

Data-driven modelling of ground-based and space-based telescope's point spread functions

*Modélisation guidée par les données des fonctions d'étalement
du point des télescopes terrestres et spatiaux*

Thèse de doctorat de l'université Paris-Saclay

École doctorale n°127 : astronomie et astrophysique d'Île-de-France (AAIF)
Spécialité de doctorat : Astronomie et Astrophysique
Graduate School : Physique, Référent : Faculté des sciences d'Orsay

Thèse préparée dans l'unité de recherche Astrophysique, Instrumentation et
Modélisation de Paris-Saclay (Université Paris-Saclay, CNRS, CEA), sous la direction
de Jean-Luc STARCK, directeur de recherche au CEA Paris-Saclay, la co-direction de
Martin KILBINGER, directeur de recherche au CEA Paris-Saclay

Thèse soutenue à Paris-Saclay, le 21 octobre 2022, par

Tobias Ignacio LIAUDAT

Composition du jury

Membres du jury avec voix délibérative

Benjamin Wandelt Professeur, Sorbonne Université & Institut d'Astrophysique de Paris	Président
Jason D. McEwen Professeur, University College London	Rapporteur & Examineur
Gabriel Peyré Directeur de recherche, CNRS, École Normale Supérieure	Rapporteur & Examineur
Laure Blanc-Féraud Directrice de recherche, CNRS, Université Côte d'Azur	Examinatrice
Julie Delon Professeure, Université Paris Cité	Examinatrice
Andrew N. Taylor Professeur, University of Edinburgh	Examineur

Titre : Modélisation guidée par les données des fonctions d'étalement du point des télescopes terrestres et spatiaux

Mots clés : Problèmes inverses, Fonction d'étalement du point, Factorisation de matrices, Différenciation automatique, Lentille gravitationnelle faible, Optique, Mission spatiale *Euclid*.

Résumé : L'effet de lentille gravitationnel est la distorsion des images de galaxies lointaines par des objets massifs et constitue une sonde puissante de la structure à grande échelle de notre Univers. Les cosmologistes utilisent la lentille (gravitationnelle) faible pour étudier la nature de la matière noire et sa distribution spatiale. Ces études nécessitent des mesures très précises des formes des galaxies, mais la réponse instrumentale du télescope, ou fonction d'étalement du point (PSF), déforme nos observations. Cette déformation peut être confondue avec des effets de lentille faible dans les images de galaxies, ce qui constitue l'une des principales sources d'erreur systématique. Par conséquent, l'estimation d'un modèle de PSF fiable et précis est cruciale pour le succès de toute mission de lentille faible. Le champ de PSF peut être interprété comme un noyau de convolution qui affecte chacune de nos observations d'intérêt et qui varie spatialement, spectralement et temporellement. Le modèle de PSF doit faire face à ces variations et est contraint par des étoiles spécifiques dans le champ observé. Ces étoiles, considérées comme des sources ponctuelles, nous fournissent des échantillons dégradés du champ de PSF. Les observations subissent différentes dégradations en fonction des propriétés du télescope, notamment un sous-échantillonnage, une intégration sur la bande passante de l'instrument et un bruit additif. Nous construisons finalement le modèle de PSF en utilisant ces observations dégradées, puis nous utilisons le modèle pour déduire les PSFs aux positions des galaxies. Cette procédure constitue le problème inverse mal posé de la modélisation de la PSF. Le cœur de cette thèse a été le développement de nouveaux modèles non-paramétriques pour estimer les PSFs à partir des étoiles observées dans les images acquises.

Nous avons développé un nouveau modèle de PSF pour les télescopes terrestres, appelé MCCD, qui peut modéliser simultanément l'ensemble du plan focal. Par conséquent, MCCD dispose de plus d'étoiles pour contraindre un modèle plus complexe. La méthode est basée sur un schéma de factorisation matricielle, les représentations parcimonieuses et une procédure d'optimisation alternée. Nous avons inclus le modèle de PSF dans un pipeline de mesure de forme à haute performance et l'avons utilisé pour traiter $\sim 3500 \text{ deg}^2$ d'observations en bande r provenant du Canada-France Imaging Survey. Un catalogue de formes a été produit et sera bientôt publié.

L'objectif principal de cette thèse a été de développer un modèle de PSF basé sur les données qui puisse répondre aux défis soulevés par l'une des missions les plus ambitieuses en matière de lentille faible, la mission spatiale *Euclid*. Les principales difficultés liées à la mission *Euclid* sont que les observations sont sous-échantillonnées et intégrées dans une large bande passante unique. Par conséquent, il est difficile de récupérer et de modéliser les variations chromatiques de la PSF à partir de ces observations. Notre principale contribution est un nouveau cadre pour la modélisation de la PSF basée sur un modèle optique différentiable permettant de construire un modèle de front d'onde basé sur les données. Le nouveau modèle, appelé WaveDiff, est basé sur un schéma de factorisation matricielle et des polynômes de Zernike. Le modèle s'appuie sur des méthodes modernes basées sur le gradient et la différenciation automatique pour l'optimisation, qui n'utilise que des observations dégradées et bruitées. Les résultats montrent que WaveDiff peut modéliser les variations chromatiques des PSF et gérer la super-résolution avec une grande précision.

Title : Data-driven modelling of ground-based and space-based telescope's point spread functions

Keywords : Inverse problems, Point spread function, Matrix factorisation, Automatic Differentiation, Weak gravitational lensing, Optics, *Euclid* space mission.

Abstract : Gravitational lensing is the distortion of the images of distant galaxies by intervening massive objects and constitutes a powerful probe of the Large Scale Structure of our Universe. Cosmologists use weak (gravitational) lensing to study the nature of dark matter and its spatial distribution. These studies require highly accurate measurements of galaxy shapes, but the telescope's instrumental response, or point spread function (PSF), deforms our observations. This deformation can be mistaken for weak lensing effects in the galaxy images, thus being one of the primary sources of systematic error when doing weak lensing science. Therefore, estimating a reliable and accurate PSF model is crucial for the success of any weak lensing mission. The PSF field can be interpreted as a convolutional kernel that affects each of our observations of interest that varies spatially, spectrally, and temporally. The PSF model needs to cope with these variations and is constrained by specific stars in the field of view. These stars, considered point sources, provide us with degraded samples of the PSF field. The observations go through different degradations depending on the properties of the telescope, including undersampling, an integration over the instrument's passband, and additive noise. We finally build the PSF model using these degraded observations and then use the model to infer the PSF at the position of galaxies. This procedure constitutes the ill-posed inverse problem of PSF modelling. The core of this thesis has been the development of new data-driven, also known as non-parametric, PSF models.

We have developed a new PSF model for ground-based telescopes, coined MCCD, which can simultaneously model the entire focal plane. Consequently, MCCD has more available stars to constrain a more complex model. The method is based on a matrix factorisation scheme, sparsity, and an alternating optimisation procedure. We have included the PSF model in a high-performance shape measurement pipeline and used it to process $\sim 3500 \text{ deg}^2$ of r-band observations from the Canada-France Imaging Survey. A shape catalogue has been produced and will be soon released.

The main goal of this thesis has been to develop a data-driven PSF model that can address the challenges raised by one of the most ambitious weak lensing missions so far, the *Euclid* space mission. The main difficulties related to the *Euclid* mission are that the observations are undersampled and integrated into a single wide passband. Therefore, it is hard to recover and model the PSF chromatic variations from such observations. Our main contribution has been a new framework for data-driven PSF modelling based on a differentiable optical forward model allowing us to build a data-driven model for the wavefront. The new model coined WaveDiff is based on a matrix factorisation scheme and Zernike polynomials. The model relies on modern gradient-based methods and automatic differentiation for optimisation, which only uses noisy broad-band in-focus observations. Results show that WaveDiff can model the PSFs' chromatic variations and handle super-resolution with high accuracy.

*Caminante, no hay camino,
se hace camino al andar.*
— Antonio Machado,
1875 - 1939

Acknowledgements

I would first like to thank Prof. Benjamin Wandelt, Prof. Jason D. McEwen, Prof. Gabriel Peyré, Prof. Laure Blanc-Féraud, Prof. Julie Delon and Prof. Andy N. Taylor for being part of the jury. I feel honoured of having a jury composed of researchers I admire, and am also grateful for the time they spent reading my thesis.

Then, I would like to deeply thank my supervisors, Jean-Luc Starck and Martin Kilbinger, which have been supportive from the first day of my PhD. You have encouraged and motivated me to throughout the whole PhD with a global pandemic in between. We managed to build a relationship with a lot of trust, which, I believe, proved to be fruitful. I have learned many things from them throughout these years and I hope I will be able to continue learning from them.

My PhD was carried out on the CosmoStat laboratory, from which its supportive atmosphere and the care taken to newcomers should be highlighted. Even though one might argue that the whole is more than the sum of its parts, the pillars of this lab are its permanent researchers. My sincere thanks to Sam Farrens, François Lanusse, Valeria Pettorino and Joana Frontera Pons, without mentioning Jean-Luc and Martin. I really appreciated the efforts taken to counteract the effects of the pandemic, like the *Learning slot*. Thanks for the advice in terms of software, research, and on how to candidate to a position. I have learned a lot from you.

I had the pleasure of sharing part of my PhD experience with several interns, PhD students and postdocs, many thanks to all of them. In particular to Virgi and Vanshi with whom we shared the 272 office. Then, I also thank the old and new journal club organisers. I would like to thank the postdocs and engineers, Santi Casas, Isabella Carucci, André Zamorano Vitorelli, Lucie Baumont and Sergio La Hera, and also to the previously-PhD-students-now-doctors colleagues Morgan Schmitz, Christophe Kervazo, Virginia Ajani, Imane El hamzaoui, Fadi Nammour, Rémi Carloni Gertosio, and Benjamin Chiche, between others. I appreciated the time spent together with the current PhD students, Benjamin Remy, Denise Lanzieri and Lisa Goh, and wish them good luck on the last parts of their PhD. Also, thanks to the DAp people for nice Friday beers, Carlos, Aoife, Boris, Rose, etc... A big thanks for Nacho, Sunayana and Koryo for the amazing post-defence concert with a fine choice of songs and a great performance. Good luck to Vilasini and Fabian on their PhDs at CosmoStat. The best of lucks to Jennifer with *Euclid*'s PSFs.

Thanks to all the CEA football people that made Thursday's evenings great, in special Tim and Fernando.

I had the beautiful opportunity to supervise three *excellent* interns throughout my thesis, Jérôme Bonnin, Aziz Ayed and Ezequiel Centofanti. This was an enriching experience for me, and I hope it was for them as well.

I want to thank the CEA for funding my PhD and for the opportunity to work in such a stimulating environment. Although I must admit I am a bit sad that you are installing ping-pong tables and table footballs¹ just after I left. I am grateful for time Pierre-Antoine Frugier spent with me explaining optic concepts always being very pedagogical. I extend my

¹*babyfoot* in French and *metegol* in Spanish.

thanks to Emeric Le Floc'h and Eric Pantin for being part of my *comité de suivi* and for following my work throughout these three years.

No podría terminar estas líneas sin agradecer a mis amigos, a ambos lados del océano, y a mi familia que siempre me apoyaron. Un gran gracias para mi madre que pudo estar presente en mi defensa de tesis.

Acknowledgements for computing resources and data

This work was granted access to the HPC resources of IDRIS under the allocations 2021-AD011011554 and 2021-AD011012983 made by GENCI. This work has made use of the CANDIDE Cluster at the Institut d'Astrophysique de Paris and made possible by grants from the PNCG and the DIM-ACAV.

This work is based on data obtained as part of the Canada-France Imaging Survey, a CFHT large program of the National Research Council of Canada and the French Centre National de la Recherche Scientifique. Based on observations obtained with MegaPrime/MegaCam, a joint project of CFHT and CEA Saclay, at the Canada-France-Hawaii Telescope (CFHT) which is operated by the National Research Council (NRC) of Canada, the Institut National des Science de l'Univers (INSU) of the Centre National de la Recherche Scientifique (CNRS) of France, and the University of Hawaii. This research used the facilities of the Canadian Astronomy Data Centre operated by the National Research Council of Canada with the support of the Canadian Space Agency.

Contents

Acknowledgements	v
Contents	vii
Acronyms	xi
Extended abstract	1
I Context	7
1 Background	9
1.1 Introduction to cosmology	10
1.1.1 General relativity and an isotropic, homogeneous Universe	10
1.1.2 The Lambda-CDM cosmological model	11
1.2 Weak gravitational lensing	13
1.2.1 Light propagation in the Universe	13
1.2.2 Shear and convergence	15
1.2.3 Estimating the shear from observations	17
1.2.4 Relations between the convergence and the shear	19
1.2.5 Relation to cosmology	21
1.3 Main ingredients of a shape measurement pipeline	23
1.3.1 Image preprocessing	24
1.3.2 Masking	26
1.3.3 Object detection and deblending	26
1.3.4 Star-galaxy separation	27
1.3.5 PSF modelling	28
1.3.6 Galaxy shape measurement	28
1.3.7 Calibration	29
1.4 The <i>Euclid</i> space mission	30
2 Introduction to the PSF	35
2.1 Introduction to optics	36
2.1.1 Scalar diffraction theory	36
2.1.2 The Fresnel approximation	38
2.1.3 The Fraunhofer approximation	39
2.1.4 Analysis of optical imaging systems	40

2.2	General observational forward model	45
2.3	PSF field contributors and related degradations	48
2.3.1	Optic-level contributors	49
2.3.2	Detector-level degradations	50
2.3.3	The atmosphere	54
2.4	Current PSF models	57
2.4.1	Parametric PSF models	57
2.4.2	Non-parametric (or data-driven) PSF models	58
2.5	General comments on PSF modelling	65
2.5.1	Desirable properties of PSF models	65
2.5.2	On the structure of the PSF field spatial variations	67
2.6	Validation of PSF models	69
2.6.1	Pixel-based metrics	70
2.6.2	Moment-based metrics	71
2.6.3	PSF error propagation and PSF requirements	73
II PSF models for ground-based telescopes		77
3	Modelling the PSF in the full focal plane at once	79
3.1	Introduction	80
3.2	The observation model	81
3.3	Resolved component analysis (RCA)	82
3.4	A new family of MCCD methods	83
3.4.1	The MCCD data model	83
3.4.2	Inverse problem and regularisation	84
3.4.3	Global model	85
3.4.4	Local model	86
3.4.5	Optimisation problem	87
3.4.6	Algorithm	88
3.4.7	PSF recovery	89
3.5	Numerical experiments with simulated data	92
3.5.1	Data	92
3.5.2	Training set	92
3.5.3	Testing data set	93
3.5.4	Quality criteria	93
3.5.5	Model parameters	95
3.5.6	Results	95
3.5.7	Comparison of computing resources	98
3.6	Numerical experiments with UNIONS/CFIS data	98
3.6.1	Dataset	98
3.6.2	Model parameters	99
3.6.3	A new metric for real data: the Q criteria	101
3.6.4	Results	103
3.7	Conclusion	104
4	Modelling the PSF for UNIONS/CFIS	107
4.1	Introduction	107
4.2	Incorporating MCCD to ShapePipe	109
4.3	Shape catalogue and PSF diagnostics	111
4.3.1	Shape measurement diagnostics	112
4.3.2	PSF diagnostics	112

4.4	Conclusion	116
III PSF models for space-based telescopes		119
5	Rethinking data-driven PSF modelling with a differentiable optical model	121
5.1	Introduction	122
5.1.1	Challenges in PSF modelling for space missions	122
5.1.2	Current approaches for PSF modelling	123
5.1.3	Our contribution	124
5.1.4	Notation	124
5.2	The PSF field and the inverse problem	124
5.2.1	The observation model	125
5.2.2	Inverse problem and regularisations	126
5.3	Related works	127
5.4	WaveDiff, a data-driven wavefront-based PSF model	128
5.4.1	Physical motivation and the wavefront error	128
5.4.2	Wavefront error PSF model	129
5.4.3	Practical optical forward model	134
5.4.4	Training	136
5.4.5	PSF recovery	137
5.5	Numerical experiments	138
5.5.1	Objectives	138
5.5.2	Setup	140
5.5.3	Results	144
5.6	Discussion and conclusion	150
6	Applying WaveDiff to <i>Euclid</i>	155
6.1	Introduction	156
6.2	Introducing prior information to the WaveDiff model	156
6.2.1	Motivation	156
6.2.2	Model description	157
6.2.3	Simulations and optical prior	158
6.2.4	Numerical experiment	161
6.2.5	Conclusion	165
6.3	Wavefront error resolution study	165
6.4	Using more exposures for the PSF model	167
6.5	Impact of SED errors in the PSF generation	170
6.5.1	Types of errors	171
6.5.2	Numerical experiments	172
6.6	Studying SED interpolation in the PSF generation	175
6.6.1	SED interpolation	175
6.6.2	Numerical experiments	176
6.7	SED errors in PSF modelling	178
6.7.1	Numerical experiment	178
6.7.2	Conclusion	179
6.8	Improving the WaveDiff optimisation procedure	180
6.8.1	Motivation	180
6.8.2	Projecting data-driven features	181
6.8.3	Introducing the WFE projections to the optimisation procedure	187
6.8.4	Numerical experiments	189
6.8.5	Results	189

6.8.6 Conclusions	193
6.9 Conclusion	193
General Conclusions and Perspectives	195
Contributions	195
Perspectives	196
Appendices	201
A MCCD optimization	203
B Publications and academic activities	207
C Résumé étendu en français (Extended abstract in French)	209
List of Figures	215
List of Tables	223
Bibliography	225

Acronyms

- 1D** 1-dimensional 81
- 2D** 2-dimensional 19, 21–23, 53, 54, 59, 62, 64, 81, 161
- 2PCF** Two-Point Correlation Function 19–21, 23, 73
- 3D** 3-dimensional 21, 22, 31, 198
- ACS** Advanced Camera for Surveys 58
- AOCS** Attitude and Orbit Control System 52
- BAO** Baryonic Acoustic Oscillations 12, 223
- BFD** Bayesian Fourier Domain 29, 64
- BFE** Brighter-fatter Effect 24, 53
- CANFAR** Canadian Advanced Network for Astronomical Research 111
- CCD** Charged Coupled Device 24–26, 31, 48, 49, 52–54, 60–63, 66, 80–87, 89, 91–93, 95, 96, 98–100, 104, 105, 110–112, 115, 122, 127, 140, 142–144, 152, 195, 198, 205, 216, 217
- Λ CDM** Lambda Cold Dark Matter 11, 12, 223
- CFHT** Canada-France-Hawaii Telescope 47, 48, 52, 57, 61, 92, 105, 107, 108, 171, 216
- CFHTLenS** Canada-France-Hawaii Telescope Lensing Survey 29, 61, 117
- CFHTLS** Canada-France-Hawaii Telescope Legacy Survey 59
- CFIS** Canada-France Imaging Survey 24–27, 29, 48, 54, 56, 60, 80–82, 92, 93, 98, 102–105, 107–110, 116, 117, 156, 171, 195–197, 208, 216, 218, 223
- CMB** Cosmological Microwave Background 10, 12, 223
- CNN** Convolutional Neural Network 65, 198
- COSMOS** Cosmological Evolution Survey 32, 58, 60, 197

- CPU** Central Processing Unit 166
- CTI** Charge Transfer Inefficiency 24, 53, 66
- DAE** Denoising Auto-Encoder 65
- DD** Data-Driven 130, 132, 133, 136, 138, 180–194, 221–223
- DE** Dark Energy 11
- DEC** Declination 62
- DECam** Dark Energy Camera 54, 62, 64
- DES** Dark Energy Survey 20, 24, 29, 30, 54, 56, 60, 62, 63, 65, 70, 80, 117, 215
- DR** Data Release 109, 116, 218
- EC** *Euclid* Consortium 50, 68–70, 171, 198, 217
- EM** Electromagnetic 11, 36, 38, 39, 50, 52
- ESA** European Space Agency 12, 24, 31, 49, 215
- Euclid*** *Euclid* space mission 62
- FFT** Fast Fourier Transform 134, 135, 166
- FLRW** Friedmann-Lemaître-Robertson-Walker 10, 11
- FOV** Field-of-View 31, 32, 40, 44–46, 49, 57, 58, 60, 61, 65, 67–70, 122–127, 129–133, 137–140, 142, 143, 149, 157–159, 165, 172, 182, 185–187, 191, 193, 198, 217, 219–222
- FT** Fourier Transform 19, 23, 38, 39, 49, 65, 129
- FWHM** Full Width Half Maximum 27, 32, 49, 54, 64, 98, 216
- GC** Galaxy Clustering 31
- GPU** Graphical Processing Unit 144, 166–168
- GR** General Relativity 10, 11, 13
- GT** Ground Truth 46, 48, 151, 156, 166, 171, 172, 174, 176, 178, 180, 181, 186–189, 191–194, 220–222
- HME** Higher Moments Errors 74, 75
- HPC** High-Performance Computing 115
- HSC** Hyper Suprime-Cam survey 24, 30, 46, 60, 65, 117
- HSM** Hirata-Seljak-Mandelbaum 71, 94, 112
- HST** Hubble Space Telescope 31, 32, 49–53, 58, 60, 123, 157, 197, 216
- JWST** James Webb Space Telescope 12, 31, 197, 215

- KiDS-1000** Kilo-Degree Survey 29, 30, 61, 117
- KS** Kaiser-Squires 20, 215
- KSB** Kaiser-Squires-Broadhurst 28, 59
- LSS** Large Scale Structure 10
- LSST** Large Synoptic Survey Telescope 18, 27, 55–57, 60, 62, 63, 67, 74, 75, 217
- MAD** Median Absolute Deviation 110
- MAMSL** Metres Above Mean Sea Level 108
- MCCD** Multi-CCD 80, 81, 83, 84, 86, 91–93, 95, 98, 104, 105, 107, 109–112, 114–117, 128, 142–144, 148, 153, 156, 195–197, 208, 223
- MDB** *Euclid*'s Mission Database 54, 217
- NASA** National Aeronautics and Space Administration 12, 31, 215, 216
- NISP** *Euclid*'s Near Infrared Spectrometer and Photometer instrument 31, 50
- OPD** Optical Path Differences 44, 67
- OT** Optimal Transport 64
- Pan-STARRS** Panoramic Survey Telescope and Rapid Response System 98
- PCA** Principal Component Analysis 58–60, 63, 64
- PDF** Probability Density Function 22
- PIFF** PSFs In the Full Field-of-View 62, 63, 65, 75, 80, 110
- PRNU** Photo Response Non-Uniformity 54
- PSF** Point Spread Function 9, 23–25, 27–30, 32, 33, 35, 36, 43, 45–52, 54, 56–67, 69–75, 80–86, 89, 91–93, 95, 96, 98–101, 103–117, 122–130, 132–153, 155–157, 161, 162, 166–168, 171–181, 186–189, 193–198, 208, 216–223
- PSFEx** PSF Extractor 57, 60–63, 65, 80, 82, 83, 86, 93, 95–99, 103–106, 109, 117, 128, 142–144, 148, 153, 217, 218, 223
- QE** Quantum Efficiency 52, 54, 217
- RA** Right Ascension 62
- RBF** Radial Basis Function 64, 81, 83, 89, 91, 92, 138
- RCA** Resolved Component Analysis 59, 61–64, 80–84, 86, 87, 95, 96, 98, 104, 105, 128, 142–144, 148–150, 153, 217, 219
- RMS** Root Mean Square 110, 139, 140, 147, 158–162, 219, 220
- RMSE** Root Mean Squared Error 70, 71, 93, 94, 96, 110, 139, 162, 164, 167–170, 173–175, 178, 179, 190, 191, 217, 220, 221

- Roman** Nancy Grace Roman Space Telescope 50
- SDSS** Sloan Digital Sky Survey 24, 59
- SED** Spectral Energy Distribution 47, 48, 125, 135, 140, 141, 156, 166, 171–180, 194, 196, 198, 219–221
- SExtractor** Source Extractor 60, 93
- SFE** Surface Errors 49, 158, 159, 220
- SNR** Signal-to-Noise Ratio 18, 25, 28, 62, 66, 72, 89, 93, 95–100, 103, 105, 108, 110, 140, 142, 143, 147, 217, 218
- SR** Super Resolution 123, 166, 190, 191
- SSS** SuperCOSMOS Sky Survey 49
- STOP** Structural-Thermal-Optical Performance 50
- SVD** Singular Value Decomposition 59
- SVO** Spanish Virtual Observatory 47
- UNIONS** Ultra-violet Near-Infrared Optical Northern Sky 98, 107–109, 171, 218
- Vera C. Rubin Observatory** Vera C. Rubin Observatory 52
- VIS** *Euclid*'s VIS Instrument 31–33, 45, 50, 52–54, 58, 82, 171, 196, 217
- WCS** World Coordinate System 24, 62
- WFE** Wavefront Errors 43, 44, 49, 55, 64, 67, 69, 128–133, 136–141, 144–146, 148, 150–152, 156–159, 161, 162, 164–168, 170, 172, 174, 178–183, 185–194, 196–198, 216, 219–223
- WL** Weak Gravitational Lensing 9, 14, 18, 21–23, 26–33, 52, 53, 58–61, 65–67, 69–71, 74, 75, 80, 105, 107, 108, 110, 116, 117, 157, 195

Extended abstract

Context

Gravitational lensing is the distortion of the images of distant galaxies by intervening massive objects and constitutes a powerful probe of the Large Scale Structure of our Universe. Cosmologists use weak (gravitational) lensing to study the nature of dark matter and its spatial distribution. These studies require highly accurate measurements of galaxy shapes, but the telescope's instrumental response, or point spread function (PSF), deforms our observations. This deformation can be mistaken for weak lensing effects in the galaxy images, thus being one of the primary sources of systematic error when doing weak lensing science. Therefore, estimating a reliable and accurate PSF model is crucial for the success of any weak lensing mission. The PSF field can be interpreted as a convolutional kernel that affects each of our observations of interest that varies spatially, spectrally, and temporally. The PSF model needs to cope with these variations and is constrained by specific stars in the field of view. These stars, considered point sources, provide us with degraded samples of the PSF field. The observations go through different degradations depending on the properties of the telescope, including undersampling, an integration over the instrument's passband, and additive noise. We finally build the PSF model using these degraded observations and then use the model to infer the PSF at the position of galaxies. This procedure constitutes the ill-posed inverse problem of PSF modelling. The core of this thesis has been the development of new data-driven PSF models.

The next generation of wide-field cosmological surveys, such as *Euclid*, will observe the Universe on an unprecedented scale and generate a vast amount of valuable information. The ever-increasing quality and amount of data pose severe challenges for the PSF model as well as very stringent requirements on its performance.

The PSF modelling problem for weak lensing studies resides in the intersection of three disciplines: optics, cosmology, and inverse problems in imaging. Therefore, the first part of this thesis introduces all the necessary concepts required to develop a comprehensive understanding of this problem. We started by describing the basic concepts of cosmology and weak lensing required to understand the final goal of a PSF model in the weak lensing analysis. We continued with an overview of the main building blocks of a shape measurement pipeline to depict the context in which the PSF model will work. The introduction continues with essential notions in optics and a review of the contributors to the PSF field. Finally, we present an overview of state-of-the-art PSF models and conclude with validation methods for PSF models.

PSF models for ground-based telescopes

Modelling the PSF in the full focal plane at once

There are several sources of spatial variations of the PSF. Some cover the entire focal plane, like the optical aberrations of the telescope and the atmosphere, while others are specific for each Charge-Coupled Device (CCD) chip. The focal plane of current wide-field imaging cameras comprises an extensive array of CCDs, constituting a CCD mosaic. Most of the current PSF models are built independently on each CCD which is a simple solution to account for both variations. However, this choice has some drawbacks. First, it limits the number of available stars to constrain the model, thus favouring simpler models. Second, spatial variations covering the entire focal plane cannot be well modelled when approximated by simple piecewise models. Consequently, PSF modelling errors arise. In order to tackle both problems simultaneously, we proposed a new PSF model, coined MCCD, that can model the full focal plane and handles both types of variations. MCCD can successfully include the camera geometry into the PSF model. The proposed PSF model is based on a matrix factorization scheme that uses different mathematical tools, such as sparse regularisations for PSF denoising and graph theory to handle localized spatial variations, among others. We developed the training algorithm by combining these concepts with block coordinate descent, efficient convex optimization methods, and proximal algorithms. We validated the model with simulations and real observations. Finally, our results show that the proposed PSF model can build more complex models, improve performance with respect to state-of-the-art methods, and be robust to handle real data.

Modelling the PSF for UNIONS/CFIS

The Canada-France Imaging Survey (CFIS) is an ongoing legacy survey at the u and r bands at the Canada-France-Hawaii Telescope (CFHT). It will cover 4800 deg^2 , and its main objective is to contribute to deep surveys like *Euclid* with its photometric redshift estimation. The excellent r -band image quality allows conducting weak lensing science. I have contributed to the high-performance shape measurement pipeline, **ShapePipe**, in several aspects of the PSF modelling and the validation tests. We produced a shape catalogue using 1700 deg^2 with the widely-used PSF model **PSFEx**.

Later, we included the MCCD PSF model into **ShapePipe** and validated its performance with real data. We have run the pipeline in a high-performance computing environment with our PSF model to produce a new shape catalogue spanning $\sim 3500 \text{ deg}^2$, one of the largest areas for weak-lensing studies at the time of writing. We then present the results of several PSF validation methods using the new catalogue. The results show that the MCCD model improves the PSF model performance with respect to the previous model and confirms the robustness and maturity of the method for handling real data.

PSF models for space-based telescopes

Rethinking data-driven PSF modelling with a differentiable optical model

The *Euclid* survey represents one of the most arduous challenges for PSF modelling. Because of the very broad passband of *Euclid*'s visible imager (VIS) ranging from 550nm to 900nm, PSF models need to capture not only the PSF field spatial variations but also its chromatic variations. Each star observation is integrated over the entire VIS passband with the object's spectral energy distribution (SED). As the observations are




Figure 01 – Different representations of a space-like PSF at a single position in the field of view. (a) Noisy polychromatic PSF at observation resolution, (b) Noiseless polychromatic observation at observation resolution, (c) High-resolution noiseless polychromatic observation, (d) Wavefront error map representing the aberrations in the optical system. Units are in μm , (e) High-resolution chromatic variations of the PSF at equally spaced wavelengths in the passband [550, 900]nm.

undersampled, a super-resolution step is also required. [Figure 01](#) present an example of an *Euclid*-like PSF, its under-sampling and its chromatic variations. Current data-driven PSF models cannot model chromatic variations. The state-of-the-art model designed for *Euclid*, the Resolved Component Analysis (RCA) model, is ~ 200 and $\sim 10^4$ times over shape and size error requirements, respectively. These models are all built in the pixel space with some dimensionality reduction method using different constraints. We propose a paradigm shift in the way data-driven PSF models are built.

The addition of a differentiable optical forward model into the proposed modelling framework changes the data-driven modelling space from the pixels to the wavefront. The optical forward model is based on Fourier optics and fundamental optic principles. Our framework, presented in [Figure 02](#), allows building powerful physically motivated interpretable models that do not require special calibration data. The new model, coined WaveDiff, relies on a matrix factorisation scheme and Zernike polynomials. The model relies on modern gradient-based methods and automatic differentiation for optimisation, which only uses noisy broad-band in-focus observations. WaveDiff is the first data-driven model that can successfully model chromatic variations on top of handling spatial variations and super-resolution (x3). It also represents a breakthrough in performance for data-driven models by decreasing over one and two orders of magnitude *Euclid*'s error requirements of size and shape, respectively. The proposed model is modular, differentiable, built in the TensorFlow framework and entirely runs on GPUs.

Applying WaveDiff to *Euclid*

The WaveDiff PSF model constitutes a promising approach for the *Euclid* mission. Nevertheless, transitioning from good performance with simulations to good performance with real data is far from trivial. This chapter addresses several emerging issues when we will deal with real *Euclid* observations. We intend that the studies and

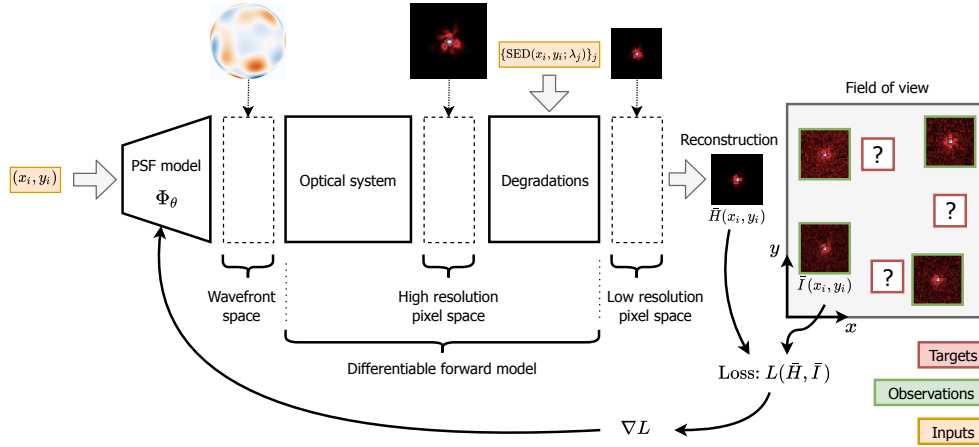


Figure 02 – A schematic of the proposed framework for data-driven wavefront PSF modelling for the *Euclid* space mission.

extensions presented in this chapter will smooth the transition from simulated to real data for the WaveDiff model. We proceed with a list of the issues addressed:

- Prior optical information about the telescope will be available. This information can come from high-fidelity simulations or complementary observations in space. However, the optical prior is expected to have some degree of error. We modify the WaveDiff model to include an optical prior and to correct it in a non-parametric way. The results show that the model can exploit the prior information and considerably improve the model performance.
- In previous studies, we assumed the same wavefront dimension to generate the observations, i.e. ground truth model, and for WaveDiff. We study how the wavefront dimension affects the PSF's modelling when the observations are generated more realistically.
- Up to this point, we have not used more than the number of stars in one exposure to constrain the WaveDiff model. The stability of *Euclid* might allow us to use subsequent exposures and therefore increase the number of stars to constrain the PSF model. Consequently, we study how the WaveDiff model exploits information from a denser stellar field.
- We have considered, in the previous studies, that the input SED information was flawless and had used 20 bins for the observations and the WaveDiff model. In a real scenario, the observations are generated without any spectral discretisation. In addition, real SED information comes in a reduced number of bins containing errors. We study how spectral discretisation affects the modelling of the PSF and ways to improve the results.
- Previous results showcased the ability of the WaveDiff model to estimate a *useful* wavefront representation that allows us to obtain a low pixel error. However, suppose we are in a new scenario where the parametric part of the WaveDiff model can reproduce the ground truth wavefront field. In that case, the current WaveDiff model cannot recover this wavefront field. It is indeed a phase retrieval problem. Nonetheless, in this new scenario, the solution to the phase retrieval

problem is also the global minimum of our PSF modelling problem. We propose a new optimisation procedure for addressing this phase retrieval problem. The procedure is based on a proposed projection from the non-parametric to the parametric part of the model. We show with numerical experiments that the WaveDiff model can estimate the ground truth wavefront field with low errors only using degraded in-focus observations.

Part I
Context

Background

Chapter Outline

1.1	Introduction to cosmology	10
1.1.1	General relativity and an isotropic, homogeneous Universe	10
1.1.2	The Lambda-CDM cosmological model	11
1.2	Weak gravitational lensing	13
1.2.1	Light propagation in the Universe	13
1.2.2	Shear and convergence	15
1.2.3	Estimating the shear from observations	17
1.2.4	Relations between the convergence and the shear	19
1.2.5	Relation to cosmology	21
1.3	Main ingredients of a shape measurement pipeline	23
1.3.1	Image preprocessing	24
1.3.2	Masking	26
1.3.3	Object detection and deblending	26
1.3.4	Star-galaxy separation	27
1.3.5	PSF modelling	28
1.3.6	Galaxy shape measurement	28
1.3.7	Calibration	29
1.4	The <i>Euclid</i> space mission	30

THIS chapter introduces the background and context of the [Point Spread Function \(PSF\)](#) modelling methods that will be later presented in this thesis. We start with a short presentation of general cosmological concepts and some current challenges in cosmology. We then continue to describe the basic theory of weak gravitational lensing, the main cosmological probe we will use to tackle such challenges. The description is narrowed down to the cornerstone of a [Weak Gravitational Lensing \(WL\)](#) analysis, the shape measurement pipeline. We will see the role of the [PSF](#) model in the pipeline and its crucial task in controlling severe systematic errors in a [WL](#) analysis. We end the chapter with a brief overview of the *Euclid* space mission and introduce the new challenges it poses for the [PSF](#) model.

1.1 Introduction to cosmology

1.1.1 General relativity and an isotropic, homogeneous Universe

Einstein's theory of **General Relativity (GR)** is at the core of our current understanding of the Universe. **GR** describes the gravitational force as the curvature of a 4-dimensional spacetime allowing us to specify how matter and radiation influence the geometry of space and time. The mathematical framework allowing this geometrical description of gravity is the Riemannian geometry, a branch of differential geometry. When applied to cosmological scales for a Universe governed by gravity, **GR** allows linking the matter and energy content of the Universe with its geometry and their dynamical temporal evolution. The field equations of **GR** allow us to study the spacetime geometry and its matter content, which we express as follows

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}\mathcal{R} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu} , \quad (1.1)$$

where $g_{\mu\nu}$ is the *metric tensor* and describes the local geometry of space-time. The *Ricci curvature tensor* is $R_{\mu\nu}$ and the *Ricci scalar* is \mathcal{R} , and can be expressed as a function of the metric tensor and its derivatives. Then, Λ is the *cosmological constant* associated with dark energy in cosmology, c is the speed of light in vacuum, G is the Newtonian constant of gravitation, and $T_{\mu\nu}$ is the energy-momentum tensor that is related to the source of the spacetime curvature.

Within this formulation and with the help of the metric tensor, we can define the line element allowing us to compute distances in the **GR** description of the Universe as follows

$$ds^2 = \sum_{\mu,\nu=0}^3 g_{\mu\nu}dx^\mu dx^\nu , \quad (1.2)$$

where x^μ is a set of coordinates with the index 0 being the time and 1-3 the indices related to space by convention.

To study the motion of particles in this theory, one has to solve the geodesic equation coming from [Equation 1.1](#) which is difficult in a general case. Solving analytically **GR**'s field equations without any simplification is in general not possible. The study of the Universe through **GR** requires further simplifying assumptions. The simplest, yet powerful, assumption we can make is that the Universe is *isotropic* and *homogeneous* at *large scales* which translates to a rotational and translational invariance of the Universe. These assumptions give rise to one of the cosmological pillars, the *Cosmological Principle*. This principle has been verified by observations of the [Cosmological Microwave Background \(CMB\)](#) and the [Large Scale Structure \(LSS\)](#). The symmetries arising from the Cosmological principle allow a great deal of simplification to be made to Einstein's field equations. The metric describing such a Universe is unique and is known as [Friedmann-Lemaître-Robertson-Walker \(FLRW\)](#), and simplifies [Equation 1.2](#) into

$$ds^2 = dt^2 - a^2(t) (d\chi^2 + f_K(\chi) (d\theta^2 + \sin^2 \theta d\phi^2)) , \quad (1.3)$$

where we expressed the line element in spherical coordinates (r, θ, ϕ) , in a comoving reference frame. Then, χ is the time-independent *comoving radial distance*, and $a(t)$ is the *scale factor* that is free to vary with time and describes how distances between points change as the Universe contracts or expands. In an expanding Universe, the distance between two comoving points, i.e. observers moving along with the expansion, is always changing (increasing). The comoving formalism allows us to define a distance

unaffected by the Universe's expansion which is described by the scale factor a . We define the comoving transverse angular distance as $f_K(\chi)$, which is the ratio of the comoving separation between two points at χ to their separation angle. The distance depends on the curvature K of the Universe. It has three expressions depending on the value of the curvature, if it is a spherical ($K > 0$), flat ($K = 0$), or hyperbolic ($K < 0$) Universe. There has been much evidence, e.g. Planck Collaboration et al. [Pla+20], that we live in a flat Universe which gives $f_K(\chi) = \chi$. See Kilbinger [Kil15, §2] for a more information on the comoving angular distance f_K .

A quantity that can be interpreted as a proxy for distance and that we will use in the next section is the *cosmological redshift* z . The expansion of space-time causes **Electromagnetic (EM)** wavelengths to increase resulting in a shift to red. The change of frequency can be expressed from the following relation

$$\frac{\lambda_{\text{ob}}}{\lambda_{\text{em}}} = \frac{a_{\text{ob}}}{a_{\text{em}}}, \quad (1.4)$$

where on the left-hand side, we have the ratio of the observed and emitted wavelengths, and on the right-hand side, the ratio of the scale factors. In an expanding Universe, we have that $a_{\text{ob}} > a_{\text{em}}$ and therefore the observed wavelength is shifted towards the red, which gives origin to the name *redshift*. Finally, the redshift can be defined as

$$z := \frac{\lambda_{\text{ob}} - \lambda_{\text{em}}}{\lambda_{\text{ob}}}. \quad (1.5)$$

1.1.2 The Lambda-CDM cosmological model

There exist several cosmological models to describe the Universe and its evolution. Nevertheless, compared to the observations, the most successful model is the **Lambda Cold Dark Matter (ΛCDM)** model. Λ is the cosmological parameter from GR field equations in Equation 1.1 and is related to a fluid known as **Dark Energy (DE)**. The model parametrises the cosmology described using GR and the FLRW metric we have seen in the previous section. ΛCDM has been very successful in describing the evolution of the Universe and predicting observations using a small set of parameters presented in Table 11.

Dark energy has been included in the cosmological model as an explanation for the accelerated expansion of the Universe [Per+99]. In turn, dark matter has been included in the model to explain several observations, for example, the mass deficit of visible matter from observations of the rotational speed of peripheral galaxies in a cluster [Zwi37]. *Dark matter* is a non-baryonic matter that does not have any electromagnetic interaction and is thus invisible.

The density parameters (Ω) seen in Table 11 shows that dark energy is the most significant contributor to the energy-matter density of the Universe with roughly the $\sim 69\%$ (see Ω_Λ). The remaining density that is encoded in the matter density Ω_m can be decomposed between a 26.2% contribution of cold dark matter (Ω_c) and a 4.8% contribution of baryonic (or ordinary) matter (Ω_b).

The nature and physics of dark energy and dark matter remain largely unknown even though they are the main contributors to the energy-matter content of our Universe. Understanding the nature of dark energy and dark matter is one of the main challenges of modern cosmology. Future surveys like *Euclid*, described in section 1.4, have been designed to enlighten our understanding of dark energy and dark matter with the help of techniques like *weak gravitational lensing*, described in section 1.2. Making sure that the mission accomplishes its goals is the main motivation of this

Parameter	Symbol	Value
Hubble parameter	H_0	$(67.66 \pm 0.42) \text{ km s}^{-1} \text{ Mpc}^{-1}$
Total matter density	Ω_m	0.3111 ± 0.0056
Dark matter density	$\Omega_c h^2$	0.11933 ± 0.00091
Baryon density	$\Omega_b h^2$	0.02242 ± 0.00014
Dark energy density	Ω_Λ	0.6889 ± 0.0056
Power spectrum normalisation	σ_8	0.8102 ± 0.0060
Spectral index	n_s	0.9665 ± 0.0038
Reionisation optical depth	τ	0.0561 ± 0.0071

Table 11 – Main cosmological parameters for Λ CDM with a 68% confidence limit from Planck Collaboration et al. [Pla+20, Tab. 2] using a combination of CMB TT, TE, EE + lowE + CMB lensing + Baryonic Acoustic Oscillations (BAO).



Figure 1.1-1 – Deep image from JWST using the Near-Infrared Camera, which shows the gravitational lensing effect of the galaxy cluster SMACS 0723. Credit: NASA and ESA.

thesis. [Figure 1.1-1](#) shows the gravitational lensing effect of foreground objects due to a massive galaxy cluster.

1.2 Weak gravitational lensing

This section summarises the most relevant aspects of weak gravitational lensing for this thesis and is based on Schneider et al. [[SEF92](#)], Seitz et al. [[SSE94](#)], Bartelmann et al. [[BS01](#)], Kilbinger [[Kil15](#)] and Mandelbaum [[Man18](#)].

1.2.1 Light propagation in the Universe

One way to describe the deflection of light in the presence of massive bodies is to use Fermat's principle of minimal light travel time. The photons travel through null geodesics given by the line element ds coming from the field equations of GR. The following equation gives the time taken by the ray to travel through the Universe

$$t = \frac{1}{c} \int \left(1 - 2\frac{\Phi}{c^2} \right) dr , \quad (1.6)$$

where Φ is the gravitational potential, dr is the light path, and we are assuming the weak-field approximation $\Phi \ll c^2$. Following an optic analogy, the gravitational potential acts as a medium with variable refractive index $n = 1 - 2\Phi/c^2$. We proceed by computing the *deflection angle* of the light ray due to the presence of massive bodies whose effect is encoded in the gravitational potential. The deflection angle is the difference in directions between the emitted and received light rays. We need to apply Fermat's principle to obtain this angle to get the Euler-Lagrange equations for the refractive index n . Then, if we integrate along the light path we obtain the deflection angle that writes

$$\hat{\alpha} = -\frac{2}{c^2} \int \nabla_{\perp} \Phi dr , \quad (1.7)$$

where $\nabla_{\perp} \Phi := \nabla \Phi - \mathbf{e}(\mathbf{e} \cdot \nabla \Phi)$ denotes the projection of $\nabla \Phi$ onto the plane orthogonal to the direction of the light ray where \mathbf{e} is the unit tangent vector of the ray.

Let us now consider a fiducial light ray presented as the bottom red light ray in [Figure 1.2-2](#). As we continue, we will consider distances in the source plane with respect to the intersection of the fiducial ray with the source plane. In addition, the angles we measure from the observer are with respect to the incoming fiducial light ray. The transverse comoving separation, x_0 , which lies in the source plane seen in [Figure 1.2-2](#), between two light rays as a function of comoving distance can be written as

$$\mathbf{x}_0(\chi) = f_K(\chi) \boldsymbol{\theta} , \quad (1.8)$$

where we consider the small angle (or 0th-order) approximation for $\boldsymbol{\theta}$ [[SEF92](#)], and it corresponds to the absence of lensing. [Figure 1.2-2](#) illustrate the angles and distances used. Following the formulation and assumptions of [Equation 1.8](#) we can also express the other transverse comoving separations as a function of angles as follows

$$\mathbf{x}(\chi) = f_K(\chi) \boldsymbol{\beta} , \quad (1.9)$$

$$d\mathbf{x}(\chi) = f_K(\chi - \chi') d\hat{\alpha} , \quad (1.10)$$

where the first equation shows the angle $\boldsymbol{\beta}$ that corresponds to the observation angle if there were no gravitational potential perturbations. The second equation

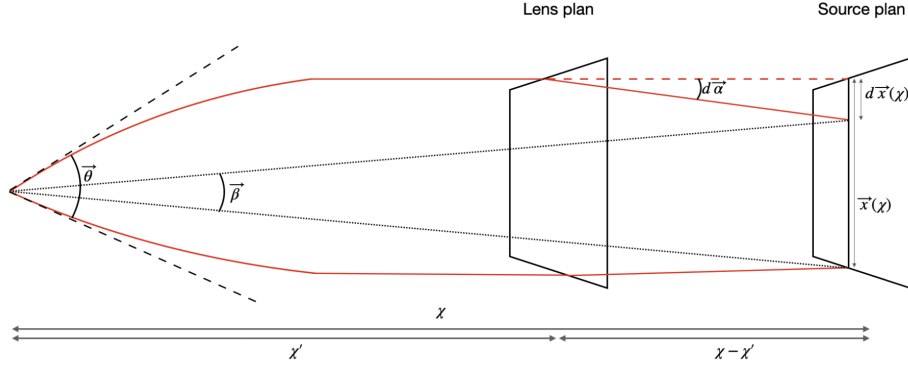


Figure 1.2-2 – Illustration of two light rays emitted from an object in the source plane, located at χ , that are then lensed by a massive object in the lens plane, located at χ' , and that finally arrive at an observer located at the origin. The bottom light ray is considered the fiducial ray. The apparent angular separation of the rays is θ , while the angle between the true positions of the source, without any lensing effect, is β . The angle α represents the difference between the two previous angles. Credit: Illustration from Guinot [Gui20].

considers an observer located in the lens plane and illustrates the angle deflection due to the gravitational potential. Let us now rewrite Equation 1.7 expressed in the comoving frame as

$$d\hat{\alpha} = -\frac{2}{c^2} \nabla_{\perp} \Phi(\mathbf{x}, \chi') d\chi', \quad (1.11)$$

which describes the deflection angle due to the presence of a potential Φ at a distance χ' from the observer. To estimate the total separation, $\mathbf{x}(\chi)$, we need to integrate over the line of sight along χ' . As the gravitational potential affects both light rays, including the fiducial ray, we need to consider the difference in the transverse gradient of the gravitational potential. The total separation writes

$$\mathbf{x}(\chi) = f_K(\chi) \boldsymbol{\theta} - \frac{2}{c^2} \int_0^{\chi} f_k(\chi - \chi') [\nabla_{\perp} \Phi(\mathbf{x}(\chi'), \chi') - \nabla_{\perp} \Phi(0, \chi')] d\chi', \quad (1.12)$$

where $\Phi(0, \chi')$ represents the potential along the fiducial light ray. The *total scaled deflection angle*, α , is difference between the apparent angle θ from Equation 1.8 and the observation angle from an unperturbed Universe from Equation 1.9. Introducing Equation 1.9 into Equation 1.12 we arrive to the expression of the *standard lens equation*

$$\alpha = \theta - \beta = \frac{2}{c^2} \int_0^{\chi} \frac{f_K(\chi - \chi')}{f_K(\chi)} [\nabla_{\perp} \Phi(\mathbf{x}(\chi'), \chi') - \nabla_{\perp} \Phi(0, \chi')] d\chi'. \quad (1.13)$$

Integrating the potential over the perturbed light path is not simple. We can assume that the value of the potential evaluated at the perturbed ray does not differ substantially from the potential on an unperturbed path. This assumption is known as the *Born approximation* and allows us to replace the separation vector \mathbf{x} by its 0-th order approximation $\mathbf{x}_0 = f_K(\chi) \boldsymbol{\theta}$ in the evaluation of the potential. See Krause, E. et al. [KH10] for a study on the impact of this approximation on cosmological analyses based on WL power spectrum. Then, the *lens equation* under the *Born approximation*

reads

$$\boldsymbol{\alpha} = \boldsymbol{\theta} - \boldsymbol{\beta} = \frac{2}{c^2} \int_0^\chi \frac{f_K(\chi - \chi')}{f_K(\chi)} [\nabla_\perp \Phi(f_K(\chi') \boldsymbol{\theta}, \chi') - \nabla_\perp \Phi(0, \chi')] d\chi'. \quad (1.14)$$

In a nutshell, Equation 1.14 expresses the difference between the apparent observation angle $\boldsymbol{\theta}$ of a source with the true angle $\boldsymbol{\beta}$ due to the perturbations in the gravitational potential between the observer and a source at a distance χ .

1.2.2 Shear and convergence

The *lens equation* gives us the relation between the observed lensed coordinates, $\boldsymbol{\theta}$, and the unlensed coordinates, $\boldsymbol{\beta}$. To study the effect of the lensing potential, we can define a linear mapping from the lensed coordinates to the unlensed coordinates by linearizing the *lens equation*. We define the *amplification matrix* as the Jacobian $\mathcal{A} := \partial\boldsymbol{\beta}/\partial\boldsymbol{\theta}$, whose elements write

$$\begin{aligned} \mathcal{A}_{ij}(\boldsymbol{\theta}, \chi) &= \frac{\partial\beta_i}{\partial\theta_j} = \delta_{ij} - \frac{\partial\alpha_i}{\partial\theta_j} \\ &= \delta_{ij} - \frac{2}{c^2} \int_0^\chi \frac{f_K(\chi - \chi') - f_K(\chi')}{f_K(\chi)} \frac{\partial^2}{\partial x_i \partial x_j} \Phi(f_K(\chi') \boldsymbol{\theta}, \chi') d\chi', \end{aligned} \quad (1.15)$$

where we have used that $\nabla_\perp = (\partial/\partial x_1, \partial/\partial x_2)$, with x_1 and x_2 the comoving transverse components, and the potential evaluated at the fiducial ray drops out as it does not depend on $\boldsymbol{\theta}$. We can extract from Equation 1.15 a useful quantity known as the *lens potential* that writes

$$\psi(\boldsymbol{\theta}, \chi) = \frac{2}{c^2} \int_0^\chi \frac{f_K(\chi - \chi')}{f_K(\chi) f_K(\chi')} \Phi(f_K(\chi') \boldsymbol{\theta}, \chi') d\chi'. \quad (1.16)$$

The expression above is helpful as it allows us to define the amplification matrix with derivatives of the lens potential as follows

$$\mathcal{A}_{ij} = \delta_{ij} - \frac{\partial^2 \psi(\boldsymbol{\theta}, \chi)}{\partial \theta_i \partial \theta_j}. \quad (1.17)$$

Let us remind that the amplification matrix provides a linear transformation between the lensed and unlensed coordinates as follows

$$\begin{pmatrix} x_u \\ y_u \end{pmatrix} = \mathcal{A} \begin{pmatrix} x_l \\ y_l \end{pmatrix}. \quad (1.18)$$

By observing Equation 1.17 and Equation 1.18 we can see some physical intuition of these quantities. The Kronecker delta δ_{ij} represents the unlensed scenario. The deviation from this scenario is quantified with the lens potential's derivatives that express the gravitational potential's impact in the light paths' directions.

The amplification matrix is parametrized using the scalar *convergence*, κ , and the two-component *shear*, $\gamma = (\gamma_1, \gamma_2)$, that is usually defined as a complex number $\gamma = \gamma_1 + i\gamma_2 = |\gamma| \exp(2i\phi)$. Both of these components are related by the *reduced shear* as follows

$$g_i := \frac{\gamma_i}{1 - \kappa}. \quad (1.19)$$

We can now express the amplification matrix as a function of the aforementioned quantities, and the new expressions writes

$$\mathcal{A} = \begin{pmatrix} 1 - \gamma_1 - \kappa & -\gamma_2 \\ -\gamma_2 & 1 + \gamma_1 - \kappa \end{pmatrix} = (1 - \kappa) \begin{pmatrix} 1 - g_1 & -g_2 \\ -g_2 & 1 + g_1 \end{pmatrix}, \quad (1.20)$$




Figure 1.2-3 – Illustration of the effect of shear on the circular object at the origin. The complex shear amplitude is left constant while the angle is changed from 0 to 2π , showing the evolution of the orientation of the ellipse. The ellipse’s eccentricity is constant throughout the angles as the shear amplitude remains unchanged. Credit: Image from Kilbinger [Kil15].

where these quantities can be expressed directly as derivatives from the lens potential as follows

$$\begin{aligned}\kappa &= \frac{1}{2} \Delta \psi = \frac{1}{2} \left(\frac{\partial^2}{\partial \theta_1 \partial \theta_1} + \frac{\partial^2}{\partial \theta_2 \partial \theta_2} \right) \psi, \\ \gamma_1 &= \frac{1}{2} \left(\frac{\partial^2}{\partial \theta_1 \partial \theta_1} - \frac{\partial^2}{\partial \theta_2 \partial \theta_2} \right) \psi, \\ \gamma_2 &= \frac{1}{2} \frac{\partial^2}{\partial \theta_1 \partial \theta_2} \psi.\end{aligned}\tag{1.21}$$

From Equation 1.20 we can disentangle the effects of convergence and shear by separating the matrix into the difference between two matrices. The first matrix is diagonal and only contains the convergence, and the second matrix is symmetric with zero trace and only contains shear elements. This matrix decomposition is written as follows

$$\mathcal{A} = \begin{pmatrix} 1 - \kappa & 0 \\ 0 & 1 - \kappa \end{pmatrix} - \begin{pmatrix} \gamma_1 & \gamma_2 \\ \gamma_2 & -\gamma_1 \end{pmatrix}.\tag{1.22}$$

The above equation helps us to obtain an intuitive understanding of the effects of the convergence and the shear to the observed images. The convergence acts as an isotropic deformation of the images. In turn, the shear causes an anisotropic deformation of the images. This second effect turns a circle into an ellipse with an orientation depending on the complex shear angle, ϕ , and an eccentricity depending on the shear amplitude, $|\gamma|$. Figure 1.2-3 illustrates the effect of shear with constant amplitude to a circle as the angle of the complex shear varies. Both effects of the amplification matrix, shear and convergence, are shown in Figure 1.2-4.

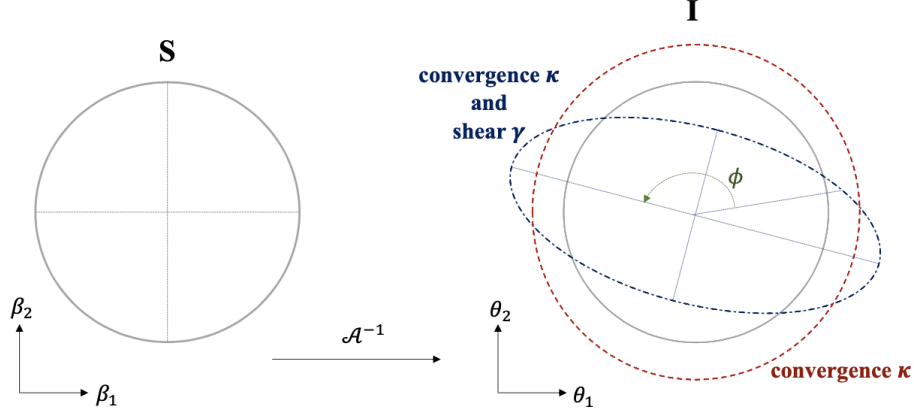


Figure 1.2-4 – Illustration of the shear and convergence effects. The circular source image is \mathcal{S} , and lensed image is \mathcal{I} . The effect of the convergence κ , shown in red, is an isotropic magnification or contraction. The effect of the shear γ , shown in blue, is an anisotropic deformation. The orientation of the ellipse ϕ depends on the complex γ phase, and the ellipse’s eccentricity depends on the complex γ amplitude. The amplification matrix \mathcal{A} relates the image coordinates, providing a linear transformation of (θ_1, θ_2) to (β_1, β_2) . Credit: Image from Ajani [Aja21].

1.2.3 Estimating the shear from observations

Let us now consider the task of estimating the shear from the observed galaxy images. We will be working in the *weak lensing regime* which assumes that the values of shear and convergence are on the order of a few per cent ($|\gamma| \ll 1$ and $|\kappa| \ll 1$). In addition, the amplification matrix is supposed to be invertible, there are no multiple images caused by lensing, and each source is mapped uniquely into one image.

From the observed galaxy images, we can estimate its shape, which is described using the complex ellipticity, which in turn is a function of the quadrupole moments of an image. These moments, $M_{\mu\nu}$, are defined as follows

$$\bar{\mu} = \frac{\int \mu \bar{I}[u, v] du dv}{\int \bar{I}[u, v] du dv}, \quad (1.23)$$

$$M_{\mu\nu} = \frac{\int \bar{I}[u, v] (\mu - \bar{\mu})(\nu - \bar{\nu}) du dv}{\int \bar{I}[u, v] du dv}, \quad (1.24)$$

where $\bar{I}[u, v]$ is the galaxy image, (\bar{u}, \bar{v}) is its centroid, and $\mu, \nu \in \{u, v\}$. There are two definitions for the complex ellipticity used throughout the literature, and both of these write

$$\epsilon = \epsilon_1 + i\epsilon_2 := \frac{(M_{uu} - M_{vv}) + i2M_{uv}}{M_{uu} + M_{vv} + 2\sqrt{M_{uu}M_{vv} - M_{uv}^2}}, \quad (1.25)$$

$$e = e_1 + ie_2 := \frac{(M_{uu} - M_{vv}) + i2M_{uv}}{M_{uu} + M_{vv}}. \quad (1.26)$$

The ellipticity can also be defined as a function of the parameters of the ellipse, both semi-axes, a and b , and the orientation ϕ . These are specific cases valid for elliptical isophotes and write

$$\epsilon = \frac{a-b}{a+b} e^{2i\phi}, \quad e = \frac{a^2 - b^2}{a^2 + b^2} e^{2i\phi}. \quad (1.27)$$

The two ellipticities definitions can be related by the following equations from Bartelmann et al. [BS01] as follows

$$\epsilon = \frac{e}{1 + \sqrt{1 - |e|^2}}, \quad e = \frac{2\epsilon}{1 + |\epsilon|^2}. \quad (1.28)$$

We can use either of those definitions, provided we do not mix their formulations. When we observe galaxy images we can measure the observed ellipticity parameters, ϵ^{obs} and e^{obs} . Let us assume that there exists an intrinsic ellipticity of the source galaxy, ϵ^{s} and e^{s} . Then, the effect of the shear field is to modify the intrinsic ellipticity to the observed one, and Seitz et al. [SS97] quantified the change for $|g| < 1$ as follows

$$\epsilon^{\text{obs}} = \frac{\epsilon^{\text{s}} + g}{1 + g^* \epsilon^{\text{s}}}, \quad (1.29)$$

$$e^{\text{obs}} = \frac{e^{\text{s}} - 2g + g^2 e^{\text{s}*}}{1 + |g|^2 - 2\text{Re}(g e^{\text{s}*})}, \quad (1.30)$$

where $*$ denotes complex conjugation, and $\text{Re}(\cdot)$ denotes the real part of its argument. In the **WL**-regime, and focusing on the first ellipticity definition, we can approximate the aforementioned relation by the following equation

$$\epsilon^{\text{obs}} \approx \epsilon^{\text{s}} + g, \quad (1.31)$$

where we approximated the denominator to unity. Measuring the ellipticity of a single galaxy does not allow us to estimate the shear. The difficulty resides in the fact that both quantities, the intrinsic ellipticity and the shear, are unknown. Nevertheless, there is a way to isolate the shear from the intrinsic ellipticity. We can assume that the intrinsic ellipticity of the galaxies is random without any preferred orientation. This assumption translates to a vanishing expectation of the intrinsic ellipticity, or $\langle \epsilon^{\text{s}} \rangle = 0$. Therefore, if we measure the ellipticity of many galaxies, we can have an estimate of the shear as follows

$$\langle \epsilon^{\text{obs}} \rangle \approx g \approx \gamma, \quad (1.32)$$

where we have also assumed that the shear applied to those galaxies is constant over the area under study and that there were no systematic errors in the galaxy shape measurement. The variance of the ellipticity can be written as $\sigma^2(\epsilon^{\text{obs}}) = \sigma^2(\epsilon^{\text{s}}) + \sigma^2(\gamma)$, where $\sigma^2(\epsilon^{\text{s}})$ is known as the *shape noise*. This noise determines the number of galaxies, N , required to measure the shear with a given **Signal-to-Noise Ratio (SNR)**. The expression of the **SNR** writes $\gamma N^{1/2} / \sigma(\epsilon^{\text{s}})$. **WL** surveys aim to have a good **SNR** in the shear estimation and, therefore, try to maximize the density of observed galaxies in the survey area.

Forthcoming **WL** surveys require high precision of the shear estimation, so some assumptions must be revisited. There is interest in the study of the assumption that the expectation of the intrinsic galaxy shapes vanishes. Indeed, there is an effect known as *intrinsic alignment*, where the shapes of galaxies are up to some degree aligned due to interactions between the different galaxies and the surrounding tidal fields without the action of gravitational lensing. This effect biases our estimation of the shear and is a source of concern for **WL** surveys like *Euclid* or **Large Synoptic Survey Telescope (LSST)**. See Troxel et al. [T15] for a review of this topic.

Finally, working in the **WL** regime and under the approximations in Equation 1.32, we have seen that we can have an unbiased estimator of the shear directly from our observations. However, the convergence, κ , is not directly measured from the observations. Inversion methods are required to estimate the convergence from the shear, as we will later see.

1.2.4 Relations between the convergence and the shear

We have seen in the previous section that we can estimate the shear from the observations of galaxies through the measurement of their ellipticities. However, we have not specified how to estimate the convergence, which will be important to relate the observations to the cosmology, as we will see in [subsection 1.2.5](#). We have seen in [Equation 1.21](#) that the shear and the convergence are related through the lensing potential. Applying the [Fourier Transform \(FT\)](#) to [Equation 1.21](#) we obtain

$$\tilde{\gamma}_1 = \frac{\ell_1^2 - \ell_2^2}{2} \tilde{\psi}, \quad \tilde{\gamma}_2 = \ell_1 \ell_2 \tilde{\psi}, \quad \tilde{\kappa} = \frac{\ell_1^2 + \ell_2^2}{2} \tilde{\psi}, \quad (1.33)$$

where \tilde{f} denotes the FT of f , and $\boldsymbol{\ell} = (\ell_1, \ell_2)$ is the **2-dimensional (2D)** Fourier wave vector, the Fourier-conjugate of $\boldsymbol{\theta}$. Working with the expressions from [Equation 1.33](#) we can find a relation between the shear and the convergence that writes

$$\tilde{\gamma}(\boldsymbol{\ell}) = \frac{(\ell_1 + i\ell_2)^2}{\ell^2} \tilde{\kappa}(\boldsymbol{\ell}) = e^{2i\beta} \tilde{\kappa}(\boldsymbol{\ell}), \quad (1.34)$$

where $\ell = \sqrt{\ell_1^2 + \ell_2^2}$, and β is the polar angle of the complex quantity $\boldsymbol{\ell}$. [Equation 1.34](#) is valid for $\ell \neq 0$.

The previous equation motivates the introduction of an important statistic highly used in cosmological analyses, the *power spectrum*. When considering cosmic shear, the first-order statistic of the shear and convergence fields is zero, as the mean of both fields vanishes. We should note that in other lensing studies, like cluster lensing, this statement does not hold. Consequently, for cosmic shear, higher-order statistics of the fields are studied, and as we will later see, they carry essential information that allows us to constrain cosmological models. The convergence power spectrum, P_κ , can be expressed as its [Two-Point Correlation Function \(2PCF\)](#) in Fourier space as follows

$$\langle \tilde{\kappa}(\boldsymbol{\ell}) \tilde{\kappa}^*(\boldsymbol{\ell}') \rangle = (2\pi)^2 \delta_{\text{D}}(\boldsymbol{\ell} - \boldsymbol{\ell}') P_\kappa(\ell), \quad (1.35)$$

where δ_{D} is the Dirac delta distribution. Note that the power spectrum only depends on the modulus of the wave vector, ℓ , which derives from the cosmological principle where κ is statistically homogeneous and isotropic. The previous assumption means that in real space the **2PCF**, $\langle \kappa(\boldsymbol{\vartheta}) \kappa(\boldsymbol{\vartheta} + \boldsymbol{\theta}) \rangle$, only depends on the modulus of the separation vector between the two angles. We recall that when using the FT to express the power spectrum, we implicitly use the *flat-sky approximation*. This approximation implied that the sky is flat and lies in two dimensions, which is as good approximation when considering small sky patches. The angles $\boldsymbol{\theta}$ represent coordinates in the celestial sphere. Therefore, the most accurate alternative would be to use spherical harmonics to describe the power spectrum.

[Equation 1.34](#) shows us that the power spectrum of the convergence coincides with the one of the shear, $P_\kappa = P_\gamma$. Because of Parseval's theorem, the **2PCF** of κ and γ are also identical.

Mass mapping methods

We have seen that our observable is the shear field that we can estimate from our observations using the galaxies' ellipticities. We have just pointed out that the shear field is tightly related to the convergence field, which carries much cosmological information, as we will see in [subsection 1.2.5](#). One way to analyse the convergence field is to use an inversion technique to estimate it from the shear field. These techniques are known as *mass mapping methods*.




Figure 1.2-5 – Mass maps using DES year 3 data release generated with the KS method [KS93], and the sparsity-based method GLIMPSE [Lan+16]. Credit: Figures from Jeffrey et al. [Jef+21].

One of the classical inversion methods is Kaiser-Squires (KS) inversion method [KS93] which combines Equation 1.33 to form an estimator of the convergence as a function of the shear as follows

$$\hat{\kappa} = \frac{\ell_1^2 - \ell_2^2}{\ell^2} \tilde{\gamma}_1 + \frac{2\ell_1 \ell_2}{\ell^2} \tilde{\gamma}_2 = \frac{\ell_1^2 + \ell_2^2}{(\ell_1 + i\ell_2)^2} \tilde{\gamma}, \quad (1.36)$$

which is not defined for $\ell = 0$. The KS method cannot estimate a constant convergence, represented by the zeroth Fourier value. The shear field is invariant to a constant convergence and, therefore, cannot be recovered by the method. This problem is known as the *mass-sheet degeneracy*. More mass mapping methods were developed throughout the years based on inpainting [Pir+09], sparsity [Lan+16], deep learning [Jef+20], and hybrid approach using sparsity and Gaussian random fields [Sta+21]. Jeffrey et al. [Jef+21] produced the largest, up to now, convergence maps from the Dark Energy Survey (DES) year 3 data release using different mass mapping methods. In Figure 1.2-5 we show two of the mass maps generated using the classical method from Kaiser et al. [KS93] and a new method from Lanusse, F. et al. [Lan+16]. The map being plotted is the E-mode of the convergence field, which corresponds to the curl-free component of the gradient field defined by the potential κ . See Kilbinger [Kil15, §3.6] for more details on E- and B-modes.

Shear two-point correlation function

Another way to study the properties of the convergence field is to exploit its relation with the shear 2PCF. The shear 2PCF can be easily computed without passing through the Fourier space. Each measured shear can be decomposed into a tangential component, γ_t , and a cross-component, γ_\times , which are defined as

$$\gamma_t = -\text{Re}(\gamma e^{-2i\phi}), \quad \gamma_\times = -\text{Im}(\gamma e^{-2i\phi}), \quad (1.37)$$

where ϕ is the polar angle of the complex quantity θ , and $\text{Re}(\cdot)$ and $\text{Im}(\cdot)$ are the real and imaginary parts of their arguments. There are three possible 2-point correlators can be built, which are $\langle \gamma_t \gamma_t \rangle$, $\langle \gamma_\times \gamma_\times \rangle$, $\langle \gamma_t \gamma_\times \rangle$. These correlator should be invariant through a parity transformation, that is exchanging $\theta \rightarrow -\theta$, which changes the sign of γ_\times , but leaves unchanged γ_t . Using the two non-zero correlators, we can

build the two components of the shear **2PCF**

$$\xi_+(\boldsymbol{\theta}) = \langle \gamma \gamma^* \rangle(\boldsymbol{\theta}) = \langle \gamma_t \gamma_t \rangle(\boldsymbol{\theta}) + \langle \gamma_\times \gamma_\times \rangle(\boldsymbol{\theta}), \quad (1.38)$$

$$\xi_-(\boldsymbol{\theta}) = \langle \gamma \gamma \rangle(\boldsymbol{\theta}) = \langle \gamma_t \gamma_t \rangle(\boldsymbol{\theta}) - \langle \gamma_\times \gamma_\times \rangle(\boldsymbol{\theta}). \quad (1.39)$$

There exist other formulations for the shear **2PCF**, for example, in Jarvis et al. [Jar+03], which exploits the fact that the positive cross-component is rotated by $\pi/4$ with respect to the tangential component. Schneider, P. et al. [Sch+02] provided a useful estimator for the shear **2PCF** using the following formula

$$\hat{\xi}_\pm(\boldsymbol{\theta}) = \frac{\sum_{i,j} w_i w_j (\epsilon_t^{\text{obs}}(\boldsymbol{\theta}_i) \epsilon_t^{\text{obs}}(\boldsymbol{\theta}_j) \pm \epsilon_\times^{\text{obs}}(\boldsymbol{\theta}_i) \epsilon_\times^{\text{obs}}(\boldsymbol{\theta}_j))}{\sum_{i,j} w_i w_j}, \quad (1.40)$$

where ϵ_t and ϵ_\times are the tangential and cross ellipticity that can be computed following the recipe in Equation 1.37, and the weights w_i and w_j correspond to the measurement uncertainty of the corresponding ellipticities ϵ_i and ϵ_j . The summation is over pairs of galaxies (i, j) with corresponding sky positions $\boldsymbol{\vartheta}_i$ and $\boldsymbol{\vartheta}_j$ whose separation $|\boldsymbol{\vartheta}_i - \boldsymbol{\vartheta}_j|$ lies in an angular distance bin around $\boldsymbol{\theta}$.

The proposed shear **2PCF** estimator uses the observed ellipticity that, as we have seen, contains an intrinsic ellipticity component in addition to the shear. Correlating the observed ellipticity would give us the sum of the correlations between these two components, for ξ_+ it would be: $\langle \epsilon_i^s \epsilon_j^{s*} \rangle$, $\langle \epsilon_i^s \gamma_j^{s*} \rangle$, $\langle \gamma_i \epsilon_j^{s*} \rangle$, and $\langle \gamma_i \gamma_j^* \rangle$. Sticking to the weak lensing regime and the assumption made in subsection 1.2.3 that the mean intrinsic ellipticity is zero, the first three correlations vanish, and we only get the last one.

1.2.5 Relation to cosmology

We have described how to estimate the shear field from our observations, i.e. galaxy images. We continued to study the relations between the shear and the convergence fields and how we can estimate the convergence from our shear estimation. Later, we have seen that the shear **2PCF** also provides us with valuable information about the 2-point statistics of the convergence field. It is now time to relate the **WL** products we studied to cosmology.

We are particularly interested in the convergence as we will see that it can be interpreted as the projected matter density distribution over the line of sight between the observer and the source. Let us recall from Equation 1.21 that the convergence κ is related to the lensing potential through a **2D** Poisson equation. Introducing the lensing potential formula in Equation 1.16 to the Poisson equation, we obtain

$$\kappa(\boldsymbol{\theta}, \chi) = \frac{1}{2} \nabla^2 \psi(\boldsymbol{\theta}, \chi) = \frac{1}{c^2} \int_0^\chi \frac{f_K(\chi - \chi')}{f_K(\chi) f_K(\chi')} \nabla^2 \Phi(f_K(\chi') \boldsymbol{\theta}, \chi') d\chi', \quad (1.41)$$

where the Laplacian, ∇^2 , is taken with respect to the angular coordinates $\boldsymbol{\theta}$. Let us now introduce the *matter density contrast* which writes

$$\delta = \frac{\rho - \bar{\rho}}{\bar{\rho}}, \quad (1.42)$$

and its relation to the gravitational potential through the Poisson equation using a **3-dimensional (3D)** Laplacian as follows

$$\nabla^2 \Phi = 4\pi G a^2 \bar{\rho} \delta, \quad (1.43)$$

where ρ is the energy-matter density, $\bar{\rho}$ is the mean energy-matter density, and the matter density contrast δ represents the energy-matter density fluctuations around its mean. The 2D angular Laplacian from Equation 1.41 can be turned into the 3D Laplacian in comoving coordinates of Equation 1.43 by adding an additional term $\partial^2\Phi/\partial\chi'^2$ to the integrand of Equation 1.41. Then, we can use the argument that the added term is expected to vanish in the line of sight integral under homogeneity arguments. Equation 1.43 is a consequence of considering that Newtonian physics suffices to describe non-relativistic perturbations in a matter-dominated era of the Universe on scales smaller than the horizon¹ [Pee80]. We refer the reader to Peebles [Pee80; Pee93] for more details on the matter density and its evolution in the Universe through time. Using Equation 1.43 and $\bar{\rho} \propto a^{-3}$ we can rewrite Equation 1.41 in terms of the density contrast δ as follows

$$\kappa(\boldsymbol{\theta}, \chi) = \frac{3H_0^2\Omega_m}{2c^2} \int_0^\chi \frac{f_K(\chi - \chi')f_K(\chi')}{f_K(\chi)} \frac{\delta(f_K(\chi')\boldsymbol{\theta}, \chi')}{a(\chi')} d\chi', \quad (1.44)$$

where we have expressed the mean matter density $\bar{\rho}$ as a function of the matter density Ω_m , H_0 is the Hubble constant, and a is the scale factor. Equation 1.44 expresses how the convergence field is the result of integrating over the line of sight of the density contrast, weighted by the scale factor a and geometrical considerations of the observer, the lens and the source comoving distances. In other words, we can consider the convergence as the *projected mass* or *projected overdensity* on the sky.

We now extend the analysis to different distances and consider a WL survey that probes the Universe until a limiting distance χ_{lim} and observes a given sample of galaxies. Let us define $n(\chi)$ as the Probability Density Function (PDF) of the galaxy samples as a function of distance. The $n(\chi)$ function informs us about how far away the observed galaxies are from us. It is common to see this PDF expressed in terms of redshift, previously explained in section 1.1, as $n(z)$ which can be related to the formulation using the comoving distance using $n(z)dz = n(\chi)d\chi$. The mean convergence of the sample of galaxies under study can be written as

$$\kappa(\boldsymbol{\theta}) = \int_0^{\chi_{\text{lim}}} n(\chi) \kappa(\boldsymbol{\theta}, \chi) d\chi. \quad (1.45)$$

For a given position $\boldsymbol{\theta}$ in the sky, we can put together Equation 1.44 and Equation 1.45, and obtain

$$\kappa(\boldsymbol{\theta}) = \frac{3H_0^2\Omega_m}{2c^2} \int_0^{\chi_{\text{lim}}} \frac{f_K(\chi)g(\chi)}{a(\chi)} \delta(f_K(\chi)\boldsymbol{\theta}, \chi) d\chi, \quad (1.46)$$

where g is the *lens efficiency* that is defined as

$$g(\chi) = \int_\chi^{\chi_{\text{lim}}} n(\chi') \frac{f_K(\chi' - \chi)}{f_K(\chi')} d\chi', \quad (1.47)$$

represents the strength of the observed lensing signal at a distance χ given some background galaxies at a distance χ' . Equation 1.45 and Equation 1.46 show the line of sight projection for the population of all observed galaxies, averaging over all the distances. In practice, performing the analysis binning the source redshift $n(z)$ is more informative. When we do averages over different redshift bins for the observed galaxies, we can study how the cosmic structure has grown with time. This WL analysis in *redshift bins* is known as *tomographic WL*, which leads to a lens

¹The distance from which one could possibly retrieve information.

efficiency per redshift bin, and it has more constraining power than the **2D** analysis from Equation 1.46. However, the tomographic analysis relies on the availability of redshift information for the observed objects, which can be estimated from *photometric redshifts* or *photo-z*. This method estimates an object's redshift from photometric measurement in different spectral bands. See Salvato et al. [SIH18] for a review on photometric redshift measurement methods.

We have seen until now that the convergence field is related to cosmological quantities in Equation 1.46. We can now extend the analysis to the **2PCF** of the convergence by applying the **FT** to Equation 1.46 and using the result in Equation 1.35 to finally obtain the convergence power spectrum

$$P_\kappa(\ell) = \frac{9}{4} \Omega_m^2 \left(\frac{H_0}{c} \right)^4 \int_0^{\chi_{\text{lim}}} \frac{g^2(\chi)}{a^2(\chi)} P_\delta \left(k = \frac{\ell}{f_K(\chi)}, \chi \right) d\chi, \quad (1.48)$$

where P_δ is the density power spectrum, and g is the lens efficiency from Equation 1.47. To obtain the previous expression, we have used several approximations: the *limber approximation* [Lim53; Kai92] where we neglect correlations along the line of sight; the *small-angle approximation* where we approximate trigonometrical functions by their first order expansion; and the *flat-sky limit* where instead of working with spherical harmonics we use Fourier modes and neglect the spherical nature of the sky. Kilbinger et al. [Kil+17] studied the impact of the Limber approximation and the flat-sky limit on the shear and the convergence power spectrum. They concluded that future **WL** surveys will require working with a spherical sky and that a second order Limber approximation is necessary.

We have introduced the building blocks that relate our observations to cosmological information. We described the relation between the gravitational potential perturbations to the deflection angle due to gravitational lensing. Then, how can these deflections be quantified in the shear and convergence fields that are finally probing our Universe, providing us with a rich source of cosmological information. Until now, we have focused on the theoretical formulation of a **WL** analysis. However, given the statistical nature of the analysis, its success relies upon controlling the systematic errors introduced in practice. The shape measurement pipeline, which is in charge of producing a **WL** shear catalogue, is in the centre of the systematic control. The latter of this chapter will focus on the practical part of the **WL** analysis.

1.3 Main ingredients of a shape measurement pipeline

In the previous section, we have seen how we can exploit the shapes of the observed galaxies to learn from and analyse our Universe. However, we have not detailed how to estimate such shapes from real observations and the many difficulties it implies.

In this section, we introduce the principal components of a *shape measurement pipeline*, whose main job is to process the observed images and create a *calibrated* shape catalogue. We emphasise the importance of the different *calibration* parts as they constitute a crucial part of the overall task. There are many sources of systematic errors in the shape measurement process, with the impact of the **PSF** in the images being the most noticeably. In the **WL** analysis, we have seen that we use the mean galaxy ellipticity as a shear estimator. This formulation is sensible to systematic errors in the measurement that will turn into estimation biases which will finally impact the cosmological analyses carried out with the shape catalogue. The overview we will present will allow us to understand better where and how the **PSF** plays a role in the shape measurement pipeline and finally in the **WL** analysis.

This section is greatly based on my contributions to **ShapePipe**, a modular shape measurement pipeline described in Farrens, S. et al. [Far+22a]. The pipeline has been used to produce the first version of the shape catalogue of **Canada-France Imaging Survey (CFIS)** in Guinot et al. [Gui+22] (see Guinot [Gui20] for a more detailed description). I contributed to the implementation of a new **PSF** modelling method described in chapter 3 as well as in its validation with real data and its compatibility with the shape measurement method used. These contributions will be described in more detail in chapter 4. Dr Martin Kilbinger led the production of a second version of the catalogue, including the new **PSF** model and doubling the area analysed with respect to the first version. **ShapePipe** was used to analyse the data, and the results will soon be published. We refer the reader to Bosch et al. [Bos+17] (Jarvis et al. [Jar+16]) for a detailed description of the shape measurement pipeline used to process **HSC (DES)** data, and Mandelbaum et al. [Man+17] (Zuntz et al. [Zun+18]) for a description of the produced shape catalogue with different validation procedures. In addition, see Mandelbaum [Man18] for a review of the steps needed to go from images to a shape catalogue.

In Figure 1.3-6 we give a simple overview of the shape measurement pipeline's building blocks, emphasising the **PSF** as it is the main topic of this thesis. In this overview, we are not making distinctions over using single exposures or stacked images in the processing steps for the sake of simplicity. In the following subsections, we will continue by describing the main processes seen in Figure 1.3-6.

1.3.1 Image preprocessing

In the image preprocessing stage, we regroup several important tasks. We start with the astrometry and photometry calibrations. *Astrometric* calibration refers to having a good knowledge of the positions of the pixels in the sky. The sky curvature and some detector effects make mapping pixel positions to sky positions non-linear and not trivial. Most astronomical studies need a good knowledge of the **World Coordinate System (WCS)** transformation, which maps local pixel position to sky positions. *Photometric* calibration refers to converting instrumental magnitudes (measurements of brightness or flux) to calibrated apparent magnitudes in a standard system. This calibration is important in order to be able to compare observations from different instruments and surveys, but also subsequent observations of the same instrument. See Hogg [Hog22] for a detailed description of what magnitudes and apparent magnitudes are. It is common to use another survey as a reference for calibration, for example, using **European Space Agency (ESA)**'s *Gaia* mission for the astrometry and **Sloan Digital Sky Survey (SDSS)** for the photometry. The **MegaPipe** pipeline was used for preprocessing the **CFIS** images used in chapter 4, and is described in Gwyn [Gwy08] (with a its new version in Gwyn [Gwy19]).

Other preprocessing steps are related to handling the **Charged Coupled Device (CCD)** technology. They include the detrending of the images (bias-subtracting, flat-fielding, etc...), handling the **CCD**-specific backgrounds, and other electronic effects like cross-talk or saturation. The **Elixir** pipeline [MCo4] was used to perform these preprocessing steps for the **CFIS** images.

Next, many detector effects need to be corrected. These effects include **Charge Transfer Inefficiency (CTI)** [Rho+10; Mas+14], and **Brighter-fatter Effect (BFE)** [Cou+18] that will be described with more detail in subsection 2.3.2. The preprocessing must also address cosmic rays and satellite trails in the images.

Finally, another aspect of preprocessing is the stacking of images. This procedure combines different single exposures of the same region of the sky to produce a single

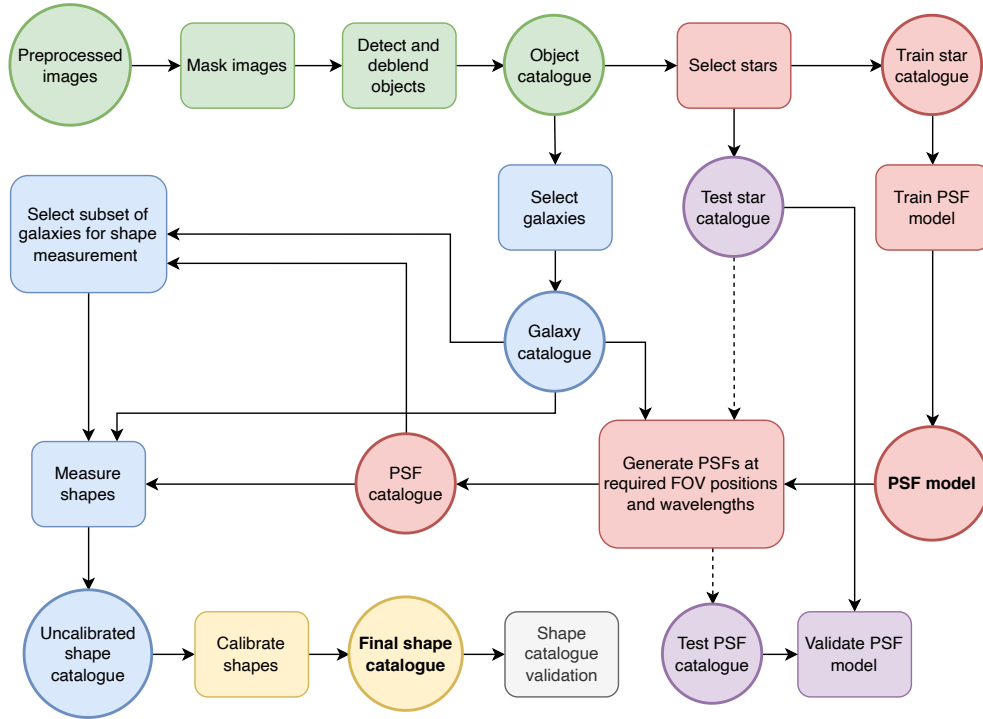


Figure 1.3-6 – Overview of the main building blocks of a shape measurement pipeline. Circles represent data or models, and rectangles represent processes. *PSF catalogue* represents the *PSFs* at the position of galaxies, while *Test PSF catalogue* represents the *PSFs* at the position of the test stars. We present the colour description of the different stages: *green* represent the common input data processing; *red* represent the *PSF* modelling part; *violet* represent the *PSF* validation part; *blue* represent the shape measurement part; *yellow* represent the shape measurement calibration; and *grey* the shape catalogue validation.

image known as a *stacked* image or *co-added* image. This procedure has advantages and disadvantages, and it depends on the image's ultimate goal. Stacking images reduce the noise level, thus increasing the *SNR*. It can also help to deal with cosmic rays and to fill the boundaries between the focal plane *CCD* chips. However, modelling the *PSF* becomes extremely challenging when we stack the images. Every single exposure has a specific *PSF* affecting the observation that depends on the properties of the atmosphere and the telescope at the time of the observation. Then, the *PSF* affecting the stacked image depends on the *PSF* of every single exposure modified by the stacking procedure. When processing the *CFIS* observations using *ShapePipe* [Gui+22] we used the single exposures to build a *PSF* model for each exposure that we then used for the shape measurement. However, we used the stacked images as input for the detection algorithm. This choice is useful, as the stacked images' high *SNR* allowed better detection results. The single exposures and the stacked images are related by the sky coordinates, that are common to both of them.




Figure 1.3-7 – Example of a [CCD](#) image and the mask produced by [ShapePipe](#). The colour bar shows the pixel mask values. Credit: Image from Farrens, S. et al. [[Far+22a](#)].

1.3.2 Masking

There are regions of the images that are unusable for a [WL](#) analysis. *Masking* refers to identifying those regions and their automated masking (i.e. mark the pixels as unusable). The principal sources of unusable regions are bright stars that saturate the pixels and have large diffraction spikes and nearby galaxies that are too close to us to be useful for [WL](#). Also, problematic pixels that are damaged or too close to the [CCD](#) edge are often masked. [Figure 1.3-7](#) provides an example of the original [CCD](#) image and the mask produced by [ShapePipe](#).

1.3.3 Object detection and deblending

Object (or source) detection refers to the task of converting a [CCD](#) image into a list of objects, each one represented as a small subset of pixels extracted from the original [CCD](#) image. For example, in the case of [CFIS](#), the [CCD](#) image is a matrix of 2048×4612 pixels, and then each object's postage stamp is an extraction of 51×51 pixels centred in the object. The size of a [WL](#) survey obliges automatizing the process. [SExtractor](#) [[BA96](#); [Ber11](#)] is the reference and widely used software and is being used in [ShapePipe](#) [[Far+22a](#)]. On top of detecting the objects, [SExtractor](#) also measures several properties of the object.

Another essential task is to identify blended objects and eventually mask or deblend them. Blended objects refer to objects whose projection in the focal plane overlaps, meaning that one pixel has a flux contribution of more than one object. [SExtractor](#) can identify blended objects and mask them. More sophisticated methods [[Mel+18](#)] were developed to deblend the objects and rely on multi-band observations and constrained matrix factorization optimization techniques. The deblending procedure refers to producing, from a single image, two or more postage stamps where each one contains a

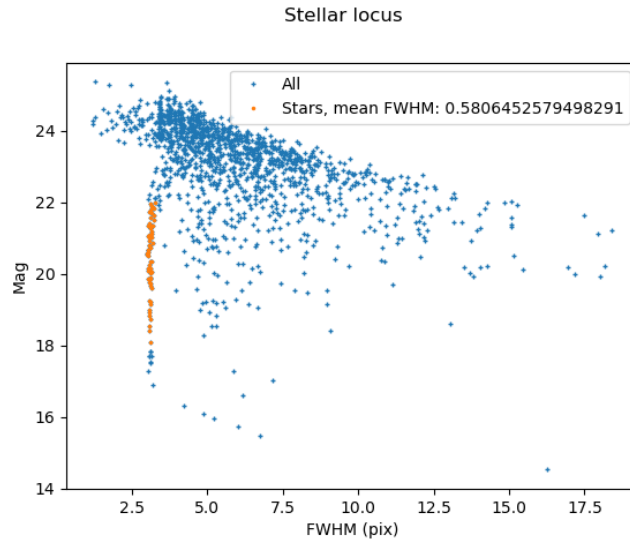


Figure 1.3-8 – ShapePipe star selection from a magnitude-size diagram. The size is represented as Full Width Half Maximum (FWHM) expressed in pixels. The FWHM expressed in the legend is in arcsec. Credit: Image from Farrens, S. et al. [Far+22a].

single object. The new era of WL surveys, with *Euclid* and LSST, will go deeper than before, being able to increase the number of observed objects. However, this increases the probability of observing blended objects, making it an even more critical task. Another approach based on deep neural networks was recently developed [Far+22b], but it still lacks the maturity to handle the difficulties of working with real data. We refer the reader to Melchior et al. [Mel+21] for a review in deblending for large sky surveys.

1.3.4 Star-galaxy separation

In the pipeline description, we are at the point where we built a list of objects from the observations. It is now the moment to classify them into stars or galaxies, use the former to build the PSF model and the latter to measure their shapes. On the one hand, it is important to have a pure star sample as we are going to consider that the observations give information of the PSF field, which on the contrary, would bias our PSF model. On the other hand, it is also important that the galaxy sample is pure. Otherwise, the WL analysis will be biased as the measured ellipticity will be the one of a star, unsheared, and not of a galaxy. We can conclude with the importance of trustworthy classification.

A common approach to identify stars is to exploit the fact that they are not resolved objects. These objects share a similar size, the size of the PSF, and do not change size regardless of their luminosity. The previous statement assumes that the images are not saturated. If we plot the observed objects in a single exposure in a size-magnitude scatter plot, we can identify a straight line of objects that vary in magnitude but not in size, which is known as the *stellar locus*. This fact is exploited to select the stars used to build the PSF model. Figure 1.3-8 shows the selection of stars in a CFIS exposure by ShapePipe.

The selection of galaxies seems trivial after the stars have been removed. Nevertheless, there are spurious detections considered as objects and objects not well suited

for WL analyses. It is then necessary to make a further selection of objects that will constitute the galaxy sample to be used for the WL analysis. A common approach is to compute the *spread model* [Moh+12; Des+12]. This quantity is built based on the ratio of the comparison of the object to a parametric PSF and an extended parametric object. Then, the galaxy sample is selected using some cuts (selection criteria) on the spread model quantity and other variables like the magnitude and SNR.

1.3.5 PSF modelling

Once a clean star sample has been selected, we can proceed to build the PSF model. The entire set of stars is usually divided into two subsets. The first subset ($\sim 80\%$), or *train star catalogue*, will be used for learning or estimating the PSF model, and the second subset ($\sim 20\%$), or *test star catalogue*, will be used to validate the PSF model. When we work with real data, we do not know the ground truth or the underlying PSF field. Therefore, we evaluate the performance of the learned PSF model when predicting the PSF field at the test positions compared to the observed test stars.

Once the PSF model has been trained, it will be used to infer the PSF at the position and wavelengths from the selected *galaxy catalogue* to produce the *PSF catalogue*. The inferred PSFs will be later used by the shape measurement algorithm.

A more detailed description of the PSF is given in chapter 2. We discuss how the PSF affects our images in section 2.1 and how we model the effect of the PSF in section 2.2. Then, we continue describing the contributors to the PSF in section 2.3, and end with the validation of the PSF model in section 2.6. An in-depth study of the modelling of the PSF is given in the following chapters of the thesis.

1.3.6 Galaxy shape measurement

The WL analysis relies on the measurement of the shear, which in turn relies on the measurement of the galaxies' shapes. The shape measurement method aims to provide the best estimate for the galaxies' shapes, expressed in terms of their ellipticity and size, taking into account the PSF effect on the images. The orders of magnitude considered for the shear in the weak-lensing regime are considerably smaller than the anisotropy of the PSF, which changes the shape of the observed galaxy. Therefore, it is crucial to properly consider the PSF in the shape measurement.

Many approaches exist for shape measurement and have been evolving throughout the years to more powerful instruments and surveys that require higher precision and tighter constraints on systematic errors. The first approaches were based on the measured moments of the images, the galaxy and the PSF. Then, a correction term was deduced as a function of the PSF shape. This constitutes the main idea of the Kaiser-Squires-Broadhurst (KSB) method [KSB95], which was later improved Kaiser [Kaioo] and Viola et al. [VMB11].

Recently, a machine-learning-based method was proposed by Tewes, M. et al. [Tew+19]. The method uses shallow neural networks trained in a supervised manner with simulations to learn to regress the shear from specific measured properties of the galaxy.

In another family of methods, the shape measurement considers the effect of the PSF on the galaxy image in the pixel space. These methods, known as *model-fitting*, use a parametric family to describe the galaxy image, for example, a Sérsic profile Sérsic [Sér63]. A forward operator then takes in the galaxy profile and the PSF, and outputs an estimated observed galaxy image. Once the pixel sampling matches the galaxy observation sampling, the estimation can be compared to the observation. The previous procedure allows to optimise over or infer the galaxy profile

parameters, representing the galaxy prior to the effects of the PSF. Then, the shape of the parametric family is estimated or computed in case there is an analytic formula relating the galaxy profile to its ellipticity. This family of methods have been used in several surveys: `im3shape`² [Zun+13] in DES [Zun+18; Gat+21]; `ngmix`³ [She15; SH17] in DES [Zun+18] and CFIS [Gui+22]; and `lensfit` Miller et al. [Mil+13] in Canada-France-Hawaii Telescope Lensing Survey (CFHTLenS) Miller et al. [Mil+13] and Kilo-Degree Survey (KiDS-1000) [Kui+15; Gib+21].

Another line of research exists that proposes to infer the ensemble shear statistic from a posterior shear estimate. The driving idea is to bypass the need to estimate each object’s ellipticity and avoid estimating the shear from averaging galaxy shapes. These methods heavily rely on Bayesian formalism and are more principled. However, they have not yet been used for a WL survey. The first of these methods is *Bayesian Fourier Domain* (BFD) [BA14; Ber+16], which builds the method in the Fourier domain. The method attempts to estimate the lensing-induced shear directly in the Fourier space instead of going through per-object shape measurements. For the second method, Schneider et al. [Sch+15] proposed a Bayesian hierarchical model approach to infer the shear. Even if the method sounds promising, it must tackle the computational burden of performing inference in high-dimensional spaces.

1.3.7 Calibration

Once we have carried out all the previous steps to have a shear estimate, we will inevitably have some bias in our estimate. The shear bias can have several sources from different pipeline parts. The most important sources are: *model bias* [Kac+14], a consequence of using parametric families of profiles for the galaxies that cannot faithfully represent the observed galaxies; *noise bias* [MV12; Kac+14], a consequence of the noise present in the observations and the non-linear relationship between the image pixels and the galaxy’s true shape; *PSF bias* [Jar+16], a consequence of PSF modelling errors and also errors in the handling of the PSF by the shape measurement algorithm; *detection bias* [She+20], a consequence of the fact that the criteria used to detect galaxies depends on the shear applied to the galaxies; and *blending bias* [Mac+21], a consequence of the presence of neighbouring galaxies impacting the shape measurement of the target galaxy. See Pujol, Arnau et al. [Puj+20] for a study on different sources of shear bias.

The shear estimate is usually parametrized using a first-order Taylor expansion. We follow Jarvis et al. [Jar+16], and we write the shear estimate as

$$\hat{g}_i \approx (1 + m_i) g_i + c_i + \alpha \epsilon_i^{\text{PSF}}, \quad (1.49)$$

where $i \in \{1, 2\}$, g_i is the error-free estimator, m_i is the *multiplicative bias*, c_i is the *additive bias*, ϵ_i^{PSF} is the PSF ellipticity, and α is the *leakage* of the PSF shape to the shear estimation. Some authors include the leakage term into the additive bias term. The main objective of the shear calibration in a WL pipeline is to estimate the terms from Equation 1.49 to be able to correct them. There are different ways to attempt to estimate these terms.

A usual way to calibrate the shear measurement, which reduces to estimating the parameters from Equation 1.49, is to use simulations. The procedure consists in building a set of simulations that should be as realistic as possible. The shape measurement pipeline is then used to process the data and estimate the shear. As

²<https://bitbucket.org/joezuntz/im3shape-git/src/master/>

³<https://github.com/esheldon/ngmix>

the ground truth shear is known, the one used to generate the simulations, the shear estimation bias can be computed. [Man+18] followed this approach to calibrate the data from [Hyper Suprime-Cam survey \(HSC\)](#).

A new approach was recently developed in Huff et al. [HM17] and Sheldon et al. [SH17]. The method is coined *metacalibration* and bypasses the need for simulations that require considerable effort to produce at a realistic level and that are computationally expensive. The main idea is to estimate with the real data a response matrix that represents the response of the shape measurement to a change in the shear. In other words, the method estimates the derivative of the measured ellipticity with respect to the shear using a finite difference approximation. The method expands the ellipticity estimator, coming from the shape measurement method, in a Taylor series around zero shear as follows

$$\hat{e} = \hat{e}|_{\gamma=0} + \left. \frac{\partial \hat{e}}{\partial \gamma} \right|_{\gamma=0} \gamma + \dots, \quad (1.50)$$

$$\hat{e} \approx \hat{e}|_{\gamma=0} + \mathbf{R}\gamma, \quad \text{with} \quad \mathbf{R} := \left. \frac{\partial \hat{e}}{\partial \gamma} \right|_{\gamma=0}, \quad (1.51)$$

where \mathbf{R} is the response matrix.

In practice, the method builds on observations, a shape measurement method and a PSF model. The galaxy observation is deconvolved by its corresponding PSF in Fourier space to account for the PSF anisotropy. This step is done without specific care about the image noise. The deconvolved galaxy is then reconvolved by an isotropic PSF which is larger in size than the original PSF. This last step helps to attenuate the high-frequency content generated in the deconvolution step, as a convolution with an isotropic PSF is equivalent to applying a low-pass filter. A positive and negative shear per component is applied to the galaxy observation, $\gamma_1^\pm, \gamma_2^\pm$, following the linear transformation from Equation 1.18 and Equation 1.20. The shapes of each the resulting images are computed giving \hat{e}_1^\pm and \hat{e}_2^\pm . The response matrix, \mathbf{R} , is then computed as follows

$$R_{i,j} \approx \frac{\hat{e}_i^+ - \hat{e}_i^-}{\Delta\gamma_j}, \quad (1.52)$$

where $\Delta\gamma_j = \gamma_j^+ - \gamma_j^-$ is the difference between the two applied shears, and $\gamma_{i,j}^+ = -\gamma_{i,j}^-$ giving $\Delta\gamma_{i,j} = 2\gamma_{i,j}$. Finally, the shear calibration is done as follows

$$\langle \gamma \rangle \approx \langle \mathbf{R} \rangle^{-1} \langle \hat{e} \rangle \approx \langle \mathbf{R} \rangle^{-1} \langle \mathbf{R} \gamma \rangle. \quad (1.53)$$

This calibration method was used in the shape measurement pipeline to process DES data in Zuntz et al. [Zun+18] and Gatti et al. [Gat+21]. There is an extension of the method, named *metadetection* [She+20], that aims to calibrate the selection bias.

1.4 The *Euclid* space mission

We have seen in the previous section the main blocks involved in a shape measurement pipeline synthesized in Figure 1.3-6. The job of the PSF model and its interaction with the other blocks of the pipeline were described. The final product of the pipeline is a calibrated shape catalogue that can be used as a fundamental input for a WL analysis, as we have seen in section 1.2. There are several successful ground-based surveys that produced WL shape catalogues like [KiDS-1000](#), [DES](#), and [HSC](#). These three surveys correspond to Stage III cosmological experiments from the definition of Albrecht et al. [Alb+06]. In order to produce an even better and more powerful




Figure 1.4-9 – A contour plot illustrating the potential perceived by an object at rest in the Sun-Earth system. The different Lagrange points can be identified. Credit: NASA.

catalogue, the [ESA](#) started to build one of the Stage IV cosmological experiments, the *Euclid* space mission [[Lau+11](#)]. *Euclid* is a medium class mission from [ESA](#)'s *Cosmic Vision* program. See Laureijs et al. [[Lau+11](#)] for *Euclid*'s study report and Racca et al. [[Rac+16](#)] for *Euclid*'s mission design.

The main objective of *Euclid* is to learn about the *Dark Universe*, the nature of dark energy and the distribution of dark matter. *Euclid* has been designed to maximize the power of its two main cosmological probes. First [Weak Gravitational Lensing \(WL\)](#), which we have introduced in [section 1.2](#). Second, [Galaxy Clustering \(GC\)](#), which refers to the 3D distribution of galaxies, measured from the positions of the sky and their distance or redshift. [GC](#) requires measuring the redshift distribution of galaxies through their $H\alpha$ emission line using slitless spectroscopy.

The *Euclid* mission has a lifespan of 6 years, and it will be launched to the L2 Sun-Earth Lagrange point. [Figure 1.4-9](#) shows a contour plot of the gravitational potential that relates to the force applied to an object at rest in the Sun-Earth system. The L2 point is a saddle point, making it stable and particularly attractive for surveys interested in the night sky. A satellite at the L2 point always keeps the Sun, Earth, and the Moon behind it, providing regular communication with the Earth and solar power. The L2 point is currently home for the *Planck* and [James Webb Space Telescope \(JWST\)](#) satellites.

Euclid uses a Korsch telescope consisting of a 1.2 m primary mirror with a focal length of 24.5 m designed to provide a large [Field-of-View \(FOV\)](#). It will carry two instruments specially designed for the two main cosmological probes. *Euclid*'s [VIS Instrument \(VIS\)](#) [[Cro+16](#)] for [WL](#) and *Euclid*'s [Near Infrared Spectrometer and Photometer instrument \(NISIP\)](#) [[Euc+22b](#)] for [GC](#). From now on, we will focus on [WL](#) and the former instrument. *Euclid*'s [VIS Instrument \(VIS\)](#) has a focal plane composed of 36 [CCDs](#), where each one is a matrix of 4096×4132 pixels. The scale of each pixel is 0.1 arcsec which gives a wide [FOV](#) of about 0.57 deg^2 (around two orders of magnitude greater than [Hubble Space Telescope \(HST\)](#) [FOV](#)). The [VIS](#) instrument has a single spectral band ranging from 550 nm to 900 nm.

Cropper et al. [[Cro+13](#)] stated that the power of a [WL](#) survey depends on five primary factors, which are the driving parameters when designing a space-based [WL](#). Let us review the five factors and see why *Euclid* will be a powerful [WL](#) survey:

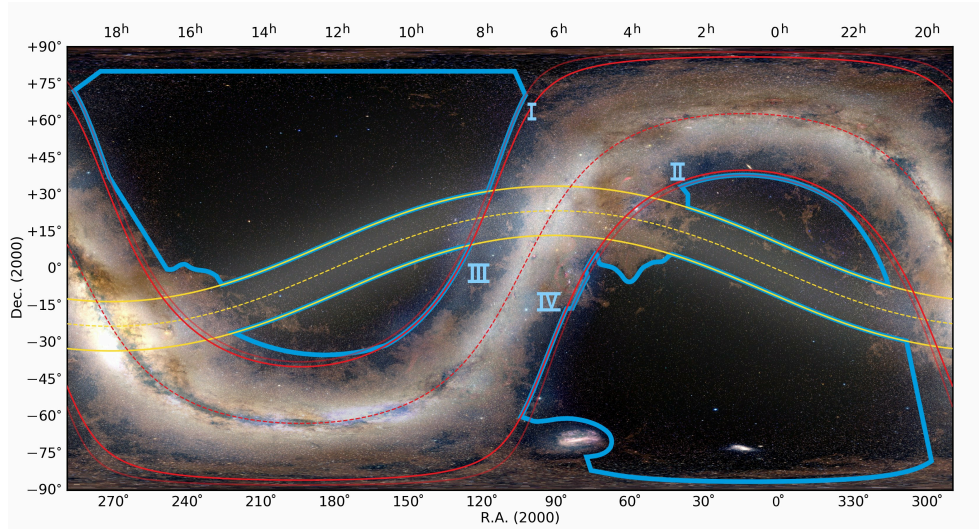


Figure 1.4-10 – The observation regions of *Euclid* as from an L2 point perspective. The four regions in blue are the regions of interest for the *Euclid* wide survey. The red and yellow regions correspond to the galactic and ecliptic planes, respectively, and will be avoided. The galactic plane contains objects from our galaxy and is thus uninteresting to probe far-away zones of the Universe. The zodiacal light from our Sun contaminates the ecliptic plane. Credit: Image from Euclid Collaboration et al. [Euc+22a].

- (i) *The size of the survey:* *Euclid* will observe $15,000 \text{ deg}^2$ of extragalactic sky, a large fraction of the celestial sphere. Figure 1.4-10 shows the part of the sky being imaged by the *Euclid* wide survey. As a reference for comparison, the current largest contiguous space-based WL catalogue corresponds to [Cosmological Evolution Survey \(COSMOS\)](#), which was observed with *HST* and spans an area of 1.64 deg^2 .
- (ii) *The limiting magnitude of the survey:* *Euclid* will be extremely sensitive reaching a lower bound of an AB magnitude of 24.5 (10σ) in the visible band. The VIS wideband allows observing very faint galaxies aiming to get a galaxy number density of $\sim 30 \text{ gal/arcmin}^2$.
- (iii) *The size and shape of the instrument PSF:* The high quality optical system provides a small PSF with a FWHM smaller than 0.18 arcsec . The satellite's location in the L2 Lagrange point gives a stable PSF.
- (iv) *How well this PSF is known:* The knowledge of the PSF principally depends on the performance of the PSF model for *Euclid*. This factor will be discussed later and constitutes the driving force of this thesis.
- (v) *How well we can correct for the sources of systematics:* This factor depends on the performance of the shape measurement pipeline, which includes the handling of the PSF model.

Some particularities of the *Euclid* mission make the PSF modelling particularly challenging. The factor (i), having a large FOV and therefore covering a large part of the sky, motivates using a larger pixel scale, which makes the observations, including

the stars, undersampled. As a consequence, the PSF model will need to super-resolve the PSF. Nevertheless, the observing strategy accounts for four dithers, consecutive observations of the same region with small pointing variations. In practice, the dithers increase the number of observed stars available to constrain the PSF model. Another difficulty coming from the factor (i) resides in the vast amount of data to process, which imposes constraints on the computing time for the PSF model. The factor (ii) motivates using a single wideband in the VIS instrument to collect as much light as possible. As a result, the PSF model needs to properly model the spectral variations of the PSF in the wideband. The design of the *Euclid* telescope includes mostly reflective elements, e.g. mirrors, from which the spectral variations follow the diffraction theory and are known. However, the optical system uses a dichroic plate which allows using both instruments simultaneously as their bandpasses are disjoint. It was recently discovered that the dichroic introduces spurious spectral variations, e.g. Venancio et al. [Ven+20, Fig. 28], introducing a great challenge to the PSF model. The fact that *Euclid* is a space mission and the factor (iii) makes the PSF to be almost diffraction-limited, making the shape of the PSF more complex. Finally, the factors (iv) and (v) impose extremely tight requirements on the performance of the PSF model.

The power of WL analyses requires tight control of the systematic errors, making the PSF model an essential component of the *Euclid* mission. The existing PSF models are not adapted for the challenges and particularities of the *Euclid* PSF. This fact motivates the development of new PSF models, a parametric approach led by Prof Lance Miller, and the novel data-driven method described later in this thesis.

* * *
* *
*

Introduction to the PSF

Chapter Outline

2.1	Introduction to optics	36
2.1.1	Scalar diffraction theory	36
2.1.2	The Fresnel approximation	38
2.1.3	The Fraunhofer approximation	39
2.1.4	Analysis of optical imaging systems	40
2.2	General observational forward model	45
2.3	PSF field contributors and related degradations	48
2.3.1	Optic-level contributors	49
2.3.2	Detector-level degradations	50
2.3.3	The atmosphere	54
2.4	Current PSF models	57
2.4.1	Parametric PSF models	57
2.4.2	Non-parametric (or data-driven) PSF models	58
2.5	General comments on PSF modelling	65
2.5.1	Desirable properties of PSF models	65
2.5.2	On the structure of the PSF field spatial variations	67
2.6	Validation of PSF models	69
2.6.1	Pixel-based metrics	70
2.6.2	Moment-based metrics	71
2.6.3	PSF error propagation and PSF requirements	73

THIS chapter discusses more in detail the [Point Spread Function \(PSF\)](#) and should help to understand the [PSF](#) modelling choices of the following chapters of this thesis. We start by introducing optical concepts required to analyse optical imaging systems that will be later used for the proposed [PSF](#) models for space-based telescopes. Then, motivated by the optical introduction, we describe the general observational forward model that we consider in this thesis. We proceed by introducing the different physical contributors to the [PSF](#). The chapter continues with an overview of state-of-the-art [PSF](#) modelling techniques that include comments on what are desirable properties of a [PSF](#) model. We end the chapter by describing different techniques for validating [PSF](#) models.




Figure 2.1-1 – Illustration of the Huygens-Fresnel principle and the modification of a wavefront due to an obstacle.

2.1 Introduction to optics

This section introduces important optical concepts for the understanding of this thesis and is based in the formalism of Goodman [Goo05]. For a profound and rigorous description of optical theory we refer the reader to the seminal book of Born et al. [BW64] or the more concise work seen in Gross [Gro05, S. 12]. For more information on practical wave propagation we refer the reader to Schmidt [Sch10].

2.1.1 Scalar diffraction theory

When studying the PSF we are studying how an optical system with a specific instrument contributes to and modifies our observations. To understand how the optical system interacts in the propagation of light, we need to dig into the nature of light, an EM wave. If one would like to make a fundamental analysis, one would need to use Maxwell's equations, solve them with the optical system under study, and obtain the electric and magnetic fields. Solving a set of coupled partial differential equations is an arduous task. Several approximations can be made, given that some conditions are met, to alleviate the mathematical burden of solving Maxwell's equations without introducing much error into the analysis.

Diffraction theory provides a fundamental framework for analysing light propagation through an optical system. It is especially the case when working with EM waves in the optical spectra when the optical image is situated close to the focus region. The *Huygens-Fresnel principle* states that every point of a wavefront may be considered as a secondary disturbance giving rise to spherical wavelets. The wavefront, at any later instant, may be regarded as the envelope of all the disturbances. Fresnel's contribution to the principle is that the secondary wavelets mutually interfere. This principle provides a powerful method of analysis of luminous wave propagation. In Figure 2.1-1, we can observe the propagation of an incident plane wavefront through an obstacle, a single slit. We can appreciate how the secondary wavelets constitute the plane




Figure 2.1-2 – Illustration of the coordinate system for the diffraction equations. Credit: figure from Ngolè Mboula, Fred Maurice [Ngo16].

wavefront before the obstacle. Then, we observe how the wavefront shape is modified due to the obstacle, following the Huygens-Fresnel principle.

The secondary waves mutually interfere, and according to their phases, they can be constructive or destructive. The analysis of the light propagation in a homogeneous medium is simple as the spherical wavelets interfere, and the total wavefront propagates spherically in the medium. However, suppose the wave encounters an obstacle. In that case, the secondary waves in the vicinity of the boundaries of the obstacle will interfere in ways that are not obvious from the incident wavefront.

Gustav Kirchhoff was a pioneer in providing a solid mathematic foundation to the *Huygens-Fresnel principle* using Green's theorem. First, by deriving the *integral theorem by Kirchhoff* which expresses the solution of the homogeneous wave equation at an arbitrary point in terms of the values of the solution and its first derivative at all points on an arbitrary closed surface surrounding the point. Then, he studied the diffraction by a planar screen, a simple setup which allows making assumptions simplifying the integral theorem. The assumptions are commonly known as the *Kirchhoff boundary conditions*. Although accurate in practice, these boundary conditions were proved inconsistent and were eliminated by Sommerfeld, who gave origin to *Rayleigh-Sommerfeld diffraction theory* that we will continue below.

Let us consider a diffractive aperture that lies in a plane (ξ, η) illuminated in the positive z direction. We will study the diffracted wave in a parallel plane (x, y) at a normal distance z from the first plane. The z axis is orthogonal to both planes and intersects both of them at their origins. Figure 2.1-2 illustrates the coordinate system described above. The diffracted wave writes

$$U(\mathbf{p}_1) = \frac{z}{j\lambda} \iint_{\Sigma} U(\mathbf{p}_0) \frac{\exp(jk r_{01})}{r_{01}^2} d\xi d\eta, \quad (2.1)$$

where λ is the wavelength, $k = 2\pi/\lambda$, $\mathbf{p}_0 = (\xi, \eta, 0)$, $\mathbf{p}_1 = (x, y, z)$, $r_{01} = \|\mathbf{p}_1 - \mathbf{p}_0\|_2$, Σ is the aperture in the (ξ, η) plane, and U is the electric field. The incident wave is $U(\mathbf{p}_0)$ and the diffracted wave is $U(\mathbf{p}_1)$.

There are two approximations in the derivation of Equation 2.1. The first approximation is that we are considering a *scalar theory of diffraction*, a scalar electric

and magnetic field, and not the fields in their full vectorial form. The scalar theory provides a full description of the EM fields in a dielectric medium that is linear, isotropic, homogeneous, and non-dispersive. However, if the medium does not verify these properties, the scalar theory is no longer exact. Even if the medium verifies these properties, if some boundary conditions are imposed on a wave, for example, an aperture, some coupling is introduced between the EM field components. In this scenario, using the scalar theory introduces a bias. Nevertheless, the EM fields are modified only at the edges of the aperture, and the effects extend over only a few wavelengths into the aperture. Therefore, if the aperture is large compared to the wavelength, the error introduced by the scalar theory is small.

Although the current formulation is powerful in representing the diffraction phenomena, it is still challenging to work with the integral from Equation 2.1. As a consequence, we will explore further approximations to the *Rayleigh-Sommerfeld diffraction theory* that will give origin to the *Fresnel diffraction* and *Fraunhofer diffraction*.

2.1.2 The Fresnel approximation

The Fresnel approximation is based on the binomial expansion of the square root in the expression $\sqrt{1+b}$ for some b^1 . The distance r_{01} can be expressed as

$$r_{01} = z\sqrt{1 + \left(\frac{x-\xi}{z}\right)^2 + \left(\frac{\eta-y}{z}\right)^2}, \quad (2.2)$$

that can be approximated, using the first two terms of the binomial expansion, as

$$r_{01} \approx z\left(1 + \frac{1}{2}\left(\frac{x-\xi}{z}\right)^2 + \frac{1}{2}\left(\frac{\eta-y}{z}\right)^2\right). \quad (2.3)$$

The r_{01} appearing in the exponential of Equation 2.1 has much more influence in the result than the r_{01}^2 in the divisor. Therefore, we use Equation 2.3 to approximate the r_{01} in the exponential, and for the divisor we approximate $r_{01}^2 \approx z^2$. Then, we can express the diffracted field as

$$U(\mathbf{p}_1) = \frac{e^{jkz}}{j\lambda z} \iint_{\Sigma} U(\mathbf{p}_0) \exp\left(j\frac{k}{2z}[(x-\xi)^2 + (\eta-y)^2]\right) d\xi d\eta, \quad (2.4)$$

and if we expand the terms in the exponential, we get

$$U(\mathbf{p}_1) = \frac{e^{jkz}}{j\lambda z} e^{j\frac{k}{2z}(x^2+y^2)} \iint_{\Sigma} \left\{U(\mathbf{p}_0) e^{j\frac{k}{2z}(\xi^2+\eta^2)}\right\} e^{-j\frac{2\pi}{\lambda z}(x\xi+\eta y)} d\xi d\eta. \quad (2.5)$$

The FT expression can be recognised in Equation 2.5 with some multiplicative factors. The diffracted wave is the FT of the product of the incident wave and a quadratic phase exponential. In this case, we have approximated the spherical secondary waves of the Huygens-Fresnel principle by parabolic wavefronts. The region where the approximation is valid is known as the *region of Fresnel diffraction*. In this region, the major contributions to the integral come from points (ξ, η) for which $\xi \approx x$ and $\eta \approx y$, meaning that the higher-order terms in the expansion that we are not considering are unimportant. The *region of Fresnel diffraction* can be seen as the (x, y, z) that verify

$$z^3 \gg \frac{\pi}{4\lambda} \left((x-\xi)^2 + (y-\eta)^2\right)^2, \quad \forall(\xi, \eta) \in \Sigma. \quad (2.6)$$

¹The binomial expansion is given by $\sqrt{1+b} = 1 + \frac{1}{2}b - \frac{1}{8}b^2 + \dots$.

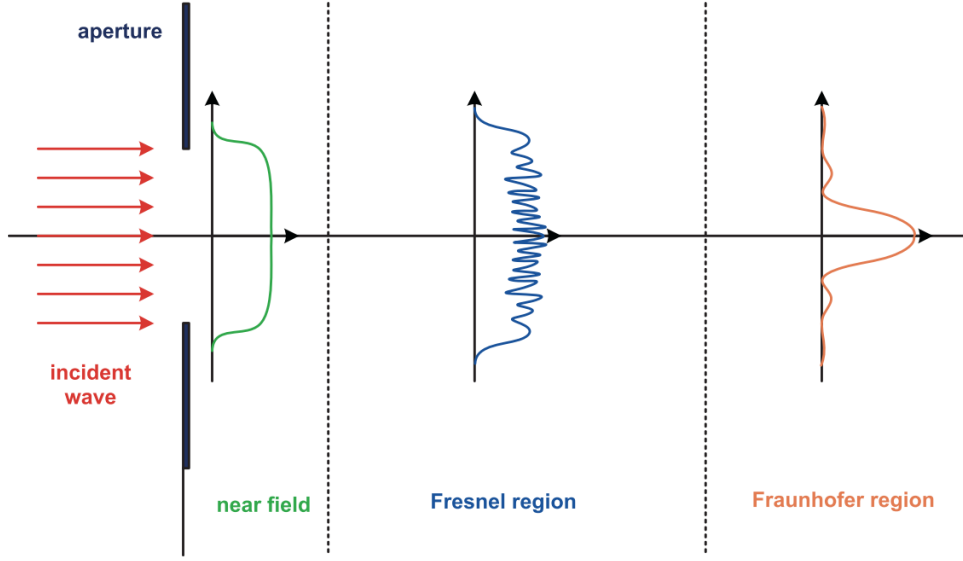


Figure 2.1-3 – Illustration of the different diffraction regions behind an aperture. Credit: Figure from Gross [Gro05].

2.1.3 The Fraunhofer approximation

We continue to present a further approximation that, if valid, can greatly simplify the calculations. The Fraunhofer approximation is based on the assumption that the exponential term with a quadratic dependence of (ξ, η) is approximately unity over the aperture. The region where the approximation is valid is the *far field* or *region of Fraunhofer diffraction*. The required condition to be in this region reads

$$z \gg \frac{k(\xi^2 + \eta^2)}{2}, \quad \forall (\xi, \eta) \in \Sigma. \quad (2.7)$$

The Fraunhofer diffraction formula is given by

$$U(\mathbf{p}_1) = \frac{e^{jkz}}{j\lambda z} e^{j\frac{k}{2z}(x^2+y^2)} \iint_{\Sigma} \{U(\mathbf{p}_0)\} e^{-j\frac{2\pi}{\lambda z}(x\xi+\eta y)} d\xi d\eta, \quad (2.8)$$

where we can reformulate the previous equation using the FT as follows

$$U(\mathbf{p}_1) = \frac{e^{jkz}}{j\lambda z} e^{j\frac{k}{2z}(x^2+y^2)} \text{FT} \{i_{\Sigma}(\mathbf{p}_0)U(\mathbf{p}_0)\}, \quad (2.9)$$

where FT is the Fourier transform, and i_{Σ} is an indicator function over the aperture. Cameras are sensitive to the light's intensity reaching their detectors. The instantaneous intensity of a EM wave is equal to its squared absolute value. Therefore, we can write the intensity of the diffracted wave as

$$I(\mathbf{p}_1) = |U(\mathbf{p}_1)|^2 = \frac{1}{\lambda^2 z^2} |\text{FT} \{i_{\Sigma}(\mathbf{p}_0)U(\mathbf{p}_0)\}|^2, \quad (2.10)$$

which is significantly simpler than the original Rayleigh-Sommerfeld expression from Equation 2.1.

Figure 2.1-3 illustrates the different types of diffraction as a function of the distance to the aperture.




Figure 2.1-4 – Illustration of the generalised *black box* concept of the imaging system and its corresponding coordinate systems. The image plane coordinates are (u, v) , the output aperture plane coordinates are (x, y) , and the object plane coordinates are (ξ, η) . Credit: Figure from Goodman [Goo05].

2.1.4 Analysis of optical imaging systems

Let us analyse a general optical imaging system composed of one or many lenses of possibly different characteristics. We treat the optical system as a *black box* which is characterized by the transformations done to an incident *object* scalar wave, U_o , into an output *image* wave, U_i . The subscript o will refer to the object wave, not the output wave. Similarly, the subscript i will refer to the image wave and not the input wave. Figure 2.1-4 illustrates the *black box* interpretation of the general optical system. We assume that the effect of the optical system between the input and output waves is well described by geometrical optics, which is an affine transformation. All the diffraction effects can be associated with one of the two pupils, input or output (see [Goo05, S. 6.1.2] for more discussion). We choose the latter one and consider the diffraction of the output wave between the output pupil and the image plane.

Let us write the output wave as a function of the input wave using the superposition integral as follows

$$U_i(u, v) = \iint_{-\infty}^{+\infty} h(u, v; \xi, \eta) U_o(\xi, \eta) d\xi d\eta, \quad (2.11)$$

where h is the field's value at image coordinates (u, v) due to a unitary point-source object at position (ξ, η) .

On the one hand, it is cumbersome to work with non-stationary systems as h in Equation 2.11. On the other hand, the study of h is the main topic of this thesis which is strongly non-stationary in systems with a large FOV. However, despite being central to this work, we can consider h to be stationary in the *isoplanatic region*. This region consists of close-by points in the FOV, where the light has travelled similar paths giving small deviations of h . We are assuming a certain regularity in h due to the physical system under study that allows the deviations to be small. In other words, we consider h to be locally stationary or stationary in patches. Figure 2.1-5 illustrates the idea of an *isoplanatic region*. This assumption limits the size of the objects being imaged, as they should have a certain size range with respect to h support so that all the objects being imaged lie within the aforementioned region. Then, for a set of positions $\{(u, v), (\xi, \eta)\}$, assuming the stationarity of h , we have




Figure 2.1-5 – Illustration of the isoplanatic region. Two rays from the same isoplanatic region travel through almost the same turbulence and suffer almost the same distortions.

that $h(u, v; \xi, \eta) = h(u - \xi, v - \eta)$. A shift in the object space translates as a shift in the image space without a change of h 's morphology.

The effect of the geometrical optics inside the *black box* on the input (object) wave writes

$$U_g(\tilde{\xi}, \tilde{\eta}) = \frac{1}{|M|} U_o\left(\frac{\tilde{\xi}}{|M|}, \frac{\tilde{\eta}}{|M|}\right), \quad \text{and} \quad \tilde{\xi} = M\xi, \quad \tilde{\eta} = M\eta, \quad (2.12)$$

where M is the magnification factor of the optical system. Following the aforementioned approximations, we can rewrite Equation 2.11 as a convolution as follows

$$U_i(u, v) = \iint_{-\infty}^{+\infty} h(u - \tilde{\xi}, v - \tilde{\eta}) U_g(\tilde{\xi}, \tilde{\eta}) d\tilde{\xi} d\tilde{\eta}, \quad (2.13)$$

where the impulse response h , under the Fraunhofer approximation, writes

$$h(u, v) = \frac{A}{\lambda z_i} \iint_{-\infty}^{+\infty} P(x, y) e^{-j \frac{2\pi}{\lambda z_i} (ux + vy)} dx dy, \quad (2.14)$$

where P is the pupil function, a binary mask indicating the obscurations in the exit pupil's aperture, A is a constant amplitude, and z_i is the distance between the exit pupil plane and the image plane. In this general case, *where the system has no aberrations*, we see that the output image is formed by a geometrical-optics transformation followed by a convolution with an impulse response that is the Fraunhofer diffraction of the exit aperture.

Polychromatic illumination: the coherent and the incoherent case

Until now, we studied a system with ideal monochromatic light. It is the turn to shift to polychromatic light as it has meaningful implications in the analysis. Even if we study the system's behaviour to light with a particular wavelength, it is practically never the case as real illumination is never perfectly chromatic, even for lasers. Therefore,

we consider a narrowband polychromatic illumination centred at a given wavelength λ . The narrowband assumption states that the bandwidth occupied is small with respect to the central wavelength. For a polychromatic light, we follow Goodman [Goo05] and consider a time-varying phasor of the field, $U_i(u, v; t)$, where its intensity is given by the time integration of its instantaneous intensity

$$I_i(u, v) = \left\langle |U_i(u, v; t)|^2 \right\rangle_t = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-T/2}^{T/2} |U_i(u, v; t)|^2 dt, \quad (2.15)$$

where we consider that the detector integration time is much greater than the optical bandwidth. We can generalise the field expression from Equation 2.13 considering polychromatic light and that the impulse response h is wavelength independent due to the narrowband assumption. Then, the field writes

$$U_i(u, v; t) = \iint_{-\infty}^{+\infty} h(u - \tilde{\xi}, v - \tilde{\eta}) U_g(\tilde{\xi}, \tilde{\eta}; t - \tau) d\tilde{\xi} d\tilde{\eta}, \quad (2.16)$$

where τ represents the delay of the wave propagation from $(\tilde{\xi}, \tilde{\eta})$ to (u, v) . Continuing with the polychromatic analysis, we rewrite the intensity from Equation 2.15 as

$$I_i(u, v) = \iint_{-\infty}^{+\infty} d\tilde{\xi}_1 d\tilde{\eta}_1 \iint_{-\infty}^{+\infty} d\tilde{\xi}_2 d\tilde{\eta}_2 h(u - \tilde{\xi}_1, v - \tilde{\eta}_1) h^*(u - \tilde{\xi}_2, v - \tilde{\eta}_2) J_g(\tilde{\xi}_1, \tilde{\eta}_1; \tilde{\xi}_2, \tilde{\eta}_2), \quad (2.17)$$

where h^* is the conjugate of h , J_g is known as the *mutual intensity* which describes the spatial coherence of U_g at two points and writes

$$J_g(\tilde{\xi}_1, \tilde{\eta}_1; \tilde{\xi}_2, \tilde{\eta}_2) = \langle U_g(\tilde{\xi}_1, \tilde{\eta}_1; t) U_g^*(\tilde{\xi}_2, \tilde{\eta}_2; t) \rangle. \quad (2.18)$$

We can distinguish two types of illuminations, *coherent* and *incoherent*. *Coherent* illumination refers to waves where their phases vary in a perfectly correlated way. This illumination is approximately the case of a laser. In *incoherent* illumination, the wave's phases vary in an uncorrelated fashion. Most natural light sources can be considered incoherent sources. The *mutual intensity* is useful to represent both types of illumination. In the case of coherent light, we obtain,

$$J_g^{\text{co}}(\tilde{\xi}_1, \tilde{\eta}_1; \tilde{\xi}_2, \tilde{\eta}_2) = U_g(\tilde{\xi}_1, \tilde{\eta}_1) U_g^*(\tilde{\xi}_2, \tilde{\eta}_2), \quad (2.19)$$

where $U_g(\tilde{\xi}_1, \tilde{\eta}_1)$ and $U_g(\tilde{\xi}_2, \tilde{\eta}_2)$ are time-independent phasor amplitudes relative to their time-varying counterpart. As both time-varying phasors are synchronized, we have taken a reference phasor and normalized them against their amplitude with respect to a reference point that can be the origin $(0, 0)$. For example,

$$U_g(\tilde{\xi}_1, \tilde{\eta}_1; t) = U_g(\tilde{\xi}_1, \tilde{\eta}_1) \frac{U_g(0, 0; t)}{\left\langle |U_g(0, 0; t)|^2 \right\rangle^{\frac{1}{2}}}. \quad (2.20)$$

Substituting Equation 2.19 into Equation 2.17 we obtain

$$I_i^{\text{co}}(u, v) = |U_i^{\text{co}}(u, v)|^2 = \left| \iint_{-\infty}^{+\infty} h(u - \tilde{\xi}, v - \tilde{\eta}) U_g(\tilde{\xi}, \tilde{\eta}) d\tilde{\xi} d\tilde{\eta} \right|^2, \quad (2.21)$$




Figure 2.1-6 – Illustration of the **Wavefront Errors (WFE)** in a one-dimensional setting where the optical system is represented as a single lens.

where we observe the *coherent illumination gives a system that is linear in the complex amplitude of the field U_g* .

If we now consider incoherent illumination, the mutual intensity writes

$$J_g^{\text{in}}(\tilde{\xi}_1, \tilde{\eta}_1; \tilde{\xi}_2, \tilde{\eta}_2) = \kappa I_g(\tilde{\xi}_1, \tilde{\eta}_1) \delta(\tilde{\xi}_1 - \tilde{\xi}_2, \tilde{\eta}_1 - \tilde{\eta}_2), \quad (2.22)$$

where κ is a real constant, δ is Dirac delta distribution, and I_g is the intensity of the U_g field. Replacing Equation 2.22 in Equation 2.17 the output (image) intensity writes

$$I_i^{\text{in}}(u, v) = |U_i^{\text{in}}(u, v)|^2 = \kappa \iint_{-\infty}^{+\infty} |h(u - \tilde{\xi}, v - \tilde{\eta})|^2 I_g(\tilde{\xi}, \tilde{\eta}) d\tilde{\xi} d\tilde{\eta}, \quad (2.23)$$

where $|h|^2$ is the *intensity impulse response*, also known as the **Point Spread Function (PSF)**. In this case, *an optical system illuminated with incoherent light is linear in intensity*. Equation 2.23 shows a commonly used fact, the output intensity is the convolution of the intensity PSF with ideal image intensity I_g .

Introducing optical aberrations

We considered in the previous development an ideal optical system without any aberrations. Geometrical optics and posterior diffraction effects well describe the system. These latter effects consider that a perfectly spherical wavefront goes through the exit lens and studies its propagation until the image plane. This type of system is known to be *diffraction-limited*. We are now interested in studying a system with aberrations where the consequence is that the output waves are not perfectly spherical.

To study this type of optical system, we start the previous diffraction-limited system and add an ideal phase-shifting plate representing the system's aberrations.




Figure 2.1-7 – Example of *Euclid*'s pupil function in the (x, y) plane for a given position in the (ξ, η) .

The plate is located in the aperture just after the exit pupil, and it deforms the phase of the output wave. To characterise the aberrations we will use the *generalised pupil function* that generalises the pupil function P from Equation 2.14 and writes

$$\mathcal{P}(x, y) = P(x, y)e^{j\frac{2\pi}{\lambda}W(x, y)}, \quad (2.24)$$

where λ is the central wavelength of the incident wave, P is the previous pupil function, a binary mask, and W represents the *Optical Path Differences (OPD)* between a perfect spherical wavefront and the real aberrated wavefront. We will also refer to the OPD as *Wavefront Errors (WFE)*. Figure 2.1-6 illustrates the concept of WFE. It is common to represent the WFE using a Zernike polynomial decomposition [Nol76] as they are orthogonal in the unit disk. A more rigorous definition of the Zernike polynomial will be introduced in section 5.4.2.

Final remarks

The aberrations represented by W , as well as the pupil function P , depend on the object's position in the FOV even though we have not introduced the dependence in the formulations. To illustrate this fact, we can consider the obscurations in the aperture as seen in Figure 2.1-7. One can notice a circular aperture with several obscurations in it, a small circular mirror and three spider arms supporting the mirror. The projection of these obscurations into the (x, y) plane depends on the object's position in the FOV making the pupil function P change with the (ξ, η) coordinates.

We have considered aberrations that only depend on the object's position in the FOV. However, depending on the optical system under study, there might be wavelength-dependent aberrations. For example, some components, such as dichroic filters², can introduce spurious wavelength dependences to the optical system's response. If this is the case, we can add a wavelength dependence to the WFE function W to account for these effects.

There are some assumptions that we will use throughout this thesis, and it is worth stating them as follows:

- The calculations we will do are done for the system following the narrowband assumption over specific central wavelengths. We assume we can work with each wavelength independently in polychromatic problems with a wide passband.

²A dichroic filter is a filter used to pass light from specific wavelengths and reflect other wavelengths selectively. The *Euclid* optical system includes a dichroic filter.

- The wave is assumed to be scalar, and we do not consider the polarisation of the field.
- We assume that the incoming light from natural sources is ideally incoherent, so the system is linear in intensity, as we have seen in Equation 2.23.
- The impulse response is considered to be spatially invariant in its isoplanatic region. In other words, we assume that the PSF is the same for a small patch of the sky.

To conclude, the usual formulation of the PSF, i.e. the intensity of the impulse response, convolving an image seen in many articles comes from the previous assumptions using the results from Equation 2.23, Equation 2.14 and Equation 2.24. We rewrite this formula as follows:

$$I_i(u, v) = (\mathcal{H} \star I_g)(u, v), \quad (2.25)$$

where we remind that the (u, v) is the image plane, we have dropped the κ term from Equation 2.23, and \mathcal{H} is the PSF that writes

$$\mathcal{H}(u, v) = \frac{A^2}{\lambda^2 z_i^2} \left| \iint_{-\infty}^{+\infty} P(x, y) e^{j\frac{2\pi}{\lambda} W(x, y)} e^{-j\frac{2\pi}{\lambda z_i} (ux + vy)} dx dy \right|^2, \quad (2.26)$$

where we are studying the PSF for a specific wavelength and FOV position.

2.2 General observational forward model

Let us focus on the main topic of this thesis, the PSF. We consider the PSF as the intensity impulse response of the imaging system under study to a point source. The concept of PSF is used throughout many imaging applications, including astronomical imaging [Lia+22; Sch19], medical imaging [DK01; JBL18], or 3D microscopy [Sou+12]. The central idea behind a PSF is that it represents transformations done to the imaged object by the imaging system. The PSF is, in a certain way, a characterisation of the imaging system. Focusing on astronomical imaging, the definition of the imaging system can vary between the different use cases and telescopes. For example, in a ground-based telescope, we will consider that the atmosphere belongs to the imaging system we are modelling. However, and naturally, in a space-based telescope, the atmosphere will not be considered. In this thesis, our focus will be on optical systems. These imaging systems work with electromagnetic radiation with a wavelength close to the visible spectrum. For example, the VIS instrument's theoretical wavelength range is from 550nm to 900nm.

The PSF describes the effects of the imaging system in the imaging process of the object of interest. As we have seen in subsection 2.1.4, the effect of the PSF to an image is a convolution, and, by definition, the PSF is the kernel of the convolution. However, this convolutional kernel varies spatially, spectrally, and temporally. We give a non-exhaustive list that motivates each of these variations:

- *Spatial variations*: The optical system presents a certain *optical axis*, which is an imaginary line where the system has some degree of rotational symmetry. In simpler words, it can be considered as the direction of the light ray that produces a PSF in the centre of the focal plane for an unaberrated optical system. The *angle of incidence* is defined as the angle between an incoming light ray and the optical axis. The main objective of the optical systems we study is to make

the incoming light rays converge in the focal plane, where there will be some measurement instrument as a camera, for example. Depending on the angle of incidence, the image will form in different positions in the focal plane. The path of the incoming light will be different for each angle of incidence, and therefore the response of the system will be different too. In other words, the PSF will change depending on the angle of incidence or spatial position in the focal plane where the image is forming. Optical systems with very wide focal planes, generally associated with very wide FOV, present more important PSF spatial variations.

- *Spectral variations:* Principally due to the diffraction phenomena and its well-known wavelength dependence covered in section 2.1. Refractive³ components of the optical system under study can add more sources of spectral variations.
- *Temporal variations:* The state of the telescope changes with respect to time; therefore, the imaged object's transformation also changes. In space-based telescopes, high-temperature gradients cause mechanical dilations and contractions that affect the optical system. In ground-based telescopes, the atmosphere composition changes with time. Consequently, it temporally affects the response of the optical system, or the PSF.

The PSF changes with space, time and wavelength. Let us refer to the PSF field, $\mathcal{H}(x, y; \lambda; t)$, as all the PSF that represent an optical system. We can define the PSF field as a varying convolutional kernel $\mathcal{H} : \mathbb{R}^2 \times \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathcal{K}$, going to some function space (for instance $\mathcal{K} = L_1(\mathbb{R}^2)$). This would be an accurate description of how the PSF affects the images considering the assumptions from section 2.1 are valid. It is, however, cumbersome to work with a convolutional kernel that varies at each point in space. We will adopt throughout this thesis an approximation that considers the PSF locally invariant in its *isoplanatic region* (see Figure 2.1-5 for an illustration). This approximation means that in a vicinity of an observed object, we will consider that the PSF only varies with time and wavelength, thus facilitating the computation of the convolution. In practice, the close vicinity, or the isoplanatic region, will be defined as the postage stamp used to image the object of interest. The typical galaxies observed for weak lensing have a comparable size with respect to the PSF size (e.g. see [Man+18, Figure 7] for a distribution of relative galaxy to PSF size in the HSC survey). Consequently, the approximation error is kept low as it is only done for small patches of the focal plane.

Let us define our object of interest, $\mathcal{I}_{\text{GT}}(x, y; \lambda)$, that is the I_g object from section 2.1, as a continuous light distribution $\mathcal{I}_{\text{GT}} : \mathbb{R}^2 \times \mathbb{R}_+ \rightarrow \mathcal{K}$. In this thesis, we are not considering transient objects, i.e. the time dependence scale of the object is comparable with the exposure time used to image it. Therefore, we can ignore the temporal dependency of the Ground Truth (GT) object, $\mathcal{I}_{\text{GT}}(x, y; \lambda) \neq f(t)$. Let us write our general observational forward model that relates our GT object of interest, our PSF and our observed image as follows

$$I_{\text{img}}(x_i, y_i; t) = \mathcal{F}_p \left\{ \int_0^{+\infty} \mathcal{T}(\lambda) (\mathcal{I}_{\text{GT}} \star \mathcal{H})(x, y; \lambda; t) d\lambda \right\} \circ N(x, y; t), \quad (2.27)$$

where \mathcal{F}_p is a degradation operator going from \mathcal{K} to $\mathbb{R}^{p \times p}$ which includes the image sampling in the instrument. Then, $I_{\text{img}}(x_i, y_i; t) \in \mathbb{R}^{p \times p}$ are the instrument's pixel

³Refraction refers to the change of direction in the propagation of a wave passing from one medium to another. Most of the wave energy is transmitted to the new medium. Reflection refers to the abrupt change of direction of the wave propagation due to a boundary between mediums. In this last case, most of the oncoming wave energy remains in the same medium.

measurements of the observed image. The variables (x_i, y_i) correspond to the centre location of the target object i and should not be mistaken with the discrete pixel coordinates of I_{img} . The instrument's transmission is represented by $\mathcal{T} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, a function with finite support, and $N(x, y; t) \in \mathbb{R}^{p \times p}$ corresponds to the noise affecting our observation and possibly a modelling error, where \circ is some composition operator.

Although Equation 2.27 provides a general observational forward model, it can be unpractical. The continuous functions \mathcal{H} , \mathcal{T} , and \mathcal{I}_{GT} are practically inaccessible. We make several assumptions to simplify the problem:

- (a) The continuous functions \mathcal{H} and \mathcal{I}_{GT} are well approximated by piece-wise constant functions over a regular grid in \mathbb{R}^2 . We assume $\mathcal{H} \approx H$ and $\mathcal{I}_{GT} \approx I_{GT}$, where $H, I_{GT} \in \mathbb{R}^{P \times P}$ with $P \geq p$. The resolution of these two variables has to be greater or equal to the observation resolution,
- (b) The noise is additive, i.e. $\circ \equiv +$,
- (c) The degradation operator is approximated by its discrete counterpart, $\mathcal{F}_p \approx F_p$, where $F_p : \mathbb{R}^{P \times P} \rightarrow \mathbb{R}^{p \times p}$. We assume that the degradation operator is *linear*, and that includes downsampling, intra-pixel shifts and linear detector effects,
- (d) We keep the approximation that the PSF is locally constant within the postage stamp of $P \times P$ values of the target image,
- (e) The integral can be well approximated by a discretised version using n_λ bins.

Taking into account the aforementioned assumptions, we can define our practical observational forward model as follows

$$I_{\text{img}}(x_i, y_i; t) = F_p \left\{ \sum_{k=1}^{n_\lambda} T(\lambda_k) (I_{GT} \star H)(x_i, y_i; \lambda_k; t) \Delta\lambda_k \right\} + N(x, y; t), \quad (2.28)$$

where T is a discretized version of \mathcal{T} , and $b^k = [b_0^k, b_1^k]$ is the k -th wavelength bin centred in λ_k , with a width of $\Delta\lambda_k = b_1^k - b_0^k$.

Particular case: a star observation

The case of star observations is of particular interest. Stars can be approximated as a spatial impulse, $\mathcal{I}_{\text{star}}(x, y; \lambda) = \delta(x, y; \lambda) = f_{(x,y)}(\lambda)$, and provide us with degraded observations of the PSF field. These observations will be crucial to constrain the PSF models. Unluckily, we do not have access to the star's spectral variation, $f_{(x,y)}(\lambda)$. However, we dispose of complementary photometric observations that can be useful to characterise the spectral variations. These observations provide us with the star's **Spectral Energy Distribution (SED)** that can be defined as follows:

‘astronomers use the term “spectral-energy distribution” to mean the spectrophotometrically calibrated flux density from a source as a function of wavelength or frequency, usually at low spectral resolution.’ [Hog22, p. 7]

These photometric observations are done in several spectral bands. Figure 2.2-8 shows the bands from the MegaCam instrument at CFHT. The filter curves can be downloaded from the Spanish Virtual Observatory (SVO) webpage⁴. For more information about SEDs and stellar photometry we refer the reader to the pedagogical document of Hogg [Hog22]. In practice, we can consider the SED as a normalised

⁴<http://svo2.cab.inta-csic.es/svo/theory/fps/index.php>




Figure 2.2-8 – The 3rd generation set of filters of the MegaCam instrument at [Canada-France-Hawaii Telescope \(CFHT\)](#) currently being used for [CFIS](#). The transmission filter response includes the full telescope and 1.25 airmasses of atmospheric attenuation. The full telescope includes the mirrors, the optics, and the [CCDs](#).

low-resolution sampling of the star’s spectral variations. We can write the [SED](#) definition we will use as

$$\text{SED}_{b^k}(\lambda_k) = \frac{1}{Z_{n_\lambda(b)}} \int_{b_0^k}^{b_1^k} f_{(x,y)}(\lambda) d\lambda, \quad (2.29)$$

where we continued to use the b^k bin definition from [Equation 2.28](#), and $Z_{n_\lambda(b)}$ is a constant used so that the [SED](#) is normalised to unity. We have that $\sum_{k=1}^{n_\lambda} \text{SED}_{b^k}(\lambda_k) = 1$. We continue by considering that the [GT](#) image in [Equation 2.28](#) is a star, and we use the spectral bins from the [SED](#) definition to discretize the spectral integration. Finally, we write the practical star observation model as

$$I_{star}(x_i, y_i; t) = F_p \left\{ \sum_{k=1}^{n_\lambda} T(\lambda_k) \text{SED}_{b^k}(\lambda_k) H(x_i, y_i; \lambda_k; t) \Delta\lambda_k \right\} + N(x, y; t), \quad (2.30)$$

where we consider the star observation $I_{star}(x_i, y_i; t) \in \mathbb{R}^{P \times P}$ as a degraded version of the [PSF](#) field $\tilde{H}(x_i, y_i; t) \in \mathbb{R}^{P \times P}$.

2.3 PSF field contributors and related degradations

So far, we have described how the [PSF](#) interacts with the images we observe and how we can model an observation. However, we did not give much information about the different contributors to the [PSF](#) field, and about the different degradations represented by F_p in [Equation 2.28](#) that can occur when modelling observations. We provide a non-exhaustive list of contributors to the [PSF](#) field, sources of known degradations, and the atmosphere’s effect in our [PSF](#) modelling problem.

2.3.1 Optic-level contributors

These contributors affect the PSF by altering the wave propagation in the optical system. In other words, they generally affect the wavefront's amplitude and phase.

- *Diffraction phenomena and the aperture size:* As we have seen in [section 2.1](#), the diffraction phenomena happening in the optical system plays an essential role in the formation of the PSF. The size of the optical system aperture and the wavelength of the light being studied are of particular interest. [Equation 2.26](#) shows us that under some approximations the PSF is the FT of the aperture. Therefore, the size of the aperture and the PSF are closely related. For example, if we consider an ideal circular aperture, the diffraction pattern associated with it is the well-known *Airy disk*. The relation between the width of the PSF and the diameter of the aperture is given by

$$\theta_{\text{FWHM}} = 1.025 \frac{\lambda}{d}, \quad (2.31)$$

where θ_{FWHM} is the FWHM expressed in radians, λ is the wavelength of the light being studied, and d is the diameter of the aperture. The width of the PSF is a fundamental property of an optical system as it defines the resolution of the system. In other words, the PSF size defines the optical system's ability to distinguish small details in the image.

- *Optical aberrations:* These aberrations are due to imperfections in the optics, for example, a not ideally spherical mirror or a not perfectly aligning of the optical components. The optical aberrations play a significant role in the morphology of the PSF and can be modelled using the WFE introduced in the generalized pupil function from [Equation 2.24](#). Some aberrations have a distinctive name, for example, coma, astigmatism, and defocus.
- *Surface errors or polishing effects:* One would ideally like perfectly smooth surfaces in mirrors and lenses. However, small imperfections are arising in the optical surfaces due to the imperfect polishing of the surfaces. [Figure 2.3-9](#) shows the measurement of Surface Errors (SFE) in the HST. We will present *Euclid* simulations with and without SFE in [chapter 6](#) and show their effect. Gross et al. [[Gro+06](#), Section 35.2] gives a more in-depth analysis of surface errors focusing on the tolerancing of SFE.
- *Obscurations:* Complex optical systems have telescope designs where some elements can obscure some part of the pupil. This is the case for *Euclid* as it can be seen in [Figure 2.1-7](#). It is also possible to notice the obscurations in the HST in [Figure 2.3-9](#).
- *Scattered light:* Optical elements, detectors, and instruments give rise to scattered light that reach the CCDs. Krist [[Kri95](#)] studied the problem for the HST. Storkey et al. [[Sto+04](#)] developed methods to clean observations with scattered light from the SuperCOSMOS Sky Survey (SSS). Sandin, Christer [[San14](#)] studied the effect of scattered light on the outer parts of the PSF.
- *Ice contamination:* Some water particles may get into the optical system during the manufacturing process. In space, this water turns into ice and produces contaminations that are very hard to model and affect the telescope's PSF. The ice contamination produces a time-variable transmission degradation that also varies with wavelength and FOV position. A notable example is the ESA

Gaia mission that suffered from ice contamination (see Gaia Collaboration et al. [Gai+16, Section 4.2.1]) and required several decontamination procedures to remove the ice from the optical system slowly. Although difficult to predict, the ice contamination could also be an issue for *Euclid*.

- *Chromatic optical components:* These components have a particular wavelength dependence, excluding the natural chromaticity due to diffraction. A particular example is a dichroic filter which ideally serves as a mirror for some wavelength and is transparent to other wavelengths. The *Euclid* optical system includes a dichroic filter which allows using both instruments, VIS and *Euclid's* Near Infrared Spectrometer and Photometer instrument (NISIP), simultaneously as their passbands are disjoint. These chromatic optical components do not have an ideal bandpass transmission curve, and therefore they introduce spurious wavelength dependences to the PSF. There is an ongoing effort in the *Euclid* Consortium (EC) to characterise the chromatic variations introduced by *Euclid's* dichroic filter. See Baron et al. [Bar+22] for recent results on *Euclid's* dichroic.
- *Light polarisation:* In the diffraction theory analysis carried out in section 2.1 we have neglected the light polarisation, and considered the EM fields as scalar fields. There are some regions where the incoming light has been polarised by Galactic foreground dust. In addition, complex chromatic optical components can introduce light polarisation. There is an ongoing effort within the EC to study the impact of light polarisation in the PSF and the biases it introduces in weak lensing analysis. Lin et al. [Lin+20] studied the impact of light polarisation on weak lensing systematics for Nancy Grace Roman Space Telescope (Roman) and found out that the systematics introduced by light polarisation are comparable to Roman's requirements.
- *Thermal variations:* The thermal variations in a telescope introduce mechanical variations in its structure that affect the performance of the optical system. The origin of the thermal variations is strong temperature gradients due to the sun's illumination. It is sometimes referred to as the *telescope's breathing* for its repetitive pattern due to the orbits. For example, thermal variations can introduce a small defocusing of the system that will change the PSF morphology. Nino et al. [Nin+07] performed a study of the thermal variations in the HST. For *Euclid*, a Structural-Thermal-Optical Performance (STOP) analysis was carried out to predict the impact of the thermal variations on the optical system. This effect is naturally more significant in space-based telescopes as the temperature gradients in space are considerably more prominent than the ones found on the ground.

2.3.2 Detector-level degradations

These degradations are related to the detectors being used and, therefore, to the intensity of the PSF. They affect the observed images through the degradation operator F_p from Equation 2.28, and as we will use star images to constrain PSF models, it is necessary to consider their effects. Some of these degradations are non-convolutional and will therefore not be well modelled by a convolutional kernel. Nevertheless, we expect the VIS images' preprocessing to correct these effects. However, the correction will not be perfect, and some modelling errors can propagate to the observations.




Figure 2.3-9 – Surface errors measured for the HST which are in the range of $\pm 30\text{nm}$. Credit: Figure from Krist et al. [KHS11].




Figure 2.3-10 – Example of two different pixelations on the same high-resolution image representing an Airy PSF. The difference between the two pixelations is an intra-pixel shift of $(\Delta x, \Delta y) = (0.35, 0.15)$ between them.

- Undersampling and pixellation:* The EM wave that arrives at the detectors is a continuous function. The discrete pixels in the detectors integrate the functions and measure the intensity of the wave in their respective area. We name this process *pixellation*, although it is also known as *sampling*. High et al. [Hig+07] performed an early study on the effects of pixellation in WL and on the choice of pixel scale for a WL space-based mission. Krist et al. [KHS11, Section 3] gives some insight on pixellation effects for the HST. There are two aspects of pixellation that play a crucial role in PSF modelling. First, the sampling is done with the same grid, but it is important to consider that the continuous function is not necessarily centred on the grid. This difference means that there can and will exist intra-pixel shifts between the different pixellations. Figure 2.3-10 shows how two pixel representations of the same light profile change due to two different pixellations. When optimising a PSF model to reproduce some observed stars, the centroids of both images must be the same. Suppose the image centroids are the same, and the underlying model is a good representation of the observation. In that case, the residual image between the two pixellated images will be close to zero. If the centroids are not the same, the residual can be far from zero even though the model is a good representation of the observation, as illustrated in the residual image in Figure 2.3-10. The second aspect is related to the Nyquist-Shannon sampling theorem. The theorem states the required number of samples that we need to use to determine a signal of a given bandwidth perfectly. In the telescopes we will study, the bandwidth and number of samples are related to the diameter of the aperture and the pixel size. Depending on the telescope’s design, the sampling may not verify the Nyquist-Shannon theorem. If the images are undersampled, meaning that the theorem is not verified, a super-resolution step is required in the PSF modelling, which, as we will later see, is the case of *Euclid*.
- Optical throughput and CCD Quantum Efficiency:* The optical throughput of the system is the combined effect of the different elements composing the optical system, such as mirrors and optical elements like coatings [Ven+16]. The filter being used in the telescope forms part of the optical throughput, as it can be seen in Figure 2.2-8 for the MegaCam set of filters. Figure 2.2-8 also includes the CCD Quantum Efficiency (QE) which describes the sensibility of the CCD to detect photons of different wavelengths. Commonly, CCDs do not have a uniform response to the different wavelengths. Therefore, we need to multiply the CCD QE with the telescope’s optical throughput to compute the total transmission. For example, figure 7 from the MegaCam page⁵ gives more detailed information for the CFHT. Figure 2.3-11 shows the CCD QE for *Euclid*’s VIS instrument.
- CCD misalignments:* Ideally, we expect that all the CCDs in the detector lie in a single plane that happens to be the focal plane of the optical system. However, it is not the case in practice, as there might be small misalignments between the CCDs introducing small defocuses that change from CCD to CCD. See Jee et al. [JT11, Figure 8] for a study of this effect for the Vera C. Rubin Observatory (Vera C. Rubin Observatory).
- Guiding errors:* Even if space telescopes are expected to be very stable when doing observations thanks to the Attitude and Orbit Control System (AOCS), there will exist a small residual motion that is called pointing jitter. The effect on the observation is introducing a small blur that can be modelled by a specific

⁵<https://www.cfht.hawaii.edu/Instruments/Imaging/Megacam/specsinformation.html>

convolutional kernel that depends on the pointing time series. Fenech Conti [Fen17, Section 4.8.3] proposes to model the effect for *Euclid* with a Gaussian kernel.

- *Charge Transfer Inefficiency:* CCD detectors are in charge of converting incoming photons to electrons and collecting them in a potential well in the pixel during an exposure. When the exposure finishes, the charge on each pixel is read. The collected electrons are transferred through a chain of pixels to the edge of the CCD where they are amplified and then read. High energy radiation above the Earth's atmosphere gradually damages CCD detectors. The silicon damage in the detectors creates traps for the electrons that are delayed during the reading procedure. This effect is known as CTI and produces trailing of bright objects and blurring of the image. Given the harsh space environment, this effect is noticeably important for space telescopes. CTI effects are expected to be corrected in the VIS image preprocessing. Rhodes et al. [Rho+10] carried out a study on the impact of CTI for WL studies. Massey et al. [Mas+09] developed a model to correct for CTI for the HST and later improved it in Massey et al. [Mas+14].
- *Brighter-fatter Effect:* The assumption that each pixel photon count is independent of its neighbours does not hold in practice. There is a photoelectron redistribution in the pixels as a function of the number of photoelectrons in each pixel. The BFE is due to the accumulation of charge in the pixels' potential wells and the build-up of a transverse electric field. The effect is stronger for bright sources. Antilogus et al. [Ant+14] studied the effect and observed that the images from the CCDs do not scale linearly with flux, so bright star sizes appear larger than fainter stars. Guyonnet et al. [Guy+15] and Coulton et al. [Cou+18] proposed methods to model and correct this effect. The preprocessing of VIS images is supposed to correct for the BFE, but there might be some residuals.
- *Wavelength Dependent Sub-Pixel response:* There exist a charge diffusion between neighbouring pixels in the CCD. Niemi et al. [Nie+15] carried out a study of this effect for an *Euclid*'s VIS CCD and modelled the response of the CCD. Niemi et al. [Nie+15] proposed to model the effect as a Gaussian convolutional kernel where the standard deviations of the 2D kernel are wavelength dependent, $\sigma_x(\lambda)$ and $\sigma_y(\lambda)$, and where measured with a reference VIS CCD. Krist [Kri03] studied the charge diffusion in HST and proposed spatially varying blur kernels to model the effect.
- *Noise:* There are several sources of noise in the measurements. *Thermal noise* [Nyq28] refers to the signal measured in the detector due to the random thermal motion of electrons which is usually modelled as Gaussian. *Readout noise* [BTM04] refers to the uncertainty in the photoelectron count due to imperfect electronics in the CCD. *Dark-current shot noise* [Bae06] refers to the random generation of electrons in the CCD, and even though it is related to the temperature, it is not Gaussian.
- *Tree rings and edge distortions:* There exist electric fields in the detector that are transverse to the surface of the CCD. The origin of these fields includes doping gradients or physical stresses on the silicon lattice. This electric field displaces charge, modifying the effective pixel area. Consequently, it changes the expected astrometric and photometric measurements. This electric field also generates concentric rings, *tree rings*, and bright stripes near the boundaries

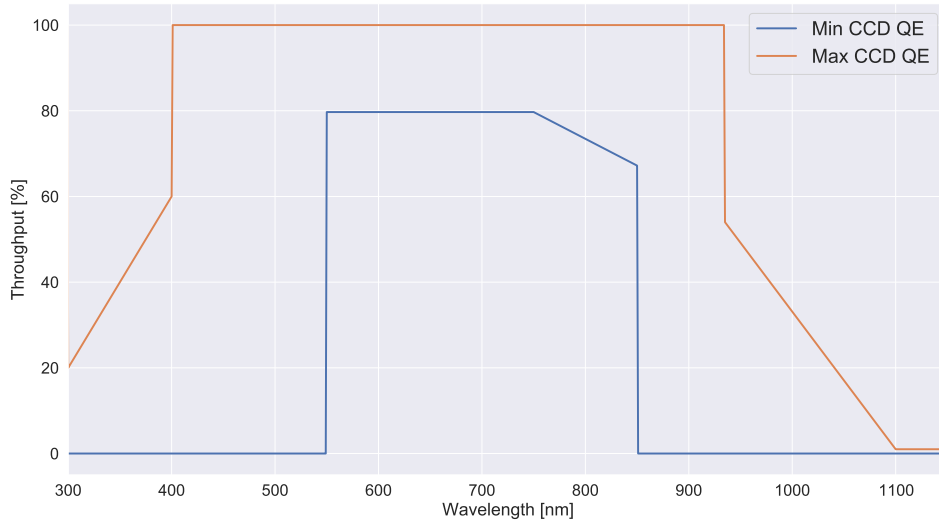


Figure 2.3-11 – *Euclid*’s requirements on the CCD QE for the VIS instrument from the *Euclid*’s Mission Database (MDB).

of the CCD, *edge distortions*. Given the close relationship between this effect and the detector, its importance depends strongly on the instrument being used. This effect was not noticed in the MegaCam used in CFIS. However, it was a major concern in the Dark Energy Camera (DECam) used in the DES. Plazas et al. [PBS14] studied and measured the effect for the DECam. Jarvis et al. [Jar+20, Figure 9] illustrates the consequence of tree rings in the PSF modelling.

- *Other effects*: These effects include *Colour-dependent Photo Response Non-Uniformity (PRNU)* and *Detector Analog-to-Digital non-linearity*, which are both expected to be corrected in the preprocessing of VIS images.

2.3.3 The atmosphere

The atmosphere is central in ground-based telescopes’ PSFs. The way the atmosphere affects our images will strongly depend on the exposure time used to image an object. The PSF induced by the atmosphere for a very short exposure will look like a speckle, while a long exposure will produce a PSF that resembles a 2D Gaussian, or more precisely, a Moffat profile [Mof69]. As a first approximation for long exposures, we can consider that the effect of the atmosphere on the PSF is that of a spatially varying low-pass filter. Therefore, broadening the PSF and limiting the telescope’s resolution. Astronomers usually use the term *seeing* to refer to the atmospheric conditions of the telescope, and it is measured as the FWHM of the PSF. The loss of resolution due to the atmosphere is one of the main motivations for building space-telescopes like *Euclid*, where the PSF is close to the diffraction limit and very stable.

The atmosphere is a heterogeneous medium whose composition changes with the three spatial dimensions and with time. The inhomogeneity of the atmosphere affects the propagation of light waves that arrive at the telescope. Instead of supposing that the light waves arriving at the telescope are plane waves, as emitted by the far away source under study, these incoming waves already have some phase lags or leads with




Figure 2.3-12 – Illustration of six von Kármán phase screen layers at different altitudes simulated for LSST. Credit: Figure from Jee et al. [JT11].

respect to an ideal plane wave. We can say that the atmosphere introduces WFE in the optical system. These effects can be resumed as an effective phase-shifting plate, $\Phi_{\text{eff}}(x, y, t)$. However, the calculation of this effective plate is cumbersome as it involves having a model of the atmosphere and integrating the altitude, (z), so that we have the spatial distribution, (x, y), of the effective WFEs. The model of the atmosphere is represented by the refractive-index structure constant C_n^2 which gives the vertical profile of the atmosphere. However, the C_n^2 is difficult to measure, and even if it is available, it is computationally expensive to exploit.

To simulate the effect of the atmosphere, we can discretise the integral over the altitude into M phase screens at different altitudes. Each phase screen will have specific properties and will be characterised by its power spectrum that can be modelled by a von Kármán model of turbulence [Kár30] that writes

$$\Psi(\nu) = 0.023 r_0^{-5/3} \left(\nu^2 + \frac{1}{L_0^2} \right)^{-11/6}, \quad (2.32)$$

where ν is the two-dimensional frequency, r_0 is the Fried parameter, and L_0 is the outer scale. Both parameters, r_0 and L_0 , are generally expressed in meters. The Fried parameter relates to the turbulence amplitude, and the outer scale relates to the correlation length. See Figure 2.3-12 for an example of atmospheric phase screens. For lengths longer than L_0 , the power of the turbulence asymptotically flattens. If we take the limit of L_0 to infinite, we converge to the Kolmogorov model of turbulence [Kol91].

Once the phase screens, $\Phi_m(x, y)$, have been simulated following Equation 2.32, the temporal variation of the screen has to be taken into account. This temporal variation is usually modelled with the wind's properties at the phase screen's reference altitude. We describe the wind with two components, v_x and v_y , where we have assumed that $v_z = 0$. We then obtain the effective phase screen by a weighted average




Figure 2.3-13 – Example of atmospheric PSFs with different exposure times. The simulation was done using the atmospheric parameters from Jee et al. [JT11] for an LSST-like scenario.

of the phase screens at the different altitudes as

$$\Phi_{\text{eff}}(x, y, t) = \sum_{m=1}^M c_m \Phi_m(x, y, t), \quad (2.33)$$

where $\{c_m\}$ are some weights. The difficulty of modelling the atmosphere is that the time scales are comparable with the exposure time. Therefore, the PSF we might estimate for a given time snapshot will change with respect to another PSF at another snapshot within the same camera exposure. This change means that to model the PSF physically, we need to integrate the instantaneous PSF over time which writes

$$I_{\text{img}}(x_i, y_i) = \int_{t_o}^{t_o + T_{\text{exp}}} I_{\text{img}}(x_i, y_i, t) dt, \quad (2.34)$$

where $I_{\text{img}}(x_i, y_i, t)$ is the instantaneous image of the object affected by the PSF $\mathcal{H}(\dots; t)$, t_o is a random initial time and T_{exp} is the exposure time.

Finally, we need to choose the time step size to discretise the integral from Equation 2.34. Each instantaneous PSF will look like a speckle. Once we add them up in the integral, the PSF starts becoming rounder and smoother. Figure 2.3-13 shows examples of atmospheric PSFs using different exposure times that were simulated using 6 phase screens using the parameters from Jee et al. [JT11] that correspond to a LSST-like scenario. It is interesting to see how the short-exposure PSF looks like a speckle, and then the profile becomes more and more smooth as the exposure time increases. As a reference, the exposure time used for the r-band observations in CFIS is 200s [col16]. Vries et al. [Vri+07] studied the change in the ellipticity of the PSF due to atmospheric turbulence as a function of the exposure time. Vries et al. observed that the ellipticity of the PSF decreases its amplitude as the exposure time increases.

Another effect that should be taken into account is the atmospheric differential chromatic refraction. This effect represents the refraction due to the change of medium from vacuum to the Earth's atmosphere, which changes as a function of zenith angle and wavelength. Meyers et al. [MB15] performed a study on the impact of the atmospheric chromatic effect on weak lensing for surveys like LSST and DES.

Heymans et al. [Hey+12a] performed a study on the impact of atmospheric distortions on weak-lensing measurement with real data from CFHT. Heymans et al. characterised the ellipticity contribution of the atmosphere to the PSF for different exposure times. To achieve this, she computed the two-point correlation function of the residual PSF ellipticity between the observations and a PSF Extractor (PSFEx)-like PSF model (described in detail in section 2.4). [Sal+09] studied the image quality and the observing environment at CFHT. Jee et al. [JT11] carried out a simulation study to evaluate the impact of atmospheric turbulence on weak-lensing measurement in LSST. Jee et al. [JT11] used the atmospheric parameters from [Ello2] that were measured in the LSST site in Cerro Pachón, Chile.

Another interesting way to simulate the atmosphere and the PSF is to use a photon Monte Carlo approach, known as *photon shooting*. This line of work was carried out in Peterson et al. [Pet+15; Pet+19; Pet+20] with a simulator available coined PhoSim⁶ that aims to model LSST. The method consists of sampling photons from astronomical sources and simulating their interactions with their models of the atmosphere, the optics and the detectors. Although PhoSim seems to be an interesting tool for simulations, it is not straightforward how it could be used in a *computationally tractable* way to model the PSF from observations.

To conclude, we have seen that it is possible to develop a physical model of the atmosphere based on the optical understanding we have from section 2.1 and the studies of atmospheric turbulence of Kármán and Kolmogorov. However, this approach has two inconveniences. First, the approach requires physical measurements of the atmosphere at the telescope’s site, which is not always available. Second, it is computationally costly as there is an integration on the altitude, as the atmospheric properties change, and a temporal integration to reach the exposure time used. An alternative would be to use complementary observations from a wavefront sensor to better estimate the effective phase screen, Φ_{eff} . A wavefront sensor provides wavefront information about the incoming light wave. These sensors are usually used in adaptive-optics settings in the real-time control loop of the mirror deformations, which aim to provide a deformation that would counteract the effect of the atmosphere on the incoming wavefront. However, not all telescopes have a wavefront sensor to be exploited.

In practice, it is required to use long exposure times to obtain deeper observations, meaning observing fainter objects that are important for weak-lensing studies. This fact, simplifies the PSF modelling task as the long temporal integration smooths the PSF profile as well as the PSF spatial variations over the FOV. Therefore, a completely data-driven approach for modelling the PSF can offer a feasible and effective solution in this scenario.

2.4 Current PSF models

Let us now discuss some of the most known PSF models. There are two main approaches, or families of methods, for PSF modelling, *parametric* and *non-parametric*, also known as *data-driven*, PSF models.

2.4.1 Parametric PSF models

This approach consists of building a physical model of the optical system that aims to be as close as possible to the actual telescope. Then, a few model parameters are fit to the star observations. This family of methods is capable of handling the chromatic

⁶https://bitbucket.org/phosim/phosim_release/wiki/Home

variations of the PSF as well as complex detector effects. These models have only been developed for space missions. The difficulty of building a physical model for the atmosphere, already discussed in [subsection 2.3.3](#), makes them unpractical for ground-based telescopes. We mention two models for two space telescopes, *HST* and *Euclid* respectively.

Tiny-Tim

The parametric model developed for the *HST* is *Tiny-Tim* [Kri93; KB95; KHS11]. It has been used to model the PSF of the *Advanced Camera for Surveys (ACS)* Wide Field Camera in the *HST* for the *COSMOS*. The *COSMOS* survey is a 2 deg^2 field was used to create the only space-based weak-lensing catalogue. The first WL catalogue used the *Tiny-Tim* model [Lea+07]. Rhodes et al. [Rho+07] studied the stability of *HST*'s PSF noticing a temporal change of focus in the images. Following WL catalogues [Sch+10] used a data-driven PSF model based on *Principal Component Analysis (PCA)*. It was later shown that a simple data-driven model [HA17] outperformed *Tiny-Tim*, exposing limitations of parametric modelling. Some events, such as launch vibrations and ice contaminations, introduce significant variations into the model. These reasons prevent a complete ground characterization from being successful.

VIS physical model

This model constitutes one of the current approaches for PSF modelling for *Euclid*. Professor Lance Miller at the University of Oxford is currently the leader in the development of this PSF model. The model has not yet been published, but there are internal *Euclid* technical documents describing the method. This approach builds a VIS simulator, which includes many physical effects and a module for fitting parameters from star observations. The PSF model has as input different measurements from the VIS instrument as well as from *Euclid*'s optical system. The optimization of the model is a degenerate problem. The model is built in the wavefront space with a few parameters and should be constrained from intensity observations which fall into the phase retrieval category of inverse problems [She+15]. There is an ongoing effort to design a campaign of phase diversity observations that will allow breaking some degeneracies in the fitting of the model. The current plan is to perform computationally-intensive calibrations, with phase diversity observations, in some specific stellar calibration fields once per month. These calibrations are expected to provide better estimates for most of the model's parameters. Then, using the available stars, the PSF model will also fit some parameters on each camera exposure. Fenech Conti [Fen17] provides the only published description of the VIS physical model, although it is an early stage of the PSF model. Even though the phase diversity observations are important for the calibration of the model, they are expensive in terms of observational time, which is a scarce resource in a space mission. These observations are, to some extent, dangerous in terms of focusing and defocusing the spacecraft in space.

2.4.2 Non-parametric (or data-driven) PSF models

The *non-parametric PSF models*, also known as *data-driven*, only rely on the observed stars to build the model in pixel space. It is blind to most of the physics of the inverse problem. The models assume regularity in the spatial variation of the PSF field across the FOV. These methods usually differ in the way they exploit this regularity. Data-driven methods can easily adapt to the current state of the optical system. However, they have difficulties modelling complex PSF shapes occurring in diffraction-limited

settings. A limitation shared by all the data-driven methods is their sensitivity to the available number of stars to constrain the model. When the number of stars falls below some threshold, the model built is usually considered unusable for WL purposes. This family of methods has been widely used for modelling ground-based telescope PSFs. Nevertheless, they are not yet capable of successfully modelling the chromatic variations in addition to the spatial variations and the super-resolution.

We proceed by describing several PSF models in chronological order. The first methods, described in more detail, were used to process real data from different surveys, except for Resolved Component Analysis (RCA). The latter methods are worth mentioning but have not been used to produce a WL catalogue with all the validation and testing it implies.

Shape interpolation

The first approaches for PSF modelling consisted in estimating some parameters of the PSF at the positions of interest. It was done for early studies in WL and is closely related to the widely-employed galaxy shape measurement method KSB [KSB95]. This precursor method only required the PSF’s ellipticity and size at the positions of the galaxies. Therefore, a full pixel image of the PSF was unnecessary. Then, the KSB method can correct the galaxy shape measurement for the effects of the PSF. The method used to interpolate the shape parameters to other positions is usually a polynomial interpolation. For example, this was the case for the early WL study of Canada-France-Hawaii Telescope Legacy Survey (CFHTLS) [Fu +08]. Gentile et al. [GCM13] reviewed the different interpolation methods and studied their performance for WL studies. Viola et al. [VMB11] performed a study showcasing different biases present in the KSB [KSB95] method that are consequence of problematic assumptions. Recent WL studies no longer use this approach. The WL studies have evolved to more sophisticated galaxy shape measurement techniques that require a full pixel image of the PSF at the position of galaxies.

Principal Component Analysis (PCA)

Principal Component Analysis is a widely-known method for multivariate data analysis and dimensionality reduction. Let us start with a set of star observations in \mathbb{R}^{p^2} that we concatenate in a matrix $\bar{\mathbf{I}} = [\bar{I}_1, \dots, \bar{I}_n]$. We have flattened the 2D images into an array to simplify expressions. One would like to represent the observations with r components $\{X_i\}_{i=1}^r$ in \mathbb{R}^{p^2} , where $r \geq n$, assuming that $p^2 > n$. The PCA method gives r orthonormal components in \mathbb{R}^{p^2} which define directions in the \mathbb{R}^{p^2} space where the variance of the dataset $\bar{\mathbf{I}}$ is maximized.

If n components are used to represent the observations, then the learned components in the PCA procedure represent a basis of the subspace spanned by the observations, or the columns of $\bar{\mathbf{I}}$. The method can be interpreted as a linear transformation to a new representation with orthogonal components. As it is usual to observe regularity in the spatial variations of the PSF, most of the dataset variability can be described with a few components. Then, one can only use the first r principal components and achieve a dimensionality reduction of the observations. The dimensionality reduction technique allows to denoise the model as the observational noise is not able to be represented with r components and only the PSF trends are well described. It is possible to observe similarities in doing a Singular Value Decomposition (SVD) of the observations and only to retain the components with the r highest singular values.

The PCA method was used to model the PSF for the SDSS [Lup+01], although it was referenced as the Karhunen-Loève transform. Jarvis et al. [JJo4] then proposed its

use in a **WL** context. Jee et al. [Jee+07] used **PCA** to model the spatial and temporal variations of the **HST PSF**. Jee et al. [JT11] also used **PCA** to model the **PSF** in **LSST** simulation. **HST**'s **COSMOS** catalogue [Sch+10] used **PCA** to model the **PSF**. The use of **PCA** showcased the utility and robustness of data-driven methods and the importance of using a pixel representation of the **PSF**.

PSFEx

PSFEx⁷ [Ber11] has been widely used in astronomy for weak-lensing surveys, for example **DES** year 1 [Zun+18], **HSC** [Man+17], and the first **WL** shape catalogue of **CFIS** [Gui+22]. It was designed to work together with the **Source Extractor** (**SExtractor**) [BA96] software which builds catalogues from astronomical images and measures several properties of the observed stars. **PSFEx** models the variability of the **PSF** in the **FOV** as a function of these measured properties. It builds independent models for each **CCD** in the focal plane and works with polychromatic observations. It cannot model the chromatic variations of the **PSF** field. The model is based on a matrix factorisation scheme, where one matrix represents **PSF** features and the other matrix represents the feature weights. Each observed **PSF** is represented as a linear combination of **PSF** features. The feature weights are defined as a polynomial law of the selected measured properties. This choice allows having an easy interpolation framework for target positions. In practice, the properties that are generally used are both components of the **PSF**'s **FOV** position. The **PSF** features are shared by all the observed **PSFs** and are learned in an optimisation problem. The **PSF** reconstruction at a **FOV** position (x_i, y_i) can be written as

$$\bar{I}_{\text{star}}^{\text{PSFEx}}(x_i, y_i) = F^{\text{PSFEx}} \underbrace{\left\{ \sum_{\substack{p, q \geq 0 \\ p+q \leq d}} x_i^p y_i^q S_{p,q} + S_0 \right\}}_{\bar{H}^{\text{PSFEx}}(x_i, y_i)}, \quad (2.35)$$

where $\bar{I}_{\text{star}}^{\text{PSFEx}}(x_i, y_i) \in \mathbb{R}^{p \times p}$ is the **PSFEx** reconstruction of the observed star $\bar{I}(x_i, y_i)$, $S_{p,q} \in \mathbb{R}^{P \times P}$ represents the learned **PSF** features or *eigenPSFs*, $S_0 \in \mathbb{R}^{P \times P}$ represents a first guess of the **PSF**, the polynomial law is defined to be of degree d , and F^{PSFEx} represents the degradations required to match the model with the observations. The model's **PSF** reconstruction is represented by \bar{H}^{PSFEx} . The first guess can be computed as a function of the median of all the observations. The dimensions p and P will be the same if no downsampling operation is included in F^{PSFEx} .

The **PSF** features are learned in an optimisation problem that aims to minimise the reconstruction error between the **PSFEx** model and the observations, which reads

$$\min_{\substack{S_{p,q} \\ \forall p, q \geq 0, p+q \leq d}} \left\{ \sum_{i=1}^{n_{\text{obs}}} \left\| \frac{\bar{I}(x_i, y_i) - \bar{I}_{\text{star}}^{\text{PSFEx}}(x_i, y_i)}{\hat{\sigma}_i^2} \right\|_F^2 + \sum_{\substack{p, q \geq 0 \\ p+q \leq d}} \|T_{p,q} S_{p,q}\|_F^2 \right\}, \quad (2.36)$$

where $\hat{\sigma}_i^2$ represent the estimated per-pixel variances, \bar{I} represents the noisy observations, and $\|\cdot\|_F$ the Frobenius norm of a matrix. The second term in Equation 2.36 corresponds to a Tikhonov regularization, where $T_{p,q}$ represents some regularization weights to favour smoother **PSF** models. The **PSF** recovery at target positions is

⁷github.com/astromatic/psfex

straightforward, one just needs to introduce new positions in the Equation 2.35 after having learned the PSF features $S_{p,q}$. The recovery at a new FOV position (x_j, y_j) simply writes

$$\bar{H}^{\text{PSFEx}}(x_j, y_j) = \sum_{\substack{p,q \geq 0 \\ p+q \leq d}} x_j^p y_j^q S_{p,q} + S_0, \quad (2.37)$$

where \bar{H}^{PSFEx} is the model's PSF reconstruction, and S_0 and $S_{p,q}$ were learned during the training procedure.

lensfit

lensfit [Mil+07; Kit+08] makes reference to a Bayesian galaxy shape measurement method for WL surveys. It also includes a data-driven PSF model that will also be referenced as *lensfit* and is sparsely described throughout the different publications involving the shape measurement *lensfit* [Mil+13; Kui+15; Gib+21]. This method has been used with real data to produce the WL shape catalogues of the CFHTLenS [Kui+15] and the KiDS-1000 [Gib+21]. The code is not publicly available.

The PSF model is different from the previous ones. PSFEx and RCA learn some features or *eigenPSFs* that are shared by all the PSFs. The *lensfit* model is fitted pixel-by-pixel. Each pixel is represented as a polynomial model of degree d of the FOV positions. The *lensfit* model can use all the observations in one exposure, meaning that it uses several CCDs at once. The model uses the low order polynomials, up to degree $n_c < d$, to be fitted independently for each CCD and the rest of the monomials are fitted using the observations from all the CCDs. This multi-CCD modelling is a significant change with respect to previous methods that built independent models for each CCD. The total number of coefficients *per pixel* is

$$N_{\text{coeff}} = \frac{1}{2} ((d+1)(d+2) + (N_{\text{CCD}} - 1)(n_c + 1)(n_c + 2)), \quad (2.38)$$

where N_{CCD} is the total number of CCDs in the camera, d represents the degree of the polynomial varying in the full FOV, and n_c the polynomial that is CCD-dependent. We can write the description of the pixel $[l, m]$ of the PSF model for a FOV position (x_j, y_j) in CCD k as follows

$$\bar{H}^{\text{lensfit}}(x_j, y_j)[l, m] = \sum_{\substack{p,q \geq 0 \\ p+q \leq n_c}} x_j^p y_j^q a_{(p,q),[l,m]}^k + \sum_{\substack{p+q > n_c \\ p+q \leq d}} x_j^p y_j^q b_{(p,q),[l,m]}, \quad (2.39)$$

where $a_{(p,q),[l,m]}^k$ is the coefficient specific of CCD k , pixel $[l, m]$, and polynomial (p, q) to be fitted to the observations. The coefficient $b_{(p,q),[l,m]}$ is shared by all the CCDs.

One thing to notice in this approach is that as the modelling of the PSF is done pixel-by-pixel, then every observation should share the same pixel grid of the PSF. There is no guarantee that an observation will have its centroid aligned with the chosen pixel grid. Therefore, the PSF model has to be aligned with the observations. Other methods, like PSFEx and RCA, interpolate the model to the observation's centroids. However, *lensfit* interpolates all the observations to the model's pixel grid with a sinc function interpolation which implies interpolating noisy images. The procedure is described in Kuijken et al. [Kui+15].

For KiDS-1000 Kuijken et al. [Kui+15], the hyperparameters used by *lensfit* are $n_c = 1$, $d = 3$, and $N_{\text{CCD}} = 32$ (from CFHT's OmegaCAM instrument), where the

images used belong to a 32×32 pixel grid. When fitting the model’s parameters, each star is given a weight that is a function of its SNR with the following empirical formula

$$w_i = \frac{s_i^2}{s_i^2 + 50^2}, \quad (2.40)$$

where s_i is the measured SNR of the star i .

Resolved Component Analysis (RCA)

RCA⁸ [Ngo+16] is a state-of-the-art data-driven method designed for the *Euclid* space mission (*Euclid*) mission [Sch+20]. The model builds an independent model for each CCD, and, like PSFEx, is based on a matrix factorisation scheme. However, there are two fundamental changes with respect to PSFEx. The first difference is that, in RCA, the feature weights are defined as a further matrix factorisation. The feature weights are also learned from the data and are constrained to be part of a dictionary built with different spatial variations based on the harmonics of a fully connected graph. The graph is built using the star positions as the nodes and a function of the inverse distance between the stars to define the edge weights. The second difference corresponds to the regularisations used in the loss function, and the optimisation algorithms [BT09; Con13]. RCA uses a positivity constraint, a denoising strategy based on a sparsity constraint in the starlet [SMF15] domain, which is a wavelet representation basis, and a constraint to learn the useful spatial variations from the graph-harmonics-based dictionary. A more detailed description can be found in section 3.3.

PSFs In the Full Field-of-View (PIFF)

PSFs In the Full Field-of-View (PIFF) [Jar+20] is a recently developed PSF model that was used for the DES year 3 [Gat+21] replacing PSFEx that was used for the DES year 1 release. The model is led by Dr Mike Jarvis, open source⁹ and still being developed. It is intended for the LSST survey. Some improvements of PIFF with respect to PSFEx are the ability to use the full focal plane to build the model, a modular and user-friendly design that will enable further improvements, and modelling the PSF in sky coordinates rather than pixel coordinates. This last element was motivated by the strong tree ring detector effect observed in DES instrument, DECam, that introduces astrometric distortions that are easier to correct in sky coordinates. Pixel coordinates refer to the coordinates defined on the pixel grid of the instrument. Sky coordinates refer to the angles in the celestial sphere, which are known as Right Ascension (RA) and Declination (DEC). The geometric transformations that allow going back and forth between the pixel and sky coordinates are known as WCS.

Being a modular PSF model, PIFF allows choosing between different options for the PSF model and the interpolation method. For example, the model can be an analytical profile like a Gaussian, a Moffat or a Kolmogorov profile, or a more general non-parametric profile called PixelGrid. The interpolation method can be a simple polynomial interpolation, K-nearest neighbours method, a Gaussian process (also known as Kriging), or a *Basis-function polynomial interpolation*. Let us clarify the difference between the first and last mentioned interpolations. The simple polynomial interpolation first fits the PSF model’s parameters \mathbf{p} for each observed star. Then, it fits the coefficients of a polynomial of the 2D star positions that will later be used to interpolate. In the *Basis-function polynomial interpolation*, the interpolation

⁸github.com/CosmoStat/rca

⁹github.com/rmjarvis/Piff

coefficients of the position polynomial are fitted at the same time that the model's parameters using all the available stars (from a single CCD or the entire exposure). If this last option is used with the PixelGrid model, it comes closer to the approaches of PSFEx and RCA without the specific characteristics of each model. We have only mentioned position polynomials, but, as in PSFEx, the interpolation polynomial can be built on any parameter of the PSF, as, for example, a colour parameter.

The PIFF PSF model includes an outlier check after the algorithm has converged. The outlier check is based on the chi-squared, χ^2 , pixel error of the model and the observation. The model implements an iterative refining approach which means that after the model has converged, one (or more) outlier star(s) is(are) removed from the observations. A new iteration then starts with the model being recomputed. Although this approach effectively removes outlier stars not representative of the PSF (because they are binary stars or have some contamination), it can be very computationally demanding. The computing time is increasing linearly with the number of iterations used, which might be problematic depending on the total area to analyse or the available computing resources. We refer the reader to Jarvis et al. [Jar+20] for more details.

The DES year 3 shape catalogue [Gat+21] used the PIFF model. The model was a PixelGrid with Basis-function polynomial interpolation using a 3rd order polynomial. Although PIFF has the *potential* to build a model across several CCDs chips, in practice each model was built independently for each CCD.

Other PSF models

- *Shapelets*: Refregier [Ref03] proposed a framework to analyse images based on a series of localised basis functions of different shapes named *shapelets*. Images can then be decomposed using these basis functions. Refregier et al. [RBo3] continued the work proposing the *shapelet* framework to be used for building shear estimates as well as modelling the PSF. The PSF modelling consists of decomposing the star images in the *shapelet* basis and then performing an interpolation of the coefficients to positions of interest. Essentially, it is an extension of the approach seen in shape interpolation. Expressing the image in *shapelet* coefficients allows denoising the star images and provides an easier framework for the galaxy-PSF deconvolution. However, capturing all the PSF structure in a finite expansion over analytical functions is not always possible, leading to information being lost. Massey et al. [MR05] extended the framework from Cartesian to polar *shapelets*.
- *Moffatlets and Gaussianlets*: Li et al. [Li+16] proposed two other series of basis functions to decompose the PSF named *Moffatlets* and *Gaussianlets*. Li et al. compared the PSF reconstruction using the aforementioned basis with a PCA-based method on LSST images. Using analytical basis functions leads to more denoised models, as expected. Nie et al. [Nie+21] continued the approach and forced the principal components being learned in the PCA-like algorithm to be built using the *Moffatlets* basis. This choice avoids the principal components of learning noise as the *Moffatlets* basis avoids it. A comparison with a reference PSF model like PSFEx to have a reference performance is missing in both analyses. In addition, the performance comparison is made at the same position as the observed stars, so the model's generalisation to other positions, a fundamental task of the PSF model, is not studied.
- *Fourier-based methods*: Zhang [Zha07] proposed a Fourier-based method for

directly measuring the cosmic shear taking into account the PSF which was further developed in several publications [Zha11; ZLF15; Lu+17; Zha+19]. The method is based on the quadrupole moments of the galaxy images (described in detail in subsection 2.6.2) but is measured in Fourier space. The handling of the PSF is also done in Fourier space. Lu et al. [Lu+17] explores different interpolation approaches for the PSF in the aforementioned Fourier framework. The 2D power spectrum of the observed PSFs are interpolated to target positions. The interpolation is done pixel-by-pixel, and the best-performing method is a well-parametrised polynomial interpolation. An advantage of the Fourier interpolation is that the 2D power spectrum is automatically centred in Fourier space, simplifying the handling of images with intra-pixel shifts or, what is the same, different centroids. Another Fourier-based shear measurement method is Bayesian Fourier Domain (BFD) [BA14; Ber+16] which is build on the Bayesian formalism. However, it does not include a specific PSF model.

- *Optimal Transport (OT)-based methods:* There exist two approaches involving Optimal Transport (OT) [PC19] to tackle the PSF modelling problem. Ngolè et al. [NS17] proposes to use Wasserstein barycenters as a non-linear geometric-aware interpolation procedure of a low-dimensional embedding representation of the PSFs. Although elegant, the performance of the approach does not seem to justify its computational burden. In the comparison method, a Radial Basis Function (RBF) interpolation of the principal components obtained through PCA achieves a similar performance. The performance of the PCA method is better in terms of ellipticity but slightly worse in terms of the pixel error and PSF FWHM. Schmitz [Sch19] worked on a data-driven PSF model based on RCA that would be able to model the chromatic variations of the PSF through the use of Wasserstein barycenters that were previously developed in Schmitz et al. [Sch+18]. The OT-based PSF model coined λ RCA was compared to RCA. The comparison showed a lower pixel and size error for λ RCA, although the ellipticity error was similar or better for RCA. This method for capturing the chromatic dependence of the PSF assumes that the chromatic variation of the PSF is smooth over all the passband. This assumption can be somehow considered as if the only chromatic effect of the PSF is due to the diffraction phenomena, which exhibits a smooth variation with the $1/\lambda$ dependence in the WFE that was already presented in Equation 2.26. However, if this is not the case and another type of non-smooth chromatic variation is present, currently occurring in *Euclid* [Ven+16], there is no straightforward way to adapt the λ RCA model to account for it.
- *Wavefront approaches:* Davis et al. [DRR16] proposed a wavefront-based PSF model for the DECam instrument. The model is based on Zernike polynomials [Nol76] that are fitted to out-of-focus images that have considerably more wavefront information than in-focus stars. Then, for each exposure, a new fit is done based on the measured quadrupole moments of the in-focus star images. It is not easy to understand at which point the quadrupole moments are constraining the proposed PSF model and at which point it is the base physical wavefront measured from the out-of-focus images that are the only part driving the performance of the model. However, the model requires out-of-focus observations that are not always available. Soulez et al. [SCU16] proposed to model the propagation of light through the mirrors of the optical system. The PSF modelling problem is recast into a phase retrieval problem. The article is a

proof-of-concept as there are only qualitative results, and many PSF-modelling difficulties remain unaddressed.

- *Deep learning approaches:* A model coined PSF-NET was proposed by Jia et al. [Jia+20c], and is based on two Convolutional Neural Networks (CNNs) trained jointly. One network has to transform high-resolution images into low-resolution images, while the other has to do the opposite. The CNNs are trained in a supervised way expecting that the first network will learn a PSF manifold. However, it is not clear how the approach handles the spatial variation of the PSF, and it has not been tested for a WL purpose. Jia et al. [Jia+20a] proposed another approach for PSF modelling based on Denoising Auto-Encoders (DAEs), but the spatial variation of the PSF remains untackled. Another approach is followed by Herbel et al. [Her+18], where the PSF profile is modelled by a parametric function consisting of a base profile of two Moffat profiles and several parametrised distortions to be more expressive. A CNN is trained in a supervised manner to predict the parameters of the parametric profile from a noisy star observation. The neural network provides a good estimation of the parameters, but the spatial variation of the PSF is, again, not addressed. Having a PSF model that can model the observations is important. However, in PSF modelling for WL analysis, a crucial part is to capture the spatial variations of the PSF and that the model outputs the PSF at different positions in the FOV. In WL studies, the objects of interest are at positions where we do not have observations of the PSF.

2.5 General comments on PSF modelling

2.5.1 Desirable properties of PSF models

In the previous section, we reviewed the most relevant PSF models developed so far up to our knowledge. After studying many models, we can conclude on desirable properties they should aim for. The PSF model should:

- (a) *Have an accurate modelling of the PSF light profile.* The smoothness, or structure, in the PSF profile is coming because the PSF comes from the FT of a particular finite-length aperture that limits the frequency content of the PSF. This frequency limitation avoids having a Dirac distribution as a PSF, as it would require an infinite frequency content to build the Dirac. The *ideal PSF* that would give us a perfect resolution would be the identity element of the convolution, therefore the Dirac distribution. In ground-based telescopes, the effect of the atmosphere can be seen, as we already studied and considering a long exposure time, as a low-pass filter for the PSF, smoothing the PSF light profile. We can exploit the smoothness of the PSF profile to build the PSF model.
- (b) *Not produce noisy estimations of the PSF.* The presence of noise in the PSF estimations is an issue for the purely data-driven models. PSF models based on fixed basis functions like Shapelets or Moffatlets will always output denoised PSFs as the noise cannot be represented by their basis functions. However, these models are not flexible enough to model the complexity of the PSF light profile from real images and generate a modelling bias. As a consequence, recent WL surveys, like HSC or DES, are relying on data-driven model like PSFEx or PIFF.
- (c) *Have a good generalisation power.* The PSF model needs to capture most of the relevant information from the observations, including the spatial variation of the

PSF field and the PSF light profile. Then, to be able to exploit the information and predict the PSF at positions where there is no direct information of the PSF. WL analysis rely on accurate galaxy shape measurements, and in those positions, there are no observations of the PSF. The WL analysis relies on the generalisation power of the PSF model.

- (d) *Be able to exploit the structure of the PSF field variations.* This desired property is related to the previous point (c). An exciting approach to obtaining a good generalisation power is to learn the structure of the PSF fields variations. This structure is a consequence of the physical properties of the telescope's optical system. The subsection below provides a physical understanding of the PSF field structure which imposes a certain smoothness to the variations. The spatial variations are also structured due to the atmosphere if we study the PSF field of a ground-based telescope. The existing structure of the PSF field should be exploited more concisely, which would reduce the high dimensionality of the PSF modelling problem. A data-driven PSF model should aim to use a low complexity representation of the PSF field, which would be able to learn its structure and spatial variations.
- (e) *Handle the CCD discontinuities of the PSF field.* The CCDs misalignments are the source of the discontinuities in the spatial variations of the PSF. Nevertheless, the PSF field is piece-wise continuous, and the borders delimiting the discontinuity are well known as the geometry of the focal plane is known with precision. A straightforward way to handle the discontinuities is to independently model the PSF for each CCD. Although this is simple to implement, it limits the number of stars available to constrain the PSF model. Another more difficult but potentially more powerful approach is to build a PSF model for the entire focal plane, taking into account the focal plane discontinuities. The last approach has only been successfully implemented in *lensfit* with a simple pixel-by-pixel modelling approach.
- (f) *Be unbiased.* We want that the residuals between the PSF model estimations and the PSF field to be uncorrelated. Avoid biases propagating from the PSF model into the shape measurement. The WL analysis relies on the study of the correlation between the galaxy shapes. Suppose a PSF modelling error introduces a systematic error, or bias, into the shape measurement. In that case, the shear estimation will be biased and propagate into the cosmological analysis.
- (g) *Be robust to variations and contaminations of the star sample.* Contamination can come from the star purity, the fact that the objects classified as stars are good representations of the PSF and are not small galaxies or binary stars [KC17]. Contamination can also come from an imperfect image preprocessing where detector effects residuals, like CTI, have remained. The model should be robust to the different spatial distributions of the observed stars, to different SNRs, and the different number of observed stars.
- (h) *Work smoothly with the shape measurement algorithm.* The PSF model is tightly related to the galaxy shape measurement algorithm used and should be developed with that in mind. At first, the algorithms relied on some summary statistic of the PSF, like the image moments that we will define in section 2.6, to correct the galaxy images for the effects of the PSF. However, the way the PSF affects the images is not straightforward in general. The summary statistics analysis of the PSF and its effects are precise for specific cases, for example, if the PSF and

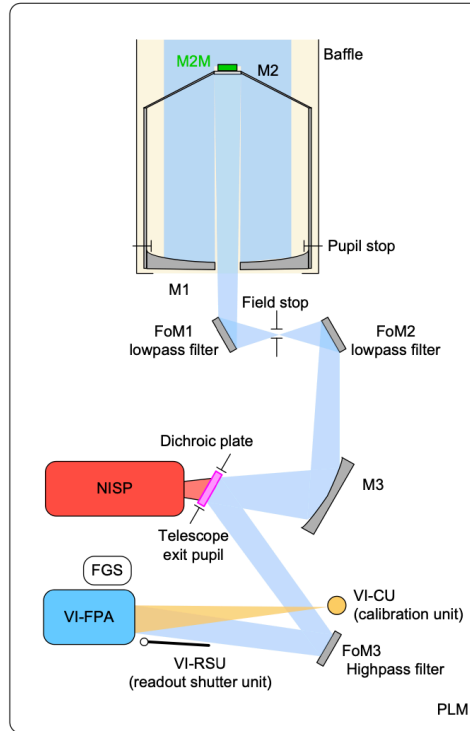


Figure 2.5-14 – Illustration of *Euclid*'s optical system. Credit: Image from Racca et al. [Rac+16].

the galaxy have Gaussian profiles. Shape measurement algorithms that require a full pixel description of the PSF are now dominating WL analyses, and it will continue to be the case for the *Euclid* mission and LSST.

We continue by giving insights on the nature of the PSF field spatial variations from an optical perspective.

2.5.2 On the structure of the PSF field spatial variations

To have a better understanding of the spatial variations of the PSF, we have to pay more attention to the telescope's optical system. We will take as example the *Euclid*'s optical system, which is presented in Figure 2.5-14. In this case, we only consider the variations due to the telescope's optics, which is the predominant source of spatial variation in a space mission. The optical system is composed of big mirrors like M₁, which have sizes closer to the aperture of the entrance pupil, and smaller mirrors like FOM2. The input wavefront acquires WFEs as it propagates through the different mirrors of the optical system. In other words, each mirror introduces specific Optical Path Differences, or WFEs, to the incoming wavefront.

Let us define a mirror's footprint as the projection of an incoming wavefront from a specific point in the FOV into the optical surface. The footprint of the different *Euclid*'s optical elements' footprints changes as the position in the FOV changes. Figure 2.5-15 shows the variations of the different *Euclid*'s optical elements' footprints. The footprints of mirrors located close to the entrance or exit pupil, like M₁, do not change much from the positions in the FOV. Thus, the WFEs acquired from these mirrors have a high spatial correlation (with

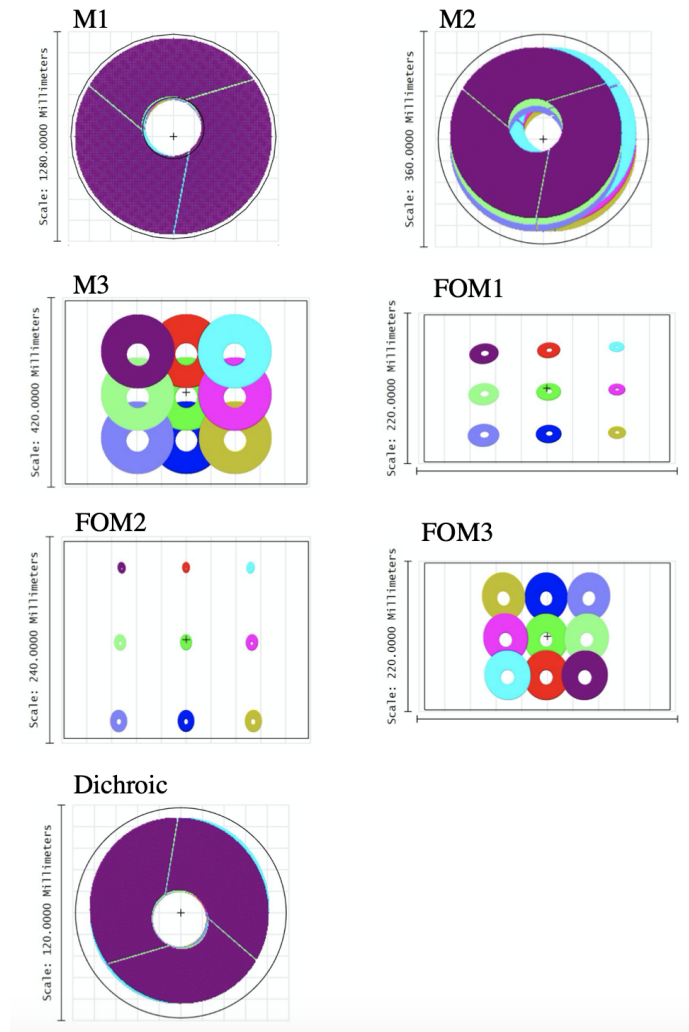


Figure 2.5-15 – Variations of the different optical elements' footprints due to 9 incoming beams from different FOV positions. The figures were computed using ray-tracing simulations. Credit: Image from the EC technical note EUCL-OXF-TN-8-012 produced by Prof Lance Miller.

respect to variations in the [FOV](#) position), or, in other words, the spatial variation of the [PSF](#) field due to these fields is low. The projection of the wavefront footprint in the smaller mirrors like [FOM2](#) spans a region that represents a small fraction of the total [FOM2](#) mirror surface. As a consequence, the projections of distant [FOV](#) positions in the [FOM2](#) mirror are not correlated, meaning that the [WFEs](#) introduced by this mirror strongly depends on the [FOV](#) position. Therefore, the [PSF](#) field spatial (with the [FOV](#)) variations are high.

It is also important to consider the frequency composition of the [WFE](#) introduced by the mirrors. The size of the mirror relative to the exit pupil impacts the Zernike composition of the [WFE](#) added by that mirror to the total [WFE](#) introduced by the system. A mirror with a large relative size will have a [WFE](#) with a higher frequency component, which also means that we would require a higher number of Zernike orders to describe the [WFE](#) contribution appropriately. This higher frequency component is due to the input rays' compression from a large surface into the exit pupil, a smaller surface. On the contrary, if the mirror has a small relative size, the rays reflected on it will finally expand to cover the entire exit pupil. As a consequence, the [WFE](#) contribution from the mirror will have a lower frequency component as the high frequencies are expanded into lower frequencies.

To summarize, we have distinguished two properties of the mirrors that affect the total [WFE](#). The first is the frequency component of the [WFE](#). The second is the spatial variation of the [WFE](#) introduced with respect to the [FOV](#) positions. While bigger mirrors introduce a [WFE](#) with a higher frequency component, they are usually placed close to the entrance or exit pupil and have a lower spatial variation. The opposite is true for smaller mirrors, where the [WFE](#) they introduce has a lower frequency component. However, as the smaller mirrors are usually located away from the aforementioned pupils, they have a higher spatial variation.

The dichroic deserves special attention as the [EC](#) has discovered it adds spurious spectral dependencies to the wavefront. [Figure 2.5-16](#) shows the beam's angle of incidence with respect to each optical element's normal. We can observe that the angle of incidence in the dichroic changes as a function of the [FOV](#) position. The spectral contribution of the dichroic depends on the incoming rays' angle of incidence. Consequently, we have a complex spectral contribution of the dichroic that also changes with respect to the [FOV](#) position and the wavelength. The nature of the dichroic and the change of angle of incidence makes the spectral contribution a non-separable function of position and wavelength. At the time of writing, modelling the dichroic is the subject of study within the [EC](#).

Once the [PSF](#) model has been developed with all the aforementioned properties in mind, it is essential to validate the model's performance. The validation should help to ensure that the expected performance of the model is achieved or identify sources of problems and drive directions for the improvement of the model. In the next section, we give an overview of [PSF](#) model validations.

2.6 Validation of PSF models

The validation of [PSF](#) models is not trivial. In order to derive a validation method, it would be necessary to quantify the impact of [PSF](#) modelling errors on the final objective of our analysis. Strictly speaking, for our [WL](#)-based cosmological analysis, the objective would be to derive constraints on the parameter of the cosmological model under analysis. This exercise is hard given the analysis's complexity and the




Figure 2.5-16 – Angle of incidence with respect to the normal of each optical element from rays incoming from one FOV position. The dichroic angles are plotted for another FOV position to illustrate its change. The colour bars are in degrees, and the x and y axes are in millimetres. The figures were computed using ray-tracing simulations. Credit: Image courtesy of Dr Pierre-Antoine Frugier from the EC.

data volume. Nevertheless, and using some assumptions, it was done to set the PSF modelling requirements for the *Euclid* mission as will be shown in subsection 2.6.3. However, some assumptions used do not always hold for the high complexity of the PSF in a space-based mission like *Euclid*. Even though it is important to derive requirements for the PSF model, these do not give much information on the nature of the errors and possible problems the PSF model has. Therefore, it is necessary to derive different diagnosis or null tests. Jarvis et al. [Jar+16] proposes a set of null tests for the science verification of the DES WL shear catalogues that includes the validation of the PSF model. Next, we will describe the most used PSF diagnosis that will help us to validate the performance of the PSF models.

2.6.1 Pixel-based metrics

The most straightforward diagnosis we can think of is to compute the pixel residual of our PSF model. We would separate the observed stars into two distinct sets, the first one will be used to constrain or train the PSF model, and the second one to validate the model. Once the PSF model has been trained, we use it to predict the PSF in the positions of the stars in the testing set. We can then compute the Root Mean Squared Error (RMSE) of the pixel reconstruction residuals. Ideally, the PSF model should be able to predict the observed star perfectly, and the reconstruction residual would only contain the observational noise. If we work with simulations, we can produce noiseless stars for our testing set, and the RMSE will directly indicate the pixel reconstruction error. When working with real data, we need to consider the

different noise levels in the data. We propose pixel-based reconstruction metrics for real data in [subsection 3.6.3](#).

2.6.2 Moment-based metrics

WL analysis are interested in measuring the shape of galaxies as the measured ellipticity is an estimator of the shear. Cosmologists have developed formulations to relate the PSF errors, expressed in terms of shape and size metrics [Mas+12], to the errors of the cosmological parameters of interest [Cro+13]. Therefore, it seems natural to have diagnosis metrics based on the ellipticity and size of the PSF. Both of these metrics are determined using the moments of the polychromatic observation $\bar{I}[u, v]$. Following Hirata et al. [HS03], we redefine the image moments that we will use in practice, including a weight function as follows

$$\bar{\mu} = \frac{\int \mu \bar{I}[u, v] w[u, v] du dv}{\int \bar{I}[u, v] w[u, v] du dv}, \quad (2.41)$$

$$M_{\mu\nu} = \frac{\int \bar{I}[u, v] (\mu - \bar{\mu}) (\nu - \bar{\nu}) w[u, v] du dv}{\int \bar{I}[u, v] w[u, v] du dv}, \quad (2.42)$$

where $\mu, \nu \in \{u, v\}$ and $w[u, v]$ is a weight window that is used to avoid the divergence of the integrals due to noise. The weight function is also useful to compute the moments from PSFs that are diffraction-limited (i.g. an Airy profile), as they prevent the integral from diverging due to the wings of the PSF. [Equation 2.41](#) defines the first-order moments, while [Equation 2.42](#) defines the second-order moments. The ellipticities, or *shape* metrics, are defined as

$$e = e_1 + ie_2 = \frac{(M_{uu} - M_{vv}) + i2M_{uv}}{M_{uu} + M_{vv}}, \quad (2.43)$$

and the size metric is defined as

$$R^2 = T = M_{uu} + M_{vv}. \quad (2.44)$$

The method used to estimate these metrics is the widely-used adaptive moment algorithm from Galsim's [Hirata-Seljak-Mandelbaum \(HSM\) module](#)¹⁰ [HS03; Man+05]. The adaptive moment algorithm measurement provides σ as size which relates to the above-defined size metric as $R^2 = 2\sigma^2$. The measurements are carried out on well-resolved polychromatic images. If the observations are undersampled, as is the case for *Euclid*, a super-resolution step is required for the model.

The measurements of the shape parameters based on the image moments are very sensitive to the image noise. If we are working with real data, we do not have access to the ground truth images and are obliged to work with noisy observations. Therefore, we need to average over many objects in order to be able to conclude from the different diagnostics. We continue by presenting different moment-based metrics.

Shape RMSE

We start with a set of test stars and their corresponding PSF estimations. Then, the most direct moment-based metric is to compute the RMSE of the ellipticities and size residuals between the observations and the model prediction.

¹⁰github.com/GalSim-developers/GalSim

Meanshapes

A good diagnostic is to compute the spatial distribution of the ellipticities and size residuals, which we call *meanshapes*. This diagnostic allows us to inspect if there are spatial correlations in the shape and size residuals, which would indicate that the PSF model is not capturing certain spatial variations from the PSF field. In order to have a finely sampled distribution, we need to average over many exposures, as the available stars from a single exposure are not enough to observe patterns. The shape measurements are also noisy, and therefore averaging over many exposures allows to smooth out the residuals to observe systematic errors. In practice, we divide the focal plane into several spatial bins, consider several exposures, and then the value of each bin is built by averaging the residuals of all the stars that are located within that bin. A ground-based survey allows us to average the ellipticity contribution of the atmosphere [Hey+12a], as it can be considered a random field with zero mean. Then, the observed ellipticity distribution over the focal plane is due to the telescope's optical system that is consistent in every exposure. It is also possible to plot the same spatial distribution but observe the positions of the stars. Such a plot will help to observe if there are regions of stellar under density that could eventually affect the PSF model.

ρ -statistics

Rowe [Row10] proposed to compute the auto- and cross-correlations of the ellipticities and their residuals as a diagnostic. The diagnostic was then expanded by Jarvis et al. [Jar+16] to a combination of ellipticities, sizes, and residuals. Following Jarvis et al. [Jar+16], we define the ρ -statistics as follows

$$\rho_1(\theta) = \langle \delta e_{\text{PSF}}^*(\boldsymbol{\theta}') \delta e_{\text{PSF}}(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle, \quad (2.45)$$

$$\rho_2(\theta) = \langle e_{\text{PSF}}^*(\boldsymbol{\theta}') \delta e_{\text{PSF}}(\boldsymbol{\theta}' + \boldsymbol{\theta}) \rangle, \quad (2.46)$$

$$\rho_3(\theta) = \left\langle \left(e_{\text{PSF}}^* \frac{\delta R_{\text{PSF}}^2}{R_{\text{PSF}}^2} \right) (\boldsymbol{\theta}') \left(e_{\text{PSF}} \frac{\delta R_{\text{PSF}}^2}{R_{\text{PSF}}^2} \right) (\boldsymbol{\theta}' + \boldsymbol{\theta}) \right\rangle, \quad (2.47)$$

$$\rho_4(\theta) = \left\langle \delta e_{\text{PSF}}^*(\boldsymbol{\theta}') \left(e_{\text{PSF}} \frac{\delta R_{\text{PSF}}^2}{R_{\text{PSF}}^2} \right) (\boldsymbol{\theta}' + \boldsymbol{\theta}) \right\rangle, \quad (2.48)$$

$$\rho_5(\theta) = \left\langle e_{\text{PSF}}^*(\boldsymbol{\theta}') \left(e_{\text{PSF}} \frac{\delta R_{\text{PSF}}^2}{R_{\text{PSF}}^2} \right) (\boldsymbol{\theta}' + \boldsymbol{\theta}) \right\rangle, \quad (2.49)$$

where * denotes complex conjugation, $\boldsymbol{\theta}$ and $\boldsymbol{\theta}'$ denote sky positions, and θ denotes the scalar distance of $\boldsymbol{\theta}$. If we assume that the ellipticities are random fields that are isotropic and statistically homogenous, we can compute the correlation $\rho(\boldsymbol{\theta}, \boldsymbol{\theta}')$ as $\rho(|\boldsymbol{\theta} - \boldsymbol{\theta}'|) = \rho(\theta)$, using a the scalar distance θ . This choice means we are assuming translational and rotational symmetry, which is what the Cosmological Principle states. We will define several θ bins in a logarithmic scale, corresponding to $\ln \theta - \Delta \ln \theta / 2 < \theta_{ij} < \ln \theta + \Delta \ln \theta / 2$, where $\theta_{ij} = |\boldsymbol{\theta}_i - \boldsymbol{\theta}_j|$ is the distance between two objects at $\boldsymbol{\theta}_i$ and $\boldsymbol{\theta}_j$. Then, the correlation function at θ can be computed using the following unbiased estimator of ρ that is

$$\hat{\rho}(\theta) = \frac{\sum_{i,j} w_i w_j e_i^A {}^* e_j^B}{\sum_{i,j} w_i w_j}, \quad (2.50)$$

where we are computing, as an example, the correlation of two ellipticities e^A and e^B , and the weights depend on the SNR of the ellipticity measurements. We carry out the weighted sum over the pairs of objects within each bin.

The former statistics are interesting as they can be propagated to the shear **2PCF** from [Equation 1.38](#). Following Jarvis et al. [[Jar+20](#)], we include the **PSF** errors into the shear **2PCF** and make appear the ρ -statistics. We can then express the systematic error in the shear **2PCF** as

$$\begin{aligned} \delta\xi_+(\theta) = & 2 \left\langle \frac{R_{\text{PSF}}^2}{R_{\text{gal}}^2} \frac{\delta R_{\text{PSF}}^2}{R_{\text{PSF}}^2} \right\rangle \xi_+(\theta) + \left\langle \frac{R_{\text{PSF}}^2}{R_{\text{gal}}^2} \right\rangle^2 \rho_1(\theta) - \alpha \left\langle \frac{R_{\text{PSF}}^2}{R_{\text{gal}}^2} \right\rangle \rho_2(\theta) \\ & + \left\langle \frac{R_{\text{PSF}}^2}{R_{\text{gal}}^2} \right\rangle^2 \rho_3(\theta) + \left\langle \frac{R_{\text{PSF}}^2}{R_{\text{gal}}^2} \right\rangle^2 \rho_4(\theta) - \alpha \left\langle \frac{R_{\text{PSF}}^2}{R_{\text{gal}}^2} \right\rangle \rho_5(\theta), \end{aligned} \quad (2.51)$$

where α is the **PSF** leakage from [Equation 1.49](#). [Equation 2.51](#) provides a way to propagate requirements from the shear **2PCF** to the **PSF** by means of the ρ -statistics.

Scale-dependent PSF leakage

The **PSF** leakage allows us to quantify how the **PSF** is affecting the shear estimation through the shape measurement. This diagnostic depends on the shape measurement method used and how this method handles the **PSF**. The leakage appears in our linear modelling of the shear bias from [Equation 1.49](#), where we have decomposed the total additive bias into **PSF**-dependent (leakage) and **PSF**-independent terms. Following Jarvis et al. [[Jar+16](#)], we can define the scale-dependent **PSF** leakage estimation as

$$\alpha(\theta) = \frac{\xi_+^{\text{gP}}(\theta) - \langle e_{\text{gal}} \rangle^* \langle e_{\text{PSF}} \rangle}{\xi_+^{\text{PP}}(\theta) - |\langle e_{\text{PSF}} \rangle|^2}, \quad (2.52)$$

where e_{gal} and e_{PSF} are the measured galaxy and **PSF** ellipticity, ξ_+^{gP} is the cross-correlation between the galaxy ellipticity and the **PSF** ellipticity, and ξ_+^{PP} is the autocorrelation of the **PSF** ellipticity. Even if the result is scale-dependent due to the correlation functions, in [Equation 1.49](#), we can see that the **PSF** leakage is not a scale-dependent quantity. Therefore, the leakage is expected to be consistent at all scales.

2.6.3 PSF error propagation and PSF requirements

The pioneer in the **PSF** error propagation was Paulin-Henriksson et al. [[Pau+08](#)] which was then extended by Massey et al. [[Mas+12](#)]. The proposed framework is based on the second-order moments of the images, i.e. complex ellipticity e and size R^2 . It expresses how the **PSF**, or some other effect, affects the observed ellipticity and size. For example, only considering the effect of **PSF** on the unweighted moments from [subsection 1.2.3](#), we obtain

$$e_{\text{obs}} = e_{\text{gal}} + \frac{R_{\text{PSF}}^2}{R_{\text{PSF}}^2 + R_{\text{gal}}^2} (e_{\text{PSF}} - e_{\text{gal}}) \quad \text{and} \quad R_{\text{obs}}^2 = R_{\text{gal}}^2 + R_{\text{PSF}}^2, \quad (2.53)$$

where the subscript $_{\text{obs}}$ refers to the quantity measured to the observed galaxy, the subscript $_{\text{gal}}$ refers to the intrinsic quantity of the galaxy, and $_{\text{PSF}}$ refers to the quantity measured to the **PSF**. There is an intrinsic assumption in the previous equation that is the observational model, $I_{\text{obs}} = I_{\text{gal}} * H_{\text{PSF}}$, and that all the moments are well defined. Then, [Equation 2.53](#) can be rewritten to express the quantity of interest, the intrinsic galaxy ellipticity, as follows

$$e_{\text{gal}} = \frac{e_{\text{obs}} R_{\text{obs}}^2 - e_{\text{PSF}} R_{\text{PSF}}^2}{R_{\text{obs}}^2 - R_{\text{PSF}}^2}. \quad (2.54)$$

The error propagation consists of expanding the previous equation in a first-order Taylor series with respect to the quantities that we are interested in propagating. In this case, it will be the shape and size of the PSF, and the propagation writes

$$\widehat{e}_{\text{gal}} \approx e_{\text{gal}} + \frac{\partial e_{\text{gal}}}{\partial (R_{\text{PSF}}^2)} \delta (R_{\text{PSF}}^2) + \frac{\partial e_{\text{gal}}}{\partial e_{\text{PSF}}} \delta e_{\text{PSF}}, \quad (2.55)$$

where δ refers to errors in the model with respect to the ground truth. It is straightforward to compute the partial derivatives in Equation 2.55 from the expression in Equation 2.54. We then obtain the following expression

$$\widehat{e}_{\text{gal}} \approx e_{\text{gal}} \left(1 + \frac{\delta (R_{\text{PSF}}^2)}{R_{\text{gal}}^2} \right) - \left(\frac{R_{\text{PSF}}^2}{R_{\text{gal}}^2} \delta e_{\text{PSF}} + \frac{\delta (R_{\text{PSF}}^2)}{R_{\text{gal}}^2} e_{\text{PSF}} \right), \quad (2.56)$$

The previous ellipticity estimator expression can be used in the assumption of Equation 1.32 to obtain a shear estimator. The obtained estimator can then be used in the shear bias parametrization from Equation 1.49. At this point, we can express the additive and multiplicative biases as a function of the elements from Equation 2.56. This analysis shows us that the multiplicative bias is related to the size of the PSF with its estimation error and the size of the galaxy. The result was expected if we paid attention to the first term of Equation 2.56.

This framework allows us to consider different types of errors. Massey et al. [Mas+12] uses it to include errors due to non-convolutional detector effects, due to imperfect shape measurement, and the fact that the shape measurement method used weighted moments, i.e. Equation 2.42, instead of the unweighted counterpart considered in the analysis. The procedure consists in adding the desired effect to the galaxy ellipticity expression, e.g. Equation 2.54, and then adding their corresponding partial derivatives to the Taylor expansion seen in Equation 2.55. Cropper et al. [Cro+13] uses this formalism to derive requirements for a WL mission in space.

The aforementioned framework was used to derive the current PSF model requirements for *Euclid*, which are:

- the requirement on the ellipticity error of the PSF model¹¹ is $\sigma(e_{\text{res}}) < 2 \times 10^{-4}$,
- the requirement on the size error for the PSF model¹² is $\sigma(R_{\text{res}}^2)/\langle R^2 \rangle < 10^{-3}$,

where $\sigma(\cdot)$ is the root means square value, and, e_{res} and R_{res}^2 refer to the residual ellipticity and size of the PSF model, respectively.

The previous framework has a fundamental limitation, noted by Schmitz et al. [Sch+20], which is that the propagation is based on second-order moments of the images. A perfect second-order moment estimation of the PSF would have a zero shear bias contribution in the formalism described. However, the PSF's Higher Moments Errors (HME) of the PSF will impact the shear biases and are not considered in the framework proposed by Paulin-Henriksson et al. [Pau+08]. The higher the contribution of higher moments on the PSF the more significant the deviations will be. A space mission like *Euclid* will have a PSF close to the diffraction limit, meaning that its shape will be complex and not well described by a Gaussian (or by its second-order moments). The important higher moment contributions are why the previous requirements, based on the second order moments, should be used *with caution*. The LSST collaboration,

¹¹ *Euclid* requirement ID: R-WL.2.1-8.

¹² *Euclid* requirement ID: R-WL.2.1-9.

concerned with the previous issue, studied the contribution to systematic biases of the HME of the PSF model on the shear measurement [ZMC21; Zha+22]. Zhang et al. [ZMC21] showed that the HME of the PSF model may be a significant source of systematics in upcoming WL analyses. Zhang et al. [Zha+22] studied the impact of moments from the 3rd to 6th order on the cosmological parameter inference, concluding that the HME of PSF models like PSFEx and PIFF should be reduced for future surveys like LSST if the WL analysis is to remain unchanged.

* * *
* *
*

Part II

PSF models for ground-based telescopes

Modelling the PSF in the full focal plane at once

Chapter Outline

3.1	Introduction	80
3.2	The observation model	81
3.3	Resolved component analysis (RCA)	82
3.4	A new family of MCCD methods	83
3.4.1	The MCCD data model	83
3.4.2	Inverse problem and regularisation	84
3.4.3	Global model	85
3.4.4	Local model	86
3.4.5	Optimisation problem	87
3.4.6	Algorithm	88
3.4.7	PSF recovery	89
3.5	Numerical experiments with simulated data	92
3.5.1	Data	92
3.5.2	Training set	92
3.5.3	Testing data set	93
3.5.4	Quality criteria	93
3.5.5	Model parameters	95
3.5.6	Results	95
3.5.7	Comparison of computing resources	98
3.6	Numerical experiments with UNIONS/CFIS data	98
3.6.1	Dataset	98
3.6.2	Model parameters	99
3.6.3	A new metric for real data: the Q criteria	101
3.6.4	Results	103
3.7	Conclusion	104

This chapter is based on a publication accepted in a peer-reviewed journal:

Liaudat, T., Bonnin, J., Starck, J.-L., Schmitz, M. A., Guinot, A., Kilbinger,

M., Gwyn, S. D. J., ‘Multi-CCD modelling of the point spread function’. In: *A&A* 646 (2021), A27

THIS chapter presents a new PSF model for ground-based telescopes coined **Multi-CCD (MCCD)**, that was previously published in Liaudat et al. [Lia+21a]. This new PSF model allows building a single model for the entire focal plane that spans several CCDs in a particular geometry. The model is first described and then tested with simulated data. We end the chapter by showing that the proposed model is robust enough to handle real data from the CFIS survey and has better performance than a state-of-the-art PSF model. The proposed **Multi-CCD (MCCD)** model is publicly available on GitHub¹ with its documentation page.

3.1 Introduction

The good quality of the incoming observations of CFIS allow us to carry out a WL analysis and provide a unique opportunity to develop and test new PSF models. We provide more information about the CFIS survey in section 4.1. In the previous chapter, we have seen the difficulty of building a physically motivated PSF model for a ground-based survey. The temporal integration seen in subsection 2.3.3 due to a fast-changing atmosphere and a long exposure time, makes it computationally intractable for the moment. Consequently, a pure-data driven approach seems the most appropriate for a ground-based survey. The narrow *r*-band filter used in CFIS that can be seen in Figure 2.2-8, allow us to neglect chromatic variations of the PSF without making much error.

We have seen in the previous chapter that most of the data-driven PSF models are build independently in a single CCD chip. This choice, simplifies the building of the PSF model but greatly limits the amount of stars that can be used to constrain the model. The first WL analysis of the CFIS data in Guinot et al. [Gui+22] used the PSFEx model. With more data being observed, a natural improvement to the upcoming version of the catalogue would be to use a better PSF model. This motivated us to tackle one of the principal limitations of the current ground-based PSF models, the number of stars. We propose to increase the number of stars to have more constraining power, which will allow us to use more complex PSF models and thus capture more complex PSFs shapes and PSF spatial variations. A straightforward way to accomplish this goal is to drop the independent modelling of each CCD chip and develop a single model for the entire focal plane. As the MegaCam instrument used in CFIS has 40 CCDs, this new choice will increase by a factor of ~ 40 the number of available stars in average.

The *lensfit* model [Mil+13] is currently working in a multi-CCD approach, but the model is built in a pixel-by-pixel basis, neglecting the correlations between pixels in the same image. The recently proposed PIFF model [Jar+20] intends to build a PSF model for the entire focal plane. However, it is not mature to work with real data, and the recent DES WL shape catalogue [Gat+21] used the CCD-independent mode for the PIFF model.

In this chapter, we present a new PSF model built on the entire focal plane coined **MCCD**. The new method is based on **RCA** and can capture large patterns spreading across several or all CCDs. We compare the results with both **RCA** and **PSFEx** based

¹github.com/CosmoStat/mccd

Variable	Description
<i>Observational model</i>	
\mathcal{H}	PSF field
\mathcal{F}	degradation operator
u_i^k	2D position of star i in CCD k
n_{star}^k	number of observed stars in CCD k
N	number of observed stars in all the CCDs
$n_{k,i}$	noise realisation of star i in CCD k
$y_{k,i}$	square star observation stamp i on CCD k
n_y	number of pixels on one dimension of $y_{k,i}$
D	downsampling factor
$\mathbf{y}_{k,i}$	1-dimensional (1D) column representation of $y_{k,i}$
Y_k	matrix stacking all the star observations $\mathbf{y}_{k,i}$
<i>PSF model</i>	
\hat{H}_k	PSF model estimation of the observed stars Y_k
r_k, \tilde{r}_k	local and global number of eigenPSFs
S_k, \tilde{S}	local and global <i>eigenPSF</i> matrices
A_k, \tilde{A}_k	local and global weight matrices
$\alpha_k, \tilde{\alpha}$	local and global spatial constraint weights
V_k^\top, Π_k	local and global spatial constraint dictionaries
$K_\sigma^{\text{Loc}}, K_\sigma^{\text{Glob}}$	local and global denoising parameters
$(e_{k,i}, a_{k,i})$	RCA graph constraint parameters
$\mathbf{w}_{k,i}, \tilde{\mathbf{w}}_i$	local and global weight vectors for the sparsity inducing term
Φ	sparsity inducing transform
<i>PSF recovery</i>	
ϕ	Radial Basis Function (RBF) kernel
N_{RBF}	number of elements used to estimate the RBF interpolant
$(\lambda_i)_{i=1}^{N_{\text{RBF}}}$	RBF interpolation weights
$A_k(N_{\text{RBF}})$	weight matrix composed by the N_{RBF} closest stars of a given target position
$A_{k,u}, \tilde{A}_{k,u}$	local and global interpolated weight columns for a target position u
$\hat{H}(u)$	recovered PSF at position u

Table 31 – Summary of important variables used in this chapter.

on simulations and real data from CFIS. section 3.3 reviews the RCA method, while the proposed MCCD methods are described in section 3.4. Experiments on simulated images are shown in section 3.5 and tests on real data are presented in section 3.6. We give our conclusions in section 3.7. In addition, Table 31 provides a glossary of variables used throughout this chapter.

3.2 The observation model

Let us define $\mathcal{H}(u)$ as the PSF field involved in our problem. It is a continuous function of a two-dimensional position, $u = (x, y)$, which, in principle, could be image coordinates based on the camera's CCD pixels or could also be celestial coordinates

such as right ascension and declination. Throughout this chapter, we assume that the PSF field accounts for the contribution of all effects from optical aberrations, diffraction, and atmospheric distortions.

Our observation model is a simplified version of the general observational model presented in section 2.2. The current model consists of images, I_k , the pixels in one CCD chip, k , which contains n_{star}^k noisy stars at positions, u_i^k . We define a ‘stamp’ as a square small image cutout centred on a single star. Each star observation stamp, i , on the CCD’s k can be written as:

$$y_{k,i} = \mathcal{F}(\mathcal{H}(u_i^k)) + n_{k,i}, \quad (3.1)$$

where $n_{i,k}$ accounts for a noise image that we will consider to be white and Gaussian, and \mathcal{F} is the degradation operator. Three main effects are taken into account in this operator: *i*) the discrete sampling into a finite number of pixels, namely an image stamp of $n_y \times n_y$ pixels; *ii*) a sub-pixel shift that depends on where the centroid of the image is placed with respect to the pixel grid; and *iii*) a downsampling that affects the pixels in the stamp by a factor of D leaving a $D n_y \times D n_y$ stamp. For example, to handle the *Euclid* mission sampling rate [Cro+13], a factor $D = 1/3$ is required to achieve a Nyquist sampling rate on all of *Euclid*’s VIS passband. In this chapter we will focus on the CFIS survey that is Nyquist sampled, so we will use a unitary value for D .

We write each of these stamps into a one-dimensional column vector and, therefore, $Y_k = [\mathbf{y}_{k,1} \cdots \mathbf{y}_{k,n_{\text{star}}^k}]$ is the matrix containing all the observed stamps in CCD’s k . It contains n_{star}^k columns and $D n_y \times D n_y$ rows. Finally, we concatenate all CCD matrices and obtain $Y = (Y_1 \cdots Y_K)$.

3.3 Resolved component analysis (RCA)

The RCA method is based on a matrix factorisation scheme. It was first presented in Ngolè et al. [Ngo+16] and later evaluated on *Euclid* image simulations in Schmitz et al. [Sch+20]. As with PSFEx, described in detail in section 2.4, this method also builds independent models for each CCD within an exposure and is able to handle under-sampled images. Any observed star i from CCD k is modelled as a linear combination of PSF features, called eigenPSFs in the following, as

$$\hat{\mathbf{h}}_{k,i}^{RCA} = S_k \mathbf{a}_{k,i}, \quad (3.2)$$

where S_k is the matrix composed of the eigenPSFs, $\mathbf{a}_{k,i}$ a vector containing the set of linear weights, and $\hat{\mathbf{h}}_{k,i}^{RCA}$ is the reconstructed PSF.

The modelling is recast into an optimisation problem where the S_k and A_k matrices are estimated simultaneously. The problem is ill-posed due to the undersampling and the noise, meaning that many PSF fields can reproduce the observed stars. In order to break this degeneracy, the RCA uses a series of regularisers during the optimisation procedure to enforce certain mild assumptions on the PSF field: *(i)* the low-rankness of the solution, enforced by setting the number of eigenPSFs learned, N , to be small; *(ii)* the positivity of the reconstructed PSFs; *(iii)* the sparsity of the PSF representation on an appropriate basis; and *(iv)* the spatial constraints that account for imposing a certain structure within the A_k matrix. This last constraint is imposed by a further factorisation of A_k into $\alpha_k V_k^T$. The computation of the V_k^T matrix is addressed in subsection 3.4.4. Finally, the PSF model reads:

$$\hat{H}_k^{RCA} = S_k \alpha_k V_k^T, \quad (3.3)$$

and the optimisation problem that the **RCA** method solves is:

$$\min_{S_k, \alpha_k} \left\{ \frac{1}{2} \|Y_k - \mathcal{F}(S_k \alpha_k V_k^T)\|_F^2 + \sum_{i=1}^N \|w_{k,i} \odot \Phi_{\mathbf{s}_{k,i}}\|_1 + \iota_+(S_k \alpha_k V_k^T) + \iota_\Omega(\alpha_k) \right\}, \quad (3.4)$$

where $w_{k,i}$ are weights, Φ represents a transformation allowing the eigenPSFs to have a sparse representation, \odot denotes the Hadamard product, ι_+ is the indicator function of the positive orthant, and ι_Ω is the indicator function over a set Ω . This last set is defined as a set of sparse vectors and is used to enforce the spatial constraints.

The **PSF** recovery at a position u_j is carried out by a **Radial Basis Function** (**RBF**) interpolation of the learned columns of the A_k matrix, issuing a vector, $\hat{\mathbf{a}}_{k,j}$. In this way, the spatial constraints encoded in the A_k matrix are preserved when estimating the **PSF** at galaxy positions. Finally, the reconstructed **PSF** is:

$$\hat{\mathbf{h}}_{k,j}^{RCA} = S_k \hat{\mathbf{a}}_{k,j}. \quad (3.5)$$

3.4 A new family of MCCD methods

The **MCCD** methods we propose in this study are aimed at exploiting all the information available in a single exposure, which requires the handling of all **CCDs** simultaneously. The main advantage of this approach is that we can build a more complex model since the number of stars available for training is considerably larger than the ones available in a model based on individual **CCDs**. We aim to construct a model that is capable of capturing **PSF** features following a global behaviour, in spite of the fact that the **PSF** field is discontinuous at **CCD** boundaries. The main reason behind this discontinuity effect is in the misalignment between various **CCDs**. Methods such as **PSFEx** or **RCA**, which process each **CCD** independently, avoid the discontinuity problem through construction, but have difficulties capturing global patterns of **PSF** variability that occur on scales larger than a single **CCD**. The main idea behind our **MCCD** approach is to include both a global model that provides a baseline estimation of the **PSF** and a local model that provides **CCD**-specific corrections.

3.4.1 The MCCD data model

In a typical wide-field setting, the **PSF** field, \mathcal{H} , exhibits a certain regularity that we translate into spatial correlations of the **PSFs**. The model we build for a specific **CCD** k is the matrix $\hat{H}_k \in \mathbb{R}^{n_y^2 \times n_{star}^k}$, which is composed by the concatenation of the estimations of the different stars encountered in that **CCD**. Each postage stamp column of length, n_y^2 , corresponds to the model for a specific flattened star from the n_{star}^k stars present in **CCD** k .

The **PSF** field at star positions is reconstructed as a linear combination of **PSF** features, called *eigenPSFs*, which are learned from the observations. As previously stated, we want to have both a global and a local component for the model, so we need different eigenPSFs for each component. Hence, the model is based on a matrix factorisation scheme as follows

$$\hat{H}_k = \underbrace{S_k A_k}_{\text{Local: } \hat{H}_k^{\text{Loc}}} + \underbrace{\tilde{S} \tilde{A}_k}_{\text{Global: } \hat{H}_k^{\text{Glob}}}, \quad (3.6)$$

where $S_k \in \mathbb{R}^{n_y^2 \times r_k}$ contains r_k local eigenPSFs and $\tilde{S} \in \mathbb{R}^{n_y^2 \times \tilde{r}}$ contains \tilde{r} global eigenPSFs. The matrices, $A_k \in \mathbb{R}^{r_k \times n_{star}^k}$ and $\tilde{A}_k \in \mathbb{R}^{\tilde{r} \times n_{star}^k}$, correspond to the local

and global weights of the linear combinations, respectively. We can see that for a given CCD k , the final model, \hat{H}_k , is made up of the sum of the contributions of the local, \hat{H}_k^{Loc} , and global, \hat{H}_k^{Glob} , models.

Now, let us build a single model for all the K CCDs in the focal plane. We start by building a single matrix containing all the PSF models by concatenating the model \hat{H}_k for each CCD as follows

$$\hat{H} = (\hat{H}_1 \quad \cdots \quad \hat{H}_K), \quad (3.7)$$

where $\hat{H} \in \mathbb{R}^{n_y^2 \times N}$ and $N = \sum_{k=0}^K n_{star}^k$ is the total number of stars in one camera exposure. Then we can concatenate the different eigenPSF matrices, k , into a single matrix

$$S = (S_1 \quad \cdots \quad S_K \quad \tilde{S}), \quad (3.8)$$

where $S \in \mathbb{R}^{n_y^2 \times r}$ and we concatenated the global eigenPSF matrix, \tilde{S} , at the end. This leaves a total of $r = \sum_{k=1}^K r_k + \tilde{r}$ columns for the S matrix. We can follow a similar procedure to define A as a block matrix

$$A = \begin{pmatrix} A_1 & 0 & \cdots & 0 \\ 0 & A_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_K \\ \tilde{A}_1 & \tilde{A}_2 & \cdots & \tilde{A}_K \end{pmatrix}, \quad (3.9)$$

where $A \in \mathbb{R}^{r \times N}$ and 0 is used for matrices made up of zeros. The last row of the A block matrix is composed by the global model weights \tilde{A}_k . Having already defined the Multi-CCD matrices, \hat{H} , S and A , we can write the final model as

$$\hat{H} = SA, \quad (3.10)$$

where we include all the CCDs. Expanding it leads to a formula, as shown in Equation 3.6, for each CCD.

3.4.2 Inverse problem and regularisation

The estimation of our model, summarised in the matrices S and A of Equation 3.10, is posed as an inverse problem. Given the observation and MCCD data models presented above, this problem amounts to the minimisation of $\|Y - \mathcal{F}(SA)\|_F^2$, where $\|\cdot\|_F$ denotes the Frobenius matrix norm. This problem is ill-posed due to the noise in the observations and to the degradation operator, \mathcal{F} , meaning that there are many PSF models that would match the star observations. In order to break this degeneracy, we enforce several constraints, based on basic knowledge of the PSF field, which we use to regularise our inverse problem. Similarly to the ones exploited in the RCA method [Sch+20], we use the following constraints:

1. *Low rank:* PSF variations can be explained by a small number of eigenPSFs. This constraint can be enforced by the proper choice of two parameters, the number of local, r_k , and global, \tilde{r} , eigenPSFs. These parameters are directly linked with the complexity of the model we are addressing for each given case and its selection naturally depends on the PSF field we will be facing. It is important to allow the model a certain complexity so that it can correctly capture the PSF field's variations, but it should not be much more complex as the model will tend to overfit the noisy observations and, therefore, lose its generalising power for estimating the PSF in galaxy positions.

2. *Positivity*: the reconstructed PSFs \hat{H} should only contain non-negative pixel values.
3. *Sparsity*: the observed PSFs are structured images; a way to promote our model to follow this structured behaviour is to enforce the sparsity of the eigenPSFs in an appropriate basis.
4. *Spatial constraints*: the regularity of the PSF field \mathcal{H} means that the smaller the distance between two PSFs positions u_i, u_j the smaller the difference between their representations should be $\mathcal{H}(u_i), \mathcal{H}(u_j)$. This regularity can be achieved by enforcing constraint in the coefficient matrices A_k, \tilde{A}_k ; for example, the line l of A_k corresponds to the contribution of eigenPSF l to the n_{star}^k stars in CCD k located in positions $(u_i)_{i=1}^{n_{star}^k}$. The closer the positions, the closer the coefficient values should be.

These constraints are used by both parts of our model, namely, the global and the local components. As mentioned above, the spatial constraint is enforced by further factorisation of the coefficient matrices, A . However, since we want to enforce different properties for the global and the local contributions, the factorisation used differs for each case.

3.4.3 Global model

We want the global component to provide a baseline estimation of the PSF and for that we propose that the coefficients follow a polynomial variation of the position. The global coefficient matrix, \tilde{A}_k , is factorised into $\tilde{A}_k = \tilde{\alpha} \Pi_k$, where $\tilde{\alpha} \in \mathbb{R}^{\tilde{r} \times \tilde{r}}$ is a weight matrix and $\Pi_k \in \mathbb{R}^{\tilde{r} \times n_{star}^k}$ contains each considered monomials evaluated at global star positions. The dimension, \tilde{r} , is determined by the maximum allowed degree in the polynomials: for all monomials of degree less than a given d , we have $\tilde{r} = \binom{d+2}{2} = \frac{(d+1)(d+2)}{2}$. For example, for $d = 2$ (i.e. $\tilde{r} = 6$), we have

$$\Pi_k = \begin{pmatrix} 1 & \cdots & 1 \\ x_{k,1} & \cdots & x_{k,n_{star}^k} \\ y_{k,1} & \cdots & y_{k,n_{star}^k} \\ x_{k,1}^2 & \cdots & x_{k,n_{star}^k}^2 \\ x_{k,1}y_{k,1} & \cdots & x_{k,n_{star}^k}y_{k,n_{star}^k} \\ y_{k,1}^2 & \cdots & y_{k,n_{star}^k}^2 \end{pmatrix}, \quad (3.11)$$

where $(x_{k,i}, y_{k,i})_{1 \leq i \leq n_{star}^k}$ are the global pixel coordinates of the observed stars distributed in the k^{th} CCD. The global component of the model for a specific CCD k are as follows

$$\hat{H}_k^{\text{Glob}} = \tilde{S} \tilde{\alpha} \Pi_k. \quad (3.12)$$

It is important to mention that despite our choice, throughout this chapter, of using position polynomials for building the global space constraint, the model is not necessarily restricted to that choice. The Π_k matrix could be constructed using other parameters of the observations in order to facilitate the capture of other dependencies and could also follow other types of functions.

3.4.4 Local model

It is possible to define different types of local models. In this paper, we discuss three options that depend on how we enforce the local spatial constraint. More specifically, they depend on how we factorise the local A_k matrix in the relation:

$$\hat{H}_k^{\text{Loc}} = S_k A_k. \quad (3.13)$$

Nevertheless, the **MCCD** framework does not restrict us to these three options and it is possible to define other local models. It is worth remarking that all the framework and optimisation procedures are maintained throughout the different flavours of the **MCCD** algorithms. The main difference is the way the spatial constraints are enforced in the local and global models.

MCCD-RCA

One motivation for the local model is to provide **CCD**-specific corrections and to do so, our first choice is **RCA**'s spatial constraint strategy which leads to the **MCCD-RCA** algorithm. The motivation for this choice is the capability of the **RCA** spatial constraint to handle different types of **PSF** variations. On the one hand, it can capture smooth variations over the **CCD** and on the other hand, it can account for localised changes that affect a reduced number of **PSFs**. If the **PSFs** were sampled on a regular grid, this would mean capturing variations occurring at different spatial frequencies. Unfortunately, the **PSF** locations do not coincide with a regular grid but on what could be seen as a fully connected undirected weighted graph where the weights can be defined as a function of the distance between the different nodes (**PSF** locations)². However, the **RCA** spatial constraint exploits the graph harmonics in order to capture the different **PSF** variations. These harmonics are represented by the eigenvectors of the graph's Laplacian matrix [Chug7], which depend on how we define the graph's weights. A parametric function of the **PSF** distances serves as the graph's weights, as in Schmitz et al. [Sch+20], and the selection of the function's parameters can be done following Ngolè et al. [Ngo+16]. For each local model, i.e. each **CCD** in the mosaic, we define r_k graphs, each corresponding to one of the r_k local parameters of the parametrized distance function of the graph's weights. For each graph, we can extract the m_k most useful eigenvectors of its Laplacian matrix, which correspond to the eigenvectors of the lowest eigenvalues. We then present all of them as the columns of the matrix $V_k^{\text{RCA}} \in \mathbb{R}^{n_{star}^k \times r_k m_k}$. In this way, we can write

$$A_k^{\text{RCA}} = \alpha_k^{\text{RCA}} V_k^{\text{RCA}\top}, \quad (3.14)$$

where $\alpha_k^{\text{RCA}} \in \mathbb{R}^{r_k \times r_k m_k}$ is a weight matrix that is used to enforce the spatial constraints. In other words, the sparsity of A_k^{RCA} 's rows in the dictionary $V_k^{\text{RCA}\top}$. Full details are available in Ngolè et al. [Ngo+16] and Schmitz et al. [Sch+20].

MCCD-POL

The second local model, referred to as **MCCD-POL**, follows a polynomial spatial constraint. Similar to **PSFEx**, we factorise the the local weights into two matrices as follows

$$A_k^{\text{POL}} = \alpha_k^{\text{POL}} \Pi_k^{\text{POL}}, \quad (3.15)$$

²A graph G can be defined as a pair (V, E) , where V is the set of vertices and E the set of edges that connects the different vertices. In our case, each star position constitutes a vertex and there is one edge for each pair of vertices. The edges have no preferred direction and its value depends on the distance between the two vertices it connects.

where Π_k^{POL} has the same form as the matrix in Equation 3.11, with the difference that in this case, the positions are represented in local coordinates of its corresponding CCD k . As with d in the global model, a parameter is chosen to define the maximum order of the polynomial used.

MCCD-HYB

The third option consists of using the two local models we presented above, namely, the RCA and polynomial, to work together in an hybrid algorithm we will refer to as MCCD-HYB. The idea behind it is that the addition of the polynomial space constraint could help the original graph constraint to capture the different features found. In this case, we factorise the local weights with block matrices as

$$A_k^{\text{HYB}} = \alpha_k^{\text{HYB}} V_k^{\top \text{HYB}} = \begin{pmatrix} \alpha_k^{\text{RCA}} & 0 \\ 0 & \alpha_k^{\text{POL}} \end{pmatrix} \begin{pmatrix} V_k^{\top} \\ \Pi_k^{\text{POL}} \end{pmatrix}, \quad (3.16)$$

where α_k^{POL} and Π_k^{POL} are the matrices defined in the polynomial version and α_k^{RCA} and V_k^{\top} are the matrices defined in the original MCCD-RCA algorithm.

Finally, generically including the spatial constraints in Equation 3.6, we get the following description of our model for a specific CCD

$$\hat{H}_k = S_k \alpha_k V_k^{\top} + \tilde{S} \tilde{\alpha} \Pi_k, \quad (3.17)$$

which we can also write in a global form, $\hat{H} = S \alpha V^{\top}$, where \hat{H} and S have already been defined in Equation 3.7 and Equation 3.8, and where α and V^{\top} are the following matrices

$$\alpha = \begin{pmatrix} \alpha_1 & 0 & \cdots & 0 \\ 0 & \ddots & & \vdots \\ \vdots & & \alpha_N & 0 \\ 0 & \cdots & 0 & \tilde{\alpha} \end{pmatrix}, \quad V^{\top} = \begin{pmatrix} V_1^{\top} & 0 & \cdots & 0 \\ 0 & V_2^{\top} & & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \cdots & 0 & V_N^{\top} \\ \Pi_1 & \Pi_2 & \cdots & \Pi_N \end{pmatrix}. \quad (3.18)$$

3.4.5 Optimisation problem

Combining the regularisations enumerated in Section 3.4.2 and the data model described in Section 3.4.1, we can construct the optimisation problem in an elegant way by reformulating Equation (3.4). However, we can split the optimisation problem into a more convenient form

$$\min_{\substack{S_1, \dots, S_N, \tilde{S} \\ \alpha_1, \dots, \alpha_N, \tilde{\alpha}}} \left\{ \sum_{k=1}^N \left(\frac{1}{2} \|Y_k - \mathcal{F}_k(S_k \alpha_k V_k^{\top} + \tilde{S} \tilde{\alpha} \Pi_k)\|_F^2 + \sum_{i=1}^{r_k} \|\mathbf{w}_{k,i} \odot \Phi \mathbf{s}_{k,i}\|_1 + \iota_+(S_k \alpha_k V_k^{\top} + \tilde{S} \tilde{\alpha} \Pi_k) + \iota_{\Omega_k}(\alpha_k) \right) + \sum_{i=1}^{\tilde{r}} \|\tilde{\mathbf{w}}_i \odot \Phi \tilde{\mathbf{s}}_i\|_1 + \iota_{\tilde{\Omega}}(\tilde{\alpha}) \right\}. \quad (3.19)$$

In the previous equation, the columns of $Y_k \in \mathbb{R}^{D^2 n_y^2 \times n_{star}^k}$ are the stars distributed in the k^{th} CCD sensor, \mathcal{F}_k is the degradation operator, $\mathbf{w}_{k,i}$ and $\tilde{\mathbf{w}}_i$ are weight vectors, Φ is a transform that allows a sparse representation of our eigenPSFs, and Ω_k and $\tilde{\Omega}$

are sets to enforce sparsity and normalisation of the rows of α_k and $\tilde{\alpha}$, respectively. The indicator function of a set \mathcal{C} is written as $\iota_{\mathcal{C}}(\cdot)$, that is equal to 0 if the argument belongs to \mathcal{C} and $+\infty$ otherwise. For example, ι_+ is the indicator function over the positive orthant. More explicitly, the sets Ω_k and $\tilde{\Omega}$ are defined the following way

$$\Omega_k = \{\alpha_k \mid \forall i \in \{1, \dots, r_k\}, \|(\alpha_k^\top)_i\|_0 \leq \eta_{k,i} \wedge \|(\alpha_k V_k^\top)_i\|_2 = 1\}, \quad (3.20)$$

$$\tilde{\Omega} = \{\tilde{\alpha} \mid \forall i \in \{1, \dots, \tilde{r}\}, \|(\tilde{\alpha}^\top)_i\|_0 \leq \tilde{\eta}_i \wedge \|(\tilde{\alpha} \Pi_k)_i\|_2 = 1\}, \quad (3.21)$$

where $(\eta_{k,i})_{1 \leq i \leq r_k}$ and $(\tilde{\eta}_i)_{1 \leq i \leq \tilde{r}}$ are appropriately chosen integers, and $\|\cdot\|_0$ is the pseudo-norm ℓ_0 that returns the number of non-zero elements of a vector. So we are enforcing, in the global case, the row $i \in \{1, \dots, \tilde{r}\}$ of $\tilde{\alpha}$ to have at most $\tilde{\eta}_i$ non-zero elements. An interpretation could be that we are forcing each eigenPSF to follow a small number of positional polynomials as \tilde{A} 's rows will be sparsely represented over the Π_k matrices.

The Φ transform used throughout this chapter is the starlet transform [SMB11]. We enforce the sparsity on the different decomposition levels excluding the coarse scale. The ℓ_1 term promotes the sparsity of the eigenPSFs with respect to Φ while the weights $\mathbf{w}_{k,i}$ and $\tilde{\mathbf{w}}_i$ regulate the sparsity penalisation against the other constraints and should be adapted throughout the optimisation algorithm depending on the noise level.

The second term in each of the Ω sets (e.g. $\|(\alpha_k V_k^\top)_i\|_2 = 1$) was not mentioned in the regularisation Section 3.4.2, but they are needed to avoid a degenerated solution, for example $\|S_k\|_F \rightarrow \infty$ and $\|A_k\|_F \rightarrow 0$, due to the usual scale indeterminacy when doing a matrix factorisation. To avoid this, we normalise the A_k and \tilde{A} columns. This translates to forcing the normalisation of the eigenPSF weights contributing to model each observed star. This does not mean that the eigenPSF weights will be the same for each star, but that the norms of the weight vectors are equal.

3.4.6 Algorithm

The optimisation in Equation 3.19 is non-convex as we are facing a matrix factorisation problem. To overcome this situation we use an alternating minimisation scheme where we optimise one variable at a time, iterating over the variables as studied in Xu et al. [XY13] and Bolte et al. [BST14]. In consequence, we can at most expect them to converge towards a local minima. The main iteration is performed over the different variables occurring in Equation 3.19, first over the global $\tilde{S}, \tilde{\alpha}$ and then over the local $S_1, \alpha_1, \dots, S_K, \alpha_K$.

The method is shown in Algorithm 1, which contains the four main optimisation problems derived from the alternating scheme. There exists a wide literature on minimisation schemes involving non-smooth terms, specifically proximal methods [PB14], which we can exploit in order to handle the four cases. Notably, we use the algorithm proposed by Condat [Con13] for problems (II), (III), and (IV). For problem (I), we use the method proposed by Liang et al. [LLS18] which is an extension of the well-known FISTA algorithm [BT09]. Even though the ℓ_0 pseudo-norm is non-convex and, therefore, not adapted to the general scenario of the aforementioned algorithms, we can alleviate this issue by combining the application of its proximal operator and a given heuristic.

With regard to the Algorithm 1's initialisation, we start by a preprocessing where we reject stars that are strong outliers in terms of shape or size. We run the shape measurement algorithm mentioned in subsection 3.5.4 on the training stars and discard those that are several sigmas away from nearby stars. At this moment, we can assign a specific weight for each training star. There are three available options:

(i) to use a unitary weight for each training star; (ii) to use a weight provided by the user; (iii) to compute a weight ω_i as a function of the star's SNR based on $\omega_i \propto \text{SNR}_i / (\text{SNR}_i + \text{median}(\text{SNR}))$ and bound to a specific interval to avoid bright stars from dominating the optimisation.

Next, we continue with all the local eigenPSFs set to zero, as seen in line 4 of Algorithm 1; and the $\tilde{\alpha}$ matrix set to the identity, favouring the specialisation of each global eigenPSF to one specific monomial. By following this procedure, we are training a global polynomial model that fits the stars as best as it can. Later on, the local models work with the residuals between the observed stars and the global model, trying to capture variations missed in the previous step.

There are four iteration loops in algorithm 1. In line 8, this is the main iteration, and in line 15, the iteration over the CCDs for the training of the local model. The other two iterations on lines 9 and 14 correspond to a refinement of the estimation. Our objective is to correctly estimate the global and the local contributions for the model and to do this, we alternate the minimisation between the global and the local contributions, which we call outer minimisation. On top of that, each of these two contributions include an inner alternating minimisation scheme as we are performing a matrix factorisation for the local and for the global models. For example, we are simultaneously minimising over S_k, α_k for the local model and over $\tilde{S}, \tilde{\alpha}$ for the global model. We want to refine this inner minimisation, meaning that the optimisation of the two variables separately approaches the joint optimisation of both variables. To accomplish this, we need to go through a small number of iterations, which are described by the n superscript variables, before continuing the iteration of the next alternating scheme. The optimisation strategy can be seen as a compound alternating minimisation scheme considering the outer and the inner alternations.

More details about the optimisation strategy, and how we solve each of the optimization problems appearing in algorithm 1, can be found in section A.

3.4.7 PSF recovery

Once the training of the model on the observed stars is complete, we can continue with the problem of estimating the PSF field at galaxy positions. We call this problem PSF recovery. Gentile et al. [GCM13] conducted a study on PSF interpolation techniques and Ngolè et al. [NS17] proposed a sophisticated approach based on optimal transport theory [PC19]. We will follow a Radial Basis Function (RBF) interpolation scheme with a thin plate kernel³, as in Schmitz et al. [Sch+20], due to its simplicity and good performance. This choice comes with the assumption that the influence of each observation does not depend on the direction but only on the distance to the target which is well described by the RBF kernel.

The RBF interpolation of a function f on a position u works by building a weighted linear combination of RBF kernels ($\phi(\cdot)$) centred in each of the available training star positions u_i . The interpolation function reads

$$\hat{f}(u) = \sum_{i=1}^{N_{\text{RBF}}} \lambda_i \phi(\|u - u_i\|), \quad (3.22)$$

where $(\lambda_i)_{i=1}^{N_{\text{RBF}}}$ are the linear weights that need to be learnt and N_{RBF} is the number of elements used to estimate the interpolant. In order to learn the weights, we force the exact reconstruction of the interpolant on the known positions, that is $\hat{f}(u_i) = f(u_i) \forall i \in \{1, \dots, N_{\text{RBF}}\}$. By fixing the aforementioned constraint we have

³Where the kernel is defined as $\phi(r) = r^2 \ln(r)$.

Algorithm 1 Multi-CCD Resolved Components Analysis

Initialisation:

- 1: Preprocessing()
- 2: **for** $k = 1$ to K **do**
- 3: Harmonic constraint parameters $(e_{k,i}, a_{k,i})_{1 \leq i \leq r_k} \rightarrow V_k^\top, \alpha_k^{(0,0)}$
- 4: $0_{n_y^2 \times r_k} \rightarrow S_k^{(0,0)}$
- 5: **end for**
- 6: Global coordinates $\rightarrow \Pi_k, \tilde{\alpha}^{(0,0)}$ ($\tilde{\alpha}^{(0,0)} = I$)
- 7: $0_{n_y^2 \times \tilde{r}} \rightarrow \tilde{S}^{(0,0)}$

Alternate minimisation:

- 8: **for** $l = 0$ to l_{max} **do** Algorithm's main iterations
 - 9: **for** $n = 0$ to n_G **do** Global alternating iterations
 - 10: Noise level, $\tilde{\alpha}^{(l,n)} \rightarrow$ update $\tilde{W}^{(l,n)}$
 - 11: $\tilde{S}^{(l+1,n+1)} = \arg \min_{\tilde{S}} \{ \sum_{k=1}^K \frac{1}{2} \| Y_k - \mathcal{F}_k(S_k^{(l,0)} \alpha_k^{(l,0)} V_k^\top + \tilde{S} \tilde{\alpha}^{(l,n)} \Pi_k) \|_F^2$
 $+ \sum_i \| \tilde{\mathbf{w}}_i^{(l,n)} \odot \Phi \tilde{\mathbf{s}}_i \|_1 \}$ (I)
 - 12: $\tilde{\alpha}^{(l+1,n+1)} = \arg \min_{\tilde{\alpha}} \{ \sum_{k=1}^K \frac{1}{2} \| Y_k - \mathcal{F}_k(S_k^{(l,0)} \alpha_k^{(l,0)} V_k^\top + \tilde{S}^{(l+1,n+1)} \tilde{\alpha} \Pi_k) \|_F^2$
 $+ \iota_{\tilde{\Omega}}(\tilde{\alpha}) \}$ (II)
 - 13: **end for**
 - 14: **for** $n = 0$ to n_L **do** Local alternating iterations
 - 15: **for** $k = 1$ to K **do** CCD iterations
 - 16: Noise level, $\alpha_k^{(l,n)} \rightarrow$ update $W_k^{(l,n)}$
 - 17: $S_k^{(l+1,n+1)} = \arg \min_{S_k} \{ \frac{1}{2} \| Y_k - \mathcal{F}_k(S_k \alpha_k^{(l,n)} V_k^\top$
 $+ \tilde{S}^{(l+1,n_G)} \tilde{\alpha}^{(l+1,n_G)} \Pi_k) \|_F^2 + \sum_i \| \mathbf{w}_{k,i}^{(l,n)} \odot \Phi \mathbf{s}_{k,i} \|_1$
 $+ \iota_+(S_k \alpha_k^{(l,n)} V_k^\top + \tilde{S}^{(l+1,n_G)} \tilde{\alpha}^{(l+1,n_G)} \Pi_k) \}$ (III)
 - 18: $\alpha_k^{(l+1,n+1)} = \arg \min_{\alpha_k} \{ \frac{1}{2} \| Y_k - \mathcal{F}_k(S_k^{(l+1,n+1)} \alpha_k V_k^\top$
 $+ \tilde{S}^{(l+1,n_G)} \tilde{\alpha}^{(l+1,n_G)} \Pi_k) \|_F^2 + \iota_{\Omega_k}(\alpha_k) \}$ (IV)
 - 19: **end for**
 - 20: **end for**
 - 21: **end for**
-




Figure 3.4-1 – Example of the interpolation procedure involved in the PSF recovery.

a system of N_{RBF} equations with N_{RBF} unknown that are the λ_i weights. Once the system is solved, it is just a matter of evaluating the interpolant on the desired position u following Equation 3.22.

At this point, we need to choose the function f over which we go on to interpolate. A straightforward choice would be to use the reconstructed PSFs at the training positions as the $f(u_i)$. Nevertheless, this would not take into account the specificities and structure of our model. Following the discussion in Schmitz et al. [Sch+20, §4.2], we use the learnt A_k and \tilde{A}_k matrices. These matrices encompass all the spatial distribution properties of the learned features, that is, our eigenPSFs; thus it is natural for our framework to use these values as the function to interpolate.

We continue with a brief explanation of the interpolation procedure. For one given target position u in CCD k , we consider the N_{RBF} closest observed stars to that position that also belong to the CCD k . We call $A_k(N_{\text{RBF}})$ to the A_k matrix composed only with the columns of the aforementioned N_{RBF} stars. We want to estimate the interpolated column vector $A_{k,u}$. For this, we use a RBF interpolation scheme for each row of the $A_k(N_{\text{RBF}})$ matrix. The elements of the row t represent the $(f^{(t)}(u_i))_{i=1}^{N_{\text{RBF}}}$ evaluations and the element $A_{k,u}^{(t)}$ represents the interpolated value $\hat{f}^{(t)}(u)$. The same procedure is repeated for each row of the $A_k(N_{\text{RBF}})$ matrix so as to obtain the column vector $A_{k,u}$. This is illustrated in Figure 3.4-1. We repeat the procedure with the global component matrix, \tilde{A}_k , in order to obtain $\tilde{A}_{k,u}$, another column vector with the interpolated values. At this point, we note that we handle the global and the local contributions independently. Once we have calculated the two interpolated vectors, the reconstructed PSF is obtained following the MCCD data model as can be seen in the next equation

$$\hat{H}(u) = \tilde{S}\tilde{A}_{k,u} + S_k A_{k,u}. \quad (3.23)$$

We found that restricting the N_{RBF} neighbours to a single CCD for the global components gives better results. This might be due to the fact that the global components are able to capture some of the discontinuities from one CCD to another and, therefore, the interpolation is degraded when using stars from different CCDs. The number of neighbours N_{RBF} should be chosen as a function of the available




Figure 3.5-2 – Star positions in CFHT’s MegaCam used for the simulated dataset. The positions were taken from a real CFIS exposure.

number of stars per CCD in the training set and as the RBF kernel chosen. Henceforth, and given the training set we handle in this study, N_{RBF} is set to 20.

3.5 Numerical experiments with simulated data

3.5.1 Data

The simulated data set we create to evaluate MCCD set is based on a Canada-France Imaging Survey (CFIS)⁴ MegaCam⁵ exposure from the Canada-France-Hawaii Telescope (CFHT). It contains 2401 stars distributed along 40 CCDs over a field of view of $\sim 1 \text{ deg}^2$ as shown in Figure 3.5-2. Each CCD consists of a matrix of 2048 by 4612 pixels with some given gaps between the different CCDs. The horizontal gap length consist of ~ 70 pixels while vertical gaps contain ~ 425 pixels.

3.5.2 Training set

Our simulation pipeline considers a Moffat PSF profile with normalised flux drawn using the Galsim software⁶ [Row+15] for each position in the exposure. To simulate the PSF shape variation, we used two radial analytic functions which define our ground truth shape ellipticities distortions. Shearing stars leads naturally to a size variation. Figure 3.5-3 shows the resulting e_1 , e_2 and size maps. Our pipeline performs the following steps:

1. Simulate Moffat stars with a size fixed to the mean size measured in the real exposure.
2. Shear the simulated stars as a function of their position using the two analytical functions.
3. Apply a random sub-pixel shift following a uniform distribution centred on zero.

⁴www.cfht.hawaii.edu/Science/CFIS/

⁵www.cfht.hawaii.edu/Instruments/Imaging/MegaPrime/

⁶github.com/GalSim-developers/GalSim

4. Apply a binning to get a 51×51 pixel image, with a pixel size equivalent to CFIS MegaCam's maps, that is, 0.187 arcsec.
5. Add a constant white Gaussian noise to the images, with standard deviation σ , derived from the desired SNR level

$$\text{SNR} = \frac{\|y\|_2^2}{\sigma^2 p^2}, \quad (3.24)$$

where y is the image postage stamp consisting of p^2 pixels. Each experience will consist of a constant SNR value, as we later see, which is drawn from the set $\{10, 30, 50, 70\}$.

Since PSFEx was designed as a companion software to SExtractor, we need to follow a different procedure to generate the simulated data. We first need to process our simulations with SExtractor, so that the catalogue produced can be used as inputs for PSFEx. To accomplish this, we mimic a complete CCD so that SExtractor is able to process it. We create star images as we already described for the MCCD method but without noise as it will be added later. Then we distribute them on a mock image of 2048×4612 pixels. The corresponding positions will be the pixel coordinates that are presented in Figure 3.5-2. Once the mock image is created, we add the noise value according to the desired SNR to the whole image. When the mock image is created, we run SExtractor in order to have a star catalogue that PSFEx can use as input.

3.5.3 Testing data set

For the testing, we want to observe how well the different models capture the ellipticity maps when trained on real star positions. Therefore, the positions in each CCD are taken from a regular grid of 20×40 and considering that the total amount of CCDs is 40, we finally obtain a total of 32000 stars with which to test our model. These stars are simulated following the same ellipticity maps (see Figure 3.5-3), without any sub-pixel shift and without any noise. The goal is now to use the training data (i.e. simulated observed stars) to learn the model, and then to predict the PSFs at positions of test stars. As we have the ground truth at these positions, without noise and sub-pixel shift, it is easy to get a robust evaluation of model predictions.

3.5.4 Quality criteria

In order to correctly assess the performance of our PSF modelling algorithm, we consider several criteria

- Pixel Root Mean Squared Error (RMSE) : calculated between the pixel images of the recovered PSFs and the noiseless test stars. The expression of the pixel RMSE is the following

$$Y_{pix}^{RMS} = \sqrt{\langle (Y^* - \hat{Y})^2 \rangle} = \left(\frac{1}{N n_y^2} \sum_{i=1}^N \sum_{j=1}^{n_y^2} (Y_{i,j} - \hat{Y}_{i,j})^2 \right)^{\frac{1}{2}}, \quad (3.25)$$

where $Y_{i,j}$ is the pixel j of test star i that has a total of n_y^2 pixels, N is the total number of test stars, $\hat{Y}_{i,j}$ is the estimation of the test star's pixel and $\langle \cdot \rangle$ denotes the mean over all the elements in the array.

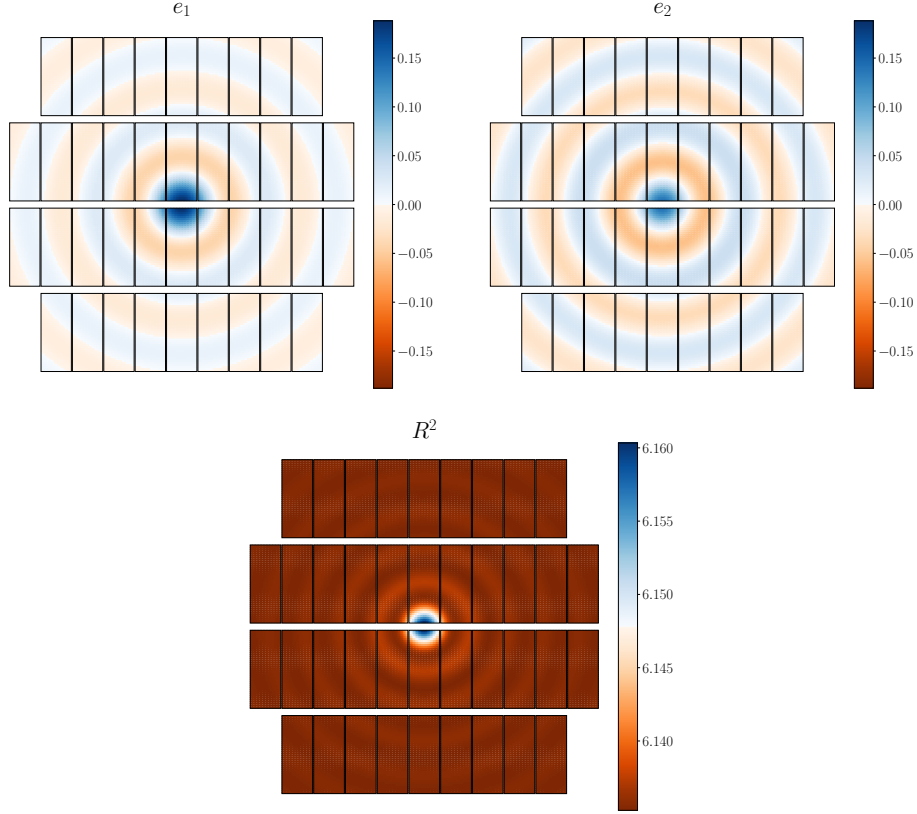


Figure 3.5-3 – Shape measurement results of the simulated test star catalogue following the analytical ellipticities.

- Shape (ellipticity) error: We estimate the ellipticities of reconstructed stars using the adaptive moments' ellipticity estimator from **Galsim**'s **HSM** module [HSo3; Man+05]. The shape and size definitions can be found in subsection 2.6.2. For each of the ellipticity components, the **RMSE** is calculated as

$$e^{RMS} = \sqrt{\langle (e^* - \hat{e})^2 \rangle} = \left(\frac{1}{N} \sum_{i=1}^N (e_i - \hat{e}_i)^2 \right)^{\frac{1}{2}}. \quad (3.26)$$

- Size error: We use the measurements from **HSM** and the definition in subsection 2.6.2 to compute the following **RMSE**

$$R^{2,RMS} = \sqrt{\langle (R^{2,*} - \hat{R}^2)^2 \rangle} = \left(\frac{1}{N} \sum_{i=1}^N (R_i^2 - \hat{R}_i^2)^2 \right)^{\frac{1}{2}}. \quad (3.27)$$

- Moment residual maps: To visualise the shape and size errors, we plot these quantities as a function of their position on the focal plane.

When comparing two methods, we define the relative gain with regard to the metric, m , of method 1 with respect to the method 2 as

$$G_{1/2}(m) = \frac{m_2 - m_1}{m_2} \times 100\%. \quad (3.28)$$

3.5.5 Model parameters

Based on experiments with simulated and real data, we have chosen the following parameters

- **PSFEx**: we use the following configuration:

```
PSF_SAMPLING  1.0
PSF_SIZE      51,51
PSFVAR_KEYS   XWIN_IMAGE,YWIN_IMAGE
PSFVAR_GROUPS 1,1
PSFVAR_DEGREES 2
```

XWIN_IMAGE and YWIN_IMAGE refer to the windowed centroid positions in pixel coordinates, and PSFVAR_DEGREES to the maximum polynomial degree. The **PSFEx** software⁷ does not include publicly an interpolation method, so we use an available **PSFEx** interpolation module⁸.

- **RCA**: we set r equal to eight local components, the denoising parameters K_σ^{RCA} to 1, and the other parameters to their default value from its official repository⁹.
- **MCCD**: we use the same parameters as **RCA** for the local component and a maximum polynomial degree of eight for the global components. The denoising parameters K_σ^{Loc} and K_σ^{Glob} are set to 1 for the local and the global contributions.

The **MCCD** parameters that most affect its behaviour are mentioned above. Their choice greatly relies on the training data set used. Depending on the number of stars available and the complexity of the instrument's **PSF**, it may be preferable to adopt a more complex model by augmenting the number of local components, r , and the maximum polynomial degree. However, if the stars are not enough to constrain the model, we may end with a model that overfits the training stars. A proper selection of the denoising parameters can control the bias-variance trade-off in the estimation. A high value of the denoising parameter, namely, 3, leads to an extremely denoised model. It will contain a high estimation bias that can be related with a model that cannot capture some spatial variations and fine details of the **PSF**. On the contrary, if the denoising parameter is close to zero, the only denoising performed by the **MCCD** is due to the low-rank constraint and therefore the estimations can be rather noisy.

3.5.6 Results

Comparison between PSFEx, RCA and MCCD-RCA

The first results can be seen in [Figure 3.5-4](#) and [Figure 3.5-5](#), where we compare the **PSFEx**, **RCA**, and **MCCD-RCA** algorithms. We observe that **MCCD-RCA** outperforms the other methods, with an average pixel RMS improvement over **PSFEx** of 51% and ellipticity RMS improvement ranging from 15% for stars with an **SNR** 10 to 36% for a **SNR** of 70. The **RCA** is almost as good as **MCCD-RCA** for the pixel error, but does not provide good results for the other metrics. This behaviour can be explained by the fact that the model strongly deteriorates for some **CCDs**, giving extreme ellipticities and sizes values. These deteriorations of the model are not strong

⁷www.astromatic.net/software/psfex

⁸github.com/esheldon/psfex

⁹github.com/CosmoStat/rca




Figure 3.5-4 – RMSE on pixels, shape, and size metrics as a function of stars SNR for the three main methods. The RMSE are plotted with solid lines and the gain of the methods with respect to PSFEX are plotted with dashed lines.

enough to produce a large pixel error but causes much more significant errors on the moments. We include in Figure 3.5-6 RCA’s R^2 residual map that shows the catastrophic failure in the modelling of some CCDs.

We can see on the right column of the residual maps in Figure 3.5-5 that PSFEX’s ellipticity residuals follow the global pattern from the dataset. This means that the ellipticity is not captured in the model, showing some difficulties found when modelling a global ellipticity pattern using independent models for each CCD. The MCCD-RCA algorithm, which builds up a model for the whole focal plane, does a better job in capturing the global ellipticity pattern. The MCCD-RCA’s residuals are smaller and less correlated with the pattern of the dataset. With regard to the third row of Figure 3.5-5, where the size of the simulated PSFs is practically constant, we observe that the MCCD-RCA has slightly larger errors when the training star density is low, as in the bottom-right corner (see Figure 3.5-2).

Comparison between MCCD-POL, MCCD-RCA, and MCCD-HYB

The comparison between the MCCD-POL, MCCD-RCA, and MCCD-HYB methods is shown in Figure 3.5-7. First, we notice that MCCD-POL presents poor performance in most of the metrics. This indicates that the local polynomial model is not able to capture the PSF variations that are left over from the difference of the global model and the observed stars. Hence, even if MCCD-POL has a lower pixel error than PSFEX (see Figure 3.5-4), it has greater ellipticity errors. Capturing these PSF variations properly is essential for obtaining good ellipticity performances. The MCCD-RCA and MCCD-HYB have similar behaviours, but MCCD-HYB uses a mixed approach

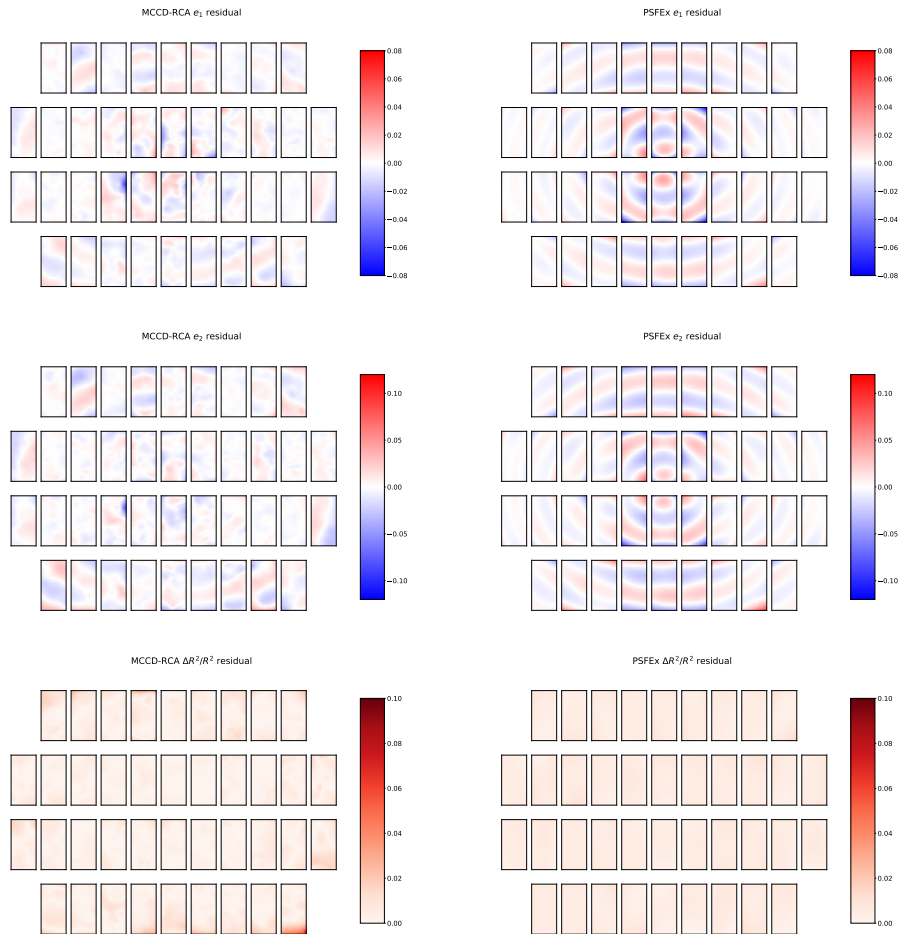


Figure 3.5-5 – Moment residual maps comparing the MCCD-RCA algorithm on the left and the **PSFEx** algorithm on the right. They are obtained by subtracting the model's and the test star's measured shape and size metrics and plotting them on their corresponding position over the focal plane. The **SNR** value of the star dataset is 50.

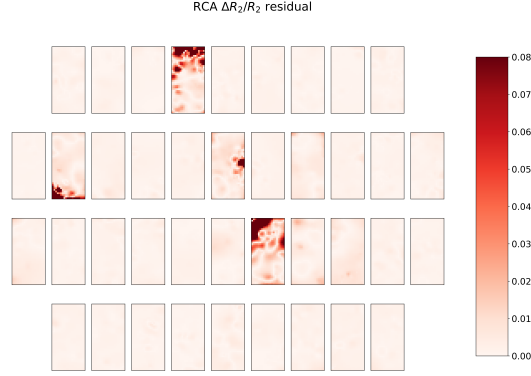


Figure 3.5-6 – Residual $\Delta R_2/R_2$ map of the RCA algorithm of stars with an SNR of 50. The CCDs where the RCA model is having degeneracies that can be clearly spotted on the map.

of a polynomial and graph-based local model outperforms the original MCCD-RCA method in terms of ellipticity components. The average gain in both components of MCCD-HYB with respect to MCCD-RCA is around 18%, proving the utility of using the hybrid approach. This suggests that there are some features related to the PSF shape that can be captured by a simple polynomial model and not by the graph-based model alone. Examples of global and local eigenPSF from the MCCD-HYB model can be seen in Figure 3.5-8.

3.5.7 Comparison of computing resources

The MCCD methods take $\sim 2.9\times$ more CPU-time than PSFEx when compared on the same machine. We evaluate it on the fitting and validation procedures, that is, the estimation of the PSF model and the recovery of PSF at test positions. A relevant detail is that the PSFEx package is coded in the C programming language, while the MCCD methods are completely coded in Python.

3.6 Numerical experiments with UNIONS/CFIS data

In this section, we compare the MCCD-HYB method with PSFEx using real data from the Ultra-violet Near-Infrared Optical Northern Sky (UNIONS) survey, which is a collaboration between the Panoramic Survey Telescope and Rapid Response System (Pan-STARRS) and CFIS. We use the r -band data from the latter.

3.6.1 Dataset

We analysed a subset of around 50 deg^2 from the whole CFIS survey area that, in total, will eventually span 5000 deg^2 . It corresponds to the subset named W₃ described in Erben et al. [Erb+13], and includes 217 exposures. Each CCD from each exposure has been processed independently with SExtractor. The stars were selected in a size-magnitude diagram, in the magnitude range between 18 and 22, and a Full Width Half Maximum (FWHM) range between 0.3 and 1.5 arcsec. In order to validate the PSF models, we randomly split the stars into a testing and a training dataset, trying





Figure 3.5-7 – Comparison of the performance of MCCD-POL, MCCD-RCA, and MCCD-HYB methods in terms of the RMSE on the pixels, the shape and the size metric of the star’s SNR. The RMSE are plotted on a solid line and the gain of the method with respect to **PSFEx** is plotted on dashed lines.


to estimate the first set of stars while constructing our model only with the second. The training dataset is composed of 80% of the detected stars and the test dataset of the remaining 20%. We consider a fixed threshold on the number of training stars per **CCD**, meaning that if the number of training stars in a given **CCD** is less than 30, we discard the **CCD**. The star density of the training dataset is presented in Figure 3.6-9. The ellipticity and the size of the training stars can be seen in Figure 3.6-10. Each bin represents the mean shape measurement over all the stars with a centroid located within the bin.

3.6.2 Model parameters

The setup of **PSFEx** for this experiment is similar to the one used for the simulated images, which can be found in Section 3.5.5. The MCCD-HYB method uses a maximum global polynomial degree of 3, with 16 local components and the denoising parameters K_σ set up to 0.1. In order to compare the star images with the different methods (**PSFEx** and MCCD-HYB), the models need to match the flux as well as the centre of the star. Hence, after estimating a **PSF** model at a given star location, the **PSF** is normalised and shifted to match the star. For this purpose, we use the same intra-pixel shift and flux estimation methods for both **PSF** models: (i) we estimate the star and the **PSF** centroids; (ii) we calculate the shift needed by the **PSF** to match the star and construct a shifting kernel; and (iii) the **PSFs** are convolved by their corresponding shifting-kernel. To match the flux, we calculate an α parameter for each test star and **PSF** that corresponds to the argument that minimises the function,



(a) Global eigenPSFs.



(b) Local eigenPSFs.

Figure 3.5-8 – Example eigenPSFs extracted from the MCCD-HYB PSF model trained on the simulated dataset with an SNR of 70. The local eigenPSFs were extracted from the graph’s spatial constraint of a central CCD. It can be seen from the eigenPSFs that the global model is specialising on the shape of the PSF while the local model specialises on capturing its ellipticity. It is also worth to mention that the first global eigenPSF found on the first row provides the baseline isotropic PSF the model uses.




Figure 3.6-9 – Star density of all the training dataset with respect to their position on the MegaCam’s focal plane. We have on average 1560 training stars per exposure.

$f(\alpha) = \|I_1 - \alpha I_2\|_2$, where I_1 and I_2 are the star and the PSF, respectively.

3.6.3 A new metric for real data: the Q criteria

Performing a comparison between two PSF models with real data is an arduous task since we do not know the shapes and pixel values of the observed stars. However, subtracting our estimated model from an observed star (i.e. pixel residual) should lead to a residual map containing only noise if the model is perfect. The probability of having our model correlated with the noise is negligible. Therefore, from this point of view, the method with the smallest pixel RMS residual error can be considered as the best. Using all the test stars y_s and our estimates \hat{y}_s , we calculate the pixel RMS residual error, $\text{Err} = \sqrt{\frac{1}{N_i N_s} \sum_s \sum_i (y_{s,i} - \hat{y}_{s,i})^2}$, where N_s is the number of stars and N_i is the number of pixels we consider in a given image when we use a 10 pixel radius circle from the centre of the residual images. The noise standard deviation σ_{noise} is calculated from the stars only using the pixels outside the aforementioned circle. For a perfect modelling, we would have $\text{Err} \approx \sigma_{\text{noise}}$, and we define the Q_{p_1} metric as

$$Q_{p_1} = (\text{Err}^2 - \sigma_{\text{noise}}^2)^{1/2}. \quad (3.29)$$

We next introduce two metrics to quantify how noisy the models are. The variance of the PSF model for the test stars s reads $\sigma_s^2 = [\text{Var}(y_s - \hat{y}_s) - \sigma_{\text{noise}}^2(y_s)]_+$, where $\text{Var}(\cdot)$ is a usual variance estimator, the operator $[\cdot]_+$ sets to zero negative values and $\sigma_{\text{noise}}^2(y_s)$ is the noise variance estimation for a single star. We present the Q_{p_2} and Q_{p_3} metrics in the following equations

$$Q_{p_2} = \left(\frac{1}{N_s} \sum_s \sigma_s^2 \right)^{1/2}, \quad Q_{p_3} = \left(\frac{1}{N_s} \sum_s (\sigma_s^2 - Q_{p_2}^2)^2 \right)^{1/4}. \quad (3.30)$$

The Q_{p_2} metric represents the modelling error expectation for a given star and the Q_{p_3} metric indicates the fluctuation of the modelling error. A perfect PSF model would give values close to zero for the three metrics.

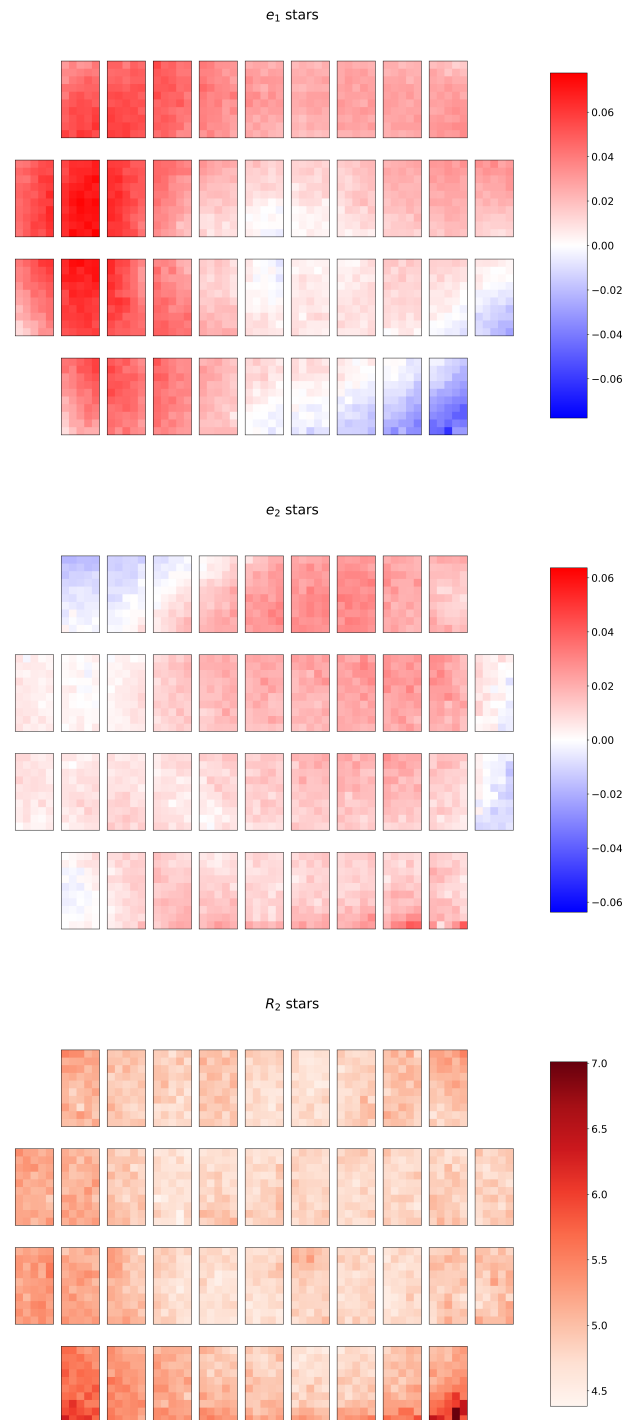


Figure 3.6-10 – Ellipticities and size *meanshapes* maps (see [subsection 2.6.2](#) for more information about the maps) measured on the training stars of our CFIS dataset.

3.6.4 Results

Method	Q_{p_1}	Q_{p_2}	Q_{p_3}
PSFEx	15.56	8.13	14.31
MCCD-HYB	12.14	6.68	10.86
Gain _{PSFEx}	22%	18%	24%
Noise Std. Dev. (σ_{noise})	15.83		

Table 32 – Q_p criterion using all test stars of the W3 dataset from CFIS. The gain of the MCCD-HYB with respect to PSFEx and the noise standard deviation are also presented.

The main results of the experiment are synthesised in Table 32, where the Q_p criteria are given. In the Q_{p_1} column, we can observe a 22% gain of the MCCD-HYB method with respect to PSFEx. From Q_{p_2} and Q_{p_3} metrics, we also conclude that the MCCD-HYB model is considerably less noisy than the one from PSFEx.

In order to explore potential remaining structure in the residuals, we stack together the residuals for all 534 test stars from a random exposure. These are shown, along with the stacking of the test stars themselves, in Figure 3.6-11. We can see that PSFEx has a sharper stacked error compared to MCCD-HYB. This could indicate that our algorithm is better at capturing the size of the PSF, as the peak of the residual is directly related to it. Considering that there is no trace of shifting errors and that we are calculating the flux optimally, a greater mismatch in the size of the PSF equals to a greater peak pixel error on the residual. The third row presents the mean of the stacked absolute value of the residuals for both of the PSF models so that the residuals can not cancel themselves. We observe the same behaviour described above with the PSFEx pixel error distribution being sharper but more centred. It is also possible to notice the higher noise PSFEx has when compared to the MCCD-HYB model.

Figure 3.6-12 presents examples of star image reconstructions by the two different PSF models, PSFEx and MCCD-HYB, and their corresponding residuals. The proposed method yields a near noiseless model when compared to PSFEx, as can clearly be seen on the top-left and bottom-right stars of Figure 3.6-12, where the stars have low SNR of 19.3 and 4.2, respectively. Both models share a good estimation of the bottom-right star, which comes a low-stellar-density region of the focal plane (the bottom-right corner, as can be seen in Figure 3.6-9). On the bottom-left star of Figure 3.6-12, we observe a similar type of error as that appearing in Figure 3.6-11.

It is difficult to derive conclusions of different PSF model performances based on the shape measurement of noisy stars due to its high stochasticity. Nevertheless, driven by the comments from DES Y1 [Zun+18] on the residual mean size offset from the PSFEx model, we conducted a study with our data. We measured the size from the training stars and from both calculated PSF models, PSFEx and MCCD-HYB, and then computed the residual. The RMS residual size of the $\Delta R^2/R^2$ value gave 4.82×10^{-2} for PSFEx and 4.02×10^{-2} for MCCD-HYB. This represents a 16% gain of our proposed algorithm.

Figure 3.6-13 presents in the left column the histogram of the residuals and in the right column the histograms of the size metrics. We notice that the MCCD-HYB




Figure 3.6-11 – Stacked star profile from all 534 test stars in a random [CFIS](#) exposure (id. 2099948) (top), corresponding stacked residuals after subtraction by the [PSFEx](#) (middle left) or [MCCD-HYB](#) (middle right) [PSF](#) models. The bottom row includes the stacked absolute value of the residuals.

algorithm has a sharper residual size around zero. The figure also includes the mean of the residuals for each [PSF](#) model. This shows that both models tend to overestimate the size of the [PSF](#). However, the [MCCD-HYB](#) model presents a 30% gain in the mean residual size with respect to [PSFEx](#), indicating a smaller bias in the shape.

3.7 Conclusion

In this chapter, we presented a family of non-parametric, or data-driven, [Point Spread Function](#) ([PSF](#)) modelling methods coined [MCCD](#), including its best-performing extension [MCCD-HYB](#), which are built upon the existing [Resolved Component Analysis](#) ([RCA](#)) method and are capable of constructing [PSF](#) models that span all the [Charged Coupled Device](#) ([CCD](#)) from an instrument’s focal plane at once. Naturally, the use of more stars for the training allows us to build more complex models that can capture evasive features. Our model is composed of global components, spanning all the [CCDs](#),




Figure 3.6-12 – Examples of real CFIS test stars, the estimations of both methods, and the pixel residuals. The exposures are the same as in Figure 3.6-11. We present four test stars with the estimated PSF models and the corresponding residuals. The top-left star corresponds to a star extracted from the top-left corner of the focal plane with an SNR of 19. The top-right star corresponds to a star leading to a high error for both methods. The bottom-left star corresponds to a star located in the centre of the focal, with a relatively high SNR of 160. The bottom-right star corresponds to a star located in the bottom-right corner of the focal plane, with a low SNR of 4.

and local components that are CCD-specific. By using this structure, we can better capture global patterns and features that might be lost when using only a local model like in RCA or, the widely used algorithm, PSFEx.

The method was first tested with a set of simulated PSFs following a real star spatial distribution over MegaCam’s focal plane, an instrument from the Canada-France-Hawaii Telescope (CFHT). Its use leads to better performance in all the evaluated metrics when compared to PSFEx. We then tested the method on a set of real CFIS images, an imaging survey based on CFHT, in order to confirm that it can handle real data. Our method achieves a smaller pixel root-mean-square residual than PSFEx and the estimated model is considerably less noisy. The proposed method can naturally handle more complex PSF profiles. Thanks to its formulation, it can also handle super-resolution, making it suitable for undersampled data.

The performance gain of the MCCD methods over PSFEx is higher when using our simulated dataset than when using the real dataset. This can be explained by the fact that our simulated dataset shows more intricate variations in the PSF than the real data does and MCCD is better at capturing such strong variations.

In the spirit of reproducible research, the MCCD algorithm is publicly available on GitHub at github.com/CosmoStat/mccd, including the material needed to reproduce the simulated experiences. A webpage with the MCCD code documentation is available at <https://cosmostat.github.io/mccd/about.html>.

In the next chapter we will detail the application of the MCCD PSF model to produce the v1 version for the CFIS WL shape catalogue that spans more than

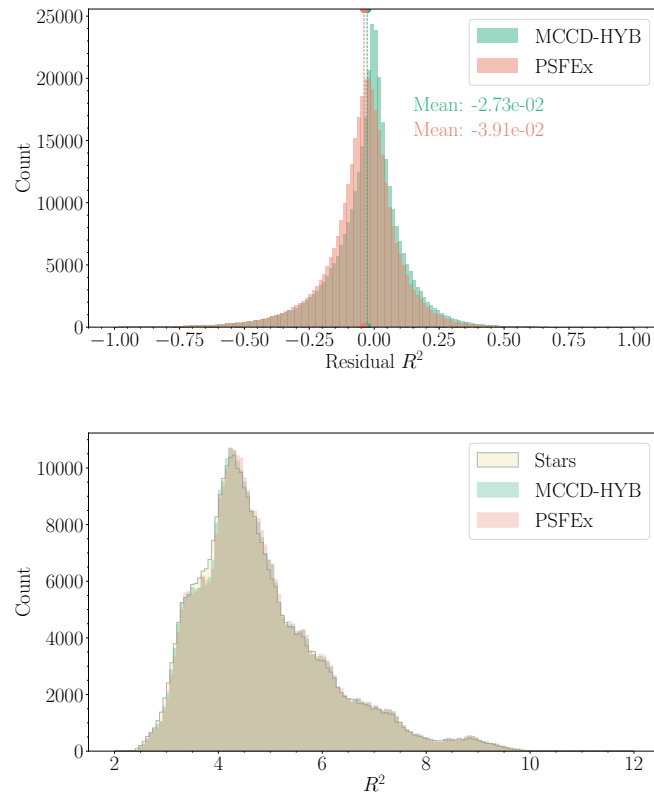


Figure 3.6-13 – Histograms showing the distribution of the size metric over the training stars and their residuals for both PSF models, PSFEx and MCCD-HYB.

3500deg². We will also present PSF validation results using the vast amount of real data available.

* * *
* * *
* * *

Modelling the PSF for UNIONS/CFIS

Chapter Outline

4.1	Introduction	107
4.2	Incorporating MCCD to ShapePipe	109
4.3	Shape catalogue and PSF diagnostics	111
4.3.1	Shape measurement diagnostics	112
4.3.2	PSF diagnostics	112
4.4	Conclusion	116

THIS chapter presents the application of the [MCCD PSF](#) model to the shape measurement pipeline [ShapePipe](#) and its processing of the most recent [CFIS](#) data release. We start with a brief description of the [CFIS](#) survey showing that its characteristics make it a good candidate for a [WL](#) analysis. We then presents the incorporation of the [MCCD PSF](#) model to the high-performance [WL](#) shape measurement pipeline [ShapePipe](#). The [MCCD PSF](#) model was used to produce the v1 [CFIS](#) shape catalogue of more than $\sim 3500 \text{ deg}^2$. We present some diagnostics of the [PSF](#) model performance and the quality of the shape catalogue before ending the chapter with some conclusions.

4.1 Introduction

[Canada-France Imaging Survey \(CFIS\)](#)¹ is a photometric survey at [Canada-France-Hawaii Telescope \(CFHT\)](#)² in the mount Mauna Kea at Hawaii and uses a 3.6 metres high-quality telescope. [CFIS](#) is part of the larger [Ultra-violet Near-Infrared Optical Northern Sky \(UNIONS\)](#)³ collaboration. One of its major objectives is to provide ground-based photometric support for the *Euclid* mission in the Northern Hemisphere. The [CFIS](#) survey targets a total area of $\sim 5000 \text{ deg}^2$. Using the wide-field imager [MegaCam](#)⁴ instrument, [CFIS](#) will observe the sky in the r- and u- bands that were presented in [Figure 2.2-8](#). [Figure 4.1-1](#) shows the sky coverage of the [CFIS](#) survey for the r-band as of the March 2022. The coverage can be compared to the *Euclid* footprint presented in [Figure 1.4-10](#).

¹<https://www.cfht.hawaii.edu/Science/CFIS/>

²<https://www.cfht.hawaii.edu/en/gallery/>

³<https://www.skysurvey.cc>

⁴<https://www.cfht.hawaii.edu/Instruments/Imaging/MegaPrime/>




Figure 4.1-1 – CFIS r-band sky coverage as of March 2022. Credit: UNIONS collaboration.




Figure 4.1-2 – CFIS delivered image quality for the r- and u-band as of March 2022. Credit: UNIONS collaboration.

The CFIS survey has excellent r-band data. Following the five factors from section 1.4 we can see that the data is a great candidate for a Weak Gravitational Lensing study as:

- it is *wide*, covering $\sim 5000\text{deg}^2$, a large part of the Northern Hemisphere that can be seen in Figure 4.1-1,
- it is *deep*, taking 3 exposures of 200 seconds and reaching a planned depth of 24.1 at an SNR of 10 for extended sources observed in the stacked images,
- it has *great image quality*, which corresponds to a small PSF thanks to the privileged CFHT location at 4200 Metres Above Mean Sea Level (MAMSL) and top performance of the MegaCam instrument that can be seen in Figure 4.1-2.

The CFIS collaboration counts around 150 scientists, mostly from Canada and France. Two major teams are working to build two independent WL shape measurement




Figure 4.1-3 – CFIS r-band DR3 used to produce the v1 shape catalogue. The total area is divided into seven patches from P1 to P7. Credit: UNIONS collaboration.

pipelines to exploit the CFIS data. Prof Ludovic Van Waerbeke leads the first one in Canada, and Dr Martin Kilbinger leads the second one in France. The first team is building a pipeline based on the *lensfit* shape measurement and PSF modelling method [Mil+13]. The second team, and the one I am contributing, has built the *ShapePipe* pipeline [Gui+22; Far+22a] which is based on the *ngmix* shape measurement method [She15] using metacalibration [HM17; SH17], a calibration method briefly described in subsection 1.3.7.

The CFIS collaboration made three *internal Data Releases (DRs)* of preprocessed r-band data on, February 2019 (DR1 with $\sim 2000\text{deg}^2$), April 2020 (DR2 with $\sim 2800\text{deg}^2$), and March 2021 (DR3 with $\sim 3600\text{deg}^2$). The first release was used to build a first shape catalogue with *ShapePipe* and has been detailed in Guinot et al. [Gui+22]. However, the catalogue has not been released to the public yet. The two pipelines have processed the DR3 and are currently being validated. A joint article with their description is expected. Figure 1.3-6 shows the area observed in the DR3 release, dividing it into 7 regions, or patches.

The *ShapePipe* version used for the DR1 in Guinot et al. [Gui+22] used the *PSFEx* PSF model. Later, we have integrated the *MCCD* PSF model into the pipeline, validated it and used it to generate the *ShapePipe* v1 catalogue version using the DR3. On what rests of this chapter we will go through the use of *MCCD* in *ShapePipe* and its processing of the CFIS r-band DR3.

4.2 Incorporating MCCD to ShapePipe

The main ingredients of the *ShapePipe* pipeline were already introduced in section 1.3 where Figure 1.3-6 shows the main building blocks. Farrens, S. et al. [Far+22a] presents a more software-related description, while Guinot et al. [Gui+22] a more weak-lensing related description. The code is available online⁵ as well as its documentation⁶.

⁵<https://github.com/CosmoStat/shapepipe>

⁶<https://cosmostat.github.io/shapepipe/>

There are two details that are specific to **ShapePipe** that change with respect to the general description in [Figure 1.3-6](#). First, the detection in **ShapePipe** is done in the stack. This means that from the three exposures that we have from each region of the sky, the images are stacked to produce a single image with a higher SNR. This stacked image is then used to perform the detection of object. Second, the shape measurement is done in a multi-epoch approach using the three exposures and not the stacked image. The **ngmix** method Sheldon [\[She15\]](#) fits a single parametric profile to the three exposures simultaneously using the corresponding PSF of each exposure. See Guinot et al. [\[Gui+22, §5.4\]](#) for a more detailed description of the method. The PSF model is built on each exposure independently of other exposures.

Having a PSF model that works on its own is a required condition to be the PSF model used to produce a WL shape catalogue. However, it is not a sufficient condition and there is still work to be done to integrate and validate that the method is working smoothly with the rest of the shape measurement pipeline. We continue by detailing some of this work.

First, the pipeline preprocessing of an exposure is done independently on each CCD. As **MCCD** builds a single model for the entire focal plane it was required to develop a preprocessing step that would gather all the CCD files from the same exposure and merge them into a new file. In the processing, a common coordinate system is generated following MegaCam’s focal plane geometry, and some outlier rejection criteria is performed. The criterion finally used for the v1 catalogue is to reject CCDs that have fewer stars than a given threshold that we set to 20.

An automated processing of many exposures requires handling outliers in the data, which was not done in [chapter 3](#). We therefore developed post-processing steps in order to identify and remove possible outliers that might impact the shape measurement. Given the amount of data we are targeting to process we can not afford to recompute the PSF model if there are some outliers like it is done in **PIFF** described in [section 2.4](#). After having trained the **MCCD** model, we check the pixel residuals between the PSFs produced by the model and the observed stars. For each star, we divide the postage stamp into two regions divided by a circular aperture. The outer region of the observation is used to estimate the noise level, $\hat{\sigma}$, using the **Median Absolute Deviation (MAD)** assuming Gaussian additive noise. We use the inner region to compute the residual between the PSF model and the observation, and then we calculate the **Root Mean Square (RMS)** value of the residual. If the quotient between the RMSE and the noise is over a certain threshold, Th_{star} , we mark the star as an outlier. Then, if the percentage of outlier stars over the total number of stars in the CCD is over another threshold, Th_{CCD} , we consider the CCD as an outlier. Consequently, we remove the outlier CCDs from the model and do not perform shape measurement on those CCDs.

The **MCCD** main parameters used for the v1 **ShapePipe** catalogue are presented in [Table 41](#). It can be surprising that the wavelet denoising is not being used in the model. Nonetheless, the wavelet denoising using starlets [\[SMB11\]](#), an isotropic undecimated wavelet, was biasing our PSF model. We were able to verify the bias by observing the ellipticity variations in the PSF model and compare them to the stars. The bias was towards rounder PSFs, which is expected as the starlets are *isotropic*. Therefore, denoising by thresholding starlet decomposition coefficients will bias the model towards smaller absolute ellipticity values. The PSF ellipticity variations in the CFIS data are faint, and therefore the small bias introduced by the wavelet denoising was not acceptable for our WL purposes.

Then, we need to make sure that the PSF model works smoothly with the shape measurement method. While this task might seem straightforward, it requires a good understanding of the PSF model *and* the shape measurement method. For example,

Parameter	Value
Model type	MCCD-HYB
Number of local eigenPSFs	$r = 8$
Local polynomial degree	$d^{\text{Loc}} = 2$
Global polynomial degree	$d^{\text{Glob}} = 8$
Local denoising level	$K_{\sigma}^{\text{Loc}} = 0$
Global denoising level	$K_{\sigma}^{\text{Glob}} = 0$
Outlier CCD rejection threshold	$T_{\text{CCD}} = 15\%$
Outlier star rejection threshold	$T_{\text{star}} = 1.25$

Table 41 – Summary of the MCCD PSF model parameters used for the v1 shape catalogue.

I needed to solve the inconsistency of the pixel definition between the PSF model and the shape measurement algorithm by careful cross-checking and examination the software implementations. The first (second) one indexes the pixels starting from zero (one), and it corresponds to the bottom-left corner (centre) of the pixel. Consequently, strong shape biases originated due to the disagreement on the required centroid of the PSF that was later used for deconvolution.

We finally included a pipeline module that is able to automatically produce PSF validation plots from the catalogue. It produces different *meanshapes* plots and the ρ -statistics that have been described in section 2.6. In the next section, we will show the validation plots from the v1 ShapePipe catalogue.

4.3 Shape catalogue and PSF diagnostics

We have successfully processed an area of 3663 deg^2 with ShapePipe. The unmasked, or effective, area corresponds to 3455 deg^2 . The processing run for ~ 3 months on Canadian Advanced Network for Astronomical Research (CANFAR), on more than 150 virtual machines using ~ 8 CPUs on each. The computing time was dominated by the shape measurement method `ngmix` with metacalibration.

We applied some selection criteria to the galaxy samples before the metacalibration, which were:

- $10 < \text{SNR} < 500$,
- $T_{\text{gal}}/T_{\text{PSF}} > 0.3$,
- $N_{\text{epoch}} \geq 2$,

where N_{epoch} is the number of epochs (or exposures) used to image an object, T_{gal} is the size of the galaxy deconvolved by the PSF, and T_{PSF} is the PSF size. The aforementioned selection criteria give us a total of $\sim 97.5 \times 10^6$ galaxies that are required to measure their shape. The weighted number density [Hey+12a] of the catalogue is $7.98 \text{ gal arcmin}^{-2}$. The amount of validation stars is $\sim 5.3 \times 10^6$ which corresponds to a random set of 20% of all the stars. These stars were not used to train the PSF model but to produce the validation plots.

We continue by presenting validation diagnostics of the shape catalogue, focusing on the PSF.

4.3.1 Shape measurement diagnostics

The first set of diagnostics will be related to the shear calibration from [subsection 1.3.7](#). We rewrite the shear bias parametrisation from [Equation 1.49](#) in matrix form as follows

$$\begin{bmatrix} \hat{g}_1 \\ \hat{g}_2 \end{bmatrix} = \underbrace{\begin{bmatrix} R_{1,1} & R_{1,2} \\ R_{2,1} & R_{2,2} \end{bmatrix}}_{\mathbf{R}} \begin{bmatrix} g_1 \\ g_2 \end{bmatrix} + \begin{bmatrix} c_1 \\ c_2 \end{bmatrix} + \alpha \begin{bmatrix} \epsilon_1^{\text{PSF}} \\ \epsilon_2^{\text{PSF}} \end{bmatrix}, \quad (4.1)$$

where we have added crossed multiplicative bias terms with $R_{1,2}$ and $R_{2,1}$. In this formulation, $R_{1,1}$ would correspond to $(1 + m_1)$ from [Equation 1.49](#). The metacalibration method allows us to compute the response matrix, as shown in [subsection 1.3.7](#).

We present the different parameters appearing in [Equation 4.1](#). The response matrix from the entire shape catalogue is

$$\mathbf{R} = \begin{bmatrix} 0.685 & 1.41 \times 10^{-4} \\ 1.39 \times 10^{-4} & 0.686 \end{bmatrix}. \quad (4.2)$$

The response matrix is almost diagonal, meaning there is no cross-contamination between the ellipticity components. The diagonal values obtained are similar to the ones obtained in the previous analysis [[Gui20](#); [Gui+22](#)].

The additive bias, $[c_1, c_2]^T$, can be computed using a weighted average. We can assume that the average intrinsic ellipticity and average shear vanish if we use a large enough area. We present the additive bias computed over the v1 catalogue, and we compare it with the one computed in Guinot et al. [[Gui+22](#)], $[c_1^{\text{old}}, c_2^{\text{old}}]^T$, in the following equations

$$\begin{bmatrix} c_1 \\ c_2 \end{bmatrix} = \begin{bmatrix} (-1.86 \pm 0.24) \times 10^{-4} \\ (1.52 \pm 0.24) \times 10^{-4} \end{bmatrix}, \quad \begin{bmatrix} c_1^{\text{old}} \\ c_2^{\text{old}} \end{bmatrix} = \begin{bmatrix} (-4.95 \pm 0.58) \times 10^{-4} \\ (4.66 \pm 0.59) \times 10^{-4} \end{bmatrix}. \quad (4.3)$$

The additive bias from the v1 catalogue is reduced, which is a positive indicator for the [MCCD PSF](#) model.

We proceed with the computation of the scale-dependent [PSF](#) leakage that is represented as α in [Equation 4.1](#). This quantity has been described in [section 2.6.2](#), and we use the estimator from [Equation 2.52](#) to compute it. The shapes of the [PSF](#) were measured with [GalSim](#)'s [HSM](#) module [[HS03](#); [Man+05](#)]. [Figure 4.3-4](#) shows the [PSF](#) leakage as a function of scale. The leakage is computed for each patch from [Figure 4.1-3](#), where the curve colours match the respective patch. The figure shows that the patch P7 seems to have some problems, although this has not yet been investigated. Besides that apparent issue, the overall leakage is around 2%, which is a small amount of leakage. The average all-scale leakage computed is 0.0278, which is lower than the previous value of 0.033 [[Gui+22](#)].

4.3.2 PSF diagnostics

In this section we focus on the moment-based [PSF](#) metrics. We start with the *meanshape* plots, where we can see a quantity averaged over the exposures binned in focal plane positions. Given the number of test stars, we can use a higher number of bins as we will not be noise dominated. We will use 20×40 bins in each [CCD](#), which means that each bin is composed of (102.4×115.3) pixels.

We first present the star density as a function of the focal plane position in [Figure 4.3-5](#). We can already see that some vertical lines in the [CCDs](#) seem not to be working correctly. There is also an underdense region in the bottom right corner of




Figure 4.3-4 – Scale-dependent PSF leakage computed over each patch of the v1 shape catalogue and the entire field.




Figure 4.3-5 – Test star density as a function of focal plane position.




Figure 4.3-6 – Meanshapes plots of both ellipticity components of the observed test stars and the PSF model. The first column corresponds to the first ellipticity component, while the second column to the second ellipticity component. The first row, figures (4.3-6a) and (4.3-6b), corresponds to the observed stars. The second row, figures (4.3-6c) and (4.3-6d), correspond to the PSF model. The third column, figures (4.3-6e) and (4.3-6f), correspond to the residual (star – PSF) between the observed stars and the PSF model.

the focal plane. This underdensity also seems to be present in the other corners of the focal plane but in a less significant way. The origin of the strong underdensity is not yet clear.

Figure 4.3-6 shows the meanshapes plots of the two ellipticity components. The scale of the residues colour bar is a tenth of the scale of the ellipticity colour bar. The MCCD PSF model is doing a good job, and the residual values are low over all the focal plane. By eye, one cannot tell the difference between the two first rows, corresponding to the observed test stars and the PSF model, respectively. However, there are still some visible residuals of low amplitude by high frequency in Figure 4.3-6e and Figure 4.3-6f. These are PSF spatial variations that the MCCD model is not able

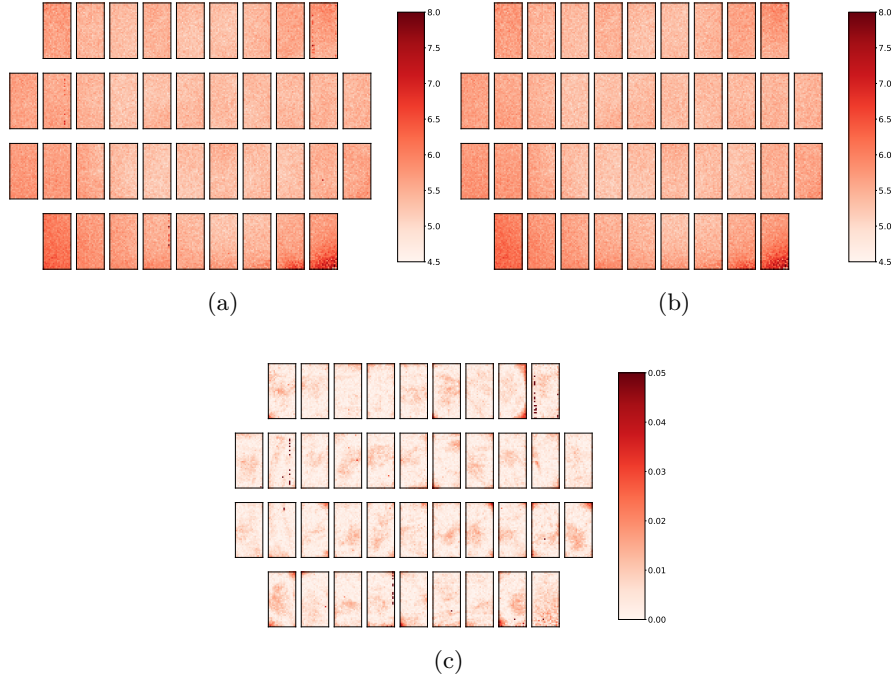


Figure 4.3-7 – Meanshapes plots of the observed stars and PSF’s sizes. The top left plot, figure (4.3-7a), shows the stars size, while the top right plot, figure (4.3-7b), shows the PSF size. The bottom plot, figure (4.3-7c), shows the residual between the stars and the PSF sizes that is computed as $|(R_{\text{star}}^2 - R_{\text{PSF}}^2)/R_{\text{star}}^2|$.

to capture. We expected that using the total number of stars in one exposure and a more complex model like MCCD would be able to capture such variations. This result indicates that one exposure does not have enough information to capture such a high-frequency low-amplitude pattern. The model should be able to reproduce the pattern due to the high number of eigenPSFs we set. However, the stars’ degradations and spatial distribution are insufficient to provide adequate information about those patterns. One difficulty worth mentioning with the modelling of these spatial variations is that to observe them, like in Figure 4.3-6, we need to process a vast amount of area. This processing implies an immense amount of computing resources and time⁷, which we do not dispose of.

The size meanshapes are shown in Figure 4.3-7. There are some vertical lines in some CCDs that seem to have problems and can be seen in the residual plot Figure 4.3-7c and more subtly on the stars sizes Figure 4.3-7a. The residuals are overall low in the focal plane. However, the highest errors are concentrated on the corners of some CCDs.

Figure 4.3-8 shows the ρ -statistics, described in section 2.6.2, from the current and past version of the ShapePipe catalogue. The v1 results are consistent with the previous catalogue. The ρ -statistics have consistently decreased for the larger scales (high θ). This result indicates that the MCCD model is better at modelling long-range spatial correlation by using the entire focal plane.

We have not included any pixel-based metric in this section. One reason is that we are not comparing the two PSF models. We are presenting shape catalogue diagnostics

⁷The v1 processing took 3 months in a High-Performance Computing (HPC) environment.

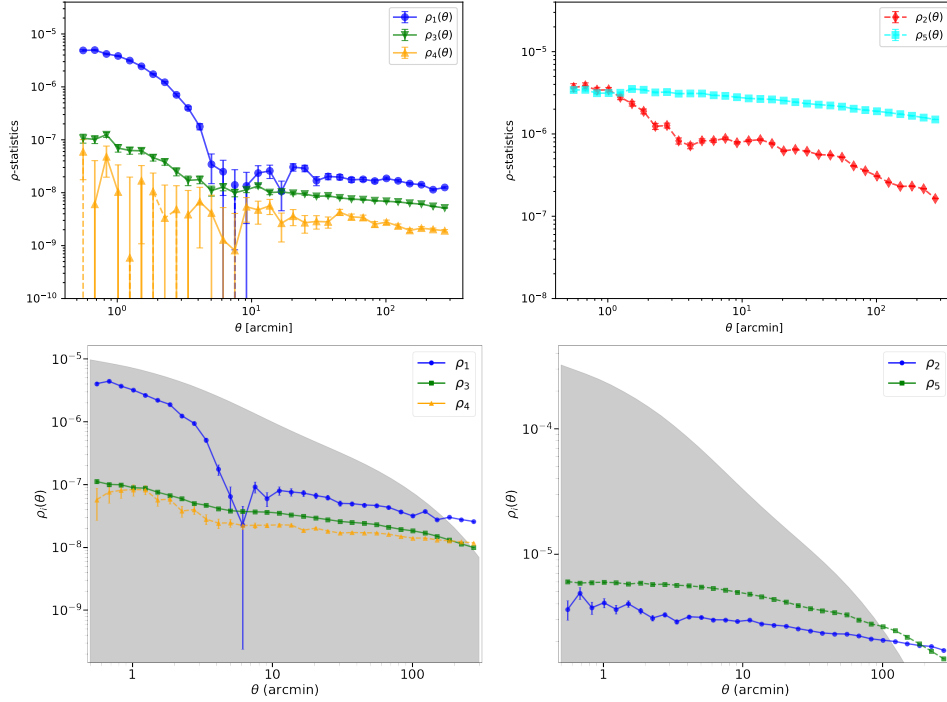


Figure 4.3-8 – The ρ -statistics computed on the PSF model and the test stars. The top row shows the results for the v1 catalogue, while the bottom row shows the results from the previous version of the catalogue from Guinot et al. [Gui+22]. Credit: second row plots from Guinot et al. [Gui+22].

showing that MCCD did a good job with modelling the PSF. The pixel-based metrics help compare two PSF models, and evaluating them using the same conditions is necessary. Ideally, it would be nice to have the CFIS data processed with both models and then choose the best performing one. However, the vast amount of data and the long processing times do not allow recomputing the entire dataset. Another reason is that the pipeline implementation of the pixel-based metrics was not finished when the CFIS data was processed. On the contrary, the meanshapes and the ρ -statistics plots were automatically generated from the PSF validation modules.

4.4 Conclusion

In this chapter we have presented the application of the MCCD PSF model in the ShapePipe shape measurement pipeline to process the r-band CFIS DR3. We briefly described the work required to incorporate the MCCD model into the existing ShapePipe pipeline. We then presented diagnostics for the shape measurement and the PSF modelling. The diagnostics gave good results with similar or lower systematics than the previous version of the catalogue from Guinot et al. [Gui+22]. We have shown that the MCCD model is robust enough to work with real data and that it can work in a high-performance shape measurement pipeline and produce a competitive WL shape catalogue.

The mere processing and handling of the $\sim 3500 \text{ deg}^2$ was a challenge. The current shape catalogue is one of the largest up to date when compared other WL shape

catalogues from the northern-sky like [CFHTLenS](#) (154 deg² [[Hey+12b](#)]), [KiDS-1000](#) (1000 deg² [[Gib+21](#)]), [HSC Y3](#) (433 deg² [[Li+22](#)]), and comparable to the southern-sky [DES Y3](#) (4143 deg² [[Gat+21](#)]). A more detailed analysis of the v1 shape catalogue and its comparison with the independent pipeline using *lensfit* is being carried out. It would certainly be interesting to compare the performances of both pipelines over the same [CFIS](#) area.

Even if the [MCCD](#) model is performing better than the previously used, [PSFEx](#), there are still visible low-amplitude high-frequency residuals in the ellipticity mean-shapes plots. This result indicates that the amount and quality of stars in one exposure do not provide enough information to model these variations accurately. We recall that each exposure is processed independently of the other exposures, and the [PSF](#) model is trained from a blind guess each time. If there is not enough information on these variations in one exposure, an improvement path would be to start using information from more exposures and start with a first guess that would allow capturing the low-amplitude high-frequency ellipticity variations. However, this is not a trivial exercise, as the atmospheric contribution of each exposure is non-negligible and independent of the others. The ellipticity patterns observed in the mean-shapes plots correspond to the ones introduced by optical aberrations from the telescope that are systematically introduced in each exposure. On the contrary, the atmospheric contribution averaged over many exposures vanishes.

The current chapter concludes this thesis's ground-based telescopes [PSF](#) modelling part. We have proposed a new [PSF](#) model, tested it, validated it, implemented it, and finally applied it to generate a [WL](#) shape catalogue with the [CFIS](#). The following part of this thesis will be related to the [PSF](#) modelling of space-based telescopes.

* * *
* *
*

Part III

PSF models for space-based telescopes

Rethinking data-driven PSF modelling with a differentiable optical model

Chapter Outline

5.1	Introduction	122
5.1.1	Challenges in PSF modelling for space missions	122
5.1.2	Current approaches for PSF modelling	123
5.1.3	Our contribution	124
5.1.4	Notation	124
5.2	The PSF field and the inverse problem	124
5.2.1	The observation model	125
5.2.2	Inverse problem and regularisations	126
5.3	Related works	127
5.4	WaveDiff, a data-driven wavefront-based PSF model	128
5.4.1	Physical motivation and the wavefront error	128
5.4.2	Wavefront error PSF model	129
5.4.3	Practical optical forward model	134
5.4.4	Training	136
5.4.5	PSF recovery	137
5.5	Numerical experiments	138
5.5.1	Objectives	138
5.5.2	Setup	140
5.5.3	Results	144
5.6	Discussion and conclusion	150

This chapter is based on an article submitted to a peer-reviewed journal:

Liaudat, T., Starck, J.-L., Kilbinger, M., Frugier, P.-A., ‘Rethinking data-driven point spread function modeling with a differentiable optical model’. In: *arXiv:2203.04908* (2022)

The work presented in this chapter was previously presented in a peer-reviewed conference:

Liaudat, T., Starck, J.-L., Kilbinger, M., Frugier, P.-A., ‘Rethinking the modeling of the instrumental response of telescopes with a differentiable optical model’. In: *NeurIPS 2021 Machine Learning for Physical sciences workshop*. Nov. 2021

FUTURE missions, like *Euclid*, impose severe requirements to the PSF model. The current state-of-the-art data-driven PSF model for *Euclid* [Sch+20] is not capable of modelling wavelength variations and is far from requirements, at a factor between 100 and 300 for the PSF shape and a factor of 10^4 for the PSF size. Furthermore, the recently discovered effect of the dichroic coating from the *Euclid*’s optical system makes PSF modelling even more challenging. The dichroic prevents us from assuming that the spectral variation of the PSF is smooth and only due to the diffraction phenomena. This effect challenges how we built data-driven models and the direction taken in Schmitz [Sch19] for the data-driven modelling of the PSF’s chromatic variations. All these facts triggered us to rethink how we build *data-driven PSF* models for space surveys.

In this chapter we present a new PSF model for space-based telescopes, coined *WaveDiff*, which was previously submitted in Liaudat et al. [Lia+22]. *WaveDiff* builds its model in the wavefront error space and is able to remain data-driven. We present the details of the new model and continue with a performance comparison with several state-of-the-art data-driven PSF models in a simulated simplified *Euclid* scenario. We end the chapter by discussing the results and with a conclusion.

5.1 Introduction

5.1.1 Challenges in PSF modelling for space missions

The main difference compared to a ground-based telescope PSF is the effect of the atmosphere, which acts as a fast-changing stochastic filter. The PSF tends to resemble a low-pass filter as the exposure time increases. Therefore, it strongly affects the observations and changes over time. The improvements in the PSF stability and quality are one of the main reasons to send a telescope into space.

As already mentioned above, telescope observations in galaxy imaging surveys include unresolved objects, which are stars, and extended objects of interest, such as galaxies, at different positions in the *Field-of-View* (FOV). The modelling of the PSF in the FOV can be seen as an inverse problem. We can consider stars as samples of the PSF field in the FOV. These samples can be used to constrain our model, which we will later use to infer the PSF at target positions. PSF field modelling for space missions encompasses several challenges, and we briefly review the three we will focus on in this chapter:

1. *The PSF varies spatially in the FOV.* The instruments are designed with a wide FOV and are composed of a large array of CCDs. These dimensions make the PSF vary substantially across the focal plane. The lack of atmosphere and the high-quality optics used in telescopes makes the PSF close to the diffraction limit. Consequently, the PSF has more complex shapes with higher spatial frequencies. The model needs to capture the different spatial frequency variations of the PSF from the stars in order to infer the PSF at target positions.
2. *The observations are undersampled.* This is generally the case for space missions, and the model needs to super-resolve the output PSFs. This scenario differs from

the usual **Super Resolution (SR)** task, as we do not have several low-resolution observations of the same object. In this case, we have several samples of the undersampled **PSF** field at *different positions* in the **FOV**.

3. *The PSF varies as a function of wavelength.* Also known as chromatic variations, they must be included in the **PSF** model for most science goals. However, the modelling difficulty resides in some instruments' very broad passband. This means that each star observation is integrated with respect to its wavelength throughout the passband (e.g. the *Euclid* passband extends from 550 to 900nm [Cro+16]). The **PSF** chromaticity will also be very likely important for the *Roman* space telescope [Spe+15] and LSST [LSS+09].

In addition to the three significant challenges mentioned above, we also need to address the fact that the telescope will change over time during its journey in space. Some reasons include strong temperature gradients or degradations of the instrument. Although capturing these variations into the **PSF** model might seem intricate, in practice, one can build independent **PSF** models at each point in time. As a consequence, we have a different model for each time snapshot, and we cope with the temporal variations. We will analyse a single point in time, or a single exposure, in the current chapter. The only limitation related to this approach is the number of observed stars available to constrain the **PSF** model. The stability of a space mission allows using several exposures to increase the available information when building the **PSF** model.

5.1.2 Current approaches for **PSF** modelling

We now review the two main approaches, or families of methods, for **PSF** modelling as follows:

Parametric PSF models This approach consists in building a parametric model of the entire optical system that aims to be as close as possible to the actual telescope. Then, a few model parameters are fit to the star observations. This family of methods is capable of handling the chromatic variations of the **PSF**. This approach has been used for the **Hubble Space Telescope (HST)**, with the **Tiny Tim** method [Kri93; KB95; KHS11]. However, it was later shown that a simple data-driven model [HA17] outperformed **Tiny Tim**, exposing limitations of parametric modelling. Errors will arise if there is a mismatch between the parametric model and the ground truth. Furthermore, even if, ideally, there was no mismatch, the optimisation of these models is a degenerate problem. It requires potentially expensive calibration information, usually in the form of out-of-focus observations, to break degeneracies. Some events, such as launch vibrations and ice contaminations, introduce significant variations into the model. These reasons prevent a complete ground characterisation from being successful.

Data-driven PSF models This approach, also known as *non-parametric PSF models*, only relies on the observed stars to build the model in pixel space. It is blind to the physics of the inverse problem. The models assume some regularity in the variation of the **PSF** field across the **FOV**. These methods differ in the way they exploit this regularity and handle the super-resolution [Lia+21a; Ber11; Sch+20; Jar+20; Mil+13]. Data-driven methods can easily adapt to the current state of the optical system. However, they have difficulties modelling complex shapes of **PSFs**. A limitation shared by all the data-driven methods is their sensitivity to the available

number of stars to constrain the model. When the number of stars falls below some threshold, the model built is usually considered unusable. This family of methods has been widely used for modelling ground-based telescope PSFs. Nevertheless, they cannot successfully model the chromatic variations in addition to the spatial variations and the super-resolution.

5.1.3 Our contribution

In this chapter, we propose a novel data-driven point spread function model coined WaveDiff. Our model can handle super-resolution and model spatial variations of wide-field telescopes. To the best of our knowledge, this is the first data-driven PSF model that can also capture the diffraction-induced chromatic variations of the PSF.

The proposed WaveDiff model represents a *paradigm shift* with respect to current state-of-the-art *data-driven PSF* models as it builds the model in the wavefront space. Although this is also the case for parametric models, WaveDiff remains data-driven. We can include the physics behind the PSF modelling inverse problem into the WaveDiff model using the *differentiable optical forward model* that we have designed. The forward model is theoretically based on optics principles and practically implemented in a modern automatically differentiable framework that runs on GPU. The optical forward model represents a simplified version of the telescope’s optical system and allows going from the wavefront to the pixel representation of the PSF. The forward model is fast and generic, as it depends on a reduced number of parameters known for optical telescopes. It could be easily adapted to model the PSF field of other telescopes. In addition, we propose a framework for training the WaveDiff model based on modern stochastic first-order optimisation methods. Even though the WaveDiff model is built in the wavefront space, it is constrained only using noisy in-focus observations without needing special calibration observations. The WaveDiff model can estimate a useful manifold in the wavefront space that approximates the underlying PSF field of the observations. We present three variations of the WaveDiff model that we refer to as its flavours.

5.1.4 Notation

In this chapter we use capital letters (e.g. H) for discrete functions where the codomain is $\mathbb{K}^{a \times b}$ for some $a, b \in \mathbb{N}$, and \mathbb{K} is \mathbb{R} or \mathbb{C} . Calligraphic capital letters (e.g. \mathcal{H}) are used for continuous functions where the codomain is \mathbb{K} . The variables (x, y) are assigned for the positions in the FOV, $[u, v]$ for positions in the focal plane, and $[\xi, \eta]$ for positions in the pupil plane.

5.2 The PSF field and the inverse problem

Let us define the PSF field for a particular image exposure as a function that has as inputs a position in the FOV and a specific wavelength and outputs a monochromatic PSF. We denote this function by $I : \mathbb{R}^2 \times \mathbb{R}_+ \rightarrow \mathbb{R}^{M \times M}$, where M^2 is the number of pixels of the desired super-resolved postage stamp. PSF modelling consists in building an estimator, H , of the PSF field, and then using it to output any monochromatic PSF at a set of target FOV positions $\{(x_j, y_j)\}_{j=1, \dots, m_{\text{target}}}$. The imaging instrument has a certain passband and several degradations that we need to consider in our model to match our predictions with the observations. We recall that a monochromatic PSF is a PSF image in a square postage stamp at a specific wavelength.




Figure 5.2-1 – Different representations of a space-like PSF at a single position in the field-of-view. The SED used corresponds to one of the templates from the numerical experiments described in subsection 5.5.2. (a) Noisy polychromatic PSF at observation resolution, (b) Noiseless polychromatic observation at observation resolution, (c) High-resolution noiseless polychromatic observation, (d) Wavefront error map representing the aberrations in the optical system. Units are in μm , (e) High-resolution chromatic variations of the PSF at equally spaced wavelengths in the passband [550, 900]nm.

5.2.1 The observation model

We consider stars as point sources. Therefore, star observations provide samples of the PSF field at their corresponding positions. The spectral energy of the point source at each wavelength is not necessarily the same, and we have to take this into account for our observations. The Spectral Energy Distribution (SED) gives us the normalised energy as a function of wavelength of a star at a FOV position (x_i, y_i) or $\int_{\text{passband}} \text{SED}(x_i, y_i; \lambda) d\lambda = 1$. From now on, we assume the SEDs to be known. The following observational model,

$$\bar{I}(x_i, y_i) = F_d \left\{ \int_{\text{passband}} \text{SED}(x_i, y_i; \lambda) I(x_i, y_i; \lambda) d\lambda \right\} + \mathbf{n}_i, \quad (5.1)$$

relates the star observation $\bar{I}(x_i, y_i) \in \mathbb{R}^{N \times N}$, where N^2 is the number of pixels in the observed postage stamp, at the FOV position $(x_i, y_i) \in \mathbb{R}^2$ to the objective PSF field $I(x_i, y_i; \lambda) \in \mathbb{R}^{M \times M}$, which is at a higher resolution with respect to the observations ($M > N$). The desired instrumental response I is integrated in the instrument's passband weighted by the star's normalised spectral energy distribution $\text{SED}(x_i, y_i; \lambda) \in \mathbb{R}_+$. Then, it is degraded with the operator $F_d : \mathbb{R}^{M \times M} \rightarrow \mathbb{R}^{N \times N}$, which accounts for: (i) downsampling that affects the pixels in the stamp by a factor D (meaning that $DN = M$), (ii) a sub-pixel shift that depends on where the position of the centroid of the object is placed with respect to the pixel grid, and (iii) other pixel-level degradations such as detector effects or guiding errors. The observational noise is modelled by \mathbf{n}_i , which we assume to be white and Gaussian for simplicity, i.e. $\mathbf{n}_i \sim \mathcal{N}(0, \sigma_i^2 \mathbf{I}_N)$. We make a distinction between a monochromatic PSF $I(x_i, y_i; \lambda)$, that is evaluated at a single wavelength, and its polychromatic counterpart $\bar{I}(x_i, y_i)$, which has been integrated in wavelength over a passband. Figure 5.2-1 presents examples of different representations of a PSF at a single FOV position.

5.2.2 Inverse problem and regularisations

The PSF modelling problem consists in building a model $H : \mathbb{R}^2 \times \mathbb{R}_+ \rightarrow \mathbb{R}^{M \times M}$ that best approximates $\{I(x_j, y_j; \lambda)\}_{j=1, \dots, n_{\text{target}}}$ for all the target positions and wavelengths. To do so, we have as input a set of noisy degraded polychromatic observations $\{\bar{I}(x_i, y_i)\}_{i=1, \dots, n_{\text{obs}}}$. The way we constrain our model with the observations is to minimise some type of reconstruction error between our model and the observations, for example, $\sum_{i=1}^{n_{\text{obs}}} \|\bar{I}(x_i, y_i) - \bar{H}(x_i, y_i)\|_F^2$, where $\|\cdot\|_F$ denotes the Frobenius norm of a matrix. We need to consider the same degradations applied to $I(x_i, y_i; \lambda)$ in Equation 5.1 to the model H in order to match the observation. This inverse problem is ill-posed due to the degradations in F_d , the broad passband integration, and the noise. Consequently, many PSF models can reproduce the observations.

To regularise the inverse problem, we can inject some prior knowledge we have about the PSF field. For example:

- *PSF field regularity*: The most used regularisation is to exploit the regularity of the PSF field in the FOV. This regularisation imposes that two close-by positions in the FOV should have a close PSF representation. The regularity translates as a correlation in space (and time) of the PSFs. Most current PSF models rely on this fact to build a PSF model that can obtain a good generalisation to target positions. Each PSF model imposes the regularity differently, e.g. Liaudat et al. [Lia+21a], Ngolè et al. [Ngo+16], Bertin [Ber11] and Jee et al. [Jee+07].
- *Positivity*: Given the way the PSF physically forms, it should not contain negative values. The presence of negative values in the observations is due to the noise. Therefore, one can then consider the positivity of the PSF to regularise the problem.
- *Smoothness of the PSF*: One can consider that the PSFs are structured, and therefore the model should output smooth or piecewise-smooth images. One way to consider this would be to use predefined parametric functions that are already smooth to build the PSF model, e.g. the Moffat function [BPZ88; Li+16]. This choice can be effective in avoiding noise contamination to the model. However, it is challenging to model PSF fields that are not well represented by the chosen base function. Another approach, proposed by Ngolè et al. [Ngo+16], would be to use sparsity to enforce structure in the PSF image. In practice, this is used to denoise the model so that it does not learn the observational noise. It can be achieved by enforcing a sparse representation of the PSF on some dictionary that can be built using, for example, isotropic undecimated wavelets, also known as starlets [SMB11]. Although this technique proved to be flexible and adaptable to the observed data, the resulting PSFs are biased towards the properties of the dictionary. In the case of starlets, the resulting PSF would be biased towards isotropic features, smoothing out anisotropies of the PSF.
- *Manifold hypothesis*: From a general perspective, the manifold hypothesis [FMN16] states that real-world data tends to lie in the vicinity of a low-dimensional manifold embedded in a higher-dimensional space. This hypothesis can be enforced in several ways depending on how the PSF model is built. For example, in Ngolè et al. [Ngo+16] the PSF model is built in a constrained matrix factorization framework. The PSF field is represented by a matrix factorisation, and the low dimensionality of the manifold is enforced by keeping a low rank in the factorisation. The rank is left as a model hyperparameter.

Up to now, the properties mentioned were mild assumptions of the PSF modelling problem. One might notice that this prior knowledge is a consequence of the physics involved in the image formation of the telescope. For example, the CCD detectors measure the intensity of the electric field, which is naturally positive. Also, the smoothness and structure of the PSF come from the fact that it is the diffraction pattern of the incoming wavefront. Instead of building our PSF model over the assumptions mentioned above, we can take a step back and inject the physics of the optical system into our model. In practice, we account for the physics of the inverse problem with the addition of a differentiable optical forward model. The PSF modelling space is shifted from the pixels to the wavefront, which is closer to the physics of the problem and a more natural space to model the PSF. With this change, we are not only considering the previously mentioned priors but adding much more physical information to solve our inverse problem. This change is fundamentally different in how data-driven PSF models are built.

The PSFs generated using the optical forward model are naturally without negative pixel values and smooth. However, as already mentioned, current data-driven models can also account for the positivity and smoothness of the PSF. Then, what would be the advantage of building a wavefront-based model? Current data-driven models only tackle consequences of the physical process involved in the inverse problem, for example, the positivity. The proposed wavefront-based model incorporates the physics involved. The reconstructed pixel PSF is computed as the intensity of the diffracted electric field in the focal plane. Furthermore, the optical forward model cannot reproduce observational noise, which allows us to avoid overfitting it without using a denoising technique.

It may seem that we are adding a complete model of the optical system as it is done in the parametric PSF models from subsection 5.1.2. Nevertheless, we remain on a general data-driven framework, as we will see in section 5.4. We include in our model the diffraction phenomena, which have been well studied, and general properties of the telescope like the focal length, pupil diameter, obscurations, and pixel spacing. The optical forward model is end-to-end differentiable, using an automatic differentiation framework. In addition, it uses the advances in stochastic first-order optimisation algorithms developed in recent years by the machine learning and deep learning community.

5.3 Related works

Data-driven efforts to model chromatic variations Previous works attempted to solve the problem of modelling the chromatic variations from the data-driven side. An approach based on optimal transport [Sch+18] was developed in Schmitz [Sch19]. The chromatic variations were modelled using the displacement interpolation [McC97], or a weighted Wasserstein barycenter [AC11], between the monochromatic PSFs at the two extreme wavelengths of the passband. However, the results were not encouraging, and the hypothesis taken by this approach might not hold for *Euclid*. The authors in Soulez et al. [SCU16] tried to model the propagation of the incoming wavefront from the stars through the different mirrors in the optical system. They propose to infer the aberrations introduced by each mirror by solving a phase retrieval type of problem. Finally, the proposed formulation is recast into a constrained optimisation problem using the star observations in the FOV. However, it is a proof-of-concept article, as many difficulties remained unaddressed, and the results were only qualitative. None of these studies proved to be a satisfying solution.

Phase retrieval with automatic differentiation Estimating the wavefront of an in-focus observation falls under the phase retrieval problems. Recent works [Wan+20; Won+21], based on a framework established by Jurling et al. [JF14] are tackling the phase retrieval problem [She+15] using automatic differentiation. Their objective is solely to estimate the wavefront of a single image. Using a first-order optimisation method over their forward model, they can obtain a good estimate of the wavefront. In our PSF modelling problem, we do not have access to the PSF pixel image at the positions of interest. In addition, estimating the wavefront is not the final goal of the PSF model. We use the wavefront as an auxiliary product useful for our modelling purposes but not as an objective. Therefore, these methods are not directly usable for our PSF modelling problem. Nevertheless, it illustrates the virtues and the utility of using automatic differentiation to tackle phase retrieval problems.

State-of-the-art data-driven PSF models Several state-of-the-art data-driven PSF models exist. In this study, we will focus on the three most relevant methods, which we will later use for comparison. These methods are PSFEx [Ber11], described in section 2.4, RCA [Ngo+16; Sch+20], described in section 3.3, and M CCD [Lia+21a], described in chapter 3. Although none of them is capable of modelling spectral variations, all of them can model spatial variations and handle super-resolution.

5.4 WaveDiff, a data-driven wavefront-based PSF model

5.4.1 Physical motivation and the wavefront error

Current and future telescope optical systems require high optical performance, thus imposing complex designs. To illustrate this, Figure 2.5-14 presents an example of the optical system of the *Euclid* space telescope from Racca et al. [Rac+16]. The PSF can be modelled with high accuracy and complexity using powerful commercial optic simulators. They depend on ray-tracing techniques to account for each optical surface. Although very accurate, these simulators are not practical for our problem due to their high complexity. They are not suited for massive PSF computations as are required in survey-type missions.

Nevertheless, we can consider a simpler optical system that is a good approximation of the actual system. This system should be able to reproduce the effects of the entire system while being more accessible to model. We use some ideas from parametric models [KHS11] and adopt a single converging lens system. Stars are well approximated by point sources at infinity and are considered as plane waves when entering the system. An ideal optical system turns the plane waves into spherical waves, and the rays converge at a single point at the focal plane of the instrument. The exact point at which the rays converge depends on the angle of incidence of the incoming rays with respect to the optical axis.

Due to imperfections and aberrations of the optical system, the outgoing wavefront is not exactly spherical, and the rays do not converge in a single point. The difference between the real aberrated wavefront and the ideal spherical wavefront is called the **Wavefront Errors (WFE)**. Figure 2.1-6 presents a simplified sketch of the WFE. The advantage of this simpler optical system is that we can take advantage of Fraunhofer's approximation of diffraction [Goo05, Chapter 4.3] and its effect on the wavefront [Goo05, Chapter 5.2] as we have seen in section 2.1. These approximations allow us to relate the electric field's propagation and diffraction between the pupil and focal planes. The field at the focal plane is the two-dimensional Fourier transform of the field at the pupil plane. The sensors measure the intensity of the electric field at

the focal plane. Therefore, we can relate the electric field at the pupil plane and the observed pixel PSF by taking the squared modulus of its Fourier transform. The monochromatic PSF at the focal plane can be written as

$$\hat{\mathcal{H}}(x_i, y_i; \lambda)[u, v] \propto \left| \iint_{\mathbb{R}^2} \mathcal{U}_p(x_i, y_i; \lambda)[\xi, \eta] \exp \left[\frac{-2\pi i}{\lambda f_L} (u\xi + v\eta) \right] d\xi d\eta \right|^2, \quad (5.2)$$

where f_L is the optical system's focal length, and $(x_i, y_i) \in \mathbb{R}^2$ is a FOV position that corresponds to a determined set of sky angular coordinates. We consider that for a specific FOV position and wavelength we have a given electric field at the pupil plane $\mathcal{U}_p(x_i, y_i; \lambda) : \mathbb{R}^2 \rightarrow \mathbb{C}$. The two coordinates of the pupil plane that we integrate over are $[\xi, \eta]$, while the coordinates at the focal plane are $[u, v]$, where we observe our pixel PSF. We can rewrite the previous equation using the Fourier Transform (FT) [Bra78] as

$$\hat{\mathcal{H}}(x_i, y_i; \lambda)[u, v] \propto \left| \text{FT} \{ \mathcal{U}_p(x_i, y_i; \lambda) \} \left[\frac{u}{\lambda f_L}, \frac{v}{\lambda f_L} \right] \right|^2. \quad (5.3)$$

Then, we need to specify the electric field at the pupil plane \mathcal{U}_p that characterizes the system's aberrations. The electric field writes

$$\mathcal{U}_p(x_i, y_i; \lambda)[\xi, \eta] = \mathcal{P}(x_i, y_i; \lambda)[\xi, \eta] \exp \left[\frac{2\pi i}{\lambda} \phi_\theta(x_i, y_i)[\xi, \eta] \right], \quad (5.4)$$

where $\mathcal{P}(x_i, y_i; \lambda) : \mathbb{R}^2 \rightarrow \mathbb{R}$ represents the obscurations encountered at the focal plane for a respective FOV position and wavelength, and $\phi_\theta(x_i, y_i) : \mathbb{R}^2 \rightarrow \mathbb{R}$ represents the WFE or the aberrations of our optical system. Obscurations at the pupil arise commonly due to the lens' shapes and the superposing of mirrors and supports. It is especially the case for telescopes due to occultation from a secondary mirror and its supporting structures. The obscurations are known and unlikely to change during the telescope's lifetime, although their figure may vary with the FOV position. They significantly affect the shape of the PSF and are added to our model. For simplicity, we neglect the obscuration's \mathcal{P} FOV variations. We consider an all-reflective-telescope designs and neglect any wavelength-dependence of the WFE ϕ_θ . We note that such wavelength dependence could be added to our model, for example, to model the dichroic filter of *Euclid*'s optical system.

To summarize, the WFE map is transformed into a complex electric field via Equation 5.4. This field is propagated to the focal plane, and its intensity is computed to provide the pixel PSF with Equation 5.3. We refer to the process of obtaining the pixel images from the WFE map as the *optical forward model*.

5.4.2 Wavefront error PSF model

In the following, we focus on how to use the WFE to model the PSF. The core of the proposed framework is the WFE model, or WFE PSF model, Φ_θ , which is given by the following equation of the electric field at the pupil plane

$$U_p(x_i, y_i; \lambda) = P \odot \exp [2\pi i \Phi_\theta(x_i, y_i; \lambda)], \quad (5.5)$$

where \odot is the Hadamard or element-wise product, $U_p(x_i, y_i; \lambda) \in \mathbb{C}^{K \times K}$, and $P, \Phi_\theta(x_i, y_i; \lambda) \in \mathbb{R}^{K \times K}$. For the sake of simplicity, we use a fixed obscuration P that does not depend on the FOV position and wavelength. The WFE PSF model depends on some parameters, represented with θ , and its goal is to compute a WFE map at any position in the FOV. To conduct this task, we have to estimate, or *learn*, the




Figure 5.4-2 – A schematic of the proposed framework for data-driven wavefront PSF modelling.

parameters θ of our model from the observed images using the full differentiable forward model that includes the degradations from the observational model in Equation 5.1. A schematic of the proposed framework is presented in Figure 5.4-2. We provide more detail of the forward model in subsection 5.4.3.

Our framework shifts a large part of the complexity from the PSF model into the forward model, which encodes all the diffraction phenomena and degradations. This change allows us to simplify the construction of the WFE PSF model block seen in Figure 5.4-2. This framework adds a physical meaning to the model as the WFE can be interpreted from a physical perspective. The WFE could even be measured directly on the telescope.

We next focus on a reliable generalisation capability of the PSF to target positions and the model’s adaptation to observed data. For this reason we propose to use a weighted sum of wavefront features (or *eigenWFE* if we draw a parallel with the notion of *eigenPSF* [Lia+21a; Sch+20; Ngo+16]). A wavefront feature is distributed across the FOV and, therefore, shared by all the PSFs at any FOV position. Then, the weight that goes with that wavefront feature will change depending on the FOV position. To facilitate the optimisation of the model, we consider two classes of wavefront features. The first features are fixed and should provide a good representation basis of our WFE. The second features are data-driven, which should be learned from the data, and smoothly adapt to the observations. We can write our WFE PSF model as follows

$$\Phi_{\theta}(x_i, y_i; \lambda) = \underbrace{\frac{1}{\lambda} \sum_{l=1}^{n_Z} f_l^Z(x_i, y_i) S_l^Z}_{\Phi^Z(x_i, y_i; \lambda)} + \underbrace{\frac{1}{\lambda} \sum_{k=1}^{n_{DD}} f_k^{DD}(x_i, y_i) S_k^{DD}}_{\Phi^{DD}(x_i, y_i; \lambda)}, \quad (5.6)$$

where $f_l^Z, f_k^{DD} : \mathbb{R}^2 \rightarrow \mathbb{R}$ and $S_l^Z, S_k^{DD} \in \mathbb{R}^{K \times K}$. The functions f_l^Z, f_k^{DD} give the weights for the wavefront features S_l^Z, S_k^{DD} that correspond to the fixed and the Data-Driven (DD) parts respectively. A fixed number of features is set for both parts, n_Z and n_{DD} , for the fixed and the DD respectively. These hyperparameters allow controlling the desired complexity of the model.

Zernike polynomials as fixed features

Zernike polynomials [Nol76] have been widely used in the optics community. These polynomials provide a good basis for modelling circular apertures as they are orthogonal on the unit disk. They are well suited to model the low-frequency variations of the WFE. We follow Noll [Nol76] definition and use polar coordinates $[\rho, \phi]$ instead of the Cartesian WFE coordinates $[\xi, \eta]$. The Zernike polynomials are defined as

$$Z_n^m[\rho, \phi] = \sqrt{n+1} R_n^m(\rho) \sqrt{2} \cos(m\phi), \quad \text{Even polynomial and } m \neq 0, \quad (5.7)$$

$$Z_n^{-m}[\rho, \phi] = \sqrt{n+1} R_n^m(\rho) \sqrt{2} \sin(m\phi), \quad \text{Odd polynomial and } m \neq 0, \quad (5.8)$$

$$Z_0^0[\rho, \phi] = \sqrt{n+1} R_n^0(\rho), \quad \text{for } m = 0, \quad (5.9)$$

where

$$R_n^m(\rho) = \sum_{s=0}^{\frac{n-m}{2}} \frac{(-1)^s (n-s)!}{s! \left(\frac{n+m}{2} - s\right)! \left(\frac{n-m}{2} - s\right)!} \rho^{n-2s}, \quad \text{with } R_n^m(1) = 1, \forall n, m, \quad (5.10)$$

where m and n are nonnegative integers that respect $n \geq m \geq 0$ and that $n - m$ is even. The Zernike polynomial coordinates are limited to $0 \leq \phi < 2\pi$ and $0 \leq \rho \leq 1$. The orthogonality relation of the Zernike polynomials can be written as

$$\int W(\rho) Z_j Z_{j'} d^2\rho = \delta_{j,j'}, \quad (5.11)$$

where j and j' represent a way of indexing the (n, m) and (n', m') values, and $W(\rho)$ is a weight function defined as

$$W(\rho) = \begin{cases} \frac{1}{\pi}, & \text{if } \rho \leq 1, \\ 0, & \text{if } \rho > 1, \end{cases} \quad (5.12)$$

Figure 5.4-3 shows the first 15 Zernike polynomials that are a function of the pupil coordinates $[\xi, \eta]$ on a circular aperture. The weights f_l^Z , for each discretised Zernike polynomial S_l^Z , are defined as independent polynomial variations of the FOV position, (x_i, y_i) . One of the hyperparameters of the parametric model is the number of Zernike polynomials, n_Z . The other hyperparameter is the maximum degree of the FOV position polynomial, d_Z , that will accompany each Zernike polynomial. The parametric model can be expressed as

$$\Phi^Z(x_i, y_i; \lambda)[\xi, \eta] = \frac{1}{\lambda} \sum_{l=1}^{n_Z} \underbrace{\pi_l^Z(x_i, y_i)^T}_{f_l^Z(x_i, y_i)} \mathbb{1}_{n_{d_Z}} S_l^Z[\xi, \eta], \quad (5.13)$$

where S_l^Z is the Zernike polynomial, $\mathbb{1}_{n_{d_Z}} \in \mathbb{R}^{n_{d_Z} \times 1}$ is a vector of ones and the position polynomial vector π_l^Z for the Zernike index l writes

$$\pi_l^Z(x, y) = \left[\pi_{l,[0,0]}^Z, \pi_{l,[1,0]}^Z x, \pi_{l,[0,1]}^Z y, \dots, \pi_{l,[0,d_Z]}^Z y^{d_Z} \right], \quad (5.14)$$

where $\{\pi_{l,[i,j]}^Z\}_{i+j \leq 1}$ are the coefficients for of the position polynomial. The total number of parameters to estimate if we use n_Z Zernike polynomials is $n_Z n_{d_Z}$, where $n_{d_Z} = (d_Z + 1)(d_Z + 2)/2$ corresponds to the number of monomials (and parameters) in the polynomial. We recall that the Zernike polynomial definition is in pupil plane coordinates $[\xi, \eta]$ (e.g. the maps seen in Figure 5.4-3). Note that in this model n_Z does not need to be equal to n_{d_Z} . On the one hand, n_Z controls the maximum frequency content of the WFE in a given FOV position. On the other hand, n_{d_Z} controls the maximum frequency in the FOV spatial variation of the WFE.




Figure 5.4-3 – The first 15 Zernike polynomials with their Noll’s single-index prefaced by the hash symbol #, and ordered vertically by their radial degree n and horizontally by their azimuthal degree m .

Data-driven features

The data-driven features, represented by the S_k^{DD} matrices, are non-parametric. We do not impose a specific structure on the features, which are entirely learned from the data. These **DD** features should adapt to capture variations that are not well modelled by the Zernike-based features and should also account for unexpected mismatches between our forward model and the ground truth model.

It is important to maintain a certain regularity in the **PSF** field and have a good generalisation capability over the full **FOV**. The underlying hypothesis of the **PSF** field regularity is a consequence of the optical system seen in Figure 2.5-14. The system is composed of mirrors of different sizes. The way the mirror’s footprints propagate to the instrument’s focal plane imposes a regularity in the **PSF** field. The simplified model we propose in Figure 2.1-6 cannot inherently account for this propagation, and it is necessary to impose the regularity in the $f_k^{\text{DD}}(x, y)$. The way we impose this regularity will define different flavours of the proposed WaveDiff model:

- **WaveDiff-original:** the feature weights are based on **FOV** position polynomials up to a degree d_{DD} . The amount of data-driven features used, n_{DD} , is fixed by the maximum polynomial degree d_{DD} , and is computed as $n_{\text{DD}} = (d_{\text{DD}} + 1)(d_{\text{DD}} + 2)/2$. In order to provide more flexibility to the model without modifying the maximum polynomial degree, we add a mixing matrix $A_{\text{mix}} = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_{n_{\text{DD}}}]$ that allows the features to contribute to more than one monomial from the **FOV** position polynomial. The data-driven contribution to the **WFE** writes

$$\Phi^{\text{DD}}(x_i, y_i; \lambda)[\xi, \eta] = \frac{1}{\lambda} \sum_{k=1}^{n_{\text{DD}}} \underbrace{\boldsymbol{\pi}^{\text{DD}}(x_i, y_i)^{\text{T}} \mathbf{a}_k}_{f_k^{\text{DD}}(x_i, y_i)} S_k^{\text{DD}}[\xi, \eta], \quad (5.15)$$

where $\boldsymbol{\pi}^{\text{DD}}(x_i, y_i) \in \mathbb{R}^{n_{\text{DD}} \times 1}$ is a vector composed with each monomial of the n_{DD} -dimensional **FOV** position polynomial, $\mathbf{a}_k \in \mathbb{R}^{n_{\text{DD}} \times 1}$ is a column from the mixing matrix $A_{\text{mix}} \in \mathbb{R}^{n_{\text{DD}} \times n_{\text{DD}}}$, and $S_k^{\text{DD}} \in \mathbb{R}^{K \times K}$ is a data-driven feature.

The polynomial vector is expressed as follows

$$\boldsymbol{\pi}^{\text{DD}}(x, y) = \left[\pi_{[0,0]}^{\text{DD}}, \pi_{[1,0]}^{\text{DD}}x, \pi_{[0,1]}^{\text{DD}}y, \pi_{[2,0]}^{\text{DD}}x^2, \pi_{[1,1]}^{\text{DD}}xy, \dots, \pi_{[0,d_{\text{DD}}]}^{\text{DD}}y^{d_{\text{DD}}} \right]^{\text{T}}, \quad (5.16)$$

where $\pi_{[i,j]}^{\text{DD}} \in \mathbb{R}$ for $i + j \leq d_{\text{DD}}$ are the polynomial coefficients. In this model, S_k^{DD} , $\boldsymbol{\alpha}_k$ and $\boldsymbol{\pi}^{\text{DD}}$ for $k = 1, \dots, n_{\text{DD}}$ will be learned from the data. The total number of parameters of the data-driven part is $n_{\text{DD}} + n_{\text{DD}}^2 + n_{\text{DD}} K^2$. Our proposition is a constraining in the WFE spatial variations, as we are regularized by a d_{DD} position polynomial. However, as the data-driven features S_k^{DD} have no imposed structure, they are able to reproduce high WFE frequencies. These choices were motivated by the discussions in section 2.5 and particularly subsection 2.5.2.

- WaveDiff-graph:** we want to add constraints to the feature weights to capture smooth and localised spatial variations of the PSF field. To achieve this goal, we follow the approach presented in Ngolè et al. [Ngo+16] and already applied in Schmitz et al. [Sch+20] and Liaudat et al. [Lia+21a]. We aim to capture variations that occur at different spatial frequencies. Unfortunately, the fact that the PSF positions in the FOV are randomly distributed makes it hard to define the different types of variations. To tackle this issue, we define a fully-connected undirected weighted graph where each node corresponds to the position of a star. The weights are built as a parametric function of parameters Λ of the distance between the different stars. Once the graph is defined, we compute its Laplacian matrix $L_{\text{p},\Lambda} \in \mathbb{R}^{n_{\text{stars}} \times n_{\text{stars}}}$ [Chu97], that encodes all the weights and connections in the graph. This matrix is real and symmetric, and we can decompose it as $L_{\text{p},\Lambda} = V_{\Lambda} \Sigma_{\Lambda} V_{\Lambda}^{\text{T}}$, with Σ_{Λ} being a diagonal matrix containing the eigenvalues. The eigenvectors of the Laplacian matrix, V_{Λ} , represent the graph harmonics or a basis for defining the graph Fourier transform [Ric+19]. From these, a dictionary is built that is used to factorise the feature weights, as follows. In order to capture more spatial variations, we use r different parameters, $\{\Lambda_1, \dots, \Lambda_r\}$ for the graph's weight function. These parameters define different graphs and subsequently different harmonics. We concatenate the graph-harmonics, or Laplacian eigenvectors, of all the graphs to build the graph dictionary that we define as $V = [V_{\Lambda_1}, \dots, V_{\Lambda_r}]$. We write $\mathbf{f}_k^{\text{DD}} = [f_k^{\text{DD}}(x_1, y_1), \dots, f_k^{\text{DD}}(x_{n_{\text{stars}}}, y_{n_{\text{stars}}})] \in \mathbb{R}^{1 \times n_{\text{stars}}}$, where the weight function of the DD feature k is expressed as a vector of all star positions. Then, we factorise the feature weight vector as $\mathbf{f}_k^{\text{DD}} = \boldsymbol{\alpha}_k V^{\text{T}}$, where $\boldsymbol{\alpha}_k \in \mathbb{R}^{1 \times r n_{\text{stars}}}$. The DD feature weights are built as a linear combination of graph-harmonics, represented as row vectors in V^{T} . The vector $\boldsymbol{\alpha}_k$ defines the weights of this combination. We impose sparsity on the representation vector so that the learned feature can specialise on a specific or a small subset of graph harmonics. The sparsity in the $\boldsymbol{\alpha}_k$ vector can be achieved by adding an ℓ_1 constraint in the final loss function. We call it the *graph constraint*. The optimisation algorithm will select the best graph-harmonics for a specific feature that allows obtaining a good representation of our observations. For a more detailed description of how we define the graph weights and the choice of the Λ parameters, see [Ngo+16, Sections 5.2, 5.5.3].
- WaveDiff-polygraph:** the feature weights simultaneously use the two previously mentioned constraints. It contains two contributions, one based on FOV position polynomials and the other based on the graph constraint.

5.4.3 Practical optical forward model

The choice of the proposed model and Fourier optics allows computing the forward modelling of the wavefront to the pixels using a Fourier Transform, as described in [subsection 5.4.1](#). Given the vast amount of data to process in upcoming missions, it is essential to develop fast methods. This requirement motivates the use of the [Fast Fourier Transform \(FFT\)](#) algorithm. [Equation 5.3](#) shows that the sampling of the wavefront space changes with wavelength (λ). This observation implies that resampling is required for each wavelength if we want to work with a fixed space $\mathbb{C}^{K \times K}$, which is highly impractical. We need to calculate several monochromatic [PSFs](#) at different wavelengths to compute a single polychromatic [PSF](#) that is integrated into the instrument's passband. Furthermore, we have to take into account that the input and output matrices of the [Fast Fourier Transform \(FFT\)](#) algorithm share dimensions. We want to be Nyquist sampled at all wavelengths to avoid any aliasing problem.

One way to accelerate the calculations is to fix the dimensions of the electric field at the pupil plane from [Equation 5.5](#) as $\mathbb{C}^{K \times K}$, and then add a wavelength-dependent zero padding. By adding zeros in the obscured part and fixing the non-obscured dimensions, we can control the sampling of the electrical field at the pupil plane. This practical method avoids resampling the electrical field for each wavelength and allows to be Nyquist sampled. We zero-pad the electrical field matrix from $\mathbb{C}^{K \times K}$ to $\mathbb{C}^{p(\lambda) \times p(\lambda)}$. The required dimensions for the input matrix of the [FFT](#) operation can be seen in the following equation

$$p(\lambda) = \frac{K \lambda f_L}{M_D \Delta} Q, \quad (5.17)$$

where K is the fixed matrix dimension of the electrical field that is equal to the pupil diameter in pixels, f_L is the focal length, M_D is the length of the pupil diameter, Δ is the length of a squared-pixel in the focal plane, λ is the wavelength, and Q is a dimensionless oversampling factor. The following [section 5.4.3](#) describes the derivation of the [Equation 5.17](#). In practice, we fix the telescope's parameters, we choose the dimension K we will work with, and then we calculate the required Q so that we respect the Nyquist criterion ($p(\lambda) \geq 2K$) for every λ in our passband. Following the variable zero padding, each computed monochromatic [PSF](#) will have different dimensions. A subsequent crop can solve this issue to a fixed size of $\mathbb{R}^{M \times M}$. The energy contribution of the cropped pixels is negligible, given that M is sufficiently large. The propagation of the wavefront [PSF](#) to the oversampled pixel [PSF](#) can be resumed as follows

$$\hat{H}(x_i, y_i; \lambda) \propto \underset{p(\lambda) \times p(\lambda) \rightarrow M \times M}{\text{crop}} \left\{ \underbrace{\left| \text{FFT} \left\{ \underset{K \times K \rightarrow p(\lambda) \times p(\lambda)}{\text{pad}} \left(\overbrace{P \odot \exp \left[\frac{2\pi i}{\lambda} \Phi_\theta(x_i, y_i) \right]}^{\text{Electric field}} \right) \right\} \right|^2}_{\text{Pixel representation}} \right\}. \quad (5.18)$$

[Equation 5.18](#) represents the optical system box in [Figure 5.4-2](#) that gives us access to the monochromatic oversampled pixel [PSF](#). A drawback of the aforementioned zero-padding approach is that we cannot compute a monochromatic [PSF](#) at an arbitrary wavelength. We are constrained to a set of finite wavelengths that correspond to integer values of $p(\lambda) \in \mathbb{N}$. If the parameter K and Q are large enough, there will be enough possible wavelength values, and it will not impact the results. In practice, we

carry out an interpolation to obtain the SED information at the specific wavelengths defined by Equation 5.17.

We still need to apply the degradations to the previous pixel PSF so that it can match our observations. This action resumes in applying the observational model seen in Equation 5.1 to the monochromatic oversampled PSF. We first need to normalize the PSF so that it has a unit pixel sum, or $\sum \sum_{l,m=1}^M \hat{H}(x_i, y_i; \lambda)[u_l, v_m] = 1$. We can then integrate over the instrument's passband and downsample to match the instrument's sampling, which is resumed as follows

$$\hat{H}(x_i, y_j) = D_D \left\{ \sum_{k=1}^{n_\lambda} \hat{H}(x_i, y_j; \lambda_k) \text{SED}(x_i, y_i; \lambda_k) \right\}, \quad (5.19)$$

where we discretised the integral in Equation 5.1 in n_λ evenly spaced bins, and D_D represents a downsampling by a factor D that can include a further crop on the postage stamp. To match the observation sampling we need to use $D = Q$. The SED in Equation 5.19 has been normalized so that for a given n_λ and position (x, y) , it verifies $\sum_{k=1}^{n_\lambda} \text{SED}(x, y; \lambda_k) = 1$. As one can see from Equation 5.19, we are not including other detector effects for the sake of simplicity. We restrict ourselves to using an oversampling factor $Q \in \mathbb{N}$ as the implementation of the downsampling operator is simplified. This constraint allows us to implement the downsampling operator with a well-parametrized average pooling layer widely available from the deep learning community.

Wavelength dependent zero padding formula

Table 51 – Optic variables used for the optical forward model.

Variable	Description
λ	Wavelength [μm]
K	Matrix dimension of the pupil
f_L	Telescope's focal length [m]
M_D	Pupil diameter [m]
Δ	Pixel size [μm]
Q	Oversampling factor [Dimensionless]
$p(\lambda)$	Matrix dimension including pupil and external obscuration

We resume the optical variables used to derive Equation 5.17 in Table 51. The diameter of the pupil aperture is by definition M_D , so the grid spacing in the pupil plane is $\delta_\xi = M_D / K$. The length of the simulated pupil plane is $M_D p(\lambda) / K$, as it has a total of $p(\lambda)$ grid elements and $p(\lambda) \geq K$. The length of the simulated focal plane part is $\Delta p(\lambda)$. This result is calculated as the length of one pixel, Δ , times the matrix dimension we will apply the FFT to, $p(\lambda)$, or the number of grid elements per dimension. Recall that we are modelling the propagation of a wave from the pupil plane, variables $[\xi, \eta]$, to the focal plane, variables $[u, v]$, under Fraunhofer diffraction.

Equation 5.3 shows that the frequencies from the pupil plane coordinates $[\xi, \eta]$ are directly mapped to the focal plane's spatial coordinates $[u, v]$. The equations write

$$f_\xi = \frac{u}{\lambda f_L}, \quad f_\eta = \frac{v}{\lambda f_L}. \quad (5.20)$$

we consider one of the two equations above, thanks to the symmetry of the problem. Replacing then the maximum frequency in the pupil plane, $1/\delta_\xi$, and the maximum

position in the focal plane, we obtain

$$\frac{K}{M_D} = \frac{p(\lambda)\Delta}{\lambda f_L}. \quad (5.21)$$

We can adjust the image’s resolution by dividing the pixel size Δ by the oversampling factor Q . In other words, with $Q > 1$, we are reducing the pixel size and increasing the system resolution. We include the oversampling factor, and we obtain

$$\frac{K}{M_D} = \frac{p(\lambda)\Delta}{\lambda f_L Q}. \quad (5.22)$$

Finally, if we solve for $p(\lambda)$ we obtain [Equation 5.17](#). For more information on optical simulations we refer the reader to [[Sch10](#), §4].

5.4.4 Training

The parameters θ we need to estimate are contained in the [WFE PSF](#) model Φ_θ from [Equation 5.6](#). For the fixed features, we optimize with respect to f_l^Z , and for the data-driven features with respect to f_k^{DD} and S_k^{DD} . We assume white Gaussian noise as in [Equation 5.1](#) and propose to use the following loss function to optimise our model

$$L(\theta) = \frac{1}{n_{\text{stars}}} \sum_{i=1}^{n_{\text{stars}}} \frac{1}{\hat{\sigma}_i} \left\| \hat{H}(x_i, y_j; \Phi_\theta) - \bar{I}(x_i, y_i) \right\|_F^2 + \mathcal{R}(\theta), \quad (5.23)$$

where $\hat{\sigma}$ is an estimation of the observed star noise standard deviation. The estimation uses the median absolute deviation (MAD) [[SMF15](#)]. We use the pixels from the outer region of the squared postage stamp, where we have masked the circular central region that contains most of the star’s energy. We represent a set of regularisers on the model’s parameters θ with $\mathcal{R}(\theta)$. Depending on which WaveDiff flavour we are using, the regulariser can take different forms. For the WaveDiff-graph model we use a $\mu \|\alpha_k\|_1$ regularization to impose sparsity on the different α_k vectors that were defined in [section 5.4.2](#), and are used to build the [DD](#) feature weights. The hyperparameter μ allows us to control the strength of the regularisation. We use the same regularisation for the graph part in the WaveDiff-polygraph model. However, we do not use any regularisation for the WaveDiff-original model, and we discuss this choice in [section 5.6](#).

The optimisation algorithm used in practice is Rectified Adam [[Liu+19](#)], a stochastic gradient-based method widely used in the machine learning community. This optimisation method proved to be more stable for the training than the classical Adam algorithm [[KB14](#)]. The data-driven features and the feature weight parameters are initialised using a uniform distribution with zero mean. The bounds of the distribution are kept close to zero, -1nm and 1nm , so that the starting [PSF](#) model contains a small amount of aberrations, close to a perfect system, and almost no spatial variations. This choice corresponds to starting with an almost perfect system, close to a vanishing [WFE](#). However, the small non-zero aberration values help to escape local minima at the start of the optimisation algorithm. The optimisation adds aberrations to the system until the observed [PSF](#) field is well represented. We have found that the best working optimisation strategy is to iterate the optimisation between the parameters of the parametric and the data-driven parts. This strategy means we first fix the data-driven parameters and optimise the parametric part. We then fix the parametric part and optimise the data-driven parameters. We repeat these iterations until convergence.

Different learning rates and number of epochs are used for each optimisation sequence. Algorithm 2 presents a summary of the training procedure.

The proposed PSF model framework using automatic differentiation is versatile. If prior information is available, it can be added as regularisations at the wavefront or pixel level. The framework can be adapted for other imaging applications, and other task-specific regularisers can be added to the loss. As the model is developed in the TensorFlow framework [Mar+15], making these adaptations is straightforward.

Algorithm 2 WaveDiff training procedure

- 1: Hyperparameters:
 - 2: Learning rates $\{\eta_m^{\text{DD}}\}_{m=1,\dots,m_{\max}}, \{\eta_m^{\text{Z}}\}_{m=1,\dots,m_{\max}}$
 - 3: Number of epochs $\{N_{\text{ep},m}^{\text{DD}}\}_{m=1,\dots,m_{\max}}, \{N_{\text{ep},m}^{\text{Z}}\}_{m=1,\dots,m_{\max}}$
 - 4: Regularisation strength μ (Graph flavours only)

 - 5: Initialisation:
 - 6: Data-driven features: $S_k^{\text{DD}} \sim \mathcal{U}[-10^{-3}, 10^{-3}]^{K \times K}, \forall k$
 - 7: Data-driven mixing matrix: $A_{\text{mix}} = I_{n_{\text{DD}}}$
 - 8: Data-driven feature weights: $\pi_{[i,j]}^{\text{DD}} \sim \mathcal{U}[-10^{-2}, 10^{-2}], \forall i + j \leq d_{\text{DD}}$
 - 9: Parametric feature weight: $\pi_{l,[s,t]}^{\text{Z}} \sim \mathcal{U}[-10^{-2}, 10^{-2}], \forall l \leq n_{\text{Z}}, s + t \leq d_{\text{Z}}$
 - 10: Estimate noise level, $\hat{\sigma}_i \leftarrow \text{MAD noise estimation}(\bar{I}(x_i, y_i))$
 - 11: Graph constraint initialisation [Ngo+16, §5.2, §5.5.3] (Graph flavours only)
 - 12: Generate Zernike polynomial maps $S_l^{\text{Z}}, \forall l$ and telescope obscuration P

 - 13: Alternate optimisation:
 - 14: **for** $m = 1$ to m_{\max} **do**
 - 15: Optimise $L(\theta)$ over $\pi_{l,[s,t]}^{\text{Z}}$, (parametric part)

$$\forall l \leq n_{\text{Z}}, s + t \leq d_{\text{Z}} \text{ using } \eta_m^{\text{Z}}, N_{\text{ep},m}^{\text{Z}}$$
 - 16: Optimise $L(\theta)$ over $\pi_{[i,j]}^{\text{DD}}, A_{\text{mix}}, S_k^{\text{DD}}$, (data-driven part)

$$\forall k \leq n_{\text{DD}}, i + j \leq d_{\text{DD}} \text{ using } \eta_m^{\text{DD}}, N_{\text{ep},m}^{\text{DD}}$$
 - 17: **end for**

 - 18: Return PSF model: Φ_θ
-

5.4.5 PSF recovery

PSF recovery consists in using the trained PSF model to estimate the PSF at other positions in the FOV where star observations are not available. The recovery requires specifying how the feature weight functions can extrapolate to different positions in the FOV. Then, it is possible to reuse the learned and fixed WFE features from Equation 5.6 to calculate the WFE at the FOV position of interest.

The procedure is straightforward when the weight functions are defined as FOV polynomials, as in the parts involving Zernike polynomials or the WaveDiff-original. It consists in evaluating the polynomial on a new position and obtaining the new feature weights.

In the case of the WaveDiff models with graph flavours, the PSF recovery is not direct. During the optimisation, it is not only the DD features being learned but also their corresponding spatial variations that depend on the graph constraint, described in section 5.4.2. We want to be able to recover the PSF at any position in the FOV while respecting the learned spatial variations. To accomplish this, we consider the different weight functions f_k^{DD} , or the vector \mathbf{f}_k^{DD} , learned for each feature S_k^{DD} . These encompass the learned spatial variations that contain a selection of useful graph harmonics. We follow the strategy from chapter 3 and use the learned f_k^{DD} at different FOV positions to build a local interpolant. Then, we use it to interpolate the feature weights to the desired position. In practice, we use a Radial Basis Function (RBF) interpolation scheme. For a new position (x_j, y_j) , we select the closest N_{RBF} training star positions $\{(x_i, y_i)\}_{i=1, \dots, N_{\text{RBF}}}$. The interpolant function of the weights corresponding to feature k is built using the following equation

$$g_k(x, y) = \sum_{i=1}^{N_{\text{RBF}}} \lambda_{k,i} K(\| [x - x_i, y - y_i]^\top \|_2), \quad (5.24)$$

where $K : \mathbb{R}_+ \rightarrow \mathbb{R}$ is the interpolation kernel or radial basis function that takes as input the distance between the desired and the training position. The weights that define the interpolant g_k are $\{\lambda_{k,i}\}_{i=1, \dots, N_{\text{RBF}}}$. We learn these weights by using a set of exact interpolation constraints in the training positions that writes $\{g_k(x_i, y_i) = f_k^{\text{DD}}(x_i, y_i)\}_{i=1, \dots, N_{\text{RBF}}}$. We follow [Lia+21a, §3.7] and use a thin plate kernel that is defined as $K(r) = r^2 \ln(r)$, and set N_{RBF} to 20. To resume, we have to build one interpolant function for each feature k and each new position (x_j, y_j) . Even though it may seem that it is a costly procedure in terms of time and resources, the interpolation is done using 1-D functions, making it fast and cheap.

5.5 Numerical experiments

In this section, we start by defining the objectives of the numerical experiments. Then, we introduce the experiment setup and the data set we use to accomplish our goals. Finally, we present and discuss the results we find.

5.5.1 Objectives

To evaluate the novel framework and show its utility, we consider a scenario where some physical phenomena introduced some unexpected complexity into the PSF field. This choice means that we are not explicitly accounting for this increase in complexity in our model. An ancillary interpretation is that we have underestimated the complexity of the PSF field in our model. The practical way we set up the experiment is to simulate the WFE of a PSF field that is more complex than what the models can reproduce in the WFE space. In this setting, none of the models can have a zero reconstruction error over the ground truth WFE, even in a noiseless scenario. After simulating the complex WFE field, we generate the observations, the pixel PSFs, using the optical forward model. The added complexity is represented in the WFE composition and its spatial variation and not in a change of the optical forward model.

The chosen setup allows comparing different model performances and observing to which degree the models can build an approximation of a complex WFE PSF field. The setup is also helpful to study at which point the built approximation is good in terms of pixel PSF reconstruction error. We use this setting to study the following items:

- (a) Pixel reconstruction performance at the observation resolution and at 3 times the observation resolution (super-resolution) for polychromatic images.
- (b) Pixel reconstruction performance at the observation resolution as a function of wavelength for the best performing model of (a).
- (c) Pixel reconstruction performance as a function of the total number of observed stars in **FOV** that are used to constrain the **PSF** models.
- (d) Error in the recovered **WFE** with respect to the ground truth **WFE** of the **PSF** field.
- (e) Errors using weak-lensing metrics.

Reconstruction metrics

The main metric to compare the performance is the residual **Root Mean Squared Error (RMSE)**. As we work with simulations, we have access to the ground truth **PSF** field. Therefore, we can calculate noiseless **PSFs** at any resolution and **FOV** position, which facilitates the performance comparison. We define the **RMS** of a matrix $A_i \in \mathbb{K}^{K \times K}$, or a vector in \mathbb{K}^{K^2} , as

$$\text{RMS}(A_i) = \left(\frac{1}{K^2} \sum_{j=1}^{K^2} |A_{i;j}|^2 \right)^{1/2}, \quad (5.25)$$

where we use a single index that corresponds to a flattened matrix. Let us consider a set of n noiseless $K \times K$ dimensional images $\{A_i\}_{i=1}^n$, and the corresponding reconstructions $\{\hat{A}_i\}_{i=1}^n$. The main pixel reconstruction errors we use are the absolute error, Err_{abs} , and the relative error, Err_{rel} , that are defined as

$$Err_{\text{abs}} = \frac{1}{n} \sum_{i=1}^n \text{RMS}(A_i - \hat{A}_i), \quad Err_{\text{rel}} = \frac{1}{n} \sum_{i=1}^n \frac{\text{RMS}(A_i - \hat{A}_i)}{\text{RMS}(A_i)} \times 100\%. \quad (5.26)$$

A constant value over a **WFE** matrix does not affect the morphology of the pixel **PSF**. Therefore, we remove the mean of each **WFE** matrix before computing the **RMS** of the residual in the **WFE** error calculation. We only use the non-obscured elements of the **WFE** matrix to compute the **RMS**.

Moment-based metrics

Weak lensing sets the most stringent requirements on the errors of the **PSF** models. Cosmologists have developed formulations to relate the **PSF** errors expressed in terms of shape and size metrics [Mas+12] to the errors of the cosmological parameters of interest [Cro+13]. Therefore, the **PSF** model requirements are usually set up in terms of these shape and size metrics, as we have seen in subsection 2.6.3. The definitions of the ellipticity and size and the way to calculate them can be found in subsection 2.6.2.

We first measure the ellipticity and size parameters of each image from the sets of ground truth images $\{A_i\}_{i=1}^n$, and reconstructions $\{\hat{A}_i\}_{i=1}^n$. The shape measurements are carried out on super-resolved polychromatic images. Then, the results are used to build vectors as for example $\mathbf{e}_{1,2} = [e_{1,2;1}, \dots, e_{1,2;n}]$. Finally, the error metrics used are the **RMSE** over each shape and size metric that can be written as

$$\text{RMSE}(\mathbf{e}_{1,2}) = \text{RMS}(\mathbf{e}_{1,2} - \hat{\mathbf{e}}_{1,2}), \quad \text{RMSE}(\mathbf{R}^2) = \frac{\text{RMS}(\mathbf{R}^2 - \hat{\mathbf{R}}^2)}{\langle \mathbf{R}^2 \rangle}, \quad (5.27)$$

where the size metric is divided by its mean, as it is usually done in the literature.

5.5.2 Setup

To demonstrate our novel framework, we simulate a simplified **FOV** with 2000 star observations for training and 400 noiseless target stars for testing. All positions are randomly distributed in the **FOV**. To evaluate the impact of the number of stars in the **PSF** field on the performance, we build three subsets \mathcal{S}_1 , \mathcal{S}_2 , and \mathcal{S}_3 containing 200, 500, and 1000 stars, respectively. The stars are assigned to the subsets so that they verify $\mathcal{S}_1 \subset \mathcal{S}_2 \subset \mathcal{S}_3 \subset \mathcal{S}$, where \mathcal{S} is the data set with all the training stars. This choice allows us to provide a fair comparison and be less sensitive to randomness. The observations have a variable **Signal-to-Noise Ratio** (**SNR**) that is uniformly distributed in the range [10, 110]. The **SNR** definition we use is

$$\text{SNR} = \frac{\|\bar{I}\|_F^2}{\sigma_{\bar{I}}^2 N^2}, \quad \|\bar{I}\|_F^2 = \sum_{i,j=1}^{N,N} |\bar{I}_{i,j}|^2, \quad (5.28)$$

where $\bar{I} \in \mathbb{R}^{N \times N}$ is an observation corrupted with white Gaussian noise with standard deviation $\sigma_{\bar{I}}$, and the operation $\|\cdot\|_F^2$ denotes the Frobenius matrix norm.

We use the parameters of an optical model close to the *Euclid*'s VIS instrument model [Lau+11; Ven+20]. We consider a broad passband of [550, 900]nm, a telescope's focal length of $f_L = 24.5\text{m}$, a telescope pupil diameter of $M_D = 1.2\text{m}$, a squared pixel size of $\Delta = 12\mu\text{m}^2$, and an oversampling factor of $Q = 3$. The obscuration used is fixed for every position in the **FOV** and is built as a superposition of simple geometrical shapes, as seen in Venancio et al. [Ven+20, Fig. 7]. A moving average filter is applied to the obscuration to smooth the transitions between the obscured and illuminated parts to avoid aliasing. We set $n_\lambda = 20$ wavelength bins for the approximation in Equation 5.19. The dimensions used are $K = 256$ for the **WFE** space, $M = 64$ for the high-resolution pixel space, and $N = 32$ for the observation, or low resolution, pixel space.

The stars need **SED** information to build polychromatic observation. We randomly choose for each star one of the 13 different types of stellar **SED** templates from [Pic98] following the work in [KTC16]. Figure 5.5-4 presents the stellar **SEDs** used to build our polychromatic **PSF** observations. The **SED** data is considered to be a perfectly-known input to the **PSF** model. The ground truth (**GT**) **PSF** field is built using 45 Zernike polynomials with a d_Z of 2 for each mode. The polynomial coefficients are randomly chosen such that the total aberration at any position in the **FOV** is close to a nominal value. The total amount of aberrations can be characterized by the **RMS** value of the **WFE** map, and the nominal value is defined as 100nm. No other detector effects are considered in this chapter. The instrument's focal plane is constituted of a squared matrix of squared **CCD** chips or **CCD** mosaic [Cro+16]. This chapter does not consider any discontinuity between the different **CCD** chips. It assumes that the **CCDs** are perfectly aligned in the focal plane and share the same properties. The diffraction phenomena are considered the only sources of chromatic variation of the **PSF** field. Figure 5.5-5 presents four **PSF** examples from the simulated data set, where the observed stars and their corresponding **WFE** is shown.

The following models are compared:

- (i) *Zernike 15*: a model using the proposed differentiable framework, but without any data-driven features. Only the fixed features, Zernike polynomials, are being used. We use $n_Z = 15$, and $d_Z = 2$. In the context of this study, this model represents a parametric approach where its parameters were badly specified. As

¹This corresponds to 0.1 arcsecond




Figure 5.5-4 – Templates for the stellar SEDs used to build the polychromatic PSFs.




Figure 5.5-5 – Example of PSFs from the simulated data set. The first row contains images of observed stars, while the second row contains the corresponding WFE map used to generate each one.

a consequence, the complexity of the model is far away from the ground truth PSF field.

- (ii) *Zernike 40*: a model similar to (i), with $d_Z = 2$, but using $n_Z = 40$. In the context of this study, this model represents a parametric model where its parameters were slightly misspecified. This model is close to the ground truth PSF field. However, it lacks complexity as the number of Zernike polynomials is lower than the ground truth.
- (iii) *WaveDiff-original*: the model described in section 5.4.2 using $n_Z = 15$, $d_Z = 2$, and $d_{DD} = 5$ that corresponds to $n_{DD} = 21$. This model uses the same badly specified parameters of the fixed features as the model (i).
- (iv) *WaveDiff-graph*: a similar model as (iii) with a badly specified number of Zernike polynomials, using $n_Z = 15$ and $d_Z = 2$. However, this model uses the previously described graph constraint for the data-driven features. In order to perform a fair comparison we use $n_{DD} = 21$.
- (v) *WaveDiff-polygraph*: a similar model as (iii) and (iv), using $n_Z = 15$ and $d_Z = 2$. For the data-driven features, this model uses $d_{DD} = 3$ that corresponds to 10 features, and 10 more features corresponding to the graph constraint.
- (vi) *PSFEx*: The choice of the model parameters is based on the numerical experiment done in [Sch+20]. The model is built independently on each CCD. The spatial variations of the PSF, PSFVAR_KEYS, are a function of the position parameters, XWIN_IMAGE, YWIN_IMAGE. The sampling and the image dimension are adjusted for our experiments. We use a PSF_SAMPLING of 1. and 0.33, and a PSF_SIZE of 32,32 and 64,64, for reconstructions at the observation resolution and at 3 times the observation resolution, respectively. We use a PSFVAR_DEGREES of 3, the degree of the polynomial used for the spatial variation function, as we found that it gave better results than other values. See section 5.3 for a detailed description.
- (vii) *RCA*: the current state-of-the-art data-driven PSF model [Sch+20] which has been specially designed for the *Euclid* mission. The model is built independently on each CCD. We verified that the parameters used in [Sch+20] were the best-performing ones. Therefore, we use 4 features, or EigenPSFs, with a denoising parameter K of 3. We adjust the upsampling parameter D to match the two different reconstruction resolutions desired that in our case is 1 and 0.33.
- (viii) *MCCD*: a recently developed data-driven PSF model [Lia+21a] that combines ideas from the two previous models. PSFEx and RCA build a PSF model per CCD chip, using fewer stars. The MCCD model is able to build a PSF model in the full FOV taking into account the geometry of the instrument. We compared models with different parameters and selected the one that was the best performing. The global polynomial degree is set to 6 with a global denoising parameter, K_GLOB, set to 3. For the local part, the hybrid model is being used with 4 graph features and a degree of 2 for the polynomial features. The local denoising parameter, K_LOC, is also set to 3.

Our setup for testing PSFEx and RCA is different from Ngolè et al. [Ngo+16] and [Sch+20]. Here, the observed PSF field has a spatially varying SNR, and we simulate more realistic polychromatic PSFs. The observed PSF positions are uniformly distributed in the FOV. The studies [Ngo+16; Sch+20] used monochromatic PSFs

Table 52 – Hyperparameters of the different PSF models that use the optical forward model. All the models use a batch size of 32. The WaveDiff-graph and the WaveDiff-polygraph use a regularization parameter μ of 1×10^{-8} .

PSF model	Data set	m_1				m_2			
		$\eta_1^Z [\times 10^{-3}]$	η_1^{DD}	$N_{ep,1}^Z$	$N_{ep,1}^{DD}$	$\eta_2^Z [\times 10^{-3}]$	η_2^{DD}	$N_{ep,2}^Z$	$N_{ep,2}^{DD}$
Zernike 15	\mathcal{S}_1	5.0	-	40	-	1.0	-	40	-
	\mathcal{S}_2			30				30	
	\mathcal{S}_3			30				30	
	\mathcal{S}			20				20	
Zernike 40	\mathcal{S}_1	5.0	-	40	-	1.0	-	40	-
	\mathcal{S}_2			30				30	
	\mathcal{S}_3			30				30	
	\mathcal{S}			20				20	
WaveDiff-original	\mathcal{S}_1	10.0	0.1	30	300	4.0	0.06	30	300
WaveDiff-polygraph	\mathcal{S}_2			30	200			30	150
	\mathcal{S}_3			20	150			20	100
	\mathcal{S}			15	100			15	50
WaveDiff-graph	\mathcal{S}_1	10.0	0.4	30	300	4.0	0.2	30	300
	\mathcal{S}_2			30	200			30	150
	\mathcal{S}_3			20	150			20	100
	\mathcal{S}			15	100			15	50

sampled at 600 nm, requiring a super-resolution factor of 2 to be Nyquist sampled. In the current experiment, we use monochromatic PSFs from 550 nm to 900 nm to build polychromatic PSFs. This choice requires an integer super-resolution factor of 3 to be Nyquist sampled at all wavelengths. As a consequence of the different experimental setup, the performance results of PSFEx and RCA are expected to be different compared to Ngolè et al. [Ngo+16] and Schmitz et al. [Sch+20].

Additional parameters of the models, such as learning rates and the number of epochs, can be found in Table 52.

Data set adaptation for PSFEx, RCA, and M CCD

The PSFEx and RCA models build independent PSF models for each CCD chip. However, we have simultaneously simulated the PSF field in the full FOV. Therefore, dividing the FOV to respect the instrument’s geometry is necessary. In this case, we follow *Euclid*’s VIS instrument geometry [Cro+16]. The focal plane consists of a 6×6 CCD matrix, where each CCD is a 4096 pixel square. Depending on the star positions, we divide the FOV into 36 subsets. As the star positions are uniformly distributed over the FOV, the star density is approximately constant on each CCD. The input of the RCA model is shared with M CCD. The latter model continues with internal processing where all the CCDs are merged into a single data set with global coordinates. We have included the geometry and the coordinate transformations required by the M CCD model.

To create the PSF model with PSFEx, we need to process the simulations with SExtractor [Ber11]. For that, we create a full CCD image of 4096×4096 pixels with the postage stamps of the stars at their corresponding positions. Each star has its already defined varying SNR. Choosing the right parameters to make sure that all the stars are detected and extracted by SExtractor can be tedious, in particular for low SNR stars. We build a second mock CCD image that we only use for the detection. In this new mock image with low-valued pixels, we place high-valued pixels on the

Table 53 – Polychromatic test star reconstruction Err_{abs} and Err_{rel} at the observation resolution (x1) and at super-resolution (x3). The presented results were obtained with models being trained on the largest training star set \mathcal{S} .

PSF model	$Err_{\text{abs}} [\times 10^{-5}]$ (Err_{rel})	
	Resolution x1	Resolution x3
(i) Zernike 15	77.5 (10.6%)	19.0 (12.8%)
(ii) Zernike 40	32.6 (4.4%)	8.7 (5.9%)
(iii) WaveDiff-original	6.4 (0.86%)	1.9 (1.3%)
(iv) WaveDiff-graph	71.8 (9.8%)	19.0 (12.7%)
(v) WaveDiff-polygraph	6.9 (0.94%)	1.9 (1.3%)
(vi) PSFEx	69.2 (9.5%)	66.3 (43.0%)
(vii) RCA	39.6 (5.4%)	85.3 (55.5%)
(viii) MCCD	43.5 (6.0%)	97.7 (63.4%)

positions of the stars, thus forcing the detection of the stars. Then, the extraction of the stars is done using the first mock CCD image. Following this procedure, we ensure that all the stars are correctly detected and extracted by SExtractor. Finally, the star observations can be used to build PSF models with PSFEx.

5.5.3 Results

The results were obtained from the set of testing stars that were not used for training the different PSF models. The WaveDiff models were trained using an Nvidia Tesla V100 Graphical Processing Unit (GPU) of 32GB. The training time scales with the number of observations in the training set. However, the number of epochs used for the training also varies. The training of the WaveDiff-original model of the 2000-star data set took 11.5 hours, while the 200-star data set took 8.2 hours. Note that the hyperparameters, especially the number of epochs, were set to maximize the performance, not the speed.

Results: (a), Polychromatic errors

Table 53 summarizes the pixel reconstruction performance results on polychromatic images. There is a significant improvement of WaveDiff-original and WaveDiff-polygraph, compared to PSFEx, RCA, and MCCD. These three comparison models do not use any forward model and are blind to the physics involved in the inverse problem we are tackling. Considering the reconstruction at the observation resolution (x1), the WaveDiff-original model reaches 6 to almost 11 times lower absolute errors. The errors are decreased by a factor between 34 and 51 when considering the super-resolution task (x3).

The Zernike 15 model, or (i), underperforms with respect to the WFE-based models at the observation resolution. It is even poorer than the state-of-the-art models, thus highlighting the lack of representation of a reduced number of Zernike polynomials. If we remove the data-driven features from the WaveDiff models, we obtain a model equivalent to model (i). This difference in the performance between model (i) and the WaveDiff models remark the importance of the data-driven features in the WFE-based PSF model. It also shows the effectiveness of the data-driven features in adapting to the observations and generalizing to the target positions. Concerning the super-resolution task, the model (i) still outperforms the pixel-based




Figure 5.5-6 – Examples of learned data-driven eigenWFEs, or WFE features, of the WaveDiff-original model. These features correspond to polynomial feature weight functions. Units are in μm .

models. There is a huge performance gap between the models using a differentiable forward model and the ones that are building their models directly in the pixel space. The super-resolution is taken into account naturally when using the optical forward model. This handling of the super-resolution is done in the image formation process that is encoded in the forward model with the subsequent downsampling to obtain a PSF at the observation resolution. The observed performance gap underlines the importance of adding prior physical information to the inverse problem to solve a challenging task such as super-resolution.

Model (ii) is performing poorly with respect to WaveDiff-original and WaveDiff-polygraph, even though the number of Zernike polynomials, $d_Z = 40$, is close to the ground truth. This minor mismatch leads to large PSF reconstruction errors, which demonstrates the rigidity of parametric models. The WaveDiff-original model with only 15 Zernike polynomials does not suffer from this restriction, showing the flexibility of the data-driven component of the model.

It is important to compare the performance of the different WaveDiff flavours. There is a remarkable difference, around an order of magnitude, between the WaveDiff-graph and the original and polygraph flavours. This fact shows that *only* using the graph constraint for the weight function is not effective in building an appropriate WFE manifold for the PSF field, in contrast with the polynomial weight functions. Figure 5.5-6 presents some examples of learned data-driven WFE features for the WaveDiff-original model. The features show that some structure was learned, indicating that we are estimating a WFE manifold. The WFE manifold learning is a consequence of the optimization through the differentiable optical model used in the proposed framework. Note that no constraint has been applied to the data-driven WFE features. We found that the polynomial weight functions that are more constrained than their graph counterparts allow learning more structure into the data-driven features.

Figure 5.5-7 presents the reconstruction of a test PSF done by the WaveDiff-original model. The first three rows show the excellent reconstruction quality of the WaveDiff-original model for the different pixel scenarios. The fourth row shows important errors in the WFE reconstruction; nonetheless, the pixel reconstructions are very similar. This fact shows the richness of the WFE space and how two very different WFE maps can recreate particularly similar pixel PSF images.

Results: (b), Monochromatic errors

Figure 5.5-8 presents the pixel reconstruction errors as a function of wavelength. We only present the WaveDiff-original model as it is the best performing in the results from the objective (a). One can see that the errors are kept low over all the




Figure 5.5-7 – Visual reconstruction of a test PSF with the WaveDiff-original model. The first column shows the ground truth noiseless PSF, the second column the reconstruction of the PSF model, and the third column the residual between the ground truth and the model. **First row:** Polychromatic pixel representation at the observation resolution. **Second row:** Polychromatic representation at 3 times the observation resolution. **Third row:** Monochromatic pixel representation at 750 nm and at 3 times the observation resolution. **Fourth row:** WFE representation of the PSF, where units are in μm . We have removed the mean of the WFE and applied the obscurations. The first three rows show the absolute value of the residual, while the last one is just the residual.




Figure 5.5-8 – Test star reconstruction error as a function of wavelength for the WaveDiff-original model at the observation resolution. **Left axis:** Absolute error, or Err_{abs} , and is plotted with a solid line. The error bars represent the variance over the RMS errors of the different test stars being used for the calculation of Err_{abs} . **Right axis:** Relative error, or Err_{rel} , and is plotted with a dashed line.

instrument’s passband. The model successfully captures the chromatic variations of the ground truth PSF field. The mean relative error over the passband is 1.46%. It is important to highlight this result, as the data used to constrain the PSF model are noisy polychromatic PSFs. The loss function we use to optimize our model is built on the reconstruction of polychromatic PSFs. Estimating chromatic variations from polychromatic observations is a difficult ill-posed problem. There are several sets of monochromatic PSFs that can represent the same polychromatic PSF, as information is lost in the spectral integration of Equation 5.1. Ensuring that the monochromatic error is kept low is an essential step in the validation of the proposed PSF model. One can also notice the error bars in Figure 5.5-8, representing the standard deviation of the reconstruction error within the set of test stars. The variance in the error is kept considerably low with respect to Err_{abs} , the mean RMS error on the test set. A visual reconstruction example of a monochromatic PSF can be seen in the third row of Figure 5.5-7.

Results: (c), Errors as a function of the total number of training stars

The PSF modelling task becomes more arduous as we reduce the number of training stars used to constrain the PSF model. This fact is expected as there is less information available to build the PSF models. Figure 5.5-9 presents the pixel reconstruction relative errors, Err_{rel} , as a function of the number of training stars. WaveDiff-original and WaveDiff-polygraph maintain excellent performance even with the smallest training set with 10% of the stars in \mathcal{S} . These results show the robustness of the proposed models with a low number of observed stars. Both models can properly handle a variable number of training stars with different SNR. WaveDiff-original and WaveDiff-polygraph can incorporate the information present in the stars as the reconstruction

errors monotonically decrease when we augment the number of training stars. **PSFEx** shows a rather consistent performance with the variation of the number of training stars. This result can be explained by the fact that it is the most constrained model. As a consequence, it has difficulties to incorporate more information into the model, but, in turn, it is more robust to a variation in the number of training stars. **RCA** exhibits rather similar behaviour to **MCCD**, with a performance improvement as the number of stars increases. **RCA** is not able to produce a result with the smallest data set \mathcal{S}_1 . However, **MCCD** can still work with a low number of training stars. The results from the largest data set, \mathcal{S} , in [Figure 5.5-9](#) are the same as the ones from the left column of [Table 53](#).

We repeat the previous experiment with the reconstruction error being evaluated at super-resolution, and we obtain the results shown in [Figure 5.5-10](#). The performance in the 2000-star training data set, \mathcal{S} , coincides with the right column of [Table 53](#). One can observe the same trend as before, with a huge performance gap between the WaveDiff family, which uses an optical forward model, and the state-of-the-art models. One can notice that the errors of the WaveDiff-original and WaveDiff-polygraph models are not monotonically decreasing as we increase the number of stars. Nevertheless, the errors are always low, never over the relative 6% error. This result shows that, for a varying number of training stars, the proposed models are satisfactorily generalizing to test positions as well as capturing the complexity of the high-resolution PSF field. Between the state-of-the-art pixel-based models, **PSFEx** is the best one performing in the super-resolution reconstruction task. In the reconstruction at the observation resolution **PSFEx** is underperforming with respect to **RCA** and **MCCD**. One explanation of this behaviour is that in our experimental setup, we perform a 3-times super-resolution, which differs from the factor 2 used in [\[Ngo+16; Sch+20\]](#) as already mentioned in [subsection 5.5.2](#). As the problem becomes more challenging and ill-posed, complex models tend to degenerate if they are not well regularized. However, simpler models like **PSFEx** tend to be more robust and give consistent results, as their limited capacity regularizes them. This fact illustrates the importance of validating the reconstruction results at both resolutions.

Results: (d), WFE reconstruction errors

The advantage of validating the PSF models with simulations is that we have access to the ground truth WFE of the PSF field. [Table 54](#) presents the performance results on the WFE recovery for the models based on WFE. It is interesting to observe that all the models obtain high errors in the WFE recovery. This result illustrates the difficulty of recovering the WFE from degraded in-focus observations. The ground truth WFE field is more complex than what the PSF models can represent, as mentioned in [subsection 5.5.2](#), so a zero WFE recovery error was not possible to achieve.

Even though WaveDiff-original and WaveDiff-polygraph achieve an outstanding pixel reconstruction performance, they have large errors in WFE recovery. The previous results indicate that the WFE manifold these models estimate is far from the ground truth WFE manifold. However, the pixel PSF fields the different WFE manifolds represent are very close, and this is what matters to our PSF model. The WFE manifold estimated by our model in these experiments is an auxiliary product allowing our models to incorporate physical information, capture chromatic variations, and achieve good super-resolution performance. [Figure 5.5-7](#) presents an example of two similar PSFs being generated by two very different WFEs. The WFE recovery can be found in the fourth row.




Figure 5.5-9 – Polychromatic PSF relative reconstruction error, Err_{rel} , at observation resolution as a function of the total number of training stars in the FOV. The training data sets correspond, in ascending order, to $\mathcal{S}_1 \subset \mathcal{S}_2 \subset \mathcal{S}_3 \subset \mathcal{S}$, described in subsection 5.5.2. The results are computed with respect to the reconstruction of test stars. The set of test stars is maintained throughout the models and the training sets, allowing a truthful comparison. The result of RCA for the \mathcal{S}_1 (200 stars) data set is not shown as it is very poor.




Figure 5.5-10 – Polychromatic PSF relative reconstruction error, Err_{rel} , at three times (x3) the observation resolution as a function of the total number of training stars in the FOV. This experiment represents the super-resolution task.

Table 54 – Target stars’ WFE recovery errors with the PSF models being trained on the \mathcal{S} data set.

PSF model	WFE Err_{abs} [nm] (Err_{rel})
(i) Zernike 15	106 (136%)
(ii) Zernike 40	102 (130%)
(iii) WaveDiff-original	101 (129%)
(iv) WaveDiff-graph	108 (138%)
(v) WaveDiff-polygraph	111 (142%)

Table 55 – Moment-based metrics of the different PSF models trained with the \mathcal{S} data set.

PSF model	RMSE		
	$e_1[\times 10^{-2}]$	$e_2[\times 10^{-2}]$	$R^2/\langle R^2 \rangle[\times 10^{-1}]$
(i) Zernike 15	3.92	2.69	0.92
(ii) Zernike 40	2.12	1.38	0.97
(iii) WaveDiff-original	0.23	0.16	0.13
(iv) WaveDiff-graph	4.19	1.64	1.39
(v) WaveDiff-polygraph	0.24	0.14	0.16
(vi) PSFEx	4.38	4.25	14.8
(vii) RCA	4.61	4.17	36.0
(viii) M CCD	9.79	7.32	47.4

Results: (e), Moment-based metrics

Table 55 presents the performance results of the different PSF models using moment-based metrics based on the shape and the size of the PSF. The results are coherent with the pixel errors in the super-resolution column from Table 53. This observation is expected as the moment-based metrics are computed on the polychromatic super-resolved reconstructed PSFs. The error reduction achieved by the WaveDiff-original model with respect to RCA constitutes a *breakthrough in performance* for data-driven PSF modelling. The e_1 and e_2 errors are 20 and 26 times lower, and the size error, $R^2/\langle R^2 \rangle$, is 276 times lower. These results highlight the importance of the proposed model for current and future weak-lensing studies.

5.6 Discussion and conclusion

Discussion

In this section, we discuss several scientific choices, the utility of the present chapter, possible improvements, and extensions.

In subsection 5.4.4, we introduced a regulariser term in the loss function. One of the two best-performing models, WaveDiff-original, does not include any regularisation. We tried using a $\nu\ell_2(\Phi_\theta)$ regularisation term that constrained the total amount of aberration that is represented by the energy of the WFE map. The physical motivation was to favour less aberrated models, as we saw that the WaveDiff-original model converged towards a WFE manifold with more aberration than the ground truth. This result is illustrated by the WFE reconstruction in the fourth row of Figure 5.5-7.

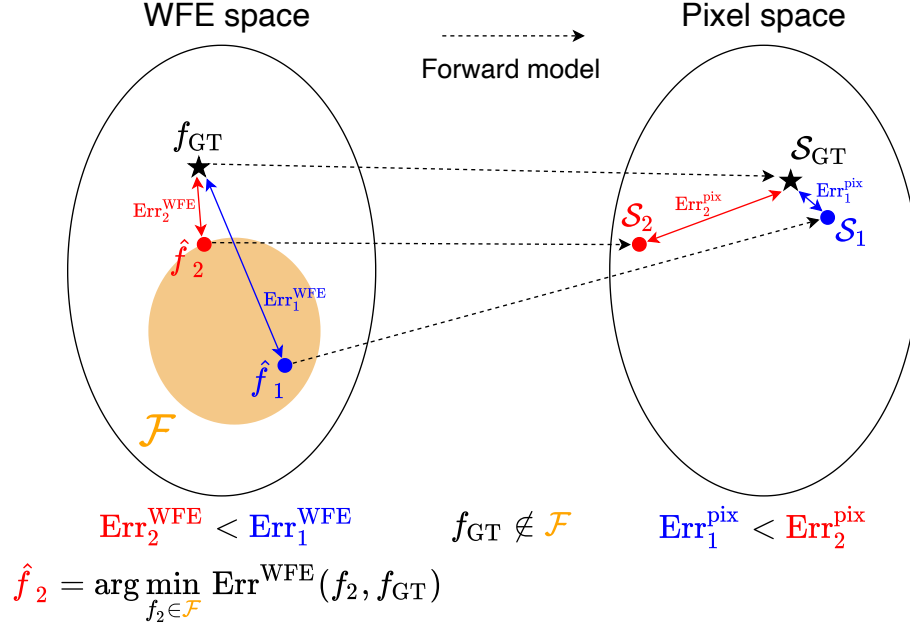


Figure 5.6-11 – Illustration of the second scenario of the WFE-based PSF modelling. The ellipse on the left represents the possible PSF fields defined in the WFE space. The ellipse on the right, the space of pixel PSFs representations. The forward model allows to go from a WFE-defined PSF model into a set of pixel PSF. The space of possible models given a set of WaveDiff hyperparameters is \mathcal{F} . If we assume that $f_{GT} \notin \mathcal{F}$, the closest PSF field in terms of WFE to the GT, \hat{f}_2 , is not the best choice in terms of pixel error.

After tuning the ν hyperparameter accompanying the regulariser and running the experiments, we noticed that all the pixel performance results were degraded with respect to the original model. However, the WFE reconstruction error decreased considerably. Even though the estimated WFE manifold got closer to the ground truth WFE manifold of the PSF field, the pixel performance worsened. This experience showcases that regularisers in the WFE space should be used with care. We identify two scenarios of the proposed WFE-based PSF modelling. The first scenario is aiming to reconstruct the *ground truth* WFE manifold. A typical example would be parametric models. The second scenario is aiming to reconstruct a *useful* WFE manifold. In this case, *useful* refers to a WFE manifold that has a similar pixel representation to the ground truth WFE manifold. The WaveDiff model we are proposing is an example of this scenario. Introducing a regulariser on the WFE would make sense on the first scenario. For example, if it is available *a priori* information about the ground truth WFE manifold, it could be used to build the regulariser. This situation could occur if complementary optical system measurements are available. The second scenario should be taken with more care as the WFE regulariser can bias the estimated useful WFE manifold. This bias represents a larger pixel space deviation even if both manifolds become closer in the WFE. Figure 5.6-11 illustrates the second scenario mentioned.

In the previous section, we presented certain advantages of the proposed data-driven approach with respect to parametric PSF models. It compared the Zernike 15 and Zernike 40 models with the WaveDiff-original model. However, it does not mean that parametric and data-driven models cannot be complimentary. The WFE output

of the parametric model can be used as input in the proposed framework. The model would shift from the second to the first scenario mentioned above, trying to reconstruct the ground-truth **WFE** rather than a useful representation. The parametric model output could be introduced to the data-driven model using a **WFE** regulariser. The task of the data-driven optimization would be to estimate the difference between the ground truth **WFE** and the **WFE** estimated by the parametric model.

This chapter finds similar performances for the WaveDiff-original and WaveDiff-polygraph models. The results presented in Table 53 showed a slightly better performance of the WaveDiff-original model. We therefore chose it to create the visual examples of Figure 5.5-6 and Figure 5.5-7. It is, however, hard to conclude which model is better with the different metrics used. The structure imposed into the Φ_θ function constrains the **WFE** manifold learned from the observations. The Φ_θ function that allows learning a better **WFE** manifold depends on the ground-truth **WFE**, which in turn depends on the telescope we are modelling. In the current setting, both models achieve great performances. A follow-up study would evaluate the models under high-fidelity optical simulations that generally involve costly ray-tracing techniques.

Even though the simulations used in this chapter are representative of a space telescope, there are several simplifications we have assumed. For example, we have considered that all the **CCD** chips are perfectly aligned in the focal plane and do not include any detector-level effects. These and other effects could be included in the proposed framework. Detector-level effects, such as guiding errors or charge transfer inefficiency, could be included in the optical forward model.

New flavours of the WaveDiff **PSF** model can be developed by taking advantage of our framework and redefining the function Φ_θ in Figure 5.4-2. Previous works that tried to leverage the power of modern neural networks into **PSF** models were blind to the physics of the inverse problem and built the model in pixel space [Jia+20b; Jia+20a; Jia+20c]. Our framework uses an end-to-end differentiable optical forward model, runs on GPU, and is entirely coded in TensorFlow [Mar+15]. This construction makes it an exciting candidate to develop *physically-motivated* neural network approaches to boost the **PSF** model.

Conclusion

This chapter introduces a novel data-driven modelling method of a point spread function field coined WaveDiff, which represents a paradigm shift with respect to current state-of-the-art data-driven methods. We propose to include a differentiable physics-based forward model that allows changing the modelling space from pixels to the wavefront. This change transfers most of the modelling complexity into the forward model and simplifies the building of the point spread function model. As the entire optical forward model is differentiable, we can use modern optimization techniques to estimate the model parameters in the wavefront space. The WaveDiff model is built using the TensorFlow framework and runs in GPU. We have verified that our model can learn from in-focus observations a *useful* wavefront-based manifold representation of the **PSF** field. We have not used special calibration data or wavefront information. The WaveDiff model is only constrained from noisy in-focus polychromatic under-sampled observations using the spectral energy distribution of each observed object. The proposed WaveDiff model is capable of super-resolving the **PSF** field and is, to the best of our knowledge, the *first data-driven PSF model* to effectively capture the wavelength dependence of the point spread function.

The proposed method was tested using a set of simulations of a space telescope which included a broad passband, different spectral energy distribution for the observed

objects, variable signal-to-noise ratios, and the undersampling of the observations. We compared our WaveDiff model against three state-of-the-art data-driven point spread function models, including [PSFEx](#), [MCCD](#), and [RCA](#), the current data-driven approach for the *Euclid* satellite. The WaveDiff model represents a *breakthrough in performance* obtaining 6 and 44 *times* lower reconstruction error than [RCA](#) at observation resolution and super-resolution (x3), respectively. Our proposed model obtains 0.86% and 1.3% relative reconstruction error at x1 and x3 the observation resolution, respectively. We validated that the WaveDiff model captures the chromatic variations of the point spread function with the monochromatic reconstruction error. The WaveDiff model achieves a mean relative error of 1.46% over the broad instrument's passband. We have further tested the WaveDiff model varying the number of observed stars used to constrain the model. Our model maintains good performance even when using the 10% of the original observed stars.

Using simulations to validate the model allows accessing and comparing the estimated and the ground truth wavefronts. We have found that the wavefront manifold estimated by the WaveDiff model and the ground truth are far from each other. However, and interestingly, both wavefront manifolds represent very similar pixel [PSF](#) fields, as was shown in the pixel errors above. These results support using the wavefront space as an auxiliary space to represent the pixel [PSF](#) field. It has not necessarily to be the objective to find the ground truth wavefront manifold, but a manifold that is *useful* for our objectives. These are defined in pixel space rather than in the wavefront space.

Finally, we have compared the different point spread function models using weak-lensing metrics such as ellipticity and size. Compared to the [RCA](#) model, the WaveDiff model reduces the ellipticity errors by a factor of 20 and 26 for each component, respectively. In addition, the size error obtained is 276 times lower than the one obtained by [RCA](#). These results highlight the utility and confirm the breakthrough in performance achieved by the proposed WaveDiff model.

In the spirit of reproducible and reusable research, all software, scripts, and analysis notebooks are publicly available at github.com/tobias-liaudat/wf-psf.

* * *
* *
*

Applying WaveDiff to *Euclid*

Chapter Outline

6.1	Introduction	156
6.2	Introducing prior information to the WaveDiff model	156
6.2.1	Motivation	156
6.2.2	Model description	157
6.2.3	Simulations and optical prior	158
6.2.4	Numerical experiment	161
6.2.5	Conclusion	165
6.3	Wavefront error resolution study	165
6.4	Using more exposures for the PSF model	167
6.5	Impact of SED errors in the PSF generation	170
6.5.1	Types of errors	171
6.5.2	Numerical experiments	172
6.6	Studying SED interpolation in the PSF generation	175
6.6.1	SED interpolation	175
6.6.2	Numerical experiments	176
6.7	SED errors in PSF modelling	178
6.7.1	Numerical experiment	178
6.7.2	Conclusion	179
6.8	Improving the WaveDiff optimisation procedure	180
6.8.1	Motivation	180
6.8.2	Projecting data-driven features	181
6.8.3	Introducing the WFE projections to the optimisation procedure	187
6.8.4	Numerical experiments	189
6.8.5	Results	189
6.8.6	Conclusions	193
6.9	Conclusion	193

THE previous chapter introduced a novel framework for building PSF models that include the WaveDiff model. The results presented were encouraging. However, unaddressed issues arise when we get closer to a real *Euclid* scenario. In this chapter, we will go through several studies addressing many of these issues and proposing extensions to the WaveDiff model.

6.1 Introduction

The WaveDiff PSF model presented in [chapter 5](#) constitutes a promising approach for the *Euclid* mission. Nevertheless, transitioning from a good performance with simulations to a good performance with real data is far from easy. We were able to verify this with the MCCD PSF model used for CFIS, described in [chapter 3](#) and [chapter 4](#). Even though the ground-based telescope represents a much simpler task than *Euclid*, we have seen the difficulties arising when dealing with real data.

This chapter addresses several issues emerging when *Euclid* observations are available. We intend that the studies presented in this chapter will smooth the transition from simulated to real data for the WaveDiff model. We have also proposed improvements to the model that will later be useful. We continue by presenting the list of issues that we will address in this chapter:

- Prior optical information about the telescope will be available. This information can come from high-fidelity simulations or complementary observations in space. The prior is expected to have some degree of error. How can we exploit this prior information to improve the WaveDiff model?
- In the experiences in [chapter 5](#) we assumed the same WFE dimension to generate the observations, the GT model, and for WaveDiff. We study how the dimension of the WFE affects the modelling of the PSF when the observations are generated more realistically.
- In [chapter 5](#), we have seen that the WaveDiff model is robust when using a fraction of the nominal number of stars per exposure. The stability of *Euclid* might allow using subsequent exposures and therefore increase the number of stars to constrain the PSF model. We study how the WaveDiff model exploits information from a denser stellar field.
- We have considered, in the previous chapter, that the input SED information was flawless and had used 20 bins to generate the observations and for the WaveDiff model. In a real scenario, the observations are generated without any spectral discretisation. In addition, the real SED information comes in a reduced number of bins containing errors. We study how spectral discretisation affects the modelling of the PSF and ways to improve the results.
- Previous results showcased the ability of the WaveDiff model to estimate a useful WFE representation that allows us to obtain a low pixel error. However, suppose we are in a new scenario where the parametric part of the WaveDiff model can reproduce the GT WFE field. In that case, the current WaveDiff model is not able to estimate this GT WFE. It is indeed a phase retrieval problem. Nonetheless, in this new scenario, the solution to the phase retrieval problem is also the global minimum of our PSF modelling problem. We study how the WaveDiff model can deal with the problem in this new scenario.

6.2 Introducing prior information to the WaveDiff model

6.2.1 Motivation

In [chapter 5](#), each time we trained a PSF model, we started without any knowledge of the optical aberrations, i.e. the WFE. We had information about the optical system that we used to build the optical forward model. However, we had no information

about the WFE distribution. The high-quality design of *Euclid* will make the PSF field very stable, i.e. a high correlation between consecutive exposures. Nonetheless, this stability is not enough to allow a PSF model trained a single time to be the solution for the expected WL analysis. We still need to train the model for each epoch, but we can exploit the current information we have about the PSF field.

We can profit from several sources of prior optical information. A first option would be to estimate the model in a crowded stellar region with a higher constraining power. The trained model can then be used as prior optical information. Another option is to use the trained parametric model from *Euclid* and exploit their WFE information. An alternative and the most interesting one is to profit from the *phase diversity* calibration observations that are expected to be conducted regularly with *Euclid*. This alternative provides us with very rich observations with much wavefront information. Phase diversity observations consist in observing the same stellar field with different amounts of defocus, e.g. $(-\Delta f, 0, -\Delta f)$. The defocusing of the telescope allows exposing more wavefront information. One of the first and more notable applications of phase diversity techniques was the determination of Hubble Space Telescope (HST) aberrations [Fie+93]. See [MBJo6] for a review on phase diversity. A final source of optical information is the on-ground characterization of optical components. However, this is only useful for elements which will not vary considerably in the telescope’s journey to space. A notable example is the spectral characterization of the dichroic element in *Euclid*. A common trait between the aforementioned sources is that the optical information is defined in the WFE space. This fact was yet another motivation for developing a physically based data-driven model.

No matter the source, the optical information we mentioned will have errors. Even if they are small, they can represent large pixel deviations due to the sensibility of the wavefront. Consequently, we cannot directly use them to compute the PSF, and we need to incorporate them into our framework. We want our data-driven model in the WFE space to be able to estimate the difference between the prior WFE information and the ground truth WFE. We will use a set of complex *Euclid* optical simulations to generate the observations to study this scenario. Then, we will generate different optical priors by adding different amounts of error to the simulations. We continue by evaluating the performance of the model under the different errors in the optical prior.

6.2.2 Model description

The addition of the optical prior into the model is straightforward. We need to rewrite the model’s WFE Equation 5.6 and add a new term with the optical prior contribution. Rewriting Equation 5.6 we obtain

$$\Phi_{\theta}(x_i, y_i; \lambda) = \Phi^Z(x_i, y_i; \lambda) + \Phi^{\text{DD}}(x_i, y_i; \lambda) + \Phi^{\text{prior}}(x_i, y_i; \lambda). \quad (6.1)$$

We assume that the prior term is a *black box* provide us with the expected WFE contribution for a FOV position and wavelength. The chromatic dependency of the optical prior does not need to follow the diffraction dependence and can add other chromatic dependences. In that case, the data-driven model, Φ^{DD} , should be modified so that it can also correct possible errors in the prior’s chromatic variations.

In this section we do not consider chromatic variations other than diffraction, so that we can reuse the definitions of Φ^{DD} and Φ^Z from chapter 5. Consequently, we have $\Phi^{\text{prior}}(x_i, y_i; \lambda) = \Phi^{\text{prior}}(x_i, y_i)/\lambda$.




Figure 6.2-1 – The figure (a) shows the simulation’s random positions in the FOV for the train and test datasets. The figure (b) shows the spatial distribution of the total amount of aberrations in the FOV.




Figure 6.2-2 – Histogram of the total amount of aberrations in WFE RMS of the optical simulations with and without Surface Errors (SFE).

6.2.3 Simulations and optical prior

The *Euclid* mission is not yet in space, so we have to work with simulations. Dr Pierre-Antoine Frugier, one of *Euclid* experts in optics, provided us with high-quality ray-tracing simulations which include all the optical complexity of the telescope except for the dichroic’s chromatic effect.

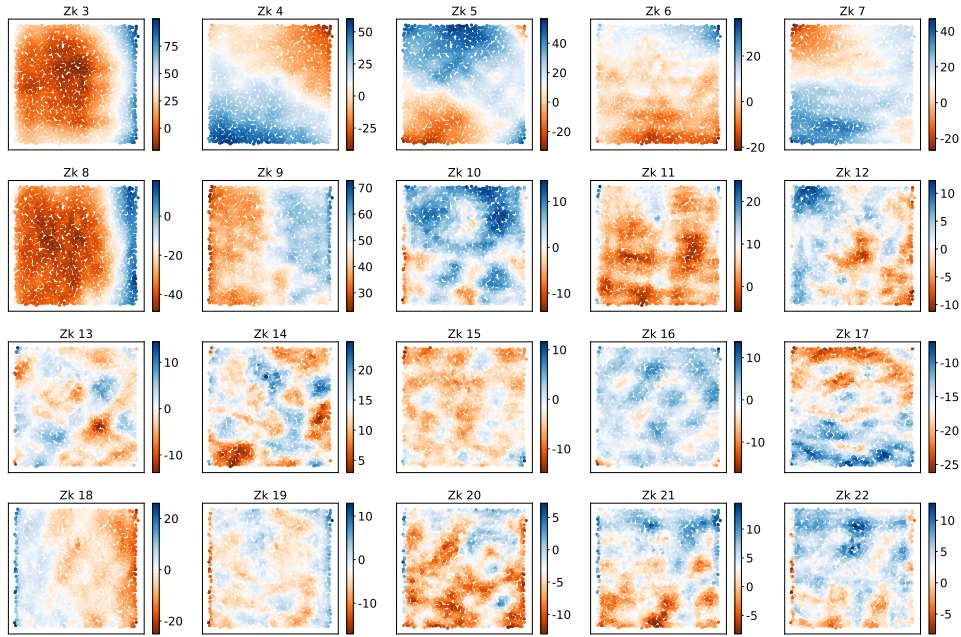


Figure 6.2-3 – Spatial distribution of the Zernike coefficients from order 3 to 22 from the *Euclid*'s optical simulations. Units are in nm RMS.

Simulations

We randomly selected 2600 FOV positions, where 2000 are used to train the model, and 600 are saved for testing. The ray-tracing simulations provide us with the WFE representation of the aberrations for these positions in the FOV. We chose to project the memory intensive WFE maps into Zernike decompositions of order 66, which facilitates its handling. The high Zernike order guarantees a good representation of the ray-tracing computed aberrations. Figure 6.2-1a illustrates the training and testing positions in the FOV with arbitrary units in $[0, 1000] \times [0, 1000]$. Figure 6.2-1b shows the spatial variations in the FOV of the total amount of aberrations that is expressed as WFE RMS. The mean amount of aberration in the field is ~ 88.4 nm RMS. Figure 6.2-2 shows a histogram of the amount of aberrations of the different FOV positions for two datasets. The first one is in orange with Surface Errors (SFE), also known as mirror polishing errors, and the second one is in light blue without SFE. We will work with first dataset, with SFE, but Figure 6.2-2 allow us to observe the impact of the SFE in the WFE aberrations.

We can observe the high complexity of the simulations by observing the spatial variation of each Zernike coefficient. Figure 6.2-3 shows the low order Zernikes, from order 3 to 22, and Figure 6.2-4 shows the Zernikes from order 23 to 42. We can observe that the value of the coefficients decreases as we increase the order. Nevertheless, the values are not negligible and cannot be ignored in the analysis.

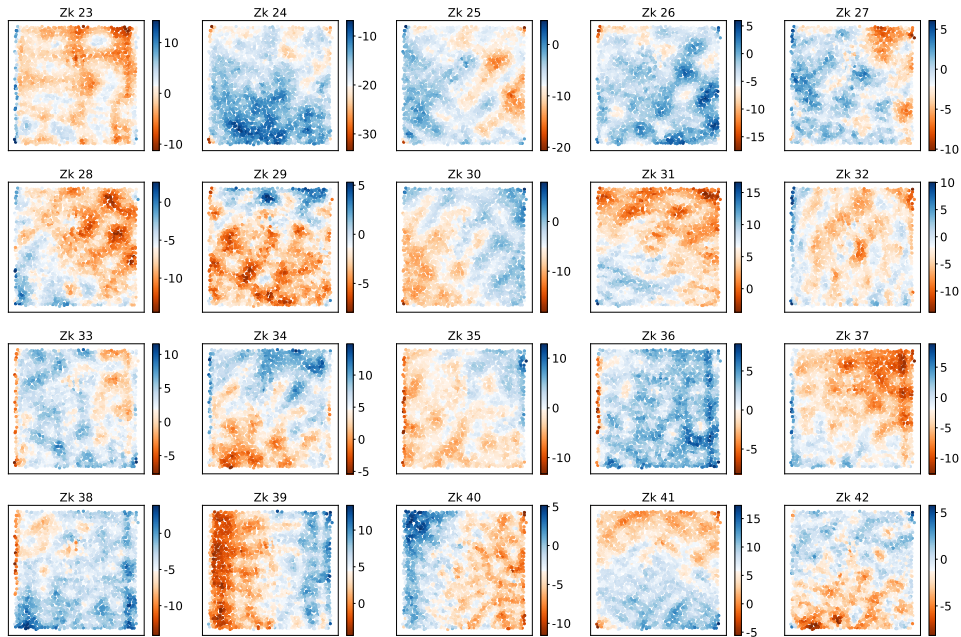


Figure 6.2-4 – Spatial distribution of the Zernike coefficients from order 23 to 42 from the *Euclid*'s optical simulations. Units are in nm RMS.




Figure 6.2-5 – Spatial distribution of the realisation of the total aberration error map, $\Delta\Phi^{\text{Err}}$, for the 10nm RMS case.




Figure 6.2-6 – Spatial distribution of some example Zernike coefficients from the error map with 10nm mean WFE RMS realisation. Units are in nm RMS.

Optical prior

In order to build the different optical priors, we need to generate error maps that will be added to the simulated WFE as follows

$$\Phi^{\text{prior}}(x_i, y_i) = \Phi^{\text{sims}}(x_i, y_i) + \Delta\Phi^{\text{Err}}(x_i, y_i), \quad (6.2)$$

where $\Delta\Phi^{\text{Err}}$ represents the WFE error map, Φ^{sims} the WFE from the *Euclid* simulation, and Φ^{prior} the prior that will be used in the WaveDiff PSF model. Similarly to subsection 5.5.2, we use a random 2D position polynomial variation for each Zernike coefficient. The polynomial degree is 4, and the Zernike order is the same as in the simulations, i.e. 66. Each Zernike order varies with its corresponding polynomial. However, we add a constraint so that the total amount of aberrations in any single position should be too far away from the chosen mean aberration level of the error map. We chose six levels of aberrations that we quantify with the overall mean WFE RMS. We go from almost no error to a high error contribution, and we use [1, 2, 5, 10, 20, 40] nm RMS. The mean WFE RMS of the optical simulation is ~ 88.4 nm. Then, the error maps expressed relative to the simulation's mean WFE RMS are written as [1.13, 2.26, 5.66, 11.31, 22.62, 45.25]%, respectively. Figure 6.2-5 presents the spatial distribution of the total aberration for the 10nm random realisation. Figure 6.2-6 shows examples of spatial variations for six Zernike coefficients from the 10nm RMS realisation.

6.2.4 Numerical experiment

In this section, we conduct a numerical experiment to compare the performance of the WaveDiff PSF model using the optical priors with different error levels. We focus on the best-performing model from the chapter chapter 5, which corresponds to a WaveDiff-original model, detailed in subsection 5.5.2. In this experiment, we change the order of Zernike polynomials to $n_Z = 66$ to match the number of Zernikes from the prior and the simulations. We compare three versions of this model, one where




Figure 6.2-7 – Relative pixel reconstruction **RMSE** at 3 times the *Euclid* resolution for different levels of error in the prior. The *x*-axis shows the amount of error added to the simulations to form the prior expressed as a percentage of the mean **WFE RMS** of the simulations.

we only use the data-driven part, one with only the parametric part, and another one using both parts (*complete*) as it was done in the WaveDiff-original model from [chapter 5](#). The three model versions use the optical prior, and their differences reside in the *trainable* part of the model. All the following results are obtained from the test set of stars that was not used for training the models.

Results

[Figure 6.2-7](#) shows the pixel **RMSE** error on the super-resolution scenario. On top of the three models mentioned above, we draw the error we would get if we only use the imperfect optical prior. We can see that the pixel error done by the prior increases monotonically with the energy of the error map added to the prior, as expected. The prior-only curve is the starting point of the **PSF** models, and the difference in performance with this curve indicates the gain or loss the model obtains with respect to the prior. The parametric-only model is only improving the results when the performance of the prior is already poor. This model cannot achieve a relative error lower than 3% even if the prior's error is already close to 1%. The difficulty of optimising the parametric model makes the performance drop when the prior is already performing well.

The two WaveDiff models containing a non-parametric, or data-driven, part have remarkable performances. More precisely, the model only using the non-parametric part is consistently obtaining better results. Both models can maintain a low error even when the prior performs poorly. The only case where both models are degrading the prior's performance is when the prior is the closest to the ground truth, i.e. the simulations.

[Figure 6.2-8](#) and [Figure 6.2-9](#) present the **PSF** models' performances in terms of ellipticity and size, respectively. The WaveDiff non-parametric is obtaining notable




Figure 6.2-8 – Ellipticity errors expressed relative to the *Euclid* requirement, which for each ellipticity component corresponds to 2×10^{-4} .

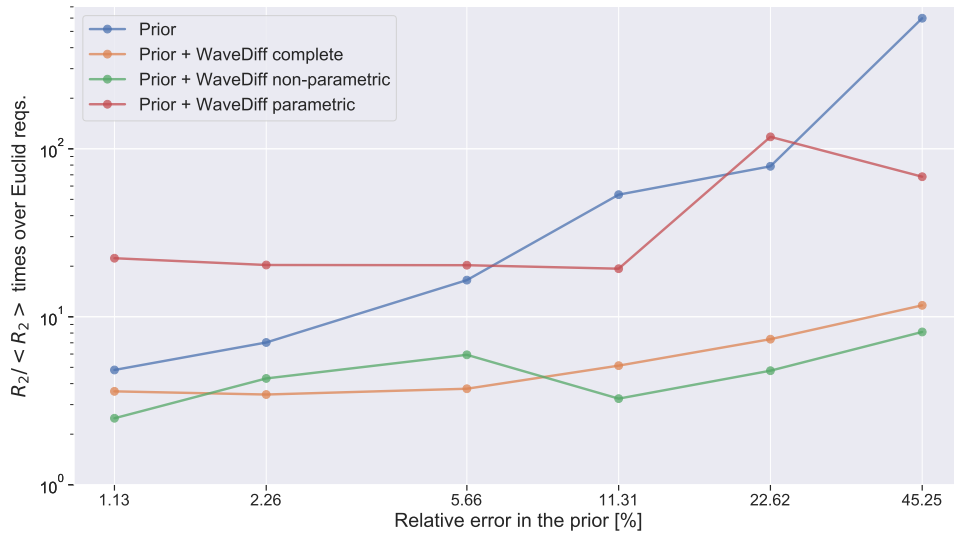


Figure 6.2-9 – Size error expressed relative to the *Euclid* requirement, which corresponds to 1×10^{-3} .




Figure 6.2-10 – Relative pixel reconstruction RMSE as a function of wavelength for the prior and the WaveDiff model using only the non-parametric, or data-driven, part. The error done by the prior is shown with dashed lines. We have omitted the curves for higher prior errors to keep the limit of the y -axis low.

results. In the three metrics, the model is never over 10 times the requirements. The figures show how the prior goes over two orders of magnitude over the requirements when the error introduced is close to 50%. However, the non-parametric model learns a useful correction for the inaccurate prior that brings it back under one order of magnitude.

The monochromatic error is presented in Figure 6.2-10, where we only plot the curves under 1.1%. We can see that the monochromatic error is stable throughout the *Euclid*'s bandpass. Only the two less flawed priors are below this arbitrary 1.1% threshold, while the WaveDiff model can obtain this performance with the 11.3% prior.

Figure 6.2-11 shows the WFE reconstruction error. The parametric-only model consistently degrades the WFE performance of the prior, driving the WFE representation away from the ground truth. The WFE allow us to note that the behaviour of the complete model changes considerably with respect to the non-parametric one. Figure 6.2-7 showed similar pixel performance of both models. However, in the WFE space, the complete model converges far from the prior while the non-parametric model converges closer. The latter model is the only one able to decrease the WFE error for some priors. The previous result can be explained by noting that the complete model starts with optimising its parametric part. This optimisation is hard with a non-convex optimisation landscape full of hard-to-escape local minima. The chosen first-order optimisation method converges in a point far away from the ground truth, which is deduced from the parametric model result in Figure 6.2-11. The complete model optimisation continues with the non-parametric part, which finds a WFE representation that gives a low pixel error, as seen in Figure 6.2-7. However, the WFE representation is far from the ground truth as its starting point, the parametric model optimisation result, was already far away. By contrast, the non-parametric-only model, with a presumed smoother optimisation landscape, can converge to a WFE representation closer to the ground truth.




Figure 6.2-11 – Relative WFE reconstruction error for different level of error in the prior.

One may notice that the prior’s relative WFE error does not match the x -axis. This mismatch is because the x and y axis represent different things. The x axis represents the relative total error’s WFE energy introduced to the prior with respect to the total simulation’s WFE energy. However, the y axis represents the relative mean prior’s error with respect to the simulation. As the prior is built from the sum of the simulation and the error terms, the x axis would match the prior error if the simulation and the error terms were orthogonal in every point of the FOV, which is not the case in our experiment.

6.2.5 Conclusion

We showed that the WaveDiff model could work with imperfect prior optical information and exploit it to improve the model’s performance. We used *Euclid*’s optical simulations that include all the optical complexity of the telescope except for the dichroic plate. Imperfect priors with different levels of error were generated from the simulations and fed to the model. Numerical experiments showed that the WaveDiff model achieves remarkable performances and can improve the prior performance even if it contains a substantial error. By only using the data-driven part of the WaveDiff model, we can achieve consistently better performances. Results suggest that the optimisation of the parametric part is converging to a WFE representation further away from the ground truth even if the prior is close to the ground truth. This result illustrates the difficulty of optimisation of the Zernike-based part from the observed in-focus stars. Consequently, with the current optimisation strategy and using an optical prior, only using the WaveDiff non-parametric part is advisable.

6.3 Wavefront error resolution study

In the previous sections and in chapter 5, we have assumed a fixed dimension, K , with a value of 256 for the WFE representation. We used the same WFE dimension K

to generate the *GT* observations and for the WaveDiff model. Nonetheless, the real observations from *Euclid* will not have a sampled wavefront. The discretisation of the data comes in the detector’s pixels when the wavefront reaches the focal plane. In our modelling *WFE* space, which corresponds to the wavefront at the output pupil before the diffraction effect is taken into account, the *WFE* is still continuous. In this section, we will study the impact of using a continuous *WFE* to generate the observations but a discrete *WFE* to model the *PSF* using WaveDiff.

We are interested in studying how the chosen WaveDiff model’s *WFE* dimension affects its performance when trained with observations generated with a continuous *WFE*. Using a continuous *WFE* to generate the data would be ideal, but it is impossible in practice. Therefore, we will set a very high *WFE* dimension so that the *WFE* used to generate the observations will be closer to a continuous function.

Numerical experiment

We simulate a *PSF* fields with the same parameters as in subsection 5.5.2. However, we change the *WFE* dimension K from 256 to 4096. The chosen *WFE* resolution requires computing *FFT* over matrices that have between $\sim 11500^2$ and $\sim 18800^2$ elements depending on the wavelength, due to the required zero-padding described in subsection 5.4.3. We train the WaveDiff-original model, using the hyperparameters of subsection 5.5.2 except for the *WFE* resolution. We use three k dimensions of 64, 128, and 256, which will resume in three models. The dimension K cannot be increased to 512 or beyond as the *GPUs* runs out of memory. The observations, or training stars, are taken from the *PSF* field simulated using $K = 4096$. This last simulation was done in a *CPU*-version of the code that can handle larger *WFE* dimensions at the expense of a much longer computing time and the loss of the automatic differentiation capability. We perform 9 independent model trainings and show the median of the realisations.

Using a higher *WFE* dimension for the model impacts the computing time of the model, as the *FFT* are computed over larger matrices. Therefore, we also compare the computing time of the models as a function of K . The *GPU* used to train the models is an Nvidia Tesla V100 of 32GB.

We also want to quantify the pixel error due to the *PSF* generation with different *WFE* dimensions. This error does not require any training of the WaveDiff model. It does require the generation of another star dataset using the same positions, *SEDs*, and *PSF* field parameters as the one before, but using another *WFE* dimension. We use a K of 256 to generate this dataset for comparison.

Results

Figure 6.3-12 shows in the left axis the relative pixel error of the different models as a function of their *WFE* resolution with the dimension K . The errors are shown for the *PSF* reconstruction at the *Euclid* resolution and the *Super Resolution* (*SR*). The errors decrease as the *WFE* dimension increases. However, the *Super Resolution* (*SR*) error difference between the models with K of 128 and 256 are close. These errors are considerably smaller than the one using a K of 64, which is strongly degraded. The errors at the *Euclid* resolution are consistent with the ones at *SR* and follow a similar trend.

The dotted lines from Figure 6.3-12 show the error between the two datasets and are of 0.52% and 0.61% for the *Euclid* resolution and the *SR*, respectively. The *PSF* generation errors are considerably lower than the model’s errors. However, these errors




Figure 6.3-12 – The left axis shows the relative pixel RMSE at one and three times the *Euclid* resolution as a function of the modelling WFE dimension. The right axis shows the computing time required to train the corresponding WaveDiff model. The dotted lines represent the error between two datasets that share all their properties except that they were generated with WFE dimensions of 4096 and 256. The blue and green lines correspond to errors computed at three and one times the *Euclid* resolution.

represent the minimum errors done if the PSF model can perfectly estimate the WFE field and are not negligible.

As expected, the computing time increases with the WFE dimension. The time increase is logarithmic, as it looks linear on an x -axis with a base-2 logarithmic scale.

Conclusion

We have studied the impact of using a high WFE dimension to generate the observations and how the WaveDiff model's WFE dimension affects its performance.

The more realistic observations, with a high WFE dimension K , add a PSF generation error, of about 0.61% in the super-resolution scenario. The WaveDiff modelling error is higher as it achieves a median error of 2.62%. The WaveDiff error is higher than the one obtained in chapter 5, indicating that the slight difference between PSFs generated with different WFE dimensions affects the training of the model.

Our results indicate that reducing the WFE dimension of the WaveDiff model to 64 degrades the performance considerably. The model's performance using a dimension of 128 is lower but close to the performance when using 256. Going above the 256 dimension is not an option due to memory limitations of the GPU. In addition, the flattening of the performance curves indicates that it will not improve the performance much. Therefore, depending on the most valuable criteria chosen, performance or time and memory footprint, we can select between a resolution of 128 or 256.

6.4 Using more exposures for the PSF model

This section will study the impact of increasing the number of exposures used to train the PSF. In chapter 5 we have studied the performance of the WaveDiff model if we reduce the number of training stars from the nominal number of 2000 stars per

exposure. However, we have not studied if the performance of the WaveDiff model improves, stagnates or deteriorates if we use more than one exposure. We will assume that the space telescope remains unchanged when using four consecutive exposures or less, which means having at most 8000 stars. This assumption implies that we are increasing the number of stars in the field of view we are using to constrain the model.

Setup

To study the impact of increasing the number of exposures, we will use a similar PSF field as in [chapter 5](#), more specifically [subsection 5.5.2](#). The PSF field WFE variations will be the same, but we modify the optical forward model. The only hyperparameters we change are the WFE dimension K , to 128 (instead of 256), and n_λ , the number of bins in the spectral integral discretisation, to 8 (instead of 20). The motivation for these changes is to reduce the memory footprint of the model. The GPU used for running our models is memory limited, and increasing the number of stars increases the memory footprint. More intelligent ways to implement the model and the dataset to reduce the memory usage exists. However, we have not followed them due to time constraints.

We use five training datasets containing 8000, 6000, 4000, 2000, and 1000 stars per exposure. Following the previous chapter, each dataset is a subset of all the datasets that are larger than it.

We compare two WaveDiff models based on the WaveDiff-original model that was best performing in [chapter 5](#). The first model uses the same hyperparameters, $n_Z = 15$, $d_Z = 2$, and $d_{DD} = 5$. For the second model we increase d_{DD} to 10. The motivation is to increase the capacity of the WFE model, as we have more data available. We run the models 5 times on each dataset. We then present the mean and median result of the experiment repetitions for the x1 and x3 *Euclid* resolution, respectively. The median for the super-resolution task is used to avoid the influence of an outlier repetition. The error bars on the plots are the standard deviation of the repetitions. The results presented are always from the test set of stars that is shared by all the training datasets.

Numerical experiment

[Figure 6.4-13](#) shows the relative pixel RMSE at the *Euclid* resolution. We observe a general trend which decreases as we increase the number of stars. However, there is considerable stochasticity in the performance, shown by the error bars representing the standard deviation of the repetitions. The relative pixel error at three times the *Euclid* resolution is presented in [Figure 6.4-14](#). In this case, we observe that the simpler model, with $d_{DD} = 5$, seems to perform better, except for the 6000-star dataset that seems to be an outlier. The errors are decreased by increasing the number of stars, but they tend to flatten. Considering the model with $d_{DD} = 5$, if we compute the relative performance gain of using 2 and 4 exposures with respect to a single exposure, we obtain 15% and 20%, respectively. [Figure 6.4-14](#) shows that the gain is considerable until the 4000-star dataset, and then the gain reduces. We have assumed that each exposure has 2000 stars. The performance values seen in [Figure 6.4-13](#) and [Figure 6.4-14](#) are lower than the ones in [section 5.5](#) because we slightly changed the optical forward model to make it less memory intensive. Consequently, these results cannot be quantitatively compared to the ones in the previous chapter.

We wonder how using more stars to constrain the model impacts the WFE reconstruction. We are maintaining the scenario from [section 5.5](#), meaning that a zero WFE error is no achievable. Using a more complex model, like $d_{DD} = 10$, and




Figure 6.4-13 – Relative pixel RMSE at the *Euclid* resolution. The error bars represent the standard deviation of the experiment repetitions.




Figure 6.4-14 – Relative pixel RMSE at three times the *Euclid* resolution. The error bars represent the standard deviation of the experiment repetitions.




Figure 6.4-15 – Relative WFE RMSE. The median results and the best performance of the repetitions are drawn in solid and dashed lines, respectively. The error bars represent the standard deviation of the experiment repetitions.

more stars might lead to an improvement in the WFE recovery or different behaviour. Nonetheless, Figure 6.4-15 shows us the relative WFE RMSE, and we observe that both models are still doing a big WFE error for all the number of training stars. By observing the best performance, we see that in none of the repetitions, we obtain a WFE error under 100%.

Conclusion

To conclude, the WaveDiff model can exploit the information from a more dense stellar field and improve the performance. However, the gain decreases as we go over 4000 stars for training the model and starts to stagnate. We tried two models, the WaveDiff-original model from chapter 5, and one with increased complexity in its data-driven part. The simpler model is the best-performing one. The more complex model cannot exploit the information from the dense stellar field to build a more complex WFE manifold. By observing the WFE recovery error, we see that the behaviour of the models is similar to the one observed in chapter 5.

Finally, the number of exposures we recommended using to constrain the model will depend on the available computing resources and time. Using more stars is not costless, as training the model requires more time and resources. The best compromise between performance and time would be using two exposures. This choice gives us around 4000 stars, and a gain of performance of about 15% in pixel error at three times the *Euclid* resolution.

6.5 Studying the impact of stellar SED errors in the PSF generation

In the previous sections, we have defined one discretisation of the spectral integration in the instrument's bandpass, defined with the parameter n_λ in Equation 5.19. We have

used the same discretisation in the optical forward model to generate the observations as in the WaveDiff model. This choice assumes a simplification of the generation of the observations. Real SEDs have details as the ones we can see in Figure 5.5-4. Once the satellite is launched, the observations we will obtain from *Euclid* will have continuous spectral integration. The observations' discretisation comes at the pixel level, and as we have seen in the general observational model from Equation 2.27, the spectral integration is continuous. We also assumed that the input SED information given to the model was flawless, which will not be the case. This section will study the impact of the spectral integration's discretisation and the errors in the input SED information to the PSF generation.

6.5.1 Types of errors

Binning errors

The SED information for *Euclid* will mostly come from photometric in-ground observations of complementary surveys, for example, UNIONS/CFIS in the Northern Hemisphere. The number of filters available is limited; thus, the number of photometric bands will be limited too. For example, Figure 2.2-8 shows the filters available for CFHT. The stellar photometric sky coverage will come from different telescopes with different properties and filters. We assume that we dispose of 8 bins in the SED discretisation¹, that have equal length, are disjoint, and that cover the entire VIS passband. The errors coming from the discretisation of the SED will be named *binning errors*.

In practice, we do not dispose of a continuous function of the SED, not even for simulation. However, we do dispose of SED templates with a high resolution. The SED information can be seen in Figure 5.5-4, where the stellar templates have a bin resolution of 1nm, i.e. 350 bins used for the *Euclid* passband. We will consider these templates as our Ground Truth (GT) SEDs. We recall the definition of the stellar SED binning from Equation 2.29 in the following equation

$$\text{SED}_{b^k}(\lambda_k) = \frac{1}{Z_{n_\lambda(b)}} \int_{b_0^k}^{b_1^k} f_{(x,y)}(\lambda) d\lambda, \quad (6.3)$$

where $f_{(x,y)}(\lambda)$ is the continuous stellar spectral dependency of an element in position (x, y) , $Z_{n_\lambda(b)}$ is a normalisation factor, $b^k = [b_0^k, b_1^k]$ is the k -th wavelength bin with a length of $|b^k| = b_1^k - b_0^k$, and λ_k is the centre of the b^k bin. Considering the $\{b^k\}_{k=1}^{n_\lambda(b)}$ bins to be the ground truth SED we generate broader bins $\{B^k\}_{k=1}^{n_\lambda(B)}$ with $B^k = [B_0^k, B_1^k]$ as follows

$$\text{SED}_{B^k}(\lambda_k) = \frac{1}{Z_{n_\lambda(B)}} \sum_j \text{SED}_{b^j}(\lambda_j) |B^k \cap b^j| \quad (6.4)$$

where the sum goes over all the passband, $n_\lambda(B)$ represents the total number of bins in the B bin set, \cap denotes the intersection of the segments denoting the bins, and $Z_{n_\lambda(B)}$ is a normalisation factor.

Uncertainty in the stellar SEDs

If we desire a more realistic scenario, we must consider the errors in the SEDs. These errors arise from the photometric measurements. The statistics of the SED errors

¹After consulting the EC.

are not yet known. Therefore, we follow the *Euclid* requirement R-EXT.2.3-2. This requirement indicates that the post-calibration relative photometric error in the ground-based observations of the wide and deep surveys shall be less than 1%. We assume that the errors in each bin are independent and identically distributed following a Gaussian distribution. Then, the noisy input SED writes

$$\widehat{\text{SED}}_{B^k}(\lambda_k) = \text{SED}_{B^k}(\lambda_k) (1 + n_\sigma(\lambda_k)) , \quad (6.5)$$

where $n_\sigma \sim \mathcal{N}(0, \sigma)$. The requirement is specified as a relative error. Therefore, we define the error as multiplicative to avoid the dependency with the number of bins used $n_\lambda(B)$.

6.5.2 Numerical experiments

We will study the impact of both SED errors mentioned to the PSF generation.

Setup

We use the 13 SED templates with 350 bins with a resolution of 1nm from Figure 5.5-4 as our GT. The 350 bins are used to generate our GT pixel PSFs. The WFE uses 45 as its Zernike order, and a dimension K of 256. We focus on one FOV position from the GT PSF field from chapter 5 and use its Zernike distribution to generate the different PSFs of this section. We will only compare the pixel errors of PSFs at 3 times the *Euclid* resolution.

Results: errors due to the binning

We first study the error *only* due to the binning of the SED in the PSF generation. Different PSFs are generated with the SED from Equation 6.4 using a different number of $n_\lambda(B)$ from $\{4, 8, 16, 32, 64\}$. We compute the errors with respect to the PSF generated using 350 bins and show the results in Figure 6.5-16. In this case, we consider a noiseless scenario; therefore, the errors are only due to the binning. The results show that the errors decrease as we increase the number of bins, as expected. We observe that the errors depend on the type of stellar SED used, where the star with the flatter SED shows the highest error. When using 8 bins, all the errors are under 0.4%. Figure 6.5-17 shows a visual example comparing the GT PSF and the ones with the highest errors that are the ones computed using only 4 bins. The images show that with 4 bins we can already obtain a PSF that resembles the GT PSF.

Results: errors due to the SED uncertainty

We continue by studying the errors *only* due to uncertainty in the SEDs that can be seen in Equation 6.5. In this case, for each number of bins we generate a set of PSFs using the corresponding SED $_{B^k}$, and another set of PSFs using its noisy version, $\widehat{\text{SED}}_{B^k}$. The binning error will vanish as we compare the PSFs generated with the same number of SED bins each time. The reference standard deviations for the noise generation will be $\sigma = \{0.05, 0.02, 0.01, 0.005\}$, where $\sigma = 0.01$ would correspond to the aforementioned *Euclid* requirement. We draw independent realizations for each bin from $\mathcal{N}(0, \sigma)$.

Figure 6.5-18 presents the errors only due to the SED uncertainty for different number of bins. When the reference σ is equal to or lower than the *Euclid* requirement, 0.01, the errors do not change much with the number of bins used and are always below 0.1%. The error is low if we compare it to the one obtained for 8 bins in Figure 6.5-16.




Figure 6.5-16 – Relative pixel PSF RMSE at three times the *Euclid* resolution for three types of stellar SEDs using a different number of bins in a noiseless scenario.




Figure 6.5-17 – Pixel errors for PSFs at three times the *Euclid* resolution. Each row corresponds to a different stellar SED type. The first column shows the PSF generated with 350 bins, the second column the PSF generated using only 4 bins, and the third column the absolute difference of the residual between the other columns.




Figure 6.5-18 – Relative pixel PSF RMSE at three times the *Euclid* resolution for different reference σ values as a function of the number of SED bins. Only the error from the SED uncertainty is considered. The errors are computed between the PSFs generated by noisy and binned SEDs and the PSFs obtained with just the binned SEDs.

The errors increase as the reference σ increases. For high values of SED uncertainty, the lower bin number shows higher errors, while the high bin numbers show greater robustness to noise.

Results: errors due to the binning and the SED uncertainty

In this experiment, we study the two previous errors simultaneously, the errors due to the binning and the uncertainty in the SED. The results are shown in Figure 6.5-19. If we stand on a *Euclid* scenario using 8 bins and a reference σ equal to the requirements, i.e. 0.01, the binning errors dominate. As we increase the number of bins, the binning error decreases until the errors due to the SED uncertainty start to dominate, even if we are on the uncertainty SED requirement.

Conclusion

We have studied the errors related to the input SED in the PSF generations. All of these errors seem to be smaller than the ones observed in the previous chapter. However, the studied errors represent a lower bound of the errors in the previous chapter. The reason is that even if the PSF model managed to estimate the GT WFE representation, we would still have the errors related to the PSF generation. These SED errors can be considered errors between our optical forward model and the GT forward model.

In an expected *Euclid* scenario, 8 bins and σ of 0.01, the errors are under 0.2%, and are dominated by the SED binning. The number of input bins is generally not modifiable, as it depends on the photometric surveys. However, given an input number of bins, we can interpolate the binned SED to increase the resolution artificially and hopefully decrease the binning error. The SED interpolation will be addressed in the following section.




Figure 6.5-19 – Relative pixel PSF RMSE at three times the *Euclid* resolution considering the errors due to the binning and to the uncertainty in the SED. The errors are drawn as a function of the number of bins used and the noise standard deviation.




Figure 6.6-20 – Illustration of the bins' wavelength centres as we add new interpolation bins. The example shows an SED with 8, 16, and 24 bins, where the input SED is the 8-bin one, and the passband goes from 550nm to 900nm.

6.6 Studying stellar SED interpolation in the PSF generation

In this section, we assume that we have 8 input SED bins as input and cannot modify it. We then study if we can reduce the PSF generation error by interpolating this input to a higher number of bins.

6.6.1 SED interpolation

To define the SED interpolation, we need first to define the centre wavelength of the new interpolated bins and the interpolation method to obtain the SED value at these wavelengths. We continue with the choice that our bins will cover the entire passband from 550nm to 900nm and that the bins will be disjoint.

The original 8 input SED values will be maintained, and we will add new bins between them. As the original 8 bins already cover the entire passbands, the new bins will cover an already covered part of the passband. Therefore, we need to reassigned the bins' length so that when adding new bins, we respect the choice that all the bins are disjoint. Figure 6.6-20 shows an example of the bins' centre wavelength positions for the interpolation. We maintain a constant bin length for each scenario, except for the border bins. If we add an odd number of centres between the original 8 bins, the border bins will have half the length of the other bins in the SED. This case is illustrated in Figure 6.6-20 by using half a circle for the added centre. More precisely,




Figure 6.6-21 – Interpolated SEDs for different reference σ values using a linear interpolant. The GT SED is drawn in black. Each colour is one of the reference σ values.

the 16 point SED has 17 bins from which two have half the full bin length. The bin length is considered in the SED normalisation.

We first experimented with cubic spline interpolation. However, we finally chose a linear interpolation as it consistently gave better results. All the results presented in this section are from the linear interpolation. Figure 6.6-21 shows examples of the different interpolation points for a given GT SED and different levels of SED uncertainty. We recall that the input 8-bin SED will have errors impacting the interpolation and the PSF generation.

6.6.2 Numerical experiments

We generate PSFs at three times the *Euclid* resolution for all the stellar SED types from Figure 5.5-4. We use different levels of SED uncertainty, and interpolate the input noisy 8-bin SED into 16, 24, and 32 bins. The errors are always computed with respect to the super-resolved PSF generated with the noiseless GT 350-bin SED.

Figure 6.6-22 shows the errors as a function of the number of bins and the reference σ . The errors shown are the average over the different types of stellar SEDs. We observe that the total error can be bin-dominated or noise-dominated, which depends on the number of bins used and the level of noise (or SED uncertainty). Suppose the noise level is within *Euclid*'s requirements, i.e. 0.01. In that case, increasing the number of bins by interpolation from 8 to 16 reduces the error by a factor of 2.5, achieving an error of 0.08%. After the 16 bins, the error does not decrease as we start to be noise dominated. If the noise is half the requirement or even lower, it is worth increasing the number of bins by interpolation to 32. Nonetheless, suppose the level of input error is twice or more than the requirements. In that case, the errors are noise dominated and increasing the number of bins does not considerably reduce the errors.

The interpolation to 24 bins does not improve the result compared to the 16 bins




Figure 6.6-22 – Relative RMSE of the super-resolved PSFs for each σ level and as a function of the number of bins up to 32 bins, using linear interpolation for the new SED points.




Figure 6.6-23 – Comparison of the interpolation with 16 and 24 bins, where we add one and two points between the input 8-bin points, respectively. The input SED errors shown have a different sign.

in any scenario. This result might be expected because the 24-bin SED adds an even number of points between two input 8-bin SED points. The middle point in the segment delimited by two consecutive input 8-bin SED points is not taken when the number of added points is even. If the input SED errors have different signs, it is expected that the middle point will have a reduced interpolation error with respect to the points set at one-third and two-thirds of the segment. Figure 6.6-23 illustrates this example where two input errors have different sign. If the input errors have the same sign, neither interpolation is favoured.

One might wonder why the PSF generation error is decreasing when we interpolate the input 8-bin SED to more bins if the interpolation is not bringing new SED information. The answer is that the PSF model introduces more information as we use more bins. The error is made by discretising the spectral integral from Equation 2.27 is being reduced. The regularity of the stellar SED studied, shown in Figure 5.5-4, allows us to have a low interpolation error. Consequently, in an SED noiseless scenario, we see the PSF generation errors decrease as we interpolate to increase the number of bins. However, suppose we have a considerable noise in the input SEDs. In that case, we will be making interpolation errors that first compensate for the error reduction due to the increase of the number of bins and finally dominate the PSF generation error.




Figure 6.7-24 – Relative pixel **PSF RMSE** at three times the *Euclid* resolution for the WaveDiff-original model using different **SED** inputs. The **GT PSFs** are generated with noiseless 350-bin **SEDs**. The input **SEDs** has 8 bins with different levels of noise that are represented in the x -axis. The WaveDiff **SED** interpolation is compared between 8 (no interpolation), 16, and 32 bins.

6.7 SED errors in PSF modelling

In this section, we study the impact of the **SED** errors, from the binning and the noise, on the WaveDiff **PSF** model. The impact on the model’s performance of the **SED** interpolation is also studied.

6.7.1 Numerical experiment

We generate a dataset that uses another realisation of the **WFE** distribution from subsection 5.5.2. The same optical parameters are used for the simulations except for the **SED**-related hyperparameter, n_λ . We will have 2000 stars available for training and 400 noiseless stars for testing. The observations and the **GT** stars are generated using the noiseless 350-bin **SEDs** from the 13 stellar types of Figure 5.5-4.

The model tested will be the WaveDiff-original model from subsection 5.5.2. The **SED** inputs will always have 8 bins, but a changing level of noise, or **SED** uncertainty. The WaveDiff model will be tested under three interpolation scenarios, without and with interpolation to 16 and 32 bins. We avoid the 24-bin interpolation following the comments from section 6.6. We run 5 realisations for each **SED** input and present the median of the results. The error bars represent the standard deviation of the results.

Figure 6.7-24 shows the super-resolved pixel errors for the WaveDiff model for different **SED** interpolations as a function of the input 8-bin **SED** noise level. The interpolation to 32 bins is consistently giving better results. The increase in performance is smaller in the noiseless scenario, which is unexpected following the results from the previous sections. In the noiseless scenario, the **PSF** generation error is dominated by the binning error, and therefore increasing the number of bins should increase the performance. Nevertheless, one should not forget that the **PSF** generation errors for the noise levels used were always below the 0.3%, as we have seen in Figure 6.6-22.




Figure 6.7-25 – Relative pixel PSF RMSE at the *Euclid* resolution for the WaveDiff-original model for different SED interpolations as a function of the input 8-bin SED noise level.

The total PSF modelling errors are not dominated by the PSF generation errors, and are between 2% and 4%.

We note that the performance of the WaveDiff model is robust to different noise levels in the SED input. Even if we use an error representing twice *Euclid* requirements, $\sigma = 0.02$, the errors are kept low. The high noise level even improves the model’s performance using the 32-bin interpolation.

Figure 6.7-25 shows the PSF modelling errors at the *Euclid* resolution. We observe a similar behaviour to the super-resolution modelling errors.

6.7.2 Conclusion

In the previous sections, we have studied the impact of the SED errors on the PSF generation with the optical forward model. We have extended the analysis to the PSF modelling in the current section. We have seen that the PSF modelling errors are more significant than the PSF generation errors, around one order of magnitude.

Analysing the impact of SED errors on the PSF model is not straightforward. The way of handling the input SED in the model affects the model’s optimisation. Having defined the model’s parameters in the WFE space makes the optimisation non-convex, and a very badly posed inverse problem, which makes the error analysis not simple. The overall reconstruction error in this SED-realistic scenario is higher with respect to the error from section 5.5 where we did not consider any SED errors. The super-resolution error went from 1.3% to between 2% and 4%. The PSF generation errors due to the SED are lower than 0.3% in the noise scenarios studied. Consequently, we can say that using a realistic SED also impacts the PSF modelling besides the effect it has on the PSF generation.

The reconstruction error decreases for the different noise levels when using a 32 bin interpolation. However, increasing the number of bins is not free as the overall computing time increases linearly with the number of bins used. The reason is that more monochromatic PSFs have to be computed for the spectral integration.

SED interpolation may be required if the optical prior contains specific spectral variations that need to be included in the model. In this case, the interpolation will allow us to include the spectral variations into the model without requiring more input **SED** information. This reflection is done with the dichroic spectral’s complexity in mind. First, this complexity needs several samples² to be well represented due to its sharp variations at some wavelengths. In contrast with the dichroic, the regularity of the stellar **SEDs** allows us to use 4 or 8 bins and still obtain a low error. Second, it is not yet known if the dichroic errors will dominate with respect to the **SED** errors. In order to answer the question, we need first to quantify the impact of the **SED** errors on the WaveDiff **PSF** modelling, which was one of the objectives of the current study.

6.8 Improving the WaveDiff optimisation procedure

We propose, in this section, a novel procedure to optimise the WaveDiff **PSF** model. The section is based on a scenario where the parametric part of the model can reproduce the **GT WFE** field. Estimating the parameters of the parametric part from degraded in-focus observations is impossible with the current optimisation procedure. We develop a novel strategy based on projecting the non-parametric estimation onto the parametric part.

6.8.1 Motivation

Let us recall some results we have seen in the previous sections regarding optimising the WaveDiff **PSF** model.

We have seen that the current optimisation procedure converges to a local minimum. This observation is expected as we tackle a non-convex problem with a gradient-based optimisation method. The optimisation consists in cycles where we optimise the parametric part and then continue with the non-parametric, or **Data-Driven (DD)**, part. See Algorithm 2 for more details. We have used two optimisation cycles in all the previous experiences with the WaveDiff model. We noted that the model was stuck in a local minimum that the optimisation could not escape. This fact is easily confirmed by plotting the validation loss, which is the loss computed on the testing datasets, as a function of the number of epochs. Therefore, even if we increase the number of cycles in the optimisation procedure, the algorithm will not escape from the local minimum.

In most of the previous sections and chapter 5, we assumed a scenario where the WaveDiff’s parametric part was misspecified with respect to the **GT WFE** field. We showed that the non-parametric part could estimate a useful representation of the **WFE** field, which had low errors in the pixel space but high errors in the **WFE** space. See Figure 5.6-11 for an illustration. Another scenario to consider is when the parametric part is well specified, meaning the model can reproduce the **GT WFE** field. We have experienced, and if we try to optimise the parametric part of the model from a random initialisation, we do not converge to the global minimum, the **GT WFE**. We converge to a local minimum, and the performance of the parametric part is poor. section 6.2 shows that even if the starting point of the only-parametric part is close to the **GT**, as it was the case of a low-error prior, the optimisation takes us away from that good starting point. The **WFE** error of the parametric-only model is increased as well as its pixel error. This behaviour shows the difficulty of optimising the parametric part using in-focus degraded star observations.

²The number is not yet known, but it is expected to be at least 24 or 32.




Figure 6.8-26 – Amplitude distribution of the first 15 Zernike coefficients. The difference between the *GT* and the estimated coefficients of the model’s parametric part is plotted in blue. The non-parametric contribution to each coefficient is plotted in green.

The *DD* part of the model has been successful in converging on a good pixel representation of the *PSF* field, even if the *WFE* error is high. We have noted that the optimisation is considerably more stable than the one of the parametric part. We explain this fact with the over-parametrisation of the *DD* part. This characteristic allows having a much smoother optimisation landscape than the landscape of the parametric part, although it is still non-convex. In order to give some context, we can compare the number of parameters of the parametric and non-parametric parts of the *WaveDiff-original* model from subsection 5.5.2. The first one has 90 parameters, and the second one $\sim 1.3 \times 10^6$.

Figure 6.8-26 gives some insight to understand better the behaviour and interaction of the parametric and non-parametric parts. We have analysed a *WaveDiff-original* model once it has converged in a scenario like the one in subsection 5.5.2. We have plotted the error of the Zernike coefficients done by the parametric part. Then, we projected the non-parametric estimation to each Zernike polynomial which gives us the contribution of the non-parametric part to each specific coefficient. We can compare the non-parametric contribution and the error done by the parametric part. We observe that the non-parametric part tries to correct the parametric error with its contribution.

Having all the aforementioned points in mind, we aim to exploit the better optimisation of the non-parametric part and the ability of the parametric part to represent the *GT WFE* field. We propose to optimise the non-parametric part of the model and then transfer the information that this part shares with the parametric part. We detail how to transfer, or project, the *DD* part onto the parametric part.

6.8.2 Projecting data-driven features

To transfer information from the non-parametric to the parametric part, we first define an inner product between the model’s features, guaranteeing the orthogonality between the Zernike features. We then show that the inner product can be used to

project a [WFE](#) map onto the space spanned by a given number of Zernike polynomial features. This projection allows us to define the contribution of the non-parametric part that will be injected into the parametric part. We finally show how to remove the previous contribution from the non-parametric part to avoid duplicating one part of the [WFE](#).

Zernike inner product

We define the inner product between the features $S_k, S_l \in \mathbb{R}^{K \times K}$ as follows

$$\langle S_k, S_l \rangle = \sum_{\xi, \eta \in D} \frac{S_k[\xi, \eta] S_l[\xi, \eta]}{|D|}, \quad (6.6)$$

where $D \in \mathbb{R}^{K \times K}$ is a circular aperture. This definition of inner product guarantees the orthonormality between the Zernike features, which reads

$$\langle S_k^Z, S_l^Z \rangle = \delta_{kl}. \quad (6.7)$$

Consequently, if a [WFE](#) map is built as a linear combination of Zernike features, which is written as

$$\text{WFE}(x, y)[\xi, \eta] = \sum_{k=1}^{nz} f^k(x, y) S_k^Z[\xi, \eta], \quad (6.8)$$

where we recall that (x, y) correspond to [FOV](#) position and $[\xi, \eta]$ to the model's [WFE](#) coordinates. Then, the Zernike coefficients can be retrieved from the [WFE](#) map employing the inner product

$$f^k(x, y) = \langle \text{WFE}(x, y), S_k^Z \rangle. \quad (6.9)$$

Parametric and non-parametric parts

Let us first recall how we have defined the parametric, [Equation 5.13](#), and non-parametric, [Equation 5.15](#), parts of the WaveDiff-original model. A more detailed description can be found in [subsection 5.4.2](#). The parametric part writes

$$\Phi^Z(x, y; \lambda)[\xi, \eta] = \frac{1}{\lambda} \sum_{l=1}^{nz} \underbrace{\pi_l^Z(x, y)^T}_{f_l^Z(x, y)} \mathbb{1}_{n_{dz}} S_l^Z[\xi, \eta], \quad (6.10)$$

where S_l^Z is a Zernike feature, $f_l^Z(x, y)$ is the polynomial of the index l , and for each l we have a different polynomial π_l^Z that writes

$$\pi_l^Z(x, y) = \left[\pi_{l, [0,0]}^Z, \pi_{l, [1,0]}^Z x, \pi_{l, [0,1]}^Z y, \dots, \pi_{l, [0, dz]}^Z y^{dz} \right], \quad (6.11)$$

where we are using a Zernike order of n_Z , and a maximum polynomial degree of d_Z that gives us n_{dz} monomials. The non-parametric, or [Data-Driven](#) (DD), part writes

$$\Phi^{\text{DD}}(x, y; \lambda)[\xi, \eta] = \frac{1}{\lambda} \sum_{k=1}^{n_{\text{DD}}} \pi^{\text{DD}}(x, y)^T \mathbf{a}_k S_k^{\text{DD}}[\xi, \eta], \quad (6.12)$$

where S_k^{DD} is a [DD](#) feature, \mathbf{a}_k is a column of the mixing matrix A , and $\pi^{\text{DD}}(x, y)$ is a [FOV](#) position polynomial shared by all the [DD](#) features. The polynomial writes

$$\boldsymbol{\pi}^{\text{DD}}(x, y) = \left[\pi_{[0,0]}^{\text{DD}}, \pi_{[1,0]}^{\text{DD}}x, \pi_{[0,1]}^{\text{DD}}y, \pi_{[2,0]}^{\text{DD}}x^2, \pi_{[1,1]}^{\text{DD}}xy, \dots, \pi_{[0,d_{\text{DD}}]}^{\text{DD}}y^{d_{\text{DD}}} \right]^{\text{T}}. \quad (6.13)$$

It is important to note that, if the mixing matrix A is diagonal, then each **DD** feature contributes to a single monomial from $\boldsymbol{\pi}^{\text{DD}}$. However, as A is optimised, there is no reason for it to be diagonal. Therefore, the **DD** features will contribute to more than one monomial.

Non-parametric WFE projection

Both parts of the model have some points in common, as they contribute to the total **WFE**. Even though they have a different nature, a piece of the **DD** contribution can be represented by the parametric part. We will now see which is this piece of the **DD** part.

The learned **DD** features contain high-frequency **WFE** variations, see [Figure 5.5-6](#) for an example. Therefore, we would need a very high Zernike order to decompose the **DD** feature without errors. However, the Zernike order of the parametric part is limited to n_Z . We use the inner product defined in [Equation 6.6](#) to decompose the **DD** part with respect to a Zernike feature of index l as follows

$$\langle \Phi^{\text{DD}}, S_l^Z \rangle = \frac{1}{\lambda} \sum_{k=1}^{n_{\text{DD}}} \boldsymbol{\pi}^{\text{DD}}(x, y)^{\text{T}} \mathbf{a}_k \langle S_k^{\text{DD}}, S_l^Z \rangle. \quad (6.14)$$

If we perform the same inner product, but with the parametric part, we obtain

$$\langle \Phi^Z, S_l^Z \rangle = \frac{1}{\lambda} \boldsymbol{\pi}_l^Z(x, y)^{\text{T}} \mathbb{1}_{n_{d_Z}} = \frac{1}{\lambda} f_l^Z(x, y). \quad (6.15)$$

We reformulate [Equation 6.14](#) as follows

$$\langle \Phi^{\text{DD}}, S_l^Z \rangle \lambda = \underbrace{\sum_{\substack{i,j \geq 0 \\ i+j \leq d_Z}} \Delta \pi_{l,[i,j]}^Z x^i y^j}_{\Delta f_l^Z(x,y)} + \underbrace{\sum_{\substack{i+j > d_Z \\ i+j \leq d_{\text{DD}}}} c_{[i,j]} x^i y^j}_{C(x,y)}, \quad (6.16)$$

where $C(x, y)$ is a polynomial containing monomials of order greater than d_Z , and $\Delta f_l^Z(x, y)$ is a polynomial of degree d_Z with parameters $\Delta \pi_{l,[i,j]}^Z$. We are assuming that $d_Z < d_{\text{DD}}$ which has been the case until now. Nevertheless, the results hold otherwise, because if $d_Z > d_{\text{DD}}$, then $C(x, y) = 0$.

We aim to find an expression for the coefficients $\Delta \pi_{l,[i,j]}^Z$, as they correspond to the part of the **DD** contribution that the parametric part can represent. These coefficients correspond to the parametric correction required by the parametric part to acquire the information from the **DD** part. Rearranging [Equation 6.14](#) using matrix multiplications we obtain

$$\begin{aligned} \langle \Phi^{\text{DD}}, S_l^Z \rangle \lambda &= \quad (6.17) \\ &= \left[\pi_{[0,0]}^{\text{DD}} \quad \dots \quad \pi_{[0,d_{\text{DD}}]}^{\text{DD}} y^{d_{\text{DD}}} \right] \left(\begin{bmatrix} | \\ \mathbf{a}_1 \\ | \end{bmatrix} \underbrace{\langle S_1^{\text{DD}}, S_l^Z \rangle}_{\in \mathbb{R}} + \dots + \begin{bmatrix} | \\ \mathbf{a}_{n_{\text{DD}}} \\ | \end{bmatrix} \underbrace{\langle S_{n_{\text{DD}}}^{\text{DD}}, S_l^Z \rangle}_{\in \mathbb{R}} \right). \end{aligned}$$




Figure 6.8-27 – Illustration of the tensor product involved in the **DD** part of the model.

Then, the non-parametric corrections for the first two terms, $\Delta\pi_{l,[0,0]}^Z$ and $\Delta\pi_{l,[1,0]}^Z$, can be expressed as

$$\begin{aligned}\Delta\pi_{l,[0,0]}^Z &= \pi_{[0,0]}^{\text{DD}} (\mathbf{a}_1[1] \langle S_1^{\text{DD}}, S_l^Z \rangle + \mathbf{a}_2[1] \langle S_2^{\text{DD}}, S_l^Z \rangle + \cdots + \mathbf{a}_{n_{\text{DD}}}[1] \langle S_{n_{\text{DD}}}^{\text{DD}}, S_l^Z \rangle) \in \mathbb{R}, \\ \Delta\pi_{l,[1,0]}^Z &= \pi_{[1,0]}^{\text{DD}} (\mathbf{a}_1[2] \langle S_1^{\text{DD}}, S_l^Z \rangle + \mathbf{a}_2[2] \langle S_2^{\text{DD}}, S_l^Z \rangle + \cdots + \mathbf{a}_{n_{\text{DD}}}[2] \langle S_{n_{\text{DD}}}^{\text{DD}}, S_l^Z \rangle) \in \mathbb{R},\end{aligned}$$

where $\mathbf{a}_1[1]$ corresponds to the first element of the first column of the matrix A . By looking carefully to [Equation 6.17](#), we can generalise the parametric correction coefficients as follows

$$\Delta\pi_{l,[i,j]}^Z = \pi_{[i,j]}^{\text{DD}} \sum_{n=1}^{n_{\text{DD}}} \mathbf{a}_n[p(i,j)] \langle S_n^{\text{DD}}, S_l^Z \rangle, \quad (6.18)$$

where $i, j \geq 0$ and $i + j \leq d_Z$, and $l = 1, \dots, n_Z$. The order of the indexes $[i, j]$ can be seen in [Equation 6.13](#). This order is required to index the vectors \mathbf{a}_n , and is given by the function $p(i, j)$ that writes

$$p(i, j) = \frac{(i+j)(i+j+1)}{2} + j + 1. \quad (6.19)$$

The required correction, $\pi_{l,[i,j]}^{Z*}$, to apply to the parametric model in order to incorporate the information from the **DD** part can be expressed as

$$\pi_{l,[i,j]}^{Z*} = \pi_{l,[i,j]}^Z + \Delta\pi_{l,[i,j]}^Z. \quad (6.20)$$

Projected data-driven subtraction

We have just seen how to modify the parametric part to account for the information learned by the **DD** part. We now need to modify the **DD** part so that the information we have just transferred to the parametric part is not replicated in the **DD** part. The sum of both parts of the model should remain constant in this projection procedure.




Figure 6.8-28 – Illustration of the tensor product using the mixed **DD** features $\tilde{\mathbf{S}}^{\text{DD}}$, involved in the **DD** part of the model.

We start by rewriting the **DD** part of the model from Equation 6.12 using a tensor product as follows

$$\Phi^{\text{DD}} = \frac{1}{\lambda} \sum_{k=1}^{n_{\text{DD}}} \underbrace{\pi^{\text{DDT}}}_{\in \mathbb{R}} \mathbf{a}_k S_k^{\text{DD}} = \frac{1}{\lambda} \pi^{\text{DDT}} A \mathbf{S}^{\text{DD}}. \quad (6.21)$$

Figure 6.8-27 illustrates the tensor product of Equation 6.21 where $\mathbf{S}^{\text{DD}} \in \mathbb{R}^{n_{\text{DD}} \times K \times K}$ is the tensor composed of all the learned **DD** features $S_i^{\text{DD}} \in \mathbb{R}^{K \times K}$.

In order to express what was transmitted to the parametric model, one would like to define it as the projection onto the first n_z Zernike polynomials of the **DD** features. Then, remove its Zernike projections from each **DD** feature. However, we would be neglecting that only the **WFE** projection associated with **FOV** spatial variations, i.e. polynomial in x and y , of order less or equal to d_z was transmitted. The polynomial degree of the **DD** part is in general bigger, $d_{\text{DD}} > d_z$. Consequently, the higher-frequency **FOV** spatial variations are not transmitted as the parametric part cannot represent them. The solution is pretty straightforward if we suppose the A matrix is the identity matrix. We have to remove the projection over the first n_z Zernike polynomials for the first n_{d_z} matrices of \mathbf{S}^{DD} .

The mixing matrix, A , will not be the identity, and it will mix the **DD** features S_k^{DD} as follows

$$\tilde{\mathbf{S}}^{\text{DD}} = A \mathbf{S}^{\text{DD}} = \sum_{i=1}^{n_{\text{DD}}} \sum_{k=1}^{n_{\text{DD}}} \mathbf{a}_k[i] S_k^{\text{DD}}. \quad (6.22)$$

The mixed **DD** feature tensor $\tilde{\mathbf{S}}^{\text{DD}}$ has a direct matching with the spatial variations from π^{DD} . Therefore, we need to remove the contribution of the first n_z Zernike polynomials to the first n_{d_z} matrices of the mixed **DD** features $\tilde{\mathbf{S}}^{\text{DD}}$. Figure 6.8-28 illustrates the matrices from the mixed **DD** features that have contributed to transmitting information to the parametric model and need to be modified.

The only parameters we can modify from the **DD** part is the \mathbf{S}^{DD} tensor. The vector π^{DD} and the matrix A are shared by **DD** features of a higher **FOV** spatial variation, and we do not want to modify them. We look for a new target tensor $\mathbf{S}^{*\text{DD}}$ such that once it is multiplied by A , its first n_{d_z} matrices are orthogonal to the first n_z Zernike polynomials. We denote $\tilde{\mathbf{S}}^{*\text{DD}} = A \mathbf{S}^{*\text{DD}}$, the mixing of the target tensor. This




Figure 6.8-29 – Comparison between the **WFE** map from the parametric part of the model before and after the projection of the **DD** part, and the **GT WFE**. The left image shows the **WFE** of the parametric part of a WaveDiff-original model once the training has finished. The image corresponds to a **PSF** reconstruction for a random point in the **FOV**. The centre image shows the new parametric part once the projection of the **DD** part was added. The right image shows the **GT WFE** of the **PSF** field projected over the first n_Z Zernike polynomials.

new tensor, $\tilde{\mathbf{S}}^{*\text{DD}}$, can be computed by subtracting the Zernike polynomial projection over the first n_{dz} matrices as follows

$$\tilde{\mathbf{S}}_k^{*\text{DD}} = \begin{cases} \tilde{\mathbf{S}}_k^{\text{DD}} - \sum_{i=1}^{n_Z} \langle \tilde{\mathbf{S}}_k^{\text{DD}}, \mathbf{S}_i^Z \rangle \mathbf{S}_i^Z & \text{if } 1 \leq k \leq n_{dz} \\ \tilde{\mathbf{S}}_k^{\text{DD}} & \text{if } n_{dz} < k \leq n_{\text{DD}} \end{cases}. \quad (6.23)$$

Once we have computed the tensor $\tilde{\mathbf{S}}^{*\text{DD}}$, we can recover the desired parameters by unmixing the tensor as follows

$$\mathbf{S}^{*\text{DD}} = \mathbf{A}^{-1} \tilde{\mathbf{S}}^{*\text{DD}}. \quad (6.24)$$

We cannot guarantee that once the optimisation algorithm has converged, the matrix \mathbf{A} is nonsingular. However, in practice, the matrix is initialised with the identity and converges to a nonsingular matrix.

Testing the projection alone

We have described how to project the information of the **DD** part of the model onto the parametric part. Then, how to subtract from the **DD** part the content that we have just projected.

Figure 6.8-29 shows the effect of the **DD** features projection on the parametric part of the model. We can see that the projected **WFE**, centre image, looks much closer to the **GT WFE**, right image, than the original parametric **WFE**, left image. These **WFE** maps were taken from a trained WaveDiff-original model with hyperparameters as in subsection 5.5.2. These **WFEs** confirm what we have seen in Figure 6.8-26, that the **DD** part is bringing the parametric part closer to the **GT**.

Figure 6.8-30 presents the Zernike coefficients amplitudes from the model of Figure 6.8-29. The value of the Zernike coefficients is obtained through the inner product with the corresponding Zernike polynomial. We can now quantitatively see the improvement of the parametric part after the **DD** projection and how it gets closer to the **GT** Zernike coefficients amplitudes.




Figure 6.8-30 – Amplitude of the first $n_Z = 15$ Zernike polynomial coefficients. The parametric part of the model before and after the DD projection are shown in blue and orange, respectively. The GT PSF field values are shown in green. We show, in red, the total WFE prediction of the WaveDiff-original model projected onto these first Zernike polynomials. The WFEs used come from a trained WaveDiff-original model and its WFE recovery of a random position in the FOV.

There is one unaddressed point from the projection procedure described above. The DD features corresponding to high-order FOV polynomials could not be transferred as we projected for $n_{dz} < n_{DD}$. The point is whether these features contain meaningful information on the Zernike polynomials of order below n_Z . Figure 6.8-30 shows the Zernike coefficients from the full WFE prediction of the WaveDiff-original model. If we compare this prediction, in red, and the corrected parametric part, in orange, we can see the contribution of these DD features from high-order FOV polynomials that were not projected. Both curves almost overlap, meaning most of the information was passed to the parametric part of the model.

6.8.3 Introducing the WFE projections to the optimisation procedure

We have defined how to project the information from the DD part to the parametric part of the model. It is now the turn to exploit this procedure to improve the optimisation algorithm. The implementation is simple. After each optimisation cycle finishes, we project the model's learned DD part. However, we need to make more changes to the original optimisation algorithm to exploit the projections better.

Let us recall that the DD projection is not changing the overall WFE representation of the model, Φ , but rather changing its two contributors, the parametric part, Φ^Z , and the DD part, Φ^{DD} . This change is helpful as the initialisation points of the gradient-based optimisation have changed. We proceed by describing modification to the optimisation algorithm required to exploit better the potential of the DD projections:

- (i) *Consider a new optimisation scenario:* The previous scenario studied in chapter 5 considered that the WFE PSF field was more complex than what the parametric part of the mode was able to represent. Figure 5.6-11 shows the previous scenario where the GT WFE could not be generated with the WaveDiff model. For this section, we will consider a new scenario where the parametric part of the model




Figure 6.8-31 – Loss function as a function of the number of epochs for the **DD** part of the model in the second optimisation cycle. One model has been reset after performing the projection, and the other model has not been reset.

can represent the **GT WFE** field. This new scenario guarantees us that there exists a global minimum for the problem and that we know which one it is, the **GT WFE**. The **DD** projection helps to improve the **WFE** estimation. Therefore, it motivates the change of scenario so that we can study how the **WFE** recovery improves with the proposed **DD** projection. We note that the new scenario falls within the category of phase retrieval problems, as our target is now the **GT WFE**.

- (ii) *Increase the number of optimisation cycles:* The **DD** projections change the parametric and the **DD** parts. These new and more convenient starting points are expected to be exploited in the next optimisation cycle of the model. Consequently, to take advantage of the projections done at the end of each cycle, we need to increase the number of optimisation cycles. In all the previous numerical experiments, we have been using 2 cycles. We will now increase it to 5 and 9 cycles.
- (iii) *Reset the **DD** part after the projection:* We have seen that the current gradient-based algorithm converges to a local minimum. The **DD** projection is a redistribution of the **WFE** between the parametric and the **DD** parts. Nevertheless, the overall **WFE** has not changed. If we start another optimisation cycle, we will still be at the local minimum, which is hard to escape. We propose to reset, or re-initialise, the **DD** part of the model after the **DD** projection. This operation helps to escape the local minimum, as the overall **WFE** changes. It also allows the **DD** part to explore more optimisation landscape while considering what has been previously learned. Suppose the number of optimisation cycles is large enough. In that case, the procedure can exploit the ability of the parametric part to represent the **GT WFE** field and the smooth optimisation landscape of the **DD** part. Figure 6.8-31 shows the optimisation loss of the **DD** part of the model in the second optimisation cycle. It is clear that the model without resetting the **DD** part is stuck in a local minimum it cannot escape, no matter the number of optimisation epochs. The model being reset starts from a higher loss error but can skip the previous local minimum and obtain a lower final error.
- (iv) *Make sure that the optimisation in each cycle has converged:* As we have seen in the previous chapter, computing time is an issue with **PSF** models. Therefore,

when we first increased the number of cycles, we decided to maintain the total number of epochs of the WaveDiff-original model from [chapter 5](#) constant. By doing this, the overall computing time remains practically unchanged. However, this choice means we are using fewer epochs for each cycle. The consequence of this choice is that the [DD](#) parts were not properly converging, and the projection was not useful. Therefore, we must maintain the original number of epochs that allow the [DD](#) part to converge in each cycle. The cost we will have to pay is the increase in the overall computing time.

- (v) *Avoid optimising the parametric part:* The optimisation of the parametric part from degraded in-focus observations is complicated and degenerate. Even if the parametric part can perfectly reproduce the [GT WFE](#) field, estimating its parameters is not possible with the gradient-based optimisation technique used. The set of parameters obtained is far from the [GT](#), and their performance is poor. Therefore, we will focus on the [DD](#) part optimisation, which is better-behaved thanks to its over-parametrisation. Until we develop a better way to optimise the parametric part, we will stop optimising it and only use it to store the projections of the [DD](#) part at the end of each optimisation cycle.

We have described several changes proposed to the optimisation algorithm based on our knowledge of the problem and numerous numerical experiments that we omit for brevity.

6.8.4 Numerical experiments

We test the proposed optimisation framework using the [DD](#) projection. We start by generating a [PSF](#) field with the parameters from [subsection 5.5.2](#). The number of training and testing stars is 2000 and 400, respectively. Following point (v) from the previous subsection, only the [DD](#) part of the model is trained and is then projected onto the parametric part. We focus on the WaveDiff-original model from [subsection 5.5.2](#), but using a n_z of 45, so that we are in the new optimisation scenario of point (i).

We vary the total number of cycles and the number of epochs per cycle to study points (ii) and (iv), respectively. Once the model has been trained, we will evaluate the performance of the full model, i.e. parametric and [DD](#), and only the parametric part. We compute the relative reconstruction error of the pixel [PSFs](#) at three times the *Euclid* resolution, and the relative [WFE](#) reconstruction error. We run 5 training realisations for each setting and show the median, the best, and the standard deviation of the results.

6.8.5 Results

[Table 61](#) shows the performance results for the WaveDiff model trained following the five points from [subsection 6.8.3](#). The first thing to observe is the effect of the [DD](#) part of the model in the results. The effect can be seen by comparing the full and parametric evaluations. The [DD](#) part reduces considerably the pixel error and increases the [WFE](#) error. The [WFE](#) contributions of the [DD](#) part cannot be represented by the parametric model but greatly improve the pixel reconstruction. On the contrary, only using the parametric model gives a better [WFE](#) reconstruction but a worse pixel error, as we have explained through the illustration in [Figure 5.6-11](#).

Then, we observe in [Table 61](#) how the pixel errors decrease for both evaluations as we increase the number of training cycles. As we use more cycles, the parametric part is getting closer to the [GT WFE](#) and the [DD](#) part has a better starting point.

Table 61 – Performance results for the WaveDiff model resetting the **DD** part after each projection, as in point (iii). Each cycle uses 75 epochs to ensure that the **DD** part optimisation converges as mentioned in point (iv).

Parameters		SR Pixel RMSE [%]	WFE RMSE [%]
Cycles	Evaluation	Median \pm Std Dev (best)	Median \pm Std Dev (best)
2	Parametric	14.8 \pm 1.5 (11.8)	86 \pm 22 (62)
2	Full	3.5 \pm 1.4 (2.2)	158 \pm 30 (91)
5	Parametric	7.9 \pm 3.4 (4.1)	98 \pm 39 (32)
5	Full	3.1 \pm 1.1 (1.7)	166 \pm 52 (37)
9	Parametric	1.9 \pm 4.5 (1.7)	14 \pm 52 (11)
9	Full	1.6 \pm 1.2 (1.5)	30 \pm 63 (20)

Table 62 – Performance results for the WaveDiff model *without* resetting the **DD** part after each projection. Each cycle uses 75 epochs to ensure that the **DD** part optimisation converges as mentioned in point (iv).

Parameters		SR Pixel RMSE [%]	WFE RMSE [%]
Cycles	Evaluation	Median \pm Std Dev (best)	Median \pm Std Dev (best)
2	Parametric	18.5 \pm 1.0 (18.2)	103 \pm 14 (91)
2	Full	4.0 \pm 1.3 (2.4)	179 \pm 28 (124)
5	Parametric	18.8 \pm 2.3 (15.2)	101 \pm 18 (66)
5	Full	4.4 \pm 0.5 (3.7)	211 \pm 8 (194)

This result shows that the model can capture new information on each cycle and successfully transfer it to the parametric part. A remarkable result is the low **WFE** reconstruction error achieved by the model with 9 cycles, that is 14%. It is one order of magnitude below the typical **WFE** errors the model had, which are represented by the full evaluation of the model with 2 cycles. We note that all the models start without knowledge of the **WFE**, making the result more meaningful.

We have included the information about the best performance in [Table 61](#). One might remark the high standard deviation of the results for the model trained with 9 cycles. This result is due to one or two models that have converged to a far away **WFE** that gives poor results, which drives the mean and standard deviation to higher values. However, in practice, it is easy to detect if the model has converged in a bad local minimum. We observe the pixel reconstruction error done on the train stars and compare it to the *usual* pixel error. Then, as the error is considerably higher, we can detect which model is an outlier and eventually retrain it. If we use some prior **WFE** information in the model as in [section 6.2](#), we expect finding an outlier to be rare.

[Table 62](#) shows the performance of the models trained using all the required epochs per cycle, but without resetting the **DD** part after projection. This configuration means that we do not respect point (iii). We note a degradation in these models' pixel and **WFE** performance. The degradation is evident when observing the best **WFE** error, confirming that resetting the **DD** part improves the new optimisation procedure.

[Table 63](#) shows the performances when we reset the **DD** part, but we use a reduced number of epochs per cycle. In this case, we do not respect point (iv). The pixel error is greatly impacted, showing poor performance. For the model trained with 9 cycles

Table 63 – Performance results for the WaveDiff model resetting the **DD** part after each projection, as in point (iii). In this case, each cycle uses only 15 epochs.

Parameters		SR Pixel RMSE [%]	WFE RMSE [%]
Cycles	Evaluation	Median \pm Std Dev (best)	Median \pm Std Dev (best)
5	Parametric	12.6 ± 1.2 (10.5)	110 ± 16 (83)
5	Full	7.0 ± 1.5 (3.6)	121 ± 17 (88)
9	Parametric	8.6 ± 2.3 (5.5)	79 ± 17 (54)
9	Full	7.1 ± 0.8 (5.3)	117 ± 13 (90)




Figure 6.8-32 – WFE representations for a given training point in the FOV. The first row shows, from left to right, the parametric, **DD**, and the full WFE prediction for the new optimisation algorithm using 9 cycles. The second row shows the residual WFE for the parametric and the full WFEs in the first and third columns, respectively, with respect to the ground truth WFE that is placed in the second column.




we went from a pixel error of 1.6% in Table 61 to 7.1% in Table 63. This result shows the importance of ensuring that the optimisation of the **DD** part converges in each cycle.

Visualising the WFE

We will show WFE reconstruction to showcase the results of the proposed optimisation algorithm. We use as example the best performing model trained with 9 cycles from Table 61 that respects all the points proposed in subsection 6.8.3.

Figure 6.8-32 shows the WFE contributions of both parts of the model and the full WFE prediction. We can see that the parametric part is very close to the **GT WFE**. Then, the non-parametric, or **DD**, part is composed of high-frequency WFE variations that the parametric part cannot represent. These high-frequency variations are the ones that greatly improve the pixel error by driving us away from the **GT WFE**.

Figure 6.8-33 presents the WFE evolution of the two parts of the model throughout the cycles. Figure 6.8-33a shows the evolution of the parametric part, and how the

(a) Parametric part **WFE** evolution at the end of each cycle.(b) Evolution of the parametric part **WFE** residual at the end of each cycle with respect to the **GT WFE**.(c) Evolution of the non-parametric, or **DD**, part before being projected at the end of each cycle.Figure 6.8-33 – Evolution of different **WFE** parts of the model as the number cycle increases. The same scale is used for all the images in each figure.

parametric part gets closer to the *GT WFE* after each projection. Figure 6.8-33b shows the residual of the parametric part with respect to the *GT WFE*, where we can see that the residual gets each time closer to zero. Finally, Figure 6.8-33c lets us see the non-parametric part at the end of each cycle before being projected. We can observe that as the cycles go through, the information projected is reduced, showing some convergence of the algorithm. If we pay attention to the details, we can see that the non-parametric part matches the patterns in the parametric residuals from Figure 6.8-33b.

6.8.6 Conclusions

In this section, we have proposed a new optimisation procedure for the WaveDiff model. The new procedure is based on a method to project *WFE* information from the *DD* to the parametric part of the model. We have given a detailed explanation of the projection method. Then, we have described several modifications to the previous optimisation procedure required to exploit the proposed projection method's potential.

We carried out numerical experiments to validate the new optimisation approach. The results were remarkable as the new approach allowed us to obtain an outstanding estimate of the underlying *GT WFE*. Starting from no knowledge of the *GT WFE* and only using degraded in-focus stars, we were able to estimate the underlying *WFE* field. The previous optimisation procedure allowed us to obtain a low pixel error but with a high *WFE* error. If the parametric part of the model is complex enough, we can now target to recover the true *GT WFE* and not just a useful *WFE* that gives us a low pixel error. The current drawback of the new procedure is the increase in computing time.

Finally, we have addressed a phase retrieval problem, where we are interested in estimating the *WFE* from in-focus intensity-only observations. However, the setup is not typical for phase retrieval, where we generally dispose of several intensity observations of the same underlying complex signal. In our case, we dispose of a single intensity observation of each star, where each one has its complex signal, i.e. its *WFE* representation. However, the complex signals belong to the same underlying *WFE* field. In our model, we are simultaneously addressing the phase retrieval problem of all the stars in the *FOV*. We exploit the fact that the *WFE* representations of the stars are related to each other by a common *WFE* field, which is the *WFE*-based *PSF* model we are estimating.

6.9 Conclusion

In this chapter, we have addressed several issues when applying WaveDiff to real *Euclid* observations.

We have shown how the WaveDiff model can incorporate and exploit prior information about the *Euclid* telescope. The WaveDiff model is robust to different levels of errors in the prior and consistently improves the prior performance. We continued by studying the use of more than one exposure to constrain the *PSF* model, which allows us to increase the number of training stars. The errors decrease as we increase the number of stars, but they stagnate as we go over two exposures. The overall computing time increases when we increment the number of training stars. Therefore, we concluded that the performance increase does not justify going over 4000 training stars, i.e. two exposures.

We studied how using an approximation of the ideal image formation model affects the *PSF* generation and the WaveDiff *PSF* modelling results. First, we studied the

impact of the [WFE](#) dimension. Then, we studied the impact of the input [SEDs](#) discretisation and its errors. We computed the error of the [PSF](#) generation in both cases, which can be considered a lower bound to the total [PSF](#) modelling error. The generation error was low. However, the [PSF](#) modelling error increased due to the discrepancy between the ideal image formation model and our approximation with the optical forward model.

In the last section, we proposed a novel optimisation procedure for the WaveDiff model considering a new scenario where the parametric model can represent the [GT WFE](#) field. The procedure is based on projecting information from the [DD](#) to the parametric part of the model. We show with numerical experiments that the WaveDiff model can estimate the [GT WFE](#) field with low errors only using degraded in-focus observations. The drawback of the approach is that the new optimisation procedure increases the computing time of the model.

* * *
* *
*

General Conclusions and Perspectives

Chapter Outline

Contributions	195
Perspectives	196

The main goal of this thesis has been the development of new [Point Spread Function \(PSF\)](#) models for ground-based and space-based telescopes. The main target application of these models is [Weak Gravitational Lensing \(WL\)](#) studies. We have focused on the [Canada-France Imaging Survey \(CFIS\)](#) and the *Euclid* space mission. We will then describe the contributions of this thesis and conclude with its perspectives.

Contributions

I have addressed the [PSF](#) modelling problem for [WL](#) studies that reside in the intersection of three disciplines, namely optics, cosmology, and inverse problems in imaging. Therefore, the first part of this thesis introduces all the necessary concepts required to develop a comprehensive understanding of the [PSF](#) modelling problem. I started describing the basic concepts of cosmology and weak gravitational lensing required to understand the final goal of the [PSF](#) model in the [WL](#) analysis. I continued with an overview of the main building blocks of a shape measurement pipeline to depict the context in which the [PSF](#) model will work. The introduction continues with essential notions in optics and a review of the contributors to the [PSF](#) field and state-of-the-art [PSF](#) models. This part concluded with validation methods for [PSF](#) models.

The second part of this thesis focused on [PSF](#) modelling for ground-based telescopes. Nearly all the state-of-the-art [PSF](#) models work independently on each [CCD](#)-chip from the instrument's focal plane. Wide field imagers, generally used for [WL](#) studies, have large focal planes with a mosaic composed of several [CCD](#) chips. The most significant limitation for building more complex and powerful [PSF](#) models is the number of stars available to constrain the model. I developed a new data-driven [PSF](#) model, coined [Multi-CCD \(MCCD\)](#), that can build more complex models by using the entire focal plane at once. The new model is based on a constrained matrix factorisation framework that exploits proximal optimisation algorithms in a composed alternating procedure. The method yielded an improved performance with respect to

the state-of-the-art. Then, I implemented the method in an existing high-performance shape measurement pipeline and used it to produce the new **CFIS** shape catalogue of more than $\sim 3500 \text{ deg}^2$. The **MCCD** model proved robust enough to work with real **CFIS** data and showed an enhanced performance with respect to the previous **PSF** model used in the pipeline.

The third part of this thesis focused on **PSF** modelling for space-based telescopes. I have concentrated on the challenges raised by the *Euclid* space mission. The single wide passband of *Euclid*'s **VIS** instrument requires the modelling of the **PSF**'s spectral variations, and the pixel size requires a super-resolution step of a factor of 3. To address both of these challenges, I have developed a data-driven **PSF** model, coined WaveDiff, based on a framework for modelling the **PSF** that is fundamentally different from previous data-driven models. The new model is physically based and relies on a differentiable optical forward model, allowing us to build a data-driven model in the **Wavefront Errors (WFE)** space instead of the pixel space. This change permits a natural handling of the **PSF** spectral variations and the super-resolution. Modern gradient-based methods, a cornerstone of the deep learning revolution, are used to optimise the model thanks to its implementation in an automatic differentiation framework. The WaveDiff model showed an outstanding performance and can model the **PSF** spectral variations. The gain in performance over the *Euclid* requirements with respect to the previous *Euclid* data-driven **PSF** model in the current setting is between two and four orders of magnitude.

In the final chapter of this thesis, the application of the WaveDiff model for *Euclid* was addressed. I proposed an extension of the **PSF** model to use complementary information in the form of an imperfect optical prior. The model proved to be able to exploit the available information and to be robust to underlying errors in the prior. The chapter continues to study a more realistic scenario with the impact of using a discrete **WFE** for the modelling and a discrete imperfect stellar **Spectral Energy Distribution (SED)** as input. Then, I studied the possible improvement of using more than one exposure to constrain the **PSF** model. Finally, I proposed a new optimisation procedure for the WaveDiff **PSF** model that is based on a projection of the estimated data-driven **WFE** to a parametric model of the **WFE**. The new procedure can provide remarkable estimations of the underlying ground truth **WFE** solely from degraded in-focus star observations.

Perspectives

Ground-based telescopes

Deep plug-and-play optical priors for ground-based PSF Models One of the most significant limitations of ground-based **PSF** models is that the fast-changing and stochastic atmosphere limits the number of stars available to constrain the model. Consequently, **PSF** models are built independently for each camera exposure. However, the different survey data releases are done in batches as the sky coverage increases. Once a good part of the survey area has been imaged, one can compute the average ellipticity of the observed stars as a function of their position on the focal plane. The atmospheric ellipticity contribution to the star observations has a zero mean if computed over a large enough area. Therefore, we can obtain a fine-sampled characterisation of the ellipticity contribution of the telescope's optical system. It does not exist a **PSF** model that is currently exploiting this information. A single exposure does not allow recovering the high-frequency variations due to the lack of constraining information on the available stars. I propose to build a data-driven prior

of the telescope’s optical aberrations with the help of deep learning techniques. The project requires realistic simulations based on [CFIS](#) stars measurement from the v1 shape catalogue. Then, the simulations are to train deep learning-based denoisers. These networks can be included as proximal operators in the optimisation framework of the [MCCD PSF](#) model. The denoisers are included as a plug-and-play approach that avoids the dangerous black-box usage of deep neural networks. The framework provides a controlled environment to exploit the power of the deep learning-based denoisers. If the method succeeds, it could be used for the next release of the shape catalogue and also motivate its adoption for other ground-based telescopes.

Space-based telescopes

WaveDiff for the James Webb Space Telescope The launch of the [JWST](#) has been a success, and it is already observing the Universe. The ongoing [COSMOS-Webb](#) survey will map a contiguous area of 0.6 deg^2 . This area will overlap with the [COSMOS-HST](#) area, which is famous for its mass map. These incoming observations will provide an invaluable test bench for the WaveDiff model, allowing us to address possible problems and improvements before the launch of *Euclid*. If the model is successful, it could be a valuable tool for many astronomers starting to exploit the power of [JWST](#). The WaveDiff model, with its optical forward model, is based on general principles and does not present any fundamental limitation to be adapted to the [JWST](#) telescope. However, the main difficulty resides in the proper handling of [JWST](#)’s segmented hexagonal mirrors that will introduce discontinuities into the [WFE](#) field.

PSF error propagation The current propagation of [PSF](#) modelling errors into the shape measurement and more deeply into the cosmological parameter estimation has fundamental limitations. The origin of these limitations is that as instruments become more powerful, the [PSF](#) becomes further away from a Gaussian. Therefore, second-order moments cease to be good shape descriptors. A new perspective is to develop better [PSF](#) error propagations. One can build a shear estimator directly from the images, bypassing the need to use second-order moments, e.g. the direction proposed by Schneider et al. [[Sch+15](#)] with a Bayesian hierarchical model. Consequently, using our differentiable [PSF](#) model, we can quantify the impact of the [PSF](#) on the shear estimation using a first-order Taylor expansion and extract the different derivatives from the differentiable models.

Non-parametric phase retrieval WaveDiff’s new optimisation procedure, presented in [section 6.8](#), shows a new way to address the challenging phase retrieval problem. The typical phase retrieval problem exploits several intensity measurements of the same complex signal. We dispose of a single intensity image for each complex signal in our setting. Nonetheless, we exploit simultaneously all the intensity images and the fact that all of their complex signals are related to the underlying [WFE](#) field. A perspective to improve the reconstruction performance is to include more than one intensity observation of each object. *Euclid* will carry out periodic phase-diversity observations as we have mentioned in [subsection 6.2.1](#). *Euclid*’s phase diversity campaign accounts for three observations of the same objects, where one observation is in focus, and the other two have different levels of defocus. These observations carry a great amount of phase information we can exploit in our model. Handling these new observations is not arduous, which amounts to adding a constant layer of defocus to the model’s [WFE](#) for the corresponding defocused observations. In [WFE](#), the defocus

is described by the Zernike S_4^Z . The WaveDiff's phase retrieval should be compared with classical parametric and non-parametric phase retrieval methods, which can only handle several observations of a single complex signal.

Euclid-specific perspectives

- *CCD-chip misalignment*: This thesis assumes that the CCD-chips are perfectly placed in the focal plane. However, CCD-chip misalignments exist, causing a spatially varying defocus that changes independently in each CCD. This issue can be addressed by assuming that each CCD is a plane segment in 3D space. The WaveDiff model can include an additive defocus contribution for each CCD to the current WFE model. A defocus is represented by a single Zernike, S_4^Z , and plane segment parameters model its spatial variation. Finally, the plane parameters can be included in the model's optimisation procedure accounting for the misalignments of the CCD. If the CCD chips cannot be assumed plane segments, a more complex parametrisation can replace the surface, and the same procedure can be followed.
- *Data-driven modelling of pixel-level effects*: This thesis has focused on the optical contributions to the PSF field and has not considered detector-level effects in all our simulations. These effects exist, affect our PSF at the pixel level, and must be handled. This issue can be addressed by including a pixel-level data-driven contribution to the WaveDiff model after the optical forward model. The modularity of the model permits a smooth implementation. Convolutional and non-convolutional effects will affect the PSF. The former effect can be handled through a convolutional layer without non-linearities. The latter effect can be handled in two ways, using a shallow CNN or a matrix factorisation approach as in the ground-based PSF models. Care must be taken in the optimisation procedure so that the pixel-level models do not learn to model optic-based contributions of the PSF that are better handled in the WFE. The final choice of the pixel-level model requires a deeper study of the detector effects in the specific *Euclid* scenario.
- *Taking into account the dichroic-filter*: The spurious spectral dependency introduced by the dichroic filter seen in Figure 2.5-14 has not been addressed yet. The dichroic can be modelled by a complementary $\Phi^{\text{DC}}(x, y; \lambda)$ to our model's WFE $\Phi(x, y; \lambda)$. Handling dichroic presents two main difficulties. First, the contribution cannot be factorised on spatial-only and spectral-only variations as follows $\Phi^{\text{DC}}(x, y; \lambda) \neq g_1(x, y) g_2(\lambda)$. Second, the high complexity of the spectral variation of the dichroic, i.e. $\Phi^{\text{DC}}(x_i, y_i; \lambda)$ for a fixed (x_i, y_i) , requires using a high number n_λ of spectral discretisation points. The spectral variation comprises peaks and valleys where their maxima and minima change as a function of the FOV position. There is a current EC effort to characterise $\Phi^{\text{DC}}(x, y; \lambda)$ with measurements to a spare dichroic filter that should be close to the filter integrated in the *Euclid* module. The dichroic model will be available but exploiting it in the model is not trivial. First, it is important to compare the impact of the dichroic with respect to errors introduced by the degraded input SED information to determine which one is dominating. Second, it is crucial to develop a spatially varying optimal spectral discretisation that will allow us to account, with a low error, for the high complexity of the dichroic spectral variations without requiring an extremely high number of spectral bins. The computing time increases linearly with the number of spectral bins, so keeping

it as low as possible is crucial. Third, the dichroic model will not be ideal and will contain some errors. Consequently, the WaveDiff model should include a data-driven correction for the dichroic model $g_{\theta}^{\text{DC}}(x, y, \lambda)$ to account for these errors.

* * *
* *
*

Appendices

MCCD optimization

Optimisation methods

In this appendix, we include details on the practical resolution of the four optimisation problems seen in Algorithm 1. For more information about proximal operators and proximal algorithms, we refer to Parikh et al. [PB14] and Beck [Bec17].

Problem (III)

As in most of the optimisation problems, the algorithm used depends on the objective function we work with. In this case, we use the primal-dual algorithm in Condat [Con13, Alg 3.1]¹. The main motivation resides in the nature of the constraints we use when optimising over S_k , as we face one smooth and two non-smooth terms, and a linear operator. The optimisation algorithm aims at solving the following problem

$$\text{Find } \hat{x} \in \arg \min_{x \in \mathcal{X}} [F(x) + G(x) + H(L(x))], \quad (\text{A.1})$$

where: (i) F is convex, differentiable and its gradient is L-Lipschitz continuous; (ii) G and H are proximable functions that should have closed form proximal operators; (iii) L is a bounded linear operator; and (iv) the set of minimisers of the aforementioned optimisation problem is non-empty. It is straightforward to identify the different functions in the optimisation of the local S_k matrix which match the formulation of Equation A.1. Following the notation we use throughout the chapter 3, let $F_k(x = (S_1, \dots, S_N, \tilde{S}, \alpha_1, \dots, \alpha_N, \tilde{\alpha})) = \frac{1}{2} \|Y_k - \mathcal{F}_k(\hat{H}_k)\|_F^2$, with $\hat{H}_k = S_k \alpha_k V_k^\top + \tilde{S} \tilde{\alpha} \Pi_k$, and $G(S_k) = \sum_i \|\mathbf{w}_{k,i} \odot \Phi_{\mathbf{s}_{k,i}}\|_1$. Let $H(S_k) = \iota_+(S_k)$ and the linear operator L be $L(S_k) = S_k \alpha_k^{(l)} V_k^\top + \tilde{S}^{(l)} \tilde{\alpha}^{(l)} \Pi_k$. For the moment, we consider Φ to be the identity.

To solve the algorithm, we need the proximal operator of H^* , the adjoint function of H , the proximal operator of G , and the gradient of F with its Lipschitz constant.

Starting with H , the proximal operator of H^* can be calculated directly using the proximal operator of the function H itself by means of the Moreau decomposition [Bec17, Theorem 6.44]. The proximal operator of an indicator function over a set \mathcal{C} is the orthogonal projection over that set. Therefore, we note $[X]_+$ the projection of $X \in \mathbb{R}^{n \times m}$ onto the positive orthant, that is,

$$\text{prox}_{\iota_+(\cdot)}(X) = [X]_+ \rightarrow [X_{i,j}]_+ = \begin{cases} X_{i,j} & \text{if } X_{i,j} \geq 0, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.2})$$

¹We use the implementation found in the python package github.com/CEA-COSMIC/ModOpt from Farrens et al. [Far+20].

Continuing with G , the proximal operator of the ℓ_1 norm is the soft thresholding operator which can be defined component-wise, for $x, \lambda \in \mathbb{R}$, as

$$\text{SoftThresh}_\lambda(x) = (|x| - \lambda)_+ \text{sign}(x) = \begin{cases} x - \lambda, & x \geq \lambda, \\ 0, & |x| < \lambda, \\ x + \lambda, & x \leq -\lambda. \end{cases} \quad (\text{A.3})$$

We name $L_{\nabla_{S_k} F(\cdot)}$ the Lipschitz constant of F 's gradient. The next equations resume what we need to use the chosen optimisation algorithm

$$\nabla_{S_k} F(S_k) = -\mathcal{F}_k^*(Y_k - \mathcal{F}_k(\hat{H}_k))V_k\alpha_k^\top, \quad (\text{A.4})$$

$$L_{\nabla_{S_k} F(\cdot)} = \rho(\mathcal{F}_k^* \circ \mathcal{F}_k)\rho(\alpha_k V_k^\top (\alpha_k V_k^\top)^\top), \quad (\text{A.5})$$

$$\text{prox}_{\tau G(\cdot)}([\mathbf{s}_{k,i}]_j) = \text{SoftThresh}_{\tau[\mathbf{w}_{k,i}]_j}([\mathbf{s}_{k,i}]_j), \quad (\text{A.6})$$

$$\text{prox}_{\sigma H^*(\cdot)}(X) = X - (X)_+, \quad (\text{A.7})$$

where the proximal operator of G is defined component-wise, the notation $[\mathbf{s}_{k,i}]_j$ represents the element j of the i column vector of matrix S_k , \mathcal{F}_k^* is the adjoint operator of \mathcal{F}_k , and $\rho(\cdot)$ is the spectral radius² that we calculate using the power method [GV96]. For the algorithm's parameters τ and σ , based on Condat [Con13, Theorem 3.1], we use

$$\tau = \frac{1}{\alpha L_{\nabla_{S_k} F(\cdot)}}, \quad \sigma = \frac{\alpha L_{\nabla_{S_k} F(\cdot)}}{2\|L\|_{op}^2}, \quad (\text{A.8})$$

where $\|\cdot\|_{op}$ is the operator norm [AB07] and α is a parameter we set to $3/2$. Being L a bounded linear operator we can calculate $\|L\|_{op}$ as $\sqrt{\rho(L^*L)}$ being L^* its adjoint operator.

We now consider the case where Φ is not the identity, but it is orthonormal, $\Phi^T \Phi = I$. We can adapt the soft thresholding operator in order to cope with the G term. This would be $\mathbf{s}_k \rightarrow \Phi^T \text{SoftThresh}_{\tau \mathbf{w}_k}(\Phi \mathbf{s}_k)$. When using undecimated wavelets as the starlets, the orthonormal condition is not met. Nevertheless, they are tight frames whose Gram matrix is close to the identity which means that the presented formulation will be a good approximation. We refer to Starck et al. [SMF15] for more information on wavelets.

The remaining optimisation problems

We deal in a similar way with the problems (II) and (IV) from Algorithm 1 using the same optimisation method proposed in Condat [Con13]. On the other hand, for problem (I), we use the optimisation algorithm in Liang et al. [LLS18]. This is due to the fact that we are neglecting the positivity constraint as we account for it when optimising over the other variables. In order to use these algorithms we need to

²The spectral radius can be defined as $\rho(B) = \max\{|\lambda_1(B)|, \dots, |\lambda_n(B)|\}$ where $\lambda_i(B)$ are the eigenvalues of the matrix B .

compute the gradients of the differentiable term of each problem as follows

$$\nabla_{S_k} F_k(x) = -\mathcal{F}_k^*(Y_k - \mathcal{F}_k(\hat{H}_k))(\alpha_k V_k^\top)^\top, \quad (\text{A.9})$$

$$\nabla_{\alpha_k} F_k(x) = -S_k^\top \mathcal{F}_k^*(Y_k - \mathcal{F}_k(\hat{H}_k))V_k, \quad (\text{A.10})$$

$$\nabla_{\tilde{S}} F(x) = \sum_{k=1}^N \nabla_{\tilde{S}} F_k(x) = \sum_{k=1}^N -\mathcal{F}_k^*(Y_k - \mathcal{F}_k(\hat{H}_k))(\tilde{\alpha} \Pi_k)^\top, \quad (\text{A.11})$$

$$\nabla_{\tilde{\alpha}} F(x) = \sum_{k=1}^N \nabla_{\tilde{\alpha}} F_k(x) = \sum_{k=1}^N -\tilde{S}^\top \mathcal{F}_k^*(Y_k - \mathcal{F}_k(\hat{H}_k))\Pi_k^\top, \quad (\text{A.12})$$

where $F = \sum_{k=1}^N F_k = \frac{1}{2} \|Y - \mathcal{F}(H + \tilde{S}\tilde{\alpha}\Pi)\|_F^2$. Concerning the global optimisation over \tilde{S} and $\tilde{\alpha}$, we need to consider all the CCDs when computing the gradient. So we can reformulate the global formulas as

$$\nabla_{\tilde{S}} F(x) = -\mathcal{F}^*(Y - \mathcal{F}(\hat{H}))(\tilde{\alpha}\Pi)^\top, \quad (\text{A.13})$$

$$\nabla_{\tilde{\alpha}} F(x) = -\tilde{S}^\top \mathcal{F}^*(Y - \mathcal{F}(\hat{H}))\Pi^\top. \quad (\text{A.14})$$

An approximation for the Lipschitz constants of the different gradients can be calculated as

$$L_{S_k} = \rho(\mathcal{F}_k^* \circ \mathcal{F}_k) \rho(\alpha_k V_k^\top (\alpha_k V_k^\top)^\top), \quad (\text{A.15})$$

$$L_{\alpha_k} = \rho(\mathcal{F}_k^* \circ \mathcal{F}_k) \rho(S_k^\top S_k) \rho(V_k^\top V_k), \quad (\text{A.16})$$

$$L_{\tilde{S}} = \rho(\mathcal{F}^* \circ \mathcal{F}) \rho(\tilde{\alpha}\Pi(\tilde{\alpha}\Pi)^\top), \quad (\text{A.17})$$

$$L_{\tilde{\alpha}} = \rho(\mathcal{F}^* \circ \mathcal{F}) \rho(\tilde{S}^\top \tilde{S}) \rho(\Pi \Pi^\top), \quad (\text{A.18})$$

where $\rho(\cdot)$ is the spectral radius.

Finally, we also need the proximal operator of the indicator function over the unit-ball $\iota_{\mathcal{B}}(\cdot)$, where $\mathcal{B} = \{x \in \mathbb{R}^n \mid \|x\|_2 = 1\}$. It can be computed as

$$\text{prox}_{\iota_{\mathcal{B}}(\cdot)}(x) = \frac{x}{\|x\|_2}. \quad (\text{A.19})$$

Sparsity enforcement parameters

There are two moments when we enforce sparsity during the optimisation. First, when we denote the eigenPSFs by the use of the ℓ_1 norm as in Equation 3.19. The \mathbf{w} weights are set depending on a noise estimation of the observed images, and the parameters K_σ^{Loc} and K_σ^{Glob} . The noise standard deviation is estimated using the median absolute deviation. The higher the K_σ parameters are set, the higher the thresholding and the denoising will be. Second, when we enforce the spatial constraints through α sparsity. In this case, we follow the sparsity enforcement proposed in Ngolè et al. [Ngo+16].

Publications and academic activities

Publications during this thesis

Accepted in peer-reviewed journals

- **Liaudat, T.**, Bonnin, J., Starck, J.-L., Schmitz, M. A., Guinot, A., Kilbinger, M., Gwyn, S. D. J., ‘Multi-CCD modelling of the point spread function’. In: *A&A* 646 (2021), A27
- Kervazo, C., **Liaudat, T.**, Bobin, J., ‘Faster and better sparse blind source separation through mini-batch optimization’. In: *Digital Signal Processing* 106 (2020), p. 102827
- Farrens, S. Guinot, A. Kilbinger, M. Liaudat, T. Baumont, L. Jimenez, X. Peel, A. Pujol, A. Schmitz, M. Starck, J.-L. Vitorelli, A. Z. ‘ShapePipe: A modular weak-lensing processing and analysis pipeline’. In: *A&A* 664 (2022), A141
- Guinot, A., Kilbinger, M., Farrens, S., Peel, A., Pujol, A., Schmitz, M., Starck, J.-L., Erben, T., Gavazzi, R., Gwyn, S., Hudson, M. J., Hildebrandt, H., **Liaudat, T.**, Miller, L., Spitzer, I., Van Waerbeke, L., Cuillandre, J.-C., Fabbro, S., McConnachie, A., ‘ShapePipe: a new shape measurement pipeline and weak-lensing application to UNIONS/CFIS data’. In: *arXiv:2204.04798* (2022)

Accepted in peer-reviewed conferences

- **Liaudat, T.**, Starck, J.-L., Kilbinger, M., Frugier, P.-A., ‘Rethinking the modeling of the instrumental response of telescopes with a differentiable optical model’. In: *NeurIPS 2021 Machine Learning for Physical sciences workshop*. Nov. 2021
- **Liaudat, T.**, Starck, J.-L., Kilbinger, M., ‘Semi-Parametric Wavefront Modeling for the Point Spread Function’. In: *52ème Journées de Statistique de la Société Française de Statistique (SFdS)*. Nice, France, June 2021

Submitted to peer-reviewed journals

- **Liaudat, T.**, Starck, J.-L., Kilbinger, M., Frugier, P.-A., ‘Rethinking data-driven point spread function modeling with a differentiable optical model’. In: *arXiv:2203.04908* (2022)

Open source contributions

- The [MCCD PSF](#) model presented in [chapter 3](#).
<https://github.com/CosmoStat/mccd>
- [ShapePipe](#): Shape measurement pipeline used to process [CFIS](#) to obtain the results shown in [chapter 4](#).
<https://github.com/CosmoStat/shapepipe>
- The [WaveDiff PSF](#) model presented in [chapter 5](#).
<https://github.com/tobias-liaudat/wf-psf>
- [WaveDiff](#) extensions and the experimentation presented in [chapter 6](#).
<https://github.com/tobias-liaudat/wf-SEDs>

Teaching

During the course of the thesis, I have been a teaching assistant in the applied mathematics department of the French *École Polytechnique (l'X)*.

I taught 64 hours¹ in the Academic year 2020-2021 in the following courses:

- Statistics [MAP433] for Master students with Prof Eric Moulines.
- Optimization and control [MAP435] for Master students with Prof Grégoire Allaire.
- Mathematical modelling [MAA107] for Bachelor students with Prof Vincent Bansaye and Prof Thibaut Mastrolia.
- Applied mathematics python projects [MAP361P] for Master students with Prof Arvind Singh.

I taught 32 hours in the Academic year 2021-2022 in the following course:

- Optimization and control [MAP435] for Master students with Prof Grégoire Allaire.

¹Legal maximum for a PhD student in an academic year.

Résumé étendu en français

Extended abstract in French

Sujet : Modélisation guidée par les données des fonctions d'étalement du point des télescopes terrestres et spatiaux.

Contexte

L'effet de lentille gravitationnel est la distorsion des images de galaxies lointaines par des objets massifs et constitue une sonde puissante de la structure à grande échelle de notre Univers. Les cosmologistes utilisent la lentille (gravitationnelle) faible pour étudier la nature de la matière noire et sa distribution spatiale. Ces études nécessitent des mesures très précises des formes des galaxies, mais la réponse instrumentale du télescope ou fonction d'étalement du point (PSF pour point spread function), déforme nos observations. Cette déformation peut être confondue avec des effets de lentille faible dans les images de galaxies, ce qui constitue l'une des principales sources d'erreur systématique lors de l'étude de la lentille faible. Par conséquent, l'estimation d'un modèle de PSF fiable et précis est cruciale pour le succès de toute mission de lentille faible. Le champ de PSF peut être interprété comme un noyau de convolution qui affecte chacune de nos observations d'intérêt variant spatialement, spectralement et temporellement. Le modèle de PSF devant prendre en compte ces variations et peut être contraint par des étoiles spécifiques dans le champ de vision. Ces étoiles, considérées comme des sources ponctuelles, nous fournissent des échantillons dégradés du champ de PSF. Les observations subissent différentes dégradations en fonction des propriétés du télescope, notamment un sous-échantillonnage, une intégration sur la bande passante de l'instrument et un bruit additif. Nous construisons finalement le modèle de PSF en utilisant ces observations dégradées d'étoiles, puis nous utilisons le modèle pour déduire la PSF aux positions des galaxies. Cette procédure constitue le problème inverse mal posé de la modélisation de la PSF. Le cœur de cette thèse a été le développement de nouveaux modèles de PSF basés sur les données, également connus sous le nom de modèles non paramétriques.

La prochaine génération de relevés cosmologiques à grand champ, tels que *Euclid*, observera l'Univers à une échelle sans précédent et générera une grande quantité d'informations précieuses. La qualité et la quantité toujours croissantes des données

posent des sérieux défis au modèle de PSF et imposent des exigences très strictes quant à ses performances.

Le problème de la modélisation de la PSF pour les études de lentilles faibles se situe à l'intersection de trois disciplines : l'optique, la cosmologie et les problèmes inverses en imagerie. Par conséquent, la première partie de cette thèse introduit tous les concepts nécessaires pour développer une compréhension globale de ce problème. Nous avons commencé par décrire les concepts de base de la cosmologie et de la lentille faible nécessaires pour comprendre l'objectif final du modèle de PSF dans l'analyse de la lentille faible. Nous avons poursuivi avec une vue d'ensemble des principaux éléments constitutifs d'un pipeline de mesure de forme pour décrire le contexte dans lequel le modèle de PSF s'inscrit. L'introduction se poursuit avec des notions essentielles en optique et une revue des contributeurs au domaine des PSF. Enfin, nous présentons un aperçu de l'état de l'art des modèles de PSF et concluons avec les méthodes de validation des modèles de PSF.

Modèles PSF pour les télescopes terrestres

Modélisation simultanée de la PSF dans le plan focal complet

Il existe plusieurs sources de variations spatiales de la PSF. Certaines couvrent l'ensemble du plan focal, comme les aberrations optiques du télescope et de l'atmosphère, tandis que d'autres sont spécifiques à chaque puce CCD (Charge-Coupled Device). Le plan focal des caméras d'imagerie à grand champ actuelles comprend un vaste réseau de CCDs, constituant une mosaïque de CCDs. La plupart des modèles de PSF actuels sont construits indépendamment sur chaque CCD, ce qui est une solution simple pour tenir compte des deux variations. Cependant, ce choix présente plusieurs inconvénients. Il limite le nombre d'étoiles disponibles pour contraindre le modèle, favorisant ainsi des modèles plus simples. En outre, les variations spatiales couvrant l'ensemble du plan focal ne peuvent pas être correctement modélisées si elles sont approximées par des modèles simples par morceaux. Par conséquent, des erreurs de modélisation de la PSF apparaissent. Afin de résoudre les deux problèmes simultanément, nous avons proposé un nouveau modèle de PSF, appelé MCCD, qui peut modéliser le plan focal complet et gérer les deux types de variations. MCCD peut inclure avec succès la géométrie de la caméra dans le modèle de PSF. Notre modèle de PSF est basé sur un schéma de factorisation matricielle qui utilise différents outils mathématiques, tels que les régularisations parcimonieuses pour le débruitage de la PSF et la théorie des graphes pour traiter les variations spatiales localisées. Nous avons développé l'algorithme d'apprentissage en combinant ces concepts avec la descente de coordonnées par blocs, des méthodes efficaces d'optimisation convexe et des algorithmes proximaux. Nous avons validé le modèle à l'aide de simulations et d'observations réelles. Enfin, nos résultats montrent que le modèle de PSF proposé peut construire des modèles complexes, améliorer les performances par rapport aux méthodes de pointe et être suffisamment robuste pour traiter des données réelles.

Modélisation de la PSF pour UNIONS/CFIS

Le Canada-France Imaging Survey (CFIS) est un relevé en cours dans les bandes u et r au télescope Canada-France-Hawaii (CFHT). Il couvrira 4800 deg^2 , et son objectif principal est de contribuer aux relevés profonds comme *Euclid* avec son estimation photométrique du redshift. L'excellente qualité des images de la bande r permet de mener des recherches sur les lentilles faibles. J'ai contribué au pipeline de mesure de




FIGURE C-1 – Différentes représentations d’une PSF spatiale à une position donnée dans le champ de vision. (a) PSF polychromatique bruitée à la résolution d’observation d’*Euclid*, (b) observation polychromatique non bruitée à la résolution d’observation d’*Euclid*, (c) observation polychromatique non bruitée à haute résolution, (d) carte d’erreur de front d’onde représentant les aberrations du système optique. Les unités sont en μm , (e) Variations chromatiques à haute résolution de la PSF à des longueurs d’onde équidistantes dans la bande passante [550, 900]nm.

forme de haute performance, **ShapePipe**, dans plusieurs aspects de la modélisation du PSF et des tests de validation. Nous avons produit un catalogue de formes utilisant 1700 deg^2 avec le modèle PSF largement utilisé **PSFEx**.

Plus tard, nous avons inclus le modèle de PSF MCCD dans **ShapePipe** et validé ses performances avec des données réelles. Nous avons exécuté le pipeline dans un environnement de calcul haute performance avec notre modèle PSF pour produire un nouveau catalogue de formes couvrant $\sim 3500 \text{ deg}^2$, l’une des plus grandes zones pour les études de lentille gravitationnelle faible au moment de la rédaction. Nous présentons ensuite les résultats de plusieurs méthodes de validation de la PSF utilisant le nouveau catalogue. Les résultats montrent que le modèle MCCD améliore les performances du modèle de PSF par rapport au modèle précédent et confirme la robustesse et la maturité de la méthode pour le traitement des données réelles.

Modèles de PSF pour les télescopes spatiaux

Repenser la modélisation PSF basée sur les données avec un modèle optique différentiable

Le relevé de *Euclid* représente l’un des défis les plus difficiles à relever pour la modélisation de la PSF. En raison de la très large bande passante de l’imageur visible (VIS) de *Euclid* allant de 550 nm à 900 nm, les modèles de PSF doivent capturer non seulement les variations spatiales du champ de PSF mais aussi ses variations chromatiques. Chaque observation d’étoile est intégrée sur l’ensemble de la bande passante VIS avec la distribution spectrale d’énergie (SED) de l’objet. Comme les observations sont sous-échantillonnées, une étape de super-résolution est également nécessaire. La [Figure C-1](#) présente un exemple de PSF de type *Euclid* ainsi que son sous-échantillonnage et ses variations chromatiques. Les modèles actuels de PSF basés sur




FIGURE C-2 – Schéma du cadre proposé pour la modélisation basée sur les données de la PSF basée sur le front d'onde pour la mission spatiale *Euclid*.

les données ne peuvent pas modéliser les variations chromatiques. Le modèle de pointe conçu pour *Euclid*, le modèle *resolved component analysis* (RCA), est respectivement ~ 200 et $\sim 10^4$ fois supérieur aux exigences d'erreur de forme et de taille. Ces modèles sont tous construits dans l'espace des pixels avec une méthode de réduction de la dimensionnalité utilisant différentes contraintes. Nous proposons un changement de paradigme dans la façon dont les modèles de PSF basés sur les données sont construits.

En intégrant un modèle optique différentiable dans le cadre de modélisation proposé, l'espace de modélisation non paramétrique passe des pixels au front d'onde. Le modèle est basé sur l'optique de Fourier et des principes fondamentaux de l'optique. Notre cadre, présenté dans [Figure C-2](#), permet de construire de puissants modèles interprétables motivés physiquement et qui ne nécessitent pas de données de calibration spéciales. Le nouveau modèle, baptisé WaveDiff, s'appuie sur un schéma de factorisation matricielle et des polynômes de Zernike. Le modèle s'appuie sur des méthodes modernes basées sur l'utilisation de gradients et la différenciation automatique pour l'optimisation, utilisant uniquement des observations bruitées à large bande et focalisées. WaveDiff est le premier modèle basé sur les données qui peut modéliser avec succès les variations chromatiques en plus de gérer les variations spatiales et la super-résolution (x3). Il représente également une avancée en termes de performances pour les modèles pilotés par les données en diminuant respectivement plus d'un et de deux ordres de grandeur les exigences d'erreur de taille et de forme d'*Euclid*. Le modèle proposé est modulaire, différentiable, construit en TensorFlow et fonctionne entièrement sur GPU.

Application de WaveDiff à *Euclid*

Le modèle de PSF WaveDiff constitue une approche prometteuse pour la mission *Euclid*. Néanmoins, passer d'une bonne performance sur des simulations à une bonne performance sur des données réelles est loin d'être trivial. Ce chapitre aborde plusieurs questions qui émergeront nécessairement lorsque les observations d'*Euclid* seront rendu disponibles. Nous espérons que les études présentées dans ce chapitre permettront de faciliter la transition entre les simulations et les données réelles pour le modèle WaveDiff. Nous commençons cette réflexion à travers une liste de questions abordées :

- Des informations optiques préalables sur le télescope *Euclid* seront disponibles.

Ces informations peuvent provenir de simulations de haute fidélité ou d'observations complémentaires dans l'espace. Cependant, on s'attend à ce que l'information préalable comporte un certain degré d'erreur. Nous modifions le modèle WaveDiff afin d'inclure cette information optique préalable et de la corriger de manière non paramétrique. Les résultats montrent que le modèle peut exploiter l'information préalable et améliorer considérablement les performances du modèle.

- Dans les études précédentes, nous avons supposé que la dimension du front d'onde était la même pour générer les observations, c'est-à-dire le modèle de vérité terrain, et pour WaveDiff. Nous étudions comment la dimension du front d'onde affecte la modélisation de la PSF lorsque les observations sont générées de manière plus réaliste.
- Jusqu'à présent, nous n'avons pas utilisé plus que le nombre d'étoiles dans une exposition pour contraindre le modèle WaveDiff. La stabilité d'*Euclid* pourrait nous permettre d'utiliser des expositions ultérieures et donc d'augmenter le nombre d'étoiles pour contraindre le modèle de PSF. Par conséquent, nous étudions comment le modèle WaveDiff exploite les informations provenant d'un champ stellaire plus dense.
- Nous avons considéré, dans les études précédentes, que les informations SED d'entrée étaient sans défaut et avons utilisé 20 bins pour générer les observations et pour le modèle WaveDiff. Dans un scénario réel, les observations sont générées sans aucune discrétisation spectrale. En outre, les informations SED réelles sont fournies dans un nombre réduit de bins contenant des erreurs. Nous étudions comment la discrétisation spectrale affecte la modélisation de la PSF et des moyens d'améliorer les résultats.
- Les résultats précédents ont montré la capacité du modèle WaveDiff à estimer une représentation utile du front d'onde qui nous permet d'obtenir une faible erreur de pixel. Cependant, supposons que nous soyons dans un nouveau scénario où la partie paramétrique du modèle WaveDiff peut reproduire le champ de front d'onde de la vérité terrain. Dans ce cas, le modèle WaveDiff actuel ne peut pas récupérer ce champ de front d'onde. Il s'agit en effet d'un problème de récupération de phase. Néanmoins, dans ce nouveau scénario, la solution au problème de récupération de phase est également le minimum global de notre problème de modélisation du PSF. Nous proposons une nouvelle procédure d'optimisation pour résoudre le problème de reconstruction de phase. Cette procédure est basée sur la projection d'informations de la partie non paramétrique vers la partie paramétrique du modèle. Nous montrons à l'aide d'expériences numériques que le modèle WaveDiff peut estimer le champ de front d'onde de la vérité terrain avec de faibles erreurs en utilisant uniquement des observations dégradées focalisées.

* * *
* *
*

List of Figures

01	Different representations of a space-like PSF at a single position in the field of view. (a) Noisy polychromatic PSF at observation resolution, (b) Noiseless polychromatic observation at observation resolution, (c) High-resolution noiseless polychromatic observation, (d) Wavefront error map representing the aberrations in the optical system. Units are in μm , (e) High-resolution chromatic variations of the PSF at equally spaced wavelengths in the passband $[550, 900]\text{nm}$	3
02	A schematic of the proposed framework for data-driven wavefront PSF modelling for the <i>Euclid</i> space mission.	4
1.1-1	Deep image from JWST using the Near-Infrared Camera, which shows the gravitational lensing effect of the galaxy cluster SMACS 0723. Credit: NASA and ESA.	12
1.2-2	Illustration of two light rays emitted from an object in the source plane, located at χ , that are then lensed by a massive object in the lens plane, located at χ' , and that finally arrive at an observer located at the origin. The bottom light ray is considered the fiducial ray. The apparent angular separation of the rays is θ , while the angle between the true positions of the source, without any lensing effect, is β . The angle α represents the difference between the two previous angles. Credit: Illustration from Guinot [Gui20].	14
1.2-3	Illustration of the effect of shear on the circular object at the origin. The complex shear amplitude is left constant while the angle is changed from 0 to 2π , showing the evolution of the orientation of the ellipse. The ellipse's eccentricity is constant throughout the angles as the shear amplitude remains unchanged. Credit: Image from Kilbinger [Kil15].	16
1.2-4	Illustration of the shear and convergence effects. The circular source image is \mathcal{S} , and lensed image is \mathcal{I} . The effect of the convergence κ , shown in red, is an isotropic magnification or contraction. The effect of the shear γ , shown in blue, is an anisotropic deformation. The orientation of the ellipse ϕ depends on the complex γ phase, and the ellipse's eccentricity depends on the complex γ amplitude. The amplification matrix \mathcal{A} relates the image coordinates, providing a linear transformation of (θ_1, θ_2) to (β_1, β_2) . Credit: Image from Ajani [Aja21].	17

1.2-5	Mass maps using DES year 3 data release generated with the KS method [KS93], and the sparsity-based method GLIMPSE [Lan+16]. Credit: Figures from Jeffrey et al. [Jef+21].	20
1.3-6	Overview of the main building blocks of a shape measurement pipeline. Circles represent data or models, and rectangles represent processes. <i>PSF catalogue</i> represents the PSFs at the position of galaxies, while <i>Test PSF catalogue</i> represents the PSFs at the position of the test stars. We present the colour description of the different stages: <i>green</i> represent the common input data processing; <i>red</i> represent the PSF modelling part; <i>violet</i> represent the PSF validation part; <i>blue</i> represent the shape measurement part; <i>yellow</i> represent the shape measurement calibration; and <i>grey</i> the shape catalogue validation.	25
1.3-7	Example of a CCD image and the mask produced by ShapePipe. The colour bar shows the pixel mask values. Credit: Image from Farrens, S. et al. [Far+22a].	26
1.3-8	ShapePipe star selection from a magnitude-size diagram. The size is represented as FWHM expressed in pixels. The FWHM expressed in the legend is in arcsec. Credit: Image from Farrens, S. et al. [Far+22a]. . . .	27
1.4-9	A contour plot illustrating the potential perceived by an object at rest in the Sun-Earth system. The different Lagrange points can be identified. Credit: NASA.	31
1.4-10	The observation regions of <i>Euclid</i> as from an L2 point perspective. The four regions in blue are the regions of interest for the <i>Euclid</i> wide survey. The red and yellow regions correspond to the galactic and ecliptic planes, respectively, and will be avoided. The galactic plane contains objects from our galaxy and is thus uninteresting to probe far-away zones of the Universe. The zodiacal light from our Sun contaminates the ecliptic plane. Credit: Image from Euclid Collaboration et al. [Euc+22a].	32
2.1-1	Illustration of the Huygens-Fresnel principle and the modification of a wavefront due to an obstacle.	36
2.1-2	Illustration of the coordinate system for the diffraction equations. Credit: figure from Ngolè Mboula, Fred Maurice [Ngo16].	37
2.1-3	Illustration of the different diffraction regions behind an aperture. Credit: Figure from Gross [Gro05].	39
2.1-4	Illustration of the generalised <i>black box</i> concept of the imaging system and its corresponding coordinate systems. The image plane coordinates are (u, v) , the output aperture plane coordinates are (x, y) , and the object plane coordinates are (ξ, η) . Credit: Figure from Goodman [Goo05]. . . .	40
2.1-5	Illustration of the isoplanatic region. Two rays from the same isoplanatic region travel through almost the same turbulence and suffer almost the same distortions.	41
2.1-6	Illustration of the Wavefront Errors (WFE) in a one-dimensional setting where the optical system is represented as a single lens.	43
2.1-7	Example of <i>Euclid</i> 's pupil function in the (x, y) plane for a given position in the (ξ, η)	44
2.2-8	The 3rd generation set of filters of the MegaCam instrument at CFHT currently being used for CFIS. The transmission filter response includes the full telescope and 1.25 airmasses of atmospheric attenuation. The full telescope includes the mirrors, the optics, and the CCDs.	48

<i>List of Figures</i>	217
2.3-9 Surface errors measured for the HST which are in the range of $\pm 30\text{nm}$. Credit: Figure from Krist et al. [KHS11].	51
2.3-10 Example of two different pixellations on the same high-resolution image representing an Airy PSF. The difference between the two pixellations is an intra-pixel shift of $(\Delta x, \Delta y) = (0.35, 0.15)$ between them.	51
2.3-11 <i>Euclid</i> 's requirements on the CCD QE for the VIS instrument from the MDB.	54
2.3-12 Illustration of six von Kármán phase screen layers at different altitudes simulated for LSST. Credit: Figure from Jee et al. [JT11].	55
2.3-13 Example of atmospheric PSFs with different exposure times. The simulation was done using the atmospheric parameters from Jee et al. [JT11] for an LSST-like scenario.	56
2.5-14 Illustration of <i>Euclid</i> 's optical system. Credit: Image from Racca et al. [Rac+16].	67
2.5-15 Variations of the different optical elements' footprints due to 9 incoming beams from different FOV positions. The figures were computed using ray-tracing simulations. Credit: Image from the EC technical note EUCL-OXF-TN-8-012 produced by Prof Lance Miller.	68
2.5-16 Angle of incidence with respect to the normal of each optical element from rays incoming from one FOV position. The dichroic angles are plotted for another FOV position to illustrate its change. The colour bars are in degrees, and the x and y axes are in millimetres. The figures were computed using ray-tracing simulations. Credit: Image courtesy of Dr Pierre-Antoine Frugier from the EC.	70
3.4-1 Example of the interpolation procedure involved in the PSF recovery. . .	91
3.5-2 Star positions in CFHT's MegaCam used for the simulated dataset. The positions were taken from a real CFIS exposure.	92
3.5-3 Shape measurement results of the simulated test star catalogue following the analytical ellipticities.	94
3.5-4 RMSE on pixels, shape, and size metrics as a function of stars SNR for the three main methods. The RMSE are plotted with solid lines and the gain of the methods with respect to PSFEx are plotted with dashed lines. . . .	96
3.5-5 Moment residual maps comparing the MCCD-RCA algorithm on the left and the PSFEx algorithm on the right. They are obtained by subtracting the model's and the test star's measured shape and size metrics and plotting them on their corresponding position over the focal plane. The SNR value of the star dataset is 50.	97
3.5-6 Residual $\Delta R_2/R_2$ map of the RCA algorithm of stars with an SNR of 50. The CCDs where the RCA model is having degeneracies that can be clearly spotted on the map.	98
3.5-7 Comparison of the performance of MCCD-POL, MCCD-RCA, and MCCD- HYB methods in terms of the RMSE on the pixels, the shape and the size metric of the star's SNR. The RMSE are plotted on a solid line and the gain of the method with respect to PSFEx is plotted on dashed lines. . .	99

3.5-8	Example eigenPSFs extracted from the MCCD-HYB PSF model trained on the simulated dataset with an SNR of 70. The local eigenPSFs were extracted from the graph’s spatial constraint of a central CCD. It can be seen from the eigenPSFs that the global model is specialising on the shape of the PSF while the local model specialises on capturing its ellipticity. It is also worth to mention that the first global eigenPSF found on the first row provides the baseline isotropic PSF the model uses.	100
3.6-9	Star density of all the training dataset with respect to their position on the MegaCam’s focal plane. We have on average 1560 training stars per exposure.	101
3.6-10	Ellipticities and size <i>meanshapes</i> maps (see subsection 2.6.2 for more information about the maps) measured on the training stars of our CFIS dataset.	102
3.6-11	Stacked star profile from all 534 test stars in a random CFIS exposure (id. 2099948) (top), corresponding stacked residuals after subtraction by the PSFEx (middle left) or MCCD-HYB (middle right) PSF models. The bottom row includes the stacked absolute value of the residuals.	104
3.6-12	Examples of real CFIS test stars, the estimations of both methods, and the pixel residuals. The exposures are the same as in Figure 3.6-11. We present four test stars with the estimated PSF models and the corresponding residuals. The top-left star corresponds to a star extracted from the top-left corner of the focal plane with an SNR of 19. The top-right star corresponds to a star leading to a high error for both methods. The bottom-left star corresponds to a star located in the centre of the focal, with a relatively high SNR of 160. The bottom-right star corresponds to a star located in the bottom-right corner of the focal plane, with a low SNR of 4.	105
3.6-13	Histograms showing the distribution of the size metric over the training stars and their residuals for both PSF models, PSFEx and MCCD-HYB.	106
4.1-1	CFIS r-band sky coverage as of March 2022. Credit: UNIONS collaboration.	108
4.1-2	CFIS delivered image quality for the r- and u-band as of March 2022. Credit: UNIONS collaboration.	108
4.1-3	CFIS r-band DR3 used to produce the v1 shape catalogue. The total area is divided into seven patches from P1 to P7. Credit: UNIONS collaboration.	109
4.3-4	Scale-dependent PSF leakage computed over each patch of the v1 shape catalogue and the entire field.	113
4.3-5	Test star density as a function of focal plane position.	113
4.3-6	Meanshapes plots of both ellipticity components of the observed test stars and the PSF model. The first column corresponds to the first ellipticity component, while the second column to the second ellipticity component. The first row, figures (4.3-6a) and (4.3-6b), corresponds to the observed stars. The second row, figures (4.3-6c) and (4.3-6d), correspond to the PSF model. The third column, figures (4.3-6e) and (4.3-6f), correspond to the residual (star – PSF) between the observed stars and the PSF model.	114
4.3-7	Meanshapes plots of the observed stars and PSF’s sizes. The top left plot, figure (4.3-7a), shows the stars size, while the top right plot, figure (4.3-7b), shows the PSF size. The bottom plot, figure (4.3-7c), shows the residual between the stars and the PSF sizes that is computed as $ (R_{\text{star}}^2 - R_{\text{PSF}}^2)/R_{\text{star}}^2 $	115

4.3-8	The ρ -statistics computed on the PSF model and the test stars. The top row shows the results for the v1 catalogue, while the bottom row shows the results from the previous version of the catalogue from Guinot et al. [Gui+22]. Credit: second row plots from Guinot et al. [Gui+22].	116
5.2-1	Different representations of a space-like PSF at a single position in the field-of-view. The SED used corresponds to one of the templates from the numerical experiments described in subsection 5.5.2. (a) Noisy polychromatic PSF at observation resolution, (b) Noiseless polychromatic observation at observation resolution, (c) High-resolution noiseless polychromatic observation, (d) Wavefront error map representing the aberrations in the optical system. Units are in μm , (e) High-resolution chromatic variations of the PSF at equally spaced wavelengths in the passband [550, 900]nm.	125
5.4-2	A schematic of the proposed framework for data-driven wavefront PSF modelling.	130
5.4-3	The first 15 Zernike polynomials with their Noll's single-index prefaced by the hash symbol #, and ordered vertically by their radial degree n and horizontally by their azimuthal degree m	132
5.5-4	Templates for the stellar SEDs used to build the polychromatic PSFs.	141
5.5-5	Example of PSFs from the simulated data set. The first row contains images of observed stars, while the second row contains the corresponding WFE map used to generate each one.	141
5.5-6	Examples of learned data-driven eigenWFEs, or WFE features, of the WaveDiff-original model. These features correspond to polynomial feature weight functions. Units are in μm	145
5.5-7	Visual reconstruction of a test PSF with the WaveDiff-original model. The first column shows the ground truth noiseless PSF, the second column the reconstruction of the PSF model, and the third column the residual between the ground truth and the model. First row: Polychromatic pixel representation at the observation resolution. Second row: Polychromatic representation at 3 times the observation resolution. Third row: Monochromatic pixel representation at 750 nm and at 3 times the observation resolution. Fourth row: WFE representation of the PSF, where units are in μm . We have removed the mean of the WFE and applied the obscurations. The first three rows show the absolute value of the residual, while the last one is just the residual.	146
5.5-8	Test star reconstruction error as a function of wavelength for the WaveDiff-original model at the observation resolution. Left axis: Absolute error, or Err_{abs} , and is plotted with a solid line. The error bars represent the variance over the RMS errors of the different test stars being used for the calculation of Err_{abs} . Right axis: Relative error, or Err_{rel} , and is plotted with a dashed line.	147
5.5-9	Polychromatic PSF relative reconstruction error, Err_{rel} , at observation resolution as a function of the total number of training stars in the FOV. The training data sets correspond, in ascending order, to $\mathcal{S}_1 \subset \mathcal{S}_2 \subset \mathcal{S}_3 \subset \mathcal{S}$, described in subsection 5.5.2. The results are computed with respect to the reconstruction of test stars. The set of test stars is maintained throughout the models and the training sets, allowing a truthful comparison. The result of RCA for the \mathcal{S}_1 (200 stars) data set is not shown as it is very poor.	149

5.5-10	Polychromatic PSF relative reconstruction error, Err_{rel} , at three times (x3) the observation resolution as a function of the total number of training stars in the FOV. This experiment represents the super-resolution task.	149
5.6-11	Illustration of the second scenario of the WFE-based PSF modelling. The ellipse on the left represents the possible PSF fields defined in the WFE space. The ellipse on the right, the space of pixel PSFs representations. The forward model allows to go from a WFE-defined PSF model into a set of pixel PSF. The space of possible models given a set of WaveDiff hyperparameters is \mathcal{F} . If we assume that $f_{GT} \notin \mathcal{F}$, the closest PSF field in terms of WFE to the GT, \hat{f}_2 , is not the best choice in terms of pixel error.	151
6.2-1	The figure (a) shows the simulation's random positions in the FOV for the train and test datasets. The figure (b) shows the spatial distribution of the total amount of aberrations in the FOV.	158
6.2-2	Histogram of the total amount of aberrations in WFE RMS of the optical simulations with and without Surface Errors (SFE).	158
6.2-3	Spatial distribution of the Zernike coefficients from order 3 to 22 from the <i>Euclid</i> 's optical simulations. Units are in nm RMS.	159
6.2-4	Spatial distribution of the Zernike coefficients from order 23 to 42 from the <i>Euclid</i> 's optical simulations. Units are in nm RMS.	160
6.2-5	Spatial distribution of the realisation of the total aberration error map, $\Delta\Phi^{Err}$, for the 10nm RMS case.	160
6.2-6	Spatial distribution of some example Zernike coefficients from the error map with 10nm mean WFE RMS realisation. Units are in nm RMS.	161
6.2-7	Relative pixel reconstruction RMSE at 3 times the <i>Euclid</i> resolution for different levels of error in the prior. The x -axis shows the amount of error added to the simulations to form the prior expressed as a percentage of the mean WFE RMS of the simulations.	162
6.2-8	Ellipticity errors expressed relative to the <i>Euclid</i> requirement, which for each ellipticity component corresponds to 2×10^{-4}	163
6.2-9	Size error expressed relative to the <i>Euclid</i> requirement, which corresponds to 1×10^{-3}	163
6.2-10	Relative pixel reconstruction RMSE as a function of wavelength for the prior and the WaveDiff model using only the non-parametric, or data-driven, part. The error done by the prior is shown with dashed lines. We have omitted the curves for higher prior errors to keep the limit of the y -axis low.	164
6.2-11	Relative WFE reconstruction error for different level of error in the prior.	165
6.3-12	The left axis shows the relative pixel RMSE at one and three times the <i>Euclid</i> resolution as a function of the modelling WFE dimension. The right axis shows the computing time required to train the corresponding WaveDiff model. The dotted lines represent the error between two datasets that share all their properties except that they were generated with WFE dimensions of 4096 and 256. The blue and green lines correspond to errors computed at three and one times the <i>Euclid</i> resolution.	167
6.4-13	Relative pixel RMSE at the <i>Euclid</i> resolution. The error bars represent the standard deviation of the experiment repetitions.	169
6.4-14	Relative pixel RMSE at three times the <i>Euclid</i> resolution. The error bars represent the standard deviation of the experiment repetitions.	169
6.4-15	Relative WFE RMSE. The median results and the best performance of the repetitions are drawn in solid and dashed lines, respectively. The error bars represent the standard deviation of the experiment repetitions.	170

6.5-1	Relative pixel PSF RMSE at three times the <i>Euclid</i> resolution for three types of stellar SEDs using a different number of bins in a noiseless scenario.	173
6.5-1	Pixel errors for PSFs at three times the <i>Euclid</i> resolution. Each row corresponds to a different stellar SED type. The first column shows the PSF generated with 350 bins, the second column the PSF generated using only 4 bins, and the third column the absolute difference of the residual between the other columns.	173
6.5-1	Relative pixel PSF RMSE at three times the <i>Euclid</i> resolution for different reference σ values as a function of the number of SED bins. Only the error from the SED uncertainty is considered. The errors are computed between the PSFs generated by noisy and binned SEDs and the PSFs obtained with just the binned SEDs.	174
6.5-1	Relative pixel PSF RMSE at three times the <i>Euclid</i> resolution considering the errors due to the binning and to the uncertainty in the SED. The errors are drawn as a function of the number of bins used and the noise standard deviation.	175
6.6-2	Illustration of the bins' wavelength centres as we add new interpolation bins. The example shows an SED with 8, 16, and 24 bins, where the input SED is the 8-bin one, and the passband goes from 550nm to 900nm.	175
6.6-2	Interpolated SEDs for different reference σ values using a linear interpolant. The GT SED is drawn in black. Each colour is one of the reference σ values.	176
6.6-2	Relative RMSE of the super-resolved PSFs for each σ level and as a function of the number of bins up to 32 bins, using linear interpolation for the new SED points.	177
6.6-2	Comparison of the interpolation with 16 and 24 bins, where we add one and two points between the input 8-bin points, respectively. The input SED errors shown have a different sign.	177
6.7-2	Relative pixel PSF RMSE at three times the <i>Euclid</i> resolution for the WaveDiff-original model using different SED inputs. The GT PSFs are generated with noiseless 350-bin SEDs. The input SEDs has 8 bins with different levels of noise that are represented in the x -axis. The WaveDiff SED interpolation is compared between 8 (no interpolation), 16, and 32 bins.	178
6.7-2	Relative pixel PSF RMSE at the <i>Euclid</i> resolution for the WaveDiff-original model for different SED interpolations as a function of the input 8-bin SED noise level.	179
6.8-2	Amplitude distribution of the first 15 Zernike coefficients. The difference between the GT and the estimated coefficients of the model's parametric part is plotted in blue. The non-parametric contribution to each coefficient is plotted in green.	181
6.8-2	Illustration of the tensor product involved in the DD part of the model.	184
6.8-2	Illustration of the tensor product using the mixed DD features $\tilde{\mathbf{S}}^{\text{DD}}$, involved in the DD part of the model.	185
6.8-2	Comparison between the WFE map from the parametric part of the model before and after the projection of the DD part, and the GT WFE. The left image shows the WFE of the parametric part of a WaveDiff-original model once the training has finished. The image corresponds to a PSF reconstruction for a random point in the FOV. The centre image shows the new parametric part once the projection of the DD part was added. The right image shows the GT WFE of the PSF field projected over the first n_Z Zernike polynomials.	186

6.8-3A	Amplitude of the first $n_z = 15$ Zernike polynomial coefficients. The parametric part of the model before and after the DD projection are shown in blue and orange, respectively. The GT PSF field values are shown in green. We show, in red, the total WFE prediction of the WaveDiff-original model projected onto these first Zernike polynomials. The WFEs used come from a trained WaveDiff-original model and its WFE recovery of a random position in the FOV.	187
6.8-3B	Loss function as a function of the number of epochs for the DD part of the model in the second optimisation cycle. One model has been reset after performing the projection, and the other model has not been reset.	188
6.8-3C	WFE representations for a given training point in the FOV. The first row shows, from left to right, the parametric, DD, and the full WFE prediction for the new optimisation algorithm using 9 cycles. The second row shows the residual WFE for the parametric and the full WFEs in the first and third columns, respectively, with respect to the ground truth WFE that is placed in the second column.	191
6.8-3D	Evolution of different WFE parts of the model as the number cycle increases. The same scale is used for all the images in each figure.	192
C-1	Différentes représentations d'une PSF spatial à une position donnée dans le champ de vision. (a) PSF polychromatique bruitée à la résolution d'observation d' <i>Euclid</i> , (b) observation polychromatique non bruitée à la résolution d'observation d' <i>Euclid</i> , (c) observation polychromatique non bruitée à haute résolution, (d) carte d'erreur de front d'onde représentant les aberrations du système optique. Les unités sont en μm , (e) Variations chromatiques à haute résolution de la PSF à des longueurs d'onde équidistantes dans la bande passante [550, 900]nm.	211
C-2	Schéma du cadre proposé pour la modélisation basée sur les données de la PSF basée sur le front d'onde pour la mission spatiale <i>Euclid</i>	212

List of Tables

11	Main cosmological parameters for Λ CDM with a 68% confidence limit from Planck Collaboration et al. [Pla+20, Tab. 2] using a combination of CMB TT, TE, EE + lowE + CMB lensing + BAO.	12
31	Summary of important variables used in this chapter.	81
32	Q_p criterion using all test stars of the W_3 dataset from CFIS. The gain of the MCCD-HYB with respect to PSFEx and the noise standard deviation are also presented.	103
41	Summary of the MCCD PSF model parameters used for the v1 shape catalogue.	111
51	Optic variables used for the optical forward model.	135
52	Hyperparameters of the different PSF models that use the optical forward model. All the models use a batch size of 32. The WaveDiff-graph and the WaveDiff-polygraph use a regularization parameter μ of 1×10^{-8}	143
53	Polychromatic test star reconstruction Err_{abs} and Err_{rel} at the observation resolution (x_1) and at super-resolution (x_3). The presented results were obtained with models being trained on the largest training star set \mathcal{S}	144
54	Target stars' WFE recovery errors with the PSF models being trained on the \mathcal{S} data set.	150
55	Moment-based metrics of the different PSF models trained with the \mathcal{S} data set.	150
61	Performance results for the WaveDiff model resetting the DD part after each projection, as in point (iii). Each cycle uses 75 epochs to ensure that the DD part optimisation converges as mentioned in point (iv).	190
62	Performance results for the WaveDiff model <i>without</i> resetting the DD part after each projection. Each cycle uses 75 epochs to ensure that the DD part optimisation converges as mentioned in point (iv).	190
63	Performance results for the WaveDiff model resetting the DD part after each projection, as in point (iii). In this case, each cycle uses only 15 epochs.	191

Bibliography

- [AB07] Aliprantis, C., Border, K., *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, 2007 (page 204).
- [AC11] Agueh, M., Carlier, G., ‘Barycenters in the Wasserstein Space’. In: *SIAM Journal on Mathematical Analysis* 43.2 (2011), pp. 904–924 (page 127).
- [Aja21] Ajani, V. ‘Higher order statistics for cosmology : likelihood development for future surveys like Euclid’. 2021UNIP7071. PhD thesis. 2021 (page 17).
- [Alb+06] Albrecht, A., Bernstein, G., Cahn, R., Freedman, W. L., Hewitt, J., Hu, W., Huth, J., Kamionkowski, M., Kolb, E. W., Knox, L., Mather, J. C., Staggs, S., Suntzeff, N. B., *Report of the Dark Energy Task Force*. 2006 (page 30).
- [Ant+14] Antilogus, P., Astier, P., Doherty, P., Guyonnet, A., Regnault, N., ‘The brighter-fatter effect and pixel correlations in CCD sensors’. In: *Journal of Instrumentation* 9.03 (Mar. 2014), pp. C03048–C03048 (page 53).
- [BA14] Bernstein, G. M., Armstrong, R., ‘Bayesian lensing shear measurement’. In: *Monthly Notices of the Royal Astronomical Society* 438.2 (Jan. 2014), pp. 1880–1893 (pages 29, 64).
- [BA96] Bertin, E., Arnouts, S. ‘SExtractor: Software for source extraction’. In: *Astron. Astrophys. Suppl. Ser.* 117.2 (1996), pp. 393–404 (pages 26, 60).
- [Bae06] Baer, R. L. ‘A model for dark current characterization and simulation’. In: *Sensors, Cameras, and Systems for Scientific/Industrial Applications VII*. Ed. by M. M. Blouke. Vol. 6068. International Society for Optics and Photonics. SPIE, 2006, pp. 37–48 (page 53).
- [Bar+22] Baron, M., Sassolas, B., Frugier, P.-A., Venancio, L. M. G., Amiaux, J., Castelnau, M., Keller, F., Dovillaire, G., Treimany, P., Juvénal, R., Miller, L., Pinard, L., Ealet, A., ‘Measurement and modelling of the chromatic dependence of a reflected wavefront on the Euclid space telescope dichroic mirror’. In: *Space Telescopes and Instrumentation 2022: Optical, Infrared, and Millimeter Wave*. Ed. by L. E. Coyle, S. Matsuura and M. D. Perrin. Vol. 12180. International Society for Optics and Photonics. SPIE, 2022, p. 121804V (page 50).
- [Bec17] Beck, A. *First-Order Methods in Optimization*. MOS-SIAM Series on Optimization. Society for Industrial and Applied Mathematics, 2017 (page 203).

- [Ber+16] Bernstein, G. M., Armstrong, R., Krawiec, C., March, M. C., ‘An accurate and practical method for inference of weak gravitational lensing from galaxy images’. In: *Monthly Notices of the Royal Astronomical Society* 459.4 (Apr. 2016), pp. 4467–4484 (pages 29, 64).
- [Ber11] Bertin, E. ‘Automated Morphometry with SExtractor and PSFEx’. In: *Astronomical Data Analysis Software and Systems XX*. Ed. by I. N. Evans, A. Accomazzi, D. J. Mink and A. H. Rots. Vol. 442. Astronomical Society of the Pacific Conference Series. July 2011, p. 435 (pages 26, 60, 123, 126, 128, 143).
- [Bos+17] Bosch, J., Armstrong, R., Bickerton, S., Furusawa, H., Ikeda, H., Koike, M., Lupton, R., Mineo, S., Price, P., Takata, T., Tanaka, M., Yasuda, N., AlSayyad, Y., Becker, A. C., Coulton, W., Coupon, J., Garmilla, J., Huang, S., Krughoff, K. S., Lang, D., Leauthaud, A., Lim, K.-T., Lust, N. B., MacArthur, L. A., Mandelbaum, R., Miyatake, H., Miyazaki, S., Murata, R., More, S., Okura, Y., Owen, R., Swinbank, J. D., Strauss, M. A., Yamada, Y., Yamanoi, H., ‘The Hyper Suprime-Cam software pipeline’. In: *Publications of the Astronomical Society of Japan* 70.SP1 (Oct. 2017). S5 (page 24).
- [BPZ88] Bendinelli, O., Parmeggiani, G., Zavatti, F., ‘CCD star images : on the determination of Moffat’s PSF shape parameters.’ In: *Journal of Astrophysics and Astronomy* 9 (1988), pp. 17–24 (page 126).
- [Bra78] Bracewell, R. *The Fourier Transform and its Applications*. Second. Tokyo: McGraw-Hill Kogakusha, Ltd., 1978 (page 129).
- [BS01] Bartelmann, M., Schneider, P., ‘Weak gravitational lensing’. In: *Physics Reports* 340.4 (2001), pp. 291–472 (pages 13, 18).
- [BST14] Bolte, J., Sabach, S., Teboulle, M., ‘Proximal Alternating Linearized Minimization for Nonconvex and Nonsmooth Problems’. In: *Math. Program.* 146.1–2 (Aug. 2014), pp. 459–494 (page 88).
- [BT09] Beck, A., Teboulle, M., ‘A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems’. In: *SIAM J. Img. Sci.* 2.1 (Mar. 2009), pp. 183–202 (pages 62, 88).
- [BTM04] Basden, A., Tubbs, B., Mackay, C., ‘L3CCD’s: Low Readout Noise CCDs in Astronomy’. In: *Scientific Detectors for Astronomy*. Ed. by P. Amico, J. W. Beletic and J. E. Beletic. Dordrecht: Springer Netherlands, 2004, pp. 599–602 (page 53).
- [BW64] Born, M., Wolf, E., *Principles of Optics: Electromagnetic Theory of Propagation, Interference, and Diffraction of Light*. 6th. Oxford: Pergamon Press, 1964 (page 36).
- [Chu97] Chung, F. R. K. *Spectral Graph Theory*. American Mathematical Society, 1997 (pages 86, 133).
- [col16] collaboration, T. C. ‘The Canada-France Imaging Survey (CFIS)’. CFHT Large Proposal 2017-2019. 2016 (page 56).
- [Con13] Condat, L. ‘A primal–dual splitting method for convex optimization involving Lipschitzian, proximable and linear composite terms’. In: *Journal of Optimization Theory and Applications* 158.2 (2013), pp. 460–479 (pages 62, 88, 203, 204).

- [Cou+18] Coulton, W. R., Armstrong, R., Smith, K. M., Lupton, R. H., Spergel, D. N., ‘Exploring the Brighter-fatter Effect with the Hyper Suprime-Cam’. In: *ApJ* 155.6 (2018), p. 258 (pages 24, 53).
- [Cro+13] Cropper, M., Hoekstra, H., Kitching, T., Massey, R., Amiaux, J., Miller, L., Mellier, Y., Rhodes, J., Rowe, B., Pires, S., ‘Defining a weak lensing experiment in space’. In: *MNRAS* 431.4 (2013), pp. 3103–3126 (pages 31, 71, 74, 82, 139).
- [Cro+16] Cropper, M., Pottinger, S., Niemi, S., Azzollini, R., Denniston, J., Szafraniec, M., Awan, S., Mellier, Y., Berthe, M., Martignac, J., Cara, C., Giorgio, A.-M. D., Sciortino, A., Bozzo, E., Genolet, L., Cole, R., Philippon, A., Hailey, M., Hunt, T., Swindells, I., Holland, A., Gow, J., Murray, N., Hall, D., Skottfelt, J., Amiaux, J., Laureijs, R., Racca, G., Salvignol, J.-C., Short, A., Alvarez, J. L., Kitching, T., Hoekstra, H., Massey, R., Israel, H., ‘VIS: the visible imager for Euclid’. In: *Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave*. Ed. by H. A. MacEwen, G. G. Fazio, M. Lystrup, N. Batalha, N. Siegler and E. C. Tong. Vol. 9904. International Society for Optics and Photonics. SPIE, 2016, pp. 269–284 (pages 31, 123, 140, 143).
- [Des+12] Desai, S., Armstrong, R., Mohr, J. J., Semler, D. R., Liu, J., Bertin, E., Allam, S. S., Barkhouse, W. A., Bazin, G., Buckley-Geer, E. J., Cooper, M. C., Hansen, S. M., High, F. W., Lin, H., Lin, Y.-T., Ngeow, C.-C., Rest, A., Song, J., Tucker, D., Zenteno, A., ‘THE BLANCO COSMOLOGY SURVEY: DATA ACQUISITION, PROCESSING, CALIBRATION, QUALITY DIAGNOSTICS, AND DATA RELEASE’. In: *The Astrophysical Journal* 757.1 (Sept. 2012), p. 83 (page 28).
- [DK01] Dougherty, G., Kawaf, Z., ‘The point spread function revisited: Image restoration using 2-D deconvolution’. In: *Radiography* 7 (Nov. 2001), pp. 255–262 (page 45).
- [DRR16] Davis, C. P., Rodriguez, J., Roodman, A., ‘Wavefront-based PSF estimation’. In: *Ground-based and Airborne Telescopes VI*. Ed. by H. J. Hall, R. Gilmozzi and H. K. Marshall. Vol. 9906. International Society for Optics and Photonics. SPIE, 2016, pp. 2156–2168 (page 64).
- [Ello2] Ellerbroek, B. L. ‘Efficient computation of minimum-variance wave-front reconstructors with sparse matrix techniques’. In: *J. Opt. Soc. Am. A* 19.9 (Sept. 2002), pp. 1803–1816 (page 57).
- [Erb+13] Erben, T., Hildebrandt, H., Miller, L., Waerbeke, L., Heymans, C., Hoekstra, H., Kitching, T. D., Mellier, Y., Benjamin, J., Blake, C., Bonnett, C., Cordes, O., Coupon, J., Fu, L., Gavazzi, R., Gillis, B., Grocutt, E., Gwyn, S. D. J., Holhjem, K., Hudson, M. J., Kilbinger, M., Kuijken, K., Milkeraitis, M., Rowe, B. T. P., Schrabback, T., Semboloni, E., Simon, P., Smit, M., Toader, O., Vafaei, S., Uitert, E., Velander, M., ‘CFHTLenS: the Canada–France–Hawaii Telescope Lensing Survey – imaging data and catalogue products’. In: *Monthly Notices of the Royal Astronomical Society* 433.3 (June 2013), pp. 2545–2563 (page 98).
- [Euc+22a] Euclid Collaboration, ‘Euclid preparation - I. The Euclid Wide Survey’. In: *A&A* 662 (2022), A112 (page 32).
- [Euc+22b] Euclid Collaboration, ‘Euclid preparation - XVIII. The NISP photometric system’. In: *A&A* 662 (2022), A92 (page 31).

- [Far+20] Farrens, S., Grigis, A., El Gueddari, L., Ramzi, Z., G.R. C., Starck, S., Sarthou, B., Cherkaoui, H., Ciuciu, P., Starck, J.-L., ‘PySAP: Python Sparse Data Analysis Package for multidisciplinary image processing’. In: *Astronomy and Computing* 32 (July 2020), p. 100402 (page 203).
- [Far+22a] Farrens, S. Guinot, A. Kilbinger, M. Liaudat, T. Baumont, L. Jimenez, X. Peel, A. Pujol, A. Schmitz, M. Starck, J.-L. Vitorelli, A. Z. ‘ShapePipe: A modular weak-lensing processing and analysis pipeline’. In: *A&A* 664 (2022), A141 (pages 24, 26, 27, 109, 207).
- [Far+22b] Farrens, S. Lacan, A. Guinot, A. Vitorelli, A. Z. ‘Deep transfer learning for blended source identification in galaxy survey data’. In: *A&A* 657 (2022), A98 (page 27).
- [Fen17] Fenech Conti, I. ‘Point Spread Function Modelling and Shear Calibration for Weak Lensing Surveys’. PhD thesis. University of Malta, Mar. 2017 (pages 53, 58).
- [Fie+93] Fienup, J. R., Marron, J. C., Schulz, T. J., Seldin, J. H., ‘Hubble Space Telescope characterized by using phase-retrieval algorithms’. In: *Appl. Opt.* 32.10 (Apr. 1993), pp. 1747–1767 (page 157).
- [FMN16] Fefferman, C., Mitter, S., Narayanan, H., ‘Testing the manifold hypothesis’. In: *Journal of the American Mathematical Society* 29.4 (Feb. 2016), pp. 983–1049 (page 126).
- [Fu +08] Fu, L. Semboloni, E. Hoekstra, H. Kilbinger, M. van Waerbeke, L. Tereno, I. Mellier, Y. Heymans, C. Coupon, J. Benabed, K. Benjamin, J. Bertin, E. Doré, O. Hudson, M. J. Ilbert, O. Maoli, R. Marmo, C. McCracken, H. J. Ménard, B. ‘Very weak lensing in the CFHTLS wide: cosmology from cosmic shear in the linear regime’. In: *A&A* 479.1 (2008), pp. 9–25 (page 59).
- [Gai+16] Gaia Collaboration, ‘The Gaia mission’. In: *A&A* 595 (2016), A1 (page 50).
- [Gat+21] Gatti, M. ‘Dark energy survey year 3 results: weak lensing shape catalogue’. In: *Monthly Notices of the Royal Astronomical Society* 504.3 (Apr. 2021), pp. 4312–4336 (pages 29, 30, 62, 63, 80, 117).
- [GCM13] Gentile, M., Courbin, F., Meylan, G., ‘Interpolating point spread function anisotropy’. In: *A&A* 549 (2013), A1 (pages 59, 89).
- [Gib+21] Giblin, Benjamin, Heymans, Catherine, Asgari, Marika, Hildebrandt, Hendrik, Hoekstra, Henk, Joachimi, Benjamin, Kannawadi, Arun, Kuijken, Konrad, Lin, Chieh-An, Miller, Lance, Tröster, Tilman, van den Busch, Jan Luca, Wright, Angus H. Bilicki, Maciej, Blake, Chris, de Jong, Jelte, Dvornik, Andrej, Erben, Thomas, Getman, Fedor, Napolitano, Nicola R. Schneider, Peter, Shan, HuanYuan, Valentijn, Edwin, ‘KiDS-1000 catalogue: Weak gravitational lensing shear measurements’. In: *A&A* 645 (2021), A105 (pages 29, 61, 117).
- [Goo05] Goodman, J. W. ‘Introduction to Fourier optics’. In: *Introduction to Fourier optics, 3rd ed., by JW Goodman. Englewood, CO: Roberts & Co. Publishers, 2005* 1 (2005) (pages 36, 40, 42, 128).
- [Gro+06] Gross, H. H., Zügge, H., Peschka, M., Blechinger, F., ‘Handbook of Optical Systems: Volume 3: Aberration Theory and Correction of Optical Systems’. In: 2006 (page 49).
- [Gro05] Gross, H. *Handbook of Optical Systems, Volume 1, Fundamentals of Technical Optics*. Wiley-VCH, 2005 (pages 36, 39).

- [Gui+22] Guinot, A., Kilbinger, M., Farrens, S., Peel, A., Pujol, A., Schmitz, M., Starck, J.-L., Erben, T., Gavazzi, R., Gwyn, S., Hudson, M. J., Hildebrandt, H., **Liaudat, T.**, Miller, L., Spitzer, I., Van Waerbeke, L., Cuillandre, J.-C., Fabbro, S., McConnachie, A., ‘ShapePipe: a new shape measurement pipeline and weak-lensing application to UNIONS/CFIS data’. In: *arXiv:2204.04798* (2022) (pages 24, 25, 29, 60, 80, 109, 110, 112, 116, 207).
- [Gui20] Guinot, A. ‘Weak lensing analysis of the Canada-France Imaging Survey : from pixels to cosmology, preparation for the Euclid mission’. 2020UNIP7106. PhD thesis. 2020 (pages 14, 24, 112).
- [Guy+15] Guyonnet, A., Astier, P., Antilogus, P., Regnault, N., Doherty, P., ‘Evidence for self-interaction of charge distribution in charge-coupled devices’. In: *Astronomy & Astrophysics* 575 (Feb. 2015), A41 (page 53).
- [GV96] Golub, G. H., Van Loan, C. F., *Matrix Computations (3rd Ed.)* USA: Johns Hopkins University Press, 1996 (page 204).
- [Gwy08] Gwyn, S. D. J. ‘MegaPipe: The MegaCam Image Stacking Pipeline at the Canadian Astronomical Data Centre’. In: *Publications of the Astronomical Society of the Pacific* 120.864 (Feb. 2008), pp. 212–223 (page 24).
- [Gwy19] Gwyn, S. ‘MegaPipe 2.0: 10000 Square Degrees of CFHT MegaCam Imaging’. In: *Astronomical Data Analysis Software and Systems XXVII*. Ed. by P. J. Teuben, M. W. Pound, B. A. Thomas and E. M. Warner. Vol. 523. Astronomical Society of the Pacific Conference Series. Oct. 2019, p. 649 (page 24).
- [HA17] Hoffmann, S. L., Anderson, J., *A Study of PSF Models for ACS/WFC*. Instrument Science Report ACS 2017-8. Oct. 2017 (pages 58, 123).
- [Her+18] Herbel, J., Kacprzak, T., Amara, A., Refregier, A., Lucchi, A., ‘Fast point spread function modeling with deep learning’. In: *Journal of Cosmology and Astroparticle Physics* 2018.07 (July 2018), pp. 054–054 (page 65).
- [Hey+12a] Heymans, C., Rowe, B., Hoekstra, H., Miller, L., Erben, T., Kitching, T., Van Waerbeke, L., ‘The impact of high spatial frequency atmospheric distortions on weak-lensing measurements’. In: *Monthly Notices of the Royal Astronomical Society* 421.1 (2012), pp. 381–389 (pages 57, 72, 111).
- [Hey+12b] Heymans, C., Waerbeke, L., Miller, L., Erben, T., Hildebrandt, H., Hoekstra, H., Kitching, T. D., Mellier, Y., Simon, P., Bonnett, C., Coupon, J., Fu, L., Harnois-Déraps, J., Hudson, M. J., Kilbinger, M., Kuijken, K., Rowe, B., Schrabback, T., Semboloni, E., Uitert, E., Vafaei, S., Velander, M., ‘CFHTLenS: the Canada–France–Hawaii Telescope Lensing Survey’. In: *Monthly Notices of the Royal Astronomical Society* 427.1 (Nov. 2012), pp. 146–166 (page 117).
- [Hig+07] High, F. W., Rhodes, J., Massey, R., Ellis, R., ‘Pixelation Effects in Weak Lensing’. In: *Publications of the Astronomical Society of the Pacific* 119.861 (2007), pp. 1295–1307 (page 52).
- [HM17] Huff, E., Mandelbaum, R., ‘Metacalibration: Direct Self-Calibration of Biases in Shear Measurement’. In: *arXiv preprint arXiv:1702.02600* (2017) (pages 30, 109).
- [Hog22] Hogg, D. W. *Magnitudes, distance moduli, bolometric corrections, and so much more*. 2022 (pages 24, 47).

- [HS03] Hirata, C., Seljak, U., ‘Shear calibration biases in weak-lensing surveys’. In: *MNRAS* 343.2 (2003), pp. 459–480 (pages 71, 94, 112).
- [Jar+03] Jarvis, M., Bernstein, G. M., Fischer, P., Smith, D., Jain, B., Tyson, J. A., Wittman, D., ‘Weak-Lensing Results from the 75 Square Degree Cerro Tololo Inter-American Observatory Survey’. In: *The Astronomical Journal* 125.3 (Mar. 2003), pp. 1014–1032 (page 21).
- [Jar+16] Jarvis, M., Sheldon, E., Zuntz, J., Kacprzak, T., Bridle, S., Amara, A., Armstrong, R., Becker, M., Bernstein, G., Bonnett, C., ‘The DES Science Verification weak lensing shear catalogues’. In: *MNRAS* 460.2 (2016), pp. 2245–2281 (pages 24, 29, 70, 72, 73).
- [Jar+20] Jarvis, M., Bernstein, G. M., Amon, A., Davis, C., Léget, P. F., Bechtol, K., Harrison, I., Gatti, M., Roodman, A., Chang, C., Chen, R., Choi, A., Desai, S., Drlica-Wagner, A., Gruen, D., Gruendl, R. A., Hernandez, A., MacCrann, N., Meyers, J., Navarro-Alsina, A., Pandey, S., Plazas, A. A., Secco, L. F., Sheldon, E., Troxel, M. A., Vorperian, S., Wei, K., Zuntz, J., Abbott, T. M. C., Aguena, M., Allam, S., Avila, S., Bhargava, S., Bridle, S. L., Brooks, D., Carnero Rosell, A., Carrasco Kind, M., Carretero, J., Costanzi, M., Costa, L. N., De Vicente, J., Diehl, H. T., Doel, P., Everett, S., Flaughner, B., Fosalba, P., Frieman, J., García-Bellido, J., Gaztanaga, E., Gerdes, D. W., Gutierrez, G., Hinton, S. R., Hollowood, D. L., Honscheid, K., James, D. J., Kent, S., Kuehn, K., Kuropatkin, N., Lahav, O., Maia, M. A. G., March, M., Marshall, J. L., Melchior, P., Menanteau, F., Miquel, R., Ogando, R. L. C., Paz-Chinchón, F., Rykoff, E. S., Sanchez, E., Scarpine, V., Schubnell, M., Serrano, S., Sevilla-Noarbe, I., Smith, M., Suchyta, E., Swanson, M. E. C., Tarle, G., Varga, T. N., Walker, A. R., Wester, W., Wilkinson, R. D., ‘Dark Energy Survey year 3 results: point spread function modelling’. In: *Monthly Notices of the Royal Astronomical Society* 501.1 (Dec. 2020), pp. 1282–1299 (pages 54, 62, 63, 73, 80, 123).
- [JBL18] Joyce, K. T., Bardsley, J. M., Luttmann, A., ‘Point Spread Function Estimation in X-Ray Imaging with Partially Collapsed Gibbs Sampling’. In: *SIAM Journal on Scientific Computing* 40.3 (2018), B766–B787 (page 45).
- [Jee+07] Jee, M. J., Blakeslee, J. P., Sirianni, M., Martel, A. R., White, R. L., Ford, H. C., ‘Principal Component Analysis of the Time- and Position-dependent Point-Spread Function of the Advanced Camera for Surveys’. In: *Publications of the Astronomical Society of the Pacific* 119.862 (Dec. 2007), pp. 1403–1419 (pages 60, 126).
- [Jef+20] Jeffrey, N., Lanusse, F., Lahav, O., Starck, J.-L., ‘Deep learning dark matter map reconstructions from DES SV weak lensing data’. In: *Monthly Notices of the Royal Astronomical Society* 492.4 (Jan. 2020), pp. 5023–5029 (page 20).
- [Jef+21] Jeffrey, N., ‘Dark Energy Survey Year 3 results: Curved-sky weak lensing mass map reconstruction’. In: *Monthly Notices of the Royal Astronomical Society* 505.3 (May 2021), pp. 4626–4645 (page 20).
- [JF14] Jurling, A. S., Fienup, J. R., ‘Applications of algorithmic differentiation to phase retrieval algorithms’. In: *Journal of the Optical Society of America A* 31.7 (2014), p. 1348 (page 128).

- [Jia+20a] Jia, P., Li, X., Li, Z., Wang, W., Cai, D., ‘Point spread function modelling for wide-field small-aperture telescopes with a denoising autoencoder’. In: *Monthly Notices of the Royal Astronomical Society* 493.1 (Feb. 2020), pp. 651–660 (pages 65, 152).
- [Jia+20b] Jia, P., Wu, X., Li, Z., Li, B., Wang, W., Liu, Q., Popowicz, A., *Modelling the Point Spread Function of Wide Field Small Aperture Telescopes With Deep Neural Networks – Applications in Point Spread Function Estimation*. 2020 (page 152).
- [Jia+20c] Jia, P., Wu, X., Yi, H., Cai, B., Cai, D., ‘PSF-NET: A Nonparametric Point-spread Function Model for Ground-based Optical Telescopes’. In: *AJ* 159.4, 183 (Apr. 2020), p. 183 (pages 65, 152).
- [JJ04] Jarvis, M., Jain, B., *Principal Component Analysis of PSF Variation in Weak Lensing Surveys*. 2004 (page 59).
- [JT11] Jee, M. J., Tyson, J. A., ‘Toward Precision LSST Weak-Lensing Measurement. I. Impacts of Atmospheric Turbulence and Optical Aberration’. In: *Publications of the Astronomical Society of the Pacific* 123.903 (May 2011), pp. 596–614 (pages 52, 55–57, 60).
- [Kac+14] Kacprzak, T., Bridle, S., Rowe, B., Voigt, L., Zuntz, J., Hirsch, M., MacCrann, N., ‘Sersic galaxy models in weak lensing shape measurement: model bias, noise bias and their interaction’. In: *Monthly Notices of the Royal Astronomical Society* 441.3 (May 2014), pp. 2528–2538 (page 29).
- [Kai00] Kaiser, N. ‘A New Shear Estimator for Weak-Lensing Observations’. In: *The Astrophysical Journal* 537.2 (July 2000), pp. 555–577 (page 28).
- [Kai92] Kaiser, N. ‘Weak Gravitational Lensing of Distant Galaxies’. In: *The Astrophysical Journal* 388 (Apr. 1992), p. 272 (page 23).
- [Kár30] Kármán, T. v. ‘Mechanische Aenlichkeit und Turbulenz’. In: *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* 1930 (1930), pp. 58–76 (pages 55, 57).
- [KB14] Kingma, D. P., Ba, J., ‘Adam: A Method for Stochastic Optimization’. In: (Dec. 2014) (page 136).
- [KB95] Krist, J. E., Burrows, C. J., ‘Phase-retrieval analysis of pre- and post-repair Hubble Space Telescope images’. In: *Appl. Opt.* 34.22 (Aug. 1995), pp. 4951–4964 (pages 58, 123).
- [KC17] Kuntzer, T., Courbin, F., ‘Detecting unresolved binary stars in Euclid VIS images’. In: *A&A* 606 (2017), A119 (page 66).
- [KH10] Krause, E. Hirata, C. M. ‘Weak lensing power spectra for precision cosmology - Multiple-deflection, reduced shear, and lensing bias corrections’. In: *A&A* 523 (2010), A28 (page 14).
- [KHS11] Krist, J. E., Hook, R. N., Stoehr, F., ‘20 years of Hubble Space Telescope optical modeling using Tiny Tim’. In: *Optical Modeling and Performance Predictions V*. Ed. by M. A. Kahan. Vol. 8127. International Society for Optics and Photonics. SPIE, 2011, pp. 166–181 (pages 51, 52, 58, 123, 128).
- [Kil+17] Kilbinger, M., Heymans, C., Asgari, M., Joudaki, S., Schneider, P., Simon, P., Van Waerbeke, L., Harnois-Déraps, J., Hildebrandt, H., Köhlinger, F., Kuijken, K., Viola, M., ‘Precision calculations of the cosmic shear power spectrum projection’. In: *Monthly Notices of the Royal Astronomical Society* 472.2 (Aug. 2017), pp. 2126–2141 (page 23).

- [Kil15] Kilbinger, M. ‘Cosmology with cosmic shear observations: a review’. In: *Reports on Progress in Physics* 78.8 (2015), p. 086901 (pages 11, 13, 16, 20).
- [Kit+08] Kitching, T. D., Miller, L., Heymans, C. E., Waerbeke, L., Heavens, A. F., ‘Bayesian galaxy shape measurement for weak lensing surveys - II. Application to simulations’. In: *Monthly Notices of the Royal Astronomical Society* 390.1 (Oct. 2008), pp. 149–167 (page 61).
- [KLB20] Kervazo, C., **Liaudat, T.**, Bobin, J., ‘Faster and better sparse blind source separation through mini-batch optimization’. In: *Digital Signal Processing* 106 (2020), p. 102827 (page 207).
- [Kol91] Kolmogorov, A. N. ‘The Local Structure of Turbulence in Incompressible Viscous Fluid for Very Large Reynolds Numbers’. In: *Proceedings of the Royal Society of London Series A* 434.1890 (July 1991), pp. 9–13 (pages 55, 57).
- [Kri03] Krist, J. ‘ACS WFC & HRC field- dependent PSF variations due to optical and charge diffusion effects’. In: *Applied Categorical Structures* (2003), p. 6 (page 53).
- [Kri93] Krist, J. ‘Tiny Tim : an HST PSF Simulator’. In: *Astronomical Data Analysis Software and Systems II*. Ed. by R. J. Hanisch, R. J. V. Brissenden and J. Barnes. Vol. 52. Astronomical Society of the Pacific Conference Series. 1993, p. 536 (pages 58, 123).
- [Kri95] Krist, J. E. ‘WFPC2 Ghosts, Scatter and PSF Field Dependence’. In: *Calibrating Hubble Space Telescope. Post Servicing Mission. Proceedings of a Workshop held at the Space Telescope Science Institute, in Baltimore, Maryland*. Ed. by A. Koratkar and C. Leitherer. May 1995 (page 49).
- [KS93] Kaiser, N., Squires, G., ‘Mapping the Dark Matter with Weak Gravitational Lensing’. In: *The Astrophysical Journal* 404 (Feb. 1993), p. 441 (page 20).
- [KSB95] Kaiser, N., Squires, G., Broadhurst, T., ‘A Method for Weak Lensing Observations’. In: *ApJ* 449 (1995), p. 460 (pages 28, 59).
- [KTC16] Kuntzer, T., Tewes, M., Courbin, F., ‘Stellar classification from single-band imaging using machine learning’. In: *A&A* 591 (2016), A54 (page 140).
- [Kui+15] Kuijken, K., Heymans, C., Hildebrandt, H., Nakajima, R., Erben, T., Jong, J. T., Viola, M., Choi, A., Hoekstra, H., Miller, L., ‘Gravitational lensing analysis of the Kilo-Degree Survey’. In: *MNRAS* 454.4 (2015), pp. 3500–3532 (pages 29, 61).
- [Lan+16] Lanusse, F., Starck, J.-L., Leonard, A., Pires, S. ‘High resolution weak lensing mass mapping combining shear and flexion’. In: *A&A* 591 (2016), A2 (page 20).
- [Lau+11] Laureijs, R., Amiaux, J., Arduini, S., Augueres, J.-L., Brinchmann, J., Cole, R., Cropper, M., Dabin, C., Duvet, L., Ealet, A., ‘Euclid definition study report’. In: *ArXiv e-prints* (2011) (pages 31, 140).
- [Lea+07] Leauthaud, A., Massey, R., Kneib, J.-P., Rhodes, J., Johnston, D. E., Capak, P., Heymans, C., Ellis, R. S., Koekemoer, A. M., Fevre, O. L., Mellier, Y., Refregier, A., Robin, A. C., Scoville, N., Tasca, L., Taylor, J. E., Waerbeke, L. V., ‘Weak Gravitational Lensing with COSMOS: Galaxy Selection and Shape Measurements’. In: *The Astrophysical Journal Supplement Series* 172.1 (Sept. 2007), pp. 219–238 (page 58).

- [Li+16] Li, B.-S., Li, G.-L., Cheng, J., Peterson, J., Cui, W., ‘The point spread function reconstruction by using Moffatlets — I’. In: *Research in Astronomy and Astrophysics* 16.9 (Sept. 2016), p. 007 (pages 63, 126).
- [Li+22] Li, X., Miyatake, H., Luo, W., More, S., Oguri, M., Hamana, T., Mandelbaum, R., Shirasaki, M., Takada, M., Armstrong, R., Kannawadi, A., Takita, S., Miyazaki, S., Nishizawa, A. J., Plazas Malagon, A. A., Strauss, M. A., Tanaka, M., Yoshida, N., ‘The three-year shear catalog of the Subaru Hyper Suprime-Cam SSP Survey’. In: *Publications of the Astronomical Society of Japan* 74.2 (Mar. 2022), pp. 421–459 (page 117).
- [Lia+21a] **Liaudat, T.**, Bonnin, J., Starck, J.-L., Schmitz, M. A., Guinot, A., Kilbinger, M., Gwyn, S. D. J., ‘Multi-CCD modelling of the point spread function’. In: *A&A* 646 (2021), A27 (pages 79, 80, 123, 126, 128, 130, 133, 138, 142, 207).
- [Lia+21b] **Liaudat, T.**, Starck, J.-L., Kilbinger, M., Frugier, P.-A., ‘Rethinking the modeling of the instrumental response of telescopes with a differentiable optical model’. In: *NeurIPS 2021 Machine Learning for Physical sciences workshop*. Nov. 2021 (pages 122, 207).
- [Lia+22] **Liaudat, T.**, Starck, J.-L., Kilbinger, M., Frugier, P.-A., ‘Rethinking data-driven point spread function modeling with a differentiable optical model’. In: *arXiv:2203.04908* (2022) (pages 45, 121, 122, 207).
- [Lim53] Limber, D. N. ‘The Analysis of Counts of the Extragalactic Nebulae in Terms of a Fluctuating Density Field.’ In: *The Astrophysical Journal* 117 (Jan. 1953), p. 134 (page 23).
- [Lin+20] Lin, C.-H., Tan, B., Mandelbaum, R., Hirata, C. M., ‘The impact of light polarization effects on weak lensing systematics’. In: *Monthly Notices of the Royal Astronomical Society* 496.1 (May 2020), pp. 532–539 (page 50).
- [Liu+19] Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., Han, J., *On the Variance of the Adaptive Learning Rate and Beyond*. Aug. 2019 (page 136).
- [LLS18] Liang, J., Luo, T., Schönlieb, C.-B., ‘Improving "Fast Iterative Shrinkage-Thresholding Algorithm": Faster, Smarter and Greedier’. In: (Nov. 2018) (pages 88, 204).
- [LSK21] **Liaudat, T.**, Starck, J.-L., Kilbinger, M., ‘Semi-Parametric Wavefront Modelling for the Point Spread Function’. In: *52ème Journées de Statistique de la Société Française de Statistique (SFdS)*. Nice, France, June 2021 (page 207).
- [LSS+09] LSST Science Collaboration, Abell, P. A., Allison, J., Anderson, S. F., Andrew, J. R., Angel, J. R. P., Armus, L., Arnett, D., Asztalos, S. J., Axelrod, T. S., al. ‘LSST Science Book, Version 2.0’. In: *arXiv e-prints* (Dec. 2009) (page 123).
- [Lu+17] Lu, T., Zhang, J., Dong, F., Li, Y., Liu, D., Fu, L., Li, G., Fan, Z., ‘Testing PSF Interpolation in Weak Lensing with Real Data’. In: *The Astronomical Journal* 153.4 (Apr. 2017), p. 197 (page 64).
- [Lup+01] Lupton, R., Gunn, J. E., Ivezić, Z., Knapp, G. R., Kent, S., Yasuda, N., ‘The SDSS Imaging Pipelines’. In: (2001) (page 59).

- [Mac+21] MacCrann, N., Becker, M. R., McCullough, J., Amon, A., Gruen, D., Jarvis, M., Choi, A., Troxel, M. A., Sheldon, E., Yanny, B., Herner, K., Dodelson, S., Zuntz, J., Eckert, K., Rollins, R. P., Varga, T. N., Bernstein, G. M., Gruendl, R. A., Harrison, I., Hartley, W. G., Sevilla-Noarbe, I., Pieres, A., Bridle, S. L., Myles, J., Alarcon, A., Everett, S., Sanchez, C., Huff, E. M., Tarsitano, F., Gatti, M., Secco, L. F., Abbott, T. M. C., Aguena, M., Allam, S., Annis, J., Bacon, D., Bertin, E., Brooks, D., Burke, D. L., Carnero Rosell, A., Carrasco Kind, M., Carretero, J., Costanzi, M., Crocce, M., Pereira, M. E. S., De Vicente, J., Desai, S., Diehl, H. T., Dietrich, J. P., Doel, P., Eifler, T. F., Ferrero, I., Ferte, A., Flaugher, B., Fosalba, P., Frieman, J., Garcia-Bellido, J., Gaztanaga, E., Gerdes, D. W., Giannantonio, T., Gschwend, J., Gutierrez, G., Hinton, S. R., Hollowood, D. L., Honscheid, K., James, D. J., Lahav, O., Lima, M., Maia, M. A. G., March, M., Marshall, J. L., Martini, P., Melchior, P., Menanteau, F., Miquel, R., Mohr, J. J., Morgan, R., Muir, J., Ogando, R. L. C., Palmese, A., Paz-Chinchon, F., Plazas, A. A., Rodriguez-Monroy, M., Roodman, A., Samuroff, S., Sanchez, E., Scarpine, V., Serrano, S., Smith, M., Soares-Santos, M., Suchyta, E., Swanson, M. E. C., Tarle, G., Thomas, D., To, C., Wilkinson, R. D., Collaboration), (‘Dark Energy Survey Y3 results: blending shear and redshift biases in image simulations’. In: *Monthly Notices of the Royal Astronomical Society* 509.3 (Oct. 2021), pp. 3371–3394 (page 29).
- [Man+05] Mandelbaum, R., Hirata, C. M., Seljak, U., Guzik, J., Padmanabhan, N., Blake, C., Blanton, M. R., Lupton, R., Brinkmann, J., ‘Systematic errors in weak lensing: application to SDSS galaxy-galaxy weak lensing’. In: *MNRAS* 361.4 (2005), pp. 1287–1322 (pages 71, 94, 112).
- [Man+17] Mandelbaum, R., Miyatake, H., Hamana, T., Oguri, M., Simet, M., Armstrong, R., Bosch, J., Murata, R., Lanusse, F., Leauthaud, A., Coupon, J., More, S., Takada, M., Miyazaki, S., Speagle, J. S., Shirasaki, M., Sifón, C., Huang, S., Nishizawa, A. J., Medezinski, E., Okura, Y., Okabe, N., Czakon, N., Takahashi, R., Coulton, W. R., Hikage, C., Komiyama, Y., Lupton, R. H., Strauss, M. A., Tanaka, M., Utsumi, Y., ‘The first-year shear catalog of the Subaru Hyper Suprime-Cam Subaru Strategic Program Survey’. In: *Publications of the Astronomical Society of Japan* 70.SP1 (Dec. 2017). S25 (pages 24, 60).
- [Man+18] Mandelbaum, R., Lanusse, F., Leauthaud, A., Armstrong, R., Simet, M., Miyatake, H., Meyers, J. E., Bosch, J., Murata, R., Miyazaki, S., Tanaka, M., ‘Weak lensing shear calibration with simulations of the HSC survey’. In: *Monthly Notices of the Royal Astronomical Society* 481.3 (Sept. 2018), pp. 3170–3195 (pages 30, 46).
- [Man18] Mandelbaum, R. ‘Weak Lensing for Precision Cosmology’. In: *Annual Review of Astronomy and Astrophysics* 56 (2018), pp. 393–433 (pages 13, 24).
- [Mar+15] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Jia, Y., Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay

- Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, Xiaoqiang Zheng, *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015 (pages 137, 152).
- [Mas+09] Massey, R., Stoughton, C., Leauthaud, A., Rhodes, J., Koekemoer, A., Ellis, R., Shaghoulain, E., ‘Pixel-based correction for Charge Transfer Inefficiency in the Hubble Space Telescope Advanced Camera for Surveys’. In: *Monthly Notices of the Royal Astronomical Society* 401.1 (Dec. 2009), pp. 371–384 (page 53).
- [Mas+12] Massey, R., Hoekstra, H., Kitching, T., Rhodes, J., Cropper, M., Amiaux, J., Harvey, D., Mellier, Y., Meneghetti, M., Miller, L., ‘Origins of weak lensing systematics, and requirements on future instrumentation (or knowledge of instrumentation)’. In: *MNRAS* 429.1 (2012), pp. 661–678 (pages 71, 73, 74, 139).
- [Mas+14] Massey, R., Schrabback, T., Cordes, O., Marggraf, O., Israel, H., Miller, L., Hall, D., Cropper, M., Prod’homme, T., Matias Niemi, S., ‘An improved model of charge transfer inefficiency and correction algorithm for the Hubble Space Telescope’. In: *Monthly Notices of the Royal Astronomical Society* 439.1 (Feb. 2014), pp. 887–907 (pages 24, 53).
- [MB15] Meyers, J. E., Burchat, P. R., ‘IMPACT OF ATMOSPHERIC CHROMATIC EFFECTS ON WEAK LENSING MEASUREMENTS’. In: *The Astrophysical Journal* 807.2 (July 2015), p. 182 (page 56).
- [MBJ06] Mugnier, L. M., Blanc, A., Jérôme Idier, ‘Phase Diversity: A Technique for Wave-Front Sensing and for Diffraction-Limited Imaging’. In: ed. by P. Hawkes. Vol. 141. *Advances in Imaging and Electron Physics*. Elsevier, 2006, pp. 1–76 (page 157).
- [MC04] Magnier, E. A., Cuillandre, J.-C., ‘The Elixir System: Data Characterization and Calibration at the Canada-France-Hawaii Telescope’. In: *Publications of the Astronomical Society of the Pacific* 116.819 (May 2004), pp. 449–464 (page 24).
- [McC97] McCann, R. J. ‘A Convexity Principle for Interacting Gases’. In: *Advances in Mathematics* 128.1 (1997), pp. 153–179 (page 127).
- [Mel+18] Melchior, P., Moolekamp, F., Jerdee, M., Armstrong, R., Sun, A.-L., Bosch, J., Lupton, R., ‘Scarlet: Source separation in multi-band images by Constrained Matrix Factorization’. In: *Astronomy and Computing* 24 (2018), pp. 129–142 (page 26).
- [Mel+21] Melchior, P., Joseph, R., Sanchez, J., MacCrann, N., Gruen, D., ‘The challenge of blending in large sky surveys’. In: *Nature Reviews Physics* 3.10 (Aug. 2021), pp. 712–718 (page 27).
- [Mil+07] Miller, L., Kitching, T. D., Heymans, C., Heavens, A. F., Van Waerbeke, L., ‘Bayesian galaxy shape measurement for weak lensing surveys – I. Methodology and a fast-fitting algorithm’. In: *Monthly Notices of the Royal Astronomical Society* 382.1 (Oct. 2007), pp. 315–324 (page 61).
- [Mil+13] Miller, L., Heymans, C., Kitching, T., Van Waerbeke, L., Erben, T., Hildebrandt, H., Hoekstra, H., Mellier, Y., Rowe, B., Coupon, J., ‘Bayesian galaxy shape measurement for weak lensing surveys–III. Application to the Canada–France–Hawaii Telescope Lensing Survey’. In: *MNRAS* 429.4 (2013), pp. 2858–2880 (pages 29, 61, 80, 109, 123).

- [Mof69] Moffat, A. F. J. ‘A Theoretical Investigation of Focal Stellar Images in the Photographic Emulsion and Application to Photographic Photometry’. In: *A&A* 3 (Dec. 1969), p. 455 (page 54).
- [Moh+12] Mohr, J. J., Armstrong, R., Bertin, E., Daues, G., Desai, S., Gower, M., Gruendl, R., Hanlon, W., Kuropatkin, N., Lin, H., Marriner, J., Petracic, D., Sevilla, I., Swanson, M., Tomashek, T., Tucker, D., Yanny, B., ‘The Dark Energy Survey data processing and calibration system’. In: *SPIE Proceedings*. Ed. by N. M. Radziwill and G. Chiozzi. SPIE, Sept. 2012 (page 28).
- [MR05] Massey, R., Refregier, A., ‘Polar shapelets’. In: *Monthly Notices of the Royal Astronomical Society* 363.1 (Oct. 2005), pp. 197–210 (page 63).
- [MV12] Melchior, P., Viola, M., ‘Means of confusion: how pixel noise affects shear estimates for weak gravitational lensing’. In: *Monthly Notices of the Royal Astronomical Society* 424.4 (Aug. 2012), pp. 2757–2769 (page 29).
- [Ngo+16] Ngolè, F., Starck, J.-L., Okumura, K., Amiaux, J., Hudelot, P., ‘Constraint matrix factorization for space variant PSFs field restoration’. In: *Inverse Problems* 32.12 (2016), p. 124001 (pages 62, 82, 86, 126, 128, 130, 133, 137, 142, 143, 148, 205).
- [Ngo16] Ngolè Mboula, Fred Maurice, ‘Méthodes et algorithmes avancés pour l’imagerie astronomique de haute précision’. 2016SACLS343. PhD thesis. 2016 (page 37).
- [Nie+15] Niemi, S.-M., Cropper, M., Szafraniec, M., Kitching, T., ‘Measuring a charge-coupled device point spread function’. In: *Experimental Astronomy* 39.2 (June 2015), pp. 207–231 (page 53).
- [Nie+21] Nie, L., Li, G., Peterson, J. R., Wei, C., ‘The point spread function reconstruction – II. The smooth PCA’. In: *Monthly Notices of the Royal Astronomical Society* 503.3 (Mar. 2021), pp. 4436–4445 (page 63).
- [Nin+07] Nino, D., Makidon, R., Lallo, M., Sahu, K., Sirianni, M., Casertano, S., ‘HST Focus Variations with Temperatures’. In: *Instrument Science Report ACS 2008-03* (Dec. 2007) (page 50).
- [Nol76] Noll, R. J. ‘Zernike polynomials and atmospheric turbulence*’. In: *J. Opt. Soc. Am.* 66.3 (1976), pp. 207–211 (pages 44, 64, 131).
- [NS17] Ngolè, F., Starck, J.-L., ‘Point Spread Function Field Learning Based on Optimal Transport Distances’. In: *SIAM Journal on Imaging Sciences* 10.3 (2017), pp. 1549–1578 (pages 64, 89).
- [Nyq28] Nyquist, H. ‘Thermal Agitation of Electric Charge in Conductors’. In: *Phys. Rev.* 32 (1 July 1928), pp. 110–113 (page 53).
- [Pau+08] Paulin-Henriksson, S., Amara, A., Voigt, L., Refregier, A., Bridle, S., ‘Point spread function calibration requirements for dark energy from cosmic shear’. In: *A&A* 484.1 (2008), pp. 67–77 (pages 73, 74).
- [PB14] Parikh, N., Boyd, S., ‘Proximal Algorithms’. In: *Found. Trends Optim.* 1.3 (Jan. 2014), pp. 127–239 (pages 88, 203).
- [PBS14] Plazas, A., Bernstein, G., Sheldon, E., ‘On-Sky Measurements of the Transverse Electric Fields’ Effects in the Dark Energy Camera CCDs’. In: *Publications of the Astronomical Society of the Pacific* (July 2014) (page 54).

- [PC19] Peyré, G., Cuturi, M., ‘Computational Optimal Transport: With Applications to Data Science’. In: *Foundations and Trends® in Machine Learning* 11.5-6 (2019), pp. 355–607 (pages 64, 89).
- [Pee80] Peebles, P. J. E. *The large-scale structure of the universe*. 1980 (page 22).
- [Pee93] Peebles, P. J. E. *Principles of Physical Cosmology*. 1993 (page 22).
- [Per+99] Perlmutter, S., Aldering, G., Goldhaber, G., Knop, R. A., Nugent, P., Castro, P. G., Deustua, S., Fabbro, S., Goobar, A., Groom, D. E., Hook, I. M., Kim, A. G., Kim, M. Y., Lee, J. C., Nunes, N. J., Pain, R., Pennypacker, C. R., Quimby, R., Lidman, C., Ellis, R. S., Irwin, M., McMahon, R. G., Ruiz-Lapuente, P., Walton, N., Schaefer, B., Boyle, B. J., Filippenko, A. V., Matheson, T., Fruchter, A. S., Panagia, N., Newberg, H. J. M., Couch, W. J., Project, T. S. C., ‘Measurements of Ω and Λ from 42 High-Redshift Supernovae’. In: *The Astrophysical Journal* 517.2 (June 1999), pp. 565–586 (page 11).
- [Pet+15] Peterson, J. R., Jernigan, J. G., Kahn, S. M., Rasmussen, A. P., Peng, E., Ahmad, Z., Bankert, J., Chang, C., Claver, C., Gilmore, D. K., Grace, E., Hannel, M., Hodge, M., Lorenz, S., Lupu, A., Meert, A., Nagarajan, S., Todd, N., Winans, A., Young, M., ‘SIMULATION OF ASTRONOMICAL IMAGES FROM OPTICAL SURVEY TELESCOPES USING A COMPREHENSIVE PHOTON MONTE CARLO APPROACH’. In: *The Astrophysical Journal Supplement Series* 218.1 (May 2015), p. 14 (page 57).
- [Pet+19] Peterson, J. R., Peng, E., Burke, C. J., Sembroski, G., Cheng, J., ‘Deformation of Optics for Photon Monte Carlo Simulations’. In: *The Astrophysical Journal* 873.1 (Mar. 2019), p. 98 (page 57).
- [Pet+20] Peterson, J. R., O’Connor, P., Nomerotski, A., Magnier, E., Jernigan, J. G., Cheng, J., Cui, W., Peng, E., Rasmussen, A., Sembroski, G., ‘Sensor Distortion Effects in Photon Monte Carlo Simulations’. In: *The Astrophysical Journal* 889.2 (Feb. 2020), p. 182 (page 57).
- [Pic98] Pickles, A. J. ‘A Stellar Spectral Flux Library: 1150 - 25000 Å’. In: *Publications of the Astronomical Society of the Pacific* 110.749 (July 1998), pp. 863–878 (page 140).
- [Pir+09] Pires, S., Starck, J.-L., Amara, A., Teyssier, R., Réfrégier, A., Fadili, J., ‘FAST STatistics for weak Lensing (FASTLens): fast method for weak lensing statistics and map making’. In: *Monthly Notices of the Royal Astronomical Society* 395.3 (May 2009), pp. 1265–1279 (page 20).
- [Pla+20] Planck Collaboration, ‘Planck 2018 results - VI. Cosmological parameters’. In: *A&A* 641 (2020), A6 (pages 11, 12).
- [Puj+20] Pujol, Arnau, Sureau, Florent, Bobin, Jerome, Courbin, Frederic, Gentile, Marc, Kilbinger, Martin, ‘Shear measurement bias - I. Dependencies on methods, simulation parameters, and measured parameters’. In: *A&A* 641 (2020), A164 (page 29).
- [Rac+16] Racca, G. D., Laureijs, R., Stagnaro, L., Salvignol, J.-C., Alvarez, J. L., Criado, G. S., Venancio, L. G., Short, A., Strada, P., Bönke, T., Colombo, C., Calvi, A., Maiorano, E., Piersanti, O., Prezelus, S., Rosato, P., Pinel, J., Rozemeijer, H., Lesna, V., Musi, P., Sias, M., Anselmi, A., Cazaubiel, V., Vaillon, L., Mellier, Y., Amiaux, J., Berthé, M., Sauvage, M., Azzollini, R., Cropper, M., Pottinger, S., Jahnke, K., Ealet, A., Maciaszek, T.,

- Pasian, F., Zacchei, A., Scaramella, R., Hoar, J., Kohley, R., Vavrek, R., Rudolph, A., Schmidt, M., ‘The Euclid mission design’. In: *Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave*. Ed. by H. A. MacEwen, G. G. Fazio, M. Lystrup, N. Batalha, N. Siegler and E. C. Tong. Vol. 9904. International Society for Optics and Photonics. SPIE, 2016, pp. 235–257 (pages 31, 67, 128).
- [RB03] Refregier, A., Bacon, D., ‘Shapelets — II. A method for weak lensing measurements’. In: *Monthly Notices of the Royal Astronomical Society* 338.1 (Jan. 2003), pp. 48–56 (page 63).
- [Ref03] Refregier, A. ‘Shapelets — I. A method for image analysis’. In: *Monthly Notices of the Royal Astronomical Society* 338.1 (Jan. 2003), pp. 35–47 (page 63).
- [Rho+07] Rhodes, J. D., Massey, R. J., Albert, J., Collins, N., Ellis, R. S., Heymans, C., Gardner, J. P., Kneib, J.-P., Koekemoer, A., Leauthaud, A., Mellier, Y., Refregier, A., Taylor, J. E., Waerbeke, L. V., ‘The Stability of the Point-Spread Function of the Advanced Camera for Surveys on the Hubble Space Telescope and Implications for Weak Gravitational Lensing’. In: *The Astrophysical Journal Supplement Series* 172.1 (Sept. 2007), pp. 203–218 (page 58).
- [Rho+10] Rhodes, J., Leauthaud, A., Stoughton, C., Massey, R., Dawson, K., Kolbe, W., Roe, N., ‘The Effects of Charge Transfer Inefficiency (CTI) on Galaxy Shape Measurements’. In: *Publications of the Astronomical Society of the Pacific* 122.890 (Apr. 2010), pp. 439–450 (pages 24, 53).
- [Ric+19] Ricaud, B., Borgnat, P., Tremblay, N., Gonçalves, P., Vandergheynst, P., ‘Fourier could be a data scientist: From graph Fourier transform to signal processing on graphs’. In: *Comptes Rendus Physique* 20.5 (2019). Fourier and the science of today / Fourier et la science d’aujourd’hui, pp. 474–488 (page 133).
- [Row+15] Rowe, B., Jarvis, M., Mandelbaum, R., Bernstein, G. M., Bosch, J., Simet, M., Meyers, J. E., Kacprzak, T., Nakajima, R., Zuntz, J., ‘GALSIM: The modular galaxy image simulation toolkit’. In: *Astronomy and Computing* 10 (2015), pp. 121–150 (page 92).
- [Row10] Rowe, B. ‘Improving PSF modelling for weak gravitational lensing using new methods in model selection’. In: *MNRAS* 404.1 (2010), pp. 350–366 (page 72).
- [Sal+09] Salmon, D., Cuillandre, J.-C., Barrick, G., Thomas, J., Ho, K., Matsushige, G., Benedict, T., Racine, R., ‘CFHT Image Quality and the Observing Environment’. In: *Publications of the Astronomical Society of the Pacific* 121.882 (2009), pp. 905–921 (page 57).
- [San14] Sandin, Christer, ‘The influence of diffuse scattered light - I. The PSF and its role in observations of the edge-on galaxy NGC 5907’. In: *A&A* 567 (2014), A97 (page 49).
- [Sch+02] Schneider, P. van Waerbeke, L. Kilbinger, M. Mellier, Y. ‘Analysis of two-point statistics of cosmic shear - I. Estimators and covariances’. In: *A&A* 396.1 (2002), pp. 1–19 (page 21).

- [Sch+10] Schrabback, T., Hartlap, J., Joachimi, B., Kilbinger, M., Simon, P., Benabed, K., Bradac, M., Eifler, T., Erben, T., Fassnacht, C. D., High, F., William, Hilbert, S., Hildebrandt, H., Hoekstra, H., Kuijken, K., Marshall, P. J., Mellier, Y., Morganson, E., Schneider, P., Semboloni, E., Van Waerbeke, L., Velander, M. ‘Evidence of the accelerated expansion of the Universe from weak lensing tomography with COSMOS*’. In: *A&A* 516 (2010), A63 (pages 58, 60).
- [Sch+15] Schneider, M. D., Hogg, D. W., Marshall, P. J., Dawson, W. A., Meyers, J., Bard, D. J., Lang, D., ‘HIERARCHICAL PROBABILISTIC INFERENCE OF COSMIC SHEAR’. In: *The Astrophysical Journal* 807.1 (July 2015), p. 87 (pages 29, 197).
- [Sch+18] Schmitz, M. A., Heitz, M., Bonneel, N., Ngole, F., Coeurjolly, D., Cuturi, M., Peyré, G., Starck, J.-L., ‘Wasserstein Dictionary Learning: Optimal Transport-Based Unsupervised Nonlinear Dictionary Learning’. In: *SIAM Journal on Imaging Sciences* 11.1 (2018), pp. 643–678 (pages 64, 127).
- [Sch+20] Schmitz, M. A., Starck, J. -.-L., Mboula, F. N., Auricchio, N., Brinckmann, J., Capobianco, R. I. V., Clédassou, R., Conversi, L., Corcione, L., Fourmanoit, N., Frailis, M., Garilli, B., Hormuth, F., Hu, D., Israel, H., Kermiche, S., Kitching, T. D., Kubik, B., Kunz, M., Ligorì, S., Lilje, P. B., Lloro, I., Mansutti, O., Marggraf, O., Massey, R. J., ‘Euclid: Non-parametric point spread function field recovery through interpolation on a graph Laplacian’. In: *A&A* 636 (2020), A78 (pages 62, 74, 82, 84, 86, 89, 91, 122, 123, 128, 130, 133, 142, 143, 148).
- [Sch10] Schmidt, J. D. *Numerical Simulation of Optical Wave Propagation with Examples in MATLAB*. SPIE, July 2010 (pages 36, 136).
- [Sch19] Schmitz, M. A. ‘Euclid weak lensing : PSF field estimation’. 2019SACLS359. PhD thesis. 2019 (pages 45, 64, 122, 127).
- [SCU16] Soulez, F., Courbin, F., Unser, M., ‘Back-propagating the light of field stars to probe telescope mirrors aberrations’. In: *Proc. SPIE 9912, Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation II, 991277* (July 22, 2016) (Aug. 2016) (pages 64, 127).
- [SEF92] Schneider, P., Ehlers, J., Falco, E. E., *Gravitational Lenses*. Springer, 1992 (page 13).
- [Sér63] Sérsic, J. L. ‘Influence of the atmospheric and instrumental dispersion on the brightness distribution in a galaxy’. In: *Boletín de la Asociación Argentina de Astronomía La Plata Argentina* 6 (Feb. 1963), pp. 41–43 (page 28).
- [SH17] Sheldon, E. S., Huff, E. M., ‘Practical Weak-lensing Shear Measurement with Metacalibration’. In: *ApJ* 841.1 (2017), p. 24 (pages 29, 30, 109).
- [She+15] Shechtman, Y., Eldar, Y. C., Cohen, O., Chapman, H. N., Miao, J., Segev, M., ‘Phase Retrieval with Application to Optical Imaging: A contemporary overview’. In: *IEEE Signal Processing Magazine* 32.3 (2015), pp. 87–109 (pages 58, 128).
- [She+20] Sheldon, E. S., Becker, M. R., MacCrann, N., Jarvis, M., ‘Mitigating Shear-dependent Object Detection Biases with Metacalibration’. In: *The Astrophysical Journal* 902.2 (Oct. 2020), p. 138 (pages 29, 30).
- [She15] Sheldon, E. *NGMIX: Gaussian mixture models for 2D images*. Aug. 2015 (pages 29, 109, 110).

- [SIH18] Salvato, M., Ilbert, O., Hoyle, B., *The many flavours of photometric redshifts*. 2018 (page 23).
- [SMB11] Starck, J.-L., Murtagh, F., Bertero, M., ‘Starlet transform in astronomical data processing’. In: *Handbook of Mathematical Methods in Imaging*. Springer, 2011, pp. 1489–1531 (pages 88, 110, 126).
- [SMF15] Starck, J.-L., Murtagh, F., Fadili, J., *Sparse Image and Signal Processing: Wavelets and Related Geometric Multiscale Analysis*. 2nd ed. Cambridge University Press, 2015 (pages 62, 136, 204).
- [Sou+12] Soulez, F., Denis, L., Tourneur, Y., Thiébaud, É., ‘Blind deconvolution of 3D data in wide field fluorescence microscopy’. In: *2012 9th IEEE International Symposium on Biomedical Imaging (ISBI)*. 2012, pp. 1735–1738 (page 45).
- [Spe+15] Spergel, D., Gehrels, N., Baltay, C., Bennett, D., Breckinridge, J., Donahue, M., Dressler, A., Gaudi, B. S., Greene, T., Guyon, O., Hirata, C., Kalirai, J., Kasdin, N. J., Macintosh, B., Moos, W., Perlmutter, S., Postman, M., Rauscher, B., Rhodes, J., Wang, Y., Weinberg, D., Benford, D., Hudson, M., Jeong, W. -S., Mellier, Y., Traub, W., Yamada, T., Capak, P., Colbert, J., Masters, D., Penny, M., Savransky, D., Stern, D., Zimmerman, N., Barry, R., Bartusek, L., Carpenter, K., Cheng, E., Content, D., Dekens, F., Demers, R., Grady, K., Jackson, C., Kuan, G., Kruk, J., Melton, M., Nemati, B., Parvin, B., Poberezhskiy, I., Peddie, C., Ruffa, J., Wallace, J. K., Whipple, A., Wollack, E., Zhao, F., *Wide-Field Infrared Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report*. 2015 (page 123).
- [SS97] Seitz, C., Schneider, P., ‘Steps towards nonlinear cluster inversion through gravitational distortions. III. Including a redshift distribution of the sources.’ In: *Astronomy and Astrophysics - A&A* 318 (Feb. 1997), pp. 687–699 (page 18).
- [SSE94] Seitz, S., Schneider, P., Ehlers, J., ‘Light propagation in arbitrary space-times and the gravitational lens approximation’. In: *Classical and Quantum Gravity* 11.9 (Sept. 1994), pp. 2345–2373 (page 13).
- [Sta+21] Starck, J.-L., Themelis, K. E., Jeffrey, N., Peel, A., Lanusse, F., ‘Weak-lensing mass reconstruction using sparsity and a Gaussian random field’. In: *A&A* 649 (2021), A99 (page 20).
- [Sto+04] Storkey, A. J., Hambly, N. C., Williams, C. K. I., Mann, R. G., ‘Cleaning sky survey data bases using Hough transform and renewal string approaches’. In: *Monthly Notices of the Royal Astronomical Society* 347.1 (Jan. 2004), pp. 36–51 (page 49).
- [Tew+19] Tewes, M., Kuntzer, T., Nakajima, R., Courbin, F., Hildebrandt, H., Schrabback, T., ‘Weak-lensing shear measurement with machine learning - Teaching artificial neural networks about feature noise’. In: *A&A* 621 (2019), A36 (page 28).
- [TH15] Troxel, M., Ishak, M., ‘The intrinsic alignment of galaxies and its impact on weak gravitational lensing in an era of precision cosmology’. In: *Physics Reports* 558 (2015). The intrinsic alignment of galaxies and its impact on weak gravitational lensing in an era of precision cosmology, pp. 1–59 (page 18).

- [Ven+16] Venancio, L. M. G., Carminati, L., Alvarez, J. L., Amiaux, J., Bonino, L., Salvignol, J.-C., Vavrek, R., Laureijs, R., Short, A., Boenke, T., Strada, P., ‘Coating induced phase shift and impact on Euclid imaging performance’. In: *Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave*. Ed. by H. A. MacEwen, G. G. Fazio, M. Lystrup, N. Batalha, N. Siegler and E. C. Tong. Vol. 9904. International Society for Optics and Photonics. SPIE, 2016, pp. 325–330 (pages 52, 64).
- [Ven+20] Venancio, L. M. G., Carminati, L., Amiaux, J., Bonino, L., Racca, G., Vavrek, R., Laureijs, R., Short, A., Boenke, T., Strada, P., ‘Status of the performance of the Euclid spacecraft’. In: *Space Telescopes and Instrumentation 2020: Optical, Infrared, and Millimeter Wave*. Ed. by M. Lystrup, M. D. Perrin, N. Batalha, N. Siegler and E. C. Tong. Vol. 11443. International Society for Optics and Photonics. SPIE, 2020, pp. 45–60 (pages 33, 140).
- [VMB11] Viola, M., Melchior, P., Bartelmann, M., ‘Biases in, and corrections to, KSB shear measurements’. In: *Monthly Notices of the Royal Astronomical Society* 410.4 (Jan. 2011), pp. 2156–2166 (pages 28, 59).
- [Vri+07] Vries, W. H., Olivier, S. S., Asztalos, S. J., Rosenberg, L. J., Baker, K. L., ‘Image Ellipticity from Atmospheric Aberrations’. In: *The Astrophysical Journal* 662.1 (June 2007), pp. 744–749 (page 56).
- [Wan+20] Wang, F., Bian, Y., Wang, H., Lyu, M., Pedrini, G., Osten, W., Barbastathis, G., Situ, G., ‘Phase imaging with an untrained neural network’. In: *Light: Science and Applications* 9.1 (2020) (page 128).
- [Won+21] Wong, A., Pope, B., Desdoigts, L., Tuthill, P., Norris, B., Betters, C., ‘Phase retrieval and design with automatic differentiation’. In: *Journal of the Optical Society of America B* 38.9 (2021), p. 2465 (page 128).
- [XY13] Xu, Y., Yin, W., ‘A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion’. English (US). In: *SIAM Journal on Imaging Sciences* 6.3 (Oct. 2013), pp. 1758–1789 (page 88).
- [Zha+19] Zhang, J., Dong, F., Li, H., Li, X., Li, Y., Liu, D., Luo, W., Fu, L., Li, G., Fan, Z., ‘Testing Shear Recovery with Field Distortion’. In: *The Astrophysical Journal* 875.1 (Apr. 2019), p. 48 (page 64).
- [Zha+22] Zhang, T., Almoubayyed, H., Mandelbaum, R., Meyers, J. E., Jarvis, M., Kannawadi, A., Schmitz, M. A., Guinot, A., Collaboration, T. L. D. E. S., *Impact of Point Spread Function Higher Moments Error on Weak Gravitational Lensing II: A Comprehensive Study*. 2022 (page 75).
- [Zha07] Zhang, J. ‘Measuring the cosmic shear in Fourier space’. In: *Monthly Notices of the Royal Astronomical Society* 383.1 (Dec. 2007), pp. 113–118 (page 63).
- [Zha11] Zhang, J. ‘Measuring the reduced shear’. In: *Journal of Cosmology and Astroparticle Physics* 2011.11 (Nov. 2011), pp. 041–041 (page 64).
- [ZLF15] Zhang, J., Luo, W., Foucaud, S., ‘Accurate shear measurement with faint sources’. In: *Journal of Cosmology and Astroparticle Physics* 2015.01 (Jan. 2015), pp. 024–024 (page 64).

- [ZMC21] Zhang, T., Mandelbaum, R., Collaboration, T. L. D. E. S., ‘Impact of point spread function higher moments error on weak gravitational lensing’. In: *Monthly Notices of the Royal Astronomical Society* 510.2 (Dec. 2021), pp. 1978–1993 (page 75).
- [Zun+13] Zuntz, J., Kacprzak, T., Voigt, L., Hirsch, M., Rowe, B., Bridle, S., ‘IM3SHAPE: a maximum likelihood galaxy shear measurement code for cosmic gravitational lensing’. In: *MNRAS* 434.2 (2013), pp. 1604–1618 (page 29).
- [Zun+18] Zuntz, J., Sheldon, E., Samuroff, S., Troxel, M. A., Jarvis, M., MacCrann, N., Gruen, D., Prat, J., Sánchez, C., Choi, A., ‘Dark Energy Survey Year 1 results: weak lensing shape catalogues’. In: *MNRAS* 481.1 (2018), pp. 1149–1182 (pages 24, 29, 30, 60, 103).
- [Zwi37] Zwicky, F. ‘On the Masses of Nebulae and of Clusters of Nebulae’. In: *The Astrophysical Journal* 86 (Oct. 1937), p. 217 (page 11).