct-imaging instruments with a Zernike phase-mask sensor. *A&A*, 555, 2013.
- M. N’Diaye, A. Vigan, K. Dohlen, et al. Calibration of quasi-static aberrations in exoplanet direct-imaging instruments with a Zernike phase-mask sensor. II. Concept validation with ZELDA on VLT/SPHERE. *A&A*, 592(A79), jun 2016.
- M. N’Diaye, F. Martinache, N. Jovanovic, et al. Calibration of the island effect: Experimental validation of closed-loop focal plane wavefront control on Subaru/SCEXAO. *A&A*, 610, feb 2018.
- Robert J Noll. Zernike polynomials and atmospheric turbulence. *Journal of the Optical Society of America*, 66(3):207–211, 1976.
- Scott W. Paine and James R. Fienup. Machine learning for improved image-based wavefront sensing. *Optics Letters*, 43(6):1235, mar 2018.
- B. Paul, J. F. Sauvage, and L. M. Mugnier. Coronagraphic phase diversity: Performance study and laboratory demonstration. *A&A*, 552, 2013.
- Richard G Paxman, Timothy J Schulz, and James R Fienup. Joint estimation of object and aberrations by using phase diversity. *Journal of the Optical Society of America A*, 9(7), 1992.

- Richard G Paxman, Brian J Thelen, and John H Seldin. Phase-diversity correction of turbulence-induced space-variant blur. *Optics Letters*, 19(16):1231, 1994.
- W. J. Pearson, L. Wang, J. W. Trayford, et al. Identifying Galaxy Mergers in Observations and Simulations with Deep Learning. *A&A*, 626(A49), feb 2019.
- C. Plantet, S. Meimon, J.-M. Conan, and T. Fusco. Experimental validation of LIFT for estimation of low-order modes in low-flux wavefront sensing. *Optics Express*, 21(14):16337, jul 2013.
- A. Potier, P. Baudoz, R. Galicher, et al. Exoplanet direct imaging in ground-based conditions on THD2 bench. In *Proceedings of Adaptive Optics for Extremely Large Telescopes VI*, oct 2019.
- Axel Potier, Pierre Baudoz, Raphaël Galicher, et al. Comparing focal plane wavefront control techniques: Numerical simulations and laboratory experiments. *A&A*, 635(A192), mar 2020.
- Xin Qi, Guohao Ju, and Shuyan Xu. Efficient solution to the stagnation problem of the particle swarm optimization algorithm for phase diversity. *Applied Optics*, 57(11):2747, apr 2018.
- René Racine, Gordon A. H. Walker, Daniel Nadeau, et al. Speckle Noise and the Detection of Faint Companions. *Publications of the Astronomical Society of the Pacific*, 111(759):587–594, may 1999.
- David M. Reiman and Brett E. Göhre. Deblending galaxy superpositions with branched generative adversarial networks. *Monthly Notices of the Royal Astronomical Society*, 485(2):2617–2627, feb 2019.
- J-B Ruffio. Master Thesis: Application of Electric Field Conjugation on SPHERE and use of Phase Diversity on the AOF. *ISAE-Superaero*, 2014.
- J.-F Sauvage, L Mugnier, B Paul, and R Villecroze. Coronagraphic phase diversity: a simple focal plane sensor for high-contrast imaging. *Optics Letters*, 100(37):23, 2012.
- Jean-François Sauvage, Thierry Fusco, Gérard Rousset, and Cyril Petit. Calibration and precompensation of noncommon path aberrations for extreme adaptive optics. *Journal of the Optical Society of America A*, 24:2334–2346, 2007.

- Jean-François Sauvage, Thierry Fusco, Masen Lamb, et al. Tackling down the low wind effect on SPHERE instrument. In *Adaptive Optics Systems V*, volume 9909, page 990916. SPIE, jul 2016.
- Lloyd S Shapley. A value for n-person games. *Contributions to the Theory of Games*, 28(2):307–317, 1953.
- Garima Singh, Olivier Guyon, Pierre Baudoz, et al. Lyot-based Low Order Wavefront Sensor: Implementation on the Subaru Coronagraphic Extreme Adaptive Optics System and its Laboratory Performance. *Adaptive Optics Systems IV*, 9148:914848, jun 2014.
- Garima Singh, Raphaël Galicher, Pierre Baudoz, et al. Active minimization of non-common path aberrations in long-exposure imaging of exoplanetary systems. *A&A*, 631(A106), oct 2019.
- Anand Sivaramakrishnan, Rémi Soummer, Laurent Pueyo, et al. Sensing Phase Aberrations Behind Lyot Coronagraphs. *The Astrophysical Journal*, 688:701–708, 2008.
- Rémi Soummer, André Ferrari, Claude Aime, and Laurent Jolissaint. Speckle Noise and Dynamic Range in Coronagraphic Images. *The Astrophysical Journal*, 669:642–656, 2007.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, and Ruslan Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014.
- Hari Subedi, Neil T. Zimmerman, N. Jeremy Kasdin, et al. Coronagraph-integrated wavefront sensing with a sparse aperture mask. *Journal of Astronomical Telescopes, Instruments, and Systems*, 1(3):039001, jun 2015.
- Matthias Tecza, Niranjana Thatte, Fraser Clarke, et al. SWIFT: An adaptive optics assisted I/z band integral field spectrograph. *New Astronomy Reviews*, 49(10-12):647–654, jan 2006.
- Matthias Tecza, Niranjana Thatte, Fraser Clarke, et al. SWIFT de-magnifying image slicer: diffraction limited image slicing at optical wavelengths. In Eli Atad-Ettinger and Dietrich Lemke, editors, *Advanced Optical and Mechanical Technologies in Telescopes and Instrumentation*, volume 7018, page 70182O. SPIE, jul 2008.

- Eric J Topol. High-performance medicine: the convergence of human and artificial intelligence. *Nature Medicine*, 25(1):44–56, 2019.
- Robert K. Tyson. *Principles of Adaptive Optics*. CRC press, 2015.
- A. Vigan, M. N'Diaye, K. Dohlen, et al. Calibration of quasi-static aberrations in exoplanet direct-imaging instruments with a Zernike phase-mask sensor. III. On-sky validation in VLT/SPHERE. *A&A*, 629(A11), jul 2019.
- Mike Walmsley, Annette M. N. Ferguson, Robert G. Mann, and Chris J. Lintott. Identification of Low Surface Brightness Tidal Features in Galaxies Using Convolutional Neural Networks. *Monthly Notices of the Royal Astronomical Society*, 483(3), nov 2018.
- Mike Walmsley, Lewis Smith, Chris Lintott, et al. Galaxy Zoo: Probabilistic morphology through Bayesian CNNs and active learning. *Monthly Notices of the Royal Astronomical Society*, 491(2):1554–1574, jan 2020.
- Michael J. Wilby, Christoph U. Keller, Frans Snik, et al. The coronagraphic Modal Wavefront Sensor: a hybrid focal-plane sensor for the high-contrast imaging of circumstellar environments. *A&A*, 597(A112), oct 2017.
- P. G. Zhang, C. L. Yang, Z. H. Xu, et al. Hybrid particle swarm global optimization algorithm for phase diversity phase retrieval. *Optics Express*, 24(22):25704, oct 2016.
- Peiguang Zhang, Chengliang Yang, Zihao Xu, et al. High-accuracy wavefront sensing by phase diversity technique with bisymmetric defocus diversity phase. *Nature Scientific Reports*, 7(1), dec 2017.