



HAL
open science

Generative modeling for weak lensing inverse problems

Benjamin Remy

► **To cite this version:**

Benjamin Remy. Generative modeling for weak lensing inverse problems. Instrumentation and Methods for Astrophysic [astro-ph.IM]. Université Paris-Saclay, 2023. English. NNT : 2023UPASP163 . tel-04594032

HAL Id: tel-04594032

<https://theses.hal.science/tel-04594032v1>

Submitted on 30 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Generative modeling for weak lensing inverse problems

*Modèles génératifs pour problèmes inverses
de lentille gravitationnelle faible*

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, et le
co-encadrement de **François LANUSSE**, chargé de recherche

Thèse soutenue à Paris-Saclay, le 27 novembre 2023, par

Benjamin Remy

Composition du jury

Membres du jury avec voix délibérative

André Ferrari Professeur, Université Côte d'Azur, CNRS and Ob- servatoire de la Côte d'Azur	Président
Gary Bernstein Professeur, University of Pennsylvania	Rapporteur & Examineur
Michel Elad Professeur, Technion	Rapporteur & Examineur
Emille E. O. Ishida Ingénieure de recherche CNRS, LPC, Université Clermont-Auvergne	Examinatrice
Laurence Perreault-Levasseur Professeure adjointe, Université de Montréal	Examinatrice

Titre : Modèles génératifs pour problèmes inverses de lentille gravitationnelle faible

Mots clés : lentillage gravitationnel, apprentissage automatique, modèles génératifs, cartes de masse, cisaillement gravitationnel

Résumé : Le lentillage gravitationnelle, qui génère un effet de déformation des images de galaxies lointaines à travers l'influence de densités de matières massives dans la ligne de visée, est très prometteur pour répondre aux questions relatives à la matière noire et à l'énergie sombre. Cet effet permet de sonder directement la distribution de matière noire dans l'Univers, qui est invisible autrement. Dans le régime où ces déformations sont faibles, il est possible de cartographier la distribution de matières projetées dans la ligne de visée, appelée carte de masse, à partir de la mesure de la déformation d'un grand nombre de galaxies. Cependant, la reconstruction de ces cartes de masse est un problème inverse qui est mal posé, à cause de données manquantes et de bruits dans le signal mesuré, et nécessite donc de l'information à priori pour être résolu.

L'objectif principal de cette thèse est d'utiliser les récentes avancées sur les modèles génératifs qui permettent de modéliser des distributions

complexes dans des espaces de très haute dimension. Nous proposons en particulier une nouvelle méthode pour résoudre les problèmes inverses de haute dimension et mal posés en caractérisant la distribution a posteriori complète. En apprenant la distribution a priori à partir de simulations cosmologiques, nous pouvons reconstruire des cartes de masse de très hautes résolutions, y compris aux petites échelles, tout en quantifiant les incertitudes associées. De plus, nous présentons une nouvelle méthode de mesure du cisaillement gravitationnel en créant un modèle décrivant les données observées au niveau des pixels. Contrairement aux méthodes standards, cette méthode ne repose pas sur la mesure d'ellipticité des galaxies et introduit donc un nouveau paradigme pour la mesure du cisaillement gravitationnel. Nous proposons en particulier un modèle hiérarchique Bayésien, avec des composantes génératives apprises et des composantes analytiques physiques. Nous montrons que cela permet de résoudre le biais de modèles dans l'estimation du cisaillement gravitationnel.

Title : Generative modeling for weak lensing inverse problems

Keywords : weak lensing, machine learning, generative models, mass-mapping, cosmic shear

Abstract : Gravitational lensing, which is the effect of the distortion of distant galaxy images through the influence of massive matter densities in the line of sight, holds significant promise in addressing questions about dark matter and dark energy. It reflects the distribution of total matter of the Universe and is therefore a promising probe for cosmological models. In the case where these distortions are small, we call it the weak gravitational lensing regime and a straightforward mapping exists between the matter distribution projected in the line of sight, called mass-map, and the measured lensing effect. However, when attempting to reconstruct matter mass-maps under conditions involving missing data and high noise corruption, this linear inverse problem becomes ill-posed and may lack a meaningful solution without additional prior knowledge.

The main objective of this thesis is to employ

recent breakthroughs in the generative modeling literature that enable the modeling of complex distribution in high-dimensional spaces. We propose in particular a novel methodology to solve high-dimensional ill-posed inverse problems, characterizing the full posterior distribution of the problem. By learning the high dimensional prior from cosmological simulations, we demonstrate that we are able to reconstruct high-resolution 2D mass-maps alongside uncertainty quantification. Additionally, we present a new method for cosmic shear estimation based on forward modeling of the observation at the pixel level. This represents a new paradigm for weak lensing measurement as it does not rely on galaxy ellipticities anymore. In particular, we propose to build a hybrid generative and physical hierarchical Bayesian model and demonstrate that we can remove the source of model bias in the estimation of the cosmic shear.

Acknowledgments

Je tiens d'abord à remercier les membres du jury de ma soutenance de thèse, pour leur considération et leur bienveillance, Professeur Gary Bernstein, Professeur Michael Elad, Professeure Laurence Perreault Levasseur, Docteure Emille E. O. Ishida et Professeur André Ferrari.

Je tiens à exprimer ma profonde gratitude envers Jean-Luc, mon directeur de thèse. Merci pour ton accueil chaleureux au sein de Cosmostat, ton accompagnement scientifique, tes précieux conseils et la liberté dont tu m'as fait bénéficier. Tout cela a été la condition de mon épanouissement en thèse. Je suis profondément reconnaissant d'avoir eu l'opportunité de travailler à tes côtés et de bénéficier de ton expertise.

Mes pensées se tournent alors vers François, mon co-directeur de thèse, en qui j'ai une reconnaissance infinie. Ton accompagnement constant, autant scientifique que moral, ton écoute et ta patience ont permis une expérience de thèse exceptionnelle. Tu as partagé avec moi tes connaissances, tes idées et ton temps sans réserve. Je suis très fier d'avoir commencé mon parcours de recherche à tes côtés et tu resteras une source d'inspiration inépuisable.

Merci également à Florent et Alexandre pour leurs précieux conseils durant nos comités de suivi, ainsi qu'à Thierry Fouchet et Laurent Verstraete pour nos échanges de mails et leur suivi via l'école doctorale.

Ma thèse fut réalisée au sein du laboratoire Cosmostat dont les membres permanents œuvrent au quotidien pour générer une ambiance de soutien aux étudiants. Au travers des learning slots, des journal clubs et des hackathons en tout genre, j'ai beaucoup appris grâce à vous Sam, Martin, Joana et Valeria. A special mention goes for Sam and Martin, for their tireless warmth and unwavering kindness that breathes life into Cosmostat.

J'ai eu le grand plaisir de faire ma thèse aux côtés d'étudiants et de postdoc, et je tiens à les remercier. En particulier à Fadi, dont il est impossible de compter le temps passé ensemble à échanger, rire, faire du sport, se serrer les coudes. Tobias qui est resté une source d'apaisement et d'inspiration tout du long et un précieux camarade d'escalade. Virginia et Rémi pour m'avoir montré la voie d'un thésard au labo. Et bien sûr à Zaccharie pour avoir été mon tout premier collaborateur, modèle en machine learning et été une très grande source d'inspiration. Je tiens aussi à remercier ces jeunes chercheurs qui m'ont tant inspiré dans l'ombre sans le savoir, Raphaël Duqué, Matthieu Gosselin, Remy Joseph et Morgan Schmitz. Et bien sûr à tous les postdocs, thésards et stagiaires du labo, merci pour nos discussions et nos rires, Justine, Denise, Sergio, Benjamin, Lisa, Fabian, Villasini, Nicolas, André, Lucie, Jennifer, Romain, Sacha, Ezequiel, Alex, Hubert, Utsav, Emma, Jingwei et Elisa.

Merci enfin à mes cousins cousines, oncles et tantes, et grands-parents pour tous ces moments de bonheur passés ensemble pendant mon parcours. Je suis tellement heureux que chacun mène son aventure tout en étant émerveillé de celle des autres.

Merci infiniment à mes parents, mon frère et ma sœur, grâce à qui tout fut possible. Je vous aime. Je vous dois tout.

Acknowledgements for funding and computing resources

My Ph.D. thesis was supported by the Centre National d'Etudes Spatiales and by a public grant overseen by the French National Research Agency (ANR) through the program UDOPIA, project funded by the ANR-20-THIA-0013-01. This work was granted access to the HPC resources of IDRIS under the allocations 2021-AD011011554R1 and 2022-102038 made by GENCI (Grand Equipement National de Calcul Intensif).

Contents

Acknowledgments	iv
Contents	vii
Acronyms	xi
1 Introduction	1
I Context	3
2 Short introduction to cosmology	5
2.1 Toolkit for modern cosmology	5
2.1.1 General Relativity and the FLRW metric	6
2.1.2 Dynamical evolution of the Universe	6
2.1.3 Λ CDM, dark matter and dark energy	8
2.2 Large scale structure	9
2.2.1 Matter density contrast	9
2.2.2 Matter power spectrum	10
2.3 Weak gravitational lensing	11
2.3.1 Gravitational potential and the lens equation	11
2.3.2 Shear and convergence fields	14
3 Cosmic shear measurement methods	17
3.1 Weak lensing galaxy images	18
3.1.1 Galaxy morphologies	18
3.1.2 Shearing transform	19
3.1.3 Point spread function	20
3.1.4 Image pixelization	22
3.1.5 Noise corruption sources	22
3.1.6 From patch to postage stamps (detection and blending)	22
3.2 Shape measurement methods	23
3.2.1 (Weighted-) Moments	24
3.2.2 Model fitting	26
3.3 PSF correction methods	29
3.3.1 KSB	29
3.3.2 BJ02 and re-Gaussianization	29

3.4	Calibration methods	31
3.4.1	Response to shear	31
3.4.2	Selection response	32
3.4.3	METADETECTION with automatic differentiation	34
3.5	Conclusion	35
II Generative modeling and Bayesian statistics		37
4	Bayesian statistics	39
4.1	Introduction to Bayesian modeling	39
4.1.1	General introduction of Bayesian modeling	40
4.1.2	Bayesian modeling in Cosmology	41
4.2	Markov Chain Monte Carlo sampling	41
4.2.1	Metropolis-Hastings	42
4.2.2	Langevin dynamics	43
4.2.3	Hamiltonian Monte Carlo	43
4.3	Variational inference	45
4.3.1	Forward and reverse KL divergences	47
4.3.2	Evidence lower bound (Evidence Lower Bound (ELBO))	48
5	Deep generative models	51
5.1	Introduction to deep learning	52
5.1.1	Multi Layer Perceptron	52
5.1.2	Optimizing neural networks	54
5.1.3	Automatic differentiation	55
5.1.4	Differentiating implicit functions	56
5.1.5	Architectures	57
5.1.6	Normalization	58
5.2	Normalizing Flows	60
5.2.1	Bijections	62
5.2.2	Optimizing a normalizing flow	63
5.2.3	Conditional normalizing flows	64
5.3	Score-based diffusion models	65
5.3.1	Stochastic differential equations	65
5.3.2	Denosing score matching	66
5.4	Annealed sampling	67
III Contributions to weak lensing inverse problems		71
6	Weak lensing mass-mapping	73
6.1	Weak lensing Mass-Mapping	74
6.1.1	Shear and convergence	74
6.1.2	Mass-mapping inverse problem in a Bayesian framework	75
6.2	COSMOS dark matter mass-map reconstruction	77
6.2.1	COSMOS survey	77
6.2.2	Simulating COSMOS lensing catalog from κ TNG suite of simulations	79
6.3	Literature review on priors	80
6.3.1	Flat prior: Kaiser-Squires Reconstruction	80
6.3.2	Gaussian Random Field: Wiener Filter	81

6.3.3	Sparse priors	82
6.3.4	Deep implicit prior	82
6.3.5	Score-matching prior	83
6.4	High-dimensional posterior sampling	84
6.4.1	The posterior annealing problem	85
6.4.2	Mass-maps posterior sampling	86
6.5	Posterior moments validation	88
6.5.1	Gaussian posterior	88
6.5.2	Full posterior moments against moments network	89
6.5.3	Point estimates validations	92
6.6	Reconstruction of the COSMOS field	93
6.6.1	HSC/ACS COSMOS posterior mass maps	93
6.6.2	Cluster detection with posterior samples	94
6.6.3	Going further	94
7	Hierarchical Bayesian modeling of cosmic shear and galaxy morphologies	101
7.1	Pixel-level weak lensing Hierarchical Bayesian Model	102
7.1.1	Hubble Space Telescope (HST)/Advanced Camera for Survey (ACS) Cosmic Evolution Survey (COSMOS) postage stamps	103
7.1.2	Galaxy Latent Variable model	104
7.1.3	Forward modeling galaxy images	107
7.2	Solving the model bias with LVMS	109
7.2.1	Sampling with Hamiltonian Monte Carlo	110
7.2.2	Mean-field Variational inference	111
7.3	Going further	112
7.3.1	Identifying the model bias with HSM	112
7.3.2	Accurate Variational inference with Normalizing Flows	113
7.3.3	Combining posterior distributions	114
7.4	Conclusion	115
	Conclusion	117
	Appendices	119
	A Résumé en français	121
	List of Figures	125
	Bibliography	127

Acronyms

ACS Advanced Camera for Survey

BAO Baryon Acoustic Oscillations

BLISS Bayesian Light Source Separator

BJ02 Bernstein and Jarvis 2002

CAMB Code for Anisotropies in the Microwave Background

CAMELS Cosmology and Astrophysics with MachinE Learning Simulations

CANDELS Cosmic Assembly Near-infrared Deep Extragalactic Legacy Survey

CCD Charge-Coupled Device

CLASS Cosmic Linear Anisotropy Solving System

CMB Cosmic Microwave Background

CNN Convolutional Neural Network

COSMOS Cosmic Evolution Survey

DSPS Differentiable Stellar Population Synthesis

ELBO Evidence Lower Bound

ESA European Space Agency

FLRW Friedmann-Lemaître-Robertson-Walker

GREAT₃ Third Gravitational Lensing Accuracy Testing

GRF Gaussian Random Field

HBM Hierarchical Bayesian Model

HMC Hamiltonian Monte Carlo

HSC Hyper Suprime-Cam

HSM Hirata Seljak and Mandelbaum

HST Hubble Space Telescope

JWST James Webb Space Telescope

KL Kullback–Leibler

KSB Kaiser-Squires-Broadhurst

LVM Latent Variable Model

LSS Large Scale Structure

LSST Large Synoptic Survey Telescope

MALA Metropolis Adjusted Langevin Algorithm

MAP Maximum a Posteriori

MCMC Markov Chain Monte Carlo

MH Metropolis-Hastings

MLE Maximum Likelihood Estimate

MLP Multilayer Perceptron

MRI Magnetic Resonance Imaging

MSE Mean Squared Error

NASA National Aeronautics and Space Administration

NF Normalizing Flow

NLL Negative Log Likelihood

NISP Near Infrared Spectro Photometer

ODE Ordinary Differential Equation

PGM Probabilistic Graphical Model

PSF Point Spread Function

PSNR Peak Signal-to-Noise Ratio

ReLU Rectified Linear Unit

RDSM Residual Denoising Score-Matching

RMSE Root Mean Square Error

RWMH Random Walk Metropolis-Hastings

S10 Schrabback 2010 catalog

SDE Stochastic Differential Equation

SDSS Sloan Digital Sky Survey

SED Spectral Energy Distribution

SN Spectral Normalization

SNR Signal-to-Noise Ratio

VAE Variational Autoencoder

VI Variational Inference

VIS *Euclid's* VIS Instrument

Introduction

The recent space-based mission *Planck* has significantly advanced our understanding of the cosmological model that describes our Universe. Through precise measurements of the Cosmic Microwave Background, this mission has firmly established the standard Λ CDM model. According to this model, 5% of the total energy density consists of baryonic matter, while substantial 27% is attributed to dark matter, with an overwhelming 68% assigned to dark energy. The standard model is now well established with its parameters measured to a remarkable level of accuracy. Nevertheless, numerous fundamental questions persist, particularly concerning the nature of dark matter and dark energy and therefore about our comprehension of the Universe. Upcoming stage IV wide-field optical galaxy surveys, including the ESA *Euclid* mission, the Vera C. Rubin Observatory Legacy Survey of Space and Time and the Roman Space Telescope are designed to push further this investigation, by probing the large-scale structure of the Universe.

Gravitational lensing, which is the effect of the distortion of distant galaxy images through the influence of massive matter densities in the line of sight, holds significant promise in addressing questions about dark matter and dark energy. It reflects the distribution of total matter of the Universe and is therefore a promising probe for cosmological models. In the case where these distortions are small, we call it the weak gravitational lensing regime and a straightforward mapping exists between the projected matter distribution, called mass-map, and the measured lensing effect. However, when attempting to reconstruct matter mass-maps under conditions involving missing data and high noise corruption, this linear inverse problem becomes ill-posed and may lack a meaningful solution without additional prior knowledge.

The signal processing literature has already extensively addressed ill-posed inverse problems, for mass-mapping and beyond astrophysics such as medical imaging or geophysics. These methods, which operate under the assumption that the signal to recover exhibits sparsity in its representation, have been demonstrated to be successful in reconstructing highly resolved mass-maps. However, to this date, these methods have fallen short in terms of quantifying the uncertainty associated with the reconstruction. From a Bayesian perspective, recovering the full posterior distribution of a high-dimensional inverse problem remains an unresolved challenge.

The effect of lensing can be measured through a field called the cosmic shear, which itself presents many challenges as well. One of the significant challenges in cosmic shear analysis from upcoming surveys requires exceedingly accurate measurements of galaxy shapes. These shape measurements, in turn, depend on exquisite control of the instrument's Point Spread Function and can be impacted by the details of galaxy morphologies, which can induce biases in the cosmic shear measurement. Consequently,

it is imperative to develop methods that can effectively address all potential sources of bias and account for the propagation of errors throughout the measurement pipeline up to the parameters of interest.

The main objective of this thesis is to employ recent breakthroughs in the generative modeling literature that enable the modeling of complex distribution in high-dimensional spaces. This will be applied to develop cutting-edge methods for mass-mapping tools and cosmic shear measurement for next-generation surveys. Ultimately, these advancements will enhance our ability to constrain cosmological parameters, while ensuring accurate estimation of their uncertainties.

The first part of the thesis aims to introduce the background tools and the concepts necessary to understand the motivation of this thesis. In particular, [chapter 2](#) minimally introduces the basic equations of General Relativity and provides an overview of the dynamical evolution of the Universe. We also present the Λ CDM model and how the large-scale structures form. We then present the weak lensing formalism, introducing the shear and convergence fields. [chapter 3](#) outlines the formation of galaxy images from background galaxies to the CCD and the state-of-the-art shape measurement and bias calibration methods.

The second part is dedicated to the presentation of the generative modeling framework and Bayesian modeling. In [chapter 4](#), we provide a general introduction to Bayesian modeling and efficient inference methods based on sampling with Markov chain Monte Carlo and optimization with variational inference. We then present the generative modeling techniques in [chapter 5](#) enabling us to learn arbitrary distribution, starting with a general introduction to deep learning, and the state-of-the-art methods such as normalizing flows and diffusion models. We finally present our contribution enabling inference in high dimensional spaces.

In the third part, I address the reconstruction of mass-maps from weak lensing measurements. The [chapter 6](#) starts by giving an overview of the problem and a literature review on the methods addressing it. I then develop a novel methodology to solve high-dimensional ill-posed inverse problem, characterizing the full posterior distribution of the problem. Finally, I demonstrate this method on real data, reconstructing the highest resolved 2D mass-map of the HST/ACS COSMOS field along uncertainty quantification. Finally, in [chapter 7](#), I present a new method for cosmic shear estimation based on forward modeling of the observation at the pixel level. This represents a new paradigm for weak lensing measurement as it does not rely on galaxy ellipticities anymore. As demonstrated on galaxy postage stamps from the HST/ACS COSMOS survey, inferring the cosmic shear from a hybrid generative and physical hierarchical Bayesian model enables us to remove the source of model bias in the estimation. This offers a promising path towards unbiased and accurate cosmic shear estimation under the probabilistic modeling framework.

Part I
Context

Short introduction to cosmology

Chapter Outline

2.1	Toolkit for modern cosmology	5
2.1.1	General Relativity and the FLRW metric	6
2.1.2	Dynamical evolution of the Universe	6
2.1.3	Λ CDM, dark matter and dark energy	8
2.2	Large scale structure	9
2.2.1	Matter density contrast	9
2.2.2	Matter power spectrum	10
2.3	Weak gravitational lensing	11
2.3.1	Gravitational potential and the lens equation	11
2.3.2	Shear and convergence fields	14

THIS chapter serves as an introduction to the fundamental mathematical description of our Universe. The purpose is to provide a contextual motivation for the physical quantities we will study and extract from data. We first introduce the main equations of cosmology describing the evolution of the Universe and its components. We then provide an overview of the standard Λ CDM model, which we will assume throughout this thesis. We then delve into the concept of weak gravitational lensing, which is the cosmological probe we will study and one of the primary motivations for the next generation of surveys *Euclid*, the Vera C. Observatory, and the Roman Space Telescope. We finally provide a concise description of the matter power spectrum which is a statistical quantity that can be derived from cosmological models and that we will use in the development of our methods.

2.1 Toolkit for modern cosmology

This section discusses the principal equations of Cosmology and notably introduces the Friedmann-Lemaître-Robertson-Walker (**FLRW**) metric, the comoving distance and Friedman equations. This discussion will help to draw the context of our studies and provide the necessary information about the physical quantities involved.

2.1.1 General Relativity and the FLRW metric

The current framework for describing the Universe and the dynamics of the objects we can observe is Einstein's theory of General Relativity. This theory provides the mathematical foundation for describing how the curvature of the four-dimensional space-time influences the dynamics of matter and radiation and conversely, how these entities affect the curvature itself. Riemannian geometry is the natural framework for articulating this geometrical description and allows us to relate the space-time metric to the components of the Universe. The principal equation of General Relativity is

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}\mathcal{R} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu} \quad (2.1)$$

which involves the metric tensor $g_{\mu\nu}$ which describes the space-time geometry, or the infinitesimal distances of a trajectory in a curved space-time, i.e.

$$ds^2 = \sum_{\mu,\nu=0}^3 g_{\mu\nu}dx^\mu dx^\nu. \quad (2.2)$$

Additionally, $R_{\mu,\nu}$ is the *Ricci curvature tensor*, function of the metric tensor and its derivatives. \mathcal{R} is the *Ricci scalar*, contraction of the Ricci tensor ($\mathcal{R} = g^{\mu\nu}R_{\mu\nu}$). $T_{\mu\nu}$ is the *energy-momentum tensor*.

Our current description of the Universe is built from the *cosmological principle*, i.e. assuming that the Universe is homogeneous and isotropic at large scales. An *homogeneous* Universe refers to the fact that at large scales, the matter distribution is equally distributed, or that the Universe appears the same from any location. An *isotropic* Universe refers to the fact that looking at any direction, the Universe appears the same as well. The first implies translational invariance and the second rotational invariance of space-time, which can be described with the Friedmann-Lemaître-Robertson-Walker (**FLRW**) metric

$$ds^2 = c^2 dt^2 - a^2(t) (d\chi^2 + f_K(\chi) (d\theta^2 + \sin^2 \theta d\psi^2)), \quad (2.3)$$

expressed in spherical coordinates (χ, θ, ϕ) , suitable for isotropic descriptions, where χ denotes the time-independent *comoving radial distance*, $f_K(\chi)$ the *comoving angular distance*, and $a(t)$ the *scale factor*, evolving as a function of time and describing how distances vary as the Universe contracts or expands. The comoving angular distance has three distinct forms depending on the curvature K :

$$f_K(\chi) = \begin{cases} K^{-1/2} \sin(K^{1/2}\chi) & \text{for } K > 0 \text{ (spherical)} \\ \chi & \text{for } K = 0 \text{ (flat)} \\ (-K)^{-1/2} \sinh[(-K)^{1/2}\chi] & \text{for } K < 0 \text{ (hyperbolic)}. \end{cases} \quad (2.4)$$

The latest cosmological results from the Planck mission [Agh+20], analyzing the Cosmic Microwave Background (Cosmic Microwave Background (**CMB**)), are in favor of a flat universe, implying a comoving angular distance of the form $f_K(\chi) = \chi$.

2.1.2 Dynamical evolution of the Universe

To describe the dynamical evolution of the Universe, we can derive the *Friedmann equations* [Fri22] equations resulting from zeroth component of equation Equation 2.1 using the **FLRW** metric and $T^{00} = \rho$ the energy density assuming the Universe is a perfect fluid.

$$\left(\frac{\dot{a}}{a}\right) + \frac{K}{a^2} = \frac{8\pi G\rho}{3}. \quad (2.5)$$

This is the first Friedmann equation, which introduces the *Hubble parameter*

$$H(t) = \frac{\dot{a}(t)}{a(t)}, \quad (2.6)$$

and states that the evolution of the scale factor a depends only on the energy density and the curvature of the Universe. This also allows us to introduce the *Hubble constant* H_0 corresponding to the Hubble parameter at the present day, i.e. $H_0 = \dot{a}(t_0)/a(t_0) = \dot{a}(t_0)$, given that $a(t_0) = 1$ by convention. This equation however does not provide information about the expansion. We need to derive the other components of [Equation 2.1](#) leading to the second Friedmann equation

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3p) = -\frac{4\pi G}{3}\rho(1 + 3w), \quad (2.7)$$

where the acceleration of the scale factor is related to the energy density ρ , the pressure of the perfect fluid p , or alternatively to the parameter $w = p/\rho$ representing the *equation of state*. An accelerated or decelerated universe would respectively correspond to $\ddot{a} > 0$ and $\ddot{a} < 0$ and would provide information about the relative dominance of its components as it implies $w < -1/3$ or $w > -1/3$.

The Friedman equations relate the scale factor to the energy of a photon and hence to its wavelength as $\lambda \propto a(t)$. Therefore we can relate the wavelength λ_e of a photon emitted at time t_e to the wavelength λ_{obs} of the same photon at the time of observation t_{obs} through the following ratio equality

$$\frac{\lambda_e}{\lambda_{\text{obs}}} = \frac{a(t_e)}{a(t_{\text{obs}})}. \quad (2.8)$$

In an expanding universe, the scale factor is larger at time of observation compared to its value at time of emission, i.e. $a(t_{\text{obs}}) > a(t_e)$ and therefore the observed wavelength is also greater at time of observation compared to time of emission, i.e. $\lambda(t_{\text{obs}}) > \lambda(t_e)$. This translates to a shift towards the larger wavelengths, so in the visible spectrum towards the red, hence the denomination of *redshift*, which is defined as

$$z = \frac{\lambda_{\text{obs}} - \lambda_e}{\lambda_e}. \quad (2.9)$$

This also enables us to express time of emission in terms of redshift through the following relation between the scale factor at emission time and the observed redshift (considering that $a(t_{\text{obs}}) = 1$)

$$a(t_e) = \frac{1}{1 + z}. \quad (2.10)$$

This is particularly useful as the measurement of redshift can be directly linked to the comoving distance of objects through the Hubble law. The comoving distance χ is the radial physical distance rescaled by the scale factor, i.e. $d\chi = cdt/a$ such that two particles initially motionless stay equidistant in terms of this distance as the [FLRW](#) metric evolves. It is therefore possible to rewrite this distance in terms of the scale factor only through the Hubble law:

$$\chi(a) = \int_a^1 \frac{cda'}{a'^2 H(a')}, \quad (2.11)$$

or with respect to the redshift using [Equation 2.10](#)

$$\chi(z) = \int_0^z \frac{cdz'}{H(z')}. \quad (2.12)$$



Figure 2.1: Patch from the James Webb Space Telescope (James Webb Space Telescope (JWST)) Near-Infrared Camera. One can observe many objects in this image such as galaxies and stars and the galaxy cluster SMACS 0723. This picture also illustrates the effect of strong gravitational lensing of distant galaxies appearing as arcs due to the presence of massive objects in the line of sight. Credit: National Aeronautics and Space Administration (NASA) and European Space Agency (ESA).

2.1.3 Λ CDM, dark matter and dark energy

Numerous cosmological models are aiming to describe our Universe, but the most successful one to date is the Λ Cold Dark Matter (Λ CDM) model. This model compresses all the observations into 6 parameters as presented in [Table 2.1](#). The Λ CDM model primarily relies on two components: a non-vanishing cosmological constant Λ and *Cold Dark Matter* (CDM), hence its name. The model is founded on the mathematical framework detailed in the previous section, parametrized by General Relativity and the [FLRW](#) metric. The evolution of the Universe is driven by various fluids that compose it, which we discuss in the following. Ordinary matter such as baryons, lepton photons, or neutrinos, actually represent only a minority of the

total energy density of the Universe. Cold dark matter exhibits gravitational behavior similarly to ordinary matter, but without electromagnetic interactions which makes it invisible. Furthermore, the accelerated expansion of the Universe is driven by a fluid referred as *dark energy*, exercising a negative pressure. In the Λ CDM model, this energy manifests as a constant vacuum energy density and is represented by the cosmological constant Λ in Einstein's equations (e.g. Equation 2.1). At our present epoch, the Universe is dominated at roughly 68% by dark energy. Additionally, cold dark matter is the second largest contributor to the total energy density by 26.2%, while ordinary matter represents only 4.8% (see respectively the dark energy, cold dark matter, and baryonic matter densities Ω_Λ , Ω_c and Ω_b in Table 2.1).

Parameter	Symbol	Value
Baryon density	$\Omega_b h^2$	0.02212 ± 0.00022
Dark matter density	$\Omega_c h^2$	0.1206 ± 0.0021
Angular size of the sound horizon at recomb	$100\theta_{\text{MC}}$	1.04077 ± 0.00047
Reionisation optical depth	τ	0.0522 ± 0.0080
Power law amplitude	$\ln(10^{10} A_s)$	3.040 ± 0.016
Spectral index	n_s	0.9626 ± 0.0057
Hubble constant	H_0	66.88 ± 0.92
Dark energy density	Ω_Λ	0.679 ± 0.013
Total matter density	Ω_m	0.321 ± 0.013
Power spectrum normalization	σ_8	0.8118 ± 0.0089

Table 2.1: Cosmological parameters table from Aghanim et al. [Agh+20, Planck 2018 results. VI. Cosmological parameters]. The first six parameters are the one inferred with Markov Chain Monte Carlo from Planck CMB Power spectra TT, TE, EE + lowE, combined with CMB lensing and Baryon Acoustic Oscillations (BAO), and the remaining parameters are derived from them.

2.2 Large scale structure

2.2.1 Matter density contrast

We have so far described the Universe under the assumptions of homogeneity and isotropy. However, we can observe that there are indeed structures around us, such as galaxies, galaxy clusters, and even larger scale structures such as the cosmic web (see Figure 2.1). These ones can be understood as emergent structures from a small deviation to the FLRW metric which propagates to the matter density field. The matter density field is then expressed as a deviation from a mean background density

$$\delta(\mathbf{r}, a) = \frac{\delta\rho_m(\mathbf{r}, a)}{\bar{\rho}_m(a)} = \frac{\rho_m(\mathbf{r}, a) - \bar{\rho}_m(a)}{\bar{\rho}_m(a)}, \quad (2.13)$$

where $\delta(\mathbf{r}, a)$ is called the *matter density contrast*, $\rho_m(\mathbf{r}, a)$ the matter density field and $\bar{\rho}_m(a)$ the background mean density field at instant t corresponding to the scale factor $a(t)$. The evolution of the density field ρ can be derived through the *Vlasov-Poisson* [Pee80] system of equations which yield the following dynamical evolution of the density contrast δ assuming a linear regime [Dodo3]:

$$\ddot{\delta} + 2H\dot{\delta} = 4\pi G\bar{\rho}_m\delta. \quad (2.14)$$

The solution of this second-order differential equation is the sum of two modes where the spatial and temporal coordinates are decoupled:

$$\delta(\mathbf{r}, a) = D_+(a)\delta_+(\mathbf{r}) + D_-(a)\delta_-(\mathbf{r}), \quad (2.15)$$

where D_+ and D_- are called the *growing* and *decaying modes*. The decaying mode will eventually vanish and therefore will not affect the structure formation, so only the first term remains

$$D_+(a) = \frac{5\Omega_m H(a)}{2H_0} \int_0^a \frac{da'}{a'} \left(\frac{H_0}{H(a')} \right)^3. \quad (2.16)$$

For convenience, we normalize the growing mode by its value at $a = 1$

$$D(a) = \frac{D_+(a)}{D_+(a=1)}, \quad (2.17)$$

so that the growing mode equals one at present epoch, i.e. $D_+(a=1) = D_+(z=0) = 1$, and so that the density contrast at location \mathbf{r} and redshift z be $\delta(\mathbf{r}, z) = D(z)\delta(\mathbf{r})$.

2.2.2 Matter power spectrum

A statistic which can be analytically derived from the density contrast is its *two-point correlation function* of the matter density field, which is given by

$$\langle \delta(\mathbf{r}')\delta(\mathbf{r} + \mathbf{r}') \rangle_{\mathbf{r}'} = \int_{\mathbb{R}^3} \delta(\mathbf{r} + \mathbf{r}')\delta(\mathbf{r}')d\mathbf{r}' \quad (2.18)$$

and its corresponding Fourier representation, known as the *3D matter power spectrum* $P_\delta(\mathbf{k})$ expressed in terms of Fourier mode \mathbf{k} , can be obtained through the 3D Fourier transform of the density field $\tilde{\delta}(\mathbf{k}) = \frac{1}{(2\pi)^3} \int e^{-i\mathbf{k}\cdot\mathbf{r}}\delta(\mathbf{r})d^3r$:

$$\langle \tilde{\delta}(\mathbf{k})\tilde{\delta}^*(\mathbf{k}') \rangle = (2\pi)^3 P_\delta(\mathbf{k})\delta_D^{(3)}(\mathbf{k} - \mathbf{k}'), \quad (2.19)$$

where $\tilde{\delta}^*$ is the complex conjugate of $\tilde{\delta}$ and $\delta_D^{(3)}$ is the Dirac delta distribution in 3D space. As we assume isotropy, we can rewrite Equation 2.19 in terms of a scalar mode $k = |\mathbf{k}|$ such that $\langle \tilde{\delta}(k)\tilde{\delta}^*(k') \rangle = (2\pi)^3 P_\delta(k)\delta_D(k - k')$. This is the function we will refer to later in the text as the *matter power spectrum*. The matter power spectrum can be related to the *primordial power spectrum* $P_p(k)$ through a *transfer function*

$$T(k) = \frac{\tilde{\delta}(k, a=1)\tilde{\delta}(k=0, a=0)}{\tilde{\delta}(k, a=0)\tilde{\delta}(k=0, a=1)}, \quad (2.20)$$

modulating each Fourier mode between the end of inflation $a = 0$, epoch where the primordial power spectrum stands, and the present epoch $a = 1$, when the density contrast can be described by the matter power spectrum. The relation is given by

$$P_\delta(k) = T^2(k)P_p(k), \quad (2.21)$$

and where the transfer function $T(k)$ is usually computed from Boltzmann codes such as Cosmic Linear Anisotropy Solving System (CLASS) [Les11] or Code for Anisotropies in the Microwave Background (CAMB) [LCL00]. While running these codes provides very accurate results, they may be computationally expensive and not

differentiable. There exist however analytical fitting formulae which provide approximate solutions such as the one from Eisenstein et al. [EH98]. The expression of the primordial power spectrum $P_p(k)$ comes from inflation models. Initially homogenous, quantum fluctuations present in the very early Universe are transformed to macroscopic perturbation by the inflation mechanism, corresponding to an exponential expansion in a very short period [Gut81; Lin82]. Therefore, the primordial power spectrum is given by the following power law

$$P_p(k) = A_s \left(\frac{k}{k_p} \right)^{n_s - 1}, \quad (2.22)$$

where k_p is a given pivot scale, n_s called the spectral index and A_s is the amplitude of the power spectrum.

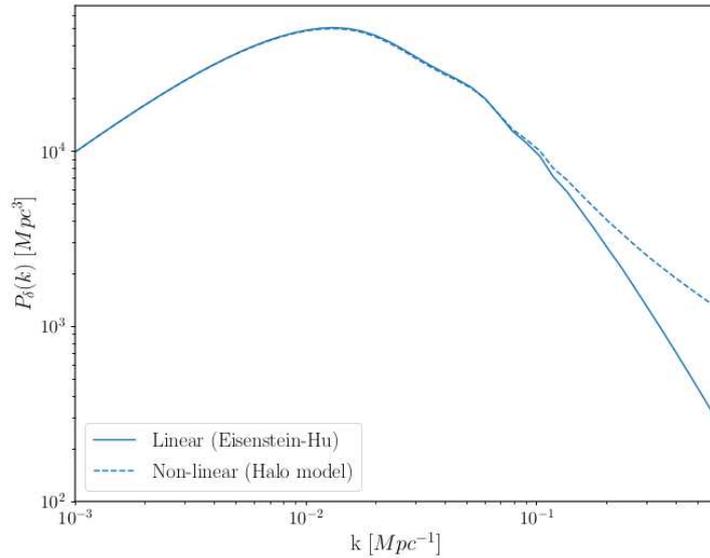


Figure 2.2: Comparison between the matter power spectrum given by Eisenstein et al. [EH98] and the HALOFIT model from Takahashi et al. [Tak+12]. Credit Virginia Ajani.

2.3 Weak gravitational lensing

This subsection is mainly based on the review on cosmic shear from Kilbinger [Kil15]. We describe here the main equations of the weak gravitational lensing probe, leading to the convergence and shear fields that we will aim to reconstruct from observations in this thesis.

2.3.1 Gravitational potential and the lens equation

Gravitational lensing refers to the effect of the inhomogeneous Universe causing a deflection of light paths. As seen in the previous section, the Universe contains fluctuations of the matter density which, for an ideal fluid of zero pressure, is related to the Newtonian gravitational potential Φ through the *Poisson* equation

$$\nabla^2\Phi = 4\pi G a^2 \bar{\rho}\delta. \quad (2.23)$$

The solution of this differential equation is given by [Equation 2.15](#). We will study cases where the effect of lensing is very small, i.e. $\Phi \ll c^2$, which is the case for most mass distribution in the Universe, especially the large scale structures (Large Scale Structure (LSS)) which consist of a network of voids, filaments, and halos. One way to derive the equations of weak gravitational lensing is to consider null geodesics on which photons propagate. However, we first need to introduce the Newtonian potential Φ in the FLRW metric which does not represent a homogeneous universe anymore:

$$ds^2 = \left(1 + \frac{2\Psi}{c^2}\right) c^2 dt^2 - a^2(t) \left(1 - \frac{2\Phi}{c^2}\right) dl^2, \quad (2.24)$$

where we contracted the notation $dl^2 = (d\chi^2 + f_K(\chi)(d\theta^2 + \sin^2\theta d\psi^2))$ and introduced the other Bardeen gravitational potential Ψ along Φ , being equals in the absence of anisotropic stress, which is the case on large scales, hence $\Psi = \Phi$. It is obvious to see that in absence of perturbation and consequently of these potentials, [Equation 2.24](#) reduces to [Equation 2.3](#). According to this metric, the travel time for a photon in proper coordinates dr describing the line path is given by the integration of the line element ds , which gives

$$t = \frac{1}{c} \int \left(1 - \frac{2\Phi}{c^2}\right) dr. \quad (2.25)$$

One can recognize in [Equation 2.25](#), the geometrical optics equation expressing the time travel of a photon going through a medium with variable refractive index $n = 1 - 2\Phi/c^2$ and changing its direction accordingly. Gravitational lensing is therefore named after this analogy, the gravitational potential acting as a medium or a *lens*. In geometrical optics, it is usual to apply Fermat's principle, also known as the *principle of least time*, to derive the *deflection angle*. Similarly, by applying $\delta t = 0$ and integrating the resulting Euler-Lagrange from [Equation 2.25](#) along the line path, we get the following deflection angle $\hat{\alpha}$, i.e. difference between the angle of emission after emission and the angle as perceived by the observer:

$$\hat{\alpha} = -\frac{2}{c^2} \int \nabla_{\perp} \Phi dr, \quad (2.26)$$

where we consider the gradient projected on the orthogonal direction to the light path, i.e. $\nabla_{\perp} = \nabla\Phi - \mathbf{e}(\mathbf{e} \cdot \nabla\Phi)$ where \mathbf{e} is a unit vector which is tangent to the light path.

We are interested in studying the deflection of light at cosmological scales, so we will now describe these phenomena in terms of comoving distance from the observer χ . Let us consider two light rays emitted by a distance source that are separated with a transverse comoving separation \mathbf{x}_0 . In absence of density perturbations, these two rays would be seen by the observer as separated by an angle θ as illustrated in [Figure 2.3](#). Therefore, we can relate the comoving separation \mathbf{x}_0 in terms of comoving angular distance f_K and angle θ :

$$\mathbf{x}_0(\chi) = f_K(\chi)\theta. \quad (2.27)$$

However in the presence of lensing, a light ray emitted by the source gets deflected by the gravitational potential Φ . Rewriting [Equation 2.25](#) in terms of comoving distance we get

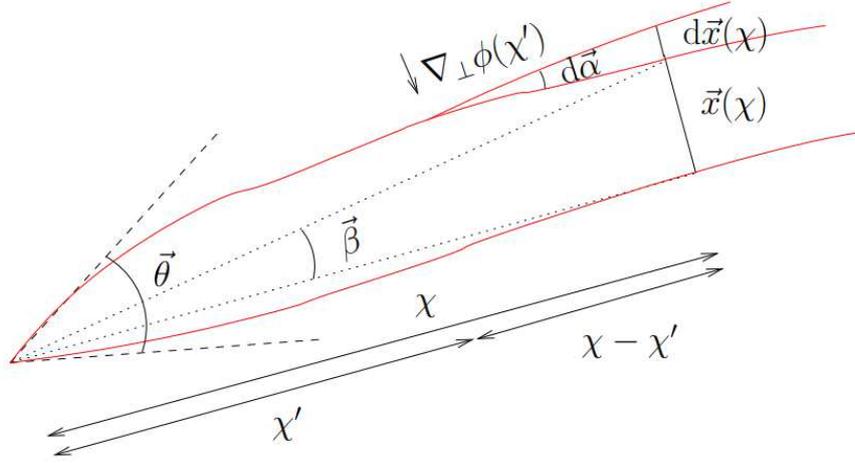


Figure 2.3: Illustration of the lensing of two light rays emitted by a distant source. The two sources have an original transverse comoving separation of $x(\chi)$. The two rays are observed with an angular separation θ , while the true angular separation between the two sources is β . This is due to the effect of gravitational lensing, deflecting the emitted ray by an angle α .

$$d\hat{\alpha} = -\frac{2}{c^2}\nabla_{\perp}\Phi(\mathbf{x}, \chi')d\chi'. \quad (2.28)$$

This allows us to quantify the change in transverse comoving distance as $d\mathbf{x} = f_K(\chi - \chi')d\hat{\alpha}$. To obtain the total change in transverse comoving separation, we need to integrate over the line of sight along χ' . Considering the deflection of the two light rays, the observed transverse comoving separation is therefore

$$\mathbf{x}(\chi) = f_K(\chi)\boldsymbol{\theta} - \frac{2}{c^2}\int_0^{\chi} f_K(\chi - \chi')\left[\nabla_{\perp}\Phi(\mathbf{x}(\chi'), \chi') - \nabla_{\perp}\Phi^{(0)}(0, \chi')\right]d\chi', \quad (2.29)$$

where $\Phi^{(0)}$ denotes the potential along the path of the second photon.

The *lens equation* can then be derived by computing the difference between the unlensed coordinates β and observed image coordinates θ . The unlensed coordinates is naturally defined by $\beta = \mathbf{x}(\chi)/f_K(\chi)$ and the *lens equation* is

$$\beta = \theta - \alpha \quad (2.30)$$

where we introduced the *deflection angle* α as

$$\alpha = \frac{2}{c^2}\int_0^{\chi} f_K(\chi - \chi')\left[\nabla_{\perp}\Phi(\mathbf{x}(\chi'), \chi') - \nabla_{\perp}\Phi^{(0)}(0, \chi')\right]d\chi'. \quad (2.31)$$

To make Equation 2.31 only dependent on the observed angle θ and comoving radial distance χ , we can substitute the separation vector \mathbf{x} by its 0th-order solution \mathbf{x}_0 defined in Equation 2.27. The potential is therefore considered the same being integrated along the lensed or unlensed path. This substitution is called the *Born approximation*.

2.3.2 Shear and convergence fields

The linear mapping between the lensed coordinates $\boldsymbol{\theta}$ to the unlensed coordinates $\boldsymbol{\beta}$ is described by the Jacobian $\mathbf{A} = \partial\boldsymbol{\beta}/\partial\boldsymbol{\theta}$, also called the *amplification matrix*. By writing the amplification matrix in terms of deflection angle $\boldsymbol{\alpha}$, we get

$$A_{ij} = \frac{\partial\beta_i}{\partial\theta_j} = \delta_{ij} - \frac{\partial\alpha_i}{\partial\theta_j} \quad (2.32)$$

$$= \delta_{ij} - \frac{2}{c^2} \int_0^{\chi} d\chi' \frac{f_K(\chi - \chi') f_K(\chi')}{f_K(\chi)} \frac{\partial^2}{\partial\theta_i \partial\theta_j} \Phi(f_K(\chi')\boldsymbol{\theta}, \chi'). \quad (2.33)$$

Inverting the derivative and the integral, we can write the amplification matrix coefficient in terms of a new 2-dimensional potential ψ called the *lensing potential*

$$A_{ij} = \delta_{ij} - \partial_i \partial_j \psi, \quad (2.34)$$

where

$$\psi(\boldsymbol{\theta}, \chi) = \frac{2}{c^2} \int_0^{\chi} d\chi' \frac{f_K(\chi - \chi') f_K(\chi')}{f_K(\chi)} \Phi(f_K(\chi')\boldsymbol{\theta}, \chi'). \quad (2.35)$$

By separating the amplification matrix can be reparametrized as

$$A = \begin{pmatrix} 1 - \kappa & 0 \\ 0 & 1 - \kappa \end{pmatrix} + \begin{pmatrix} -\gamma_1 & -\gamma_2 \\ -\gamma_2 & \gamma_1 \end{pmatrix} \quad (2.36)$$

where we introduced a scalar field κ called the *convergence* and the 2-spin fields γ_1 and γ_2 called the *cosmic shear*. It is clear by identification to [Equation 2.34](#) that the expression of these fields with respect to the lensing potential is

$$\kappa = \frac{1}{2}(\partial_1 \partial_1 + \partial_2 \partial_2)\psi = \frac{1}{2}\nabla^2\psi, \quad (2.37)$$

$$\gamma_1 = \frac{1}{2}(\partial_1 \partial_1 - \partial_2 \partial_2)\psi, \quad (2.38)$$

$$\gamma_2 = \partial_1 \partial_2 \psi, \quad (2.39)$$

It is much more convenient to work with the Fourier space representation of the convergence and shear fields as the partial derivations of the lensing potential transform to multiplications

$$\tilde{\kappa} = \frac{1}{2}(k_1^2 + k_2^2)\tilde{\psi} \quad (2.40)$$

$$\tilde{\gamma}_1 = \frac{1}{2}(k_1^2 - k_2^2)\tilde{\psi}, \quad (2.41)$$

$$\tilde{\gamma}_2 = k_1 k_2 \tilde{\psi}, \quad (2.42)$$

And therefore deriving a linear mapping between the shear and convergence in Fourier space, which is the base of the Kaiser-Squires transformation presented in Kaiser et al. [[KS93](#)] and further studied in [chapter 6](#)

$$\tilde{\gamma}_1 + i\tilde{\gamma}_2 = \left(\frac{k_1^2 - k_2^2}{k^2} + i \frac{2k_1 k_2}{k^2} \right) \tilde{\kappa}, \quad (2.43)$$

where $k^2 = k_1^2 + k_2^2$. Note that the solution is not defined for $k = 0$, which means that the mean of the convergence field cannot be directly constrained from shear, which is usually known as the mass-sheet degeneracy.

It is important to notice that in practice, we cannot distinguish between the convergence field, which applies a magnification, and the cosmic shear, which applies a distortion to background objects. The amplification matrix can be factorized by the $1 - \kappa$ term to give

$$A = (1 - \kappa) \left(\mathbb{I}_2 - \begin{pmatrix} g_1 & g_2 \\ g_2 & -g_1 \end{pmatrix} \right), \quad (2.44)$$

where \mathbb{I}_2 is the identity matrix and where we introduced the reduced shear $g = \frac{\gamma}{1 - \kappa}$, which also has a spin-two properties similarly to the shear. However, in the weak lensing regime, the convergence κ is very small, i.e. $\kappa \ll 1$, which translates into a well motivated approximation of the shear γ to the measured reduced shear g (i.e. $\gamma \approx g$).

Additionally, the convergence can be interpreted as the 2-dimensional projection of the 3-dimensional matter density field over the line of sight. Indeed, by combining the lensing potential definition in Equation 2.35, and its second-order derivative to defining the convergence in Equation 2.37, we get

$$\kappa(\boldsymbol{\theta}, \chi) = \frac{1}{2} \nabla^2 \Psi \quad (2.45)$$

$$= \frac{1}{c^2} \int_0^\chi d\chi' \frac{f_K(\chi - \chi') f_K(\chi')}{f_K(\chi)} \nabla^2 \Phi(f_K(\chi') \boldsymbol{\theta}, \chi'). \quad (2.46)$$

In the latest equation, the laplacian ∇^2 is taken with respect to the angular coordinates, and can be transformed to a 3D Laplacian in comoving coordinates. The transformation generates an additional term which is a second-order derivative along the comoving coordinate $\partial^2/\partial\chi^2$, which vanishes by integration along the line of sight by assuming that the sum of positive and negative contributions of the potential gives zero. Using the Poisson equation (Equation 2.23), we can express the Laplacian of the gravitational potential in terms of density contrast

$$\kappa(\boldsymbol{\theta}, \chi) = \frac{3H_0^2 \Omega_m}{2c^2} \int_0^\chi \frac{d\chi'}{a(\chi')} \frac{f_K(\chi - \chi') f_K(\chi')}{f_K(\chi)} \delta(f_K(\chi') \boldsymbol{\theta}, \chi'), \quad (2.47)$$

assuming that $\bar{\rho} \propto a^{-3}$. Considering a population of sources of density $n(\chi)$, the mean convergence can be computed by weighting the latest equation such that

$$\kappa(\boldsymbol{\theta}) = \int_0^{\chi_{\text{lim}}} d\chi n(\chi) \kappa(\boldsymbol{\theta}, \chi). \quad (2.48)$$

This gives the following expression for the convergence map

$$\kappa(\boldsymbol{\theta}) = \frac{3H_0^2 \Omega_m}{2c^2} \int_0^{\chi_{\text{lim}}} \frac{d\chi'}{a(\chi')} q(\chi) f_K(\chi') \delta(f_K(\chi') \boldsymbol{\theta}, \chi'), \quad (2.49)$$

where we define the lensing efficiency q such that

$$q(\chi) = \int_\chi^{\chi_{\text{lim}}} d\chi' n(\chi') \frac{f_K(\chi - \chi')}{f_K(\chi)}. \quad (2.50)$$

In this chapter, we have therefore presented the cosmological background that is useful for the understanding of this thesis. A primary aspect for the analysis the weak gravitational lensing probe is the measurement of the cosmic shear. The next chapter is therefore dedicated to the presentation of standard methods for these measurement from galaxy optical images.

* * *
* * *
* * *

Cosmic shear measurement methods

Chapter Outline

3.1	Weak lensing galaxy images	18
3.1.1	Galaxy morphologies	18
3.1.2	Shearing transform	19
3.1.3	Point spread function	20
3.1.4	Image pixelization	22
3.1.5	Noise corruption sources	22
3.1.6	From patch to postage stamps (detection and blending)	22
3.2	Shape measurement methods	23
3.2.1	(Weighed-) Moments	24
3.2.2	Model fitting	26
3.3	PSF correction methods	29
3.3.1	KSB	29
3.3.2	BJ02 and re-Gaussianization	29
3.4	Calibration methods	31
3.4.1	Response to shear	31
3.4.2	Selection response	32
3.4.3	METADETECTION with automatic differentiation	34
3.5	Conclusion	35

This chapter covers a project presented at a conference:

Vitorelli, A. Z., **Remy, B.**, Guinot, A., Lanusse, F., Martins, T., ‘AutoMetaCal: self-calibration of shear biases with automatic differentiation’. In: *Rencontres de Moriond, proceedings (2022)*

THIS chapter introduces the physical forward model generating observed optical images of galaxies and state-of-the-art methods to extract cosmic shear information. We start by discussing all the physical processes involved in the journey of a photon from its emission from a distant galaxy to the Charge-Coupled Device (CCD) of our instrument. This allows us to understand how the original image of a galaxy is

transformed and how to estimate cosmic shear by shape measurement with the least bias possible. We also introduce PSF correction and calibration methods to mitigate remaining biases in shear estimation.

3.1 Weak lensing galaxy images

In the following subsections, we describe how the continuous galaxy image generated by a distant galaxy is first sheared, then convolved by the instrument Point Spread Function (PSF), discretized on the CCD, and corrupted by noise [Figure 3.1](#).

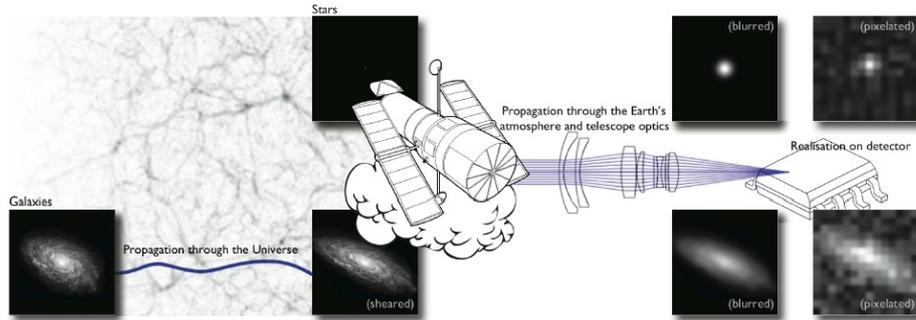


Figure 3.1: Journey of a photon from galaxy emission through the [LSS](#) of the Universe and the optical instrument ending up on the [CCD](#). Credit: Figure from Mandelbaum et al. [[Man+14b](#)], adapted from Kitching et al. [[KHM11](#)].

3.1.1 Galaxy morphologies

We have seen in the previous chapter that the [LSS](#) of the Universe acts as a weak gravitational lens. More specifically, we have seen that in this regime the lensed coordinates of an object are related to its true coordinates by a linear mapping involving the shear and convergence fields and that cosmological information can be extracted from them. However, to measure these fields, we need to analyze their effect on background objects and especially on galaxy shapes. As we will see in this chapter, the effect of cosmic shear on galaxy shapes can be seen as an alteration of their ellipticities. Therefore, we need to have a good knowledge of galaxy images before their light goes through the [LSS](#) of the Universe. It is often assumed in this context that a galaxy can be classified as elliptical or spheroid (bulge) plus disk systems [[VB10](#)]. Many analytical profiles are well suited to fit such galaxy morphologies, such as the de Vaucouleurs [[de 48](#)], or Sersic profiles [[Ser68](#)], even though they cannot necessarily describe them in detail [[VB10](#); [Mel+10](#)]. Mandelbaum et al. [[Man+14a](#)] identified that this task can be exceedingly challenging for modern surveys, mentioning that for instance, only 20% of galaxies observed by the [HST](#) can be perfectly fit by such models, while other 50% can be well modeled by them even though not capturing the fine details, and the remaining 30% are completely irregular and cannot be fit by such models. [Figure 3.2](#) illustrates this challenge with real galaxy images from the [HST](#) observed by the Advanced Camera for Surveys ([ACS](#)) in the [COSMOS](#) survey. This issue is particularly prevalent in space-based surveys like [HST](#), [Euclid](#), and the Roman space telescope. These surveys will uncover the intricate substructure of galaxies. Meanwhile, ground-based surveys like Large Synoptic Survey Telescope ([LSST](#)) will

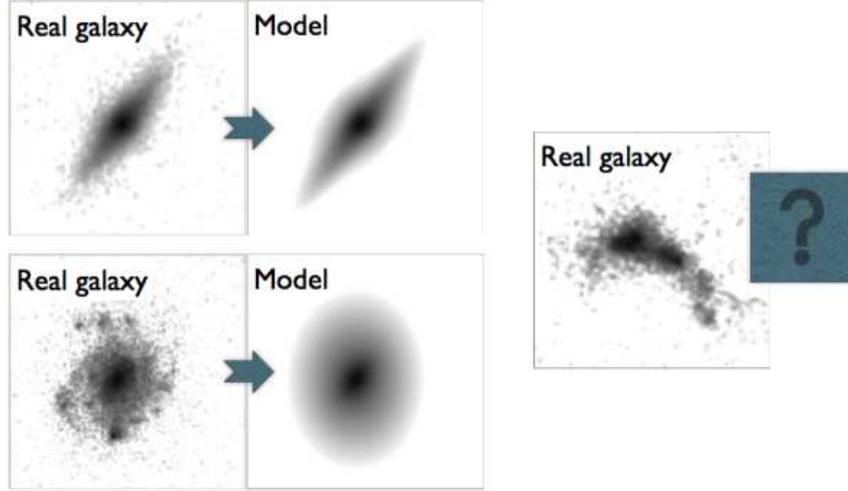


Figure 3.2: Illustration of three different types of real galaxies from *HST* observed by the Advanced Camera for Surveys (*ACS*) in the *COSMOS* survey. The *top left* image shows a galaxy which corresponds almost perfectly to a simple parametric profile. The *bottom left* image shows a galaxy which can also be fit by a parametric profile, but missing the finer details of its substructure. The image on the *right*, however, is irregular and cannot be fit by a simple parametric model. The presence of these irregular galaxy morphologies makes the measurement of cosmic shear highly challenging. Credit Mandelbaum et al. [Man+14a].

observe fainter galaxies that appear as elliptical objects, which our parametric models can adequately characterize.

3.1.2 Shearing transform

The light emitted by a distant galaxy is lensed by the *LSS* of the Universe, which results in a distorted version of the original image. According to Equation 2.44, we can define the shearing operator $s_{g_1, g_2}[\cdot]$ of reduced shear components g_1 and g_2 as a linear mapping between the coordinates of the original image \mathcal{I} and the lensed image as

$$s_{(g_1, g_2)}[\mathcal{I}(\mathbf{x})] = \mathcal{I}(T(g_1, g_2)\mathbf{x}), \quad (3.1)$$

where \mathbf{x} are the image coordinates and $T(g_1, g_2)$ is the shearing matrix defined as

$$T(g_1, g_2) = \frac{1}{\sqrt{1 - g_1^2 - g_2^2}} \begin{pmatrix} 1 + g_1 & g_2 \\ g_2 & 1 - g_1 \end{pmatrix}, \quad (3.2)$$

which we normalized with the term $\sqrt{1 - g_1^2 - g_2^2}$ so that it conserves the total area of the image. Similarly, we can define the shearing transform in Fourier space:

$$s_{(g_1, g_2)}[\tilde{\mathcal{I}}(\mathbf{k})] = \tilde{\mathcal{I}}(T(g_1, g_2)\mathbf{k}). \quad (3.3)$$

The shearing of an image is a distortion of the grid \mathcal{G} , on which is sampled the image. This therefore requires to sample the image on a new coordinate grid given by the transformation. The new coordinate grid $\{T(g_1, g_1)\mathbf{x}\}_{\mathbf{x} \in \mathcal{G}}$ does not necessarily

necessarily coincide with the former $\{\mathbf{x}\}_{\mathbf{x} \in \mathcal{G}}$ and the values of the new image need to be interpolated. The interpolation of an image \mathcal{I}_1 of size $N \times N$ with an interpolation kernel K is defined as:

$$\mathcal{I}_2(x, y) = \sum_{i, j=-N/2}^{N/2-1} \mathcal{I}_1(i, j)K(x - i, y - j). \quad (3.4)$$

Bernstein et al. [BG14] presents multiple interpolation kernels and their relative efficiency both in real and Fourier space

- *Linear* interpolation kernel

$$K(x) = \begin{cases} 1 - |x| & |x| \leq 1 \\ 0 & \text{otherwise} \end{cases} \quad (3.5)$$

- *Cubic* interpolation kernel

$$K(x) = \begin{cases} \frac{3}{2}|x^3| - \frac{5}{2}x^2 + 1 & |x| \leq 1 \\ -\frac{1}{2}|x^3| + \frac{5}{2}x^2 - 4|x| + 2 & 1 \leq |x| \leq 2 \\ 0 & |x| \geq 2 \end{cases} \quad (3.6)$$

- *Quintic* interpolation kernel

$$K(x) = \begin{cases} 1 + \frac{x^3}{12}(-95 + 138x - 55x^2) & |x| \leq 1 \\ \frac{(x-1)(x-2)}{24}(-138 + 348x - 249x^2 + 55x^3) & 1 \leq |x| \leq 2 \\ \frac{(x-2)(x-3)^2}{24}(-54 + 50x - 11x^2) & 2 \leq |x| \leq 3 \\ 0 & |x| \geq 3 \end{cases} \quad (3.7)$$

- *Lanczos* interpolation kernel at order m

$$K(x) = \text{sinc}(x)\text{sinc}(x/m)\Pi(x/2m), \quad (3.8)$$

where

$$\Pi(x) = \begin{cases} 1 & |x| < 0.5 \\ 0.5 & |x| = 0.5 \\ 0 & |x| > 0.5 \end{cases} \quad (3.9)$$

The final recommendation of Bernstein et al. [BG14] is to perform the shearing and convolution with the **PSF** in Fourier space after zero padding and discrete Fourier transform. Zero padding seems to be mandatory to prevent *ghost* appearance, resulting from the folding of periodical repetition of the discrete Fourier transform, which biases the quadrupole computation for shape measurement.

3.1.3 Point spread function

The Point Spread Function (**PSF**) definition. This effect can be mathematically described as a pixel-wise deconvolution between the galaxy image and the **PSF** image, called kernel. As this kernel varies spatially, temporally, and potentially spectrally, we refer as the **PSF** field. Several sources can generate a **PSF** effect on the image when

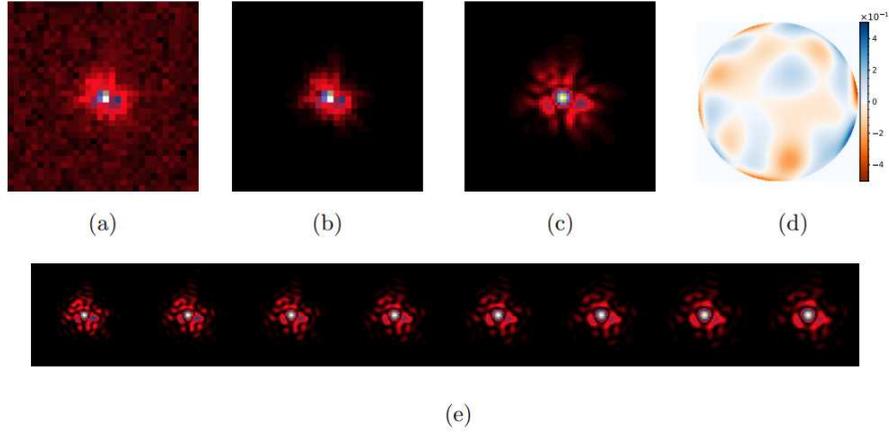


Figure 3.3: The figure shows different representations of a space-like PSF at a single position in the field of view estimated from a star. (a) Noisy polychromatic PSF at observation resolution, (b) Noiseless polychromatic observation at observation resolution, (c) High-resolution noiseless polychromatic 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. Credit: Credit: Liaudat et al. [Lia+23]

imaging the sky. For space-based telescopes, the optical system has many contributions to the PSF. First, the lens and aperture introduce a diffraction of the light coming from every direction of the sky and mostly dictate the width of the PSF. As the effect of the PSF is a convolution of the galaxy image, the width of the PSF will dictate the maximal resolution available from the survey. Then, there are several other contributions at the optic level such as optical aberrations, polishing effects (smoothness of the mirrors), external sources of obscurations, scattering light, ice contamination, light polarization, or thermal variation. Additionally, some optical components may induce a chromatic variation of the PSF, e.g. the dichroic filter initially designed to reflect some wavelengths and to be transparent for others used in the *Euclid*'s VIS Instrument (VIS) and Near Infrared Spectro Photometer (NISIP) instruments of the Euclid survey. For ground-based surveys, the atmosphere has an enormous contribution to the PSF field. As the atmosphere varies with time, this generates a temporal variation in addition to the spatial variation of the PSF field. In this thesis, we will focus on space-based optical measurement so we can neglect the temporal variation of the PSF.

Because the effect of the PSF on the galaxy image is non-negligible and affects its shape, it is crucial to have a good model. To model the PSF field, we use the fact that the light profile of stars can be considered as spatial impulse, i.e. $\mathcal{I}_{\text{star}} = \delta(x, y, \lambda) = f_{(x,y)}(\lambda)$. The variation with respect to the wavelength can be neglected if the instrument does not generate chromatic dependence, e.g. for the COSMOS survey considered later in this thesis. For a survey like Euclid, we would need complementary observation informing about the spectral energy distributions of both stars and galaxies. The convolution of the PSF with such spacial impulses therefore provides the exact PSF image at the star positions. This results as a PSF field evaluated at all star positions, which can then be interpolated at galaxy positions for a full PSF field. For a comprehensive study on how to model the PSF for space-based

and ground-based survey, we refer the read to Liaudat [Lia22].

3.1.4 Image pixelization

The light emitted by galaxies going through the LSS of the universe and the optical instrument can be considered as a continuous electromagnetic wave. For optical survey, the light finally encounters the CCD, discretizing the galaxy image. This discretization of the continuous galaxy image is the result of the integration of multiple photons falling into the same pixel. Therefore, the final observed image is not a *sampling* of the continuous image at the pixel coordinate, but the integration of all the photons in the pixel area. This can be seen as a convolution between the sampling of the center pixel coordinate and the area of the pixel, also known as the *pixel response*. It also happens that the pixel grid does not respect the Nyquist-Shannon theorem. Indeed, the optics of the telescope, and especially the diameter of the aperture, dictates the maximum frequency that we could possibly measure with our detector. However, when the pixel size, dictating the maximum wavelength measurable by the detector does not verify the Nyquist-Shannon theorem, then the image is *undersampled*. It is particularly the case for the *Euclid* survey.

3.1.5 Noise corruption sources

The image taken by the CCD is corrupted by various sources of noise, which appear as random fluctuation in the final image. The level of noise is roughly the same for the pointing directions of the survey, therefore the galaxy magnitude will set the Signal-to-Noise Ratio (Signal-to-Noise Ratio (SNR)). The CCD receives light not only from the object of interest like stars and galaxies, but also from every direction of the sky and non-resolved objects. This is called *photon noise* or *shot noise* and is primary in the image noise contamination. The Poisson law is often associated to this noise source as it measures the probability of photons arriving to the CCD in a fixed interval of time. Non-astrophysical sources of noise also affect the CCD. Indeed, the thermal variation of the CCD structure generates electrons, independently from any signal received, and would be observable in complete darkness. This source of noise is therefore called *thermal noise* [Nyq28] or *Dark-current shot noise* [Bae06]. Lastly, the physics involved in the CCD electronics, such as the analog-to-digital conversion, and amplification, may generate noise in the produced image called *read noise* [BTM04].

3.1.6 From patch to postage stamps (detection and blending)

Weak lensing surveys are very wide and require the procedure of source detection to be automated. SEXTRACTOR [BA96] was proposed for this purpose and has been widely used for the detection of stars and galaxies. It is today the software of reference for detecting sources and is for instance used in the SHAPEPIPE [Gui+22; Far+22]. After an estimation of the sky background, sources are detected using a thresholding algorithm. As opposed to stars, galaxies are extended objects, which require a small threshold for detection. This task is agnostic to the type of object detected, so is followed by a star-galaxy separation algorithm. On the one hand, it is essential to have a pure sample of stars for the estimation of a precise PSF model. On the other hand, the galaxy sample also needs to be pure otherwise the shape measurement estimator will be biased as it would measure the shape of non-elliptical objects. To that end, a classifier is usually used and neural networks are the state-of-the-art algorithm for classification tasks.

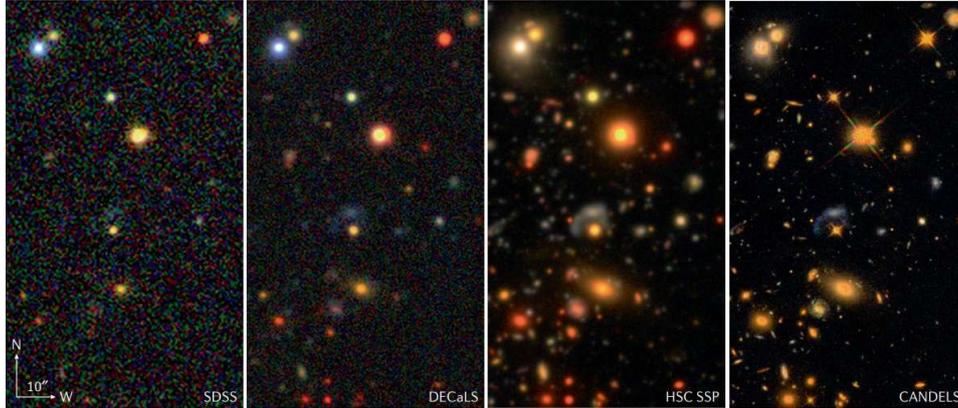


Figure 3.4: Different survey observing the same region of the sky. From left to right the Sloan Digital Sky Survey (Sloan Digital Sky Survey (**SDSS**), 0.396 arcsec pixel scale), the DECam Legacy Survey (DECaLS, 0.262 arcsec pixel scale), the Hyper Suprime-Cam Subaru Strategic Program (Hyper Suprime-Cam (**HSC**) SPP, 0.168 arcsec) and the Cosmic Assembly Near-infrared Deep Extragalactic Legacy Survey (Cosmic Assembly Near-infrared Deep Extragalactic Legacy Survey (**CANDELS**)) programme with the Hubble Space Telescope (**HST**, 0.06 arcsec pixel scale). This illustrates how increasing the resolution of a survey with the same depth helps to distinguishing objects. Credit Melchior et al. [Mel+21].

An additional source of bias in galaxy survey is the *blending* of sources. It occurs when multiple objects overlap in the same region of the sky. Such blend systems can be observed in Figure 3.4. A classical detection algorithm would be likely to detect only one instead of two isolated sources and this error can propagate bias to the shape measurement and then to the cosmological analysis. Melchior et al. [Mel+21] reviews the different challenges and potential solutions to this issue. Modern surveys like the Dark Energy Survey, **LSST**, or Euclid are designed to maximize the number of detected galaxies, with a *deep* observation strategy, which therefore increases the probability of observing overlapping objects. Dawson et al. [Daw+15] estimates that about 14% of the population of detected galaxies for the **LSST** survey will be blended, resulting in a significant systematic bias for shear estimation. In fact, this figure is way below the current estimates. **HSC** for instance sees about 60% of blended objects Bosch et al. [Bos+18]. Several solutions exist to tackle the blending bias. Firstly, blended sources can be discarded from the sample at the cost of reducing the statistical power of the analysis. Secondly, a deblending effort can result in multiple images of isolated sources. This is a very hard challenge as retrieving images of the isolated objects is an ill-posed inverse problem, meaning several solutions can fit the observations. There are several attempts to tackle this problem, for instance, the SCARLET algorithm proposed in Melchior et al. [Mel+18] estimates the maximum a posteriori solution under non-negative matrix factorization constraints. Liu et al. [LMR23] and Hansen et al. [Han+22] with the BLISS method propose to model the full probabilistic catalog of objects using generative models and variational inference techniques.

3.2 Shape measurement methods

Classical methods for cosmic shear estimation are based on the measurement of galaxy shapes, or ellipticities as they serve as a proxy for the shear. Indeed, we have seen in

subsection 3.1.2 that the shear applies a linear transformation to the coordinate grid of galaxy images. This results in an observed galaxy image whose ellipticity e_{obs} is the sum of the ellipticity of the original galaxy image e_{orig} and of the cosmic shear γ , i.e. $e_{\text{obs}} = e_{\text{orig}} + \gamma$. In the following, we present how to measure galaxy shapes by computing the moments of images or by using model-fitting methods.

3.2.1 (Weighted-) Moments

Weighted moments of an image of intensity I are given by the following equations:

$$F = \sum w(x, y)I(x, y) \quad (3.10)$$

$$x_0 = \sum w(x, y)xI(x, y) \quad (3.11)$$

$$\sigma^2 = \sum w(x, y)^2 \sigma^2(I(x, y)) \quad (3.12)$$

$$T = \sum w(x, y)I(x, y)(x^2 + y^2) \quad (3.13)$$

$$M_1 = \sum w(x, y)I(x, y)(x^2 - y^2) \quad (3.14)$$

$$M_2 = \sum w(x, y)I(x, y)2xy \quad (3.15)$$

$$e_1 = M_1/T \quad (3.16)$$

$$e_2 = M_2/T \quad (3.17)$$

$$S/N = F/\sigma(F), \quad (3.18)$$

where $w : \mathcal{X} \rightarrow \mathbb{R}$ is a window function. It is simpler to compute the moments of an image without applying the weights as it does not require choosing an adapted window function, but it is much more sensitive to the image noise [KSB94, KSB]. We often use a Gaussian kernel for the window function w . Not only do these moments capture summaries of a given image, but they are also related to typical characteristics of the galaxy within. We detail below the physical understanding of galaxy image moments.

Flux Integrating all pixel intensity of an image gives the flux F of this image.

Size T can be seen as an estimate of this object size. However, we rarely rely on galaxy size in our analysis, but rather on the ratio between galaxy and PSF size, often referred as T/T_{PSF} .

Ellipticity The ellipticity components e_1 and e_2 of a galaxy are of particular interest for cosmological analysis as they serve as an estimator for the cosmic shear. We also often refer to this summary as the galaxy shape. Assuming that the galaxy has an elliptical profile, meaning that lines of constant galaxy profile intensity (isophotes), i.e. $\{(x, y)\}$ s.t. $I(x, y) = \text{constant}$ follows the equation of an ellipse:

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} = \text{constant}, \quad (3.19)$$

where $0 \leq b, 0 \leq a$. Assuming $b \leq a$, a and b are called the major and minor axis of the ellipse. It is not necessarily the case for an arbitrary choice of coordinate, but it does not matter as we will work with absolute differences in the following. As the galaxy will likely not be exactly aligned with the coordinates grid, we need to introduce

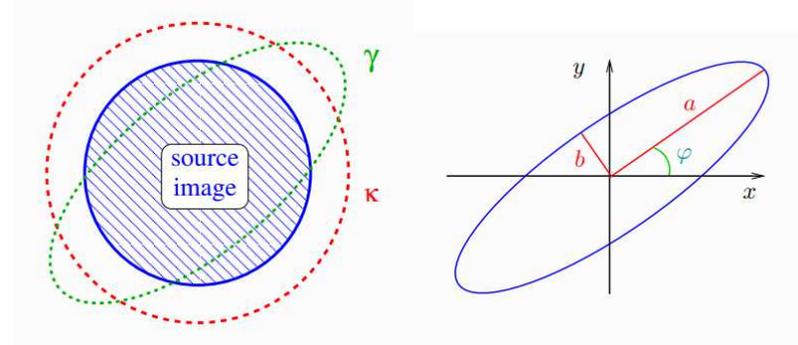


Figure 3.5: The shear contributes to an anisotropic stretch of the source, transforming a circle to an ellipse, while the convergence κ contributes to an isotropic deformation (or magnification). Credit: Martin Kilbinger.

its orientation β and begin the angle between the major axis and the coordinates grid x -axis. See Figure 3.5 for an illustration. There are multiple definitions of the galaxy shape. In this thesis, we will only refer to the *ellipticity* (or *distortion*) e as:

$$e = |e|e^{2i\beta} \quad \text{with} \quad |e| = \frac{|a^2 - b^2|}{a^2 + b^2}, \quad (3.20)$$

or its *reduced* version

$$g = |g|e^{2i\beta} \quad \text{with} \quad |g| = \frac{|a - b|}{a + b}. \quad (3.21)$$

Note the factor 2 in the phases of those complex values. It is due to the fact that rotating an ellipse by 180 degrees yields the same ellipse, $e = |e|e^{2i\beta} = |e|e^{2i(\beta+\pi)}$.

We can therefore relate the ellipticity to the shapes measured by the weighted moment as

$$e_1 = |e| \cos(2\beta), \quad (3.22)$$

$$e_2 = |e| \sin(2\beta). \quad (3.23)$$

The ellipticity can be also defined in terms of tangential components and cross-component with respect to a given angle ϕ :

$$\gamma_t = \text{Re}(\gamma e^{-2i\phi}) \quad (3.24)$$

$$\gamma_\times = \text{Im}(\gamma e^{-2i\phi}) \quad (3.25)$$

This terminology is particularly useful when working on galaxy lensing, i.e. when we study the correlation between the galaxy shapes and the position of a particular object, such as galaxy positions at lower redshifts (galaxy-galaxy lensing).

Signal-to-Noise ratio (SNR) The SNR quantifies how much of a galaxy light profile's signal is perceived in the image compared to the background noise. In Equation 3.18, it is defined as the ratio between the galaxy flux and the noise standard deviation, but there exist several other definitions in the literature.

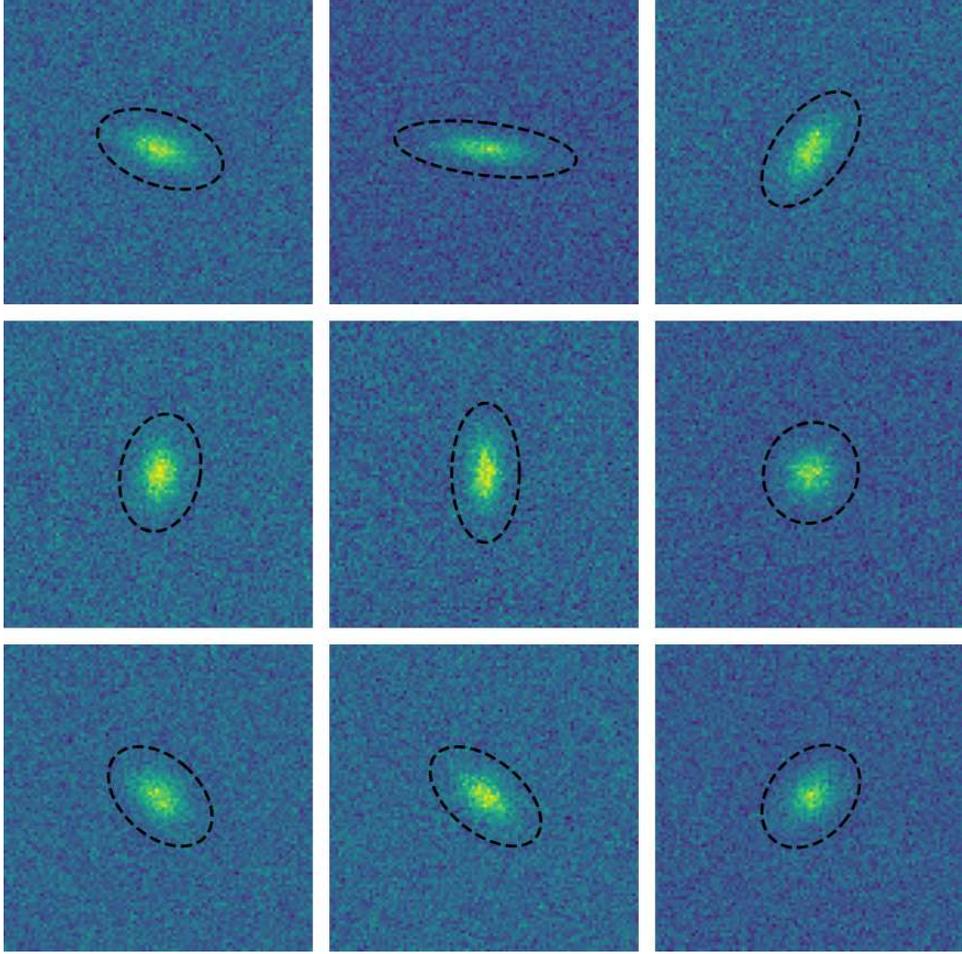


Figure 3.6: Galaxy ellipticities can be estimated by fitting a parametric model, here a Gaussian light profile. The covariance matrix contains information on the major-to-minor axis ratio and galaxy orientation. It is therefore possible to extract the information of the galaxy shape and to draw the associated ellipse (here in dotted black).

3.2.2 Model fitting

Another way of extracting physical quantities from a galaxy image is to fit a parametric model to the light profile. In this section, we discuss how model fitting techniques work and how to extract elliptical components from fitted profiles. We first introduce models which present isotropic light profiles (Exponential, Gaussian, and Sersic models). These models will be able to fit the size and the flux of the profile. The generated image is then sheared to model the orientation and elongation of the profile. Note that some models can directly incorporate the ellipticity in their parametrization, such as the Gaussian model with its covariance matrix. The latter will be particularly useful for characterizing the ellipticity of galaxies and PSF in the later subsections. Anisotropic models are parametrized with the two-dimensional vector $\mathbf{x} = [x, y]$ and for isotropic models, we will distance modulus to the center $r = \sqrt{x^2 + y^2}$.

Exponential profile

$$I(r) \propto e^{-\frac{r}{r_0}}, \quad (3.26)$$

where r_0 is called the scale radius.

Isotropic Gaussian profile The 2-dimensional isotropic Gaussian profile is given by this equation:

$$I(r) = F \times \frac{1}{2\pi\sigma^2} e^{-\frac{r^2}{2\sigma^2}}, \quad (3.27)$$

where σ is the Gaussian standard deviation of the light profile and F is the total flux.

Anisotropic Gaussian profile Another way to parametrize a Gaussian light profile and directly modeling the galaxy ellipticity is to take into account the full covariance matrix.

$$I(x, y) = F \times \frac{1}{2\pi|\Sigma|^{\frac{1}{2}}} e^{-(\mathbf{x}-\mathbf{x}_0)^T \Sigma (\mathbf{x}-\mathbf{x}_0)}, \quad (3.28)$$

where $\mathbf{x} = [x, y]$ is the coordinate grid, $\mathbf{x}_0 = [x_0, y_0]$ is the galaxy center offset and F is the total flux. This way we directly model the anisotropy of the profile through the covariance matrix Σ . The ellipticity parameters can be extracted from the eigenvalues of the covariance matrix:

Sersic profile

$$I(r) = F \times I_0 \exp\left(-b_n \left[\left(\frac{r}{r_{\text{hlr}}}\right)^{\frac{1}{n}} - 1\right]\right), \quad (3.29)$$

where n is called the *Sersic index* and drives how much picky is the profile. For instance, with $n = 0.5$, we recover the Gaussian profile, with $n = 1$ the exponential profile. In [Figure 3.7](#), we illustrate that the higher n is, the more the intensity is concentrated in the center of the profile. r_{hlr} is the half-light-radius, I_0 is a normalization factor such that the integration of I over all the space is F , which corresponds to the flux of the galaxy. Finally, b_n is a dimensionless constant depending on n . Ciotti et al. [[CB99](#)] shows by integrating the luminosity inside a given radius that this parameter is the solution to the equation:

$$\gamma(2n, b_n) = \Gamma(2n)/2, \quad (3.30)$$

where $\gamma(\alpha, x) = \int_0^x e^{-t} t^{\alpha-1} dt$ is the (left) incomplete gamma function and $\Gamma(\alpha) = \gamma(\alpha, +\infty)$ is the complete gamma function. This equation cannot be solved analytically with a closed-form expression and therefore requires numerical approximation. One simple way to approximate b_n is to rely on the asymptotic expansion computed in Ciotti et al. [[CB99](#)].

$$b_n \approx 2n - \frac{1}{3} + \frac{4}{405n} + \frac{46}{25515n^2} + \frac{131}{1148175n^3} + \mathcal{O}(n^{-4}). \quad (3.31)$$

This method remains limited by the maximum polynomial order, here n^{-4} . Another way is to reformulate [Equation 3.30](#) as a root-finding problem (see [subsection 5.1.4](#)), and use a quasi-Newton method to find a solution. When targeting the root of this

problem, we used the Broyden method [Bro65] which approximates efficiently the Jacobian and its inverse.

Note that if we want to perform gradient-based inference with such a model, the computation of the model gradient with respect to the Sersic index n will be spatially computationally expensive because it will require computing gradient through all of the iterations necessary to solve Equation 3.30. In our inference pipelines, we used a machine learning trick based on *implicit optimization* enabling us to store the gradient with the space complexity of only one iteration. See subsection 5.1.4 for more details.

The choice of parametric light profile is often motivated by prior knowledge of the galaxy population of study. As illustrated in Figure 3.7 a Sersic model can capture both exponential and Gaussian profiles, but is much more computationally expensive to compute due to the Sersic index normalization factor. A motivated choice to work with exponential or Gaussian profile would therefore make the fitting much quicker.

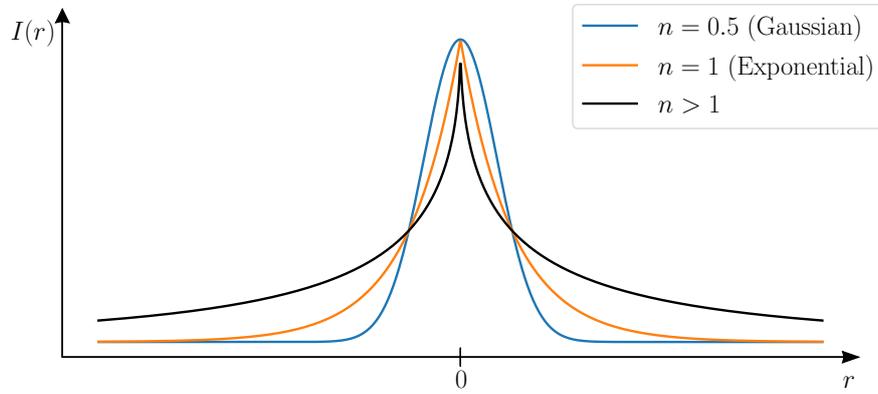


Figure 3.7: Illustration of one-dimensional projection of the Sersic light profile for different values of Sersic indices n . The higher n , the more the brightness is concentrated in the center of the galaxy profile.

Then, to take into account the ellipticity of the galaxy profile, the image is sheared as discussed in subsection 3.1.2, conserving the area of the profile. The additional parameters are then the minor-to-major axis ratio q and the position angle β .

χ^2 minimization This constitutes a model for galaxy light profile f_{model} which depends on a set of parameters $(\alpha_1, \dots, \alpha_n)$, such as the radius r_0 , the sersic index n , shear parameters q and β , etc.. This model can then be fitted to an observed image X by χ^2 minimization:

$$\chi^2 = \sum_{x=1}^{N_x} \sum_{y=1}^{N_y} \frac{(X(x, y) - f_{\text{model}}(x, y | \alpha_1, \dots, \alpha_n))^2}{\sigma(x, y)^2}, \quad (3.32)$$

where X is of size (N_x, N_y) and $\sigma(x, y)$ is a noise image. This is the purpose of the GALFIT software¹ [Pen+02; Pen+10], where the χ^2 objective in Equation 3.32 is minimized using the efficient Levenberg-Marquardt algorithm [Mar63].

¹<https://users.obs.carnegiescience.edu/peng/work/galfit/galfit.html>

3.3 PSF correction methods

The point spread function (PSF) isotropy is a major source of shear measurement bias, especially for methods where shapes are estimated from the observed image, i.e. when the galaxy image is convolved with the PSF image. The weak lensing literature contains several attempts to correct the weighted moments of galaxy images taking into account the bias due to the PSF ellipticity. As in this manuscript, we will only compare our contributions to the state-of-the-art PSF correction method Hirata Seljak and Mandelbaum (HSM), we only briefly describe the first method proposed for this task Kaiser-Squires-Broadhurst (KSB) [KSB95].

3.3.1 KSB

The KSB (named after the authors Kaiser et al. [KSB95]) and its improvements [LK97; Hoe+98, KSB+] are based on the assumption that the contribution of the PSF is highly anisotropic but small distortion convolved with a large circularly symmetric window, then corrected galaxy ellipticity can be expressed as

$$e_{\text{cor}} = e_{\text{obs}} - P_{\alpha\beta}^{\text{sm}} p_{\beta}, \quad (3.33)$$

where p is a vector measuring the PSF anisotropy and P^{sm} is the smear polarisability tensor given in Hoekstra et al. [Hoe+98]. These two terms are actually computed in several different ways in the literature, depending on their physical interpretation. We refer the reader to Heymans et al. [Hey+06] for more description on the method and its various versions.

3.3.2 BJo2 and re-Gaussianization

HSM (named after C. Hirata, U. Seljak and R. Mandelbaum, authors of Hirata et al. [HS03] and Mandelbaum et al. [Man+05]) is a well established method used for shear estimation in several surveys [Man+05; Man+18]. To measure the cosmic shear HSM first measures the weighted moments of postage stamps in order to estimate the galaxy shapes e_1 and e_2 and then apply a shape correction in order to account for the error due to PSF intrinsic ellipticity.

[HS03] consider alternative definitions of an object shape in an image I as

$$e_+^{(I)} = (M_{xx} - Myy)/T^{(I)} \quad (3.34)$$

$$e_{\times}^{(I)} = (2M_{xy})/T^{(I)} \quad (3.35)$$

$$T^{(I)} = (M_{xx} + Myy), \quad (3.36)$$

where $M_{i,j} = \sum u_i u_j w(x,y) I(x,y)$ and $u = [x, y]$. The shape of the PSF can be measured in the exact same way, i.e. computing the weighted moments of the PSF image, which we refer to g . The same window function is used to compute the weighted moments of the observed image and of the PSF image. The principal idea of [HS03] is that perform a first correction under the assumption that the profile of the PSF is Gaussian, and then perform a second correction to the non-Gaussianity of the PSF. Assuming that the PSF image g is Gaussian, we can describe the PSF image similarly to Equation 3.28, as:

$$\gamma(x, y) \approx G(x, y) = \frac{1}{2\pi(|\Sigma_G|)^{\frac{1}{2}}} e^{-\frac{1}{2}u^T(\Sigma^{(f)})^{-1}u}, \quad (3.37)$$

where again $u = [x, y]$ and σ_G is the covariance matrix of the Gaussian PSF.

[HS03, (HS03)] showed that for a Gaussian PSF, we can apply to prescription from [BJ02] to correct from the ellipticity of the PSF. The ellipticity of the processing galaxy f can be obtained from the measurement of the ellipticity of the galaxy image I after correction of the Gaussian PSF such as, assuming circular Gaussian:

$$e^{(f)} = e^{(I)}/R, \quad (3.38)$$

$$R = 1 - \frac{T_g(1 - \beta_{22}^{(g)})(1 + \beta_{22}^{(I)})}{T_I(1 + \beta_{22}^{(g)})(1 - \beta_{22}^{(I)}), \quad (3.39)$$

where β_{22} is a ratio of Laguerre inverse transform coefficient $\beta_{22} = b_{22}/b_{00}$ and

$$b_{pq}^{(I)} = \int_{\mathbb{R}^2} \psi_{pq}^*(x, y; \Sigma^{(I)}) I(x) dx dy, \quad (3.40)$$

$\Sigma^{(I)}$ being the covariance matrix of the Gaussian profile fitted to the object in image I and $\psi(x, y; \Sigma)$ the Laguerre basis functions. This correction proposed by Bernstein and Jarvis 2002 (BJ02) is perfect for Gaussian PSF, however as demonstrated in HS03, this correction can lead to significant errors in the measurement of the processing ellipticity for non-Gaussian PSF. HS03 proposed therefore to transform the image to remove the effect of the non-Gaussianity of the PSF, hence the *re-Gaussianization* method, and then apply Equation 3.38 to the transformed image.

In order to Gaussianize the PSF image let us first define the residual image to a the Gaussian profile in Equation 3.37 as $\epsilon(x, y) = g(x, y) - G(x, y)$. This way, the observed image is the convolution of the preseeing image f with the two-component of the PSF, i.e. $I = (G + \epsilon) \otimes f$, which leads to the regaussianized image

$$\tilde{I} = G \otimes f = I - \epsilon f \quad (3.41)$$

. In this equation, both I and ϵ are known, but not f as it is still the image of interest. However as noted by HS03, because the residual image ϵ is very small, this equation would still stand if we use a Gaussian approximation to f . Therefore, similarly to the PSF, the pressing image can be approximated with a Gaussian profile such that:

$$f(x, y) \approx \frac{1}{2\pi(|\Sigma(f)|)^{\frac{1}{2}}} e^{-\frac{1}{2}u^T(\Sigma(f))^{-1}u} \quad (3.42)$$

This way, in order to accurately compute the shape of the galaxy $e^{(f)}$ in an image I , one can first re-Gaussianize the image according to Equation 3.41, then computing the weighted moment on the resulting image \tilde{I} , and applying the correction to a Gaussian PSF prescribed by BJ02 in Equation 3.38.

Since this method was established in [HS03], it has been tested and used on real data. It was first used in [Man+05] to generate an accurate shape catalog and perform galaxy-galaxy lensing cosmological analysis and later in [Man+18] to perform weak lensing shear calibration for the HSC survey.

Finally, once the adapted corrections have been made to provide galaxy shape catalog, the cosmic shear $\hat{\gamma}$ can be estimated as the average of these shapes and one estimate of the error bars through the standard deviation of the measured shapes, divided by the square root of the number of galaxies.

$$\hat{\gamma}_{\text{HSM}} = \frac{1}{N} \sum_{i=1}^N e \quad (3.43)$$

$$\sigma(\hat{\gamma}_{\text{HSM}})^2 = \frac{1}{N} \sum_{i=1}^N (\hat{\gamma}_{\text{HSM}} - e)^2 \quad (3.44)$$

Later in this manuscript, we perform cosmic shear measurements on simulated [HST](#) images, and we refer to this method as [HSM](#) [[HS03](#); [Man+05](#), (for Hirata, Seljak and Mandelbaum)] and [Equation 3.43](#) and [Equation 3.44](#) are used to provide the estimations along their uncertainties.

3.4 Calibration methods

To quantify the error of the measurement method, we can relate the observed shear \hat{g} to the ground truth shear g with a first-order Taylor expansion [[Hey+06](#); [Hut+06](#); [Mas+07b](#)].

$$\hat{g}_i - g_i = m_i g_i + c_i, \quad (3.45)$$

where $i \in \{1, 2\}$ denotes the shear component and m_i and c_i are respectively called the multiplicative and additive biases. Very often, we consider that the additive bias is due to the [PSF](#) ellipticity and is therefore handled by the correction described in the former section. However, there almost always remains a multiplicative bias m after correction because real galaxies cannot be perfectly described with elliptical profiles. The multiplicative bias can be computed by simulating the observations with a given shear. As we have access to the ground truth shear and galaxy shapes in the simulations, it is straightforward to compute the difference between the expectation and the estimation from our shape measurement algorithm. This however requires very precise simulations of the sky which is an arduous task to perform for modern surveys. For instance, such multiplicative bias calibration was used in [[Man+18](#)] to calibrate the weak lensing measurements for the [HSC](#) survey.

[[HM17](#); [SH17](#)] proposed to perform this calibration directly from data, characterizing how the shape measurement method respond to the addition of a very small shear. Indeed, assuming that the additive bias has been corrected, we can write the estimated shape of a galaxy a Taylor expansion around the true shape such as:

$$e = e|_{\gamma=0} + \left. \frac{\partial e}{\partial \gamma} \right|_{\gamma=0} \gamma + o(\gamma) \quad (3.46)$$

Assuming that there is no preferential orientation of the galaxies in the sky, this translates into a zero average if the true galaxy shapes, i.e. $\langle e \rangle|_{\gamma=0} = 0$. Therefore, [Equation 3.46](#) is particularly useful because we can recover [Equation 3.45](#) by averaging the measured shapes e and identify the multiplicative bias m to the response matrix $R = \left. \frac{\partial e}{\partial \gamma} \right|_{\gamma=0}$.

3.4.1 Response to shear

Estimating the multiplicative bias therefore only necessitates computing the response matrix. This task is however challenging as this requires simulating a very small shear to the pre-seeing galaxy image in order to compute the derivative. In [[HM17](#); [SH17](#)],

one can find an extensive description of the processing of the galaxy images to simulate observations with such a small shearing transform. The response matrix can finally be computed for each image using finite differences:

$$R_{ij} = \frac{e_i^+ - e_i^-}{\Delta\gamma_j}, \quad (3.47)$$

where $\Delta\gamma_j = \gamma_j^+ - \gamma_j^-$, with $\gamma_j^- = -\gamma_j^+$, e_i^+ is the measured shape component i on the image were we simulated a positive shear γ_j^+ , and respectively e_i^- a negative shear.

In order to measure e^+ and e^- , we need to simulate a small shear in the observed galaxies. As the shear is a transformation applied on the pre-seeing galaxy image, and not on the convolved observed image, we need to be careful in the processing of these images. [HM17] refer to the image processed by the metacalibration pipeline as *counterfactual images*. An image with a small shear γ is built as follows:

$$I'_\gamma = \Gamma \otimes [S(\gamma) ((\Pi^{-1} \otimes I))], \quad (3.48)$$

where $S(\gamma)$ is the shearing operator, Π the PSF image, Γ another PSF that we will discuss below, and I the observed image. The observed postage stamps are first deconvolved with all the observational responses (such as the PSF, pixel response, etc.), then a small shear is applied to the deconvolved image and the resulting image is convolved again with a PSF Γ . As the convolution of images is computationally expensive, we perform this operation in Fourier space instead, resulting in elementwise multiplications:

$$I' = \tilde{\Pi} \cdot S(\gamma) (\tilde{I}/\tilde{\Pi}) \quad (3.49)$$

There is no need for the reconvolution PSF Γ to be the same as the PSF of the instrument Π , but there are some constraints. First, deconvolving the observed image by the PSF is performed in Fourier space, which involves dividing the image spectrum by the PSF image spectrum. At high frequencies, the PSF spectrum is very small or zero which leads to a deconvolved image with very high or infinite power due to the high-frequency noise amplification. We therefore need to choose a PSF Γ which will cancel this noise amplification. [HM17] proposed to use a dilated version of the instrument PSF:

$$\Gamma(x) = \Pi((1 + 2|\gamma|)x), \quad (3.50)$$

for $\|\tilde{\Pi}(k)\|$ monotonically decreasing with k . Once the counterfactual images are built, one can use the shape measurement algorithms to compute e^+ and e^- . Note that to mitigate the bias due to the ellipticity shape, we could also use an isotropic PSF for reconvolution. One choice would be to sum the four 90 deg rotations of the PSF Π together. For image deconvolution, however, it is mandatory to have a good model for Π .

3.4.2 Selection response

We have so far defined the response matrix for an individual galaxy, but the computation of the response matrix is too noisy to be applied to individual shape measurement. The final shear is estimated by averaging individual shape measurement, so in a similar spirit, we can average the quantities of Equation 3.46:

$$\langle e \rangle = \langle e|_{\gamma=0} \rangle + \left\langle \frac{\partial e}{\partial \gamma} \Big|_{\gamma=0} \gamma \right\rangle + \mathcal{O}(\gamma^2) \quad (3.51)$$

$$\approx 0 + \left\langle \frac{\partial e}{\partial \gamma} \Big|_{\gamma=0} \right\rangle \langle \gamma \rangle + \mathcal{O}(\gamma^2) \quad (3.52)$$

where $\langle \cdot \rangle$ refers to the averaging of quantities over a selected catalog of galaxies. This leads to the new estimator for the cosmic shear

$$\langle \gamma \rangle = \langle R^{-1} \rangle \langle e \rangle, \quad (3.53)$$

where $\langle R^{-1} \rangle$ is the averaging of all the response matrices computed on individual galaxy images.

In addition to the response of the shape measurement algorithm responsivity to the shear, Sheldon et al. [SH17] noted that the detection of sources also responds to the shear and as a consequence the catalog on which is performed as the average in Equation 3.53. Figure 3.8 illustrates that due to PSF convolution, two close sources can be detected as a single source by a detection algorithm based on surface brightness thresholding. However, if shear is applied before the convolution with the PSF, as it is in the observed sky, the detection algorithm can identify the two distinct sources. In order to mitigate the detection bias, Sheldon et al. [SH17] proposed to invert the averaging and the derivative of the response matrix:

$$\langle R \rangle = \frac{\partial \langle e \rangle}{\partial \gamma} \Big|_{\gamma=0}; \quad \langle R_{ij} \rangle = \frac{\langle e_i^+ \rangle - \langle e_i^- \rangle}{\Delta \gamma_j}. \quad (3.54)$$

This metacalibration method handling the detection bias was dubbed *metacalibration*. It is intended to be used for estimating galaxy shape catalog for the Vera C. Rubin Observatory LSST observations, as tested on simulations in [She+23].

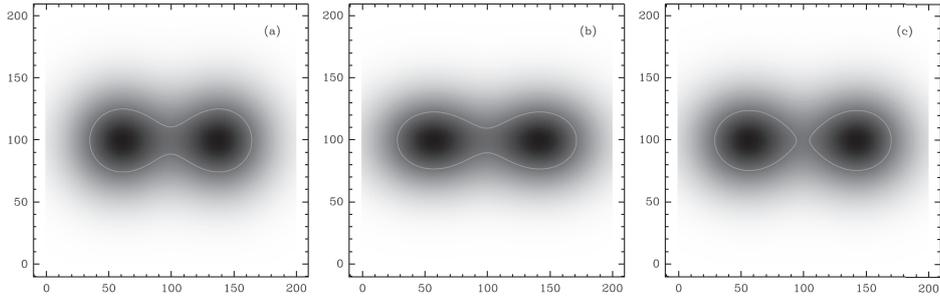


Figure 3.8: Illustration of shear-dependent object detection in the presence of a PSF. In this example, two sources are present and convolved with a PSF. Contours represent constant brightness. In the *left* panel, no shear is added and a detection algorithm finds only one source as highlighted by the single contour. A shear is added before convolution by the PSF in the *right* panel and the detection algorithm finds now two sources as highlighted by the discontinuous iso-contour. This effect is mainly due to the PSF convolution as compared with the *middle* panel where a shear is applied after the PSF convolution and the detection algorithm finds again a single source. Figure from Sheldon et al. [SH17].

3.4.3 Metadetection with automatic differentiation

While metadetection has proven to be very efficient on simulations, reducing the multiplicative bias under stage IV surveys requirements, it is still subject to approximation errors and limited to the first-order Taylor expansion. In [Vit+22] we propose to solve for the approximation error due to the finite difference approximation of the response matrix in Equation 3.47 by using automatic differentiation. As it will be explained in subsection 5.1.3, automatic differentiation allows us to compute exact derivatives by computing the chain rule of a differentiable program.

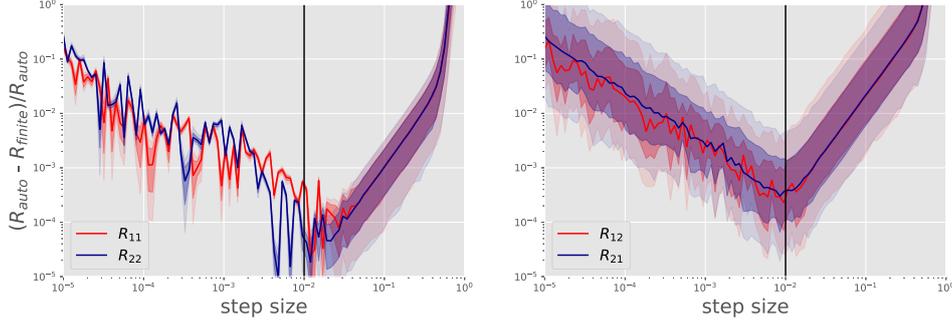


Figure 3.9: Relative difference between response matrix coefficients R_{ij} computed by automatic differentiation and finite differences. We represent the median, 68% and 95% of the responses computed on individual galaxies. As the usual step size for finite differences is 10^{-2} (e.g. used in NGMIX by default), we indicate this choice by a vertical line in both panels.

The averaging in Equation 3.54 can be rewritten in terms of weighted average of galaxy ellipticities

$$\langle e \rangle = \frac{1}{N} \sum_{j=0}^N w_j(x, t) e_{i,j}, \quad (3.55)$$

where the weights w are function of an object characteristic x and a threshold t . The selection is often performed on the size ratio between the galaxy and the PSF T/T_{PSF} or the SNR of the galaxy. Usually, a hard threshold is performed to compute the weights w such that

$$w(x, t) = \begin{cases} 1 & \text{if } x \leq t \\ 0 & \text{otherwise.} \end{cases} \quad (3.56)$$

However, such selection function cannot be used with automatic differentiation since the derivate of this function is zero everywhere but at $x = t$, where it is not defined. This is why we proposed to use a smooth monotonically increasing function instead:

$$w(x, t, \sigma) = \frac{1}{2} \left(1 + \operatorname{erf} \left(\frac{\log(x) - \log(t)}{\sigma} \right) \right), \quad (3.57)$$

with $\operatorname{erf}(x) = \frac{2}{\pi} \int_0^x e^{-t^2} dt$ and where σ drives the smoothness of the function. The smaller σ , the sharper is the weight function, until looking like hard thresholding for $\sigma = 0$.

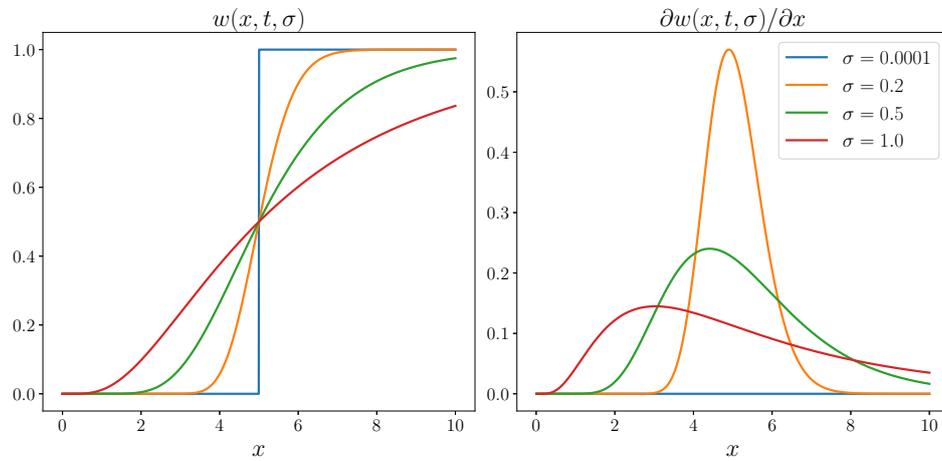


Figure 3.10: Plot of the selection function $w(x, t, \sigma)$ in Equation 3.57 for different smoothing factors. The x-axis is arbitrary and here the threshold is $t = 5$. The *left* panel shows the smooth selection function and the *right* panel shows their derivative.

3.5 Conclusion

In this chapter, we have seen how galaxy images are formed and how to extract cosmic shear information. We have covered the state-of-the-art methods for PSF ellipticity correction and calibration to recover unbiased shear estimation.

* * *
* *
*

Part II

Generative modeling and Bayesian statistics

Bayesian statistics

Chapter Outline

4.1	Introduction to Bayesian modeling	39
4.1.1	General introduction of Bayesian modeling	40
4.1.2	Bayesian modeling in Cosmology	41
4.2	Markov Chain Monte Carlo sampling	41
4.2.1	Metropolis-Hastings	42
4.2.2	Langevin dynamics	43
4.2.3	Hamiltonian Monte Carlo	43
4.3	Variational inference	45
4.3.1	Forward and reverse KL divergences	47
4.3.2	Evidence lower bound (ELBO)	48

THIS chapter introduces the different concepts of Bayesian modeling. We aim to reconstruct a signal from measurement or to infer model parameters from observations. We therefore need efficient methods enabling us to handle high-dimensional parameter space and to quantify the uncertainty inherent to the problem. We start this chapter by identifying the key ingredients of Bayesian modeling, and the Bayes’s identity which is at the core of the probabilistic framework. We then present sampling algorithms, based on Markov Chain Monte Carlo (Markov Chain Monte Carlo (**MCMC**)) enabling us to model exactly posterior distributions, and our contribution to make them completely score-based. We finally present variational inference (Variational Inference (**VI**)) algorithms which rely on optimization and as a consequence, are much faster than **MCMC**s but approximate.

4.1 Introduction to Bayesian modeling

Bayesian modeling is a handy tool for quantifying the uncertainty of the solutions to a problem where the observations are corrupted or may have been generated from a non-unique set of parameters. In this section, we address the following questions:

- *How to model uncertainties of the reconstruction of a signal which generates observations?*

- How to sample a distribution from which only the evaluation of its log probability density function is possible?
- How to learn the probability density of a distribution from which we only have access to samples?
- How can we perform this in the most computationally efficient way?
- How to fit a family of functions to a distribution, when only either samples are available or probability density evaluation?

We describe in this section notations and important theorem of Bayesian statistics

4.1.1 General introduction of Bayesian modeling

Let us introduce a random variable, denoted as X and its associated probability density function (p.d.f.) $X \sim p_X(x)$. Throughout this thesis, we will interchangeably refer to the random variable X to its realization $x \sim X$. Similarly, we will use the shorthand $p(x)$ to refer to the density function $p_X(x)$.

Let us consider another random variable y which is observed and has been generated by another random variable x according to a *forward model*. If the process that maps the random variable x to y is not deterministic –for instance, involving stochastic processes or simply not invertible (e.g. in cases of underdetermined systems where there are more unknowns than equations)– then the inverse mapping from y to x is not unique and can be expressed under a conditional probability distribution $p(x|y)$ called the *posterior distribution*. This conditional distribution is most often not known and we need to rely on the *Bayes' identity* to characterize it

$$p(x|y) = \frac{p(x, y)}{p(y)} = \frac{p(y|x)p(x)}{p(y)} \quad (4.1)$$

This identity involves several distributions that we describe below

Posterior distribution $p(x|y)$ The *posterior distribution* plays a central role in Bayesian modeling as it characterizes the possible solutions conditioned on observed data. This p.d.f. is often unknown analytically and therefore requires the Bayes' identity to be characterized.

Likelihood distribution $p(y|x)$ The *likelihood distribution* quantifies the probability of observing data y conditioned on an original signal x . It is most often given by the forward model mapping x to y .

Prior $p(x)$ The *prior* distribution encapsulates all the prior information we have about the signal x we intend to reconstruct. It reflects our knowledge about the possible values of x before considering any observation.

Joint distribution $p(x, y)$ In many cases, we have access to expressions for both the likelihood and the prior distributions. We are therefore able to evaluate the *joint distribution*. It is crucial to notice that the joint distribution is proportional to the posterior distribution when only x varies (i.e. $p(x|y) \propto p(x, y)$ w.r.t. x).

Bayesian evidence $p(y)$ It is the marginal density of the observations. The evidence is a constant if we assume a given model. This term is usually intractable as computing it requires to marginalize over all possible signal x . While not necessarily useful for parameter inference as we only require evaluating the joint distribution, proportional to the posterior. However, when working with Bayesian modeling, we may also be interested in model comparison. Bayesian evidence is therefore essential for such tests as it will be discussed later.

4.1.2 Bayesian modeling in Cosmology

Notations in the cosmological literature will often differ from the ones of signal processing and machine learning. In cosmology, one is particularly interested in recovering the uncertainties of cosmological model parameters, denoted θ , from observations denoted d . The posterior distribution would then be denoted $p(\theta|d)$, which can be confusing if compared to the notation we use in this manuscript. Because in this case, the unknown are both the cosmological parameters and the right model, we condition the distribution by the choice of model, denoted \mathcal{M} . The Bayes' identity is therefore expressed as

$$p(\theta|d, \mathcal{M}) = \frac{p(d|\theta, \mathcal{M})p(\theta, \mathcal{M})}{p(d|\mathcal{M})}. \quad (4.2)$$

In some context, the evidence $p(d|\mathcal{M})$ can become the distribution of interest. It is still not important for parameter inference, but is the key distribution to determine what is the best model given the observations. When comparing two models \mathcal{M}_0 and \mathcal{M}_1 under observations d , we can determine which one is most likely condition to the observations by computing the evidence ratio K :

$$K = \frac{p(\mathcal{M}_1, |d)}{p(\mathcal{M}_0|d)} = \frac{p(d|\mathcal{M}_1) p(\mathcal{M}_1)}{p(d|\mathcal{M}_0) p(\mathcal{M}_0)}. \quad (4.3)$$

Under the assumption that the two models are of equal probability, i.e. $p(\mathcal{M}_1) = p(\mathcal{M}_0)$, the evidence ratio becomes

$$K = \frac{p(d|\mathcal{M}_1)}{p(d|\mathcal{M}_0)} = \frac{p(\mathcal{M}_1, |d)}{p(\mathcal{M}_0|d)}. \quad (4.4)$$

Computing this ratio therefore enables to discriminate between the two models under observing conditions d .

In practice, computing Bayesian evidence is intractable to compute as it requires to marginalise over all possible parameters $p(d|\mathcal{M}) = \int_{\Theta} p(d|\theta, \mathcal{M})p(\theta, \mathcal{M})d\theta$. This intractability can be mitigated with an amortized neural network, as done in Jeffrey et al. [JW23].

4.2 Markov Chain Monte Carlo sampling

In this section, we will use the signal processing notations x and y because of the historicity of the discussed methods, which have been mostly developed by this community. The way of modeling a posterior distribution can differ on whether we have access to the likelihood or not. In some cases, one can evaluate the likelihood p.d.f., for instance when the likelihood results from an analytical forward model, but cannot draw samples as easily. The cosmological community usually faces this problem for cosmological inference. In other cases, we can acquire samples from cosmological

simulations, without being able to formulate a probability density function of the underlying implicit distribution. In this section, we will discuss how to sample from a distribution when only density evaluation is possible. Let us denote by $p(x)$ the target distribution from which we aim to sample from. And, by $\pi(x)$ another distribution proportional to $p(x)$.

4.2.1 Metropolis-Hastings

We need an algorithm to sample from a distribution $\pi(x) \propto p(x)$, while only being able to evaluate probability densities. This can be performed using a class of algorithms called Markov Chain Monte Carlo (**MCMC**). **MCMC** generates a suite of point $\{x_i\}$ where each element is sampled from a distribution (hence *Monte Carlo*), only conditioned on the previous element (hence the *Markov Chain*). For this reason, we often say that Markov chains are processes without memory.

One key ingredient of **MCMC** algorithms is the choice of a *proposal* distribution q , such that the element x_t is drawn from this distribution, conditionally to the previous element x_{t-1} . This can be denoted as $x_t \sim q(\cdot|x_{t-1})$. Once the proposal is drawn, an accept-reject criteria is performed ensuring that the proposal is indeed sampled from the target distribution. The acceptance threshold is defined as

$$a(x', x_t) = \min \left(1, \frac{\pi(x') q(x_t|x')}{\pi(x_t) q(x'|x_t)} \right). \quad (4.5)$$

Sampling from a uniform random variable $u \sim \mathcal{U}[0, 1]$, the proposal is accepted if $u < a(x', x_t)$. In order to maximize the number of accepted sampled, we therefore need to design a proposal distribution that maximizes the acceptance ratio a . When using a centered Gaussian distribution of fixed variance for the proposal distribution, the **MCMC** algorithm is called the *Random Walk Metropolis Hastings* and is further described in Algorithm 1.

Algorithm 1 Metropolis-Hastings

Require: target distribution $\pi \propto p$, initialization sample x_0 , proposal distribution q , $t = 0$, size of chain N

```

1: while  $t < N$  do
2:    $x' \sim q(\cdot|x_t)$  ▷ Draw from proposal distribution
3:    $a(x', x_t) = \min \left( 1, \frac{\pi(x') q(x_t|x')}{\pi(x_t) q(x'|x_t)} \right)$  ▷ Compute acceptance probability
4:    $u \sim \mathcal{U}[0, 1]$  ▷ Sample a uniform random variable
5:   if  $u < a(x', x_t)$ 
6:      $x_{t+1} \leftarrow x'$  ▷ Add proposal to chain
7:      $t \leftarrow t + 1$ 
8:   else
9:     end if
10: end while

```

While this method asymptotically ensures convergence of the Markov chain, it might take infinite time to reach full coverage of the target distribution. Therefore, most of the improvements for making **MCMC** algorithms more efficient were based on designing a more efficient proposal distribution q in Equation 4.5. One venue for such improvement is to leverage the first-order derivative of the log probability density function. We will discuss some of these techniques in the next paragraphs.

4.2.2 Langevin dynamics

In this section, we will introduce the *Langevin dynamics*. Let $\nabla_x \log \pi(x) = \nabla_x \log p(x)$ denotes the gradient of our target distribution. The random walk discussed in the former paragraph can be seen in a continuous process described by the stochastic differential equation $dX = dW$, where W is a standard Brownian motion and is called the *diffusion* term of the process, and X a continuous random variable. Langevin dynamics adds a *drift* term to the process as

$$dX = \nabla_x \log \pi(X)dt + \sqrt{2}dW. \quad (4.6)$$

In practice, we need to discretize this stochastic differential equation (SDE) in order to generate a discrete Markov process. There are several methods to discretize differential equations (Euler, 4th order Runge-Kutta, etc.), we use here the Euler-Maruyama discretization method:

$$x_{t+1} = x_t + \tau \nabla_x \log \pi(x_t) + \sqrt{2\tau} \epsilon_t, \quad (4.7)$$

where τ is a small step size and ϵ is sampled from a centered reduced normal distribution, i.e. $\epsilon \sim \mathcal{N}(0, \mathbb{I}_d)$. This generate a new proposal distribution $q(\cdot|x_t) = \mathcal{N}(x_t + \tau \nabla_x \log \pi(x_t)|2\tau)$.

The drift term in Equation 4.7 makes the sampling more efficient thanks to the drift term that pushes back the proposal towards the distribution manifold. Discretization of the Langevin process can lead to numerical approximation error, this is where calibrating the process with the accept-reject Metropolis-Hastings (MH) step will ensure sampling ergodicity. This sampling method is named Metropolis adjusted Langevin algorithm and is presented in Algorithm 2.

Algorithm 2 Metropolis Ajusted Langevin Algorithm (Metropolis Adjusted Langevin Algorithm (MALA))

Require: target distribution $\pi \propto p$, initialization sample x_0 , $t = 0$, size of chain N , step size τ

while $t < N$ **do**

$q(\cdot|x_t) = \mathcal{N}(x_t + \tau \nabla_x \log \pi(x_t)|2\tau)$ ▷ Langevin proposal (Equation 4.7)

$x' \sim \mathcal{N}(x_t + \tau \nabla_x \log \pi(x_t)|2\tau)$ ▷ Draw from proposal distribution

$a(x', x_t) = \min\left(1, \frac{\pi(x') q(x_t|x')}{\pi(x_t) q(x'|x_t)}\right)$ ▷ Compute acceptance probability

$u \sim \mathcal{U}[0, 1]$ ▷ Sample a uniform random variable

if $u < a(x', x_t)$

$x_{t+1} \leftarrow x'$ ▷ Add proposal to chain

$t \leftarrow t + 1$

end if

end while

4.2.3 Hamiltonian Monte Carlo

Similarly to Langevin dynamics, Hamiltonian Monte Carlo (Hamiltonian Monte Carlo (HMC)) proposes to design a way to explore efficiently parameters space leveraging the gradients of the log p.d.f. (also known as the *score function*) using Hamiltonian dynamics

[Nea11; Bet18]. Seen from an energy-based model perspective, the distribution $\pi(x)$ can be written as

$$\pi(x) = \frac{1}{\mathcal{Z}} e^{-E(x)}, \quad (4.8)$$

where \mathcal{Z} is a normalizing constant so that π integrates to 1 over the parameter space, and E is an *energy function*. In practice, we may not have access to the normalized density function because the normalization factor \mathcal{Z} is often intractable to compute, but rather to a function $-E$ proportional to π with respect to x

$$\log \pi(x) \propto -E(x). \quad (4.9)$$

In the **HMC** algorithm, the idea is to augment the energy function $E(x)$ with momentum parameters m :

$$\begin{aligned} H(x, m) &= E(x) + K(x, m) \\ &= -\log \pi(x) + \log \pi(m|x), \end{aligned} \quad (4.10)$$

where H is called the Hamiltonian. The dynamic of the parameters x and momentum m is then described by the Hamiltonian equations:

$$\begin{aligned} \frac{dx}{dt} &= +\frac{\partial H}{\partial m} = \frac{\partial K}{\partial m} \\ \frac{dm}{dt} &= -\frac{\partial H}{\partial x} = -\frac{\partial K}{\partial x} - \frac{\partial E}{\partial x} \end{aligned} \quad (4.11)$$

These are continuous equations and can be discretized with a leapfrog integrator such as

$$\begin{aligned} m_{t+\frac{\alpha}{2}} &= m_t + \frac{\alpha}{2} \nabla \log p(x_t) \\ x_{t+\alpha} &= x_t + \alpha \mathbf{M}^{-1} m_{t+\frac{\alpha}{2}} \\ m_{t+\alpha} &= m_{t+\frac{\alpha}{2}} + \frac{\alpha}{2} \nabla \log p(x_{t+\alpha}) \end{aligned} \quad (4.12)$$

where α is the step size, m is the auxiliary momentum and \mathbf{M} is a preconditioning matrix that could take into account the space metric, but in our case the identity matrix.

Following this procedure, **HMC** is supposed to sample from $p(\mathbf{x})$, but as explained in [Bet18], the discretization induces a small error that will bias the resulting transition and requires a correction. In order to correct this bias, every sample is considered as a Metropolis-Hastings (MH) [Met+53; Has70] proposal and is accepted or rejected according to an acceptance probability. This acceptance probability is designed from the Hamiltonian transition and is also only score-dependent as we explain below.

The discretized implementation of the Hamiltonian Monte Carlo algorithm requires correcting for the discretization error. The chain update computed with Equation 4.12 is a proposal which is accepted with a probability of the form

$$\alpha = \min\{1, p(x_*)q(x_n|x_*)/p(x_n)q(x_*|x_n)\}, \quad (4.13)$$

where x_n is the last sample of the chain, x_* is the **HMC** proposal, p is the target density and q is a proposal distribution from a random walk, i.e. $q(x_n|x_*) = \mathcal{N}(x_n|x_*, \mathbf{M})$ with \mathbf{M} the **HMC** preconditioning matrix. In our approach, we do not directly have access to the distribution p , but to its score function $\nabla \log p$. However, we

can still approximate the log ratio needed to compute the [MH](#) acceptance probability from the scores using the path integral:

$$\log p(x_*) - \log p(x_n) = \int_0^1 \nabla_x \log p(t * (x_* - x_n) + x_n) \cdot (x_* - x_n) dt, \quad (4.14)$$

which we evaluate in practice with a simple 4 points Simpson integration rule. This integral could be approximated to any precision at the cost of additional score evaluations.

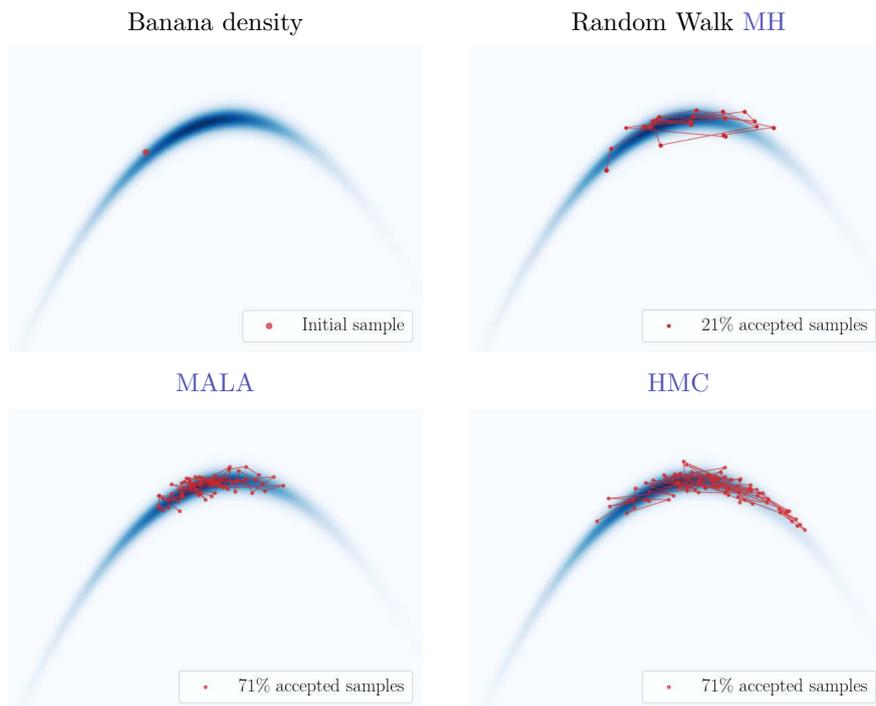


Figure 4.1: Comparison between the samplers on the banana distribution (*top left*). On the *top right* is the Random Walk Metropolis Hastings, on the *bottom left* the Metropolis Adjusted Langevin Algorithm and on the *bottom right* the Hamiltonian Monte Carlo.

Thus, with [Equation 4.14](#) at hand, we are able to implement a large class of [MCMC](#) algorithms such as [HMC](#) or Metropolis Adjusted Langevin Algorithm, using the score function only.

4.3 Variational inference

Sampling distributions can take a very long time, even for efficient gradient-based samplers like [HMC](#). Complementary to [MCMC](#) sampling algorithm, the field of statistics brought another class of inference algorithms, called *Variational inference* ([VI](#)). This class of method differs from [MCMC](#) in many senses. The aim of [VI](#) is to fit a family of parametric distributions to the distribution of interest. Therefore, if there

Algorithm 3 Hamiltonian Monte Carlo algorithm [Bet18]

Require: target distribution $\pi \propto p$, initialization sample x_0 , $t = 0$, size of chain N , step size α , leapfrog steps T

```

1: while  $t < N$  do
2:    $m_0 \sim \mathcal{N}(0, \mathbb{I}_d)$  ▷ Draw random momentum
3:    $x \leftarrow x_t$ 
4:   for  $0 \leq n \leq \lfloor T/\alpha \rfloor$  do ▷ Leapfrog integration
5:      $m \leftarrow m + \frac{\alpha}{2} \nabla \log p(x)$ 
6:      $x \leftarrow x + \alpha \mathbf{M}^{-1} m$ 
7:      $m \leftarrow m + \frac{\alpha}{2} \nabla \log p(x)$ 
8:   end for
9:    $a(x, x_t) = \min(1, \exp(H(m_t, x_t) - H(m, x)))$  ▷ Acceptance probability
10:   $u \sim \mathcal{U}[0, 1]$  ▷ Sample a uniform random variable
11:  if  $u < a(x', x_t)$ 
12:     $x_{t+1} \leftarrow x$  ▷ Add proposal to chain
13:     $t \leftarrow t + 1$ 
14:  end if
15: end while

```

Algorithm 4 Hamiltonian Monte Carlo algorithm [Bet18], score-only version [Rem+20; Ram+20]

Require: score of the target distribution $\nabla \log p(x)$, initialization sample x_0 , $t = 0$, size of chain N , step size α , leapfrog steps T

```

1: while  $t < N$  do
2:    $m_0 \sim \mathcal{N}(0, \mathbb{I}_d)$  ▷ Draw random momentum
3:    $x \leftarrow x_t$ 
4:   for  $0 \leq n \leq \lfloor T/\alpha \rfloor$  do ▷ Leapfrog integration
5:      $m \leftarrow m + \frac{\alpha}{2} \nabla \log p(x)$ 
6:      $x \leftarrow x + \alpha \mathbf{M}^{-1} m$ 
7:      $m \leftarrow m + \frac{\alpha}{2} \nabla \log p(x)$ 
8:   end for
9:    $a(x, x_t) = \min(1, \exp(\log p(x_t) - \log p(x)))$  ▷ Acceptance probability
   according to Equation 4.14
10:   $u \sim \mathcal{U}[0, 1]$  ▷ Sample a uniform random variable
11:  if  $u < a(x', x_t)$ 
12:     $x_{t+1} \leftarrow x$  ▷ Add proposal to chain
13:     $t \leftarrow t + 1$ 
14:  end if
15: end while

```

is no guarantee that the distribution of interest is in the family, the inference will be subject to *approximation error*. In addition, this method does not rely anymore on sampling Markov chains, but rather on optimization algorithms. It can therefore also be subject to *optimization error*. Finally, as we are aiming to build an estimator of the target distribution, it can also be subject to *statistical error* due to statistical bias of the estimator. We will use the notation \mathcal{F} to denote *families* of distributions, i.e. a given set of distributions, often sharing the same parametrization. Different families of distributions exist in the literature, these include the *mixture of Gaussians*, which are characterized by means, variances, and relative weights of the Gaussians; *exponential families*; and neural-based families such as *normalizing flows*, which are parametrized by neural network weights. VI is therefore the task of optimizing the parameters in order to find the closest distribution within the family to the target distribution. We will review here the objects necessary to the understanding of the three next chapters of this thesis. For a more in-depth review, we refer the reader to Blei et al. [BKM17] from which this section is widely inspired.

4.3.1 Forward and reverse KL divergences

In the following, we will call q the surrogate distribution and p the target distribution we aim to approximate. In order to compare distributions, we cannot define a proper *distance* as we can naturally do with vectors for instance. However, we can rely on *divergences* which quite similarly quantifies how far a distribution is from another one, but as opposed to distances, divergences are not symmetric. The most widely used divergence in the VI literature is the *Kullback-Leibler divergence* [KL51], or Kullback–Leibler (KL)-divergence for short. As we will use the KL divergence to quantify how close our surrogate distribution is from the target, we distinguish between the *forward* or the *reverse* KL-divergences, whether we measure the divergence between p and q or the reverse.

Forward KL-divergence The Forward KL-divergence is defined as the expectancy of the ratio between the log probability density functions p and q , with respect to a random variable x sampled from p .

$$D_{\text{KL}}(p||q) = \mathbb{E}_{x \sim p} \left[\frac{\log p(x)}{\log q(x)} \right]. \quad (4.15)$$

In order to evaluate the forward KL-divergence, we not only need to be able to evaluate the log p.d.f. of the target distribution but also to sample (or have access to samples) from it. We usually build a surrogate distribution q which for which it is easy to evaluate its log-probability.

Reverse KL-divergence The reverse KL-divergence is defined as the expectancy of the ratio between the log probability density functions q and p , with respect to a random variable x sampled from q .

$$D_{\text{KL}}(q||p) = \mathbb{E}_{x \sim q} \left[\frac{\log q(x)}{\log p(x)} \right] \quad (4.16)$$

In order to evaluate the reverse KL-divergence, we only need to be able to evaluate the p.d.f. of the target distribution. We therefore need to build a surrogate distribution from which it is easy to both evaluate the log-probability and to sample from it.

As the KL-divergence is not symmetric, the behavior of the surrogate distribution is not the same when it is optimized to minimize the forward or the reverse KL.

We illustrate in [Figure 4.2](#) that a distribution q optimized so that the forward **KL**-divergence is minimal, results in a wider distribution that matches the entire support of the target distribution. Minimizing the reverse **KL** however, the surrogate distribution will try to match the target, while ensuring to be zero where the target is zero which may lead to missing modes.

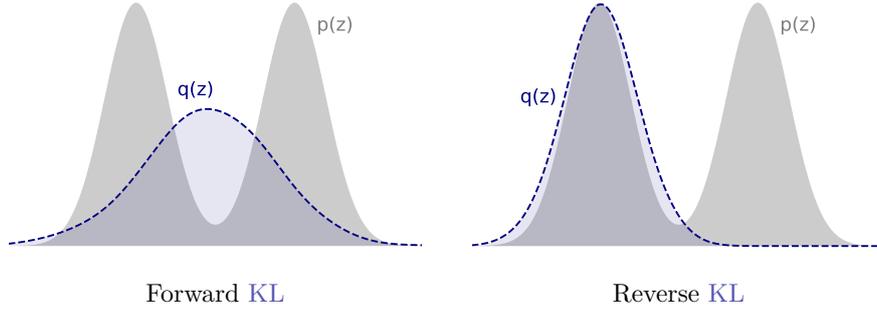


Figure 4.2: Comparison of the different behaviors between the minimization of the forward and the reverse **KL** divergences.

4.3.2 Evidence lower bound (ELBO)

In the context of Bayesian inference, our goal is to fit the surrogate distribution q to a posterior distribution $p(x|y)$, where y corresponds to the observations. As we have seen before, evaluating $p(x|y)$ is often not possible, but we can rely on the likelihood and prior through the Bayes' identity as follows

$$q^* = \arg \min_{q \in \mathcal{F}} D_{\text{KL}}(q || p(x|y)) \quad (4.17)$$

$$= \arg \min_{q \in \mathcal{F}} \mathbb{E}_{x \sim q} [\log q(x) - \log p(x|y)] \quad (4.18)$$

$$= \arg \min_{q \in \mathcal{F}} \mathbb{E}_{x \sim q} [\log q(x) - \log p(x, y) + \log p(y)]. \quad (4.19)$$

Evaluating the reverse **KL**-divergence requires therefore to evaluate the Bayesian evidence as rewritten in [Equation 4.19](#), which is often intractable to compute. However, as the expectancy is taken over the random variable x drawn from the surrogate distribution, the evidence $p(y)$ can be seen as a constant of the optimization problem. Therefore, when minimizing the reverse **KL**, we can drop the evidence from the expression to minimize a new loss, i.e.

$$q^* = \arg \min_{q \in \mathcal{F}} \mathbb{E}_{x \sim q} [\log q(x) - \log p(x, y)]. \quad (4.20)$$

This loss is commonly known in the **VI** literature as the negative *evidence lower bound*. Indeed, by inverting the terms of [Equation 4.19](#), we can rewrite the Bayesian evidence $p(y)$ as

$$p(y) = \underbrace{\mathbb{E}_{x \sim q} [\log p(x, y) - \log q(x)]}_{\text{ELBO}(y)} + \underbrace{D_{\text{KL}}(q \| p(x|y))}_{\geq 0}. \quad (4.21)$$

The **KL** term is necessarily positive by definition, therefore the second term constitutes a lower bound of the Bayesian evidence. The algorithm approximating the posterior distribution by minimizing the **ELBO** is called the Stochastic Variational Inference and is described in 5.

Algorithm 5 Stochastic Variational Inference algorithm

Require: joint distribution $p(x, y)$, observations y , variational distribution q_φ , step size α

- 1: **while** ELBO has not converged **do**
 - 2: $\mathcal{L}(\varphi) = \mathbb{E}_{x \sim q_\varphi} [\log p(x, y) - \log q_\varphi(x)]$ \triangleright Compute the ELBO
 - 3: $\varphi \leftarrow \varphi - \alpha \frac{\partial \mathcal{L}}{\partial \varphi}$ \triangleright compute gradient descent according to [subsection 5.1.2](#)
 - 4: **end while**
-

* * *
* * *
* * *

Deep generative models

Chapter Outline

5.1	Introduction to deep learning	52
5.1.1	Multi Layer Perceptron	52
5.1.2	Optimizing neural networks	54
5.1.3	Automatic differentiation	55
5.1.4	Differentiating implicit functions	56
5.1.5	Architectures	57
5.1.6	Normalization	58
5.2	Normalizing Flows	60
5.2.1	Bijections	62
5.2.2	Optimizing a normalizing flow	63
5.2.3	Conditional normalizing flows	64
5.3	Score-based diffusion models	65
5.3.1	Stochastic differential equations	65
5.3.2	Denoising score matching	66
5.4	Annealed sampling	67

This chapter covers the content of an article accepted to a peer-reviewed journal:

Remy, B., Lanusse, F., Jeffrey, N., Liu, J., Starck, J.-L., Osato, K., Schrabback, T., ‘Probabilistic mass-mapping with neural score estimation’. In: *Astronomy & Astrophysics* 672 (2023), A51

and presented in a peer-reviewed conferences:

Remy, B., Lanusse, F., Ramzi, Z., Liu, J., Jeffrey, N., Starck, J.-L., ‘Probabilistic Mapping of Dark Matter by Neural Score Matching’. In: *3rd Machine Learning and the Physical Science workshop at NeurIPS* (2020)

Zeghal, J., Lanusse, F., Boucaud, A., **Remy, B.**, Aubourg, E., ‘Neural Posterior Estimation with Differentiable Simulators’. In: *Machine Learning for Astrophysics workshop at ICML* (2022)

THIS chapter introduces the tools for modeling arbitrary distributions using deep learning. The first part of this chapter is dedicated to the introduction of the building blocks of deep neural networks and various algorithms to optimize them. In particular, we focus on linear and convolutional layers, explicit and implicit automatic differentiation for their optimization. We then present various architectures useful for tackling our challenges. In the second part of this chapter, we present state-of-the-art generative models, which aim to model our distributions of interest. Among them, we will cover normalizing flows, latent variable models, score-based diffusion models, and finally discuss the annealed HMC sampler that we first proposed in Remy et al. [Rem+20] and Ramzi et al. [Ram+20].

5.1 Introduction to deep learning

For a comprehensive introduction to machine learning and deep learning, we refer the reader to the amazing books *Probabilistic Machine Learning* from Kevin P. Murphy [Mur22], reviewing machine learning under the lens of probabilistic modeling and the *Deep learning textbook* from I. Goodfellow, Y. Bengio and A. Courville [GBC16].

Neural networks are designed to approximate arbitrary functions f^* mapping an ensemble \mathcal{X} called the input space, to another ensemble \mathcal{Y} called the output space. In the following, we will assume that these ensembles are Euclidian spaces, composed of vectors (or tensors), and we will denote $\mathcal{X} = \mathbb{R}^d$ and $\mathcal{Y} = \mathbb{R}^p$. The fundamental characteristic of neural networks is that they are parametric functions, whose parameter θ are to be optimized to fit the desired function f^* . In this manuscript, we consider only *feedforward networks* which are suites of non-linear functions f_i mapping the input x to the output y . In that respect, there is no back feedback of information along the program. These intermediate functions f_i are called the *layers* of the network, and the more layers there are, the more f is able to model complex functions. The layers in a feedforward neural network are composed as follows

$$f = f_1 \circ f_2 \circ \dots \circ f_N. \quad (5.1)$$

A neural network is therefore a function of input x and parameters θ and output y , i.e. $f(x; \theta) = y$, which we will also denote as $f_\theta(x) = y$. The task we aim to solve is so challenging that we need to involve a large number of layers, making the neural network deep, hence their denomination *deep neural networks*. In the following, we will review the most widely used types of layers and neural network architectures that we used to design our algorithms. We will then present the optimization programs to find the optimal parameters and fit a desired arbitrary function. We will finally cover various architectures, explaining how they were especially useful for our goals.

5.1.1 Multi Layer Perceptron

The Multilayer Perceptron (MLP) is the most simple neural network architecture, as its layers are composed of *linear* functions composed with *non-linear activation functions* as follows

$$y' = Wx + b \quad (5.2)$$

$$y = a(y'), \quad (5.3)$$

where we called a the non-linear activation function, W a matrix of dimension $\mathbb{R}^{p \times d}$ called the *weight matrix* and $b \in \mathbb{R}^p$ called the *bias* parameters. In this case the neural network parameters are the ensemble $\theta = \{W_i, b_i\}_{i \in [1, N]}$ if there are N layers.

Activation functions Activation functions are essential to neural network layers as they contribute to their non-linearity. The activation function in Equation 5.2 is an element-wise function, which means that it applied independently on each element of

the input vector, i.e. $a \left(\begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_d \end{bmatrix} \right) = \begin{bmatrix} a(x_1) \\ a(x_2) \\ \vdots \\ a(x_d) \end{bmatrix}$.

We list below several of these non-linear functions and illustrate them in Figure 5.1.

- *Rectified Linear Unit (Rectified Linear Unit (ReLU))*

$$f(x) = \max(0, x) \quad (5.4)$$

- *Sigmoid*

$$f(x) = \frac{1}{1 + e^{-x}} \quad (5.5)$$

- *Hyperbolic tangent*

$$f(x) = \tanh(x) \quad (5.6)$$

- *Sinus*. Even though the sinus activation function is rarely used in neural network designs, it has the particularity to be infinitely differentiable, with analytical derivatives.

$$f(x) = \sin(x) \quad (5.7)$$

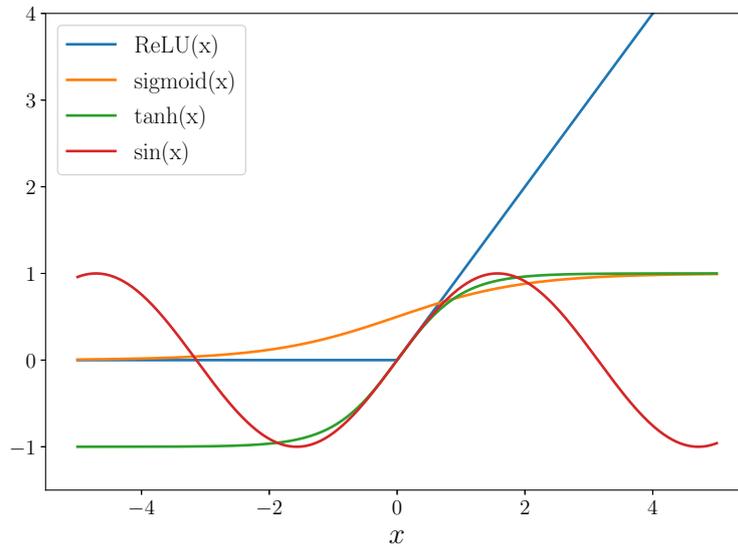


Figure 5.1: Illustration of various activation functions. All these functions have support on \mathbb{R} , and are non-linear and differentiable.

5.1.2 Optimizing neural networks

Now that we have designed a feedforward neural network parameterized by their weights θ , we need to optimize them to fit the function of interest. We first need to specify a *loss function* \mathcal{L} which will be the criteria quantifying how close our function is to the target. The optimization problem is thus a loss minimization problem, minimization because the goal is to reduce the differences between the target distribution and the approximated one. All neural network optimization algorithms are based on *gradient descent* iterations. The loss function \mathcal{L} is a landscape in the space of parameters θ and to find the minimum of this landscape we iteratively move a set of parameters towards the direction of the inverse gradients of the loss. Given a training dataset of pairs of input and output $\mathcal{D}_{\text{train}} = \{(x, y)\}_{x \in \mathcal{X}, y \in \mathcal{Y}}$ the gradient descent iteration is given by

$$\theta_{t+1} = \theta_t - \frac{\alpha_t}{|\mathcal{D}_{\text{train}}|} \sum_{(x,y) \in \mathcal{D}_{\text{train}}} \frac{\partial \mathcal{L}(\theta_t, x, y)}{\partial \theta}, \quad (5.8)$$

where α_t is a small step size, θ_t are the neural network weights to update, t the current time step of the optimization, and $|\mathcal{D}_{\text{train}}|$ being the cardinal (size) of the dataset. In this update, the direction of the update is computed over the entire training set $\mathcal{D}_{\text{train}}$ which is intractable in practice due to the immense amount of data we need. Equation 5.8 can be approximated using small subsets of the training set, called *mini-batches*, asymptotically leading to the same results.

Stochastic Gradient Descent Working with mini-batch instead of evaluating the loss gradient over the entire training set makes the optimization much more tractable. With uniform sampling of examples $\{(x_i, y_i)\}_{i=1 \dots N}$, stochastic gradients are equal in expectancy to the full gradients, i.e. stochastic gradients are an unbiased estimator of the full gradient, i.e.

$$\mathbb{E}_{(x,y) \sim \mathcal{D}_{\text{train}}} \left[\frac{\partial \mathcal{L}(\theta, x, y)}{\partial \theta} \right] = \frac{1}{|\mathcal{D}_{\text{train}}|} \sum_{(x,y) \in \mathcal{D}_{\text{train}}} \frac{\partial \mathcal{L}(\theta, x, y)}{\partial \theta}, \quad (5.9)$$

but introduces variance in the gradient evaluation and therefore stochasticity in the optimization process. The update equation becomes:

$$\theta_{t+1} = \theta_t - \frac{\alpha_t}{N} \sum_{i=1}^N \frac{\partial \mathcal{L}(\theta_t, x_i, y_i)}{\partial \theta}, \quad (5.10)$$

where N is the size of the mini-batch.

Momentum & Adam Adaptation of the step size along the optimization process. Momentum uses information from previous gradient evaluations to compute a more efficient direction for the current update. For simplicity, we will note $\nabla \mathcal{L}(\theta)$ the gradient evaluation of the loss function, with respect to parameters θ , evaluated with mini-batches.

$$\begin{aligned} m_{t+1} &= \beta m_t + \nabla \mathcal{L}(\theta) \\ \theta_{t+1} &= \theta_t - \alpha_t m_{t+1}. \end{aligned} \quad (5.11)$$

β is a scalar parameter, often chosen to $\beta = 0.99$. One can notice that taking $\beta = 0$, we recover Equation 5.8. m_t is the momentum parameter at step t .

Slightly more advanced, Kingma et al. [KB14] propose the Adam optimizer. This one is the most widely used in the machine learning community. It has the advantage of adapting the direction of the optimization step with momentum and also of adapting the learning rate along the training.

$$\begin{aligned} m_t &= \frac{1}{1 - \beta_1^t} \left(\beta m_{t-1} + (1 - \beta_1) \nabla \mathcal{L}(\theta_t) \right) \\ g_t &= \frac{1}{1 - \beta_2^t} \left(\beta_2 g_{t-1} + (1 - \beta_2) \nabla \mathcal{L}(\theta_t)^2 \right) \\ \theta_{t+1} &= \theta_t - \frac{\eta}{\sqrt{g_t + \epsilon}} m_t. \end{aligned} \tag{5.12}$$

m_t is a moving average of the gradients, like in Equation 5.11 and g_t is a moving average of the squared gradients. β_1 and β_2 are scalar parameters (often chosen as $\beta_1 = 0.9$ and $\beta_2 = 0.999$), and ϵ is a very small scalar to avoid numerical problems (e.g. $\epsilon = 10^{-8}$).

5.1.3 Automatic differentiation

All of the precedent methods for training a neural network require evaluating gradients of a loss function with respect to the neural network weights. Because neural networks might be deep and have complex architecture, we do not want to derive analytically its gradient. We therefore rely on *automatic differentiation*. Assessing the loss function involves assessing multiple subfunctions that can be differentiated individually. By applying the chain rule, it is possible to calculate the gradient of the loss function by utilizing the intermediate functions' gradients. For instance, evaluating the derivative of function $h = f \circ g$ at x_0 is given by the equation:

$$\frac{\partial h}{\partial x}(x_0) = \frac{\partial(f \circ g)}{\partial x}(x_0) = \frac{\partial g}{\partial x}(x_0) \left(\frac{\partial f}{\partial g} \circ g \right)(x_0). \tag{5.13}$$

With this equation, we can compute exactly the gradient of the loss function. Estimating this gradient could be done using *numerical (or finite) differences*, but would suffer from approximation error. A neural network is a composition of many functions and evaluating the chain rule can be quite costly. There are different ways of evaluating Equation 5.13, the *forward mode* and *backward mode* automatic differentiation. We will not detail the process of these modes in this text, but we refer to the reader to Roger Grosse's lecture notes on backpropagation and automatic differentiation¹. Using the forward or backward mode for automatic differentiation will depend on the relative dimensionality of the input and output spaces.

Forward mode Most efficient when $d \ll p$, i.e. low dimensional input, high dimensional output.

Backward mode Most efficient when $p \ll d$, i.e. low dimensional output, high dimensional input. This is typically the case when optimizing a neural network. Indeed, the loss output is a scalar, i.e. of dimension 1 and we aim to differentiate with respect to a high number of parameters, e.g. hundreds of thousands, millions, or even more.

¹https://www.cs.toronto.edu/~rgrosse/courses/csc421_2019/

5.1.4 Differentiating implicit functions

Computing the derivative of a function with automatic differentiation will scale not only in time, with respect to the computational graph size, but also spatially (in memory size). For a deep computational graph, one will need to compute and keep in memory all the intermediate gradient evaluations necessary to evaluate the chain rule. There are however special cases where there are better methods for gradient computation than standard forward or backward automatic differentiation. It is the case of computing derivative of *implicit functions*.

Implicit functions An implicit function f can be defined by its *fixed point equation*:

$$z = f(a, z), \quad (5.14)$$

where $z \in \mathbb{R}^d$ are the parameters with respect to which we aim to compute the derivative, and $a \in \mathbb{R}^p$ are auxiliary parameters. The gradients of this function defined implicitly are well-defined and can be derived according to the implicit function theorem.

Let us consider $f : \mathbb{R}^d \times \mathbb{R}^n \rightarrow \mathbb{R}^n$, $a_0 \in \mathbb{R}^p$ and $z_0 \in \mathbb{R}^n$ such that $f(a_0, z_0) = 0$ and f is continuously differentiable with non-singular Jacobian $\frac{\partial f}{\partial a}(a_0, z_0) \in \mathbb{R}^{n \times n}$.

Then, there exist open sets $S_{a_0} \subset \mathbb{R}^p$ and $S_{z_0} \subset \mathbb{R}^n$ respectively containing a_0 and z_0 , and a unique continuous function $z^* : S_{a_0} \rightarrow S_{z_0}$ such that

- $z_0 = z^*(a_0)$
- $f(a, z^*(a)) = 0, \quad \forall a \in S_{a_0}$
- z^* is differentiable on S_{z_0} .

For an intuitive understanding of this theorem will illustrative implementations, we refer the reader to the implicit-layers website².

The solution of an implicit function can either be seen fixed point equation solution or as the root of another function, and both approaches are equivalent. Let us write the derivative computation of an implicit function defined by root finding. Let a^p be the input of the algorithm and $z(a)$ its output. We aim to compute both $z^*(a)$ solution of

$$f(a, z^*(a)) = 0, \quad (5.15)$$

where f is a function defining implicitly the solution z^* , and the derivative $\frac{\partial z}{\partial a}(a_0)$ for arbitrary $a_0 \in S_{a_0}$. Differentiating Equation 5.15 with respect to a yields:

$$\frac{\partial f}{\partial a}(a_0, z_0) + \frac{\partial f}{\partial z}(a_0, z_0) \frac{\partial z^*}{\partial a} = 0. \quad (5.16)$$

Inverting this equation, we get:

$$\frac{\partial z^*}{\partial a} = - \left[\frac{\partial f}{\partial z}(a_0, z_0) \right]^{-1} \frac{\partial f}{\partial a}(a_0, z_0) \quad (5.17)$$

This last equation defines how to implicitly compute the derivative of the optimal solution z^* to the problem formulated in Equation 5.15, from the derivative of the function f .

²https://implicit-layers-tutorial.org/implicit_functions

Implicit inverse function In [Zeg+22], we designed a bijective function f with expressive high order derivatives. The optimization procedure of this project required multiple gradient computations of the neural network architecture. Using standard normalizing flows architecture (see section 5.2) such as RealNVP [DSB17] normalizing flow with affine coupling layers. It is clear that differentiating an affine function more gives a function which will return only zeros. We therefore designed a non-decreasing and infinitely differentiable function as

$$\begin{cases} f_y^\varphi(x) &= \frac{\sigma(x) - \sigma(x_0)}{\sigma(x_1) - \sigma(x_0)} (1 - c_\varphi(y)) + c_\varphi(y)x \\ \sigma(x) &= c_\varphi(y)x + \frac{1 - c_\varphi(y)}{1 + e^{-\rho(x)}} \\ \rho(x) &= a_\varphi(y) \left(\log \left(\frac{x}{1-x} \right) + b_\varphi(y) \right) \end{cases}, \quad (5.18)$$

where $a_\varphi(y)$, $b_\varphi(y)$ and $c_\varphi(y)$ are estimated from a neural network of parameters φ .

While being non-decreasing on a compact support, therefore invertible, there is no analytical expression of the inverse of function ρ . However, in order to optimize and perform inference with a normalizing flow, one is required to evaluate bijective functions in both directions. We relied on the implicit function theorem to evaluate the inverse function and compute its gradients. Indeed one can define implicitly the inverse of function ρ through a root finding problem as

$$\rho(x) - y = 0, \quad (5.19)$$

where y is known and we aim to find the corresponding x . More details on normalizing flows and RealNVP can be found in section 5.2.

5.1.5 Architectures

Convolutional layers

Convolutional layers have been developed for neural networks dealing with images as inputs. Instead of a weight matrix mapping between the input and the output of the layer, the convolutional layers involve a convolutional kernel $W^{n \times n \times c}$, where n is the width and height of the kernel and c a channel dimension. The convolutional layer is defined as

$$y = W \otimes x + b. \quad (5.20)$$

It is particularly suited to image-processing tasks because of its translational equivariance property. Indeed, if the input x is spatially translated, the output y follows the same translation, i.e. $S(y) = f(S(x))$ if S is a translation operator and f the convolutional layer defined in Equation 5.20. To understand why this property is useful, let us take the example of a classification task. If we were to classify images of cats and dogs with a deep learning model, the activation of the neural network weights should not change whether the cat or the dog is positioned on the top left or bottom right of the image, and most importantly the final answer should be the same.

Residual layers

The deeper the neural network, the harder it is to train. In the same time, it is known that complex tasks require very deep neural networks to have good performance. He et al. [He+15] proposed to use residual connection between neural network layers to ease the optimization. A classical layer, e.g. linear as in Equation 5.2 or convolutional as in Equation 5.20 can be expressed as a function f parametrized by weights θ taking an input tensor x and outputting a tensor y , i.e.

$$y = f_{\theta}(x). \quad (5.21)$$

Such layer needs to learn the full mapping between $x \in \mathbb{R}^d$ and $y \in \mathbb{R}^p$. In deep neural networks, many of these functions are stacked and each one needs to learn a complete representation. Residual connections enable to give a reference to the layer, so that it only needs to learn the residual information between the previous layer and the output, i.e.

$$y = f_{\theta}(x) + x. \quad (5.22)$$

Often, dimensions p and d are different, and one can perform a simple linear mapping to match the dimensions

$$y = f_{\theta}(x) + Wx, \quad (5.23)$$

where $W \in \mathbb{R}^{p \times d}$ is a weight matrix whose parameters are optimized jointly with the one of f_{θ} .

5.1.6 Normalization

Batch normalization Training deep neural networks can be fastidious. As discussed in subsection 5.1.2, data need to be preprocessed to be standardized and make the training more stable. However, it is not only the statistics of the neural network input which are likely to change over the training process. Indeed, after each optimization step, neural network weights are updated, and therefore the statistics of each layer output as well. Ioffe et al. [IS15] refer to this phenomenon as *internal covariate shift* and propose to add *batch normalization* layers to regularize the amplitude of neural network hidden layers. A batch normalization layer is placed after the activation function of the previous layer and evaluated over the current mini-batch x

$$\phi_{\text{BN}}(x) = \left(\frac{\gamma}{\sqrt{\text{std}(x)}} \right) (x - \bar{x}) + \beta, \quad (5.24)$$

where \bar{x} and $\text{std}(x)$ refer to the mean and standard deviation computed over the batch dimension of x . γ and β are parameters learned over the neural network optimization process. They are also often referred as **state** parameters in neural network libraries such as **JAX**, to distinguish them from neural network weights which we refer as **params** variables.

This way the input of each layer following a batch normalization layer is recentered and rescaled to variance equals one, which makes the training more stable and often faster.

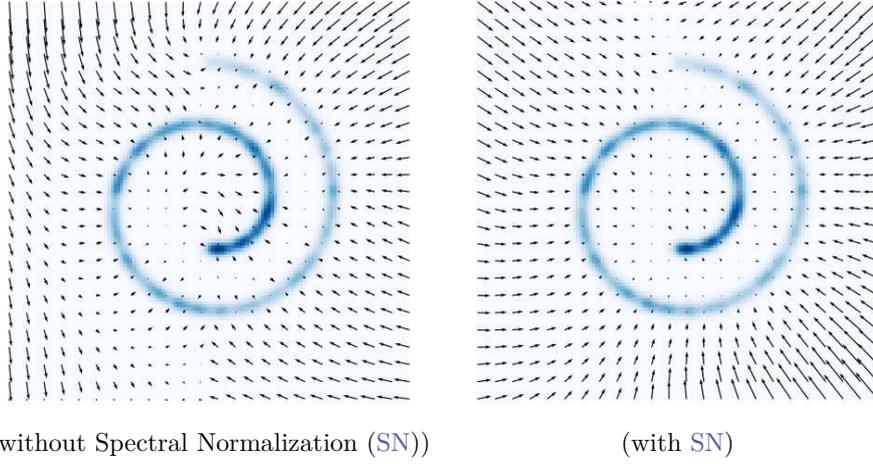


Figure 5.2: Comparison of the score function learning with and without spectral normalization. In blue is the density of the Swiss roll distribution and the black vector field is the learned score function evaluated on a grid.

Spectral normalization It is also possible to perform normalization at the layer level, this is called *layer normalization*. This type of normalization will enforce constraints on the regularity of the network. One example is enforcing the lipschitzness of a neural network with *spectral normalization* (SN).

A function $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is k -Lipschitz continuous if

$$\forall (x, y) \in \mathbb{R}^d \times \mathbb{R}^d, \quad \|f(x) - f(y)\|_2^2 \leq k\|x - y\|_2^2. \quad (5.25)$$

The lipschitzness of a function also accepts a more general form of definitions, with various distances in the input and output space, but Equation 5.25 is sufficient for our description.

Spectral Normalization on the network regularity in those regions. Indeed, as it can be seen in figure Figure 5.2, spectral normalization smooths the learned gradient map far from the high densities. Regularizing the spectral normal of a network lowers its Lipschitz constant, which prevents high variation between close points, thus aligning unconstrained gradients. Controlling the exact Lipschitz constant of a neural network cannot be done directly, but we can control the Lipschitzness of the network layer by layer. One way to do it is to use *spectral normalization*. It can indeed be shown that for linear and convolutional layers, the Lipschitz constant is given by the spectral norm of the neural network weights. Let us make the derivation for a linear layer. A similar demonstration can be done for other kinds of layers such as convolutional, pooling, attention and can be found in [Gou+21]. A linear layer can be expressed with a weight matrix $W \in \mathbb{R}^{d \times p}$ and bias vector $b \in \mathbb{R}^p$, therefore the Lipschitz continuity of a linear layer can be characterized by finding an upper bound of the ratio, for $(x, y) \in \mathbb{R}^d \times \mathbb{R}^d$ and $x \neq y$,

$$\frac{\|Wx + b - Wy - b\|_2^2}{\|x - y\|_2^2} = \frac{\|W(x - y)\|_2^2}{\|x - y\|_2^2} \leq \sigma(W), \quad (5.26)$$

where $\sigma(W) \triangleq \sup_{a \neq 0} \frac{\|Wa\|_2^2}{\|a\|_2^2}$ is the spectral norm of the weight matrix W . Because we chose to use the ℓ_2 norm in Equation 5.25, the spectral norm is given by the

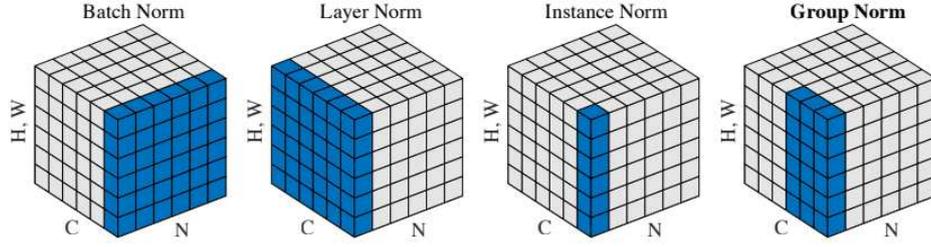


Figure 5.3: Illustration of the different normalization strategies. N corresponds to the batch size, C is the number of channels, H and W are respectively the height and width of the spatial axis. Credit: Wu et al. [WH18]

first singular value of the matrix W . Therefore, it is clear that if we control σ in Equation 5.26, we control an upper bound to the Lipschitz constant of the layer.

As a neural network can be seen as a composition of many functions ϕ_i , such that $f(x) = (\phi_1 \circ \phi_2 \circ \dots \circ \phi_N)(x)$, we can compute the upper bound of the Lipschitz constant of the neural network as the product of Lipschitz constant of each subfunction, i.e. $L(f) \leq \prod_{i=1}^N L(\phi_i)$. As the Lipschitz constant of standard activation function, such as Rectified Linear Unit or the sigmoid function, is 1, the upper bound is given by the product of the spectral norm of weight matrices of the network.

In practice, weight matrices' spectral norm can be normalized by dividing their highest spectral value, which can be estimated using the Power iteration algorithm proposed in Gouk et al. [Gou+21] and describe in Algorithm 6.

Algorithm 6 Power iteration method [Gou+21], returning the highest singular value of a matrix W

```

1: for  $i = 0$  to  $n$  do
2:    $x_i \leftarrow W^T W x_{i-1}$ 
3: end for
4:  $\sigma_{\max}(W) \leftarrow \frac{\|W x_n\|_2^2}{\|x_n\|_2^2}$ 

```

This operation can be done after each layer evaluation of the network, or at once over the whole tree of parameters of a neural network³

5.2 Normalizing Flows

Modeling complex probability densities is a task the machine learning community has been trying to address for years. Most known examples are Mixture of analytic distributions such as mixture of Gaussians or Kernel Density Estimators. However, even though the literature contains a lot of mathematical derivations and interpretations of these models, they lack a lot of flexibility and require prior knowledge about the distribution (number of modes in the distribution of interest for instance).

Recently, Rezende et al. [RM16] proposed a class of model much more flexible, called *normalizing flows* (Normalizing Flow (NF)). A normalizing flow is a composition

³Neural network libraries such as DeepMind's Haiku propose layer-wise and tree-wise implementation of spectral normalization

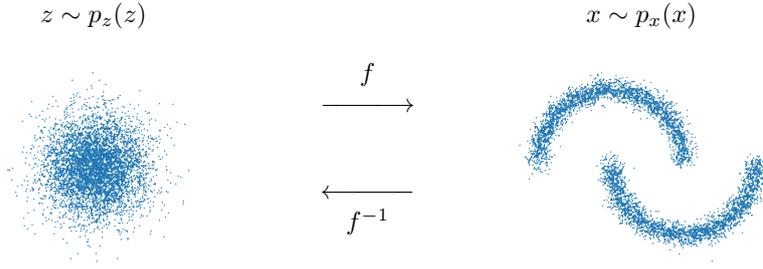


Figure 5.4: Flow f mapping a two-dimensional multivariate Gaussian on the *left* to a two-moons distribution on the *right*. This mapping is bijective, for each sample from the two moons distribution corresponds a unique sample from the multivariate Gaussian through f^{-1} .

of continuous and differentiable functions f , also called *diffeomorphisms*, mapping a random variable z with associated p.d.f $p(z)$ to another random variable x with associated p.d.f. $p_x(x)$. In other words, every samples $z \sim p_z(z)$ will have one and only one corresponding x such that $x = f(z)$. Inversely, for all $x \sim p_x(x)$ there exist a unique $z \sim p(z)$ such that $z = f^{-1}(x)$. Sampling a very complex distribution can therefore become an easy task if one knows the diffeomorphism that maps it to a distribution from which it is easy to sample, such as a normal distribution. With such mapping, it is not only straightforward to sample complex distribution, but also to evaluate its probability density thanks to the change of variable theorem:

The probability density of $p_x(x)$ can also be evaluated from the probability of the corresponding sample z under the change of variable theorem:

$$p_x(x) = p_z(z) \left| \det \left(\frac{\partial f}{\partial x}(z) \right) \right|^{-1}, \quad (5.27)$$

where $x \sim p_x(x)$, $z \sim p_z(z)$ are such that $x = f(z)$, and $\frac{\partial f}{\partial x}(z)$ denotes the Jacobian matrix of f evaluated at z .

Indeed, we can substitute z by $f^{-1}(x)$ in [Equation 5.27](#) which gives

$$p_x(x) = p_z(f^{-1}(x)) \left| \det \left(\frac{\partial f}{\partial x}(f^{-1}(x)) \right) \right|^{-1} \quad (5.28)$$

$$= p_z(f^{-1}(x)) \left| \det \left(\frac{\partial f^{-1}}{\partial x}(x) \right) \right|, \quad (5.29)$$

where $\frac{\partial f^{-1}}{\partial x}(x)$ denotes the Jacobian of the inverse transformation f^{-1} evaluated at x .

As one can see in [Equation 5.29](#), in order to evaluate log probability densities using a normalizing flow, we need to compute the determinant of the Jacobian of the bijective transformation. This step can be very computationally expensive and is the bottleneck of normalizing flows. To minimize the computational cost of this evaluation, efficient flows [[DSB17](#); [PPM17](#); [Kin+16](#); [KD18](#); [Dur+19](#)] are built such that the Jacobian of the transformations is represented as a triangular matrix, i.e. such that the determinant becomes the multiplication of all the diagonal elements of the Jacobian matrix.

In addition, parts of the mapping function f are modeled with neural networks, which is why they are called neural density estimators in the astrophysics literature. We review below different architectures of normalizing flows and in particular the ones used in our contributions discussed in [chapter 7](#).

5.2.1 Bijections

Normalizing flows are based on the design of bijective functions, with the constraint to have a tractable Jacobian determinant computation. In the experiments discussed in this manuscript, we used two flexible bijectors, one proposed by Dinh et al. [[DSB17](#)], called the *coupling flow*, and the other proposed by Papamakarios et al. [[PPM17](#)] called *masked autoregressive flow*. We present a few details of how these to build these bijectors in the following.

Coupling Flows In coupling flows [[DSB17](#)], the total dimension of the vector space is split into two parts or dimensions d and $D - d$. The bijector applied to the first part is the identity function. For the second part, a function composing $z_{1:d}$ and $z_{d+1:D}$ as follows

$$\begin{aligned} x_{1:d} &= z_{1:d} \\ x_{d+1:D} &= g(z_{d+1:D}, h(z_{1:d})). \end{aligned} \quad (5.30)$$

The choice of the functions g and h is free, as long as they remain bijective functions. One choice for these functions has been proposed in the same paper, called Real NVP.

Real NVP The coupling equations for Real NVP are as follows

$$x_{1:d} = z_{1:d} \quad (5.31)$$

$$x_{d+1:D} = z_{d+1:D} \odot \exp(s(z_{1:d})) + t(z_{1:d}), \quad (5.32)$$

The Jacobian of the transformation given by [Equation 5.31](#) is

$$\frac{\partial f}{\partial x}(x) = \begin{bmatrix} \mathbb{I}_d & 0 \\ \mathbf{C} & \mathbf{D} \end{bmatrix}, \quad (5.33)$$

where

- $\mathbf{C} = \frac{\partial f}{\partial z_{[1:d]}}(x[d+1 : D])$ is the Jacobian of the function f with respect to parameters $z_{[1,d]}$ and evaluated at $x_{[d+1:D]}$.
- $\mathbf{D} = \text{diag}(\exp(s(z_{[1,d]})))$ is the diagonal matrix where the diagonal elements are the parameters of $\exp(s(z_{[1,d]}))$.

This way, the full Jacobian $\partial f / \partial x$ is lower triangular and the determinant term in [Equation 5.29](#) is:

$$\det \left(\frac{\partial f}{\partial x} \right) = \exp \left(\sum_j s(z_{1:d})_j \right), \quad (5.34)$$

which is much less costly computationally to compute than the determinant of a full matrix. Moreover, it is worth noticing that $\partial f / \partial x$ only depends on the evaluation of the function s but not on the Jacobians of s or t . This means that these two functions

can be parametrized by very flexible models that have fast-forward evaluation but arbitrary complex Jacobians, which is the case of neural networks.

Finally, computing the reverse transformation f^{-1} has the same characteristic as the forward transformation f . The inverse transformation equation is:

$$z_{1:d} = x_{1:d}, \quad (5.35)$$

$$z_{d+1:D} = (x_{d+1:D} - t(x_{1:d})) \cdot (-s(x_{1:d})). \quad (5.36)$$

Masked Autoregressive Flow Papamakarios et al. [PPM17] proposed to factorize the distribution of a variable x in an autoregressive manner, i.e. as:

$$p(x) = \prod_{i=1}^d p(x_i | x_{1:i-1}), \quad (5.37)$$

where $p(x_i | x_{1:i-1})$ is a one dimensional distribution over x_i , the i -th element of the random variable x , condition on the previous elements $[1 : i - 1]$. The conditional distribution is parametrized by a Gaussian distribution whose mean and variance are conditioned on previous variables:

$$p(x_i | x_{1:i-1}) = \mathcal{N}(x_i | \mu_i, \sigma_i), \quad (5.38)$$

$$x_i = u_i \sigma_i + \mu_i, \quad \text{with} \quad \begin{aligned} \mu_i &= f_{\mu_i}(x_{1:i-1}) \\ \sigma_i &= f_{\sigma_i}(x_{1:i-1}) \end{aligned} \quad (5.39)$$

The inverse transformation

$$u_i = (x_i - \mu_i) \sigma_i^{-1}, \quad \text{with} \quad \begin{aligned} \mu_i &= f_{\mu_i}(x_{1:i-1}) \\ \sigma_i &= f_{\sigma_i}(x_{1:i-1}) \end{aligned} \quad (5.40)$$

The Jacobian determinant

$$\left| \det \left(\frac{\partial f^{-1}}{\partial x} \right) \right| = \prod_i \sigma_i^{-1}, \quad \text{with} \quad \sigma_i = f_{\sigma_i}(x_{1:i-1}) \quad (5.41)$$

5.2.2 Optimizing a normalizing flow

Bijection transformations of a normalizing flow are parametrized by neural networks of weights φ . We denote as q_φ its probability density function. When using a normalizing flow to fit a distribution p , we need to optimize parameters φ such that q_φ becomes very close to p . This closeness can be measured with either the forward or reverse KL-divergences as presented in Equation 4.15 and Equation 4.16.

Maximizing the ELBO To fit a normalizing flow to a distribution from which we are only able to evaluate the probability density, we need to minimize the reverse KL-divergence. The optimization objective becomes:

$$\varphi^* = \arg \min_{\varphi} D_{\text{KL}}(q_\varphi || p) \quad (5.42)$$

$$= \arg \min_{\varphi} \mathbb{E}_{x \sim q_\varphi} [\log q_\varphi(x)] - \mathbb{E}_{x \sim q_\varphi} [\log p(x)] \quad (5.43)$$

$$= \arg \min_{\varphi} -\text{ELBO}(q_\varphi) \quad (5.44)$$

Therefore, fitting a normalizing flow to a target distribution minimizing the reverse [KL](#)-divergence is equivalent to minimizing the evidence lower bound of the data under the surrogate model.

Minimizing the negative log likelihood When having access to samples $x \sim p$ of the distribution we aim to fit, one can use the forward [KL](#)-divergence. The optimization objective for the normalizing flow is, therefore:

$$\varphi^* = \arg \min_{\varphi} D_{KL}(p||q_{\varphi}) \quad (5.45)$$

$$= \arg \min_{\varphi} \mathbb{E}_{x \sim p} [\log p(x)] - \mathbb{E}_{x \sim p} [\log q_{\varphi}(x)] \quad (5.46)$$

$$= \arg \min_{\varphi} \underbrace{-\mathbb{E}_{x \sim p} [\log q_{\varphi}(x)]}_{\text{NLL}} \quad (5.47)$$

The first term of [Equation 5.47](#) is removed as it does not depend on parameters φ . This objective function is called the *negative log-likelihood* (*Negative Log Likelihood (NLL)*). The $\log q_{\varphi}$ term can be evaluated thanks to the change of variable formula in [Equation 5.27](#) and the expectancy is taken with respect to samples from the target distribution p . This loss function is in particular used in [subsection 7.1.2](#)

5.2.3 Conditional normalizing flows

As discussed in the previous subsection, a normalizing flow enables us to model the density distribution of a random variable x . Noticing that all the presented bijections are formulated as conditional distribution, it is natural to extend the conditioning with another random variable y to model conditional probabilities $p(x|y)$. It is indeed straightforward to model the distribution of a random variable x conditioned on the observation of another random variable y by stacking the conditions. For example the autoregressive factorization in [Equation 5.37](#) can be extended in a conditional manner as:

$$p(x|y) = \prod_{i=1}^d p(x_i|x_{1:i-1}, y), \quad (5.48)$$

which can be implemented feeding y to every function f_{μ_i} and f_{σ_i} . This also applies to RealNVPs as they can be seen as a special case of MAFs.

Training a conditional normalizing flow is almost exactly similar to training an unconditional one, except we need to consider the conditional variable in the loss function. Later in this manuscript, we will describe the generative model for galaxy morphologies, in which we used a conditional normalizing flow to model the latent space distribution. As we aimed to train the [NF](#) from samples of the latent space distribution, with the corresponding conditional variable, we could use the negative log-likelihood loss as described above. In the case of a conditional [NF](#), the loss is also an [NLL](#), but where the expectancy is taken with respect to the joint distribution $p(x, y)$. Indeed, starting from minimizing the [KL](#)-divergence between the posterior $p(x|y)$ and the surrogate model $q_{\varphi}(x|y)$ for all condition y we get:

$$\varphi^* = \arg \min_{\varphi} \mathbb{E}_{y \sim p(y)} [D_{\text{KL}}(p(x|y) || q_{\varphi}(x|y))] \quad (5.49)$$

$$= \arg \min_{\varphi} \mathbb{E}_{y \sim p(y)} \left[\mathbb{E}_{x \sim p(x|y)} [\log p(x|y)] - \mathbb{E}_{x \sim p(x|y)} [\log q_{\varphi}(x|y)] \right] \quad (5.50)$$

$$= \arg \min_{\varphi} - \mathbb{E}_{y \sim p(y)} \left[\mathbb{E}_{x \sim p(x|y)} [\log q_{\varphi}(x|y)] \right] \quad (5.51)$$

$$= \arg \min_{\varphi} - \mathbb{E}_{x, y \sim p(x, y)} [\log q_{\varphi}(x|y)]. \quad (5.52)$$

5.3 Score-based diffusion models

In this section, we present another class of generative models, called *score-based* or *diffusion* models. This kind of generative model was first introduced by Sohl-Dickstein et al. [Soh+15] and later pushed to the state of the art in image generation tasks by Song et al. [SE19] and Ho et al. [HJA20]. Score-based and diffusion model differ slightly, notably in their training objective. However, they share the same fundamental idea of learning a revert corruption process, which leads to the mapping of a target distribution to a well-characterized distribution such as a multivariate Gaussian. In this section, we will only cover score-based generative model as presented in Song et al. [SE19; Son+20b] as it is the specific class we used to develop our methods.

5.3.1 Stochastic differential equations

Diffusion models share a similar spirit with normalizing flows in the sense that they map an arbitrary distribution to a well-known distribution such as a multivariate Gaussian. Diffusion models work by randomly corrupting data with noise, up to having a purely well-characterized noise and learning the reverse process. The forward process can be expressed as a Stochastic Differential Equation (Stochastic Differential Equation (SDE))

$$dx = f(x, t)dt + g(t)dW, \quad (5.53)$$

where x is the variable of interest, evolving with time t , such that $x(0)$ corresponds to the data point of the distribution of interest p_0 , and $x(T) \sim p_T$ the end of the forward process. W is a stochastic process, often chosen to be the Wiener process, describing Brownian motion. $f(\cdot, t) : \mathcal{X} \rightarrow \mathcal{X}$ is called the *drift* function and $g(t) : \mathbb{R} \rightarrow \mathbb{R}$ is called the *diffusion* function. When discretized, this can also be seen as a Markov process, where each iteration $x(t)$ is sampled from a distribution conditioned on the previous time step $t - 1$, i.e. $x(t) \sim q(\cdot | x(t - 1))$.

The core principle of diffusion models is to learn how to invert this forward process in order to map data randomly sample from the distribution that follows $x(T)$ in order to get new samples from the target distribution $x(0)$. As there are several ways to formulate the forward process, reversing the process involves different tasks and objects to be modeled. Considering the SDE formulation in Equation 5.53, Anderson [And82] showed that the reverse process also follows an SDE of form:

$$dx = [f(x, t) - g(t)^2 \nabla_x \log p_t(x)] dt + g(t)d\bar{W}, \quad (5.54)$$

where \bar{W} is also a stochastic process with time going from T to 0, and dt a infinitesimal negative time step. This reverse SDE involves the same term as the

forward, but also the *score function* $\nabla_x \log p_t(x)$. Learning the score function and running an **SDE** to sample new data points is called *score-based generative model*.

In our work [Rem+20; Ram+20; Rem+23], we opted to utilize the score-based model paradigm due to its ability to evaluate the exact likelihood and score functions. This feature makes it well-suited for integration with other likelihoods that incorporate analytical physical knowledge and to be used as part of **MCMC** algorithms.

5.3.2 Denoising score matching

The aim of this section is to describe how one can learn the score function $\nabla_x \log p(x)$. Approximating the score can be done either directly with a function $s_\theta(x)$ or approximating the score of an approximate probability density function $\nabla_x \log p_\theta(x)$ ($= s_\theta(x)$).

$$\mathcal{L}(\theta) \propto \mathbb{E}_{x \sim p_{\text{data}}} \left[\|\nabla_x \log p_{\text{data}}(x) - s_\theta(x)\|^2 \right] = D_F(p_{\text{data}} \| p_\theta). \quad (5.55)$$

This expression is also known as the *Fisher divergence* D_F . This loss function necessitates evaluating $\nabla_x \log p_{\text{data}}(x)$, which is currently unknown. The only available information we have are samples x , implicitly drawn from p_{data} from which we aim to learn to score.

Hyvärinen [Hyv05] proposed to reformulate this loss function, depending only on the approximated score s_θ and samples $x \sim p_{\text{data}}$. By performing an integration by part, one can get the following loss function:

$$\mathcal{L}(\theta) \propto \mathbb{E}_{x \sim p_{\text{data}}} \left[\text{tr}(\nabla_x s_\theta(x)) - \|s_\theta(x)\|^2 \right]. \quad (5.56)$$

One can, in principle, learn the score function of any arbitrary distribution with this loss function. However, this requires evaluating the trace of the score Jacobian which in practice is heavily computationally demanding with respect to the dimension of the data. This term can at best be computed with a number of backpropagation proportional to the dimension of the data $\mathcal{O}(d)$ [MSS12].

To solve this complexity issue, Song et al. [Son+20a] proposed to approximate the Jacobian trace using random projection of the score, by introducing the sliced score matching loss:

$$\mathcal{L}(\theta) \propto \mathbb{E}_{\substack{v \sim p_v \\ x \sim p_{\text{data}}}} \left[v^T (\nabla_x s_\theta(x)) v - \frac{1}{2} (v^T s_\theta(x))^2 \right]. \quad (5.57)$$

p_v can be a arbitrary distribution with the conditions $\mathbb{E}_{p_v}[vv^T]$ positive definite and $\mathbb{E}_{p_v}[\|v\|_2^2] < +\infty$. Multivariate standard normal ($\mathcal{N}(0, \mathbb{I}_d)$), a multivariate Rademacher distributions (the uniform distribution over $\{\pm 1\}^d$) are such distributions. Equation 5.57 is also called the *sliced Fisher divergence*.

In the next section, we will show that for efficient sampling of the data distribution it is much more convenient to evaluate the score of the distribution which has been convolved with a Gaussian kernel.

In order to learn directly the score of the convolved distribution, Vincent [Vin11] and Alain et al. [AB13] demonstrated that one can change the training objectives from fitting the scores to denoising data at different amount of noise corruption levels. Let $r(x) : \mathcal{X} \rightarrow \mathcal{X}$ denote a Gaussian denoiser, and ϵ be a random variable sampled from a multivariate Gaussian $\mathcal{N}(0, \sigma^2 \cdot \mathbb{I}_d)$.

If the denoiser r_θ is trained under the \mathcal{L}_2 (or mean-squared-error loss):

$$\mathcal{L}_{\text{DAE}}(\theta) = \mathbb{E}_{\substack{x \sim p_{\text{data}} \\ \epsilon \sim \mathcal{N}(0, \sigma^2 \mathbb{I}_d)}} \left[\|r(x + \epsilon) - x\|^2 \right] \quad (5.58)$$

then, we can relate the optimal denoiser r_{θ^*} to the score function $\nabla_x \log p(x)$ and the noise standard deviation σ with the following expression:

$$r_{\theta^*}(x) = x + \sigma^2 \nabla_x \log p_\sigma(x) + \mathcal{O}(\sigma^2), \quad (5.59)$$

as $\sigma \rightarrow 0$. As it is formulated, the learned score model in practice may not be optimal for all noise level. Lim et al. [Lim+20] identified several sources of error and proposed to reformulate Equation 5.58 to mitigate them.

5.4 Annealed sampling

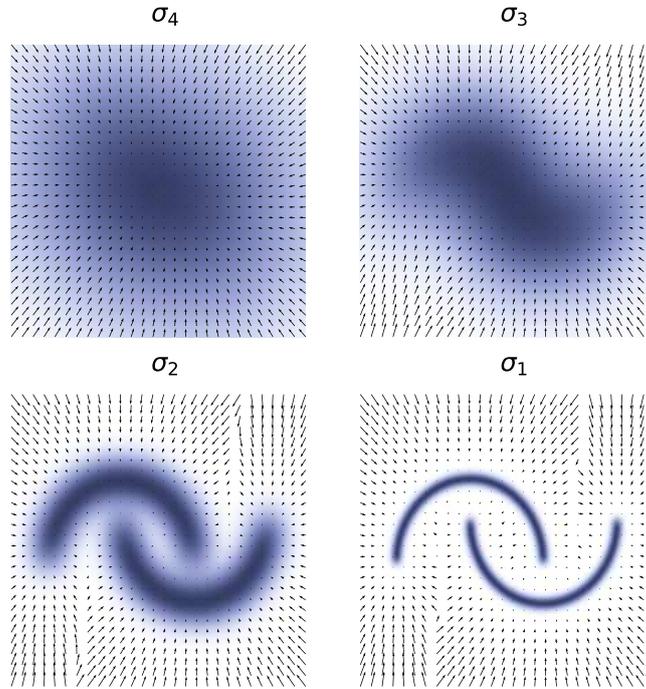


Figure 5.5: Example of annealing on the two moons distribution. The distribution is convolved with a multivariate Gaussian of variance σ . This variance decreases over the sampling procedure so that the MCMC is always over high-density regions, i.e. $\sigma_4 > \sigma_3 > \sigma_2 > \sigma_1$. The chain is initialized from a wide multivariate Gaussian, as on the top left panel, and converges to the data distribution, as on the bottom right panel. In blue is the density of the convolved distribution and the arrows represent its score function.

Following our preliminary work [Rem+20], we adopt in Remy et al. [Rem+23] a two-step sampling procedure. The first step is based on an *annealed* version the Hamiltonian Monte Carlo (HMC) [Nea11; Bet18], similar to the annealed Langevin Dynamics proposed by [SE19]. The second step corresponds to a projection of the

target distribution using an Ordinary Differential Equation (ODE) to reach zero temperature.

Sampling with annealed HMC In this step, we initialize a chain using white Gaussian noise with a high temperature σ_T^2 . Leveraging the conditional noise property of the score function described in section [subsection 5.3.2](#) we have direct access to the score function of the convolved distribution $\nabla \log p_{\sigma_T^2}$, which we can use in a score-based HMC. We then let the chain evolve under Hamiltonian dynamics, and progressively lower the temperature σ^2 of the conditional score with a geometric schedule of common ratio equal to 0.98. Each time the temperature σ^2 is decreased, the MCMC chain thermalizes to the new temperature in a few HMC steps. As described in Song et al. [SE19], small enough steps are important to ensure non-zero transition probabilities between adjacent temperature values. Once the chain has reached sufficiently low temperature (ideally $\sigma^2 = 0$) we stop the chain and *only retrieve the last sample*. Multiple independent samples are obtained by running multiple independent annealed HMC chains in parallel. One must not think that having one chain per sample makes the process much longer than having multiple samples per chain. It is indeed much more efficient because in practice it is very long to ensure sample independence within the same chain.

Projection to zero temperature In practice, however, it has proven to be difficult to anneal chains all the way to zero temperature. We find that at low temperatures, the chains do not properly thermalize at each step and residual noise remains in our samples. This was also observed by other authors using annealed Langevin Dynamics instead of HMC [Jol+20; SE19]. In the application to mass-mapping presented in [chapter 6](#), we can reach $\sigma = 10^{-3}$ by fine-tuning our annealing scheme.

Therefore, in the second step of our sampling procedure, we propose a strategy to transport the annealed HMC samples obtained in step 1 at finite temperature all the way to $\sigma^2 = 0$. To achieve this, we use remarkable results from Song et al. [Son+20b] which establishes a parallel between denoising diffusion models (generative models based on random walks) and Stochastic Differential Equations. We refer the interested reader to [Son+20b] for more mathematical details and derivations, but the main result that we use from this paper is the following Ordinary Differential Equation:

$$dx = -\frac{1}{2} \sqrt{\frac{d}{dt}(\sigma^2(t))} \nabla_x \log p_t(\mathbf{x}) dt, \quad (5.60)$$

where $\sigma^2(t)$ is the increasing variance schedule of the Gaussian distribution we convolve the data distribution with. In our particular implementation, we used a linear schedule $t \mapsto \sigma^2(t) = t$. This ODE describes a deterministic process $\{\mathbf{x}(t)\}_{t=0}^T$ indexed by a continuous time variable $t \in [0, T]$, such that $\mathbf{x}(0) \sim p_0$ and $\mathbf{x}(T) \sim p_T$, where p_0 denotes the data distribution and p_T the convolution between the data distribution and a multivariate Gaussian of variance T . This ODE can be solved by any black-box ODE solver provided that the convolved score is available. In particular, it means that if we start the ODE at a point in p_t , where t is the intermediate temperature at which we have stopped our HMC chains, we can transport that point to p_0 . This procedure very effectively removes residual noise while making sure we reach a point in the target distribution at zero temperature.

* * *
* *
*

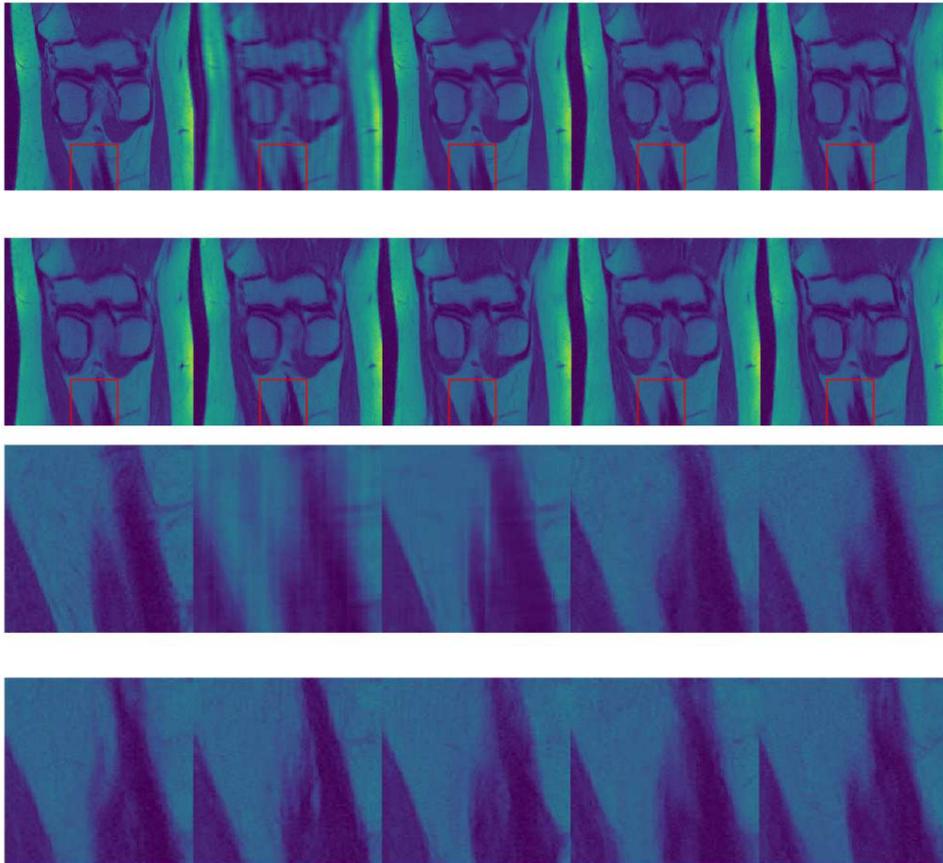


Figure 5.6: Bayesian posterior sampling for Magnetic Resonance Imaging (MRI) reconstruction. The top leftmost image is the ground truth image. The top second to the left image is the zero-filled retrospectively under-sampled image. The top third to the left image is the reconstruction of the under-sampled image by the UPDNet architecture. All the other images are denoised samples from the estimated posterior distribution obtained by a tempered HMC. The zero-filling achieves a Peak Signal-to-Noise Ratio (PSNR) of 25.55 dB, each sample 27.63 dB on average, the mean of the samples 30.04 dB, and the neural network 32.15 dB. A zoom of the region in the red square is provided in Appendix A.4.

Part III

Contributions to weak lensing inverse problems

Weak lensing mass-mapping

Chapter Outline

6.1	Weak lensing Mass-Mapping	74
6.1.1	Shear and convergence	74
6.1.2	Mass-mapping inverse problem in a Bayesian framework	75
6.2	COSMOS dark matter mass-map reconstruction	77
6.2.1	COSMOS survey	77
6.2.2	Simulating COSMOS lensing catalog from κ TNG suite of simulations	79
6.3	Literature review on priors	80
6.3.1	Flat prior: Kaiser-Squires Reconstruction	80
6.3.2	Gaussian Random Field: Wiener Filter	81
6.3.3	Sparse priors	82
6.3.4	Deep implicit prior	82
6.3.5	Score-matching prior	83
6.4	High-dimensional posterior sampling	84
6.4.1	The posterior annealing problem	85
6.4.2	Mass-maps posterior sampling	86
6.5	Posterior moments validation	88
6.5.1	Gaussian posterior	88
6.5.2	Full posterior moments against moments network	89
6.5.3	Point estimates validations	92
6.6	Reconstruction of the COSMOS field	93
6.6.1	HSC/ACS COSMOS posterior mass maps	93
6.6.2	Cluster detection with posterior samples	94
6.6.3	Going further	94

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

Remy, B., Lanusse, F., Jeffrey, N., Liu, J., Starck, J.-L., Osato, K., Schrabback, T., ‘Probabilistic mass-mapping with neural score estimation’. In: *Astronomy & Astrophysics* 672 (2023), A51

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

Remy, B., Lanusse, F., Ramzi, Z., Liu, J., Jeffrey, N., Starck, J.-L., ‘Probabilistic Mapping of Dark Matter by Neural Score Matching’. In: *3rd Machine Learning and the Physical Science workshop at NeurIPS (2020)*

Ramzi, Z., **Remy, B.**, Lanusse, F., Starck, J.-L., Ciuciu, P., ‘Denoising score-matching for uncertainty quantification in inverse problems’. In: *Deep inverse workshop at NeurIPS (2020)*

THIS chapter introduces our novel method for reconstructing dark matter mass-maps. We begin by framing mass-mapping as an ill-posed inverse problem, highlighting the linear relationship between convergence and shear fields with masked data and shape noise corruption. We then reframe this problem within a Bayesian framework, focusing on the posterior distribution of the convergence, conditioned on shear field measurements. To achieve this, we must select an informative prior distribution for the convergence, and the following section delves into a literature-based discussion of these priors. While classical methods offer point estimates to solve the problem, our approach is to infer the complete posterior distribution, thereby capturing the reconstruction’s uncertainty. Inference for such a high-dimensional posterior distribution is very challenging. Therefore, we present how annealed Hamiltonian Monte Carlo can be employed to address this inference problem. Lastly, we demonstrate our method by applying it to reconstruct the highest-resolution mass-map of the [HST/ACS COSMOS](#) field, complete with uncertainty quantification

6.1 Weak lensing Mass-Mapping

6.1.1 Shear and convergence

Galaxy surveys like *Euclid*, the Vera C. Rubin Observatory [LSST](#), or the Roman Space Telescope, will measure billions of galaxy shapes. This measure will serve as an estimate of the cosmic shear field γ . The cosmic shear field is a tracer of the large-scale-structures of the Universe and can be linearly related to the projected matter density field κ , called the convergence field through the equation

$$\tilde{\gamma}_1 + i\tilde{\gamma}_2 = \left(\frac{k_1^2 - k_2^2}{k^2} + i \frac{2k_1k_2}{k^2} \right) (\tilde{\kappa}_E + i\tilde{\kappa}_B), \quad (6.1)$$

or in a vectorized form

$$\begin{pmatrix} \tilde{\gamma}_1 \\ \tilde{\gamma}_2 \end{pmatrix} = \frac{1}{k^2} \underbrace{\begin{pmatrix} k_1^2 - k_2^2 & -2k_1k_2 \\ 2k_1k_2 & k_1^2 - k_2^2 \end{pmatrix}}_{\mathbf{P}} \begin{pmatrix} \tilde{\kappa}_1 \\ \tilde{\kappa}_2 \end{pmatrix} \quad (6.2)$$

where κ_E and κ_B are respectively the E- and B- mode of the convergence field, γ_1 and γ_2 the two components of the shear field adopting the complex representation. We also recall that the operator \mathbf{P} is unitary, i.e. $\mathbf{P}^*\mathbf{P} = \mathbb{I}_d$. In the following, we will refer to \mathbf{P} as the *forward operator*.

In the following paragraphs, we shed light on the various challenges the process of reconstructing the underlying convergence field from the observed weak gravitational

lensing shear field arises. Furthermore, we present our novel approach, which addresses and resolves these inherent issues, leading to a more accurate and physically motivated mass-map reconstruction.

Survey mask All surveys have a delimited sky area for their measurements. For instance, the [HST/ACS COSMOS](#) field that we study in this chapter has a sky area of 1.6 deg^2 . To this date, the mass-map covering the largest sky fraction, about 4143 deg^2 is reconstructed from the Dark Energy Survey third year (Y3) weak lensing data [[Jef+21](#)]. When solving the inverse problem, we therefore need to take into consideration the fact that the survey area is limited, with potentially an irregular contour. Moreover, some of the inner sections of the survey area are not usable due to very bright stars saturating the [CCDs](#) or image artifacts preventing the reliable measurement of galaxy shapes in some regions of the survey.

Vanishing B mode. It is physically motivated to assume that there is no leakage between the E- and B- modes and that there is no contribution to the convergence B-mode component. Nevertheless, intrinsic alignment can contribute to non-zero B-mode [[Bla+19](#); [Sam+19](#)], but is neglected in this work.

Uncertainty quantification Mass-map reconstruction is an ill-posed inverse problem, meaning that multiple convergence maps can fit to a given shear map measurement. This is primarily due to the fact that the presence of the survey mask induces missing data, therefore there are regions of the sky where there is no data to drive the reconstruction. Additionally, the shear measurements rely on galaxy shapes average making the measurements subject to significant shape noise which introduces substantial uncertainties in the reconstruction. State-of-the-art mass mapping algorithms are designed to provide the maximum a posteriori solution, so maximizing the probability of the solution given measurements. However, they do not provide uncertainty quantification associated to the solution. There is, therefore, a need for a method which is able to not only recover a precise convergence map but also able to quantify the uncertainty of the reconstruction. From a Bayesian perspective, this would mean being able to estimate the full posterior distribution of the mass-mapping inverse problem.

Full cosmological prior There is no full analytical description of convergence maps. As we will see in [section 6.3](#), we can build informative prior from the expected knowledge of the 2pt statistics of the field but only model the Gaussian linear structures of the field. To model the non-linear gravitational collapse of the matter density field, we can however rely on N-body simulations which can be seen as *samples* draw from an implicit underlying prior distribution.

6.1.2 Mass-mapping inverse problem in a Bayesian framework

The ill-posed mass-mapping inverse problem can be reformulated as a probabilistic inference problem from a Bayesian perspective as explained in [section 4.1](#):

$$p(\kappa|\gamma) = \frac{p(\gamma|\kappa)p(\kappa)}{p(\gamma)}, \quad (6.3)$$

We detail each term of this equation according to the mass-mapping inverse problem.

Posterior $p(\kappa|\gamma)$ It models the probability of a convergence map κ conditioned on the measured shear field γ . This is the distribution of interest as it models the uncertainties of the inverse problem solutions.

Likelihood $p(\gamma|\kappa)$ It encodes the forward model from the convergence to the shear in equation Equation 6.1. In our work, the forward process returns a binned shear map γ , where each pixel corresponds to the shear field averaged over the pixel area and takes as input a pixelized convergence map κ . Working with real data, it is assumed that the measurement for each pixel is degraded by shape noise n_s due to the finite average of galaxy intrinsic ellipticities in the bin. This noise is assumed to be white Gaussian, i.e. $n_s \sim \mathcal{N}(0, \Sigma_n)$, where Σ_n is the noise covariance matrix of the shear map. Moreover, regions of the map containing missing data need to be taken into account in the likelihood. One obvious way of incorporating a mask within the likelihood is to multiply the data and the forward operator by a binary matrix \mathbf{M} composed of ones and zeros whether the pixel does or does not have data. In this case, the likelihood operator is no longer unitary and loses nice mathematical properties needed for our annealed derivations in section 5.4. The other way to incorporate a mask in the likelihood is to consider infinite variance for pixels with missing data. In practice, we set a very high variance, such as 10^{10} , for these pixels in the covariance matrix Σ_n . More specific information on the strategy we follow to emulate the COSMOS shape catalog can be found in section 6.2.

Thus, the log-likelihood takes the following form:

$$\log p(\gamma|\kappa) = -\frac{1}{2}(\gamma - \mathbf{F}^* \mathbf{P} \mathbf{F} \kappa)^\dagger \Sigma_n^{-1} (\gamma - \mathbf{F}^* \mathbf{P} \mathbf{F} \kappa) + \text{constant}, \quad (6.4)$$

where \mathbf{F} and \mathbf{F}^* are respectively direct and inverse Fourier transform. Because the Fourier operator \mathbf{F} is unitary, i.e. $\mathbf{F}^* \mathbf{F} = \mathbb{I}_d$, it is equivalent to refer to the likelihood in real or in Fourier space:

$$\log p(\gamma|\kappa) \propto -\frac{1}{2}(\gamma - \mathbf{F}^* \mathbf{P} \mathbf{F} \kappa)^\dagger \Sigma_n^{-1} (\gamma - \mathbf{F}^* \mathbf{P} \mathbf{F} \kappa) \quad (6.5)$$

$$= -\frac{1}{2} \mathbf{F} (\gamma - \mathbf{F}^* \mathbf{P} \tilde{\kappa})^\dagger (\mathbf{F}^* \Sigma_n \mathbf{F})^{-1} \mathbf{F} (\gamma - \mathbf{F}^* \mathbf{P} \tilde{\kappa}) \quad (6.6)$$

$$= -\frac{1}{2} (\tilde{\gamma} - \mathbf{P} \tilde{\kappa})^\dagger \tilde{\Sigma}_n^{-1} (\tilde{\gamma} - \mathbf{P} \tilde{\kappa}), \quad (6.7)$$

where $\tilde{\Sigma}_n$ is no longer diagonal in Fourier space.

Prior $p(\kappa)$ It encodes the knowledge about the convergence κ . There is no full analytical description of the convergence map prior. We will discuss various choices of priors in the next section.

Bayesian evidence $p(\gamma)$ It is the marginal density of the observations. The evidence is a constant if we assume a given model, and will be ignored in the rest of this work as we do not consider Bayesian model comparison. In this work, we assume $p(\gamma) = p(\gamma|\mathcal{M})$, where \mathcal{M} is Λ CMD with Planck parameters described in subsection 2.1.3.

All existing mass-mapping techniques can be understood under the lens of this Bayesian formulation and will mostly differ in their choice of prior, and in the specific algorithm used to recover a point estimate of the convergence map. As in practice,

this problem is ill-posed due to noise corruption and missing data in Equation 6.1, the posterior $p(\kappa|\gamma)$ can be both wide and heavily prior dependent, which explains why all these different techniques yield different answers.

6.2 COSMOS dark matter mass-map reconstruction

In this chapter, we will demonstrate our mass-mapping algorithm and compare it to existing methods in the literature on the reconstruction of the *HST/ACS* COSMOS field. In the following, we discuss various information about this field such as the choice of the shear catalog, survey mask, and the field resolution. To validate our methods, we will also mock the COSMOS field with hydrodynamical N-body simulations.

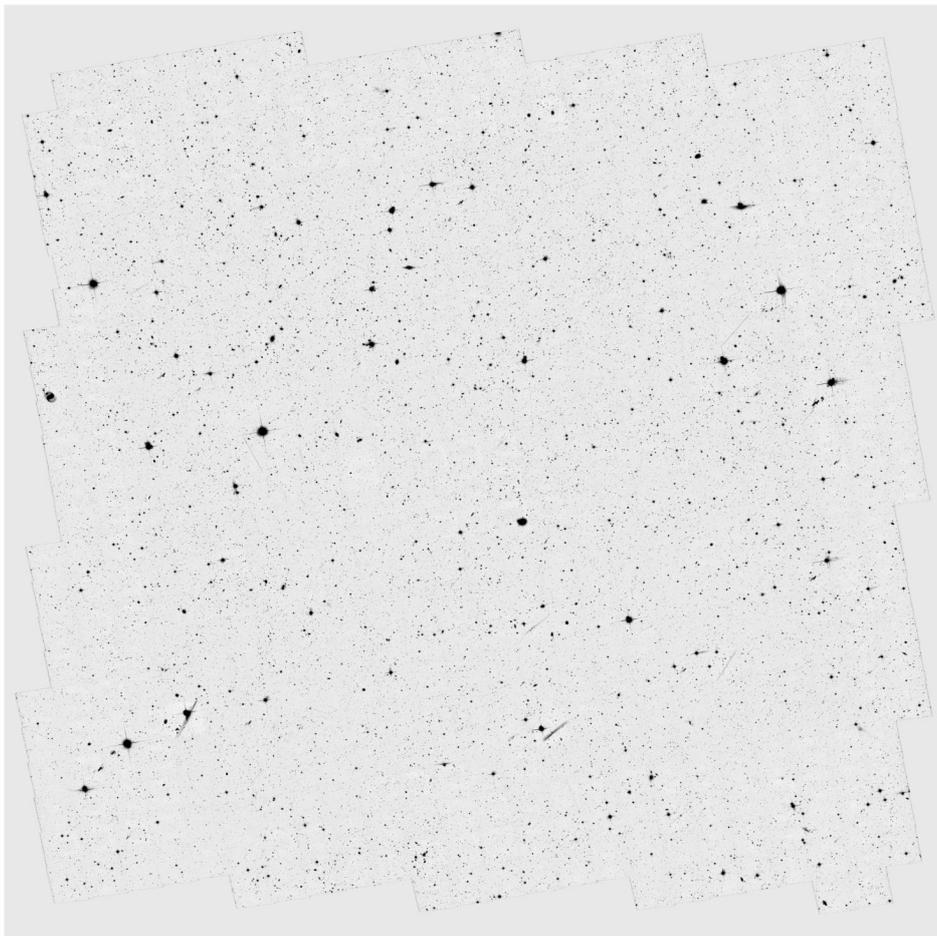


Figure 6.1: COSMOS *HST/ACS* field [Sco+07c]. This mosaic, representing the combination of all the *ACS* pointings, is taken from [Koe+07].

6.2.1 COSMOS survey

The Cosmic Evolution Survey (*COSMOS*) [Sco+07c] covers a field measuring 1.64 square degrees and has been imaged using the Advanced Camera for Surveys (*ACS*)

aboard the Hubble Space Telescope (HST). The ACS has captured a high-resolution image of this field, which can be observed in Figure 6.1. Since the release of this field, numerous shape measurement methods discussed in section 3.2 have been employed, resulting in the creation of catalogs galaxy galaxy shapes [e.g. Mas+07c]. In our study, we used the shape catalog obtained in Schrabback et al. [Sch+10] (Schrabback 2010 catalog (S10)), using the KSB+ method [KSB95; Hoe+98; Sch+07]. The spatially and temporally varying PSF of the ACS camera was modeled using principal component analysis interpolation. A step of multiplicative bias calibration (see section 3.4) was performed employing a power law with respect to the SNR fitted on simulations.

The catalog from [Sch+10] is divided into a bright $i^+ \leq 25$ (Subaru SExtractor MAG_AUTO magnitude) and a faint $i^+ > 25$ galaxy sample. To estimate galaxy redshifts for the bright sample, individual high-quality photometric redshifts are available by cross-matching against the COSMOS-30 catalog [Ilb+08]. For the faint catalog, Schrabback et al. [Sch+10] proposed a functional form for the overall $n(z)$ distribution, based on an extrapolation to fainter magnitudes of the i_{814} -redshift relation observed in the $23 \leq i_{814} \leq 25$ range. In our analysis, we combine both bright and faint galaxy samples into a single-shape catalog. The redshift distribution of the bright, faint, and combined samples are illustrated in Figure 6.2. The only cut we applied was rejecting galaxies in the bright sample with $z_{\text{phot}} < 0.6$ and $i^+ > 24$. The photometric redshift cut is motivated by S10 finding indications that many of these galaxies are truly at high redshifts. Thus, their inclusion would imply that the used estimate of the redshift distribution is inaccurate. This yields a total of 417117 galaxies. This redshift distribution is particularly useful to create mocks of lensing catalogs from dark matter simulations (see subsection 6.2.2).

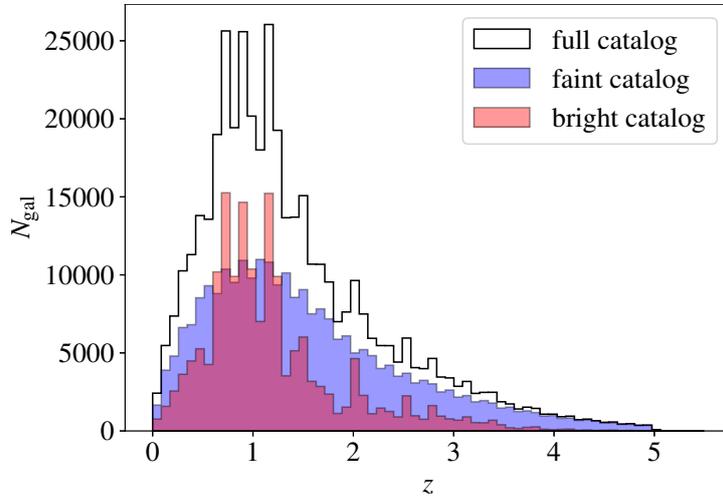


Figure 6.2: Redshift distribution of the COSMOS catalog from Schrabback et al. [Sch+10]. The redshifts for bright galaxies ($i^+ < 25$) are derived from the COSMOS-30 catalog [Ilb+08], while the $n(z)$ distribution for faint galaxies ($i^+ > 25$) is based on an extrapolation to fainter magnitudes as described in Schrabback et al. [Sch+10].

To match the resolution of the κ TNG convergence maps simulations discussed in the next section, we binned the catalog into pixels of resolution 0.29 arcmin per pixel, which translates into 64.2 galaxies per arcmin² on average. Due to insufficient resolution caused by factors such as very bright stars saturating the CCDs or image

artifacts that prevent the reliable measurement of galaxy shapes in some regions of the survey, we defined a binary mask capturing the limits of the survey as well as the missing data. One can see the catalog mask in [Figure 6.10](#) (bottom left), used in the mock and [COSMOS S10](#) analyses. This binning also yields a shape noise variance map. The variance of each pixel can be defined from the overall variance of shapes in the catalog σ_e , rescaled by the number of galaxies per pixel N_i . For empty pixels, we assumed a very large variance to express the lack of data to constrain the likelihood. This gives:

$$\Sigma_n = \text{diag}(\sigma_1^2, \dots, \sigma_d^2), \quad \text{where} \quad \sigma_i^2 = \begin{cases} \frac{\sigma_e^2}{N_i} & \text{if } N_i > 0 \\ 10^{10} & \text{otherwise} \end{cases} \quad (6.8)$$

Figure 6.3: Mask from [S10](#) catalog of the [HST/ACS COSMOS](#) field.

6.2.2 Simulating COSMOS lensing catalog from κ TNG suite of simulations

We present in this section the κ TNG simulations we used to learn a mass-map prior and to create mock data for validation of the method. The κ TNG suite of ray-traced weak lensing mocks maps generated from the IllustrisTNG (TNG) hydrodynamical simulations [[Nel+19](#); [Pil+18](#); [Nel+18](#); [Spr+18](#); [Nai+18](#); [Mar+18](#)]. The TNG simulations are a set of cosmological, large-scale gravity and magneto-hydrodynamical simulations, where baryonic processes such as stellar evolution, chemical enrichment, gas cooling, supernovae, and black hole feedback are incorporated as subgrid models. Therefore, These simulations go beyond Gaussian Random Fields by accurately resolving the small-scale, nonlinear gravitational evolution of the matter density field, providing a wealth of additional information.

κ TNG maps are generated through ray tracing of the trajectories of light-rays from redshift $z = 0$ to the target source redshift. A large number of realizations are generated by randomly translating and rotating the snapshots. Eventually, the full set includes 10,000 realizations of $5 \times 5 \text{ deg}^2$ convergence maps for 40 source redshift up to $z_s = 2.6$, at a 0.29 arcmin pixel resolution. We combined these source planes to match the [COSMOS](#) redshift distribution in [Figure 6.2](#). All of the source redshifts higher than $z_s = 2.6$ were recycled to the last source plane $z = 2.6$ with the appropriate weight designed to match the expected lensing kernel. This recycling procedure is described in [Figure 6.4](#). All of TNG snapshots and κ TNG convergence maps were generated assuming a flat Λ -cold dark matter cosmological model at the Planck 2015 [[Ade+16](#)] cosmology, i.e. with Hubble constant $H_0 = 67.73 \text{ km s}^{-1} \text{ Mpc}^{-1}$ baryonic density $\Omega_b = 0.0486$, matter density $\Omega_m = 0.3089$, spectral index of scalar perturbations $n_s = 0.9667$, and amplitude of matter fluctuations at $8 \text{ h}^{-1} \text{ Mpc}$ $\sigma_8 = 0.8159$, assuming massless neutrinos, with an effective number of neutrino species $N_{\text{eff}} = 3.046$. Because κ TNG maps are computed by ray-tracing through the IllustrisTNG hydrodynamic simulations, there is no need for post-Born correction. Though the current κ TNG do not incorporate intrinsic alignment, we can address the effect using halo or galaxy shape [[Shi+23](#); [Kur+21](#)].

According to this procedure, we built a dataset of convergence maps to train our score network, and we also built mocked lensing catalog for validation study

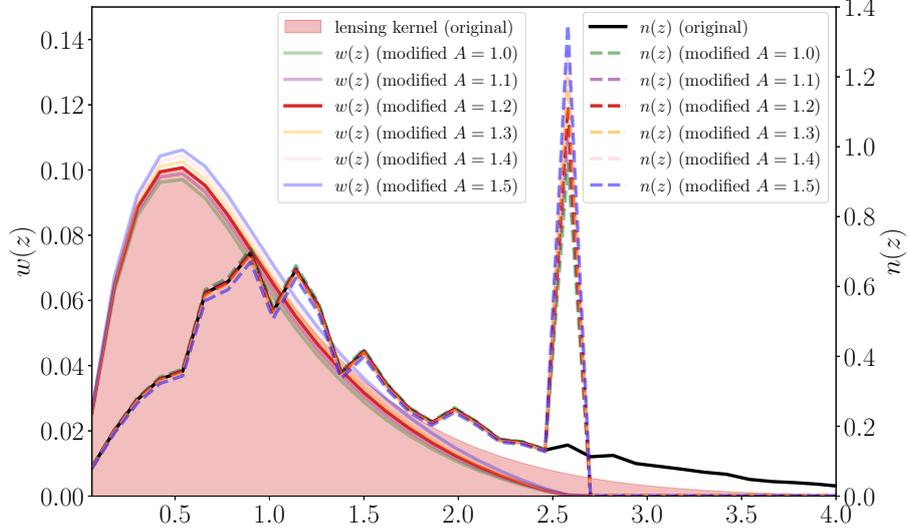


Figure 6.4: This figure shows how the convergence maps above redshift $z_s = 2.6$ are recycled to match the expected lensing kernel assuming the full redshift distribution. The original redshift distribution $n(z)$ is shown in black and the recycled redshift distributions in dashed lines. The redshift density at $z = 2.6$ where scaled by a factor A , and the corresponding lensing kernels $w(z)$ are plotted in plain lines. From this study $A = 1.2$ is the best scaling factor fitting the expected lensing kernel.

before applying our method on the **S10** catalog. We first applied the Kaiser-Squires transformation to transform the convergence maps to shear maps and then corrupted the field with artificial shape noise adding Gaussian noise to each pixel bin according to Equation 6.8.

6.3 Literature review on priors

In the peculiar case of dark matter convergence maps, there is no analytical formulation for the prior. Indeed, generating accurate convergence maps requires to solve N-body simulations such as Liu et al. [Liu+18, MassiveNuS] or Osato et al. [OLH21, κ TNG]¹ and then to use ray-tracing to project the 3D matter density field in a 2D convergence map. The purpose of this section is therefore to discuss various prior assumptions in the literature such as flat (Kaiser-Squires), Gaussian Random Field (Wiener filter), sparse (MCA-lens, GLIMPSE) or learned from N-body simulations with deep learning methods (Deep-mass or our contribution based on neural score-matching).

6.3.1 Flat prior: Kaiser-Squires Reconstruction

The Kaiser-Squires method [KS93] can be seen as a simple maximum likelihood estimate (Maximum Likelihood Estimate (MLE)) of the convergence map, typically

¹<http://columbialensing.org/>

followed by a certain amount of Gaussian smoothing.

$$\check{\kappa}_{ks} = \arg \min_{\kappa} \|\gamma - \mathbf{F}^* \mathbf{P} \mathbf{F} \kappa\|_2^2 = (\mathbf{F}^* \mathbf{P} \mathbf{F})^\dagger \gamma \quad (6.9)$$

$$\kappa_{ks} = s * \check{\kappa}_{ks} \quad (6.10)$$

where s is a Gaussian smoothing kernel of a given scale, and $(\mathbf{F}^* \mathbf{P} \mathbf{F})^\dagger$ is a pseudo-inverse of the operator $\mathbf{F}^* \mathbf{P} \mathbf{F}$. While this method is the fastest, it does not take into account masks and leads to leakage between E and B modes of the convergence field. Applying the pseudo inverse operator to the shear map yields a noisy convergence map, this is why we apply a smoothing kernel. For Kaiser-Squires, the heteroscedasticity does not impact the solution, whereas it can for certain extensions such as the Generalized Kaiser-Squires method [Sta+21, Appendix B.1], in which an iterative approach with little regularization that takes the mask and noise heteroscedasticity into account. This approach remains very limited to reconstruct a precise convergence map because the likelihood is highly dominated by the noise. This motivates the use of a physically motivated prior, which would help in reconstructing the features hidden by the likelihood noise.

6.3.2 Gaussian Random Field: Wiener Filter

The Wiener filter [Wie49] solution to the mass-mapping inverse problem assumes a Gaussian Random Field prior on κ and takes advantage of the fact that the power spectrum of the convergence can be analytically predicted from cosmological models, and accurately describes the field on large scales. This prior on the convergence can be expressed as a Gaussian distribution with a diagonal covariance matrix \mathbf{S} in Fourier space:

$$p_{\text{GRF}}(\kappa) = \frac{1}{\sqrt{\det 2\pi \mathbf{S}}} \exp\left(-\frac{1}{2} \tilde{\kappa}^\dagger \mathbf{S}^{-1} \tilde{\kappa}\right), \quad (6.11)$$

where \mathbf{S} is the convergence power spectrum, and GRF stands for Gaussian Random Field.

The solution of the inverse problem can be formulated as:

$$\hat{\kappa}_{\text{wiener}} = \arg \min_{\kappa} \|\Sigma_n^{-1/2}(\gamma - \mathbf{F}^* \mathbf{P} \mathbf{F} \kappa)\|_2^2 - \log p_{\text{GRF}}(\kappa). \quad (6.12)$$

This Wiener solution corresponds to the *maximum a posteriori* (Maximum a Posteriori (MAP)) solution under this Gaussian prior, and also matches the mean of the Gaussian posterior. An appealing property of this estimator is that the solution can be easily recovered analytically in the case where the noise is homoscedastic, as both signal and noise covariance matrices become diagonal in Fourier space. The Wiener Filter reconstruction [Lah+94; Zar+95] is given in Fourier space by:

$$\tilde{\kappa} = \mathbf{S} \mathbf{P}^\dagger \left[\mathbf{P} \mathbf{S} \mathbf{P}^\dagger + \tilde{\Sigma}_n \right]^{-1} \tilde{\gamma}, \quad (6.13)$$

where \mathbf{S} and $\tilde{\Sigma}_n$ are respectively the signal and noise covariance matrix in Fourier space.

In more complex cases, where the noise covariance is not diagonal in Fourier space (for instance because of a mask in pixel space), the solution can still be recovered efficiently using optimization algorithms, e.g. with proximal methods [Bob+12; Sta+21], or by means of a messenger field [EW13]. One can also draw samples from the Wiener posterior with the messenger field algorithm [Jef+18].

6.3.3 Sparse priors

Convergence maps contain non-Gaussian features, that are not well recovered with the methods described above. Several mass-mapping algorithms have been proposed, relying on a wavelet sparsity prior [SPR06; LDS12; Lan+16; Pri+18; Sta+21], which can be formulated as:

$$\log p(x) = \|\Phi^t x\|_p \quad (6.14)$$

with $p < 2$, and where Φ is a wavelet dictionary and $\|\cdot\|_p$ is a sparsity promoting ℓ_p norm.

The convergence map is the solution of the following sparse recovery optimization problem:

$$\hat{\kappa} = \arg \min_{\kappa} \|\Sigma^{-1/2}(\gamma - \mathbf{F}^* \mathbf{P} \mathbf{F} \kappa)\|_2^2 + \lambda \|\Phi^t \kappa\|_p, \quad (6.15)$$

where λ is the regularisation parameter, weighting the sparse regularisation constraint. The GLIMPSE method [Lan+16], allows in addition to take into account masks, non-uniform noise, flexion data if they are available, and also does not require to transform the shear catalog on pixelized map. For further details on the design of sparse prior for weak lensing mass-mapping, we refer the reader to François Lanusse’s thesis [Lan15].

6.3.4 Deep implicit prior

While all the priors described above have closed-form expressions, it is also possible to design a mass-mapping method where the prior is defined *implicitly*. The first dark matter map reconstruction from weak lensing observational data using deep learning was shown in [Jef+20a]. The method, called DEEPMASS, is a Convolutional Neural Network (Convolutional Neural Network (CNN)) trained on pairs of simulated pixelized shear and convergence maps. One can show that under some assumptions, a deep learning model can estimate the mean of the posterior distribution. In a nutshell, the network needs to be trained to minimize the mean squared error (Mean Squared Error (MSE)) of the output convergence κ and the training convergence and shear maps must be drawn respectively from the prior $p(\kappa)$ and likelihood distribution $p(\gamma|\kappa)$. We make the derivation of this result later in subsection 6.5.2.

While the Wiener Filter assumes a Gaussian prior over the convergence, simulated training data for DEEPMASS are drawn from the *full* non-Gaussian prior $p(\kappa)$ and thus improve the accuracy of the reconstruction. DEEPMASS is therefore able to recover the non-linear structures of the convergence better than the Wiener filter and to reduce the MSE of the reconstruction.

Even if DEEPMASS reconstructs high quality convergence maps, DEEPMASS alone *only provides the mean posterior and cannot quantify the uncertainties of the reconstruction*. Furthermore, as any direct inversion method based on neural networks, the likelihood Equation 6.4 is learned implicitly by the model and does not explicitly constrain the solution at inference time. Meaning that although we have not found in our experiments obvious failures, the CNN *may* in theory fails in ways that would lead to a map not consistent with observations, for instance, creating spurious artifacts, or missing structures present in the true map. Another side effect of implicitly learning the likelihood during training is that the model is trained for a specific survey configuration, and retraining the network is required if either the survey mask or the noise covariance matrix is different.

In this work, we propose a new approach which estimates the full posterior distribution $p(\kappa|\gamma)$, being able to not only recover the posterior mean, but also to quantify the uncertainties of the reconstruction. In addition, in our method the likelihood is explicit, meaning that it does not require a retraining for a new survey configuration.

6.3.5 Score-matching prior

Hybrid denoiser As explained in [subsection 5.3.2](#), it is possible to learn the score of the underlying distribution of dark matter mass maps from cosmological simulations. In theory, this prior can be arbitrarily complex and the denoising score matching method should be able to model it. In practice, we are limited by the constraining power of the architecture and the neural network can have difficulties to learn the prior in all of its complexity. We propose therefore to use the physically motivated GRF prior as a *first guess* and let the neural network learn the residual model. As the GRF prior is supposed to be accurate at large scales and capturing the 2pt statistics of the field, the residual prior should then focus all of its modeling power to capture the non-linear scales of the field.

As we aim to learn the residual score with denoising score matching, we need to learn a residual denoiser where the first guess would be performed by the Gaussian Random Field (GRF) prior. To this end, we need to feed the neural network the GRF prior score $\nabla_{\kappa} \log p_{\text{GRF}}(\tilde{\kappa})$ along with the noisy $\tilde{\kappa}$ and train it the the denoising score matching loss in [Equation 5.58](#). We call this loss Residual Denoising Score-Matching (Residual Denoising Score-Matching (RDMSM))

$$\mathcal{L}_{\text{RDMSM}}(\theta) = \mathbb{E}_{\substack{u \sim \mathcal{N}(0, \mathbb{I}_d) \\ \sigma \sim \mathcal{N}(0, s)}} \|u + \sigma r_{\theta}(\kappa + \sigma u, \nabla_{\kappa} \log p_{\text{GRF}}(\kappa + \sigma u))\|_2^2 \quad (6.16)$$

In addition, it is interesting to see how the GRF prior acts as a denoiser. We can define a Gaussian denoiser from the GRF prior defined in [Equation 6.11](#) and the relation between scores and denoisers in [Equation 5.59](#).

$$r_{\text{GRF}}(\kappa, \sigma) = \kappa + \sigma^2 \nabla_{\kappa} \log p_{\text{GRF}, \sigma}(\kappa), \quad (6.17)$$

where $r_{\text{GRF}}(\tilde{\kappa}, \sigma)$ denotes de GRF denoiser with input $\tilde{\kappa}$ with the knowledge of the noise level σ and $\log p_{\text{GRF}, \sigma}(\tilde{\kappa}) = \int p_{\text{GRF}}(\kappa) q_{\sigma}(\tilde{\kappa}|\kappa) d\kappa$, $q_{\sigma}(\tilde{\kappa}|\kappa)$ being a multivariate Gaussian distribution of variance σ^2 and $\tilde{\kappa} = \kappa + \sigma \epsilon$ with $\epsilon \sim \mathcal{N}(0, \mathbb{I}_d)$.

[Figure 6.5](#) is an illustration of the comparison between the GRF prior and hybrid prior on denoising tasks. We can observe that the GRF denoiser performs well at removing the noise but is only able to reconstruct the large scales of the map. This behavior is expected since this prior is completely analytical and encoded by the power spectrum of the map which does not capture the statistics of the non-linear scales.

Training details The general approach to training score models with Denoising Score Matching was described in [subsection 5.3.2](#). We describe in this paragraph the actual implementation details which enable us to learn the parametric denoiser r_{θ} . We used a 3-scale U-net architecture inspired from [\[RFB15\]](#), which broadly speaking is an autoencoding neural network with skip connections as schemed in [Figure 6.6](#). It is composed of Resnet [\[He+15\]](#) blocks of 2 convolutional layers, each block being followed by batch normalization layers. The downsampling is performed by average pooling and upsampling by interpolation using nearest neighbor. The sequence of channels of

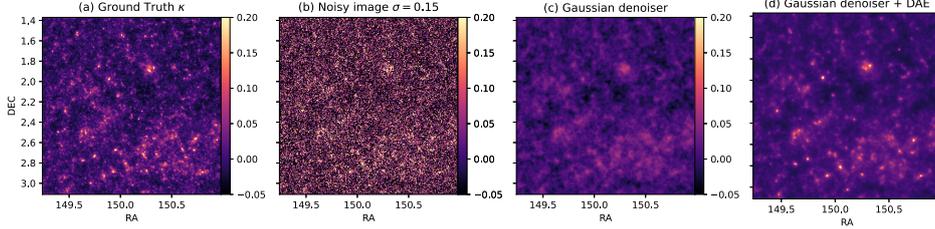


Figure 6.5: Comparison between (a) ground truth convergence maps κ , (b) noisy convergence map $\tilde{\kappa} = \kappa + \sigma u$, (c) GRF denoiser estimation $r_{\text{GRF}}(\tilde{\kappa}, \sigma)$ and the hybrid denoiser.

the the different scales of the Unet was $[3, 32, 64, 128, 128]$, where 3 corresponds to the input of the network. As we are learning a score model, there are many regions in the high-dimensional space \mathbb{R}^d where there is no training data, but where we will be likely to explore during the sampling procedure. This is why we used tree-based spectral normalization (SN) as described in subsection 5.1.6 to enforce the lipschitzness of the network and therefore constrain the estimation in empty regions to align with well-constrained scores towards regions of high densities of the target distribution.

Our architecture takes multiple input images of size 360×360 which are stacked into channels. First the noisy convergence map $\tilde{\kappa}$ 360×360 pixels² corresponds to 1.64 deg^2 , which means a resolution of $0.29 \text{ arcmin/pixel}$. Second, the noise level of the map σ multiplied elementwise to a 360×360 array of ones. It is important to notice here that the nature of the noise is isotropic Gaussian and can be summarized in its standard deviation which we inform the neural network. Astrophysical noise like shape noise is involved in the likelihood covariance matrix, not in the prior. Third, the GRF prior score estimate from computed from Equation 6.11 $\nabla_{\kappa} \log p_{\text{GRF}, \sigma}(\tilde{\kappa})$. As advised in [SE19], we divided the output of the network by the the noise level σ . This is because it is observed that the norm of the score function is inversely proportional, which the network suffers to rescale automatically and is necessary when dealing with large orders of magnitude of noise during the annealing.

Regarding the optimization procedure, we used the Adam optimizer [KB14] with a ramp learning schedule, starting from 10^{-4} and decreasing up to 10^{-7} . The loss was minimized using mini-batch of size 32 for 4×10^5 optimization steps.

6.4 High-dimensional posterior sampling

In this section, we describe the algorithm to sample the mass-mapping posterior distribution. As we aim to perform sampling in a very high dimensional space, we propose to use a state-of-the-art sampling scheme inspired by diffusion-based image generation models. We first show that while these methods are very efficient at sampling from prior distributions, they fail at sampling posteriors when we consider an analytical likelihood distribution, even in very simplistic cases. We then describe how to transform the likelihood distribution by annealing to perform annealed HMC and solving the high-dimensional mass-mapping posterior inference.

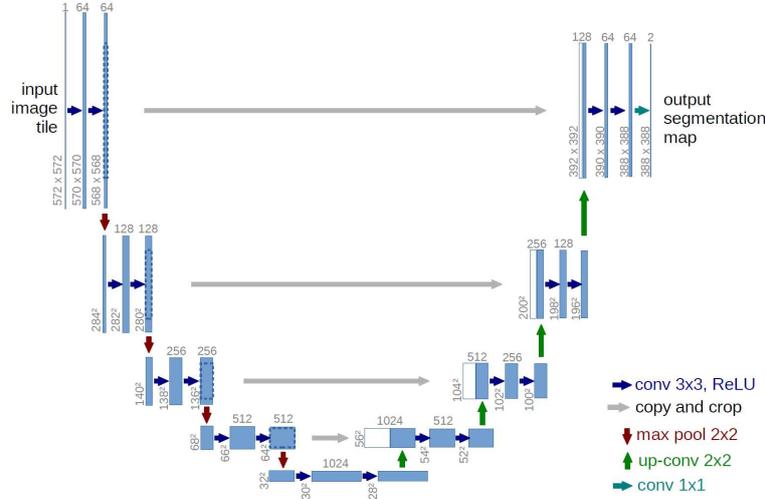


Figure 6.6: Scheme of the U-net architecture. Credit: Ronneberger et al. [RFB15].

6.4.1 The posterior annealing problem

It is natural to decompose the score of the posterior as the sum of the likelihood and prior score according to Bayes' rule:

$$\nabla_x \log p(x|y) = \nabla_x \log p(y|x) + \nabla_x \log p(x). \quad (6.18)$$

It is however much less trivial to define the annealed posterior score for arbitrary temperature t such that

$$\nabla_{x(t)} \log p(x(t)|y) = \nabla_{x(t)} \log p(y|x(t)) + \nabla_{x(t)} \log p(x(t)). \quad (6.19)$$

Indeed using the reverse-SDE formulation in Equation 5.54 requires particular properties of the distribution $p(x(t)|y)$. Assuming a Gaussian diffusion process, we need to verify that $p(x(t)|y) = \int p(x|y)p(x|x(t))dx$, where $p(x|x(t))$ is a Gaussian distribution. However, we only have access to the analytical likelihood $p(y|x)$ and to the annealed prior score $\nabla_{x(t)} \log p(x(t))$. One trivial operation would be to anneal the likelihood with the same kernel as for the prior, i.e. $p(y|x(t)) = \int p(y|x)p(x|x(t))dx$, but this has the major issue of leading to the annealed posterior with a Gaussian kernel.

This phenomenon appears for any type of likelihood and prior even when they are Gaussian, because the convolution of a product of distribution densities is not the product of the convolution of the densities, i.e. for an arbitrary distributions $p(x) = p_1(x)p_2(x)$, and g :

$$(p * g)(x) \neq (p_1 \otimes g)(x)(p_2 \otimes g)(x) \quad (6.20)$$

or in score terms

$$\nabla_x \log(p \otimes g)(x) \neq \nabla_x \log(p_1 \otimes g)(x) + \nabla_x \log(p_2 \otimes g)(x). \quad (6.21)$$

As a consequence, we cannot use the annealed posterior score $\nabla_{x(t)} \log p(x(t)|y)$ as defined in Equation 6.19 in the reverse SDE or ODE presented in subsection 5.3.1 to sample the posterior distribution. We show in Figure 6.7 that running an ODE

with the annealed score leads to a wrong estimation of a 1-dimensional posterior distribution, and it also fails to estimate the correct posterior of mass-maps assuming a GRF prior in subsection 6.5.1.

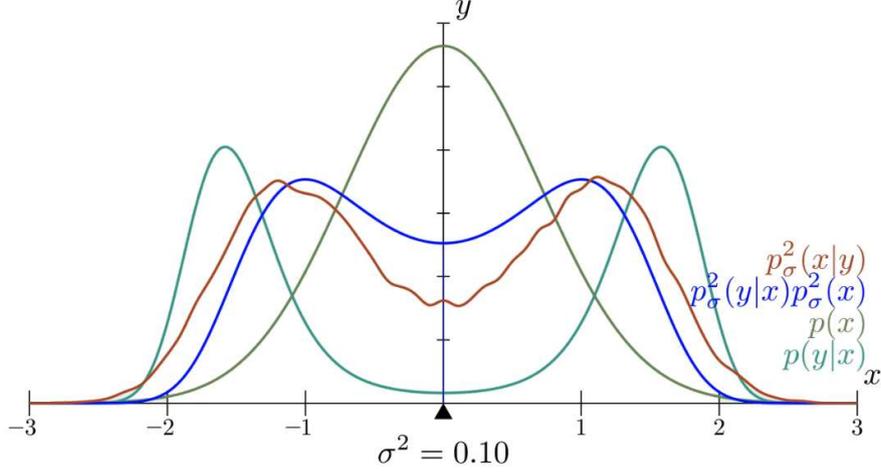


Figure 6.7: Annealing jointly the likelihood and the prior is not equivalent to annealing the posterior distribution. In this example, the prior distribution $p(x)$ is in green and the likelihood $p(y|x)$ in grey. Sampling with the reverse-SDE (Equation 5.54) by annealing both the likelihood and the prior yields the distribution in blue, while a correct annealing of the posterior should yield the the distribution in red.

6.4.2 Mass-maps posterior sampling

According to the previous section, to perform annealed HMC on the posterior distribution, we need to anneal both the likelihood and the prior. The denoising score matching prior allows to have access to the score of the annealed prior by construction. It therefore remains to anneal the likelihood with a Gaussian kernel, i.e. such that $p(y|x(t)) = \int p(y|x)p(x|x(t))dx$, where $p(y|x)$ is defined in Equation 6.4. As both the likelihood and the convolutional kernel are Gaussians, the convolved density is also Gaussian. The presence of the operator $\mathbf{F}^*\mathbf{P}\mathbf{F}$ makes it a bit more subtle, but we can show the following proposition:

Proposition: let $x \in \mathbb{R}^d$, $y \in \mathbb{R}^d$, $p_{\sigma_1}(y|x) \triangleq \mathcal{N}(y | \mathbf{P}x, \sigma_1^2 \mathbb{I}_d)$ being the likelihood and $p_{\sigma_2}(x) \triangleq \mathcal{N}(x | 0, \sigma_2^2 \mathbb{I}_d)$ a centered multivariate Gaussian distribution. Let us assume that the operator \mathbf{P} is unitary, i.e. verifies $\mathbf{P}^\dagger \mathbf{P} = \mathbb{I}_d$.

Then, the convolution of the likelihood with the centered Gaussian is:

$$p_{\sigma_1} \otimes p_{\sigma_2}(y|x) = \mathcal{N}(y | \mathbf{P}x, (\sigma_1^2 + \sigma_2^2) \mathbb{I}_d). \quad (6.22)$$

Proof: According to the definition of p_{σ_1} and p_{σ_2} , we have:

$$\begin{aligned} p_{\sigma_1} \otimes p_{\sigma_2}(y|x) &= \int p_{\sigma_1}(y|x-t)p_{\sigma_2}(t)dt \\ &= \frac{1}{(2\pi\sigma_1^2\sigma_2^2)^d} \int \exp\left(-\frac{1}{2\sigma_1^2}((y-\mathbf{P}(x-t))^\dagger(y-\mathbf{P}(x-t)))\right) \\ &\quad \times \exp\left(\frac{t^\dagger t}{2\sigma_2^2}\right) dt, \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{(2\pi\sigma_1^2\sigma_2^2)^d} \int \exp \left(-\frac{1}{2\sigma_1^2} \left(\|y - \mathbf{P}x\|_2^2 + 2y^\dagger \mathbf{P}t \right. \right. \\
&\quad \left. \left. - 2x^\dagger \mathbf{P}^\dagger \mathbf{P}t + t^\dagger \mathbf{P}^\dagger \mathbf{P}t \right) \right) \exp \left(\frac{t^\dagger t}{2\sigma_2^2} \right) dt \\
&= \frac{1}{(2\pi\sigma_1^2\sigma_2^2)^d} \int \exp \left(-\frac{1}{2\sigma_1^2} \left(\|y - \mathbf{P}x\|_2^2 + 2y^\dagger u \right. \right. \\
&\quad \left. \left. - 2x^\dagger \mathbf{P}^\dagger u + u^\dagger u \right) \right) \exp \left(\frac{t^\dagger t}{2\sigma_2^2} \right) dt.
\end{aligned}$$

We can use the change of variable $u = \mathbf{P}t$, and notice that $du = |\det \mathbf{P}| dt = dt$ and $u^\dagger u = t^\dagger \mathbf{P}^\dagger \mathbf{P}t = t^\dagger t$, since \mathbf{P} is unitary, so that:

$$\begin{aligned}
&= \frac{1}{(2\pi\sigma_1^2\sigma_2^2)^d} \int \exp \left(-\frac{1}{2\sigma_1^2} \left(\|y - \mathbf{P}x\|_2^2 + 2u^\dagger (y - \mathbf{P}x) + u^\dagger u \right) \right) \\
&\quad \times \exp \left(\frac{u^\dagger u}{2\sigma_2^2} \right) du \\
&= \frac{1}{(2\pi\sigma_1^2\sigma_2^2)^d} \exp \left(-\frac{\|y - \mathbf{P}x\|_2^2}{2\sigma_1^2} \right) \\
&\quad \times \int \exp \left(-\frac{u^\dagger (y - \mathbf{P}x)}{\sigma_1^2} - u^\dagger u \left(\frac{1}{2\sigma_1^2} + \frac{1}{2\sigma_2^2} \right) \right) du \\
&= \frac{1}{(2\pi\sigma_1^2\sigma_2^2)^d} \exp \left(-\frac{\|y - \mathbf{P}x\|_2^2}{2\sigma_1^2} \right) \times \\
&\quad \int \exp \left(-\frac{1}{2} \left(\frac{1}{2\sigma_1^2} + \frac{1}{2\sigma_2^2} \right) \left[\frac{2u^\dagger (y - \mathbf{P}x)}{2\sigma_1^2} \left(\frac{1}{2\sigma_1^2} + \frac{1}{2\sigma_2^2} \right)^{-1} \right. \right. \\
&\quad \left. \left. + u^\dagger u \right] \right) du \\
&= \underbrace{\exp \left(-\frac{\|y - \mathbf{P}x\|_2^2}{2\sigma_1^2} \right) \exp \left(\frac{\|y - \mathbf{P}x\|_2^2}{2\sigma_1^4} \left(\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2} \right)^{-1} \right)}_{(*)} \times \\
&\quad \underbrace{\int \exp \left(-\frac{1}{2} \left(\frac{1}{2\sigma_1^2} + \frac{1}{2\sigma_2^2} \right) \left[u + \frac{y - \mathbf{P}x}{2\sigma_1} \left(\left(\frac{1}{2\sigma_1^2} + \frac{1}{2\sigma_2^2} \right)^{-1} \right)^2 \right] \right)}_{(**)} du
\end{aligned}$$

$$\begin{aligned}
&= \frac{(2\pi)^{d/2} \left(\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}\right)^{d/2}}{(2\pi\sigma_1^2\sigma_2^2)^d} \exp\left(-\frac{\|y - \mathbf{P}x\|_2^2}{2(\sigma_1^2 + \sigma_2^2)}\right) \\
&= (2\pi)^{-d/2} (\sigma_1^2 + \sigma_2^2)^{-d/2} \exp\left(-\frac{\|y - \mathbf{P}x\|_2^2}{2(\sigma_1^2 + \sigma_2^2)}\right) \\
&= (2\pi)^{-d/2} (\sigma_1^2 + \sigma_2^2)^{-d/2} \exp\left(-\frac{\|y - \mathbf{P}x\|_2^2}{2(\sigma_1^2 + \sigma_2^2)}\right) \\
&= \mathcal{N}(y \mid \mathbf{P}x, (\sigma_1^2 + \sigma_2^2)\mathbf{I}_d)
\end{aligned}$$

□

$$\begin{aligned}
(*) &= \exp\left(-\frac{\|y - \mathbf{P}x\|_2^2}{2} \left(\frac{1}{\sigma_1^2} - \frac{1}{\sigma_1^4} \left(\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}\right)\right)\right) \\
&= \exp\left(-\frac{\|y - \mathbf{P}x\|_2^2}{2(\sigma_1^2 + \sigma_2^2)}\right).
\end{aligned}$$

$$\begin{aligned}
(**) &= (2\pi)^{d/2} \left(\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}\right)^{d/2} \times \\
&\quad \underbrace{\frac{1}{(2\pi)^{d/2} \left(\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}\right)^{d/2}} \int \exp\left(-\frac{1}{2} \left(\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}\right) [u + \mu]\right) du}_{=1}.
\end{aligned}$$

6.5 Posterior moments validation

In the following, we used the HALOFIT matter power spectrum from `jax-cosmo`², using Ade et al. [Ade+16] (Table 4 final column) results, to build the Gaussian prior covariance matrix. In the following, we also refer to it as the *fiducial* power spectrum.

6.5.1 Gaussian posterior

To validate our posterior sampling algorithm we show that we can recover the Wiener Filter by averaging the posterior samples. Indeed, the Wiener Filter is the Maximum A Posteriori (MAP) solution assuming a Gaussian Random Field prior. The likelihood of the problem being Gaussian as well, the posterior is in this case also Gaussian, and therefore, the MAP matches the Posterior mean. There are several methods to recover the Wiener Filter solution to the mass-mapping inverse problem. In Figure 6.9, we compare three of these methods. First, we compute using the messenger field

²https://github.com/DifferentiableUniverseInitiative/jax_cosmo

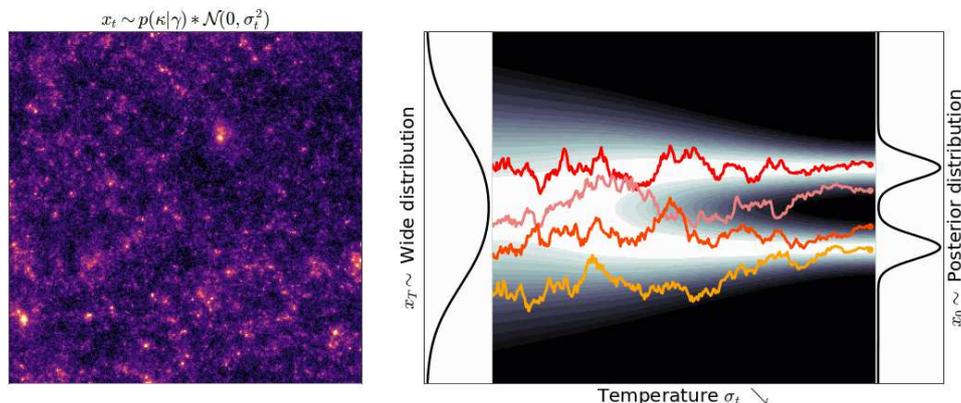


Figure 6.8: **MCMC** chains start from a wide Gaussian distribution and converge to the target multimodal distribution while we anneal the temperature of the convolved Gaussian distribution. This enables the chains to always remain in high-density regions of the distribution. This has two benefits, first, all the final samples from the target distribution are independent of each other due to the independent initial sampling. Second, all the modes of the target distribution are cached, without the classical need to jump from one to the other with traditional samplers. The *left* panel shows the final sample of an annealed **HMC** chain and the *right* panel shows a one-dimensional analogy of our sampling scheme.

algorithm from Elsner et al. [EW13] and use it as a reference to validate the two others. Second, we run a gradient descent optimization (see subsection 5.1.2) method on the posterior distribution to recover the **MAP**. Finally, we sample the posterior using the annealed **HMC** algorithm and average them to recover the posterior mean. To validate that all of these maps match, we compare them both in terms of Root Mean Square Error (**RMSE**) and in terms of power spectrum.

In this setting, we have access to analytical descriptions of the likelihood, prior, and posterior distributions so we are able to check if a posterior sampling algorithm is able or not to recover expected statistics of the field. For example, running the backward SDE fails to sample correctly the posterior distribution in this setting. This is due to the posterior annealing problem discussed in subsection 6.4.1. We illustrate in subsection 6.4.1 that while running an SDE enables to recover a mass-map with the correct power spectrum, averaging the SDE samples does not lead to the expected posterior mean mass-map. We can observe visually that the SDE-mean share similar features with the correct result, but does not have zero signal outside of the survey mask, where only the prior should drive the result.

6.5.2 Full posterior moments against moments network

While our approach allows us to get a point estimate of the posterior through the samples by computing averages or higher moments, it is also possible to use amortized methods and train a neural network as a direct estimator of these moments. It is always what is done when training a neural network on a regression task. Indeed let us take the example of training a neural network under the minimization of an ℓ_2 loss to show how to build an estimator of the posterior mean.

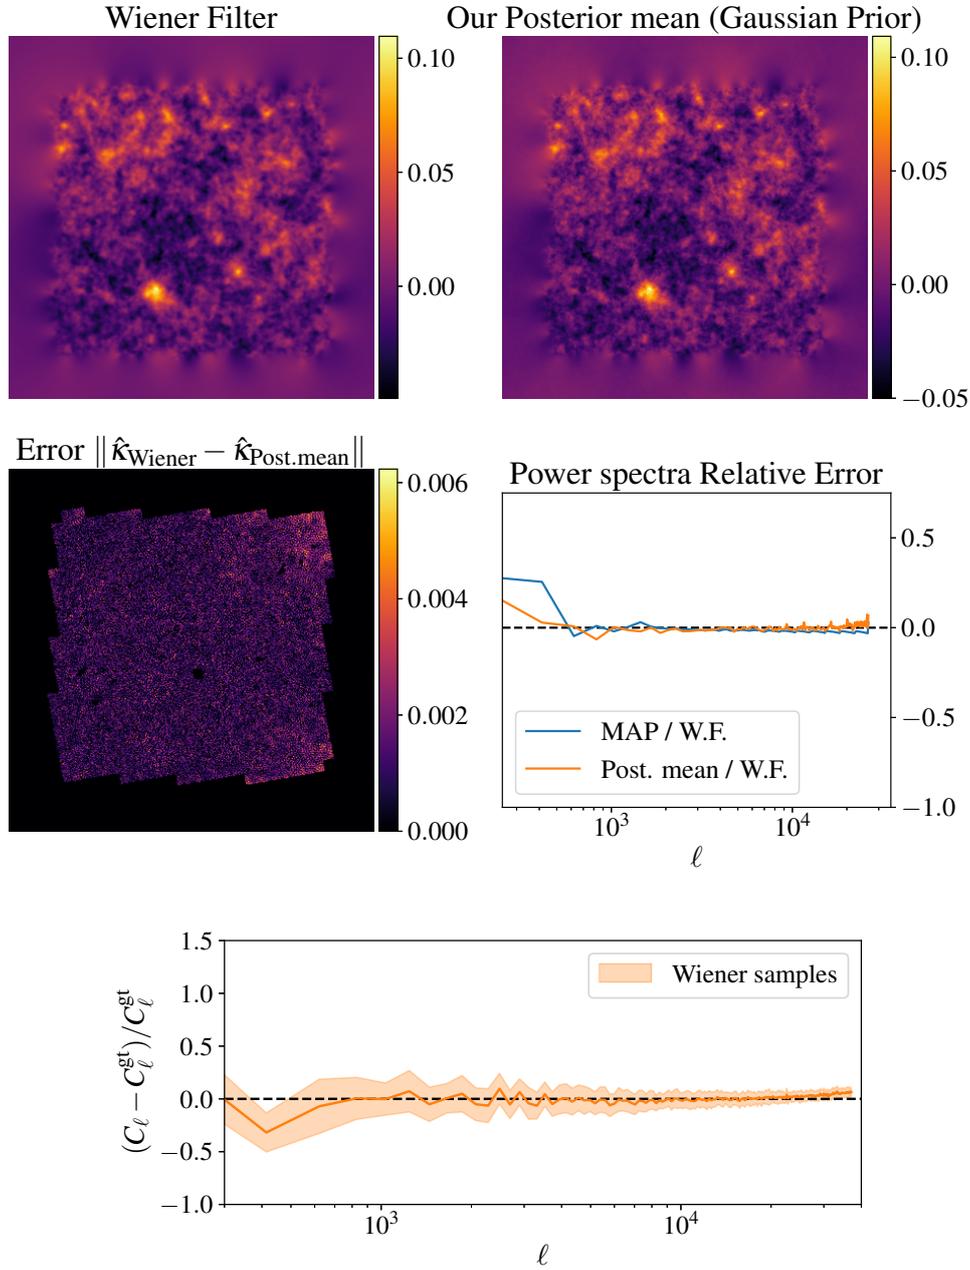


Figure 6.9: Comparison of three different methods to recover the [MAP](#) assuming a [GRF](#) prior. The messenger field algorithm solution is shown on the *top left* and the average of our posterior samples (*top right*). The *bottom left* panel shows the absolute difference between the messenger field reconstruction and our posterior mean. The *bottom right* panel shows the ratio shows the power spectrum relative error of our posterior mean and [MAP](#) using gradient-based optimization against the messenger field algorithm solution. This shows that all of these three methods are in agreement. We also show the power spectrum relative error between our posterior samples and the fiducial power spectrum used to build the [GRF](#) prior.

Let $x \in \mathbb{R}^d$ be the unknown signal we aim to reconstruct and $y \in \mathbb{R}^d$ the observation. Let us introduce a neural network f_θ , where θ are the neural network weights to optimize.

This amortized approach was formalized for inverse problem-solving in [JW20], under the name of *moments network*. This result comes from a known probability theory theorem which can also be found in Adler et al. [AÖ18].

Proposition: if

$$\theta^* = \arg \min_{\theta} \mathbb{E}_{(x,y) \sim p(x,y)} \|x - f_\theta(y)\|_2^2, \quad (6.23)$$

then the neural network is an estimator of the posterior mean

$$\hat{x} = f_{\theta^*}(y) = \int xp(x|y)dx. \quad (6.24)$$

Proof: we derive the proof for measurable Euclidian spaces \mathbb{R}^d .

Let us denote $\mathbb{E}[x|y] = \int xp(x|y)dx$. As training samples are drawn from the joint distribution, i.e. $(x, y) \sim p(x, y)$, the optimization objective is written as

$$\min_{\theta} \mathcal{L}(\theta) = \min_{\theta} \mathbb{E}_{(x,y) \sim p(x,y)} \|x - f_\theta(y)\|_2^2 \quad (6.25)$$

The minimum is found at $\nabla_{\theta} \mathcal{L} = 0$, i.e.

$$\nabla_{\theta} \mathcal{L} = 0 \Leftrightarrow \nabla_{\theta} \int \|x - f_\theta(y)\|_2^2 p(x, y) dx dy = 0 \quad (6.26)$$

$$\Leftrightarrow \int \left(\int \underbrace{\nabla_{\theta} \|x - f_\theta(y)\|_2^2}_{\nabla_{\theta} f_{\theta} \nabla_{f_{\theta}} \|x - f_{\theta}(y)\|_2^2} p(x|y) dx \right) p(y) dy = 0 \quad (6.27)$$

$$\Leftrightarrow \underbrace{\nabla_{\theta} f_{\theta} \nabla_{f_{\theta}}}_{(1)} \underbrace{\int \left(\int \|x - f_\theta(y)\|_2^2 p(x|y) dx \right) p(y) dy}_{(2)} = 0 \quad (6.28)$$

Either (1) or (2) or both are zero. (1) is likely not zero for all architecture so (2) is zero. Moreover $f \mapsto \int \left(\int \|x - f\|_2^2 p(x|y) dx \right) p(y) dy$ is a convex function in f , so has a unique minimum. Let us show that this minimum is the posterior mean.

$$(2) = \nabla_{f_{\theta}} \int \left(\int \|x - f_{\theta}(y)\|_2^2 p(x|y) dx \right) p(y) dy \quad (6.29)$$

$$= \int \left(\int (-2x + 2f) p(x|y) dx \right) p(y) dy \quad (6.30)$$

$$= \int \left(- \int 2xp(x|y) dx + 2f \right) p(y) dy \quad (6.31)$$

$$(f = \mathbb{E}[x|y]) = 0 \quad (6.32)$$

Therefore, the posterior mean is the unique minimizer of the mean square error loss function.

6.5.3 Point estimates validations

In this section, we compare the accuracy of different point estimates compared to a mocked density field mimicking COSMOS. The different approaches of dark matter mass-map reconstruction were described in section 6.3 and are summarized in the following list:

- Kaiser-Squires reconstruction [KS93]
- Wiener Filter [wiener] [EW13; JHF18], maximum/mean a posteriori assuming a Gaussian Random Field prior
- GLIMPSE (3) [Lan+16], MAP assuming a sparse prior
- MCALENS (5) [Sta+21], MAP assuming a sparse prior
- DEEPMASS Jeffrey et al. [Jef+20b], posterior mean fully amortized method
- Score-based posterior mean [Rem+23]

Table 6.1 shows a quantitative comparison of the Kaiser-Squires, the Wiener Filter, MCALENS, GLIMPSE, DEEPMASS and our mass-mapping method, based on two metrics:

- The root mean square error (RMSE) is defined as:

$$\text{RMSE}(\hat{\kappa}, \kappa^{gt}) = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{\kappa}_i - \kappa_i^{gt})^2}, \quad (6.33)$$

where i is the pixel index, $\hat{\kappa}$ the mean-subtracted estimated convergence map, κ^{gt} the mean-subtracted ground truth convergence map we aim to recover. Note that we use here the mean-subtracted convergence map. This is motivated by the mass-sheet degeneracy, which does not constrain the mean of the convergence map.

- The Pearson correlation coefficient r defined as:

$$r(\hat{\kappa}, \kappa^{gt}) = \frac{\text{Cov}(\hat{\kappa}, \kappa^{gt})}{\sigma_{\hat{\kappa}} \sigma_{\kappa^{gt}}}, \quad (6.34)$$

where Cov is the covariance and σ_{κ} is the standard deviation of the convergence map κ .

It turns out that MCALENS, DEEPMASS and our posterior mean both reach the lowest RMSEs and highest Pearson correlations. These results can be expected since both DEEPMASS and DLPOSTERIOR mean are data-driven. DEEPMASS is explicitly trained to minimize the pixel reconstruction error leading to an estimate of the posterior mean, which our method is also targeting. This result shows that our method reaches the state-of-the-art reconstruction of convergence maps.

The fact that MCALENS leads to metrics similar to DEEPMASS and DLPOSTERIOR mean could be explained by the fact that MCALENS modeling, (the convergence map being modeled as the sum of a Gaussian Random Field and a non-Gaussian one being sparse in the wavelet domain) is able to capture well enough the main properties of the convergence map.

Method	RMSE ↓	r ↑
KS ($\sigma_{\text{smooth}}=1$ arcmin)	2.40×10^{-2}	0.57
Wiener filter	2.31×10^{-2}	0.61
GLIMPSE (3)	2.84×10^{-2}	0.42
MCALENS (5)	2.19×10^{-2}	0.67
DEEPMASS	2.18×10^{-2}	0.68
std=1.54e-03 DLPOSTERIOR mean	2.16×10^{-2}	0.68

Table 6.1: Metrics comparison between different mass mapping methods. **RMSE** (the lower the better) is computed according to eq. 6.33 and the Pearson Correlation coefficient r (the higher the better) according to eq. 6.34. For a fair comparison to Kaiser-Squires, we kept the smoothing coefficient that minimized the **RMSE**.

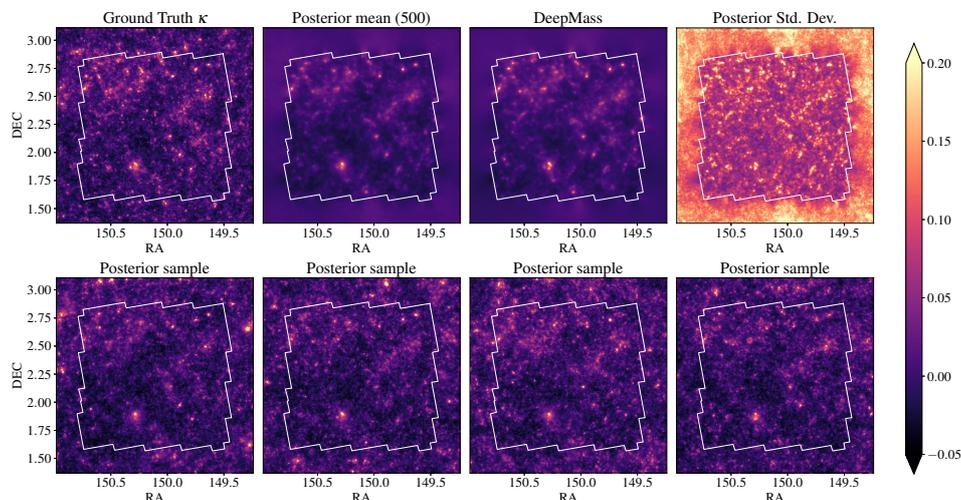


Figure 6.10: Comparison of deep learning-based mass-mapping methods on one mocked **COSMOS** field. *First row*: the ground truth κ TNG convergence map, the mean of our posterior samples (over 400 samples), DeepMass and the standard deviation of our posterior samples. All the maps use the same colorbar, except for the standard deviation, which is displayed in the range $[0, 0.035]$. *Second row*: samples from the posterior distribution, using the hybrid Gaussian-neural score prior. The mask contours corresponds to the **COSMOS** survey boundary mask.

6.6 Reconstruction of the COSMOS field

6.6.1 HSC/ACS COSMOS posterior mass maps

In this section, we now apply our full methodology to the reconstruction of the **COSMOS** field, using the catalog described in section 6.2. The likelihood covariance and the simulation-based prior remain the same as in the validation section.

As a baseline, we show the Kaiser-Squires reconstruction of the **COSMOS** convergence field in the *top-left* panel of Figure 6.12. We applied a Gaussian-smoothing with a variance $\sigma_{\text{smooth}} = 1$ arcmin, chosen such that the RMSE is minimized and Pearson correlation coefficient is maximized on simulated data. Although large scale and small scale structures can already be observed on this map, its power spectrum does not correspond to the fiducial matter power spectrum and no feature at scales smaller

than 1 arcmin can be observed. Moreover, there is not any uncertainty quantification on the reconstruction.

In [Figure 6.12](#) we also present our reconstruction of the COSMOS convergence field, alongside uncertainty quantification. The *top-right* panel shows a posterior sample that looks very similar to a posterior sample from simulated input from κ TNG simulations. Notice that although there is only input shear within the white contours, i.e. in the survey footprint, a complete convergence map is sampled from the posterior distribution. We furthermore validate, in [Figure 6.13](#), the quality of COSMOS posterior samples reconstruction by showing that their power spectra are in good agreement with those of κ TNG convergence maps.

As proposed in [subsection 6.3.4](#) we choose the posterior mean as our estimate for the convergence reconstruction. The *Bottom-left* panel of [Figure 6.12](#) shows the average of 400 samples making the DLPOSTERIOR mean of the COSMOS field. One can visually observe the similar LSS of the field between the Kaiser-Squires and the DLPOSTERIOR mean, while the latter contains much more resolved features, especially concerning the cluster's shape.

6.6.2 Cluster detection with posterior samples

Another way to examine our COSMOS convergence map is to compare to another probe for cluster mass detection. So in the *bottom-left* and *top-right* panels of [Figure 6.12](#), we overlay a subset of X-ray clusters from the Finoguenov et al. [[Fin+06](#)] catalog. Most of the X-ray clusters match to a resolved peak in the convergence field, which is in a way expected since we selected the most massive X-ray clusters.

Alongside the DLPOSTERIOR mean, we also provide the DLPOSTERIOR standard deviation, in the *bottom-right* panel of [Figure 6.12](#). One can observe again that the variance is the highest outside the survey contours since there is no data. Another interesting behavior of the posterior is that the location where the uncertainty is also high is where the signal is the most intense.

The former mass-map of the COSMOS field was published in Massey et al. [[Mas+07a](#)], using a generalized version of the Kaiser-Squires method. When comparing the two maps, although we do not share the same shape catalog or redshift distribution, they share similar mass distributions. However, the COSMOS-DLPOSTERIOR mean is much more resolved. In particular, we can identify several clusters in what looks like one coarse cluster in the Massey et al. [[Mas+07a](#)] mass-map at coordinates RA=150.4, DEC=2.5.

For this detection experiment, we added an NFW profile to the input shear map. We then ran our cluster detection procedure on the associated DLPOSTERIOR samples. [Figure 6.14](#) shows cutouts of those samples around the cluster location. The upper plot shows the samples that were selected with a 3σ threshold and the bottom plot the ones that were not selected. We can recognize the shape of a cluster among the samples with positive detection, i.e. with maximum coefficient above 3σ .

6.6.3 Going further

Having presented the method and results, we propose in this section a discussion on some important points including scientifically relevant use cases, limitations, and possible extensions.

As mentioned earlier in this chapter, taking the average over posterior samples should reduce to the DEEPMASS results. One may wonder about the tradeoff between

the two approaches if the results should be the same. We would argue that the two methods are complementary. DEEPMASS is much faster in producing a convergence map as it only requires a forward pass of the U-net. However, DEEPMASS remains an *amortized* solution, with no strong guarantees on the solution it recovers in practice, due to not having an explicit likelihood. This also means that the entire model needs to be retrained for any change in the lensing catalog (to account for variations in the mask or noise).

We expect that the method presented in this chapter will find its most compelling applications in the study of localized structures through the weak lensing effect, where the benefits of a full pixel-level posterior are the strongest. As a particularly relevant example, we will mention the discussion surrounding the potential detection in the Abell 520 (A520) cluster of a dark clump [Jee+12; Clo+12], a localized peak visible in mass-maps but with no optical counterpart. Quantifying the significance of such a structure using non-linear mass-mapping algorithms targeting a Maximum A Posterior solution (for instance based on sparse regularization) is a difficult task. It was attempted for this particular field using different techniques in Peel et al. [PLS17] and Price et al. [Pri+18], but without strong quantitative statements. The method developed in this chapter however would be able to access the full posterior of the problem.

For cosmological applications, it is however likely that the use cases of the method presented in this work will remain limited. Indeed, for Higher Order Statistics relying on simulations to evaluate their likelihood, the particular mass-mapping technique used is not critical as any systematics due to reconstruction errors induced by the algorithm are calibrated on simulations. A simple Kaiser-Squires inversion should in principle suffice in most cases of interest. From an information-theoretic point of view, the information is preserved by a Kaiser-Squires inversion (in the presence of masks, keeping both E- and B- modes) while any posterior summary may discard some cosmological information. Nevertheless, the ability to sample constrained realizations may find very useful applications such as inpainting masked regions to facilitate the computation of 3pt functions, or void detection algorithms.

We also want to highlight that the method presented in this chapter can be extended in multiple directions. In this configuration, we would be limited to analyzing small-volume surveys. However, we could train a prior on spherical maps, having a spherical likelihood and run analysis and then run large-volume surveys analyses. We are considering here a prior at a fixed cosmology, primarily due to the high computational cost of high-resolution hydrodynamical simulations at different points in cosmological parameter space. Concretely, this means that our solution is heavily biased towards the fiducial cosmology on scales poorly constrained by the data. Note however that the hybrid deep learning + physical Gaussian prior we have introduced in this work provides a natural framework for extending our method to include cosmological dependence. On large scales, the prior is mostly driven by the analytic power spectrum and thus easy on condition on cosmology, while on small scales, only the non-Gaussian residuals are captured by the neural network. This implies that the cosmology-dependent part of the model that needs to be learned on simulations is mostly on small scales. This makes it very likely that in the near future, one could develop such a cosmology-dependent residual prior from suites of numerical simulations of smaller cosmological volume, but spanning a range of cosmological models. The recent multifield Cosmology and Astrophysics with Machine Learning Simulations (CAMELS) simulations [Vil+22] would be an ideal dataset for this purpose. Another avenue for future work would be extending the method to the sphere, which is simply a matter of defining a U-net

for instance using a DeepSphere [Per+19] approach for convolutions on a spherical domain.

* * *
* * *
* * *

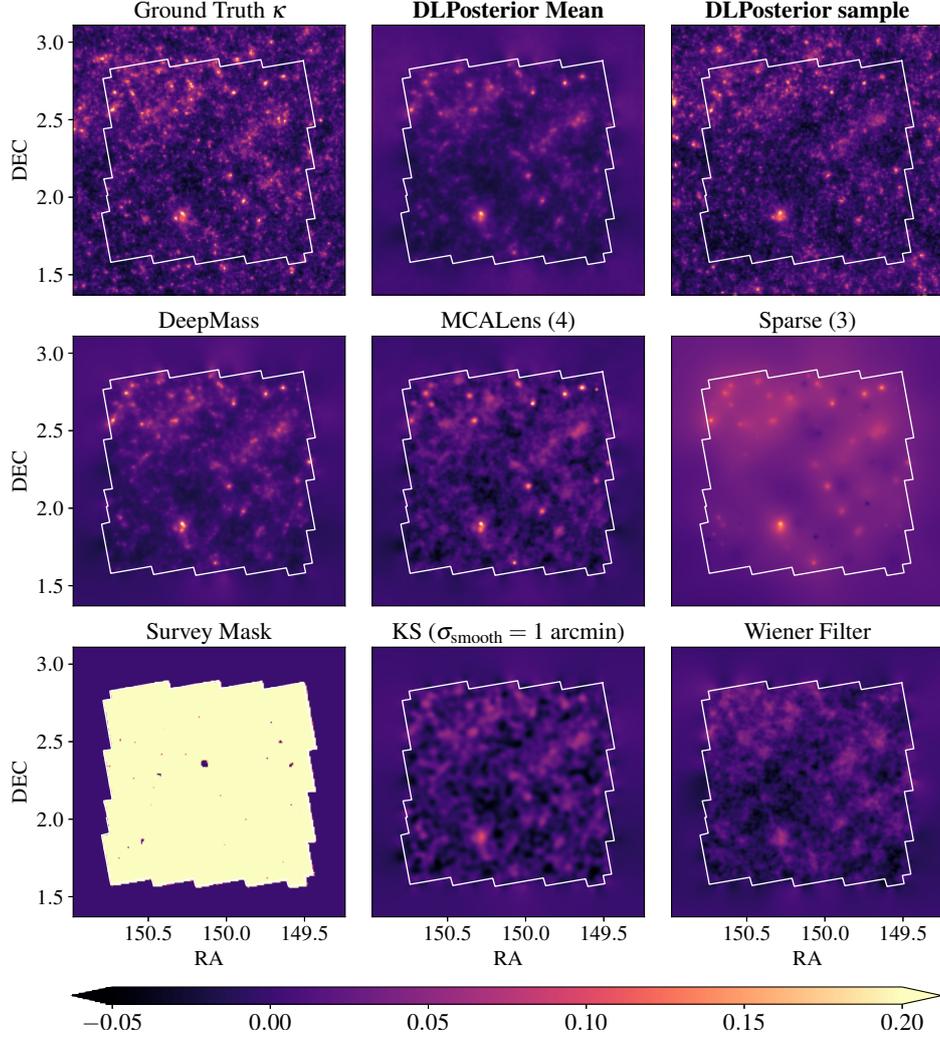


Figure 6.11: Mass-mapping methods comparison. *Ground truth* corresponds to the convergence map κ that we aim to recover, taken from Osato et al. [OLH21], DLPOSTERIOR mean and DLPOSTERIOR sample are the results discussed in section 6.4, DEEPMASS is from Jeffrey et al. [Jef+20b], MCALENS with $\lambda = 4$ is from Starck et al. [Sta+21], GLIMPSE method with $\lambda = 3$ is from Lanusse et al. [Lan+16], *Survey Mask* is the COSMOS catalog binary mask, *Kaiser-Squires* with Gaussian smoothing ($\sigma_{\text{smooth}} = 1$ arcmin) is from Kaiser et al. [KSB95] and the *Wiener Filter* is from Elsner et al. [EW13] and Jeffrey et al. [JHF18].

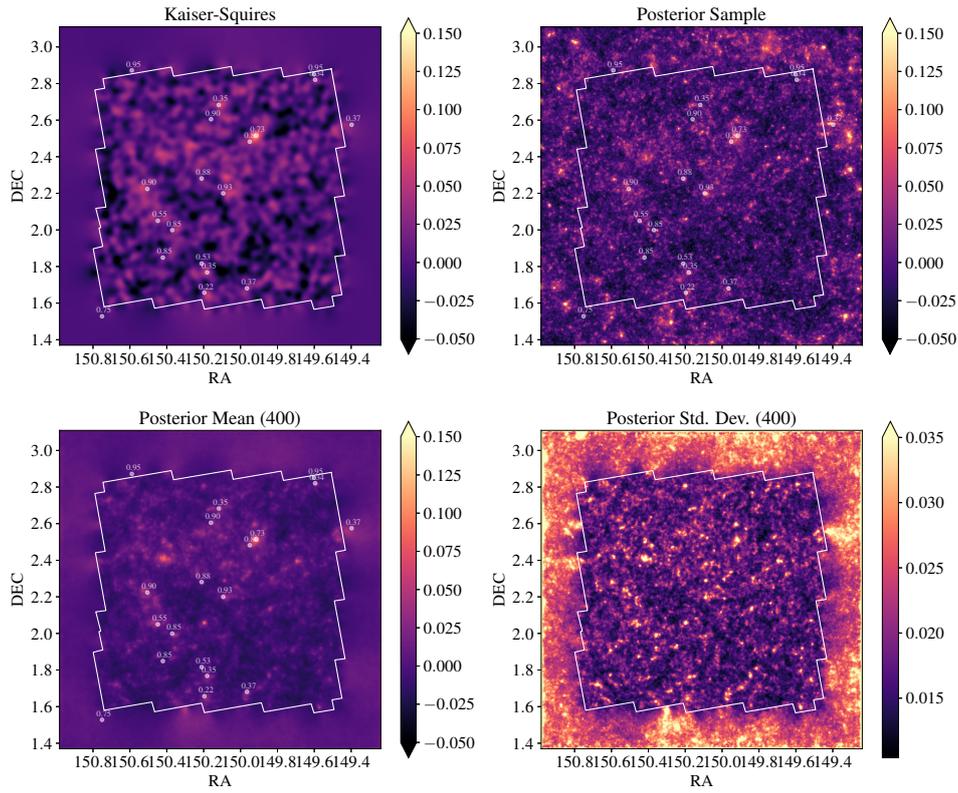


Figure 6.12: *HST/ACS COSMOS* field reconstructions along known X-ray clusters from the XMM-Newton survey. *Top-left*: Kaiser-Squires (with a Gaussian smoothing of $\sigma = 1$ arcmin), *top-right*: sample from the posterior distribution, *bottom-left*: mean of the posterior distribution, *bottom-right*: standard deviation of the posterior distribution (over 400 samples, shown in the clipped range $[0, 0.035]$).

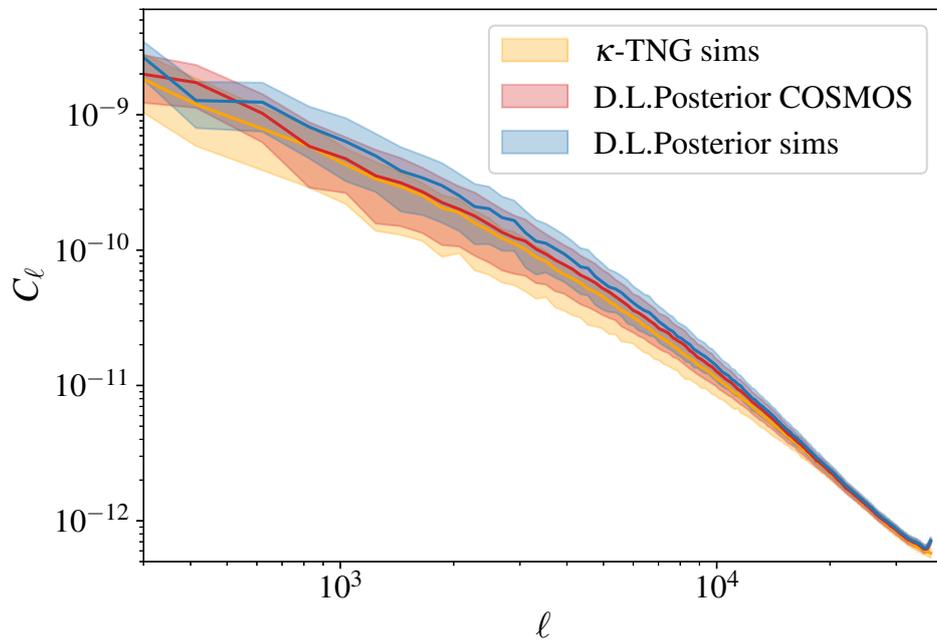


Figure 6.13: Power spectra comparison between the [COSMOS](#) field reconstruction and simulations. We compare the power spectra of posterior samples for the [COSMOS](#) field in *blue* and for a simulated field in *red*. We also display the power spectrum of samples from the training set in *orange* to show the matching.

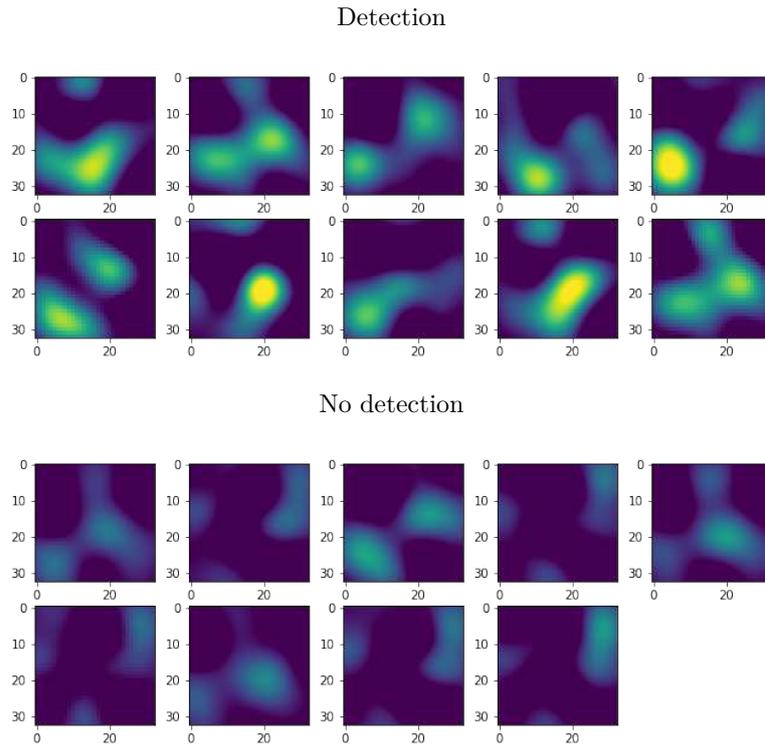


Figure 6.14: This figure shows filtered maps of the convergence posterior samples estimated by our method. Every cutout corresponds to the same region of the map. In the upper panel, a cluster was added in the input shear, while not below. The above cutouts of the filtered maps show the recovered cluster shape in the center of the map, while the below cutouts show coefficient from structure nearby. Selection was done with a 3σ threshold. The axes indicate pixel numbers.

Hierarchical Bayesian modeling of cosmic shear and galaxy morphologies

Chapter Outline

7.1	Pixel-level weak lensing Hierarchical Bayesian Model	102
7.1.1	HST/ACS COSMOS postage stamps	103
7.1.2	Galaxy Latent Variable model	104
7.1.3	Forward modeling galaxy images	107
7.2	Solving the model bias with LVMs	109
7.2.1	Sampling with Hamiltonian Monte Carlo	110
7.2.2	Mean-field Variational inference	111
7.3	Going further	112
7.3.1	Identifying the model bias with HSM	112
7.3.2	Accurate Variational inference with Normalizing Flows	113
7.3.3	Combining posterior distributions	114
7.4	Conclusion	115

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

Remy, B., Lanusse, F., Starck, J.-L., ‘Towards solving model bias in cosmic shear forward modeling’. In: *5th Machine Learning and the Physical Science workshop at NeurIPS (2022)*

THIS chapter presents our contribution to a novel shear estimation method from isolated galaxy images, without relying on shape measurement. Classical shear measurement methods suffer from the fact that the ellipticity is not a well-defined quantity for arbitrary galaxy, especially for irregular morphologies, and therefore subject to a source of error named model bias. We propose to build a forward model of the sky at the pixel level, restricting ourselves to postage stamps of isolated and centered galaxies for a simple proof of concept. We use a latent variable model able to generate arbitrary galaxy light profile corresponding to the survey of interest and a physical pipeline motivated by our understanding of the processes involved in the

generation of galaxy images discussed in [chapter 3](#). This forward model is built under the lens of probabilistic modeling, defining a Hierarchical Bayesian Model (HBM), and consequently a likelihood distribution of the observations under the condition of latent parameters, including the cosmic shear and galaxy morphologies latent variables. We show that by performing the joint inference of the cosmic shear and galaxy morphologies with an accurate light profile model, we are able to recover an unbiased estimate of the cosmic shear affecting the observed galaxy images, when parametric light profile or shape measurement methods fail.

We have seen in [chapter 3](#) how images of distant galaxies are formed and all the different physical processes involved. First, photons are emitted by a background galaxy, generating a light profile image which we consider to be the *original image*. Then they travel through the LSS of the Universe, bending their trajectory. The LSS acts as a lens and affects the original image, shearing it to a distorted observed image. Our goal is therefore to recover this shear field. Classical methods such as Kaiser et al. [KSB95], Hirata et al. [HS03] and Mandelbaum et al. [Man+05] are based on measuring the correlation between galaxy shapes, to recover the contribution of the hidden cosmic shear field. Indeed, as we can assume that there is no preferential intrinsic galaxy orientation, galaxy shape should average to zero, i.e. $\langle e_{\text{galaxy}} \rangle = 0$. However, the cosmic shear contributes to the observed galaxy ellipticity linearly as $e_{\text{obs}} = e_{\text{galaxy}} + \gamma$. Averaging over the observed ellipticities should therefore yield an estimation of the cosmic shear $\langle e_{\text{obs}} \rangle = \hat{\gamma}$. This estimator is very sensitive to the shape measurement method and to systematic error due to the PSF shape and requires calibrations as presented in [chapter 3](#). One major source of bias of this estimator has been identified by Voigt et al. [VB10] and Melchior et al. [Mel+10] when galaxy shapes are measured with simple parametric light profiles such as the de Vaucouleur or the Sérsic models that do not capture all galaxy morphological properties. As a consequence, this is a source of bias in the cosmic shear estimation method, when based on fitting such models. As this bias is due to the lack of efficiency of the light profile model, it is called *model bias*. It is exactly the source of bias we demonstrate to solve in this chapter and that we first proposed in Remy et al. [RLS22].

7.1 Pixel-level weak lensing Hierarchical Bayesian Model

In the following, we propose a new paradigm for cosmic shear estimation where we no longer rely on measuring galaxy shapes (or ellipticities) anymore. We propose to forward model the observations at the pixel level through a Hierarchical Bayesian Model (HBM) where all the physical processes involved in the generation of the observations are expressed explicitly and conditionally related under the framework of probabilistic modeling. Such approaches have already been addressed by Schneider et al. [Sch+15]. However in this paper, the author used parametric light profiles to model galaxy morphologies and are therefore subject to model bias, similarly to classical shape measurement methods. Additionally, it is not clear how to perform shear calibration from simulations, or self-calibration such as the METACALIBRATION method, to mitigate this bias in the context of Bayesian forward modeling. Alternatively, we propose to model galaxy light profile with a Latent Variable Model (Latent Variable Model (LVM)) able to generate light profile images of arbitrary morphologies. This LVM is a neural network variational autoencoder whose parameters are optimized to fit the distribution of the observed galaxy of a survey of interest (we will detail in [subsection 7.1.2](#) how to build such model and their performance on mocking real galaxy images). The unprecedented approach proposed by Lanusse et al. [Lan+21] allows us to train a generative model of galaxy light profile directly from the observations, without

relying on galaxy image simulations. This way, our model of galaxy morphologies does not consider intrinsic ellipticity anymore which differs from Schneider et al. [Sch+15]. According to the weak lensing galaxy image generation pipeline presented in section 3.1, the generated light profile image is then slightly sheared, convolved with the PSF image and corrupted by Gaussian noise. Cosmic shear, noise, and galaxy morphologies have their own latent variables following known or learned distributions, which dictates a probabilistic model structured hierarchically. This hierarchical probabilistic model presented more in depth in subsection 7.1.3. We finally infer these variables jointly from observations and demonstrate that we are able to recover an unbiased estimate of the marginal cosmic shear posterior distribution in section 7.2.

7.1.1 HST/ACS COSMOS postage stamps

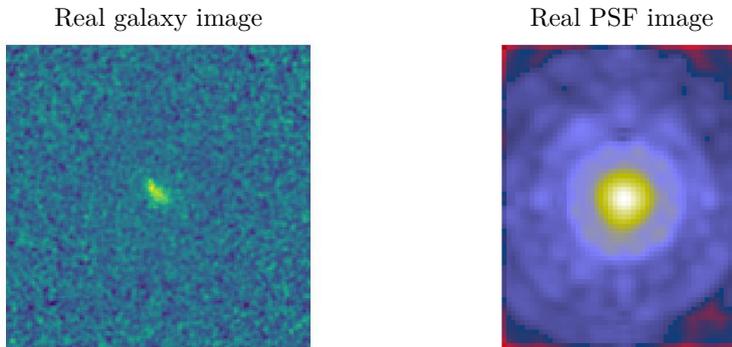


Figure 7.1: On the *left* is a postage stamp from the *Real Galaxy Dataset* Leauthaud et al. [Lea+07] and Mandelbaum et al. [Man+14a] as observed by the ACS from the COSMOS survey and on the *right* is the associated PSF image, in log scale, interpolated at the galaxy position. One can see that the noise of the observed image is correlated, which is due to the co-addition of multiple images. Moreover, one can see that the PSF image is not perfectly isotropic, which can lead to bias in the final shape measurement which requires to be calibrated.

To demonstrate our cosmic shear estimation method, we used the galaxy images taken by the Advanced Camera for Surveys (ACS) for the Cosmic Evolution Survey (COSMOS) [Sco+07c] already presented in section 6.2. A complete description of the galaxy images we used for this work can be found in Leauthaud et al. [Lea+07]. In brief, final postage stamps result from a co-addition of dithered images, ensuring that every area is covered by at least 3 exposures, with a pixel scale of 0.03 arcmin per pixel. While co-addition has the advantage of facilitating the removal of cosmic rays and filling gaps between chips, it implies more strongly correlated pixel noise in the final image. Sources are detected using the SExtractor photometry package Bertin et al. [BA96]. A dataset of deblended postage stamps was therefore compiled for Mandelbaum et al. [Man+14a] where one can find further details on the dataset, such as the magnitude or photometric cuts in the *Real Galaxy Dataset* appendix. Along with galaxy images, this dataset provides the associate PSF image at the galaxy position. The Tiny Tim PSF modeling software was used to create gris of undistorted stars, minimizing aliasing and accounting for temporal instability over the time scale of the COSMOS observation. More details on the PSF modeling can be found in Rhodes

et al. [Rho+07]. In Figure 7.1 we display a real galaxy image from the COSMOS survey and the associated PSF image.

7.1.2 Galaxy Latent Variable model

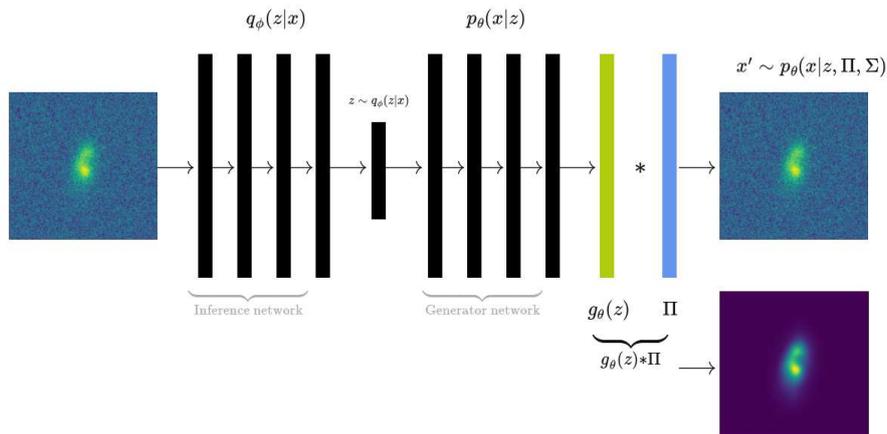


Figure 7.2: Variational Autoencoder with an additional PSF forward model layer. The observed galaxy image x is fed to the inference network, yielding distribution parameters of the surrogate posterior $q_\varphi(z|x)$. A latent variable z is sampled from this posterior distribution, and fed to the generator network generating a light profile image $g_\theta(z)$. The generated image is then convolved with the associated PSF Π , which allows us to define the data likelihood $p_\theta(x|z, \Pi, \Sigma)$.

In this work, we used the Latent Variable Model (LVM) presented in Lanusse et al. [Lan+21]. This model was developed to simulate galaxy postage stamps of the HST/ACS Cosmos survey but could be trained again to mock postage stamps from any survey in principle. The principle virtue of this work was to develop a generative model. As largely explained in the deep generative model chapter (chapter 5), a generative model can generate any data from a learned distribution, given a latent vector randomly sampled from a well-known distribution, say a multivariate Gaussian. The main problem of learning a generative model of galaxy light profiles is that we do not have direct access to datasets to train the models. Indeed, we only have access to corrupted versions of these images, as we observe the galaxy light profile convolved with the PSF of the instrument and with additive noise. However, with the knowledge of a well-calibrated PSF and type of noise, Lanusse et al. [Lan+21] demonstrated that it is possible to train a generative model from corrupted data.

The PSF of the ACS camera is well known and calibrated as explained in Rhodes et al. [Rho+07]. Moreover, we can assume the noise corruption of the postage stamps to be Gaussian, potentially with a known correlation.

By denoting the COSMOS postage stamps as a random variable x , we build the LVM to learn the distribution of the distribution $p(x)$ of the COSMOS catalog. As we are particularly interested in learning the galaxy light profile, we make explicit the knowledge of the corruption process, involving the convolution with the PSF Π_i and noise realisation n_i on the light profile image I_i for the observed pixelized image x_i such that:

$$x_i = \Pi_i \otimes I_i + n_i. \quad (7.1)$$

This way, we can introduce the **LVM** mapping a latent random variable z_i to the light profile image I_i in the observed image description such that:

$$x_i = \Pi_i \otimes g_\theta(z_i) + n_i, \quad (7.2)$$

where g_θ is the **LVM** parameterized by θ . In the following, g is called a generator and θ are neural network weights. Note that in Equation 7.2, all the variables x_i , z_i , P_i , and n_i are specific to a given galaxy i , while the **LVM** generator parameters remain the same for all galaxies. This equation defines the data likelihood $p_\theta(x_i|z_i, \Pi_i, n_i)$. For the COSMOS postage stamps, we assume the noise to be Gaussian correlated with a known covariance matrix per galaxy Σ_i . This gives the following expression for the log-likelihood of the data:

$$\log p_\theta(x_i|z_i, \Pi_i, \Sigma_i) = -\frac{1}{2} \mathbf{F}^\dagger (x_i - \Pi_i \otimes g_\theta(z_i))^\dagger \Sigma_i^{-1} \mathbf{F} (x_i - \Pi_i \otimes g_\theta(z_i)) + cst. \quad (7.3)$$

The normalization constant does not depend on the **LVM** weight θ so will not be involved in the optimization process. One can note the presence of the Fourier operator \mathbf{F} which is used in order to represent the data in a space such that the noise covariance matrix Σ_i is diagonal, and therefore evaluating the likelihood much faster.

We now have a surrogate model for the marginal likelihood of the COSMOS data $p(x)$ that we call $p_\theta(x)$ as it is parameterized by the **LVM** weights θ :

$$p_\theta(x) = \int p_\theta(x|z)p(z)dz \quad (7.4)$$

The goal to train the **LVM** would be to maximize this marginal likelihood p_θ , but this would require to marginalize over all latent variables z , which is intractable. Alternatively, the work of Lanusse et al. [Lan+21] relied on a Variation AutoEncoder (Variational Autoencoder (**VAE**)) proposed in Kingma et al. [KW13]. The idea of the **VAE**, originally known as AutoEncoding Variational Bayes, is to build a generative model with random latent variables $z \sim p(z)$ generating the observed variables x under the likelihood $p_\theta(x|z)$. Similarly to our problem, the **VAE** requires to maximize the marginal likelihood $\int p_\theta(x|z)p(z)$ with respect to the weights θ to be trained, which is intractable. Kingma et al. [KW13] therefore proposed an alternative approach, introducing a parametric distribution $q_\varphi(z|x)$, modeling the posterior distribution of the latent variable z for each observed image x . This distribution is parametrized by another neural network of weights φ . Typically, the posterior $q_\varphi(z|x)$ can be described by a parametric distribution such as a Gaussian, whose mean μ and variance σ^2 are predicted by neural networks, i.e. such that

$$q_\varphi(z|x) = \mathcal{N}(\mu_\varphi(x), \sigma_\varphi^2(x)). \quad (7.5)$$

It appears that the definition of this surrogate posterior distribution q_φ creates a proxy to the maximization of the marginal likelihood and computing its **KL** divergence with respect to the true posterior distribution $p(z|x)$:

$$D_{\text{KL}}(q_\varphi(z|x)|p(z|x)) = \mathbb{E}_{q_\varphi} [\log q_\varphi(z|x) - \log p(z|x)] \quad (7.6)$$

$$\text{(Bayes' theorem)} = \mathbb{E}_{q_\varphi} [\log q_\varphi(z|x) - \log p(z) - \log p(x|z)] + \log p(x) \quad (7.7)$$

$$\text{(Introducing } p_\theta(x|z)) = \mathbb{E}_{q_\varphi} [\log q_\varphi(z|x) - \log p(z) - \log p_\theta(x|z)] + \log p(x) \quad (7.8)$$

$$= D_{\text{KL}}(q_\varphi(z|x)||p(z)) + \log p(x) - \mathbb{E}_{q_\varphi} [\log p_\theta(x|z)]. \quad (7.9)$$

Thus, we can reorder Equation 7.9 to define a lower bound the the marginal likelihood:

$$\log p(x) = \underbrace{\mathbb{E}_{q_\varphi} [\log p_\theta(x|z)] - D_{\text{KL}}(q_\varphi(z|x)||p(z))}_{\text{ELBO}} + \underbrace{D_{\text{KL}}(q_\varphi(z|x)||p(z|x))}_{\geq 0}. \quad (7.10)$$

The last term $D_{\text{KL}}(q_\varphi(z|x)||p(z|x))$ is always positive, by definition of the KL divergence, and one can recognize the ELBO defined in subsection 4.3.2. The ELBO becomes therefore the objective function to optimize as all of the terms involved are tractable as opposed to the marginal likelihood. We illustrate the VAE in Figure 7.2. In this figure, one can observe that the observed image is fed to an inference network, that we call an *encoder*, yielding the mean $\mu_\varphi(x)$ and variance $\sigma_\varphi^2(x)$ of the surrogate posterior $q_\varphi(z|x)$. A latent variable z is then sampled from this posterior distribution and fed to a *generator* network $g_\theta(z)$ yielding the generated light profile image. The generated image is then convolved by the PSF Π associated to the input galaxy image, to match the image, hence the *auto-encoder* denomination of the network.

So far we have defined how to train the auto-encoder with the variational Bayes approach. This model is already a generative model as we can sample from the latent space prior $p(z)$ and generate images going through the decoder. However, the two-fold loss makes it hard for the optimization task to perform well on both tasks: maximizing the likelihood and minimizing the KL divergence between the data posterior and latent space prior. The maximum likelihood term forces the high quality of the reconstructed images, the better optimized, the more reconstructed images are similar to the data. The KL term ensures that the data posterior distribution $q_\varphi(z|x)$ matches a pre-defined prior distribution $p(z)$. This term is essential for the generative aspect of the network because we aim at sampling latent space variables from a simple distribution, e.g. a multivariate Gaussian distribution to generate random images. However, the data posterior is usually much more complicated than a simple Gaussian, and therefore, enforcing this KL divergence to be small can result in a very poor image quality generation.

To alleviate this issue, Lanusse et al. [Lan+21] proposed to train the VAE in a two-step manner. First training the VAE as discussed above, but with *free-bits* of information, and then training a second generative model over the latent space prior. The *free-bits* approach strengthens the importance of the quality of the generated images by penalizing the KL divergence up to some degree of information λ . The loss function is therefore:

$$\mathcal{L}_\lambda = \mathbb{E}_{q_\varphi} [\log p_\theta(x|z)] - \max(\lambda, D_{\text{KL}}(q_\varphi(z|x)||p(z))). \quad (7.11)$$

Because the KL divergence is much less regularized, the generator network is much more able to generate high-quality images, but the data posterior much less constrained and had so much freedom during the optimization that it does not have any similarity with the latent prior. To alleviate this issue, Engel et al. [EHR17] proposed to train a second generative model to learn the latent data posterior distribution. In their work Lanusse et al. [Lan+21] proposed to use a Normalizing Flow (see section 5.2) to learn a generative model of the latent space distribution. This kind of model has in principle the ability to learn arbitrary distributions by learning a bijective matching between a target distribution, here the latent posterior $p_\varphi(z|x)$ and a well-behaved distribution, easy to evaluate its probability density and to sample from, e.g. a multivariate Gaussian distribution. In practice normalizing flows are only very accurate in small

dimension spaces, but the VAE approach has the virtue of reducing the dimension of the data space into a small dimensional latent space, where the flow operates.

7.1.3 Forward modeling galaxy images

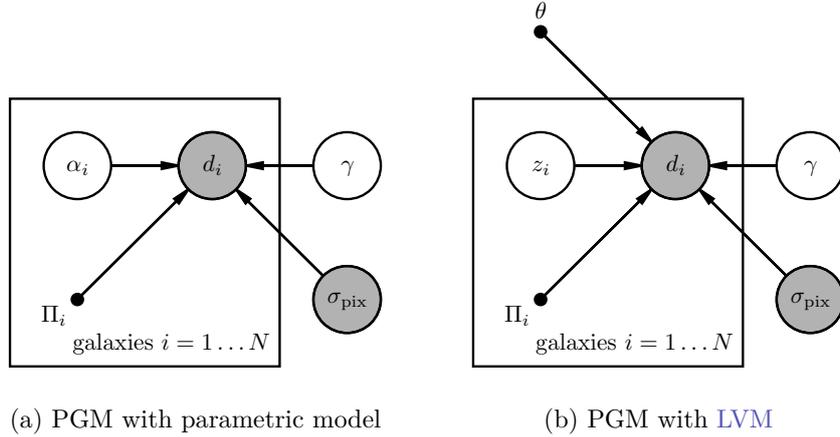


Figure 7.3: Probabilistic Graphical Model (PGM). (a) PGM involving a parametric model for galaxy light profiles with parameters α . (b) PGM involving a latent variable model for galaxy light profiles, where z are parameters of the generative model latent space, and θ are the decoder parameters, such that $\mathcal{D}_\theta(z)$ is a galaxy image. Π_i corresponds to the PSF image, assumed to be known, σ_{pix} the noise standard deviation of the postage stamp noise, and γ the cosmic shear shared by all the images. Shaded circles correspond to observed random variables, white circles correspond to the random variable we aim to infer and dots correspond to fixed variables.

From a Bayesian perspective, making the joint inference of cosmic shear and galaxy morphologies means estimating the *joint posterior* distribution $p(\mathcal{G}, \gamma | \mathcal{D})$, where $\mathcal{G} = \{z_i\}_{i=1\dots N}$ is the set of morphology parameters, $\mathcal{D} = \{d_i\}_{i=1\dots N}$ is the set of observed galaxy postage stamps and γ the cosmic shear. Hereafter, we will mention the parameters of the light profile mode, parametric or of an LVM, either as light profile parameters or morphology parameters. The Bayes identity tells us that:

$$p(\mathcal{G}, \gamma | \mathcal{D}) \propto p(\mathcal{D} | \mathcal{G}, \gamma) p(\mathcal{G}, \gamma), \quad (7.12)$$

where $p(\mathcal{D} | \mathcal{G}, \gamma)$ is the data *likelihood* given by the forward model that we described below. Since the intrinsic morphology of galaxies are independent of the cosmic shear, we can factorize the joint prior as the product of independent priors, i.e. $p(\mathcal{G}, \gamma) = p(\mathcal{G})p(\gamma)$, where $p(\mathcal{G})$ and $p(\gamma)$ are respectively the *prior* distributions of the morphological parameters and of the cosmic shear. The joint distribution has therefore the following factorization:

$$p(\mathcal{G}, \gamma | \mathcal{D}) \propto p(\mathcal{D} | \mathcal{G}, \gamma) p(\mathcal{G}) p(\gamma), \quad (7.13)$$

where we dropped the proportional normalization constant factor $p(\mathcal{D})$ which is not involved during the inference.

The probabilistic graphical model defining the likelihood is illustrated in Figure 7.3. In these diagrams, the nodes denote random variables, and arrows conditional relation

between them. White nodes correspond to the random variables on which we do the inference, conditioned on the shaded nodes corresponding to observed variables. Dot nodes correspond to fixed variables, e.g. neural network weights θ trained a priori or known physical quantities such as the PSF Π . In this work, we compare two forward models. The first one (a) using Sérsic law and the second one (b) using the LVM model described above, both for generating galaxy light profiles. Once the galaxy light profiles are drawn, we shear images with a constant value γ shared by all the galaxies. Each sheared galaxy i is then convolved by its associated Point Spread Function (PSF) Π_i . Assuming Gaussian noise corruption, the resulting image is therefore evaluated by a Gaussian likelihood of variance σ_{pix}^2 .

Sérsic model The light profile of the Sérsic model is given by the following parametrization:

$$I(r) = F \times I_0 \exp \left(-b_n \left[\left(\frac{r}{r_{\text{hlr}}} \right)^{\frac{1}{n}} - 1 \right] \right), \quad (7.14)$$

where r_{hlr} is the half-light-radius, I_0 is a normalization factor such that the integration of I over all the space is F , which corresponds to the flux of the galaxy. n is called the Sérsic index and b_n is a dimensionless constant depending on n . The galaxy is then sheared with parameters $e = (e_1, e_2)$ corresponding to its ellipticity (see, subsection 3.2.2 for further explanation on how to compute these different parameters). Finally, since the real galaxy may not be perfectly centered in the postage stamp, we add two shift parameters s_x and s_y . Therefore, each galaxy i has a set of parameter $z_i = \{n, r_{\text{hlr}}, F, e_1, e_2, s_x, s_y\}_i$. Prior distributions and parameters follow empirical distributions of the COSMOS shape catalog and are described in Table 7.1.

Sérsic parameter	Prior distribution
$\log_{10} r_{\text{hlr}}$	$\mathcal{N}(-0.68, 0.30)$
$\log_{10} n$	$\mathcal{N}(0.1, 0.39)$
$\log_{10} F$	$\mathcal{N}(-1.97, 0.53)$
$e_{1/2}$	$\mathcal{N}(0, 0.28)$
$\gamma_{1/2}$	$\mathcal{N}(0, 0.09)$
$s_{x/y}$	$\mathcal{N}(0, 1)$

Table 7.1: COSMOS sercic parameters prior. These distributions were fitted from the COSMOS shape catalog of [Lea+07; Man+12].

LVM model This VAE-based model, was proposed in [Lan+21], generating complex galaxy light profiles. In this model, the observations are generated following a random process mapping a latent space representation $z_i \in \mathbb{R}^{16}$ to the observed image d_i following the probabilistic model illustrated in Figure 7.3. The neural network weights θ of the VAE are trained maximizing the Evidence Lower Bound (ELBO) of the data. Moreover, another generative model was trained on top of the VAE to learn latent space distribution using Masked Auto Regressive Flows [PPM17]. This model was trained on the HST/ACS COSMOS survey postage stamps.

In the weak lensing regime, the shear is very small, therefore we can assume that it does not change the morphological statistics of the training sample. Thus, the model can be trained once, a priori, over a large catalog of galaxies and be used to perform the joint inference of the latent variables and the cosmic shear. For both light profile models, we denote the generated image \mathcal{I}_0 .

Cosmic shear In this model every galaxy image is sheared by a constant value $\gamma = \gamma_1 + i\gamma_2$ shared by all the galaxies. The image shearing transform, denoted $\mathcal{I}_0 \otimes \gamma$, is given by the following pixel coordinate change:

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = S(\gamma) \begin{pmatrix} x \\ y \end{pmatrix}, \quad (7.15)$$

such that

$$S(\gamma) = \frac{1}{\sqrt{1 - |\gamma|^2}} \begin{pmatrix} 1 + \gamma_1 & \gamma_2 \\ \gamma_2 & 1 - \gamma_1 \end{pmatrix} \quad (7.16)$$

Convolution with the PSF The Point Spread Function (PSF) model is assumed to be known and exact for each galaxy. We draw the PSF evaluated at the galaxy i position, denoted by Π_i , in Fourier space and then multiply it by the Fourier image of the light profile before returning in the real space, resulting in an image \mathcal{I} , such that $\tilde{\mathcal{I}} = (\tilde{\mathcal{I}}_0 \otimes \gamma) \cdot \tilde{\Pi}$.

Noise corruption We assume a known fixed Gaussian noise identical for all the galaxies. For simplicity, all of the galaxies in our observations are corrupted by a constant Gaussian noise of variance σ_{pix}^2 , as well as for our forward model. In a more realistic scenario where we would have information about the noise correlation properties for each postage stamp, one could consider a noise model per galaxy. The likelihood is therefore given by

$$p(\mathcal{D}|\mathcal{G}, \gamma) \propto \prod_{i=1, \dots, N} \mathcal{N}\left(d_i \left| \mathbf{F}^\dagger [(\tilde{\mathcal{I}}_{0,i} \otimes \gamma) \cdot \tilde{\Pi}_i], \sigma^2 \right.\right), \quad (7.17)$$

where \mathbf{F} is the inverse Fourier operator

7.2 Solving the model bias with LVMS

In this section, we investigate and validate our novel shear measurement against a well-established method developed by Hirata et al. [HS03] and Mandelbaum et al. [Man+05], **HSM** for short. We compare the different methods of **HST/ACS COSMOS** postage stamps. To build our simulations, we applied a constant shear on all of the unconvolved galaxy images and reconvolved with the **COSMOS PSF**. In this setting, we assume that there is no selection bias and that the postage stamps contain only a single and centered galaxy.

We designed a toy model where galaxies are assumed to be already detected, deblended, and centered in postage stamps so that we only need to simulate the light profile, shearing transform, and convolution with the **PSF**. To demonstrate the model bias, we distinguish two sets of observations. (1) generated with parametric galaxy light profiles using a Sérsic model and (2) using realistic galaxies based on the galaxy postage stamps from the COSMOS **HST** Advanced Camera for Surveys (**ACS**) field [Koe+07; Sco+07b; Sco+07a]. We used the postage stamps from [Lea+07] and Sérsic fit parameters compiled in [Man+12]. Because the unconvolved galaxies are very noisy, we used the Auto-Encoder from [Lan+21] to get a denoised version of unconvolved galaxies, therefore this is not a demonstration on real galaxies yet. However, the purpose of this work is to demonstrate the model bias, so we find that realistic galaxy

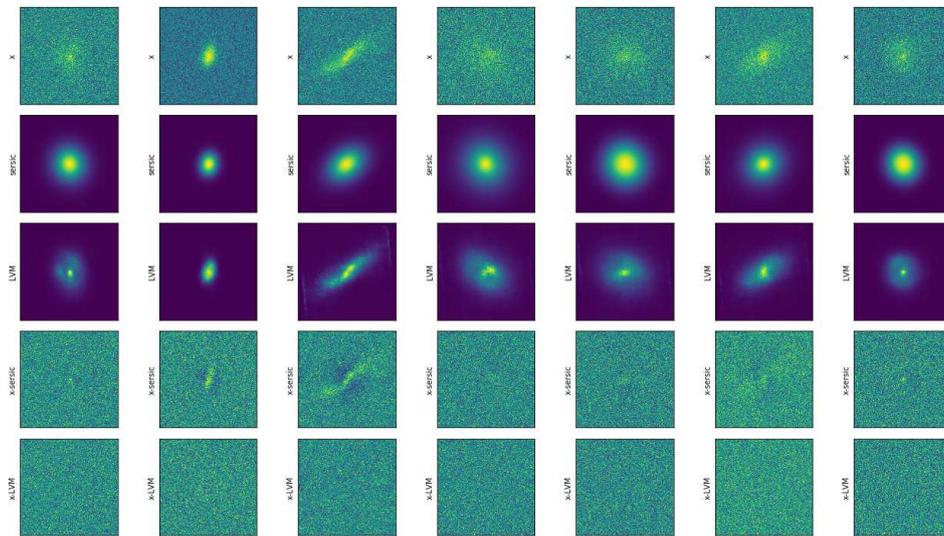


Figure 7.4: Maximum a posteriori fits comparison between a Sérsic light profile and the Latent Variable Model. The first row represents the observed postage stamps by the ACS camera of the HST Cosmos survey. The second row is the Maximum a Posteriori (MAP) fit of the Figure 7.3 (a) model, i.e. assuming a Sérsic light profile, and the third row the MAP fit assuming the Figure 7.3 (b) model, i.e. using the learned LVM. The last two rows represent the difference between the observation and the respective two fits. To highlight the error between the different models, all of the images are visually rescaled with $\text{arcsinh}(s \cdot x)/s$ (with $s = 50$) to saturate the color scale.

morphologies that cannot be captured by analytic profiles are sufficient to illustrate the argument.

In Figure 7.4, we reconstruct the MAP solution using the two PGM defined in subsection 7.1.3. The MAP is obtained running the Adam algorithm presented in subsection 5.1.2 to minimize the negative log joint posterior distribution, i.e. $-\log p(\gamma, \mathcal{G}|\mathcal{D})$. One can observe in this figure that the Sérsic parametric light profiles can fit the overall shape of the COSMOS galaxies, but not their finer substructure. However, the LVM model is able to fit the full light profile of the COSMOS galaxies as we can see in the residual images which have no longer residual patterns.

7.2.1 Sampling with Hamiltonian Monte Carlo

Each set of observations is a field of 256 galaxies, generated from the COSMOS catalog, with the following cuts: $n > 0.4$, $r_{\text{hr}} > 0.2$ and $23.5 < i^+ < 25.2$ so rather faint galaxies. That way the entire galaxy in the postage stamp is not too peaky and is not too resolved to be close to a Sérsic model. Every postage stamp was corrupted by a Gaussian noise of standard deviation $\sigma_{\text{pix}} = 0.01$. To simulate shearing, we removed any shear present in the selected sample from the COSMOS catalog, by adopting the same strategy as in the Third Gravitational Lensing Accuracy Testing (GREAT₃) Challenge [Man+14a], i.e. doubling each galaxy, and rotating the doubled galaxies by 90° , as well as their PSF, to cancel the overall sample shear and then apply the same constant shear to every galaxy $\gamma = (-0.03, 0.03)$.

In [Figure 7.5](#) we perform the inference by sampling with the [HMC](#) sampler the joint distribution $p(\gamma, \mathcal{G} | \mathcal{D}) \propto p(\mathcal{D} | \mathcal{G}, \gamma) p(\mathcal{G}) p(\gamma)$, where the likelihood is defined in [Equation 6.4](#), the prior over the galaxy morphological parameters is learned with a normalizing flow as described in [subsection 7.1.2](#) and the prior for the cosmic shear is given in [Table 7.1](#). We then marginalize over the latent parameters of galaxy morphologies in order to get the shear posterior conditioned on the observations. This shows that using the latent variable model to model the galaxy light profile and with our hybrid generative and physical [HBM](#), we can recover an unbiased marginal posterior distribution on the cosmic shear.

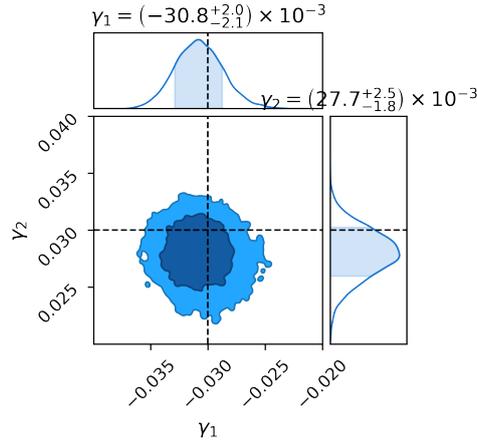


Figure 7.5: Joint inference using Hamiltonian Monte Carlo for galaxy morphologies and cosmic shear using the hybrid generative and physical HMB described in [subsection 7.1.3](#) on COSMOS galaxies.

7.2.2 Mean-field Variational inference

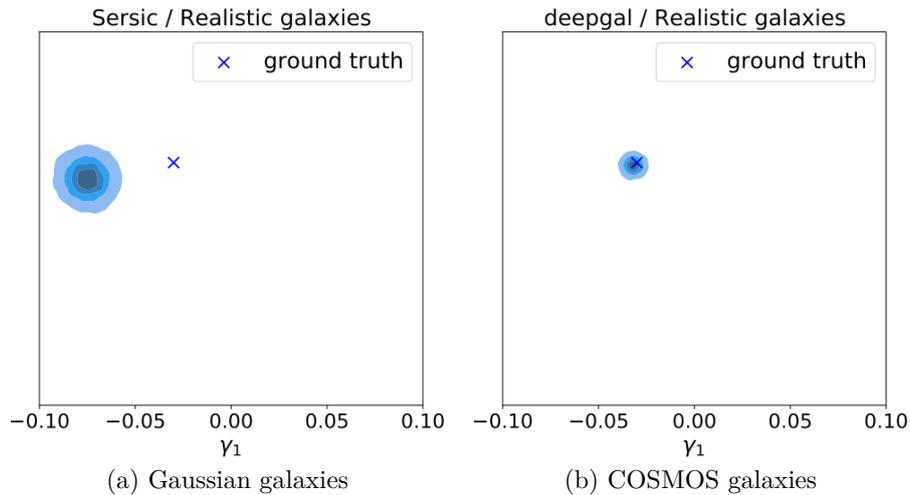


Figure 7.6: Joint inference using mean field Variational Inference for galaxy morphologies and cosmic shear applied to (a) Gaussian galaxies and (b) COSMOS galaxies.

While sampling with **HMC** enables us to get an accurate marginal posterior on the shear, the inference takes a long time before the chains converge. In order to have an inference method that scales with the number of galaxy observed, we propose therefore to use variational inference. In order to estimate the posterior distribution $p(\mathcal{G}, \gamma | \mathcal{D})$, we propose to rely on approximate inference using mean-field Variational Inference (**VI**), where we approximate the posterior distribution using a parametric distribution

$$q_\phi(\mathcal{G}, \gamma) = q_{\phi_\gamma}(\gamma) \prod_i \prod_j q_{\phi_{i,j}}(z_{i,j}), \quad (7.18)$$

parameterized by $\phi = \{\phi_\gamma, \phi_{i,j}\}$, and all the parameters are conditionally independent from each other. We used Normal distributions $q_{\phi_{i,j}}$ to approximate the posterior of parameters i of the j^{th} galaxy, so the parameters $\phi_{i,j} = (\mu_{i,j}, \sigma_{i,j})$ correspond to means and standard deviations. Similarly, q_{ϕ_γ} is a 2d-diagonal multivariate Gaussian. The parameters $\phi_{i,j}$ and ϕ_γ are then optimized maximizing the **ELBO**:

$$\phi^* = \arg \max_{\phi} \mathbb{E}_{(\mathcal{G}, \gamma) \sim q_\phi} [\log p(\mathcal{G}, \gamma | \mathcal{D}) - \log q_\phi(\mathcal{G}, \gamma)]. \quad (7.19)$$

This remains a very simplistic approach, which does not assume any correlation between variables. A more accurate model would condition the morphological parameters on the shear. In principle, the most accurate method would be to run a Markov Chain Monte Carlo, but several experiments showed that sampling-based methods are too computationally expensive for such high-dimensional problem.

We ran three sets of experiments to demonstrate the model bias. First fitting the fully parametric model to the parametric observations. Then fitting the same parametric model to the realistic galaxies and then fitting the hybrid model to the realistic galaxies. We used a Sérsic law as light profile as proposed in [Sch+15]. One could argue that there exist more complex parametric profiles such as the bulge+disk model which could reduce the aforementioned bias, but we propose here a model that does not rely on ellipticities and in principle can capture any complexity of galaxy morphologies.

First, as illustrated in **Figure 7.6** (a) shows that the parametric model induces a bias in the shear estimation when the galaxy morphologies are no longer elliptical. Besides, one can observe the non-zero residuals in the **Figure 7.4** fit to real galaxies, illustrating the model discrepancy when applied to complex morphologies. Additionally, **Figure 7.6** (b) shows that when we substitute the parametric light profile model with the LVM model, the shear estimation becomes unbiased. One can notice that the uncertainties using the generative model are smaller than for the Sérsic model, but these contours are not representative of the error bars due to the mean-field approximation.

7.3 Going further

7.3.1 Identifying the model bias with HSM

In this section, we study how a state-of-the-art shear measurement method performs on our dataset. **HSM** is a well-established method used for shear estimation in several surveys. To measure the cosmic shear **HSM** first measures the weighted-moments of postage stamps to estimate the galaxy shapes e_1 and e_2 and then applies a shape correction to account for the error due to **PSF** intrinsic ellipticity. The fact that **HSM** is based on measuring galaxy ellipticities implies that for complex galaxy morphologies, this method is likely to yield a biased estimation of the shear because galaxy shapes are

not necessarily well described by ellipticities. This source of cosmic shear estimation bias is called *model bias*.

In Figure 7.7, we illustrate where such a bias can be identified for a large number of galaxies. We iteratively increase the number of galaxies on which we apply HSM to estimate the shear. The method is applied on the same dataset studied throughout this section, i.e. HST/ACS Cosmos postage stamps on which we simulated a constant shear of value $\gamma = (+0.03, -0.03)$. The HSM method output for each postage stamp the ellipticity (e_1, e_2) of the galaxy. We then estimate the shear $\hat{\gamma}$ as the average of these galaxy ellipticities and can provide error bars through the standard deviation of the measured shapes, divided by the square root of the number of galaxies.

$$\hat{\gamma}_{\text{HSM}} = \frac{1}{N} \sum_{i=1}^N e \quad (7.20)$$

$$\sigma(\hat{\gamma}_{\text{HSM}})^2 = \frac{1}{N} \sum_{i=1}^N (\hat{\gamma}_{\text{HSM}} - e)^2 \quad (7.21)$$

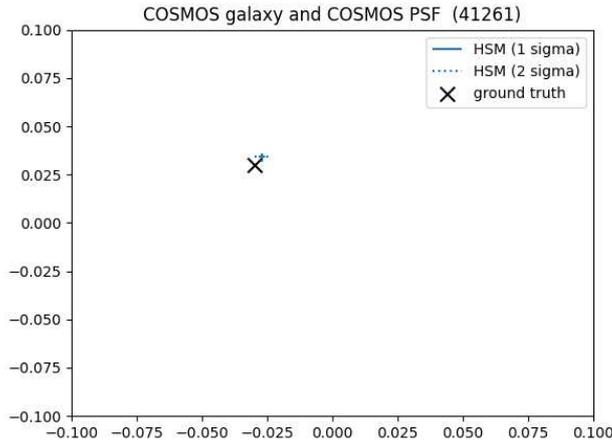


Figure 7.7: HSM applied to COSMOS galaxies & COSMOS PSF. One can observe that the shear estimation is unbiased when applying HSM on Gaussian galaxies. However, when estimating from more than 4×10^5 COSMOS images, the expected shear in the simulations are outside of the 2σ error bars of the HSM estimation and therefore considered biased.

7.3.2 Accurate Variational inference with Normalizing Flows

In subsection 7.2.2 we have identified the presence of bias in the marginal shear posterior estimation when using a parametric light profile on the HST/ACS COSMOS postage stamps. We have also shown that we can mitigate this bias by using a Latent Variable Model trained directly from the observation. However, due to the mean-field approximation, i.e. assuming that each parameter is conditionally independent from each other $p(\gamma, \mathcal{G}|\mathcal{D}) = p(\gamma) \prod_{i,j} p(z_{i,j}|\mathcal{D})$, the posterior distribution may not quantify accurately the correct uncertainties.

One venue for quantifying the full uncertainty of the cosmic shear under the observations, would be to have a surrogate posterior distribution taking into account the potential conditionality between the parameters. We could in principle decompose our surrogate posterior as follows

$$q_\phi(\mathcal{G}, \gamma) = q(\gamma|\mathcal{G})p(\mathcal{G}) \quad (7.22)$$

$$= q_{\phi_\gamma}(\gamma|\mathcal{G}) \prod_i \prod_j q_{\phi_{i,j}}(z_{i,j}). \quad (7.23)$$

This way, $q_{\phi_\gamma}(\gamma|\mathcal{G})$ could be a conditional normalizing flow parameters by the weights ϕ_γ and $q_{\phi_{i,j}}$ could a normalizing flow per galaxy of $\phi_{i,j}$ (see [section 5.2](#)).

7.3.3 Combining posterior distributions

The analysis of cosmic shear field for a wide and deep survey will involve billions of galaxies (Ivezić et al. [Ive+19] indicates that the Vera C. Rubin Observatory LSST will observe 20 billions of galaxies after 10 years of observation). Therefore, it seems intractable to perform the hierarchical inference on all the observed galaxies at once. However, we propose here to split the observations into patches and to combine the shear posterior from each patch to get the overall posterior distribution. Assuming a constant shear on all the sky, we can compute the full posterior distribution $P(\gamma|\mathcal{D})$ from independent posterior distributions inferred from different patches d_n such that $\cup_{n=1}^{n_{\text{patches}}} d_n = \mathcal{D}$ with the following expression:

$$p(\gamma|\mathcal{D}) = \frac{1}{p(\gamma)^{n_{\text{patches}}-1}} \prod_{n=1}^{n_{\text{patches}}} \int de_n p(\gamma, e_n|d_n) \quad (7.24)$$

[Equation 7.24](#) can be derived by noticing that:

1. The total likelihood is separable since realization of patches are independent from each other. Additionally, the likelihood of a set of the observed postage stamps is conditionally dependent only on the parameters of the objects within, meaning mathematically

$$p(d|e, \gamma) = \prod_{n=1}^{n_{\text{patches}}} p(d_n|e, \gamma) = \prod_{n=1}^{n_{\text{patches}}} p(d_n|e_n, \gamma) \quad (7.25)$$

2. According to [Equation 7.25](#), the total posterior distribution can also be factorized as:

$$p(\gamma|\mathcal{D}) = \int de \frac{p(\mathcal{D}|e, \gamma)p(e)p(\gamma)}{p(\mathcal{D})} \quad (7.26)$$

$$(7.27)$$

$$= \prod_{n=1}^{n_{\text{patches}}} \left[\int de_n \frac{p(d_n|e_n, \gamma)p(e_n)}{p(d_n)} \right] p(\gamma), \quad (7.28)$$

because $p(\mathcal{D}) = \prod_{n=1}^{n_{\text{patches}}} p(d_n)$ and $p(e) = \prod_{n=1}^{n_{\text{patches}}} p(e_n)$. Therefore, factorizing again by $p(\gamma)$,

$$p(\gamma|\mathcal{D}) = \prod_{n=1}^{n_{\text{patches}}} \left[\int de_n \frac{p(d_n|e_n, \gamma)p(e_n)p(\gamma)}{p(d_n)} \right] p(\gamma)^{-n_{\text{patches}}} p(\gamma) \quad (7.29)$$

(7.30)

$$= \prod_{n=1}^{n_{\text{patches}}} \left[\int de_n p(\gamma, e_n|\mathcal{D}) \right] p(\gamma)^{-n_{\text{patches}}+1}, \quad (7.31)$$

which gives [Equation 7.24](#).

7.4 Conclusion

We have presented in this work the necessity of an accurate light profile model for forward modeling of the cosmic shear. We have shown that a hybrid physical and deep learning [HBM](#) has the potential to solve the model bias, being able to capture the complex morphology of galaxies and simulating physical processes such as shearing transform and convolution by the [PSF](#). We identify many directions of research towards a real cosmological analysis by forward modeling. On the one hand, many assumptions could be relaxed by building a more complex hierarchical model, e.g. parametrizing the [PSF](#) model and jointly inferring it along with the shear. Concerning the detection and deblending of galaxies, similar work have already been proposed, e.g. Bayesian Light Source Separator ([BLISS](#)) [[Han+22](#)]. On the other hand, accurate uncertainty quantification could be enabled by scaling [MCMCs](#) for [HBMs](#) or using neural density estimator to capture the shape of the joint posterior distribution.

* * *
* * *
* * *

Conclusion

The main goal of this thesis has been the investigation and development of probabilistic methods for ill-posed inverse problems with a particular focus on weak lensing mass-map reconstruction and unbiased cosmic shear measurement.

The first part of the thesis introduced the basic concepts of weak gravitational lensing and classical methods for cosmic shear measurement. As the estimation of cosmic shear relies on the measurement of galaxy shapes, which is not a well-defined quantity for arbitrary galaxy profiles, I presented the correction and calibration methods proposed in the literature to mitigate this issue. I also presented our contribution to leverage automatic differentiation techniques and improve calibration efficiency.

The second part of this thesis focused on the generative modeling and Bayesian modeling framework. In particular, I introduced different neural-based architectures enabling us to learn very complex distributions in high-dimensional spaces. I eventually presented a diffusion-based sampler, the annealed HMC, enabling us to sample high-dimensional distribution from the knowledge of the score of its density function.

Then I presented a unified view of mass-map reconstruction as a Bayesian posterior estimation problem. Up to the present time, existing methods either relied on simple prior, such as assuming a Gaussian Random Field or only returning a point estimate of the posterior distribution. Instead, I have proposed a framework which allows us to use numerical simulations to provide a full non-Gaussian prior on the convergence field and to sample from the full high-dimensional posterior of the mass-mapping problem.

Finally, in the last chapter, I presented the issues of classical shape measurement methods introducing a bias due to the assumption of elliptical shape of galaxy light profiles. Next-generation surveys will photograph the sky with an unprecedented resolution and we will observe much more galaxies with an irregular morphology. I therefore introduced a new paradigm for cosmic shear measurement, no longer relying on galaxy shape. To overcome the model bias and provide uncertainty quantification, I developed a hierarchical Bayesian model where galaxy morphologies are modeled with deep generative techniques, and the physical processes involved in the generation of galaxy images are modeled analytically. As demonstrated on isolated HST/ACS COSMOS galaxy postage stamps, this new method is not only capable of recovering an unbiased estimate of the cosmic shear but also of quantifying the uncertainty of the estimate through its posterior distribution. This offers a promising path toward fully unbiased cosmic shear estimation under the probabilistic modeling framework.

Outlook

Solving arbitrary inverse problems with diffusion samplers Solving arbitrary inverse problems with diffusion samplers In this thesis, we have demonstrated how to efficiently sample from a high-dimensional posterior distribution with an annealing strategy. However, we have restricted our assumption on the annealing of a likelihood involving a linear forward model. Further study is therefore necessary to extend this method to arbitrary inverse problems. To this end, Adam et al. [Ada+22] has already extended the annealing to a non-linear forward model for galaxy source reconstruction from strong-lens systems. However, this paper presents the same assumption concerning the convolution of the posterior, which is not the convolution of both the likelihood and the prior. We therefore need to figure out a correct way to anneal the posterior distribution to be able to use SDE or ODE sampler properly. Another possibility is to rely on the fully amortized posterior distribution as in Legin et al. [Leg+23] where they directly learn the score of the full posterior distribution instead of the prior only. This way, the posterior is correctly annealed and one can apply any SDE or ODE sampler without concern. This however has the drawbacks of losing the analytical likelihood and fully relying on the accuracy of the simulations.

Hierarchical Bayesian Model for *Euclid* Weak lensing measurements with upcoming space-based surveys such as *Euclid* and *Roman* will be very sensitive to their PSF measurement errors. In particular, their PSF is not only varying spatially but also chromatically. Thus, PSF models will require precise measurement of the galaxies Spectral Energy Distribution (Spectral Energy Distribution (SED)) of the galaxies. We know that especially for *Euclid*, which has a wide photometric passband, information about galaxy SEDs will be very heterogeneous, coming from many different sources like ground-based surveys. It is not clear today how the error on galaxy SED measurements propagates to the PSF error and to cosmic shear measurement uncertainties and bias. I therefore propose to extend the Hierarchical Bayesian Model I am currently working on, taking into account a galaxy SED model. There already exist a differentiable physical model of galaxy SED, such as Differentiable Stellar Population Synthesis (DSPS) [Hea+21, DIFFSTAR], and of the PSF [Lia+22, WAVEDIFF]. Such a hierarchical model would perfectly handle the heterogeneous sources of SED measurements, and enable us to infer galaxy properties such as the metallicity, mass, etc., useful to condition the galaxy morphological model. I have already demonstrated that shear model bias can be solved with the correct galaxy morphological model [RLS22]. This study will enlighten us about the impact of SEDs and PSF measurement errors on weak lensing measurements. Again, inference on such a complex HBM is known to be very difficult, and it is precisely where the community would benefit from the efficient sampler.

Appendices

Résumé en français

Reconstruction de carte de masse avec prior appris sur des simulations

De nombreux problèmes inverses en astrophysique sont mal posés en raison de données manquantes et de mesures dominées par le bruit. C'est le cas de la reconstruction de cartes de masse par lentille gravitationnelle faible, où l'on cherche à reconstruire le champ de densité de matière à partir de catalogues de mesures de formes de galaxies. Dans Remy et al. [Rem+20], j'ai proposé de reformuler cette classe de problèmes inverses dans une perspective bayésienne, où la distribution de vraisemblance est décrite par un modèle direct physique et d'apprendre ensuite une distribution a priori apprise sur des simulations, contenant des informations sur le champ de densité de matière, agissant comme un régularisateur.

Il n'existe pas de modèle analytique qui capture la nature non-gaussienne du champ de densité de matière de l'Univers tardif. Néanmoins, les simulations hydrodynamiques à N-corps permettent d'échantillonner des cartes à partir d'un prior implicite complet, prenant en compte l'effondrement gravitationnel non-linéaire et les effets baryoniques. J'ai donc proposé un modèle de prior hybride combinant un Champ Gaussien Aléatoire avec un modèle d'apprentissage profond émulant les échelles non-gaussiennes du prior. Le réseau de neurones a été entraîné avec un objectif de *score-matching* afin d'obtenir les gradients de la distribution a priori, nécessaires pour des méthodes d'échantillonnage efficaces MCMC basées sur les gradients. Cette méthode d'inférence fait partie des récemment introduit *modèles de diffusion*, devenue l'état de l'art en modélisation générative.

Dans Remy et al. [Rem+23], j'ai finalement appliqué cette méthode à la reconstruction des cartes de masse les plus résolues du champ COSMOS HST/ACS à ce jour (voir Figure A.1), avec une quantification robuste des incertitudes.

Inférence en grande dimension pour les problèmes inverses

Les méthodes de Monte Carlo par Chaînes de Markov (MCMC) sont largement utilisées dans la communauté cosmologique pour échantillonner les distributions a posteriori. Cependant, lorsque la dimension de l'espace des paramètres devient trop élevée (dimension $d > 10^5$ pour une carte de masse de 360×360 pixel²), le coût de calcul pour les MCMC standards sont gigantesque à cause de la *malédiction de la dimension* (*curse of dimensionality* en anglais), le coût croissant exponentiellement avec la dimension. Cette difficulté peut être levée si l'on a accès aux gradients

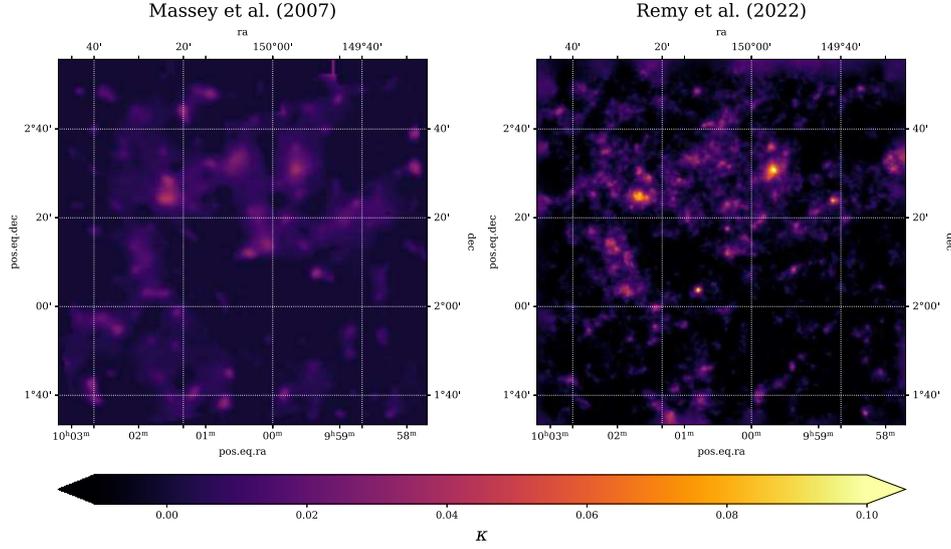


Figure A.1: Reconstruction du champs HST/ACS COSMOS. À *gauche* se trouve la reconstruction présentée dans Massey et al. [Mas+07a] utilisant une variante de Kaiser et al. [KSB95]. À *droite* se trouve la moyenne de la distribution posterior calculée avec notre méthode, utilisant un prior appris sur la suite de simulation κ TNG.

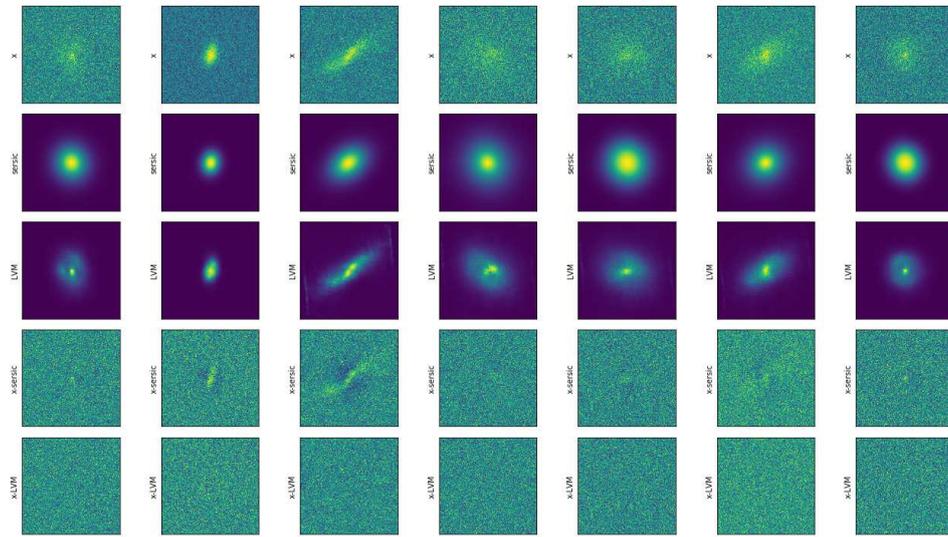
de la distribution a posteriori, qui peuvent être utilisés pour faire des propositions plus efficaces, comme dans l'algorithme de Hamiltonian Monte Carlo (HMC). J'ai démontré l'efficacité des MCMC basés sur les gradients lorsqu'ils sont combinés avec une stratégie de *recuit* pour la cosmologie dans Remy et al. [Rem+20; Rem+23] et pour la reconstruction d'images médicales dans [Ram+20]. Avec ces projets, j'ai été un pionnier dans l'utilisation du HMC avec recuit pour les problèmes inverses en astrophysique. Contrairement aux méthodes d'échantillonnage standards, l'HMC avec recuit garantit d'obtenir des échantillons de la distribution a posteriori indépendants, car seul le dernier échantillon des chaînes initialisées aléatoirement est considéré. De plus, la stratégie de recuit permet de capturer tous les modes de la distribution a posteriori car elle est initialisée par une distribution unimodale large, et converge aléatoirement vers un mode du posterior.

Mesure de cisaillement gravitationnel probabiliste et modèle génératif profond de morphologies de galaxy

L'analyse du cisaillement cosmique, au cœur de la mission *Euclid* de l'ESA, de l'*Observatoire Vera C. Rubin* Legacy Survey of Space and Time, et du télescope spatial *Nancy Grace Roman*, nécessite des mesures de formes de galaxies extrêmement précises. Les mesures de formes dépendent à leur tour d'un contrôle précis de la Fonction d'Étalement du Point (PSF) de l'instrument, et sont également impactées par les détails des morphologies des galaxies qui peuvent induire des biais dans la mesure du cisaillement gravitationnel.

Dans ce projet, j'ai développé un *modèle bayésien hiérarchique* (HBM) d'images observées de galaxies, où la plupart du modèle est physiquement motivé, comme la prescription de la PSF et l'application du cisaillement gravitationnel, mais d'autres

composantes telles que le modèle de morphologie des galaxies est modélisé par un modèle génératif à variables latentes, basé sur des réseaux de neurones profonds. Les profils de lumière paramétriques (par exemple gaussien, Sersic ou bulbe+disque) génèrent un *biais de modèle* dans le processus de mesure de cisaillement car ils ne peuvent pas capturer tous les détails de morphologie des galaxies ou les corrélations entre la morphologie et d'autres paramètres tels que la magnitude et le décalage vers le rouge. L'ajustement de ce modèle direct aux observations signifie inférer une grande quantité de paramètres, ce qui est intractable pour les échantillonneurs standards. J'ai donc développé un modèle des images de galaxies observées incluant la prescription de la PSF, et du cisaillement gravitationnel, ainsi qu'un modèle génératif pour les profils de lumières des galaxies dans [RLS22]. Les méthodes standards d'estimation du cisaillement basées sur les mesures de forme nécessitent de nombreuses études de l'impact des erreurs des moments d'ordre supérieur de la PSF, ce qui est un effet non-négligeable pour les relevés à venir telles que LSST ou *Euclid*. Avec la méthode proposée, le cisaillement gravitationnel est non seulement estimé avec précision, mais la quantification des incertitudes est permise en estimant la distribution a posteriori (voir [Figure A.2](#)).



(a)

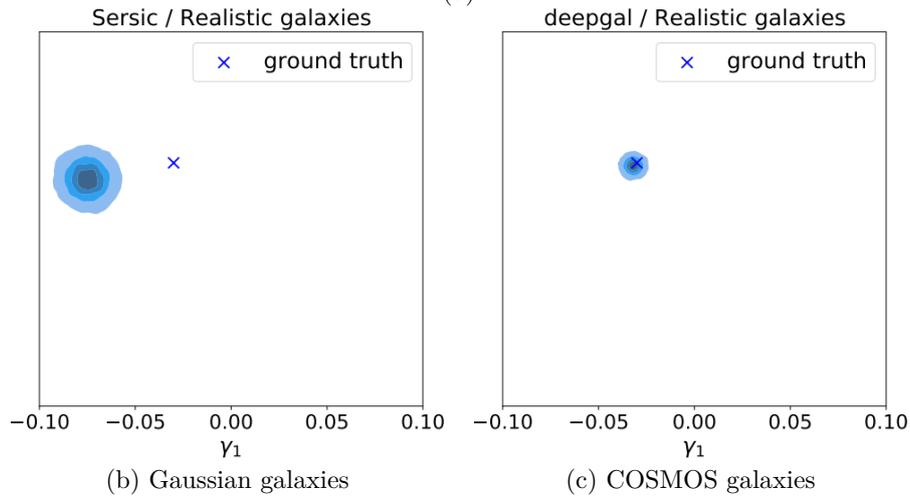


Figure A.2: (a) Reconstruction du **MAP** utilisant un profil de Sérsic ou un modèle génératif à variable latente. La première ligne représente les galaxies observées par la caméra ACS du relevé HST COSMOS. La deuxième ligne représente la reconstruction **MAP** utilisant un profil de lumière Sérsic, et la troisième ligne utilisant un modèle génératif à variable latente. Les deux dernières lignes représentent la différence entre les observations et les reconstructions. (b) et (c) comparent les contours du cisaillement gravitationnel inférés conjointement avec la morphologies des galaxies présentées dans (a). (b) présente un biais car un profil de Sérsic est utilisé, tandis que (c) est non-biaisé grâce à l'utilisation d'un modèle génératif à variables latentes capturant parfaitement la forme des galaxies.

List of Figures

2.1	Patch from the James Webb Space Telescope (JWST) Near-Infrared Camera. One can observe many objects in this image such as galaxies and stars and the galaxy cluster SMACS 0723. This picture also illustrates the effect of strong gravitational lensing of distant galaxies appearing as arcs due to the presence of massive objects in the line of sight. Credit: NASA and ESA.	8
2.2	Comparison between Eisenstein & Hu and HALOFIT matter power spectra	11
2.3	Scheme of weak gravitational lensing deviation of light ray	13
3.1	Journey of a photon from galaxy emission through the LSS to the CCD	18
3.2	Parametric models against different galaxy morphologies	19
3.3	Illustration of a PSF varying spatially and chromatically	21
3.4	Illustration of the survey depth's influence on galaxy blending	23
3.5	Contribution of the shear and convergence to the ellipticity	25
3.6	Fit of ellipses to galaxy images	26
3.7	Illustration of one-dimensional projection of the Sersic light profile for different values of Sersic indices n	28
3.8	Illustration of shear-dependent object detection in the presence of a PSF	33
3.9	Relative difference between response matrix coefficients R_{ij} computed by automatic differentiation and finite differences	34
3.10	Plot of the selection function $w(x, t, \sigma)$ for different smoothing factors	35
4.1	Comparison between the Random Walk Metropolis-Hastings (RWMH), MALA and HMC samplers on the banana distribution	45
4.2	Comparison of the different behaviors between the minimization of the forward and the reverse KL divergences.	48
5.1	Illustration of various activation functions. All these functions have support on \mathbb{R} , and are non-linear and differentiable.	53
5.2	Comparison of the score function learning with and without spectral normalization	59
5.3	Illustration of the different neural network normalization strategies	60
5.4	Flow f mapping a two-dimensional multivariate Gaussian to a two-moons distribution	61
5.5	Example of annealing on the two moons distribution	67
5.6	Bayesian posterior sampling for MRI reconstruction	69

6.1	Mosaic of the COSMOS HST/ACS field	77
6.2	Redshift distribution of the COSMOS catalog from Schrabback et al. [Sch+10]	78
6.3	Mask from S10 catalog of the HST/ACS COSMOS field.	79
6.4	Effective lensing kernel	80
6.5	Illustration of the hybrid denoiser	84
6.6	Scheme of the U-net architecture. Credit: Ronneberger et al. [RFB15].	85
6.7	Annealing jointly the likelihood and the prior is not equivalent to annealing the posterior distribution	86
6.8	Annealing sampling 1-dimensional illustration	89
6.9	Comparison of three different methods to recover the MAP assuming a GRF prior	90
6.10	Comparison of deep learning-based mass-mapping methods on one mocked COSMOS field	93
6.11	Mass-mapping methods comparison	97
6.12	HST/ACS COSMOS field reconstructions along known X-ray clusters from the XMM-Newton survey.	98
6.13	Power spectra comparison between the COSMOS field reconstruction and simulations	99
6.14	Filtered maps of the convergence posterior	100
7.1	Galaxy postage stamp and PSF image from HST/ACS COSMOS	103
7.2	Variational Autoencoder with an additional PSF forward model layer	104
7.3	Probabilistic graphical model generating galaxy images	107
7.4	MAP fits using a Sérsic or LVM light profile	110
7.5	Joint inference using Hamiltonian Monte Carlo for galaxy morphologies and cosmic shear using the hybrid generative and physical HMB	111
7.6	Joint inference using mean field Variational Inference for galaxy morphologies and cosmic shear	111
7.7	HSM applied to COSMOS galaxies	113
A.1	Reconstruction du champs HST/ACS COSMOS. À <i>gauche</i> se trouve la reconstruction présentée dans Massey et al. [Mas+07a] utilisant une variante de Kaiser et al. [KSB95]. À <i>droite</i> se trouve la moyenne de la distribution posterior calculée avec notre méthode, utilisant un prior appris sur la suite de simulation κ TNG.	122
A.2	Reconstruction du MAP utilisant un profil de Sérsic ou un modèle génératif à variable latente	124

Bibliography

- [AB13] Alain, G., Bengio, Y., ‘What regularized auto-encoders learn from the data generating distribution’. In: *1st International Conference on Learning Representations, ICLR 2013 - Conference Track Proceedings*. Vol. 15. 2013, pp. 3743–3773 (page 66).
- [Ada+22] Adam, A., Coogan, A., Malkin, N., Legin, R., Perreault-Levasseur, L., Hezaveh, Y., Bengio, Y., ‘Posterior samples of source galaxies in strong gravitational lenses with score-based priors’. In: *arXiv preprint arXiv:2211.03812* (2022) (page 118).
- [Ade+16] Ade, P. A., Aghanim, N., Arnaud, M., Ashdown, M., Aumont, J., Baccigalupi, C., Banday, A., Barreiro, R., Bartlett, J., Bartolo, N., ‘Planck 2015 results-xiii. cosmological parameters’. In: *Astronomy & Astrophysics* 594 (2016), A13 (pages 79, 88).
- [Agh+20] Aghanim, N., Akrami, Y., Ashdown, M., Aumont, J., Baccigalupi, C., Ballardini, M., Banday, A., Barreiro, R., Bartolo, N., Basak, S., ‘Planck 2018 results-VI. Cosmological parameters’. In: *Astronomy & Astrophysics* 641 (2020), A6 (pages 6, 9).
- [And82] Anderson, B. D. ‘Reverse-time diffusion equation models’. In: *Stochastic Processes and their Applications* 12.3 (1982), pp. 313–326 (page 65).
- [AÖ18] Adler, J., Öktem, O., ‘Deep bayesian inversion’. In: *arXiv preprint arXiv:1811.05910* (2018) (page 91).
- [BA96] Bertin, E., Arnouts, S., ‘SExtractor: Software for source extraction.’ In: *Astronomy & Astrophysics* 117 (1996), pp. 393–404 (pages 22, 103).
- [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. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. Feb. 2006, pp. 37–48 (page 22).
- [Bet18] Betancourt, M. *A Conceptual Introduction to Hamiltonian Monte Carlo*. 2018 (pages 44, 46, 67).
- [BG14] Bernstein, G. M., Gruen, D., ‘Resampling Images in Fourier Domain’. In: *Publications of the Astronomical Society of the Pacific* 126.937 (Mar. 2014), p. 287 (page 20).

- [BJ02] Bernstein, G., Jarvis, M., ‘Shapes and shears, stars and smears: optimal measurements for weak lensing’. In: *The Astronomical Journal* 123.2 (2002), p. 583 (page 30).
- [BKM17] Blei, D. M., Kucukelbir, A., McAuliffe, J. D., ‘Variational Inference: A Review for Statisticians’. In: *Journal of the American Statistical Association* 112.518 (Apr. 2017), pp. 859–877 (page 47).
- [Bla+19] Blazek, J. A., MacCrann, N., Troxel, M., Fang, X., ‘Beyond linear galaxy alignments’. In: *Physical Review D* 100.10 (2019), p. 103506 (page 75).
- [Bob+12] Bobin, J., Starck, J. .-, Sureau, F., Fadili, J., ‘CMB Map Restoration’. In: *Advances in Astronomy* 2012, 703217 (Jan. 2012), p. 703217 (page 81).
- [Bos+18] Bosch, J., Armstrong, R., Bickerton, S., Furusawa, H., Ikeda, H., Koike, M., Lupton, R., Mineo, S., Price, P., Takata, T., ‘The hyper supprime-cam software pipeline’. In: *Publications of the Astronomical Society of Japan* 70.SP1 (2018), S5 (page 23).
- [Bro65] Broyden, C. G. ‘A Class of Methods for Solving Nonlinear Simultaneous Equations’. In: *Mathematics of Computation* 19 (1965), pp. 577–593 (page 28).
- [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 22).
- [CB99] Ciotti, L., Bertin, G., ‘Analytical properties of the $R^{1/m}$ luminosity law’. In: *arXiv preprint astro-ph/9911078* (1999) (page 27).
- [Clo+12] Clowe, D., Markevitch, M., Bradač, M., Gonzalez, A. H., Chung, S. M., Massey, R., Zaritsky, D., ‘On dark peaks and missing mass: a weak-lensing mass reconstruction of the merging cluster system A520’. In: *The Astrophysical Journal* 758.2 (2012), p. 128 (page 95).
- [Daw+15] Dawson, W. A., Schneider, M. D., Tyson, J. A., Jee, M. J., ‘THE ELLIPTICITY DISTRIBUTION OF AMBIGUOUSLY BLENDED OBJECTS’. In: *The Astrophysical Journal* 816.1 (2015), p. 11 (page 23).
- [de 48] de Vaucouleurs, G. ‘Recherches sur les Nebuleuses Extragalactiques’. In: *Annales d’Astrophysique* 11 (Jan. 1948), p. 247 (page 18).
- [Dod03] Dodelson, S. *Modern Cosmology*. 2003 (page 9).
- [DSB17] Dinh, L., Sohl-Dickstein, J., Bengio, S., *Density estimation using Real NVP*. 2017 (pages 57, 61, 62).
- [Dur+19] Durkan, C., Bekasov, A., Murray, I., Papamakarios, G., ‘Neural spline flows’. In: *Advances in neural information processing systems* 32 (2019) (page 61).
- [EH98] Eisenstein, D. J., Hu, W., ‘Baryonic features in the matter transfer function’. In: *The Astrophysical Journal* 496.2 (1998), p. 605 (page 11).
- [EHR17] Engel, J., Hoffman, M., Roberts, A., ‘Latent constraints: Learning to generate conditionally from unconditional generative models’. In: *arXiv preprint arXiv:1711.05772* (2017) (page 106).
- [EW13] Elsner, F., Wandelt, B. D., ‘Efficient Wiener filtering without preconditioning’. In: *Astronomy & Astrophysics* 549 (2013), A111 (pages 81, 89, 92, 97).

- [Far+22] 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: *Astronomy & Astrophysics* 664 (2022), A141 (page 22).
- [Fin+06] Finoguenov, A., Guzzo, L., Hasinger, G., Scoville, N. Z., Aussel, H., Boehringer, H., Brusa, M., Capak, P., Cappelluti, N., Comastri, A., Giardini, S., Griffiths, R. E., Impey, C., Koekemoer, A. M., Kneib, J. -., Leauthaud, A., Fevre, O. L., Lilly, S., Mainieri, V., Massey, R., McCracken, H. J., Mobasher, B., Murayama, T., Peacock, J. A., Sakelliou, I., Schinnerer, E., Silverman, J. D., Smolcic, V., Taniguchi, Y., Tasca, L., Taylor, J. E., Trump, J. R., Zamorani, G., ‘The XMM-Newton wide-field survey in the COSMOS field: VI. Statistical properties of clusters of galaxies’. In: (Dec. 2006) (page 94).
- [Fri22] Friedmann, A. ‘Über die Krümmung des Raumes’. In: *Zeitschrift für Physik* 10 (Jan. 1922), pp. 377–386 (page 6).
- [GBC16] Goodfellow, I., Bengio, Y., Courville, A., *Deep Learning*. <http://www.deeplearningbook.org>. MIT Press, 2016 (page 52).
- [Gou+21] Gouk, H., Frank, E., Pfahringer, B., Cree, M. J., ‘Regularisation of neural networks by enforcing lipschitz continuity’. In: *Machine Learning* 110 (2021), pp. 393–416 (pages 59, 60).
- [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., Tobias, L., Miller, L., Spitzer, I., Waerbeke, L. V., Cuillandre, J.-C., Fabbro, S., McConnachie, A., Mellier, Y., ‘ShapePipe: A new shape measurement pipeline and weak-lensing application to UNIONS/CFIS data’. In: *Astronomy & Astrophysics* 666 (2022), A162 (page 22).
- [Gut81] Guth, A. H. ‘Inflationary universe: A possible solution to the horizon and flatness problems’. In: *Physical Review D* 23.2 (1981), p. 347 (page 11).
- [Han+22] Hansen, D., Mendoza, I., Liu, R., Pang, Z., Zhao, Z., Avestruz, C., Regier, J., *Scalable Bayesian Inference for Detection and Deblending in Astronomical Images*. 2022 (pages 23, 115).
- [Has70] Hastings, W. K. ‘Monte carlo sampling methods using Markov chains and their applications’. In: *Biometrika* 57.1 (1970), pp. 97–109 (page 44).
- [He+15] He, K., Zhang, X., Ren, S., Sun, J., ‘Deep Residual Learning for Image Recognition’. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015), pp. 770–778 (pages 58, 83).
- [Hea+21] Hearin, A. P., Chaves-Montero, J., Alarcon, A., Becker, M. R., Benson, A., ‘DSPS: Differentiable Stellar Population Synthesis’. In: *arXiv preprint arXiv:2112.06830* (2021) (page 118).
- [Hey+06] Heymans, C., Van Waerbeke, L., Bacon, D., Berge, J., Bernstein, G., Bertin, E., Bridle, S., Brown, M. L., Clowe, D., Dahle, H., ‘The Shear Testing Programme—I. Weak lensing analysis of simulated ground-based observations’. In: *Monthly Notices of the Royal Astronomical Society* 368.3 (2006), pp. 1323–1339 (pages 29, 31).
- [HJA20] Ho, J., Jain, A., Abbeel, P., ‘Denoising diffusion probabilistic models’. In: *Advances in Neural Information Processing Systems* 33 (2020), pp. 6840–6851 (page 65).

- [HM17] Huff, E., Mandelbaum, R., ‘Metacalibration: direct self-calibration of biases in shear measurement’. In: *arXiv preprint arXiv:1702.02600* (2017) (pages 31, 32).
- [Hoe+98] Hoekstra, H., Franx, M., Kuijken, K., Squires, G., ‘Weak lensing analysis of Cl 1358+ 62’. In: *New Astronomy Reviews* 42.2 (1998), pp. 137–140 (pages 29, 78).
- [HS03] Hirata, C., Seljak, U., ‘Shear calibration biases in weak-lensing surveys’. In: *Monthly Notices of the Royal Astronomical Society* 343.2 (2003), pp. 459–480 (pages 29–31, 102, 109).
- [Hut+06] Huterer, D., Takada, M., Bernstein, G., Jain, B., ‘Systematic errors in future weak-lensing surveys: requirements and prospects for self-calibration’. In: *Monthly Notices of the Royal Astronomical Society* 366.1 (2006), pp. 101–114 (page 31).
- [Hyv05] Hyvärinen, A. ‘Estimation of Non-Normalized Statistical Models by Score Matching’. In: *Journal of Machine Learning Research* 6 (2005), pp. 695–708 (page 66).
- [Ilb+08] Ilbert, O., Capak, P., Salvato, M., Aussel, H., McCracken, H., Sanders, D., Scoville, N., Kartaltepe, J., Arnouts, S., Le Floch, E., ‘Cosmos photometric redshifts with 30-bands for 2-deg²’. In: *The Astrophysical Journal* 690.2 (2008), p. 1236 (page 78).
- [IS15] Ioffe, S., Szegedy, C., ‘Batch normalization: Accelerating deep network training by reducing internal covariate shift’. In: *International conference on machine learning*. pmlr. 2015, pp. 448–456 (page 58).
- [Ive+19] Ivezić, Ž., Kahn, S. M., Tyson, J. A., Abel, B., Acosta, E., Allsman, R., Alonso, D., AlSayyad, Y., Anderson, S. F., Andrew, J., ‘LSST: from science drivers to reference design and anticipated data products’. In: *The Astrophysical Journal* 873.2 (2019), p. 111 (page 114).
- [Jee+12] Jee, M., Mahdavi, A., Hoekstra, H., Babul, A., Dalcanton, J., Carroll, P., Capak, P., ‘A study of the dark core in A520 with the Hubble Space Telescope: The mystery deepens’. In: *The Astrophysical Journal* 747.2 (2012), p. 96 (page 95).
- [Jef+18] Jeffrey, N., Abdalla, F. B., Lahav, O., Lanusse, F., Starck, J.-L., Leonard, A., Kirk, D., Chang, C., Baxter, E., Kacprzak, T., ‘Improving weak lensing mass map reconstructions using Gaussian and sparsity priors: application to DES SV’. In: *Monthly Notices of the Royal Astronomical Society* 479.3 (2018), pp. 2871–2888 (page 81).
- [Jef+20a] 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 (2020), pp. 5023–5029 (page 82).
- [Jef+20b] 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 (2020), pp. 5023–5029 (pages 92, 97).

- [Jef+21] Jeffrey, N., Gatti, M., Chang, C., Whiteway, L., Demirbozan, U., Kovács, A., Pollina, G., Bacon, D., Hamaus, N., Kacprzak, T., ‘Dark Energy Survey Year 3 results: Curved-sky weak lensing mass map reconstruction’. In: *Monthly Notices of the Royal Astronomical Society* 505.3 (2021), pp. 4626–4645 (page 75).
- [JHF18] Jeffrey, N., Heavens, A. F., Fortio, P. D., ‘Fast sampling from Wiener posteriors for image data with dataflow engines’. In: *Astronomy and Computing* 25 (Oct. 2018), pp. 230–237 (pages 92, 97).
- [Jol+20] Jolicoeur-Martineau, A., Piché-Taillefer, R., Combes, R. T. d., Mitliagkas, I., ‘Adversarial score matching and improved sampling for image generation’. In: (2020), pp. 1–29 (page 68).
- [JW20] Jeffrey, N., Wandelt, B. D., ‘Solving high-dimensional parameter inference: marginal posterior densities & Moment Networks’. In: *arXiv preprint arXiv:2011.05991* (2020) (page 91).
- [JW23] Jeffrey, N., Wandelt, B. D., *Evidence Networks: simple losses for fast, amortized, neural Bayesian model comparison*. 2023 (page 41).
- [KB14] Kingma, D. P., Ba, J., ‘Adam: A Method for Stochastic Optimization’. In: *CoRR* abs/1412.6980 (2014) (pages 55, 84).
- [KD18] Kingma, D. P., Dhariwal, P., ‘Glow: Generative flow with invertible 1x1 convolutions’. In: *Advances in neural information processing systems* 31 (2018) (page 61).
- [KHM11] Kitching, T. D., Heavens, A. F., Miller, L., ‘3D photometric cosmic shear’. In: *Monthly Notices of the Royal Astronomical Society* 413.4 (2011), pp. 2923–2934 (page 18).
- [Kil15] Kilbinger, M. ‘Cosmology with cosmic shear observations: a review’. In: *Reports on Progress in Physics* 78.8 (2015), p. 086901 (page 11).
- [Kin+16] Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., Welling, M., ‘Improved variational inference with inverse autoregressive flow’. In: *Advances in neural information processing systems* 29 (2016) (page 61).
- [KL51] Kullback, S., Leibler, R. A., ‘On information and sufficiency’. In: *The annals of mathematical statistics* 22.1 (1951), pp. 79–86 (page 47).
- [Koe+07] Koekemoer, A. M., Aussel, H., Calzetti, D., Capak, P., Giavalisco, M., Kneib, J.-P., Leauthaud, A., Le Fevre, O., McCracken, H., Massey, R., ‘The cosmos survey: Hubble space telescope advanced camera for surveys observations and data processing’. In: *The Astrophysical Journal Supplement Series* 172.1 (2007), p. 196 (pages 77, 109).
- [KS93] Kaiser, N., Squires, G., ‘Mapping the dark matter with weak gravitational lensing’. In: *Astrophysical Journal, Part 1 (ISSN 0004-637X), vol. 404, no. 2, p. 441-450*. 404 (1993), pp. 441–450 (pages 14, 80, 92).
- [KSB94] Kaiser, N., Squires, G., Broadhurst, T., ‘A method for weak lensing observations’. In: *arXiv preprint astro-ph/9411005* (1994) (page 24).
- [KSB95] Kaiser, N., Squires, G., Broadhurst, T., ‘A Method for Weak Lensing Observations’. In: *The Astrophysical Journal* 449 (Aug. 1995), p. 460 (pages 29, 78, 97, 102, 122, 126).

- [Kur+21] Kurita, T., Takada, M., Nishimichi, T., Takahashi, R., Osato, K., Kobayashi, Y., ‘Power spectrum of halo intrinsic alignments in simulations’. In: *Monthly Notices of the Royal Astronomical Society* 501.1 (2021), pp. 833–852 (page 79).
- [KW13] Kingma, D. P., Welling, M., ‘Auto-encoding variational bayes’. In: *arXiv preprint arXiv:1312.6114* (2013) (page 105).
- [Lah+94] Lahav, O., Fisher, K. B., Hoffman, Y., Scharf, C. A., Zaroubi, S., ‘Wiener Reconstruction of All-Sky Galaxy Surveys in Spherical Harmonics’. In: *The Astrophysical Journal* 423.1 (Mar. 1994), p. L93 (page 81).
- [Lan+16] Lanusse, F., Starck, J. L., Leonard, A., Pires, S., ‘High resolution weak lensing mass mapping combining shear and flexion’. In: *Astronomy and Astrophysics* 591 (2016) (pages 82, 92, 97).
- [Lan+21] Lanusse, F., Mandelbaum, R., Ravanbakhsh, S., Li, C.-L., Freeman, P., Póczos, B., ‘Deep generative models for galaxy image simulations’. In: *Monthly Notices of the Royal Astronomical Society* 504.4 (2021), pp. 5543–5555 (pages 102, 104–106, 108, 109).
- [Lan15] Lanusse, F. ‘Reconstruction parcimonieuse de la carte de masse de matière noire par effet de lentille gravitationnelle’. 2015SACLS102. PhD thesis. 2015 (page 82).
- [LCL00] Lewis, A., Challinor, A., Lasenby, A., ‘Efficient computation of cosmic microwave background anisotropies in closed Friedmann-Robertson-Walker models’. In: *The Astrophysical Journal* 538.2 (2000), p. 473 (page 10).
- [LDS12] Leonard, A., Dupé, F.-X., Starck, J.-L., ‘A compressed sensing approach to 3D weak lensing’. In: *Astronomy & Astrophysics* 539 (2012), A85 (page 82).
- [Lea+07] Leauthaud, A., Massey, R., Kneib, J.-P., Rhodes, J., Johnston, D. E., Capak, P., Heymans, C., Ellis, R. S., Koekemoer, A. M., Le Fevre, O., ‘Weak gravitational lensing with COSMOS: Galaxy selection and shape measurements’. In: *The Astrophysical Journal Supplement Series* 172.1 (2007), p. 219 (pages 103, 108, 109).
- [Leg+23] Legin, R., Ho, M., Lemos, P., Perreault-Levasseur, L., Ho, S., Hezaveh, Y., Wandelt, B., ‘Posterior Sampling of the Initial Conditions of the Universe from Non-linear Large Scale Structures using Score-Based Generative Models’. In: *arXiv preprint arXiv:2304.03788* (2023) (page 118).
- [Les11] Lesgourgues, J. ‘The cosmic linear anisotropy solving system (CLASS) I: overview’. In: *arXiv preprint arXiv:1104.2932* (2011) (page 10).
- [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 preprint arXiv:2203.04908* (2022) (page 118).
- [Lia+23] Liaudat, T., Starck, J.-L., Kilbinger, M., Frugier, P.-A., ‘Rethinking data-driven point spread function modeling with a differentiable optical model’. In: *Inverse Problems* 39.3 (2023), p. 035008 (page 21).
- [Lia22] Liaudat, T. I. ‘Data-driven modelling of ground-based and space-based telescope’s point spread functions’. 2022UPASP118. PhD thesis. 2022 (page 22).
- [Lim+20] Lim, J. H., Courville, A. C., Pal, C. J., Huang, C.-W., ‘AR-DAE: Towards Unbiased Neural Entropy Gradient Estimation’. In: *International Conference on Machine Learning*. 2020 (page 67).

- [Lin82] Linde, A. D. ‘A new inflationary universe scenario: a possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems’. In: *Physics Letters B* 108.6 (1982), pp. 389–393 (page 11).
- [Liu+18] Liu, J., Bird, S., Matilla, J. M. Z., Hill, J. C., Haiman, Z., Madhavacheril, M. S., Petri, A., Spergel, D. N., ‘MassiveNuS: cosmological massive neutrino simulations’. In: *Journal of Cosmology and Astroparticle Physics* 2018.03 (2018), p. 049 (page 80).
- [LK97] Luppino, G., Kaiser, N., ‘Detection of weak lensing by a cluster of galaxies at $z = 0.83$ ’. In: *The Astrophysical Journal* 475.1 (1997), p. 20 (page 29).
- [LMR23] Liu, R., McAuliffe, J. D., Regier, J., *Variational Inference for Deblending Crowded Starfields*. 2023 (page 23).
- [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: *Monthly Notices of the Royal Astronomical Society* 361.4 (2005), pp. 1287–1322 (pages 29–31, 102, 109).
- [Man+12] Mandelbaum, R., Lackner, C., Leauthaud, A., Rowe, B., ‘COSMOS real galaxy dataset’. In: *Zenodo* (Jan. 2012) (pages 108, 109).
- [Man+14a] Mandelbaum, R., Rowe, B., Bosch, J., Chang, C., Courbin, F., Gill, M., Jarvis, M., Kannawadi, A., Kacprzak, T., Lackner, C., ‘The third gravitational lensing accuracy testing (GREAT₃) challenge handbook’. In: *The Astrophysical Journal Supplement Series* 212.1 (2014), p. 5 (pages 18, 19, 103, 110).
- [Man+14b] Mandelbaum, R., Rowe, B., Bosch, J., Chang, C., Courbin, F., Gill, M., Jarvis, M., Kannawadi, A., Kacprzak, T., Lackner, C., Leauthaud, A., Miyatake, H., Nakajima, R., Rhodes, J., Simet, M., Zuntz, J., Armstrong, B., Bridle, S., Coupon, J., Dietrich, J. P., Gentile, M., Heymans, C., Jurling, A. S., Kent, S. M., Kirkby, D., Margala, D., Massey, R., Melchior, P., Peterson, J., Roodman, A., Schrabback, T., ‘THE THIRD GRAVITATIONAL LENSING ACCURACY TESTING (GREAT₃) CHALLENGE HANDBOOK’. In: *The Astrophysical Journal Supplement Series* 212.1 (2014), p. 5 (page 18).
- [Man+18] Mandelbaum, R., Lanusse, F., Leauthaud, A., Armstrong, R., Simet, M., Miyatake, H., Meyers, J. E., Bosch, J., Murata, R., Miyazaki, S., ‘Weak lensing shear calibration with simulations of the HSC survey’. In: *Monthly Notices of the Royal Astronomical Society* 481.3 (2018), pp. 3170–3195 (pages 29–31).
- [Mar+18] Marinacci, F., Vogelsberger, M., Pakmor, R., Torrey, P., Springel, V., Hernquist, L., Nelson, D., Weinberger, R., Pillepich, A., Naiman, J., ‘First results from the IllustrisTNG simulations: radio haloes and magnetic fields’. In: *Monthly Notices of the Royal Astronomical Society* 480.4 (2018), pp. 5113–5139 (page 79).
- [Mar63] Marquardt, D. W. ‘An algorithm for least-squares estimation of nonlinear parameters’. In: *Journal of the society for Industrial and Applied Mathematics* 11.2 (1963), pp. 431–441 (page 28).

- [Mas+07a] Massey, R., Rhodes, J., Ellis, R., Scoville, N., Leauthaud, A., Finoguenov, A., Capak, P., Bacon, D., Aussel, H., Kneib, J.-P., Koekemoer, A., McCracken, H., Mobasher, B., Pires, S., Refregier, A., Sasaki, S., Starck, J.-L., Taniguchi, Y., Taylor, A., Taylor, J., ‘Dark matter maps reveal cosmic scaffolding’. In: *Nature* 445 (Jan. 2007). 286–290 (pages 94, 122, 126).
- [Mas+07b] Massey, R., Heymans, C., Bergé, J., Bernstein, G., Bridle, S., Clowe, D., Dahle, H., Ellis, R., Erben, T., Hettterscheidt, M., ‘The Shear Testing Programme 2: Factors affecting high-precision weak-lensing analyses’. In: *Monthly Notices of the Royal Astronomical Society* 376.1 (2007), pp. 13–38 (page 31).
- [Mas+07c] Massey, R., Rhodes, J., Leauthaud, A., Capak, P., Ellis, R., Koekemoer, A., Réfrégier, A., Scoville, N., Taylor, J. E., Albert, J., ‘COSMOS: three-dimensional weak lensing and the growth of structure’. In: *The Astrophysical Journal Supplement Series* 172.1 (2007), p. 239 (page 78).
- [Mel+10] Melchior, P., Böhnert, A., Lombardi, M., Bartelmann, M., ‘Limitations on shapelet-based weak-lensing measurements’. In: *Astronomy & Astrophysics* 510 (2010), A75 (pages 18, 102).
- [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 23).
- [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 23).
- [Met+53] Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., Teller, E., ‘Equation of state calculations by fast computing machines’. In: *The Journal of Chemical Physics* 21.6 (1953), pp. 1087–1092 (page 44).
- [MSS12] Martens, J., Sutskever, I., Swersky, K., ‘Estimating the Hessian by Back-Propagating Curvature’. In: *Proceedings of the 29th International Conference on International Conference on Machine Learning, ICML’12*. Edinburgh, Scotland: Omnipress, 2012, pp. 963–970 (page 66).
- [Mur22] Murphy, K. P. *Probabilistic Machine Learning: An introduction*. MIT Press, 2022 (page 52).
- [Nai+18] Naiman, J. P., Pillepich, A., Springel, V., Ramirez-Ruiz, E., Torrey, P., Vogelsberger, M., Pakmor, R., Nelson, D., Marinacci, F., Hernquist, L., ‘First results from the IllustrisTNG simulations: A tale of two elements—chemical evolution of magnesium and europium’. In: *Monthly Notices of the Royal Astronomical Society* 477.1 (2018), pp. 1206–1224 (page 79).
- [Nea11] Neal, R. M. ‘MCMC using hamiltonian dynamics’. In: *Handbook of Markov Chain Monte Carlo* (2011), pp. 113–162 (pages 44, 67).
- [Nel+18] Nelson, D., Pillepich, A., Springel, V., Weinberger, R., Hernquist, L., Pakmor, R., Genel, S., Torrey, P., Vogelsberger, M., Kauffmann, G., ‘First results from the IllustrisTNG simulations: the galaxy colour bimodality’. In: *Monthly Notices of the Royal Astronomical Society* 475.1 (2018), pp. 624–647 (page 79).

- [Nel+19] Nelson, D., Springel, V., Pillepich, A., Rodriguez-Gomez, V., Torrey, P., Genel, S., Vogelsberger, M., Pakmor, R., Marinacci, F., Weinberger, R., Kelley, L., Lovell, M., Diemer, B., Hernquist, L., ‘The IllustrisTNG simulations: public data release’. In: *Computational Astrophysics and Cosmology* 6.1, 2 (May 2019), p. 2 (page 79).
- [Nyq28] Nyquist, H. ‘Thermal Agitation of Electric Charge in Conductors’. In: *Phys. Rev.* 32 (1 July 1928), pp. 110–113 (page 22).
- [OLH21] Osato, K., Liu, J., Haiman, Z., ‘ κ TNG: effect of baryonic processes on weak lensing with IllustrisTNG simulations’. In: *Monthly Notices of the Royal Astronomical Society* 502.4 (2021), pp. 5593–5602 (pages 80, 97).
- [Pee80] Peebles, P. J. E. *The large-scale structure of the universe*. 1980 (page 9).
- [Pen+02] Peng, C. Y., Ho, L. C., Impey, C. D., Rix, H.-W., ‘Detailed structural decomposition of galaxy images’. In: *The Astronomical Journal* 124.1 (2002), p. 266 (page 28).
- [Pen+10] Peng, C. Y., Ho, L. C., Impey, C. D., Rix, H.-W., ‘Detailed decomposition of galaxy images. II. Beyond axisymmetric models’. In: *The Astronomical Journal* 139.6 (2010), p. 2097 (page 28).
- [Per+19] Perraudin, N., Defferrard, M., Kacprzak, T., Sgier, R., ‘DeepSphere: Efficient spherical convolutional neural network with HEALPix sampling for cosmological applications’. In: *Astronomy and Computing* 27 (2019), pp. 130–146 (page 96).
- [Pil+18] Pillepich, A., Nelson, D., Hernquist, L., Springel, V., Pakmor, R., Torrey, P., Weinberger, R., Genel, S., Naiman, J. P., Marinacci, F., ‘First results from the IllustrisTNG simulations: the stellar mass content of groups and clusters of galaxies’. In: *Monthly Notices of the Royal Astronomical Society* 475.1 (2018), pp. 648–675 (page 79).
- [PLS17] Peel, A., Lanusse, F., Starck, J.-L., ‘Sparse reconstruction of the merging a520 cluster system’. In: *The Astrophysical Journal* 847.1 (2017), p. 23 (page 95).
- [PPM17] Papamakarios, G., Pavlakou, T., Murray, I., ‘Masked autoregressive flow for density estimation’. In: *Advances in neural information processing systems* 30 (2017) (pages 61–63, 108).
- [Pri+18] Price, M. A., Cai, X., McEwen, J. D., Pereyra, M., Kitching, T. D., ‘Sparse Bayesian mass-mapping with uncertainties: local credible intervals’. In: (Dec. 2018) (pages 82, 95).
- [Ram+20] Ramzi, Z., **Remy, B.**, Lanusse, F., Starck, J.-L., Ciuciu, P., ‘Denoising score-matching for uncertainty quantification in inverse problems’. In: *Deep inverse workshop at NeurIPS* (2020) (pages 46, 52, 66, 74, 122).
- [Rem+20] **Remy, B.**, Lanusse, F., Ramzi, Z., Liu, J., Jeffrey, N., Starck, J.-L., ‘Probabilistic Mapping of Dark Matter by Neural Score Matching’. In: *3rd Machine Learning and the Physical Science workshop at NeurIPS* (2020) (pages 46, 51, 52, 66, 67, 74, 121, 122).
- [Rem+23] **Remy, B.**, Lanusse, F., Jeffrey, N., Liu, J., Starck, J.-L., Osato, K., Schrabback, T., ‘Probabilistic mass-mapping with neural score estimation’. In: *Astronomy & Astrophysics* 672 (2023), A51 (pages 51, 66, 67, 73, 92, 121, 122).

- [RFB15] Ronneberger, O., Fischer, P., Brox, T., ‘U-net: Convolutional networks for biomedical image segmentation’. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III* 18. Springer. 2015, pp. 234–241 (pages 83, 85, 126).
- [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., ‘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 (2007), p. 203 (pages 103, 104).
- [RLS22] **Remy, B.**, Lanusse, F., Starck, J.-L., ‘Towards solving model bias in cosmic shear forward modeling’. In: *5th Machine Learning and the Physical Science workshop at NeurIPS* (2022) (pages 101, 102, 118, 123).
- [RM16] Rezende, D. J., Mohamed, S., *Variational Inference with Normalizing Flows*. 2016 (page 60).
- [Sam+19] Samuroff, S., Blazek, J., Troxel, M., MacCrann, N., Krause, E., Leonard, C., Prat, J., Gruen, D., Dodelson, S., Eifler, T., ‘Dark Energy Survey Year 1 results: constraints on intrinsic alignments and their colour dependence from galaxy clustering and weak lensing’. In: *Monthly Notices of the Royal Astronomical Society* 489.4 (2019), pp. 5453–5482 (page 75).
- [Sch+07] Schrabback, T., Erben, T., Simon, P., Miralles, J.-M., Schneider, P., Heymans, C., Eifler, T., Fosbury, R. A. E., Freudling, W., Hetterscheidt, M., Hildebrandt, H., Pirzkal, N., ‘Cosmic shear analysis of archival HST/ACS data’. In: *Astronomy & Astrophysics* 468.3 (June 2007), pp. 823–847 (page 78).
- [Sch+10] Schrabback, T., Hartlap, J., Joachimi, B., Kilbinger, M., Simon, P., Benabed, K., Bradač, M., Eifler, T., Erben, T., Fassnacht, C. D., High, F. W., 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: *Astronomy and Astrophysics* 516 (June 2010), A63 (pages 78, 126).
- [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 (2015), p. 87 (pages 102, 103, 112).
- [Sco+07a] Scoville, N., Abraham, R., Aussel, H., Barnes, J., Benson, A., Blain, A., Calzetti, D., Comastri, A., Capak, P., Carilli, C., ‘COSMOS: Hubble space telescope observations’. In: *The Astrophysical Journal Supplement Series* 172.1 (2007), p. 38 (page 109).
- [Sco+07b] Scoville, N., Aussel, H., Brusa, M., Capak, P., Carollo, C. M., Elvis, M., Giavalisco, M., Guzzo, L., Hasinger, G., Impey, C., ‘The cosmic evolution survey (COSMOS): overview’. In: *The Astrophysical Journal Supplement Series* 172.1 (2007), p. 1 (page 109).
- [Sco+07c] Scoville, N., Aussel, H., Brusa, M., Capak, P., Carollo, C. M., Elvis, M., Giavalisco, M., Guzzo, L., Hasinger, G., Impey, C., ‘The cosmic evolution survey (COSMOS): overview’. In: *The Astrophysical Journal Supplement Series* 172.1 (2007), p. 1 (pages 77, 103).

- [SE19] Song, Y., Ermon, S., ‘Generative modeling by estimating gradients of the data distribution’. In: *Advances in neural information processing systems* 32 (2019) (pages 65, 67, 68, 84).
- [Ser68] Sersic, J. L. *Atlas de Galaxias Australes*. 1968 (page 18).
- [SH17] Sheldon, E. S., Huff, E. M., ‘Practical weak-lensing shear measurement with metacalibration’. In: *The Astrophysical Journal* 841.1 (2017), p. 24 (pages 31, 33).
- [She+23] Sheldon, E. S., Becker, M. R., Jarvis, M., Armstrong, R., Collaboration, T. L., ‘Metadetection Weak Lensing for the Vera C. Rubin Observatory’. In: *arXiv preprint arXiv:2303.03947* (2023) (page 33).
- [Shi+23] Shi, J., Sunayama, T., Kurita, T., Takada, M., Sugiyama, S., Mandelbaum, R., Miyatake, H., More, S., Nishimichi, T., Johnston, H., ‘The Intrinsic Alignment of Galaxy Clusters and Impact of Projection Effects’. In: *arXiv preprint arXiv:2306.09661* (2023) (page 79).
- [Soh+15] Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S., ‘Deep unsupervised learning using nonequilibrium thermodynamics’. In: *International conference on machine learning*. PMLR. 2015, pp. 2256–2265 (page 65).
- [Son+20a] Song, Y., Garg, S., Shi, J., Ermon, S., ‘Sliced Score Matching: A Scalable Approach to Density and Score Estimation’. In: *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*. Ed. by R. P. Adams and V. Gogate. Vol. 115. Proceedings of Machine Learning Research. PMLR, 22–25 Jul 2020, pp. 574–584 (page 66).
- [Son+20b] Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., Poole, B., ‘Score-based generative modeling through stochastic differential equations’. In: *arXiv preprint arXiv:2011.13456* (2020) (pages 65, 68).
- [Spr+18] Springel, V., Pakmor, R., Pillepich, A., Weinberger, R., Nelson, D., Hernquist, L., Vogelsberger, M., Genel, S., Torrey, P., Marinacci, F., ‘First results from the IllustrisTNG simulations: matter and galaxy clustering’. In: *Monthly Notices of the Royal Astronomical Society* 475.1 (2018), pp. 676–698 (page 79).
- [SPR06] Starck, J.-L., Pires, S., Réfrégier, A., ‘Weak lensing mass reconstruction using wavelets’. In: *Astronomy & Astrophysics* 451.3 (2006), pp. 1139–1150 (page 82).
- [Sta+21] Starck, J.-L., Themelis, K., Jeffrey, N., Peel, A., Lanusse, F., ‘Weak-lensing mass reconstruction using sparsity and a Gaussian random field’. In: *Astronomy & Astrophysics* 649 (2021), A99 (pages 81, 82, 92, 97).
- [Tak+12] Takahashi, R., Sato, M., Nishimichi, T., Taruya, A., Oguri, M., ‘Revising the halofit model for the nonlinear matter power spectrum’. In: *The Astrophysical Journal* 761.2 (2012), p. 152 (page 11).
- [VB10] Voigt, L., Bridle, S., ‘Limitations of model-fitting methods for lensing shear estimation’. In: *Monthly Notices of the Royal Astronomical Society* 404.1 (2010), pp. 458–467 (pages 18, 102).
- [Vil+22] Villaescusa-Navarro, F., Genel, S., Angles-Alcazar, D., Thiele, L., Dave, R., Narayanan, D., Nicola, A., Li, Y., Villanueva-Domingo, P., Wandelt, B., ‘The camels multifield data set: Learning the universe’s fundamental parameters with artificial intelligence’. In: *The Astrophysical Journal Supplement Series* 259.2 (2022), p. 61 (page 95).

- [Vin11] Vincent, P. ‘A connection between scorematching and denoising autoencoders’. In: *Neural Computation* 23.7 (2011), pp. 1661–1674 (page 66).
- [Vit+22] Vitorelli, A. Z., **Remy, B.**, Guinot, A., Lanusse, F., Martins, T., ‘Auto-MetaCal: self-calibration of shear biases with automatic differentiation’. In: *Rencontres de Moriond, proceedings* (2022) (pages 17, 34).
- [WH18] Wu, Y., He, K., *Group Normalization*. 2018 (page 60).
- [Wie49] Wiener, N. *The Extrapolation, Interpolation and Smoothing of Stationary Time Series with engineering applications*. New York: John Wiley & Sons, Inc., 1949 (page 81).
- [Zar+95] Zaroubi, S., Hoffman, Y., Fisher, K. B., Lahav, O., ‘Wiener Reconstruction of the Large-Scale Structure’. In: *The Astrophysical Journal* 449 (Aug. 1995), p. 446 (page 81).
- [Zeg+22] Zeghal, J., Lanusse, F., Boucaud, A., **Remy, B.**, Aubourg, E., ‘Neural Posterior Estimation with Differentiable Simulators’. In: *Machine Learning for Astrophysics workshop at ICML* (2022) (pages 51, 57).