



MONTEFIORE INSTITUTE OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE
SPACE SCIENCES, TECHNOLOGIES AND ASTROPHYSICS RESEARCH (STAR) INSTITUTE

Deep learning-based focal plane wavefront sensing for exoplanet imaging

A PhD thesis by
Maxime QUESNEL

Supervisors:

Prof. Gilles LOUPPE

Dr. Gilles ORBAN DE XIVRY

Successfully defended on the 16th of October 2024

Doctoral Committee

President

Prof. Pierre Geurts¹

Jury members

Prof. Marc Van Droogenbroeck¹

Dr. Olivier Absil²

Dr. Johan Mazoyer³

Dr. Markus Kasper⁴

Supervisors

Prof. Gilles Louppe¹

Dr. Gilles Orban de Xivry²

¹Montefiore Institute, University of Liège (Belgium)

²STAR Institute, University of Liège (Belgium)

³LESIA, Paris Observatory (France)

⁴European Southern Observatory (Germany)

Abstract

Exoplanet imaging presents significant challenges due to the extreme contrast and close angular separation between exoplanets and their host stars. Although coronagraphs and adaptive optics largely address these issues, high-contrast imaging instruments are still limited by residual wavefront errors, primarily from instrumental phase aberrations. Focal-plane wavefront sensing is appropriate to handle pupil-plane phase errors because it measures the aberrations directly at the science focal plane. The current sensing approaches can nonetheless be slow, unstable, and often designed for specific instruments.

In this thesis, new deep learning-based methods are explored to provide a fast, robust, and flexible solution for focal-plane wavefront sensing that can be applied to various instrumental designs. State-of-the-art deep convolutional neural networks are implemented and trained in a supervised way, achieving high performance and showing good robustness to changing aberration regimes and noise content. The problem of phase retrieval behind vortex phase masks is notably revisited with deep neural networks and simulated data. A new CNN-based approach that achieves a 100% science duty cycle using only focal-plane images is developed and tested with both scalar and vector vortex coronagraphs.

Supervised learning techniques require ground-truth phase aberrations for training the neural networks, that cannot be obtained accurately in real systems. To address this limitation, an autoencoder architecture is proposed, that is composed of a convolutional neural network as the encoder and a differentiable simulator as the decoder. As a result, the models are trained exclusively on observed images, enforcing the latent space to represent phase aberrations. Additionally, a variational component is introduced to predict probability distributions and assess phase uncertainties.

The deep learning-based methods developed in this thesis are successfully tested in-laboratory using the Subaru Coronagraphic Extreme Adaptive Optics (SCEXAO) instrument. Pre-training on large simulated datasets enhances the performance on small in-lab sets, and closed-loop experiments demonstrate robust convergence, effectively correcting introduced aberrations. The simulator-based autoencoder, tested across different wavefront aberration regimes, delivers very good PSF reconstructions and phase estimations. Validating the method with a simple optical propagation model shows great promise, and further developments of the simulator should improve its robustness against the varying conditions occurring in real systems.

This thesis demonstrates that methods based on deep artificial neural networks can offer an accurate, fast and robust solution for focal plane wavefront sensing, paving the way for future developments and on-sky applications.

Résumé

L'imagerie d'exoplanètes est complexe en raison des contrastes et séparations angulaires entre les exoplanètes et leurs étoiles hôtes. Bien que les coronographes et l'optique adaptative atténuent grandement ces problèmes, les instruments à haut contraste sont encore limités par des erreurs résiduelles de front d'onde dues aux aberrations de phase instrumentales. La détection du front d'onde au plan focal est efficace pour traiter les erreurs de phase du plan pupille, car elle mesure directement les aberrations au niveau des images scientifiques. Cependant, les méthodes actuelles peuvent être lentes, instables, et souvent adaptées à des instruments spécifiques.

Cette thèse explore de nouvelles méthodes d'apprentissage profond pour offrir une solution rapide, robuste et flexible à la détection du front d'onde au plan focal, adaptable à divers designs instrumentaux. Des réseaux neuronaux convolutifs, entraînés sur des données simulées, atteignent des performances élevées et montrent une bonne robustesse face aux aberrations et au bruit. Une nouvelle méthode basée sur les coronographes vortex est développée, permettant de résoudre l'ambiguïté de signe de la phase sans perte de temps d'observation.

Les techniques d'apprentissage supervisé nécessitent une valeur de référence pour les aberrations de phase afin d'entraîner les réseaux neuronaux, ce qui ne peut pas être obtenu de manière précise dans les systèmes réels. Une architecture d'auto-encodeur est alors proposée, composée d'un réseau neuronal convolutif suivi d'un simulateur différentiable. Les modèles sont par conséquent entraînés exclusivement sur les images observées, ce qui contraint l'espace latent à représenter les aberrations de phase. Une approche variationnelle est aussi introduite afin de prédire des distributions de probabilité et évaluer les incertitudes sur l'estimation de phase.

Les méthodes d'apprentissage profond développées dans cette thèse sont testées avec succès en laboratoire sur l'instrument Subaru/SCEXAO. L'avantage de pré-entraîner les modèles sur de nombreuses données simulées est d'augmenter considérablement les capacités de prédictions sur de petits jeux de données de laboratoire. Les expériences en boucle fermée ont aussi démontré une convergence robuste, corrigeant efficacement les aberrations introduites. La méthode d'autoencoder est testée à travers divers régimes d'aberration, offrant de bonnes capacités de reconstruction d'images et d'estimation de la phase. Développer davantage le simulateur devrait par ailleurs améliorer sa robustesse face aux conditions variables qu'un instrument peut connaître.

Cette thèse démontre que les méthodes basées sur les réseaux de neurones artificiels peuvent offrir une solution précise, rapide et robuste pour la détection de front d'onde au plan focal, ouvrant la voie à de futurs développements et à des applications sur le ciel.

Contents

Abstract	v
Publications	xxi
Acknowledgments	xxiii
Preface	1
I Introduction	3
1 High-contrast imaging	5
1.1 Science goals	6
1.1.1 Indirect detection methods	6
1.1.2 Direct imaging	8
1.2 Coronagraphy	12
1.2.1 Coronagraphic concepts	12
1.2.2 Vortex phase masks	14
1.3 Adaptive optics	18
1.3.1 Optical aberrations	18
1.3.2 Atmospheric turbulence	22
1.3.3 Adaptive optics	23
1.3.4 Non-common path aberrations	25
1.4 Focal-plane wavefront sensing	27
1.4.1 General problem statement	27
1.4.2 Phase sign ambiguity	28
1.4.3 Types of sensing techniques	29
1.5 Science data processing	33
1.6 Current limitations to high-contrast imaging	35
2 Deep learning for wavefront sensing	37
2.1 Deep neural networks	37
2.1.1 General concept	37
2.1.2 Convolutional neural networks	40
2.1.3 Motivations and challenges	43
2.2 Application to wavefront sensing	44

CONTENTS

2.2.1	Supervised phase retrieval	44
2.2.2	Physics-based retrieval	45
Scope and outline of the thesis		48
II Simulations		51
3	General supervised learning approach	53
3.1	Introduction	54
3.2	Synthetic data generation	54
3.2.1	General procedure	54
3.2.2	Implementation of the Vector Vortex coronagraph	56
3.2.3	Data pre-processing	57
3.3	Deep learning models	58
3.3.1	Convolutional Neural Networks	58
3.3.2	Mixture Density Layer	59
3.3.3	Training procedure	60
3.4	Experiments and results	61
3.4.1	Comparison between CNN architectures	62
3.4.2	Predicted uncertainties with mixture density models	62
3.4.3	Influence of the Vector Vortex coronagraph	64
3.5	Conclusions	66
4	Phase diversity with vortex masks	68
4.1	Vortex phase diversity concept	69
4.2	Deep learning approach	71
4.2.1	Data generation	71
4.2.2	Model architecture	72
4.2.3	Training procedure	73
4.3	Results and analysis	75
4.3.1	Phase sign determination	76
4.3.2	Comparison to classical phase diversity	76
4.3.3	Companion subtraction with dual-polarization	78
4.3.4	Theoretical limit	80
4.3.5	Model robustness	80
4.3.6	Closed-loop experiments	84
4.4	Conclusions	87

5	Physics-informed learning with autoencoders	89
5.1	Introduction	90
5.1.1	Differentiable optical systems	90
5.1.2	Potential physics-based architectures	91
5.2	General simulator-based autoencoder method	93
5.2.1	Data generation	93
5.2.2	Autoencoder architecture	94
5.2.3	Experiments and results	95
5.3	Unsupervised vortex phase diversity	101
5.3.1	Models and results	101
5.3.2	Rejection factor optimization	102
5.4	Variational inference approach	103
5.4.1	Definition	105
5.4.2	Training the SimVAE	107
5.4.3	Results and discussions	108
5.5	Conclusions	111
III	Application to real data	114
6	In-lab experiments on SCE_xAO	116
6.1	The Subaru/SCE _x AO instrument	117
6.2	SCE _x AO data management	117
6.2.1	In-lab data acquisition	119
6.2.2	Pre-processing	122
6.2.3	Simulations and in-lab data comparison	122
6.3	Supervised learning tests	124
6.3.1	Performance in three aberration regimes	124
6.3.2	Transfer learning from simulations	126
6.3.3	Closed-loop performance	129
6.4	Simulator-based autoencoder	129
6.4.1	Model performance	129
6.4.2	Two-step approach: learning pre-existing aberrations	135
6.4.3	Variational inference	136
6.4.4	Method refinements	137
6.5	Discussion	139

CONTENTS

IV	Conclusions	141
7	Conclusions	143
7.1	Summary	143
7.2	Perspectives	145
	Bibliography	148
V	Appendices	161
A	Pupil impact on the Vortex performance	163
B	Asymmetric pupil for the Vortex	165

List of Figures

1.1	Contrasts versus angular separations of exoplanets with direct imaging.	9
1.2	Examples of directly imaged exoplanets.	11
1.3	Ray diagram of a typical focal-plane coronagraph, based on the Lyot design.	13
1.4	Relationships between Lyot, vortex and apodized phase plate coronagraphs.	15
1.5	Structure of a vector vortex coronagraph with topological charge $\ell_p = 2$	16
1.6	Charge-6 scalar vortex phase mask.	18
1.7	Effect of the atmospheric turbulence on a point source image and AO correction.	20
1.8	Ordering of the Zernike polynomials according to the Noll convention.	21
1.9	General representation of an adaptive optics system.	24
1.10	PSF flux profile and Strehl ratio.	25
1.11	Images containing speckles obtained from a high-contrast imaging instrument.	26
1.12	Image formation and inverse problem to solve in focal-plane wavefront sensing.	28
1.13	Illustration of the sign ambiguity problem.	30
1.14	PSF estimation using an ADI sequence and median subtraction.	35
2.1	Neural network working principle, as a multilayered perceptron to be trained in a supervised fashion.	39
2.2	The LeNet-5 convolutional architecture applied to the task of handwritten digit recognition.	41
2.3	The ResNet structure.	42
3.1	Example of a generated Zernike coefficients set and WFE distribution.	55
3.2	Example of phase maps and pre-processed PSFs.	56
3.3	Transmitted flux after the vortex mask and the downstream Lyot stop as a function of the level of wavefront aberration.	57
3.4	The U-Net architecture, adapted for a regression task.	58
3.5	Illustration of training and validation losses.	61
3.6	Empirical distributions of RMS WFE on the test dataset.	63
3.7	Examples of predicted distributions with a mixture model.	64
3.8	Absolute difference of the predicted weights and entropy of the Gaussian mixture.	65
3.9	Residual RMS WFE and corresponding SNR behind a VVC.	67

4.1	Comparison of simulated PSFs for two conjugated phase maps.	70
4.2	Examples of PSFs at different SNR levels for $+\ell_p$	73
4.3	The EfficientNet principle and its performance.	74
4.4	RMSE per Zernike mode starting from the defocus mode.	77
4.5	Performance in terms of RMS WFE on the phase residuals at different SNR levels.	79
4.6	Phase prediction errors at different SNR levels with atmospheric turbulence residuals.	81
4.7	Performance with altered SNR levels during evaluation.	82
4.8	Performance with different input WFE levels defined during evaluation.	82
4.9	Robustness on higher-order aberrations.	83
4.10	Iterative approach for a few test samples with post-VVC PSFs.	85
4.11	Evolution of post-VVC PSFs through iterations.	86
4.12	Contrast curves on a PSF sample at different close-loop state.	87
5.1	Two proposed approaches for phase retrieval with physics-informed learning.	91
5.2	The proposed simulator-based autoencoder architecture (SimAE).	94
5.3	Examples of training and validation losses over the epochs for each approach.	96
5.4	Performance per Zernike mode, for low and high aberration regimes.	96
5.5	Performance at different SNR levels for the SimAE and the standard CNN models.	97
5.6	Performance with changing the SNR and the input WFE at test time.	99
5.7	Performance of models fine-tuned on each test sample individually.	100
5.8	Residual errors as a function of the input WFE.	100
5.9	WFE distributions and robustness assessment for the SimAE vortex phase diversity.	102
5.10	Simulated PSF with different rejection factors and retrieval of the factor during training.	104
5.11	Diagram of the simulator-based variational autoencoder architecture (SimVAE).	106
5.12	Generated Zernike distribution, following a centered Gaussian with a decreasing standard deviation over the modes.	108
5.13	Metrics over the training epochs with the SimVAE.	108
5.14	Predicted phase maps and reconstructed PSFs with the SimVAE method.	109
5.15	Residual RMS WFE for each test sample with the SimVAE.	110
6.1	The Subaru telescope and its SCEXAO instrument.	118
6.2	SCEXAO pupil stop and DM flat map.	119

LIST OF FIGURES

6.3	WFE distributions for each of the three early datasets, in the low, medium and high aberration regimes.	120
6.4	Raw in-lab SCExAO samples taken from the early datasets, at low, medium, and high aberration regimes.	121
6.5	Examples of in-focus PSFs taken from each aberration dataset.	123
6.6	One-shot phase prediction errors on the test samples with SCExAO in-lab datasets in three aberration regimes	125
6.7	Phase residual errors for models trained only with in-lab data, with simulated data, and fine-tuning the simulation-based model with small amount of lab data.	126
6.8	Comparison of model performance lab data-only models, simulation-based models, and fine-tuned models with small amount of lab data.	127
6.9	Closed-loop PSF corrections on the SCExAO bench over ten iterations and for ten different input WFEs.	130
6.10	Residuals errors with respect to the input WFE for each test sample.	131
6.11	Example of predicted phase and reconstructed PSFs with the SimAE on SCExAO lab data.	132
6.12	Performance on the SimAE on SCExAO lab data with a median input WFE of 260 nm RMS.	133
6.13	Median phase maps computed from the test prediction set	134
6.14	Learned NCPAs during the first training procedure.	135
A.1	Contrast curves for central obstructions of 15% and 0%.	163
A.2	Evolution of post-VVC PSFs through iterations.	164
B.1	Aperture asymmetries of different sizes in the Lyot plane.	165
B.2	RMS WFE for each Zernike mode with classical phase diversity and with pupil asymmetry.	166

Acronyms

AD	Automatic Differentiation
ADI	Angular Differential Imaging
AE	Autoencoder
AGPM	Annular Groove Phase Mask
AI	Artificial intelligence
AO	Adaptive Optics
AU	Astronomical Unit
CNN	Convolutional Neural Network
DL	Deep Learning
DM	Deformable Mirror
DNN	Deep Neural Network
EFC	Electric Field Conjugation
ELT	Extremely Large Telescope
FC	Fully Connected
FoV	Field of View
FPWFS	Focal Plane Wavefront Sensing
FQPM	Four Quadrant Phase Mask
GPU	Graphics processing unit
HCI	High Contrast Imaging
IR	Infrared
IWA	Inner Working Angle
JWST	James Webb Space Telescope
MDN	Mixture density network
ML	Machine Learning
MLE	Maximum Likelihood Estimation
MLP	Multilayer Perceptron
MSE	Mean Squared Error
NCPA	Non-Common Path Aberration
NN	Neural Network
PCA	Principal Component Analysis

LIST OF FIGURES

PDF Probability Density Function
PSD Power Spectral Density
PSF Point Spread Function
PyWFS Pyramid Wavefront Sensor
RMS Root Mean Square
RMSE Root Mean Square Error
SNR Signal-to-Noise Ratio
SR Strehl Ratio
SVC Scalar Vortex Coronagraph
VAE Variational Autoencoder
VC Vortex Coronagraph
VGP Vanishing Gradient Problem
VLT Very Large Telescope
VVC Vector Vortex Coronagraph
WFE Wavefront Error
WFS Wavefront Sensor
XAO Extreme Adaptive Optics

Publications

During this thesis, I have published three articles as first author: one journal paper in *Astronomy and Astrophysics* (A&A, peer-reviewed), and two SPIE conference proceedings. I was also co-author of one journal paper published in *Monthly Notices of the Royal Astronomical Society* (MNRAS, peer-reviewed). Those contributions are listed below:

- Journal articles:
 1. [Quesnel et al. \(2022b\)](#): A deep learning approach for focal-plane wavefront sensing using vortex phase diversity
 2. [Orban de Xivry et al. \(2021\)](#): Focal plane wavefront sensing using machine learning: performance of convolutional neural networks compared to fundamental limits
- Conference proceedings:
 1. [Quesnel et al. \(2022a\)](#): A simulator-based autoencoder for focal plane wavefront sensing
 2. [Quesnel et al. \(2020\)](#): Deep learning-based focal plane wavefront sensing for classical and coronagraphic imaging

The content of these articles is included in the thesis, along with further studies and results. On top of the published papers, I have also given eight talks and presented two posters during international conferences and workshops. I have also delivered multiple seminar talks at the University of Liège and at research institutes abroad.

Acknowledgments

I would first like to thank my supervisors Gilles Louppe and Gilles Orban de Xivry. I know that PhD supervisors play such an important role in the success (or failure) of a PhD, and I feel very privileged to have had such kind, patient, intelligent and professional supervisors. You were always here for me, helping me to make my PhD a great research and human experience. I am also very grateful to Olivier Absil, who has provided invaluable support and guidance throughout my PhD journey. Thank you to the president of the jury, Pierre Geurts, and to the jury members, Marc Van Droogenbroeck, Johan Mazoyer, and Markus Kasper, for taking the time to read my thesis and providing valuable feedback. I also extend my gratitude to the University of Liège for welcoming me and providing all I needed to carry out my PhD.

During those five years, I have had the chance to encounter many incredible people. Within PSILab and across the astrophysics department, it would be impossible to name everyone, but I would like to especially acknowledge some of those closest to me, with whom I shared countless laughs and unforgettable moments: Mariam, Lorenzo, Jahanvi, Carles, Malavika, Rakesh, Valentin, Jyotirmay, Elena, Nikolay, and many more. I have also had wonderful times with Niyati and Tomeu during conferences in Munich, Montreal, Avignon and Tokyo. Thanks to Anaïs for hosting Carles and me in your office when we first arrived at the Montefiore institute. It was always a pleasure interacting with you and other Montefiore mates. Matthia, Omer, Sébastien, Antoine, Joeri, Arnaud, Pascal, Romain are among the great people I had the chance to encounter. I will always remember the fun times playing Rikiki during lunch breaks. Special recognition goes to the sloths in front of Montefiore, who have been hanging there for years, giving inspiration for resilience but also, occasionally, for laziness. And how can I not mention our dear 48 and 58 buses? Rarely on time, sometimes nowhere to be found, but absolutely essential to go to work. Thanks to Alan, our loyal GPU cluster, which greatly contributed to me obtaining my results on time. I am also grateful to the Subaru/SCEXAO team for allowing me to test my methods on their instrument, and for providing help during my research works. I want to express very special thoughts for Matt, whom we lost under horrible circumstances. Matt worked with us in PSILab as a postdoc researcher, and I deeply appreciated his sharp sense of humor and enthusiasm. This tragedy is a reminder that life is precious, fragile and also quite absurd. We really should enjoy this special time we have on Earth, and never sacrifice moments we can spend with people we love.

I am deeply grateful to everyone who supported me throughout this PhD journey, both in the office and in my personal life. The encouragement, kindness and love you provided guided me through every phase of this experience. Thank you.

Preface

Are we alone in the Universe? Are there any planet outside our Solar System that harbors life? These are among the most exciting questions we are asking ourselves today. Searching for exoplanets has been carried out for a few decades now, and thanks to impressive technological development, a few thousand of them have already been detected. While no signs of life has yet been found, this exciting journey is just starting.

The large majority of exoplanetary detections are obtained using indirect methods, observing disturbances on the parent star caused by the planets. Getting direct images of exoplanets remains of great interest because of the possibility of studying their atmosphere. It is however a real challenge, because planets are extremely faint compared to their star, which means the star glow overwhelms that of the planetary companion. On top of it, from the Earth point of view, exoplanets appear as orbiting very close to their star due to the interstellar distances. Exoplanet imaging can be compared to trying to spot a firefly next to a lighthouse, more than one thousand kilometers away. In addition to this contrast problem, the Earth's atmosphere also disturbs the light coming from these exoplanets and their host stars. This phenomenon, which makes stars twinkle in the night sky, hinders ground-based observations. Fortunately, highly capable telescopes and instruments have been built to tackle these challenges, with impressive results and discoveries. To this day, more than sixty exoplanets have been directly observed from the ground or space.

Because of the limits set by the constraints described above, imaged exoplanets are very massive, gaseous, located far from their star, and much larger and younger than our Jupiter, for instance. Although such celestial bodies have great scientific interest, Earth-like exoplanets remain the most interesting targets. To aim for such planets with direct imaging, it is necessary to improve the instruments performance and the image processing techniques used to extract planetary signals. Telescopes are made of imperfect optical components, causing science images to contain not only potential planets but also noise, which can share similar characteristics to that of exoplanets. Challenges are especially important in the vicinity of the star, where the star glow is stronger and the noise due to optical errors prominent. It is in these areas that faint planets orbiting closer to their star can be found. Some optical defects that decrease telescope capabilities are not easily corrected, and they represent one of the main obstacles to imaging small rocky planets. Addressing the problem of optical errors with the help of artificial intelligence (AI) is the focus of this thesis. AI has grown tremendously these last years, showing incredible capabilities. Existing deep neural networks in particular are very relevant to our image-based problem, and this thesis proposes to implement such architectures to mitigate optical aberrations and help exoplanet imaging reach new heights.

PART I

Introduction

High-contrast imaging

To give proper context to this thesis, the case for exoplanet imaging is first made, before presenting the instrumentation that is essential to detect exoplanets: coronagraphy and adaptive optics. Wavefront aberrations and the methods to correct them are then described, as well as data processing techniques used to extract exoplanet after the observations. To conclude, the limitations of high-contrast imaging and the importance of focal-plane wavefront sensing to correct non-common path aberrations are highlighted.

Contents

1.1	Science goals	6
1.1.1	Indirect detection methods	6
1.1.2	Direct imaging	8
1.2	Coronagraphy	12
1.2.1	Coronagraphic concepts	12
1.2.2	Vortex phase masks	14
1.3	Adaptive optics	18
1.3.1	Optical aberrations	18
1.3.2	Atmospheric turbulence	22
1.3.3	Adaptive optics	23
1.3.4	Non-common path aberrations	25
1.4	Focal-plane wavefront sensing	27
1.4.1	General problem statement	27
1.4.2	Phase sign ambiguity	28
1.4.3	Types of sensing techniques	29
1.5	Science data processing	33
1.6	Current limitations to high-contrast imaging	35

1.1 Science goals

By targeting very faint exoplanets close to extremely bright stars, high-contrast imaging is full of challenges, but offers considerable scientific benefits. Before presenting the case for exoplanet imaging, indirect detection methods are first described.

1.1.1 Indirect detection methods

In the last few decades, detecting and characterizing exoplanets has been achieved following complementary approaches, as shown in Table 1.1. Most of them are not based on observing the exoplanet itself, but are rather looking at the planet’s effect on its parent star. Among the indirect approaches that exist today, two in particular have been very successful in detecting a large number of planetary companions: radial velocity and transit photometry. The microlensing and astrometry methods are also briefly presented in this section.

Radial velocity

Because the center of mass of a planetary system is always different from the geometrical center of the parent star, during a planet’s orbit the star also experiences orbital motion. This movement towards and away from the observer creates a tiny wavelength shift of the star’s emission spectrum, explained by the Doppler effect. The variation of the star’s velocity can be measured, and depending on its strength, a lower bound for the orbiting planet’s mass can be deduced (Table 1.1). The radial velocity method has led to the first exoplanet detection around a Sun-like star: a Jupiter-like planet named 51 Pegasi b (Mayor & Queloz, 1995). Since then, more than a thousand exoplanets have been detected with this approach (Table 1.1). Spectrographs with high spectral resolution and dispersion are required for radial velocity: to detect a hot Jupiter-like planet for instance, a resolution $R = \frac{\lambda}{\Delta\lambda} \sim 100\,000$ is needed, with λ the observing wavelength. Among the most advanced instruments, HARPS at La Silla Observatory (Mayor et al., 2003) has been particularly successful (Bonfils et al., 2013) and characterized systems such as Gliese 581, which contains two terrestrial planets in the so-called “habitable zone” (Selsis et al., 2007).

Transit photometry

When an exoplanet passes in front of its parent star, it occults part of the starlight. In some conditions, this eclipse can be measured by transit photometry (Charbonneau et al., 2000). This happens only if the orbital plane is edge-on with respect to the observer’s

Discovery method	Number of detections	Observational biases	Characterization			
			Orbit	Mass	Size	Spectra
Transit	4146	close-in edge-on	a, ϵ, i, T		ratio	
Radial velocity	1071	massive, close-in edge-on	a, ϵ, ω, T	min		
Microlensing	204	low-mass mid-separation	a	ratio		
Direct imaging	69	young, massive far-out	a, ϵ, i, T			
Astrometry	3	massive, far-out pole-on	$a, \epsilon, i, \Omega, T$	ratio		

Table 1.1 Main exoplanet finding methods (11 exist in total), their number of detections as of December 12, 2023 (source: [NASA Exoplanet Archive](#)), observational biases and parameters that can be derived with characterization. “Close-in”, “mid-separation” and “far-out” refer to the distance of the planet from its parent star, while “edge-on” and “pole-on” indicate the preferred inclination of the orbit with respect to the observer. The six orbital parameters are the semi-major axis a , the eccentricity ϵ , the inclination i , the longitude of the ascending node Ω , the argument of the periapsis ω , and the time of periapsis passage T . Some of these parameters may need additional data to be recovered (e.g., the inclination i with direct imaging). Other planetary characterization parameters are indicated with a color code: *green* for direct estimations (often depending on stellar information) or using “reliable” models, *orange* when the parameter cannot be directly retrieved or when several models compete, and *red* when the parameter cannot be estimated. For instance, it is never possible to directly find the exoplanet’s mass, but rather the mass ratio between the planet and its star, or a minimal value ($M \sin(i)$ with radial velocity). The only methods that allow for atmospheric spectra measurements are direct imaging and transits.

vantage point, which therefore bias transit detections to planets with short orbital periods. Thousands of exoplanets have been detected transiting stars (Table 1.1). Confirming detections is not always easy and false positives are frequent (Fressin et al., 2013), e.g., due to another object passing in front or behind the star. Several space telescopes have focused on transits, such as the Kepler mission (Koch et al., 2010; Batalha et al., 2013). Ground observations have also been successful, e.g., with the TRAPPIST program (Jehin et al.,

2011). One of the discovered planetary system is particularly exceptional: seven rocky exoplanets around TRAPPIST-1 (Gillon et al., 2017). Beyond detection, the size of the planets can also be inferred from how much the star decreases in brightness, and some orbital parameters can be derived from the transit periodicity (Batalha et al., 2013). More importantly, since a portion of the starlight passes through the planet’s atmosphere before reaching the observer, the atmospheric composition can be analyzed (Charbonneau et al., 2002).

Gravitational microlensing

When a star passes in the line of sight to another, much more distant, star, its gravitational field magnifies the light of the background object. If a planet is also present around the lensing (foreground) star, this brightening effect is amplified. These events are short (in the order of days or weeks) and occur only once, making detections particularly difficult to confirm or further analyze. Unlike with other methods, low-mass planets in medium orbits (1 to 10 astronomical units (AU)) are the most likely to be detected by microlensing (Tsapras, 2018).

Astrometry

As explained above with radial velocity, a star experience motion due to the gravitational pull of its planets. Instead of analyzing the star’s spectra, the astrometry method focuses on finding the spatial displacements of the star in the sky. Planets with large orbits are naturally more likely to be detected, which is complementary to the radial velocity approach. Astrometry is extremely challenging and has provided only three detections so far. Datasets from the Gaia mission, which contain astrometry measurements of nearly two billion stars, are expected to yield tens of thousands of discoveries in the next few years (Perryman et al., 2014).

1.1.2 Direct imaging

Obtaining a direct image of a planet is extremely challenging, essentially because two large gaps exist between exoplanets and their parent star: high contrasts and small angular separations.

The contrast is the flux ratio between the planet and the star: $C = \frac{\mathcal{F}_{\text{planet}}}{\mathcal{F}_{\text{star}}}$. Contrast values can be found from 10^{-3} to 10^{-10} , corresponding to very massive planets (thermal emission observed in the infrared) and Earth-like planets (reflected light from the star in the visible), respectively. On the other hand, the angular separation is the apparent distance between

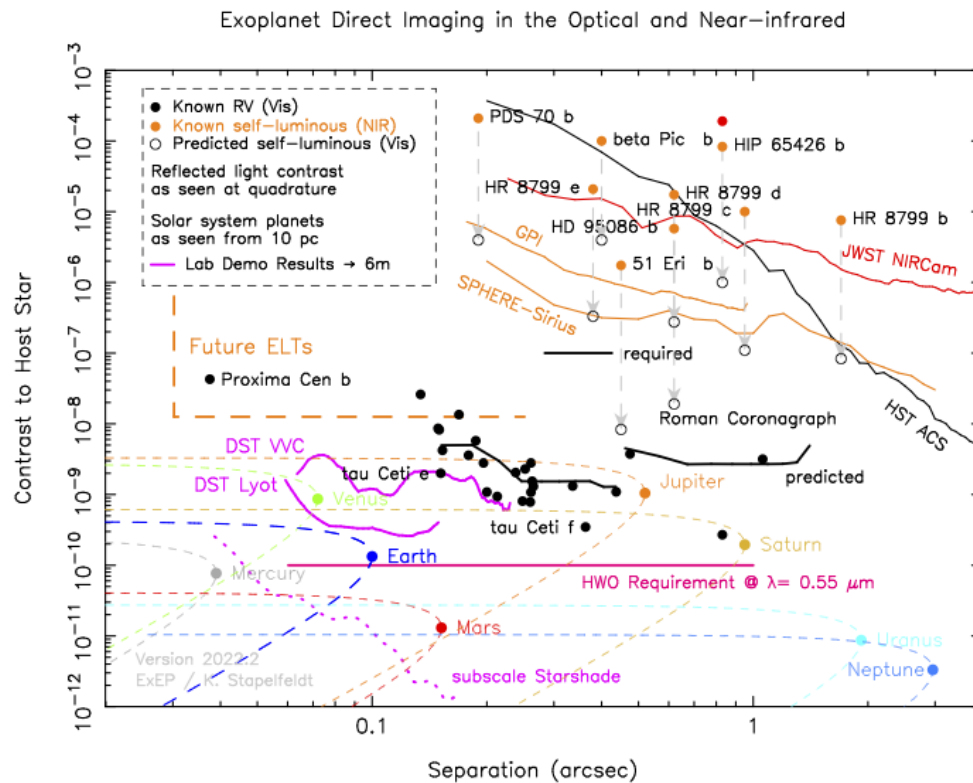



Figure 1.1 Contrasts versus angular separations of exoplanets with direct imaging. Already detected exoplanets in the near-infrared (orange dots) can be compared with their predicted contrasts in the visible (gray circles). Exoplanets detected with the radial velocity method (black dots) and Solar system planets (multicolor dots; dashed line orbits) are also indicated for reference. Current contrasts reached by space-based instruments inside the Hubble Space Telescope (HST; black curve) and the James Webb Space Telescope (JWST; red curve) are shown. Ground-based instruments (orange curves) perform better so far, with the best results achieved with the VLT/SPHERE on Sirius. Future instruments will push the contrast limits: ground-based extremely large telescopes (ELTs; not limited to the European ELT) will reach 10^{-8} contrasts and angular separations down to $0.03''$ (dashed orange line), the Roman Space Telescope Coronagraph instrument is predicted to dig contrasts up to 2×10^{-9} (solid black line), while the Habitable Worlds Observatory mission concept should allow detecting Earth-like exoplanets at 10^{-10} contrasts (pink line). In-lab results are also shown: the Decadal Survey Testbed (DST) helped to obtain impressive performance with vector vortex (VVC) and Lyot coronagraphs (magenta lines). Credits: [NASA/JPL/Caltech](#) (version: November 2023).

the planet and the star, and is often expressed in arc seconds (arcsec), which is equal to 1/3600th of a degree. Typical encountered separations range from 0.1 to 1 arcsec.


The contrast and separations encountered in exoplanet imaging are illustrated in Figure 1.1. Contrasts between our Solar System’s planets and the Sun are very large, and these planets would not be detected with current technology. Detections so far have rather concerned young, hot and massive exoplanets. In the infrared (IR) wavelengths, these planets have enough thermal emission compared to their star to give contrasts in the order of 10^{-3} to 10^{-6} . This is why IR passbands are used in HCI, especially the bands that have a good transmittance through the atmosphere (J, H, K and L). While reflected light from the planet can also be targeted, contrasts are much stronger and only future projects will be able to reach such values (Figure 1.1). All detected exoplanets are also orbiting far away from their parent star, because it is then possible to spatially distinguish them. As shown in Figure 1.1, future telescopes will be able to reach much smaller angular separations. Because of these important challenges, direct imaging has only yielded about 1% of all the exoplanet detections so far (Table 1.1).

One of the main advantages of the direct imaging technique is that it allows to study the exoplanet’s atmosphere by retrieving its spectra. Some molecules in planetary atmospheres can be indications of biological presence or activities, such as water H_2O , methane CH_4 , molecular oxygen O_2 , ozone O_3 and nitrous oxide N_2O (Seager & Deming, 2010; Schwieterman et al., 2018). Unlike the transit photometry method, which also allows the retrieval of planetary spectra, direct imaging can probe planets in wider orbits. With the next generation of extremely large telescopes soon to operate, combined with future space missions, detecting rocky planets and probing their atmospheres for signs of biological activity will become possible.

Thus far only giant exoplanets have been discovered with direct imaging. The first imaged exoplanet is a “super-Jupiter” orbiting the brown dwarf 2M1207 (Chauvin et al., 2004, 2005). This detection, shown in Figure 1.2a, was achieved with the NACO instrument at the Very Large Telescope (VLT) observing in different wavelength bands. The HR 8799 system is one of the most important discoveries made with direct imaging: four young super-Jupiters were directly imaged around a bright and young star (Marois et al., 2008, 2010), as it can be seen in Figure 1.2c. The orbital motion of these exoplanets was also observed over several years, which allowed to derive precise orbital parameters (Soummer et al., 2011; Wang et al., 2018). Another famous detection is that of β Pictoris b (Lagrange et al., 2009), a giant planet whose spectra was later retrieved with a molecular mapping technique, showing the presence of H_2O and CO in its atmosphere (Hoeijmakers et al., 2018). Exoplanet surveys allow probing hundreds of potential planetary systems and get a census of young giant exoplanets. Using the VLT/SPHERE instrument, the cur-



(a) 2M1207 b from [Chauvin et al. \(2004\)](#). (b) HIP 65426 b from [Chauvin et al. \(2017\)](#).



(c) HR 8799 system from [Marois et al. \(2010\)](#).

Figure 1.2 Examples of directly imaged exoplanets. (a): 2M1207 b is clearly visible positioned to the lower left of the central star. (b): HIP 65426 b appears in the lower left corner of the image. (c): The four companions of HR 8799 are indicated with their respective letters.

rent SHINE program has already yielded multiple detections that way, including two new discoveries (Vigan et al., 2021). One of them is HIP 65426 b (Figure 1.2b), notably re-imaged using the James Webb Telescope (Carter et al., 2023). Despite its young age, no debris disk has been detected around HIP 65426 b, challenging current planetary formation models (Chauvin et al., 2017). Contrasts and angular separation of some of those directly imaged exoplanets are indicated in Figure 1.1.

Other privileged targets with the direct imaging method are circumstellar disks, in which planetary formation occurs. Studying these disks allows understanding better the origin of exoplanets and how they form. The disk of PDS 70 and its recently discovered exoplanet are a prominent case (Keppler et al., 2018).

1.2 Coronagraphy

As explained in the previous section, contrasts between planets and stars are significant, and this represents a major obstacle for detecting planetary companions. The coronagraph is a fundamental optical system for HCI, designed to block out the starlight while keeping any off-axis exoplanet light mostly unaltered. Originally developed to observe the corona of the Sun, coronagraphs have since been implemented to detect exoplanets and circumstellar disks around other stars. Different varieties of coronagraphs now exist, and Section 1.2.1 offers a non-exhaustive presentation. Vortex coronagraphs in particular are of interest for this thesis, and are thus described in detail in Section 1.2.2.

1.2.1 Coronagraphic concepts

Various families of coronagraphs exist (Guyon, 2011; Belikov et al., 2023) and this section focuses on a few relevant focal and pupil plane coronagraphs.

Focal-plane coronagraphs

The first type of coronagraph was invented by Bernard Lyot to study the Sun’s corona (Lyot, 1939), and it later appeared that Lyot coronagraphs can be very beneficial for directly imaging exoplanets as well (Vilas & Smith, 1987; Watson et al., 1991). The original Lyot coronagraph is composed of two main elements:

- A focal plane mask: It blocks the on-axis light from the star.
- A Lyot stop: Downstream of the mask and at the exit pupil plane, it is a diaphragm that blocks the light of the Airy rings beyond the focal plane mask.




Figure 1.3 Ray diagram of a typical focal-plane coronagraph, based on the Lyot design. Most of the on-axis light (star; dark orange) is cancelled by a focal-plane mask, while some residual light due to optical imperfections and wavefront aberrations (pale orange) still reach the detector. On the other hand, any off-axis light (exoplanet; blue) is unaffected by the mask.

A ray diagram of such a focal-plane coronagraph encountering starlight and exoplanet light is represented in Figure 1.3. This simple design can be further improved by placing an apodizer in the entrance pupil plane to reduce the diffraction effects of the telescope aperture (Soummer, 2005), which are also due to the central obscuration and the secondary mirror supports (‘spiders’). An apodized pupil Lyot coronagraph is for instance used in the SPHERE instrument (Carillet et al., 2011; Guerri et al., 2011).

Other types of focal-plane coronagraphs include phase mask coronagraphs. The most prominent designs are the four-quadrant and optical vortex masks, which are presented in Section 1.2.2.

Pupil-plane coronagraphs

There are also notable pupil-plane coronagraphs, such as the Apodizing Phase Plate (APP) coronagraph. It is a transparent mask that suppresses the starlight by creating a “dark hole” covering 180 degrees in the science camera (Codona et al., 2006), i.e., the starlight is suppressed to a significant degree across the image. The APP was first implemented on-sky using the MMT Observatory (Kenworthy et al., 2007). The APP has several advantages, such as requiring only one optical element and being insensitive to tip-tilt aberrations. It faces some drawbacks, however, since it cannot cover the whole field of view (FoV) at once, and it does not perform well in broadband (multiple wavelengths at once). The

vector-phase design of the APP (vAPP), is a halfwave plate with variable fast-axis orientation (Snik et al., 2012). The coronagraph therefore applies opposite phase ramps to each circular polarization. The polarization can then be split by using a quarter-waveplate and a polarizing beam-splitter. The grating-vAPP implementation, instead, applies a tip-tilt phase on top of the apodizing phase to separate the polarization states (Doelman et al., 2021). Both methods then produce two PSFs with opposite dark hole regions, which allows to cover the entire FoV, at the expense of losing half the light in each region. As illustrated in Figure 1.4, the vAPP and grating-vAPP are also based on the geometric phase.

Several studies are under development to improve APP-based coronagraphic designs, in particular using optimization methods. We can for instance mention the attempt to pair together the apodized-pupil Lyot and vAPP coronagraphs (Por, 2020). By deriving an optimization problem, the best designs are found for specific instruments and pupil aperture configurations. In parallel, to leverage the vortex coronagraphs and vAPP performance, Doelman et al. (2020) propose to incorporate polarization grating elements. This prevents polarization leakage while keeping great broadband capabilities. In this study, vAPP designs are determined with the global optimizer from Por (2017). Besides works on the APP, non-linear optimization techniques have also helped improve shaped pupil Lyot coronagraph designs (Zimmerman et al., 2016), for instance. The optimization solvers used in these studies perform generally well, but as discussed in Section 5.1.1, automatic differentiation can be a much more powerful and efficient tool to optimize coronagraphs and optical parameters in general.

1.2.2 Vortex phase masks

Vortex coronagraphs (VC) are particularly relevant to this thesis due to their performance and polarization properties, which can benefit focal-plane wavefront sensing (see Chapter 4). VCs are first presented in details in this section.

Origin

As illustrated in Figure 1.4, the origin of the VC can be traced back to the four-quadrant phase mask (FQPM; Rouan et al., 2000), which is itself based on the principle of the phase mask (Rodier & Rodier, 1997). The FQPM works with destructive interference on the incoming light, by applying a π phase shift on two of the four quadrants. Although this concept is supposed to be achromatic, i.e., the performance should be the same for every observing wavelength, the standard implementation with physical steps shows chromatic behaviour. Using achromatic half-wave plates, which are based on the geometric phase,

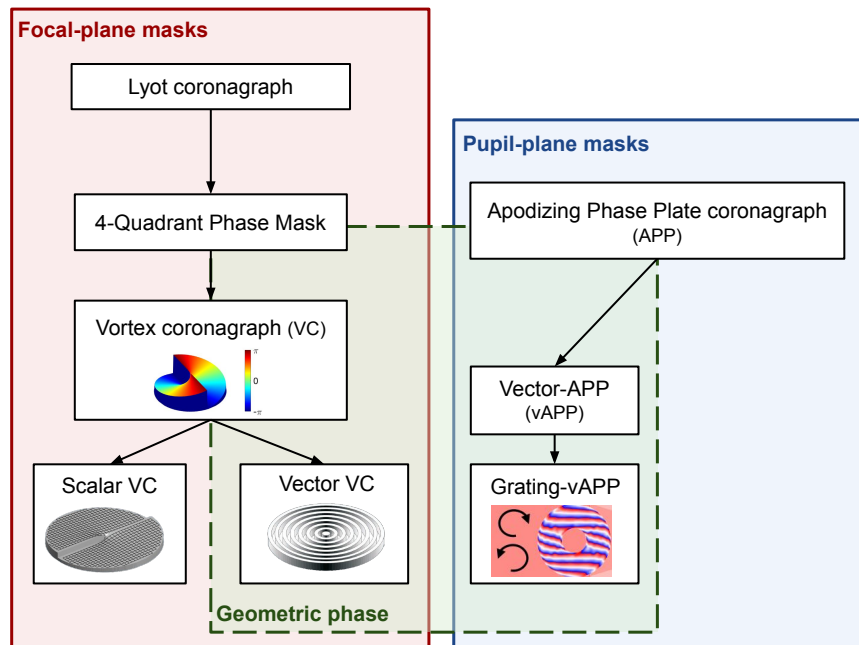


Figure 1.4 Relationships between Lyot, vortex and apodized phase plate coronagraphs, as focal or pupil-plane masks, and whether they have the geometric phase property.

have been notably proposed (Mawet et al., 2006). One main issue for the FQPM, however, remains: the transitions that exist between the four quadrants can cause exoplanet light subtraction. It is by focusing on solving this issue that the annular groove phase mask (AGPM) was developed by Mawet et al. (2005), which is composed of subwavelength gratings to produce an optical vortex on the incoming light. An AGPM can be seen as a FQPM with a very large number of segments, leading to a continuous helical phase ramp (König, 2023). In parallel, and coincidentally, Foo et al. (2005) also proposed a concept of vortex coronagraph based on a helical piece of glass.

Principle of the Vector VC

Generally speaking, the vortex coronagraph uses a transparent phase mask at the focal plane that diffracts the on-axis light outside the pupil (Mawet et al., 2005). A downstream pupil-plane Lyot stop then blocks this light. As a consequence, the starlight is suppressed, and only off-axis light, such as of an exoplanet, appears in the detector plane. The vortex coronagraph thus also follows Figure 1.3, except that the planetary signal also passes through the mask, and the on-axis light from the star that has been diffracted is blocked by the Lyot stop. The companion's light is also slightly reduced by the mask the smaller




Figure 1.5 Structure of a vector vortex coronagraph with topological charge $\ell_p = 2$ in the L band (3–4 μm). *Left*: Helical phase ramp applied to the incoming wavefront (from [Mawet et al., 2005](#)). *Right*: Image of the center obtained with a scanning electron microscope (from [Delacroix et al., 2013](#)).

its off-axis angle. In comparison to the Lyot design, the vortex coronagraph has a smaller inner working angle (IWA) but is more sensitive to low-order phase aberrations.

The vector flavor of the vortex coronagraph (VVC) was specifically designed to reach the best achromatic performance possible. The VVC generates two conjugated phase ramps with a complex transmission:

$$\tau = \exp(i \pm \ell_p \theta), \quad (1.1)$$

with ℓ_p the topological charge and the θ the azimuthal coordinate. These conjugated phase ramps are applied to both circular polarization of the incoming light.

The choice of the topological charge of the VVC is important for several reasons. First, only even charges can in theory perfectly cancel on-axis light ([Mawet et al., 2005](#)). Also, lower charges produce smaller IWAs and are more sensitive to low-order aberrations than higher charges. Moreover, as shown in Figure 1.5, charge-2 AGPMs have a circular grating design and are therefore easier to manufacture. This is why vortex coronagraphs with $\ell_p = 2$ have been essentially used and are thus exclusively considered in this thesis.

For a perfect wavefront and a circular pupil, and neglecting the leakage term due to imperfect manufacturing and chromatic effects, the VVC totally cancels the on-axis light of the star. As explained later in Section 1.3, the wavefront is often distorted by atmospheric turbulence and optical imperfections. The resulting phase aberrations decrease the coronagraph's efficiency, allowing some starlight to leak through. The VVC is particularly sensitive to low-order aberrations such as tip-tilt, which causes the starlight not to be centered anymore. Precise low-order wavefront control is therefore important for the VVC to work properly ([Huby et al., 2015](#)).

Overall the VVC has multiple advantages, such as a small IWA (e.g., $0.9 \lambda/D$ for $\ell_p = 2$), quasi-achromaticity across a wide bandwidth, high throughput, high contrast sensitivity and a 360° discovery space (Mawet et al., 2010).

The scalar VC

An optical vortex can be vectorial or scalar in nature, depending on how it affects the electromagnetic field. Early studies (Swartzlander, 2006; Swartzlander et al., 2008) paved the way for a scalar flavor of vortex coronagraphs (SVC). But it is with the development of space telescopes for HCI that the interest on them really emerged (Ruane et al., 2019).

SVCs differ from VVCs in the effect they have on the circular polarization of the incoming light. With the SVC, the same phase ramp $\exp(i \ell_p \theta)$ is applied to both circular polarizations. This is important when wavefront control needs to be very accurate: the differential aberrations in each polarization are not averaged and can be retrieved and corrected with wavefront control. Unlike VVCs, the scalar phase masks typically suffer from chromaticity (Desai et al., 2022) due to the way they are manufactured. Recent works have investigated how to make SVCs achromatic. Optimizing optical components or stacking multiple materials can improve the bandwidth (Ruane et al., 2019), while phase wrapping can be another approach to achieve this goal (Galicher et al., 2020).

Manufacturing

VVCs can be made of various materials, but it is the developments in diamond micro-optics (Karlsson & Nikolajeff, 2003) that offered a unique opportunity to produce VVCs with transmissivity at all visible and IR wavelengths: thanks to the birefringent properties of diamonds, geometric phase and thus achromaticity is achieved. Manufacturing a vortex on a diamond works well in mid-infrared wavelengths, in particular in the L and N bands (Delacroix et al., 2012). In the near-IR and in the visible, however, the grooves of the sub-wavelength gratings become too small to be made on diamonds. In this case, liquid crystal polymers are usually preferred for short observing wavelengths (Mawet et al., 2009). As illustrated in Figure 1.6, SVCs are typically made of glass substrates forming a thin geometrical pattern shaped as a staircase (Delacroix et al., 2013). Recent developments have indicated that SVCs could also be made of metasurfaces (König et al., 2023).

Applications

Since 2012, VVCs have been integrated with success in several high-contrast imaging instruments in the L' band: first was for VLT/NACO (Mawet et al., 2013), where β Pictoris was successfully re-imaged (Absil et al., 2013), followed by LBT/LMIRCam (Defrère




Figure 1.6 Structure of a charge-6 scalar vortex phase mask. Such single-layer design is fundamentally chromatic. *Left*: 6-level staircase design (from [Ruane et al., 2019](#)). *Middle and Right*: Picture of a prototype and microscopic image around the central area (from [Desai et al., 2021](#)).

[et al., 2014](#)) and Keck/NIRC2 ([Femenía Castellá et al., 2016](#)). N and M-band VVCs have also been implemented on VLT/NEAR ([Maire et al., 2020](#)) and VLT/ERIS ([Kenworthy et al., 2018](#); [Davies et al., 2023](#)), respectively. Developments for VVCs are also under way regarding future projects, in particular for ELT/METIS ([Carlomagno et al., 2020](#)).

On the other hand, no SVC has yet been implemented on an observing instrument, but several projects with this flavor of vortex phase masks are underway. Current studies target space applications, in particular for the Habitable Worlds Observatory ([König et al., 2023](#); [Desai et al., 2024](#)).

1.3 Adaptive optics

Adaptive Optics (AO) systems are built to correct the atmospheric turbulence in real-time, which makes direct imaging possible for ground-based observations. Wavefront correction is also essential for focal-plane coronagraphs to work, to ensure that the starlight is well focused on the coronagraphic mask. Before diving into atmospheric turbulence and AO (Sections 1.3.2 and 1.3.3, respectively), it is first important to introduce basic concepts of optical aberrations.

1.3.1 Optical aberrations

Diffraction-limited system

Optical elements in an observing system modify the shape and direction of the incoming wavefront, as illustrated in Figure 1.3. A converging lens or mirror follows the Fraunhofer

propagation, i.e., the approximation of the Fresnel-Kirchhoff diffraction for large propagation distances. This means Fourier transforms describe well the transformation of electric fields between pupil and focal planes. In an optical system without aberrations, the electric field at the focal-plane of a converging lens can thus be written as

$$E_{\text{focal}} = \mathcal{F}[A_{\text{pupil}}], \quad (1.2)$$

where A_{pupil} is the amplitude of the pupil-plane aperture, \mathcal{F} is the Fourier transform operator, and E_{focal} is the electric field in the focal-plane.

The corresponding point-spread function (PSF) is the intensity in the focal plane and is expressed as

$$\mathcal{I} = |E_{\text{focal}}|^2. \quad (1.3)$$

For a circular aperture and a perfect optical system, such a PSF is called an Airy pattern (Goodman, 2005). The angular resolution Θ of the system is in this case defined as

$$\Theta \simeq 1.22 \frac{\lambda}{D}, \quad (1.4)$$

where D is the diameter of the aperture and λ is the observing wavelength. Such an angular resolution would be achieved if the telescope was diffraction-limited: the diffraction caused by the aperture D of the telescope will always limit the resolution. The minimum separation needed to be able to distinguish two objects, such as a star and a planet, is thus defined by the Equation 1.4. It is called the Rayleigh criterion. For instance, if we consider an 8-meter telescope observing at $2 \mu\text{m}$, any planet needs to lie at least 0.06 arcsec away from the star to be distinguishable.

Phase aberrations

In the real world, optical systems are never perfect, and changes in the index of refraction in the medium or optical surface defects produce differential optical paths. Phase aberrations make the overall wavefront deviate from an ideal shape (flat or spherical) during propagation inside the observing instrument. They correspond to phase shifts with respect to a reference unaberrated wavefront. In the presence of phase aberrations, the expression for the PSF becomes:

$$\mathcal{I} = \left| \mathcal{F}[A_{\text{pupil}} e^{i\Phi}] \right|^2, \quad (1.5)$$

where \mathcal{I} is the PSF or intensity in the focal plane, and Φ are the phase aberrations in the pupil plane.

The resulting PSF therefore contains what we call speckles. Those originate from the interference between wavefronts that follow different optical paths due to the optical




Figure 1.7 Effect of the atmospheric turbulence on a non-coronagraphic PSF and AO correction. *Left*: Short exposure PSF showing the presence of speckles distributed around the center. *Middle*: Long exposure PSF appearing as a large seeing halo. *Right*: A closed-loop AO correction yielding a diffraction-limited PSF. From [Cantalloube \(2016\)](#).

imperfections. Speckles appear as blob-like structures, sharing the same focal-plane signature as exoplanets. An image containing such structure is shown in Figure 1.7 (left), where multiple speckles are distributed around the center of the image. They represent the main source of noise in HCI images, preventing the detection of exoplanets ([Racine et al., 1999](#)).

To describe and analyze phase aberrations, it is practical to use an orthogonal basis, because it allows to decompose the wavefront into independent components. The most common basis is made of Zernike polynomials, that allow to express the phase as

$$\Phi = \sum_{j=1}^{N+1} c_j Z_j, \quad (1.6)$$

with Φ the phase, c_j the Zernike coefficients, Z_j the Zernike polynomials and N the number of modes considered.

The Zernike modes Z_j are defined on a circular aperture and are characterized by their radial and azimuthal orders. In this thesis, the Noll convention ([Noll, 1976](#)) is followed to represent the Zernike modes. In such a representation, the modes are arranged as shown in Figure 1.8 and their expressions can be found in Table 1.2. The modes are ordered in terms of their frequency content and symmetry and can be defined with sine and cosine functions using polar coordinates. The symmetric modes notably depend on $\cos(m\theta)$ and the antisymmetric ones on $\sin(m\theta)$, with m the azimuthal order and θ the azimuthal angle.

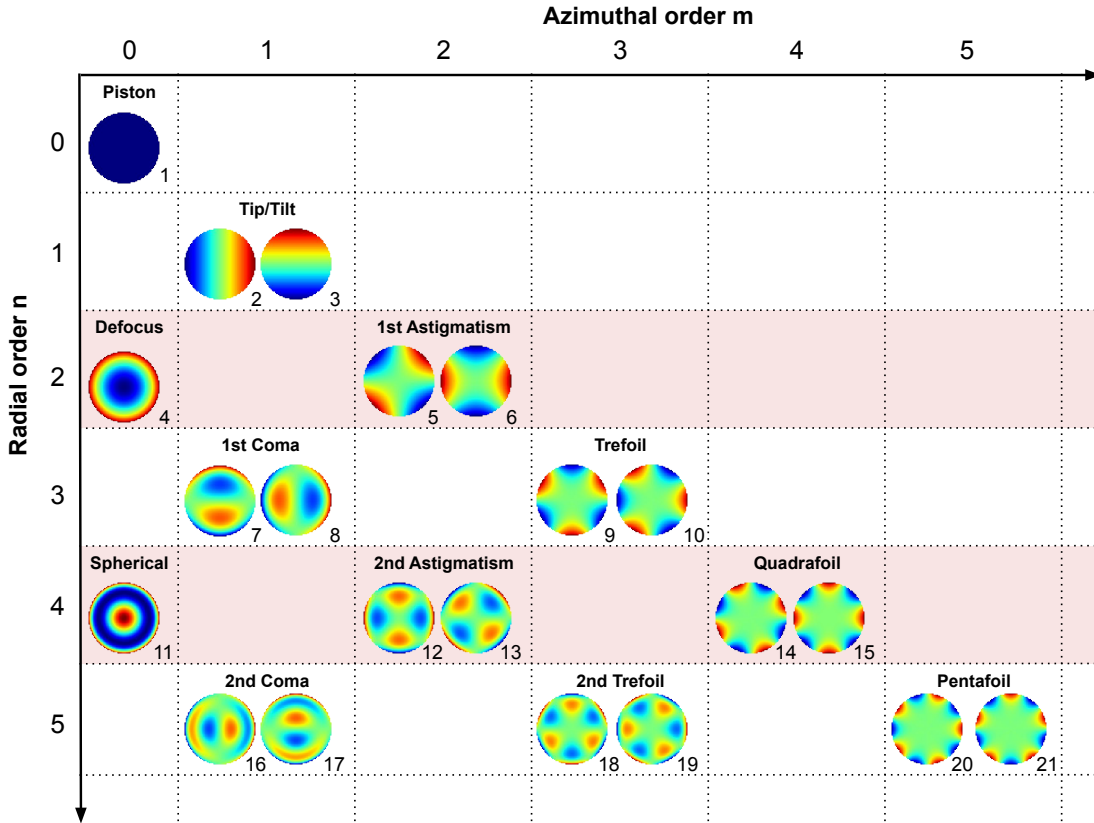


Figure 1.8 Ordering of the first 21 Zernike polynomials with their indices j according to the Noll convention. The even radial order are highlighted in red.

j	n	m	Z_j
1	0	0	1
2	1	1	$2\rho \cos(\theta)$
3		1	$2\rho \sin(\theta)$
4	2	0	$\sqrt{3}(2\rho^2 - 1)$
5		2	$\sqrt{6}\rho^2 \cos(2\theta)$
6	2	2	$\sqrt{6}\rho^2 \sin(2\theta)$
7	3	1	$\sqrt{8}(3\rho^3 - 2\rho) \sin(\theta)$
8		1	$\sqrt{8}(3\rho^3 - 2\rho) \cos(\theta)$
9		3	$\sqrt{8}\rho^3 \sin(3\theta)$
10	3	3	$\sqrt{8}\rho^3 \cos(3\theta)$
11	4	0	$\sqrt{5}(6\rho^4 - 6\rho^2 + 1)$
12		2	$\sqrt{10}(4\rho^4 - 3\rho^2) \cos(2\theta)$
13		2	$\sqrt{10}(4\rho^4 - 3\rho^2) \sin(2\theta)$
14	4	4	$\sqrt{10}\rho^4 \cos(4\theta)$
15		4	$\sqrt{10}\rho^4 \sin(4\theta)$
16	5	1	$\sqrt{12}(10\rho^5 - 12\rho^3 + 3\rho) \cos(\theta)$
17		1	$\sqrt{12}(10\rho^5 - 12\rho^3 + 3\rho) \sin(\theta)$
18		3	$\sqrt{12}(5\rho^5 - 4\rho^3) \cos(3\theta)$
19		3	$\sqrt{12}(5\rho^5 - 4\rho^3) \sin(3\theta)$
20		5	$\sqrt{12}\rho^5 \cos(5\theta)$
21	5	$\sqrt{12}\rho^5 \sin(5\theta)$	

Table 1.2 Expressions for the first 21 Zernike polynomials Z_j in polar coordinates (ρ, θ) with their radial order n and azimuthal order m (even modes in red).

Amplitude aberrations

In addition to phase errors on the wavefront, there are also amplitude aberrations that impact the intensity in the science camera. Amplitude errors can be due to several factors, such as irregularities in optical surfaces or phase errors in intermediate planes (outside the pupil plane), which translates into amplitude errors in the next plane due to Fresnel propagation.

Some methods have been developed to measure amplitude errors; for instance the coronagraphic phase diversity approach proposed by [Herscovici-Schiller, O. et al. \(2018\)](#). Two deformable mirrors (DMs, see Section 1.3.3) are usually required to correct amplitude aberrations ([Pueyo et al., 2009](#)), although it is also possible to use only one DM, at the expense of correcting only half the focal-plane ([Mazoyer et al., 2013](#)). Amplitude aberrations may cause a loss of contrast, but they have a minor impact compared to phase aberrations, which definitely need to be mitigated for high-contrast imaging.

1.3.2 Atmospheric turbulence

The Earth’s atmosphere is a dynamic medium. Turbulence, i.e., rapid fluctuations of temperature and density, arise in such medium. This impacts the refraction index of the atmosphere, and local variations of this index cause the incoming star wavefront to be distorted before reaching ground-based telescopes. Resulting short-exposure images thereby contain speckles, as it can be seen in Figure 1.7 (left). Their lifetime, as quantified by the coherence time τ_0 , is in the order of milliseconds ([Macintosh et al., 2005](#)). With long exposures, these speckles average out over time, producing blurred PSFs with a full width at half maximum (FWHM) that measures the so-called seeing (Figure 1.7, middle). The coherence length of the atmosphere is quantified by the Fried parameter r_0 ([Fried, 1966](#)), which depends on the observing wavelength: $r_0 \propto \lambda^{6/5}$. Although the angular resolution θ of a telescope increases with larger apertures D (see Equation 1.4), it is in fact limited by r_0 in the presence of turbulence. If $D \geq r_0$, the effective resolution is then:

$$\theta \approx \frac{\lambda}{r_0}. \quad (1.7)$$

The telescope is in this case seeing-limited because it has the same resolution as a telescope of diameter r_0 . In good astronomical sites at $\lambda = 500$ nm, r_0 takes values between 10 and 20 cm. Following Equation 1.4, it roughly translates to a seeing between 1 and 0.5”.

Ground-based telescopes can have very large mirrors. For instance, the four VLTs have primary mirrors of 8.2 m, and the future ELT will have a 39 m segmented primary mirror. To be able to reach the extreme spatial resolution they are capable of, it is essential

to correct the atmospheric turbulence. Adaptive optics offers a powerful solution, and it is described in Section 1.3.3.

1.3.3 Adaptive optics

As described in the previous section, the atmospheric turbulence considerably reduces the spatial resolution of telescopes. To mitigate this major issue, adaptive optics systems have been developed. An AO system is built to perform real-time corrections of the wavefront, most commonly in a closed-loop fashion (Figure 1.9). It is based on three main components:

- **Wavefront sensor:** It measures the optical aberrations present in the incoming wavefront. The Shack-Hartmann (SHWFS, [Shack & Platt, 1971](#)) and the Pyramid wavefront sensors (PyWFS, [Ragazzoni, 1996](#)) are the two main designs used today. The PyWFS is notably more sensitive to low-order aberrations in the presence of noise compared to the SHWFS, which is a strong advantage for HCI ([Guyon, 2005](#)).
- **Real-time controller:** Based on the information obtained from the wavefront sensor (WFS), it computes the command to be sent to the deformable mirror in terms of actuator displacements. This is achieved not only from the current wavefront measurement, but also integrating the previous steps.
- **Deformable mirror:** It changes its shape in real-time according to the command sent by the controller, in order to modify and correct the incoming wavefront.

For a point source such as a star, AO corrects most of the speckles and the energy is gathered around the center of the image, as illustrated in Figure 1.7 (right). In this example, the diffraction rings shows that the image is diffraction-limited.

The quality of an AO system is determined by its capability to properly analyze and correct the incoming wavefront, and Extreme Adaptive Optics (XAO) have been developed to meet the requirements of HCI ([Guyon, 2018](#)). A typical XAO system operates at 1 kHz. The primary mirror is usually too large to be altered at such a rate, so the DM is then a much smaller mirror that can be common to all instruments (e.g., on the ELT) or installed in each one of them (e.g., SPHERE on the VLT, [Beuzit et al., 2019](#)).

The performance of an AO system is generally quantified by the Strehl ratio SR:

$$\begin{aligned} \text{SR} &= \frac{I_{\text{corr}}(0)}{I_{\text{diff}}(0)} \\ &\simeq e^{-\sigma_{\Phi}^2}, \end{aligned} \tag{1.8}$$

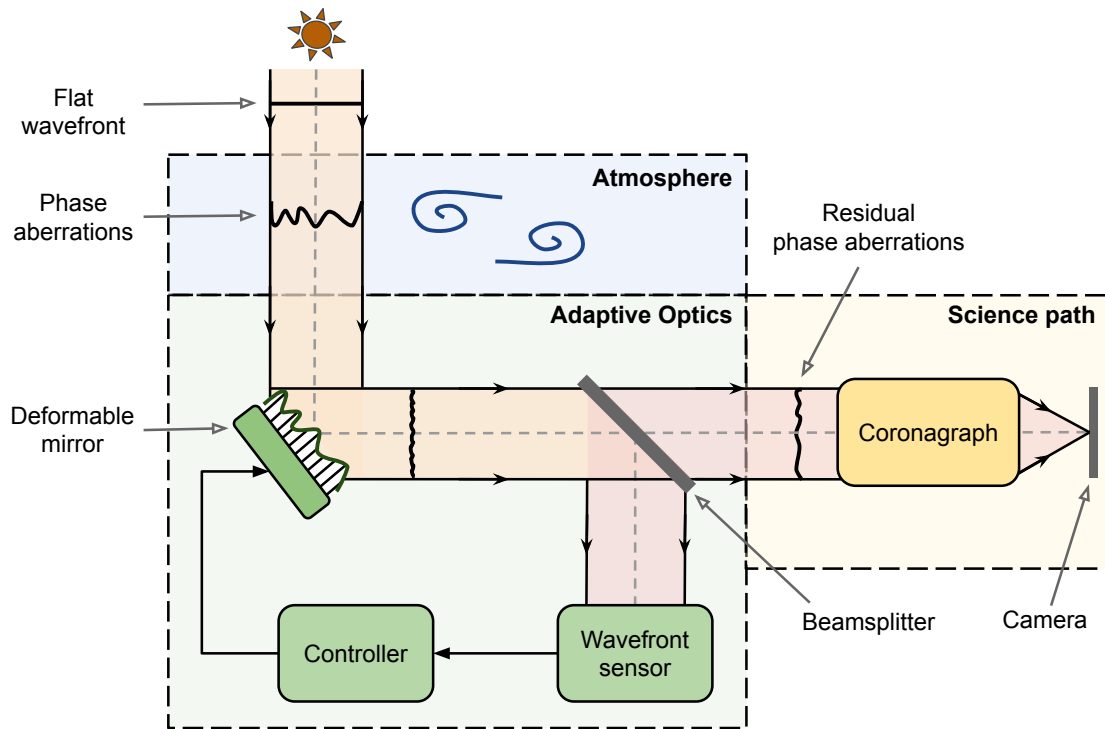


Figure 1.9 General representation of an adaptive optics system. The atmosphere (blue) distorts the incoming wavefront (orange), producing phase aberrations. The AO system corrects the aberrations in real-time using a wavefront sensor, a controller and a deformable mirror (green), leaving mainly high-order residual aberrations behind. A beamsplitter allows to provide a fraction of the starlight to a coronagraph (yellow) and the science camera. The non-common path (red) also gives rise to NCPAs on top of the atmospheric residuals, resulting in the speckles observed in the science image.

with $I_{\text{corr}}(0)$ the peak intensity of the corrected PSF, $I_{\text{diff}}(0)$ the non-aberrated peak intensity (diffraction limited), Φ the phase aberrations and σ_{Φ} the root-mean-square (RMS) wavefront error (WFE) of the phase in the pupil plane. The Strehl ratio can be expressed as a function of σ_{Φ} in the small aberration regime, i.e., if $\sigma_{\Phi} < \lambda/2\pi$. An illustration of the SR is shown in Figure 1.10, where the SR corresponds to the peak intensity of the aberrated PSF since the diffraction-limited PSF is normalized. In a real-life system, it is not possible to directly get an exact reference for the intensity, because of the intrinsic aberrations that are always present in the instrument. Different methods can thus be used to obtain a SR value, which can make comparisons between systems difficult (Roberts et al., 2004).

The Strehl ratio is an important metric in the context of HCI because a high SR means




Figure 1.10 PSF flux profile and corresponding Strehl ratio (SR). An horizontal cut is taken from the PSFs (*left*) and the corresponding 1D flux is plotted (*right*). The flux is normalized with regard to the diffraction-limited PSF (red), and the SR is therefore equal to the relative flux of the aberrated PSF (orange). In this example, we have $SR \approx 70\%$.

more energy in the PSF core and therefore a stronger exoplanet signal. As illustrated in Figure 1.1, however, the contrast remain the most fundamental metric to assess HCI performance. The raw contrast, in particular, is often used to showcase coronagraphic capabilities (Delacroix et al., 2022), potentially operating behind an AO system. The raw contrast is defined as the ratio between the intensity of the post-coronagraphic PSF and an off-axis (non-coronagraphic) PSF, and is computed at every separation from the star. The $5\text{-}\sigma$ contrast is also often used to assess detection limits (Mawet et al., 2014; Carlomagno et al., 2020) for post-processing techniques (see Section 1.5).

1.3.4 Non-common path aberrations

Not all phase aberrations can be corrected by the AO system. As it can be seen in Figure 1.9, the optical path leading to the science camera and the one reaching the wavefront sensor diverge. As a consequence, aberrations arising between the beamsplitter and the science camera cannot be seen by the WFS. Additionally, any aberration present between the beamsplitter and the WFS will be measured, but the correction will introduce new aberrations in the path leading to the science camera. AO systems are therefore not able to detect and correct these so-called non-common path aberrations (NCPAs), which can manifest themselves as residual speckles in the focal plane. The amount of NCPA depends on the instrument, but typical values are around 50 nm RMS such as in SPHERE (Vigan et al., 2019).




Figure 1.11 Example of two images obtained with the SPHERE instrument behind a coronagraph. *Left*: Image acquired using an internal source and containing only NCPAs. *Right*: On-sky image containing NCPAs and dynamic AO residuals (smooth halo at the center). In both cases, many residual speckles are present. The circular structure hints at the area of AO correction while the cross-like pattern is an artifact generated by the piezo actuators of the DM. The colors are arbitrary. From [Potier et al. \(2020\)](#); [Galicher & Mazoyer \(2023\)](#).

Optical aberrations can have different dynamics, depending on their source. We can define two types of NCPAs ([Martinez et al., 2012](#)):

- **Static**: caused by optical surface errors and misalignment in the instrument optical train.
- **Quasi-static**: originating from thermal and optomechanical deformations, including moving optics, occurring during observation sequences.

Fast varying aberrations due to atmospheric turbulence residuals produce speckles that are averaged over the integration time of the detector, resulting in a smooth halo in the image that can be removed in post-processing. On the other hand, static NCPAs can be corrected with a single calibration of the instrument before the observation. Quasi-static NCPAs are, however, particularly problematic because they can evolve during the observation sequence: they have typical timescales in the order of minutes or hours, depending on the instrument and the observing conditions ([Martinez et al., 2013](#); [Milli et al., 2016](#)). Some NCPA speckles have also been found to vary on the order of a few seconds and are likely due to turbulence within the instrument ([Vigan et al., 2022](#)). Quasi-static NCPAs are therefore neither smoothed nor can they be easily calibrated. Typical science images after coronagraphy and adaptive optics can be found in Figure 1.11, where remaining quasi-static speckles around the center of the images can be clearly seen.

1.4 Focal-plane wavefront sensing

As explained above, it is essential to correct wavefront aberrations in high-contrast imaging to be able to distinguish planetary companions. Low-order aberrations in particular can strongly degrade the coronagraphic performance. Even with an efficient adaptive optics system, NCPAs allow light artifacts from the star to reach the science camera. Focal-plane wavefront sensing (FPWFS), by utilizing information from the focal-plane (where the science camera is positioned), aims at estimating and correcting NCPAs to enhance the performance of high-contrast imaging (HCI). This section first defines the inverse problem faced by FPWFS, before explaining the sign ambiguity prevalent to this problem, before finally presenting current FPWFS methods.

1.4.1 General problem statement

The simplified expression of Equation 1.5 states that the PSF intensity \mathcal{I} is a function of phase aberrations Φ . In reality, the relationship between the phase and the PSF is not deterministic. Noise sources also impact the PSF, such as photon, detector and AO residual noise. As a consequence, for a given Φ , \mathcal{I} is rather a realization of a probability distribution

$$\mathcal{I} \sim p(\mathcal{I}|\Phi). \quad (1.9)$$

This process is illustrated in Figure 1.12 (top). In Bayesian term, and as represented in Figure 1.12 (bottom), the problem faced by FPWFS is to estimate the posterior $p(\Phi|\mathcal{I})$ from the observed \mathcal{I} as

$$p(\Phi|\mathcal{I}) \propto p(\mathcal{I}|\Phi) p(\Phi), \quad (1.10)$$

where $p(\mathcal{I}|\Phi)$ is the likelihood and $p(\Phi)$ the prior. The evidence $p(\mathcal{I})$ is not shown in the expression because it is only a regularization term and does not depend on Φ .

Several aspects of phase aberrations can help define the prior distribution. For example, we know that phase aberrations created by optical systems usually follows a power spectral density (PSD) profile $S \approx 1/f^2$, with f the spatial frequency (Dohlen et al., 2011; Orban de Xivry et al., 2021). A decomposition into Zernike components (Equation 1.6) therefore means low-order modes have stronger values than high-order modes. The global amounts of phase aberrations are also often constrained within a range, with RMS WFE around known levels for HCI instruments (see Section 1.3.4).

Extracting an estimate of Φ from the posterior $p(\Phi|\mathcal{I})$ is most commonly done with

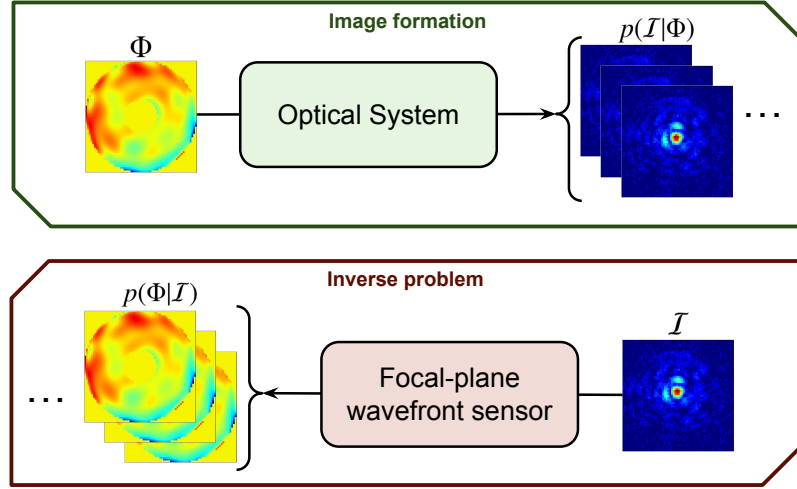


Figure 1.12 Image formation (top) and inverse problem (bottom) to solve with focal-plane wavefront sensing, where phase map distributions ($p(\Phi|I)$; left) are to be reconstructed from images (I ; right).

the Maximum a posteriori (MAP) estimate

$$\widehat{\Phi} = \arg \max_{\Phi} p(\Phi|I). \quad (1.11)$$

Other ways to compute $\widehat{\Phi}$ exist, such as with the posterior mean or Monte Carlo methods. While the primary target of FPWFS are NCPAs, all residual phase aberrations present in the system can potentially be estimated.

Solving the inverse problem defined in Equation 1.10 is challenging due to the loss of phase information that occurs when the signal is converted into intensity. This is caused by the absolute value operation applied to the electric field by the detector (Equation 1.5). Additionally, beyond the presence of noise in the images, properly estimating $\widehat{\Phi}$ from the posterior is also difficult because of a phase sign ambiguity, which will be explained in the next section.

1.4.2 Phase sign ambiguity

One important hindrance to phase retrieval from PSF measurements is the so-called twin-image problem (Gonsalves, 1982; Guizar-Sicairos & Fienup, 2012). It manifests a sign ambiguity for Zernike modes of even radial order (vertical axis in Figure 1.8), such as defocus, astigmatism or spherical modes.

The twin-image problem can be explained by the Hermitian symmetry existing in Fraunhofer propagation

$$\mathcal{F}[E_{\text{pupil}}(x)] = \mathcal{F}[E_{\text{pupil}}^*(-x)], \quad (1.12)$$

where E_{pupil} denotes the electric field in the pupil plane and E_{pupil}^* its conjugate.

Even Zernike modes are symmetric in horizontal and vertical axis, unlike odd modes, as it can be seen in Figure 1.8. This means flipping any axis does not have any effect on the even modes

$$\Phi_{\text{even}}(x) = \Phi_{\text{even}}(-x), \quad (1.13)$$

where Φ_{even} is the phase containing only even Zernike modes.

By combining Equations 1.5, 1.12 and 1.13, it can be deduced that the same PSFs are obtained from $\Phi_{\text{even}}(x)$ and $-\Phi_{\text{even}}(x)$:

$$\left| \mathcal{F}[A e^{i\Phi_{\text{even}}(x)}] \right|^2 = \left| \mathcal{F}[A e^{-i\Phi_{\text{even}}(x)}] \right|^2. \quad (1.14)$$

This degeneracy of the phase sign means the phase retrieval solution $p(\Phi|I)$ is bimodal. Figure 1.13 illustrates this ambiguity for single modes and a combination of several low-order modes. One of the standard way to solve this problem is to incorporate phase diversity (Gonsalves, 1982). We can observe in Figure 1.13c, that if the astigmatism mode is fixed, for two opposite defocus maps, the resulting PSFs are different. This means that knowing the sign of only one mode of even radial order is enough to lift the ambiguity. The traditional phase diversity approach is thus to introduce a known even mode aberration, typically a defocus (Vievard et al., 2019). The resulting PSF can then be used together with the science image to perform focal-plane wavefront sensing. Because this problem exists only with centrosymmetric pupils, it is also possible to use an asymmetric pupil to solve the ambiguity (Martinache et al., 2016; Bos et al., 2019).

1.4.3 Types of sensing techniques

Several approaches have been developed for focal-plane wavefront sensing. The most natural strategy is to directly estimate the phase aberrations from the science image or PSF. This is, however, rarely possible due to the phase sign ambiguity presented in Section 1.4.2. Temporal or spatial modulation is thus often necessary (Jovanovic et al., 2018), although direct retrieval remains possible in some cases.

Temporal modulation

To solve the sign ambiguity, temporal diversity can be exploited by injecting known phase errors and using multiple focal-plane images.

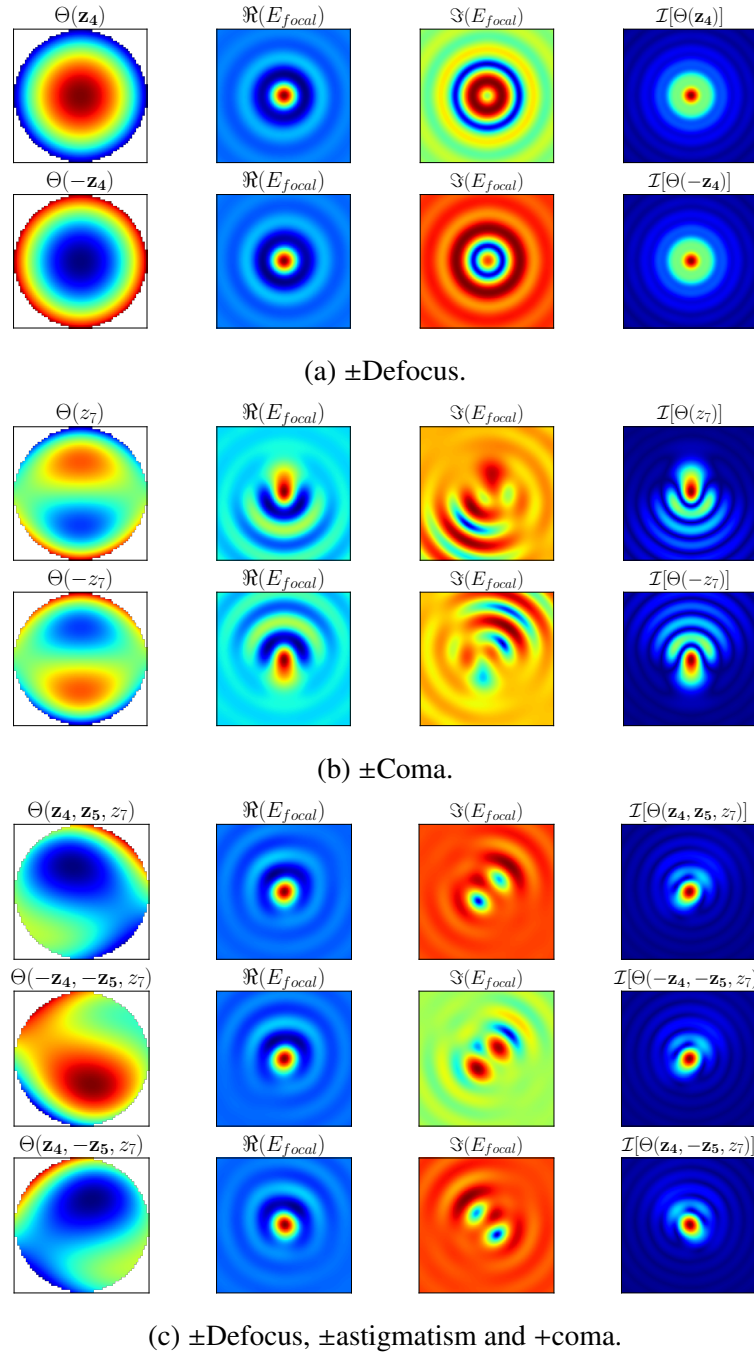


Figure 1.13 Illustration of the sign ambiguity problem. Phase maps Φ are shown with the resulting focal-plane electric fields (real \Re and imaginary \Im parts) and images \mathcal{I} after a pupil-plane to focal-plane wavefront propagation using a circular aperture. The even modes are indicated in bold font. (a) Even mode: defocus (z_4) with opposite signs. (b) Odd mode: coma (z_7) with opposite signs. (c) Composition of even and odd modes: defocus, astigmatism (z_5), and coma with three different sign configurations.

One of the most standard approaches for FPWFS is the Gerchberg-Saxton algorithm (G-S, [Gerchberg, 1972](#)). It iteratively reconstructs an estimate of the phase aberrations using constraints on the amplitude. The electric field in the pupil plane is first computed with a random phase map and the known aperture. The electric field in the focal-plane is then obtained after a Fourier transform, and its amplitude is replaced with the real PSF amplitude ($A = \sqrt{\text{PSF}}$). An estimate of the phase in the pupil-plane is finally obtained with an inverse Fourier transform. This process is repeated until the reconstructed PSF is close enough to the real one, according to a defined error criterion. In its original implementation, the G-S algorithm cannot reconstruct the exact phase due to the twin-image problem. This is why the standard G-S algorithm should be modified to include phase diversity, for instance using a defocused PSF ([Orban de Xivry et al., 2021](#)). Beyond FPWFS, the algorithm was notably applied to reconstruct the wavefront from PyWFS measurements on the SEAL testbed ([Chambouleyron et al., 2024](#)). The G-S algorithm has several limitations, such as the time it takes to converge to a solution, the dependence on the initial state of the phase, the lack of noise representation and the phase wrapping occurring during the process.

Among the numerous approaches proposed since then, the Fast and Furious algorithm (F&F, [Keller et al., 2012](#); [Korkiakoski et al., 2014](#)) was developed based on sequential phase diversity. The focal-plane image is numerically split into even and odd components and the DM command at the previous step is used to lift the sign ambiguity. F&F was applied on-sky with SCEXAO to correct the island and low-wind effects in particular ([Bos et al., 2020](#)).

Temporal diversity can also be achieved with coronagraphic systems, as done by the COFFEE algorithm ([Sauvage et al., 2012](#); [Paul et al., 2013](#)). It estimates the quasi-static aberrations using a physical model of the optical system and two focal plane images. The method was tested on the SPHERE instrument ([Paul et al., 2014](#)) and an extension also allows for amplitude aberrations measurement ([Herscovici-Schiller, O. et al., 2018](#)).

It is also relevant to measure the focal-plane electric field, which contains information about the speckles. Pair-wise probing (PWP, [Give'on et al., 2011](#)) is a widely used technique that introduces phase probes on the DM to help estimate the focal-plane electric field to be cancelled. The control approach commonly used with PWP aims to create a “dark hole” — a region of significantly reduced contrast in the focal-plane image where speckles are minimized. Rather than correcting the NCPAs directly, the DM is used to introduce specific aberrations that destructively interfere with the speckles. The most common technique for creating such dark hole is Electric Field Conjugation (EFC, [Give'on et al., 2007](#)), which uses a model of the instrument to compute the optimal DM command for cancelling the electric field in the focal plane. The PWP+EFC combination has notably been used to

create dark holes on sky with SPHERE (Potier et al., 2022). Recognizing the robustness limitations of model-dependent methods, a model-free approach was proposed, using an interaction matrix calibrated directly on-sky (iEFC, Haffert et al., 2023). Dark hole digging with both scalar and vector vortex coronagraphs has been performed on the IACT testbed (Desai et al., 2023), comparing various configurations of EFC: with pairwise probing, iEFC, and the self-coherent camera technique (discussed in the following paragraph).

Spatial modulation

Spatial diversity relies on a single PSF, for instance by making the focal-plane image interfere with an arbitrary probe.

The self-coherent camera (SCC, Baudoz et al., 2006) technique works on this principle. A small amount of the light is allowed to pass through a “pinhole” in the Lyot plane of the coronagraph. This light then interferes with any on-axis starlight that is not blocked by the coronagraph, i.e., the aberrated part of the light. The resulting Fizeau fringes in the focal plane encode the NCPAs, from which the electric field in the focal plane can be computed to estimate and correct the aberrations. Among its various implementations, the SCC is the baseline for the upgrade of the Gemini Planet Imager (Chilcote et al., 2020).

Asymmetric pupil masks are also a way to perform spatial modulation and lift the sign ambiguity. The asymmetric pupil Fourier wavefront sensor (APF-WFS, Martinache, 2013), for instance, is an interferometric approach working in the small aberration regime that takes a non-coronagraphic image in the Fourier plane to reconstruct the NCPAs. The vAPP coronagraph (see Section 1.2.1) can be implemented to work as a wavefront sensor to estimate NCPAs (Bos et al., 2019): the dual-PSF measurements obtained behind the vAPP and an asymmetric pupil are fed into a non-linear algorithm to reconstruct the phase aberrations. In the same paper, the method was demonstrated on-sky with SCExAO.

Direct phase retrieval

Some methods have been developed without requiring modulation of the wavefront. The Quadrant Analysis of Coronagraphic Images for Tip-tilt Sensing (QACITS) technique was proposed by Huby et al. (2015) for the vortex coronagraph. Tip-tilt errors create asymmetries in a post-VVC PSF. By splitting the image into four quadrants and using an analytical model, the amount of aberrations for these two Zernike modes can then be retrieved. One major advantage of QACITS is that it does not require any modification of the optical system. No phase diversity is needed because tip-tilt odd radial order modes. QACITS was tested on-sky with the Keck/NIRC2 instrument (Huby et al., 2017). The VVC’s properties for phase retrieval can also be leveraged with the Nijboer-Zernike Phase Retrieval method

(Riaud et al., 2012b,a). With this approach, the circular polarization are split to provide two complementary focal-plane signatures, and an analytical model can then reconstruct both low and higher-order aberrations.

Another way to work without modulation is to use AO residuals as phase diversity. The phase sorting interferometry (PSI, Codona et al., 2008) method achieves this by using the wavefront sensor slopes together with the science images to reconstruct the NCPAs. Linear dark field control (LDFC, Miller et al., 2017) monitor speckles changes observed outside the dark hole region to reconstruct the pupil-plane aberrations. LDFC cannot create a dark hole itself and provides a relative measurement of the NCPAs.

Challenges

Most of the approaches described above either introduce modulation of the wavefront that reduces the science duty cycle or require hardware modifications. The few techniques that can directly reconstruct the pupil-plane phase only work for specific coronagraph or instrumental designs, and sometimes make assumptions about the wavefront aberration level. This is why there is no unique FPWFS method used on-sky today. Current methods also lack generalization by providing point estimates rather than full posterior distributions, as defined in Equation 1.10. Such an estimate does not quantify the level of confidence in the prediction.

Having an approach that is valid for different conditions is thus strongly motivated. Deep learning methods are an interesting option that can offer generic models with minimal hardware requirements, and that have the potential of performing direct phase retrieval. Such approaches are also suitable to predict posterior distributions, which incorporate prediction uncertainties. Deep learning methods are introduced in Section 2.2.

1.5 Science data processing

After wavefront sensing and control, any remaining speckles may hinder the capacity to extract exoplanetary signal. Image processing techniques applied on the acquired science data, i.e., the PSF plus any potential exoplanet in the FoV, represent a fundamental step to achieve detection. Post-processing techniques allow for characterizing exoplanets as well, notably by extracting their flux and spectra to study their atmospheric content.

Most post-processing techniques are based on the angular differential imaging (ADI, Marois et al., 2006) observing and processing strategy. While following a star in the sky along its trajectory, the pupil of the instrument is kept in position (i.e., the field derotator is switched off), so the aberrations that produce the speckles stay static with regard to the

science camera. Meanwhile, the relative rotation of the sky makes any companion of the on-axis star move around it. This allows to discriminate apparent rotating planets from (quasi-) static speckles, as represented in Figure 1.14. Beyond ADI, there are other observing methods for HCI: Spectral, Polarimetric, and Reference-star Differential Imaging (SDI, PDI, and RDI, respectively).

From the data cubes obtained with such observations (frames at different times, wavelengths, or polarizations), a reference PSF \mathcal{I}_{ref} is computed. It corresponds to the speckle field that can then be subtracted from the data \mathcal{I}_{raw} to remove as much as possible the aberrations and reach higher contrasts in the post-processed image $\mathcal{I}_{\text{post}}$:

$$\mathcal{I}_{\text{post}} = \mathcal{I}_{\text{raw}} - \mathcal{I}_{\text{ref}}. \quad (1.15)$$

There exist several ways to get this reference image. The most standard one is to compute the median frame of the cube $\mathcal{I}_{\text{ref}} = \text{median}(\mathcal{I}(\theta_0), \dots, \mathcal{I}(\theta_n))$, where $\mathcal{I}(\theta_i)$ is the image at parallactic angle θ_i . The median operation removes the moving companions while retaining most of the static speckle field in \mathcal{I}_{ref} . More advanced approaches such as Principal Component Analysis (PCA, Soummer et al., 2012; Amara & Quanz, 2012) and Locally Optimized Combination of Images (LOCI, Lafrenière et al., 2007) perform better than the median combination approach. Maximum-likelihood methods such as ANDROMEDA (Mugnier et al., 2009; Cantalloube et al., 2015), the KLIP forward-model matched filter (Pueyo, 2016; Ruffio et al., 2017), or PACO (Flasseur et al., 2018) go much further by making assumptions about the noise distribution on the data to better extract planetary signals.

Approaches based on machine learning have also been introduced recently. For instance, SODINN (Gomez Gonzalez et al., 2018) is a method that classifies pixels containing exoplanet signals from those containing noise only in ADI sequences. An extension to SODINN works annular-wise and accounts for the different noise distributions in the data (Cantero et al., 2023). Using half-sibling regression to handle noise in the data (Gebhard et al., 2022) or Generative Adversarial Networks (Yip et al., 2020) have also been proposed to help increase detections. A deep learning version of the PACO algorithm (deep PACO, Flasseur et al., 2024) has recently been introduced.

Adaptive optics combined with coronagraphy, NCPA correction, and data processing have allowed the detection of more than 60 exoplanets (Table 1.1). Examples of such processed images are shown in Figure 1.2. To obtain these images, a myopic deconvolution algorithm, the Template LOCI (TLOCI, Marois et al., 2014) and the LOCI approaches were used (Figures 1.2a, 1.2b and 1.2c, respectively). The four unveiled planets of the HR8799 system in particular illustrate well what direct imaging can accomplish.

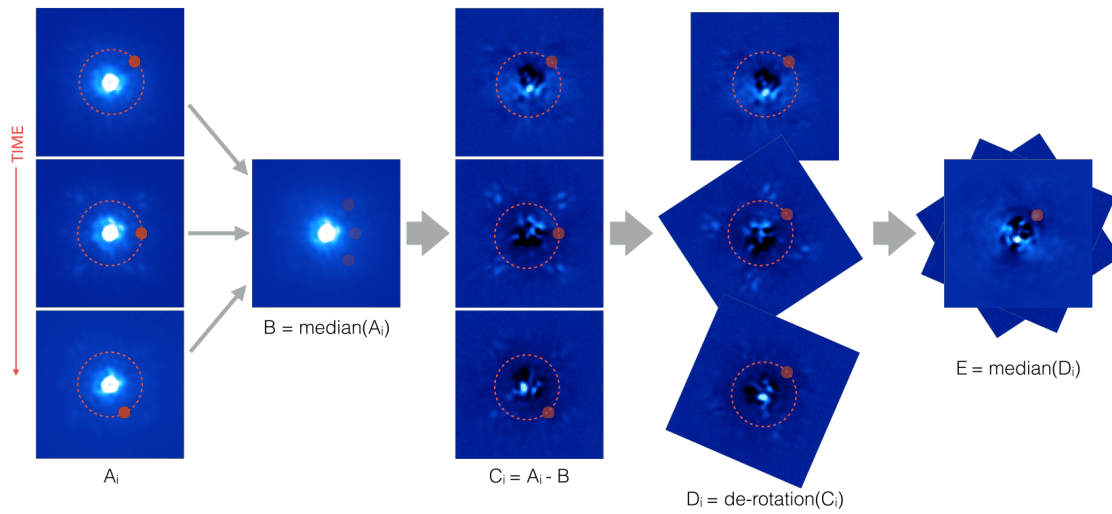


Figure 1.14 ADI sequence (A) with rotation planets over time, computation of the reference PSF using the median frame (B), new sequence with most of the speckles subtracted (C), derotated images according to their parallactic angles (D), and final post-processed image with a prominent exoplanetary signature (E). From [Gómez González \(2017\)](#)

1.6 Current limitations to high-contrast imaging

As explained in the previous sections, high-contrast imaging is an extremely challenging field. Only exoplanets presenting substantial brightness, size and distance from their star have been found so far. The contrasts needed to detect a larger range of exoplanets are difficult to reach for various reasons.

Apart from the inherent high contrasts and small angular separations between stars and their exoplanets (Section 1.1), which are addressed with coronagraphs (Section 1.2) and adaptive optics for ground-based telescopes (Section 1.3.3), remaining optical aberrations represent obstacles to detect a large variety of planetary companions. For ground-based observations, the AO system cannot correct all phase aberrations, which cause starlight leakage from the coronagraph. Temporal errors remain present because of the delay between the wavefront measurement and the correction by the AO system. This corresponds to the servolag error, and it notably produces an asymmetric halo in the focal plane, called the wind driven halo ([Cantalloube et al., 2018, 2020](#)). Another source of error is the low wind effect, which causes a gradient of temperature between the spiders of the telescope, therefore disturbing the phase in the pupil plane ([Sauvage et al., 2015; Milli et al., 2018](#)). Low-order residuals such as tip-tilt aberrations, which can be due to atmospheric residuals

or vibrations in the system (Lozi et al., 2018a), also strongly limit focal-plane coronagraphs. Additionally, the limited resolution of the wavefront sensor and the finite number of actuators for the DM do not allow to correct all high-order spatial modes. This does not represent, however, an important constraint to exoplanet imaging since high-order phase residuals are weak compared to low-order aberrations in the close vicinity of the star. All the mentioned wavefront errors play their role in limiting HCI, but one of the most important contribution today remain non-common path aberrations, which cannot be handled by the AO system. As explained in Section 1.3.4, NCPAs are particularly problematic when they have a quasi-static lifetime. Phase aberrations are also dominant in space-based observatories, making wavefront sensing and control of prime importance.

As explained in Section 1.4, focal-plane wavefront sensing offers the most natural approach to correct NCPAs. Traditional FPWFS methods work iteratively, which means predictions can be time-consuming since several steps are needed before reaching convergence (Section 1.4.3). Some approaches are also dependent on the instrument and restricted to low-order errors or small aberration regimes, for instance. Deep learning techniques can potentially offer a compelling solution to these limitations, by building performant and robust models, suited to different aberration contents and levels, offering fast predictions while being adaptable to the instrument and observing conditions. This thesis thereby focuses on exploring how deep learning approaches can leverage FPWFS to correct NCPAs in AO systems. The next chapter introduces deep neural networks, in particular how they are suited for image-based reconstructions and wavefront sensing.

Deep learning for wavefront sensing

This chapter introduces some fundamental concepts of deep neural networks, focusing on convolutional neural networks due to their importance in this thesis and presenting the motivations behind these approaches. Methods developed for wavefront sensing are then discussed, primarily within a supervised learning framework, and the motivation behind adopting physics-based learning approaches is finally highlighted.

Contents

2.1	Deep neural networks	37
2.1.1	General concept	37
2.1.2	Convolutional neural networks	40
2.1.3	Motivations and challenges	43
2.2	Application to wavefront sensing	44
2.2.1	Supervised phase retrieval	44
2.2.2	Physics-based retrieval	45

2.1 Deep neural networks

2.1.1 General concept

Data-driven models built to make predictions from observations are rising with machine learning. More specifically, deep artificial neural networks (NN) combined with vast amounts of data have allowed building accurate and robust models. With the availability of large-scale datasets, enhanced computing power facilitated by graphics processing units (GPUs), as well as the development of advanced algorithms, “deep” learning (DL) has recently made breakthroughs in a wide range of fields and applications, including virtual assistants and chatbots, language translation, self-driving vehicles, or robotics. ([Alzubaidi et al., 2021](#); [Sarker, 2021](#)).

A neural network architecture allows to build a differentiable model to approximate a function $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ parameterized by θ , with \mathcal{X} and \mathcal{Y} the input and output spaces, respectively. This function therefore yields predictions $y \in \mathcal{Y}$ from input data $x \in \mathcal{X}$, so that $y \approx f_\theta(x)$. A probabilistic representation with predictions $p(y|x) := f_\theta(x)$ can also be preferred, for instance in focal-plane wavefront sensing, as defined in Equation 1.10 and applied in Chapters 3 and 5.

The most basic constituent used to approximate a function f_θ is the Perceptron (Rosenblatt, 1958). It is an artificial neuron that computes a feature h from a linear combination of inputs x_i and weights θ_i ,

$$h = \theta_0 + \sum_i x_i \theta_i. \quad (2.1)$$

A non-linear activation function σ can then be applied on the output features as $o = \sigma(h)$. Various activation functions have been developed over the years. The rectified linear unit (ReLU), defined as $\sigma(x) = \max\{0, x\}$, is one of the most used today (Agarap, 2018).

Several of these artificial neurons can be combined in a neural network layer. In turn, several of these layers can be stacked one after the other to form a multilayer perceptron (MLP). MLPs are known to be universal approximators (Hornik et al., 1989) and are capable of learning non-linear representations. To train the parameters of a neural network, two fundamental components are needed: a loss function that computes a prediction error and a gradient descent algorithm, commonly referred to as backpropagation (Rumelhart et al., 1986), which updates the model parameters using the prediction error and derivatives obtained from the chain rule.

In supervised learning, a model is built using labelled data, i.e., the target's ground truth is known for every training sample. The gap between the predictions $\widehat{y}_i \approx f_\theta(x_i)$ and the labels y_i is therefore assessed with a loss function $\mathcal{L}_{\text{sup}} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathcal{R}$. The empirical risk for f_θ , which accounts for the average error on a training dataset containing N samples, is then

$$R(f_\theta) = \frac{1}{N} \sum_{i=1}^N \mathcal{L}(y_i, f_\theta(x_i)). \quad (2.2)$$

Empirical risk minimization is what drives the model training, aiming to find the optimal model f_θ^* from a class of functions \mathcal{F}

$$f_\theta^* = \arg \min_{f \in \mathcal{F}} R(f_\theta). \quad (2.3)$$

This solution can be usually found using stochastic gradient descent, which allows to iteratively update the model parameters θ_i as

$$\theta_i \leftarrow \theta_i - \eta \frac{\partial \mathcal{L}}{\partial \theta_i}, \quad (2.4)$$

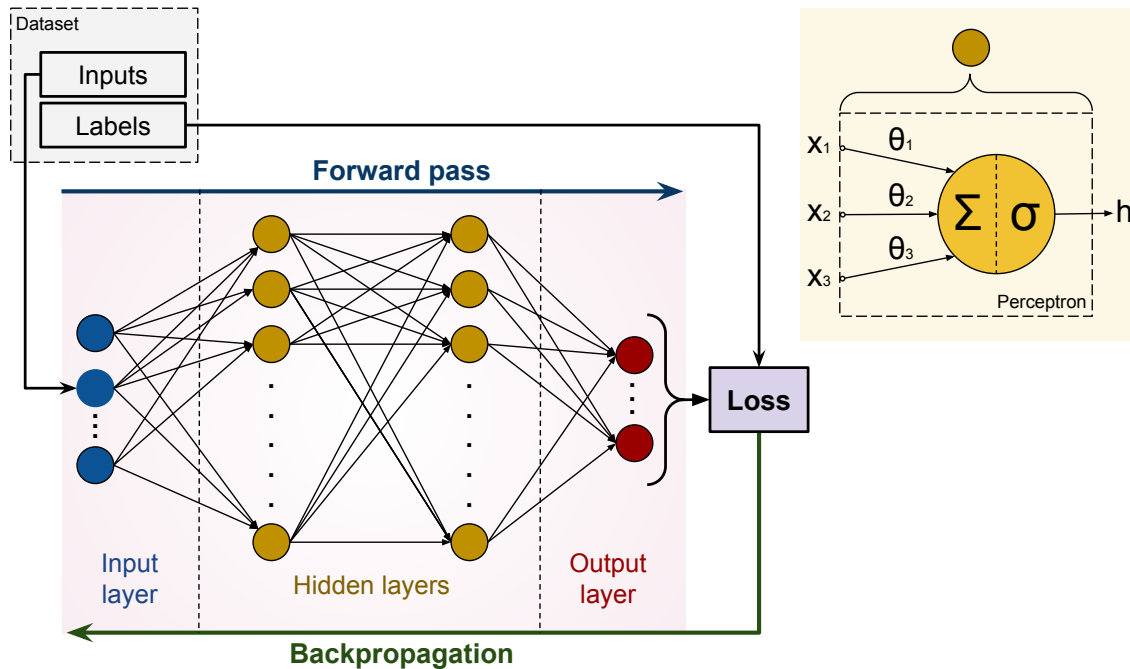


Figure 2.1 Neural network working principle, as a multilayered perceptron to be trained in a supervised fashion. The architecture is composed of several parameterized hidden layers (orange) that take the input data (blue) and yields predictions (red). One perceptron is shown in more detailed (top right): it computes a linear combination of the inputs before applying an activation function σ . For each training sample, a loss is computed from the predictions and the labels, which allows to optimize the parameters using an algorithm called backpropagation. This process is performed iteratively until the loss converges to a minimum.

where η is the learning rate and $\frac{\partial \mathcal{L}}{\partial \theta_i}$ is the gradient obtained using the chain rule. An illustration of an MLP in a supervised learning setup is shown in Figure 2.1.

In unsupervised learning, models training does not require a label for each training sample. Such models can be built for tasks such as ordering data into clusters (“clustering”), denoising, generating new data or reconstructing the input signal. For this latter task, autoencoder architecture can be used, for instance, with a loss function $\mathcal{L}_{\text{unsup}} : \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{R}$ (see Chapter 5).

For a thorough description of the working principle of deep neural networks (DNNs), the reader is encouraged to look into Prince (2023).

2.1.2 Convolutional neural networks

Multilayer perceptrons do not handle well spatial information because each input feature is processed independently and neurons are fully connected. This hinders the ability of MLPs to capture the structure in images and variations that might occur, such as spatial shifts. MLPs are also cumbersome models with a lot of redundant neural connections, making them ineffective with high-dimensional input spaces. Convolutional neural networks (CNNs) represent a much better option for image prediction tasks, including wavefront reconstruction.

CNN building blocks

The idea behind CNNs is to add inductive biases, i.e., assumptions about the data, to neural networks that are dealing with images (Cohen & Shashua, 2017). Three such biases exist with CNNs:

- **Locality:** The network should analyze the data at various scales to capture local features.
- **Translation invariance:** When parts of the image are shifted to different pixel locations, the model should produce consistent responses.
- **Hierarchical composition:** Detected patterns can be described as compositions of simpler patterns.

These properties make CNNs much more effective and performant with images than MLPs. CNN architectures are typically composed of four main building blocks: convolutional layers, pooling operations, activation functions and fully connected (FC) layers.

A convolutional layer applies successive convolutions by “sliding” a kernel, i.e, a matrix composed of trainable parameters, over the input data. For a 3D input $x \in \mathbb{R}^{C \times M \times N}$, e.g., an RGB image ($C = 3$) of $M \times N$ pixels, and a kernel $\theta \in \mathbb{R}^{C \times K \times L}$, the result of the convolution operation for each element of a hidden feature h can be written as

$$\begin{aligned} h_{i,j} &= \beta_{i,j} + \sum_{c=0}^{C-1} (\theta_c * x_c)[i, j] \\ &= \beta_{i,j} + \sum_{c=0}^{C-1} \sum_{k=0}^{K-1} \sum_{l=0}^{L-1} \theta_{c,k,l} x_{c,i+k,j+l}, \end{aligned} \tag{2.5}$$

where β is a bias that is learned together with θ . The output $h \in \mathbb{R}^{M \times N}$ is a feature map matching the input dimensions, typically followed by a non-linear activation. Equation 2.5




Figure 2.2 The LeNet-5 convolutional architecture applied to the task of handwritten digit recognition. From [Lecun et al. \(1998\)](#).


is a simplified version of the convolution operation, as other parameters such as the stride and dilation rate are ignored (the reader is referred to [Prince \(2023\)](#) for a description of these parameters).

Convolutional layers allow the extraction of spatial features at different scales from the data. Pooling layers then progressively reduce the spatial size of the representation, which decreases the number of parameters and the computational cost, while also supporting hierarchical feature composition. FC layers are finally placed before the output of the CNN to learn non-linear combinations of the high-level features provided by the final convolutional layer.


Types of architectures

One of the first successfully developed CNN is LeNet ([Lecun et al., 1998](#)). It was specifically designed for the task of handwritten digit recognition: composed of a succession of convolutional, pooling and fully connected layers, the network took as input a drawn number from 0 to 9 and was successfully able to classify these input images to the correct digit. LeNet was groundbreaking at the time of its introduction because it demonstrated the effectiveness of artificial neural network approaches for image recognition tasks. It paved the way for subsequent advancements in CNNs. The principles and architectural concepts of LeNet have been since widely adapted for various other tasks, such as object recognition, facial recognition, and image classification in general. A representation of the LeNet-5 architecture is shown in [Figure 2.2](#).

Deeper architectures have emerged since then with increasing complexity, which has also raised new challenges. One such challenge is the vanishing gradient problem ([Hochreiter, 1998](#)), which occurs when the gradients become increasingly small with the chain rule of the backpropagation algorithm. This can lead to meager or even no updates of the



(a) The ResNet-34 architecture.



(b) Building block of ResNet.

Figure 2.3 The ResNet structure, based on successive residual learning blocks. From [He et al. \(2015\)](#)

weights and therefore prevent model learning. Additionally, very deep CNNs often experience accuracy saturation or even degradation as more layers are added. This phenomenon is observed with the training data and thus cannot be a result of overfitting.

Residual neural networks (ResNets), introduced by [He et al. \(2015\)](#), address both vanishing gradient and degradation problems. ResNets contain batch normalization layers, that not only accelerate and stabilize training but also reduce the probability of facing vanishing gradients. By normalizing activations, batch normalization prevents them from saturating and keeps the gradients from diminishing to zero. To address the degradation problem, ResNets use skip connections by summation, as illustrated with the 34-layer version of ResNet in Figure 2.3a. These skip connections use identity mappings, i.e., the input of a block is added to its output (Figure 2.3b). This approach improves the gradient’s flow during backpropagation, which leads to easier optimization and prevents decreasing accuracy with the addition of more layers.

Many deep CNNs have been proposed beside ResNets. The U-Net and EfficientNet architectures have notably been used in this thesis, and their specificities are described in Sections 3.3.1 and 4.2.2, respectively. During the timeframe of this thesis, new types of CNN architectures have emerged. Among them, we can cite the Vision Transformer (ViT, [Dosovitskiy et al., 2021](#)), an application of Transformer architecture to computer vision, as well as ConvNext ([Liu et al., 2022](#)), which is inspired by ViT.

2.1.3 Motivations and challenges

Deep neural networks are very powerful and capable of approximating extremely complex functions. Many physical problems are difficult to simulate and inverse problems are often ill-posed and tedious to solve. Deep learning allows finding solutions only using observed information, and if data is sufficient and diverse enough, competitive models can be built. Depending on the architecture complexity and the amount of input data, training a DNN can take hours or even days, but predictions are very fast, usually in the order of 10 milliseconds per forward pass on a standard GPU. This means that when DNNs are deployed in production, some can achieve real-time predictions. DNN are particularly good at learning complex visual patterns and predict unseen information. We live in a world where data is abundant and where computing power is dramatically increasing. In such context DNNs are working tremendously well.

Choosing the number of neurons and layers for an NN architecture has a strong impact on the model performance. Based on the universal approximation theorem ([Hornik et al., 1989](#)), a single layer should be sufficient to represent any function, but the size of the layer could then prevent proper learning and generalization ([Goodfellow et al., 2016](#)). Multiple layers are better at capturing more complex representations because they learn all the intermediate features, i.e., the various levels of abstraction. But very deep NNs are susceptible to experience two main issues. One is overfitting, which occurs when the NN performs well on training data but fails to give consistent results on new data ([Ying, 2019](#)). Increasing the size of the training dataset can help, but regularization techniques are usually needed to improve generalization, for instance with dropout layers, loss penalty terms or early stopping. Another common issue for DNNs is the vanishing gradient problem, which is mitigated by architectures such as ResNets (see Section [2.1.2](#)).

DNNs contain millions of parameters to be fitted to the data. Because of their complexity, DNNs are often seen as a “black box”, similarly to the intricate neural connections of the brain. Extracting physical meaning from deep learning models is therefore not always straightforward, although there are ways to interpret DNNs ([Montavon et al., 2018](#); [Zhang et al., 2020](#)). Supervised learning approaches are often the easiest and fastest way to train models, but creating realistic labelled datasets can be time-consuming and is not always possible. For certain tasks, unsupervised learning and transfer learning can thus represent the best approach to achieve results with proper accuracy and robustness. Methods developed in these areas are proposed in this thesis (Chapters [5](#) and [6](#)).

2.2 Application to wavefront sensing

Artificial neural networks and other deep learning methods are being transformative for astronomy and optical applications, such as wavefront sensing and control. This section presents supervised learning methods that have been developed for wavefront sensing in astronomy, and then discusses the potential of physics-based models.

2.2.1 Supervised phase retrieval

Enhancing existing wavefront sensors

Wavefront sensing is crucial in adaptive optics systems for ground-based exoplanet imaging, as outlined in Section 1.3. Neural networks trained in a supervised way have been applied to address these challenges. The first application goes back over thirty years ago: a perceptron composed of 150 hidden units was trained on focused and defocused images from the Multiple Mirror Telescope (Angel et al., 1990) to infer the piston and tip-tilt aberrations and perform co-phasing between the telescope mirrors. This corrected the telescope’s point spread function (PSF) near the diffraction limit, first with simulations and then on a real star (Sandler et al., 1991). Early NNs were also able to infer atmospheric parameters on top of wavefront slopes (Montera et al., 1996), although small NNs have shown limitations in performance in the presence of noise (McGuire et al., 1999).

This success with early neural networks paved the way for further developments. Traditional wavefront sensors like the Shack-Hartmann and Pyramid wavefront sensors have limitations in terms of non-linearities and dynamic range. This motivates the exploration of deep neural networks to address these issues. A CNN can be used to infer the non-linear term of the reconstruction, as proposed by Landman & Haffert (2020). Combined with a linear model, the method can then retrieve the full aberrated wavefront with improved dynamic range and without noise amplification. WFS dynamic range has been improved with CNNs for other types of wavefront sensors as well (Allan et al., 2020a,b), and deep CNNs have also been developed to function as image-based wavefront sensors (Nishizaki et al., 2019). One limitation of AO systems are the inherent temporal delays that produce wavefront estimation errors (see Section 1.6). This has prompted research on predictive control, from early NNs (Jorgenson & Aitken, 1992) to deep convolutional long short-term memory (LSTM) networks (Liu et al., 2020; Swanson et al., 2021). Beyond supervised learning, reinforcement learning techniques have also been implemented for predictive control (Nousiainen et al., 2024; Pou et al., 2024; Gutierrez et al., 2024), and such approaches are briefly discussed in the thesis perspectives (Section 7.2).

Focal-plane sensing approaches

As demonstrated in Section 1.3.4, non-common path aberrations remain in AO systems and are an important part of current HCI research. Focal-plane wavefront sensing, presented in Section 1.4, allows measuring NCPAs, and deep neural networks represent a legitimate avenue for addressing the non-linearity of the problem. Phase aberration estimation from PSFs has been leveraged using standard CNN architectures such as LeNet (Wu et al., 2020; Naik et al., 2020), Inception (Andersen et al., 2019) or ResNets (Guo et al., 2019; Andersen et al., 2020). Modifications in the architectures, for instance replacing the fully-connected layers with deconvolution layers can in some cases improve speed and precision (Guo et al., 2019). Data pre-processing techniques, including principal component analysis (PCA) (Terrerri et al., 2022b,a) and Fourier-transformations (Naik et al., 2020), are also valuable for extracting relevant PSF features and reducing dimensionality before training models. Recent advancements also include the use of deep neural networks to leverage photonic fiber devices for focal-plane wavefront sensing (Norris et al., 2022; Wei et al., 2023).

Space telescopes, while free from atmospheric turbulence, still suffer from instrument-based aberrations and can thus benefit from focal-plane wavefront sensing. Early work applied neural networks to predict low-order Zernike modes for the Hubble Space Telescope (Barrett & Sandler, 1993). Harnessing advances in deep CNNs, Paine & Fienup (2018) applied the Inception v3 architecture to provide initial estimates of the wavefront from simulated James Webb Space Telescope PSFs. This approach extends the capture range beyond what standard iterative methods offer, particularly in the presence of large aberrations. Additionally, neural networks are being explored for correcting aberrations in Earth observation satellites (Dumont et al., 2024).

Challenges persist in focal-plane wavefront sensing using deep learning. Supervised methods require labeled datasets, which are difficult to acquire or may lack accuracy. Simulations offer a practical way to build these datasets, but trained models on such data often lack robustness with real data. Therefore, learning methods that do not rely on labeled data are appealing, as they can be trained on observed data only.

2.2.2 Physics-based retrieval

The integration of physical models into deep learning frameworks for focal-plane wavefront sensing is relatively underexplored. However, recent advances in other optical sensing fields demonstrate that incorporating physical models can enhance the robustness of wavefront sensing.

For instance, in the context of holographic imaging, Peng et al. (2020) have proposed

a method that combines neural networks with an optical model to learn object phase information, while incorporating known Zernike modes to enhance the phase retrieval process.

When constructing robust models is challenging, models optimized directly on observed images can be highly effective. Such an approach has been developed by [Bostan et al. \(2020\)](#) using two neural networks: one to encode Zernike modes and estimate phase aberrations, and the other to generate the object’s phase from a random tensor. These outputs are combined to produce the wavefront, which is then simulated to reconstruct the images. The networks’ weights are iteratively optimized by minimizing the least-square loss between observed and reconstructed images. Similarly, [Wang et al. \(2020\)](#) propose an optical propagation model alongside an NN inspired by U-Net. This method incorporates a “deep image prior”, which leverages prior information about the object’s phase regularities to regularize the parameter optimization process. The reconstructed images are computed from the phase predicted by the NN. These “untrained” neural networks are a great way to bypass the training part while having a model fitted to the currently observed data. They are however slower in operation since they need to converge to a solution through multiple iterations.

Deep learning methods that incorporate physical models do not rely on labeled data and can be considered unsupervised learning approaches. However, the incorporation of a simulator in the architecture implies that some level of supervision is inherent in their design¹. Additionally, the methods described above use fixed optical models, which may not perfectly match real-world data. It may thus be beneficial to address the question of optimizing them, as it is done in this thesis in Chapters 5 and 6. Additionally, the deterministic nature of phase estimation in these methods could be improved by considering estimation uncertainties. This is explored through probabilistic predictions in Chapters 3 and 5.

¹For simplicity, the term “unsupervised learning” will be used in Chapters 5 and 6 to differentiate such methods from supervised learning approaches that use labels.

Scope and outline of the thesis

As discussed in Chapter 1, exoplanet imaging is of prime importance and improving instrumental capabilities is needed. In particular, the residual phase aberrations that exist in high-contrast imaging systems are to be mitigated to reach new discovery limits.

Wavefront sensing and control is especially motivated for future projects, such as the Extremely Large Telescope with its instrument METIS, that will be able to probe at very small angular separations from the star thanks to the increased resolution. This region is polluted by low-order aberrations, producing flux leakage from focal-plane coronagraphs such as vortex phase masks. Space-based telescopes are also limited by optical aberrations, especially because they aim at much higher contrasts. The James Webb Space Telescope is currently applying co-phasing of its primary mirror to mitigate low-order optical errors. The upcoming Roman Space Telescope will be the first space-based telescope to use active wavefront sensing, which will be essential to leverage its coronagraphic capabilities. Future space projects, such as the Habitable Worlds Observatory, will surely benefit from advanced wavefront sensing and control strategies to correct instrumental aberrations.

This PhD thesis aims at improving the estimation of phase aberrations in various amplitude regimes and spatial frequency contents. Focal-plane wavefront sensing is the chosen approach because it also measures non-common path aberrations, which represent one of the main limitation for HCI systems today. Recent developments in deep neural networks, as presented in Chapter 2, show great promises to help wavefront sensing reach new limits. This thesis therefore aims at exploring how recent deep neural networks perform and how they can pave the way for future deep learning-based methods to be implemented on-sky.

In Chapter 3, a new deep learning-based focal plane wavefront sensing method is first presented with simulated data. Recent deep convolutional neural networks are implemented and the potential benefits of predicting distribution parameters is assessed. This thesis notably investigates how a new phase diversity method with the vortex coronagraph can be used to retain a 100% science duty cycle (Chapter 4). Applying the method on a real instrument can be strongly limiting since the exact wavefront state is never known, and labels used to train models in a supervised way are therefore not accurate. Label-free or “unsupervised” learning approaches are thus very interesting to explore, as described in Chapter 5. In this regard, an autoencoder architecture that uses a differentiable simulator as a decoder is proposed, allowing to build unsupervised learning models as well as to optimize physical parameters within the simulator.

Thanks to a collaboration initiated with the SCEXAO team at the Subaru Telescope, studies on a real instrument are carried out in Chapter 6. Models are trained with data acquired in the lab, i.e., using an in-house laser as light source and a deformable mirror

to generate and correct phase aberrations. Both supervised and physics-based learning approaches are tested on SCExAO in visible light. Phase map estimations from generated datasets are performed, as well as closed-loop corrections to converge to even lower phase residuals, yielding point-spread functions close to the diffraction limit. Transfer learning strategies are also investigated, in order to quickly adapt pre-trained models on new data as well as to fine-tune instrumental parameters with the physics-based autoencoder. The main points demonstrated in this thesis work as well as the different prospects that arise from it are finally summed up in Chapter 7.

PART II

Simulations

General supervised learning approach

This chapter presents the general focal-plane wavefront sensing approach of this thesis, implementing deep convolutional neural networks trained in a supervised manner. The data generation procedure to simulate point-spread functions is presented here, both in classical imaging and behind vortex phase masks. The supervised learning framework is then described. The performance between several deep CNNs is compared, as well as the impact the vortex coronagraph has on the results. Additionally, mixture density models are also trained and analyzed. Most of the content of this chapter has been published in the SPIE conference proceedings ([Quesnel et al., 2020](#)).

Contents

3.1	Introduction	54
3.2	Synthetic data generation	54
3.2.1	General procedure	54
3.2.2	Implementation of the Vector Vortex coronagraph	56
3.2.3	Data pre-processing	57
3.3	Deep learning models	58
3.3.1	Convolutional Neural Networks	58
3.3.2	Mixture Density Layer	59
3.3.3	Training procedure	60
3.4	Experiments and results	61
3.4.1	Comparison between CNN architectures	62
3.4.2	Predicted uncertainties with mixture density models	62
3.4.3	Influence of the Vector Vortex coronagraph	64
3.5	Conclusions	66

3.1 Introduction

As described in Chapter 1, residual wavefront errors, and non-common path aberrations in particular, represent key limitations to high-contrast imaging. Estimating phase aberrations with focal-plane wavefront sensing is particularly motivated, and CNNs trained in a supervised way can constitute a powerful solution (see Chapter 2). This chapter aims to leverage supervised CNNs to estimate phase maps from PSFs using simulated data. The importance of phase diversity is highlighted and an additional out-of-focus PSF is used to solve the sign ambiguity, as explained in Section 1.4.2. An implementation of mixture density networks (MDNs) is proposed in order to emphasize the sign ambiguity with probability distributions, and the method is tested with the Vector Vortex coronagraph (VVC).

The data generation procedure and deep learning approach are first explained in Sections 3.2 and 3.3. Results are then showcased in Section 3.4, with a performance comparison between the U-Net and ResNet-50 architectures as a baseline, before evaluating the MDN and studying the impact of simulated post-VVC PSFs on the performance.

3.2 Synthetic data generation

This section presents how the data is generated with simulations for classical and coronagraphic imaging, as well as the transformations made on the PSFs before model training.

3.2.1 General procedure

For focal-plane wavefront sensing in a deep learning framework, a labeled dataset composed of point-spread functions and phase maps is constructed. In classical imaging, i.e., with an optical propagation between pupil and focal planes using a single Fourier transform, the PSF is expressed from the phase as in Equation 1.5.

The sets of Zernike coefficients are first randomly generated within the range $[-1, 1]$ and each coefficient is then divided by its corresponding radial order to approximate a $1/f^2$ power spectral density profile, typically encountered with good quality optics (Dohlen et al., 2011). The entire set of coefficients for each phase map is re-scaled according to a given median root-mean-square (RMS) wavefront error $\overline{\text{WFE}}$:

$$\widehat{c}_k = \frac{c_k}{\overline{\text{WFE}} \times \text{median} \left(\sqrt{\sum_i c_{k,i}^2} \right)}, \quad (3.1)$$

where \widehat{c}_k is the normalized set of coefficients for the phase map k (see an example in Figure 3.1; left).

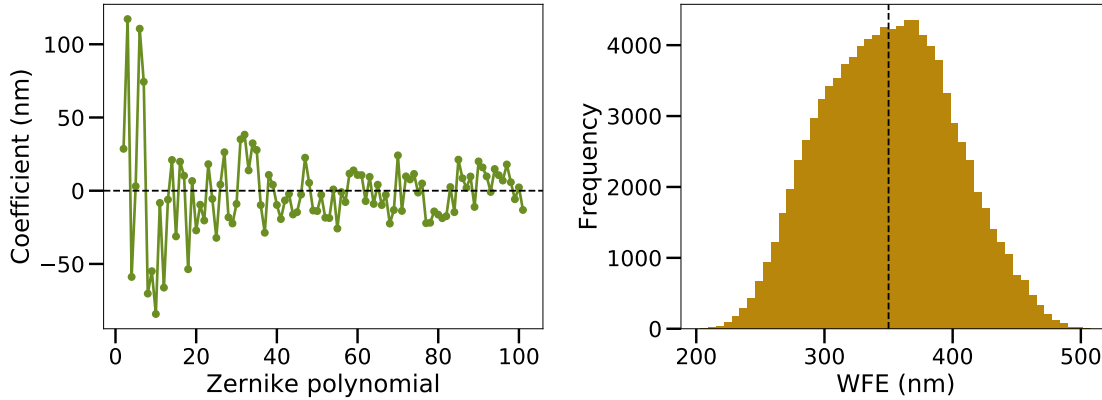


Figure 3.1 *Left*: Example of a generated set of Zernike coefficients. *Right*: Example of an RMS WFE distribution in a dataset composed of 10^5 samples. The median RMS WFE considered here is equal to 350 nm.

One can find that the RMS WFE of the phase is equal to the standard deviation on the Zernike coefficients used for the normalization (Schmidt, 2010):

$$\text{WFE}(\Phi_k) = \frac{1}{N_{\text{pix}}} \sqrt{\sum_i^{N_{\text{pix}}} \Phi_{k,i}^2} = \sqrt{\sum_i^{N_{\text{modes}}} c_{k,i}^2}, \quad (3.2)$$

with N_{pix} the number of phase map pixels that are within the entrance aperture, while N_{modes} is the number of Zernike modes considered.

The input wavefront error is, therefore, a distribution centered around $\overline{\text{WFE}}$, as represented with the histogram in Figure 3.1 (right). A higher $\overline{\text{WFE}}$ leads to a larger distribution. The phase aberrations are then simulated using Zernike polynomials as defined in Equation 1.6. Regarding the ordering of the Zernike polynomials, the Noll convention (Noll, 1976) is used starting from the tip mode. A circular entrance aperture is considered, i.e., not including the central obstruction and support structures due to a secondary mirror.

The PSFs are computed using the PROPER (Krist, 2007) optical propagation package. The data is generated in the K band at $\lambda = 2200$ nm, while an aperture diameter of 10 m, a pixel scale of $0.25 \lambda/D/\text{pix} \simeq 11$ mas/pix, and a field-of-view of $32.3 \lambda/D \simeq 1.47''$ are considered. The phase maps and final PSFs are cropped to contain 129×129 pixels. Each sample from the dataset is composed of in-focus and out-of-focus PSFs, necessary for phase diversity. A defocus of $\lambda/4 = 550$ nm is introduced in the pupil plane to obtain the out-of-focus PSFs. Figure 3.2 illustrates phase maps and PSFs generated with a median RMS WFE of 350 nm (1 radian) distributed over 20 and 100 modes.

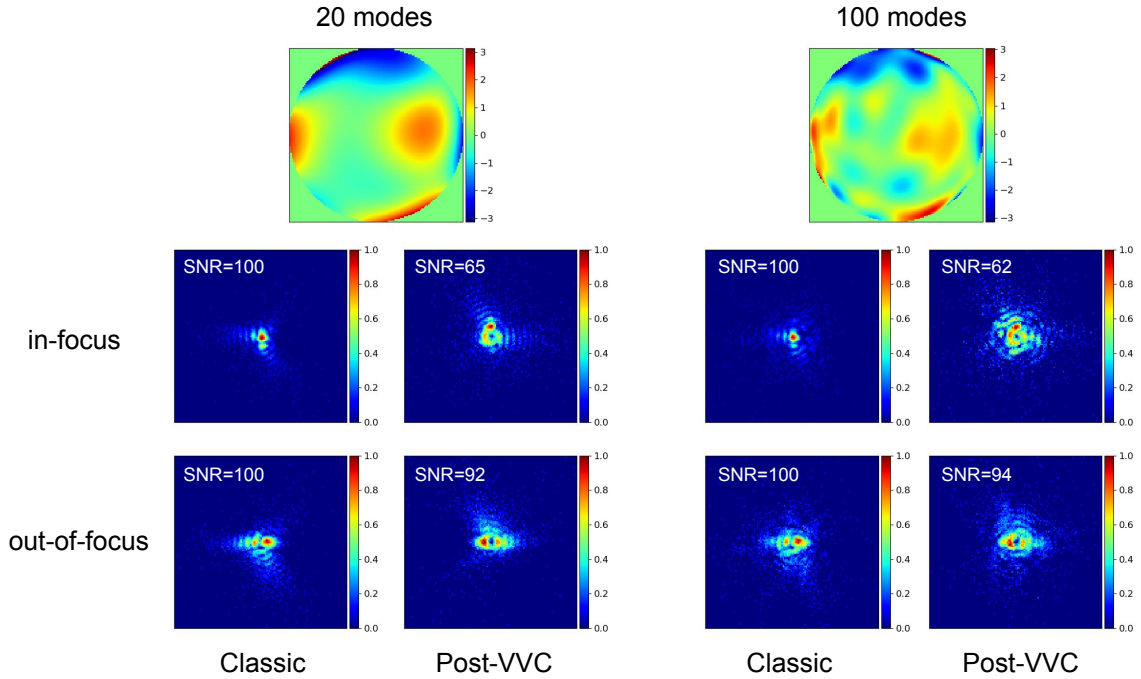


Figure 3.2 Example of phase maps and pre-processed PSFs with a median RMS WFE of 350 nm and an additional defocus of 550 nm. A fixed stellar flux of 10^4 photons (SNR = 100) is considered. *Top*: Phase maps with 20 (left) and 100 modes (right). *Middle*: Corresponding normalized in-focus PSFs. *Bottom*: Normalized out-of-focus PSFs. For each case, classical and post-VVC PSFs are shown.

3.2.2 Implementation of the Vector Vortex coronagraph

As described in Section 1.2.2, the Vector Vortex coronagraph (VVC, [Mawet et al., 2005](#)) is an optical device designed to suppress the direct starlight in order to reach higher contrasts in the star’s surroundings. It uses a transparent phase mask at the focal plane that diffracts the on-axis light outside the pupil, which is then blocked by a Lyot stop. As expressed in Equation 1.1, the VVC generates two conjugated phase ramps for each circular polarization of the incoming light. For a flat wavefront and a circular pupil, the VVC perfectly cancels the on-axis PSF. If the wavefront contains phase aberrations, however, some light will leak. Figure 3.3 illustrates the influence of the incoming level of wavefront error on the transmitted flux ratio. One can notice that regardless of the additional defocus, the transmitted flux tends to converge towards a similar fraction for very high input wavefront errors. In this aberration regime, some light also leaves the field of view, reducing the measured transmitted flux. Additionally, for a fixed integration time, the signal-to-noise

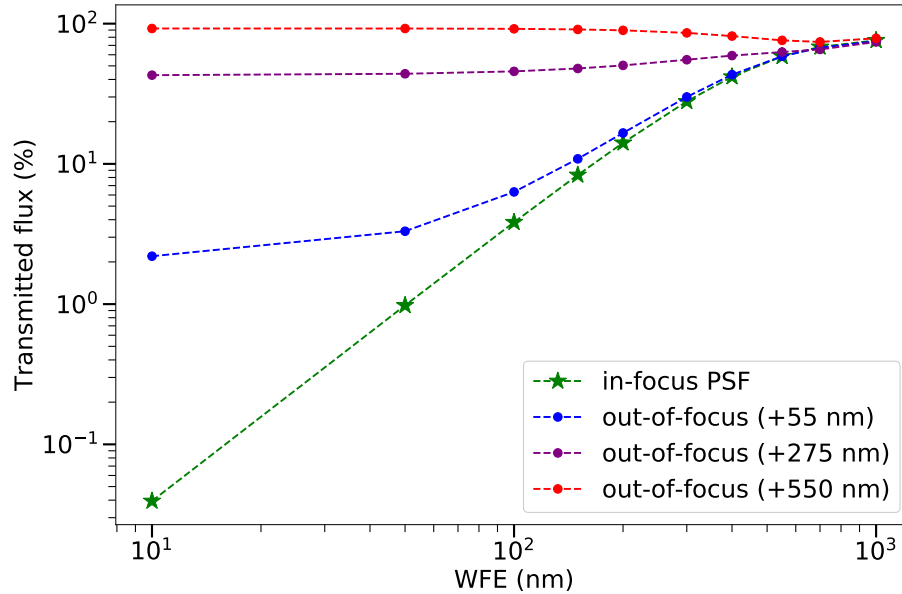


Figure 3.3 Transmitted flux after the vortex mask and the downstream Lyot stop as a function of the level of wavefront aberration, for different additional defocus. For each case, a single data sample generated with 20 Zernike modes is used.

ratio (SNR) will be lower than in classical imaging. This has an impact on focal-plane wavefront sensing compared to a case without coronagraph (see Section 3.4.3 for some comparison results).

The post-VVC PSFs are generated using the end-to-end HCI simulator HEEPS (Carlo-magno et al., 2020), which incorporates the optical propagation package PROPER (Krist, 2007). A topological charge $\ell_p = \pm 2$ is considered throughout the study. In order to simulate the PSFs, two images are generated with each opposite topological charge ℓ_p before being added together. We also perform some experiments with only one sign of the phase screw ($\ell_p = +2$) to investigate whether selecting one circular polarization could lift the sign ambiguity (Section 3.4.3). The diffracted light outside the geometric pupil is blocked by a Lyot stop, undersized to 98% of the pupil diameter. To generate the out-of-focus PSFs, the defocus is introduced in the entrance pupil plane. Examples of post-VVC PSF pairs can be found in Figure 3.2.

3.2.3 Data pre-processing

In the simulations, photon noise is considered, which means the SNR is defined as $\text{SNR} = \sqrt{N_{\text{ph}}}$, where N_{ph} is the total number of photons in the entrance pupil plane. An SNR of

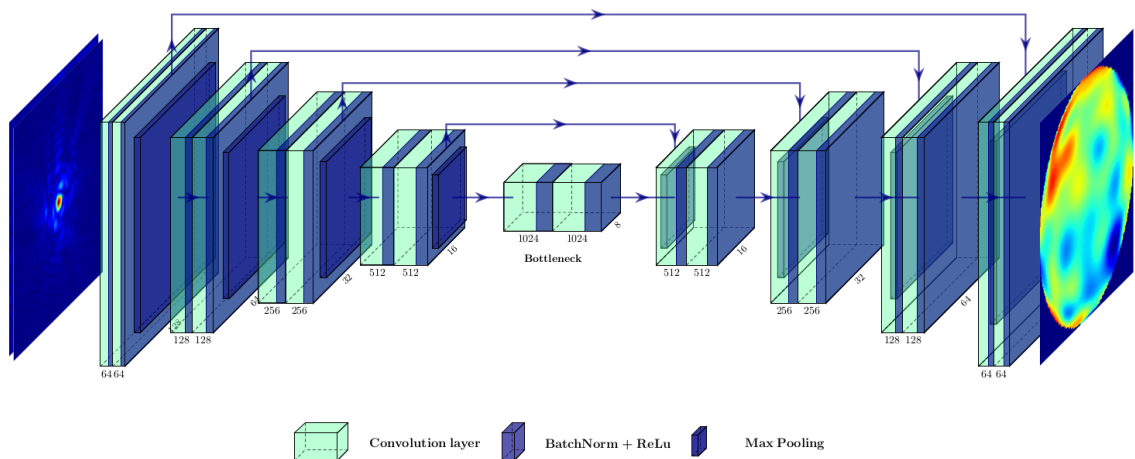


Figure 3.4 The U-Net architecture, adapted for a regression task. Credits: P-O Vanberg.

100 is chosen for most of the experiments. With the VVC, the flux suppression induced by the mask is also accounted for, reducing the SNR in the focal plane (Section 3.4.3).

A square-root stretching operation is then applied to the PSFs to assist the CNN in identifying the speckle distribution. This scaling method is preferred because it offers a good balance between revealing faint structures and avoiding saturation in bright areas, which can occur, for instance, with a log-stretching.

Finally, the PSFs are normalized with a min-max scaling to obtain fluxes in the range $[0,1]$, which ensures same-scale quantities are fed to the CNN. This has been shown to provide better stability and performance during the training of deep neural networks (Bishop, 1995).

3.3 Deep learning models

In this section, the convolutional neural networks and the mixture layer used in this chapter are first presented. The components allowing to train the models are then explained, in particular the loss functions and some key hyperparameters.

3.3.1 Convolutional Neural Networks

The goal is to build models to map phase aberrations y from PSFs x , i.e., to approximate a non-linear function f_θ parameterized by θ , such that $y \approx f_\theta(x)$. In order to do so, two deep neural network architectures are considered: ResNet-50 (He et al., 2015) and U-Net (Ron-

neberger et al., 2015). Two other architectures, VGG-16 (Simonyan & Zisserman, 2015) and Inception-v3 (Szegedy et al., 2015), were also tested but produced poorer results on the datasets and are thus not presented in this study.

ResNet

ResNet-50 is an architecture based on residual neural networks, which are described in Section 2.1.2. ResNet-50 (He et al., 2015) has been adapted to our regression task by modifying the first and last layers of the network. This ResNet-50 predicts the Zernike coefficients from the input PSFs and the initialized weights are pre-trained with the ImageNet dataset.

U-Net

U-Net was introduced by Ronneberger et al. (2015) for biomedical image segmentation. U-Net is composed of a downsampling path followed by an upsampling of the feature maps, producing outputs at the same dimension as the input. Upsampling the feature maps is performed with transposed convolutions. It also uses skip connections by concatenating feature maps from previous layers to retain local information during the forward pass. To apply U-Net for our task, the softmax layer at the end of the architecture was removed to perform a regression task instead of segmentation. As illustrated in Figure 3.4, phase maps are directly inferred from PSFs and U-Net is trained from scratch.

3.3.2 Mixture Density Layer

Mixture density networks (MDNs), introduced by Bishop (1994), predict weighted sums of probability distributions instead of single quantities, allowing to estimate uncertainties in the data. A mixture of Gaussian distributions is selected for the models. We only need to modify the last layer of ResNet-50 to produce three parameters per Gaussian k and Zernike mode z : the mean $\mu_{z,k}(x)$, the standard deviation $\sigma_{z,k}(x)$ and the weight $\alpha_{z,k}(x)$ given to each Gaussian. From these outputs, we reconstruct the likelihood of a coefficient y_z given an input x as

$$\begin{aligned}
 p(y_z|x) &= \sum_k \alpha_{z,k}(x) \mathcal{N}(y_z | \mu_{z,k}(x), \sigma_{z,k}^2(x)) \\
 &= \sum_{k=1}^{N_{\text{gauss}}} \frac{\alpha_{z,k}(x)}{\sqrt{2\pi} \sigma_{z,k}(x)} \exp\left(-\frac{(y_z - \mu_{z,k}(x))^2}{2 \sigma_{z,k}^2(x)}\right).
 \end{aligned} \tag{3.3}$$

For this study, two Gaussian distributions ($N_{\text{gauss}} = 2$) are used, which means that six targets for each Zernike mode are predicted instead of a single one. The weights are normalized such that $\sum_k \alpha_{z,k}(x) = 1$, with $\alpha_{z,k}(x) \in [0, 1]$. This architecture allows us to measure prediction uncertainties and assess the sign ambiguity observed with even Zernike modes. We expect to predict one Gaussian close to the true value of the Zernike coefficient and another one around the opposite value when there is sign degeneracy, i.e., only for even modes. Such a method could for instance be exploited to measure the capability of diversity techniques to lift the sign ambiguity. In the following sections, “ResNet-MDN” indicates the ResNet-50 architecture combined with the mixture density layer.

3.3.3 Training procedure

To train the models, a loss function needs to be defined. An optimizer will minimize this loss by successive updates of the neural network’s weights, allowing to fit the model to the data. In the standard case, i.e., without the mixture density layer, the chosen loss function is the root-mean-square error of the phase residuals

$$\mathcal{L}_{\text{cnn}}(\Phi, \widehat{\Phi}(x; \theta)) = \sqrt{\frac{1}{N} \sum_{i,j} (\Phi_{i,j} - \widehat{\Phi}_{i,j}(x; \theta))^2}, \quad (3.4)$$

where $\Phi_{i,j}$ and $\widehat{\Phi}_{i,j}(x; \theta)$ are the true and estimated phase respectively, θ represents the neural network weights, x corresponds to the input PSFs and N is the total number of pixels in each phase map. On the other hand, with ResNet-MDN architecture we want to maximize the likelihood $p(y_z|x)$ of the true Zernike coefficients y_z defined in Equation 3.3. This is the same as minimizing

$$\mathcal{L}_{\text{mdn}}(y_z, x; \theta) = -\log p_\theta(y_z|x). \quad (3.5)$$

A small offset $\epsilon = 0.005$ rad is also added to the predicted $\sigma_{z,k}(x)$ to improve the stability of the minimization. During training the loss is averaged over the elements of the current batch before being fed to the back-propagation algorithm.

A dataset composed of 10^5 PSF pairs is randomly split into training (90%) and validation (10%) sets. A batch of 64 data samples is fed to the neural network at each iteration and the Adam optimizer (Kingma & Ba, 2017) is used, which is a well-known flavor of stochastic gradient descent. A learning rate of 10^{-3} is set for the CNN while a learning rate of 10^{-5} is preferred for ResNet-MDN to avoid training instabilities, which has been observed with its loss function. The learning rate is then decreased by a factor of 2 each time the validation loss stagnates over 20 epochs, resulting in immediate drops of the loss,

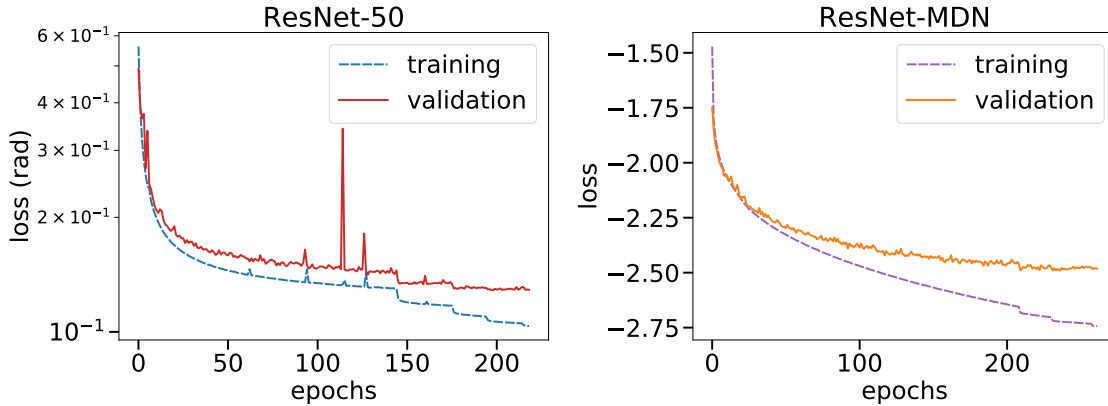


Figure 3.5 Illustration of training and validation losses obtained through the epochs, with phase diversity and on a dataset comprising 100 modes, a median input RMS WFE of 350 nm and classical imaging PSFs. *Left*: Using ResNet-50 (with weight decay). *Right*: Using ResNet-MDN.

as it can be noticed in Figure 3.5. Only the model defined at the epoch with the lowest validation loss is kept for evaluation.

Since ResNet-50 tends to be particularly prone to overfitting in our case, regularization is applied with this architecture using a technique called weight decay. It works by adding a penalty to the loss function, encouraging smaller weights. A factor of 10^{-7} was found to yield the best performance. More generally, deep neural networks trained with Adam and/or batch normalization have a specific response to weight decay, and multiple discussions can be found about this topic in the literature ([van Laarhoven, 2017](#); [Zhang et al., 2018](#); [Loshchilov & Hutter, 2019](#)).

3.4 Experiments and results

In this section, performance is assessed with both classical and coronagraphic imaging, and the way mixture density models can exhibit prediction uncertainties is also illustrated. To evaluate the models, test datasets are used, which are composed of 1000 samples generated the same way as the training and validation sets, except for a different random seed on the Zernike coefficients.

Number of modes	20		100	
	70	350	70	350
U-Net	8.7	13.6	20.2	56.9
ResNet	8.8	12.5	19.9	38.9
ResNet-MDN	9.2	15.0	20.7	71.8

Table 3.1 Comparison of RMS WFE on the phase residuals (in nm) averaged over 10^3 test samples. The models were trained on 10^5 PSF pairs with phase diversity and an SNR of 100.

3.4.1 Comparison between CNN architectures

An overview of the performance of the ResNet-50 and U-Net architectures is given with our approach. A more detailed analysis is done by [Orban de Xivry et al. \(2021\)](#).

Here the models are trained with datasets containing classical imaging PSFs (without the VVC). Models trained with and without phase diversity are also compared. Predictions on the test dataset are made, and the root-mean-square error (RMSE) on the phase residuals, defined the same way as the loss function metric in Equation 3.4, is computed. The performance is illustrated in Figure 3.6 for a median input RMSE of 350 nm (1 radian) distributed over 100 Zernike modes and with an SNR of 100. The histograms represent the RMSEs computed from each test sample. Using diversity clearly reduces the error on the phase residuals, due to having two inputs instead of a single one (the effective SNR increases by a factor $\sqrt{2}$), but most importantly because the predictions are especially improved on the even modes, since the sign ambiguity is overcome. ResNet-50 is able to learn 100 modes slightly better than U-Net but, as shown in Table 3.1, almost identical RMSEs are obtained between the two architectures when trained with 20 modes or an input WFE of 70 nm. In the following experiments, ResNet-50 and ResNet-MDN are exclusively used.

3.4.2 Predicted uncertainties with mixture density models

From mixture model predictions, the RMS WFE can be obtained from the weighted mean $\widehat{\mu}_z^* = \widehat{\alpha}_{z,1}\widehat{\mu}_{z,1} + \widehat{\alpha}_{z,2}\widehat{\mu}_{z,2}$, corresponding to the predicted Zernike coefficient for each mode z . The RMSE of the residual phase can then be computed with Equations 1.6 and 3.4. We observe a degradation in performance with the mixture density models in terms of RMSE compared to the standard CNN models (Table 3.1). This is especially true for 100 Zernike modes and an input WFE of 350 nm. Training on 5×10^5 images compensates for the

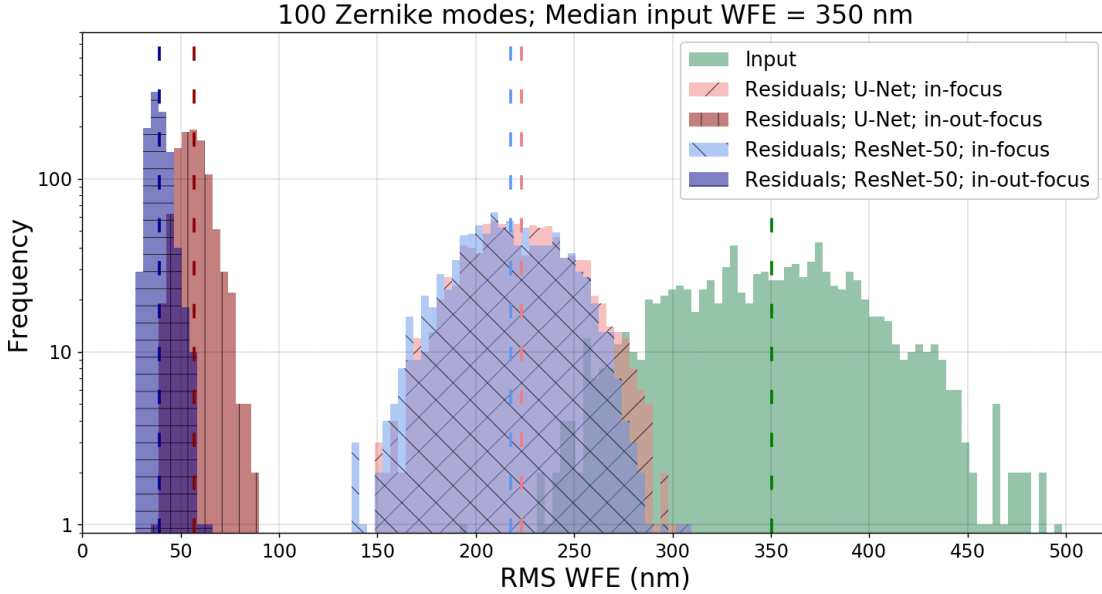


Figure 3.6 Empirical distributions of RMS WFE on the test dataset for models trained over 100 modes and with a median WFE of 350 nm (green). Performance with the ResNet-50 and U-Net architectures are compared, as are the cases with (in-out-focus; darker colors) and without (in-focus; lighter) phase diversity.

difference in performance, and the following results are thus obtained from models trained with this dataset size.

As illustrated with reconstructed distributions from the predicted parameters in Figure 3.7, ResNet-MDN is well capable of recognizing the sign ambiguity when fed only with in-focus PSFs, by estimating two Gaussian distributions centered near the true coefficient and its opposite value. This works particularly well for low order Zernike modes, while the estimations degrade for higher order modes: the predicted absolute weight difference $\Delta\hat{\alpha}_z = |\hat{\alpha}_{z,1} - \hat{\alpha}_{z,2}|$ starts to be equal to 1 for even modes from $z > 20$, which means only one Gaussian is predicted instead of two (Figure 3.8; left). In these cases the predicted $\hat{\mu}_z$ of the single Gaussian is approximately equal to zero, which is what the standard CNN predicts for even mode when facing the twin-image problem. More training data helps improve the predictions for high-order even modes. On the other hand, with phase diversity, the MDN predicts a single Gaussian for every Zernike mode as expected.

The metric we use here to measure the sign ambiguity is the entropy (Shannon, 1948)

$$H = - \sum_i p(y_i|x) \log_2 p(y_i|x), \quad (3.6)$$

where y_i are Zernike coefficient values (x-axis of Figure 3.7). The advantage of this metric

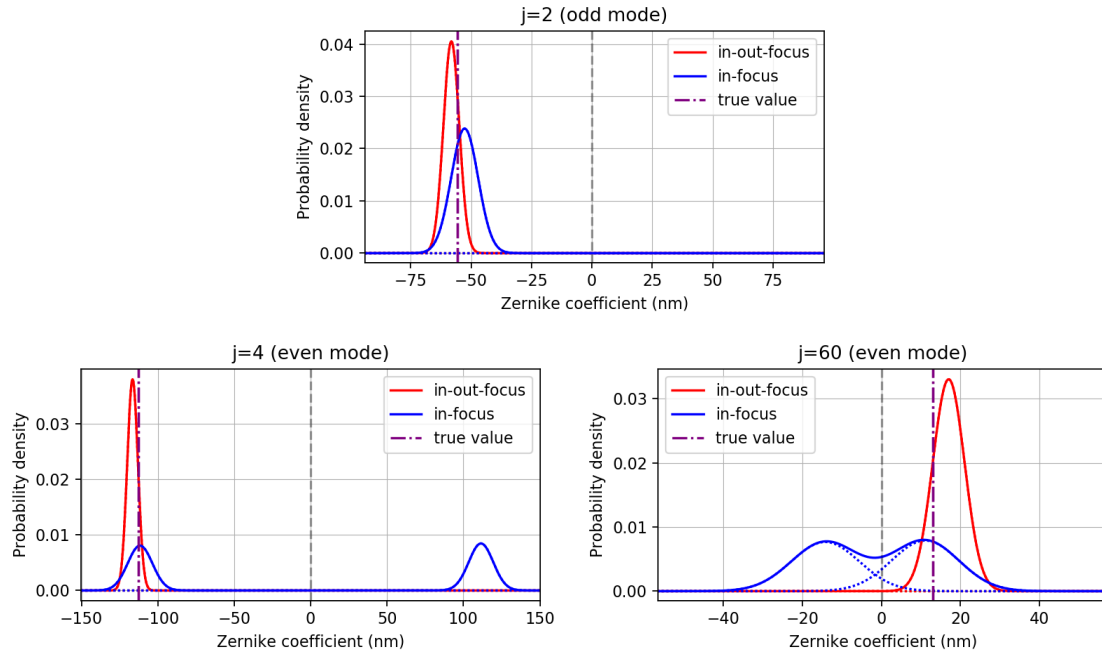


Figure 3.7 Examples of predicted distributions with a mixture model for an odd mode (tip; top), a low order even mode (defocus; bottom left), and a higher order even mode (bottom right).

is that it depends both on the number of Gaussians and their widths, i.e., it considers the total uncertainty in the prediction. The entropy computed for each Zernike mode is shown in Figure 3.8 (right). The entropy is consistently higher for even modes compared to odd modes without phase diversity, which showcases the twin-image problem. One can also notice that for the high-order modes for which ResNet-MDN is not able to predict two Gaussians (Figure 3.8; left), the entropy is still important because the predicted $\widehat{\sigma}_z$ is sufficiently large in these cases. Regarding odd Zernike modes, the entropy is always higher when trained only with in-focus PSFs since it lacks the additional out-of-focus PSFs that naturally improve the performance.

3.4.3 Influence of the Vector Vortex coronagraph

In this section the performance using post-coronagraphic PSFs is assessed, i.e., using a dataset generated with the VVC described in Section 3.2.2, comparing the residual errors obtained with models trained on classical imaging PSFs. In particular, three types of models are studied: trained using classical PSFs and a fixed SNR of 100, using post-VVC PSFs at the same SNR of 100 (for both in and out-of-focus PSFs), and using post-VVC

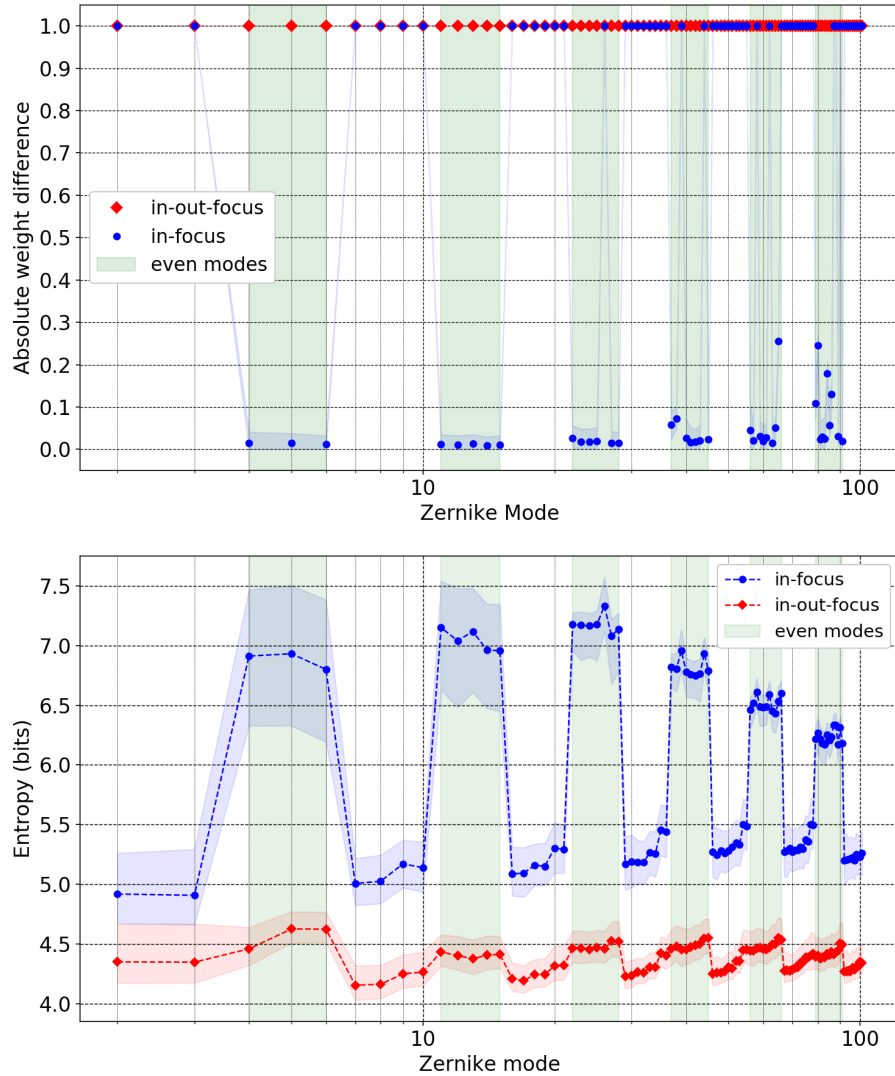


Figure 3.8 *Top*: Absolute difference of the predicted weights. A value of 1 indicates only a single Gaussian is predicted while 0 shows that two Gaussians with equal weights are obtained. *Bottom*: The entropy of the predicted Gaussian mixture for each Zernike mode. In both plots, the median values are represented, and the shaded regions extend to the 5–95th percentiles.

PSFs at a fixed stellar flux, resulting in lower SNRs in the focal plane because of the flux suppression caused by the vortex mask.

The results are displayed in Figure 3.9. At low input RMS WFE, the performance for each case reaches a plateau set by the photon noise limit (Orban de Xivry et al., 2021). On the other hand, at high RMS WFE, the performance degrades and this deviation from the photon noise limit relates to an additional model error. Models trained with post-VVC PSFs at a fixed stellar flux (blue curve) give higher residual RMS WFE for any input WFE which is essentially due to their lower SNR in the focal plane (see Figure 3.9; right). Finally, models trained on post-VVC PSFs at a fixed SNR of 100 (purple curve) give better performance in the low aberration regime (≤ 125 nm) compared to classical PSFs at the same SNR, but poorer performance at higher aberration levels. Since the vortex mask removes the Airy disk, we would expect the aberrations to be better revealed in the focal plane, potentially improving learning the mapping between the phase and PSFs. This could explain the gain in performance obtained with the VVC at low levels of aberrations. Still, the stronger degradation in performance with higher input RMS WFE is not fully understood. One hypothesis is that this could be due to the non-linearity of the VVC at high levels of aberrations, which would make the models more difficult to train. This trend has also been observed with models trained on other SNRs and levels of additional defocus. However, the differences in performance are relatively small and mostly within the error bars.

3.5 Conclusions

In this chapter, a deep learning-based method has been developed to predict wavefront aberrations typical of NCPAs from simulated PSFs, notably using the ResNet-50 architecture. The approach has been tested on both classical and post-coronagraphic PSFs with vector vortex phase masks. Models trained with both types of PSFs yield similar performance at the same level of photon noise, albeit with slight differences. Mixture density models have also been implemented to predict the ambiguity on the phase sign. It has proven to work well in our setup, even though some limitations regarding higher order Zernike modes have been observed. A more detailed study of the method applied to classical imaging can be found in Orban de Xivry et al. (2021). In particular, it discusses the robustness of the models, how the performance relate to the theoretical limit and a comparison is made with the iterative Gerchberg-Saxton algorithm. The initial experiments with post-VVC PSFs presented in this chapter represent a starting point to explore further developments. Specifically, the unique properties of vortex coronagraphs can be harnessed for focal-plane wavefront sensing. This is the focus of the next chapter.

3.5. CONCLUSIONS

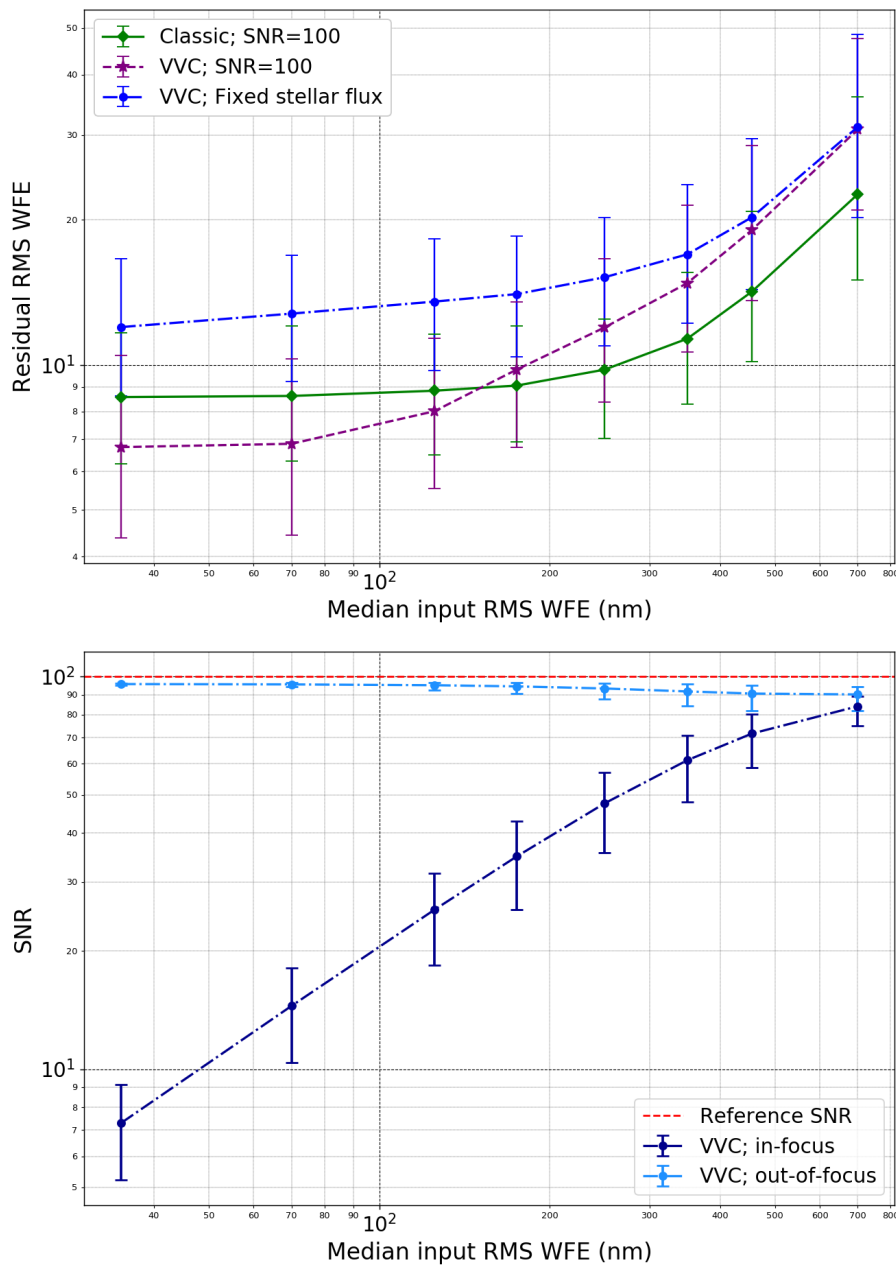


Figure 3.9 *Top*: Comparison of residual RMS WFE between models trained on post-VVC and classical imaging PSFs for different input WFE, each reported point corresponding to a specific model. *Bottom*: SNR of the in and out-of-focus PSFs for each model with the VVC, considering a fixed stellar flux (blue curve in the top plot). In both plots the median values are represented and the error bars correspond to the 5–95th percentiles.

Phase diversity with vortex masks

This chapter explored a novel phase diversity method that achieved a 100% science duty cycle with scalar and vector vortex phase masks. The properties of this types of coronagraphs and the effects they have on light polarization allows to solve the phase sign ambiguity and perform focal-plane wavefront sensing. Using a state-of-the-art CNN, phase aberrations are successfully inferred from simulated focal-plane images behind vortex coronagraphs. Most of this work appears in refereed journal paper A&A ([Quesnel et al., 2022a](#)). Additionally, Section 4.3.6 includes non-published experiments on closed-loop simulations and Appendix B presents the effect an asymmetric Lyot stop as an alternative approach for addressing sign ambiguities.

Contents

4.1	Vortex phase diversity concept	69
4.2	Deep learning approach	71
4.2.1	Data generation	71
4.2.2	Model architecture	72
4.2.3	Training procedure	73
4.3	Results and analysis	75
4.3.1	Phase sign determination	76
4.3.2	Comparison to classical phase diversity	76
4.3.3	Companion subtraction with dual-polarization	78
4.3.4	Theoretical limit	80
4.3.5	Model robustness	80
4.3.6	Closed-loop experiments	84
4.4	Conclusions	87

4.1 Vortex phase diversity concept

As described in the previous chapters, the phase sign ambiguity is typically solved with phase diversity using, for instance, an additional defocused PSF (Gonsalves, 1982), or an asymmetric pupil mask (Martinache, 2013). This, however, reduces the science duty cycle because some observing time, and/or part of the science beam, has to be dedicated to wavefront measurements exclusively. Based on the properties of the vector vortex coronagraph (VVC, Mawet et al., 2005), a Nijboer-Zernike phase retrieval approach tailored to the post-VVC PSF was formulated in Riaud et al. (2012a,b). They proposed to split the two circular polarization states to exploit the phase diversity introduced by the two opposite topological charges associated with the VVC.

An illustration of the sign ambiguity is shown in Figure 4.1, where two phase maps with opposite signs for their even Zernike modes are generated. After propagation through a VVC, the in-focus PSFs are the same in both cases (Figure 4.1d and f), showcasing the twin-image problem. The out-of-focus PSFs, however, are different (Figure 4.1e and g) because the added defocus has the same sign in both cases, which allows the ambiguity to be lifted. Now, if the two orthogonal circular polarization states are split downstream of the VVC to separate the conjugated phase ramps ($-\ell_p$ and $+\ell_p$), or if the case of the scalar vortex coronagraph (SVC) is considered, the in-focus PSFs are not identical anymore (Figure 4.1h and j, or Figure 4.1i and k). The resulting PSFs are actually switched between the two circular polarization states. This indicates that the sign ambiguity can potentially be lifted when using either the two PSFs obtained from the separate circular polarization states or the single PSF behind the SVC independently of the polarization state. This illustrates the fact the VC provides an azimuthal phase diversity, which can be used instead of the radial phase diversity provided by an additional defocus (Riaud et al., 2012a). In the case of the SVC, the sign ambiguity would then be lifted similarly to using only an out-of-focus PSF in classical phase diversity (e.g., Lamb et al., 2021).

This chapter revisits the problem of phase retrieval behind a vortex coronagraph using deep learning techniques. Unlike an analytical approach, which could show limitations regarding its formulation, deep learning models can be trained regardless of the instruments and observing conditions. In Section 4.2, the deep learning approach is presented, based on convolutional neural networks (CNNs), which have the advantage of being flexible and easy to implement, and they have already been shown to be capable of reaching fundamental noise limits in previous works (Quesnel et al., 2020; Orban de Xivry et al., 2021). In Section 4.3, quantitative results on simulated data are provided. The performance of the vortex phase diversity method is compared to a classical approach and the robustness of the models is assessed, notably in the presence of representative atmospheric turbulence




Figure 4.1 Comparison of simulated PSFs for two conjugated phase maps ϕ (left) and ϕ' (right): for ϕ' , we set opposite Zernike coefficients to those of ϕ only for the even modes, with a total of 18 modes starting from defocus. (a): Entrance annular pupil. (b, c): The conjugated phase maps. (d, e): In-focus and out-of-focus PSFs obtained from propagating (b) with both polarization states together. (f, g): The same as (d, e) but using (c) for propagation instead. (h, i): In-focus PSFs obtained from (b) with $-\ell_p$ and $+\ell_p$ used separately. (j, k): The same as (h, i) but using (c) for propagation instead.

residuals. The performance of the models in a closed-loop process is finally highlighted in Section 4.3.6, showcasing the impact it has on the contrast. The potential of using an asymmetric Lyot stop and a single in-focus PSF with the VVC is also investigated in Appendix B.

4.2 Deep learning approach

4.2.1 Data generation

In the simulations, an annular entrance pupil is considered with a diameter of 8 m and a central obstruction of 30%. An observed bandwidth of 0.2 μm was defined around 2.2 μm (K band), by simulating a total of five wavelengths. A pixel scale of $0.25 \lambda/D/\text{pix}$ was set with a detector containing 64×64 pixels, giving a field-of-view of $16 \lambda/D$. The most relevant simulation parameters are listed in Table 4.1.

As in Chapter 3, the phase aberrations are produced using annular Zernike polynomials (see Equation 1.6). The generated datasets are composed of 18 or 88 Zernike modes, up to the fifth and 12th radial orders, respectively, excluding the piston, tip, and tilt modes. The set of Zernike coefficients for each sample was first randomly generated within the range $[-1, 1]$ before each coefficient was divided by its corresponding radial order to approximate a $1/f^2$ power spectral density profile, typically encountered with good quality optics (Dohlen et al., 2011). Low and high aberration levels, represented by wavefront error (WFE) distributions centered at a 70 and 350 nm root-mean-square (RMS), respectively, are considered by normalizing the Zernike coefficients accordingly (see Equation 3.1). An example of such a distribution can be seen in Figure 3.1 (right). For classical phase diversity, which is used as a reference to compare the performance of various setups, the additional defocus was set to $\lambda/5$, that is, 440 nm RMS. In our case, this amount of diversity is close to the optimal value measured in terms of phase retrieval performance, and other defocus values may provide better results in different experimental configurations. The defocus was added in the entrance pupil plane as if done by the deformable mirror of an adaptive optics system, which means that the resulting defocused PSFs contain more flux than the in-focus PSFs as the coronagraphic performance of the VC is degraded.

To increase the representativeness of the simulations and to test the robustness of the approach, atmospheric turbulence residuals are added to the phase maps. A state-of-the-art extreme adaptive optics (AO) was simulated using the COMPASS library (Ferreira et al., 2018), assuming a loop frequency of 3.5 kHz, 2-frame delay, a 50×50 deformable mirror (i.e., 2040 modes/valid actuators), and a pyramid sensor with $5 \lambda/D$ of modulation (without noise). This has yielded a Strehl ratio of about 98% at 2.2 μm , corresponding to a

Parameter	Value
Central obstruction	30 %
Topological charge	2
Pixel scale	$0.25 \lambda/D$ / px
Field-of-view	$16 \lambda/D$

Table 4.1 Data generation parameters.

WFE of about 50 nm RMS. The AO residuals are sampled at 10 Hz and a sequence of ten consecutive phase screens is used by summing up the corresponding PSFs. A 1-s exposure in the presence of a given amount of static NCPAs is therefore simulated. The results with data containing these AO residuals are shown in Section 4.3.5.

To simulate a PSF obtained behind a VVC, two propagations are performed, one with $+\ell_p$ and the other with $-\ell_p$, to consider each circular polarization state. The downstream Lyot stop blocked 2% of the outer pupil area (but the central obstruction was not oversized). The resulting PSFs were then either summed up to reproduce the nonpolarized case, or they were kept separate to consider the dual-polarization case. To simulate the SVC, only one such PSF was taken. The optical propagation was handled by the HEEPS package¹ (Carlomagno et al., 2020), which makes use of PROPER (Krist, 2007). Examples of generated phase maps and PSFs can be found in Figure 4.1. Photon noise is then added to the PSFs, so that the signal-to-noise ratio (SNR) was defined as $\text{SNR} = \sqrt{N_{\text{ph}}}$, where N_{ph} is the number of photons. Figure 4.2 illustrates PSFs at increasing amounts of SNR. A square-root stretching operation was applied to the PSFs to help the CNN identify the speckle patterns. Finally, the PSFs are normalized with a min-max scaling to obtain flux in the range [0,1], which ensured the CNN was fed with same-scale quantities.

4.2.2 Model architecture

As in the previous chapter, deep neural network models are trained to approximate a non-linear function that maps Zernike coefficients from PSFs. In this chapter, the state-of-the-art deep CNN called EfficientNet (Tan & Le, 2019) is implemented. EfficientNet stands out from other deep CNNs by using a new scaling technique: all dimensions of the CNN (depth, width, and resolution) are scaled by the same compound coefficient, as shown in Figure 4.3a. This allows us to leverage the correlation that exists between the different network dimensions, not only to increase accuracy but also the efficiency of the models in terms of FLOPS (Floating Point Operations Per Second). With the compound scaling

¹<https://github.com/vortex-exoplanet/HEEPS>


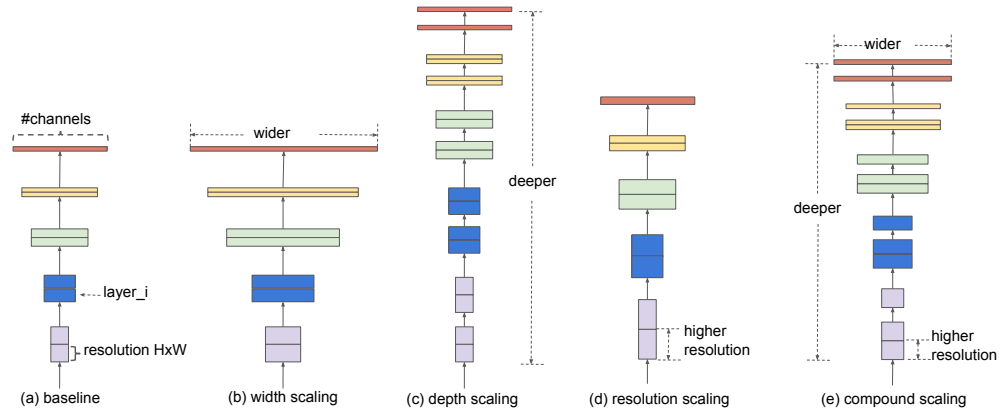


Figure 4.2 Examples of PSFs at different SNR levels (defined in the entrance pupil plane) for $+\ell_p$. The resulting SNR in the detector place is reduced due to the extinction factor introduced by the coronagraph and by the beam splitting between the two polarization channels. The level of NCPA is equal to 70 nm RMS distributed over 18 modes (*top*) and 350 nm RMS over 88 modes (*bottom*). AO residuals are also present: each PSF is the result of combining ten PSFs, with each containing a different AO residual phase screen.


method, the architecture is therefore optimized to the dimensions of the input. The parameters are inferred from the original model or baseline EfficientNet-B0, and there are thus several models available (from B0 to B7). The performance of EfficientNet on a standard dataset is shown in Figure 4.3b. Among the different models available, EfficientNet-B4 is chosen because it offers best trade-off between model performance and runtime for our task. EfficientNet-B4 has a total of 1.9×10^7 parameters and 4.2×10^9 FLOPS. It has about the same number of parameters as the ResNet-50 architecture, which was used in Quesnel et al. (2020) and Orban de Xivry et al. (2021).

4.2.3 Training procedure

General parameters Similarly to Chapter 3, a dataset composed of 10^5 PSFs (or PSF pairs for the cases with two input channels) was randomly split into training (90%) and validation (10%) sets for each training. Each sample also contains the true NCPA phase maps as labels, while the AO phase screens are never given. Batches composed of 64 data samples were then consecutively fed to the neural network. The loss function is defined as the root-mean-square error (RMSE) of the phase residuals. Weight updates based on the loss were handled by the Adam optimizer (Kingma & Ba, 2017). Pre-trained models on ImageNet were used to initialize the weights. The training of the model was stopped if no



(a) The scaling method of EfficientNet (d).



(b) Performance of EfficientNet on ImageNet compared to standard CNNs.

Figure 4.3 The EfficientNet principle and its performance compared to other deep CNNs. From Tan & Le (2019).

Designation	Charge	Defocus	Inputs
VVC in-focus	$\pm\ell_p$	no	1
VVC [in, out]-focus	$\pm\ell_p$	yes	2
VVC dual-polar	$[\ell_p, -\ell_p]$	no	2
SVC	$+\ell_p$	no	1
no vortex [in, out]-focus	0	yes	2
VVC dual-polar; diff PSFs	$[\ell_p, -\ell_p]$	no	1

Table 4.2 Configurations considered for phase retrieval.

improvement of the validation loss was observed over 25 epochs. This results in training procedures lasting between 50 and 250 epochs.

The use of weight decay To improve the performance, a penalty on the loss (“weight decay”) of 10^{-7} is set for the low aberration regime and 10^{-6} for the higher aberration regime. An initial learning rate of 10^{-3} is also set, which was decreased by a factor of two as soon as the validation loss reached a plateau over 15 epochs. This results in sudden loss drops, allowing the performance to be greatly improved.

4.3 Results and analysis

Here we compare the capacity of different configurations to lift the sign ambiguity as well as their performance. The designation of these configurations, together with some of their parameters, can be found in Table 4.2: we consider the cases of the VVC with or without classical phase diversity (“VVC [in, out]-focus” and “VVC in-focus”, respectively), which are compared to the new approaches presented in this chapter (“VVC dual-polar” and “SVC”). The non-coronagraphic case (“no vortex [in, out]-focus”) is evaluated as well. The possibility of working with differential PSFs obtained by subtracting the separate circular polarization states (“VVC dual-polar; diff PSFs”) is also investigated. In the last part of this section, atmospheric turbulence residuals are added and the robustness of the models regarding variations in the SNR levels, input wavefront errors, and Zernike polynomial orders are assessed. All models are evaluated using 1000 test samples.

4.3.1 Phase sign determination

To determine whether the models predict the correct sign, we looked at the performance per Zernike mode. The metric used is the RMSE per mode:

$$\sigma_z = \sqrt{\frac{1}{N_{\text{test}}} \sum_i^{N_{\text{test}}} (\hat{c}_i - c_i)^2}, \quad (4.1)$$

where N_{test} is the number of test samples, while \hat{c} and c are the estimated and true Zernike coefficients, respectively.

In Figure 4.4 the performance per mode between four cases for two different aberration contents are compared. A network using only in-focus PSFs in the non-polarized case with the VVC yields no correction for even Zernike modes because the model tends to predict zero for the coefficients facing the ambiguity (due to the l_2 -norm training loss). For odd modes, the model is able to provide some correction, even though its quality is limited by the loss function, which does not discriminate between even and odd modes. Adding defocused PSFs as input solves the problem as expected (Quesnel et al., 2020). In the dual-polarization case, a network using either one or both circular polarization states separately as input (SVC and VVC, respectively) also yields good performance for even modes as well as for odd modes. This indicates that the sign ambiguity is properly lifted with these two approaches.

It is noteworthy that the performance marginally depends on the Zernike mode: the error tends to increase for larger angular azimuthal orders at a given radial order. The interpretation is that since the phase information is of higher spatial frequency and located closer to the edge of the pupil in these cases, it is more difficult for the CNN model to identify those features.

4.3.2 Comparison to classical phase diversity

We now compare our method to the classical phase diversity approach in terms of overall phase retrieval performance. The RMS WFE on the phase residuals is used as a metric, and it is defined for each test sample as:

$$\sigma_\phi = \sqrt{\frac{1}{N_{\text{pix}}} \sum_i^{N_{\text{pix}}} (\hat{\phi}_i - \phi_i)^2}, \quad (4.2)$$

where N_{pix} is the number of pixels, while $\hat{\phi}$ and ϕ are the estimated and true pupil phases, respectively.

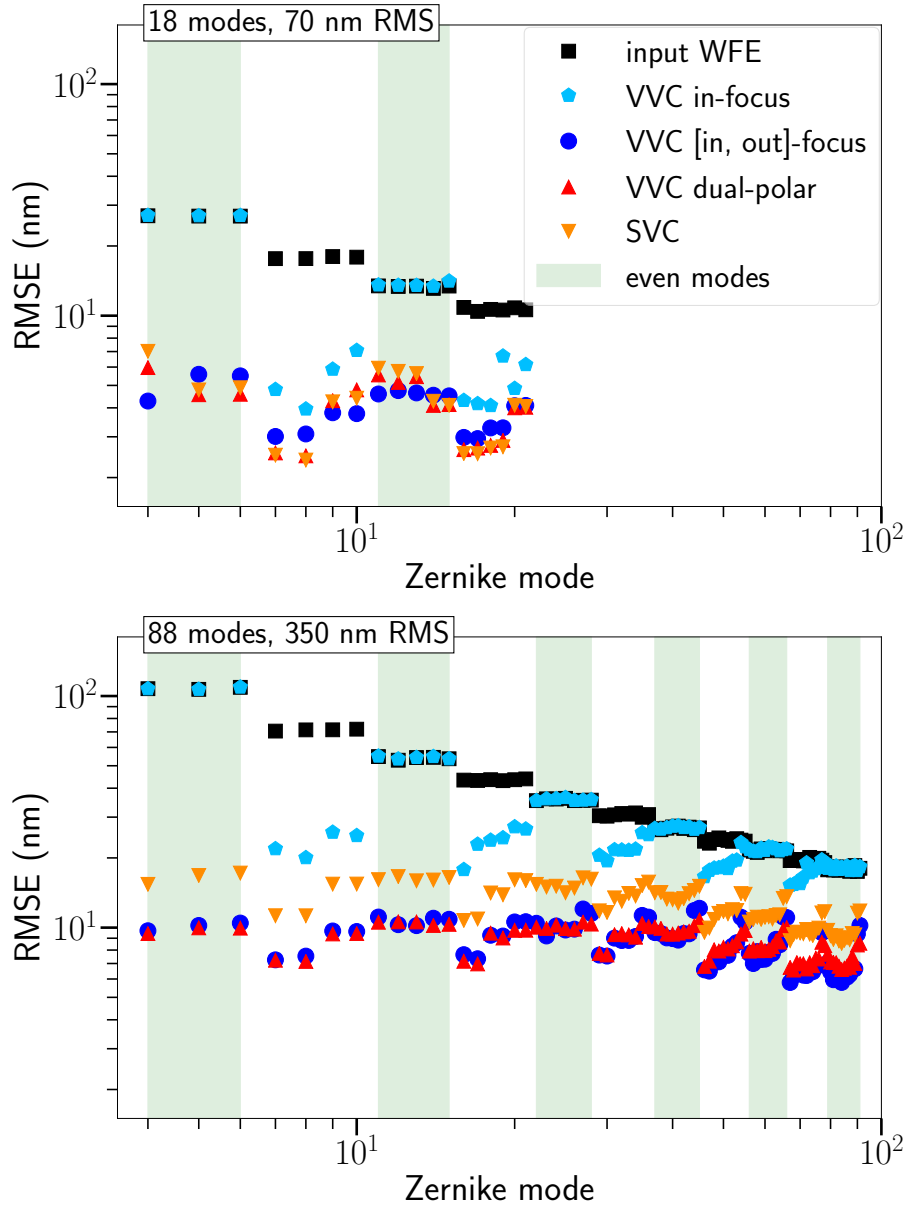


Figure 4.4 RMSE per Zernike mode, following the Noll convention, starting from the defocus mode. Four cases were compared (see Table 4.2 for notations), using a single in-focus post-VVC PSF without splitting the polarization states (cyan), two post-VVC PSFs with additional defocus (dark blue), the two post-VVC PSFs associated with each polarization state (red), and a single PSF after the SVC (orange). The RMSE of the input phase maps is represented in black and the even modes are indicated by the green areas. *Top*: Input WFE of 70 nm RMS distributed over 18 modes. *Bottom*: 350 nm over 88 modes. In both examples, the SNR in the entrance pupil plane is equal to 100.

In our simulations, the fact that the vortex coronagraphs block out most of the starlight is considered and that for a given stellar magnitude, the resulting flux in the detector plane is reduced. The flux is also equally split between each PSF for all the cases with two channels, while for the configurations with a single one, the PSF receives the total flux remaining after the vortex mask. The performance of the trained models at different SNR levels defined in the entrance pupil plane is shown in Figure 4.5. In our case, SNRs between 10^1 and 3×10^3 correspond to stars of apparent magnitudes in the range from 18.6 to 6.2.² For a median input WFE of 70 nm with 18 modes (Figure 4.5, top), the simulated performance is almost identical for the classical, SVC, and VVC dual-polarization approaches, even though the additional defocus increases the overall SNR at the focal plane for the classical method. For a median input WFE of 350 nm with 88 modes (Figure 4.5, bottom), the phase residuals are distinctly higher for all the configurations, and a plateau is reached for SNRs above 1000. We can especially notice that the sole PSF behind the SVC somewhat limits the performance in this case. The main hypothesis for this discrepancy is that, in a high aberration regime, the effects of the nonlinear nature of the problem are greater. The extra information given by having two input channels is therefore favorable and makes the models easier to train. In general, it is more difficult to train datasets containing strong aberrations, and this can typically be improved by using more data, more complex architectures, and/or stronger weight decay.

4.3.3 Companion subtraction with dual-polarization

The possible presence of planetary companions in the detected images is also considered. This additional, off-axis source of light is largely unaffected by the vortex phase ramp and therefore adds the same signature in both circular polarization states. This additional light source may bias the phase retrieval process, and lead to unwanted planetary signal subtraction. A possible workaround is to subtract one polarization image from the other, in an attempt to remove the signature of any off-axis light source. The phase retrieval capabilities using the difference between both polarization states after the VVC were thus assessed. The results are shown in Figure 4.5 and are compared with the other configurations. We only obtain a marginal increase in the error at high SNRs, especially in the low aberration regime, which can be explained by the loss of information produced by subtracting one PSF from the other.

²with an integration time of 1s, a transmission and quantum efficiency equal to 50%, a telescope diameter of 8 m, and a filter bandwidth of 50 nm.

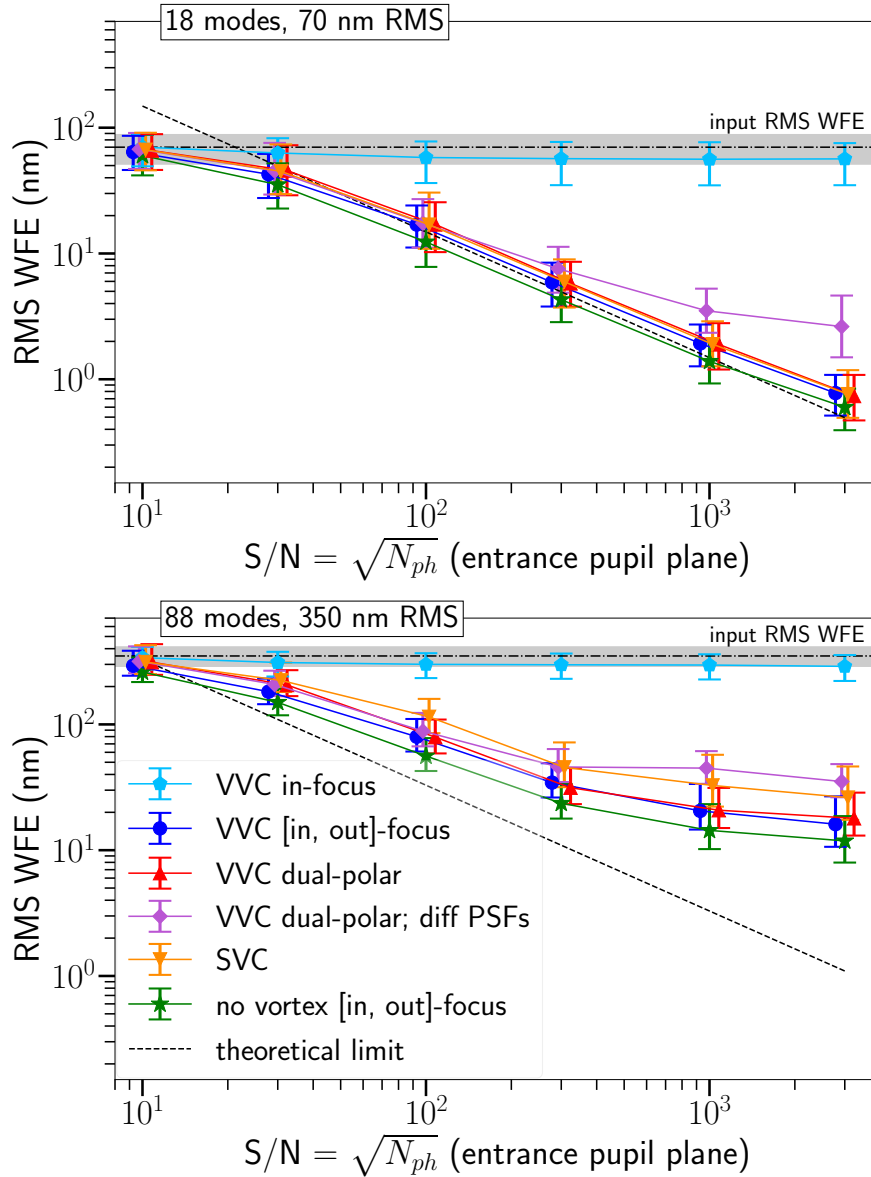


Figure 4.5 Performance in terms of RMS WFE on the phase residuals at different SNR levels. Each point corresponds to a model trained and evaluated on the indicated SNR (six SNRs are considered, and slight horizontal shifts were applied to be able to discern each point). The same colors in Figure 4.4 are used, with the addition of the performance with classical imaging (green), using differential post-VVC PSFs (violet), as well as the theoretical limit (black dashed line). The median values are represented and the error bars correspond to the 2–98th percentiles. The indicated SNRs are defined in the entrance pupil plane, accounting for flux suppression by the vortex mask. *Top*: Input WFE of 70 nm RMS distributed over 18 modes. *Bottom*: 350 nm over 88 modes.

4.3.4 Theoretical limit

The performance of the various configurations is finally compared to the theoretical limit in Figure 4.5. This limit is discussed in [Orban de Xivry et al. \(2021\)](#) for non-coronagraphic imaging. For both the non-coronagraphic and vortex imaging cases, the residual errors reach the fundamental limit in the low aberration regime (Figure 4.5, top). In a higher aberration regime, the performance does not reach the fundamental limit, and the gap increases toward higher SNRs (Figure 4.5, bottom). This could potentially be improved with more robust training, for instance by increasing the number of training data samples, as demonstrated by [Orban de Xivry et al. \(2021\)](#). It is worth noting that the residual errors are constrained by the WFE distribution in the data toward lower SNRs, while the theoretical limit is independent of the input WFE distribution and continues to increase for lower SNRs, thus yielding residual WFE below the limit.

4.3.5 Model robustness

To test how the method handles more realistic ground-based observations, atmospheric turbulence residuals are added in addition to the NCPAs, as described in Section 4.2.1. This represents an additional source of noise since the AO residuals are not included in the labels for training. Examples of input PSFs at the different flux levels can be found in Figure 4.2. The performance now starts to reach a plateau of a few nm RMS in the low aberration regime at high SNRs (Figure 4.6, top), due to the presence of these atmospheric turbulence residuals. In the high NCPA regime (Figure 4.6, bottom), the AO residuals however become negligible, and the performance is almost identical to the case without turbulence (Figure 4.5, right).

We finally study the robustness of the models regarding a variation in the data during evaluation. First, we may encounter different flux levels than those considered during training. Figure 4.7 illustrates how models in the VVC dual-polar configuration trained on data containing 70 nm RMS behave in such conditions. Whether the training SNR is low or high, models only show good robustness to other flux levels within a limited range, outside which the performance is strongly degraded. If a more robust model is required, it is also possible to train with various flux levels. We investigated this by using a training dataset covering the entire test SNR range, without increasing its size. The median performance is much more consistent at every SNR; although, the variation in the residual error between samples is greater, and a small degradation can naturally be seen compared to using identical training and testing SNR (as shown in Figure 4.6).

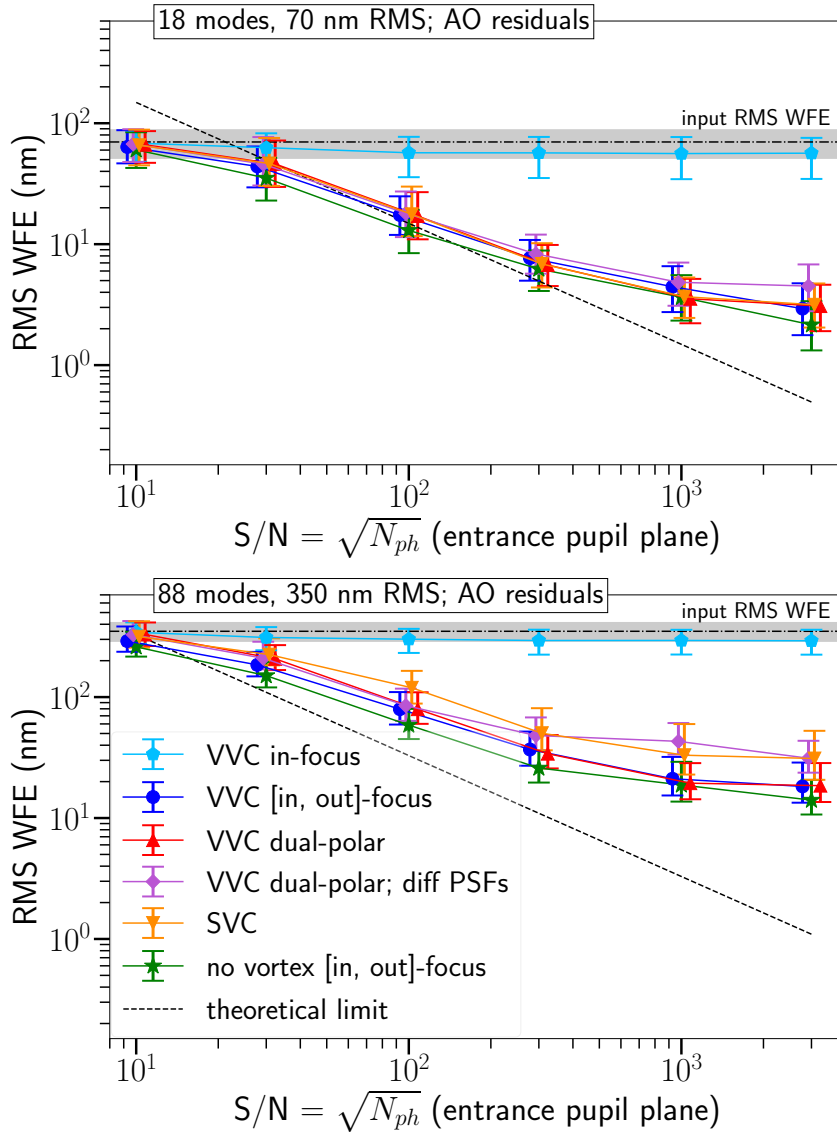


Figure 4.6 Phase prediction errors at different SNR levels, presented the same way as in Figure 4.5, but this time also including atmospheric turbulence residuals in the PSFs during both training and testing.




Figure 4.7 Performance with altered SNR levels during evaluation for three models trained on data with a median RMS WFE of 70 nm over 18 modes, with an SNR of 30 (purple), 1000 (blue), and with SNRs uniformly distributed over the entire SNR range (green). Each point is obtained from a testing batch composed of 1000 samples (the median value together with the 2–98th percentiles are shown).

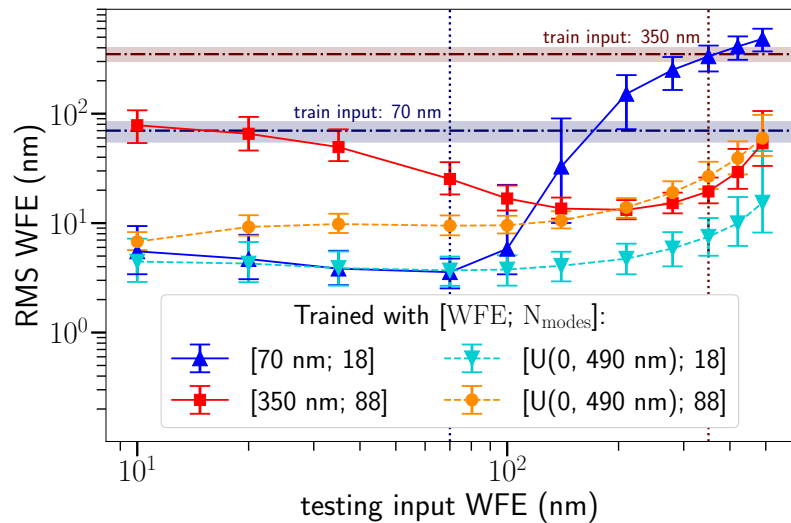


Figure 4.8 Performance with different input WFE levels defined during evaluation for models trained on data with a median RMS WFE of 70 nm over 18 modes (blue), and 350 nm RMS over 88 modes (red). Models were also trained on data following a uniform distribution covering the whole input WFE range, using both spatial frequency regimes (cyan and orange). The SNR is 1000 and each training dataset contains 10^5 samples.

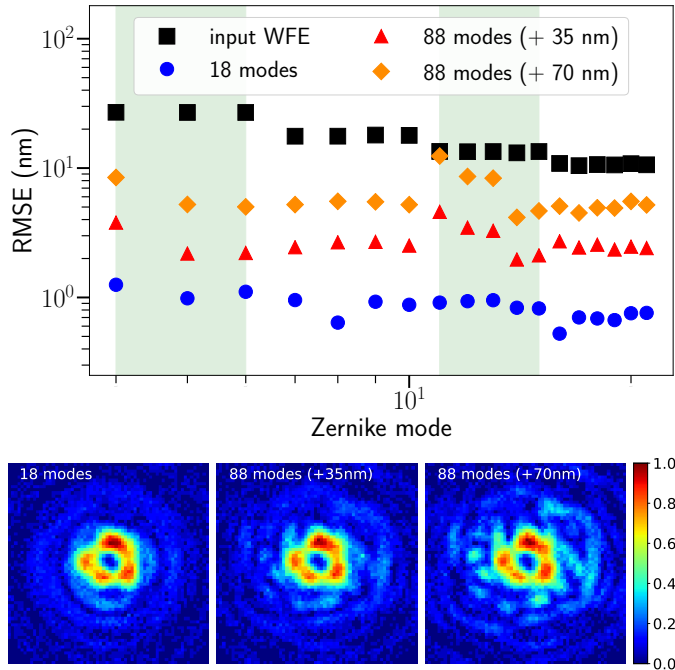


Figure 4.9 Robustness on higher-order aberrations. *Top*: Performance per Zernike mode on test data following the distribution of the training data (red), adding, to the test data, 70 modes containing 35 nm RMS (purple) and 70 nm RMS (blue) of NCPAs. *Bottom*: Example of post-VVC PSFs for each case ($+\ell_p$).

The change in performance when evaluating the model outside the input WFE training range is also studied. Figure 4.8 shows the robustness of models trained on the two aberration regimes studied in this chapter. Data containing more aberrations rapidly deteriorate the reconstruction. The models perform better when evaluated at lower aberration levels, but they have limited performance when trained in the high aberration regime. To overcome these limitations, two models are trained over the entire test WFE range for each of the Zernike mode contents considered in the chapter. Such models show excellent robustness, with minimal degradation compared to models with identical training and testing WFE distributions. This suggests that these models could be robustly used in closed-loop operations, even with the aberration level decreasing with time. Regarding the varying spatial power spectral density of the wavefront, the residuals are generally constant along the Zernike modes, as seen in Figure 4.4. When giving the reconstructed PSFs as input to the same trained model, we have observed that most residual RMS WFE stay below 10 nm for a model trained on 70 nm RMS as input and an SNR of 1000. Some closed-loop experiments will be presented in Section 4.3.6.

It can also be expected to have observations containing higher-order NCPAs (in addition to the changing atmospheric residuals) than considered during training. For a model trained on 18 modes at 70 nm RMS (SNR = 1000), 70 higher-order Zernike modes are added in the test data. In Figure 4.9, we observe a moderate degradation for the 18 modes when increasing the wavefront error contained in these additional modes, because the central PSF signature is mostly preserved.

4.3.6 Closed-loop experiments

In order to demonstrate the capabilities of the models to perform iteratively, closed-loop tests are performed. This is achieved by successively simulating the PSF \mathcal{I} with the generator g at step t using the predicted and true phase maps obtained at the previous step ($\widehat{\Phi}_{t-1}$ and Φ_{t-1}), and injecting \mathcal{I} to the model f for the next predictions $\widehat{\Phi}_t$:

$$\begin{aligned} \widehat{\Phi}_t &= f(\mathcal{I}_t) \\ \text{with } \mathcal{I}_t &= g(\widehat{\Phi}_{t-1} - \Phi_{t-1}). \end{aligned} \quad (4.3)$$

Results are shown in Figure 4.10 for ten test samples taken from the 70, 350 and [0, 490] nm RMS datasets, through ten iterations. Except for the uniformly distributed dataset, the best performance is obtained for the first iteration, and the following iterations show nonetheless restrained degradation, with an RMS WFE that converges to 60 nm and 5 nm RMS in the high and low aberration regimes, respectively. The WFE distributions of these two datasets approximate a normal distribution, which is not ideal for closed-loop application, since lower WFEs quickly become predominant. This is why a uniform or a logarithmic distribution are more relevant in this case. Results using the uniformly distributed dataset show much better closed-loop stability, with residuals WFE decreasing over the few first iterations to converge to 10 nm RMS on average.

Corrected PSFs for one sample taken from each of the three datasets are shown in Figure 4.11. With low aberrations, the PSF after 10 iterations remains very close to the diffraction limit. In the high aberration regime, some degradation can be observed between the first and the 10th iterations, although it remains within a reasonable range. The improvement of wavefront quality is quantified with the raw contrast at each separation from the center of the image, which is defined as

$$C_r = \frac{\langle \mathcal{I}_r(\theta) \rangle_\theta}{\mathcal{I}_0^{Airy}}, \quad (4.4)$$

where $\langle \mathcal{I}_r(\theta) \rangle_\theta$ is the average flux, at a specific distance r over the angular position θ

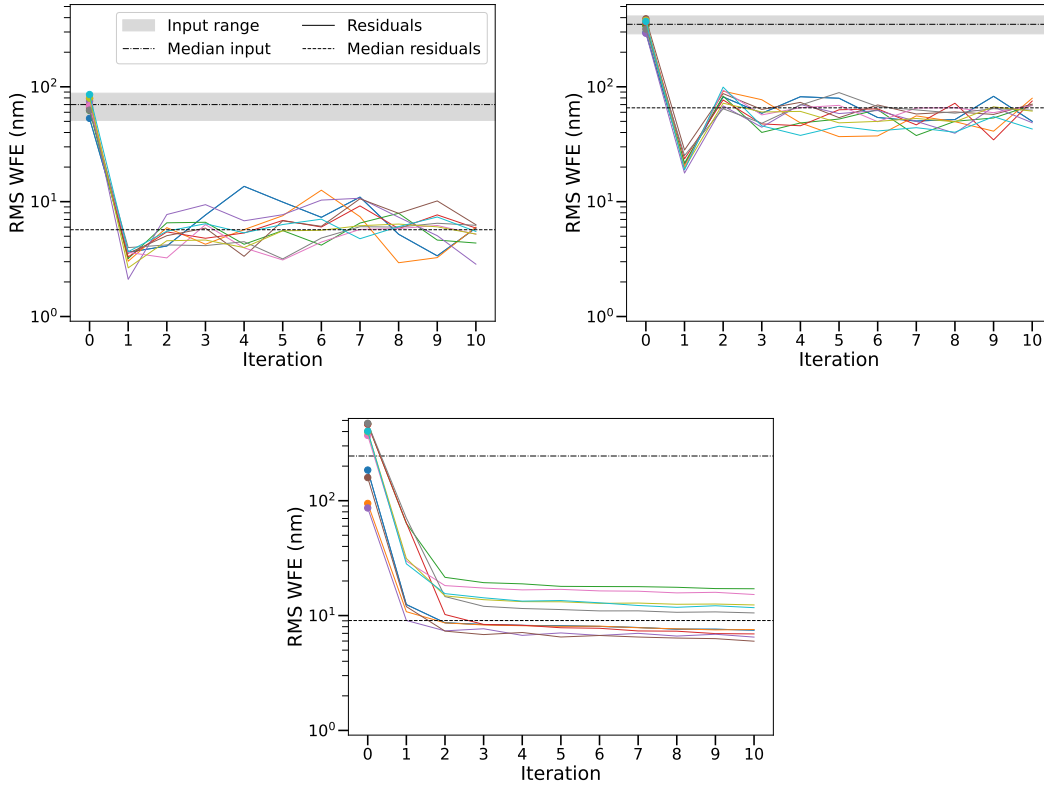


Figure 4.10 Iterative approach for ten test samples with post-VVC PSFs (multicolor lines), over ten iterations. The results for the three datasets are shown: input WFE centered around 70 nm RMS (top left), 350 nm RMS (top right), and in the range [0, 490] nm RMS (bottom). The models used are trained on Efficientnet-b4 with a entrance SNR of 1000.

within the corresponding annulus, and I_0^{Airy} is the flux at the central pixel of the non-coronagraphic and non-aberrated image (Airy disk).

The contrast curves are shown in Figure 4.12, where the separations are converted to the λ/D unit. It can be noticed that there is a small gain in contrast at most separation from the star, and that the contrast gets particularly deep at close separations because of the PSFs “donut” shape, due to the vortex phase singularity. For the obscuration of 30% that is used in this chapter, a relatively important fraction of the light passes through the phase mask, producing bright residuals in the images (Figure 4.11). This prevents reaching deep contrasts at separations of a few λ/D . Results with a smaller central subtraction or a circular pupil are shown in Appendix A, with models trained on the uniformly distributed dataset (WFE $\in [0, 490]$ nm RMS). This dataset provides much better robustness than the normally distributed datasets, especially for the case with the circular aperture where in

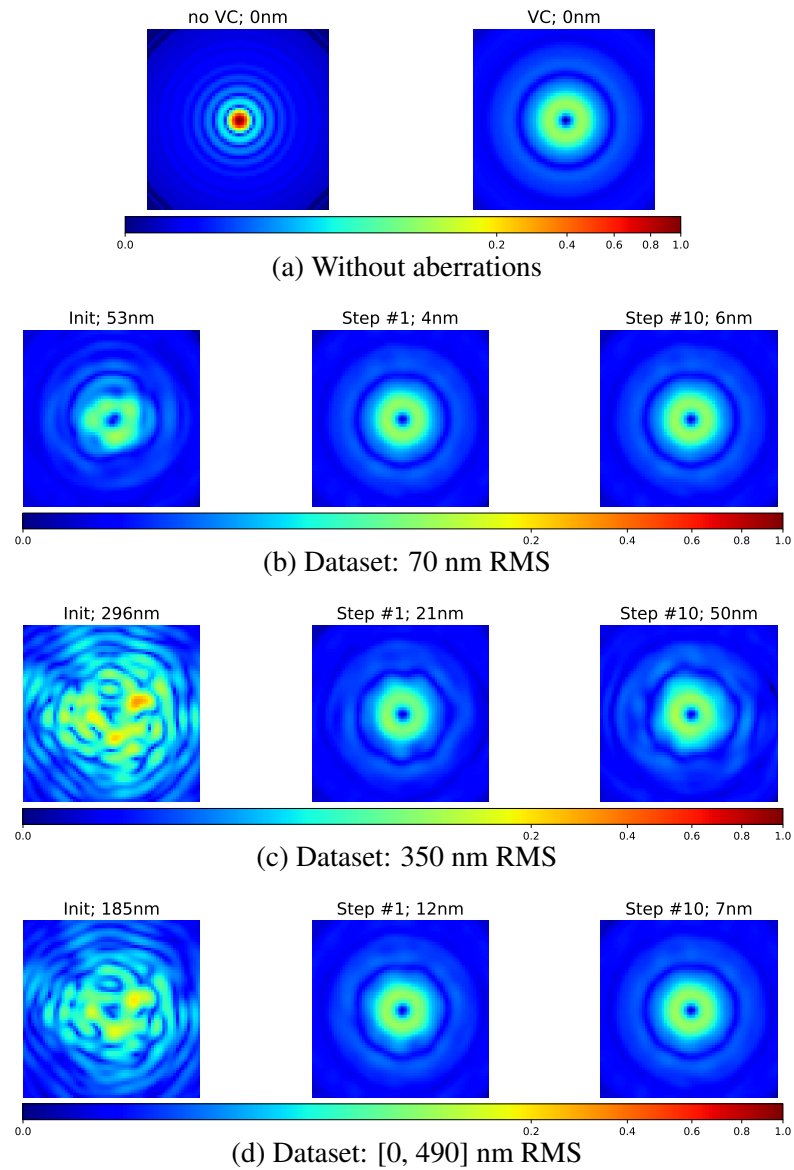


Figure 4.11 Evolution of post-VVC PSFs through iterations with the 70 nm RMS (b), 350 nm RMS (c), and [0, 490] nm RMS (d) datasets. The classical imaging and post-VVC PSFs (a), both diffraction-limited, are displayed for comparison. From left to right: Initial aberrated PSF, image after first correction/iteration, and after 10 corrections. The PSFs are all normalized by the peak intensity of the NCPA-free classical imaging PSF. Each PSF is shown after a power-law normalization to emphasize the residual speckles. The wavefront aberration level of the PSFs is indicated at the top of each corresponding image.

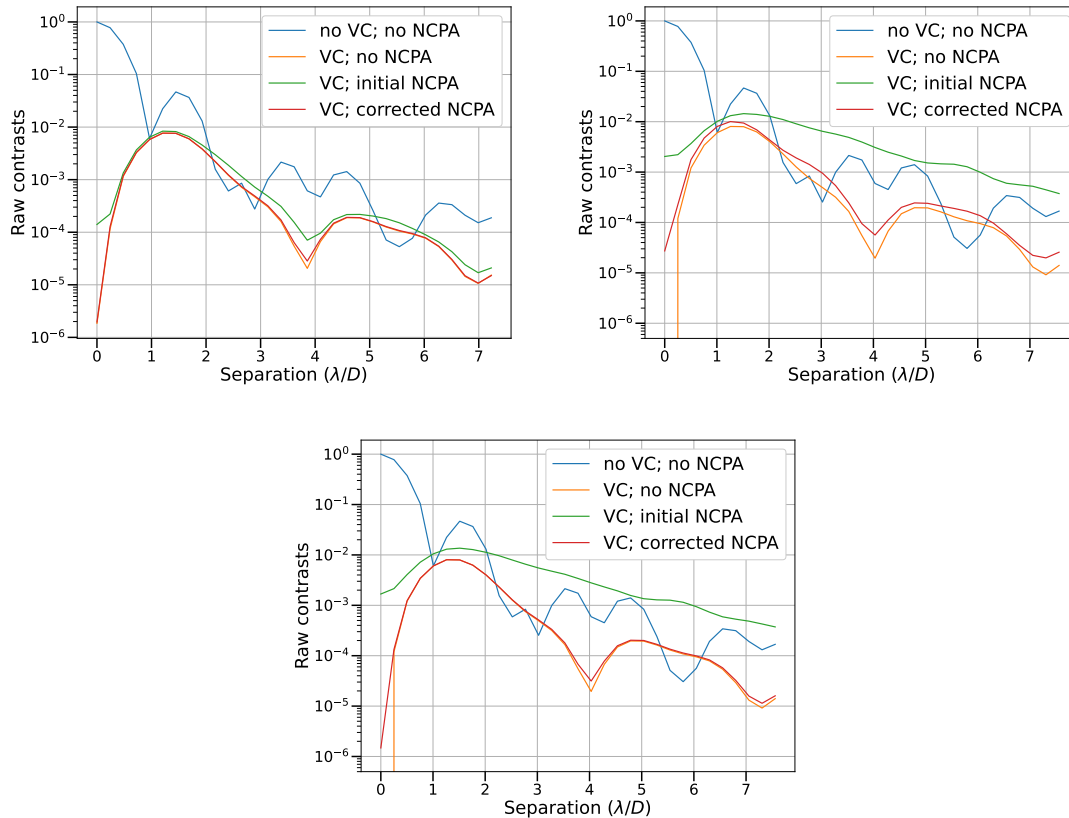


Figure 4.12 Contrast curves after 10 iterations at each separation from the center of the image. The average over 10 PSF samples is taken. The case without coronagraph and aberrations (blue) is compared to adding a vortex phase mask without aberrations (orange), with the initial NCPAs (green), and after correction (red). The three datasets used contain median input WFEs of 70 nm (top left), 350 nm (top right), and [0, 490] nm RMS (bottom).

the low aberration regime the light leakage through the vortex coronagraph is particularly small, which necessitates a model trained on very low aberrations as well.

4.4 Conclusions

In this chapter, a novel approach to performing focal-plane wavefront sensing with vortex coronagraphs has been explored. Using an EfficientNet deep CNN, the modulation introduced by the vortex coronagraph (either scalar or vectorial after splitting the circular polarization states) is leveraged to solve the sign ambiguity and perform FPWFS for

various signal-to-noise ratios, input wavefront errors, and spatial frequency contents. The dual-polarization method with the vector vortex coronagraph offers very similar performance to the classical phase diversity method, which uses additional defocused PSFs, even though the level of light is largely reduced after filtering by the vortex coronagraph.

For instance, considering a star of magnitude 6.2 observed at a wavelength of 2200 nm, a residual of 0.73 nm RMS is obtained from an input WFE of 70 nm RMS. With the SVC, which provides a single focal-plane image, a loss in performance is only observed for high aberration levels. In the presence of high orders and high levels of aberrations (e.g., 88 Zernike modes at 350 nm RMS of phase error), training a CNN is generally challenging. In this case, the performance reaches a plateau of approximately 20 nm RMS, even with bright stars. Under these circumstances, increasing the amount of training data, employing larger and deeper CNN architectures, and implementing regularization techniques can further improve the phase retrieval accuracy. Atmospheric turbulence residuals that are expected in ground-based observations only produce minor degradation in performance in a low NCPA regime, and they should not be a concern in practice. Models trained on data containing particularly wide WFE and SNR distributions provide very good robustness.

Closed-loop tests have also been carried out by successively feeding the corrected PSF to the model to obtain the next prediction. The results further highlight the stability of the prediction process as well as the gains in terms of contrasts achieved with those residual NCPAs.

It is difficult to obtain reliable and very precise NCPA labels for model training in supervised learning. Employing unsupervised learning techniques such as autoencoder-based architectures, is another interesting approach that is explored in the next chapter.

Physics-informed learning with autoencoders

In this chapter, a physics-based approach is proposed by incorporating a differentiable optical simulator into the deep learning framework, resulting in an autoencoder architecture that estimates phase aberrations in the latent space. This unsupervised method therefore eliminates the need for labeled data during training. The chapter begins with an overview of automatic differentiation and discuss two possible physics-based learning approaches. Section 5.2, which has been published in the SPIE conference proceedings (Quesnel et al., 2022b), then details the autoencoder method and present the results. Further developments include the incorporation of vortex phase diversity into the autoencoder and the exploration of a variational inference approach.

Contents

5.1	Introduction	90
5.1.1	Differentiable optical systems	90
5.1.2	Potential physics-based architectures	91
5.2	General simulator-based autoencoder method	93
5.2.1	Data generation	93
5.2.2	Autoencoder architecture	94
5.2.3	Experiments and results	95
5.3	Unsupervised vortex phase diversity	101
5.3.1	Models and results	101
5.3.2	Rejection factor optimization	102
5.4	Variational inference approach	103
5.4.1	Definition	105
5.4.2	Training the SimVAE	107
5.4.3	Results and discussions	108
5.5	Conclusions	111

5.1 Introduction

One important drawback for standard deep learning approaches is the lack of physical principle within their architectures. This can cause a poor interpretability of the models, making it potentially difficult to validate the results. The generalization capabilities of models can also be an issue when the physical process is not taken into account, because real-world data that is distributed differently than the training data may lead to non-realistic predictions. Training datasets should then be very large to present proper robustness, which may be cumbersome to generate, especially for supervised learning.

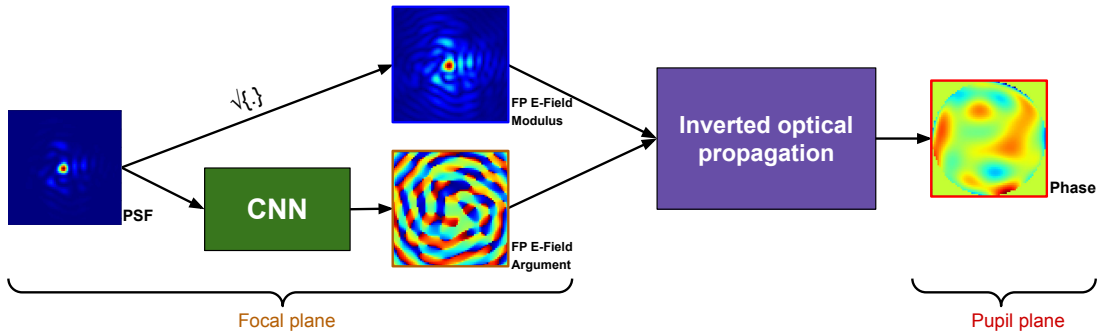
5.1.1 Differentiable optical systems

As explained above, incorporating physical models into the learning procedure would be highly beneficial. Such an integration nonetheless requires the parameters to be differentiable, so gradients can be computed during the backward pass through the model. These gradients are needed to optimize a physics-based deep learning model, as it is done to learn neural network weights (see Equation 2.4).


Automatic differentiation (AD) offers an efficient method to evaluate derivatives. There are two ways to apply the chain rule for the backpropagation algorithm using AD: forward and reverse accumulations. Choosing one mode or the other determines in which order the derivatives of the chain rule are multiplied, and impacts computational efficiency, depending on the dimensions of the inputs and outputs (Baydin et al., 2017). For artificial neural networks, reverse-mode AD is more computationally efficient compared to forward-mode AD, since the output is typically a scalar value (the loss), while the input is usually of very high dimension (e.g., batches of multichannel images). For this reason, reverse-mode AD is commonly used in deep learning and is also relevant for the optimization of physical systems.

Several studies have demonstrated the effectiveness of AD for phase retrieval (Jurling & Fienup, 2014) and telescope optics optimization (Liaudat et al., 2021) with reverse-mode accumulation. Numerical wavefront sensors can also be built using AD. For instance, Landman et al. (2022) proposed a joint optimization procedure where a differentiable Fourier-filtering wavefront sensor is optimized while a CNN is trained to reconstruct the wavefront aberrations. Additionally, AD has been used to optimize coronagraphic masks (Wong et al., 2021; Por et al., 2022). Another noteworthy application is kernel phase analysis, initially applied to interferometry (Martinache, 2010), and later adapted for instrumental self-calibration with forward-mode AD (Pope et al., 2021).

Libraries have been developed to perform AD on optical systems, notably Morphine



(a) Inverted simulator-based architecture



(b) Autoencoder-based architecture

Figure 5.1 Two proposed approaches for phase retrieval with physics-informed learning. (a): Using an inverted simulator and a CNN to learn the non-invertible part. (b): An autoencoder architecture containing a simulator to perform unsupervised learning.

(Wong et al., 2021), which combines the optical propagation package Poppy with JAX to make the code differentiable and allow gradient computation. Morphine was later replaced by δ Lux (Desdoigts et al., 2023), which is much more efficient and allows for high dimensional optimization of optical parameters.

5.1.2 Potential physics-based architectures

This section explores how we can harness the concept of physics-informed modeling with AD for our problem using an optical simulator. In particular, two approaches are investigated: a supervised learning method containing an inverted simulator (Figure 5.1a), and a label-free method based on a simulator-based autoencoder (Figure 5.1b). Advantages and drawbacks of both architectures are summarized in Table 5.1.

	Inverted simulator	Autoencoder
Structure	Backward simulator + CNN	CNN + forward simulator
Advantages	Learn only what is needed	Unsupervised learning
Drawbacks	<ul style="list-style-type: none"> – Limited invertible optics – Supervised learning 	<ul style="list-style-type: none"> – Imperfect simulators

Table 5.1 Specific features of the two envisioned physics-based approaches.

Inverted simulator

The idea behind the first approach is to perform part of the work previously carried out by the CNN, using a backward simulator for the optical transformations that are invertible, i.e., the optical propagation of the electric field between the pupil and focal planes. Only the transformation of the focal-plane electric field (FP E-field) into intensities is not invertible, and therefore can be handled by a CNN. Specifically, only the argument of the FP E-field needs to be inferred by the CNN (Figure 5.1a), since the FP E-field modulus is directly obtained as the square root of the PSF intensity. Although this inverted approach only learns a limited optical process, experiments have shown that it is neither easier nor faster to learn the FP E-field compared to the pupil-plane one, and the overall phase retrieval performance does not improve either. A “lighter” version of U-Net with four times fewer parameters than the original was also tested, showing consistent phase retrieval performance. One potential issue is phase wrapping, which occurs for aberrations with peak-to-valley WFEs outside the range $[-\pi, \pi]$, similar to what iterative methods such as the Gerchberg-Saxton may encounter (see Section 1.4.3).

Simulator-based autoencoder

The motivation for the second approach is to make the FPWFS approach “unsupervised”, i.e., not relying on labels to train the models. This is justified as building models with in-lab and on-sky data is particularly challenging due to the lack of precise knowledge about the NCPAs themselves, it may therefore be difficult to build reliable labeled datasets. This is where CNN-based supervised learning show limitations. By reconstructing the input data, deep learning autoencoders are a well-proven unsupervised approach (Baldi, 2012). For focal-plane wavefront sensing the targets are the NCPAs instead of the point spread functions (PSF), and one way to extract this phase information with an autoencoder is to add physical knowledge to the model. This is achieved by including a simulator in the decoder that is forcing the latent space to be a representation of the phase aberrations. This approach therefore overcomes the limitations of the inverted simulator, and offers a label-

free training. The main requirement for such a method to work efficiently is, however, to have a reliable differentiable simulator.

Since the inverted simulator concept do not offer as many benefits as the simulator-based autoencoder, the following sections focus on the latter approach.

5.2 General simulator-based autoencoder method

In this section, the simulator-based autoencoder method is designed and tested. As in previous chapters, the data generation process is described as well as the training procedure. Experiments in the classical imaging case are then carried out.

5.2.1 Data generation

For this study, we consider an optical propagation between a pupil plane and a focal plane, as defined in Equation 1.5. As before, the phase maps are built using Zernike polynomials following the Noll convention, starting from the tip mode, and based on an annular aperture with a central obstruction of 30% of the total diameter. The sets of Zernike coefficients are randomly generated to approximate a $1/f^2$ power spectral density profile and scaled to a given median root-mean-square wavefront error (Equation 3.1). In the simulations, two aberration regimes are considered: NCPAs around 70 nm RMS covering 20 Zernike modes, and 350 nm RMS over 100 modes. The PSFs are computed using the PROPER optical propagation package. Again, the data is generated in the K band at $\lambda = 2200$ nm, while an aperture diameter of 10 m, a pixel scale of about 11 mas/pix, and a field-of-view of about 1.5'' are considered. The phase maps as well as the PSFs contain 64×64 pixels. Out-of-focus PSFs with a defocus of $\lambda/4 = 550$ nm RMS are also generated for phase diversity. Only photon noise is added to the PSFs, and a square-root stretching operation is performed to the PSFs, before normalizing the flux in the range [0,1]. This operation is done inside the encoder for the autoencoder architecture.

For the experiments done in Section 5.2.3, data containing atmospheric turbulence residuals are also generated. This is achieved by simulating an extreme AO system using the COMPASS library (Ferreira et al., 2018), giving a wavefront error of about 50 nm RMS. To simulate a 1-s exposure in the presence of a given amount of static NCPA, a sequence of 10 consecutive AO phase screens that are 100ms apart is used, and the corresponding PSFs are summed up. These AO simulations are described in more detail in Chapter 4.2.1.




Figure 5.2 The proposed simulator-based autoencoder architecture (SimAE), uses a CNN as the encoder and a differentiable optical simulator as the decoder. The latent space contains the Zernike coefficients which are used during training to reconstruct noiseless PSFs, while during inference these coefficients can be applied for NCPA correction.

5.2.2 Autoencoder architecture

In this work, an autoencoder architecture is proposed, as presented in Section 5.1.2. As in previous chapters, the encoder is a neural network that approximates a function f that maps Zernike coefficients z from PSFs x : $z \approx f(x)$. The deep CNN EfficientNet (Tan & Le, 2019), presented in Section 4.2.2 and previously used in Chapter 4, is employed for this task.

The decoder corresponds to a differential optical simulator, i.e., it reconstructs the input PSFs based on the predicted Zernike coefficients, allowing it to work without labels. This decoder is the same as the data generator defined in Equation 1.5, with the exception that it produces noiseless PSFs. In this section, all the simulator parameters are fixed, and learning optical parameters is notably investigated in Section 5.3.2. Regarding the latent space that exists between the encoder and the decoder, the number of Zernike modes has to be predefined. During inference on a trained model, the encoder containing the CNN allows us to obtain the Zernike coefficients that correspond to the NCPAs of interest. Figure 5.2 represents the simulator-based autoencoder architecture, which is called “SimAE” in the following sections.

The loss function, which is used to optimize the neural network’s weights in order to fit the model to the data, needs to be adapted to our SimAE architecture. Since we are working in a case where only photon noise is present in the input images, i.e., each pixel follows a Poisson distribution, the typical mean-square error loss that makes the assumption of Gaussian noise is not appropriate for the PSF residuals. We can rather define

the loss so that the log-likelihood of the Poisson probability distribution is maximized, giving the expression

$$\mathcal{L}_{\text{SimAE}}(x; \theta) = -\frac{1}{N_{\text{pix}}} \sum_{i=1}^{N_{\text{pix}}} \log \left(\frac{\lambda_i(x_i; \theta)^{x_i}}{x_i!} \exp(-\lambda_i(x_i; \theta)) \right), \quad (5.1)$$

where $\lambda_i(x_i; \theta)$ is the rate of the Poisson distribution, representing the photon flux at pixel i . This corresponds to the noiseless PSF pixel reconstructed by the decoder, whose flux is scaled to the measured flux of the input PSF pixel x_i . The parameters θ of the CNN define the output of the encoder, i.e., the latent space, which represents the aberrations affecting the photon flux at each pixel. The final loss value is the average of the log-likelihoods over the whole image of size N_{pix} .

The baseline on which the SimAE models are compared is a standard supervised CNN, as applied in Chapters 3 and 4. As a reminder, the loss function is in this case defined as the RMS error of the phase residuals. It can be expressed using the Zernike coefficients as:

$$\mathcal{L}_{\text{CNN}}(z, \hat{z}(x; \theta)) = \sqrt{\frac{1}{N_{\text{modes}}} \sum_i^{N_{\text{modes}}} (z_i - \hat{z}_i(x; \theta))^2}, \quad (5.2)$$

where z and \hat{z} correspond to the true and predicted Zernike coefficients respectively, while N_{modes} is the number of Zernike modes (20 or 100).

5.2.3 Experiments and results

Experimental protocol

The hyperparameters used to train the models are as presented in Section 4.2.3. No weight decay is used this time. Efficientnet-b4 weights are initialized with models pre-trained on ImageNet. An example of the resulting training and validation losses for each method can be found in Figure 5.3. The models are then systematically evaluated on 10^3 new samples.

Performance compared to a standard deep CNN

The performance of each Zernike mode is first compared. The metric used is the RMSE per mode, computed over the entire test set:

$$\sigma_z = \sqrt{\frac{1}{N_{\text{test}}} \sum_i^{N_{\text{test}}} (\hat{c}_i - c_i)^2}, \quad (5.3)$$




Figure 5.3 Examples of training and validation losses over the epochs for each approach: CNN (*left*) and SimAE (*right*). In these examples, the models are trained on datasets containing 350 nm RMS of aberrations distributed over 100 modes, and an SNR equal to 300.

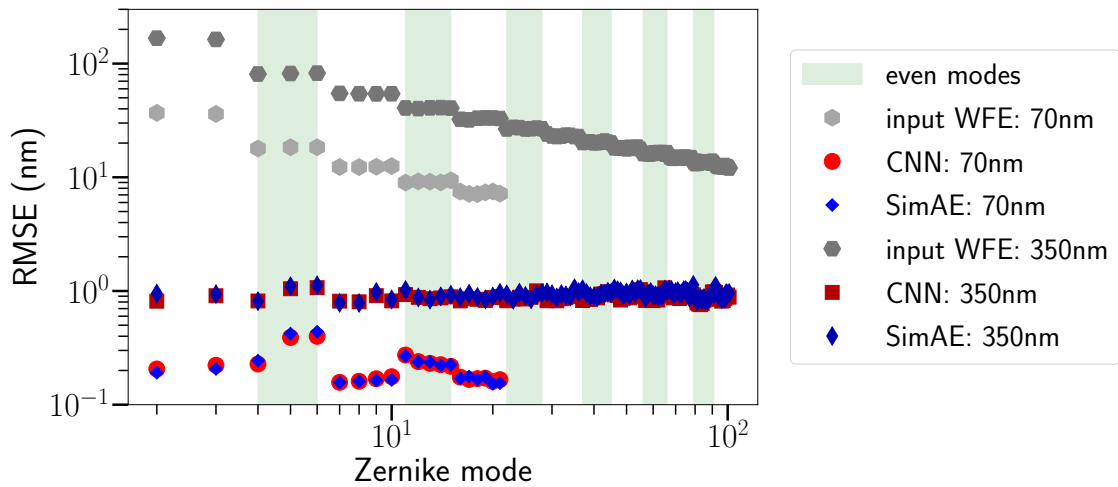


Figure 5.4 Performance per Zernike mode, for a low aberration regime (70 nm over 20 modes, light shades) and a high aberration regime (350 nm over 100 modes, darker shades), at an SNR of 1000. The CNN and SimAE architectures are compared (blue and red, respectively).

where N_{test} is the number of test samples, while \hat{c} and c are the estimated and true Zernike coefficients, respectively.

As shown in Figure 5.4, the Zernike coefficients are very well reconstructed, even at high aberrations contents (350 nm RMS over 100 modes). There is no particular distinction between even and odd modes, and the SimAE and CNN manifest very similar




Figure 5.5 Performance at different SNR levels for the SimAE (blue) and the standard CNN (red) models, at low and high aberration regimes (darker and brighter shades respectively). The points correspond to the median value and the error bars to the 2–98th percentiles.

performance.

Since the SimAE is defined to take into account photon noise through its loss function, it is worth assessing the performance for different SNR levels. To do so, we compare the residual errors in terms of RMS WFE on the total phase residuals. The metric is therefore defined for each test sample as:

$$\sigma_{\Phi} = \sqrt{\frac{1}{N_{\text{pix}}} \sum_i^{N_{\text{pix}}} (\widehat{\Phi}_i - \Phi_i)^2}, \quad (5.4)$$

where N_{pix} is the number of pixels within the pupil area, while $\widehat{\Phi}$ and Φ are the estimated and true pupil phases, respectively.

The errors on the phase retrieval for a range of SNR levels are shown in Figure 5.5, where each point corresponds to a model trained and evaluated at the same given SNR. The two approaches show almost identical performance and a plateau is reached around an SNR of 3000 in the high aberration regime. This has been observed and discussed in previous works (Orban de Xivry et al., 2021; Quesnel et al., 2022a) as well. Only at very low SNRs, does the SimAE appear to differ from the CNN approach. This difference

could be due to the loss function of the SimAE that is defined on the PSFs and not on the NCPAs, so that there is no implicit regularization occurring with supersized CNNs, which has been described by [Orban de Xivry et al. \(2021\)](#). The residual errors thus tend to follow the theoretical limit for lower SNRs with the SimAE.

Robustness of the trained models

Deep learning models may show limitations when evaluated on data lying outside the training data distribution. It is therefore important to assess the stability of the models in such conditions, which would probably occur when applying the algorithm on sky.

Here we compare the SimAE method with the CNN, this time varying the testing SNR (Figure 5.6, top), on a model trained at a median WFE of 70 nm RMS distributed over 20 modes. Giving data containing less noise to the models always provides stable predictions, and good improvement is even obtained with the SimAE when trained at an SNR of 30. If more noise is introduced to the test data, the error quickly increases, at about the same rate with both approaches. The testing input WFE present in the data is also modified (Figure 5.6, bottom). Here the CNN and SimAE showcase identical behavior, with constant performance for lower testing WFE than the training one, and strong degradation for higher testing WFE, although with the SimAE the performance does not degrade as quickly as with the CNN. Similar trends regarding the CNN have been found in previous works ([Orban de Xivry et al., 2021](#); [Quesnel et al., 2022a](#)). Overall the SimAE method offers some improvement in terms of robustness compared to the CNN.

Transfer learning

Instead of directly applying a trained model on observed data, we propose to fine-tune the model using a single “on-sky” image, initializing with the trained weights of the original model based on synthetic data. This is done to illustrate how we could quickly adapt the model from simulated to on-sky data. To test this on the SimAE method, a model trained on data containing 70 nm RMS and an SNR of 100 is used. The levels of NCPA and SNR are changed on simulated test data, and the results are shown in Figure 5.7. The models perform much better when transfer learning is done on each test sample, in particular when their SNR and WFE are higher than for the original training data. This second training on a single test sample is performed over 200 epochs, to make sure all the models have time to converge. The training process takes about 20 seconds to complete on one GPU (Nvidia GeForce RTX 2080 Ti). For data relatively close to the ones used for pre-training, convergence is naturally achieved after much fewer epochs and the final prediction can be made in only a few seconds. These early results have shown that transfer learning on real

5.2. GENERAL SIMULATOR-BASED AUTOENCODER METHOD




Figure 5.6 *Top*: Performance with changing the SNR of the test data for two models trained on SNRs of 30 and 300 (indicated by the vertical dotted lines). The median input WFE is 70 nm with 20 modes. *Bottom*: Performance for different testing input WFE, on models trained at 70 nm and 350 nm (20 and 100 modes respectively).

data is particularly motivated. This is investigated in Chapter 6 with supervised learning models and the SimAE.

Atmospheric turbulence residuals

In the previous experiments, the decoder, i.e., the differentiable optical simulator, was able to exactly reconstruct what the data generator produced to build the training and test sets,

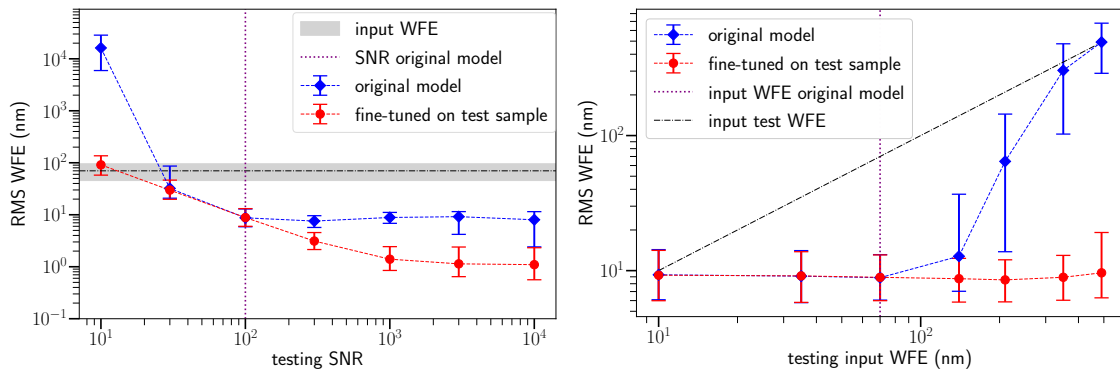


Figure 5.7 Performance of models fine-tuned on each test sample individually, for varying SNR (*left*) and input WFE levels (*right*). The model used for both plots was initially trained on data with 70 nm RMS, 20 modes, and an SNR of 100, as indicated by the vertical dotted lines.




Figure 5.8 Residual errors as a function of the input WFE, for a model trained on a dataset with WFEs distributed around 70 nm (20 modes). A SimAE model with data containing only photon noise (cyan), is compared to CNN and SimAE models (purple and blue, respectively) trained on data containing AO residuals, which are not provided to the decoder during training and evaluation.

except for photon noise, but for which the loss function was adapted. In order to test the method on data the decoder cannot totally reconstruct, atmospheric turbulence residuals are added to the input data, without giving this information to the decoder. These AO residuals amounts to 50 nm RMS of aberrations, and the way they are included in the data

is described in Section 5.2.1. Even with the presence of this additional source of noise, the training is very stable, and the evaluation results shown in Figure 5.8 indicate that the SimAE suffers a reasonable degradation in performance compared to a case without turbulence (from 1 to 4 nm RMS on average). Using a CNN with labels that do not account for these AO residuals gives the same results.

5.3 Unsupervised vortex phase diversity

In this section, the dual-polarization approach behind a vortex coronagraph is tested on the SimAE. The simulator is thus upgraded to include differentiable propagation through a vortex coronagraph.

5.3.1 Models and results

For the experiment, the data is generated with a uniform WFE distribution between 0 and 350 nm RMS, and two spatial frequency contents are selected: 20 and 100 Zernike modes (Figure 5.9a). Photon noise is included to give an entrance SNR of 1000. As explained before, the effective SNR is reduced in the focal-plane due to the vortex mask, also depending on the level of aberrations. The two datasets, composed of 10^5 samples each, are then used to train the SimAE, as well as a CNN in a supervised learning way for comparison. The ResNet-50 architecture is used for both approaches, since it yielded better results than EfficientNet in this case.

The distribution of residual phase errors on the 1000 test samples with the SimAE is shown in Figure 5.9a. We can see that the amount of Zernike modes considered has a strong impact on the overall performance. As done with classical imaging (Figure 5.6, bottom), the robustness with varying testing WFE is evaluated. The results are shown in Figure 5.9b. The SimAE provides some improvements in the low spatial frequency content compared to the CNN. This is notably true for testing WFEs that are close to the average training WFE (training WFE $\in [0, 350]$ nm RMS). At a very low testing WFE of 10 nm RMS, the CNN performs relatively well on average but contains a few very poor predictions. Overall, the SimAE has better constrained results with lower error bars. In the large frequency content range (100 Zernike modes), both approaches do not give very good predictions. This is explained by the fact that not only the training WFE range is very large, so is the number of Zernike modes to predict, and the SNR in the images is also changing depending on the level of aberrations. A more complex and robust architecture would be needed in this case.

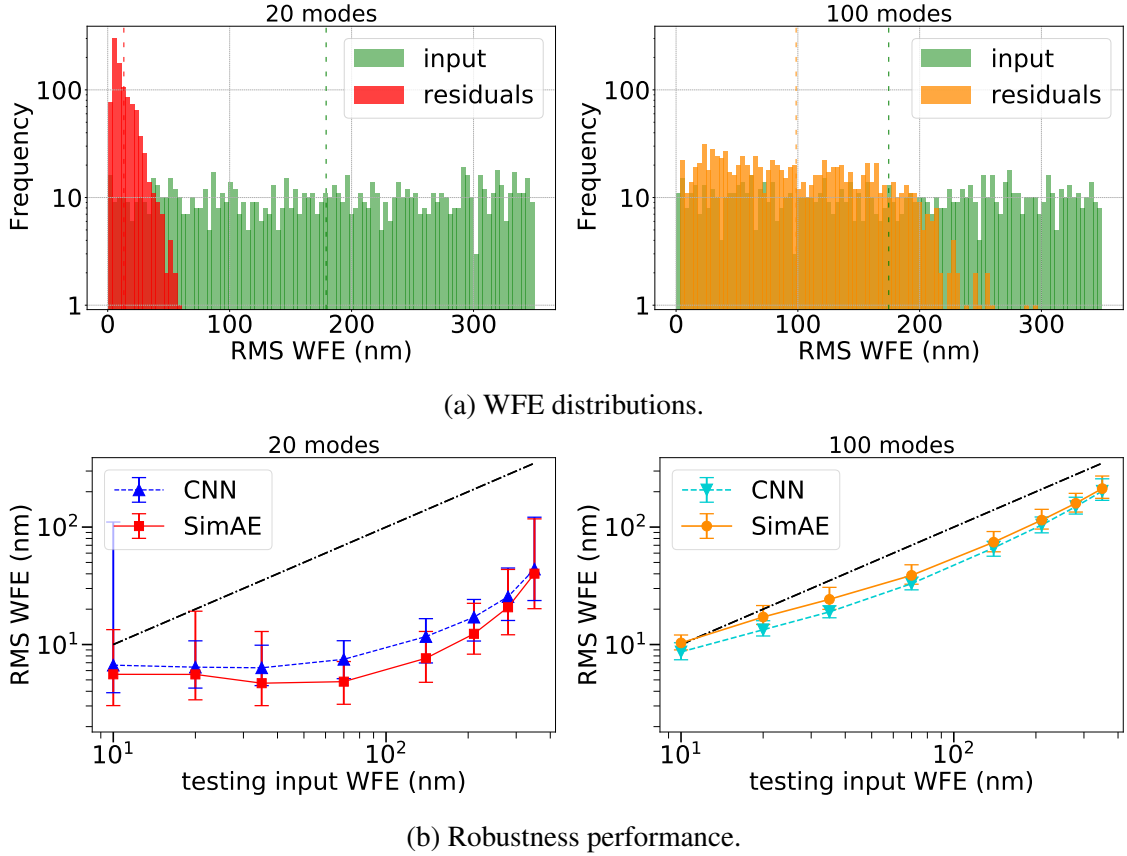


Figure 5.9 (a): Input WFEs in the range [0, 350] nm RMS (green), and phase residuals with the SimAE (red and orange). 1000 test samples are used to produce the results. (b): Performance for different input WFE on the SimAE and supervised CNN models trained in the range [0, 350] nm RMS. Vortex phase diversity data is used at an entrance SNR of 1000. Results are with aberrations decomposed into 20 and 100 Zernike modes (left and right, respectively). The error bars are obtained with the 2–98th percentiles.

5.3.2 Rejection factor optimization

One important advantage of a differentiable simulator is that it is possible to learn some of its parameters. One relevant parameter that is explored here is the rejection factor, which quantifies the capacity of the vortex mask to cancel the on-axis incoming light. It corresponds to the inverse quantity of the flux leakage, and it can be used to assess the performance of the vortex mask. In the vortex simulator, it implies that the final PSF can be defined as:

$$\mathcal{I}_{\text{true}} = \left(1 - \frac{1}{R}\right) \mathcal{I}_{\text{perfect}} + \frac{1}{R} \mathcal{I}_{\text{classic}}, \quad (5.5)$$

with $\mathcal{I}_{\text{perfect}}$ the PSF behind a perfect vortex coronagraph, $\mathcal{I}_{\text{classic}}$ the non-coronagraphic PSF, and R the rejection factor.

Typical values for a VVC that are suitable to work on-sky are in the order of $10^2 - 10^3$ (Defrère et al., 2014; Jolivet et al., 2019). It may not be straightforward to know the exact value for R after manufacturing the phase mask. Parametrizing the model with such factor can therefore reduce reconstruction uncertainties and make the SimAE perform better.

To assess the feasibility of learning R in a deep learning context, datasets containing VVCs with R comprised between 10 and 10^3 are simulated. SimAE models are then trained with an initial factor set to $R = 100$. As shown in Figure 5.10 (bottom), the learning is tracked over the epochs for seven different datasets and models. A relatively small dataset composed of 10^3 training samples is used, and the phase is composed of 20 modes at aberration levels around 70 nm RMS, while the SNR is set to 1000 (defined in the pupil-plane, upstream of the vortex mask). The models perfectly reconstruct the rejection factor up to $R = 50$, and the worst reconstruction still yields 86% accuracy (for $R = 10^3$). As it can be seen in Figure 5.10 (top), for $R = 10^3$, the difference between the true image and a PSF behind a perfect vortex is barely distinguishable. The noise present in the data is considered to be the sole limitation to the rejection factor reconstruction.

The residual wavefront error also remains quite stable (between 35 and 40 nm RMS) for all rejection factors, showing that the slight errors on the reconstruction \widehat{R} for rejection factors above 50 do not affect the prediction capabilities of the encoder. Using more data and having more input flux decreases the residual WFE but do not show a particular improvement on \widehat{R} .

Training custom parameters alongside a neural network may require using multiple learning rates, because the custom parameters can have values at a different order of magnitude than the neural network weights. For simplicity, the inverse of the rejection factor R^{-1} was trained instead, so a common (initial) learning rate of 10^{-3} could be used. Optimizing directly R with this relatively small learning rate works as well, but a much larger amount of epochs is then needed to converge to the best solution.

Depending on the instrument in use, other simulator parameters can be potentially optimized, such as camera rotations, flips, amplitude maps, and pixel scale. This is discussed in Section 6.4.4 with the Subaru/SCEXAO instrument.

5.4 Variational inference approach

The models trained so far have performed deterministic predictions. Probability distributions via the posterior provide a more general representation of the predictions, as presented in Equation 1.10. In this section, the application of a variational approach to the

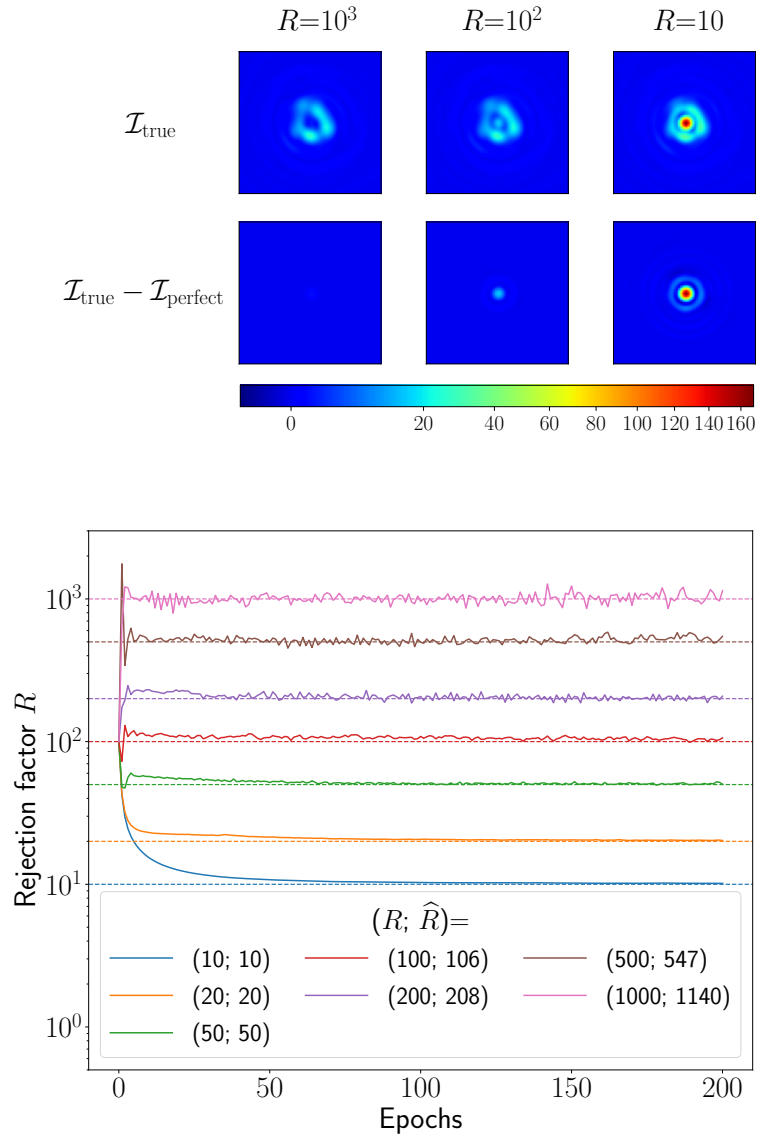


Figure 5.10 *Top*: PSFs $\mathcal{I}_{\text{true}}$ simulated with three different rejection factors R are shown next to the difference between $\mathcal{I}_{\text{true}}$ and a PSF behind a perfect vortex coronagraph $\mathcal{I}_{\text{perfect}}$. The mean square error of the difference is indicated above each corresponding frame. The flux is normalized to be equal to 10^6 and the images are without photon noise. *Bottom*: Learned values for the vortex rejection factor through the training epochs. Seven different R are simulated and the initialized factor for the training is always 100. The final estimation \hat{R} is written in the legend. The dataset used in this example contains 10^3 samples composed of 20 Zernike modes and a median WFE of 70 nm RMS, while the entrance SNR level is 1000. The encoder is composed of the ResNet-50 architecture.

simulator-based autoencoder is explored.

The models trained so far have produced deterministic predictions. Incorporating probability distributions through the posterior offers a more comprehensive representation of the phase predictions, as illustrated in Equation 1.10. In this section, the application of a variational approach to the simulator-based autoencoder is explored.

5.4.1 Definition

Variational inference aims to predict probability distributions instead of deterministic values. The standard way to achieve this is to predict distribution parameters, such as the mean and the variance of a Gaussian if such posterior is considered. This approach aligns with the formalism of Equation 1.10 and works similarly to mixture models implemented in Chapter 3. A variational autoencoder (VAE) is an unsupervised approach that uses distributions in the latent space, incorporating a prior distribution. The final encoding is then obtained by sampling the posterior distribution. A schematic representation of a VAE is provided in Figure 5.11.

Because it is not possible to backpropagate through a stochastic node, direct sampling from the posterior distribution is not feasible. VAEs address this issue using a “reparametrization trick”, which separates randomness from the parameters by introducing an external random variable, ϵ . For a Gaussian posterior, the expression of the latent variables is

$$z = \mu + \sigma \cdot \epsilon, \quad (5.6)$$

where μ is the mean, σ the standard deviation predicted by the encoder, and z the sample drawn in the latent space.

The loss function of a VAE contains two key terms. Firstly, there is the reconstruction error, which computes the discrepancy between the input data and the output of the decoder as

$$\mathcal{L}_{\text{reco}} = -\frac{1}{N_{\text{pix}}} \sum_p^{N_{\text{pix}}} \log(p_{\psi}(x|z)), \quad (5.7)$$

where $p_{\psi}(x|z)$ is the likelihood obtained from the decoder with parameters ψ , $q_{\theta}(z|x)$ the posterior predicted by the encoder with parameters θ , and N_{pix} the number of pixels in the image. The reconstruction term has been the loss function used in this chapter so far, as represented in Equation 5.1.

The second term is the Kullback-Leibler (KL) divergence (Kullback & Leibler, 1951),

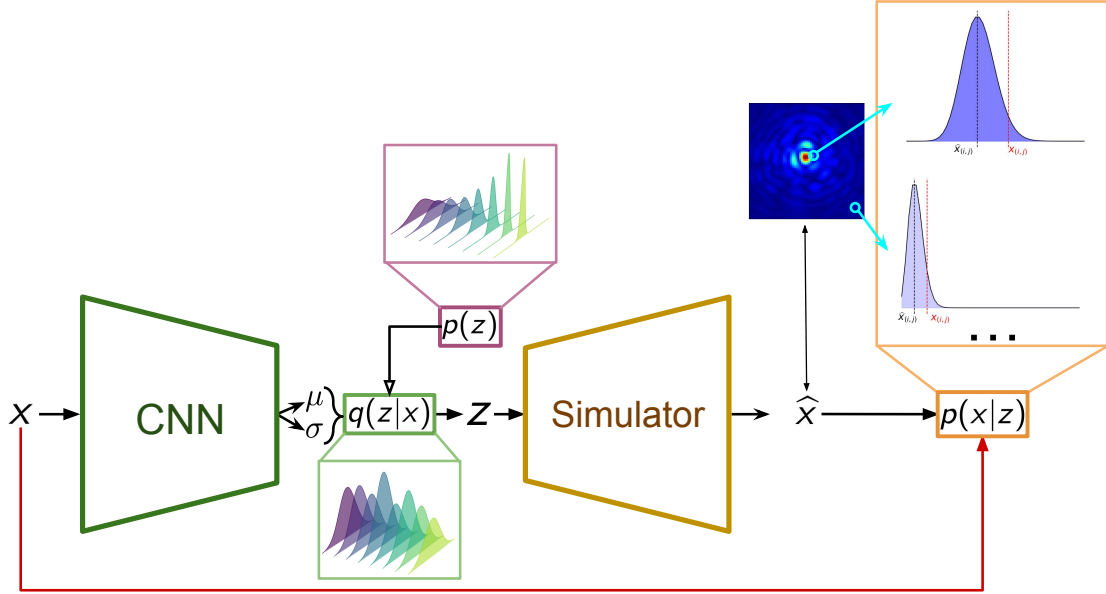


Figure 5.11 Diagram of the simulator-based variational autoencoder architecture (SimVAE). The encoder (green) estimates the parameters μ and σ for each Zernike modes, allowing to build a Gaussian posterior $q(z|x)$ parametrized by μ and σ . The Zernike coefficients z are then sampled from $q(z|x)$ and fed into the decoder (yellow), which estimates the noiseless reconstruction \hat{x} . This allows to compute the likelihood $p(x|z)$ for the input x , following the specified noise distribution, which, in our case, is a Poisson distribution parametrized by \hat{x} as the mean value. The SimVAE model is trained by maximizing $p(x|z)$ while simultaneously minimizing the gap between $q(z|x)$ and a prior $p(z)$, which contains information about the coefficient distribution for each Zernike mode.

which regularizes the training by taking into account prior assumption about the data:

$$\begin{aligned} \mathcal{L}_{\text{KL}} &= \text{KL}(q_{\theta}(z|x) \parallel p(z)) \\ &\simeq \sum_i^{N_{\text{modes}}} q_{\theta}(z_i|x) \log \left(\frac{q_{\theta}(z_i|x)}{p(z_i)} \right), \end{aligned} \quad (5.8)$$

where $p(z)$ is the prior distribution and N_{modes} the size of the latent space, corresponding to the number of Zernike modes. Minimizing this loss term forces the posterior distribution to be as close as possible to the prior distribution.

The total loss for a given batch of data sampled from $p(x)$, thus becomes

$$\mathcal{L}_{\text{total}} = \frac{1}{N_{\text{samples}}} \sum_{i=1}^{N_{\text{samples}}} (\mathcal{L}_{\text{reco}} + \beta \mathcal{L}_{\text{KL}}), \quad (5.9)$$

where N_{samples} is the batch size and β is a regularization factor that can be applied to reduce the impact of the KL divergence term during training, if needed (Higgins et al., 2017).

Unlike standard autoencoders, which only penalize poor reconstructions by the decoder, a variational autoencoder also accounts for the quality of the encoding and the latent space through the KL loss term. VAEs are especially advantageous for generating new data, because the latent space becomes interpretable thanks to the prior. In our simulator-based autoencoder, the decoder already provides an interpretable latent space, and generating new data is not the focus of this thesis. Another advantage of VAEs is that they provide prediction uncertainty with the posterior, which can help to assess the quality of the predictions. VAE models are also more likely to generalize well and are therefore less prone to overfitting.

5.4.2 Training the SimVAE

In our phase retrieval application, the mean $\widehat{\mu}_z$ and standard deviation $\widehat{\sigma}_z$ for each Zernike mode z are predicted by the encoder. This allows us to construct a posterior distribution $p(z|x) = \mathcal{N}(\widehat{\mu}_z, \widehat{\sigma}_z)$. For the reparametrization trick defined in Equation 5.6, which enables the sampling of Zernike coefficients from the posterior $p(z|x)$, the random variable ϵ is sampled from a uniform distribution defined in the range $[0, 0.01]$. This prevents the sampled coefficients z from becoming excessively large and improves stability during training.

Gaussian distributions are considered here as they allow to apply the reparametrization trick defined in Equation 5.6, and it represents a good starting point to assess the capabilities of the SimVAE. Experiments with uniformly distributed coefficients are done using real data in Section 6.4.3. To test the relevance of a Gaussian prior, Zernike coefficients are generated following a centered Gaussian distribution instead of a uniform distribution (before normalization by the radial orders and the chosen input WFE; see Equation 3.1). This gives the distribution per Zernike modes shown in Figure 5.12, with the $1/f^2$ power law over the modes. The corresponding standard deviation of the generated distribution σ_z is then used to define the prior for the training as $p(z) = \mathcal{N}(0, \sigma_z)$. The KL divergence is finally computed using Equation 5.8 and a factor $\beta = 10^{-2}$ is set to balance the two loss terms (see Equation 5.9).

For the experiments, three datasets each are employed, that contain dual-polarization post-VVC PSFs in different spatial frequency and aberration regimes: 20 and 100 modes with 70 nm RMS, as well as 50 modes with 350 nm RMS. As in the previous section, the SNR is set to 1000 and is defined in the entrance pupil plane, while the training set is composed of 10^5 samples. Training and validation losses for the low regime model




Figure 5.12 Generated Zernike distribution, following a centered Gaussian with a decreasing standard deviation over the modes. The colors show the sample count at the indicated WFE on the y-axis.




Figure 5.13 Metrics over the training epochs. *Left*: VAE loss for the training and validation sets. *Middle*: Evolution of the phase residuals (on the training and validation sets). *Right*: Learning rate values decreased when reaching validation loss leveling off.

are shown in Figure 5.13 (left), as well as the monitored phase WFE and the decreasing learning rate through the epochs (Figure 5.13, middle and right, respectively).

5.4.3 Results and discussions

For each test data sample (100 in total), the models are evaluated with 100 forward passes to sample the posterior distribution. This ensures representative results and allows quantifying prediction uncertainties. If the final residual WFE needs to be expressed, it can be defined as the mean of the RMS WFE on the phase residuals over both the test data samples and the distribution samples:




Figure 5.14 Predicted phase maps (top) and reconstructed PSFs (middle and bottom) with the SimVAE method with post-VVC PSFs at an entrance SNR of 1000 and 350 nm RMS WFE over 50 modes. The reconstructions shown are the mean values (i.e., without photon noise) except for the input. One PSF sample is represented, while five resulting predictions are shown, corresponding to the 0th, 25th, 50th, 75th and 100th percentiles of the RMS WFE.

$$\langle \text{WFE}(\Phi) \rangle = \frac{1}{N_{\text{test}} N_{\text{dist}} N_{\text{pix}}} \sum_t \sum_d \sqrt{\sum_p^{N_{\text{pix}}} (\widehat{\Phi}_{t,d,p} - \Phi_{t,p})^2}, \quad (5.10)$$

with N_{test} , N_{dist} and N_{pix} the numbers of test samples, distribution samples and pixels, respectively. $\widehat{\Phi}$ represents the predicted phase maps from the latent space and Φ the true phase maps used for the simulations.

Examples of five predictions for a given data sample is shown in Figure 5.14. These predictions are selected according to their residual WFEs: different percentiles are chosen, from the worst prediction to the best one (from left to right). We can see that the predictions are visually almost identical to each other, showing that the predicted $\widehat{\sigma}_z$ are very small. This can be explained by the fact that the constraint put on ϵ to be sampled from $[0, 0.01]$

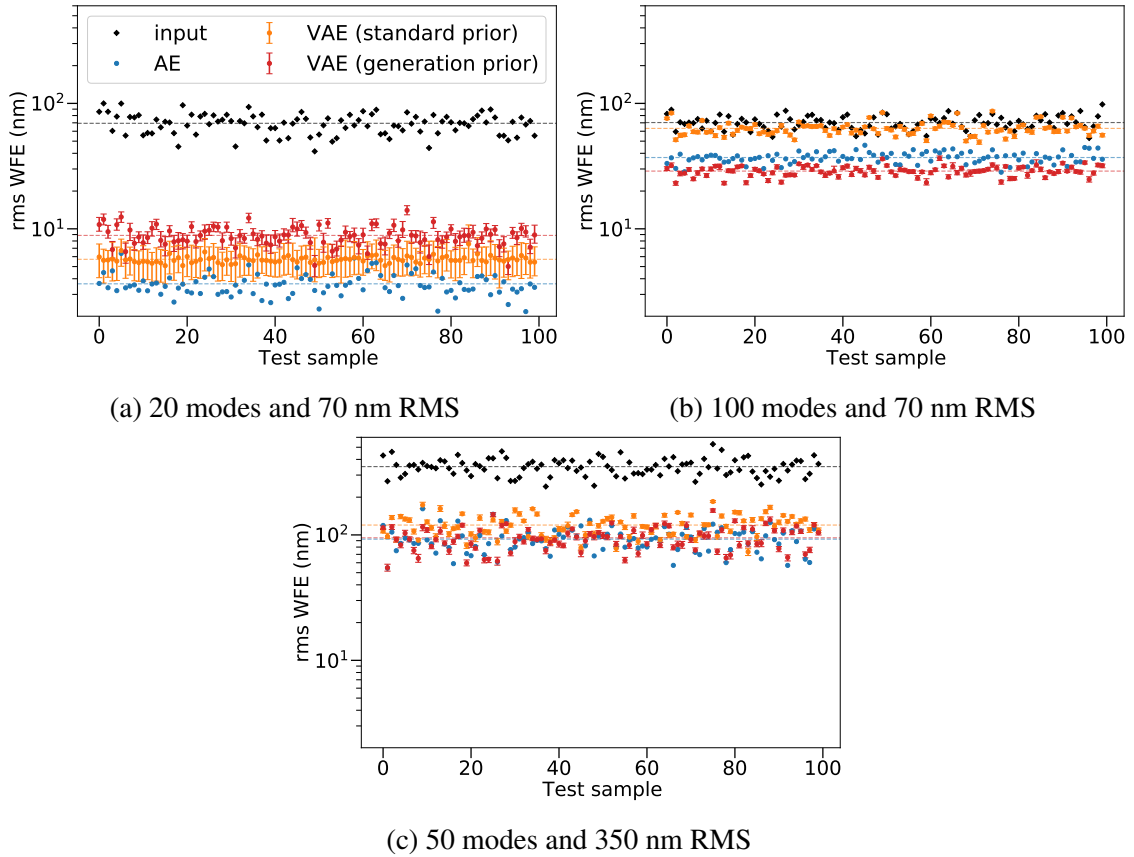


Figure 5.15 Residual RMS WFE for 100 test samples with the SimVAE, in different aberration regimes (a, b and c). Two prior distributions are considered: standard Gaussian distribution (zero mean and unit variance; orange) and the same Gaussian distribution as used to generate the dataset (red). The results with the deterministic approach (SimAE) are also shown (blue), as well as the input WFE (black). The median values with the 2–98 percentiles are displayed.

prevents much variability in the predictions. Sampling from a larger uniform range quickly deteriorates the performance though, especially in high aberration regimes.

The residual WFEs for each test sample are represented in Figure 5.15, where the median values as well as error bars are shown for the SimVAE models with two different priors: a standard prior $p_{\text{std}} = \mathcal{N}(0, 1)$ (orange) and a prior equal to the generated distribution $p_{\text{gen}} = \mathcal{N}(0, \sigma_{\text{gen}})$ (red). The standard deviation σ_{gen} is obtained from the generated distribution shown in Figure 5.12. Results with the deterministic autoencoder are also displayed (blue). Different behaviors are observed. With low input aberrations and low-order modes (Figure 5.15a), the prior always degrades the performance, especially if

it is close to the data generation distribution p_{gen} . In this aberration regime, the SimAE performance is already good to a point that the prior included in the KL divergence term only limits the posterior to achieve very accurate predictions. Considering higher order modes (Figure 5.15b), the performance strongly deteriorates, although the SimVAE with p_{gen} does not degrade as much as the other approaches. With this prior, the SimVAE benefits from the latent space regularization, since it is much harder to predict such complex phase maps without a priori guidance. Finally, in a large aberration regime (Figure 5.15c), the SimVAE with p_{gen} and the deterministic approach perform equally. This can be explained by the constraints put on ϵ , which makes the standard deviation low compared to the Zernike coefficient range in this regime. Increasing the range of ϵ prevents convergence during training, though. The three different types of models have also been trained on data containing 100 modes at 350 nm RMS, but all models failed to converge to a solution. Experiments performed on a fixed SNR in the focal-plane have nonetheless managed to provide some improvement in terms of residual WFE in this extreme aberration regime. It has been observed multiple times that the varying SNR on the PSFs within the dataset complicate the training process.

This early study of a variational approach show some good promises. In our case, the natural choice for the Zernike coefficients to select after inference and to apply for correction would be the predicted mean value, because it is associated with the highest probability. Alternatively, performing multiple inferences and selecting the median value could also be considered for increased robustness. Moreover, an application of such a variational model could prove to be particularly advantageous in a closed-loop system. Here, the control gain could be dynamically adjusted according to the prediction uncertainty, offering better stability during closed-loop correction. For instance, if the posterior distribution has a large standard deviation, the control gain could be minimized to prevent instability and ensure convergence.

5.5 Conclusions

In this chapter, a physics-informed deep learning method for focal-plane wavefront sensing has been developed, by incorporating an optical simulator into a CNN-based autoencoder architecture (SimAE). We observe very similar performance compared to using a simple CNN, both approaches reaching the expected theoretical limit over a large range of conditions. The SimAE, however, has the major advantage of not requiring labels during the training of the models.

One promising opportunity with an approach like the SimAE is that the model can be quickly fitted to the currently observed PSF using a pre-trained model. As shown

in this chapter, such fine-tuning of the models can greatly improve the results for data showing different noise or NCPA levels. Dealing with data the decoder cannot reproduce is naturally a potential limiting factor for the method. Models have been trained with atmospheric turbulence residuals present on top of NCPAs, and the loss of performance that we obtain stays well constrained. In practice these turbulence residuals may even be given to the decoder, by using the wavefront sensor telemetry to learn what the AO residuals effectively are, with the help of a small neural network for example.

To apply the technique on-sky with a real instrument, a more comprehensive simulator may be needed. Including a coronagraph for high-contrast imaging is notably motivated. In this direction, a vortex coronagraph simulator is incorporated into the architecture, demonstrating consistent performance with classical imaging. Optical parameters present in the decoder can also be trained together with the encoder weights. Tests with the vortex rejection factor show very good retrieval for a large range of simulated values.

A variational SimAE is also implemented in this chapter. The approach is notably tested with a prior distribution matching the data generation distribution. While it achieves better average performance in a high aberration regime, degradation is observed in the low aberration regime. Overall, such a variational approach offers prediction uncertainties, unlike deterministic approaches, and this could be exploited to provide better robustness in a wavefront control scenario.

The studies on simulations carried out in this chapter have demonstrated that a simulator-based autoencoder approach represents a compelling alternative to supervised learning models. As the approach performs very well across different aberration and noise regimes, the next step is to conduct tests on real data. This is proposed in the next chapter (Section 6.4).

PART III

Application to real data

In-lab experiments on SCE_xAO

After demonstrating deep learning-based focal plane wavefront sensing methods on simulated data, this chapter extends their application to real-world data. The SCE_xAO instrument from the Subaru Telescope is operated to build in-laboratory datasets and apply real-time wavefront corrections. A deep CNN is trained with such datasets in a supervised manner, while transfer learning from simulations illustrates its effectiveness on limited in-lab data. Closed-loop experiments are conducted, achieving consistently low and stable residual aberrations. Additionally, the simulator-based autoencoder is trained with SCE_xAO data, validating the method while raising remaining limitations. These results are expected to be included in a forthcoming journal publication.

Contents

6.1	The Subaru/SCE _x AO instrument	117
6.2	SCE _x AO data management	117
6.2.1	In-lab data acquisition	119
6.2.2	Pre-processing	122
6.2.3	Simulations and in-lab data comparison	122
6.3	Supervised learning tests	124
6.3.1	Performance in three aberration regimes	124
6.3.2	Transfer learning from simulations	126
6.3.3	Closed-loop performance	129
6.4	Simulator-based autoencoder	129
6.4.1	Model performance	129
6.4.2	Two-step approach: learning pre-existing aberrations	135
6.4.3	Variational inference	136
6.4.4	Method refinements	137
6.5	Discussion	139

6.1 The Subaru/SCEXAO instrument

The Subaru Coronagraphic Extreme Adaptive Optics (SCEXAO) instrument is installed on the infrared (IR) Nasmyth focus of the Subaru Telescope in Mauna Kea, Hawaii (Figure 6.1a and 6.1b). The Subaru Telescope has a primary mirror of 8.2 m and contains a facility adaptive optics system to correct the atmospheric turbulence. Additionally, SCEXAO comprises a second-stage AO system with a 2000-actuator DM to reach extreme AO correction (Lozi et al., 2018b). The instrument is dedicated to science, i.e., to direct imaging of exoplanets and protoplanetary disks, but also to technology development. It thereby allows external groups to test new technologies in-laboratory and on-sky. Thanks to this opportunity, a collaboration with the SCEXAO team has been initiated during this thesis to test the deep learning-based methods for focal-plane wavefront sensing in the lab.

The work presented in this chapter is carried-out using the visible module VAMPIRES. The optical setup of VAMPIRES as integrated in SCEXAO is shown in Figure 6.1c (dashed box). A super continuum laser is used as the light source. It is fed to the IR bench using an optical fiber (bottom left). The light is directed towards a DM used to control phase aberrations. A dichroic then only reflects the visible light, and a periscope redirects the beam to the upper visible bench. In VAMPIRES, the signal is then captured by two cameras simultaneously using a polarization beamsplitter. During science observations with differential polarization or H- α imaging, both VAMPIRES cameras are in-focus. For our FPWFS task, however, one camera is set out-of-focus in order to obtain phase diversity (see previous chapters).

A pupil stop is applied on SCEXAO, as shown in Figure 6.2 (left). It masks the central area covering the obstruction of the secondary mirror ($\sim 30\%$ of the total pupil area), as well as the spider structure holding the mirror. Two small masks also cover dead actuators, with one of them being supported by an additional spider segment. This pupil stop is located downstream of the DM, as shown in Figure 6.1c (red vertical bar).

6.2 SCEXAO data management

Several datasets have been acquired on SCEXAO, beginning in 2021 and continuing into 2024 for further experimentation. This section details the data acquisition process and provides descriptions of the various datasets.

CHAPTER 6. IN-LAB EXPERIMENTS ON SCE_xAO

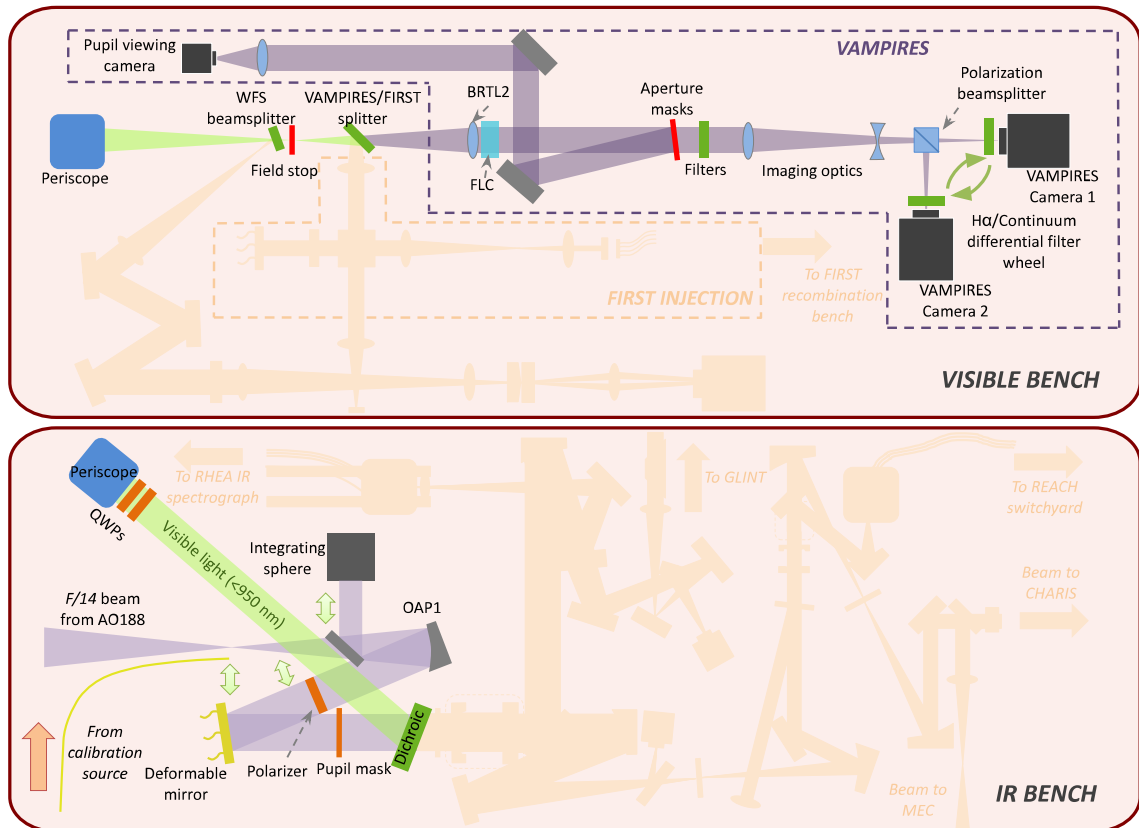



Figure 6.1 (a): View of the Subaru telescope with its dome open. (b): The SCE_xAO instrument as it is installed on Subaru. (c): Optical diagram of the SCE_xAO visible and IR benches. The non-relevant modules are faded-out. Credits: NAOJ (a, b) and [Lucas et al. \(2022\)](#) (c).




Figure 6.2 *Left*: SCEXAO pupil stop installed downstream of the deformable mirror. *Right*: Flat map applied on the DM for calibration.

6.2.1 In-lab data acquisition

The data is acquired on SCEXAO remotely using a VPN-secured access. To interact within the VAMPIRES cameras and the DM, the pyMILK package¹ is employed. This package enables interaction between custom Python scripts and the shared memory structure of the Multi-purpose Imaging Libraries toolKit (MILK). The communication with the DM can be done through different channels. For the experiments described in this chapter, one DM channel is dedicated to applying the phase maps on the DM, allowing custom aberrations to be introduced in the system. Another channel is used to apply the corrections after the phase map predictions, as described in Section 6.3.3. Ultimately, all channels interact with the same DM, which is located in the SCEXAO IR bench (Figure 6.1c, bottom).

After switching on the light source, the PSFs are inspected to ensure they are properly imaged on the two cameras (dichroic in place, adequate optical alignment). To correct most of the static aberrations on the instrument, a flat map is applied to the DM before data acquisition (Figure 6.2, right). One of the two cameras is then displaced out-of-focus by a given amount (e.g, 3 mm for the recent datasets). The integration time is set (typically 10 ms) to ensure both a sufficient signal-to-noise ratio and a rapid data generation procedure.

Early in-lab datasets

Three datasets in different aberration regimes have been first used (in early 2021), obtained by injecting 10^4 different random phase aberrations maps, with a median RMS WFE of 50, 130, and 260 nm for each of the three datasets (Figure 6.3). Each phase map is distributed uniformly over 20 Zernike modes, starting from the defocus mode (i.e., discarding tip and

¹<https://github.com/milk-org/pyMilk>




Figure 6.3 WFE distributions for each of the three early datasets, in the low, medium and high aberration regimes (blue, orange and green, respectively).

tilt aberrations). The phase map generation from Zernike coefficients follows Equation 1.6 and a distribution around a given median WFE for the full dataset is achieved with Equation 3.1. These phase maps are adjusted to fit the DM orientation and shape (Figure 6.4, left). The light source is in the visible centered at 750 nm, with a 50 nm bandwidth.

The second camera is moved out-of-focus to a corresponding -155 nm RMS of pupil-plane phase aberrations. The cameras installed back then were electron-multiplying CCDs with frame rates ranging between 1 and 100 Hz and low read-out noise levels (Lucas et al., 2022). The cameras presented pixel scales of 5.82 milliarcseconds per pixels (mas/pix) and 8.20 mas/pix, for the in-focus and out-of-focus PSFs, respectively. Additionally, there are orientation angles between pupil and focal plane on SCE_xAO: -0.55 and 2.59 radians for the in-focus and out-of-focus PSFs, respectively. The in-focus PSFs are also flipped horizontally with respect to the pupil plane.

After acquisition, the datasets are split so that 9×10^3 samples are used for training and validation, and 10^3 samples are dedicated to testing. A sample from each of the aberration regime are shown in Figure 6.4. The three datasets are used throughout the chapter, except for the closed-loop experiments of Section 6.3.3.

More recent datasets

By the time of the acquisition of the new datasets (mid-2024), the cameras installed on SCE_xAO had been upgraded as CMOS detectors with ultra-low read noise. The pixel scales are now almost identical between the two cameras: 6.04 mas/pix and 6.03 mas/pix (in- and out-of-focus, respectively, with an uncertainty of 0.29 mas/pix for each). The rotation between the pupil and focal planes are also different: -41.4 ± 3.1 and -40.6 ± 3.1 de-

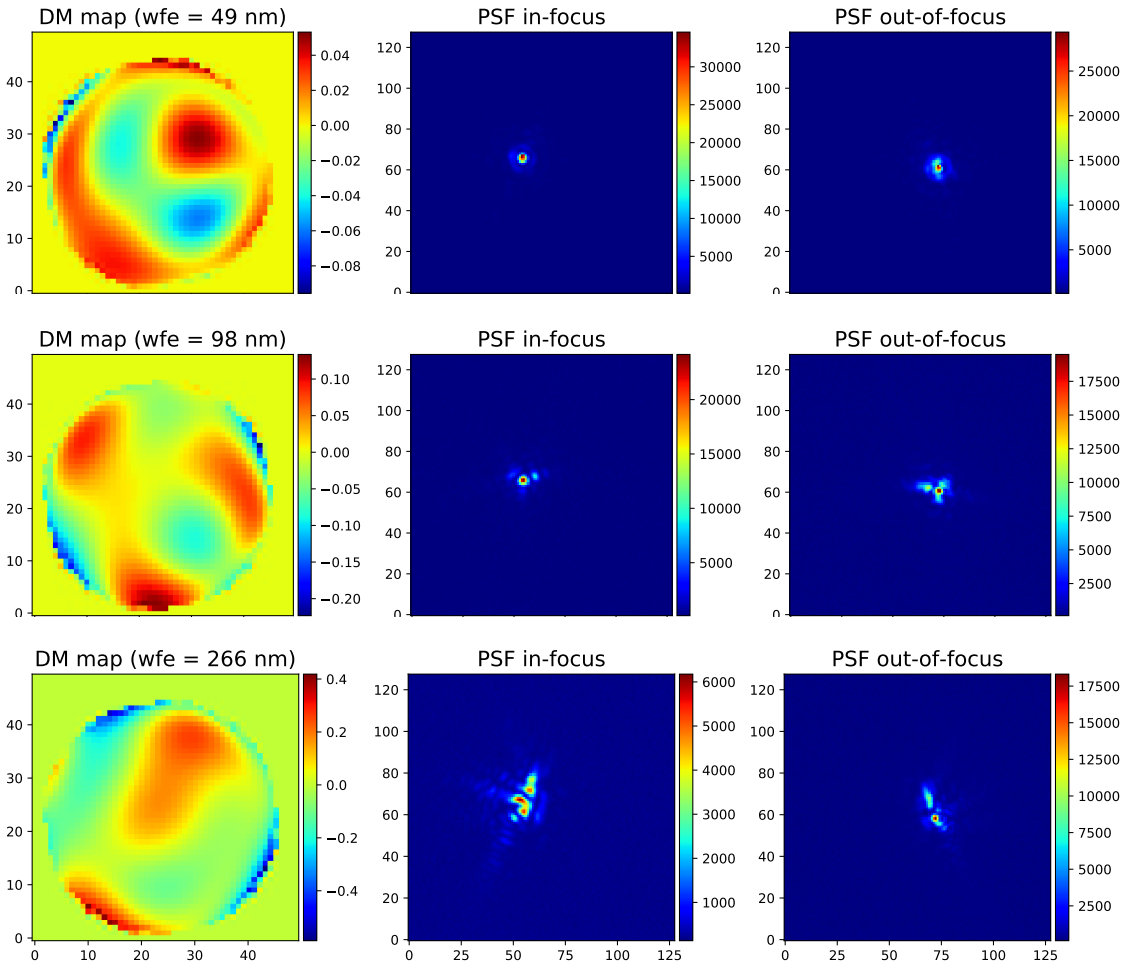


Figure 6.4 Raw in-lab SCEXAO samples at $\lambda = 750$ nm taken from the early datasets, at low, medium, and high aberration regimes (top, middle and bottom, respectively). The DM commands (left) are shown next to the resulting in-focus and out-of-focus PSFs (middle and right, respectively). The PSFs only underwent an initial cropping of 128×128 .

grees. Additionally, not only is the second camera flipped on the horizontal axis, but both cameras are also flipped along the vertical axis. The defocus applied is by a displacement of 3 mm, which translates to -245 nm RMS of pupil-plane aberrations. The observing wavelength and pupil mask remain unchanged.

Datasets are gathered in different aberrations regimes, one of them following a uniform WFE distribution between 0 and 350 nm RMS. Datasets with large WFE range have proven to be particularly robust in simulations (Sections 4.3.5 and 5.3). The aberrations

are also generated using 20 Zernike modes, while 10^4 training and 10^3 test samples are gathered in total. The uniformly distributed dataset is used for the closed-loop experiments in Section 6.3.3 and a dataset centered at 130 nm RMS is also used to train the SimAE and compare the results with the early dataset (Section 6.4.1).

6.2.2 Pre-processing

Pre-processing steps on the raw in-lab PSFs can be important for several reasons. First, removing instrumental effects helps to reduce the gap between lab data and simulations. This is important to attest the fidelity of the simulator and to prepare the implementation on the SimAE approach (Section 6.4). Secondly, reducing the dimensions of the data allows training models more efficiently.

The first calibration step is to subtract the dark current from the PSFs. There is one for each image channel, because two cameras are used in VAMPIRES to capture the in- and out-of-focus PSFs. Dark frames represent a bias on the images, containing electronic signal from the cameras even when they are not exposed to light. Subtracting the dark frame removes the corresponding dark current but not its noise. The dark noise is due to thermal agitation within the sensor, as well as shot noise due to the fluctuation of the dark current, and readout noise which occurs when the signal is read by the sensor. The combination of these random processes typically results in the dark noise following a Gaussian distribution.

Additionally, the PSFs from the two cameras are rotated and flipped respects to the pupil plane following the orientations provided in the previous section.

The standard VAMPIRES frames typically contain 500×500 pixels, while the PSF signature covers only a fraction of the field of view (Figure 6.4). Cropping the images to decrease their dimensions is thus possible without impacting the phase retrieval capabilities. The “center of mass” of the images is first found to crop around the center of the PSFs. As in the previous chapters, a 64×64 cropping is performed.

6.2.3 Simulations and in-lab data comparison

Beyond dark removal, rotating, flipping, centering and cropping the in-lab PSFs, adding noise content on simulations can improve the similarity of simulations compared to real data. Regarding the photon noise, since it is considered as dominant, the SNR can be approximated as $\text{SNR} = \sqrt{\sum_i N_i}$, with N_i the flux at pixel i . The SNR computed among all the lab samples of the early datasets is 862 ± 9 and 977 ± 4 for the in and out-of-focus PSFs, respectively. Adding shot noise (Poisson distributed) with an SNR of 1000 is therefore a good approximation. Although the dark current has been subtracted, dark noise should

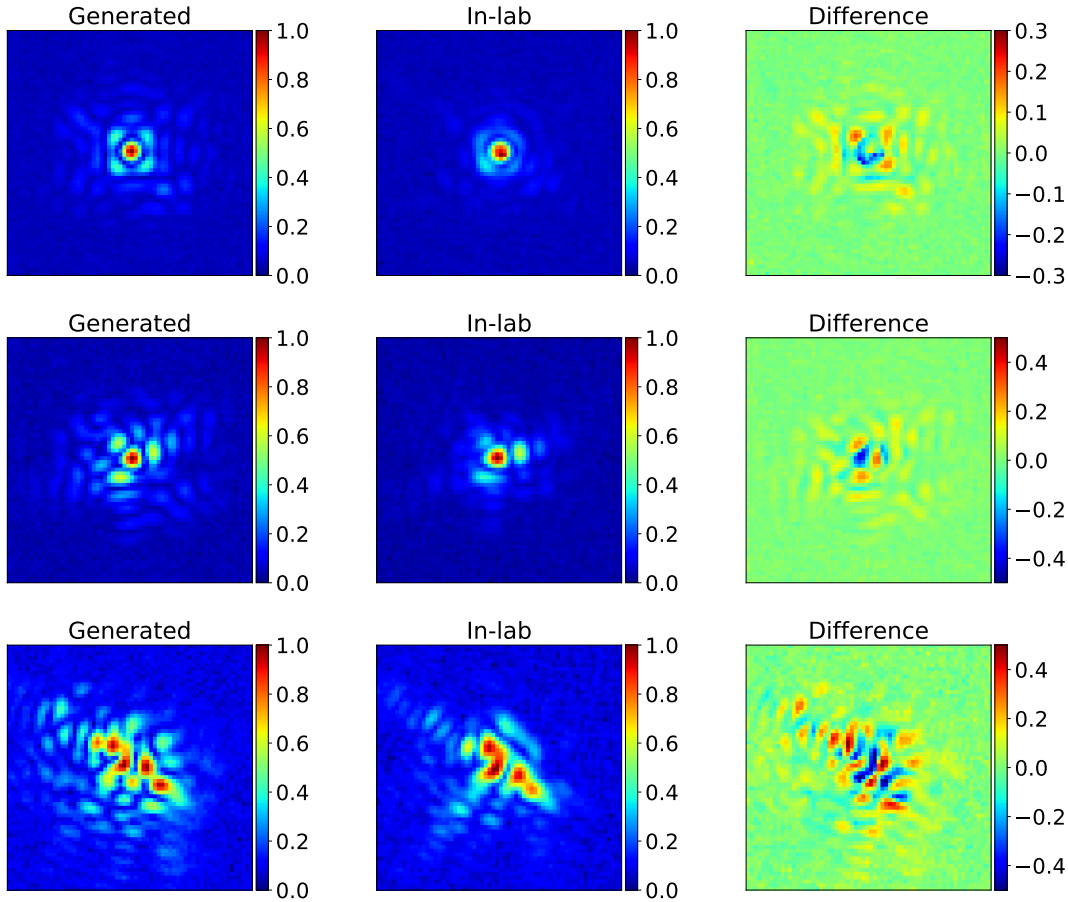


Figure 6.5 Examples of in-focus PSFs taken from each aberration dataset: 50, 130 and 260 nm RMS (top, middle, and bottom row, respectively). Simulated samples (left) are compared to their corresponding pre-processed in-lab PSF (middle). The difference between them is also shown (right).

also be considered. It is obtained as $\sigma_{\text{dark}} = \sqrt{N_{\text{dark}}}$, where N_{dark} is the average flux of the dark frames. It allows computing two Gaussian maps, zero-centered with a standard deviation equal to $\sigma_{\text{dark}-1}$ and $\sigma_{\text{dark}-2}$ for each camera. These noise maps are then added to the in-lab PSFs.

Simulations and laboratory data are compared in Figure 6.5, for the same samples as the ones presented in Figure 6.4. The phase maps used to produce the simulated PSFs are built from the same Zernike coefficients that were employed to create the DM commands responsible for the in-lab PSFs. All the images are scaled with a min-max normalization for consistency. While the simulations display similar structure to the in-laboratory data,

some discrepancies remain. There are several factors creating these differences.

Firstly, the in-lab data provided have as labels the Zernike coefficients used for the DM commands. Since some level of phase aberrations are also present in the system itself, the total phase aberrations which produce the final PSF do not exactly correspond to these commands. For the SimAE approach, for instance, it is necessary to know the difference between command and these additional aberrations if we want to properly assess the performance of the models in terms of phase reconstruction.

The polychromaticity may also be considered in the simulations, although tests with the SimAE did not impact the reconstruction much when considering the light source bandwidth. Also, the pupil may not be uniformly illuminated. We can expect a Gaussian distributed intensity, and such pattern can be trained in the simulator. Another potential effect is the level of DM calibration performed, as the injected aberrations may not correspond exactly to the ones finally produced by the DM. The pixel scales also contain some uncertainty (which are provided for the recent datasets, see Section 6.2.1). Finally, there is a small decentering between the simulations and the pre-processed SCEXAO data. This is handled directly during training with the SimAE with additional tip-tilt modes predicted by the encoder. Some of these aspects are addressed in Section 6.4.

Ultimately, the various analysis performed seem to indicate that the pre-existing aberrations play the most important role in the discrepancy between simulations and in-lab data. The fact that a single pupil plane is considered also puts some limits to the completeness of the simulator and the optimization that can be performed.

6.3 Supervised learning tests

In this section, ResNet models are first trained in a supervised way on each of the three early datasets presented previously. As done in other chapters, the models take in-focus and out-of-focus PSFs as input to predict Zernike coefficients. Results obtained with the recent datasets align with those presented in this section.

6.3.1 Performance in three aberration regimes

The results in terms of RMS WFE for the three datasets are illustrated in Figure 6.6, for 10^3 test samples. Residual WFEs are on average around 15, 25, and 170 nm RMS for the low, medium and high aberration regimes, respectively. The performance with the supervised CNN approach appear relatively good on SCEXAO data, and more training data would improve the accuracy of the predictions.

As discussed in Section 6.2.3, there is a discrepancy between the DM command labels

6.3. SUPERVISED LEARNING TESTS

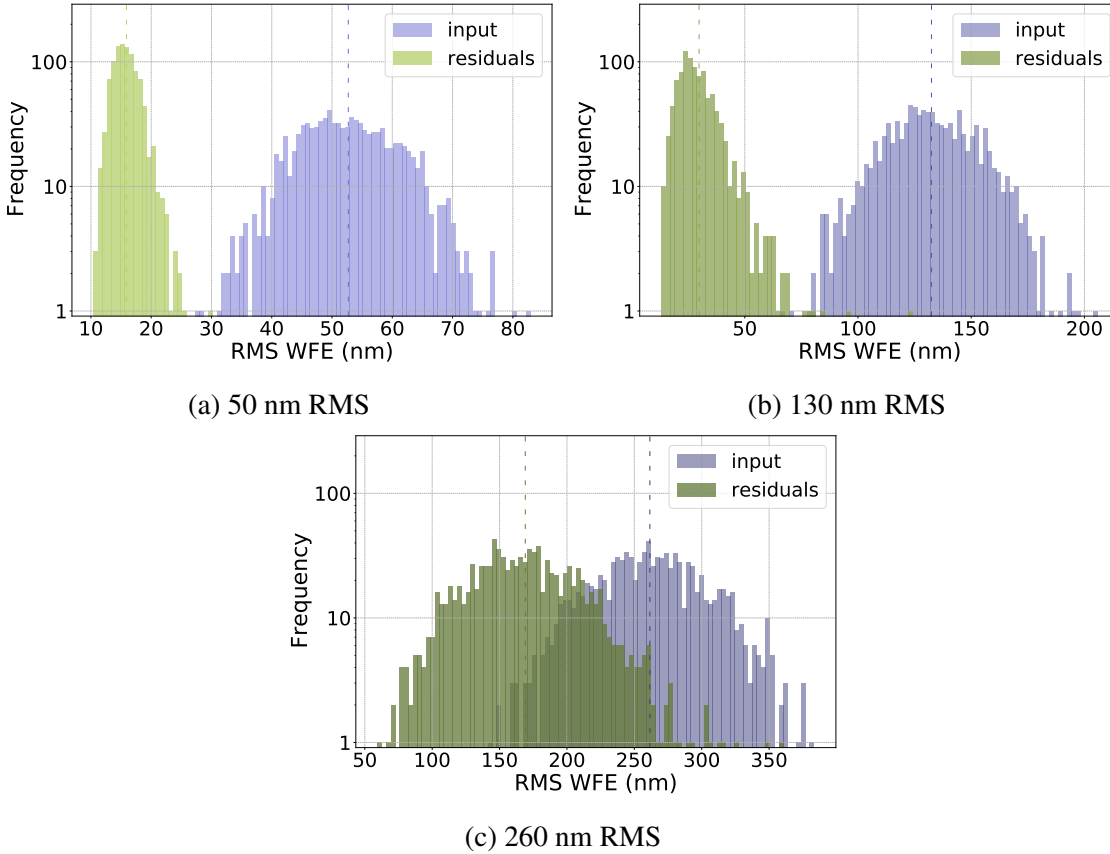


Figure 6.6 One-shot phase prediction errors on the test samples with SCE_xAO in-lab datasets in three aberration regimes: 50, 130 and 260 nm RMS. One model for each dataset is trained in a supervised manner.

and the true phase maps. During supervised training, the neural network effectively models this discrepancy, so it does not impact performance during model evaluation. However, this implies that while the predictions are accurate relative to the commands, they do not accurately represent the true wavefront present in the system. As a consequence, during control operations, the model predictions from the observed PSFs would present inaccuracies. For instance, in a situation where the instrument's state would be the same as when the datasets were gathered, the model would predict a zero correction, making it ineffective. Supervised models would therefore limit the correction capabilities on the instrument, hence the motivation for an unsupervised approach, as explored in Section 6.4.

The possibility to quickly re-train a model on a new data domain can also offer increased performance when few data are available. This type of approach is addressed in the next section.




Figure 6.7 Phase residual errors for models trained only with in-lab data (9×10^3 samples; blue), with simulated data (5×10^5 samples; red), and fine-tuning the simulation-based model with 9×10^3 in-lab frames (purple). A test set composed of 10^3 samples is used to make the predictions and compute the residual phase errors.

6.3.2 Transfer learning from simulations

Transfer learning techniques aims at re-training a pre-trained model on a different data domain than the previous one. This approach enhance models that have been pre-trained on large datasets, which is especially beneficial when the target domain data is limited and labels are difficult to obtain. As a result, significant improvements in performance and reductions in training time can be potentially achieved.

In this section, transfer learning between simulations and SCE_xAO data is explored, to demonstrate how simulations can help improve the performance of models based on small in-lab datasets.

Updating the entire network

Experiments are performed by re-training all the CNN layers at once. A first model is trained on a large dataset containing 5×10^5 simulated samples. As a second step, the model is initialized with the obtained pre-trained weights, and is then fine-tuned on in-lab datasets of several sizes. Figure 6.7 illustrates the obtained performance on the in-lab test sets, with a fined-tuned model re-trained on 9×10^3 SCE_xAO lab data. With a median

6.3. SUPERVISED LEARNING TESTS

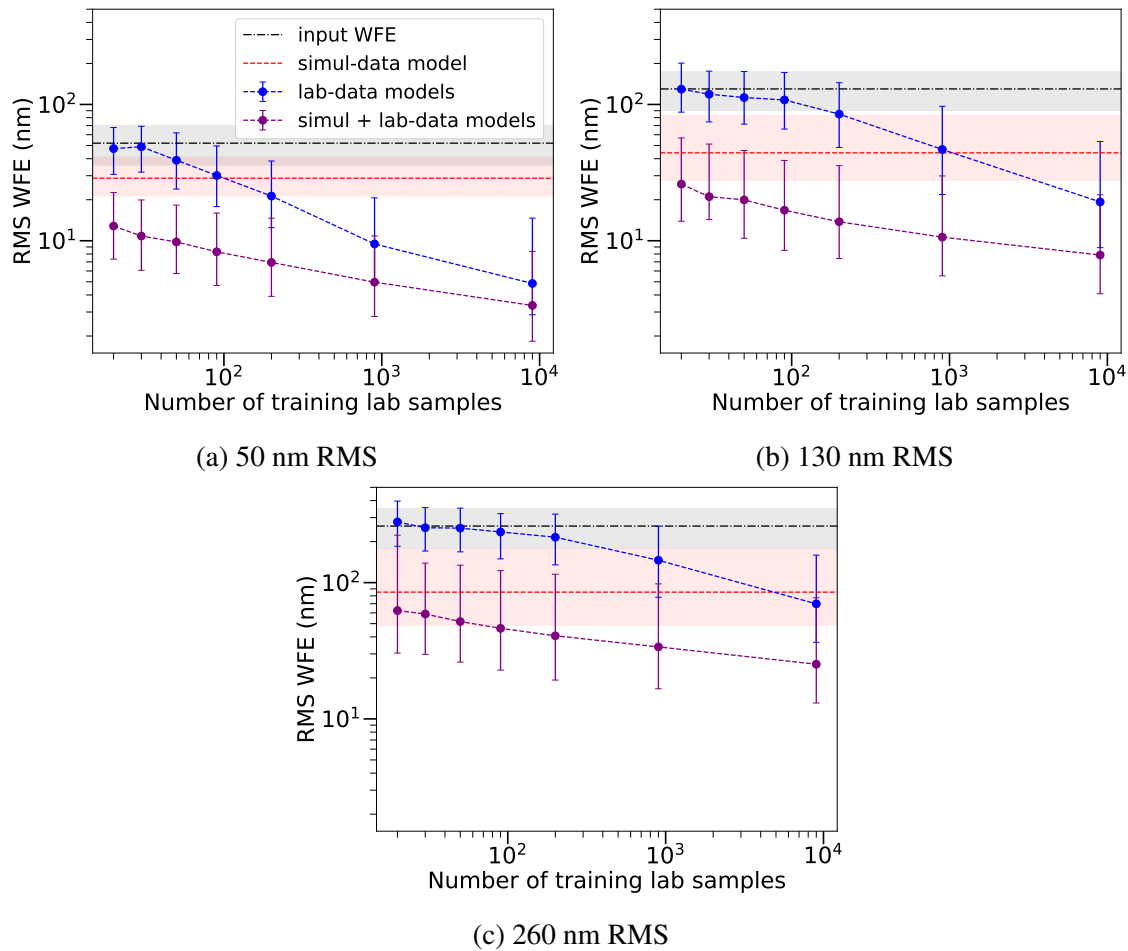


Figure 6.8 Performance comparison of models trained on lab data only (blue), trained on simulated data only (red), and fine-tuned on smaller amounts of lab data (purple). Models are trained on increasing dataset sizes along the x-axis, with median input WFE of 50 (a), 130 (b) and 260 (c) nm RMS. The evaluation is performed on 100 test samples and the 2–98th percentiles define the error bars.

input WFE of 130 nm RMS (gray), this new model allows reaching median residuals of 10 nm RMS (purple) instead of 19 nm RMS (blue). The performance is also compared to using the model trained on 5×10^5 simulated samples and performing tests with the in-lab dataset. Much higher phase residuals are then obtained (around 50 nm RMS, red).

To prove the benefits of such approach on smaller in-lab datasets, varying training dataset sizes are used. The evaluation is always performed on the in-lab test sets composed of 10^3 samples, and the results are shown in Figure 6.8 for the three aberration regimes. When the model is only trained on lab data, no phase correction is achieved us-

ing datasets of 10 and 20 lab samples, and the performance improves with larger datasets. As mentioned above, prediction errors are quite high when only training the models on the simulated data. However, by pre-training on the large simulated dataset and fine-tuning the model on lab data, important improvement is obtained when few in-lab data are used. For instance, in the medium aberration regime, re-training on 10 samples results in residuals around 30 nm RMS.

The trend is consistent in both low and high aberration regimes, as shown in Figure 6.8. This underscores the potential of fine-tuning a model directly on the instrument using only 10 acquired PSFs, for instance, each produced with different DM-generated aberrations.

Layer-wise tuning

Fine-tuning all the CNN layers at once may be the most straightforward way to transfer learning. However, selectively fine-tuning a subset of layers could lead to faster training times while maintaining similar performance. The Resnet-34 architecture, for instance, is organized into four layer groups that perform convolutions, batch normalization and ReLU activations. These layer groups are highlighted with different colors in Figure 2.3a. As an example, the first layer group of ResNet-34 contains about 6000 parameters, while the entire architecture comprises 13 millions parameters. To investigate the benefits of selectively fine-tuning specific layers, different sets of layer groups are re-trained from previously models that had been pre-trained on simulations, as described earlier.

Overall, loss of performance is observed when re-training on specific layers, although the drop is minor if only the first layer group is fined-tuned. This may be explained by the fact that the difference between data domains should be accounted for close to the input as it corresponds to small scale differences within the PSFs. Not updating the second layer group may still be limiting because its connection with the first group may no longer be smooth enough.

Besides, the gain in training time is negligible with Resnet-34 because the time taken by the backward pass through the model is under the time variation occurring between epochs. However, this approach could potentially help reduce training time when using very large architectures containing billions of parameters, although such extensive training may offer limited benefits for NCPA predictions using CNNs. Re-training the first layers would nonetheless be the strategy to adopt in this case, as the difference between data domains would arise in the focal-plane images at small spatial scales.

6.3.3 Closed-loop performance

As presented in Section 4.3.6, closed-loop corrections can help achieve increased performance. Experiments are performed on a recent SCExAO in-lab dataset (described in Section 6.2.1) to test the stability of the method. A dataset composed of 10^4 training samples is generated on the instrument with aberrations following a uniform WFE distribution between 0 and 350 nm RMS. A ResNet-50 model is then trained on this dataset directly on SCExAO. During the testing procedure, a phase map sampled from the same WFE distribution is applied to the DM to inject aberrations in the system via one channel of the SCExAO shared memory. The resulting PSFs obtained from the VAMPIRES cameras are fed successively to the trained model and each prediction is applied to the DM via another shared memory channel.


As illustrated in Figure 6.9a and 6.9b for a given input aberration case, the in-focus and out-of-focus PSFs are mostly free from speckles. The differences with the PSFs without added aberration (“original”) are very small. The residual WFEs for ten different cases (Figure 6.9c) reach a low and stable state at around 18 nm RMS on average. The benefits of using datasets containing WFEs uniformly distributed is highly motivated, and these experiments confirm the results obtained before with simulations.

6.4 Simulator-based autoencoder


The simulator-based autoencoder method presented and tested on simulations in Chapter 5 showed promising features, as it does not require labels to train the models, and it incorporates a known physical process into a deep learning architecture. In this section, the method is applied on SCExAO data to investigate whether the differentiable simulator is capable of reproducing the data with enough fidelity to achieve proper phase retrieval.

6.4.1 Model performance


ResNet-50 models are trained on the early datasets. The Poisson loss as defined in Equation 5.1 is used with in-lab data, since photon noise is for now considered to be dominant. The Figure 6.10 (top) showcases the results for the low and medium aberration regimes. The RMS WFE for each test sample is plotted with respect to its corresponding input WFE. The residual WFEs rise with the input aberrations, and poor predictions arise for a few outliers (up to ~ 100 nm RMS), while the minimum error is 23 nm RMS. Any imprecision in the simulator or the presence of pre-existing NCPAs may limit the performance, especially in the low aberration regime. The prediction error per Zernike mode is also



(a) In-focus PSFs



(b) Out-of-focus PSFs



(c) Evolution of phase residual errors for ten samples.

Figure 6.9 Closed-loop PSF corrections on the SCE_xAO bench over ten iterations. (a): In-focus PSFs through the successive correction steps for a randomly selected sample. The “initial” frame corresponds to the aberration state generated with the DM, and the “original” frame is the PSF without added aberrations (i.e., only pre-existing instrumental phase errors are present). (b): Same for the out-of-focus PSF. (c): The evolving RMS WFE over the iterations for ten different input WFEs. The iteration 0 corresponds to the input WFE.

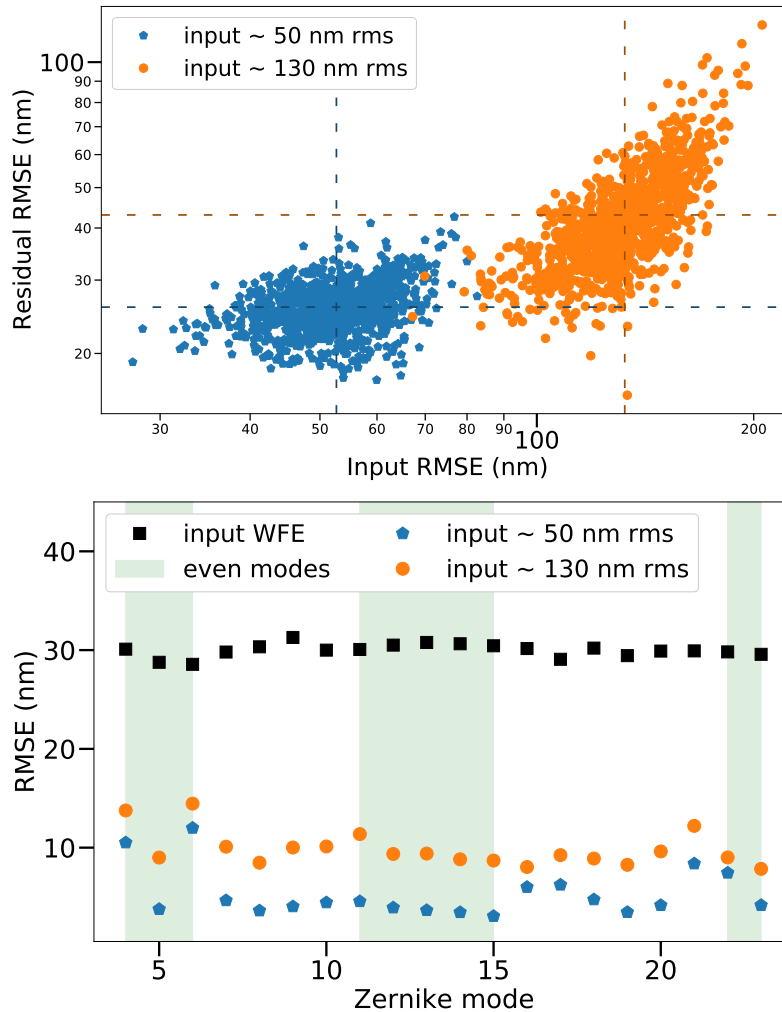


Figure 6.10 *Top*: Residuals errors with respect to the input WFE for each test sample. The mean value for the residuals and the input WFE are indicated in dashed horizontal and vertical lines, respectively. *Bottom*: Prediction error per Zernike mode. The results are obtained with the model trained on data at around 50 and 130 nm RMS (blue and orange, respectively).

shown in Figure 6.10 (bottom). The performance along the modes appear mostly constant, although two low-order modes and one high-order mode perform worse for both datasets. An example of predicted phase and reconstructed PSFs is shown in Figure 6.11.

The 260 nm RMS dataset, however, gives very poor phase reconstruction. A solution found to solve the problem is to use a model pre-trained on data containing fewer aberrations. A model initialized with pre-trained weights from the 130 nm RMS model yields

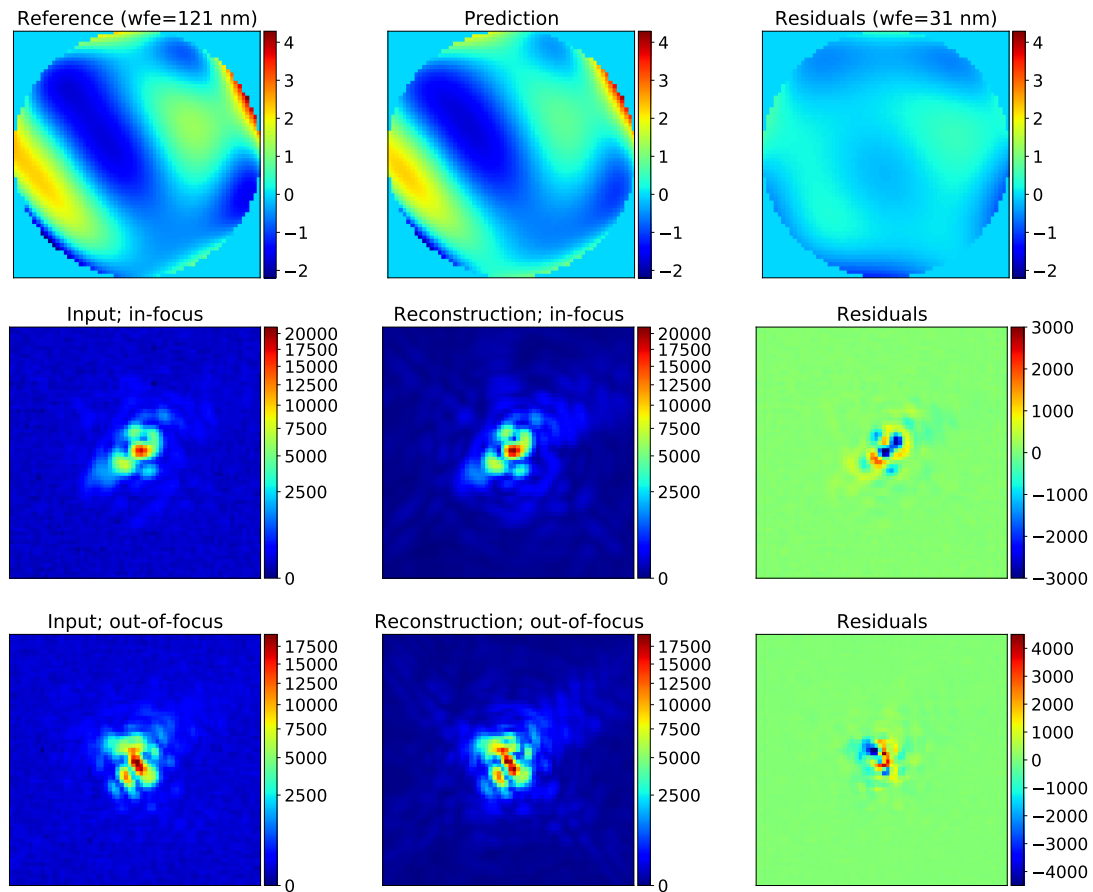



Figure 6.11 Example of predictions by the SimAE on SCE_xAO lab data for a given test sample. The dataset used to train the model contains a median 130 nm RMS of aberrations. *Top*: Phase map computed from the predicted coefficients by the encoder, together with the reference phase, which is not given to the SimAE during training. *Middle and Bottom*: Reconstructed in-focus and out-of-focus PSFs, respectively. These reconstructions are noiseless.

about 97 nm RMS of residuals, using 9×10^3 training and validation samples as well. Results are shown in Figure 6.12 between the model trained from scratch and the pre-trained one. The prospects for improvement in such aberration regime are very high, and more training data are required to leverage the method.

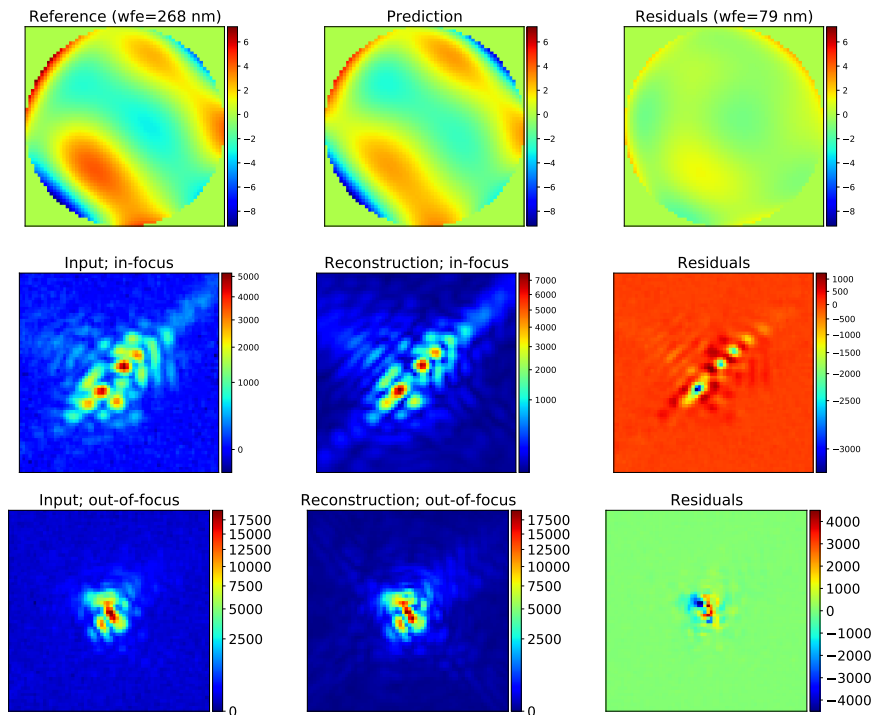
Regarding the recent data acquired, the model performance is not as good. The gain in terms of RMS WFE observed on a recent dataset centered at 130 nm RMS is only 13% on average, compared to a 67% decrease in residual WFE observed in the early dataset at

6.4. SIMULATOR-BASED AUTOENCODER



(a) Phase residuals

(b) Errors on the Zernike modes



(c) Input vs. reconstructions

Figure 6.12 Performance of the SimAE on SCE_xAO in-lab data with a median input WFE of 260 nm RMS. (a): Histogram of the phase residuals among the 1000 test samples. The performance is compared with and without using the model trained on 130 nm RMS to initialize the weights (red and orange, respectively). (b): Mean residual errors on each Zernike mode. (c): Inputs and predictions for one test sample: phase maps (top), in-focus and out-of-focus PSFs (middle and bottom, respectively).




Figure 6.13 Median phase maps computed from the test prediction set, for the low and medium regime datasets (left and middle, respectively). The difference between the two is also shown on the right. The color bars indicate values in nanometers.

the same aberration regime. Since the SimAE is trained on the PSFs, it is worth assessing the amount of PSF reconstruction error. To quantify this, the mean-square error (MSE) between the reconstructed and input PSFs is computed, with both PSFs scaled using min-max normalization for comparability across datasets. For the recent dataset, an average MSE of 9.6% is observed on the in-focus PSFs, with a relatively high standard deviation among the test samples (8%). In contrast, the model trained on the early dataset achieves an MSE of 1.8%, with a standard deviation almost equal to the average MSE. This clearly indicates that the SimAE cannot reconstruct the recent PSFs with as much fidelity as the early PSFs. To achieve more consistent results across datasets, further developments on the simulator’s robustness are therefore needed.

Even for the early datasets on which the SimAE performs well, the evaluation may be limited by the inaccuracy of the Zernike coefficient labels. It may therefore be beneficial to estimate the pre-existing phase aberrations present the dataset, i.e., the aberrations without adding ones with the DM. One potential way to achieve this is to recover the median of the predictions, assuming the pre-existing phase remained constant throughout the data acquisition process. Figure 6.13 shows the computed median phase for the low and medium regimes. The similarity between the phase maps suggest a common pattern across both datasets, corresponding to a phase error of 17 nm RMS. By subtracting this median map from the predicted NCPAs, the residual RMS WFE is reduced by 33% and 11% for the 50 nm and 130 nm RMS dataset, respectively.

Another way to account for these pre-existing phase errors is to learn them as parameters within the simulator in a two-step approach.




Figure 6.14 Learned pre-existing NCPAs during the pre-training procedure. *Left*: Learned phase aberrations for the in and out-of-focus paths (blue and red, respectively). *Right*: commands sent to the DM to produce the aberrations on the instrument (green) and total aberrations (blue and red for each optical paths).

6.4.2 Two-step approach: learning pre-existing aberrations

Because of the inherent NCPAs present in the system during data collection procedure, the SimAE would not be able to produce in the latent space the Zernike coefficients that corresponds to the DM commands. To be able to effectively evaluate the performance with labels, it is useful to first try to estimate these pre-existing aberrations.

One possible way to do that is to first train the simulator of the SimAE to learn the pre-existing NCPAs in the SCEXAO system. To do that, the known Zernike coefficients corresponding to the DM commands are defined as the input of the simulator. A second-phase training then consists in training the full SimAE (encoder + decoder) to predict these commands and have a proper evaluation using the labels.

The learned NCPAs are shown in Figure 6.14. The estimated amounts are very small compared to the DM commands, as expected. However, low-order aberrations appear as important as the higher-order ones, while we expect modes such as defocus, astigmatism or spherical aberrations to be stronger in an optical system. Moreover, by comparing the in-focus and out-of-focus paths, we can see that there are no consistent structure, while the differential path between the two is very limited on SCEXAO (one lens). It is therefore unlikely that these predicted NCPAs exactly represent the pre-existing aberrations on SCEXAO. Differences between the simulator and the real instrument beyond phase aberration most likely play an important role here. These learned NCPAs are probably compen-

sating for the difference between the simulator and the real instrument to a certain extent. Using the learned NCPAs in the evaluation process provides only for slight improvement in the estimated phase accuracy: an 11% decrease in terms of RMS WFE for the medium aberration regime dataset, for instance.

6.4.3 Variational inference

As described in Section 5.4, the variational approach allows predicting phase uncertainties by estimating distribution parameters. The early datasets in the three aberration regimes used in this chapter have been generated following a uniform distribution for each Zernike mode. The predicted distribution parameters from the encoder should therefore be the minimum a and maximum b of the uniform distribution. Since the Zernike coefficients used to generate the datasets were sampled from a uniform distribution, the reparametrization trick defined in Equation 5.6 is not valid in this case. The reparametrization trick rather becomes

$$z = a + (b - a) \odot \epsilon, \quad (6.1)$$

where ϵ is a random variable sampled from a uniform distribution defined in the range $[0, 0.01]$.

The prior distribution is also updated to follow a uniform one. The results are compared to the deterministic autoencoder in Figure 6.15 for the medium aberration regime. The residual WFE is shown for 100 different test samples, and 100 predictions are performed to obtain a median value and a standard deviation used to make error bars. As with simulations (Section 5.4), the median performance is similar between the two approaches (around 40 nm RMS).

It can be noticed that the median predictions from the VAE per test sample only rarely align with the AE predictions, as it has also been observed with simulations (Figure 5.15). The KL divergence has an impact on the model training, depending on the prior distribution considered. In some aberration regimes, this leads to reduced retrieval accuracy, while in other regimes, it results in more accurate outcomes (see Section 5.4). As a consequence, the VAE model behave differently from the AE when processing a given input sample, leading to variations in the results.

It can also be noticed that the error bars in the case of the SimVAE vary between samples. This added information is useful to account for the quality of the prediction, which could be taken into account to apply wavefront correction, as explained in Section 7.2.




Figure 6.15 Performance of the simulator-based autoencoder on a dataset centered around 130 nm RMS for 100 test samples. The variational approach (red) is compared to the deterministic one (blue).

6.4.4 Method refinements

The application of the SimAE to in-lab data has shown that the method can be successfully applied beyond simulations. However, performance can vary significantly across datasets. Several components of the method have started to be explored, and further enhancements are needed to increase the robustness of the SimAE.

Gaussian likelihood loss

Using the Poisson loss, as defined in Equation 5.1, may not be the most appropriate choice for real data due to the presence of additional noise sources. A more suitable approach can be to model the noise as having a global Gaussian distribution, considering not only photon noise but also thermal and readout noise from the cameras. As a consequence, a Gaussian negative log-likelihood loss would be appropriate. This loss function is defined

as

$$\mathcal{L}_{\text{SimAE}} = \frac{1}{2 N_{\text{pix}}} \sum_{i=1}^{N_{\text{pix}}} \left[\log(\widehat{\sigma}_i^2) + \frac{(x_i - \widehat{\mu}_i)^2}{\widehat{\sigma}_i^2} \right], \quad (6.2)$$

where $\widehat{\mu}_i$ is the flux per pixel given by the decoder and $\widehat{\sigma}_i$ is the standard deviation per pixel predicted by the encoder as an additional output.

Tests conducted on the in-lab data using this loss function reveal that the morphology of the predicted $\widehat{\sigma}$ closely resembles the reconstructed PSFs $\widehat{\mu}$. Moreover, the pixel values of the $\widehat{\sigma}$ maps are in the order of the square root of $\widehat{\mu}$. This suggests that the data predominantly contains Poisson noise. For images acquired with cameras producing more significant Gaussian-distributed noise, the loss function of Equation 6.2 would be more relevant.

Simulator parameter tuning

As previously discussed, a more realistic simulator than the one used in this chapter might be necessary to improve the SimAE performance. Several components of the simulator have already been explored and implemented in the method, including trainable simulator parameters. Firstly, inconsistencies in the centering of the PSFs between simulations and real data have been addressed by learning the tip-tilt modes in addition to the NCPAs being reconstructed. Moreover, the angle of rotation between the pupil and focal planes, although known, still presents some uncertainty. Differentiable rotation with a learnable angle has been implemented, ensuring that the reconstructed are properly aligned with the inputs.

Other parameters have been explored without much success. For instance, the non-uniformity of the pupil amplitude has been optimized. A plausible model for this pupil non-uniformity is a 2D Gaussian profile, since laser sources often display a Gaussian beam profile. Nonetheless, learning the mean and covariance of this Gaussian pattern within the simulator has resulted in only modest improvements in a few cases. An alternative is to acquire an image of the instrument’s pupil, which could be achieved in SCEXAO using the pupil viewing camera (see Figure 6.1c, top left). Furthermore, the uncertainty on the camera pixel scales can also affect reconstruction accuracy. These parameters cannot be directly optimized because they influence the frame padding size during propagation, which must be a floating-point value rather than a PyTorch tensor. As an alternative, adding minor blurring to the reconstructed PSF has been considered. This added blurring can nonetheless limit the performance due to the loss of fine details in the PSFs, and early tests did not yield noticeable improvements.

6.5 Discussion

The deep learning approaches developed in this thesis have been successfully applied to the SCEXAO instrument. Supervised CNNs have demonstrated good capabilities, with models being quickly trainable within the instrument system itself. However, the performance remains constrained due to the relatively small dataset sizes used in these studies. While generating larger datasets is feasible on SCEXAO, it makes experiments more cumbersome to replicate. Transfer learning has started to be investigated, and further studies are necessary to fully harness its potential. It will be crucial to explore how much it can be gained to transfer simulator-based autoencoder models from simulations to laboratory settings, and subsequently from the lab to on-sky scenarios. For closed-loop corrections, building in-lab datasets with a uniform distribution of wavefront errors has shown excellent results. Another WFE distribution being considered is a logarithmic distribution, which leans towards low aberration levels. This is justified because, in closed-loop corrections, residual aberrations quickly reach low aberration levels, and having models training to handle predominantly these WFE levels could enhance the corrected PSFs.

Assessing the capacities of the SimAE compared to a CNN should be further explored with the instrument. Because the SimAE has for reference an ideal simulator rather than an inaccurate baseline, it may be able to correct the pre-existing phase itself. However, this will ultimately depend on the simulator's fidelity. Experiments presented in this chapter have shown that PSF reconstruction should still be improved to obtain more consistent results over time. Future on-sky applications will be challenging given the dynamical nature of the atmospheric turbulence residuals combined with other sources of disturbance. To increase the robustness of the models, including synthetic AO residuals during dataset acquisition has notably started to be addressed. These AO screens could notably be integrated in the decoder of the SimAE. The deep neural network approaches developed in this thesis have the potential to efficiently adapt to changing conditions, making them highly valuable for achieving robust on-sky results.

PART IV

Conclusions

Conclusions

7.1 Summary

Instrumental phase aberrations, and in particular, non-common path aberrations, represent an inherent limitation for the detection of exoplanets. These phase errors deteriorate the performance of focal-plane coronagraphs, and the resulting speckles in the science images hinder the detection of faint companions close to their star. This thesis thereby explores new ways to estimate phase aberrations from focal-plane images using deep learning techniques. The idea behind deep learning with artificial neural networks is to build non-linear models that are efficient, fast, and applicable to any instruments.

It is first demonstrated in Chapter 3 that deep convolutional neural networks can be effectively trained to retrieve phase maps from simulated point-spread functions. These models perform well under different regimes, and they exhibit strong robustness to varying input phase aberrations, photon noise and spatial frequency contents. A particularly promising CNN variant examined in this chapter is the mixture density network (MDN), which provides prediction uncertainties and helps in monitoring the phase sign ambiguity.

Compared to standard iterative approaches, neural networks offer advantages such as increased robustness, faster predictions, and independence from initial guesses. A comparison with the Gerchberg-Saxton algorithm ([Orban de Xivry et al., 2021](#)) underscores the advantages of this deep learning approach. While other focal-plane wavefront sensing methods might offer superior solutions in some cases, making proper comparisons is challenging since they are usually tailored to specific instruments.

An approach leveraging the modulation introduced by vortex coronagraphs is notably developed in Chapter 4, allowing to solve the sign ambiguity while maintaining a 100% science duty cycle. This vortex phase diversity approach is tested on both scalar and vector vortex coronagraphs, across various aberration regimes and photon noise levels. In terms of performance, phase residuals under 1 nm RMS are achieved for input WFEs around 70 nm RMS when observing a magnitude-6.2 star in the K band ($\sim 10\%$ bandwidth). Atmospheric turbulence residuals included in the data only reduce the performance by a few nm RMS. These studies also reveal that training CNNs can be challenging with highly aberrated PSFs (e.g., 350 nm RMS over 88 Zernike modes).

In a real system, obtaining reliable and accurate labels for phase aberrations is challenging due to the presence of unknown pre-existing aberrations in the instrument. Unsupervised models are relevant to mitigate this problem because they can be trained solely on the observed data. A simulator-based autoencoder (SimAE) architecture developed in this thesis demonstrates that such models can readily be applied for FPWFS. The architecture contains a CNN as an encoder and a differentiable simulator as a decoder, reconstructing the observed PSFs and predicting the aberrations as Zernike coefficients in the latent space between the encoder and the decoder. Experiments with simulated datasets are presented in Chapter 5 using both classical and vortex phase diversity. The performance obtained with the SimAE is comparable to that of a standard CNN, with both approaches reaching the expected theoretical limit across a large range of conditions. However, the SimAE has the major advantage of not requiring ground truth phase information during the training of the models. Variational inference is also explored and implemented in the SimAE, by predicting posterior distribution for Zernike coefficients as well as incorporating prior knowledge about them. This provides uncertainties for each prediction, influencing potential decisions for NCPA corrections. While early results with the simulator-based variational autoencoder models are very encouraging, further studies are needed to fully harness the approach.

The application of these deep learning methods on a real instrument has been carried out, thanks to a collaboration with the Subaru/SCEXAO team started during this PhD thesis. Both supervised and unsupervised approaches have thus been tested on the VAMPIRES module at visible wavelengths, as presented in Chapter 6. Offline predictions have been successfully performed, by training models using datasets gathered in-laboratory, showing good phase retrieval performance. Pre-training models on large simulated datasets notably allows reaching better performance on small in-lab datasets. This represents great prospects to quickly fine-tune models on the bench and then on-sky. Closed-loop experiments on SCEXAO also display robust convergence, achieving residual WFEs of 15 nm RMS on average. The SimAE approach performs well in terms of PSF reconstruction and phase estimation. Since the decoder part of the architecture guides the learning instead of relying on labels, pre-existing aberrations do not affect phase estimation. Comparing the results with the supervised method is not straightforward, as the predictions in this case reflect the labels, which correspond to the DM commands, and not the actual aberrations present in the instrument. The SimAE approach still faces some limitations due to the simulator being imperfect. Recently gathered datasets have shown degraded performance, which motivates further studies on improving the robustness of the architecture and the simulator in particular. SCEXAO is a collaborative instrument that is regularly modified, causing the PSF quality to fluctuate over time. As a result, even

robust models might need to be retrained or at least fine-tuned based on previously trained models.

Deep learning models offer a flexible framework and fast inference speeds, which are appreciable features for on-sky applications. The requirement on speed is, however, not very stringent as we expect the lifetime of NCPAs that produce quasi-static speckles to be on the order of minutes. On-sky applications will naturally come with their own challenges and discrepancies unpredicted by simulations and laboratory setups. Future work should focus on enhancing the robustness of these models, refining transfer learning techniques, and exploring their integration with more complex simulators and real-time adaptive optics systems.

7.2 Perspectives

The results presented in this thesis highlight the capabilities of deep learning for focal-plane wavefront sensing. From these advancements, several promising avenues should be explored. On-sky validation of the methods is particularly crucial, and such tests are expected to be carried out on Subaru/SCE_xAO. There are several challenges arising with on-sky observations, as they introduce atmospheric turbulence residuals which dynamic can disturb NCPA sensing. Additionally, target stars are generally much dimmer than calibration sources used in laboratory settings, increasing the level of photon noise in the images. Beyond NCPA correction, addressing the low-wind effect (LWE) is also important. The LWE represent a significant limitation for SCE_xAO, producing low-order aberrations on the images that standard wavefront sensors cannot detect and correct ([Vievard et al., 2019](#)). The deep learning-based FPWFS techniques developed in this thesis are very promising to address this issue as well.

The vortex phase diversity approach developed in this thesis will also undergo laboratory tests in the near future. Potential applications of this proposed method could rely on including a polarizing beam splitter downstream of the VVC to collect both circular polarization states separately, either on a single or on two distinct sensors. Work in this direction has already been initiated by the PSILab team on the Vortex Optical Demonstrator for Coronagraphic Application (VODCA) bench of the University of Liège, using a linear polarizer and a quarter-wave plate (QWP) to obtain polarized light behind a VVC. Upcoming in-lab experiments will offer a clearer perspective on potential challenges and limitations. Additionally, a new vector vortex phase mask has been recently installed in SCE_xAO. The future inclusion of a QWP downstream of the vortex phase mask, coupled with the already existing polarizing beam splitter in VAMPIRES, would allow to test the vortex phase diversity approach with a VVC in-lab and on-sky. Since the simulation-based

FPWFS experiments work well even with a single image obtained behind an SVC, it appears that this flavor of vortex coronagraph offers an interesting alternative to the VVC, notably because it would operate without any additional optical components. Advancements in scalar vortex technology with newly manufactured masks will offer a great opportunity to enhance the proposed phase diversity approach thanks to this polarization-independent vortex coronagraph flavor.

From the deep learning methodologies developed in this thesis, several promising avenues emerge. Regarding the variational approach, further enhancements can be made by improving robustness and incorporating uncertainties in the wavefront correction strategy. The SimVAE has been implemented in a static scenario, which has not provided much additional benefit compared to a deterministic approach. However, for control applications in an adaptive optics system, inferring distributions might allow to select the best solution based on the smallest standard deviation among predictions. The control gain could then be adjusted according to the level of uncertainty. Additionally, upgrading the simulator to incorporate all optical layers and components present in SCExAO would make the reconstructed PSFs more realistic, thereby improving the performance of this autoencoder approach. Furthermore, using a differentiable optical package such as δ Lux (Desdoigts et al., 2023), could facilitate broader adoption and customization of the architecture within the community.

Beyond VAEs, recent generative modelling techniques could also be exploited for FPWFS. Most particularly, diffusion models have shown great prospects for solving inverse problems (Ho et al., 2020; Yang et al., 2023). By transforming simple distributions into complex images, diffusion models can process noisy inputs into probabilistic predictions in a robust way. The potential benefits of using a reinforcement learning (RL) approach could also be explored. RL has already been successfully implemented for adaptive optics applications (Nousiainen et al., 2024; Pou et al., 2024; Gutierrez et al., 2024), employing different strategies to address non-linearities in the wavefront sensor and correct temporal errors. Applying RL to focal-plane wavefront sensing and NCPA correction could be explored, and it would be interesting to assess whether RL offers performance improvements over the methods developed in this thesis. However, unlike CNNs, RL models are difficult to train and highly sensitive to hyperparameters, which can make their implementation challenging.

Ultimately, an essential area for investigation is the potential presence of exoplanets in focal-plane images, as the phase inference may be impacted by such signals. Having robust models is therefore important, and training the models with data containing exoplanet signals could be necessary. Another solution involves the dual-polarization VVC approach presented in Chapter 4. By subtracting the two post-VVC images, the planetary signal

would then be cancelled, preventing any interference in the predictions.

Forthcoming extremely large ground-based telescopes and space observatories will present new challenges for focal-plane wavefront sensing. The substantial costs and increasing complexity of these instruments will complicate the implementation of standard model-based algorithms. In this context, the advantages offered by deep learning methods could be invaluable, unlocking new capabilities to directly detect and characterize faint exoplanets near their host stars.

Bibliography

- Absil, O., Milli, J., Mawet, D., et al. 2013, *A&A*, 559, L12
- Agarap, A. F. 2018, arXiv e-prints, arXiv:1803.08375
- Allan, G., Kang, I., Douglas, E. S., Barbastathis, G., & Cahoy, K. 2020a, *Optics Express*, 28, 26267
- Allan, G., Kang, I., Douglas, E. S., et al. 2020b, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 11443, *Space Telescopes and Instrumentation 2020: Optical, Infrared, and Millimeter Wave*, ed. M. Lystrup & M. D. Perrin, 1144349
- Alzubaidi, L., Zhang, J., Humaidi, A. J., et al. 2021, *Journal of Big Data*, 8
- Amara, A., & Quanz, S. P. 2012, *MNRAS*, 427, 948
- Andersen, T., Owner-Petersen, M., & Enmark, A. 2019, *Optics Letters*, 44, 4618
- . 2020, *Journal of Astronomical Telescopes, Instruments, and Systems*, 6, 034002
- Angel, J., Wizinowich, P., Lloyd-Hart, M., & Sandler, D. 1990, *Nature*, 348, 221
- Baldi, P. 2012, in *Proceedings of Machine Learning Research*, Vol. 27, *Proceedings of ICML Workshop on Unsupervised and Transfer Learning*, ed. I. Guyon, G. Dror, V. Lemaire, G. Taylor, & D. Silver (Bellevue, Washington, USA: PMLR), 37–49
- Barrett, T. K., & Sandler, D. G. 1993, *Appl. Opt.*, 32, 1720
- Batalha, N. M., Rowe, J. F., Bryson, S. T., et al. 2013, *ApJS*, 204, 24
- Baudoz, P., Boccaletti, A., Baudrand, J., & Rouan, D. 2006, in *IAU Colloq. 200: Direct Imaging of Exoplanets: Science & Techniques*, ed. C. Aime & F. Vakili, 553–558
- Baydin, A. G., Pearlmutter, B. A., Radul, A. A., & Siskind, J. M. 2017, *J. Mach. Learn. Res.*, 18, 5595–5637
- Belikov, R., Stark, C., Siegler, N., et al. 2023, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 12680, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 126802G
- Beuzit, J. L., Vigan, A., Mouillet, D., et al. 2019, *A&A*, 631, A155

- Bishop, C. 1994, Mixture density networks, Workingpaper, Aston University
- . 1995, Neural networks for pattern recognition (Oxford University Press, USA)
- Bonfils, X., Delfosse, X., Udry, S., et al. 2013, *A&A*, 549, A109
- Bos, S. P., Doelman, D. S., Lozi, J., et al. 2019, *A&A*, 632, A48
- Bos, S. P., Vievard, S., Wilby, M. J., et al. 2020, *A&A*, 639, A52
- Bostan, E., Heckel, R., Chen, M., Kellman, M., & Waller, L. 2020, *Optica*, 7, 559
- Cantalloube, F. 2016, Theses, Université Grenoble Alpes
- Cantalloube, F., Mouillet, D., Mugnier, L. M., et al. 2015, *A&A*, 582, A89
- Cantalloube, F., Por, E. H., Dohlen, K., et al. 2018, *A&A*, 620, L10
- Cantalloube, F., Farley, O. J. D., Milli, J., et al. 2020, *A&A*, 638, A98
- Cantero, C., Absil, O., Dahlgvist, C.-H., & Van Droogenbroeck, M. 2023, arXiv e-prints, arXiv:2302.02854
- Carbillet, M., Bendjoya, P., Abe, L., et al. 2011, *Experimental Astronomy*, 30, 39
- Carlomagno, B., Delacroix, C., Absil, O., et al. 2020, *Journal of Astronomical Telescopes, Instruments, and Systems*, 6, 035005
- Carter, A. L., Hinkley, S., Kammerer, J., et al. 2023, *ApJ*, 951, L20
- Chambouleyron, V., Sengupta, A., Salama, M., et al. 2024, *A&A*, 681, A48
- Charbonneau, D., Brown, T. M., Latham, D. W., & Mayor, M. 2000, *ApJ*, 529, L45
- Charbonneau, D., Brown, T. M., Noyes, R. W., & Gilliland, R. L. 2002, *ApJ*, 568, 377
- Chauvin, G., Lagrange, A. M., Dumas, C., et al. 2004, *A&A*, 425, L29
- . 2005, *A&A*, 438, L25
- Chauvin, G., Desidera, S., Lagrange, A. M., et al. 2017, *A&A*, 605, L9
- Chilcote, J., Konopacky, Q., De Rosa, R. J., et al. 2020, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 11447, *Ground-based and Airborne Instrumentation for Astronomy VIII*, ed. C. J. Evans, J. J. Bryant, & K. Motohara, 114471S
- Codona, J. L., Kenworthy, M. A., Hinz, P. M., Angel, J. R. P., & Woolf, N. J. 2006, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 6269, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, ed. I. S. McLean & M. Iye, 62691N
- Codona, J. L., Kenworthy, M. A., & Lloyd-Hart, M. 2008, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7015, *Adaptive Optics Systems*, ed. N. Hubin, C. E. Max, & P. L. Wizinowich, 70155D

BIBLIOGRAPHY

- Cohen, N., & Shashua, A. 2017, Inductive Bias of Deep Convolutional Networks through Pooling Geometry, arXiv:1605.06743
- Davies, R., Absil, O., Agapito, G., et al. 2023, *A&A*, 674, A207
- Defrère, D., Absil, O., Hinz, P., et al. 2014, in *Adaptive Optics Systems IV*, ed. E. Marchetti, L. M. Close, & J.-P. Véran, Vol. 9148, International Society for Optics and Photonics (SPIE), 91483X
- Delacroix, C., Forsberg, P., Karlsson, M., et al. 2012, *Appl. Opt.*, 51, 5897
- Delacroix, C., Absil, O., Forsberg, P., et al. 2013, *A&A*, 553, A98
- Delacroix, C., Absil, O., Orban de Xivry, G., et al. 2022, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 12187, Modeling, Systems Engineering, and Project Management for Astronomy X, ed. G. Z. Angeli & P. Dierickx, 121870F
- Desai, N., Llop-Sayson, J., Bertrou-Cantou, A., et al. 2022, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 12180, Space Telescopes and Instrumentation 2022: Optical, Infrared, and Millimeter Wave, ed. L. E. Coyle, S. Matsuura, & M. D. Perrin, 121805H
- Desai, N., Llop-Sayson, J., Jovanovic, N., et al. 2021, in *Techniques and Instrumentation for Detection of Exoplanets X*, ed. S. B. Shaklan & G. J. Ruane, Vol. 11823, International Society for Optics and Photonics (SPIE), 238 – 246
- Desai, N., Mawet, D., Bertrou-Cantou, A., et al. 2024, in *Space Telescopes and Instrumentation 2024: Optical, Infrared, and Millimeter Wave*, ed. L. E. Coyle, S. Matsuura, & M. D. Perrin, Vol. 13092, International Society for Optics and Photonics (SPIE), 1309221
- Desai, N., Potier, A., Ruane, G., et al. 2023, arXiv e-prints, arXiv:2309.04920
- Desdoigts, L., Pope, B. J. S., Dennis, J., & Tuthill, P. G. 2023, *Journal of Astronomical Telescopes, Instruments, and Systems*, 9, 028007
- Doelman, D. S., Por, E. H., Ruane, G., Escuti, M. J., & Snik, F. 2020, *PASP*, 132, 045002
- Doelman, D. S., Snik, F., Por, E. H., et al. 2021, *Appl. Opt.*, 60, D52
- Dohlen, K., Wildi, F. P., Puget, P., Mouillet, D., & Beuzit, J.-L. 2011, in *Second International Conference on Adaptive Optics for Extremely Large Telescopes*. Online at http://ao4elt2.lesia.obspm.fr, 75
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., et al. 2021, in *International Conference on Learning Representations*
- Dumont, M., Correia, C. M., Sauvage, J.-F., et al. 2024, arXiv e-prints, arXiv:2403.18712

- Femenía Castellá, B., Serabyn, E., Mawet, D., et al. 2016, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9909, Adaptive Optics Systems V, ed. E. Marchetti, L. M. Close, & J.-P. Véran, 990922
- Ferreira, F., Gratadour, D., Sevin, A., & Doucet, N. 2018, 2018 International Conference on High Performance Computing & Simulation (HPCS), 180
- Flasseur, O., Bodrito, T., Mairal, J., et al. 2024, MNRAS, 527, 1534
- Flasseur, O., Denis, L., Thiébaud, É., & Langlois, M. 2018, A&A, 618, A138
- Foo, G., Palacios, D. M., & Swartzlander, Grover A., J. 2005, Optics Letters, 30, 3308
- Fressin, F., Torres, G., Charbonneau, D., et al. 2013, ApJ, 766, 81
- Fried, D. L. 1966, Journal of the Optical Society of America (1917-1983), 56, 1380
- Galicher, R., Huby, E., Baudoz, P., & Dupuis, O. 2020, A&A, 635, A11
- Galicher, R., & Mazoyer, J. 2023, arXiv e-prints, arXiv:2302.10833
- Gebhard, T. D., Bonse, M. J., Quanz, S. P., & Schölkopf, B. 2022, A&A, 666, A9
- Gerchberg, R. W. 1972, Optik, 35, 237
- Gillon, M., Triaud, A. H. M. J., Demory, B.-O., et al. 2017, Nature, 542, 456
- Give'on, A., Kern, B., Shaklan, S., Moody, D. C., & Pueyo, L. 2007, in Astronomical Adaptive Optics Systems and Applications III, ed. R. K. Tyson & M. Lloyd-Hart, Vol. 6691, International Society for Optics and Photonics (SPIE), 66910A
- Give'on, A., Kern, B. D., & Shaklan, S. 2011, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 8151, Techniques and Instrumentation for Detection of Exoplanets V, ed. S. Shaklan, 815110
- Gomez Gonzalez, C. A., Absil, O., & Van Droogenbroeck, M. 2018, A&A, 613, A71
- Gonsalves, R. A. 1982, Optical Engineering, 21, 829
- Goodfellow, I., Bengio, Y., & Courville, A. 2016, Deep Learning (MIT Press), <http://www.deeplearningbook.org>
- Goodman, J. 2005, Introduction to Fourier Optics, McGraw-Hill physical and quantum electronics series (W. H. Freeman)
- Guerri, G., Daban, J.-B., Robbe-Dubois, S., et al. 2011, Experimental Astronomy, 30, 59
- Guizar-Sicairos, M., & Fienup, J. R. 2012, J. Opt. Soc. Am. A, 29, 2367
- Guo, H., Xu, Y., Li, Q., et al. 2019, Sensors, 19, 3533
- Gutierrez, Y., Mazoyer, J., Mugnier, L. M., Herscovici-Schiller, O., & Abeloos, B. 2024, arXiv e-prints, arXiv:2406.18143
- Guyon, O. 2005, ApJ, 629, 592

BIBLIOGRAPHY

- Guyon, O. 2011, in *European Physical Journal Web of Conferences*, Vol. 16, *European Physical Journal Web of Conferences*, 03001
- Guyon, O. 2018, *Annual Review of Astronomy and Astrophysics*, 56, 315
- Gómez González, C. 2017, PhD thesis, ULiège - Université de Liège
- Haffert, S. Y., Males, J. R., Ahn, K., et al. 2023, *A&A*, 673, A28
- He, K., Zhang, X., Ren, S., & Sun, J. 2015, *Deep Residual Learning for Image Recognition*, arXiv:1512.03385
- Herscovici-Schiller, O., Mugnier, L. M., Baudoz, P., et al. 2018, *A&A*, 614, A142
- Higgins, I., Matthey, L., Pal, A., et al. 2017, in *International Conference on Learning Representations*
- Ho, J., Jain, A., & Abbeel, P. 2020, in *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20* (Red Hook, NY, USA: Curran Associates Inc.)
- Hochreiter, S. 1998, *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 6, 107
- Hoeijmakers, H. J., Schwarz, H., Snellen, I. A. G., et al. 2018, *A&A*, 617, A144
- Hornik, K., Stinchcombe, M., & White, H. 1989, *Neural Networks*, 2, 359
- Huby, E., Baudoz, P., Mawet, D., & Absil, O. 2015, *A&A*, 584, A74
- Huby, E., Bottom, M., Femenia, B., et al. 2017, *A&A*, 600, A46
- Jehin, E., Gillon, M., Queloz, D., et al. 2011, *The Messenger*, 145, 2
- Jolivet, A., Orban de Xivry, G., Huby, E., et al. 2019, *Journal of Astronomical Telescopes, Instruments, and Systems*, 5, 025001
- Jorgenson, M. B., & Aitken, G. J. M. 1992, in *Adaptive and Learning Systems*, ed. F. A. Sadjadi, Vol. 1706, *International Society for Optics and Photonics (SPIE)*, 113 – 121
- Jovanovic, N., Absil, O., Baudoz, P., et al. 2018, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 10703, *Adaptive Optics Systems VI*, ed. L. M. Close, L. Schreiber, & D. Schmidt, 107031U
- Jurling, A. S., & Fienup, J. R. 2014, *J. Opt. Soc. Am. A*, 31, 1348
- Karlsson, M., & Nikolajeff, F. 2003, *Opt. Express*, 11, 502
- Keller, C. U., Korkiakoski, V., Doelman, N., et al. 2012, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 8447, *Adaptive Optics Systems III*, ed. B. L. Ellerbroek, E. Marchetti, & J.-P. Véran, 844721
- Kenworthy, M. A., Codona, J. L., Hinz, P. M., et al. 2007, *ApJ*, 660, 762

- Kenworthy, M. A., Snik, F., Keller, C. U., et al. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10702, Ground-based and Airborne Instrumentation for Astronomy VII, ed. C. J. Evans, L. Simard, & H. Takami, 1070246
- Keppler, M., Benisty, M., Müller, A., et al. 2018, *A&A*, 617, A44
- Kingma, D. P., & Ba, J. 2017, Adam: A Method for Stochastic Optimization, arXiv:1412.6980
- Koch, D. G., Borucki, W. J., Basri, G., et al. 2010, *ApJ*, 713, L79
- König, L. 2023, PhD thesis, ULiège - Université de Liège [Sciences], Belgium
- König, L., Palatnick, S., Desai, N., et al. 2023, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12680, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 126800Q
- König, L., Palatnick, S., Desai, N., et al. 2023, in Techniques and Instrumentation for Detection of Exoplanets XI, ed. G. J. Ruane, Vol. 12680, International Society for Optics and Photonics (SPIE), 126800Q
- Korkiakoski, V., Keller, C. U., Doelman, N., et al. 2014, *Appl. Opt.*, 53, 4565
- Krist, J. E. 2007, in Optical Modeling and Performance Predictions III, ed. M. A. Kahan, Vol. 6675, International Society for Optics and Photonics (SPIE), 250 – 258
- Kullback, S., & Leibler, R. A. 1951, *The Annals of Mathematical Statistics*, 22, 79
- Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. 2007, *ApJ*, 660, 770
- Lagrange, A. M., Gratadour, D., Chauvin, G., et al. 2009, *A&A*, 493, L21
- Lamb, M. P., Correia, C., Sivanandam, S., Swanson, R., & Zavyalova, P. 2021, *MNRAS*, 505, 3347
- Landman, R., & Haffert, S. Y. 2020, *Optics Express*, 28, 16644
- Landman, R., Keller, C., Por, E. H., et al. 2022, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12185, Adaptive Optics Systems VIII, ed. L. Schreiber, D. Schmidt, & E. Vernet, 1218589
- Lecun, Y., Bottou, L., Bengio, Y., & Haffner, P. 1998, *Proceedings of the IEEE*, 86, 2278
- Liaudat, T., Starck, J.-L., Kilbinger, M., & Frugier, P.-A. 2021, Rethinking the modeling of the instrumental response of telescopes with a differentiable optical model, arXiv:2111.12541
- Liu, X., Morris, T., Saunter, C., et al. 2020, *MNRAS*, 496, 456
- Liu, Z., Mao, H., Wu, C., et al. 2022, in Proceedings - 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (IEEE Com-

BIBLIOGRAPHY

- puter Society), 11966–11976, publisher Copyright: © 2022 IEEE.; 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022 ; Conference date: 19-06-2022 Through 24-06-2022
- Loshchilov, I., & Hutter, F. 2019, Decoupled Weight Decay Regularization, arXiv:1711.05101
- Lozi, J., Guyon, O., Jovanovic, N., et al. 2018a, *Journal of Astronomical Telescopes, Instruments, and Systems*, 4, 049001
- Lozi, J., Guyon, O., Jovanovic, N., et al. 2018b, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 10703, *Adaptive Optics Systems VI*, ed. L. M. Close, L. Schreiber, & D. Schmidt, 1070359
- Lucas, M., Bottom, M., Guyon, O., et al. 2022, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 12184, *Ground-based and Airborne Instrumentation for Astronomy IX*, ed. C. J. Evans, J. J. Bryant, & K. Motohara, 121844E
- Lyt, B. 1939, *MNRAS*, 99, 580
- Macintosh, B., Poyneer, L., Sivaramakrishnan, A., & Marois, C. 2005, in *Astronomical Adaptive Optics Systems and Applications II*, ed. R. K. Tyson & M. Lloyd-Hart, Vol. 5903, *International Society for Optics and Photonics (SPIE)*, 170 – 177
- Maire, A.-L., Huby, E., Absil, O., et al. 2020, *Journal of Astronomical Telescopes, Instruments, and Systems*, 6, 035003
- Marois, C., Correia, C., Véran, J.-P., & Currie, T. 2014, in *Exploring the Formation and Evolution of Planetary Systems*, ed. M. Booth, B. C. Matthews, & J. R. Graham, Vol. 299, 48–49
- Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *ApJ*, 641, 556
- Marois, C., Macintosh, B., Barman, T., et al. 2008, *Science*, 322, 1348
- Marois, C., Zuckerman, B., Konopacky, Q. M., Macintosh, B., & Barman, T. 2010, *Nature*, 468, 1080
- Martinache, F. 2010, *ApJ*, 724, 464
- Martinache, F. 2013, *Publications of the Astronomical Society of the Pacific*, 125, 422
- Martinache, F., Jovanovic, N., & Guyon, O. 2016, *A&A*, 593, A33
- Martinez, P., Kasper, M., Costille, A., et al. 2013, *A&A*, 554, A41
- Martinez, P., Loose, C., Aller Carpentier, E., & Kasper, M. 2012, *A&A*, 541, A136
- Mawet, D., Pueyo, L., Moody, D., Krist, J., & Serabyn, E. 2010, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7739, *Modern Tech-*

- nologies in Space- and Ground-based Telescopes and Instrumentation, ed. E. Atad-Ettedgui & D. Lemke, 773914
- Mawet, D., Riaud, P., Absil, O., & Surdej, a. 2005, *The Astrophysical Journal*, 633, 1191
- Mawet, D., Riaud, P., Baudrand, J., et al. 2006, *A&A*, 448, 801
- Mawet, D., Serabyn, E., Liewer, K., et al. 2009, *Opt. Express*, 17, 1902
- Mawet, D., Absil, O., Delacroix, C., et al. 2013, *A&A*, 552, L13
- Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, *ApJ*, 792, 97
- Mayor, M., & Queloz, D. 1995, *Nature*, 378, 355
- Mayor, M., Pepe, F., Queloz, D., et al. 2003, *The Messenger*, 114, 20
- Mazoyer, J., Baudoz, P., Galicher, R., Mas, M., & Rousset, G. 2013, *A&A*, 557, A9
- McGuire, P. C., Sandler, D. G., Lloyd-Hart, M., & Rhoadarmer, T. A. 1999, in *Scientific Applications of Neural Nets*, ed. J. W. Clark, T. Lindenau, & M. L. Ristig (Berlin, Heidelberg: Springer Berlin Heidelberg), 97–138
- Miller, K., Guyon, O., & Males, J. 2017, *Journal of Astronomical Telescopes, Instruments, and Systems*, 3, 049002
- Milli, J., Banas, T., Mouillet, D., et al. 2016, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9909, *Adaptive Optics Systems V*, ed. E. Marchetti, L. M. Close, & J.-P. Véran, 99094Z
- Milli, J., Kasper, M., Bourget, P., et al. 2018, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 10703, *Adaptive Optics Systems VI*, ed. L. M. Close, L. Schreiber, & D. Schmidt, 107032A
- Montavon, G., Samek, W., & Müller, K.-R. 2018, *Digital Signal Processing*, 73, 1
- Montera, D. A., Welsh, B. M., Ruck, D. W., & Roggemann, M. C. 1996, *Appl. Opt.*, 35, 4238
- Mugnier, L. M., Cornia, A., Sauvage, J.-F., et al. 2009, *Journal of the Optical Society of America A*, 26, 1326
- Naik, K. R., Wright, R. H., Claveau, D. D., Acton, D. S., & Knight, J. S. 2020, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 11448, *Adaptive Optics Systems VII*, ed. L. Schreiber, D. Schmidt, & E. Vernet, 114481H
- Nishizaki, Y., Valdivia, M., Horisaki, R., et al. 2019, *Opt. Express*, 27, 240
- Noll, R. J. 1976, *Journal of the Optical Society of America (1917-1983)*, 66, 207
- Norris, B. R. M., Wei, J., Betters, C. H., et al. 2022, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 12185, *Adaptive Optics Systems VIII*, ed. L. Schreiber, D. Schmidt, & E. Vernet, 1218530

BIBLIOGRAPHY

- Nousiainen, J., Engler, B., Kasper, M., et al. 2024, *Journal of Astronomical Telescopes, Instruments, and Systems*, 10, 019001
- Orban de Xivry, G., Quesnel, M., Vanberg, P. O., Absil, O., & Louppe, G. 2021, *MNRAS*, 505, 5702
- Paine, S. W., & Fienup, J. R. 2018, *Opt. Lett.*, 43, 1235
- Paul, B., Mugnier, L. M., Sauvage, J. F., Ferrari, M., & Dohlen, K. 2013, *Optics Express*, 21, 31751
- Paul, B., Sauvage, J. F., Mugnier, L. M., et al. 2014, *A&A*, 572, A32
- Peng, Y., Choi, S., Padmanaban, N., & Wetzstein, G. 2020, *ACM Trans. Graph.*, 39, doi:10.1145/3414685.3417802
- Perryman, M., Hartman, J., Bakos, G. Á., & Lindegren, L. 2014, *ApJ*, 797, 14
- Pope, B. J. S., Pueyo, L., Xin, Y., & Tuthill, P. G. 2021, *The Astrophysical Journal*, 907, 40
- Por, E. H. 2017, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 10400, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, ed. S. Shaklan, 104000V
- Por, E. H. 2020, *ApJ*, 888, 127
- Por, E. H., Pueyo, L., & Soummer, R. 2022, in *Space Telescopes and Instrumentation 2022: Optical, Infrared, and Millimeter Wave*, ed. L. E. Coyle, S. Matsuura, & M. D. Perrin, Vol. 12180, *International Society for Optics and Photonics (SPIE)*, 121805J
- Potier, A., Mazoyer, J., Wahhaj, Z., et al. 2022, *A&A*, 665, A136
- Potier, A., Galicher, R., Baudoz, P., et al. 2020, *A&A*, 638, A117
- Pou, B., Smith, J., Quinones, E., Martin, M., & Gratadour, D. 2024, *arXiv e-prints*, arXiv:2405.13610
- Prince, S. J. 2023, *Understanding Deep Learning* (MIT Press)
- Pueyo, L. 2016, *ApJ*, 824, 117
- Pueyo, L., Kay, J., Kasdin, N. J., et al. 2009, *Appl. Opt.*, 48, 6296
- Quesnel, M., Orban de Xivry, G., Absil, O., & Louppe, G. 2022a, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 12185, *Adaptive Optics Systems VIII*, ed. L. Schreiber, D. Schmidt, & E. Vernet, 1218532
- Quesnel, M., Orban de Xivry, G., Louppe, G., & Absil, O. 2020, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 11448, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 114481G
- Quesnel, M., Orban de Xivry, G., Louppe, G., & Absil, O. 2022b, *A&A*, 668, A36

- Racine, R., Walker, G. A. H., Nadeau, D., Doyon, R., & Marois, C. 1999, *PASP*, 111, 587
- Ragazzoni, R. 1996, *Journal of Modern Optics*, 43, 289
- Riaud, P., Mawet, D., & Magette, A. 2012a, *A&A*, 545, A151
- . 2012b, *A&A*, 545, A150
- Roberts, Lewis C., J., Perrin, M. D., Marchis, F., et al. 2004, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 5490, *Advancements in Adaptive Optics*, ed. D. Bonaccini Calia, B. L. Ellerbroek, & R. Ragazzoni, 504–515
- Roddier, F., & Roddier, C. 1997, *PASP*, 109, 815
- Ronneberger, O., Fischer, P., & Brox, T. 2015, *U-Net: Convolutional Networks for Biomedical Image Segmentation*, arXiv:1505.04597
- Rosenblatt, F. 1958, *Psychological Review*, 65, 386
- Rouan, D., Riaud, P., Boccaletti, A., Clénet, Y., & Labeyrie, A. 2000, *PASP*, 112, 1479
- Ruane, G., Mawet, D., Riggs, A. E., & Serabyn, E. 2019, in *Techniques and Instrumentation for Detection of Exoplanets IX*, ed. S. B. Shaklan, Vol. 11117, *International Society for Optics and Photonics (SPIE)*, 454 – 469
- Ruffio, J.-B., Macintosh, B., Wang, J. J., et al. 2017, *The Astrophysical Journal*, 842, 14
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. 1986, *nature*, 323, 533
- Sandler, D. G., Barrett, T. K., Palmer, D. A., Fugate, R. Q., & Wild, W. J. 1991, *Nature*, 351, 300
- Sarker, I. H. 2021, *SN Comput. Sci.*, 2, doi:10.1007/s42979-021-00815-1
- Sauvage, J. F., Mugnier, L., Paul, B., & Villicroze, R. 2012, *Optics Letters*, 37, 4808
- Sauvage, J.-F., Fusco, T., Guesalaga, A., et al. 2015, in *Adaptive Optics for Extremely Large Telescopes IV (AO4ELT4)*, E9
- Schmidt, J. D. 2010, *Numerical Simulation of Optical Wave Propagation (SPIE)*
- Schwieterman, E. W., Kiang, N. Y., Parenteau, M. N., et al. 2018, *Astrobiology*, 18, 663
- Seager, S., & Deming, D. 2010, *ARA&A*, 48, 631
- Selsis, F., Kasting, J. F., Levrard, B., et al. 2007, *A&A*, 476, 1373
- Shack, R., & Platt, B. 1971, in *Spring Meeting of the Optical Society of America*. Ed. by DS Chairman. *Optical Society of America*. Tucson, Arizona: *Optical Society of America*
- Shannon, C. E. 1948, *Bell System Technical Journal*, 27, 623
- Simonyan, K., & Zisserman, A. 2015, *CoRR*, abs/1409.1556
- Snik, F., Otten, G., Kenworthy, M., et al. 2012, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 8450, *Modern Technologies in Space-*

BIBLIOGRAPHY

- and Ground-based Telescopes and Instrumentation II, ed. R. Navarro, C. R. Cunningham, & E. Prieto, 84500M
- Soummer, R. 2005, *ApJ*, 618, L161
- Soummer, R., Hagan, J. B., Pueyo, L., et al. 2011, *ApJ*, 741, 55
- Soummer, R., Pueyo, L., & Larkin, J. 2012, *ApJ*, 755, L28
- Swanson, R., Lamb, M., Correia, C. M., Sivanandam, S., & Kutulakos, K. 2021, *MNRAS*, 503, 2944
- Swartzlander, Grover A., J. 2006, *Optics Letters*, 31, 2042
- Swartzlander, Grover A., J., Ford, E. L., Abdul-Malik, R. S., et al. 2008, *Optics Express*, 16, 10200
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. 2015, *CoRR*, abs/1512.00567, arXiv:1512.00567
- Tan, M., & Le, Q. 2019, in *Proceedings of Machine Learning Research*, Vol. 97, *Proceedings of the 36th International Conference on Machine Learning*, ed. K. Chaudhuri & R. Salakhutdinov (PMLR), 6105–6114
- Terreri, A., Pedichini, F., Del Moro, D., et al. 2022a, *A&A*, 666, A70
- Terreri, A., Pedichini, F., Kasper, M., et al. 2022b, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 12185, *Adaptive Optics Systems VIII*, ed. L. Schreiber, D. Schmidt, & E. Vernet, 1218587
- Tsapras, Y. 2018, *Geosciences*, 8, 365
- van Laarhoven, T. 2017, *L2 Regularization versus Batch and Weight Normalization*, arXiv:1706.05350
- Vievard, S., Bos, S. P., Cassaing, F., et al. 2019, in *6th International Conference on Adaptive Optics for Extremely Large Telescopes, AO4ELT 2019*, Quebec City, Canada
- Vigan, A., N'Diaye, M., Dohlen, K., et al. 2019, *A&A*, 629, A11
- Vigan, A., Fontanive, C., Meyer, M., et al. 2021, *A&A*, 651, A72
- Vigan, A., Dohlen, K., N'Diaye, M., et al. 2022, *A&A*, 660, A140
- Vilas, F., & Smith, B. A. 1987, *Appl. Opt.*, 26, 664
- Wang, F., Bian, Y., Haichao, W., et al. 2020, *Light: Science & Applications*, 9, 77
- Wang, J. J., Graham, J. R., Dawson, R., et al. 2018, *AJ*, 156, 192
- Watson, S. M., Mills, J. P., Gaiser, S. L., & Diner, D. J. 1991, *Appl. Opt.*, 30, 3253
- Wei, J., Norris, B., Betters, C., & Leon-Saval, S. 2023, arXiv e-prints, arXiv:2311.01716

BIBLIOGRAPHY

- Wong, A., Pope, B., Desdoigts, L., et al. 2021, *Journal of the Optical Society of America B Optical Physics*, 38, 2465
- Wu, Y., Guo, Y., Bao, H., & Rao, C. 2020, *Sensors*, 20, 4877
- Yang, L., Zhang, Z., Song, Y., et al. 2023, *ACM Comput. Surv.*, 56, doi:10.1145/3626235
- Ying, X. 2019, *Journal of Physics: Conference Series*, 1168, 022022
- Yip, K. H., Nikolaou, N., Coronica, P., et al. 2020, in *Machine Learning and Knowledge Discovery in Databases*, ed. U. Brefeld, E. Fromont, A. Hotho, A. Knobbe, M. Maathuis, & C. Robardet (Cham: Springer International Publishing), 322–338
- Zhang, G., Wang, C., Xu, B., & Grosse, R. B. 2018, *CoRR*, abs/1810.12281, arXiv:1810.12281
- Zhang, Y., Tino, P., Leonardis, A., & Tang, K. 2020, *IEEE Transactions on Emerging Topics in Computational Intelligence*, 5, 726
- Zimmerman, N. T., Eldorado Riggs, A. J., Jeremy Kasdin, N., Carlotti, A., & Vanderbei, R. J. 2016, *Journal of Astronomical Telescopes, Instruments, and Systems*, 2, 011012

PART V

Appendices

Pupil impact on the Vortex performance

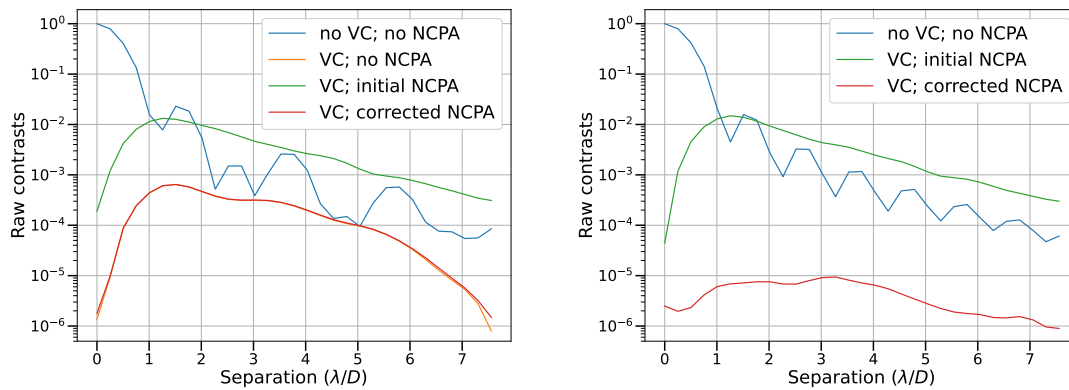


Figure A.1 Contrast curves for central obstructions of 15% (left) and 0% (circular aperture, right). The case without coronagraph and aberrations (blue) is compared to adding a vortex phase mask without aberrations (orange), at the initial NCPA level (green), and after NCPA correction (red).

As discussed in Section 4.12, the contrast achieved after wavefront control is limited for large central obscurations, which cause important flux leak through the vortex phase masks even if the wavefront is flattened. For instance, the secondary mirror of the Subaru Telescope masks 30% of the aperture, and contrasts obtained after closed-loop simulations with such an annular pupil design are shown in Figure 4.12. Smaller central obscurations are, however, more often encountered in telescopes, such as a 15% obscuration at the VLTs. Contrast results with such a pupil aperture are shown in Figure A.1 (left), using the same configurations as in the experiments in Section 4.12. The contrast gets much lower compared to Figure 4.12, although the drop in contrast at 4 λ/D is not present anymore. For a circular entrance pupil (Figure A.1, right), the contrast at intermediate separations reach values below 10^{-5} . The contrast for a non-aberrated wavefront is not displayed because all light is cancelled by the VVC (only numerical errors are present). The corresponding images after correction for both pupils are shown in Figure A.2.

APPENDIX A. PUPIL IMPACT ON THE VORTEX PERFORMANCE

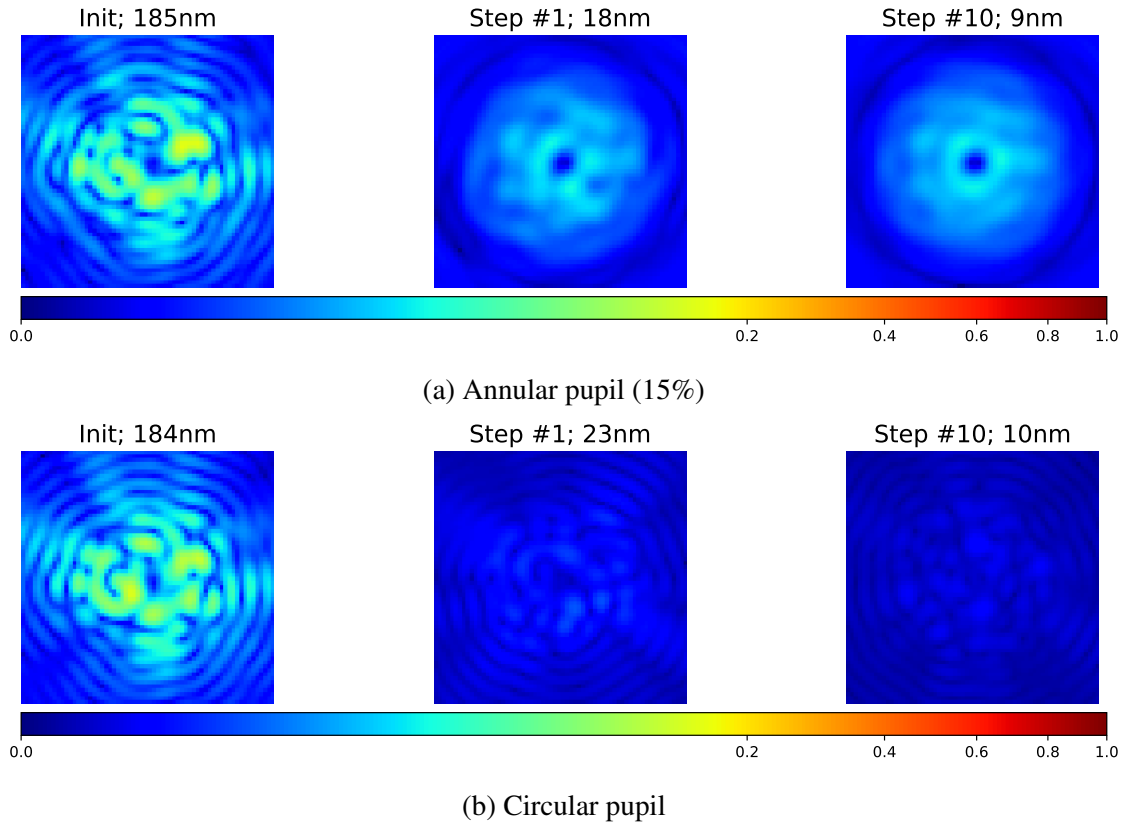


Figure A.2 Evolution of post-VVC PSFs through iterations with central obstruction of 15% (top) and 0% (bottom). *Left:* Initial PSF. *Middle:* PSF after first correction/iteration. *Right:* PSF after 10 iterations. *Top:* 70 nm RMS. *Bottom:* 350 nm RMS. The PSFs are shown after a power-law normalization to emphasize the residual speckles, and the wavefront aberration level of the PSFs is indicated at the top of each corresponding image.

Asymmetric pupil for the Vortex




Figure B.1 Aperture asymmetries of different sizes in the Lyot plane. *Left*: No obstruction. *Middle*: 10% obstruction. *Right*: 20% obstruction.

In the context of a project for the METIS instrument, which is being built for the Extremely Large Telescope, the CNN-based focal-plane wavefront sensing approach is tested behind a VVC with an asymmetric pupil defined in the Lyot plane. An annular pupil in the Lyot plane (before the Lyot stop) is first defined, with a 30% central obstruction. From this setup, three cases are tested: without asymmetry, with an asymmetric bar taking 10% of the pupil, and finally with a bar taking 20% of the pupil (Figure B.1, from left to right). Models are then trained on datasets generated with these Lyot pupils. The following data and parameters are considered: 20 and 100 modes, an input WFE of 300 nm RMS, SNRs of 100 and 1000 (detector plane), the Efficientnet-b0 architecture, 10^4 data samples, a batch size of 32 and a learning rate of 10^{-3} .

Results on each Zernike mode, using the metric of Equation 4.1 are shown in Figure B.2 for 20 and 100 modes. All modes are relatively well reconstructed, although the odd modes in the case of the asymmetric pupil show strong degradation when trained on 100 modes. Performance comparison on the phase residuals in the different cases are displayed in Table B.1. Using one PSF with an asymmetric pupil allows performing FPWFS. There is, however, a substantial degradation of the performance compared to using out-of-focus phase diversity. Increasing the obstruction to 20% yields slightly better performance in the case of 100 modes, although the optical throughput will be then decreased due to the larger mask. It can be also noticed that with 20 modes, increasing the obstruction reduces the performance, making a 10% mask more relevant overall.

APPENDIX B. ASYMMETRIC PUPIL FOR THE VORTEX

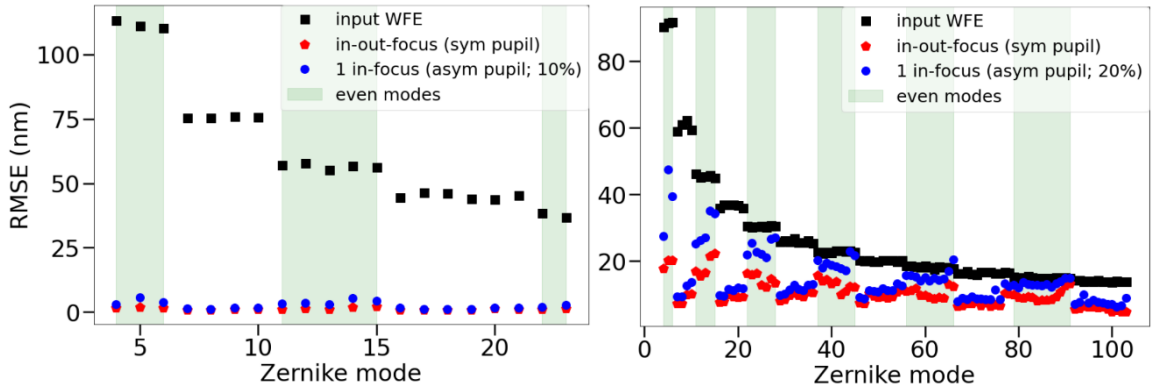


Figure B.2 RMS WFE for each Zernike mode with classical phase diversity (red) and with pupil asymmetry (blue). *Left*: 20 modes. *Right*: 100 modes. Results with 10% and 20% obstructions are shown in the case of 20 and 100 modes, respectively.

Dataset	Sym	Sym	Asym: 10%	Asym: 20%
	in-focus	in-out-focus	in-focus	in-focus
20 modes; SNR=1000	243	6	12	42
20 modes; SNR=100	242	32	53	45
100 modes; SNR=1000	252	30	94	80
100 modes; SNR=100	262	108	162	151

Table B.1 Residual RMS WFE on the phase residuals (in nm) with or without an asymmetric pupil. The median input WFE is 300 nm RMS.