DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD
DASH.HARVARD.EDU

HARVARD LIBRARY

Statistically Exploring Cracks in the Lambda Cold Dark Matter Model

Citation

Diaz Rivero, Ana. 2020. Statistically Exploring Cracks in the Lambda Cold Dark Matter Model. Doctoral dissertation, Harvard University Graduate School of Arts and Sciences.

Link

https://nrs.harvard.edu/URN-3:HUL.INSTREPOS:37369445

Terms of use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material (LAA), as set forth at

https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFlOWI2M2EwYTg

Accessibility

https://accessibility.huit.harvard.edu/digital-accessibility-policy

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. Submit a story

HARVARD UNIVERSITY
Graduate School of Arts and Sciences

DISSERTATION ACCEPTANCE CERTIFICATE

The undersigned, appointed by the

Department of Physics

have examined a dissertation entitled

Statistically Exploring Cracks in the Lambda Cold Dark Matter Model

presented by Ana Diaz Rivero

candidate for the degree of Doctor of Philosophy and hereby certify that it is worthy of acceptance.

Signature Cora Dvorkin

Typed name: Professor Cora Dvorkin, Chair

Signature Daniel Eisenstein

Typed name: Professor Daniel Eisenstein

Signature Douglas Finkbeiner

Typed name: Professor Douglas Finkbeiner

Signature Matthew B. Reece

Typed name: Professor Matthew Reece

Date: October 7, 2020

Statistically Exploring Cracks in the Lambda Cold Dark Matter Model

a dissertation presented

by

Ana Diaz Rivero

to

The Department of Physics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Physics

Harvard University

Cambridge, Massachusetts

October 2020

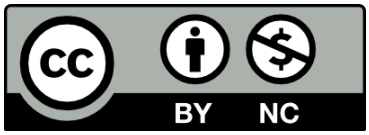

© 2020 Ana Diaz Rivero

This work is licensed under a Creative Commons Attribution Non-Commercial 4.0

(CC-BY-NC) International License

Dissertation Advisor: Cora Dvorkin Ana Diaz Rivero

Statistically Exploring Cracks in the Lambda Cold Dark Matter Model

ABSTRACT

The Lambda Cold Dark Matter (ΛCDM) model has long been lauded for its exquisite fit to large-scale cosmological observations. Recent years, however, have revealed cracks in this model and the theories that underpin it. Tensions have arisen between the values of some cosmological parameters inferred from high-redshift and low-redshift independent observables. Little progress has been made in confirming, or at least improving, our understanding about some of the most important theoretical foundations the model rests on, including inflation, dark energy, and dark matter.

This dissertation aims to shed light on some of the most important outstanding issues in the ΛCDM picture of the universe and develop statistical tools to (1) extract more and better information from cosmological and astrophysical observables, and to (2) improve the accuracy of statistical inference pipelines, questioning established modeling choices and assumptions that can give us unfair confidence in our model of the universe, skew our understanding of the universe, or both. In the first part of the dissertation we focus on probes of cold dark matter on the kiloparsec scale, a regime that remains untested and is very sensitive to dark matter microphysics. We present novel ways of mapping the distribution of dark matter on these scales that offer a variety of advantages compared to traditional methods: sensitivity to lower masses (smaller scales), much more model-independence, and a significant speed-up. In the second part of the dissertation we address the impact of canonical modeling choices in cosmological data analyses and how this relates to the important tensions we observe between high-precision datasets. We use and develop data-driven methods that can incorporate our ignorance into statistical inference pipelines, removing important assumptions that are usually baked into analyses.

We conclude by assessing the reach of our work and placing it in the context of the broader field, envisioning how it will shape research in the years to come.

Acknowledgements

First and foremost I want to thank my family. It's impossible to put into words just how much I owe my parents, Ana and Ángel, but this dissertation wouldn't exist if it weren't for everything they sacrified to get me to where I am today. To Jesús, my brother, and Núria, who is like a sister to me, I want to thank for their patience with my competitiveness in endless games of Mario Kart, Settlers of Catan, and Terraforming Mars during graduate school. To my husband, Andrew, I am indebted for making every one of my days for the last five years better and brighter.

I also want to thank my advisor, Cora, for all the opportunities and doors she has opened up for me. Her guidance and support when I was starting my Ph.D. gave me the confidence and tools necessary to eventually feel less like a student and more like a researcher. The freedom and flexibility she progressively offered me allowed me to explore my interests and ultimately made my years in graduate school genuinely enjoyable, fulfilling, and exciting. Today I feel like a real cosmologist in no small part due to all the time and effort she invested in me.

Para Abue

Contents

1 Introduction **1**

Strong Gravitational Lensing as a Dark Matter Probe **6**

2 Probing ΛCDM on Small Scales with Strong Gravitational Lensing **7**

3 Direct Detection of Dark Matter Substructure in Strong Lens Images with Convolutional Neural Networks **11**

3.1 Data and Methods 14

3.1.1 Simulating strongly-lensed images 14

3.1.2 Neural Networks 24

3.1.3 Training, validation and test sets 25

3.1.4 CNN architecture and optimization strategy 26

3.2 Results 28

3.3 Discussion and Conclusion 31

4 Extracting the Subhalo Mass Function from Strong Lens Images with Image Segmentation **37**

4.1 Data generation 39

4.2 Model architecture and training 43

4.3 The effects of noise 47

4.4 Characterizing the network performance 51

4.4.1 Example output 51

4.4.2 Testing on a single subhalo 54

4.4.3 Null tests . . . 60
4.4.4 Testing on images with many subhalos . . . 61
4.5 Determination of the subhalo mass function . . . 63
4.5.1 Expectation and variance . . . 65
4.5.2 Inferring the power-law . . . 68
4.6 Discussion and outlook . . . 70

5 Power Spectrum of Dark Matter Substructure in Strong Gravitational Lenses 74

5.1 Substructure statistics within the halo model . . . 76
5.1.1 Preliminaries: Subhalo statistics . . . 76
5.1.2 Ensemble-averaged substructure convergence . . . 80
5.1.3 The power spectrum of the convergence field . . . 81
5.2 Truncated Navarro-Frenk-White subhalo population . . . 87
5.2.1 Characteristics of the subhalo population . . . 87
5.2.2 Power spectrum: 1-subhalo term . . . 92
5.2.3 Power spectrum: 2-subhalo term . . . 98
5.3 Truncated Cored subhalo population . . . 100
5.4 Discussion and Conclusion . . . 104

6 Gravitational Lensing and the Power Spectrum of Dark Matter Substructure: Insights from the ETHOS N-body Simulations 108

6.1 The Power Spectrum of Dark Matter Substructure within galaxies . . . 109
6.2 Simulations . . . 110
6.3 Methods . . . 111
6.3.1 Power spectra from subhalo catalogs . . . 112
6.3.2 Power spectra from simulation snapshots . . . 116
6.4 Results . . . 117
6.4.1 Subhalo catalog . . . 117
6.4.2 Snapshot particle data . . . 124

6.5 Discussion and Conclusion 127

7 Quantifying the Line-of-Sight Halo Contribution to the Dark Matter Convergence Power Spectrum from Strong Gravitational Lenses 131

7.1 Analytical Calculation 134

7.1.1 Multi-plane Lens Equation 134

7.1.2 Effective Convergence for Interlopers 136

7.1.3 Power Spectrum of Interlopers 143

7.1.4 Effective Convergence Power Spectrum for a Population of tNFW Perturbers 151

7.2 Numerical Calculation 156

7.2.1 A Simulated Strong-Lensing System with Interlopers and Subhalos 158

7.2.2 Obtaining the Convergence Power Spectrum from a Simulation . . 159

7.3 Conclusions 164

Statistical Inference in the Era of Cosmological Tensions 168

8 Observable Predictions for Massive-Neutrino Cosmologies with Model-Independent Dark Energy 170

8.1 Data and Methods 172

8.2 Results 177

8.3 Discussion 182

9 Flow-Based Likelihoods for Non-Gaussian Inference 187

9.1 Data-Driven Likelihoods 191

9.1.1 Gaussian Mixture Model 192

9.1.2 Independent Component Analysis 193

9.1.3 Flow-Based Likelihood 194

9.2 Measuring Non-Gaussianity 197

9.2.1 t-statistic of Skewness and Excess Kurtosis 198

9.2.2 Pairwise Non-Gaussianity of Data Points 199

9.2.3 Nonparametric Kullback-Leibler (KL) divergence 201
9.3 Mock Weak Lensing Convergence Power Spectra 202
9.3.1 Background 202
9.3.2 Mock Data 203
9.3.3 Weak Lensing Likelihood Non-Gaussianity 204
9.4 Learning Flow-Based Likelihoods (FBL) 208
9.4.1 Network architecture and training procedure 208
9.4.2 Results 209
9.5 Discussion 212
9.6 Conclusion 215

Final Remarks **218**

Appendix **223**

A Direct Detection of Dark Matter Substructure in Strong Lens Images with Convolutional Neural Networks **224**
A.1 Subhalo population characteristics 224

B Extracting the Subhalo Mass Function from Strong Lens Images using Image Segmentation **228**
B.1 Single subhalo pixel predictions 228
B.2 Confusion matrix 230
B.3 Subhalo predictions with probability threshold 231
B.4 How realistic is 1% white noise? 233
B.5 Domain adaptation 235

C Power Spectrum of Dark Matter Substructure in Strong Gravitational Lenses **238**
C.1 Deriving $\mathcal{P}_{\rm t}$ 238
C.2 SIDM convergence profile 240
C.3 Table of Constants and Variables 241

D Gravitational Lensing and the Power Spectrum of Dark Matter Substructure:
Insights from the ETHOS N-body Simulations **242**
D.1 Truncated convergence profiles 242
D.2 Features of the convergence maps 244
D.2.1 CDM 245
D.2.2 ETHOS4 246

E Quantifying the Line-of-Sight Halo Contribution to the Dark Matter Convergence Power Spectrum from Strong Gravitational Lenses **248**
E.1 Scaling Relations for tNFW Projected Density 248
E.2 2D Fourier Transform of a Projection 249
E.3 The Area of the Double Cone 250
E.4 Number density of subhalos 251
E.5 $k \to 0$ Limit of the Perturber Power Spectrum 253
E.6 Curl and Divergence Components 254

F Observable Predictions for Massive-Neutrino Cosmologies with Model-Independent Dark Energy **259**
F.1 Results for the *Reduced* dataset 259

G Flow-Based Likelihoods for Non-Gaussian Inference **261**
G.1 Likelihood versus Sampling density 261
G.2 Toy Problems 263
G.2.1 In two dimensions 263
G.2.2 In higher dimensions 263
G.3 Data-driven Likelihoods in Data-Limited Regimes 267
G.4 Robustness of the nonparametric KL Divergence Estimator 269
G.5 Galaxy Power Spectrum Non-Gaussianity 270
G.6 Mock Weak Lensing Convergence Maps 274
G.6.1 Generating the mocks 274

G.6.2 Convergence Power Spectrum Distribution per Bin 275

References **277**

List of Figures

3.1 Realistic strong lens image simulation pipeline 15
3.2 Examples of simulated strong lens images 15
3.3 Noise realizations with varying correlation length 24
3.4 Network classification accuracy 30

4.1 Examples of segmented strong lens images 38
4.2 U-Net network architecture 44
4.3 Cross-entropy loss and classification accuracy as a function of epoch . . . 46
4.4 Classification accuracy as a function of epoch with varying noise 48
4.5 Fractional change in intensity as a function of subhalo mass 50
4.6 Probability maps for every class in a segmented image with a single subhalo 52
4.7 Single-subhalo predictions for every class 56
4.8 False-positive rate 60
4.9 Many-subhalo test set examples 62
4.10 Average number of detected subhalos per subhalo mass class for different values of the SMF logarithmic slope 65
4.11 Parametric fit to the average number of subhalos per subhalo mass bin as a function of the SMF logarithmic slope 67
4.12 Constraints on the SMF logarithmic slope for varying numbers of stacked images 70

5.1 NFW density profile with and without truncation 89
5.2 Subhalo mass function with varying logarithmic slope 91

5.3 Properties of the 1-subhalo term 96
5.4 1-subhalo term when ensemble-averaging subhalo population parameters 97
5.5 Full substructure convergence power spectrum 99
5.6 Truncated NFW and Burkert density profiles 102
5.7 1-subhalo term for populations of truncated NFW and Burkert halos . . 103

6.1 ETHOS snapshot and subhalo catalog convergence fields 113
6.2 Subhalo catalog substructure convergence power spectrum from the CDM and ETHOS4 simulations 118
6.3 Redshift and mass dependance of the subhalo catalog convergence power spectrum 119
6.4 Mass bin decomposition of the CDM substructure convergence power spectrum 123
6.5 Mass bin decomposition of the ETHOS4 substructure convergence power spectrum 124
6.6 Host dark matter convergence field 125
6.7 Simulation snapshot power spectra 126

7.1 Multi-lens plane sketch of a system with interlopers 135
7.2 Effective functions and lensing kernel for interloper halos 144
7.3 Line-of-sight integration volume 145
7.4 Interloper and subhalo power spectrum plateaus and number densities . . 157
7.5 Simulated effective convergence field 160
7.6 Numerical and analytical perturber power spectra 161
7.7 Curl and divergence components of the effective convergence field 163

8.1 The first five principal components of the dark energy equation of state . 174
8.2 Tomographic distribution of lenses $n(z)$ in the Dark Energy Survey . . . 177
8.3 Constraints on the sum of neutrino masses and S_8 from the νwCDM and $\nu\Lambda$CDM scenarios 179

8.4 Fractional difference for $D_L(z)$ and $H(z)$ for νwCDM and $\nu\Lambda$CDM cosmologies 180
8.5 Fractional difference for $P(k)$ and P_l^κ for νwCDM and $\nu\Lambda$CDM cosmologies 180

9.1 Schematic depiction of a flow-based generative model 195
9.2 Simulated weak lensing data 204
9.3 Weak lensing convergence power spectrum non-Gaussianity 207
9.4 Log-likelihood per sample and loss as a function of training epoch 210
9.5 ICA, GMM and FBL samples and log-likelihood 211
9.6 ICA, GMM and FBL non-Gaussianity comparison 213

A.1 Mass fraction in subhalos and radial subhalo distribution in simulated images with a single subhalo 224
A.2 Mass fraction in subhalos and radial subhalo distribution in simulated images with with $N_{\rm sub}$ constrained 225
A.3 Mass fraction in subhalos and radial subhalo distribution in simulated images with with $f_{\rm sub}$ constrained 227

B.1 Distribution of the true labels for pixels predicted in each class for the single-subhalo test set 229
B.2 Confusion matrix for the single-subhalo test set 231
B.3 Comparison to *Hubble Space Telescope*-like noise 234
B.4 Example of domain adaptation: subhalo population 235
B.5 Example of domain adaptation: two main lenses 236

E.6.1 Numerical curl from a single lens plane convergence field 258

F.1.1 Fractional difference for $D_L(z)$ and $H(z)$ for νwCDM and $\nu\Lambda$CDM cosmologies with the *Reduced* dataset 260
F.1.2 Fractional difference for $P(k)$ for νwCDM and $\nu\Lambda$CDM cosmologies with the *Reduced* dataset 260

G.2.1 Applying FFJORD on 2D toy data 264

G.2.2 Applying FFJORD on Gaussian 30-dimensional data with a full-rank covariance matrix 266
G.2.3 Applying FFJORD on Gaussian 30-dimensional data with a singular covariance matrix 268
G.3.1 ICA non-Gaussianity with fewer mocks 269
G.4.1 KL divergence estimator variability 270
G.5.1 Non-Gaussianity in BOSS DR12 mock galaxy power spectra 272
G.5.2 Density of individual bins in the BOSS DR12 mock galaxy power spectra 273
G.6.1 Density of individual bins in the mock weak lensing convergence power spectra 276

List of Tables

3.1 CNN architecture 28

4.1 Subhalo classification accuracy for the single-subhalo test set 59

8.1 Datasets that define the *All* and *Reduced* sets of chains 173

8.2 Mean and 95% C.L. errors for M_ν and S_8 in the νwCDM chains 178

8.3 Mean and 95% C.L. errors for M_ν and S_8 in the $\nu\Lambda$CDM chains 179

B.1 Subhalo classification accuracy for the single-subhalo test set when imposing a minimum probability threshold 232

C.1 Important constants and variables 241

D.1 Relevant properties of the CDM subhalo population at $z = 0$ 245

D.2 Relevant properties of the CDM subhalo population at $z = 0.5$ 245

D.3 Relevant properties of the CDM subhalo population at $z = 1$ 246

D.4 Relevant properties of the ETHOS4 subhalo population at $z = 0$ 246

D.5 Relevant properties of the ETHOS4 subhalo population at $z = 0.5$ 247

D.6 Relevant properties of the ETHOS4 subhalo population at $z = 1$ 247

List of Abstracts

Chapter 3: Direct Detection of Dark Matter Substructure in Strong Lens Images with Convolutional Neural Networks [1]

Strong gravitational lensing is a promising way of uncovering the nature of dark matter, by finding perturbations to images that cannot be well accounted for by modeling the lens galaxy without additional structure, be it subhalos (smaller halos within the smooth lens) or line-of-sight (LOS) halos. We present results attempting to infer the presence of substructure from images without requiring an intermediate step in which a smooth model has to be subtracted, using a simple convolutional neural network (CNN). We find that the network is only able to infer the presence of subhalos with $> 75\%$ accuracy when they have masses of $\geq 5 \times 10^9$ $M_\odot$ if they lie within the main lens galaxy. Since less massive foreground LOS halos can have the same effect as higher mass subhalos, the CNN can probe lower masses in the halo mass function. The accuracy does not improve significantly if we add a population of less massive subhalos. With the expectation of experiments such as HST and Euclid yielding thousands of high-quality strong lensing images in the next years, having a way of analyzing images quickly to identify candidates that merit further analysis to determine individual subhalo properties while preventing extensive resources being used for images that would yield null detections could be very useful. By understanding the sensitivity as a function of substructure mass, non-detections could be combined with the information from images with substructure to constrain the cold dark matter scenario, in particular if the sensitivity can be pushed to lower masses.

Chapter 4: Extracting the Subhalo Mass Function from Strong Lens Images with Image Segmentation [2, 3]

Detecting substructure within strongly lensed images is a promising route to shed light on the nature of dark matter. However, it is a challenging task, which traditionally requires detailed lens modeling and source reconstruction, taking weeks to analyze each system. We use machine learning to circumvent the need for lens and source modeling and develop a neural network to both locate subhalos in an image as well as determine their mass using the technique of image segmentation. The network is trained on images with a single subhalo located near the Einstein ring. Training in this way allows the network to learn the gravitational lensing of light and, remarkably, it is then able to accurately detect entire populations of substructure, even far from the Einstein ring. In images with a single subhalo and without noise, the network detects subhalos of mass 10^6 $M_\odot$ 62% of the time and 78% of these detected subhalos are predicted in the correct mass bin. The detection accuracy increases for heavier masses. When random noise at the level of 1% of the mean brightness of the image is included (which is a realistic approximation for the *Hubble Space Telescope*, for sources brighter than magnitude 20), the network loses sensitivity to the low-mass subhalos; with noise, the $10^{8.5}$ $M_\odot$ subhalos are detected 86% of the time, but the 10^8 $M_\odot$ subhalos are only detected 38% of the time. The false-positive rate is around 2 false subhalos per 100 images with and without noise, coming mostly from masses $m \leq 10^8$ $M_\odot$. With good accuracy and a low false-positive rate, counting the number of pixels assigned to each subhalo class over multiple images allows for a measurement of the subhalo mass function (SMF). When measured over five mass bins from 10^8 $M_\odot$ to 10^{10} $M_\odot$ the SMF slope is recovered with an error of 14.2 (16.3)% for 10 images, and this improves to 2.1 (2.6)% for 1000 images without (with 1%) noise.

Chapter 5: Power Spectrum of Dark Matter Substructure in Strong Gravitational Lenses [4]

Studying the smallest self-bound dark matter structure in our Universe can yield important clues about the fundamental particle nature of dark matter. Galaxy-scale strong gravitational lensing provides a unique way to detect and characterize dark matter substructures at cosmological distances from the Milky Way. Within the cold dark matter (CDM) paradigm, the number of low-mass subhalos within lens galaxies is expected to be large, implying that their contribution to the lensing convergence field is approximately Gaussian and could thus be described by their power spectrum. We develop here a general formalism to compute from first principles the substructure convergence power spectrum for different populations of dark matter subhalos. As an example, we apply our framework to two distinct subhalo populations: a truncated Navarro-Frenk-White subhalo population motivated by standard CDM, and a truncated cored subhalo population motivated by self-interacting dark matter (SIDM). We study in detail how the subhalo abundance, mass function, internal density profile, and concentration affect the amplitude and shape of the substructure power spectrum. We determine that the power spectrum is mostly sensitive to a specific combination of the subhalo abundance and moments of the mass function, as well as to the average tidal truncation scale of the largest subhalos included in the analysis. Interestingly, we show that the asymptotic slope of the substructure power spectrum at large wave number reflects the internal density profile of the subhalos. In particular, the SIDM power spectrum exhibits a characteristic steepening at large wave number absent in the CDM power spectrum, opening the possibility of using this observable, if at all measurable, to discern between these two scenarios.

Chapter 6: Gravitational Lensing and the Power Spectrum of Dark Matter Substructure: Insights from the ETHOS N-body Simulations [5]

Strong gravitational lensing has been identified as a promising astrophysical probe to study the particle nature of dark matter. We present a detailed study of the power

spectrum of the projected mass density (convergence) field of substructure in a Milky Way-sized halo. This power spectrum has been suggested as a key observable that can be extracted from strongly lensed images and yield important clues about the matter distribution within the lens galaxy. We use two different N-body simulations from the ETHOS framework: one with cold dark matter and another with self-interacting dark matter and a cutoff in the initial power spectrum. Despite earlier works that identified $k \gtrsim 100$ kpc^{-1} as the most promising scales to learn about the particle nature of dark matter we find that even at lower wavenumbers - which are actually within reach of observations in the near future - we can gain important information about dark matter. Comparing the amplitude and slope of the power spectrum on scales $0.1 \lesssim k/\text{kpc}^{-1} \lesssim 10$ from lenses at different redshifts can help us distinguish between cold dark matter and other exotic dark matter scenarios that alter the abundance and central densities of subhalos. Furthermore, by considering the contribution of different mass bins to the power spectrum we find that subhalos in the mass range $10^7 - 10^8$ M$_\odot$ are on average the largest contributors to the power spectrum signal on scales $2 \lesssim k/\text{kpc}^{-1} \lesssim 15$, despite the numerous subhalos with masses $> 10^8$ M$_\odot$ in a typical lens galaxy. Finally, by comparing the power spectra obtained from the subhalo catalogs to those from the particle data in the simulation snapshots we find that the seemingly-too-simple halo model is in fact a fairly good approximation to the much more complex array of substructure in the lens.

Chapter 7: Quantifying the Line-of-Sight Halo Contribution to the Dark Matter Convergence Power Spectrum from Strong Gravitational Lenses [6]

Galaxy-galaxy strong gravitational lenses have become a popular probe of dark matter (DM) by providing a window into structure formation on the smallest scales. In particular, the convergence power spectrum of subhalos within lensing galaxies has been suggested as a promising observable to study DM. However, the distances involved in strong-lensing systems are vast, and we expect the relevant volume to contain line-of-sight (LOS) halos that are not associated with the main lens. We develop a formalism

to calculate the effect of LOS halos as an effective convergence power spectrum. The multi-lens plane equation couples the angular deflections of consecutive lens planes, but by assuming that the perturbations due to the LOS halos are small, we show that they can be projected onto the main-lens plane as effective subhalos. We test our formalism by simulating lensing systems using the full multi-plane lens equation and find excellent agreement. We show how the relative contribution of LOS halos and subhalos depends on the source and lens redshift, as well as the assumed halo and subhalo mass functions. For a fiducial system with fraction of DM halo mass in substructure $f_{\rm sub} = 0.4\%$ for subhalo masses $[10^5 - 10^8]$ $M_\odot$, the interloper contribution to the power spectrum is at least several times greater than that of subhalos for source redshifts $z_s \gtrsim 0.5$. Furthermore, it is likely that for the SLACS and BELLS lenses the interloper contribution dominates: $f_{\rm sub} \gtrsim 2\%$ (4%) is needed for subhalos to dominate in SLACS (BELLS), which is higher than current upper bounds on $f_{\rm sub}$ for our mass range. Since the halo mass function is better understood from first principles, the dominance of interlopers in galaxy-galaxy lenses with high-quality imaging can be seen as a significant advantage when translating this observable into a constraint on DM.

Chapter 8: Observable Predictions for Massive-Neutrino Cosmologies with Model-Independent Dark Energy [7]

We investigate the bounds on the sum of neutrino masses in a cosmic-acceleration scenario where the equation of state $w(z)$ of dark energy (DE) is constructed in a model-independent way, using a basis of principal components (PCs) that are allowed to cross the phantom barrier $w(z) = -1$. We find that the additional freedom provided to $w(z)$ means the DE can undo changes in the background expansion induced by massive neutrinos at low redshifts. This has two significant consequences: (1) it leads to a substantial increase in the upper bound for the sum of the neutrino masses ($M_\nu < 0.33 - 0.55$ eV (95% C.L.) depending on the datasets and number of PCs included) compared to studies that choose a specific parametrization for $w(z)$; and (2) it causes $\sim 1\sigma$ deviations from ΛCDM in the luminosity distance and the Hubble expansion rate at higher redshifts ($z \gtrsim 2$), where

the contribution of DE is subdominant and there is little constraining data. The second point consequently means that there are also observable deviations in the shear power spectrum and in the matter power spectrum at low redshift, since the clustering of matter throughout cosmic time depends on the expansion rate. This provides a compelling case to pursue high-z BAO and SN measurements as a way of disentangling the effects of neutrinos and dark energy. Finally, we find that the additional freedom given to the dark energy component has the effect of lowering S_8 with respect to ΛCDM.

Chapter 9: Flow-Based Likelihoods for Non-Gaussian Inference [8]

We investigate the use of data-driven likelihoods to bypass a key assumption made in many scientific analyses, which is that the true likelihood of the data is Gaussian. In particular, we suggest using the optimization targets of flow-based generative models, a class of models that can capture complex distributions by transforming a simple base distribution through layers of nonlinearities. We call these flow-based likelihoods (FBL). We analyze the accuracy and precision of the reconstructed likelihoods on mock Gaussian data, and show that simply gauging the quality of samples drawn from the trained model is not a sufficient indicator that the true likelihood has been learned. We nevertheless demonstrate that the likelihood can be reconstructed to a precision equal to that of sampling error due to a finite sample size. We then apply FBLs to mock weak lensing convergence power spectra, a cosmological observable that is significantly non-Gaussian (NG). We find that the FBL captures the NG signatures in the data extremely well, while other commonly-used data-driven likelihoods, such as Gaussian mixture models and independent component analysis, fail to do so. This suggests that works that have found small posterior shifts in NG data with data-driven likelihoods such as these could be underestimating the impact of non-Gaussianity in parameter constraints. By introducing a suite of tests that can capture different levels of NG in the data, we show that the success or failure of traditional data-driven likelihoods can be tied back to the structure of the NG in the data. Unlike other methods, the flexibility of the FBL makes it successful

at tackling different types of NG simultaneously. Because of this, and consequently their likely applicability across datasets and domains, we encourage their use for inference when sufficient mock data are available for training.

Chapter 1

Introduction

A century ago we did not yet know that there was anything in the universe other than our own galaxy. A series of astounding theoretical and observational breakthroughs throughout the twentieth century, starting with Einstein's theory of General Relativity in 1915 and Hubble's discoveries that there are other galaxies (1924) and that the universe is expanding (1929), paved the way to establishing the standard cosmological model, called the Λ Cold Dark Matter (ΛCDM) model. It is a relatively simple model, governed primarily by six parameters: Ω_b (the energy density in baryons), Ω_c (the energy density in cold dark matter), Ω_Λ (the energy density in dark energy, assumed to be the cosmological constant), τ (the reionization optical depth), n_s (the scalar spectral index), and A_s (the curvature fluctuation amplitude).[1] Except for the scalar spectral index, which is predicted to be slightly smaller than 1 by inflation [9], the values of the ΛCDM parameters are not predicted but rather inferred from observations.

Cosmology is a unique science. Cosmologists cannot run controlled experiments, we have to rely on observations. To be more specific, we have to rely on observations of a single universe. Determining the parameters of a cosmological model therefore requires a Bayesian perspective: the data is fixed and we have to vary our models to find the best fit to observed data. Despite this limitation, we can obtain independent estimates for the cosmological parameters from different observables in our single universe. We can

[1] There are of course many extensions to this base model; some of the most popular ones include a redshift-dependant equation of state for dark energy, massive neutrinos, and a nonzero tensor-to-scalar ratio.

combine different measurements to enhance their statistical power when inferring model parameters, or we can compare the parameter values they yield to check for concordance or discordance with the model. In particular, parameters that are not inputs into the ΛCDM model but can be derived from them (and in some cases checked against measurements that are cosmology-independent), such as the Hubble constant, H_0, or the amplitude of the linear power spectrum on the scale of $8h^{-1}$ Mpc, σ_8, are immensely important to check the consistency between our universe and the ΛCDM model.

As large-scale cosmological data became available, the ΛCDM model was lauded for successfully matching observed signals. These successes spanned very different observations that probe vastly different times in cosmic history − from the primary anisotropies in the cosmic microwave background (CMB) [10, 11], to the expansion of the universe [12–14] and the growth and clustering of matter [15, 16]. Since the turn of the century, the best-fit ΛCDM model has painted a mysterious picture of the universe, with $\sim 95\%$ of the energy content in the form of unknown, invisible components called dark matter and dark energy.

However, the last few years have cast a more critical lens on the ΛCDM model as cracks in the model have started to surface. For one, we have entered the era of precision cosmology. And with high precision comes great power: comparing the parameters inferred by high-precision independent datasets has turned the era of precision cosmology into the era of cosmological tensions. Where before different datasets saw agreement in the best-fit ΛCDM model parameters, or derived parameters, they now see tension. The two most notorious tensions are the H_0 and S_8 tensions,[2] which correspond to the expansion rate of the universe and the growth of structure in the universe, respectively. The values of these parameters inferred from local measurements differ from those measured by the CMB at the level of a few standard deviations [17–21]. Whether this points to unknown new physics requires we revisit fundamental modeling choices and assumptions we've had to make.

Furthermore, many of the theoretical foundations on which the ΛCDM model rests

[2] $S_8 \equiv \sigma_8\sqrt{\Omega_m}$, where $\Omega_m = \Omega_b + \Omega_c$.

are showing signs of faltering. Despite our best efforts for several decades, direct detection experiments have failed to detect any hint of dark matter. There has also been no progress in reconciling the observed abundance of dark energy with the expectation of vacuum energy from quantum field theory, nor advances in obtaining conclusive evidence in favor of an alternative candidate to the cosmological constant as the origin of the accelerated expansion of the universe. Definitive proof for inflation, in the form of primordial B-modes imprinted in the CMB, has remained elusive. While structure formation on large scales based on the cold dark matter paradigm has worked remarkably well, probing it on small scales has proved to be a very challenging task, marred by the complexity of disentangling cosmology and dark matter physics from complex astrophysical processes that are relevant on these scales (both theoretically and observationally).

This dissertation aims to shed light on some of the most important outstanding issues with the ΛCDM model by developing statistical tools to (1) extract more and better information from cosmological and astrophysical observables, and to (2) improve the accuracy of statistical inference pipelines, questioning established modeling choices and assumptions that can give us unfair confidence in our model of the universe, skew our understanding of the universe, or both.

The first part of the dissertation, Strong Gravitational Lensing as a Dark Matter Probe, is concerned with pioneering new methods to understand the particle nature of dark matter using the distorted light of galaxies that have been strongly lensed. Strong lens images give us a window into the distribution of dark matter on the kiloparsec scale, a regime where the ΛCDM model remains untested. Chapter 2 provides background information to understand the importance of these images and how they can be used for dark matter science. We then present methods that can improve and expand the way that inference is done with these images, driven by innovative statistical and machine learning techniques. In Chapters 3 [1] and 4 [2, 3] we introduce machine learning models that can extract information on the distribution of dark matter on sub-galactic scales by analyzing strong lens images in a much faster and more model-independent way than traditional methods. In Chapters 5 [4], 6 [5], and 7 [6] we pioneer and extensively develop the

use of the strong lensing convergence power spectrum, an observable that has imprints of the statistical properties of the dark matter distribution on these scales. We show that the power spectrum's sensitivity peaks at lower masses (i.e., smaller scales) than methods that are traditionally used, making it a very promising target to either falsify or corroborate the CDM paradigm. Finally, we also assess how astrophysical processes − which can be very relevant on these scales − can lead to wrong conclusions being drawn on the particle nature of dark matter, but remain optimistic by showing that, for the high-quality imaging data we have, they are subdominant.

The second part of the dissertation, Statistical Inference in the Era of Cosmological Tensions, is concerned with improving the quality of parameter constraints obtained from data analysis pipelines. The overall goal is to emphasize the much-needed overlap between *precision* cosmology and *accuracy* cosmology. Without ensuring the parameters inferred from our data are as accurate as they are precise, we cannot know whether the persistent tensions we are seeing in the data are truly "smoking guns" for new physics or not. We scrutinize canonical modeling choices and assumptions, and replace user-imposed choices in favor of data-driven methods. In Chapter 8 we discuss the important degeneracy between dark energy and neutrinos, and how traditionally-used, restrictive parametrizations for dark energy (i.e. assuming it is the cosmological constant) can significantly underestimate the uncertainty on neutrino masses. Given the considerable gaps in our knowledge regarding the nature and behavior of dark energy, we opt for building the dark energy equation of state using principal components and let the data constrain their amplitudes. We discuss the relevance of our findings for the S_8 tension. In Chapter 9 we introduce a novel likelihood method, called flow-based likelihood (FBL), that is uniquely flexible, accurate, and general. Despite the ubiquity of Gaussian likelihoods in cosmology, many cosmological datasets are related to distributions of galaxies, which are the result of the highly non-linear process of structure formation due to gravity, and are therefore highly non-Gaussian. FBLs can remove the bias introduced in parameter constraints with the use of a wrong likelihood. By removing the inaccuracy of the likelihood as a possible source of systematic bias artificially creating tensions, FBLs could either

dispel or reinforce them.

We conclude in the section Final Remarks, discussing the impact of this work in the broader field as well as how we envision it will shape research in the years to come.

Strong Gravitational Lensing as a Dark Matter Probe

Chapter 2

Probing ΛCDM on Small Scales with Strong Gravitational Lensing

Following the many successes of the Λ Cold Dark Matter standard cosmological model at explaining the universe we observe on large scales [22–29], many have turned their sights to sub-galactic scales as a way of either reaffirming or falsifying it (see e.g. Refs. [30–41]). Many dark matter theories that are consistent on large scales have very different behaviors on sub-galactic scales. For instance, significant dark matter free streaming [42–46] or possible interactions with relativistic species [47–57] at early times can substantially reduce the number of low-mass subhalos orbiting a typical galaxy [58–60]. In addition, dark matter self-interaction [61–64] could modify the density profile of dark matter halos [60, 65–70] away from the standard CDM prediction [71]. Other dark matter particle candidates such as ultralight axions [72, 73] might also lead to interesting phenomenology on small scales (see e.g. Refs. [74, 75]). Therefore, gaining insight into the small-scale distribution and abundance of dark matter can be used to check for consistency with predictions of the CDM paradigm and, if falsified, it can offer clues as to what exotic microphysical properties it might have. Precisely because of this probing small-scale structure has become one of the most promising ways of deciphering the particle nature of dark matter.

However, assessing whether CDM provides a good fit to observations on these scales

can be significantly more difficult than it is for large-scale observations. These small-scale modes are deep in the nonlinear regime at low redshifts, and baryonic processes can play an important role on these scales [76–84], thus significantly affecting the dark matter distribution inside galaxies and their satellites. These factors make it difficult to compute robust theoretical predictions that can be compared to observations. Consequently, high resolution hydrodynamical simulations are necessary to make predictions and test observations. Not only are these simulations very computationally expensive, but how to model astrophysical phenomena accurately remains an open problem (see Ref. [85] for an overview of different approaches to modeling baryonic physics in cosmological hydrodynamical simulations).

While it is never entirely possible to neglect the influence of baryonic structures on the evolution of the small-scale dark matter distribution (see, e.g. Ref. [84]), it can be minimized by focusing our attention on the lowest mass subhalos present in galaxies. Star formation becomes increasingly inefficient as halo mass decreases [86, 87], which makes them less susceptible to baryonic feedback effects, while their abundance and internal structure are quite sensitive to the particle nature of dark matter. This combination makes them an important laboratory to test the consistency of the CDM paradigm on small scales. The flip side of the coin is that low-mass halos can be difficult − or impossible − to directly observe and characterize, even within the Local Group [88].

The fact that disentangling the impact of dark matter physics on structure formation from that of baryons is key to probing the fundamental nature of dark matter motivates the use of a gravitational method for substructure detection. Within the Local Group, methods such as tidal streams [33, 35, 89, 90] and looking at the motions of stars within the Milky Way [91], attempt to look for dark halos. Beyond the Local Group, however, galaxy-scale strong lensing systems in which a massive foreground galaxy is multiply-imaging a background source (such as another galaxy or a quasar) constitute ideal environments to study the cosmological population of low-mass (meaning sub-galactic masses) dark matter halos. The main idea behind this method is that sufficiently massive dark matter overdensities that lie close in projection to the lensed arcs or images can cause

distortions that deviate from predictions based on a *smooth lens* model (also called *macro model*). By comparing the model fit with and without clumps we can localize overdensities and attribute them a mass. They can be detected by flux-ratio anomalies in strongly lensed quasars, or in surface brightness variations in extended lensed arcs and rings (called gravitational imaging). In this dissertation we focus on galaxy-galaxy lensing systems because they are particularly well suited for dark matter science: the extended arcs/rings cover many image pixels, meaning there is a large surface area for perturbations, and mitigate the impact of microlensing by stars.

Until very recently, detections of these localized overdensities were discussed in the context of substructure/subhalos − i.e. with the assumption that the image-distorting clumps were gravitationally bound to the main lens. Constraints were phrased in terms of subhalo masses and positions, substructure abundance, and the subhalo mass function. This includes measurements from quasars [92–102], as well as those from gravitational imaging [103–109] and spatially-resolved spectroscopy [110–112]. Overall, there have been a handful of claimed detections using these methods, with masses as low as a few times 10^8 $M_\odot$.

However, there is no reason to believe that line-of-sight (LOS) field halos cannot be the source of these perturbations, since the volume between the observer and the lens, and the lens and the source, is far larger than that spanned by the main lens itself. This idea has gained traction in the literature and there has been a shift away from assuming that detections are of substructure. Evaluating, and wherever necessary, correcting this assumption is crucial: whether the sources of perturbations are subhalos or LOS halos can have profound implications for the conclusions drawn about the particle nature of dark matter.

Regardless of the origin of the signal, the examples outlined above constitute *direct* detection efforts, where individual halos are identified and characterized, and different detections can be subsequently combined to constrain the nature of dark matter (by their abundance and mass distribution). In Chapters 3 and 4 we will present work done on using machine learning to significantly speed up and expand the capabilities of direct

detection pipelines.

The direct detection approach to substructure detection is intuitive and easy to visualize − you can point at the pixels in an image and claim the presence of a dark matter overdensity at that specific location. However, phrasing the problem in this manner is not necessarily ideal. The CDM prediction that the (sub)halo mass function rises steeply towards lower masses motivates using *statistical* detection efforts to constrain the properties of entire populations of unresolvable dark overdensities [4–6, 100, 113–116]. The idea is that, although the majority of dark matter halos perturbing a strong lens image are not individually detectable, their collective perturbations can be statistically identified. In terms of connecting observations to models of dark matter, this is also a more intuitive way of phrasing the problem because dark matter theories do not make predictions about individual dark matter clumps but rather statistical predictions about the characteristics of dark matter halo populations.

In Chapters 5, 6 and 7 we discuss using the power spectrum of the projected dark matter mass density field, called the convergence power spectrum, as a statistical observable to constrain dark matter. This idea was put forth by Ref. [117] (see also Ref. [118]) and expanded upon extensively in our work [4–6]. We start off assuming that all the perturbers are subhalos, because in this regime the convergence is well defined and the problem is considerably simpler. This will allow us to gain an intuition for how to calculate this observable, and how dark matter microphysics and statistical properties of the perturber population get imprinted on it. We will also be able to compare its sensitivity on detectable scales to direct detection methods. With this foundation we will move on to the harder problem of a three-dimensional distribution of perturbers.

Due to the length of this dissertation, we redefine key equations whenever it is convenient for the reader, instead of referencing equations across chapters.

Chapter 3

Direct Detection of Dark Matter Substructure in Strong Lens Images with Convolutional Neural Networks

As discussed in Chapter 2, direct detection efforts are a popular approach to constrain the properties of dark matter from strong lens images by individually locating and characterizing small-scale ($\sim$ kpc) dark matter clumps. An important prerequisite for direct detection methods is modeling the smooth (or macro) component of the lens galaxy and reconstructing the source before (or at the same time as) inferring the presence and properties of substructure.

Modeling strong lensing systems is a complicated process that can take $\mathcal{O}$(weeks) for a single system. Different approaches to modeling a same system can yield different results (see e.g. [109]). Furthermore, mistakes in the macro model can translate into false positive substructure detections, although steps can be taken to minimize the likelihood of this happening (see e.g. [119]). Substructure is also highly degenerate with structure in the source, although in principle the presence of multiple images in the strong lensing regime can break this degeneracy, given sufficient image resolution [120, 121].

One of our primary objectives is to gauge the feasibility of sidestepping the crucial steps of modeling the smooth component of the lens and reconstructing the source, both

because of their potential to bias detections and because of their time cost. The process of inferring subhalo properties, even after a fit for the smooth lens model is obtained, is also computationally expensive and, more often than not, detailed analyses find no compelling evidence for the presence of substructure [109, 119].

Until now, there have been $\ll 100$ high-quality galaxy-galaxy lensing images available for studies of this type, which means that the time cost has not been prohibitive. In the coming years, however, hundreds or thousands of new strong lensing systems are expected to be found with optical imaging data [122–124] by experiments such as the Wide Field Infrared Survey Telescope (WFIRST), the Hubble Space Telescope (HST), the Vera Rubin Observatory, the Dark Energy Survey (DES), and *Euclid*, vastly increasing the number of images that can be used for dark matter science.[1] In an ideal world, each of these images could be analyzed individually but, in practice, having a fast method to find interesting candidates to focus resources on could accelerate the capacity of gravitational lensing-based methods to truly constrain dark matter properties.

To this end, we present results using a convolutional neural network (CNN) [125] to analyze strong lens images with varying sources, macro model parameters, and substructure populations to determine whether they are likely to contain detectable massive substructures in the vicinity of the Einstein radius of the lens. In particular, we do the inference directly on the (simulated) data, instead of first fitting a smooth model to the images, reconstructing the source, and inferring the presence of dark substructures from the residual between the smooth model image and the data.

We seek to answer the question: *what is the minimum mass that subhalos in the vicinity of the Einstein ring have to have for the neural network to identify a feature on the image as being due to their presence instead of noise or the smooth model*? Note that it is not obvious that a neural network would be able to do this at all. Traditionally, classification done with NNs is based on some of the largest or most obvious features in images they are trained on, for example with the canonical MNIST [126] or CIFAR10 [127]

[1] There are other surveys, such as the interferometer Atacama Large Millimeter/submillimeter Array (ALMA), that are expected to find many new lenses as well, but in this work we focus more on optical imaging-type data instead of, for example, working with *uv* visibilities.

datasets. Indeed, previous works that have applied CNNs to do regression or classification directly on images in the context of strong lensing have mostly focused on macro model parameters. Refs. [128] and [129] proposed using them to identify images with strong lenses in photometric data; Refs. [130] and [131] applied this deep learning method to determine the parameters (and uncertainties) of the smooth lens model using optical imaging data; and [132] used CNNs and recurrent neural networks to extract smooth lens parameters from interferometric data.

The focus of this work is fundamentally different: our goal is to address whether the image processing capability of a CNN is powerful enough to classify images based on minute differences (whether a perturber is present) even when the large-scale features of the images vary as well (the macro model parameters and sources vary from image to image). We therefore tackle the substructure detection problem as a binary classification task. Such an approach could be used to analyze any number of observed images as a filter to identify candidates that are likely to have a detectable subhalo somewhere in the image, such that traditional analyses can be carried out on them to determine substructure properties (i.e., masses and positions), while avoiding doing the same for images that would yield null detections. We emphasize that are not implying that images devoid of detectable substructure are not interesting, quite the opposite: null detections are a crucial ingredient to constrain the subhalo mass function and test the CDM paradigm (indeed, to date, constraints on subhalo properties using strong lens images are driven more by non-detections than detections). Rather, the point is that a pipeline such as the one we are suggesting in this work could determine what images have null detections in a fraction of a second instead of requiring detailed analyses for a much longer time. In essence, if the CNN's sensitivity to substructure were understood well enough, it could serve as a proxy for the detailed sensitivity function that has to be obtained on an image-by-image basis for gravitational imaging, e.g. [106, 107, 109, 119], meaning that the information from non-detections could be leveraged to constrain CDM as well. In this work, we focus in particular on galaxy-galaxy lensing systems, although this could be done for point-like sources such as quasars as well.

This chapter is organized as follows. In Section 3.1 we describe how we simulated strongly lensed images (Section 3.1.1), briefly review neural networks (Section 3.1.2), specify how we constructed our training, validation, and test sets (Section 3.1.3), and present the specific architecture and optimization parameters employed in this chapter (Section 3.1.4). In Section 3.2 we present our results and in Section 3.3 we discuss the implications of this work and conclude.

3.1 Data and Methods

3.1.1 Simulating strongly-lensed images

We employ a neural network as a *supervised* machine learning technique, meaning that, in order to learn, the algorithm requires the training data to be labelled. To train and evaluate the neural network we therefore use simulated strong lens images. We use the publicly available software package `lenstronomy` [133] to generate the images. We simulate images with 79×79 pixels that correspond to a field of view of $5.0'' \times 5.0''$, meaning they have a resolution of $0.06''$/pixel.

Each image has five different ingredients (or three, in the case of macro-only images that contain no substructure): a smooth component, a stochastic population of subhalos, a negative mass sheet to compensate for the surface mass density added in subhalos, a simulated source of light, and instrumental effects and noise. Figure 3.1 shows an example of all the components that go into simulating an image that contains substructure (fourth column), except the mass sheet since it is just constant across the image: the source (first column), smooth model (second column), and subhalo population (third column). The fifth column shows what the image looks like once it is convolved with a point spread function (PSF) kernel and noise is added to it. More details about each of these steps are provided below and in the figure caption. Figure 3.2 shows several more examples of simulated images to illustrate that the width, completeness, and shape of the Einstein ring, vary from image to image.

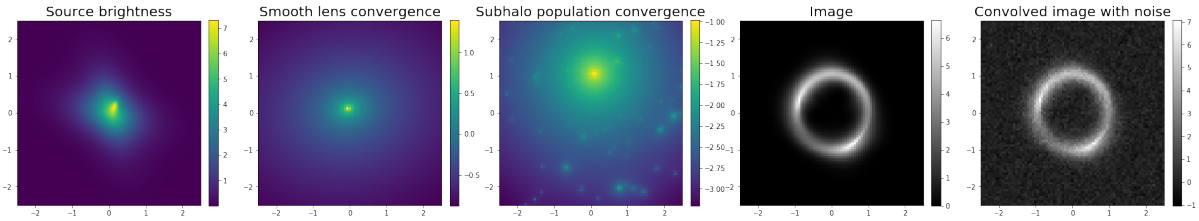

Figure 3.1: An example of the simulation pipeline (minus the negative mass sheet) for an image with a complete Einstein ring: the source brightness (first column), the convergence field of the smooth model in log units (second column), the convergence field of the subhalo population in log units (third column), the image resulting from the smooth model plus subhalos (fourth column), the image after convolving it with a PSF and adding noise (fifth column). The units of the panels are in arcseconds ($''$), each has a field of view of $5'' \times 5''$ and 79×79 pixels. The convergence field is simply the surface mass density normalized by the critical density for lensing, $\Sigma_{\rm crit} = c^2 D_{\rm os}/(4\pi G D_{\rm ol} D_{\rm ls})$, where c is the speed of light, $D_{\rm os}$ the angular diameter distance from the observer to the source, $D_{\rm ls}$ between the lens and the source, and $D_{\rm ol}$ between the observer and the lens. The image brightness is in arbitrary units of surface brightness integrated over units of an angle squared. The source has $N_{\rm clumps} = 3$, the center of the lens is at $(x, y) = (-0.05'', 0.12'')$, and its ellipticity is $(\epsilon_x, \epsilon_y) = (0.09, 0.04)$, there are 52 subhalos and the highest mass is $m_{\rm high} = 9.9 \times 10^9$ $M_\odot$. The Gaussian PSF kernel has a size of $0.07''$ and the images have Poisson noise corresponding to an exposure of 1000 seconds and 10% white noise.

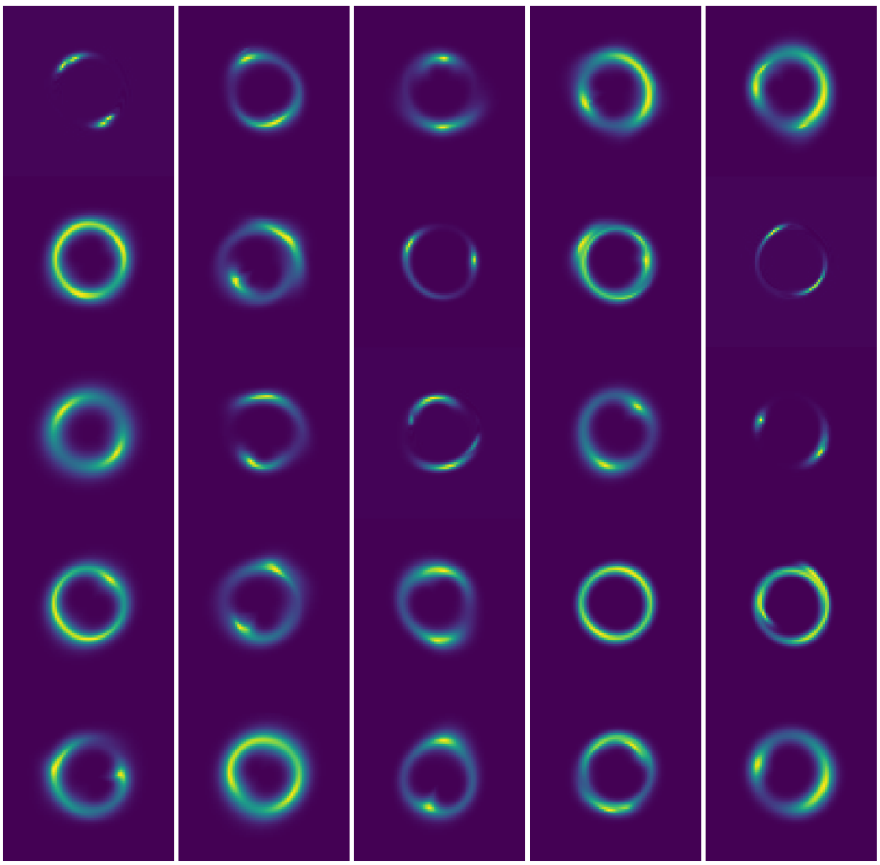

Figure 3.2: Examples of images generated with the simulation pipeline detailed in Section 3.1. The axes have been omitted for clarity, but all of these images correspond to a field of view of $5'' \times 5''$ (79×79 pixels).

Smooth lens model

We model the smooth component of the main lens as a singular isothermal ellipsoid (SIE), with surface mass density given by [134]:

$$\Sigma(\boldsymbol{r}) = \frac{f\sigma_v^2}{2G\sqrt{x^2 + f^2y^2}}, \tag{3.1}$$

where $\boldsymbol{r} = (x, y)$ is the projected two-dimensional position on the lens plane, σ_v is the velocity dispersion of the host along the line-of-sight, f the axis ratio and G the gravitational constant.

The SIE profile in `lenstronomy` is parametrized in an equivalent but slightly differently way: by the Einstein radius $\theta_{\rm E}$, which relates to the velocity dispersion by

$$\sigma_v = \sqrt{\frac{\theta_{\rm E}}{4\pi}\frac{D_{\rm os}}{D_{\rm ls}}}, \tag{3.2}$$

where $D_{\rm os}$ is the angular diameter distance from the observer to the source, and $D_{\rm ls}$ the angular diameter distance from the lens to the source; and the x and y components of the ellipticity $\{\epsilon_x, \epsilon_y\}$, which are related to the axis ratio as

$$f = 1 - \sqrt{\epsilon_x^2 + \epsilon_y^2}. \tag{3.3}$$

Finally, one can also specify the center of the lens $\{x, y\}$.

In this work, we vary x, y, ϵ_x and ϵ_y from image to image and keep the Einstein radius fixed to $1.0''$[2] for all the images (although its position does change due to the offset between the source and the center of the lens). This corresponds roughly to the typical angular size of galaxy-galaxy lenses. These parameters are chosen randomly from uniform

[2]In reality, due to the dependence of the surface mass density on the axis ratio, see Eq. (3.1), and the fact that the host has non-zero ellipticity, the Einstein radius of the images is never actually equal to one. Furthermore, the addition of substructure can change its value. See Section 3.1.3 for a discussion of the steps we took to ensure that the Einstein radii of the images with substructure was consistent with that of the macro-only images.

distributions U with ranges:

$$x, y \sim U[-0.25'', 0.25''] \tag{3.4}$$

$$\epsilon_x, \epsilon_y \sim U[0, 0.1]. \tag{3.5}$$

The lenses are placed at $z_l = 0.2$. Although it is not an intrinsic parameter of the host lens, the external shear (due to, for example, large-scale structure) is generally bunched into the parameters of the macro model. Here we do not include an external shear component, however, leaving this for future work.

Subhalo population

Our goal is to gauge the sensitivity of a neural network to perturbations caused by substructure. In the simplest case, we can therefore add a single subhalo to each image, changing its mass and position from image to image. Due to the abundance of substructure in the CDM paradigm, however, we expect there to be many subhalos in a projected area like the one under consideration in this chapter. So we can instead add a stochastic population of subhalos, where the numbers, positions and masses of the subhalos vary from image to image.

To populate our images we consider expectations from CDM and constraints from observations on the number of subhalos $N_{\rm sub}$ and the mass fraction in substructure, defined as $f_{\rm sub} = \sum_{i=1}^{N_{\rm sub}} m_i / M_{\rm host}$. Since these are functions of the range of subhalo masses considered, the host redshift, the host mass, etc., we compiled constraints for systems that are similar to our ensemble of lenses. From the high-resolution N-body simulation of a Milky Way-like halo ETHOS [60], Ref. [5] finds that between redshifts of 0 and 0.5, there are be between 25 - 35 subhalos in a projected area corresponding to the field of view of the images we are considering, by averaging over many different lines-of-sight. Note, however, that this halo is about an order of magnitude smaller than the typical masses of massive elliptical galaxies, so this can be seen as a lower bound: Ref. [135] showed that a dark matter halo eight times more massive than a Milky Way

halo can contain almost a factor of 2 more substructures with larger circular velocities. Ref. [136] finds that for a 10^{13} $M_\odot$ host, $0.1\% \lesssim f_{\rm sub} \lesssim 1\%$ between $10^6 - 10^{11}$ $M_\odot$. Using the lower limit from N-body simulations of $f_{\rm sub} = 0.3\%$ [137], Ref. [108] estimates that in the CDM paradigm, we can expect there to be 6.46 ± 0.95 substructures with masses between $4 \times 10^6 - 4 \times 10^9$ $M_\odot$ within an annulus of $0.6''$centered on the Einstein radius [137–139].

Attempts to measure $f_{\rm sub}$ from substructure detections and non-detections have yielded values that vary considerably, but all seem to roughly agree with the expectations from N-body simulations. We cite in particular constraints where the lenses were at redshifts similar to the one we are using in this work. Ref. [106] put a constraint of $f_{\rm sub} = 2.15^{+2.05}_{-1.25}\%$ or $f_{\rm sub} = 2.56^{+3.26}_{-1.50}\%$ between $4 \times 10^6 - 4 \times 10^9$ $M_\odot$, depending on the choice of prior, with a single lens at $z_l = 0.222$. More recently, using 11 SLACS lenses, Ref. [109] put a constraint of $f_{\rm sub} = 0.76^{+2.08}_{-0.52}\%$ or $f_{\rm sub} = 0.64^{+0.8}_{-0.42}\%$, depending on the choice of prior, between $4 \times 10^6 - 4 \times 10^9$ $M_\odot$ in a sample of lenses with mean $\langle z_l \rangle = 0.2$. Using the same sample of lenses, Ref. [140] finds a higher value of $f_{\rm sub} < 8.7\%$ (68% C.L.) between $10^5 - 10^{11}$ $M_\odot$, which the authors attribute to a different definition of substructure mass and mass limits, and a different shape of the substructure mass function.

To populate the host lenses with substructure we therefore consider the following points. First, we certainly want to test the networks on images that are considered "realistic" by the above guidelines. However, as can be glimpsed by the plethora of different values we gathered above, it is not necessarily clear what "realistic" means. Furthermore, the small number of actual detections and systems that have been analyzed to date means that we do not yet have a sufficient grasp of what these subhalo populations actually look like in real lenses, outside the idealized scenarios of N-body simulations that do not take into account the impact of baryonic physics. Therefore, we also want to gauge the performance of the network on a broader range of types of subhalo populations, to see how the network could fare if real subhalo populations deviate either slightly, or significantly, from the expectations based on N-body simulations.

Taking all these factors into consideration, we devise two different schemes for pop-

ulating lenses with substructure. In our first approach, which we call the $N_{\rm sub}$-bound approach, we impose a constraint on the number of subhalos that lie in the area covered by the image, and vary the highest subhalo mass $m_{\rm high}$ (Section 3.1.1). This has the consequence of having images with very different values of $f_{\rm sub}$, since it will be highly dependent on the value of $m_{\rm high}$ in an image. In the second approach, which we call $f_{\rm sub}$-bound, we instead constrain the value of $f_{\rm sub}$. This has the consequence of having very different numbers of subhalos depending on what $m_{\rm high}$ is. For instance, for $m_{\rm high} << f_{\rm sub} M_{\rm host}$, thousands of subhalos are required to satisfy the bound on $f_{\rm sub}$, while for $m_{\rm high} \lesssim f_{\rm sub} M_{\rm host}$, a handful of subhalos saturate it.

These two approaches, together with the single-subhalo case, are complimentary and shed light on different aspects of the network's sensitivity. For example, we can understand whether it is more sensitive to a larger number of lower mass subhalos, or a lesser number of more massive halos. For a given value of $m_{\rm high}$ we can also understand how the presence (or absence) of other lower mass perturbers affects the network's sensitivity. Furthermore, it allows us to explore the network's behavior among a broad range of different subhalo population characteristics.

The perturbers are placed within the lens itself, as opposed to considering line-of-sight halos that lie in the vicinity of the lensed arcs in projection. All the subhalos are always modeled with Navarro-Frenk-White (NFW) [141] profiles, and their concentrations fixed to $c = 15$. This choice is intended to be representative of the average concentration one might expect for substructure. In reality, concentration is a non-trivial function of mass and redshift. Concentration-mass relations extracted from cosmological N-body simulations show that it is a decreasing function of mass and an increasing function of redshift. In the mass ranges we are considering, the concentration is predicted to be $\gtrsim 10$. When considering subhalos, however, there is the added complication that they are tidally stripped as they move within their host's potential, so the concentration is not necessarily always well defined. Instead, subhalo profiles are sometimes characterized with an analogous parameter which is the ratio between the tidal radius and the scale radius. Using the same phenomenological equations as Ref. [4], we find that this parameter is

roughly $\lesssim 29$ (with the upper bound corresponding to subhalos at the outskirts of the host), and an average of ≈ 20. Because the tidal radius is smaller than the virial radius [138] the corresponding values of the concentration are expected to be somewhat smaller than this.

Single subhalo: We first draw a mass from a log-uniform distribution between $10^8 - 10^{11}$ $M_\odot$. `Lenstronomy` parametrizes the deflection due to NFW profiles using the scale radius in angular units and the radial deflection angle at the scale radius, and to convert physical NFW masses to these parameters we need to specify a cosmology and source and lens redshifts. We place the source at $z_s = 0.6$, the lens at $z_l = 0.2$, and use the *Planck* 2015 cosmology [142].

We allow the subhalo to be at any position in the image that satisfies two constraints:

1. It has to lie within $0.25''$ of the Einstein radius of the host, i.e. $|\theta_E - r_{\rm high}| \leq 0.25''$, where $r_{\rm high} = \sqrt{x^2_{\rm high} + y^2_{\rm high}}$.

2. The intensity at $r_{\rm high}$ must be greater than or equal to a minimum intensity threshold, unique to each image, determined by generating a macro-only image with the same macro model parameters, masking the annulus encompassed by $\theta_E \pm 0.35''$, and obtaining the maximum of the masked image.

The reason behind these constraints is that direct detection methods only have sensitivity to substructure in the vicinity of the Einstein ring or arcs in the image (e.g. [103, 104]). Therefore, if we want to know whether perturbations caused by subhalos of a given mass can be detected by the CNN, the substructure with that mass must be close to the ring/arcs. In particular, the reason for the second constraint is to make sure that, in images where the ring is largely incomplete, the subhalo still lies near an area with non-negligible intensity.

Constraining the number of subhalos: In the $N_{\rm sub}$-bound approach we generate images where the number of subhalos is drawn from a Normal distribution with mean $\mu = 60$ and standard deviation $\sigma = 15$. Once $N_{\rm sub}$ is drawn, we sample $N_{\rm sub}$ masses from a

subhalo mass function consistent with CDM, taken to be a power law with slope $\beta = -1.9$ [143] between $m_{\rm min} = 10^6$ $\mathrm{M}_\odot$ and $m_{\rm high}$, with $m_{\rm high}$ anywhere between $10^8 - 10^{11}$ $\mathrm{M}_\odot$. We again convert the masses into the scale radius in angular units and the radial deflection angle at the scale radius as required by `lenstronomy`, using the same cosmology and source and lens redshifts as above.

All the subhalo positions are chosen randomly to lie within the full area of the image[3], except for that of the most massive subhalo. The position of the most massive subhalo $r_{\rm high}$ is modified to obey the same two constraints as those detailed in Section 3.1.1 above.

Constraining the fraction of mass in substructure: In the $f_{\rm sub}$-bound approach, we fix the mass fraction in substructure to be $1 \pm 0.05\%$. For a given value of the highest subhalo mass $m_{\rm high}$, we generate draws of the subhalo mass function until the $f_{\rm sub}$ constraint is satisfied. In the cases where $m_{\rm high}$ is high enough that it already saturates the bound on $f_{\rm sub}$, we instead draw a different set of masses from the subhalo mass function that obeys the bound, and append the most massive subhalo. Alternatively we could have added a single subhalo to the image in this regime, but since we already had a dataset comprised of images with a single subhalo, we opted for this alternative approach here. In this way, we could see if there is any difference in the network's sensitivity in this regime due to an additional population of low mass subhalos.

The position of the most massive subhalo in a given image is again constrained by the two conditions described in Section 3.1.1.

More details on some of the relevant properties of the subhalo populations in the images used for training/validating/testing are provided in Appendix A.

Negative mass sheet

Since we fix the Einstein radius of the host to $1''$, we add a negative mass sheet to the substructure images to ensure that the convergence in the image is the same as for the macro-only images (since the Einstein radius depends on the convergence). The

[3]Due to projection effects, and the fact that the area probed by lensing transverse to the line-of-sight is very small, the subhalo distribution is essentially isotropic in the region of interest. See e.g. [5].

convergence field is simply the surface mass density normalized by the critical density for lensing, $\Sigma_{\rm crit} = c^2 D_{\rm os}/(4\pi G D_{\rm ol} D_{\rm ls})$, where c is the speed of light, $D_{\rm os}$ the angular diameter distance from the observer to the source, $D_{\rm ls}$ between the lens and the source, and $D_{\rm ol}$ between the observer and the lens. For each image with substructure we generate a macro-only image with the same macro-model parameters, find the total difference in convergence between the two images and add a negative convergence field that cancels this difference. All the results we present correspond to substructure and macro-model images that have the same effective Einstein radius, defined as the radius from the host center at which the convergence decreases below one.

Source

An image's sensitivity to substructure, measured as the surface brightness change $\delta I_{\rm sub}$ due to a potential perturbation caused by a subhalo $\delta\psi_{\rm sub}$, is proportional to the gradient of the source ∇S [103, 104, 120]:

$$\delta I_{\rm sub}(\boldsymbol{y}) = -\nabla S(\boldsymbol{x})|_{\boldsymbol{x}=\boldsymbol{y}-\nabla\psi_0(\boldsymbol{y})} \cdot \nabla\delta\psi_{\rm sub}(\boldsymbol{y}), \tag{3.6}$$

where $\boldsymbol{y}$ are the coordinates on the image plane and $\boldsymbol{x}$ the coordinates on the source plane. The gradient of the source brightness distribution evaluated on the source plane is translated into the image plane with the lens equation evaluated with the smooth component of the lens ψ_0. This is why highly structured sources, for example dusty star-forming galaxies that are very clumpy, are considered prime candidates to find subhalos (see e.g. [111]).

Here we focus on extended sources instead of point-like sources like quasars. We simulate sources with some degree of structure but not so much that it would be unlikely to be resolved by typical optical imaging surveys (i.e. much less structured than the simulated sources in Ref. [111] used to forecast the sensitivity of ALMA). We model the source S as one or more discrete but very close-by clumps of light s_i each modelled as a Sérsic ellipse, determined by five parameters for the ith clump: the amplitude of the intensity I_i, the half-light radius $R_{{\rm ser},i}$, the Sersic index n_i, and the x and y components

of the ellipticity $\{\epsilon_{x,i},\epsilon_{y,i}\}$.

The source is different in each image. For a given image, the number of clumps is drawn from a uniform distribution $N_{\rm clumps} \sim U[1,4]$. For each clump, the amplitude of the intensity is always fixed to unity (in arbitrary units of surface brightness integrated over units of an angle squared), and the remaining four parameters of a Sérsic ellipse are drawn from the following uniform distributions:

$$R_{{\rm ser},i} \sim U[0.1{\rm kpc}, 1{\rm kpc}] \tag{3.7}$$

$$\epsilon_{x,i}, \epsilon_{y,i} \sim U[-0.5, 0.5]. \tag{3.8}$$

$$\tag{3.9}$$

$\epsilon_{x,i}$ and $\epsilon_{y,i}$ are subject to the additional constraint that the magnitude of the ellipticity is $\epsilon_i = \sqrt{\epsilon_{x,i}^2 + \epsilon_{y,i}^2} \leq 0.4$. If there is a single clump, it is chosen to lie at the center of the image. For $N_{\rm clumps} > 1$, the relative positions of the clumps are drawn from a multivariate Normal distribution with mean $\boldsymbol{\mu} = (0,0)$ and covariance matrix with diagonal entries $\sigma_{xx}^2 = \sigma_{yy}^2 = 0.01$ and off-diagonal entries $\sigma_{xy}^2 = \sigma_{yx}^2 \sim U[-0.25, 0.25]$. The final source S is a sum of all the individual clumps, $S = \sum_{i=0}^{N_{\rm clumps}} s_i$. As mentioned above, we place the source at $z_{\rm s} = 0.6$.

Instrumental effects and noise

After the lensed image has been generated, it is convolved with a Gaussian point spread function (PSF) kernel with a full-width-half-max (FWHM) of $0.07''$ (roughly equivalent to that of HST). Then, Poisson shot noise for an exposure of 1000 seconds and Gaussian noise with a standard deviation given by some fraction p of the mean signal in the Einstein ring/arcs are added to the image; we showcase $p = \{0.01, 0.1, 0.33\}$.

Generally, the Gaussian noise added to simulated lensed images in the literature is uncorrelated and independent in each pixel. However, real data could have more complicated, correlated noise among nearby pixels due to, for example, drizzling [144]. We thus also test the performance of the NN when the noise added to the image is

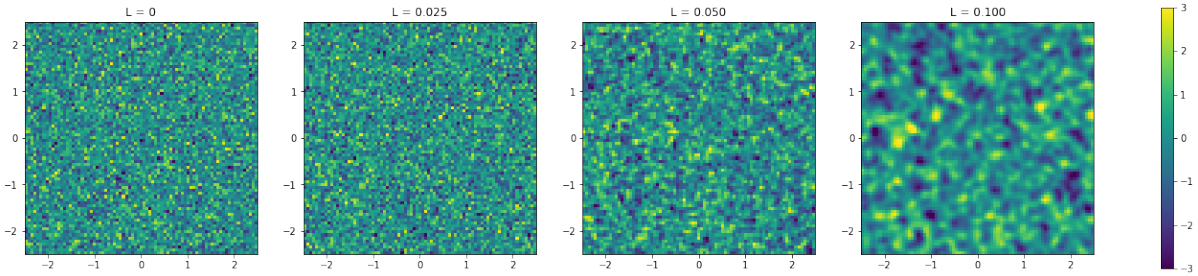

Figure 3.3: Examples of noise with varying correlation length L, obtained using the covariance matrix in Eq. (3.10). The standard deviation was fixed to $\sigma = 1$ for this Figure, and all the panels share the same colorbar (units are arbitrary). Starting from the left, the first panel corresponds to no correlation ($L = 0$). The second panel corresponds to a very small correlation length ($L = 0.025$), which is indistinguishable from the uncorrelated case. The third and fourth columns correspond to $L = 0.05$ and $L = 0.1$, respectively. The increase in correlation between the pixel values as L increases is readily apparent.

correlated. We expect that this will degrade the classification accuracy, since correlated noise could replicate the effect of subhalos more closely. We use a Gaussian Process with a squared exponential kernel K,

$$K(x, x') = \sigma^2 \exp\left(-\frac{(x - x')^2}{2L^2}\right), \tag{3.10}$$

to generate the correlated noise. σ is the standard deviation and L is the lengthscale, which determines the distance over which the pixels are correlated. We use the same value of σ as in the uncorrelated case ($L = 0$) and vary $L = \{0.05, 0.1\}$. To speed up sampling from this multivariate Gaussian distribution we use the reparametrization trick [145] with the Cholesky decomposition of the covariance matrix. Figure 3.3 shows examples of the noise with varying values of L.

3.1.2 Neural Networks

For readers unfamiliar with neural networks, we provide here a brief, high-level overview of how they work.

Neural networks are an extremely powerful tool when one has a dataset that consists of input-output pairs, $(\boldsymbol{X}_{\{i\}}, \boldsymbol{y}_{\{i\}})$ and wants to be able to obtain outputs $\hat{\boldsymbol{y}}_j$ given inputs $\boldsymbol{X}_j$ for $j \notin \{i\}$. The outputs are also referred to as (class) labels in the case of classification. Neural networks simply act as extremely complicated functions $f_{\rm NN}$ that are taught how to map an input $\boldsymbol{X}_k$ to an output $\hat{\boldsymbol{y}}_k$, $\hat{\boldsymbol{y}}_k = f_{\rm NN}(\boldsymbol{X}_k)$.

Neural networks are arranged into *layers*, and each layer contains a number of *neurons*. There are two different types of layers that are relevant to this work. The first type are *fully-connected*, also called dense, layers. These are one-dimensional layers where the neurons in a given layer are connected to all the neurons in the previous and following layers. The other type of layers are *convolutional* layers, which tend to be two- or three-dimensional. In these layers, the inputs are convolved with a *filter* that is slid over the neurons. Regardless of the type of layer, each neuron in a layer takes a linear combination of its inputs and applies a nonlinear function to them (called *activation function*).

The process of learning involves feeding the neural network many thousands (or more) of samples $\boldsymbol{X}_i$ for which their true label $\boldsymbol{y}_i$ is known, and optimizing all the parameters in the network (e.g. the weights and biases used to make linear combinations of inputs at each neuron) to minimize the *loss function*, which quantifies the difference between the predicted $\hat{\boldsymbol{y}}_i$ by the network and the true value $\boldsymbol{y_i}$.

In general, the optimization is done numerically with *stochastic gradient descent* (SGD), or some variant of it. The general idea is to use the chain rule to find the gradient of the loss with respect to every parameter in the network (called *backpropagation* in the ML jargon), and update the parameters after each iteration to minimize the loss until convergence is reached.

3.1.3 Training, validation and test sets

To convert simulated images into training/validation/test sets that can be fed into a neural network, they need to be given class labels C. Since we are trying to understand down to what subhalo mass is the CNN sensitive, we turn the problem into a binary classification task. We train a neural network on images with no substructure (macro-only images), labeled with a zero ($C = 0$), and images with substructure, with the highest subhalo mass anywhere between $m_{\rm high} = 10^8 - 10^{11}$ $\mathrm{M}_\odot$, all labeled with a one ($C = 1$).

We carry out this test with datasets that have different subhalo populations, different levels of noise, and different amounts of noise correlation. The number of samples in the training/test/validation sets can vary slightly from table to table and row to row, but

we ensure that there are at least 10^5 images for each. For cases in which there is an uneven number of images for a given class $N_{\text{samples},C=i}$, we compensate by weighing the loss function by the inverse of the fraction of training samples in a given class $f_{C=i}$:

$$w_{C=i} = \frac{1}{f_{C=i}} = \frac{N_{\text{samples}}}{N_{\text{samples},C=i}}, \tag{3.11}$$

for $i = \{0, 1\}$. When a given class has images with substructure, we ensure that there are an equal number of images with $10^8 < m_{\text{high}}/M_\odot < 10^9$ and $10^9 < m_{\text{high}}/M_\odot < 10^{11}$. We use 80% of the images for training, and 10% each for validating and testing. Furthermore, the training set is always augmented on-the-fly: each image is rotated by a random angle before going through the network, meaning that the network never sees exactly the same image twice. This helps prevent overfitting and also teaches the CNN rotational invariance.

3.1.4 CNN architecture and optimization strategy

We used `pytorch` [146] to implement the CNN. The results presented in this work are the result of a non-exhaustive grid search (see below for details) carried out using the N_{sub}-bound training/validation sets that had 1% uncorrelated noise. Our goal is to have our results serve as a proof-of-concept, showing that CNNs can become a valuable tool to help tackle an extremely complicated problem, not to spend many extra GPU hours squeezing every last point of accuracy, especially since our simulated images are not geared to replicate any one particular experiment. In reality, if one wanted to apply a pipeline like the one we are suggesting here to images taken by one (or several different) experiment(s), then a more exhaustive grid search could be carried out to improve the accuracy further.

Along these lines, we emphasize that once a good CNN architecture was found using this training/validation set, the same architecture was used for the other training sets considered. For example, the training set we used to do the grid search had images with 1% noise, but when we train a network on images with a different amount of noise, we

do not carry out a new grid search to re-optimize the network architecture. It is likely that the architecture could be fine tuned further to improve the results for the different training/validation sets.

The grid search consisted of running the networks for 100 epochs and using the accuracy over the validation set to rank the networks' performance. The parameters that were fixed or varied in the grid search are as follows. The network was forced to have two convolutional layers, each with varying filter size, stride, and number of channels. There was no zero padding, batch normalization was imposed, but maxpooling was optional. The number and width of fully connected layers was allowed to vary, from zero to five layers, and 25 to 100 nodes each. The network weights were initialized using a Normal Xavier initialization [147], whereby the weights are drawn from a Normal distribution with mean zero and variance σ^2 determined by $\sigma = \sqrt{2/(N_{\text{in}} + N_{\text{out}})}$, where N_{in} (N_{out}) is the number of input (output) neurons. Preliminary tests indicated that augmenting the data did an excellent job of preventing overfitting, so we did not implement dropout in any layer nor a regularization term in the loss function. The activation function for all layers (except for the last layer) was a ReLU function. The final outputs were passed through a softmax function, which transforms them into probabilities, by restricting them to lie between 0 and 1 and together sum to 1. We use these to assign classes: for a given image, if the probability of class 0 (1) is greater than the probability of class 1 (0), then it is assigned to class 0 (1). The loss function we used was the negative log-likelihood applied to the outputs of the softmax function, and we trained the CNN using the Adam optimizer. We set the learning rate to 0.001 and did not use a learning rate scheduler. The batch size was fixed to 64 samples per graphics processing unit (GPU), and 4 NVIDIA Tesla 2xK80s GPUs were used to train the network.

During the grid search we found that several different architectures had the same classification accuracy after 100 epochs, so out of these we picked the one with the least amount of parameters (the number of parameters varied by more than two orders of magnitude). Opting for the minimal number of parameters is beneficial to speed up training and prevent overfitting. Table 3.1 shows the network architecture that we

ultimately used. The network has 195,103 parameters.

Table 3.1: Network architecture: this network has 195,103 trainable parameters. Note that we have omitted the batch size in the dimensions of the inputs/outputs for clarity.

Number	Layer Type	Features	Dimension
1	2D Convolution	Filter size: 7 Depth: 16 Stride: 2 Maxpool: False	Input : $1 \times 79 \times 79$ Output : $16 \times 37 \times 37$
1	2D Convolution	Filter size: 4 Depth: 16 Stride: 3 Maxpool: False	Input : $16 \times 37 \times 37$ Output : $16 \times 12 \times 12$
1	Fully-connected		2304
4	Fully-connected		75
1	Fully-connected		2

3.2 Results

All the numbers given in this section correspond to having trained the network for 100 epochs with a learning rate of 10^{-3}, and an additional maximum of 30 epochs with a learning rate of 10^{-4}. All these results correspond to substructure and macro-only images that have the same effective Einstein radius.

Generally, the results for classification networks are given in terms of the overall accuracy. In this context, however, that number does not give us any insight into the NN's capacity to identify subhalos based on their mass. Furthermore, it is going to be strongly dependent on the distribution of m_{high} in the test set: we know that images with low m_{high} will be misclassified, since to the network they will be indistinguishable from macro-only images, while those with high m_{high} can be classified with more ease. This is true also of the ROC curve, which is oftentimes used to gauge the performance of a binary classifier. Due to this, we instead consider separately the accuracy for images with substructure and for macro-only images. In particular, for images with substructure we quantify the classification accuracy as a function of the highest subhalo mass.

The results for all three datasets are shown in Figure 3.4. The mass range between $10^8 - 10^9$ $M_\odot$ is omitted because its accuracy lies below 50%. Typically, in binary classification problems, the worst job a classifier can do is having around a 50% accuracy, since this corresponds to classifying randomly. Here, we can see that for low values of m_{high}, the accuracy in fact lies well below 50%. This reflects the fact that, to the network, these images are indistinguishable from macro-only images: the perturbations created by subhalos with these masses are not significant enough to create features the network can identify and distinguish from features of the source and/or macro model (or the noise, in the case of images with noise).

It is immediately apparent that the results are very similar in all three datasets for the three levels of uncorrelated noise considered. This shows us that the network is not being aided significantly by the presence of additional lower mass perturbers in the N_{sub}-bound and f_{sub}-bound datasets, and the classification is still mainly driven by the single most massive subhalo. Furthermore, while in the lowest mass bin shown the accuracy consistently surpasses 50% for low noise, we can see that subhalos have to be quite massive in order to have a classification accuracy considerably greater than 50%: the mass bin $(0.5 - 1) \times 10^{10}$ $M_\odot$ is where we start seeing accuracy $\geq 75\%$. Unsurprisingly, increasing the level of noise decreases the classification accuracy.

The addition of correlated noise, shown in the top right panel, decreases the accuracy in all the mass bins and for the macro-only images. For large correlations, the perturbers have to be somewhat more massive before they can be identified with $\geq 75\%$ accuracy, between $(1 - 5) \times 10^{10}$ $M_\odot$.

As a null test, we attempted to train a network using the same images as those in the single-subhalo dataset but without having imposed the positional constraint on the most massive subhalo, meaning it was allowed to lie anywhere in the image. Since the area of the full image is much larger than the area covered by the lensed arcs/rings, it is much more likely for the highest mass subhalo to lie somewhere where it cannot have a significant impact. Our expectation was that the network should not be able to learn to distinguish these from the macro-only images. Indeed this is what we observed.

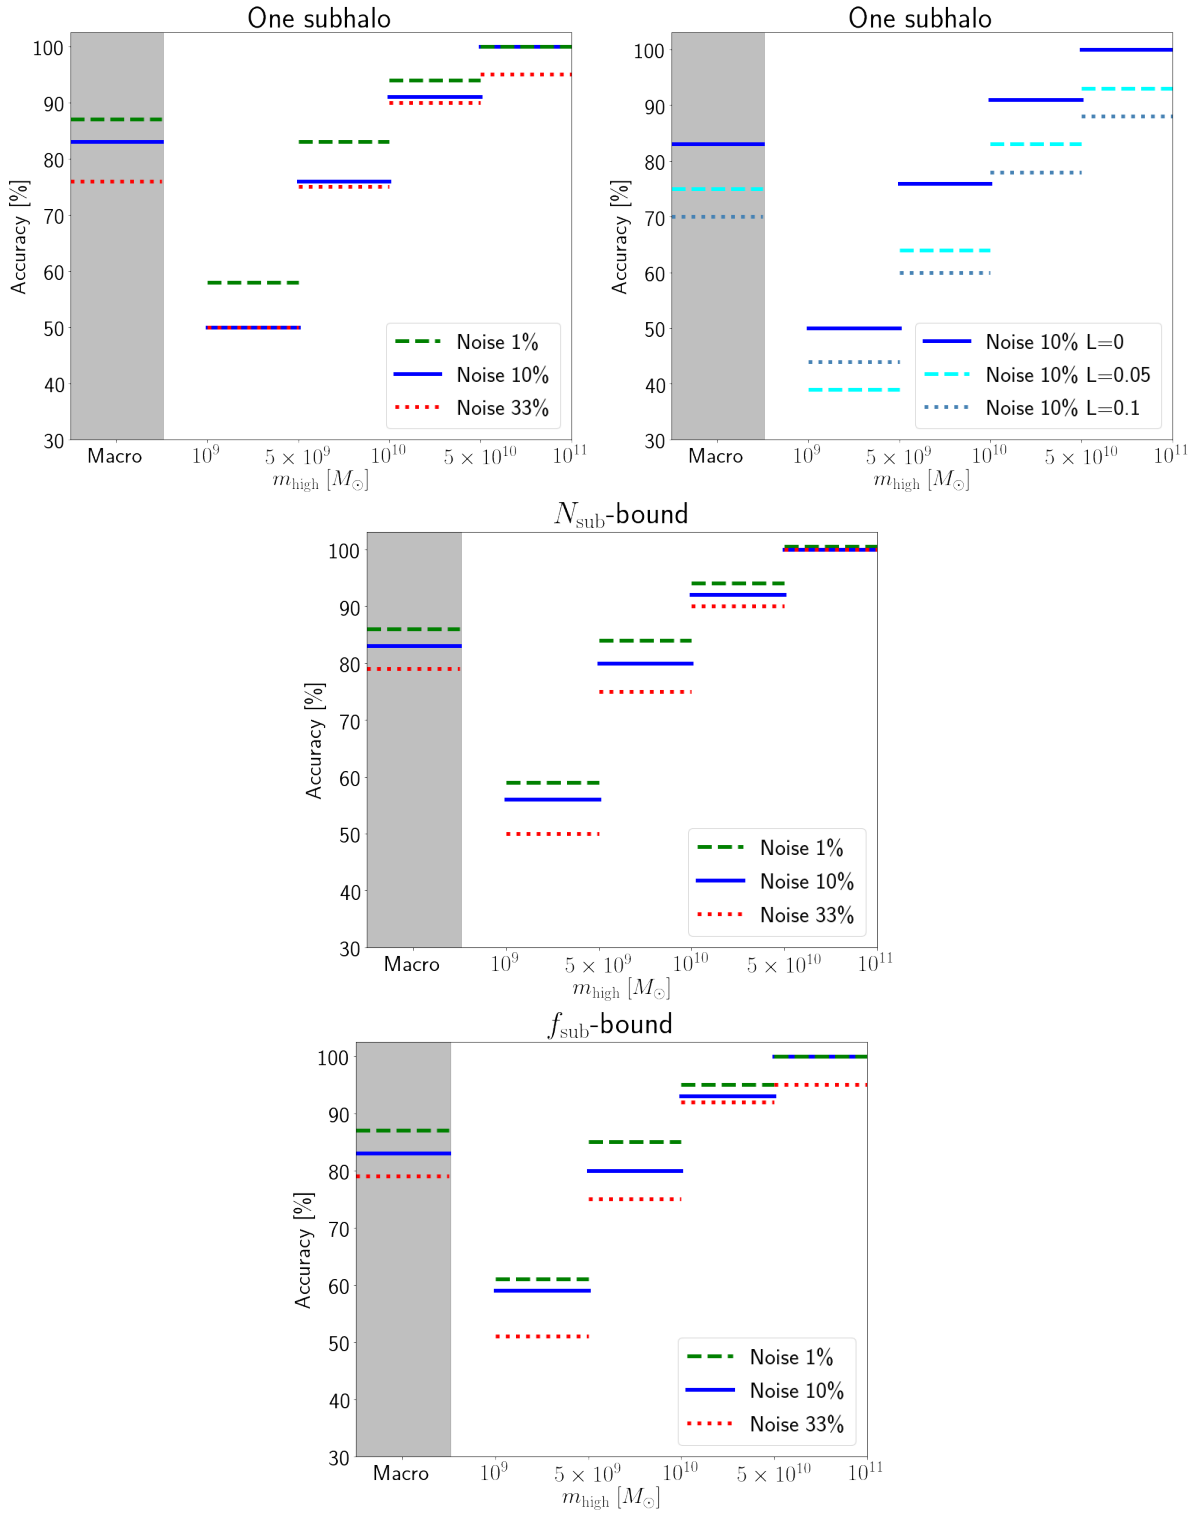

Figure 3.4: Percentage classification accuracy for the one-subhalo (top), $N_{\rm sub}$-bound (middle) and $f_{\rm sub}$-bound (bottom) test set images, as a function of the subhalo mass and for different levels of noise. In the middle panel, the accuracy line for the 1% case in the mass bin $(0.5-1) \times 10^{11}$ $M_\odot$ has been shifted upwards slightly for it to be visible, because the classification accuracy for all three levels of noise was the same. The top right panel shows the classification accuracy when the level of noise is fixed to 10% and instead the correlation lengthscale (as defined in Eq. (3.10)) is increased. All images have the same effective Einstein radius.

3.3 Discussion and Conclusion

In this chapter we have set out to explore the usability of convolutional neural networks − a machine learning technique whose image recognition capability has achieved astonishing results in many different fields over the last few years − to infer the presence of dark matter substructure in strong lens images directly, without having to carry out the inference at the level of the residual between a smooth model image and the observed image. Model fitting is troublesome in its own right: modeling choices and conventions for smooth lenses vary drastically in the literature, and it is possible to find quite different results for a same lens. Reconstructing the source is also highly non-trivial, and the strong degeneracy between source structure and subhalos (which can in theory be broken by the presence of multiple images in sufficiently high-resolution data) is a concerning source of uncertainty. Since finding a model for the smooth lens and reconstructing the source are (generally) prerequisites to finding substructure, any mistakes in these steps can trickle down into inference on the presence of substructure and its properties. Methods such as gravitational imaging mitigate the likelihood of such an eventuality by carrying out a pixel-based reconstruction of the potential instead of relying solely on an analytic fit to minimize the residual between a smooth model and a model that also has one (or more) clump(s). Such analyses have the additional advantage of being Bayesian; however, they are very computationally expensive and time consuming. Most importantly, more often than not they result in null detections.

With the expectation of thousands of new high-quality strong lens images becoming available in the near future, these factors thus motivate the development of fast, model-independent techniques to analyze strong lens images and find substructure, or come up with principled ways of choosing how to divert resources to where they can have the largest scientific impact. To this end, we train a CNN to classify images based on whether they have substructure or not. We emphasize that this classification problem is non-trivial because we are asking the network to classify images based on minute features while introducing huge variations in the large-scale characteristics from image to image, since the macro and source model parameters (which are highly degenerate with

the substructure) vary from image to image. We believe that phrasing the substructure problem in this binary way could be advantageous because it means that the images that are found to contain no detectable substructure (i.e. classified as indistinguishable from macro-only images) will not have to see more resources diverted towards them just to return null detections. Instead, resources and time can be spent analyzing images that the CNN finds are likely to have detectable substructure. If the network's sensitivity is understood sufficiently well (analogously to the sensitivity function in gravitational imaging), detections could be leveraged with non-detections to constrain the cold dark matter scenario.

We found, however, that subhalos have to be very massive, $m_{\rm sub} \gtrsim 5 \times 10^9$ $\mathrm{M}_\odot$, in the vicinity of the Einstein ring in order to be recognized with an accuracy $> 75\%$. Furthermore, the sensitivity does not seem to improve noticeably due to the presence of a larger population of lower mass subhalos, meaning that the classification is essentially driven by single very massive perturbers.

Comparisons between the sensitivity of this network and that of different methods to detect substructure is not straightforward, since sensitivity to substructure is a function of many different variables, such as the image resolution, the noise, and the source structure. Furthermore, we have not taken into account additional complicating factors in our simulated data, such as the host (or other sources of) brightness, or cosmic rays/bad pixels.

However, we can attempt to put the capacity of this CNN into context, keeping in mind the simplified nature of our simulated data. In Refs. [103, 104], the gravitational imaging technique applied to HST-like simulated images (with a resolution of $0.05''$/pixel and a signal-to-noise of *at least* 3 per pixel) was shown to have a sensitivity to subhalo masses as low as a few times 10^8 $\mathrm{M}_\odot$ for an NFW profile when the substructure is on the Einstein ring, and quickly increases with distance from the lensed images.

In terms of actual detections, to date two systems with compelling evidence for substructure have been found with gravitational imaging. One of these, SDSSJ0946+1006, was found to have a subhalo with mass $(3.51 \pm 0.15) \times 10^9$ $\mathrm{M}_\odot$ [106], while the other,

JVAS B1938+666, was found to have a subhalo with mass $(1.9 \pm 0.1) \times 10^8$ $M_\odot$ [107]. In both cases the subhalos were modelled as truncated pseudo-Jaffe profiles [99] to obtain mass estimates; mass estimates done with NFW profiles tend to recover masses that can be significantly higher (for instance around $\sim 10^{10}$ $M_\odot$ for the 3.51×10^9 $M_\odot$ Pseudo-Jaffe subhalo in SDSSJ0946+1006 [140]).

It therefore seems like the sensitivity of this CNN might be sufficient to find the perturber in SDSSJ0946+1006; it is less likely that it could find the one in JVAS B1938+666. It is worth keeping in mind, however, that the values for the accuracy as a function of subhalo mas cited in this work are the *true* subhalo masses, while direct detection efforts are sensitive only to the *effective* subhalo mass: Ref. [148] showed that the true subhalo mass can be biased by up to an order of magnitude higher than what is actually measured with strong lensing, the effective mass, meaning it is possible that claimed substructure detections actually have higher true masses than the numbers that are given.

If one did want to apply a method such as this to real data, much work would have to be done to understand the generalizability of the network's classification capacity. We explicitly avoided fine-tuning our simulation pipeline to emulate observations of a particular experiment and instead remained agnostic to serve as a proof-of-concept that could be relevant to any survey that produces similar images (i.e. it would not be valid for interferometric data, for example), so this would require matching the experimental specifications of the images in each survey as well as certainly adding complexity to the simulated data. Furthermore, while in this work we fixed the radius of the Einstein ring to the typical value of galaxy-galaxy lenses, and did not add an external shear component, we have to understand whether the network's capacity is robust to varying these two macro-model components, since known galaxy-scale strong lens images span a range of values for both of these parameters and they can be degenerate with subhalo properties. Another important aspect of the generalizability of the network relates to the implicit modelling of the smooth component of the lens. Although we have not done any explicit smooth modeling to determine the presence or absence of substructure, the network has implicitly learned about the SIE density profile since all the samples in our training data

had smooth SIE components. Although many galaxy-scale lenses seem to be well fit with SIE profiles, gauging the network's performance when trained on images simulated with a variety of density profiles would be advantageous before testing on real images. Similarly, training the network on images where the source and lens redshifts vary, the PSF is allowed to be anisotropic, and additional models for the source components are included, could also reveal important information about the network's applicability to real data. Understanding all of these factors would be crucial in order to be able to leverage the images classified as non-detections to constrain CDM together with the information extracted from images with substructure. Finally, it would be advantageous to develop a method to quantify confidence in classification, for instance using Gaussian Processes, so that images that are likely to be false positives/negatives can be identified. We leave these to future work.

In addition, we emphasize that there is no reason to believe that an approach such as this one would not be valid for images from an experiment such as ALMA, if the network were trained with an appropriate dataset. This could be an interesting extension of this work, since lens modeling in configuration (uv-visibility) space is even more time consuming than in real space, and experiments such as ALMA are expected to produce very high-quality strong lens images.

The results presented in this work were produced with a CNN whose architecture was the result of a grid search using training/validation samples from the $N_{\rm sub}$-bound dataset in which the image had 1% uncorrelated noise. It is therefore possible that the architecture could be optimized further for the different datasets, improving these results. More generally, if one had a specific experiment in mind, a grid search could be carried out with images that contained the expected levels of noise and any other relevant experimental details, such as the (possibly anisotropic) PSF.

Furthermore, the network architecture used in this work is very simple, and has few parameters compared to many typical convolutional net architectures used in the literature. For instance, well-known networks such as AlexNet [149], GoogLeNet [150], ResNet [151] and DenseNet [152] can have tens of millions of parameters. These typically

have many more convolutional layers and use additional tricks for training; for example, DenseNet has "dense blocks", within which the feature maps at each layer are concatenated to the input of every successive layer within the block, allowing later layers to leverage information from earlier layers. For instance, our results seem to suggest that the network cannot leverage information from the collective perturbations of lower mass halos, while the recent results from Refs. [153] and [154], which both used CNNs based on ResNet to infer different aspects of substructure populations from strong lenses (the subhalo mass function normalization and slope in the former, and to distinguish between substructure pertaining to two very different dark matter scenarios in the latter), show that machine learning methods can in fact be used to probe more than the single most massive halos in lenses.

We are therefore optimistic that there is room for improvement with respect to the classification capacity of the network we are using here by using a more complex CNN architecture. This is important since, if such an approach were to be pursued moving forward, uncertainties derived on, for example, substructure population parameters or likelihood of CDM, would be inextricably tied to the network's classification accuracy.

A final remark we want to bring up is with regard to the perturbations from substructure versus from line-of-sight (LOS) halos outside of the main lens halo. Original studies about perturbations to lensed images focused on subhalos as the perturbers. However, it has been pointed out that the contribution of the latter is actually likely to be comparable to, or even greater than, that of the substructure within the lens [155, 156], meaning that any attempt to use strong lensing images to constrain the particle nature of dark matter must take into account both contributions. This is particularly relevant because the effect of a LOS halo between the observer and the lens is larger than that of a subhalo of the same mass [156]. What this means is that in fact the sensitivity limits we give here for the *subhalo* mass function substructure can actually translate into sensitivity to lower masses in the full *halo* mass function if we consider LOS halos.

Strong gravitational lensing as a probe of the particle nature of dark matter has harnessed much interest over the last few years. This, together with the advent of a

huge increase in the amount of high-quality strong lens images available for dark matter science, has led to an explosion of research into methods of extracting information from strong lens images. Previous work has used ML to infer strong lens parameters [130–132], to reconstruct the sources from strong lens images [157], and most recently to infer properties of the substructure population [153, 154]. Now, this work is another step forwards towards understanding the usability of deep learning methods to speed up the analysis of strong lens images for dark matter science. With the considerable momentum that this subfield is gaining, it is possible that in the near future strong lensing will consolidate itself as one of the premier ways to uncover the nature of dark matter.

Chapter 4

Extracting the Subhalo Mass Function from Strong Lens Images with Image Segmentation

In the previous chapter we developed a technique to identify promising strong lens candidates (in the sense that they were likely to not yield null substructure detections). Thus, our ML approach would fit in with traditional inference pipelines, where detections and non-detections can ultimately be combined with a joint likelihood to infer population-level parameters. In this chapter, we improve upon this work and build a fully ML-based pipeline for substructure detection and inference on population-level characteristics (the logarithmic slope of the subhalo mass function). We do so using a technique for object detection called image segmentation. Much like in the previous chapter, we use simulated strong lens images to train and test our model.

Image segmentation seeks to classify every pixel in an image. We use a U-Net [158] architecture for this task, which consists of many convolutional layers, along with a down-sampling portion and an up-sampling portion to help the network detect features at different scales. The U-Net was designed to track cells in biological images but it has since emerged as one of the best architectures for image segmentation more broadly. We use it to classify each pixel in an image as belonging to one of several predetermined

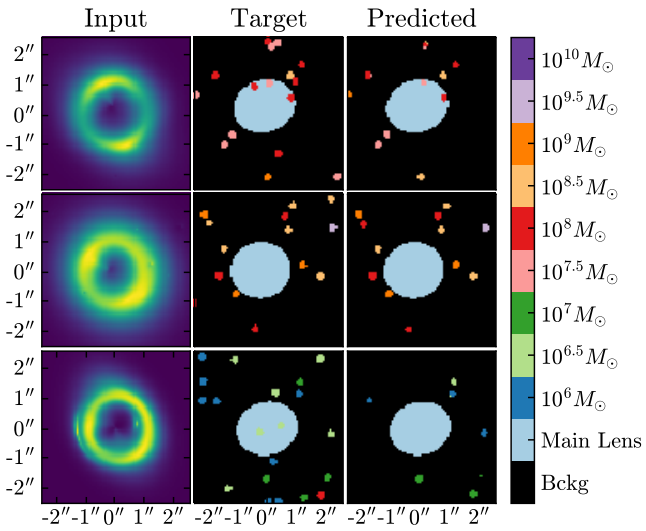

Figure 4.1: Examples of segmenting gravitationally-lensed images. The left panels show simulated images which are fed into the neural network. Each pixel in an image is mapped to a label as either being the main lens within the Einstein radius, having a subhalo of a given mass, or as none of these (background). The middle panels show the true labels for the corresponding input image. The effects of only a few of the subhalos are visible by eye. The right panels show the corresponding output of our neural network. The network was trained only on images with a single subhalo, but it is still able to find substructure in images with a rich population of subhalos. Subhalos in dimmer pixels (either in the center of the ring or the edge of the image) are more likely to be missed. Subhalos that are close together sometimes get detected as a single subhalo with the combined mass. As an example, these images do not contain noise. The effect of noise is to reduce the sensitivity to low-mass subhalos.

classes the network is trained to identify. Each pixel in our simulated images can fall into one of eleven different classes: part of the main lens, a subhalo with a mass within one of nine mass bins, or neither (background). At the pixel level, this is a classification task. However, it allows us to both locate and get the mass of substructure in the gravitational lens. Furthermore, as opposed to the traditional direct detection methods, the mass predicted by the network corresponds to the true (simulated) subhalo mass, not the effective mass.

An example of our image segmentation is shown in Fig. 4.1. The network takes in as input a lensed image as given by the left column and labels each pixel, quickly identifying both the main lens and the substructure, as shown by the right panel. The mass of the subhalos is denoted by the color of the pixel. In the middle column, we have used truth knowledge of the lens to label the area within the Einstein radius of the main lens and the different subhalos. The network has good accuracy, but does struggle to detect substructure near the edges of the image. We also see that it does not appear to add in

spurious subhalos. These two factors (good accuracy and low false-positive rate) allow us to apply this method to incorporate the step of combining images to put constraints on the subhalo mass function into a single analysis pipeline.

This chapter is organized as follows. In Sec. 4.1, we discuss our image generation pipeline; it shares many features with the pipeline in Chapter 3 but there are some key differences and additions. A detailed discussion of the network setup and the training regiment is contained in Sec. 4.2. The effect of noise in the network training is explored in Sec. 4.3. We study the accuracy of the network and the false-positive rate in Sec. 4.4. In Sec. 4.5, we apply the network to images with multiple subhalos and infer the subhalo mass function. We conclude in Sec. 4.6.

4.1 Data generation

The goal of our work is to detect dark substructure within strong lens images using the technique of image segmentation. This is a supervised learning problem, so in order to train the network, a set of training data, including the target labels, is needed. We generate strongly-lensed images using the software package `lenstronomy` [133]. We use images with 80×80 pixels with a field of view of $5'' \times 5''$. This corresponds to a resolution of $0.06''$ per pixel. Each image contains a background source light, a smooth lens, possibly substructure in the lens, possibly noise, and it is always convolved with a point spread function (PSF) of $0.07''$. In each image, the gravitational lens (main halo and subhalos) as well as the source light are unique. Each step in the simulation pipeline is detailed below.

Smooth Lens: The halo of the main lens is chosen as a singular isothermal ellipsoid (SIE) [134]. In `lenstronomy`, the SIE is parametrized by the Einstein radius (θ_E) and the ellipticity moduli. We choose the size of the Einstein radius to be typical of observed strongly-lensed galaxy-galaxy systems, drawn from a uniform distribution

$$\theta_E \in U[0.95, 1.05]'' \ , \tag{4.1}$$

and an ellipticity drawn from

$$\epsilon_{1,2} \in U[0, 0.1] \ . \tag{4.2}$$

The center is chosen to be near the middle of the image so that the images/arcs lie within the field of view. The x and y positions are drawn randomly as

$$x, y \in U[-0.25, 0.25]'' \ . \tag{4.3}$$

In this work, we fix the distance to the lens at a redshift of $z_{\rm lens} = 0.2$. Our fiducial cosmology is given by *Planck* 2015 results in Ref. [159]. These lens parameters (and the location of the source light) were chosen such that the main lens has a mass of order 10^{13} $\mathrm{M}_\odot$ (depending on the specific Einstein radius in a given image).

Subhalos: When we add substructure to the lens, it is modeled as a truncated NFW profile [160] with a concentration parameter $c = 15$. The subhalos are truncated at five times the scale radius. The network is trained on images that have either zero or one subhalo. The subhalo masses are chosen to be log-uniform over the mass range $[10^{5.75} - 10^{10.25}]$ $\mathrm{M}_\odot$, such that we obtain equal numbers of images in each mass bin. The subhalos are placed near the Einstein radius, where their effects are largest, defined by pixels which are at least 20% as bright as the brightest pixel in the image.

In Sec. 4.5, we test a trained network on images with many subhalos. For these images, the masses are drawn according to a power-law given by

$$\frac{dN}{dM} = a_0 \left(\frac{M}{m_0} \right)^{\beta} , \tag{4.4}$$

which was found to be a good fit of the subhalo population in the Aquarius simulation [143] with power-law index $\beta = -1.9$, amplitude $a_0 = 3.26 \times 10^{-5}$ $\mathrm{M}_\odot^{-1}$, and pivot point of $m_0 = 2.52 \times 10^7$ $\mathrm{M}_\odot$ although this specific normalization doesn't necessarily apply outside of the Aquarius simulation, the general form of Eq. (4.4) is universally found in N-body CDM simulations. In such simulations, the three-dimensional distribution of subhalo positions is nearly spherically symmetric, with a strong dependence on the

radius. However, because their positions are projected onto a single plane, and the fact that strong lens images have a small field of view compared to the full extent of the halo perpendicular to the line of sight, a uniform distribution for subhalo positions is a good approximation (for example, see Ref. [5]). Thus, in our images with many subhalos, we will draw their locations uniformly across the whole image.

Target Labels: The target labels, which the network is trying to predict, are generated only from the smooth lens and substructure, and do not use information from the source light or the observed image. Pixels which are inside the (possibly elliptical) Einstein radius are identified as the main lens class. For each subhalo, we draw a circle with a radius of 2 pixels centered on its location and assign all the pixels within the circle as belonging to a given subhalo mass bin class. Any pixel which has not been labeled as the main lens or a subhalo is denoted as background. This method of identifying subhalos treats all subhalo masses as identical in that the more (less) massive subhalos do not get larger (smaller) circles, even though their effects are larger (smaller). We choose to do this for two reasons. The first is that it leads to more stable training of the network. When the pixel labels change size with different masses, there are many more training pixels for the heavier classes than the light classes. This creates an imbalance that would need to be corrected for, be it by weighting subhalo classes differently in the loss function or having different numbers of training images for each subhalo mass bin. The second reason is that it makes counting subhalos easier. The predicted subhalo count can be obtained by dividing the total number of pixels predicted to be part of a subhalo mass bin by the expected area per subhalo (4π pixels). Furthermore, as we will discuss extensively in Sec. 4.5, we are interested in extracting the subhalo mass function from an ensemble of images, for which we simply need the number of subhalos in each mass bin, and therefore do not need to faithfully reconstruct the surface mass density on the lens plane.

Source Light: This work is focused on images of lensed extended objects because they offer better chances for detecting substructure than point source-like objects [111]. We do not use galaxy images (simulated or real), but allow for some structure in the source, placing between one and four clumps of light, as done in Ref. [1]. Each clump is

modeled as a Sérsic ellipse with a radius randomly drawn from a uniform distribution

$$R_{\mathrm{ser},i} \in U[0.1 \text{ kpc}, 1.0 \text{ kpc}]$$

and ellipticity

$$\epsilon_{x,i}, \epsilon_{y,i} \in U[-0.5, 0.5],$$

subject to the constraint $\sqrt{\epsilon_{x,i}^2 + \epsilon_{y,i}^2} \leq 0.4$. The location of each clump is drawn from a multivariate Gaussian with a covariance matrix with diagonal elements $\sigma_{xx}^2 = 0.01^2$, $\sigma_{yy}^2 = 0.01^2$, and off-diagonal element σ_{xy}^2 randomly chosen for each image, with a range of $[-0.25, 0.25]$. The source is placed at a distance of redshift $z_{\mathrm{source}} = 0.6$. These parameters give images that resemble true strongly-lensed galaxies rather than 4 very far-away clumps of light that are lensed into four different images. The amplitude of the intensity of each clump is set to unity (in arbitrary units of surface brightness integrated over units of angle squared).

Noise and detector effects: We consider images both with and without noise to understand the strengths and limitations of our image segmentation-based model. When noise is added, it is drawn, per-pixel, from a Gaussian with mean $\mu = 0$. The standard deviation of the Gaussian is set to 1% of the mean brightness near the Einstein ring, unless otherwise stated. This sounds like an optimistic level of noise considering the current quality of galaxy-galaxy lens images (Ref. [109] looks for substructure in images that have a signal-to-noise ratio of at least 3). However, we show in Appendix B.4 that our 1% Gaussian noise approximation results in similar signal-to-noise ratio profiles for sources brighter than magnitude 20 compared to a more detailed *Hubble Space Telescope* (HST)-like noise simulation. This choice also allows us to accurately quantify how the network is affected by noise. We do not include detector effects other than convolving the image with a Gaussian PSF kernel with a full-width half-maximum of $0.07''$, which is comparable to that of the HST. The PSF is applied to all images, regardless of the noise.

With the input images and the target labels, we are able to train the network. The specific model setup and the training details are described in more detail in the next

section.

4.2 Model architecture and training

Much of the work of the machine learning community is about object detection. For instance, a self-driving car needs to be able to identify cars, pedestrians, traffic lights, and so forth. One common method is placing boxes inside the image and then classifying what is inside the boxes, see Ref. [161] for a review. An alternative method is to classify all of the individual pixels in an image. This technique is referred to as image segmentation and a recent review of various methods can be found in Ref. [162]. While there are many models for segmenting images, the best-performing ones have some similar features. First, they are fully convolutional (in that there are no fully-connected layers). The best models also tend to have an encoder-decoder structure. This allows the models to extract features across different scales and return a high-resolution segmentation map. In particular, our network is based on the U-Net architecture [158], which has excellent sensitivity to small objects in images.

In this work, we use a U-Net to classify each pixel in a strongly lensed image into one of 11 classes. The classes are broken down as: belonging to the main lens, a subhalo with mass $\{10^6, 10^{6.5}, 10^7, 10^{7.5}, 10^8, 10^{8.5}, 10^9, 10^{9.5}, 10^{10}\}$ $M_\odot$, or none of the above (which we will refer to as background throughout the chapter). While the goal of the network is classification, it allows us to both locate subhalos and obtain their mass.

Before an image is put through the network, the image is pre-processed by dividing by the maximum pixel value. This normalization helps by forcing the brightness in all the images to have similar ranges, since despite the fact that all our source-light clumps have the same intensity, different lens parameters lead to differing amounts of magnification. Consequently, the network cannot base its classification on the absolute brightness of an image.

Once the image has been normalized, it is ready to be segmented. Our U-Net model architecture is implemented in `PyTorch` [146], and Fig. 4.2 depicts our specific set up.

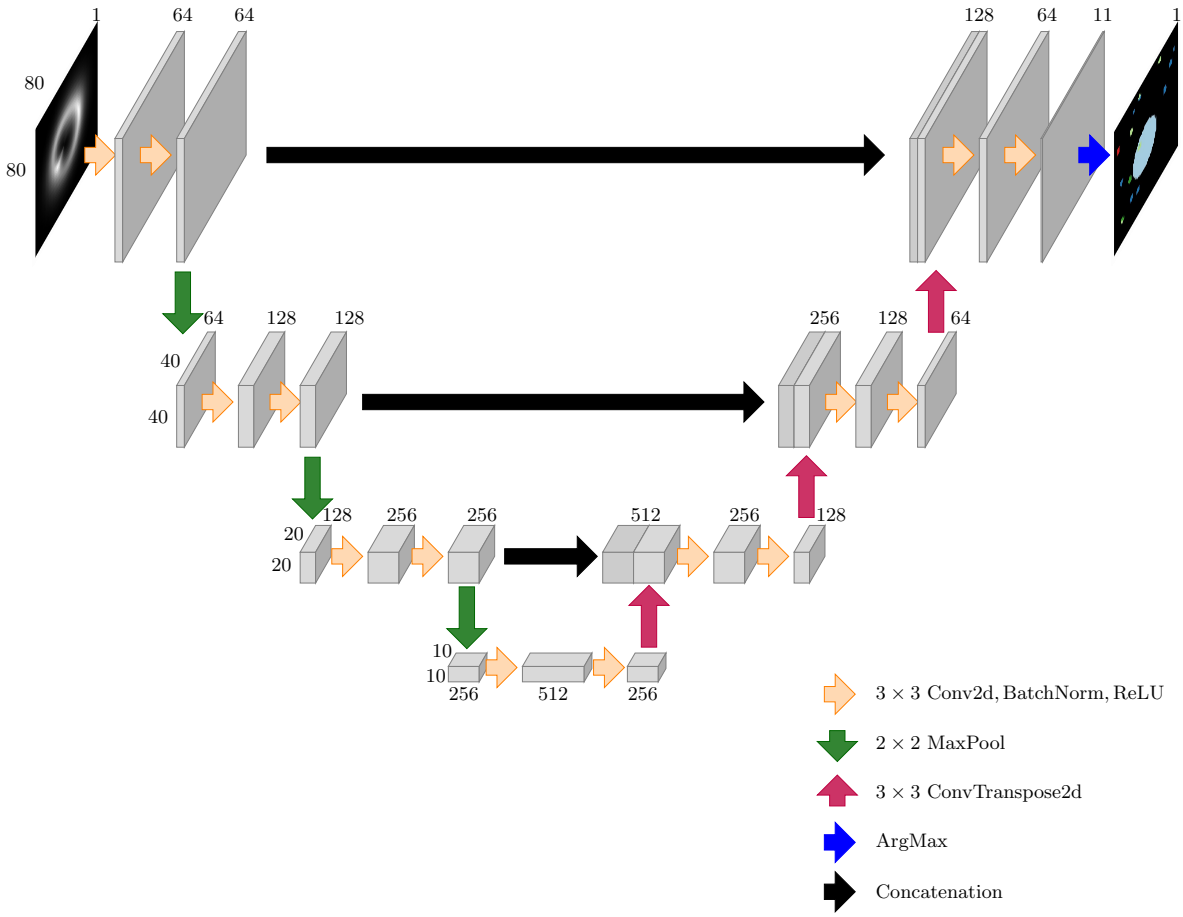

Figure 4.2: Network architecture. It takes in an 80×80 pixels image with a single layer and returns an image of the same size. The pixel values in the output correspond to the predicted class.

Each orange arrow represents three operations. The first one is a 2D convolution in which a number of 3×3 pixels filters with learnable weights are convolved with the image. The number of filters for each convolutional layer is denoted above each layer in the figure. The second operation is batch normalization [163], which normalizes the data after the convolution, leads to faster training, and helps regularize the network. The final operation represented by the orange arrows is applying the rectified linear unit (ReLU) activation function to the normalized data. This is given by

$$\mathrm{ReLU}(x) = \begin{cases} 0, & x < 0 \\ x, & x \geq 0 \end{cases} . \tag{4.5}$$

The convolutions in a given block are padded to preserve the number of pixels.

The green and red arrows depict the down- and up-sampling procedures, which cut

the number of pixels in half and double the pixel count, respectively. The down-sampling is done with a 2×2 maximum pooling operation. The up-sampling is done with a transposed convolution operation. Note that the height and width of the data at each stage is marked at the beginning of each row in the figure. Repeating the convolutional blocks (orange arrows) on the down sampled data with the same filter size allows the network to detect features at larger scales. The up-sampling transmits the information from these other scales back to the previous scale. After the up-sampling, the layer is concatenated with the last layer of the same height and width before down sampling (shown by the black arrows). This allows the network to localize the new features and to avoid losing pattern information.

After the last convolution, our images have a depth of 11 channels corresponding to the 11 classes. We apply the Softmax function along the channel such that the sum of a given pixel across all 11 channels is unity, and therefore its value for each channel can be thought of as a probability of belonging to the corresponding class. Explicitly, this is given by

$$\text{Softmax}\left(z_i\right) = \frac{e^{z_i}}{\sum_{k=1}^{K} e^{z_k}} \equiv \widehat{p}_i, \tag{4.6}$$

where z is the output for a pixel, the subscript denotes the pixel channel, and K is the total number of channels, 11 for the problem at hand. In this way, we interpret the channel to represent the predicted probability of belonging to a given class, denoted by $\widehat{p}_i$.

We train the network using a set of 9×10^4 images. Of these images, 9×10^3 have only the source light and a smooth lens. The remaining training images additionally contain exactly one subhalo. There are 9×10^3 images for each of the nine mass bins. We use an independent set of 10^4 images, with 10^3 from each of the sets mentioned above, to validate the model.

As a classification problem, the cross-entropy loss per pixel is used. This is given by

$$L = \frac{-1}{n \times p} \sum_{i=1}^{n} \sum_{j=1}^{p} \sum_{k=1}^{K} y_k^{(i,j)} \ \log\left(\widehat{p}_k^{(i,j)}\right), \tag{4.7}$$

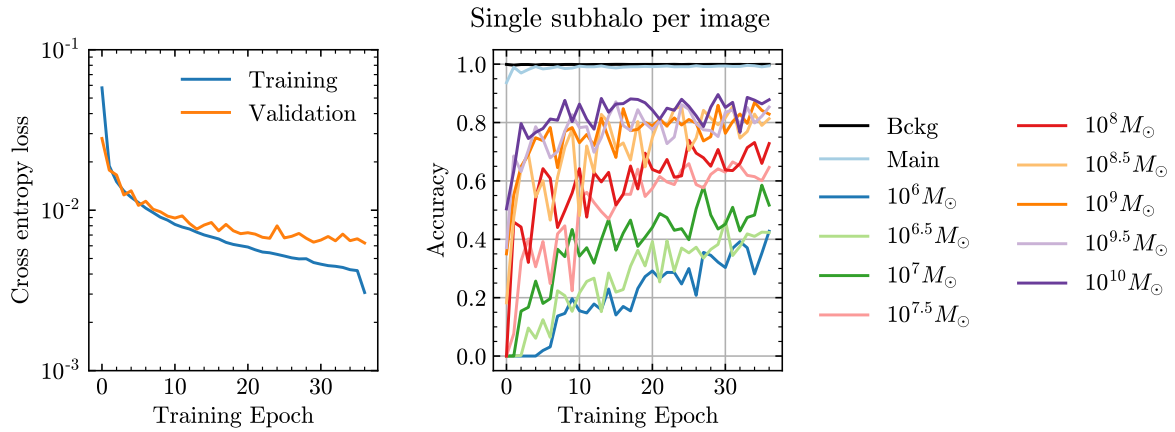

Figure 4.3: Example of training on images with no noise. The left panel displays the categorical cross entropy loss as a function of the training epoch. The lowest validation loss occurs at after epoch 36, but the network is allowed to continue training until the validation loss has not improved for 15 epochs. The learned parameters from the epoch with the lowest validation loss are used when applying the network to new data. The right panel shows the per-pixel accuracy of the validation data for each of the 11 classes. The background and smooth lens pixels are predicted correctly nearly 100% of the time. The subhalos pixels are located and assigned the correct mass with accuracy between $40\% - 85\%$.

where the sum over i goes over the n images, the sum over j runs over all of the p pixels in an image, the sum over k is the different possible classes, $y_k^{(i,j)}$ represents the true probability of pixel j in image i to belong to class k. As the true pixel is either in a given class or not, $y_k^{(i,j)}$ is either 0 or 1. Finally, $\widehat{p}$ is the probability predicted by the model.

We minimize the loss using the Adam optimizer [164] with a learning rate of 10^{-3} and the default β values. The batch size is set to 100 images. When the loss evaluated on the validation set has not improved for 5 epochs, the learning rate is dropped by a factor of 10, with a minimum rate of 10^{-6}. The training procedure is stopped when the validation loss has not improved for 15 epochs.

An example of the training is shown in Fig. 4.3 for images with no noise. The left panel shows the cross entropy loss as a function of the training epoch, where the blue and orange lines denote the training and validation sets, respectively. In addition to tracking the loss during training, we also compute the per-pixel accuracy of the validation data. We define the per-pixel accuracy as

$$\text{Pixel accuracy for class } k = \frac{\text{Number of pixels correctly predicted as class } k}{\text{Total number of truth-level class } k \text{ pixels}} \,. \tag{4.8}$$

For this, we define the a *correct* pixel assignment when the class with the largest prob-

ability $\left(\max_k \widehat{p}_k\right)$ for a given pixel matches the true assignment. There are 11 possible class assignments, so if the network is unsure of a given pixel's identity, all of the predicted probabilities could be around $1/11 \sim 9\%$. For now we still define a correct pixel assignment if the class with the largest probability matches the true class, even if the probability is low. We will examine setting a threshold on this assignment in a later section. We emphasize that the pixel accuracy for the subhalo classes requires getting both the location (the pixel) and the mass of the subhalo correct.

The pixel accuracy for each class is shown as a function of the training epoch for the validation images with no noise in the right panel of Fig. 4.3. In the first epochs, the model quickly learns to distinguish the main smooth lens from the background. After this, the effects of the subhalos are recognized, starting with the heaviest, which have the largest effects on the image. The accuracy for the 10^6 $M_\odot$ subhalo pixels reaches 40% by the end of training. We note that at this stage we are only tagging pixels as belong to a subhalo (or not), but have not discussed the detection of a subhalo as a whole. The eventual goal, however, is to build a catalog of subhalos with their positions and masses. This is done later in Sec. 4.4.

With the training procedure defined, we now move on to how noise in the images affects the training.

4.3 The effects of noise

The last section showcased our training procedure and showed examples from the noiseless images. In general, we find that applying a network that was trained with less noise to images with more noise results in very poor performance. This is not surprising, but it implies that we need to train on images with noise if we want the network to detect subhalos in noisy images. To do this, we generate 9×10^4 training images and 10^4 validation images for different fixed levels of noise, as discussed in Sec. 4.1. The noise is Gaussian with zero mean and standard deviation of a factor of $(10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2})$ of the mean of the pixel brightness for the pixels that are at least as bright as 10% of

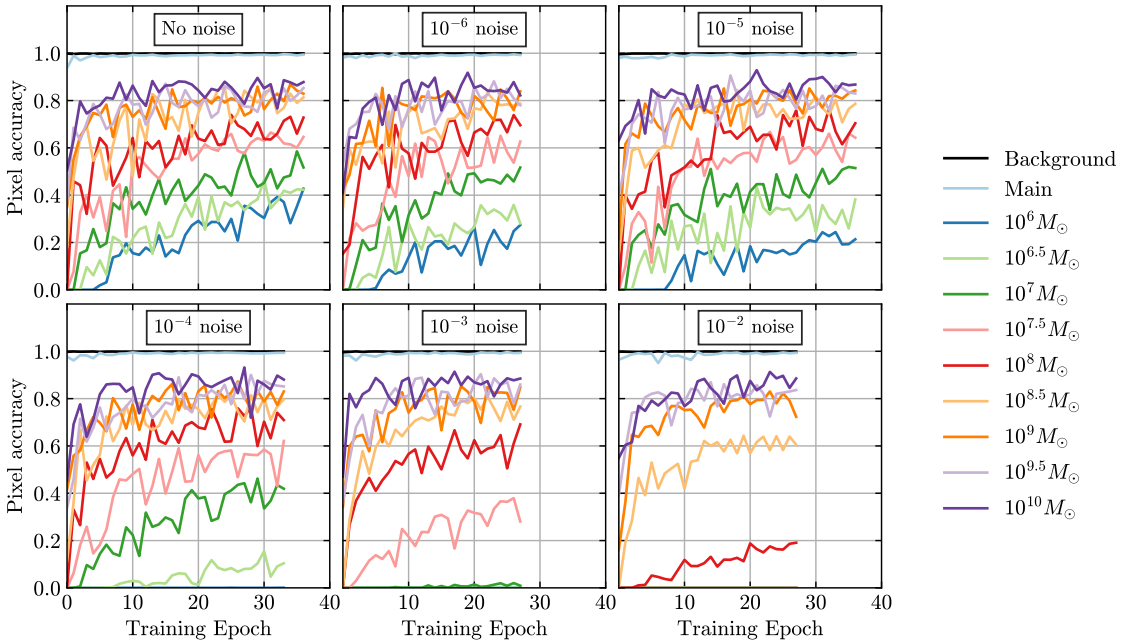

Figure 4.4: Accuracy for the different pixel classes as a function of the training epoch. The label in each panel denotes the amount of noise included in the images. The network performance on the heavy subhalos is unaffected by noise while it loses sensitivity to the lower mass subhalos with increasing noise. Networks trained on less noise and applied to more noise perform very bad, even for the main lens and heavy subhalos. Conversely, training on more noise and applying to cleaner images performs according to the noise level the network was trained on.

the maximum pixel. The smallest levels of noise are unrealistic for currently observed strong lensing systems, although 1% results in similar signal-to-noise ratios as HST (see App. B.4). The noise levels are chosen to illustrate the correlation between the network's ability to identify pixels from a subhalo and the mean fractional change in intensity due to subhalos of a given mass. This helps us to understand what the network is doing and what causes it to fail.

Fig. 4.4 shows the per-pixel accuracy as a function of the training epoch for each of the classes. The amount of noise in each panel is indicated by the label. For reference, the upper-left panel shows the results with no noise, which were also shown in Fig. 4.3. The upper-middle and upper-right panels have very small amounts of noise, 10^{-6} and 10^{-5}, respectively. Despite these small levels of noise, the pixel accuracy for the 10^{6} $M_{\odot}$ and the $10^{6.5}$ $M_{\odot}$ classes starts to decrease. The lower-left panel, with a noise at the level of 10^{-4}, has enough noise that the accuracy for the 10^{6} $M_{\odot}$ subhalos is almost zero. Similarly, the accuracy for the pixels of the $10^{6.5}$ $M_{\odot}$ class is very low. The pixel accuracy for the heavier subhalos is not significantly changed.

As the noise is further increased to 10^{-3} (in the bottom-middle panel), the network loses sensitivity to subhalos with mass below 10^7 $M_\odot$. However, we again see that the accuracy for the subhalos with $m > 10^8$ $M_\odot$ is not affected. In the final panel (lower-right), the noise level is 1%. With this much noise, the accuracy for the pixels coming from 10^8 $M_\odot$ subhalos is significantly decreased (by 70%) and the accuracy of the $10^{8.5}$ $M_\odot$ class pixels is decreased by around 25%. Again, the pixel accuracy for the heaviest subhalos do not change.

The noise level at which the network loses sensitivity to subhalos of a certain mass can be explained by the size of the perturbations caused by the subhalo. To demonstrate this, we examine 10^3 images before and after adding in a subhalo. As we have done so far, the subhalo is placed near the Einstein radius (defined by the pixels which are at least 20% as bright as the brightest pixel). We then compute the relative change in the observed image brightness for the pixels within a circle with a radius of 2 pixels centered on the location of the subhalo.[1] We plot the distribution for the mean relative change in each circle for each subhalo mass in Fig. 4.5. The bands mark 1 and 2 standard deviations.

Detecting a subhalo in the 10^6 $M_\odot$ bin requires noticing changes in intensity in clusters of pixels of order 0.01%. This highlights the enormous potential of the U-Net architecture. Without noise, it is able to correctly identify around 40% of the pixels for the 10^6 $M_\odot$ bin which cause tiny changes in the image. When the images contain noise at this same 10^{-4} level, it washes away the effects of these subhalos, as shown in the lower-left panel of Fig. 4.4. Similarly, Fig. 4.5 shows that subhalos in the 10^7 $M_\odot$ and 10^8 $M_\odot$ bins cause changes of order 0.1% and 1%, respectively. Fig. 4.4 likewise shows that the network loses sensitivity to these subhalos with these corresponding levels of noise. Throughout the rest of this work, we will compare the model with no noise or 1% noise. This amount of noise is realistic for strong lens images with apparent magnitudes brighter than 20.

Before moving on, we raise two points to the attention of the reader. The first is that the network needs to be trained on similar levels of noise to the data it will be used on. For instance, applying the network trained without noise to the noisy images

[1]This corresponds to the pixels which are labeled as the subhalo in the target data of the training set.

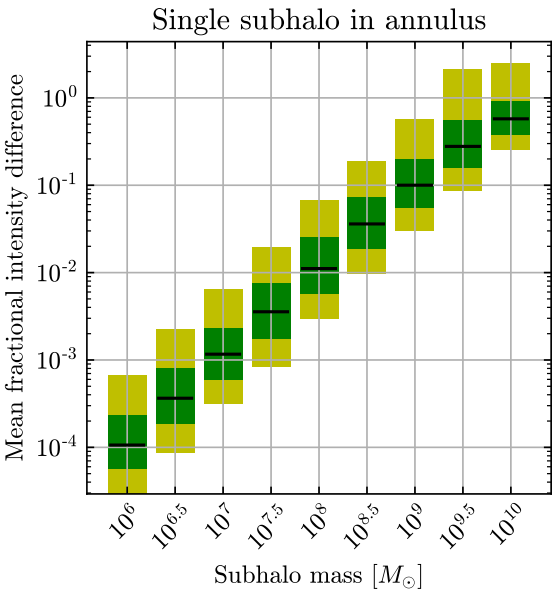

Figure 4.5: Detecting subhalos require methods that are sensitive to small changes in pixel intensities. Images are generate without substructure. A single subhalo is then added, and the mean fractional change in the circle with a radius of 2 pixels around the subhalo are recorded. The bars show the mean from repeating this process 10^3 times. The green and yellow bands show the range of the 1 and 2 standard deviations.

leads to very poor performance. However, one cannot just be conservative and train on images with too much noise. We tested this explicitly by taking the network trained on images with 10^{-2} noise and applying it to images with no noise. The performance of the subhalos with masses greater than 10^8 $M_{\odot}$ was very similar to that shown in the last panel of Fig. 4.4 for the images with 10^{-2} noise. Even though the images now have no noise, the network is not able to detect the low-mass subhalos. Training on images with varying levels of noise could alleviate this concern.

The second point we want to raise is that it is most likely possible to achieve sensitivity to higher levels of noise that we show here. Unlike random noise, the effects from perturbations to the lens are correlated across several pixels. Changing the size of the convolutional filters or the number of filters could help detect these small correlations on top of the noise. Additionally, we did not implement any class weights into the loss function. It is possible to make the network place greater emphasis on identifying certain classes more than others. Our current setup has orders of magnitude more background pixels than subhalo pixels in the target data. Forcing the network to place greater emphasis on learning the more subtle subhalos, especially in the noisy images, could help.

As a proof-of-principle, designing an optimal network for higher levels of noise is beyond the scope of this work. To this point, we have only shown results from the training and validation images. We now move onto to completely independent data to characterize the network's performance on images that it has not seen during either training or validation.

4.4 Characterizing the network performance

After the network has been trained, we apply it to a series of images that the network has not seen during training or validation to evaluate its out-of-sample performance. We do this both for a network trained on images with no noise and a network trained on images with 1% noise. First, we show an example output of the network, which helps to visualize the different channel probabilities and see common ways for the network to mislabel pixels. We then compare the true target pixels to the predicted pixels to quantify the amount of correct and mislabeled pixels. Finally, we run the network on images without substructure to determine the rate at which the network will claim to detect subhalos when they are not there (false positives).

4.4.1 Example output

In Fig. 4.6, we show a detailed example of the network without noise applied to a noiseless image with a very light subhalo, in the 10^6 $M_\odot$ bin. The upper row displays the observed image, which serves as input to the network (left), the truth-level target labels (middle), and prediction (right). We have assigned each pixel in the prediction image according to whichever class had the largest probability for that pixel. The lower three rows display the individual class probabilities, with red representing low probabilities (the network is certain that the pixel does not belong to that class) and blue is high probabilities (the network thinks this pixel belongs to the class). We chose a color map such that probabilities near 50% are white, showing that the network is unsure of those pixels.

In examining the target and prediction images of the top row, we see that the network successfully identified the subhalo in the image, despite its extremely low mass and its

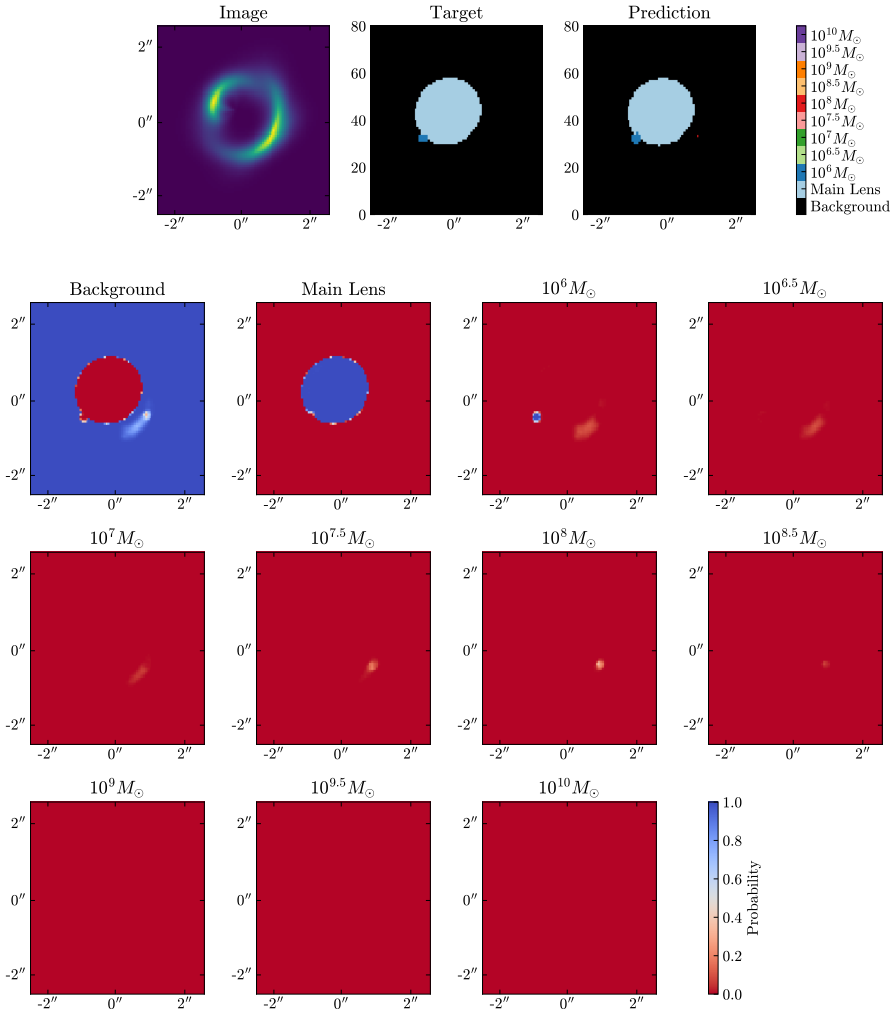

Figure 4.6: The top row shows (from left to right) an image with no noise which is input to the network, the target labels for each pixel, and the prediction from the network. The lower three rows show the probability assigned to each class for each pixel. The network prediction for each pixel is determined by the class with the maximum probability for that pixel. The true subhalo has a mass in the 10^6 $M_\odot$ bin and is detected in the prediction. A few pixels around the edge of the subhalo or main lens get misclassified. A single pixel is incorrectly predicted to belong to the 10^8 $M_\odot$ class, but this prediction can be corrected by imposing a probability threshold.

overlap with the main lens. The mass of the subhalo is indicated by the pixel colors (dark blue) representing a subhalo in the 10^6 $M_\odot$ bin. In addition, we see that the predicted size is also very close to the target, although the shape is not exactly the same. For instance, the pixel on the top of the subhalo in the predicted image extends further into the main lens than it does in the true label. When examining the per-pixel accuracy in the next subsection, this will show up as a pixel which should be predicted to be part of the main lens, but is predicted to be in a subhalo class. Similarly, there are a couple of pixels along the bottom of the predicted subhalo that are not part of the target. In the discussion of the next section, these will show up as pixels which should be predicted as background, but are instead predicted as belonging to a subhalo class. In both of these cases, although the pixels were misidentified, the subhalo was still found and they do not represent the network introducing spurious substructure. If we examine the probability map for the 10^6 $M_\odot$ class (2nd row, 3rd column), we see that the pixels around the edge of the subhalo are white, indicating that the network was not confident in these assignments. In general, we find that sometimes the subhalos are predicted to be a few pixels too large and sometimes a few pixels too small, but just as in the preceding discussion, if a single pixel that should be assigned to a subhalo class is misclassified into the main lens or background, this does not indicate that the subhalo was not found.

Similar features can be seen along the edge of the main lens. The shape and size of the lens is very similar between the target image and the predicted image. However, there are a few instances along the edge where the pixel assignments are incorrect. In the discussion below, these will be pixels that should be classified as background (main lens) but are misclassified as main lens (background). Looking at the probability maps in the 2nd row for the background and main lens, we can see that the regions where the network had classification errors have class probabilities near 50%.

This example contains another interesting feature; there is a single pixel that was predicted to belong to a 10^8 $M_\odot$ subhalo. It is challenging to find the pixel on the predicted image (it is on the lower-right side of the main lens) but the probability maps shed light on it. The background probability map has a swath of pixels beneath the

main lens that the network is unsure about. We see that these pixels are assigned some probability to belong to all of the subhalos classes with $m \leq 10^{8.5}$ $M_\odot$, although for nearly all of them the background has the highest probability. However, one of the pixels has a probability in the 10^8 $M_\odot$ bin that is slightly larger, so it gets assigned to that class despite the probability being low. Often these spurious pixels can be removed by imposing a probability threshold. For instance, the probability for these pixels is spread among many classes, and in no class is the probability at least 50%. If we default all pixels to the background class unless a subhalo probability is greater than 50%, the wrong prediction will be removed. We note that most images do not contain random pixels like this one.

While instituting a probability threshold can remove spurious pixels, it also leads to a reduction in the overall accuracy for the true subhalo pixels, which is why we did not implement a probability threshold for our fiducial results. Recall that in Eq. (4.8) the accuracy was defined as the correctly assigned pixels for a class divided by the total true pixels for that class. Some pixels along the edge of subhalos will now get misclassified as background, reducing the accuracy.

4.4.2 Testing on a single subhalo

Now that the output of the network is better understood, we move on to quantify the network's predictions. The purpose of this is to determine what the network is predicting for the pixels of any class: *are most of them correct? And if they are predicted wrongly, what class are they assigned to?* To assess this, we generate a new set of 10^4 images with the same amount of images with and without substructure as in the validation set: 10^3 have no substructure and the remaining images have a single subhalo in each and are evenly split between the nine mass bins, meaning there are 10^3 images in each.

In Fig. 4.7, we examine the per-pixel predictions on the images in the test set.[2] The title in each panel states the true label of the pixels, while the x-axis denotes the predicted class. The solid (dashed) lines represent a prediction with no probability (50%

[2]In Appendix B.2 we show the full confusion matrix. Note that each panel in Fig. 4.7 is a cross-section of the confusion matrix.

probability) threshold. The blue line corresponds to images without noise while the orange line corresponds to images with 1% Gaussian noise. Each panel is normalized such that the sum of all the classes is unity.

The first two panels (at the top, from left to right) show the pixels that at truth-level correspond to the background or main lens classes, respectively. The network often misclassifies a few pixels around the edge of the main lens (e.g. Fig. 4.6 in the previous subsection). This can either be the main lens being reconstructed as slightly too large or too small, or getting the shape slightly off, although it is typically only a handful of pixels. Because there are so many background pixels, this corresponds to 0.1% of the background pixels getting misclassified as the main lens over the entire test set; this is the most common type of misclassification for the background pixels. Similarly, these errors around the edges of the main lens lead to $\lesssim 1\%$ of the pixels which should be predicted as the main lens getting misclassified as the background class.

The rest of the class assignments in these two panels are roughly uniform for the images without noise. In the images with noise, the same effect happens for subhalos with masses $> 10^8$ $M_\odot$. The explanation is similar to that of the main lens itself: when the network locates a subhalo, some of the pixels around the edge can get mislabeled. It is rare for the network to get the exact shape of the subhalo correct. This was also shown in Fig. 4.6. We emphasize that most of the pixels that are supposed to be background or the main lens, but are predicted to be a subhalo, do not represent additional false subhalos but are rather edge effects of this type. More evidence of this is given in the next subsection. Recall that, with 1% noise, the network does not detect subhalos with $m < 10^{7.5}$ $M_\odot$. This means that there are no edge effects associated to these classes (and consequently no pixels getting misclassified as a subhalo with $m < 10^{7.5}$ $M_\odot$), as can be seen by the cutoff in the orange histograms in both of these panels.

In these first two panels, we also see the effect of including a probability threshold when making a prediction. Specifically for the images without noise, the background and main lens pixels that get predicted into the $10^{6.5}$ $M_\odot$-$10^{7.5}$ $M_\odot$ classes are reduced by a factor of 1.5-2. This shows that network is uncertain of some of these pixels that are

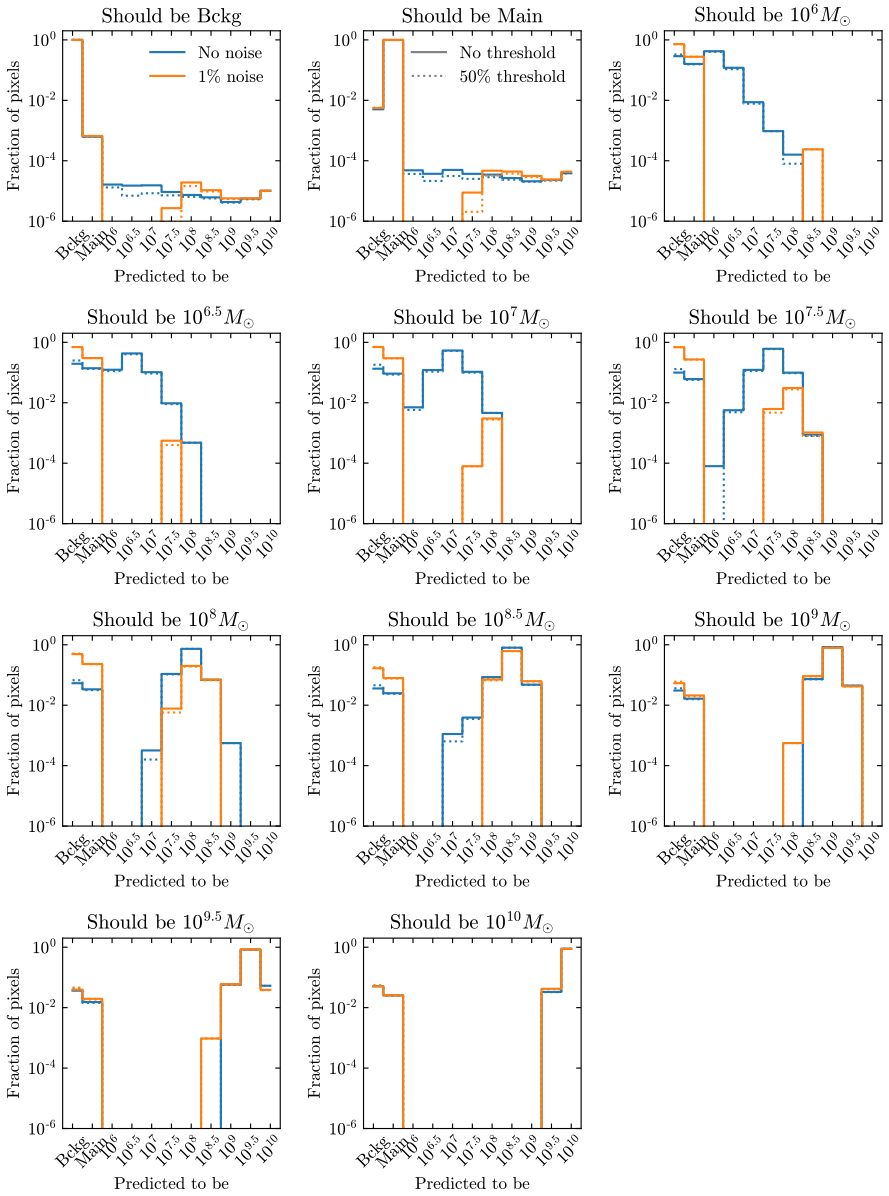

Figure 4.7: Each panel corresponds to pixels which truth-level are of the indicated class. The x-axis denotes the class that the pixels are predicted to be. The solid and dashed lines represent the predictions with no probability threshold and a 50% probability threshold, respectively. The blue lines denote the network predictions on images without noise and the orange is for images with 1% Gaussian noise. Each line in each panel is normalized to unity.

misclassified. The probability threshold does not have as large of an effect on the heavier subhalos because their effects are larger and allow the network to be more certain in its predictions pertaining to these mass classes.

The edge effects can also be seen in the other panels. For example, in the bottom-right panel, most of the pixels are correctly identified as belonging to a very massive subhalo. However, some are incorrectly marked as the main lens or background. These are from pixels around the edge of the subhalo. In addition to edge effects, the network sometimes gets the mass wrong by one mass bin. There is little difference in the results with and without noise, or with and without the probability threshold, for the heavy subhalos.

For subhalos with mass larger than $\gtrsim 10^{7.5}$ $M_\odot$ for no noise, and $10^{8.5}$ $M_\odot$ with noise, the most typical way for the prediction to be wrong is to predict the wrong mass bin. Progressing towards the panels with the lighter subhalos, we notice that the fraction of pixels incorrectly labelled as background and/or main lens increases. This makes sense as the magnitude of the deflection angles decreases with decreasing halo mass, so their effects are easier for the network to miss.

Fig. B.1 in App. B.1 shows the distribution of true labels for pixels predicted to belong to a given class, from which we see that in fact the predicted class is very likely to be correct. Thus, if the network predicts a group of pixels to have the same class, it is very likely that a subhalo is present there. The most common type of error for a pixel predicted to belong to a subhalo is that it should belong to a different (adjacent) mass bin.

Up to this point, we have only been discussing the per-pixel predictions. This makes sense from a machine learning perspective, but it does not necessarily address the physics goal of detecting subhalos themselves. This brings to light one potential challenge of using image segmentation to detect substructure in images of strong lensing: how does one go from pixels to subhalos? We have found that it is possible to either add extra pixels or miss pixels from a subhalo, especially around the edge. However the per-pixel accuracies and the example shown in Fig. 4.6 suggest that, on average, the size should be correct. This means that we can get a subhalo count by summing the number of pixels predicted

in each subhalo class and dividing by 4π pixels/subhalo (because we defined the target pixels such that the subhalos are comprised of a circle with a radius of 2 pixels).

We can then examine the subhalo detection accuracy over the test set. To define this accuracy, we take the pixels that at truth-level belong to a given subhalo class and count the number of these pixels assigned to each of the eleven possible classes. We then label the subhalo as belonging to the class with the largest count. Using this notion of a subhalo detection, Table 4.1 shows how the 10^3 subhalos of each mass bin were reconstructed. These numbers are for the predictions without a probability threshold (see App. B.3 for an analogous table including the probability threshold). The results for the network on images without noise and with 1% noise are indicated by the number without and with brackets, respectively. For the noiseless images, the network finds 616 (480 + 127 + 8 + 1) of the 1000 subhalos with a mass of 10^6 $M_\odot$ (62%), of which 78% are in the correct mass bin. The subhalo detection accuracy, as well as the probability of getting the mass correct, increase for heavier subhalo masses. Subhalos with $m \geq 10^8$ $M_\odot$ are detected more than 97% of the time in images with no noise.

For images with 1% noise, nearly all of the subhalos lighter than $m \leq 10^{7.5}$ $M_\odot$ are missed. Even though the 1% noise is on the same level as the effects on the brightness due to a 10^8 $M_\odot$ subhalo, 39% of them are detected. Once the subhalos are heavy enough to produce effects larger than the noise level, the accuracy is similar to the noiseless images. We find that subhalos with $m \geq 10^{8.5}$ $M_\odot$ ($m \geq 10^9$ $M_\odot$) are detected more than 86% (98%) of the time with 1% noise.

Table 4.1: The number of subhalos predicted to each class. For each row, there are 1000 images with a single subhalo in the indicated mass bin. The results for the network on images without noise and with noise are indicated by the number without and with brackets.

True Class	Not Detected	10^6 $M_\odot$	$10^{6.5}$ $M_\odot$	10^7 $M_\odot$	$10^{7.5}$ $M_\odot$	10^8 $M_\odot$	$10^{8.5}$ $M_\odot$	10^9 $M_\odot$	$10^{9.5}$ $M_\odot$	10^{10} $M_\odot$
10^6 $M_\odot$	384 [1000]	480 [0]	127 [0]	8 [0]	1 [0]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]
$10^{6.5}$ $M_\odot$	268 [1000]	141 [0]	475 [0]	105 [0]	11 [0]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]
10^7 $M_\odot$	157 [996]	7 [0]	138 [0]	585 [0]	108 [1]	5 [3]	0 [0]	0 [0]	0 [0]	0 [0]
$10^{7.5}$ $M_\odot$	90 [958]	0 [0]	7 [0]	138 [0]	659 [7]	105 [33]	1 [2]	0 [0]	0 [0]	0 [0]
10^8 $M_\odot$	29 [608]	0 [0]	0 [0]	0 [0]	117 [5]	780 [307]	73 [80]	1 [0]	0 [0]	0 [0]
$10^{8.5}$ $M_\odot$	16 [134]	0 [0]	0 [0]	1 [0]	4 [0]	91 [75]	841 [717]	47 [74]	0 [0]	0 [0]
10^9 $M_\odot$	1 [12]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]	78 [81]	874 [842]	47 [65]	0 [0]
$10^{9.5}$ $M_\odot$	1 [4]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]	62 [52]	879 [893]	58 [51]
10^{10} $M_\odot$	1 [1]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]	37 [41]	962 [958]

4.4.3 Null tests

The claim of the last section is that the pixels that should have been predicted as background, but were instead classified as a subhalo pixel, were dominantly from edge effects. Fig. 4.6 provided an example of one such case. In this section, we provide evidence of this claim systematically by determining the rate at which the network finds spurious substructure in images where there is only a smooth lens.

We generates a set of 10^3 images that contain only the source light and the smooth main lens, and another set that also contained 1% Gaussian noise. The trained networks were then applied to each of the images. We did not use a probability threshold for the predictions here to allow low-probability substructure to appear in the count which provides a more conservative estimate. The resulting counts are shown in Fig. 4.8 for the noiseless images in the left panel and the images with 1% noise in the right panel. The error bars were estimated using the square root of the estimated subhalo count, $\sqrt{N_{\text{subhalo}}}$.

While the images do not contain subhalos, the network incorrectly predicts a few subhalos with low masses. The rate for this to happen is less than 10 subhalo per 10^3 images for each class separately, for the images without noise. The total rate of false subhalos is around 20 subhalos per 10^3 images, coming entirely from subhalos with

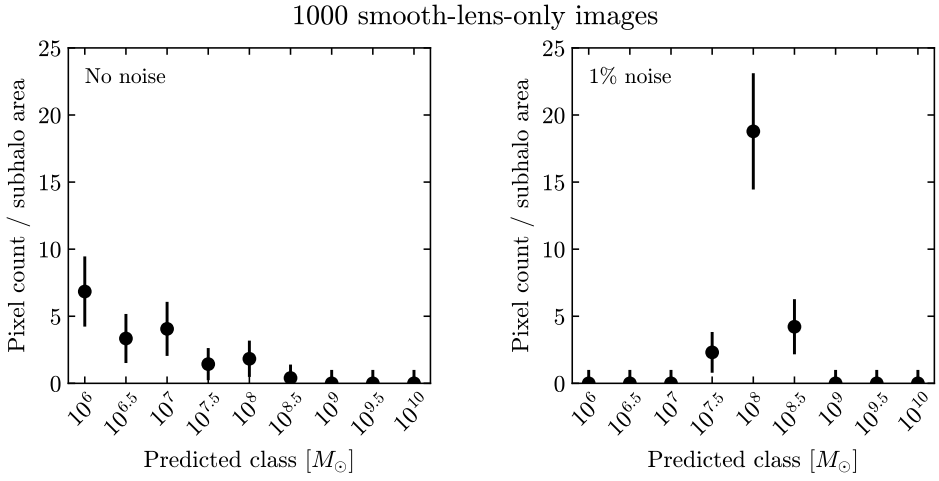

Figure 4.8: Approximate predicted subhalo counts from 10^3 images that do not contain substructure in the lens. Uncertainties are estimated using $\sqrt{N_{\text{subhalo}}}$. The left and right panels correspond to noiseless images and images with 1% noise, respectively. For noiseless images, in each individual class the false-positive rate is less than 10 subhalos per 10^3 images, and the combined rate is around 20 subhalos per 10^3 images. The total false-positive rate in images with noise is around 25 subhalos per 10^3 images, most of which are in the 10^8 $M_\odot$ mass bin.

masses $m \leq 10^8$ $M_\odot$. This gives us strong confidence that there is substructure in the lens when the network predicts a subhalo.

When noise is included, the total number of spurious subhalos increases slightly, to 25 per 10^3 images. Most of these are in the 10^8 $M_\odot$ bin. This makes sense because subhalos with this mass cause changes in the pixel brightness on the order of 1%.

4.4.4 Testing on images with many subhalos

Up to this point, we have presented results from networks trained on images with at most a single subhalo near the brightest pixels. However, galaxies are expected to have a large population of subhalos distributed throughout the main dark matter halo. We therefore investigate whether, without knowing this true distribution of subhalos, our network can be used to identify a population of substructure (i.e. more than a single subhalo per image). This is an example of *domain adaptation*, in the ML jargon. We describe this process applied to our model in Appendix B.5.

To test this, we generate a set of images with up to 25 subhalos (with the actual number in each image drawn from a uniform distribution) with random masses and locations. Four of these images are shown in the left panels of Fig. 4.9. The middle panels show the true pixel labels and the right panels show the network output. We denote the number of subhalos in each of the target and predicted images. All of these images have no noise, but the results generalize to images with noise (losing efficiency for the lower masses, as explained above).

The network predictions for these images illustrate that light substructure far from the images is hard for the network to capture. Each row also illustrates different interesting properties. In the top row, on the lower-right side of the main lens, where there are two overlapping subhalos and the network only picks up on the heavier one of the two, which makes sense. The second row contains an example of a subhalo in which the network is unsure of the mass bin. The subhalo at the top of the main lens gets half of its pixels reconstructed in the 10^6 $M_\odot$ class while the other half are put in to the $10^{6.5}$ $M_\odot$ class, although the truth-level subhalo belongs to the 10^6 $M_\odot$ class. In the third row, five

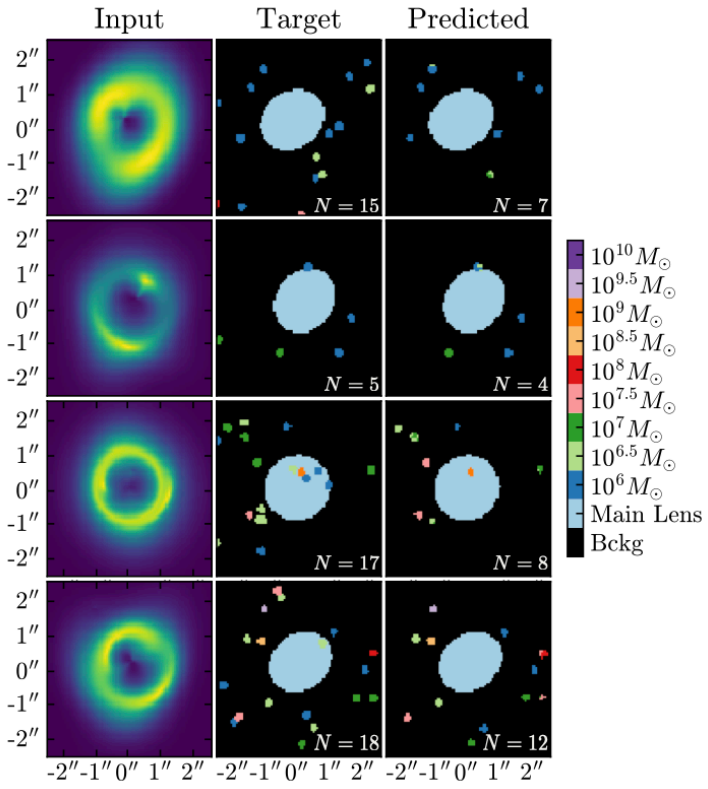

Figure 4.9: The network trained on noiseless images with a single subhalo is now applied to noiseless images with rich distributions of substructure. The network has low sensitivity to subhalos that are far from the bright pixels around he Einstein ring. Subhalos that are near each other can get reconstructed as a single subhalo with their combined mass. The number of true/predicted subhalos are denoted in the figure.

subhalos are inside of the Einstein radius. The heaviest of these is reconstructed, but the others are missed. The interior of the ring is also far from the light, making detection harder. Additionally, it seems like the network could have compensated for these subhalos by slightly increasing the main lens, but instead it does find the 10^9 $M_\odot$ subhalo, which is impressive. Finally, in the last row, we see examples of heavier subhalos, which are able to be detected closer to the edge of the image than the lighter subhalos.

From these examples, we see that the network trained on images with a single subhalo is able to detect a population of substructure. As the subhalos get further away from the bright pixels, the network loses sensitivity to them. The presence of subhalos close together is also a challenge for the network. When there is a hierarchy between their masses, the heavier subhalo will wash out the effects of the smaller one. This results in only a detection of the heavy subhalo. If nearby subhalos are similar in mass, they may get reconstructed as a single subhalo in the bin corresponding to the sum of their masses.

The variety of subhalo populations that were detected in Fig. 4.9 (both in terms of raw number of subhalos as well as their mass distribution) also highlights one of the strengths of our single-subhalo training methodology. We found empirically that when we instead trained a network on images with many subhalos drawn from a power-law mass function and then tested it on images with populations of subhalos drawn from a mass function with a different power-law index, the detection accuracy became biased. For instance, if the network is trained on images with many subhalos drawn from a steep mass function, the few images with heavy subhalos will almost certainly have lots of light substructure as well. When we applied the network to images generated with a shallower mass function, the false-positive rate was higher for images with heavy subhalos: when it saw a heavy subhalo, it expected more light subhalos than there were, thus introducing false substructure. By only having a single subhalo in the training images, our fiducial network never learns to make decisions based on population characteristics, and thus generalizes extremely well to images with very different subhalo populations.

We have now shown the our network detects subhalos near the Einstein ring with high accuracy, and is capable of generalizing to images with different subhalo populations. Additionally, it has a very low false positive subhalo rate of 2.5 subhalos per 100 images. In the next section we demonstrate how such a network could be used to extract the subhalo mass function from an ensemble of strong lens images.

4.5 Determination of the subhalo mass function

The population of subhalos under a CDM scenario is found to be well described by a power-law of the form

$$\frac{dN}{dm} \propto m^{\beta}, \tag{4.9}$$

with $\beta = -1.9$ [143]. However, models beyond CDM can affect this subhalo mass function. In previous sections, we have described how our U-Net model is able to accurately detect subhalos in simulated images. In this section, we show that we can determine the subhalo mass function using the network outputs from many images.

We first generate a mock catalog of images that can have more than a single subhalo. In each image, we draw the number of subhalos from a uniform distribution between 0 and 25, and place them uniformly throughout the image. The masses of the subhalos are drawn from a power-law with fixed $\beta = -1.9$. As the network with noise is not able to detect substructure with $m < 10^8$ $M_\odot$, and we want to compare the performance of the networks trained on noiseless and 1% noise images, we do not generate subhalos with masses that would end up in the $m \leq 10^{7.5}$ $M_\odot$ bins. Our goal will be to infer the value of β from this catalog.

If we were able to perfectly reconstruct every subhalo, we could just fit the counts-per-bin to the functional form in Eq. (4.9). However, the accuracy is worse for lower mass subhalos, effectively changing the shape of the extracted mass function. An example of this is shown in Fig. 4.10. Here we simulated 10^5 images, using the same prescription as in our mock catalog, for β =-1.7, -1.8, -1.9, and -2.0. We then applied the network to each image and counted how many pixels were assigned to each class. This pixel count was converted to an approximate subhalo count by dividing by 4π. The left panel shows the average number of detected subhalos for the images with no noise, and the right panel shows the average counts for the images with 1% Gaussian noise. While the curves in the noiseless images are nearly straight (in log-log space), there is still some curvature. This is exaggerated even more in the images with noise, where in fact the subhalo counts plateau below $\sim 10^8$ $M_\odot$.

Rather than trying to fit the data with Eq. (4.9), we instead infer the most probable power-law index to have generated the observed subhalo counts extracted from a set of images with our U-Net model. To do so, we build a likelihood function which is the product of Gaussian likelihoods for each mass bin. Namely, the likelihood is given by

$$L(\beta) \equiv \prod_{i \in \text{mass bins}} G\Big(o_i \Big| \mu_i(\beta), \sigma_i(\beta)\Big), \tag{4.10}$$

where o_i is the observed number of detected subhalos in bin i, $\mu_i(\beta)$ is the expected value, and $\sigma_i(\beta)$ is the estimated standard deviation given β. To use the likelihood to infer β,

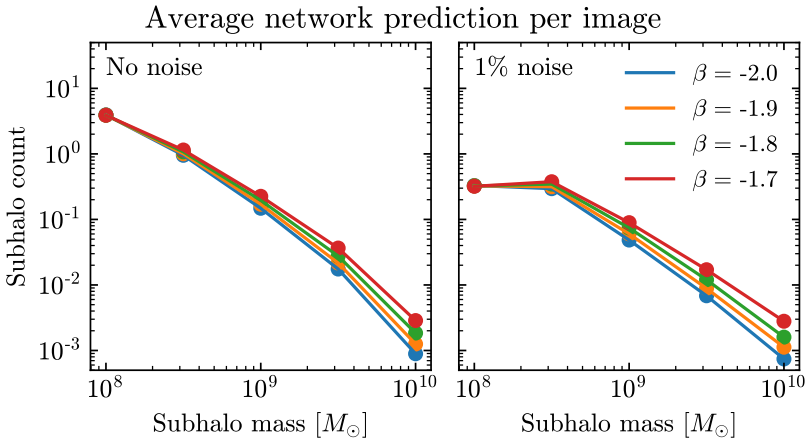

Figure 4.10: The left and right panels display the average number of detected subhalos per class for 10^5 images with an without noise, respectively. The images can have between 0 and 25 truth-level subhalo uniformly distributed throughout the image. The masses of the subhalos are drawn from a power-law with an index denoted by β. The network has less sensitivity to lower mass subhalos, resulting in the curved shape. To infer the power-law index of an independent dataset, we find which value of β yields a curve closest to that observed in the data.

we need to derive $\mu_i(\beta)$ and $\sigma_i(\beta)$. This is done in the next subsection. We emphasize that these derivations will include details that are specific to our mock catalog of images, and may not generalize to the real world, but we will discuss how to generalize to other catalogs.

4.5.1 Expectation and variance

Here, we derive the expectation and variance used in the likelihood function given in Eq. (4.10). First, we define the efficiency to tag a subhalo as

$$\varepsilon_i = \frac{\text{Number of subhalos predicted in class } i}{\text{Number of true subhalos in class } i}. \tag{4.11}$$

Because we train our network on images with a single subhalo, the network does not know about population level statistics, which is why the efficiency to zeroth order is not a function of β. If an image has N real subhalos, we define φ as the true fraction of subhalos in each mass bin, N_i. Note that φ is necessarily a function of β and is given by

$$N_i(\beta) = \varphi_i(\beta)\ N, \tag{4.12}$$

where N_i is the true number of subhalos in the ith bin. Thus, the number of subhalos expected to be predicted in a given class for an image with N total true subhalos can be written as

$$\langle n_i(\beta)|N\rangle = \varepsilon_i \; \varphi_i(\beta) \; N. \tag{4.13}$$

where n_i is the number of predicted subhalos in bin i.

From running a large number of images, we obtain a good estimate for

$$\begin{aligned} \epsilon_i(\beta) &= \varepsilon_i \; \varphi_i(\beta) \\ &= \langle n_i(\beta)|N\rangle/N \; , \end{aligned} \tag{4.14}$$

where $\epsilon_i(\beta)$ is essentially the rate at which subhalos are detected in mass bin i given the value of β. For example, in Fig. 4.11 we applied the network to count the number of predicted subhalos in 10^5 images for various fixed values of β with an average of 12.5 subhalos per image. Rather than showing the results as a function of the subhalo mass bin (as done in Fig. 4.10), we show the average count per image as a function of β, with each mass bin shown in a different panel. We find that the data can be fit well with an exponential function of the form

$$\epsilon_i(\beta) = a_i + b_i \, e^{c_i \beta} \; . \tag{4.15}$$

Given the rate (ϵ_i) and the true number of subhalos (N) in a single image, the number of predicted pixels in a given class should be be Poisson-distributed about the expectation. Thus, the probability (p) to predict n_i subhalos is given by

$$p\big(n_i(\beta)\big) = P\big(n_i|\epsilon_i(\beta) \; N\big) = \frac{\big(\epsilon_i(\beta) \; N\big)^{n_i} e^{-\epsilon_i(\beta) \; N}}{n_i!} \; . \tag{4.16}$$

The expected number of detected subhalos in a single image with N true subhalos from Eq. (4.13) can then be rewritten in terms of the weighted sum over the individual prob-

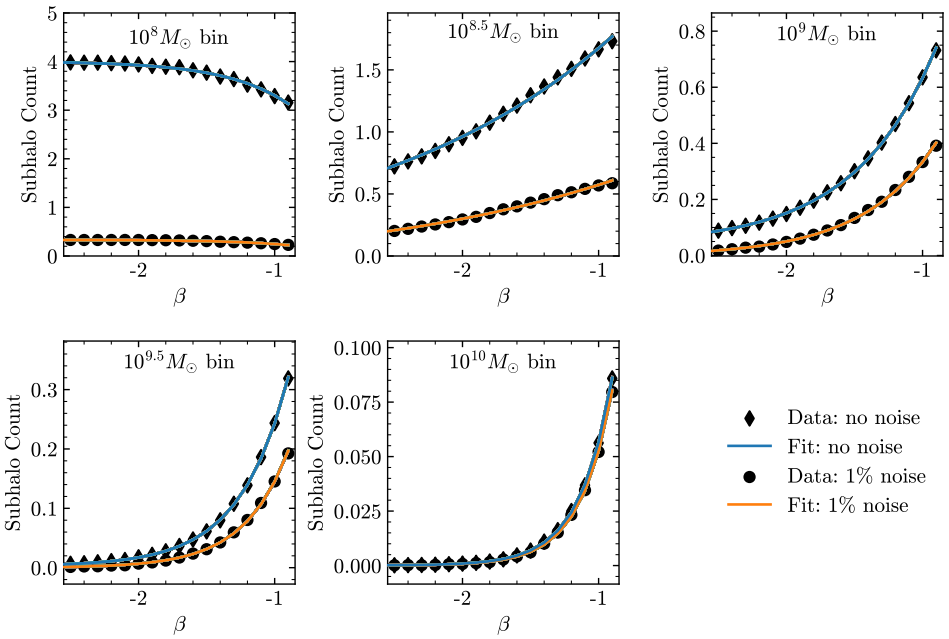

Figure 4.11: Data points show the average number of subhalos detected by our network per image over 10^5 sample images with and without noise. The lines show exponential fits to the data as described in the text. These are used to define the expected subhalo counts and the standard deviations used in the likelihood function. The downward trend for increasing β in the 10^8 $M_\odot$ panel is because of the fixed range for the number of subhalos.

abilities

$$\langle n_i(\beta)|N\rangle = \sum_{n_i} n_i \; p\big(n_i(\beta)\big) \; . \tag{4.17}$$

We now build in the assumptions about our mock data. In each image, we have placed between 0 and $N_{\rm max} = 25$ true subhalos, with the number drawn from a uniform distribution. The expected number of detected subhalos in class i for any image is then the average over the $N_{\rm max} + 1$ possible values (including 0 subhalos),

$$\langle n_i(\beta)\rangle = \sum_{N=0}^{N_{\rm max}} \frac{1}{N_{\rm max}+1} \sum_{n_i} n_i \; p\big(n_i(\beta)\big) \; . \tag{4.18}$$

In an analysis of real data, the actual distribution of subhalos would need to be obtained through detailed N-body simulations, which is beyond the scope of this work.

In any given image, the number of subhalos in a bin can vary widely. To compute the variance, we calculate $\langle n_i^2\rangle$ as a function of β, given by

$$\langle n_i^2(\beta)\rangle = \sum_{N=0}^{N_{\rm max}} \frac{1}{N_{\rm max}+1} \sum_{n_i=0}^{N} n_i^2 \; p\big(n_i(\beta)\big) \; . \tag{4.19}$$

Then the variance is

$$\sigma^2_{n_i}(\beta) = \langle n_i^2(\beta)\rangle - \langle n_i(\beta)\rangle^2 \ . \tag{4.20}$$

The range of possible true subhalos makes this variance large.

Up to this point, we have examined the expectation and variance for the number of detected subhalos per bin in a single image. However, to determine the subhalo mass function we will analyze many images together. The expected number of pixels and the variance for N_{images} independent images is given by

$$\langle n_i(\beta)|N_{\text{images}}\rangle = N_{\text{images}}\langle n_i(\beta)\rangle; \tag{4.21}$$

$$\sigma^2_{n_i(\beta)|N_{\text{images}}} = N_{\text{images}}\sigma^2_{n_i}(\beta). \tag{4.22}$$

The value of the mean and standard deviations for each bin included in the likelihood function are then given by

$$\mu_i(\beta) = N_{\text{images}}\langle n_i(\beta)\rangle \tag{4.23}$$

and

$$\sigma_i(\beta) = \sqrt{N_{\text{images}}\sigma^2_{n_i}(\beta)} \ . \tag{4.24}$$

4.5.2 Inferring the power-law

At this stage, given a subhalo mass function power-law index and some number of images, we can compute the expected number of detected subhalos and the standard deviation for each bin. When we apply our network to a set of images, we can then determine which value of β yields expectations closest to the observed detections using the likelihood function defined in Eq. (4.10). Recalling that the Gaussian expectation and variance are a function of β and that the observed data does not change, the likelihood is then a function only of β,

$$L = L(\beta). \tag{4.25}$$

To test a hypothesized value of β, the likelihood ratio is used,

$$\ell(\beta) = \frac{L\big(\beta\big)}{L\big(\widehat{\beta}\,\big)}, \tag{4.26}$$

where $\widehat{\beta}$ is the value of the slope which maximizes the likelihood. To find $\widehat{\beta}$, we use the Nelder-Mead [165] algorithm implemented in `SciPy` [166]. We then compute the test statistic t_β, defined as

$$t_\beta = -2 \log \ell\,(\beta)\,, \tag{4.27}$$

with which we can determine confidence intervals.

Fig. 4.12 does this for images containing subhalo populations drawn from a subhalo mass function with $\beta_{\rm true} = -1.9$, for different numbers of images in each panel. The best fit is the value of β that minimizes t_β, while the $1\sigma(2\sigma)$ uncertainty is given by the range between $t_\beta = 1(4)$. The orange and blue lines denote the results from images with and without noise, respectively. The last panel shows the spread of the confidence intervals as a function of the number of images used for the fit.

The first panel uses only 10 images, and the resulting inferred values have large uncertainties. This is primarily due to the large variance in the number of subhalos per image, which is incorporated into the uncertainty. While the central value for the noiseless images is not close to the true value of $\beta_{\rm true} = -1.9$, it is within the 1σ confidence interval. Increasing the number to 50 images, the uncertainty drops by a factor of 1/3. The best fit values also start to get closer to the true value. The trend continues as more images are added, with the best fit converging to the true value and the uncertainty decreasing.

We again note that noise affects our ability to tag low-mass substructure. As shown in Fig. 4.10, the number of detected subhalos does not resemble a power-law for the noisy images. Despite this, we find that the uncertainty on the inferred value of β is only around 20% worse. With 500 images, the uncertainty on the power-law index is less than 0.1. Overall, we find that, when measured over five mass bins from 10^8 $\mathrm{M}_\odot$ to 10^{10} $\mathrm{M}_\odot$, the SMF slope is recovered with an error of 14.2 (16.3)% for 10 images, and this improves to 2.1 (2.6)% for 1000 images without (with 1%) noise

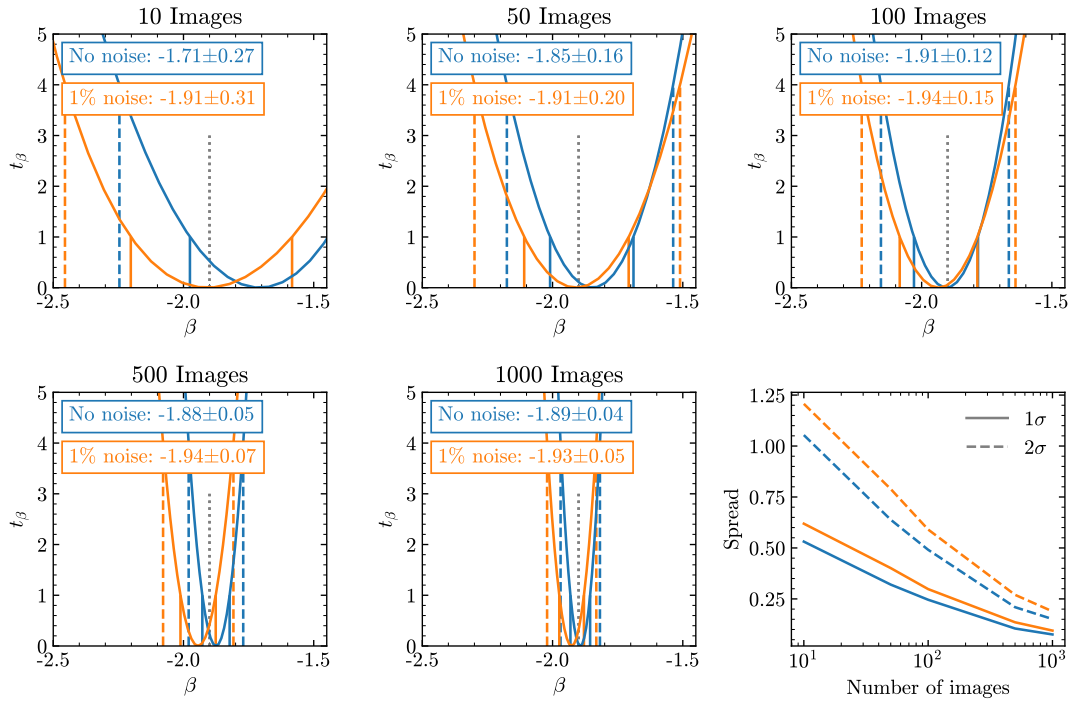

Figure 4.12: Result of scanning the test statistic t_β when scanning over the power-law index β. The test statistic is defined in Eq. (4.27). The value of β that minimizes t_β produces expected count most similar to that in the mock observed data. The true value of β in the mock data is -1.9 and is marked by the grey dotted line. The 1 and 2 σ confidence intervals are the values of β between $t_\beta = 1$ and 4, respectively. The last panel shows that the uncertainty decreases as the number of images increases. The blue and orange lines shown the results for image with no noise or 1% Gaussian noise, respectively. The last panel shows the spread (maximum - minimum) of the 1- and 2-σ confidence bands. The uncertainty on the best fit for the images with noise is about 20% larger than the noiseless images.

The results presented here strongly depend on our mock catalog. Specifically, for the images in our catalog, the number of subhalos in an image is drawn from a uniform distribution between 0-25 for the detectable mass bins ($m > 10^8$ $M_\odot$). To generalize this procedure to real lensing data, we would need to estimate the expected number of detectable subhalos per mass bin and the variance from N-body simulations.

4.6 Discussion and outlook

We developed a method to detect subhalos in strong gravitational lens images. The method is based on image segmentation: we classify each pixel in an image as belonging to either the main lens, a subhalo within a given mass range, or neither. We trained a convolutional neural network with a U-Net architecture on images with either no substructure or a single subhalo near the lensed images. When the network is applied to an

independent test set, it performs remarkably well.

We find that there are three common ways for the network to misclassify a pixel.

1. The subhalo is not detected and all of the pixels are assigned to the background or main lens classes. This happens more for light subhalos than large subhalos.

2. The pixel is on the edge of a subhalo and it is labeled as background instead of belonging to a subhalo. In these cases, the network finds the subhalo, but it is predicted to be a few pixels too large or too small.

3. Misidentifying the mass bin, generally by assigning the pixel to an adjacent mass bin. In these cases, the subhalo is still located correctly, but the mass is shifted up or down by a bin.

While the network is trained specifically looking pixel-by-pixel, we need to cluster the detected pixels into subhalos to extract some physical meaning to the pixel-based accuracies. On average, the subhalos have an area of 4π pixels, allowing for an easy conversion between the number of tagged pixels and the number of subhalos. Pixels around the edge of a subhalo can be missed while allowing for a detection of the subhalo. Because of this, the subhalo detection accuracy is better than the per-pixel accuracy.

Without noise, the class with the worst accuracy is that of subhalos in the 10^6 $M_\odot$ mass bin, but 62% of these subhalos are located, of which 78% get assigned the correct mass. The accuracy improves for heavier subhalos. At a mass of 10^8 $M_\odot$, the subhalos are detected 97% of the time, with 80% in the correct mass bin. The heaviest subhalos we considered have masses in the 10^{10} $M_\odot$ bin, and 99.9% of these subhalos are detected, with 96% getting assigned to the correct mass bin. When noise is included, the network loses sensitivity to the low-mass subhalos. For instance, with 1% noise, only 39% of the 10^8 $M_\odot$ subhalos are detected. However, the accuracy is similar for $m \geq 10^9$ $M_\odot$ with and without noise.

To put this in perspective, Refs. [103, 104] showed that gravitational imaging can find subhalos with masses of a few times 10^8 $M_\odot$ for a signal to noise ratio of as low as 3 when the subhalo has an NFW profile and on the Einstein ring. Additionally, these methods

have detected subhalos with masses of $(3.51 \pm 0.15) \times 10^9$ $M_\odot$ [106] and $(1.9 \pm 0.1) \times 10^8$ $M_\odot$ [107]. We reach good accuracy for the same range of masses detected by traditional methods using similar amounts of noise; see App. B.4 for a comparison of noise profiles. An advantage of our method is that we do not need to initially model the smooth lens to detect substructure in the system, which can take $\mathcal{O}$(weeks) to analyze in real systems. Instead, it takes us less than a second to run an image through the network. It is also important to keep in mind that our accuracies are presented for the *true* subhalo mass which was simulated. Direct detection methods are sensitive to the *effective* subhalo mass. Ref. [148] showed that the true mass can be biased by up to an order of magnitude compared to what is measured.

The network was also tested on images that do not contain substructure. False substructure was predicted at a rate of around 2.5 subhalos per 100 images. Most of these fall into the lightest mass bins that a network is sensitive to. The good detection accuracy and low false-positive rate implies that if the network predicts substructure is present an image, it is mostly likely truly due to a subhalo being present.

We also applied the network trained on images with single subhalos to images with many subhalos. We showed that the network has out-of-sample adaptability and can generalize to identify an abundance of substructure in a single image, although it does not see any such data during training. With this, we examined a method to determine the subhalo mass function power-law index from the network output of multiple images. The subhalo mass function is a key target for dark matter science because it can diagnose deviations from the standard cold dark matter scenario.

The technique uses a likelihood ratio for the count of detected subhalos in all of the bins, taking into account the expected counts given the power-law index and the network's detection accuracy for each class. We estimate that a 10% determination of the slope of the subhalo mass function will require around 50 (100) images without (with) noise. However, this is strongly dependent on astrophysical assumptions concerning the total mass of a halo contained in substructure. For instance, our catalog contained images with [0,25] subhalos with masses $m \geq 10^8$ $M_\odot$. If the true distribution of subhalos produces

significantly more or less subhalos near the Einstein ring, the number of images necessary for an accurate determination of the power-law index of the subhalo mass function could change.

We also note that Ref. [153] uses likelihood-free inference to extract population-level substructure properties about the dark matter subahalo population. They use a neural network to approximate the likelihood ratio function to obtain both the slope and the amplitude of the subhalo mass function with $\mathcal{O}(100)$ images. In contrast, we infer about the subhalo mass function by explicitly resolving individual subhalo, which can then be further studied.

In the near future, the Vera Rubin Observatory is expected to find 10^4 galaxy-galaxy strong lenses [167] and Euclid expects 10^5 galaxies lensed by galaxies in the field-of-view [168]. The ability to quickly detect substructure in these lenses could dramatically improve our understanding of dark matter. Image segmentation is a promising method to study dark matter subhalos. We stress that work will be needed to increase the robustness to noisier images (dimmer sources). Additionally, we have ignored the possibility of extra perturbations to the lens along the line-of-sight [6, 155, 156, 169], but not part of the main lens halo. These are both interesting problems that we save for future work.

Chapter 5

Power Spectrum of Dark Matter Substructure in Strong Gravitational Lenses

The main limitation of direct subhalo detection efforts is that only the most massive substructures lying within or very close to lensed arcs can be detected with large statistical significance. Even for our method in Chapter 4, the lowest mass that can be reliably obtained is $\sim 10^{8.5}$ $M_\odot$. While not directly detectable, smaller-mass substructures or those lying further away from lensed images could still potentially lead to observable effects on the lensing signal of extended arcs. Within the CDM paradigm, the (sub)halo mass function is expected to rise rapidly toward smaller masses [143], implying that typical lensed images could be perturbed by a fairly large number of unresolved low-mass substructures. In this limit, it becomes somewhat impractical to phrase the perturbations to lensed images in terms of individual subhalos. For instance, Refs. [114–116] proposed statistical techniques to harness the constraining power from these marginal detections for the properties and abundance of dark matter subhalos within lens galaxies.

The approach we propose is to describe the substructure convergence field in terms of its n-point correlation functions. For CDM, the large number of small-mass subhalos contributing to the total substructure convergence field implies that the statistics of

the latter should be nearly Gaussian. In this case, we expect the two-point correlation function (or its Fourier transform, the power spectrum) to dominate the statistical description of the substructure field. This last point was put forth in Ref. [117] to motivate an exploratory study of the detectability of the substructure convergence power spectrum within lens galaxies using the Atacama Large Millimeter/submillimeter Array (ALMA). In practice, given that strong lensing is probing the matter density field deep in the nonlinear regime, we do not expect the substructure density field to be entirely Gaussian. Nevertheless, measuring the substructure power spectrum might still lead to important insights about the abundance and internal structure of subhalos within lens galaxies.

Interestingly, Ref. [117] showed that it is possible, in principle, to measure the substructure convergence power spectrum by looking at the correlations of lensed image residuals, once a model image obtained from a purely smooth lens potential is subtracted from the data. They further showed that deep observations of strong gravitational lenses with ALMA could lead to 3σ detection of the nonvanishing amplitude of the substructure power spectrum (at least if there is abundant substructure, which is the case in CDM). Given that such measurements might be possible in the near future, the immediate question that comes to mind is: *What will we learn about low-mass subhalos from measuring the substructure convergence power spectrum?*

In this chapter, we present some much-needed answers to this question. Using the standard halo model [170] as our framework, we first develop a general formalism to compute the power spectrum of the convergence field on the lens plane due to substructure. We extend the initial approach presented in Ref. [117] to include subhalo populations that are not necessarily isotropic and homogeneous, and also take into account the 2-subhalo term. This formalism is developed in a way that makes it easy to change the statistical properties of the population as well as the intrinsic properties of subhalos, in order to facilitate its application to different dark matter scenarios. As an example we apply it to two different subhalo populations: one in which subhalos are modeled as truncated Navarro-Frenk-White (NFW) halos as would occur in standard CDM, and another one in which they are modeled as truncated cored halos as would happen in the presence of

self-interacting dark matter (SIDM). We choose the latter because of there is evidence of cored density profiles in at least some of the Local Group satellites (e.g. Refs. [171–173]). More generally, comparing these two models is interesting since they roughly bracket the range of possibilities for the inner density profiles of low-mass subhalos in a broad range of dark matter theories. We use these two examples as a springboard to discuss how the internal structure, statistical properties, and abundance of low-mass subhalos affect the shape and amplitude of the substructure convergence power spectrum.

This chapter is organized as follows. In Sec. 5.1 we present our halo model-based formalism to compute the substructure convergence power spectrum from first principles. In Sec. 5.2 we apply this formalism to study the 1- and 2-subhalo contributions to the substructure power spectrum from a population of truncated NFW subhalos. In Sec. 5.3 we turn our attention to the substructure power spectrum in the presence of a population of truncated cored subhalos, highlighting along the way the differences from the NFW case. We finally discuss our findings and conclude in Sec. 5.4.

5.1 Substructure statistics within the halo model

We work within the framework of the halo model [170], where all the dark matter is bound in roughly spherical halos. Within this model, the dark matter content of a typical lens galaxy is comprised of a smooth dark matter halo containing most the galaxy's mass, as well as a certain number of subhalos orbiting within the smooth halo. In the following, we will be concerned with these subhalos.

5.1.1 Preliminaries: Subhalo statistics

We work in projected two-dimensional (2D) space, with $\mathbf{r}$ denoting the projected 2D vector in the plane of the sky. In the limit that most of the lensing is caused by a single massive galaxy we can use the thin-lens approximation, so the relevant quantity is the surface mass density Σ of the lens galaxy in units of the critical density for lensing $\Sigma_{\rm crit}$, which is usually referred to as the convergence. $\Sigma_{\rm crit}$ depends on the angular diameter

distance between the observer and the source $D_{\rm os}$, the observer and the lens $D_{\rm ol}$ and the lens and the source $D_{\rm ls}$:

$$\Sigma_{\rm crit} = \frac{c^2 D_{\rm os}}{4\pi G D_{\rm ol} D_{\rm ls}}. \tag{5.1}$$

Here, G is the gravitational constant and c the speed of light. The convergence is also related to the projected Newtonian gravitational potential ϕ via the Poisson equation: $\nabla^2\phi = 2\kappa$. The strong lensing regime occurs when κ becomes of order unity.

The total convergence at a given point $\mathbf{r}$ on the lens plane is

$$\kappa_{\rm tot}(\mathbf{r}) = \kappa_0(\mathbf{r}) + \kappa_{\rm sub}(\mathbf{r}), \tag{5.2}$$

where κ_0 denotes the contribution from the smooth lens model (dark matter + baryons) and $\kappa_{\rm sub}$ denotes that from the subhalos. According to the standard CDM model, a typical lens galaxy will contain a large population of subhalos, all of which contribute to $\kappa_{\rm sub}$ as:

$$\kappa_{\rm sub}(\mathbf{r}) = \sum_{i=1}^{N_{\rm sub}} \kappa_i(\mathbf{r} - \mathbf{r}_i, m_i, \mathbf{q}_i), \tag{5.3}$$

where κ_i and $\mathbf{r}_i$ are the convergence and the position of the ith subhalo, respectively, m_i is the total mass of the ith subhalo, and the $\mathbf{q}_i$'s are sets of parameters that determine the internal properties of the ith subhalo. $N_{\rm sub}$ is the total number of subhalos contributing to the lensing convergence at position $\mathbf{r}$. Note that in Eq. (5.3) we have taken advantage of the fact that the overall contribution of the subhalo population is equivalent to the sum of the effect of each subhalo, which follows from the linearity of Poisson's equation. Since the convergence profile of a subhalo is always directly proportional to the subhalo mass m_i, it is useful to define $\hat{\kappa}_i \equiv \Sigma_{\rm crit}\kappa_i/m_i$. The advantage of this notation is that $\hat{\kappa}_i$ obeys a very simple normalization condition

$$\int d^2\mathbf{r}_i\, \hat{\kappa}_i(\mathbf{r}_i, \mathbf{q}_i) = 1, \tag{5.4}$$

independent of the value of $\mathbf{q}_i$. Here, the integral runs over the whole lens plane.

In general, it is impossible to know the mass, position, and internal properties of every subhalo within a lens galaxy. Instead, we would like to determine the "ensemble-averaged" properties of gravitational lensing observables given the statistical properties of subhalos, such as their mass function and spatial distribution. We shall denote by $\langle X \rangle$ the ensemble average of quantity X over all possible realizations of the subhalo density field within a lens galaxy. On the other hand, the notation $\bar{X}$ will be used to denote the "spatial" average of X over a given area of the lens plane.

Let us assume that all the statistical properties of subhalos within a lens galaxy are captured by a probability distribution function $\mathcal{P}(\mathbf{r}, m, \mathbf{q})$. It is, in general, a very good approximation (see Refs. [143, 174]) to assume that the mass and projected position of a subhalo are uncorrelated. This allows us to write the overall distribution as a product of a mass and position probability distributions as follows:

$$\mathcal{P}(\mathbf{r}, m, \mathbf{q}) = \mathcal{P}_{\rm r}(\mathbf{r})\mathcal{P}_{\rm m}(m)\mathcal{P}_{\rm q}(\mathbf{q}|m, \mathbf{r}), \tag{5.5}$$

where we have taken into account that the intrinsic properties of a given subhalo likely depend on its mass and position within the lens galaxy. The distribution $\mathcal{P}_{\rm r}(\mathbf{r})$ contains all the information about the projected spatial distribution of subhalos within the host galaxy. Given a projected number density $n_{\rm sub}(\mathbf{r})$ of subhalos, the probability of finding a subhalo within an area $d^2\mathbf{r}$ centered at position $\mathbf{r}$ is

$$\mathcal{P}_{\rm r}(\mathbf{r})d^2\mathbf{r} = \frac{n_{\rm sub}(\mathbf{r})d^2\mathbf{r}}{\int_A d^2\mathbf{r}\, n_{\rm sub}(\mathbf{r})}, \tag{5.6}$$

where A is the area of the lens plane where we have sensitivity to substructures (see below). The denominator in Eq. (5.6) is just the total number of subhalos within the area A

$$\int_A d^2\mathbf{r}\, n_{\rm sub}(\mathbf{r}) = N_{\rm sub} \equiv A\, \bar{n}_{\rm sub}, \tag{5.7}$$

where $\bar{n}_{\rm sub}$ is the average number density of subhalos averaged over the whole area A. It

is useful to write the subhalo number density as

$$n_{\rm sub}(\mathbf{r}) = \bar{n}_{\rm sub}\left(1 + \delta(\mathbf{r})\right), \tag{5.8}$$

where $\delta(\mathbf{r})$ is a stochastic random variable with $\langle\delta(\mathbf{r})\rangle = 0$. Here, the $\delta(\mathbf{r})$ field describes the fractional excess probability (compared to $\bar{n}_{\rm sub}$) of finding a subhalo at position $\mathbf{r}$. While any choice of $\delta(\mathbf{r})$ fully specifies the probability density function $\mathcal{P}_{\rm r}(\mathbf{r})$ as per Eq. (5.6) statistically independent, we will, in general, be interested in ensemble-averaging over realizations of the $\delta(\mathbf{r})$ field.

Numerical studies [143, 174] indicate that the 3D spatial distribution of subhalos near the central part of the host has a rather weak radial dependence. Taking into account projection effects and the fact that galaxy-scale strong lensing is mostly probing a small region near the projected center of the host, it is usually an excellent approximation to take $\langle n_{\rm sub}(\mathbf{r})\rangle = \bar{n}_{\rm sub} = \text{constant}$.

The subhalo mass probability distribution can be written as

$$\mathcal{P}_{\rm m}(m) \equiv \frac{1}{N_{\rm sub}}\frac{dN_{\rm sub}}{dm}, \tag{5.9}$$

where $dN_{\rm sub}/dm$ is the standard subhalo mass function. While our results are easily generalizable to any choice of mass function, we restrict ourselves to a power law mass function, $\mathcal{P}_{\rm m} \propto m^{\beta}$, for $m_{\rm low} < m < m_{\rm high}$. In the following, we assume that $\mathcal{P}(\mathbf{r}, m, \mathbf{q})$ is normalized such that

$$\int dm\, d^2\mathbf{r}\, d\mathbf{q}\, \mathcal{P}(\mathbf{r}, m, \mathbf{q}) = 1, \tag{5.10}$$

which is trivially satisfied by Eqs. (5.6) and (5.9).

As in most lensing calculations in the literature, the calculations presented in the remainder of this chapter assume that each subhalo represents an *independent* draw from the $\mathcal{P}(\mathbf{r}, m, \mathbf{q})$ probability distribution. We emphasize though that this does *not* mean that we neglect spatial correlations between subhalos; these are fully encoded in our choice of $\mathcal{P}_{\rm r}(\mathbf{r})$. In this case, the probability distribution describing the properties of

the whole subhalo population $\mathcal{P}_{\rm pop}$ can be factored out as a product of the probability distribution for single subhalos

$$\mathcal{P}_{\rm pop} = \prod_{i=1}^{N_{\rm sub}} \mathcal{P}(\mathbf{r}_i, m_i, \mathbf{q}_i). \tag{5.11}$$

We now have all the ingredients to perform ensemble averages over all possible realizations of a subhalo population.

5.1.2 Ensemble-averaged substructure convergence

It is instructive to first compute the mean ensemble-averaged substructure convergence on the lens plane $\bar{\kappa}_{\rm sub}$. It is given by

$$\begin{aligned} \bar{\kappa}_{\rm sub} &= \frac{1}{A}\int d^2\mathbf{s}\,\langle\kappa_{\rm sub}(\mathbf{s})\rangle \\ &= \frac{N_{\rm sub}}{A}\int dm_i\,d\mathbf{q}_i\,\mathcal{P}_{\rm m}(m_i)\mathcal{P}_{\rm q}(\mathbf{q}_i) \\ &\qquad \times \int d^2\mathbf{s}\,d^2\mathbf{r}_i\,\kappa_i(\mathbf{s}-\mathbf{r}_i, m_i, \mathbf{q}_i)\mathcal{P}_{\rm r}(\mathbf{r}_i), \end{aligned} \tag{5.12}$$

where we used the fact that every term in the sum in Eq. (5.3) contributes equally to $\bar{\kappa}_{\rm sub}$. The result is not surprising since it just states that the average convergence for the whole population of (statistically independent) subhalos is just $N_{\rm sub}$ times the average convergence of a single subhalo. Next, we note that the $\mathbf{r}_i$ integral above is nothing more than the convolution of the subhalo density profile κ_i with the spatial distribution $\mathcal{P}_{\rm r}$. Using the general result for the integral of a convolution,

$$\int d^2\mathbf{s}\,(f*g)(\mathbf{s}) = \int d^2\mathbf{s}\,f(\mathbf{s})\int d^2\mathbf{r}\,g(\mathbf{r}), \tag{5.13}$$

we obtain,

$$\begin{aligned} \bar{\kappa}_{\rm sub} &= \frac{N_{\rm sub}}{A\Sigma_{\rm crit}}\int dm_i\,\mathcal{P}_{\rm m}(m_i)m_i \\ &= \frac{N_{\rm sub}\langle m\rangle}{A\Sigma_{\rm crit}}, \end{aligned} \tag{5.14}$$

where we used Eq. (5.4). In the above, we have introduced the notation

$$\langle m \rangle \equiv \int dm_i \mathcal{P}_{\rm m}(m_i) m_i \tag{5.15}$$

to denote the average subhalo mass. We note that Eq. (5.14) is useful to relate $N_{\rm sub}$ and A to the physically relevant quantities $\langle m \rangle$ and $\bar{\kappa}_{\rm sub}$.

5.1.3 The power spectrum of the convergence field

We now turn our attention to the computation of the two-point correlation function of the substructure density field, or its Fourier transform, the substructure power spectrum. We emphasize that we do not assume here that the substructure convergence field is necessarily Gaussian. As such, we do not expect the power spectrum to characterize the substructure density field completely, and expect higher-point correlation functions to also contain nontrivial information. Nevertheless, the rapidly rising subhalo mass function toward the low-mass end in CDM models ensures that Gaussianity is a good first approximation [113]. Importantly, the main contributors of non-Gaussianities to the substructure field are the most massive subhalos within the lens galaxy [117]. Since we expect them to be directly detectable [106, 107, 111, 112], we can limit their influence on the statistics of the $\kappa_{\rm sub}$ field by absorbing the most massive subhalos within the macrolens mass model κ_0.

To obtain a general expression for the substructure power spectrum $P_{\rm sub}(k)$, we first compute the lens plane-averaged connected two-point correlation function $\xi_{\rm sub}(\mathbf{r})$ of the substructure convergence field $\kappa_{\rm sub}$. To simplify the derivation and avoid clutter, we first focus exclusively on performing the spatial averages encoded in the probability distribution $\mathcal{P}_{\rm r}(\mathbf{r})$. The averages over the subhalo mass and internal properties will be restored at the end of the calculation. The substructure convergence two-point function takes the

form

$$\xi_{\rm sub}(\mathbf{r}) \equiv \frac{1}{A}\int d^2\mathbf{s}\int\prod_i d^2\mathbf{r}_i\mathcal{P}_{\rm r}(\mathbf{r}_i) \times(\kappa_{\rm sub}(\mathbf{s})-\bar{\kappa}_{\rm sub})(\kappa_{\rm sub}(\mathbf{s}+\mathbf{r})-\bar{\kappa}_{\rm sub}). \tag{5.16}$$

Substituting Eq. (5.3) in the above and using the normalization condition given in Eq. (5.10), we obtain

$$\begin{aligned} A\xi_{\rm sub}(\mathbf{r}) = &\sum_i\int d^2\mathbf{s}\,d^2\mathbf{r}_i\kappa_i(\mathbf{s}-\mathbf{r}_i)\kappa_i(\mathbf{s}+\mathbf{r}-\mathbf{r}_i)\mathcal{P}_{\rm r}(\mathbf{r}_i)\\ &+\sum_i\sum_{j\neq i}\int d^2\mathbf{s}\,d^2\mathbf{r}_i\,d^2\mathbf{r}_j\mathcal{P}_{\rm r}(\mathbf{r}_i)\mathcal{P}_{\rm r}(\mathbf{r}_j)\\ &\qquad\times\kappa_i(\mathbf{s}-\mathbf{r}_i)\kappa_j(\mathbf{s}+\mathbf{r}-\mathbf{r}_j)\\ &-\bar{\kappa}_{\rm sub}\sum_i\int d^2\mathbf{s}\,d^2\mathbf{r}_i\kappa_i(\mathbf{s}-\mathbf{r}_i)\mathcal{P}_{\rm r}(\mathbf{r}_i)\\ &-\bar{\kappa}_{\rm sub}\sum_i\int d^2\mathbf{s}\,d^2\mathbf{r}_i\kappa_i(\mathbf{s}+\mathbf{r}-\mathbf{r}_i)\mathcal{P}_{\rm r}(\mathbf{r}_i)\\ &+\bar{\kappa}^2_{\rm sub}\int d^2\mathbf{s}. \end{aligned} \tag{5.17}$$

The first term arises from ensemble-averaging over the spatial distribution of a single subhalo (the "1-subhalo" term), the second term arises from averaging over pairs of distinct subhalos (the "2-subhalo" term), while the last three terms ensure that we are computing only the connected part of the two-point function. In the language of the halo model, the 1-subhalo term refers to particles or mass elements within a same subhalo, while the 2-subhalo term is due to those in distinct subhalos. The 1-subhalo term is nothing else than the convolution of the subhalo density profile with itself

$$\begin{aligned} &\int d^2\mathbf{s}\,d^2\mathbf{r}_i\kappa_i(\mathbf{s}-\mathbf{r}_i)\kappa_i(\mathbf{s}+\mathbf{r}-\mathbf{r}_i)\mathcal{P}_{\rm r}(\mathbf{r}_i)\\ &\quad=\int d^2\mathbf{x}\,\kappa_i(\mathbf{x})\kappa_i(\mathbf{x}+\mathbf{r})\\ &\quad=(\kappa_i*\kappa_i)(\mathbf{r}). \end{aligned} \tag{5.18}$$

The 2-subhalo contribution contains $N_{\rm sub}(N_{\rm sub}-1)$ identical terms which have the following form [175]

$$\begin{aligned}&\int d^2\mathbf{s}\, d^2\mathbf{r}_i\, d^2\mathbf{r}_j \mathcal{P}_{\rm r}(\mathbf{r}_i)\mathcal{P}_{\rm r}(\mathbf{r}_j)\kappa_i(\mathbf{s}-\mathbf{r}_i)\kappa_j(\mathbf{s}+\mathbf{r}-\mathbf{r}_j)\\ &\qquad = \int d^2\mathbf{x}\, d^2\mathbf{y}\, \kappa_i(\mathbf{x})\kappa_j(\mathbf{y})(\mathcal{P}_{\rm r}*\mathcal{P}_{\rm r})(\mathbf{y}-\mathbf{x}-\mathbf{r}).\end{aligned} \tag{5.19}$$

Using Eqs. (5.6) and (5.8), the convolution of the subhalo's spatial distribution is

$$\begin{aligned}(\mathcal{P}_{\rm r}*\mathcal{P}_{\rm r})(\mathbf{r}) &= \int d^2\mathbf{s}\, \mathcal{P}_{\rm r}(\mathbf{s})\mathcal{P}_{\rm r}(\mathbf{s}+\mathbf{r})\\ &= \frac{\bar{n}_{\rm sub}^2}{N_{\rm sub}^2}\int d^2\mathbf{s}\,(1+\delta(\mathbf{s}))(1+\delta(\mathbf{s}+\mathbf{r}))\\ &= \frac{\bar{n}_{\rm sub}}{N_{\rm sub}^2}\left(N_{\rm sub}+\bar{n}_{\rm sub}\int d^2\mathbf{s}\,\delta(\mathbf{s})\delta(\mathbf{s}+\mathbf{r})\right)\\ &= \frac{\bar{n}_{\rm sub}}{N_{\rm sub}}\left(1+\xi_{\rm ss}(\mathbf{r})\right),\end{aligned} \tag{5.20}$$

where we have identified the two-point subhalo correlation function $\xi_{\rm ss}(\mathbf{r})$, which encodes spatial correlation between pairs of distinct subhalos. Finally, the three last terms of Eq. (5.17) all have the same form and lead to a net contribution of $-\bar{\kappa}_{\rm sub}^2 A$. The connected two-point correlation function of the substructure convergence field thus takes the form

$$\begin{aligned}\xi_{\rm sub}(\mathbf{r}) = \frac{N_{\rm sub}}{A}(\kappa_i*\kappa_i)(\mathbf{r}) &\\ + \frac{\bar{n}_{\rm sub}N_{\rm sub}(N_{\rm sub}-1)}{AN_{\rm sub}}\int d^2\mathbf{x}\, d^2\mathbf{y}\, \kappa_i(\mathbf{x})\kappa_j(\mathbf{y})&\\ \times\left(1+\xi_{\rm ss}(\mathbf{y}-\mathbf{x}-\mathbf{r})\right)&\\ - \bar{\kappa}_{\rm sub}^2.&\end{aligned} \tag{5.21}$$

Noting that some of the integrals not involving $\xi_{\rm ss}$ in the second term exactly cancel the

third term, we are left with

$$\begin{aligned}\xi_{\rm sub}(\mathbf{r}) &= \bar{n}_{\rm sub}(\kappa_i * \kappa_i)(\mathbf{r}) \\ &+ \bar{n}_{\rm sub}^2 \int d^2\mathbf{x}\, d^2\mathbf{y}\, \kappa_i(\mathbf{x})\kappa_j(\mathbf{y})\xi_{\rm ss}(\mathbf{y}-\mathbf{x}-\mathbf{r}) \\ &- \frac{\bar{n}_{\rm sub}^2}{N_{\rm sub}} \int d^2\mathbf{x}\, d^2\mathbf{y}\, \kappa_i(\mathbf{x})\kappa_j(\mathbf{y})\left(1+\xi_{\rm ss}(\mathbf{y}-\mathbf{x}-\mathbf{r})\right).\end{aligned} \tag{5.22}$$

The first two terms correspond to the 1-subhalo and 2-subhalo terms, respectively, while the last term, suppressed by an extra factor of $N_{\rm sub}$, corresponds to the shot noise term, which only becomes important if the number of subhalos within the area of interest in the lens plane is small.

It is now straightforward to compute the convergence power spectrum by Fourier transforming Eq. (5.22). Using the following Fourier transform conventions:

$$\tilde{\kappa}(\mathbf{k}) = \int d^2\mathbf{r}\, e^{-i\mathbf{k}\cdot\mathbf{r}}\hat{\kappa}(\mathbf{r}), \tag{5.23}$$

$$\hat{\kappa}(\mathbf{r}) = \int \frac{d^2\mathbf{k}}{(2\pi)^2}\, e^{i\mathbf{k}\cdot\mathbf{r}}\tilde{\kappa}(\mathbf{k}), \tag{5.24}$$

the convergence power spectrum takes the form

$$\begin{aligned}P_{\rm sub}(\mathbf{k}) &= \int d^2\mathbf{r}\, e^{-i\mathbf{k}\cdot\mathbf{r}}\xi_{\rm sub}(\mathbf{r}) \\ &= \bar{n}_{\rm sub}|\tilde{\kappa}_i(\mathbf{k})|^2 \\ &\quad + \bar{n}_{\rm sub}^2(1-\frac{1}{N_{\rm sub}})\tilde{\kappa}_i(\mathbf{k})\tilde{\kappa}_j^*(\mathbf{k})P_{\rm ss}(\mathbf{k}),\end{aligned} \tag{5.25}$$

where $\mathbf{k}$ is the wavevector, and where we have used the convolution theorem to perform the Fourier transform. We note that the $\mathbf{r}$-independent part of the last term in Eq. (5.22) contributes an unobservable zero-mode, which we dropped in the above. Here, $P_{\rm ss}(\mathbf{k})$ is the Fourier transform of the subhalo two-point correlation function $\xi_{\rm ss}(\mathbf{r})$. In the remainder of the chapter we neglect the $1/N_{\rm sub}$ term in Eq. (5.25).

Up to this point, the only assumptions underpinning our calculation of the substructure convergence power spectrum are the statistical independence of each subhalo within

a lens galaxy, and the fact that the subhalo internal properties $\mathbf{q}_i$ do not depend on the subhalo position $\mathbf{r}_i$. We now introduce two simplifying assumptions:

- We take the subhalo convergence profile to be circularly symmetric, implying that $\tilde{\kappa}_i(\mathbf{k}) = \tilde{\kappa}_i(k)$.
- We assume that the subhalo two-point correlation function $\xi_{\rm ss}$ is homogeneous and isotropic, hence leading to $P_{\rm ss}(\mathbf{k}) = P_{\rm ss}(k)$.

Here, $k \equiv |\mathbf{k}|$. While subhalos are generally triaxial, projection effects and ensemble-averaging over all possible orientations and sizes of the subhalos' ellipticity imply that the average convergence profile is close to circularly symmetric, hence our first assumption. Our second point amounts to assuming that the small area of the lens plane probed by strong lensing images is typical of other nearby lines of sight. With these assumptions, the Fourier transform of the subhalo convergence profile is

$$\begin{aligned}\tilde{\kappa}(k) &= \int d^2\mathbf{r}\, e^{-i\mathbf{k}\cdot\mathbf{r}}\hat{\kappa}(r) \\ &= 2\pi \int dr\, r\, J_0(k\, r)\hat{\kappa}(r),\end{aligned} \tag{5.26}$$

where $J_0(x)$ is the 0th order Bessel function.

The last step of the calculation is to reinstate the averages over subhalo mass and internal properties. We can write the total substructure convergence power spectrum as the sum of the 1-subhalo and 2-subhalo terms,

$$P_{\rm sub}(k) = P_{\rm 1sh}(k) + P_{\rm 2sh}(k), \tag{5.27}$$

where the 1-subhalo term $P_{\rm 1sh}(k)$ takes the form

$$\begin{aligned}P_{\rm 1sh}(k) = \frac{(2\pi)^2\bar{\kappa}_{\rm sub}}{\langle m\rangle\Sigma_{\rm crit}} &\int dm\, d\mathbf{q}\, m^2\, \mathcal{P}_{\rm m}(m)\, \mathcal{P}_{\rm q}(\mathbf{q}|m) \\ &\times \left[\int dr\, r J_0(k\, r)\hat{\kappa}(r,\mathbf{q})\right]^2\end{aligned} \tag{5.28}$$

(the subscript i has been dropped since it is now superfluous) and the 2-subhalo term

takes the form

$$P_{\rm 2sh}(k) = \frac{(2\pi)^2 \bar{\kappa}_{\rm sub}^2}{\langle m \rangle^2} P_{\rm ss}(k) \Bigg[\int dm\, d\mathbf{q}\, m\, \mathcal{P}_{\rm m}(m)\, \mathcal{P}_{\rm q}(\mathbf{q}|m) \times \int dr\, r J_0(k\, r) \hat{\kappa}(r, \mathbf{q}) \Bigg]^2 . \tag{5.29}$$

The amplitude of the 1-subhalo term is approximately given by $P_{\rm 1sh}(k) \propto \bar{\kappa}_{\rm sub} m_{\rm eff}$, where the quantity $m_{\rm eff} \equiv \langle m^2 \rangle / \langle m \rangle$ has been referred to as the "effective mass" in the lensing literature [176–178]. This specific mass scale constitutes the primary dependence of the substructure power spectrum on the subhalo mass function, so we expect it to be one of the most constrained quantities with actual observations. The amplitude of the 1-subhalo term can be approximated as $P_{\rm 1sh}(k) \approx \bar{\kappa}_{\rm sub} m_{\rm eff} / \Sigma_{\rm crit}$. For a typical gravitational lens with $0.003 < \bar{\kappa}_{\rm sub} < 0.03$ [99], $m_{\rm eff} \sim 10^7\ {\rm M}_\odot$, and $\Sigma_{\rm crit} \sim 3 \times 10^9\ {\rm M}_\odot/{\rm kpc}^2$ (given our choices for the source and lens redshift), we thus expect

$$P_{\rm 1sh}(k) \sim 10^{-5} - 10^{-4}\, {\rm kpc}^2 \tag{5.30}$$

for scales larger than the typical size of a subhalo. On the other hand, the amplitude of the 2-subhalo term is approximately $P_{\rm 2sh}(k) \propto \bar{\kappa}_{\rm sub}^2 P_{\rm ss}(k)$, with very little dependence on the subhalo mass function. Given that typically $\bar{\kappa}_{\rm sub} \ll 1$ and that $P_{\rm ss}(k)$ can be important only on scales larger than the typical subhalo spatial separation, this term is generally subdominant compared to the 1-subhalo term, except maybe on larger scales, depending on the size of $P_{\rm ss}(k)$.

Having derived the general expression for the lens plane-averaged substructure power spectrum, we can now apply it to realistic subhalo populations by specifying the probability distributions $\mathcal{P}(\mathbf{r}, m, \mathbf{q})$ and the subhalo convergence profile $\kappa(\mathbf{r}, m, \mathbf{q})$. For definiteness, we make the following choices throughout the rest of this chapter whenever we present numerical results: we assume a lens galaxy at redshift $z = 0.5$ with virial mass and radius $M_{\rm vir} = 1.8 \times 10^{12}\ {\rm M}_\odot$, $R_{\rm max} = 409$ kpc, and Einstein radius $b = 6.3$ kpc. We take the source to be at $z = 1$.

5.2 Truncated Navarro-Frenk-White subhalo population

5.2.1 Characteristics of the subhalo population

In this section we compute the substructure power spectrum for a realistic population of smoothly truncated Navarro-Frenk-White subhalos. We are particularly interested in the strong lensing region, namely the region bounded more or less by the Einstein radius of the lens. Reference [113] performed a detailed analysis of the statistics of subhalo populations in strong lenses by looking at both the "local" (close to the Einstein radius of the host) and "distributed" (extending past the host virial radius) populations of subhalos and looking at their relative effects on lensing observables such as the lensing potential, deflection, shear and convergence. They found that the substructure contribution at a typical image position is largely dominated by the local subhalos.

The NFW density profile [71] has been found to provide a good fit to simulated CDM halos and is widely used to model the distribution of dark matter within galaxies and their satellites. This density profile (see Fig. 5.1) has an inner slope that goes as R^{-1} until it reaches the scale radius $r_{\rm s}$, where the slope steepens to R^{-3}. Formally, the NFW density profile leads to a divergent total subhalo mass. However, we expect tidal interactions to provide a finite truncation radius for a realistic subhalo orbiting within its host galaxy, hence leading to a finite subhalo mass. Here, we adopt the following truncated NFW profile (tNFW) [160] for our subhalos:

$$\rho_{\rm tNFW}(R) = \frac{m_{\rm NFW}}{4\pi R(R+r_{\rm s})^2}\left(\frac{r_{\rm t}^2}{R^2+r_{\rm t}^2}\right), \tag{5.31}$$

which is also shown in Fig. 5.1. Here, R is the three-dimensional distance from the center of the subhalo and $r_{\rm t}$ is the tidal radius. Observe that for $R \gg r_{\rm t}$, the density profile decays quickly as R^{-5}. Basically, our truncation scheme is meant to reflect that any dark matter particles outside $r_{\rm t}$ are tidally stripped as the subhalo undergoes a full orbit within its host. The tidal radius thus evolves in time, generally getting smaller as the

subhalo orbits within the tidal field of the host.

Projecting Eq. (5.31) along the line of sight leads to the following convergence profile for a tNFW subhalo [160]

$$\kappa_{\rm tNFW}(x) = \frac{m_{\rm NFW}}{\Sigma_{\rm crit} r_{\rm s}^2} \frac{\tau^2}{2\pi(\tau^2+1)^2} \Bigg[\frac{\tau^2+1}{x^2-1}(1-F(x)) + 2F(x) - \frac{\pi}{\sqrt{\tau^2+x^2}} + \frac{\tau^2-1}{\tau\sqrt{\tau^2+x^2}} L(x) \Bigg], \tag{5.32}$$

where

$$x = \frac{r}{r_{\rm s}}, \quad \tau = \frac{r_{\rm t}}{r_{\rm s}}, \tag{5.33}$$

$$F(x) = \frac{\cos^{-1}(1/x)}{\sqrt{x^2-1}}, \tag{5.34}$$

$$L(x) = \ln\left(\frac{x}{\sqrt{\tau^2+x^2}+\tau} \right). \tag{5.35}$$

The scale mass $m_{\rm NFW}$ is related to the total subhalo mass m via the relation [160]

$$m = \frac{m_{\rm NFW}\tau^2}{(\tau^2+1)^2} \left[(\tau^2-1)\ln(\tau) + \tau\pi - (\tau^2+1) \right]. \tag{5.36}$$

The parameter τ is similar to the concentration parameter, $c_{\rm vir} = R_{\rm vir}/r_{\rm s}$, which measures how concentrated the mass of a halo is since most of the mass is contained within $r_{\rm s}$. The tidal radius and virial radius are not necessarily the same however, so $c_{\rm vir} \neq \tau$.

In the notation of Sec. 5.1, the internal structure parameters for a truncated NFW subhalo are simply $\mathbf{q} = \{r_{\rm s}, r_{\rm t}\}$. Here, we adopt the following phenomenological relations

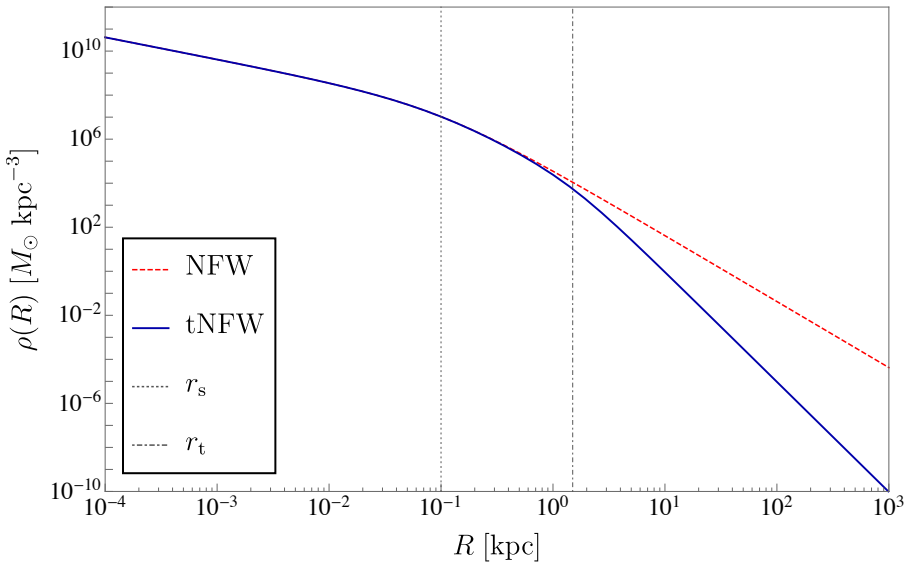

Figure 5.1: Density profile for a regular NFW profile (dashed red) and a truncated NFW profile (blue) for $\tau = 15$ and $m = 10^6$ M$_\odot$. The dotted and dashed-dotted gray lines represent the scale and tidal radius, respectively.

between the internal structure parameters and the subhalo mass and position [113]:

$$r_s = r_{s,0} \left(\frac{m}{m_0} \right)^{\gamma}, \tag{5.37}$$

$$r_t = r_{t,0} \left(\frac{m}{m_0} \right)^{1/3} \left(\frac{r_{3D}}{r_{3D,0}} \right)^{\nu}, \tag{5.38}$$

where we adopt below a fiducial value of $\gamma = 1/3$ [179, 180], and ν is a parameter that depends on the density profile of the host; for an isothermal profile $\nu = 2/3$, while $\nu = 1$ for a subhalo outside the scale radius of an NFW host [113]. The quantity r_{3D} is the three-dimensional distance between the subhalo and the center of the host galaxy, and $r_{s,0}$ and $r_{t,0}$ are, respectively, the scale and truncation radii for a subhalo of mass m_0 at position $r_{3D,0}$. For a pivot mass $m_0 = 10^6$ M$_\odot$, we adopt $r_{s,0} = 0.1$ kpc [179], $r_{t,0} = 1$ kpc, and $r_{3D,0} = 100$ kpc [139, 143].

In order to apply the result from the previous section, we need to know the distribution $\mathcal{P}_q(r_s, r_t|m)$, which we assume can be written as

$$\mathcal{P}_q(r_s, r_t|m) = \mathcal{P}_s(r_s|m)\, \mathcal{P}_t(r_t|m). \tag{5.39}$$

We model the distribution for scale radii assuming that the scatter in the scale radius-

mass relation Eq. (5.37) is normally distributed such that

$$\mathcal{P}_{\rm s}(r_{\rm s}|m) = \mathcal{N}\left(r_{\rm s,0}\left(\frac{m}{m_0}\right)^{\gamma}, \sigma_{r_{\rm s}} r_{\rm s,0}\left(\frac{m}{m_0}\right)^{\gamma}\right), \tag{5.40}$$

where $\mathcal{N}(\mu, \sigma)$ is a Gaussian probability distribution with mean μ and standard deviation σ, and $\sigma_{r_{\rm s}}$ is the fractional scatter about the scale radius-mass relation given in Eq. (5.37). We take $\sigma_{r_{\rm s}} = 0.2$ throughout the rest of this chapter, but we note that this specific choice has very little impact on our results.

Noting that $r_{\rm 3D}^2 = r^2 + h^2$, where h is the projection of $r_{\rm 3D}$ along the line of sight and r is the projection onto the lens plane, the distribution of tidal radii marginalized over h can be written as

$$\begin{aligned}\mathcal{P}_{\rm t}(r_{\rm t}|m, r) = \frac{1}{Z}\int dh\, \mathcal{P}_{\rm 3D}\left(\sqrt{r^2+h^2}\right) \\ \delta\left(r_{\rm t} - r_{\rm t,0}\left(\frac{m}{m_0}\right)^{1/3}\left(\frac{\sqrt{r^2+h^2}}{r_{\rm 3D,0}}\right)^{\nu}\right),\end{aligned} \tag{5.41}$$

where $\mathcal{P}_{\rm 3D}$ is the three-dimensional distribution of subhalos within the lens galaxy and Z is a normalization factor equal to the projection integral,

$$Z \equiv \int dh\, \mathcal{P}_{\rm 3D}(\sqrt{r^2+h^2}) = \mathcal{P}_{\rm r}(r). \tag{5.42}$$

Under the assumption that the projected distribution of subhalos is uniform, the radial distribution of subhalos is simply equal to the inverse area of the strong lensing region, $\mathcal{P}_{\rm r} = 1/A$. The choice of $\mathcal{P}_{\rm 3D}$ to obtain this is not unique. However, in the limit that the strong lensing region is probing only a small projected area of the host lens galaxy, we can obtain a unique expression for $\mathcal{P}_{\rm t}$ even when the distribution of subhalos is nonuniform. As shown in Appendix C.1, the integral in Eq. (5.41) can be performed in this limit to yield

$$\mathcal{P}_{\rm t}(r_{\rm t}|m) = \frac{1}{\nu R_{\rm max}}\frac{r_{\rm 3D,0}}{r_{\rm t}}\left[\left(\frac{m_0}{m}\right)^{1/3}\frac{r_{\rm t}}{r_{\rm t,0}}\right]^{1/\nu}. \tag{5.43}$$

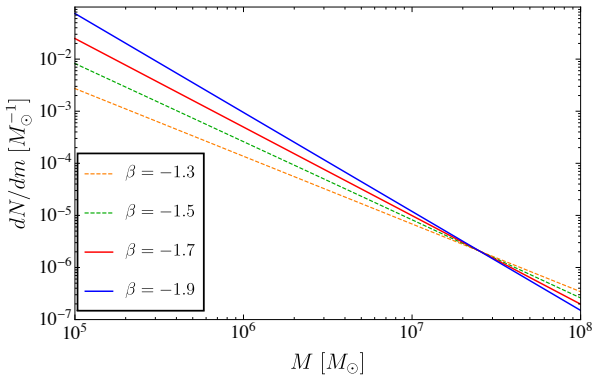

Figure 5.2: Subhalo mass function (Eq. 5.44) for different values of the power-law index β.

In the following we model the host as being isothermal, for which $\nu = 2/3$ as stated above. Note that $\mathcal{P}_\mathrm{t}$ has no dependence on the subhalo position within the host, consistent with the assumptions used in Sec. 5.1.

Lastly, we express the mass probability distribution as a power-law function [143]

$$\frac{dN_\mathrm{sub}}{dm} = a_0 \left(\frac{m}{m_*}\right)^\beta , \tag{5.44}$$

where $\beta = -1.9$ and $m_* = 2.52 \times 10^7$ $\mathrm{M}_\odot$. This mass function is illustrated in Fig. 5.2 for different choices of β. We note that the constant a_0, which normalizes the subhalo mass function, and the average convergence $\bar{\kappa}_\mathrm{sub}$ are proportional to one another as per Eq. (5.14). Typical gravitational lenses have an average convergence in the range $0.003 < \bar{\kappa}_\mathrm{sub} < 0.03$ [99], so we normalize the subhalo mass function such that $\bar{\kappa}_\mathrm{sub} = 0.02$.

Although we do not require the convergence field to be Gaussian, nor do we assume it, we do limit the large non-Gaussian contributions from the few most massive subhalos by setting an appropriate upper bound on the subhalo mass range included in our analysis. In practice, the maximum subhalo mass to include in the substructure convergence power spectrum calculation should be dictated by the dataset used to measure it. Indeed, the spatial resolution, pixel size, and the signal-to-noise ratio of the data specifies a subhalo mass sensitivity threshold below which a statistically significant direct detection of a subhalo is unlikely. For high-quality space-based optical data, this threshold could be as low as $\sim 10^8$ $\mathrm{M}_\odot$ [109], while for interferometric data it could reach $\sim 10^7$ $M_\odot$ [111]. Here, we adopt a fiducial value of $m_\mathrm{high} = 10^8$ $\mathrm{M}_\odot$. The minimum subhalo mass we

consider is $m_{\rm low} = 10^5\ {\rm M}_\odot$. As we will show below, the specific choice of $m_{\rm low}$ is largely inconsequential as long as $m_{\rm low} \ll m_{\rm high}$.

5.2.2 Power spectrum: 1-subhalo term

We can now apply the formalism developed in Sec. II to a population of tNFW subhalos to study how the abundance, density profile, radial distribution, and subhalo sizes affect the convergence power spectrum. In this case, Eq. (5.28) for the 1-subhalo term becomes

$$P_{\rm 1sh}(k) = \frac{\bar{\kappa}_{\rm sub}}{\langle m \rangle \Sigma_{\rm crit}} \int dm\ m^2\ \mathcal{P}_{\rm m}(m) \int dr_{\rm s}\ dr_{\rm t}\ \mathcal{P}_{\rm s}(r_{\rm s}|m) \times \mathcal{P}_{\rm t}(r_{\rm t}|m)\ |\tilde{\kappa}(k, r_{\rm s}, r_{\rm t})|^2. \tag{5.45}$$

Analytical discussion

For typical Poisson realizations of a population of spherically symmetric tNFW subhalos, we expect the behavior of the 1-subhalo term to depend mostly on three quantities: a low-k power spectrum amplitude, a turnover scale $k_{\rm trunc}$ corresponding approximately to the size of the largest subhalos, and an asymptotic high-k slope dictated by the small-r behavior of the subhalo density profile, which takes over for $k \gg k_{\rm scale}$ (defined below).

For small k values, we expect the 1-subhalo contribution to the power spectrum to plateau to a constant value since taking $k \to 0$ makes $J_0(kr) \to 1$ in Eq. (5.28), in which case κ and $P_{\rm 1sh}$ are k-independent. Another way to understand this low-k plateau is to realize that subhalos can be modeled as point masses, i.e. $\hat{\kappa}_i = \delta^{(2)}(\mathbf{r} - \mathbf{r}_i)$, on scales larger than the biggest subhalo's truncation radius, hence leading to $P_{\rm 1sh}(k) = \bar{\kappa}_{\rm sub}\langle m^2 \rangle / (\langle m \rangle \Sigma_{\rm crit})$. With $\bar{\kappa}_{\rm sub} = 0.02$ and our choice for the mass function parameters described above, we expect a low-k amplitude of $\sim 10^{-4}$ kpc^2.

As k is increased, the power spectrum begins probing the actual density profile of the subhalos, leading to a suppression of the power compared to the pure point-mass case. This turnover scale is determined by the truncated size of the largest subhalos, since this is the largest scale in the problem relevant to the 1-subhalo term. We therefore expect that this turnover is going to occur near a scale that corresponds to the inverse of the

tidal radius of the largest subhalo: $k_{\rm trunc} \equiv 1/r_{\rm t,max}$.

As k is further increased, the 1-subhalo term probes the intermediate scales between the typical truncation and scale radii of the tNFW subhalo population. Finally, we expect the convergence power spectrum to asymptote to a power-law behavior at large k where it is probing scales deep within the NFW scale radius. This power law can be determined by finding the small-x limit of the convergence profile given in Eq. (5.32),

$$\kappa_{\rm tNFW}(x) \to \frac{m_{\rm NFW}}{2\pi r_{\rm s}^2 \Sigma_{\rm crit}} \left(\ln\left(\frac{2}{x}\right) - 1 \right), \quad x \ll 1, \tag{5.46}$$

and taking the (2D) Fourier transform, which leads to

$$\tilde{\kappa}_{\rm tNFW}(k) \to \frac{1}{(k\, r_{\rm s})^2}, \quad k r_{\rm s} \gg 1. \tag{5.47}$$

This implies that $P_{\rm 1sh}(k) \propto 1/k^4$ for $k r_{\rm s} \gg 1$. We expect the power spectrum to reach this slope at a scale below that of the smallest scale radii in the population. It is therefore useful to define the wave number $k_{\rm scale} \equiv 1/r_{\rm s,min}$ beyond which the convergence power spectrum is a simple power law determined by the inner density profile of the subhalos.

Numerical results

Before ensemble-averaging over $\mathcal{P}_{\rm s}$ and $\mathcal{P}_{\rm t}$, it is informative to consider the shape of the convergence power spectrum for specific values of $r_{\rm s}(m)$ and $r_{\rm t}(m)$. Making the following choices:

$$\mathcal{P}_{\rm s}(r_{\rm s}|m) = \delta\left(r_{\rm s} - r_{\rm s,0} \left(\frac{m}{10^9\ {\rm M}_\odot} \right)^\gamma \right), \tag{5.48}$$

$$\mathcal{P}_{\rm t}(r_{\rm t}|m) = \delta(r_{\rm t} - 15 r_{\rm s}), \tag{5.49}$$

the 1-subhalo term takes the simple form

$$P_{\rm 1sh}(k) = \frac{\bar{\kappa}_{\rm sub}}{\langle m \rangle \Sigma_{\rm crit}} \int dm\, m^2\, \mathcal{P}_{\rm m}(m)\, |\tilde{\kappa}(k,m)|^2. \tag{5.50}$$

Note that Eq. (5.49) is equivalent to having a constant ratio for $\tau = r_{\rm t}/r_{\rm s} = 15$, which is not generally the case. From our expressions for the scale and tidal radius we expect τ to lie in the range $\approx 1 - 25$, depending on subhalo mass and position.

Figure 5.3 shows the power spectrum defined in Eq. (5.50). Panel (a) displays the features discussed in the preceding section, which have the expected behavior. The asymptotic low-k amplitude is 1.2×10^{-4} kpc^2 and matches the amplitude of the power spectrum of a population of point masses (black) with the same mass function. The truncation scale, which for $r_{\rm t,max} \simeq 7$ kpc is $k_{\rm trunc} = 0.14$ kpc^{-1} (dashed-dotted gray), very closely matches the scale at which the power spectrum turns over, consistent with the fact that this scale corresponds to the sizes of the largest subhalos. Furthermore, past $k_{\rm scale} = 21.5$ kpc^{-1} (gray) the large-k behavior matches a power law $1/k^4$ (dashed red), which again matches our expectation since in this regime we are within the scale radius of even the smallest subhalos i.e., where the tNFW convergence goes as Eq. (5.46).

In the remaining panels we vary several parameters of relevance to the power spectrum. Panel (b) shows the effect of changing the density profiles of subhalos by changing τ. When we increase τ, we are keeping $r_{\rm s}$ and m fixed while increasing $r_{\rm t}$, which means that the subhalo size is increasing and subhalos are becoming less concentrated toward the center. This has the effect of decreasing power on small scales and decreasing $k_{\rm trunc}$.

Panels (c) and (d) both reflect changes in the subhalo mass function: the former shows the result of varying $m_{\rm high}$ and $m_{\rm low}$, and the latter, the effect of making the power law shallower. Both changes affect the low-k amplitude as well as the distribution of power and slope on scales larger than $k_{\rm trunc}$; to disentangle these two effects we keep the quantity $\bar{\kappa}_{\rm sub} m_{\rm eff} = \bar{\kappa}_{\rm sub} \langle m^2 \rangle / \langle m \rangle$ fixed while changing the mass function, which makes the low-k amplitude remain the same. In this manner, we can isolate the effects of the subhalo mass function on the shape of the convergence power spectrum at high k. In Panel (c) we see that decreasing $m_{\rm high}$ by an order of magnitude adds power on small

scales. Indeed, removing the largest subhalos and redistributing their mass among smaller subhalos causes an increase in $k_{\rm trunc}$, which adds power on small scales. Panel (c) also illustrates the impact of increasing $m_{\rm low}$ from 10^5 $\mathrm{M}_\odot$ to 10^6 $\mathrm{M}_\odot$. The resulting change to the convergence power spectrum is rather small, reflecting the fact that the more massive subhalos tend to dominate the behavior of the power spectrum. This also implies that the convergence power spectrum shows little sensitivity to the low-mass cutoff of the mass function. Finally, Panel (d) shows that, by making the power law shallower, we are reducing power on small scales. To understand this effect, we refer the reader to Fig. 5.2, where one can see that by making the logarithmic slope shallower, we are decreasing the number of low-mass subhalos and are in fact increasing the number of subhalos more massive than the pivot mass. Note that despite the change in the shape of the power spectrum on intermediate scales, the spectra still match the $1/k^4$ power law of the fiducial case at $k \gtrsim k_{\rm scale}$.

Having gained some intuition into how different parameters in our model affect the power spectrum, we can move on to the more general case where we perform ensemble averages over the two intrinsic subhalo parameters: $r_{\rm s}$ and $r_{\rm t}$. The 1-subhalo power spectrum in this case is shown in Fig. 5.4. The fiducial model – shown in black in both panels – corresponds to the parameter values for $\mathcal{P}_{\rm t}$ and $\mathcal{P}_{\rm s}$, given in Eqs. (5.40) and (5.43)), $\nu = 2/3$ (isothermal lens) and $\gamma = 1/3$.

In each panel we show the effect of changing one of these parameters. Panels (a) and (b) reflect changes in ν and γ, respectively. It is immediately obvious from Panel (a) that changing the index ν has little impact on the convergence power spectrum, beside from a slight redistribution of power at intermediate and small scales. This means that the power spectrum will have limited sensitivity to the host galaxy's density profile; on the other hand, it also means that uncertainties on the density profile of the host will not prevent the power spectrum from being an effective tool to study subhalo populations.

Panel (b) of Fig. 5.4 demonstrates that the power law in the scale radius-mass relation can have a significant impact on the small-scale substructure convergence power spectrum. As we increase γ, the minimum scale radius decreases quickly, and so $k_{\rm scale}$ increases. In

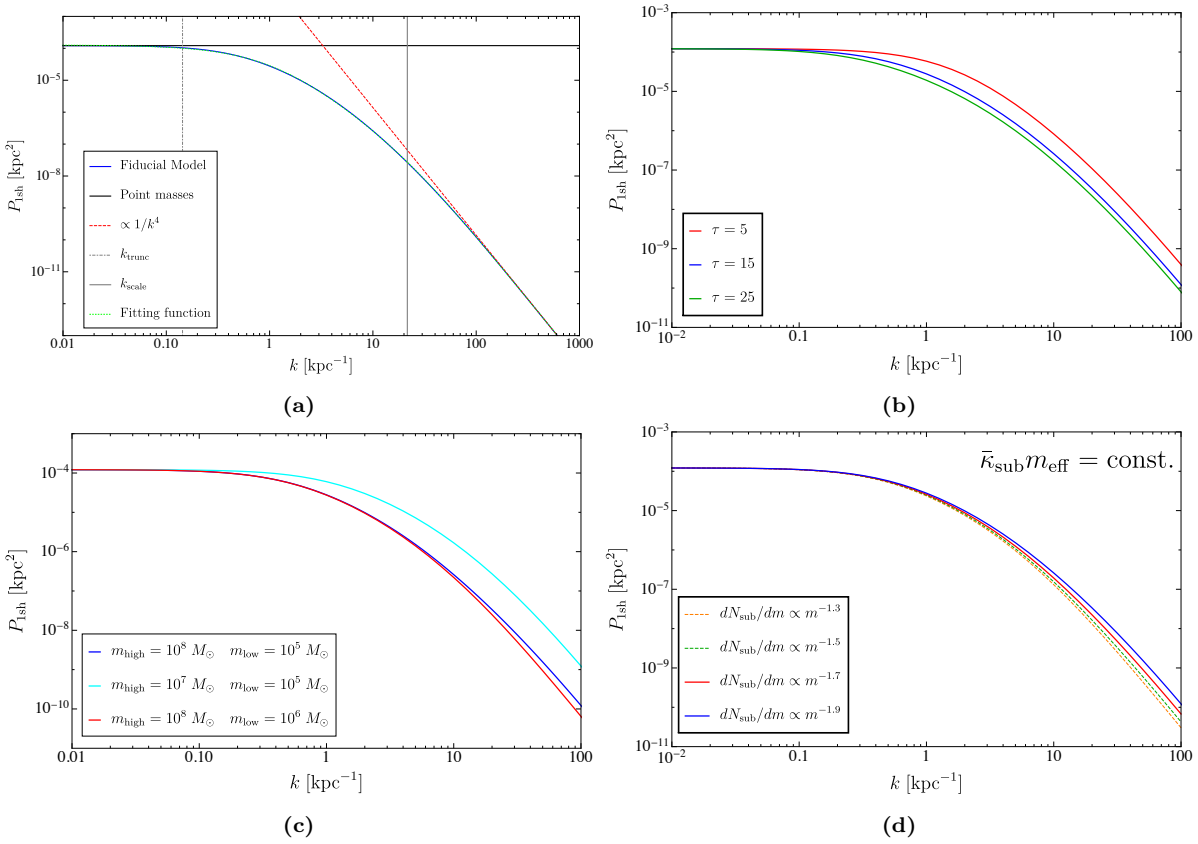

Figure 5.3: The 1-subhalo term of the convergence power spectrum of a population of truncated NFW subhalos. The solid blue line that appears in every subfigure represents the fiducial model with $\tau = 15$, 10^5 M$_\odot$ $\leq m \leq$ 10^8 M$_\odot$, $r_{\rm s}$ given by Eq. (5.37) with $\gamma = 1/3$, and $dN_{\rm sub}/dm$ given by Eq. (5.44) with $\beta = -1.9$. Panel (a) shows the features outlined in Sec. 5.2.2: the low-k amplitude of the power spectrum matches that of a population of point masses (solid black); the high-k slope is proportional to $1/k^4$ (dashed red); $k_{\rm trunc} \equiv 1/r_{\rm t,max} = 0.14$ kpc^{-1} (dotted-dashed gray); and $k_{\rm scale} \equiv 1/r_{\rm s,min} = 21.5$ kpc^{-1} (solid gray). The dotted green line corresponds to the fitting function described by Eqs. (5.51) – (5.56). In Panels (b) - (d) we change one parameter in the fiducial model while leaving the others unchanged. (b): changing τ by keeping $r_{\rm s}$ unchanged but increasing $r_{\rm t}$. (c): decreasing (increasing) $m_{\rm high}$ ($m_{\rm low}$) by an order of magnitude. (d): decreasing the logarithmic slope of the mass function down to $\beta = -1.3$. In Panels (c) and (d) $\bar{\kappa}_{\rm sub} m_{\rm eff}$ is held constant as the parameters are varied, where $m_{\rm eff} \equiv \langle m^2 \rangle / \langle m \rangle$. Note the different horizontal axis in Panel (a) and Panels (b) - (d).

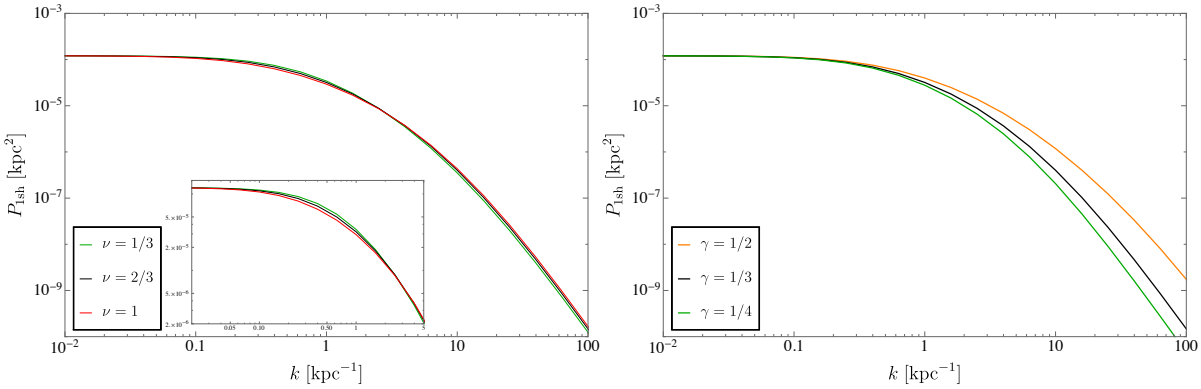

(a) **(b)**

Figure 5.4: Ensemble-averaged 1-subhalo term for a population of truncated NFW halos. The black line that appears in both panels has parameter values equal to the fiducial model in Fig. 5.3 (except for τ, which we do not fix). There are two additional parameters: $\nu = 2/3$ and $\sigma_{r_s} = 0.2$. Panel (a) varies the power-law dependence of the tidal radius on r_{3D}, Eq. (5.38). Panel (b) varies the power law of the scale radius-mass relation, Eq. (5.37).

fact $r_{s,min}$ decreases by an order of magnitude as we change γ from 1/4 to 1/2. This has the effect of adding power on small scales, as discussed in Sec. 5.2.2.

Another natural parameter to vary would be the scatter in the scale radius-mass relationship, σ_{r_s}. However, for a scatter of 20% or less, the impact on the convergence power spectrum is much smaller than the change associated with varying the index γ, and we therefore do not show it here. We also note that for a scatter larger than $\sim 20\%$, the approximate model presented in Eq. (5.40) likely breaks down at small subhalo masses, and should be replaced by a more realistic distribution of $\mathcal{P}_s(r_s|m)$.

We find that the 1-subhalo term for a population of tNFW halos is well fit by a function of the form

$$P_{1sh}(k) = \frac{g_0}{1 + g_1 k + (g_2 k)^2 + (g_3 k)^3 + (g_4 k)^4}, \tag{5.51}$$

where

$$g_0 = \frac{\bar{\kappa}_{sub} \langle m^2 \rangle}{\Sigma_{crit} \langle m \rangle}, \tag{5.52}$$

$$g_1 = \frac{(1/3)}{\gamma} \frac{\langle \tau \rangle r_{s,max}}{2\pi}, \tag{5.53}$$

$$g_2 = \left(\frac{(1/3)}{\gamma}\right)^2 \frac{\langle\tau\rangle r_{\rm s,max}}{2\pi}, \tag{5.54}$$

$$g_3 = r_{\rm s,max}, \tag{5.55}$$

$$g_4 = \frac{\langle m^2\rangle}{\int \frac{dm\, dr_{\rm t}\, dr_{\rm s}\, m^2 \mathcal{P}_{\rm m}(m)\, \mathcal{P}_{\rm s}(r_{\rm s}|m)\, \mathcal{P}_{\rm t}(r_{\rm t}|m)}{r_{\rm s}^4 \left(\frac{\tau^2}{(\tau^2+1)^2}[(\tau^2-1)\ln(\tau)+\tau\pi-(\tau^2+1)]\right)^2}}. \tag{5.56}$$

As shown, the parameters g_i are determined by the truncation, the scale radius, the mass function, and the mass-concentration relation. We note that this fit works best for values of $\gamma \leq 1/3$, and starts deviating from the "true" curve for higher values of γ. In the above, we have defined

$$\langle\tau\rangle \equiv \int dm\, dr_{\rm t}\, dr_{\rm s} \mathcal{P}_{\rm m}(m)\, \mathcal{P}_{\rm s}(r_{\rm s}|m)\, \mathcal{P}_{\rm t}(r_{\rm t}|m) \frac{r_{\rm t}}{r_{\rm s}}. \tag{5.57}$$

The fitting function is shown as a dotted green line in Panel (a) of Fig. 5.3.

5.2.3 Power spectrum: 2-subhalo term

To find the total power spectrum we have to include the contribution of the 2-subhalo term, given by Eq. (5.29). As explained in Ref. [181], the 2-subhalo term receives contributions from two distinct effects. First, subhalos have, in general, a nonuniform spatial distribution ($\mathcal{P}_{\rm r}(\mathbf{r})$ from Eq. (5.6)) due to their interaction with the potential well of their host halo. This so-called "host" contribution simply reflects the fact that subhalos can be gravitationally bound to their host lens galaxy, hence leading to a local enhancement of the convergence's two-point function. Second, subhalos can form self-bound groups orbiting their host galaxy. Due to tidal interactions with the latter, however, these subhalo groups are not expected to survive for more than a few dynamical times, [181] and we thus foresee their contribution to be subdominant. So far, this contribution to $\xi_{\rm ss}(\mathbf{r})$ has not been measured nor extracted from simulations, at least at the mass scale of interest (see

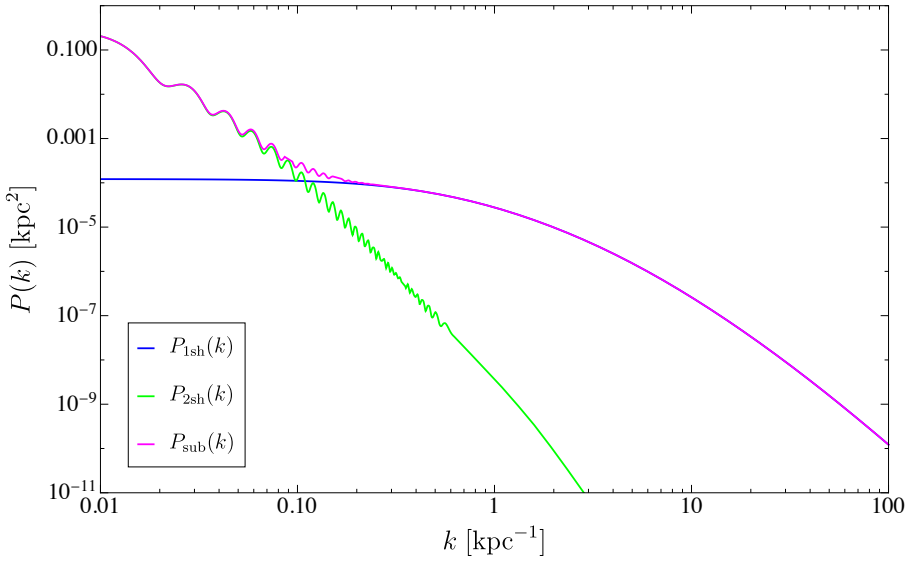

Figure 5.5: Full convergence power spectrum (magenta) and individual contributions from the 1-subhalo (blue) and 2-subhalo (green) terms, where the radial subhalo distribution used to calculate the 2-subhalo term is given by Eq. (5.58).

Ref. [182] for a measurement on cluster scales.). Due to this, we focus below on the host contribution, but the reader should keep in mind that the subhalo group contribution should be added in order to get a fully accurate estimate of the 2-subhalo term.

As an illustrative example, we choose a radial distribution of subhalos that is cored and decays as $1/r$ for large r,

$$\mathcal{P}_{\mathrm{r}}(r) = \frac{1}{2\pi(a+r)\left(R_{\mathrm{max}} + a\log\left(\frac{a}{a+R_{\mathrm{max}}}\right)\right)}, \tag{5.58}$$

where $a = 10$ kpc correponds to the core size. The total power spectrum $P_{\mathrm{sub}}(k)$ is shown in Fig. 5.5, together with the individual contribution of the 1- and 2-subhalo terms. On large scales, for $k \lesssim 0.1$ kpc$^{-1} = 1/a$, the 2-subhalo term dominates, adding power and changing the low-k slope from a constant to a power law. On small scales, however, the 1-subhalo term dominates (as expected), and the addition of the 2-subhalo term leaves the power spectrum unchanged. Note that the oscillations at small k come from having $\mathcal{P}_{\mathrm{r}}(r)$ nonzero over a finite region in the lens plane.

5.3 Truncated Cored subhalo population

In Sec. 5.1 we applied the convergence power spectrum formalism to a population of truncated NFW subhalos, since CDM halos in simulations seem to universally have NFW density profiles. We now apply the same methodology to a population of subhalos whose density profiles approximate what we expect SIDM subhalos to look like: cored at the center and with a large-r behavior similar to NFW. The idea is to gauge the extent to which the power spectrum differs for NFW and cored profiles, which could be indicative of the utility of this observable in discerning between CDM and a different dark matter scenario in which halos are predicted to have cores instead, like SIDM. As we have emphasized in preceding sections, there are essentially two types of ingredients that go into the convergence power spectrum: the statistical properties of the subhalo population and the internal subhalo parameters, which determine the surface mass density profile.

With respect to the first point, SIDM N-body simulations have shown that, at least in the case of elastic scattering with cross section $\sigma/m \lesssim 1$ cm^2/g, the spatial distribution and number density of subhalos are largely unchanged [65–68]. Indeed, we expect that the subhalo distribution on the lens plane will be largely intact with respect to the CDM case since the volume occupied from the outskirts of the lens galaxy to the edge of its central region, where dark matter self-interactions can play a role, is many orders of magnitude larger than the volume occupied by the host's core itself; in fact the latter makes up about $\sim 2\%$ of the total line-of-sight volume. Furthermore, simulations find that there is essentially no change to the subhalo mass function for moderate dark matter self-interaction cross sections (at least down to 10^6 M$_\odot$; refer to Fig. 6 of Ref. [65] to see both of these points).

With respect to the second point, there is a stark contrast between CDM and SIDM dark matter halos due to the appearance of a central core in the latter. A common cored density profile is the Burkert profile [183],

$$\rho_{\rm b}(R) = \frac{m_{\rm b}}{4\pi(R + r_{\rm b})(R^2 + r_{\rm b}^2)}, \tag{5.59}$$

where $r_{\rm b}$ is the core radius, and the scale mass $m_{\rm b}$ is the mass within the core. Here we set $r_{\rm b} = p\, r_{\rm s}$, where p is a constant that represents the size of the core as a fraction of the scale radius. Furthermore, we also add a smooth truncation term, resulting in a profile of the form

$$\rho_{\rm tBurk}(R) = \frac{m_{\rm b}}{4\pi(R + p\, r_{\rm s})(R^2 + p^2 r_{\rm s}^2)} \left(\frac{r_{\rm t}^2}{R^2 + r_{\rm t}^2}\right), \tag{5.60}$$

where the total mass of the subhalo with this profile is given by

$$m = m_{\rm b} \frac{\tau^2 \left(\pi(p-\tau)^2 + 4\tau^2 \log\left[\frac{p}{\tau}\right]\right)}{4(p^4 - \tau^4)}. \tag{5.61}$$

We call this a truncated Burkert (tBurk) profile. Note that for a given p, the intrinsic parameters for the tBurk subhalos are the same as for the tNFW ones: $\mathbf{q} = \{r_{\rm s}, r_{\rm t}\}$. This profile is shown in Fig. 5.6, where we show the tNFW profile and tBurk profile for $p = 0.7$. This choice for p is motivated by the fact that Ref. [66] finds that for them, $r_{\rm b}$ in Eq. (5.59) corresponds to the CDM $r_{\rm s}$ value of $r_{\rm b} = 0.7 r_{\rm s}$. The tBurk profile exhibits a characteristic bump expected in SIDM halos, which is due to the redistribution of mass at the halo center caused by injecting kinetic energy from the outskirts of the halo towards the inner regions [60].

Using Eq. (5.60) we find an analytic expression for the convergence:

$$\begin{aligned}\kappa_{\rm tBurk}(x) = \frac{m_{\rm b}}{8\pi\Sigma_{\rm crit} r_{\rm s}^2}\, \tau^2 \Bigg\{ \pi \Bigg(& \frac{2p\sqrt{\frac{1}{\tau^2+x^2}}}{p^4 - \tau^4} - \frac{\sqrt{\frac{1}{x^2-p^2}}}{p(\tau^2+p^2)} \\ & - \frac{\sqrt{\frac{1}{x^2+p^2}}}{p^3 - p\tau^2} \Bigg) + \frac{2\arctan\left[\frac{p}{\sqrt{x^2-p^2}}\right]}{\sqrt{x^2-p^2}(p^3+p\tau^2)} - \\ & \frac{2\tanh^{-1}\left[\frac{p}{\sqrt{p^2+x^2}}\right]}{\sqrt{x^2+p^2}(p^3-p\tau^2)} + \frac{4\tau\tanh^{-1}\left[\frac{\tau}{\sqrt{x^2+\tau^2}}\right]}{\sqrt{x^2+\tau^2}(p^4-\tau^4)} \Bigg\},\end{aligned} \tag{5.62}$$

where again $x = r/r_{\rm s}$ and $\tau = r_{\rm t}/r_{\rm s}$ (refer to Appendix C.2 for details).

As stated above, we are assuming that the spatial distribution of subhalos within the host dark matter halo remains essentially intact in going from CDM to SIDM. Under this assumption, the 2-subhalo term should remain unchanged in going from one dark matter

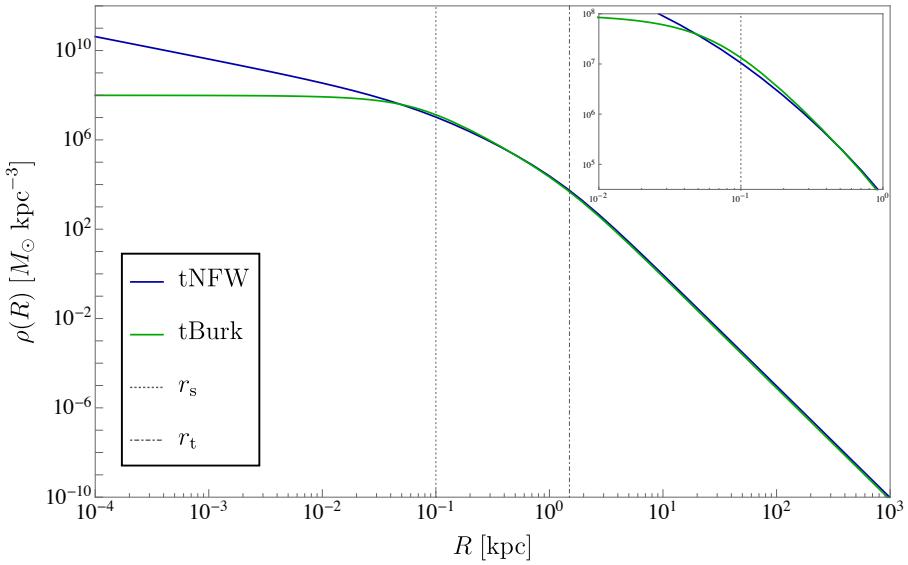

Figure 5.6: Density profile for a truncated NFW profile (solid blue) and a truncated Burkert profile (solid green) for $\tau = 15$, $p = 0.7$, and $m = 10^6$ $M_\odot$. The gray dotted and dashed-dotted lines represent the scale and tidal radius, respectively.

scenario to the other. Of course, realistically it is likely that the 2-subhalo term would actually be different to some extent: as subhalos orbit the host, the friction felt between the parent halo and the smaller subhalos would have an effect on the correlation of subhalo positions, especially since this effect would affect different subhalo orbits asymmetrically.

Assuming the 2-subhalo term to be the essentially same as in the tNFW case, we focus the rest of this section on the expected redistribution of power on small scales in the 1-subhalo term. In the forthcoming discussion we will therefore explore the extent of this high-k difference between the two density profiles we've chosen to be representative of each dark matter scenario.

We follow an identical procedure to the tNFW case to determine the 1-subhalo term of the power spectrum, which is shown in Fig. 5.7. We also show, for reference, the fiducial tNFW case shown in blue in Fig. 5.3. There is a slight increase in power with respect to the tNFW population on intermediate scales due to the redistribution of mass as the core forms, followed by the expected decrease in power on small scales due to the actual core. Despite these differences, we note that the changes of the substructure convergence power spectrum on scales $k_{\rm trunc} \lesssim k \lesssim k_{\rm scale}$ in going from the tNFW to the tBurk case

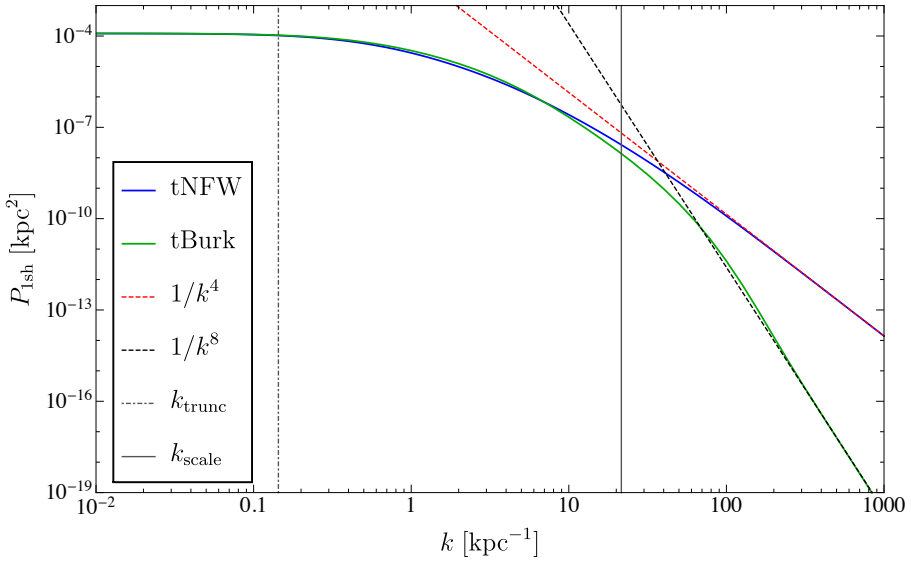

Figure 5.7: 1-subhalo power spectrum for a population of tNFW subhalos (solid blue; same fiducial model as in Fig. 5.3) and tBurk subhalos (solid green). We also show $k_{\rm trunc}$ (dotted-dashed gray) and $k_{\rm scale}$ (solid gray), as well as the $k \gg k_{\rm scale}$ behavior of both power spectra.

is well within the variation allowed by varying the statistical properties of the subhalo population, i.e., the different effects shown across Figs. 5.3 and 5.4. This implies that measurements of the power spectrum on these scales are unlikely to distinguish between a cored or cusped subhalo profile.

On even smaller scales $k \gg k_{\rm scale}$, the tBurk power spectrum $P_{\rm 1sh}(k)$ begins to significantly deviate from its tNFW counterpart. Indeed, since the Fourier transform of the truncated Burkert profile behaves as

$$\tilde{\kappa}_{\rm tBurk}(k) \to \frac{8(p^4-\tau^4)}{\tau^2\left(\pi(p-\tau)^2+4\tau^2\log\left[\frac{p}{\tau}\right]\right)}\frac{1}{(k\,p\,r_{\rm s})^4}, \tag{5.63}$$

for $k\,p\,r_{\rm s} \gg 1$, the 1-subhalo term for a population of cored subhalos goes as $P_{\rm 1sh}(k) \propto 1/k^8$ for large k, much steeper than the $1/k^4$ expected for NFW subhalos. Therefore, if at all measurable (see discussion below), the slope of the power spectrum on these scales could be decisive in determining the inner density profile of subhalos, which in turn could shed light on the particle nature of dark matter.

5.4 Discussion and Conclusion

In this chapter we have introduced a general formalism to study the two-point correlation function of the convergence field due to subhalo populations in strong gravitational lenses, keeping in mind that the observables for these types of problems tend to be photon count or surface brightness maps that exhibit multiple images due to the light from a background source (e.g. a quasar or a galaxy) having been warped by a massive foreground object, namely the gravitational lens. Fundamentally, the crucial insight to be gained from considering the convergence power spectrum is the ability to describe the effect of substructure in a language that is closer to what strong lensing observations are directly measuring. Indeed, while substructure lensing is often phrased in terms of the subhalo mass function, gravitational lensing observations primarily constrain the length of the deflection vectors at different positions on the lens plane. Since the power spectrum directly describes on which length scales the substructure contributes most to the deflection field, it allows a more direct connection to the actual observations without introducing an intermediate mass function. For the purpose of using lensing observations to extract information about dark matter physics, it is nevertheless important to connect the power spectrum language to the perhaps more familiar halo model of structure formation, for which predictions for different dark matter theories are more readily available.

We have explored in depth how different subhalo population properties affect the substructure convergence field, as well as how it differs for two alternative dark matter scenarios: CDM, which we have represented as a population of tNFW subhalos, and SIDM, where we used a truncated generalized Burkert profile to represent the subhalo population.

Using the CDM scenario as our baseline, we found that the form of the 1-subhalo term is largely determined by three key quantities: a low-k amplitude proportional to $\bar{\kappa}_{\rm sub}\langle m^2\rangle/\langle m\rangle$, a turnover scale $k_{\rm trunc}$ where the power spectrum starts probing the density profile of the largest subhalos, and the wave number $k_{\rm scale}$ corresponding to the smallest scale radii beyond which the slope of the power spectrum reflects the inner density profile of the subhalos. We have shown that the first of these is directly related to subhalo

abundance and specific statistical moments of the subhalo mass function. On the other hand, the turnover scale is determined by the average truncation radius of the largest subhalo included in the power spectrum calculation. On scales $k_{\rm trunc} \lesssim k \lesssim k_{\rm scale}$, there is significant variability depending on the statistical properties of subhalos - i.e. changes to the tidal truncation, to parameters pertaining to the subhalo mass function, or to the scale radius-mass relation can shift the distribution of power and slope on these scales in a rather degenerate manner (see Figs. 5.3 and 5.4). This indicates that measurements of the substructure convergence power spectrum might not be able to distinguish between changes to these different subhalo statistical properties.

For SIDM-like subhalos with a truncated Burkert profile, much of the same discussion applies. While in general the difference between the tNFW and tBurk power spectra is well within the range allowed by varying subhalo population parameters (such as the mass function), there is one defining characteristic that could set both scenarios apart: the high-k slope. For a population of cored, tBurk subhalos, the high-k slope is much steeper than for tNFW, and goes as $1/k^8$ as opposed to the $1/k^4$ behavior of tNFW. While not discussed in this chapter, we note that a population of truncated isothermal ("pseudo-Jaffe" [184]) subhalos would lead to a shallower substructure convergence power spectrum going as $1/k^2$ at large wave numbers. Remarkably, the high-k ($k \gtrsim k_{\rm scale}$) slope appears robust to changes in other parameters that govern the statistical properties of the subhalo population, despite the variation at intermediate wave numbers. Note that this is true even when taking into account our lack of knowledge about the 2-subhalo term, since it will not have a noticeable contribution on such small scales. Therefore, determining the high-k slope of the power spectrum would be of particular interest since it would allow us to distinguish between cusped and cored profiles, and more generally, to determine the average small-r behavior of the subhalo density profile.

The Fisher forecast estimates of Ref. [117] (Fig. 5 in their paper) seem to imply that $\sim 10 - 40$ hour long ALMA observations would be able to measure the amplitude of the power spectrum as well as $k_{\rm trunc}$. However, based on their results, it seems unlikely that these observations would be able to constrain the high-k slope of the power spectrum.

Therefore, although we may characterize the abundance of subhalos and the average size of the largest unresolved subhalos with ALMA, it appears unlikely that we will be able to fully constrain the average density profile of subhalos. A measurement of the latter would require a ~ 10 pc-level resolution within an object that is cosmologically distant from the Milky Way, a very difficult observation indeed, but not necessarily out of reach of very long baseline interferometry. Even if such a measurement could be made, however, it is likely that baryonic structures such as giant molecular clouds [185] and globular clusters [186, 187] would contribute to the convergence power spectrum on these scales and could contaminate the signal on scales $k \gtrsim 10$ kpc^{-1}.

There are several potential future directions to the work presented here. An immediate next step would be to compare our analytical results to the substructure convergence power spectrum extracted from high-resolution simulations. Such a comparison could also allow us to obtain a better estimate of the magnitude of the 2-subhalo term, and help us determine whether it can become more important than the 1-subhalo term on larger scales. This is done in Chapter 6. It would also be interesting to estimate the contribution to the convergence power spectrum from baryonic structures and line-of-sight subhalos [155, 188, 189]. The line-of-sight contribution is the subject of Chapter 7. Our analysis could also be improved by allowing the internal shape of the subhalo density profile to vary as a function of mass to take into account the fact that more massive subhalos may be more affected by baryonic feedback (and thus allowing them to form cores) than less massive subhalos. In order to combine measurements from different strong lenses, it will also be of primary importance to understand how the substructure power spectrum depends on the properties (e.g. redshift, concentration, stellar content, etc.) of the host lens galaxy [190].

In this chapter, we have computed the lens plane-averaged (that is, the monopole) substructure convergence power spectrum since it is the quantity that is most readily extracted from observations. However, since lens galaxies are generally not spherically symmetric (see e.g. Refs. [191, 192]), it is entirely possible that the substructure power spectrum is not isotropic, and it might be fruitful to also consider the higher multipoles of

the power spectrum, as it is done, for instance, in the case of the galaxy power spectrum in large-scale structure surveys (see e.g. Ref. [193]). By breaking rotational symmetry a new relevant scale could arise in the power spectrum, potentially breaking some of the degeneracy between different astrophysical parameters that was exhibited in the power spectra we considered in this chapter. In addition, it is possible that non-Gaussian signatures encoded in the higher n-point correlation functions could also contain important information about mass substructures within lens galaxies.

In conclusion, we have performed a detailed study of the amplitude and shape of the substructure convergence power spectrum within lens galaxies. We have shown how important features of the subhalo population get imprinted on the power spectrum. Based on the sensitivity and resolution of near-future observations, it appears unlikely that substructure power spectrum measurements would be able to probe the inner density profile of dark matter subhalos. Nevertheless, such measurements will provide some constraints on the abundance, mass function, and tidal truncation of low-mass subhalos within lens galaxies, and thus constitute a key consistency test of the standard CDM paradigm. In the event that the measured substructure power spectrum significantly deviates from our CDM expectations, they may even shed new light on the particle nature of dark matter.

Chapter 6

Gravitational Lensing and the Power Spectrum of Dark Matter Substructure: Insights from the ETHOS N-body Simulations

In this chapter, we present the first in-depth analysis of the dark matter substructure power spectrum in zoom-in N-body simulations of galactic halos at redshifts relevant to galaxy-scale strong lensing. We consider two high-resolution simulations of a Milky Way-sized halo, one in which the simulation particles are modeled as being CDM (namely they only interact gravitationally) and another in which they are allowed to self-interact and a cutoff is imposed in the initial cosmological matter power spectrum. We use these simulations to compute the substructure power spectrum and study its behavior as a function of redshift and of dark matter microphysics.

We focus here exclusively on the contribution from the subhalos orbiting the main lens galaxy to avoid complications related to multi-plane lensing. Since the line-of-sight contribution [155, 156, 194, 195] to the power spectrum is unlikely to be correlated with the galactic contribution we study here, it can be computed separately. This is precisely the subject of the following chapter.

This chapter is organized as follows. In Section 6.1 we succinctly summarize and rewrite some of the key equations from Chapter 5 in a format that will be beneficial here. In Section 6.2 we present the main features of the simulations used in this chapter, and in Section 6.3 we introduce our methodology. In particular, we calculate the power spectrum in two different ways: one is akin to the halo model prescription, where matter overdensities are assumed to be neatly packed into halos, and another where we make no such assumption. Comparing these two methods is instructive because the substructure within a lens is varied, and all contributions to the convergence that cannot be ascribed to the main lens galaxy do not necessarily come from neatly distinguishable subhalos. In Section 6.4 we present our results and we conclude in Section 6.5.

6.1 The Power Spectrum of Dark Matter Substructure within galaxies

Recall that the power spectrum of the substructure convergence field can be written as the sum of a one-subhalo and two-subhalo contributions

$$P_{\rm sub}(\mathbf{k}) = P_{\rm 1sh}(\mathbf{k}) + P_{\rm 2sh}(\mathbf{k}), \tag{6.1}$$

where

$$P_{\rm 1sh}(\mathbf{k}) = \frac{\bar{\kappa}_{\rm sub}\Sigma_{\rm crit}}{\langle m\rangle}\langle\widetilde{\kappa}(\mathbf{k})^2\rangle, \tag{6.2}$$

and

$$P_{\rm 2sh}(\mathbf{k}) = \left(\frac{\bar{\kappa}_{\rm sub}\Sigma_{\rm crit}}{\langle m\rangle}\right)^2\langle\widetilde{\kappa}(\mathbf{k})\rangle^2 P_{\rm ss}(\mathbf{k}). \tag{6.3}$$

$\widetilde{\kappa}(\mathbf{k})$ is the Fourier transform of the subhalo convergence profile, $P_{\rm ss}(\mathbf{k})$ is the Fourier transform of the subhalo spatial two-point correlation function (describing subhalo clustering), and the angular brackets represent an ensemble average over subhalo properties

such as their mass, truncation radius, and scale radius.

While the substructure power spectrum given in Eq. (6.1) is in principle anisotropic due to the complex structure of a typical galaxy, we expect the isotropic (monopole) contribution to dominate the signal within the small region probed by strong lensing. This monopole power spectrum is simply given by

$$P_{\rm sub}(k) = \frac{1}{2\pi}\int_0^{2\pi} P_{\rm sub}(\mathbf{k})\, d\phi, \tag{6.4}$$

where ϕ is the polar angle of the $\mathbf{k}$ vector. We will focus on this isotropic contribution in the remainder of this chapter.

6.2 Simulations

The N-body simulations used in this work are the ETHOS (Effective Theory of Structure Formation) simulations, originally presented in Ref. [60]. The goal of the ETHOS project [196, 197] is to understand how the fundamental dark matter microphysics affects structure formation on a broad range of scales. To this end, five different dark matter models were investigated: a cold dark matter (CDM) scenario and four scenarios that explore the dark matter parameter space that includes dark matter-dark radiation (DM-DR) interactions, which are responsible for a primordial cutoff in the power spectrum, and self-interacting dark matter (SIDM), labeled ETHOS1-4 depending on the choice of parameter values. In this work we focus on the CDM simulation together with the ETHOS4 model, which has been chosen to reproduce the observed kinematics and properties of Milky Way (MW) dwarf spheroidals.[1] We refer the reader to Refs. [60, 196] for more details about the ETHOS4 dark matter model, including the values of the particle physics parameters used in the simulations.

The simulations are initialized at $z = 127$ within a $100h^{-1}$ Mpc periodic box, from which a MW-sized halo (1.6×10^{12} $M_\odot$) is chosen to be resimulated. The parent simula-

[1]Note that the halos in the CDM and ETHOS4 simulations have the same initial conditions (but of course differ in the initial power spectrum and dark matter microphysics).

tion has 1024^3 particles, a mass resolution of $7.8 \times 10^7 h^{-1}$ $M_\odot$, and a spatial resolution (Plummer-equivalent softening length) of $\epsilon = 2h^{-1}$ kpc. They are thus able to resolve halos down to $\sim 2.5 \times 10^9 h^{-1}$ $M_\odot$ with 32 particles. They re-simulate the MW-sized halo to different resolution levels. For this work we use the highest resolution simulation (level 1), where the dark matter particle mass is $m_{\rm DM} = 2.756 \times 10^4$ $M_\odot$, $\epsilon = 72.4$ pc and there are approximately 4.44×10^8 high-resolution particles in each zoomed-in simulation.

The cosmological parameters used in the simulations are: $\Omega_{\rm m} = 0.302$, $\Omega_\Lambda = 0.698$, $\Omega_{\rm b} = 0.046$, $h = 0.69$, $\sigma_8 = 0.839$ and $n_s = 0.967$. The Ω_i are the density parameters for matter (m), dark energy (Λ) and baryons (b). h is defined as $h \equiv H_0/100$, where H_0 is the Hubble constant. σ_8 is the amplitude of fluctuations on a scale of $8h^{-1}$ Mpc, and n_s the primordial index of scalar fluctuations.

6.3 Methods

We extract the substructure power spectrum from the simulations in two ways. We first do so by using the subhalo catalogs obtained using the SUBFIND algorithm [198] after applying a friends-of-friends (FoF) halo finder with linking length $b = 0.2$. This procedure yields positions for all the detected subhalos together with several subhalo properties, such as the mass, half-light radius, maximum circular velocity, etc. This method is particularly advantageous because it is easier to compare to theoretical predictions, since it closely matches the notion of substructure in the halo model. Furthermore, it allows us to study novel properties of the convergence power spectrum, such as the contribution and detectability of different mass bins. We also extract the power spectrum directly from the particle data in the simulation snapshots (which we shall henceforth refer to as simulation snapshots for brevity) of the zoomed-in MW-like halo. The advantage of this method is that we do not impose any notion of how a subhalo is defined, meaning that all substructure within the simulated galactic halo is captured.

In our fiducial analysis we use the simulation snapshot (and its derived subhalo catalog) at $z = 0.5$ (a typical redshift for a lens galaxy), and place the background source

at $z = 1.5$. With the cosmological parameter values used in the simulations this yields a critical density for lensing $\Sigma_{\rm crit} = 2.35 \times 10^9$ M$_\odot$/kpc^2.

Figure 6.1 shows the projected density field obtained from the simulation snapshots (top two panels) and built from the subhalo catalogs (bottom two panels) with no mass/resolution threshold imposed. In the middle two panels we have superimposed a host profile on the convergence field obtained from the subhalo catalogs, to serve as a direct comparison to the simulation snapshots. See the section below for details on how this was done. Note that, despite the fact that Milky Way-like halos are generally less massive than typical galaxy-scale gravitational lenses, the two halos we are considering here are not far from being critical, with their convergence fields peaking around 0.3.

6.3.1 Power spectra from subhalo catalogs

We first extract the three-dimensional (3D) subhalo positions from the subhalo catalogs. We only keep subhalos within a comoving cube with side $L = 300$ kpc centered on the main lens galaxy, and those that have more than 50 particles, which corresponds to a minimum mass of 1.38×10^6 M$_\odot$. In our fiducial case we limit the highest subhalo mass to 10^8 M$_\odot$, since direct detection methods are expected to be able to detect subhalos above this mass in strong lensing images [107, 111, 199]. For completeness, we will also display power spectra that include these more massive subhalos.

To emulate ensemble averaging we project the 3D positions onto $N_{\rm proj}$ different lens planes, which replicates observing different lines-of-sight. We thus end up with $N_{\rm proj}$ two-dimensional (2D) maps of projected positions $\{\mathbf{H}_p\}$, where the index p reflects which projection the map corresponds to. We emphasize that considering $N_{\rm proj}$ different projections of the same galaxy can underestimate the variance of the power spectrum: Ref. [200] compared the variance with 1000 projections of a same subhalo population and that of 1000 independent subhalo populations, and found that the latter was significantly larger. They did however find that the difference between both scenarios was much smaller when the largest subhalos are removed (they imposed $m_{\rm high} = 10^9$ M$_\odot$), meaning that for our fiducial case we don't expect to be underestimating the variance so drastically.

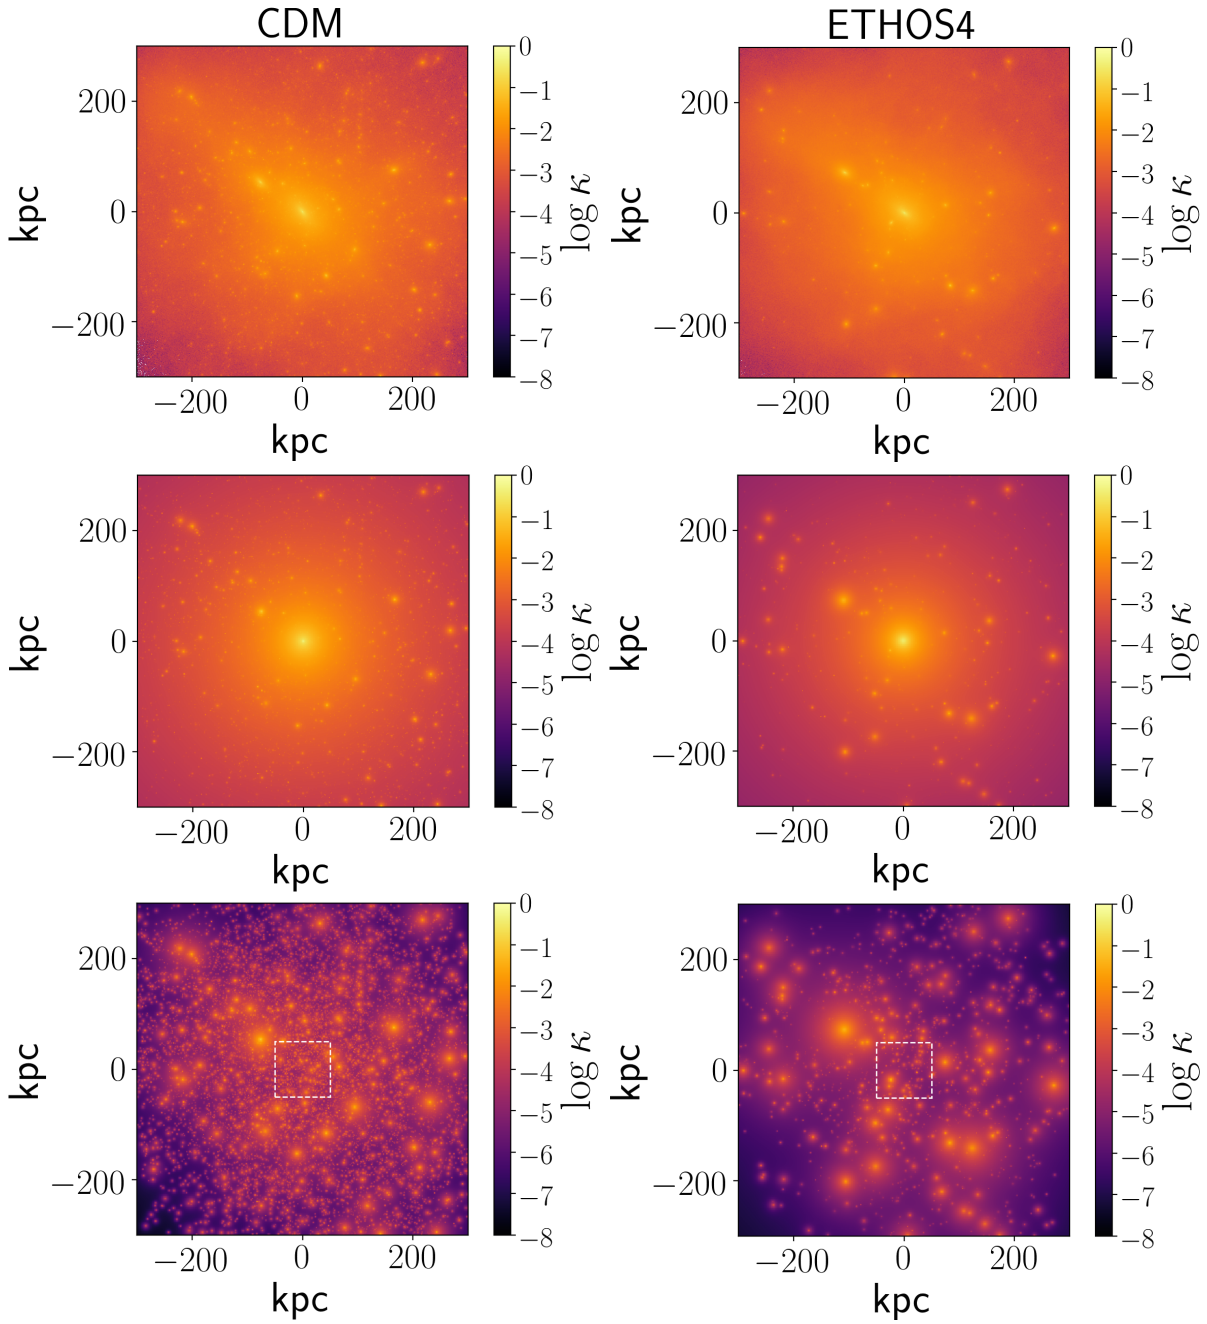

Figure 6.1: *Top left*: convergence field from the particle data for the CDM simulation at $z = 0.5$. *Top right*: convergence field from the particle data for the ETHOS4 simulation at $z = 0.5$. *Middle left*: convergence field from the subhalo catalog for the CDM simulation at $z = 0.5$ with a truncated NFW fit to the host superimposed. *Middle right:* convergence field from the subhalo catalog for the ETHOS4 simulation at $z = 0.5$ with a truncated Burkert fit to the host superimposed. *Bottom left*: convergence field from the subhalo catalog for the CDM simulation at $z = 0.5$. *Bottom right*: convergence field from the subhalo catalog for the ETHOS4 simulation at $z = 0.5$. The white square in the bottom two panels is centered at the origin and has a size of $L = 100$ kpc, therefore it represents the region under consideration in the fiducial case.

Subhalos in the CDM simulations are shown to be well fit by NFW profiles, so we fit a (truncated) NFW convergence profile to each subhalo in the projected map (see Appendix D.1). This profile is determined by three subhalo parameters: the total mass

m, the scale radius $r_{\rm s}$ and the tidal truncation radius $r_{\rm t}$. Note that truncating the NFW profile ensures that each subhalo has a total finite mass.

The subhalo finder assigns a gravitationally-bound mass to each subhalo, which we identify with the total mass parameter m of a truncated NFW subhalo. We obtain the scale radius of our subhalos using the well-known relation $r_{\rm max}/r_{\rm s} = 2.1626$ [201] for the NFW profile, where $r_{\rm max}$ is the radius at which the maximum circular velocity of the subhalo is attained, which the subhalo finder computes. We finally set the value of the tidal radius $r_{\rm t}$ by numerically solving the nonlinear relation $2\,m(< r_{\rm half}) = m$, where $r_{\rm half}$ is the radius containing half the subhalo mass, which is also reported by the subhalo finder.[2]

Due to the presence of dark matter self-interaction, the subhalos in the ETHOS4 simulation are instead fit with truncated Burkert profiles (see Appendix D.1), which can be fully specified by three parameters, namely the total mass m, the Burkert radius $r_{\rm b}$, and the tidal truncation radius $r_{\rm t}$. It is useful to write the Burkert radius as $r_{\rm b} = p\,r_{\rm s}$, where p defines the core size. We use a similar procedure as above to obtain the values of $r_{\rm s}$ and $r_{\rm t}$ from the subhalo catalog, fixing $p = 0.666$ to ensure that the standard kinematic relation $r_{\rm max}/r_{\rm s} = 2.1626$ is preserved. As a check of our calibration procedure, we compute the predicted values of $v_{\rm max}$ from our Burkert fits and compare those to the corresponding catalog entries, finding at most a 20% scatter between these values.

Notice that, although we have included all the subhalos within a cube with side $L = 300$ kpc, strong lensing cannot probe such a large area transverse to the line of sight (LOS). Therefore, after projecting we limit the box size to either $L = 100$ kpc, i.e. ± 50 kpc from the host center, or $L = 200$ kpc, depending on the scales we want to probe. Conversely, strong lens images do give us access to the entire LOS volume of the main halo, which is why it is important to first allow all the subhalos within the host to be projected before limiting the box size transverse to the LOS to a more realistic[3]

[2]We note that we could have simultaneously solved for both $r_{\rm s}$ and $r_{\rm t}$ using the nonlinear relations $2\,m(< r_{\rm half}) = m$ and $(dv^2/dr)|_{r_{\rm max}} = 0$, where v is the circular velocity profile of the subhalo. Our tests show that doing so leads to differences in the substructure power spectrum that are smaller than the scatter between different projections.

[3]We note that our projected area with sides of comoving length $L = 100$ kpc is still larger than a typical galaxy-scale strong lensing region. This allows us to capture the impact of subhalos that are on

observable size.

Applying this procedure to the 2D position maps $\{\mathbf{H}_p\}$, we obtain $N_{\rm proj}$ 2D convergence maps $\{\boldsymbol{\kappa}_p\}$, which we Fourier transform and square to obtain an estimate of the 2D power spectrum for each individual map $|\widetilde{\boldsymbol{\kappa}}_p(\mathbf{k})|^2$. A factor of $A_{\rm pix}^2/A_{\rm box}$ is necessary to normalize each power spectrum, where $A_{\rm pix}$ is the pixel area and $A_{\rm box}$ is the box area. An estimate of the monopole substructure power spectrum $P_{{\rm sub},p}(k)$ from the pth convergence map is finally computed by azimuthally averaging $|\widetilde{\boldsymbol{\kappa}}_p(\mathbf{k})|^2$,

$$P_{{\rm sub},p}(k) = \frac{1}{2\pi}\int_0^{2\pi} |\widetilde{\boldsymbol{\kappa}}_p(\mathbf{k})|^2 \, d\phi. \tag{6.5}$$

Repeating this procedure for our $N_{\rm proj}$ maps, we can compute the average substructure power spectrum $\bar{P}_{\rm sub}(k)$ as well as the 68th and 90th percentiles characterizing the distribution of power spectrum values at each wavenumber. We generally find that for a given k bin the power spectrum values are not Gaussian-distributed.

Carrying out this procedure we obtain the *total* subhalo power spectrum, including both the one- and two-subhalo contributions (see Eq. (6.1)). It is however possible to isolate the two-subhalo term by simply carrying out the procedure above directly from the position projections. For each map $\mathbf{H}_p$ we create a 2D map $\mathbf{S}$ of

$$\mathbf{S}_{p,j} = \frac{N_{p,j} - \bar{N}_p}{\bar{N}_p}, \tag{6.6}$$

where for the pth projection $N_{p,j}$ is the number of subhalos in the jth spatial pixel and $\bar{N}_p$ is the average number of subhalos per pixel. We can then follow the same procedure to Fourier transform and azimuthally average to obtain $P_{\rm ss}(k)$. The two-subhalo contribution can then be computed according to Eq. (6.3) given a choice of the subhalo convergence profile. It is important to avoid over-counting the subhalo clustering, since it contributes both in the $P_{\rm ss}(\mathbf{k})$ term and the $\langle\widetilde{\kappa}(\mathbf{k})\rangle$ term. To avoid this issue when isolating the two-subhalo term we randomize all the subhalo positions within a given projection before making the convergence maps.

the outskirts of the strong lensing region but can still influence the lensed images.

Finally, we want to point out that the smallest k mode accessible is determined by the box size as $k_{\rm min} = 2\pi/L_{\rm box}$, while the largest k mode accessible is determined by the pixel size, $k_{\rm max} = 2\pi/L_{\rm pix} = 2\pi N_{\rm pix}/L_{\rm box}$, where $N_{\rm pix}$ is the number of pixels on a given size of the box. Unless otherwise mentioned, we limit the box size to $L_{\rm box} = 100$ kpc (symmetrically centered about the host center). For computational efficiency we limit the image resolution to be 501×501 pixels. Thus, $k_{\rm min} \approx 0.06$ kpc^{-1} and $k_{\rm max} \approx 30$ kpc^{-1}.

6.3.2 Power spectra from simulation snapshots

The level-1 ETHOS simulations we use in this work contain almost half a billion particles, meaning that it can be quite costly to carry out this analysis at the N-body particle level. We use the publicly available code `nbodykit` [202] to perform parts of our analysis.[4] All its algorithms are parallel, which greatly expedites the analysis procedure.

Starting from particle catalogs, `nbodykit` can build a density mesh equal to $1 + \delta(\mathbf{x})$, meaning that to obtain the convergence field we have to rescale the mesh with factors of the average number density of particles $\bar{n}$, $\Sigma_{\rm crit}$, and the N-body particle mass $m_{\rm part}$:

$$\kappa(\mathbf{x}) = \frac{\bar{n}\ m_{\rm part}}{\Sigma_{\rm crit}}(1 + \delta(\mathbf{x})). \tag{6.7}$$

Much like our analysis based on the subhalo catalogs, we limit the particles out to 300 kpc from the host center, but we do not impose any resolution/mass thresholds for inclusion.

Unlike our catalog-based analysis where we were able to isolate the substructure contribution $\kappa_{\rm sub}$ in Eq. (5.2), we instead directly obtain the total convergence κ from the simulation snapshot. To isolate the substructure signal we are interested in, we therefore have to subtract the main host halo contribution κ_0. Our approach to remove this contribution consists of averaging many different projections to approximate the host profile,

$$\kappa_{\rm host}(\mathbf{r}) \approx \langle \kappa_{\rm box}(\mathbf{r}) \rangle, \tag{6.8}$$

[4] `nbodykit` is an open source large-scale structure toolkit written in Python.

and then subtracting this average map from a given projection p to obtain our estimate of the 2D substructure power spectrum

$$|\widetilde{\boldsymbol{\kappa}}_p(\mathbf{k})|^2 = \left[\int d^2\mathbf{r}\; e^{-i\mathbf{k}\cdot\mathbf{r}}\left(\kappa_{\mathrm{box},p}(\mathbf{r}) - \kappa_{\mathrm{host}}(\mathbf{r})\right)\right]^2, \tag{6.9}$$

before performing the angular averaging as in Eq. (6.5).

Due to the discrete nature of the simulation particles, we impose a conservative k_{max} cut beyond which we do not trust the results. For our choice of box size and $N_{\mathrm{mesh}} = 1024$, we impose $k_{\mathrm{max}} = 3\ \mathrm{kpc}^{-1}$.

6.4 Results

6.4.1 Subhalo catalog

Figure 6.2 shows the convergence power spectrum for the fiducial CDM (blue) and ETHOS4 (cyan) simulations for the larger box size, $L = 200$ kpc. This larger projected area allows us to be sensitive to the two-subhalo term on sufficiently large scales. For the CDM case, the two-subhalo term appears as an upturn in the power spectrum for $k \lesssim 0.06\ \mathrm{kpc}^{-1}$. To show that this upturn is indeed due to subhalo clustering we have overlaid the isolated two-subhalo contribution in dashed red, obtained with the method outlined in 6.3.1. As explained in Refs. [4, 181], this two-subhalo term corresponds to the so-called "host" contribution arising because all subhalos are gravitationally bound to their host galaxy. Of course, such small values of k are unobservable with the small field of view of typical strongly lensed images.

This figure has been made with 90 different projections. The solid lines correspond to the median and the shaded regions to the 68% and 90% confidence level areas. The vertical dashed line corresponds to the truncation wavenumber for the CDM simulation, defined as $k_{\mathrm{trunc}} \equiv 1/r_{\mathrm{t,max}}$. As predicted in Ref. [4], the break in the power spectrum is related to the size of the largest subhalos and the two-subhalo term becomes dominant below k_{trunc}. We can see that the amplitude of the one-subhalo term is well approximated

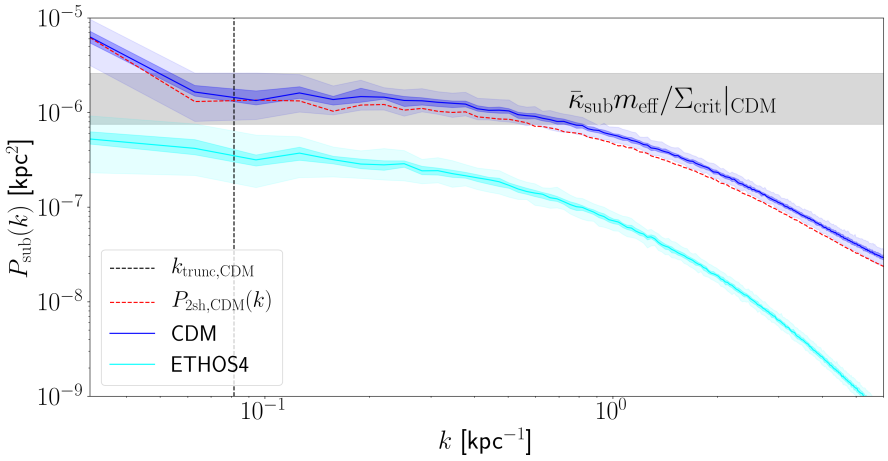

Figure 6.2: Substructure convergence power spectrum from the subhalo catalog at $z = 0.5$ and $m_{\rm high} = 10^8$ M$_\odot$ for both the CDM (blue) simulation and the ETHOS4 (cyan) simulation for a box with side $L = 200$ kpc. The shaded gray horizontal region shows the predicted amplitude from Ref. [4] with $\bar{\kappa}_{\rm sub}$ and $m_{\rm eff}$ (and their associated errors) obtained from the CDM simulations, and the vertical dashed line is the median $k_{\rm trunc} \equiv 1/r_{\rm trunc,max}$. The red dashed line is the isolated two-subhalo contribution for the CDM simulation. The wavenumbers k are in comoving coordinates.

by $\bar{\kappa}_{\rm sub} m_{\rm eff}/\Sigma_{\rm crit}$ (just as for $k_{\rm trunc}$ we only show this for the CDM simulation for clarity, but the same applies to the power spectrum obtained from the ETHOS4 simulation). The amplitude of the power spectrum is noticeably lower in ETHOS4, since there are many fewer subhalos. This dearth of small-mass subhalos is also responsible for the steeper slope at $k \gtrsim 1$ kpc^{-1} in ETHOS4. The power spectrum slope on these scales appears as a key observable that can probe the abundance of small-mass subhalos in lens galaxies. Finally, we can see that the two-subhalo term does not appear to contribute significantly to the ETHOS4 power spectrum on large scales. Indeed, the small overall number of subhalos in this case makes it difficult to probe the subhalo clustering signal.

In Figure 6.3 we show how the power spectrum shape and amplitude change as a function of redshift (top) and highest subhalo mass included (bottom) for the CDM (left) and ETHOS4 (right) simulations. The fiducial cases are kept in the same color as in Figure 6.2 (but notice that with $L = 100$ kpc the two-subhalo term is no longer clearly discernible in the CDM simulation). For the redshit evolution, we consider three different epochs: $z = \{0, 0.5, 1\}$. These redshifts correspond to the redshift of the simulation snapshot from which the subhalo catalog was obtained. For all cases, the source is assumed to be at $z = 1.5$.

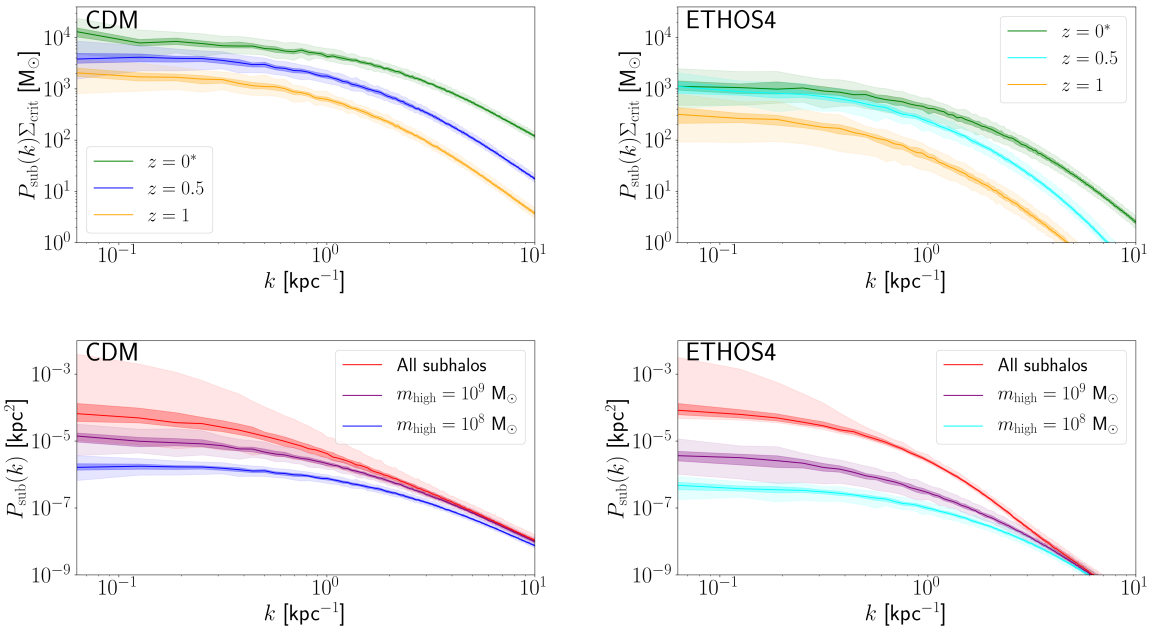

Figure 6.3: *Top Left*: redshift dependence of the convergence power spectrum times the critical density for the CDM simulation. *Top Right*: redshift dependence of the convergence power spectrum times the critical density for the ETHOS4 simulation. The critical density has been factored out to isolate the redshift evolution of the host. *Bottom Left*: mass dependence of the convergence power spectrum for the CDM simulation. *Bottom Right*: mass dependence of the convergence power spectrum for the ETHOS4 simulation. Note that the y-axis is the same for a given row but differs between rows. The wavenumbers k are in comoving coordinates. *As discussed in the text, the $z = 0$ power spectra are computed using the subhalo catalog at $z = 0$ but the distance between the observer and the lens $D_{\rm ol}$ is fixed to be the same as for a lens at $z = 0.5$ because $\Sigma_{\rm crit}$ diverges as $z \to 0$.

Since the convergence and the Einstein radius become ill-defined quantities as $z_{\rm lens} \to 0$, we artificially put our simulated $z = 0$ lens galaxy at a redshift $z_{\rm lens} = 0.5$ in order to compute their convergence field. In order words, we use the critical density for lensing $\Sigma_{\rm crit}$ corresponding to having a lens at $z = 0.5$ and source at $z = 1.5$ to compute the substructure convergence field of our simulated $z = 0$ galactic halo. For the other two epochs ($z = \{0.5, 1\}$), $\Sigma_{\rm crit}$ is computed self-consistently using the redshift of the simulated halo as the lens redshift.

Note, however, that to isolate the effect of the redshift evolution in the top two panels of Figure 6.3 we plot the product $P_{\rm sub}\Sigma_{\rm crit}$, since the value of the critical density changes between $z = 1$ and $z = 0.5$ by about a factor ~ 2. This means that any differences in the three curves are telling us something about subhalo accretion and evolution within the lens halo.

The redshift dependence shown in Figure 6.3 qualitatively agrees with what one would expect within the standard cosmological evolution: as we approach $z = 0$, more subhalos are accreted into the host halos, implying that the amplitude of the power spectrum increases. This increase is more pronounced in the CDM case as more subhalos with $m < 10^8$ $\mathrm{M}_\odot$ are accreted between $z = 0.5$ and $z = 0$ in this model. Also, as subhalos are accreted and move closer to the host center, mass loss due to tidal interaction becomes important. For the ETHOS4 simulation, we find that this leads on average to a reduction of the effective subhalo mass $m_{\rm eff}$ between $z = 0.5$ and $z = 0$, which partially compensates the slight increase in $\bar{\kappa}_{\rm sub}$ to leave the low-k amplitude nearly unchanged[5]. Furthermore, the much larger total number of subhalos in the CDM case also means that the two-subhalo term makes a non-negligible contribution at $z = 0$, which tends to increase the magnitude of the redshift evolution in this case. In contrast, the ETHOS4 model does not get a significant two-subhalo contribution at $z = 0$.

Another important aspect of the redshift evolution is the difference in the power spectrum slope for $k \gtrsim 2$ kpc^{-1}. Again, this effect is more apparent in the ETHOS4 simulation than in the CDM simulation due to the lower central densities of subhalos in

[5]Refer to Appendix D.2 for $m_{\rm eff}$ and $\bar{\kappa}_{\rm sub}$ values at $z = \{0, 0.5, 1\}$ for the two different simulation suites: Tables D.2.1, D.2.1, and D.2.1 for CDM; and Tables D.2.2, D.2.2, and D.2.2 for ETHOS4.

the former, making them more susceptible to tidal effects and mass loss. We indeed find that the substructure mass fraction in subhalos with $m < 10^7$ $M_\odot$ grows more rapidly between $z = 0.5$ and $z = 0$ in the ETHOS4 case compared to the CDM case, hence leading to a net transfer of power from larger to smaller scales in the power spectrum. This in turn results in a shallower slope for $k \gtrsim 2$ kpc^{-1} at $z = 0$ as compared to $z = 0.5$.

The three different upper mass thresholds we consider in the lower panels of Figure 6.3 are: $m_{\rm high} = \{10^8\ M_\odot, 10^9\ M_\odot\}$ and "All subhalos", where "All subhalos" means we include all subhalos above the resolution threshold. The behavior as a function of mass similarly shows the intuitive notion that, as we increase $m_{\rm high}$, the amplitude increases due to the fact that both $m_{\rm eff}$ and $N_{\rm sub}$ increase. The error bars are much larger for the case where all the subhalos are included because there are only a handful of subhalos with mass $> 10^9$ $M_\odot$, and they do not always get projected into the center-most region of the host. In the projections where even a single one of these subhalos is projected into the region of interest the amplitude is higher by about an order of magnitude. This shows that the low-k amplitude ($k \lesssim 1$ kpc^{-1}) is largely determined by the largest subhalos, as described in Refs. [4, 117, 200]. Note that for the "All subhalos" ETHOS4 power spectrum the lower bound is very small. This is simply due to the fact that, except for the rare cases when a very massive subhalo gets projected into the center-most region, the highest subhalo mass across projections is nearly constant for this simulation. On the other hand, for the CDM simulation the upper mass bound displays more variation, which is why the lower bound is larger.

A question that often comes to mind when discussing the substructure power spectrum is which mass scale is this observable most sensitive to. It is generally assumed that the largest subhalos within the strong lensing region dominate the observable power spectrum signal, since subhalos of higher mass generally warp images more. However, what we find here is more subtle. Figure 6.4 shows the decomposition of the dimensionless[6] convergence power spectrum into four different mass bins together with the power

[6]The dimensionless power spectrum is defined (in 2 dimensions) as usual:

$$\Delta^2_{\rm sub}(k) \equiv \frac{k^2 P_{\rm sub}(k)}{2\pi}.$$

spectrum including all subhalos with masses below 10^{10} $M_\odot$ for four different projections in the CDM simulation (one in each panel). Surprisingly, it can be seen that the $10^7 - 10^8$ $M_\odot$ subhalos dominate the signal almost entirely on scales 2 kpc$^{-1} \lesssim k \lesssim 15$ kpc^{-1}. Subhalos with masses between $10^9 - 10^{10}$ $M_\odot$ are quite rare, and in fact sometimes are not even present (e.g. lower right panel) in the strong lensing region. When present, they can of course dominate the signal at the lowest wavenumbers as discussed above, but they generally do not make the largest contribution to power spectrum on all observable scales. Another somewhat surprising element shown in Figure 6.4 is the relatively small contribution that the $10^8 - 10^9$ $M_\odot$ subhalos make to the overall power spectrum. Despite being quite numerous and fairly massive, they have a lesser contribution to the overall signal than their less massive counterparts, except possibly at the lowest wavenumbers.

A similar decomposition is done for the ETHOS4 simulations, and it is shown in Figure 6.5. It can be seen that the ETHOS4 projections display more variability than their CDM counterparts, due to the fact that there are many fewer subhalos. Even in this case, subhalos with mass $m < 10^8$ $M_\odot$ seem to make on average a sizable contribution on scales 2 kpc$^{-1} \lesssim k \lesssim 15$ kpc^{-1}.

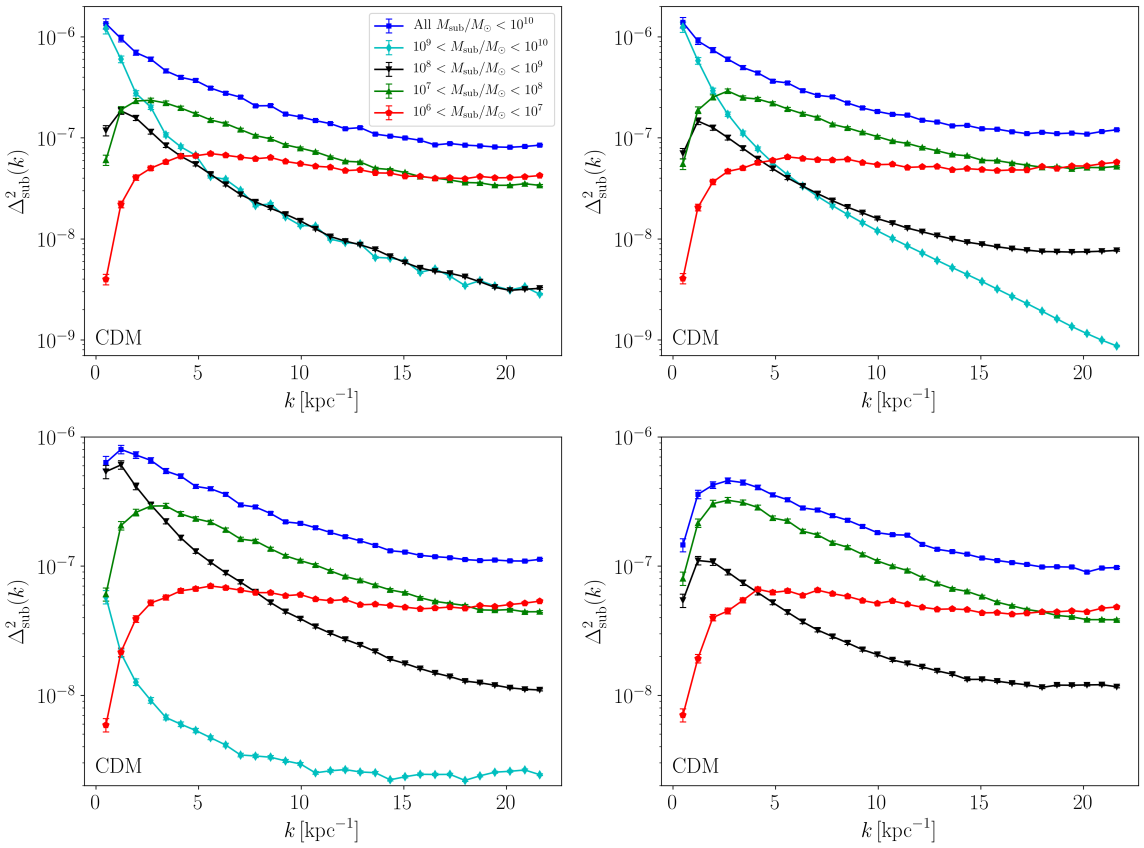

Figure 6.4: Decomposition of the CDM substructure power spectrum into its contributions from subhalos in different mass ranges. Note that the wavenumber axis is shown here on a linear scale. The four panels show different projections of the CDM subhalo populations. The blue squares show the substructure power spectrum including all subhalos with masses less than 10^{10} $\mathrm{M}_\odot$, while the other point types show the contributions from separate mass bins. We note that the contribution from the most massive subhalos included here (10^9 $\mathrm{M}_\odot < M_{\rm sub} < 10^{10}$ $\mathrm{M}_\odot$) varies significantly between different projections, with them making no contribution in the lower right panel.

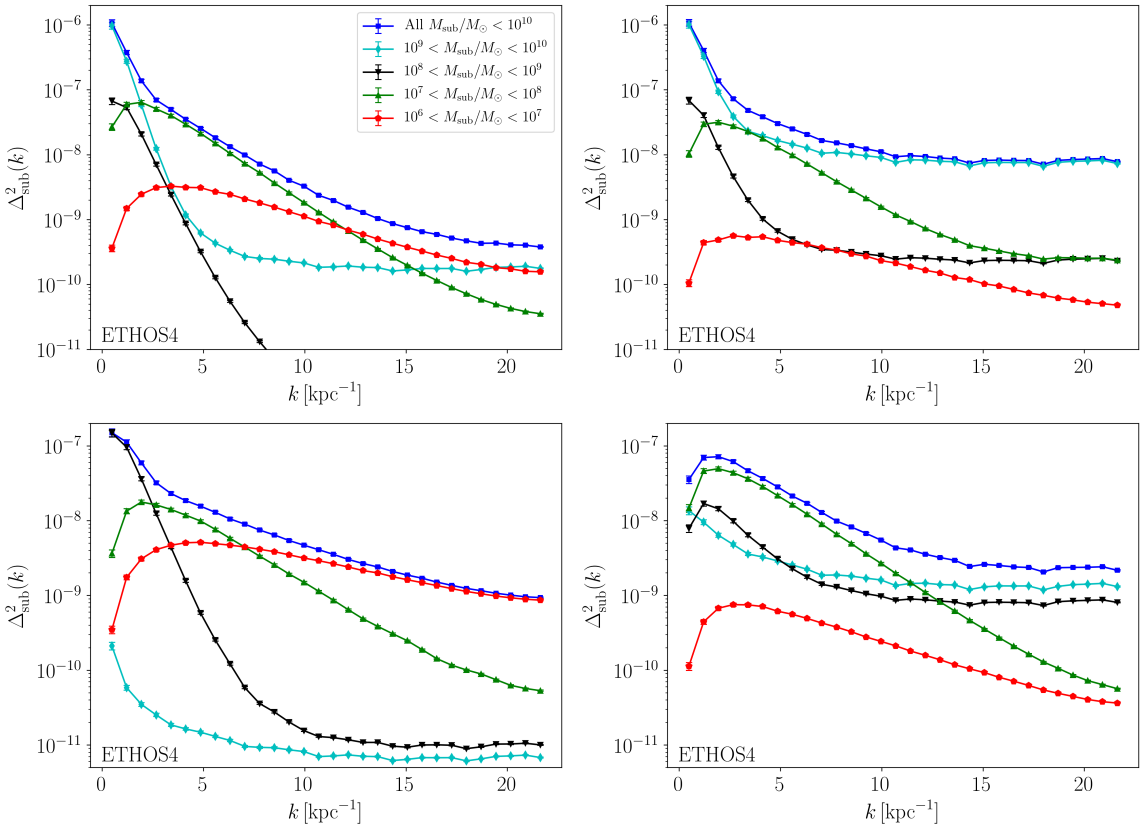

Figure 6.5: Same as Figure 6.4 but for the ETHOS4 simulation.

6.4.2 Snapshot particle data

The host convergence field obtained by averaging projections along different LOS in the CDM simulation is shown in Figure 6.6. The power spectra for both simulation suites are displayed in Figure 6.7. The top (bottom) panel corresponds to the CDM (ETHOS4) simulation. The solid blue line is the power spectrum obtained from a single projection of the N-body particles, without having performed any host subtraction. The dashed blue line is the power spectrum after removing the host contribution from a *single* projection map, thus approximating the power spectrum due to the substructure, as per Eq. (6.9). The green line is that of the average convergence map, i.e. approximately the host (the Fourier transform of Figure 6.6).

There are several notable features in these figures. First of all, notice the suppression in power at high k of the green lines compared to the solid blue lines, which shows that the averaging procedure is indeed removing the contribution from substructure on these scales. Furthermore, when the host is subtracted (dashed lines), a lot of power is lost

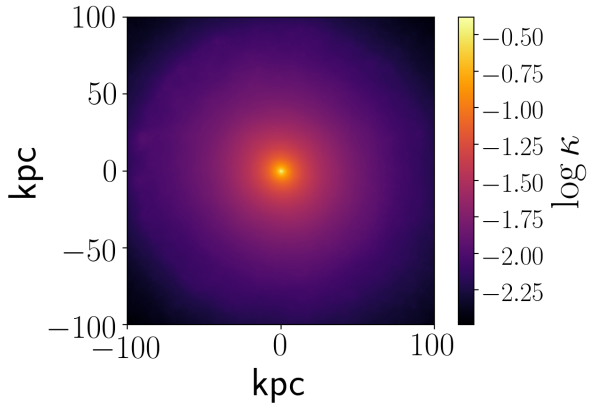

Figure 6.6: Convergence field of the host in the CDM simulation found by averaging many projections along different LOS, as per Eq. (6.8).

at low k but conversely we regain the power on small scales, which corresponds to the substructure convergence field remaining after the host is removed.

The overall amplitude is the same for both simulations, since both simulations have roughly the same number of particles in the region of interest. However, in the host and in the substructure power spectra at high k (> 1 kpc^{-1}) we can still see the suppression of power of ETHOS4 with respect to CDM due to the cutoff in the initial matter power spectrum (and to the self-interactions, albeit to a lesser extent), which causes both a suppressed number of small-mass subhalos and reduced central densities for the remaining ones.

In both figures we also overlay the catalog power spectra with no high mass cut (i.e. the same two red lines as in Figure 6.3). Unexpectedly, at high k the amplitude of the power spectra derived from the catalogs is higher than that obtained from the corresponding particle snapshot when looking at the CDM simulation (and comparable in the ETHOS4 simulation). We expected the opposite since when we Fourier transform the full simulation box we are capturing all the substructure (e.g.tidal debris), not just objects found by the halo finder. However this can be understood by considering the discrete nature of the simulation particles, and the fact that when we reach very small scales (i.e. around the scale radius of subhalos) there are in fact only a handful of particles. By instead imposing a smooth, truncated NFW profile at the catalog level we are artificially boosting the high-k signal with respect to the particle power spectrum.

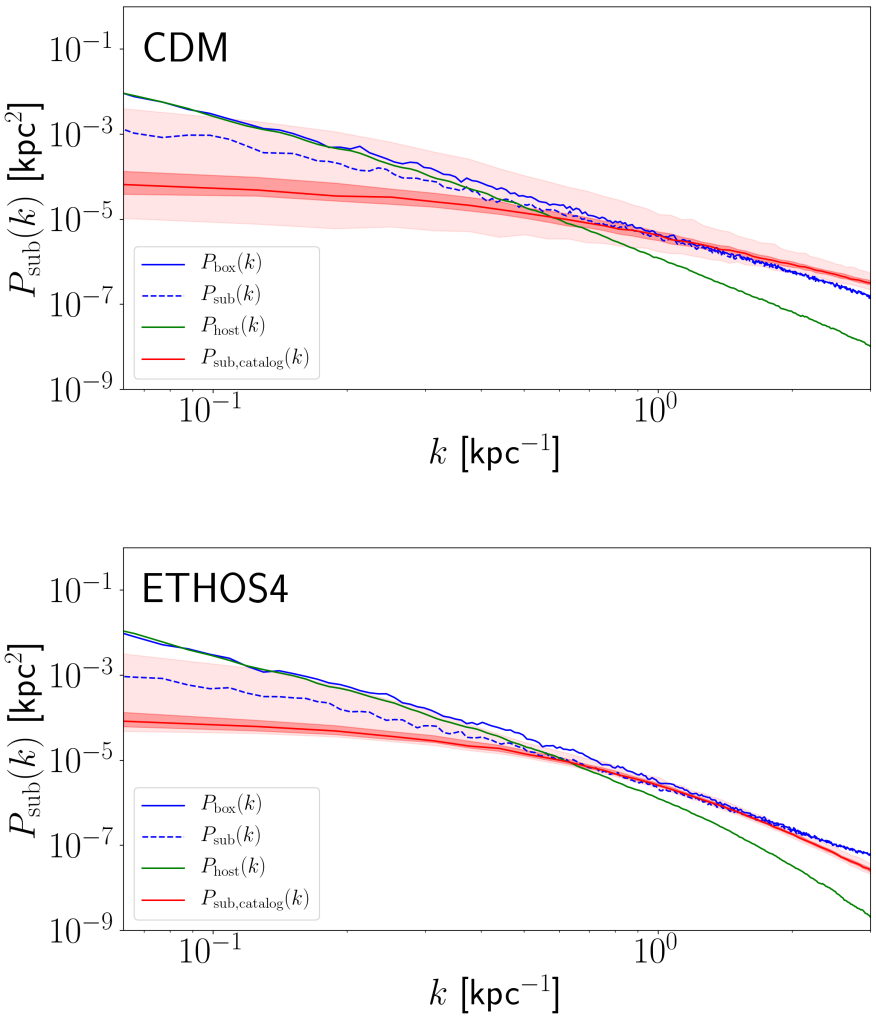

Figure 6.7: Power spectrum of the full simulation snapshots at $z = 0.5$. The CDM (ETHOS4) simulation is in the top (bottom). The solid blue line is the power spectrum of the full projected field. The green line is the power spectrum of many projections averaged together, which approximates that of the host halo (as per Eq. (6.8)). The dashed blue line is the power spectrum of a *single* projection with the average map subtracted, yielding the power spectrum due to substructure (as per Eq. (6.9)). In each plot we have overlaid the substructure power spectrum obtained from the catalogs in red, when all the subhalos are included (i.e. the same red lines as in the bottom two panels of Figure 6.3).

This effect is not quite as palpable in the ETHOS4 simulation since the truncated Burkert fit is cored in the central regions. By looking at Figure 6.1 and comparing the projections from the simulation snapshots and those built from the subhalo catalogs it becomes apparent that a part of this discrepancy might also be due to the loss of ellipticity when imposing spherically-symmetric convergence profiles.

Finally, notice that at small k the dashed lines lie within the 90% confidence band of the catalog power spectrum. This is indicative of the fact that the very large substructure in the lens is well captured by the halo finder, and since said structure dominates the amplitude (as shown in the previous section), the particle and catalog power spectra are similar on these scales. Note that because the dashed blue line is the power spectrum obtained from a single projection we do not necessarily expect it to match the solid red line. Nevertheless, it is also possible that there is an additional contribution coming from the residual host halo profile, since we are approximating the host as being spherically symmetric (Figure 6.6) but in the top panels of Figure 6.1 the host appears to have some ellipticity.

6.5 Discussion and Conclusion

In this chapter we have provided a comprehensive study of the substructure convergence power spectrum in N-body simulations. By comparing this observable in two simulations within the ETHOS framework [60] that differ in their treatment of dark matter micro-physics (both at early and late times, through dark matter-dark radiation interactions in the former and DM-DM self-interactions in the latter) we have been able to identify different ways in which details in the particle nature of dark matter can come to light through this observable.

We chose to carry out our analysis in two different, but complimentary, ways. On one hand, we have an idealized scenario in which all substructure is perfectly spherical and identifiable, and can be fit with simple density profiles. On the other hand, we have a scenario in which there are no assumptions or definitions built into what is considered

to be a subhalo − instead we capture all the structure within the host halo. The former method of course benefits from its simplicity: it allows us to clearly disentangle different subhalo properties and their impact on the shape and amplitude of the power spectrum. The latter, however, more closely approximates reality, where one cannot choose what perturbs an image or an arc, and has to carefully think about how to account for the mass distribution of the main lens galaxy itself.

By comparing the power spectrum obtained directly from the N-body particles with that from the subhalo catalogs we saw that the halo model-based computation (as used in Ref. [4]) is in fact an excellent approximation to the more detailed density field (as captured by the simulation snapshots). As we mentioned above, this does break down at high k, where the imposition of a smooth convergence profile leads to an overestimation with respect to the N-body particle power spectrum on those same scales due to the finite spatial resolution of the mass particles. On intermediate scales, the difference between the catalog and particle power spectra is well within forecasted error bars for the convergence power spectrum [120]. This result lends weight to the robustness of this observable to study substructure populations at cosmological distances from the Milky Way.

Doing the catalog analysis we confirmed many of the properties outlined in Chapter 5 and brought to light several new ones. We were able to show how the amplitude and shape of the power spectrum are related to the abundance, sizes, and masses of subhalos. Furthermore, we showed the redshift evolution of the power spectrum, and saw a difference in the standard CDM vs. DM-DR+SIDM scenarios: in the former we observed an expected increase in the amplitude of the power spectrum as more substructure was accreted; conversely, in the latter, between $z = 0$ and $z = 0.5$ there was nearly no change in the amplitude in these two redshift bins. This was partly the result of the lower total number of subhalos accreted during that time span in ETHOS4, as well the higher susceptibility of ETHOS4 subhalos to tidal disruption which caused $m_{\rm eff}$ to shrink within the strong lensing region. A non-negligible two-subhalo contribution at $z = 0$ for CDM also helps explain the faster growth of the overall power spectrum amplitude in this latter case.

The other interesting effect that came to light when comparing the ETHOS4 redshift evolution between $z = 0$ and $z = 0.5$ was the difference in slope on scales $k \gtrsim 2$ kpc^{-1}, which reflected the changing subhalo mass function as the host evolves. Both of these effects that appear in the redshift evolution of ETHOS4 - the amplitude and the slope - offer exciting possibilities. In Ref. [4] the highest k values ($\gtrsim 100$ kpc^{-1}) were identified as the most interesting region in the power spectrum to study the particle nature of dark matter. This was unfortunate since it is unlikely that we will be able to measure modes past $k \sim 100$ kpc^{-1} in the near future, and baryonic structures of comparable sizes would interfere with the isolation of the dark matter power spectrum slope on these scales. However, here we have identified other ways of probing dark matter microphysics that involve scales that can in fact be probed with current and future observations ($0.1 \leq k/\text{kpc}^{-1} < 100$).

Furthermore, by looking into the mass decomposition of the power spectrum we found that it is not exclusively sensitive to the most massive subhalos. For instance, we found that the mass range $10^7 - 10^8$ M$_\odot$ tends to dominate the power spectrum on intermediate scales (2 kpc$^{-1} \lesssim k \lesssim 15$ kpc^{-1}), particularly in the CDM simulation. In more standard gravitational imaging searches for substructure, sensitivity is assumed to be an increasing function of mass (and proximity to the images/arcs; see also e.g. Ref. [179] to see how other parameters, like concentration, can affect distortions). In this different, statistical approach we can see that this is no longer necessarily the case, and observations could probe lower masses in the highly coveted subhalo mass function.

Our results highlight several important issues in the quest to constrain the particle nature of dark matter with strong gravitational lensing. First, it confirms the statement that statistical detection methods are sensitive to an unresolvable but numerous population of subhalos. Furthermore, they show that that combining different lenses in order to boost the signal-to-noise of a substructure power spectrum measurement is highly non-trivial. Indeed, detailed models for the redshift evolution of the subhalo population for different host properties would have to be included in the fit. The other side of that coin is that when (if) observations are good enough to measure the power spectrum with a

single lens, comparing the low-k amplitude with lenses at different redshifts can serve as a diagnostic tool for dark matter deviations from standard CDM, be it an effect at early time that is imprinted on the initial power spectrum and consequently delays structure formation, or an effect at late times, like self-interactions that are strong (or weak but inelastic [203]) enough to cause subhalo stripping and/or disruption. One would of course also have to consider how the presence of baryons can disrupt substructure. Carrying this analysis out with a hydrodynamical simulation would allow for more robust predictions that can be compared with observations, since we know that on these scales baryonic processes can have quite significant effects on the dark matter distribution. Some work has been carried out to study the difference in distortions due to a population of globular clusters versus dark matter subhalos [187], showing that milliarcsecond resolution images could distinguish between these in direct detection efforts. But the impact on the power spectrum has yet to be addressed.

We want to point out that the MW-like halos we have considered in this work are not typical lens galaxies (recall that both halos are sub-critical, i.e. $\kappa < 1$). Gravitational lenses at cosmological distances from our own galaxy in general have to be more massive and dense to act as strong lenses (see, e.g., Ref. [204]). Such galaxies are expected to have more substructure (since substructure content scales with host mass). On the other hand, the increased central density in more typical strong lenses could increase subhalo tidal disruption [174]. Quantifying these effects as a function of host halo density is an interesting future step in understanding this observable.

We conclude this chapter by emphasizing that these statements are only strictly true if the main contribution to the convergence power spectrum indeed comes from small dark matter halos within the lens galaxy (what we have called substructure or subhalos), which is a key assumption we've made. If it instead comes from field halos along the line of sight, the conclusions to be drawn are different. We address this crucial distinction in Chapter 7.

Chapter 7

Quantifying the Line-of-Sight Halo Contribution to the Dark Matter Convergence Power Spectrum from Strong Gravitational Lenses

In the previous four chapters we have discussed extensively how to extract and how to interpret information pertinent to dark matter science on small scales from galaxy-galaxy strong lensing systems. Despite showcasing different approaches, all four chapters were based on the same fundamental premise that perturbations to the strong lens images are due to substructure/subhalos, i.e. dark matter halos gravitationally bound to the lens galaxy.

However, it has been claimed that a large (most likely larger) contribution to the perturbations in strong lenses comes from free dark matter clumps along the line of sight (LOS) [155]. These halos are commonly referred to as LOS halos or interlopers. This claim is not surprising nor hard to believe. The entire line-of-sight volume, comprised of a double-cone with base at the lens and ends at the observer and the source, is huge (far larger than the LOS volume spanned by the lens, which is why generally the thin-lens approximation is used in single-plane lensing). Their contribution to lensing observables

is an active area of study [156, 169, 205, 206] and recent analyses of strong gravitational lenses have begun to take it into account when placing constraints on ΛCDM [21, 140, 207–211].

In fact, it is crucial to take the LOS contribution into account before making any claim about dark matter; otherwise, we risk wrongfully falsifying or reinforcing the standard ΛCDM scenario. If, as expected, the contribution of interlopers really is greater than that of substructure, this could be good news for the ability of lensing observations to constrain the properties of dark matter: while subhalos are subject to messy, ill-understood processes as they travel through the main lens halo (such as tidal disruption), by virtue of being in the field, LOS interlopers are much less subject to environmental effects that might cause them to be disrupted.

Let us present a toy example that illustrates both the benefit gained from having the LOS contribution supersede that of subhalos and also how failing to take into account the LOS contribution can bias any inference about the particle nature of dark matter. Let us assume that the cold dark matter (CDM) paradigm really is the true dark matter model in our universe. In this scenario, the subhalo mass function rises steeply at the low-mass end and we expect a very high number of subhalos. However, if by traveling within their host's halo a large number of them are tidally disrupted, effectively disappearing, the observable number of subhalos might be a lot smaller than the expected number of subhalos under the CDM assumption. If the subhalo contribution is dominant, so we only consider subhalos, we might wrongfully falsify CDM if we do not observe a certain number of subhalos in a given mass range: for example, we may attribute the lack of halos to warm or self-interacting dark matter. If the LOS contribution really is dominant, then the lack of detection of halos in a given mass range is a much more faithful reflection of the fact that there may be some exotic dark matter physics reducing the number of halos with respect to the CDM expectation.

The importance of distinguishing between the subhalo and the interloper contributions

can be succinctly illustrated with the following schematic equation:

$$SMF = HMF(M < M_{\rm host}) * f(M_{\rm host}, z_{\rm host}, \{m_{\rm sub}\}, \{z_{\rm acc}\}, M_*, SFR...), \tag{7.1}$$

where SMF stands for subhalo mass function, HMF for halo mass function, and the second term on the right hand side represents a function of the host lens mass ($M_{\rm host}$), the host redshift ($z_{\rm host}$), the subhalo population masses ($\{m_{\rm sub}\}$), and the redshifts at which they were accreted into the host ($\{z_{\rm acc}\}$), the stellar mass in the host (M_*), and the star formation rate (SFR), among others. What this equation represents is the crucial insight that *the low-mass end of the HMF is not equal to the subhalo mass function.* The SMF is a processed version of the HMF, modified by a complicated interplay between host properties, subhalo properties, and astrophysical/baryonic effects. This is relevant to the convergence power spectrum because if it is dominated by substructure, it probes the subhalo mass function, while if it is dominated by interlopers, it probes the halo mass function. Clearly the latter is a much cleaner probe of the underlying dark matter theory than the former.

In this chapter, we develop the formalism of the LOS contribution to the convergence power spectrum, in contrast to the substructure contribution analyzed in the previous two chapters. Because previous analyses have neglected the LOS contribution, some features that have been deemed significant in past works may not be if the interlopers are included in the analysis. Furthermore, the dominance of the interloper contribution could facilitate deriving constraints on DM from the convergence power spectrum.

This chapter is organized as follows. In Section 7.1, we quantify the LOS contribution to the convergence power spectrum analytically by deriving an effective convergence for the LOS halos. In Section 7.2, we quantify the same contribution numerically by simulating a multi-plane lens system populated by LOS halos, then solving the multi-plane lens equation without any approximation. We discuss our findings and conclude in Section 7.3.

We shall refer to the halo that dominates the strong lensing as the *main lens*, the LOS halos as *interlopers* and the halos within the main lens as *subhalos* throughout the rest

of this chapter. When we are agnostic to whether a perturbation is due to subhalos or interlopers we refer to them as *perturbers*. Since we incorporate the LOS contribution to the power spectrum formalism, we will exclusively refer to this observable as the (effective) convergence power spectrum, instead of the substructure power spectrum, throughout the remainder of this chapter. We assume flat ΛCDM cosmology when calculating distances and the halo mass function.

7.1 Analytical Calculation

In the case where perturbations to strong-lens images are assumed to be caused by subhalos, the relative length scales in the problem are the physical size of the main lens along the line of sight and the distance traveled by light rays from the source to the observer. Obviously, the former is orders of magnitude smaller than the latter. Therefore, all the mass that is doing the lensing can be thought of as being on a single thin-lens plane (aptly called the thin-lens approximation). The convergence power spectrum calculation in this case is relatively straightforward since the convergence field is well defined. The addition of interlopers complicates the calculation since there is no well-defined convergence for a case with multiple consecutive thin-lens planes where each one deflects the light rays before they go onto the next plane. The angular deflections are not only added as vector fields but also are coupled to each other.

In this section we circumvent these problems by defining an effective convergence for a special case with a massive main lens coupled to low-mass interlopers. We will first go over some fundamentals of multi-plane lensing before deriving this effective convergence and, ultimately, arriving at expressions for the convergence power spectrum in the presence of interlopers in front of and behind the main lens.

7.1.1 Multi-plane Lens Equation

We model the main lens and interlopers as N consecutive thin-lens planes at redshifts z_i, where $i = 1, 2, ..., l, ..., N$ and $i > j$ implies $z_i > z_j$ (see Fig. 7.1). The main-lens plane

is indexed by l and the source plane by $s = N + 1$.

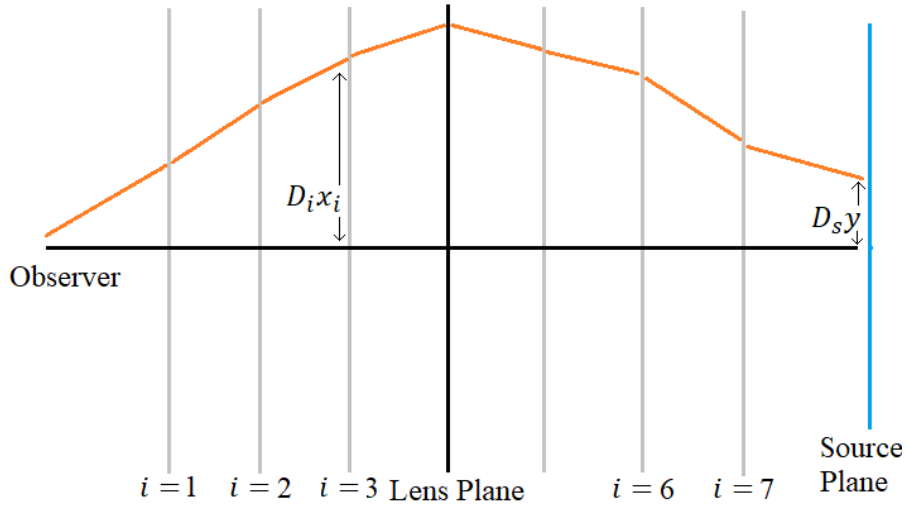

Figure 7.1: A simplified 2D sketch of the strong-lens system with interlopers. This is an example of a multi-lens plane system with $N = 7$ and $l = 4$. The orange line depicts the path that a single light ray travels from the source to the observer. The vertical gray lines correspond to the interloper planes, the vertical black line corresponds to the main-lens plane, and the vertical blue line to the source plane. $D_i x_i$ is the physical distance between the point where the light ray intersects the ith plane and the main axis that connects the center of the source plane to the observer. $D_s y = D_{N+1} x_{N+1}$ is the physical distance between the origin of the light ray and the center of the source plane.

Let us start with the multi-plane lens equation,

$$\vec{y} = \vec{x}_1 - \sum_{i=1}^{N} \vec{\alpha}_i(\vec{x}_i), \tag{7.2}$$

where $\vec{y}$, $\vec{x}_1 \in \mathbb{R}^2$ are the positions on the source plane and image plane, respectively. $\vec{x}_i$ is the angular position and $\vec{\alpha}_i$ the *deflection angle* of the light ray at the ith lens plane. Recall that, in strong gravitational lensing, the deflection angle is given by

$$\vec{\alpha}_i(\vec{x}_i) = \frac{1}{\pi} \int_{\mathbb{R}^2} d^2\vec{x} \frac{\vec{x}_i - \vec{x}}{|\vec{x}_i - \vec{x}|^2} \kappa_i(\vec{x}). \tag{7.3}$$

κ_i is the *convergence* of lens plane i, defined as the *projected mass density* Σ_i weighted by the *critical surface density* $\Sigma_{\text{crit},i}$,

$$\kappa_i(\vec{x}) \equiv \frac{\Sigma_i(D_i\vec{x})}{\Sigma_{\mathrm{crit},i}}, \tag{7.4}$$

$$\Sigma_{\mathrm{crit},i} \equiv \frac{c^2 D_s}{4\pi G D_i D_{is}}, \tag{7.5}$$

where c is the speed of light and G is the gravitational constant. D_i and D_{ij} are the angular diameter distances from the observer to the lens plane i and from lens plane i to lens plane j, respectively. For each lens plane, the derivative of the deflection angle can be written as

$$\frac{\partial\vec{\alpha}_i(\vec{x}_i)}{\partial\vec{x}_i} = \begin{pmatrix} \kappa_i + \gamma_{1,i} & \gamma_{2,i} \\ \gamma_{2,i} & \kappa_i - \gamma_{1,i} \end{pmatrix} (\vec{x}_i), \tag{7.6}$$

which is a symmetric matrix because the deflection angle at each lens plane is the gradient of the lensing potential, which is a scalar function [212]. Here γ_1 and γ_2 are shear terms that determine the amount that the image is elliptically distorted.

The intermediate lens-plane positions can be obtained by the following recursive equation:

$$\vec{x}_j = \vec{x}_1 - \sum_{i=1}^{j-1} \beta_{ij}\vec{\alpha}_i(\vec{x}_i), \quad \text{where} \quad \beta_{ij} \equiv \frac{D_{ij}D_s}{D_j D_{is}}. \tag{7.7}$$

Note that for $j = s$, $\vec{x}_s = \vec{y}$, and $\beta_{is} = 1$, we recover Eq. (7.2).

7.1.2 Effective Convergence for Interlopers

We want to define a single *effective convergence* that gives rise to images that are indistinguishable from those from a system with a main lens and interlopers. In general, angular deflections of multiple lens planes cannot be recreated by a single lens plane, so we will need to make some approximations in order to do so. In Section 7.2.2, we will show that the error introduced by these approximations is minimal.

We consider interlopers that are multiple orders of magnitude less massive than the main lens (an interloper with mass comparable to the main lens would distort the images enough to be detected and modeled directly). In strong-lensing systems, images are

formed near the Einstein radius θ_E of the main lens. We therefore use the Critical Sheet Born (CSB) approximation [213], where the intermediate lens-plane positions are assumed to be

$$\vec{x}_i = \begin{cases} \vec{x}_1 & z_i < z_l \\ (1-\beta_{li})\,\vec{x}_1 & z_i > z_l. \end{cases} \tag{7.8}$$

The underlying assumption is that the main lens does most of the lensing and the interlopers only slightly perturb the image. More specifically, first it is assumed that the light follows a path that is only deflected by the main lens. Then, the gravitational pull of an interloper on the light is integrated over this path as a perturbation. These perturbations are added to the angular deflections caused only by the main lens to get the final angular deflections.

We take the derivative of the multi-plane lens equation (Eq. 7.2),

$$\frac{\partial \vec{y}}{\partial \vec{x}_1} = I - \underbrace{\sum_{i=1}^{l-1} \frac{\partial \vec{\alpha}_i(\vec{x}_i)}{\partial \vec{x}_1}}_{\text{foreground}} \underbrace{-\frac{\partial \vec{\alpha}_l(\vec{x}_l)}{\partial \vec{x}_1}}_{\text{main-lens coupling}} \underbrace{-\sum_{i=l+1}^{s-1} \frac{\partial \vec{\alpha}_i(\vec{x}_i)}{\partial \vec{x}_1}}_{\text{background}}, \tag{7.9}$$

and define the effective convergence as

$$\frac{\partial \vec{\alpha}(\vec{x}_1)}{\partial \vec{x}_1} = \begin{pmatrix} \kappa_{\text{eff}} + \gamma_{1,\text{eff}} & \gamma_{2,\text{eff}} \\ \gamma_{2,\text{eff}} & \kappa_{\text{eff}} - \gamma_{1,\text{eff}} \end{pmatrix} (\vec{x}_1), \tag{7.10}$$

where

$$\vec{\alpha}(\vec{x}_1) \equiv \sum_{i=1}^{s-1} \vec{\alpha}_i(\vec{x}_i) \tag{7.11}$$

is the total deflection angle. In Eq. (7.9) we decompose the sum over lens planes of Eq. (7.2) into three separate terms: one for the foreground interlopers, one for the coupling to the main lens, and one for the background interlopers. Due to the recursive nature of the multi-plane lensing equation, and thus the different effect that interlopers have whether

they are in front of or behind the main lens, these terms will each lead to different effective convergences. In the remainder of this subsection we consider each term separately.

Foreground Interlopers

We derive the effective convergence for foreground interlopers by rewriting the second term on the RHS of Eq. (7.9) using the CSB approximation (Eq. 7.8):

$$\begin{aligned}\sum_{i=1}^{l-1} \frac{\partial \vec{\alpha}_i(\vec{x}_i)}{\partial \vec{x}_1} &= \sum_{i=1}^{l-1} \frac{\partial \vec{\alpha}_i(\vec{x}_1)}{\partial \vec{x}_1} \\ &= \sum_{i=1}^{l-1} \begin{pmatrix} \kappa_i + \gamma_{1,i} & \gamma_{2,i} \\ \gamma_{2,i} & \kappa_i - \gamma_{1,i} \end{pmatrix} (\vec{x}_1). \end{aligned} \tag{7.12}$$

Thus, the effective convergence for foreground halos is simply the sum of each interloper's convergence up to the main lens:

$$\kappa_{\text{eff,fg}}(\vec{x}_1) = \sum_{i=1}^{l-1} \kappa_i(\vec{x}_1) = \sum_{i=1}^{l-1} \frac{\Sigma_i(D_i \vec{x}_1)}{\Sigma_{\text{crit},i}}. \tag{7.13}$$

By taking the continuum limit, we can write this as an integral over comoving distance χ,

$$\kappa_{\text{eff,fg}}(\vec{x}_1) = \int_0^{\chi_l} d\chi \, \frac{a(\chi)\bar{\rho}_m(\chi)\delta(\chi, \vec{x}_1)}{\Sigma_{\text{crit}}(\chi)}, \tag{7.14}$$

where a is the scale factor, $\bar{\rho}_m$ is the average matter density of the universe, δ is the overdensity, and χ_l is the comoving distance to the main-lens plane. We can rewrite this more concisely using the definition of the average matter density,

$$\bar{\rho}_{\text{m}} \equiv \frac{3H_0^2}{8\pi G} \frac{\Omega_{0,\text{m}}}{a^3}, \tag{7.15}$$

where H_0 is the Hubble constant and $\Omega_{0,\text{m}}$ the matter density parameter:

$$\kappa_{\text{eff,fg}}(\vec{x}_1) = C \int_0^{\chi_s} d\chi \, W_{\text{fg}}(\chi) \, \delta(\chi, \vec{x}_1), \tag{7.16}$$

with $C = 3\Omega_{0,\mathrm{m}} H_0^2/2c^2$ and

$$W_{\mathrm{fg}}(\chi) = \begin{cases} \dfrac{\chi}{a(\chi)}\left(1 - \dfrac{\chi}{\chi_s}\right) & \chi \leq \chi_l \\ \\ 0 & \chi > \chi_l. \end{cases} \tag{7.17}$$

Note that $W_{\mathrm{fg}}(\chi)$ plays the role of a selection function for the foreground interlopers.

Main-Lens Coupling

The third term on the RHS of Eq. (7.9) corresponds to the coupling between the foreground interlopers and the main lens. This term arises because a small angular deflection by a foreground interloper not only shifts the image, but also shifts the location where the light ray intersects with the main-lens plane, consequently altering the deflection caused by the main lens.

The angular position at the lens plane can be found evaluating Eq. (7.7) at $j = l$:

$$\begin{aligned} \frac{\partial \vec{x}_l}{\partial \vec{x}_1} &= I - \sum_{i=1}^{l-1} \beta_{il} \frac{\partial \vec{\alpha}_i(\vec{x}_1)}{\partial \vec{x}_1} \\ &= I - \sum_{i=1}^{l-1} \beta_{il} \begin{pmatrix} \kappa_i + \gamma_{1,i} & \gamma_{2,i} \\ \\ \gamma_{2,i} & \kappa_i - \gamma_{1,i} \end{pmatrix} (\vec{x}_1). \end{aligned} \tag{7.18}$$

We can thus see that the effective convergence corresponding to this coupling term is a weighted sum of the foreground interlopers:

$$\kappa_{\mathrm{eff,cp}}(\vec{x}_1) = -\sum_{i=1}^{l-1} \beta_{il} \kappa_i(\vec{x}_1). \tag{7.19}$$

Following the same procedure as in Section 7.1.2, we take the continuum limit and write this as an integral over the comoving distance,

$$\kappa_{\mathrm{eff,cp}}(\vec{x}_1) = C \int_0^{\chi_s} d\chi \, W_{\mathrm{cp}}(\chi) \, \delta(\chi, \vec{x}_1), \tag{7.20}$$

where the selection function is now given by

$$W_{\rm cp}(\chi) = \begin{cases} -\dfrac{\chi\beta_{\chi l}}{a(\chi)}\left(1-\dfrac{\chi}{\chi_s}\right) & \chi \le \chi_l \\ \\ 0 & \chi > \chi_l. \end{cases} \tag{7.21}$$

Background Interlopers

Finally, we derive an effective convergence for background interlopers by rewriting the last term on the RHS of Eq. (7.9),

$$\begin{aligned} \sum_{i=l+1}^{s-1} \frac{\partial\vec{\alpha}_i(\vec{x}_i)}{\partial\vec{x}_1} &= \sum_{i=l+1}^{s-1} \frac{\partial\vec{x}_i}{\partial\vec{x}_1}\frac{\partial\vec{\alpha}_i(\vec{x}_i)}{\partial\vec{x}_i} \\ &= \sum_{i=l+1}^{s-1} (1-\beta_{li})\frac{\partial\vec{\alpha}_i((1-\beta_{li})\,\vec{x}_1)}{\partial\vec{x}_i}, \end{aligned} \tag{7.22}$$

where on the last line we have used the CSB approximation. The effective convergence corresponding to the background interlopers is then

$$\kappa_{\rm eff,bg} = \sum_{i=l+1}^{s-1} (1-\beta_{li})\kappa_i((1-\beta_{li})\vec{x}_1). \tag{7.23}$$

In the continuum limit, we get

$$\kappa_{\rm eff,bg}(\vec{x}_1) = C\int_0^{\chi_s} d\chi\, W_{\rm bg}(\chi)\,\delta(\chi,(1-\beta_{l\chi})\vec{x}_1), \tag{7.24}$$

where the selection function is

$$W_{\rm bg}(\chi) = \begin{cases} 0 & \chi \le \chi_l \\ \\ \dfrac{\chi(1-\beta_{l\chi})}{a(\chi)}\left(1-\dfrac{\chi}{\chi_s}\right) & \chi > \chi_l. \end{cases} \tag{7.25}$$

Interlopers as effective subhalos

Combining these results, we see that for a strong-lensing system with some foreground and background perturbers we can write a single effective convergence that characterizes the effect of the interlopers as

$$\kappa_{\rm eff}(\vec{x}) = \underbrace{\sum_{i=1}^{l-1}(1-\beta_{il})\kappa_i(\vec{x})}_{\text{foreground + coupling}} + \underbrace{\sum_{i=l+1}^{N}(1-\beta_{li})\,\kappa_i((1-\beta_{li})\vec{x})}_{\text{background}}\,. \tag{7.26}$$

We can write this in the continuum limit as

$$\kappa_{\rm eff}(\vec{x}) = C\int_0^{\chi_s} d\chi\, W_{\rm I}(\chi)\delta(\chi, g(\chi)\vec{x}_1), \tag{7.27}$$

where

$$W_{\rm I} \equiv W_{\rm fg} + W_{\rm cp} + W_{\rm bg} = \frac{f(\chi)D_{\chi s}\chi^2}{D_\chi D_s}, \tag{7.28}$$

and $f(\chi)$ and $g(\chi)$ are piecewise functions of the comoving distance:

$$f(\chi) = \begin{cases} 1-\beta_{\chi l} & \chi \le \chi_l \\ 1-\beta_{l\chi} & \chi > \chi_l \end{cases} \tag{7.29}$$

$$g(\chi) = \begin{cases} 1 & \chi \le \chi_l \\ 1-\beta_{l\chi} & \chi > \chi_l. \end{cases} \tag{7.30}$$

We can think of Eq. (7.26) as a projection that takes interlopers at some plane i and projects them onto the lens plane l with an effective convergence. Under this approximation, we calculate the deflection of the interlopers assuming that the light ray travels a path that is only deflected by the main lens. Since β_{il} and β_{li} go to 0 as the distance between the planes i and l goes to 0, the interlopers that are sufficiently close to the main lens are unchanged by this projection. Furthermore, since β_{il} goes to 1

as the distance between plane i and the observer goes to 0, and β_{li} goes to 1 as the distance between plane i and the source goes to 0, the interlopers that are sufficiently close to the observer and the source become insignificant after projection. At intermediate positions, where neither of these approximations hold, one could imagine that the error introduced by this projection could be significant. We discuss this further and quantify the error in Section 7.2.2, showing that it is in fact an excellent approximation even in this intermediate regime.

For the remainder of this chapter, when we have to specify a density profile for a halo we will do so with a truncated NFW profile [160] (tNFW). The form of this profile will be shown explicitly later on in Eq. (7.45); for now, the relevant aspect of this profile is that it is fully determined by two parameters, the scale radius $r_{\rm s}$ and the dimensionless truncation parameter τ, defined in terms of the truncation radius $r_{\rm t}$: $\tau \equiv r_{\rm s}/r_{\rm t}$.

Eq. (7.26) implies that an interloper with mass m and convergence κ at comoving distance χ has an effective convergence

$$\kappa_{\chi,\rm eff}(\vec{x}\,;\, m, r_{\rm s}, \tau) = f(\chi)\, \kappa(g(\chi)\vec{x}\,;\, m, r_{\rm s}, \tau). \tag{7.31}$$

Therefore, we can think of the interlopers as subhalos on the main-lens plane with a modified scale radius and mass. To obtain these scaling relations, we express the convergence in terms of the projected mass density (Eq. 7.4),

$$\frac{\Sigma_{\chi,\rm eff}(D_l\vec{x}\,;\, m, r_{\rm s}, \tau)}{\Sigma_{{\rm crit},l}} = f(\chi)\frac{\Sigma(g(\chi)D_\chi\vec{x}\,;\, m, r_{\rm s}, \tau)}{\Sigma_{{\rm crit},\chi}}, \tag{7.32}$$

and make use of the following rules:

$$\epsilon\Sigma(\vec{r}\,;\, m, r_{\rm s}, \tau) = \Sigma(\vec{r}\,;\, \epsilon m, r_{\rm s}, \tau) \tag{7.33}$$

$$\Sigma(\eta\vec{r}\,;\, m, r_{\rm s}, \tau) = \Sigma(\vec{r}\,;\, \frac{m}{\eta^2}, \frac{r_{\rm s}}{\eta}, \tau), \tag{7.34}$$

where ϵ and η are scaling constants. These are derived in Appendix E.1. With these in

hand we can rewrite Eq. (7.32) as

$$\Sigma_{\chi,\mathrm{eff}}(D_l\vec{x}\,;\, m, r_\mathrm{s}, \tau) = \Sigma(D_l\vec{x}\,;\, m_\mathrm{eff}, r_{\mathrm{s,eff}}, \tau), \tag{7.35}$$

where

$$r_{\mathrm{s,eff}}(\chi) = \frac{D_l}{g(\chi)D_\chi} r_\mathrm{s} \tag{7.36}$$

and

$$m_\mathrm{eff}(\chi) = f(\chi)\frac{\Sigma_{\mathrm{crit},l}}{\Sigma_{\mathrm{crit},\chi}}\left(\frac{D_l}{g(\chi)D_\chi}\right)^2 m \tag{7.37}$$

are the effective scale radius and the effective mass for interlopers, respectively (shown in Fig. 7.2). The effective scale radius is larger than the true scale radius for both foreground and the background interlopers. The effective mass, on the other hand, is smaller for foreground and larger for background interlopers.

7.1.3 Power Spectrum of Interlopers

In this section, we calculate the power spectrum of the interlopers using two different approaches. In the first approach we use the Limber approximation [214], which neglects the Fourier modes of the 3D matter distribution parallel to the line of sight. For this, we will express the 2D line-of-sight density as an integral over the 3D density with a window function that weighs the integral over distance, as we did in Section 7.1.2. The second approach consists of replacing the interlopers with effective subhalos that live on the lens plane, using Eq. (7.31), and calculating the Fourier transform of their 2D two-point correlation function. The interlopers that are projected onto the lens plane will be within a volume with the shape of a double cone, shown in Fig. 7.3.

While the former approach might be more familiar to cosmologists due to its similarity to the weak lensing literature, the latter is more akin to previous works on the statistics of dark matter substructure in strong gravitational lenses. In particular, this calculation closely follows that of Ref. [4] for the subhalo power spectrum. The main difference is the addition of the comoving distance χ as a parameter that an interloper has as an effective

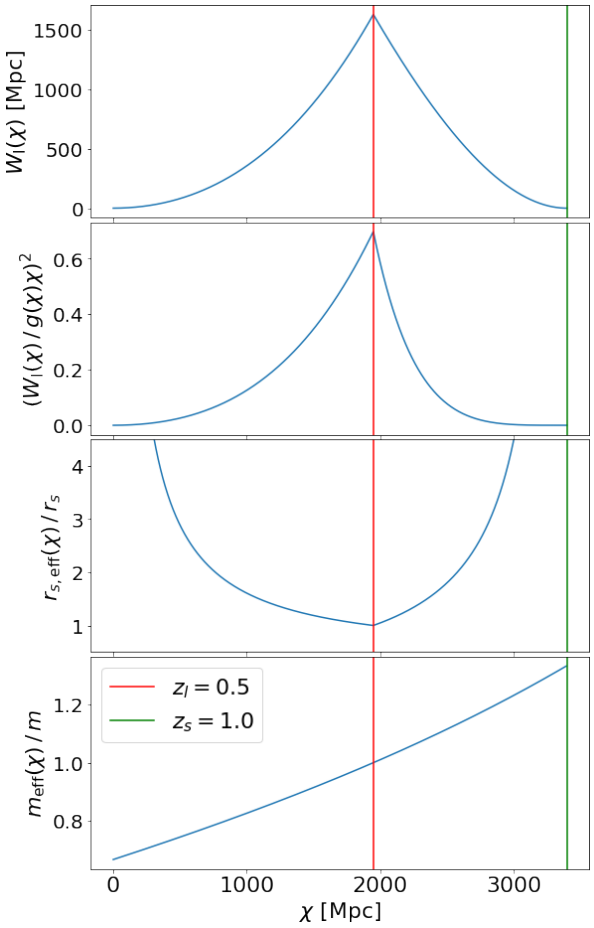

Figure 7.2: The comoving distance dependence of (from the top) **1:** the selection function W_I which weighs the overdensities at different distances to get the effective convergence in Eq. (7.27), **2:** $(W_\mathrm{I}/g(\chi)\chi)^2$ which weighs the 3D power to get the 2D power spectrum in Eq. (7.39), **3:** the ratio of effective scale radius to the scale radius of the interlopers (Eq. 7.36), **4:** the ratio of effective mass to the mass of the interlopers (Eq. 7.37). These functions also depend on the source and lens distances. Here, $z_l = 0.5$ and $z_s = 1.0$ for the lens and source redshifts, respectively. Their comoving distances are shown as red and green vertical lines.

subhalo. We will see that ultimately both approaches are in agreement.

Power Spectrum of the Interlopers using the Limber Approximation

For a 2D projected density that can be written in the form [215]

$$\delta_{\mathrm{2D}}(\vec{x}) = \int_0^{\chi_s} d\chi\, W(\chi)\delta_{\mathrm{3D}}(\chi, g(\chi)\chi\vec{x}), \tag{7.38}$$

where δ_{3D} is the 3D matter overdensity, $W(\chi)$ is the selection function that weighs distances, and $g(\chi)$ is a function that represents how the 3D density is projected down to

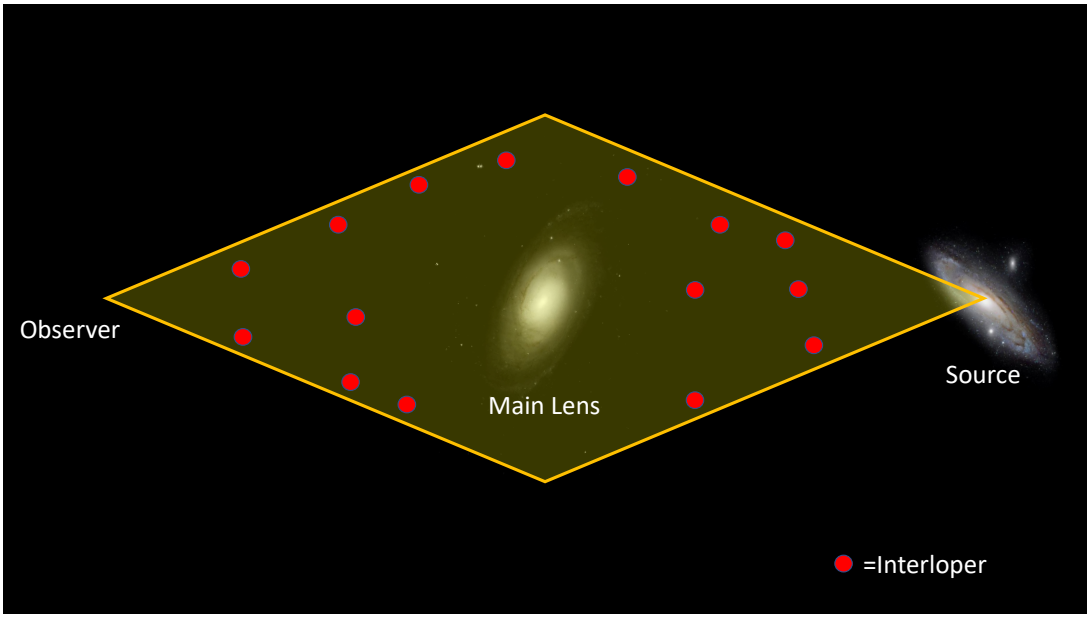

Figure 7.3: Double cone (Eq. E.17) volume integrated over for the analytical calculation of the effective convergence power spectrum due to interlopers (Eq. 7.67).

2D, the 2D power spectrum can be written in terms of the 3D power spectrum as

$$P_{\rm 2D}(\vec{l}) = \int_0^{\chi_s} d\chi \left(\frac{W(\chi)}{g(\chi)\chi}\right)^2 P_{\rm 3D}\left(\chi, \frac{\vec{l}}{g(\chi)\chi}\right), \tag{7.39}$$

where $\vec{l}$ is the Fourier conjugate to the angular position $\vec{x}$.

Combining the selection functions in Eqs. (7.17), (7.21), and (7.25), and noticing that $\vec{k} = \vec{l}/D_l$ at the lens plane, we can write the interloper power spectrum as

$$P_{\rm I}(\vec{k}) = C^2 D_l^2 \int_0^{\chi_s} d\chi \, \frac{W_{\rm I}^2(\chi)}{g^2(\chi)\chi^2} P_{\rm 3D}\left(\chi, \frac{\vec{k}D_l}{g(\chi)\chi}\right). \tag{7.40}$$

Our goal is to write this power spectrum as a function of the density profiles and ensemble properties of interlopers. Therefore, we will expand the 3D matter power spectrum in order to express it in terms of these quantities.

Let us write the 2-point correlation function of the matter overdensities,

$$\varepsilon(\vec{r}) = \frac{1}{V}\int d^3\vec{s}\,\delta(\vec{s})\delta(\vec{s}+\vec{r}), \tag{7.41}$$

where $\vec{s}$ and $\vec{r}$ are positions in comoving coordinates, and V is the comoving volume

over which these correlations are integrated. Writing the overdensities as the sum of overdensities of each individual interloper, $\delta(\vec{x}) = \sum_i \delta_i(\vec{x})$, we can rewrite the 2-point correlation function as the sum of the correlation of each interloper with itself and the correlation of each interloper with other interlopers,

$$\begin{aligned} \varepsilon(\vec{r}) &= \underbrace{\frac{1}{V}\sum_i \int d^3\vec{s}\,\delta_i(\vec{s})\delta_i(\vec{s}+\vec{r})}_{1-\text{halo term}} \\ &+ \underbrace{\frac{1}{V}\sum_i \sum_{j\neq i} \int d^3\vec{s}\,\delta_i(\vec{s})\delta_j(\vec{s}+\vec{r})}_{2-\text{halo term}}, \end{aligned} \tag{7.42}$$

where the summation is now over interlopers rather than lens planes.

We assume that the interloper positions are uncorrelated, which makes the 2-halo term vanish. Therefore, we can write the 3D matter power spectrum as the Fourier transform of the 1-halo term,

$$\begin{aligned} P_{\rm 3D}(\vec{k}) &\equiv \mathcal{F}[\varepsilon](\vec{k}) \qquad (7.43) \\ &= \frac{1}{V}\int d^3\vec{r}\,\exp[-i\vec{k}\cdot\vec{r}] \\ &\times \sum_i \int d^3\vec{s}\,\delta_i(\vec{s})\delta_i(\vec{s}+\vec{r}) \\ &= \frac{1}{V}\sum_i \left|\mathcal{F}[\delta_i](\vec{k})\right|^2. \qquad (7.44) \end{aligned}$$

For convenience, we denote the Fourier transform of a function as $\tilde{\delta}(\vec{k}) \equiv \mathcal{F}[\delta](\vec{k})$ in the remainder of the text.

We now assume the interlopers to be tNFW halos, which have a spherically symmetric density profile (in physical coordinates) given by

$$\rho(r;\, m, r_{\rm s}, \tau) = \frac{m}{4\pi r(r+r_{\rm s})^2 u(\tau)}\left(\frac{\tau^2 r_{\rm s}^2}{r^2+\tau^2 r_{\rm s}^2}\right), \tag{7.45}$$

where

$$u(\tau) \equiv \frac{\tau^2}{(\tau^2+1)^2}\left[(\tau^2-1)\ln(\tau) + \tau\pi - (\tau^2+1)\right]. \tag{7.46}$$

The density profile can be cast in dimensionless form by defining $\xi = r/r_{\rm s}$ and

$$\phi(\xi;\,\tau) \equiv \frac{1}{4\pi\xi(\xi+1)^2 u(\tau)}\left(\frac{\tau^2}{\xi^2+\tau^2}\right), \tag{7.47}$$

such that

$$\rho(r;\,m,r_{\rm s},\tau) = \frac{m}{r_{\rm s}^3}\phi\left(\frac{r}{r_{\rm s}};\,\tau\right). \tag{7.48}$$

Near halos, the overdensity is much larger than the average density $\bar{\rho}$, so we can write the overdensity due to a single tNFW interloper i as

$$\begin{aligned}\delta_i(s,\chi;\,m_i,r_{{\rm s},i},\tau_i) &= \frac{\rho(a(\chi)s;\,m_i,r_{{\rm s},i},\tau_i)-\bar{\rho}(\chi)}{\bar{\rho}(\chi)} \\ &\cong \frac{m_i}{r_{{\rm s},i}^3\bar{\rho}(\chi)}\phi\left(\frac{a(\chi)}{r_{{\rm s},i}}s;\,\tau_i\right),\end{aligned} \tag{7.49}$$

and its Fourier transform as

$$\tilde{\delta}(k,\chi;\,m_i,r_{{\rm s},i},\tau_i) = \frac{m_i}{\rho_0}\tilde{\phi}\left(\frac{r_{{\rm s},i}}{a(\chi)}k;\,\tau_i\right), \tag{7.50}$$

where $\rho_0 = \bar{\rho}(\chi=0)a^3(\chi=0)$. Spherical symmetry of the tNFW profile implies $\tilde{\delta}(\vec{k}) = \tilde{\delta}(k)$, where $k \equiv |\vec{k}|$.

Substituting this into Eq. (7.43) we obtain an expression for the 3D power spectrum in terms of the density profile of interlopers:

$$P_{\rm 3D}(\chi,k) = \frac{1}{V}\sum_i \frac{m_i^2}{\rho_0^2}\left|\tilde{\phi}\left(\frac{r_{{\rm s},i}}{a(\chi)}k;\,\tau_i\right)\right|^2. \tag{7.51}$$

Following the procedure in Ref. [4], we can convert this sum into an integral over the probability distribution of the interloper parameters,

$$\begin{aligned}P_{\rm 3D}(\chi,k) = \frac{N}{V}\int dm\, d^2\vec{q}\,\mathcal{P}(m,\vec{q}\,|\,\chi) \\ \times \frac{m^2}{\rho_0^2}\left|\tilde{\phi}\left(\frac{r_{\rm s}}{a(\chi)}k;\,\tau\right)\right|^2,\end{aligned} \tag{7.52}$$

where $\vec{q} = (r_s,\tau)$ are the intrinsic halo parameters, N is the total number of interlopers

in the ensemble, and $\mathcal{P}(m, \vec{q}\,|\,\chi)$ is the probability that an interloper has mass m and intrinsic parameters $\vec{q}$, given that it is at comoving distance χ. We now separate the probabilities,

$$\mathcal{P}(m, \vec{q}\,|\chi) = \mathcal{P}(m\,|\,\chi)\mathcal{P}(\vec{q}\,|\,m, \chi), \tag{7.53}$$

since $\mathcal{P}(A, B) = \mathcal{P}(A\,|\,B)\mathcal{P}(B)$ for any two propositions A and B. For the mass, we assume

$$\mathcal{P}(m\,|\,\chi) = \frac{V}{N} n(m, \chi), \tag{7.54}$$

where $n(m, \chi)$ is the Sheth-Tormen mass function [216],

$$\begin{aligned} n(m, z) = -B\frac{\rho_0}{m}\frac{2g(z)\delta_\mathrm{c}}{\sigma_m^2}\frac{d\sigma_m}{dm}\left(\frac{2a}{\pi}\right)^{1/2} \\ \times \left[1 + a^{-p}\left(\frac{g(z)\delta_\mathrm{c}}{\sigma_m}\right)^{-2p}\right] \\ \times \exp\left[-\frac{a(g(z)\delta_\mathrm{c})^2}{2\sigma_m^2}\right], \end{aligned} \tag{7.55}$$

with the free parameters set to $a = 0.707$ and $p = 0.3$ [217], where $g(z)$ is the growth function, $B \equiv (1 + (2^p\sqrt{\pi})^{-1}\Gamma(-p + 1/2))^{-1}$, σ_m is the standard deviation of the matter fluctuations smoothed with a top-hat filter of size $\sim (m/\rho_0)^{1/3}$, and $\delta_\mathrm{c} = 1.686$ is the collapse threshold. This probability density is already normalized since $\int dm\, n(m, \chi) = N/V$.

With this assumption, the 3D power spectrum has the form

$$\begin{aligned} P_\mathrm{3D}(\chi, k) = \frac{1}{\rho_0^2}\int_{m_\mathrm{low}}^{m_\mathrm{high}} dm\, n(m, \chi)\, m^2 \\ \times \int d^2\vec{q}\,\mathcal{P}(\vec{q}\,|\,m, \chi)\left|\tilde{\phi}\left(\frac{r_\mathrm{s}}{a(\chi)}k;\, \tau\right)\right|^2, \end{aligned} \tag{7.56}$$

where m_low and m_high are the lower and upper bounds of the mass range of the interlopers.

Using Eq. (7.40) we can finally write the power spectrum of the effective convergence

of the interlopers as a function of the density profile of interlopers:

$$P_{\rm I}(k) = \left(\frac{4\pi G}{c^2}\right)^2 D_l^2 \int_0^{\chi_s} d\chi \, \frac{W_{\rm I}^2(\chi)}{g^2(\chi)\chi^2} \times \int dm \, n(m,\chi)\, m^2 \times \int d^2\vec{q}\, \mathcal{P}(\vec{q}\,|\, m,\chi) \left|\tilde{\phi}\left(\frac{D_l r_{\rm s}}{g(\chi) D_\chi} k;\, \tau\right)\right|^2 . \tag{7.57}$$

Power Spectrum of the Interlopers as Effective Subhalos

We start with the 2-point correlation function of the effective convergence,

$$\varepsilon_{\rm 2D}(\vec{r}) = \frac{1}{A} \int d^2\vec{s}\, \kappa_{\rm eff}(\vec{s}) \kappa_{\rm eff}(\vec{s}+\vec{r}), \tag{7.58}$$

where $\vec{r}$ and $\vec{s}$ are physical coordinates on the lens plane, and A is the lens plane area. The subscript 2D is added to differentiate this from the 3D 2-point correlation function of the matter overdensities in Eq. (7.41). Just like in the 3D case, we write the total effective convergence as the sum of the effective convergence of each interloper, $\kappa_{\rm eff}(\vec{r}) = \sum_i \kappa_{i,\rm eff}(\vec{r})$. Again, we assume that the interloper positions are uncorrelated which makes the 2-halo term in Eq. (7.42) vanish and lets us write

$$\varepsilon_{\rm 2D}(\vec{r}) = \frac{1}{A} \sum_i \int d^2\vec{s}\, \kappa_{i,\rm eff}(\vec{s}) \kappa_{i,\rm eff}(\vec{s}+\vec{r}), \tag{7.59}$$

and consequently

$$P_{\rm 2D}(k) \equiv \tilde{\varepsilon}_{\rm 2D}(k) = \frac{1}{A} \sum_i \left|\tilde{\kappa}_{i,\rm eff}(k)\right|^2 , \tag{7.60}$$

where we again dropped the vector notation in $\vec{k}$ due to the radial symmetry of the effective convergence.

Using Eqs. (7.35) and (7.47), we can write

$$\begin{aligned}\kappa_{i,\mathrm{eff}}(s) &= \frac{\Sigma(s; m_{\mathrm{eff},i}, r_{\mathrm{s,eff},i}, \tau_i)}{\Sigma_{\mathrm{crit},l}} \\ &= \frac{1}{\Sigma_{\mathrm{crit},l}} \frac{m_{\mathrm{eff},i}}{r^3_{\mathrm{s,eff},i}} \int dz\ \phi\left(\frac{\sqrt{s^2+z^2}}{r_{\mathrm{s,eff},i}}\,;\,\tau_i\right). \end{aligned} \tag{7.61}$$

In Appendix E.2, we show that the Fourier transform of $\kappa_{i.\mathrm{eff}}$ can be expressed in terms of the Fourier transform of the dimensionless density profile ϕ and, consequently, we can rewrite Eq. (7.60) as

$$P_{\mathrm{2D}}(k) = \frac{1}{A} \sum_i \frac{m^2_{\mathrm{eff},i}}{\Sigma^2_{\mathrm{crit},l}} \left|\phi\left(r_{\mathrm{s,eff},i} k\,;\,\tau_i\right)\right|^2. \tag{7.62}$$

Analogously to the procedure in §7.1.3, we convert the sum into an integral over the probability distribution of the interloper parameters:

$$\begin{aligned} P_{\mathrm{2D}}(k) = \frac{N}{A} &\int d\chi\, dm\, d^2\vec{q}\, \mathcal{P}(\chi, m, \vec{q}) \\ &\times \frac{m^2_{\mathrm{eff}}(\chi)}{\Sigma^2_{\mathrm{crit},l}} \left|\phi\left(r_{\mathrm{s,eff}}(\chi) k\,;\,\tau\right)\right|^2, \end{aligned} \tag{7.63}$$

where $\mathcal{P}(\chi, m, \vec{q})$ is the probability of an interloper being at comoving distance χ and having mass m and intrinsic parameters $\vec{q}$. We substitute the effective scale radius and mass from Eqs. (7.36) and (7.37), and again separate the probabilities $\mathcal{P}(\chi, m, \vec{q}) = \mathcal{P}(\vec{q}\,|\,m, \chi)\mathcal{P}(m, \chi)$, which gives

$$\begin{aligned} P_{\mathrm{2D}}(k) = \frac{N}{A} &\int d\chi \left[f(\chi) \frac{1}{\Sigma_{\mathrm{crit},\chi}} \left(\frac{D_l}{g(\chi) D_\chi}\right)^2\right]^2 \\ &\times \int dm\, \mathcal{P}(m, \chi) m^2 \\ &\times \int d^2\vec{q}\, \mathcal{P}(\vec{q}\,|\,m, \chi) \left|\phi\left(\frac{D_l r_{\mathrm{s}}}{g(\chi) D_\chi} k\,;\,\tau\right)\right|^2. \end{aligned} \tag{7.64}$$

The probability that an interloper has mass m and is at comoving distance χ is proportional to the halo mass function $n(m, \chi)$ and the cross section of the double cone

with the lens plane as the base (Fig. 7.3). Thus, we can write[1]

$$\mathcal{P}(m,\chi) = \frac{S(\chi)}{N} n(m,\chi), \tag{7.65}$$

where $S(\chi)$ is the cross section of the double cone in comoving units at comoving distance χ and is given by

$$\frac{S(\chi)}{A} = \frac{\chi^2}{D_l^2} g^2(\chi), \tag{7.66}$$

where A is the physical area of the main lens and $g(\chi)$ was defined in Eq. (7.30) (see Appendix E.3 for a careful derivation).

Substituting Eqs. (7.65) and (7.66) into Eq. (7.64) gives

$$\begin{aligned} P_{\rm 2D}(k) = &\left(\frac{4\pi G}{c^2}\right)^2 D_l^2 \int d\chi \left[\frac{f(\chi) D_{\chi s} \chi^2}{D_s D_\chi}\right]^2 \frac{1}{g^2(\chi)\chi^2} \\ &\times \int dm\, n(m,\chi)\, m^2 \\ &\times \int d^2\vec{q}\, \mathcal{P}(\vec{q}\,|\, m,\chi) \left|\phi\left(\frac{D_l r_{\rm s}}{g(\chi) D_\chi} k\,;\, \tau\right)\right|^2 . \end{aligned} \tag{7.67}$$

We see that the factor in square brackets is exactly the selection function in Eq. (7.28). Therefore, this equation is identical to Eq. (7.57). We conclude that calculating the power spectrum of the effective convergence on the lens plane after projecting the interlopers onto the lens plane as effective subhalos is equivalent to calculating it using the Limber approximation from the 3D matter power spectrum with the selection function derived in Section 7.1.2.

7.1.4 Effective Convergence Power Spectrum for a Population of tNFW Perturbers

In this section, we compare the convergence power spectrum of interlopers to that of subhalos, referring to both collectively as perturbers. We calculate both contributions to the convergence power spectrum for a fiducial system and show, independently of profile,

[1] This is already normalized because $\int d\chi\, S(\chi) \int dm\, n(m,\chi) = N$.

how each contribution varies differently as a function of source and lens redshift.

Since the power spectrum depends on the perturber profile parameters, we now specify the probability distribution $\mathcal{P}(\vec{q}|m,\chi)$ of the intrinsic halo parameters $\vec{q} = (r_{\rm s}, \tau)$. For both interlopers and subhalos, we assume the following form:

$$\mathcal{P}(\vec{q}\,|\,m,\chi) = \delta\left(r_{\rm s} - r_{\rm s}(m)\right)\delta\left(\tau - 20\right) \tag{7.68}$$

$$r_{\rm s}(m) = r_{\rm s,0}\left[\frac{m}{m_0}\right]^{\gamma}, \tag{7.69}$$

where $r_{\rm s,0} = 0.1\,{\rm kpc}$, $m_0 = 10^6\ {\rm M}_\odot$, and $\gamma = 1/3$ [113, 179, 180]. These parameters are chosen so our results are directly comparable to the convergence power spectrum from only substructure in Ref. [4].

We carry out the integral over $\vec{q}$ in Eq. (7.67), which fixes $\tau = 20$ and $r_{\rm s} = r_{\rm s}(m)$ (given in Eqs. 7.68 and 7.69). After these choices, the effective convergence power spectrum for interlopers is

$$\begin{aligned} P_{\rm I}(k) = &\left(\frac{4\pi G}{c^2}\right)^2 D_l^2 \int_0^{\chi_s} d\chi\, \frac{W_{\rm I}^2(\chi)}{g^2(\chi)\chi^2} \\ &\times \int_{m_{\rm low}}^{m_{\rm high}} dm\, n(m,\chi) m^2 \left|\tilde{\phi}\left(\frac{r_{\rm s}(m) D_l}{g(\chi)a(\chi)\chi}k;\,\tau\right)\right|^2, \end{aligned} \tag{7.70}$$

where $n(m,\chi)$ is the Sheth-Tormen mass function [216]. We use a perturber mass range from $m_{\rm low} = 10^5\ {\rm M}_\odot$ to $m_{\rm high} = 10^8\ {\rm M}_\odot$, since lower-mass perturbers contribute little to the total power and higher-mass perturbers can be modeled directly. For subhalos, the convergence power spectrum is [4]

$$P_{\rm S}(k) = \frac{1}{\Sigma_{\rm crit}^2}\int_{m_{\rm low}}^{m_{\rm high}} dm\, m^2 n_{\rm sub}(m)\left|\tilde{\phi}(r_{\rm s}(m)k;\tau)\right|^2, \tag{7.71}$$

where $n_{\rm sub}(m)$ is the number of subhalos per physical area per mass. We use (see Appendix E.4)

$$n_{\rm sub}(m, z_l) = \frac{0.3\Sigma_{\rm crit,0.5} f_{\rm sub,0.5}(2+\beta)}{(m_{\rm high}^{2+\beta} - m_{\rm low}^{2+\beta})}\frac{(1+z_l)^{5/2}}{(1+0.5)^{5/2}} m^{\beta}, \tag{7.72}$$

where $\beta = -1.9$ and f_{sub,z_l} is the fraction of halo mass in substructure within the mass range at redshift z_l. The factor of $(1 + z_l)^{5/2}$ accounts for the redshift evolution of the subhalo mass function as the subhalos travel within their host [218]. A value of $\beta = -1.9 \pm 0.1$ is fairly well agreed-upon in the literature, both in observations [219] and simulations [135, 220–223], but f_{sub,z_l} is much less constrained. This is because the population of subhalos evolves as it travels within the host and is subject to tidal stripping.

There is no consensus of the extent to which tidal stripping happens − both with and without baryons − as a function of redshift and host mass. Different N-body simulations have found $f_{\mathrm{sub},0}$ on the order of 10^{-3} to 10^{-2} using host halo masses of $\sim 10^{12}$ $\mathrm{M}_\odot$ [5, 60, 137, 143], and we expect baryons to decrease these values. Observations for similar lens redshifts and masses seem to be consistent with a wide range of possible values. Ref. [109] found that for a sample of SLACS galaxies of similar masses (mass within Einstein radius $\sim 10^{11.4}$ $\mathrm{M}_\odot$ [224]) and lens redshifts ($z \sim 0.2$), $f_{\mathrm{sub},0.2} = 0.0076^{+.0208}_{-.0052}$ for subhalos in $[4 \times 10^6 - 4 \times 10^9]$ $\mathrm{M}_\odot$. Ref. [119] found that for BELLS lenses, the upper bound on $f_{\mathrm{sub},0.5}$ is 7% with an upper subhalo mass bound of 10^{11} $\mathrm{M}_\odot$. This constraint includes both interlopers and subhalos. Because our upper mass bound is 3 orders of magnitude below the one cited, we expect the upper bound on $f_{\mathrm{sub},0.5}$ from BELLS to be significantly below that for our mass range.

Taking into account the considerable uncertainty in these observations, and the wide range of plausible values extracted from simulations, we settle on a fiducial value of $f_{\mathrm{sub},0.5} = 4 \times 10^{-3}$ for our mass range $[10^5 - 10^8]$ $\mathrm{M}_\odot$ and host redshift ($z_l = 0.5$), for typical galaxy-scale lenses. We will nevertheless discuss in detail the dependence of our results on $f_{\mathrm{sub},0.5}$ below.

Fig. 7.6 displays the convergence power spectrum due only to interlopers (blue), due only to subhalos (green), and due to both (red) for a fiducial lensing system with $z_l = 0.5$, $z_s = 1$, and mass functions as described above. The numerical results in this figure will be described in detail in Section 7.2. For this lensing system, the power spectrum amplitude due to interlopers is 7.4 times larger than that of subhalos, meaning the former would be

the dominant contribution to any measured signal.

A signal known to be dominated by interlopers would be especially useful for constraining the low-mass end of the halo mass function, which is considered a key way of distinguishing between vanilla CDM and exotic dark matter scenarios that can lead to low-mass cutoffs. This is both because interlopers are simpler to model, as they are generally not subject to the same degree of astrophysics and tidal effects as subhalos, and because the density of interlopers is much better understood. Indeed, while $f_{\rm sub,0.5}$ may range between several orders of magnitude, the two commonly-used mass functions that would affect the interloper amplitude, Sheth-Tormen [216] and Press-Schechter [225], only differ by about a factor of two. Simulations agree with Sheth-Tormen to roughly a 10% level [226, 227], and future observations could in principle measure the halo mass function to percent-level accuracy [228].

Refs. [4, 5] pointed to several features of the power spectrum that could be used to constrain the particle nature of dark matter, such as the slope at $k > 2$ kpc^{-1}; however in the remainder of this section we focus on the amplitude at small values of k, i.e. the $k \to 0$ limit,[2] which primarily provides information about the overall abundance of perturbers within a given mass range. We will refer to this regime as the *plateau* (due to the fact that the power spectrum is constant on these scales), whose amplitude we define as $P_0 \equiv \lim_{k\to 0} P(k)$. Expressions for the interloper and subhalo plateau, $P_{\rm I,0}$ and $P_{\rm S,0}$, are derived in Appendix E.5.

We focus on the low-k scales for two main reasons. First, they are the most readily observable ones. Second, they neatly illustrate the importance of taking into account the contribution of interlopers in order to use strong-lens measurements to draw conclusions about dark matter, without having to worry about the specific details of how the interlopers and subhalos are modeled (which affect the power spectrum at higher wavenumbers [4]), since the amplitude depends only on their mass functions. Due to the fact that both the halo and subhalo mass functions evolve with redshift, P_0 depends on the geometry of

[2]To be precise, our definition of the plateau only matches the $k \to 0$ limit of power when we neglect the 2-halo term from Eq. (7.42), which would contribute an additional term to Eqs. (7.70) and (7.71). However, this term is expected to be small (particularly in the presence of baryons) and only becomes relevant at $k \lesssim 10^{-1}$ kpc^{-1} [5], so it is safe to neglect here.

the lensing system. At higher source redshifts, there are more interlopers along the line of sight, which produces a higher interloper power. The subhalo power spectrum depends on source and lens redshifts through $\Sigma_{\mathrm{crit},l}$ and n_{sub}, resulting in a somewhat different redshift dependence.

The first row of Fig. 7.4 shows the results for the power spectrum plateau due to interlopers (left), subhalos (center), and the ratio between the two (right), for our fiducial choice of $f_{\mathrm{sub},0.5}$. As expected, the relative interloper contribution generally increases with increasing source redshift, so interlopers contribute a greater fraction to the total power spectrum for the higher source redshift BELLS systems compared to SLACS: interlopers dominate over subhalos by a factor of a few for SLACS and by just over an order of magnitude for BELLS. The third row shows the same three panels but for a higher value of $f_{\mathrm{sub},0.5} = 0.02$. We can see that, for this value of the fraction of dark matter in substructure, the subhalo and interloper contributions for SLACS become roughly equal.

To better understand this turnover, and the relative contribution of perturbers more generally, we plot the number of perturbers per solid angle, as well as the ratio of subhalos to interlopers, in the second and fourth rows of Fig. 7.4. We see that the redshift dependence of these numbers is quite different from that of the plateaus,[3] but the ratios share a similar pattern. They are slightly different because a factor of $m^2 n(m)$ goes into the integral for the plateau. Compared to the plateau ratios at a given redshift, the number density ratios tend to be slightly larger, which means that to a rough approximation, we can think of the plateau turnover as the place where there are slightly more subhalos than interlopers in the field of view.

Because these results are strongly dependent on the choice of subhalo and halo mass functions, we provide an interactive version of the power spectrum plots at `https://arthur-tsang.github.io/interloper_widget.html`, where the reader can adjust the value of $f_{\mathrm{sub},0.5}$ as well as the halo mass function (Sheth-Tormen or Press-Schechter) to see how the results are affected. For SLACS lenses up to $f_{\mathrm{sub},0.5} \sim 2\%$, the interloper contribution dominates. For larger values of $f_{\mathrm{sub},0.5}$, however, the subhalo contribution

[3] In particular, the number of interlopers has a strong dependence on lens redshift, which comes from converting into angular dimensions.

takes over (albeit by less than an order of magnitude). For the BELLS lenses, a higher value of $f_{\rm sub,0.5} \sim 4\%$ is necessary for the subhalos to dominate due to their higher source redshifts.

7.2 Numerical Calculation

To verify our analytic results and, in particular, the validity of our approximations, we independently calculate the (effective) convergence power spectrum numerically. The main concern is the error introduced by approximating the interlopers as effective subhalos at the redshift of the main lens.[4] To understand the error introduced by this approximation, we calculate the convergence power spectrum from a lensing system simulated using the full multi-plane lens equation.

In short, our procedure is to first use ray-tracing to generate a map of the total angular deflection $\vec{\alpha}(\vec{x}_1) \equiv \sum_{i=1}^{N} \vec{\alpha}_i(\vec{x}_i)$, and then calculate

$$\kappa_{\rm eff,total} \equiv \frac{1}{2} \nabla \cdot \vec{\alpha}, \tag{7.73}$$

which is the multi-plane equivalent of the traditional convergence for single-plane lensing [206]. To single out the effective convergence of the interlopers (or subhalos), we subtract the convergence of the main lens:

$$\kappa_{\rm eff} = \kappa_{\rm eff,total} - \kappa_l. \tag{7.74}$$

We then convert $\kappa_{\rm eff}$ into a power spectrum using a 2D Fourier transform (squared) [5]. Since we are interested in the monopole term, we perform an azimuthal average.

In Section 7.2.1 we describe the lensing system we simulate, while in Section 7.2.2 we detail the procedure used to obtain the effective convergence power spectrum from it.

[4] Note that other works have made this approximation as well [155, 156].

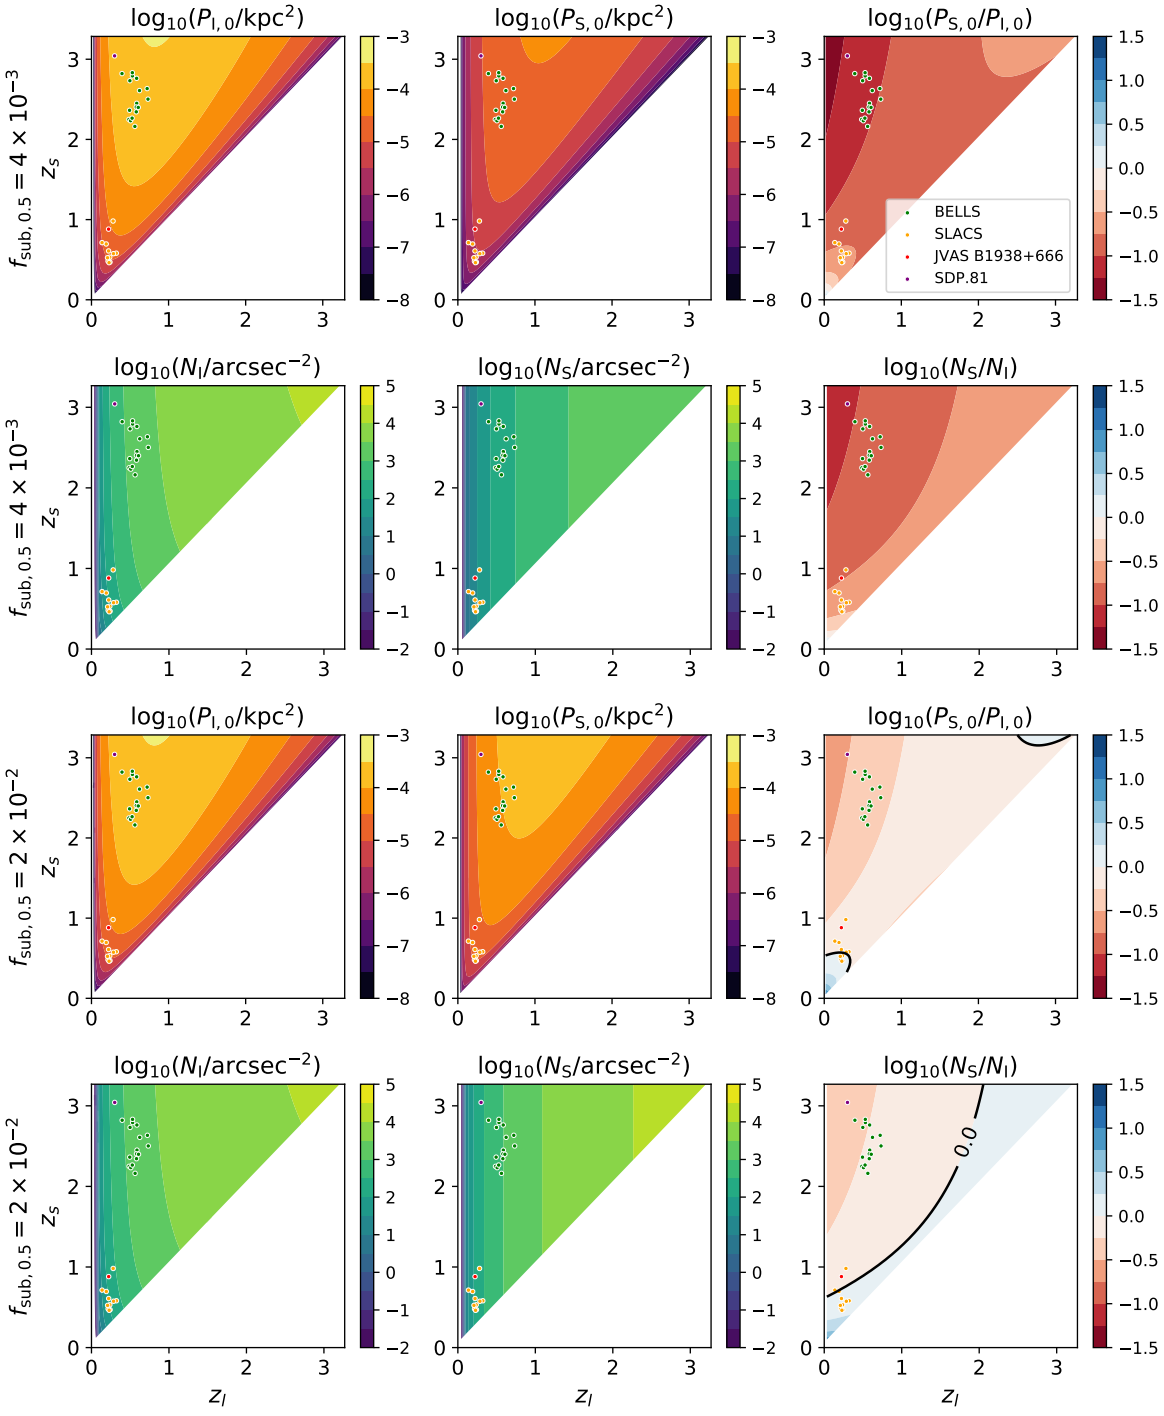

Figure 7.4: *Rows 1 and 3:* Plateau, i.e. $k \to 0$ limit, of the convergence power spectrum of perturbers in the range [10^5-10^8] $\mathrm{M}_\odot$, for interlopers (*left*), subhalos (*center*), and the ratio of subhalos to interlopers (*right*). *Rows 2 and 4:* Number of perturbers per arcsec^2, for perturbers in mass range [10^5-10^8] $\mathrm{M}_\odot$, for interlopers (*left*), subhalos (*center*), and the ratio (*right*). *Top two rows:* Fiducial subhalo normalization, $f_{\mathrm{sub},0.5} = 4 \times 10^{-3}$. *Bottom two rows:* Higher subhalo normalization, $f_{\mathrm{sub},0.5} = 2 \times 10^{-2}$, which gives a roughly equal subhalo and interloper contribution for SLACS systems. The line of equal contribution is shown in black. The dots represent some of the galaxy-galaxy lensing systems that have been studied in the literature [107, 112, 119, 140]. See `https://arthur-tsang.github.io/interloper_widget.html` for an interactive version of the power spectrum plateau plots, with adjustable $f_{\mathrm{sub},0.5}$ and halo mass function.

7.2.1 A Simulated Strong-Lensing System with Interlopers and Subhalos

We set the lens at $z_l = 0.5$ and the source at $z_s = 1.0$, and we assume a flat ΛCDM cosmology with $h = 0.675$ and $\Omega_{\rm m} = 0.316$ [11]. The lens is a singular isothermal sphere (SIS) with an Einstein radius of $\theta_{\rm E} = 7''$, which corresponds to a virial mass of $\sim 10^{14}$ $\mathrm{M}_\odot$. This is larger than the typical $\sim [0.1 - 1]''$ or $\sim [10^{12} - 10^{13}]$ $\mathrm{M}_\odot$ of galaxy-galaxy lenses [229, 230], although note that some systems may go up to $10^{14.5}$ $\mathrm{M}_\odot$ [231]. We use a somewhat larger lens in order to probe a broader range of wavenumbers; nevertheless, our results are fully applicable to galaxy-scale lenses because we use values of $f_{\rm sub,0.5}$ consistent with typical galaxy lenses, instead of scaling it up for a larger halos mass. Furthermore, the power spectrum is otherwise independent of lens size in our formalism. Note that because we calculate the power spectrum from deflection angles rather than from an observed image, the only relevant property of the source is its redshift.

We randomly populate interlopers on 100 planes, evenly spaced in redshift between the source and observer.[5] The number of interlopers on each plane is chosen as a Poisson random variable whose expected value is the number of interlopers between adjacent redshift planes according to the Sheth-Tormen mass function [216] (same as in Section 7.1.3). The interloper masses are randomly chosen between $[10^5 - 10^8]$ $\mathrm{M}_\odot$, again according to Sheth-Tormen. Their positions within each plane are uncorrelated and uniformly distributed within the double pyramid of visible structure (analogous to the double cone in Fig. 7.3, but now instead we use a double pyramid, since we simulate a square field of view). The interloper profiles are modeled as in Section 7.1.4. While the true profile of interlopers may be slightly different, note that the low-k limit of the power spectrum is sensitive to the abundance of interlopers and not to the intrinsic profile parameters. In addition to the interlopers, we add a negative mass sheet to each redshift plane to cancel out the net mass of the interlopers, which captures the fact that the underdense regions along the line of sight effectively lens as negative masses.

For completeness, we also simulate a lensing system that only has subhalos, and

[5] Note that the result is unchanged if we increase the number of planes.

one that has both subhalos and interlopers. To populate the lens with subhalos, we assume that subhalos are uncorrelated, uniformly spatially distributed, and follow the mass function of Eq. (7.72). Note that it is possible to simulate the two-subhalo term by modifying the spatial distribution, however, in a realistic lensing galaxy with baryons we expect this term to be subdominant [5].

7.2.2 Obtaining the Convergence Power Spectrum from a Simulation

We calculate the total deflection vector $\vec{\alpha}(\vec{x}_1)$ using the full multi-plane lens equation for the simulated lensing systems using `lenstronomy` [133], which is a publicly available `Python` package. We then calculate κ_{eff} by taking the divergence of $\vec{\alpha}$ using the five-point stencil method (see Appendix E.6) to limit the numerical error.

We run two different simulations, both with (500×500) pixels, but with different fields of view: $(1.6 \times 1.6)''$ (small) and $(16 \times 16)''$ (large). This is to sample a wide range of wavenumbers that would otherwise require a much larger number of pixels and would thus be computationally intractable. These two different fields of view require different treatments to be processed into the convergence power spectrum. This is because when all parts of the κ_{eff} map are statistically equivalent, the Fourier transform squared of κ_{eff} is the two-dimensional interloper power spectrum. However, our analytic approximation from Section 7.1 only applies near the Einstein radius (since the derivation relied on the CSB approximation).

Due to this, for the large field of view we filter κ_{eff} with an annular mask centered on the main lens, setting $\kappa_{\text{eff}} = 0$ outside the mask (see Fig. 7.5).[6] For the small field of view, we center the image on a point on the Einstein ring in order to remain in a regime where the CSB approximation is valid, so we can compare with our analytic results.

Having to impose a mask for the large field of view has several limitations. First, while it is not desirable to use points too far from the Einstein ring (the CSB approximation

[6]Using a mask affects the normalization of the Fourier transform, so to correct for it we divide the power by the fraction of the image covered by the mask.

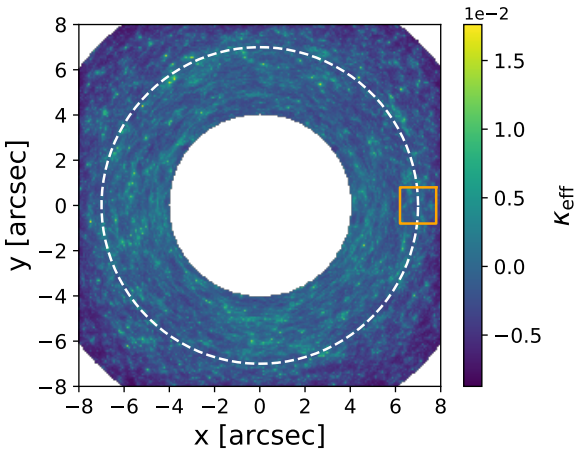

Figure 7.5: Illustration of the effective convergence field $\kappa_{\rm eff}$ used for the full multi-plane lensing calculation with $z_l = 0.5$ and $z_s = 1.0$. The annular mask can be seen for the large field of view ($16''$). The small field of view ($1.6''$) is represented with an orange square on the right. The white dashed line shows the Einstein radius.

gets progressively worse with increasing distance), a narrow mask does not estimate well the lowest-k modes since they correspond to sizes larger than the annulus width. Furthermore, the Fourier transform of the mask can itself give rise to unphysical oscillations. Ultimately, we opted for a mask width of $\pm(3/7)\theta_{\rm E}$. An annulus of this width is sufficient to smooth out the oscillations and probe relatively low-k modes. Furthermore, because it is quite wide, it allows us to be conservative when comparing it to the analytical results: we know that as the mask becomes wider, the validity of the CSB approximation decreases.

Comparison to Analytical Results

In Fig. 7.6, we compare the effective convergence power spectrum obtained following the procedure outlined above to the analytical predictions derived in Section 7.1.4. We show the contribution due solely to interlopers (blue), solely to subhalos (green), and their combination (red).[7] The analytical results are shown as dotted lines, and the numerical results are shown as solid lines. We see that the two independent estimates of the effective convergence power spectrum show excellent agreement,[8] even though the annular mask

[7] We used the same subhalo and interloper population characteristics, as described in Eqs. (7.68) to (7.72).

[8] Note that, for images without masks (subhalo only and all small field of view images), the minimum k we plot is $2\pi/L$, where L is the width of the image. For the images with masks (interloper and combined

used was quite wide. This agreement shows that treating the interlopers as effective subhalos using our framework introduces a very small error compared to the full ray tracing results, even in a regime far from the Einstein radius. Furthermore, we note that the small difference between the analytical and numerical results is much smaller than even the most optimistic expected error bars from near-future surveys [120].

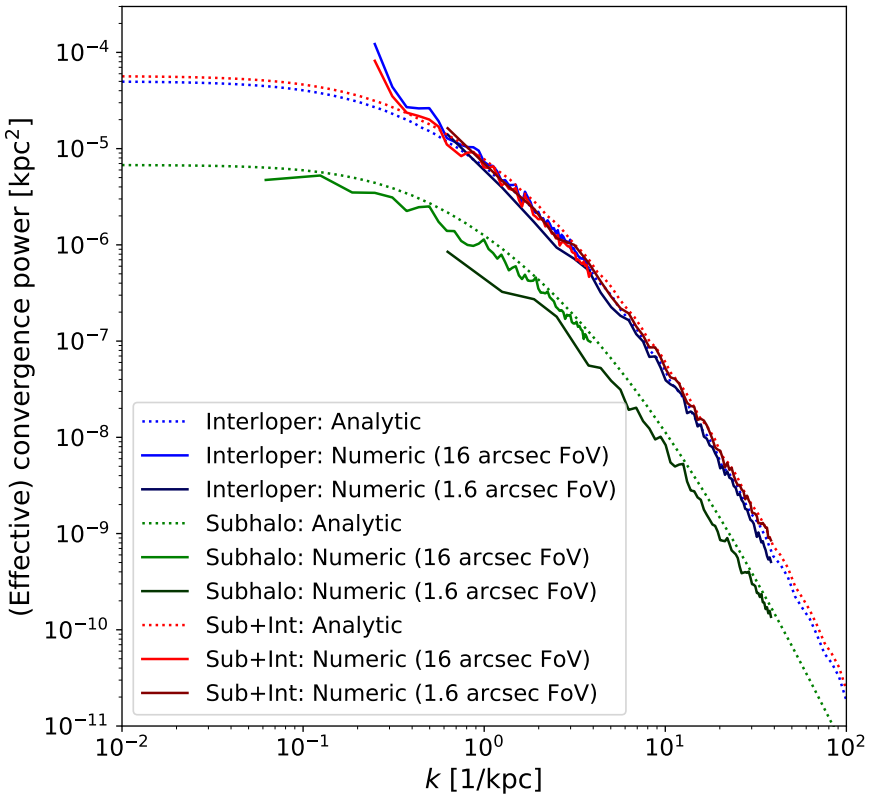

Figure 7.6: Power spectra for interlopers (blue), subhalos (green), and both (red). We show analytic results (dotted) and numerical results (solid) for the simulations with field of view (FoV) of 1.6″(darker) and 16″(lighter). All numerical curves are cut off at high k due to loss of power near a pixel size, and the masked simulations are cut off at low k, corresponding to the mask width.

Quantifying the Error in the Numerical Effective Convergence

The advantage of calculating the effective convergence as in Eq. (7.73) is that we do not have to make any approximations, such as the CSB approximation that was necessary for the analytical derivation in Section 7.1. However, this single-plane effective convergence, and the one derived in Section 7.1, do not reproduce the deflection angles exactly. To see

for large field of view), the minimum k is $2\pi/L'$, where $L' = L/4$ is the width of the annulus mask (see Fig. 7.5).

why this is the case let us separate the deflection angles into a curl-free and a divergence-free component:

$$\vec{\alpha} = \vec{\alpha}_{\rm div} + \vec{\alpha}_{\rm curl} \tag{7.75}$$

$$\nabla \times \vec{\alpha}_{\rm div} = 0 \;\; \& \;\; \nabla \cdot \vec{\alpha}_{\rm curl} = 0. \tag{7.76}$$

We call $\vec{\alpha}_{\rm div}$ the divergence component and $\vec{\alpha}_{\rm curl}$ the curl component.

In single-plane lensing, $\vec{\alpha}_{\rm curl}$ vanishes because the deflection angle can be written as the gradient of the lensing potential Ψ [212]:

$$\vec{\alpha}(\vec{x}) = \nabla_x \Psi(\vec{x}) \quad \Rightarrow \quad \nabla \times \vec{\alpha} = 0. \tag{7.77}$$

However, in the multi-plane lensing case, the coupling between the successive lens planes introduces a curl component [207]. In our calculation it is a nuisance since it indicates that the single-plane effective convergence is not fully capturing the multi-plane lensing of the interlopers. In terms of the total deflection angle, the two components can be written as (see Appendix E.6)

$$\vec{\alpha}_{\rm div}(\vec{x}) = \frac{1}{\pi} \int d^2x' \, \frac{\vec{x} - \vec{x}'}{|\vec{x} - \vec{x}'|^2} \left[\frac{1}{2} \nabla_{\vec{x}'} \cdot \vec{\alpha} \right] \tag{7.78}$$

$$\vec{\alpha}_{\rm curl}(\vec{x}) = \hat{z} \times \frac{1}{\pi} \int d^2x' \, \frac{\vec{x} - \vec{x}'}{|\vec{x} - \vec{x}'|^2} \left[\frac{1}{2} \nabla_{\vec{x}'} \times \vec{\alpha} \right]. \tag{7.79}$$

Here $\hat{z}$ is the unit vector that is orthogonal to the lens-plane and pointing towards the observer. We see that what we defined as $\kappa_{\rm eff}$ in Eq. (7.73) sources the divergence component, i.e.

$$\kappa_{\rm div} = \kappa_{\rm eff} \equiv \frac{1}{2} \nabla \cdot \vec{\alpha}, \tag{7.80}$$

and

$$\kappa_{\rm curl} \equiv \frac{1}{2} \nabla \times \vec{\alpha} \tag{7.81}$$

sources the curl component. So we can compare $\kappa_{\rm curl}$ to $\kappa_{\rm eff}$ in order to gain an under-

standing of the error introduced by treating interlopers as effective subhalos.

We see in Fig. 7.7 that for our simulation, $\kappa_{\rm curl} \ll \kappa_{\rm eff}$,[9] especially near the Einstein radius. We thus conclude that the curl component of the angular deflection is also much smaller than the divergence component, meaning that the coupling between the interlopers and the main lens is small enough to justify the projection of interlopers as effective subhalos in the main lens. To test whether the $\kappa_{\rm curl}$ in Fig. 7.7 is simply a numerical artifact, we simulate a system with only a main lens and subhalos (no interlopers) in Appendix E.6, and show in Fig. E.6.1 that the numerical error is more than two orders of magnitude smaller than the curl observed in Fig. 7.7, showing that the curl term sourced by the interlopers is physical.

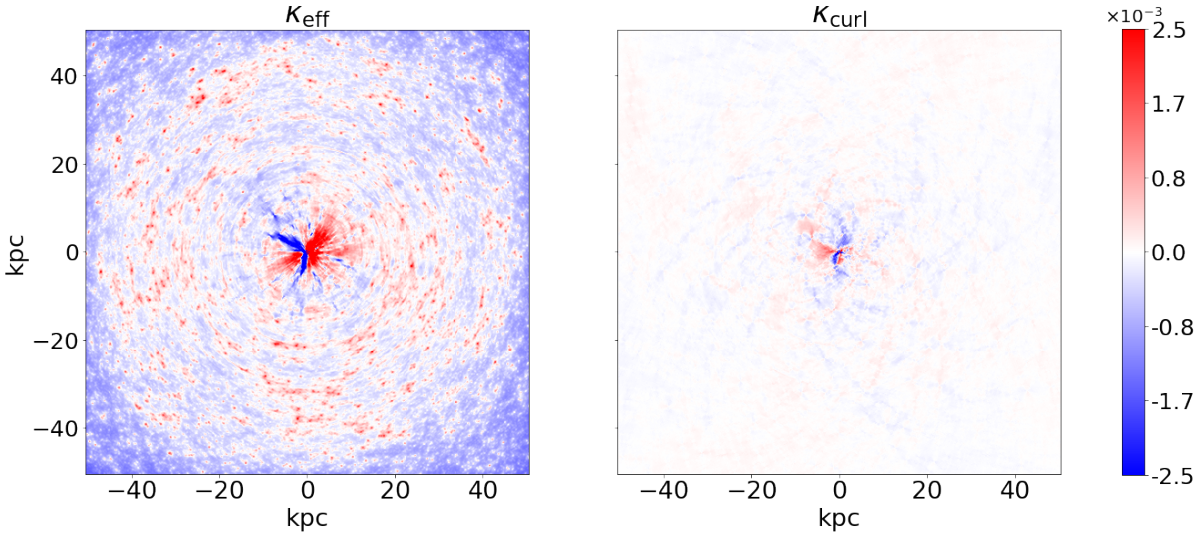

Figure 7.7: Two types of effective convergences calculated from the divergence and the curl of angular deflections in a simulated strong-lensing system with interlopers and subhalos. *Left:* Effective convergence of the interlopers + subhalos, defined in Eq. (7.74), which sources the divergence component of the angular deflections. The central dipole is caused by the coupling between the uneven distribution of interlopers and subhalos and the main lens. We can safely ignore the dipole as we are interested in a small annulus around the Einstein ring where the strong-lensing images are produced. *Right:* $\kappa_{\rm curl}$, defined in Eq. (7.81), which sources the curl component of the angular deflections. The interlopers that are far from the lens plane, either towards the observer or towards the source, contribute more to the curl component.

[9]We know from Eq. (7.77) that the curl component has to be zero when there are no interlopers, i.e. when the lensing is caused by mass on a single lens plane. In Appendix E.6 we show that the curl component that we measure is not a numerical artifact but indeed a result of the coupling of the lensing effect of multiple lens planes at various redshifts.

7.3 Conclusions

Strong gravitational lensing has long provided some of the most decisive evidence for the existence of dark matter in our universe. Within the past couple of decades, the use of gravitational lensing for dark matter science has expanded considerably, and significant theoretical and observational advances have turned it into one of the most promising probes of the nature of dark matter. In particular, there is great interest in using strongly lensed images to constrain the very low-mass end of the halo mass function ($\lesssim 10^8$ $M_\odot$), since this regime can distinguish between vanilla CDM and more exotic models, such as warm dark matter or self-interacting dark matter, that can lead to low-mass cutoffs.

The canonical approach followed has been to try to directly detect individual dark matter clumps, such as in gravitational imaging [103]. This approach has claimed detections of substructures with masses as small as a few times 10^8 $M_\odot$ [106, 107, 109, 111, 119, 140], but reliably reaching lower masses in galaxy-galaxy lenses with this approach has remained elusive.

The idea of using the subhalo convergence power spectrum [4, 5, 117, 200] was developed as a statistical detection method to obtain population-level constraints without having to individually resolve clumps. Unlike in direct detection efforts where by construction the sensitivity is maximal for the most massive clump close to the lensed images, in a power spectrum approach the higher number of lower-mass halos can actually make the sensitivity peak for the mass range of $10^7 - 10^8$ $M_\odot$ for a CDM population of subhalos, and still maintain some sensitivity at lower masses [5].

However, while perturbations due to line-of-sight halos have been studied in the context of lensed galaxies [155, 156] and lensed quasars [169, 206], and has become standard practice in direct detection pipelines [140, 207–209], its contribution to the convergence power spectrum had not yet been quantified. In this work, we set out to do so.

We show that it is possible to define an effective convergence for multi-plane lensing systems with a dominant main lens coupled to lower-mass interlopers. One can think of this effective convergence as mapping an interloper at any point along the line of sight

onto the lens plane as an effective subhalo with a modified mass and scale radius.[10] It is then possible to analytically calculate its power spectrum, incorporating the relative effect of interlopers at different redshifts as a lensing kernel. We find that the interlopers that are closest to the main-lens plane have the largest contribution to the power spectrum, while those close to the observer or source are negligible.

Because the halo and subhalo mass functions evolve with redshift, and in fact there are still considerable unknowns with respect to the subhalo mass function's normalization, we do not expect that a global statement can be made about the importance of one versus the other. Instead, we opt to thoroughly quantify each contribution as a function of both source and lens redshift, and for different choices of mass function and subhalo mass function normalizations. We show specifically what we expect the ratio of power between subhalos and interlopers to be for galaxy-galaxy lensing systems for which we currently have high-resolution imaging (Fig. 7.4).

For our fiducial choice of $f_{\rm sub,0.5} = 4 \times 10^{-3}$, we find that for both the BELLS and SLACS lenses the interloper contribution dominates, albeit to different extents. Due to the higher redshift of the BELLS sources, the interloper contribution is larger by over an order of magnitude, while the lower redshift of the SLACS sources lead the interloper plateau to only be larger by a factor of a few. As we increase or decrease $f_{\rm sub,0.5}$, each group of lenses is affected differently: for SLACS, the subhalos dominate for as little as $f_{\rm sub,0.5} \gtrsim 2\%$, while for BELLS they do so for $f_{\rm sub,0.5} \gtrsim 4\%$. This can be intuitively understood: with increasing source redshift, the LOS volume increases, overwhelming the subhalo signal.

Let us put these bounds into context by discussing the expected values of $f_{\rm sub,0.5}$ in these systems. As we discussed in Section 7.1.4, there is considerable uncertainty both on the simulation side and on the observational side. Nevertheless, even with this uncertainty, it seems unlikely that $f_{\rm sub,0.2}$ and $f_{\rm sub,0.5}$ would reach these values, especially with an upper subhalo mass of 10^8 $M_\odot$. Ref. [109] found that for SLACS lenses, the upper bound on $f_{\rm sub,0.2}$ was about 2.7% with an upper mass bound of 4×10^9 $M_\odot$. Ref.

[10] The scale radius is the relevant lensing length scale in our density profile of choice (truncated NFW). Other density profiles would see an analogous rescaling of relevant intrinsic parameters.

[119] found that for BELLS lenses, the upper bound on $f_{\rm sub,0.5}$ is 7% with an upper mass bound of 10^{11} $M_\odot$, which corresponds to $f_{\rm sub,0.5} < 2.3\%$ for an upper mass bound of 10^8 $M_\odot$. Thus, we conclude that it is likely for the interloper contribution to dominate in these two ensembles of lenses.

One worry about this approach might be that treating the interlopers as effective subhalos is overly simplistic since it neglects the recursive nature of the multi-plane lens equation, which couples the deflection angles of successive lens planes. To study this, we tested the analytical calculation with mock lensing simulations obtained by doing ray tracing with the multi-lens plane equation (without any approximations). We find that the power spectrum from the simulations matches the analytical prediction extremely well.

Furthermore, we note that defining the effective convergence as the divergence of the deflection angle does not capture the divergence-free part (what we call the curl component in Section 7.2.2) of the angular deflections sourced by interlopers. For our analysis, which projects the interlopers onto the main lens as effective subhalos, the curl part quantifies the error in doing such a projection since a single-plane effective convergence cannot create such a term (See Appendix E.6). We show in Fig. 7.7 that the curl term for low-mass interlopers is very small compared to the divergence term, meaning that the projection introduces minimal error in calculating the angular deflections for multi-plane lensing. We point out that, for more massive perturbers, the fact that multiple lens planes source a curl term suggests a novel way of identifying multi-plane lensing and therefore distinguish interlopers from subhalos, by for example measuring a B-mode power spectrum.

Our results on the importance of incorporating interlopers into the analysis of strong-lensing systems are broadly consistent with Refs. [155] and [156], which did so in the context of direct detection efforts. An important point to keep in mind is that the mass ranges and subhalo mass function normalizations they consider are different to ours; their perturber mass range spans $\sim 10^6$ $M_\odot - 10^{11}$ $M_\odot$, so converting to our definition of $f_{\rm sub,0.5}$ gives a lower value than what they cite. Ref. [155] considered a single lensing

configuration ($z_l = 0.2$ and $z_s = 1$) and found that the number of interlopers is roughly four times higher than that of subhalos. Ref. [156] found that the number of line-of-sight perturbers is comparable to subhalo perturbers for low-redshift lenses (e.g. SLACS lenses) and dominant over subhalo perturbers for high-redshift lenses (e.g. BELLS lenses). The difference between our results and these can be understood because of the different mass range: although they have a smaller subhalo mass function normalization, the extra subhalos at larger masses make the subhalo contribution comparable for SLACS instead of subdominant, as is the case in our fiducial results.[11]

As the LOS contribution has gained recognition as an integral ingredient in lens modeling, new systems have been analyzed, and older systems reanalyzed, taking it into account. To date, a single real lens has been analyzed through a power spectrum approach. Ref. [232] placed an upper bound on the convergence power spectrum due to subhalos using lens system SDSS J0252+0039 ($z_l = 0.280$, $z_s = 0.982$)[12]. Their upper bound on the power spectrum was significantly higher than the expected amplitude due to subhalos in a CDM scenario, but interestingly, according to our results, for this redshift combination we expect the line-of-sight contribution to dominate.

The fact that under many lensing configurations and reasonable subhalo population assumptions the interloper contribution dominates the signal can be good news for the capacity of strong gravitational lenses to constrain the particle nature of dark matter. The amplitude of the convergence power spectrum can essentially be tied back to a mass function (halo mass function for interlopers, subhalo mass function for subhalos), so in order to translate a power spectrum amplitude to a dark matter theory, it is paramount to understand the relevant mass function(s). The subhalo mass function is very hard to pin down. It depends on the host mass and inevitably evolves with redshift. How it is affected by subhalos traveling within the host's dark matter halo, as well as due to

[11]While these two references developed the notion of using interlopers as effective subhalos, we note that our projection prescription is inherently different. For example, Ref. [156] relied on first projecting the interloper positions onto the lens plane by ensuring they remained on the same line of sight, and then varied their mass to minimize the residual of the angular deflections between the projected interloper and a subhalo of a given mass. The downside of this projection prescription is that the minimization is done for an unobservable lensing quantity.

[12]This system had been analyzed with gravitational imaging in Ref. [109] and no evidence of a substructure above the mass-detection threshold was found.

any baryons in the host, remains an open problem: neither theory nor simulations have yet converged on a satisfying answer to these questions. In comparison, halos that are only subject to large-scale tidal fields have relatively calmer lives, and their evolution is better understood. Therefore, having a window into the smallest dark matter scales in the universe without relying on subhalos can make gravitational lensing a much more powerful (and reliable) probe of dark matter.

Furthermore, we note that the advantages of statistical detection efforts compared to direct detection ones that were introduced in the subhalo context, namely taking advantage of the much more numerous population of low-mass halos that are individually below the detection threshold, are undeniably advantageous in the interloper context as well. For instance, the number of interlopers that are massive enough for direct detection was shown to be roughly unity for the 17 BELLS lenses shown in Fig. 7.4 [119], which prevents lack of detections in the ensemble of lenses to be used to rule out CDM. Since the lower-mass interlopers are expected to be much more abundant, if the power spectrum can be measured (which Refs. [117, 120] claim can be done using near future observations), the lack of power at high redshift lenses can more decisively rule out the CDM scenario.

There is considerable momentum being harnessed by gravitational lensing as a cosmological probe. Much progress has been made over the course of the last decade regarding how to model these systems. Furthermore, in the last couple of years new methods that harness the image recognition power of machine learning methods have started being developed to accelerate the indirect and direct detection of perturbers in strong-lens images [1, 153, 154, 233]. Between these advances in detecting perturbers in optical imaging data, and the fact that we expect thousands of new high-quality optical imaging strong-lens systems to become available in the near future [122–124], we expect to have a treasure trove of data for dark matter science soon. In order for strong gravitational lensing to establish itself as a premier way of constraining dark matter, however, we need to ensure that the mapping from observations to theory is done correctly, which undoubtedly involves accounting for the line-of-sight contribution.

Statistical Inference in the Era of Cosmological Tensions

Chapter 8

Observable Predictions for Massive-Neutrino Cosmologies with Model-Independent Dark Energy

The source of the observed accelerated expansion of the universe (*dark energy* from now on) has remained an elusive component since its discovery, two decades ago [12, 234]. The standard paradigm in cosmology, the ΛCDM model, assumes that the cosmological constant Λ is the source of acceleration, and its six free parameters are constrained to exquisite precision by current Cosmic Microwave Background (CMB; e.g. [11]), Baryon Acoustic Oscillations (BAO; e.g. [235]) and weak lensing (WL) measurements (e.g. [236]). Despite the success of the ΛCDM model in providing a solid statistical fit to these various probes, no satisfying theoretical model elucidates the microphysical origin of the cosmological constant.

Cosmological probes have become precise enough that they can be used to look for physics beyond the Standard Model of particle physics; for example, to study neutrino properties. In the early universe, neutrinos are relativistic, while at late times they become non-relativistic, with their mass constituting a non-negligible fraction of the total dark matter component. The conversion of radiation to hot dark matter affects the Hubble expansion, and the residual streaming velocities are still significant enough

at low redshifts to slow down the growth of structure. This means that neutrinos affect cosmology at both the background and perturbation level.

Direct detection experiments on Earth have measured the three mixing angles and the two mass-squared splittings of the three neutrino mass eigenstates with high precision [237]. However, so far cosmology appears to be the most sensitive probe to the absolute mass scale of neutrinos. To date, the most stringent upper limit of the sum of neutrino masses ($M_\nu \equiv \sum m_\nu < 0.12$ eV) is given by the combination of CMB TT,TE,EE power spectra and lensing from *Planck*, together with BAO data, assuming the standard ΛCDM scenario [11].

When probing the presence of additional light particles (such as neutrinos) with cosmological datasets, dark energy acts as a source of systematic uncertainty that needs to be marginalized over. Extensions to the base ΛCDM model in terms of dark energy are generally implemented by allowing the equation of state w to be a function of redshift z. The most common parametrization for $w(z)$ is known as the Chevallier-Polarski-Linder (CPL) parametrization, given by $w(z) = w_0 + w_a z/(1+z)$, i.e. it has two free parameters $\{w_0, w_a\}$. However, the fact that there is no firm theoretical backbone to support any particular parametrization for $w(z)$ motivates the study of the effects of dark energy using nonparametric methods, such as Principal Components Analysis (PCA).

We consider a dark energy scenario known as the Smooth Dark Energy Paradigm [238], which makes a set of assumptions about the microphysical nature of dark energy. It assumes that the source of dark energy: (1) does not cluster inside the horizon, (2) interacts only gravitationally with dark matter and baryons, and that (3) gravity is set by General Relativity. An analysis of such a scenario has been performed in the past, including cross-checks between the background expansion of the universe and the growth of linear perturbations [239], as a smoking gun to falsify models where dark energy follows these assumptions [240]. In fact this chapter partially generalizes Ref. [240] with the inclusion of the sum of neutrino masses, except that in that work 20 principal components (PCs) were used, and here we restrict ourselves to considering 1, 3 or 5 PCs. Ref. [240] required a complete basis of PCs, and the authors determined empirically that 20 PCs

were sufficient to satisfy this requirement. Here, however, we do not require completeness but rather modes with a non-negligible signal-to-noise, and we seek to contrast our results to specific parametrizations of $w(z)$, which generally have 1-2 parameters.

There are two distinct regimes when studying extensions to ΛCDM concerning the dark energy equation of state. In the first, and most commonly considered regime, the equation of state has a lower bound set by the value of the cosmological constant, i.e. $w(z) \geq -1$. Alternatively, models that do not satisfy this criterion are in the "phantom dark energy" regime. Ref. [241] showed that the bound on the sum of the neutrino masses is more stringent with a CPL non-phantom dark energy source than in a standard ΛCDM cosmology. Indeed, the expansion rate is faster for non-phantom dynamical dark energy than for ΛCDM. The sum of neutrino masses is consequently pushed downwards to keep the angular diameter distance to the surface of the last scattering fixed. Conversely, in phantom dark energy scenarios the opposite is true and thus the bounds on M_ν degrade. Similar results have been found by other extensions to ΛCDM that include CPL-based dark energy [242, 243]. Note that the general argument is independent of the chosen CPL parametrization for the dark energy equation of state.

In this work, we aim to generalize these findings by studying constraints on the sum of neutrino masses in a cosmology with a model-independent, time-varying dark energy component, without forbidding the cross into the phantom regime. In Section 8.1 we introduce our methodology and the datasets used. In Section 8.2 we present our main results, and in Section 8.3 we discuss the implications of our findings.

8.1 Data and Methods

The datasets considered in this work include: BAO data from BOSS DR12 [244], 6dF Galaxy Survey (6dF) [245] and SDSS DR7 Main Galaxy Sample (MGS) [246]); the full 2015 lensed *Planck* CMB temperature and polarization data [247] and lensing reconstruction [248]; Dark Energy Survey (DES) four-bin tomographic weak lensing data [249, 250], shown in Figure 8.2; and the Pantheon supernovae (SN) sample [14]. The latter covers

a redshift range $0.01 < z < 2.3$, and the BAO measurements lie within this range. To account for the non-linear scales in the matter power spectrum we adopt the HALOFIT fit [251], and a mapping between arbitrary $w(z)$ and a constant dark energy equation of state is implemented as in Ref. [252].

We perform a Markov Chain Monte Carlo (MCMC) likelihood analysis with a modified version of the CosmoMC code [253–256]. We ran two different sets of chains. The first of these, which we will refer to as the *All* chains, includes all the aforementioned datasets. The second one, which we refer to as the *Reduced* chains, does not include the DES weak lensing. The datasets used for each of these chains are summarized in Table 8.1. Note that there is a known (2.4σ) tension between the value of S_8 from *Planck* ($0.848^{+0.024}_{-0.024}$) and DES ($0.782^{+0.027}_{-0.027}$) [19], and we will discuss this in Sections 8.2 and 8.3. Nevertheless, note that CMB lensing reconstruction (included in both datasets) favors a lower S_8 that than inferred from the CMB temperature and polarization under the ΛCDM model [142], and thus more in accordance with the value of S_8 measured with WL.

Table 8.1: Datasets that define the *All* and *Reduced* sets of chains

	All	*Reduced*
CMB	Full *Planck* 2015	Full *Planck* 2015
BAO	BOSS DR12 + 6dF + MGS	BOSS DR12 + 6dF + MGS
SN	Pantheon	Pantheon
WL	DES	-

We represent the dark energy equation of state as

$$w(z) = w_{\rm fiducial} + \sum_{i=1}^{N_{\rm PC}} \alpha_i e_i(z), \tag{8.1}$$

where $e_i(z)$, with $i = 1, ..., N_{\rm PC}$, are the principal components of a covariance matrix of perturbations around the fiducial model $w_{\rm fiducial} = -1$. The principal components used here have support in the range $0 < z < z_{\rm max}$, and for $z > z_{\rm max}$ we extrapolate assuming $w = -1$.

We construct the PC basis from the eigenvectors of the Wide Field Infrared Survey Telescope (WFIRST) experiment [257] supernovae Fisher matrix, which has SN data

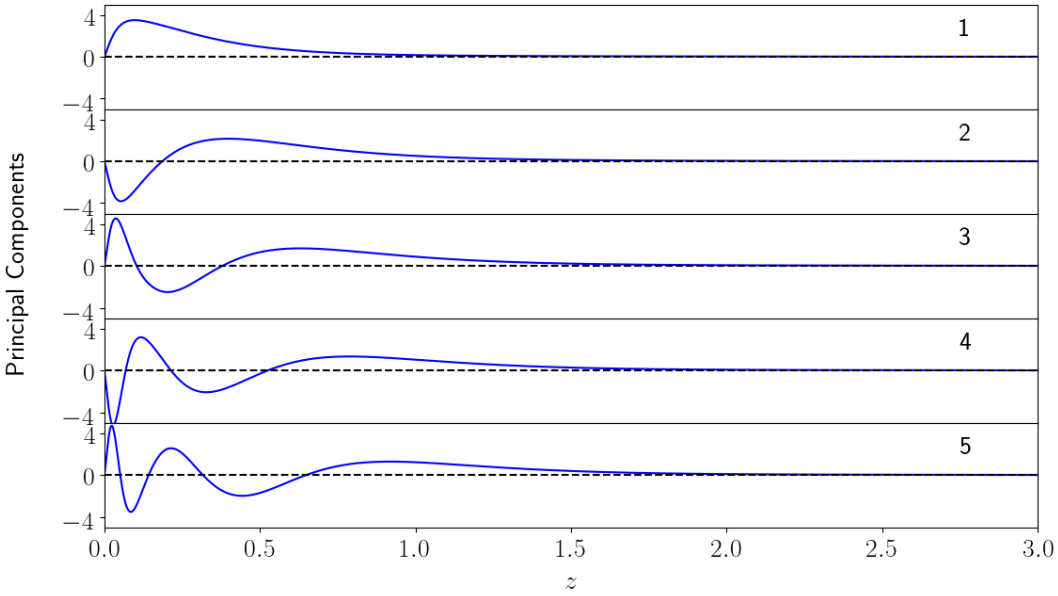

Figure 8.1: The first five principal components of the dark energy equation of state obtained from the WFIRST supernovae Fisher matrix, as described in Section 8.1 and Ref. [240].

up to $z_{\rm max} = 3$ [257]. Note that we do not include simulated WFIRST data in our likelihood analysis; instead, our goal is to use current datasets to make predictions for future experiments, such as WFIRST, that will gather SN and WL data. Finally, a *Planck*-like likelihood is also added to the total Fisher matrix. The details on how the basis was constructed can be found in equations 6 - 14 in Ref. [240]. The shape of the 5 PCs used in this work are shown in Figure 8.1.

For this work we ran three different sets of chains for each of the *All* and *Reduced* datasets with varying numbers of principal components, $N_{\rm PC} = \{1, 3, 5\}$. The choice of number of principal components is based on the following considerations. A single PC is equivalent to the usual cosmological constant model for dark energy, and serves as a baseline. Using 3 PCs is useful to contrast to results with two-parameter models for $w(z)$ (such as CPL), illustrating the consequences of such restrictive models of Dark Energy, since there is usually still a significant amount of useful information in the third PC [258]. The case of 5 PCs is chosen to showcase the effect of including PCs that have low signal-to-noise, i.e. they are relatively unconstrained by the data, on the M_ν constraints. In Section 8.3 we discuss the posteriors of the PC amplitudes to show that indeed, past

the third PC they cannot be constrained to better than ten percent.

Note that the minimum supernova redshift from the Pantheon sample imposes $z_{\rm min} = 0.01$, which means that in fact for $0 < z < z_{\rm min}$ the PCs can oscillate significantly. There is no fundamental reason why $w(z)$ must not change arbitrarily at ultra-low redshift [259], and the PCs capture that possibility.

The vector of model parameters for the chains is given by

$$\vec{\theta} = \{\Omega_c, \theta_A, \Omega_b, \tau, n_s, \ln A_s, \tau, \alpha_1, ..., \alpha_{N_{\rm PC}}, M_\nu\}. \tag{8.2}$$

Here, Ω_c is the cold dark matter density, θ_A is the angular size of the horizon at the time of recombination, Ω_b is the baryon density, τ is the reionization optical depth, and A_s and n_s are the initial curvature power spectrum amplitude and tilt. We define $M_\nu \equiv \sum m_\nu$ as the sum of the neutrino masses and assume the so-called degenerate hierarchy, where all three neutrino eigenstates are equally massive (i.e. the mass of the ith neutrino eigenstate is $m_{\nu,i} = M_\nu/3$ for $i = \{1, 2, 3\}$). Finally, the α_i parameters are the amplitudes of the PCs $e_i(z)$, as per Eq. (8.1).

To evaluate the joint effect of massive neutrinos and a model-independent equation of state for dark energy, we consider a variety of cosmological probes of geometry (the Hubble expansion rate as a function of redshift $H(z)$ and the luminosity distance $D_L(z)$) and probes of the growth (the matter power spectrum $P(k)$, the shear power spectrum P_ℓ^κ, and $S_8 \equiv \sigma_8\sqrt{\Omega_m/0.3}$, where σ_8 is the amplitude of the linear matter power spectrum at a scale of $8h^{-1}$ Mpc[1]).

Since conventions for the calculation of the shear power spectrum vary across the literature, we briefly outline the procedure used in this chapter. The shear power spectrum

[1]Generally, the parameter combination $f(z)\sigma_8$ (where $f(z)$ is the logarithmic derivative of the linear growth rate of matter fluctuations) is of interest to measure the growth of structure because it is insensitive to galaxy bias. In the ΛCDM cosmology, $f(z)$ is well approximated by $f(z) = \Omega_{\rm m}(z)^{0.545}$. Different conventions exist for the definition of S_8 however: the value of the exponent varies and sometimes Ω_m is divided by a fiducial value $\Omega_{\rm m,fid}$. Here we choose 0.5 as our exponent and use $\Omega_{m,\rm fid} = 0.3$ for a direct comparison of S_8 with values quoted in Ref. [236].

is given by

$$P_\ell^\kappa = \frac{2\pi^2}{\ell^3}\int_0^{z_s} dz \frac{W^2(z)\chi(z)}{H(z)}\Delta^2\left(k,z\right) \tag{8.3}$$

$$= \frac{2\pi^2}{\ell^3}\int_0^{z_s} dz\ F(z)\ \Delta^2\left(k,z\right), \tag{8.4}$$

where Δ^2 is the dimensionless matter power spectrum,

$$\Delta^2(k,z) = \frac{k^3 P(k,z)}{2\pi^2} \tag{8.5}$$

and

$$F(z) = \frac{W^2(z)\chi(z)}{H(z)}\frac{1}{c^3} \tag{8.6}$$

is known as the lensing weight function. Note that we use the Limber approximation [214], whereby $k \approx \ell/\chi(z)$, to evaluate the integral.

For a dark energy model that has an equation of state that is a function of redshift (and a flat universe),

$$H(z) = H_0\Big[\Omega_m(1+z)^3 + \Omega_\nu(z) + \Omega_{\rm DE} \times \exp\big[3\int_0^z d\ln(1+z')(1+w(z'))\big]\Big]^{1/2}, \tag{8.7}$$

where $\Omega_\nu(z)h^2 \approx M_\nu(1+z)^3/93.14$ eV when neutrinos become non-relativistic, χ is the comoving distance:

$$\chi(z) = c\int dz \frac{1}{H(z)}, \tag{8.8}$$

and $W(z)$ is the weight function,

$$W(z) = \frac{3}{2}\Omega_m H_0^2 g(z)(1+z). \tag{8.9}$$

Here $g(z)$ is known as the efficiency factor, and it is defined as

$$g(z) \equiv \chi(z)\int_z^\infty dz' n(z')\frac{\chi(z')-\chi(z)}{\chi(z')}, \tag{8.10}$$

where $n(z)$ is the distribution of lenses (normalized such that $\int n(z)dz = 1$). The four

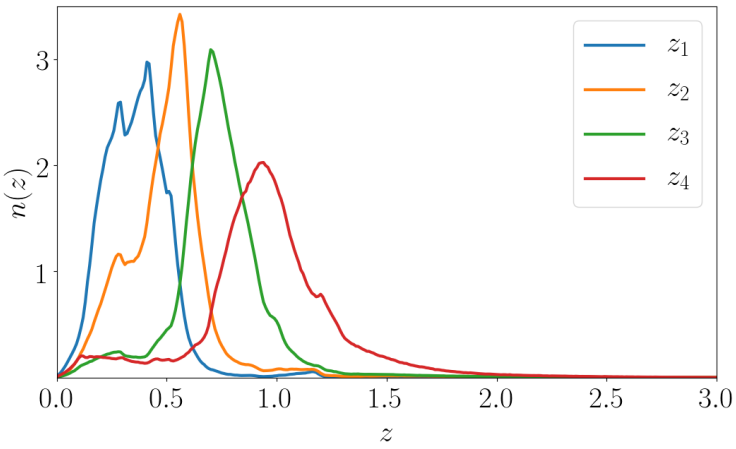

Figure 8.2: Tomographic distribution of lenses $n(z)$ in the Dark Energy Survey [250].

tomographic bins for DES [250] are shown in Figure 8.2.

Throughout the remainder of the chapter, we refer to the set of chains that vary neutrino mass and $w(z)$ constructed with PCs as νwCDM. The chains that vary neutrino mass but $w = -1$ are referred to as $\nu\Lambda$CDM, and those with the usual cosmology (i.e. with $w = -1$ and $M_\nu = 0.06$ eV) are referred to as ΛCDM.

8.2 Results

We have stated that the first three PCs are generally generally well constrained, while higher PCs are not. For our PC basis and the datasets we have employed, we find that, indeed, the first three PCs can be constrained to better than 10% while the 4th and 5th PCs cannot. We therefore emphasize that the results with 3 PCs are out fiducial results while the ones for 5 PCs should be seen as an extreme example that illustrates the effect of including modes that are ill-constrained by the data.

Tables 8.2 and 8.3 show the best-fit and 95% C.L. bounds on several cosmological parameters, including M_ν and S_8, marginalized over principal component amplitudes, for both models and datasets. For a given dataset combination, the posteriors in the νwCDM chains are significantly degraded with respect to ΛCDM, which is expected since we are marginalizing over several more parameters[2]. More important, however, is the fact that

[2]Ref. [241] points out that since the expansion rate in models where the dark energy is exclusively

the extra freedom given to the dark energy in the νwCDM chains is able to undo some of the changes induced by neutrinos, which consequently means that larger neutrino masses can be accommodated within the data: notice the difference in the νwCDM posteriors shown in Figure 8.3. As we increase the number of PCs - and thus give the dark energy more freedom - the allowed neutrino masses increase considerably, with an upper bound of $\lesssim 0.38$ (0.33) eV (95% C.L.) with 3 PCs and $\lesssim 0.55$ (0.42) eV (95% C.L.) with 5 PCs in the *All* (*Reduced*) dataset. Compare this to the results in Ref. [241], where the allowed parameter space for massive neutrinos in a cosmology with a phantom, CPL-parametrized $w(z)$ was (slightly) broader than ΛCDM but still quite limited, with $M_\nu < 0.19$ eV. Clearly, letting the behavior of dark energy be dictated by the data instead of imposing a specific parametrization can significantly open up the allowed parameter space for M_ν. Furthermore, comparing the two different dataset combinations for a given model reveals that low-redshift growth information has a slight preference for cosmologies with massive neutrinos: for both the $\nu\Lambda$CDM and νwCDM chains, M_ν is pushed upwards when going from the *Reduced* to the *All* dataset. This makes sense considering that WL surveys, such as DES, tend to have a values of S_8 that are $2-3\sigma$ lower than that of *Planck* [236], and S_8 and M_ν are anti-correlated. Note that this anti-correlation is also why in going from $\nu\Lambda$CDM to νwCDM the value of S_8 decreases.

Table 8.2: Mean and 95% C.L. errors for M_ν and S_8 in the $\boldsymbol{\nu}w$CDM chains, for the *All* and *Reduced* group of chains, and varying number of PCs. Entries with no subscript correspond to chains where only an upper bound was obtained for that parameter.

	All			*Reduced*		
$N_{\rm PC}$	1	3	5	1	3	5
M_ν [eV]	< 0.23	< 0.38	< 0.55	< 0.21	< 0.33	< 0.42
S_8	$0.81^{+0.02}_{-0.02}$	$0.81^{+0.02}_{-0.02}$	$0.80^{+0.02}_{-0.03}$	$0.82^{+0.02}_{-0.02}$	$0.82^{+0.02}_{-0.03}$	$0.82^{+0.03}_{-0.03}$

phantom is higher than that of the cosmological constant, this degrades constraints on M_ν to maintain the distance to the last scattering surface fixed (and the converse is true in non-phantom dark energy models, despite the larger parameter space). Note, however, that while this contributes to our wider posteriors to some extent, our dark energy is not forced to be phantom − and indeed Figure 8.4 shows that the fractional difference between $H(z)_{\nu w\rm CDM}$ and $H(z)_{\Lambda\rm CDM}$ is not always positive.

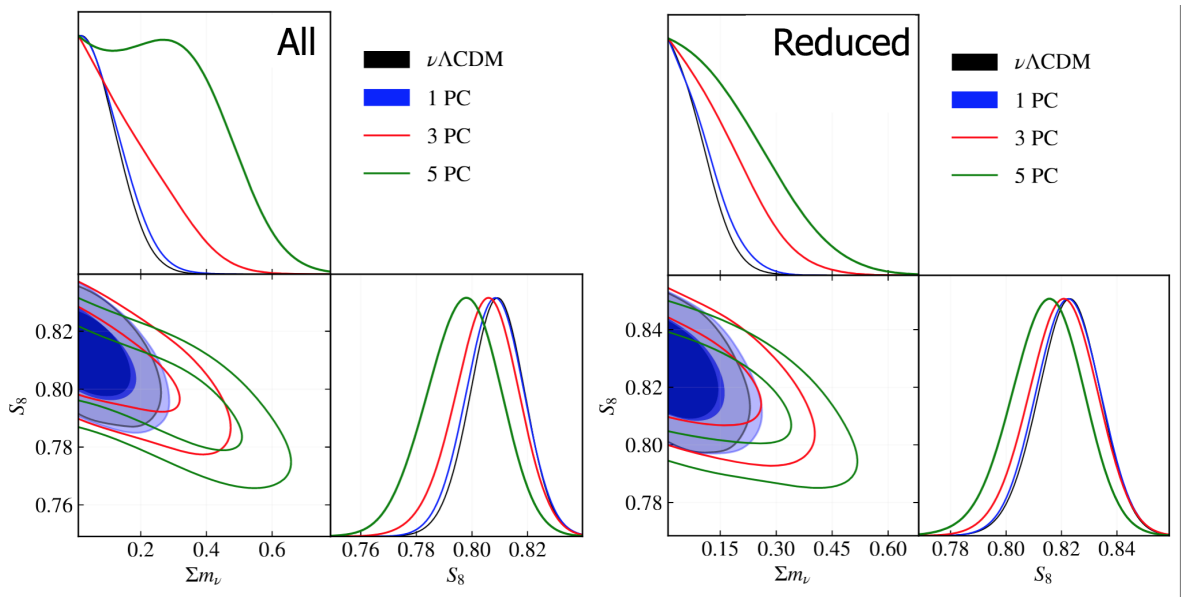

Figure 8.3: Constraints on the sum of neutrino masses and $S_8 \equiv \sigma_8\sqrt{\Omega_m/0.3}$ from the νwCDM and $\nu\Lambda$CDM scenarios. Top: *All* dataset. Bottom: *Reduced* dataset.

Table 8.3: Mean and 95% C.L. errors for M_ν and S_8 in the $\nu\Lambda$CDM chains, for the *All* and *Reduced* group of chains. Entries with no subscript correspond to chains where only an upper bound was obtained for that parameter.

	All	*Reduced*
M_ν [eV]	< 0.21	< 0.20
S_8	$0.81^{+0.02}_{-0.02}$	$0.82^{+0.02}_{-0.02}$

For the remainder of the chapter we focus on the chains with the *All* dataset; we found that the differences in the remaining posteriors between the *All* and the *Reduced* datasets were marginal, so to avoid clutter we choose to show results for the more comprehensive of the two. For reference, Appendix F.1 shows the results for the *Reduced* dataset.

Figure 8.4 shows the fractional difference of νwCDM (left) and $\nu\Lambda$CDM (right) with respect to ΛCDM for the luminosity distance $D_L(z)$ (top row) and the Hubble expansion rate as a function of redshift $H(z)$ (bottom row), for the chains with 1, 3, and 5 PCs. Note that the bands shown correspond to the 1σ confidence levels. Looking at the top left panel of Figure 8.4, it can be seen that the luminosity distance in the νwCDM and ΛCDM models agree at small redshifts ($z \lesssim 1$), but at higher redshifts there is an amplitude difference between them, reaching a $\sim 1\sigma$ difference for the case with 3 PCs, and $> 1\sigma$

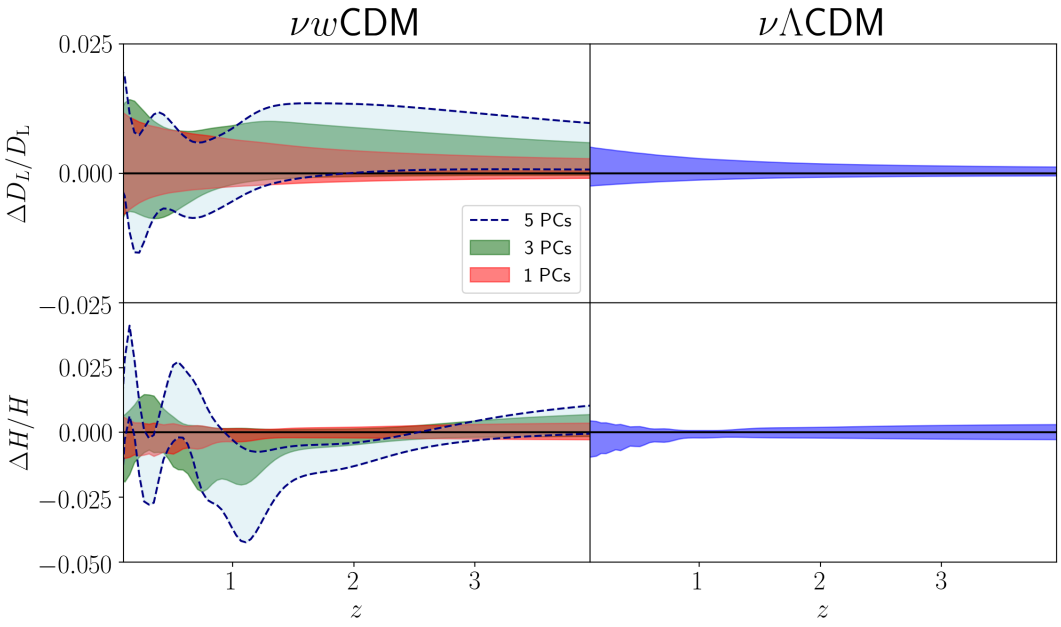

Figure 8.4: Fractional difference for the luminosity distance $D_L(z)$ (top row) and the Hubble expansion rate as a function of redshift $H(z)$ (bottom row) for νwCDM (left) and $\nu\Lambda$CDM (right), with respect to ΛCDM, for the *All* dataset. Bands correspond to 1σ confidence levels.

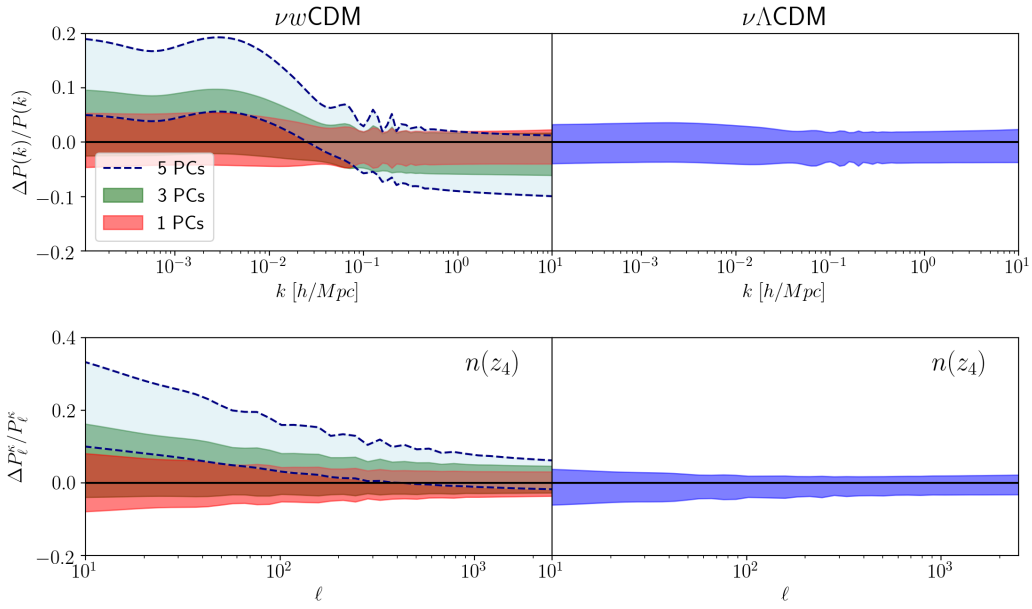

Figure 8.5: Fractional difference for the matter power spectrum $P(k)$ at $z = 0$ (top row) and the shear power spectrum P_l^κ (bottom row) for νwCDM (left) and $\nu\Lambda$CDM (right) with respect to ΛCDM, for the *All* dataset. Bands correspond to 1σ confidence levels.

with 5 PCs. This is indicative of the fact that neutrinos change the background evolution, and while at low redshifts the dynamical dark energy can counteract these changes, the same is not true for redshifts $z \gtrsim 2$, where it becomes subdominant. Furthermore, we

do not have support from BAO and SN data in this redshift range, so we find that at these redshifts the changes massive neutrinos induce in the background expansion are disfavored by the ΛCDM model. The fact that at low redshifts the dark energy is able to counteract the effects induced by neutrinos is what allows for the much larger values of M_ν in νwCDM than in ΛCDM, as shown in Figure 8.3.

This effect is also visible in the posterior for $H(z)$ (naturally, since $H(z)$ and $D_{\rm L}(z)$ are related by an integral). Between $z \approx 1 - 2.5$, the expansion rate in the νwCDM cosmology is lower than in ΛCDM, which is indicative of the dark energy behaving as a phantom component (indeed we see the $w(z)$ posteriors having values lower than -1 in this range, consistent with what is found in Ref. [260]). This behavior is driven by the dark energy attempting to appease the aforementioned "discrepancy" in the background favored by BAO data versus massive neutrinos.

Figure 8.5 shows the fractional difference in the matter power spectrum (at $z = 0$) and the cosmic shear power spectrum for the same combination of models as the ones used in Figure 8.4. On small scales we can see the characteristic suppression of power on scales below the free-streaming length of neutrinos ($\mathcal{O}(10^{-2}\ h/\text{Mpc})$ for the allowed neutrino masses in the νwCDM chains) in the matter power spectrum. Since quite large neutrino masses can be accommodated as we increase the number of PCs, this suppression becomes more marked with increasing PC number. Furthermore, on large scales (small k), there is an overall amplitude shift upwards, with the shift constituting a $> 1\sigma$ deviation from ΛCDM for the chains with 5 PCs. This is expected since the amplitude of the power spectrum increases as we increase the amount of dark energy in the universe (or decrease the amount of matter).

For the shear power spectrum, we show the fractional difference in the fourth source tomographic redshift bin $n(z_4)$ shown in Figure 8.2 (it is nearly identical in all redshift bins so we chose to show a single one for clarity). The deviation away from ΛCDM of the shear power spectrum on large scales ($\ell \lesssim 100$) is larger than that of the matter power spectrum. Recall from Eq. (8.3) that the integral to calculate the shear power spectrum contains two distinct terms: the dimensionless matter power spectrum, and the

lensing weight function, which encodes the background expansion. It is this additional dependence on the background which makes the difference with respect to ΛCDM larger in the shear power spectrum than the matter power spectrum, since as discussed above, neutrinos induce large differences in the background at high redshift.

The potentially large deviation from ΛCDM suggests this observable as an exciting candidate to falsify ΛCDM, assuming weak lensing surveys can observe a wide enough ℓ range to mitigate systematics such as the multiplicity bias, which is a nuisance parameter that shifts the overall amplitude of the cosmic shear signal. Upcoming data from the Wide Field Infrared Survey Telescope (WFIRST) [261, 262], the Vera Rubin Observatory [167, 263] and Euclid [264], will soon have better constraining power at large scales, although they will only be able to reach scales of about $\ell_{\rm min} \approx 30$.

8.3 Discussion

We have studied the constraints on the sum of the neutrino masses when marginalizing over principal components of the equation of state of a dynamical dark energy component that is allowed to cross the phantom barrier $w(z) = -1$. Exploring cosmological constraints on M_ν in the context of a general dark energy scenario is necessary because the background expansion is critical to probe the sum of the neutrino masses, meaning that there is a degeneracy between the dynamics of the dark sector and our ability to provide strong constraints on M_ν.

Our ability to constrain neutrinos with cosmology is further complicated by the fact that typical SN and (most) BAO measurements correspond to redshifts where dark energy cannot be ignored ($z \lesssim 0.7$). At higher redshifts, where dark energy is a subdominant component in the universe, the ability to compensate for neutrinos with large masses is diminished, and we found that this leads to observable deviations from ΛCDM.

We investigated the effect of massive neutrinos on a variety of cosmological probes of geometry and growth within our cosmic-acceleration scenario and found that, by giving the dark energy equation of state more freedom than in traditionally-used parametriza-

tions, dark energy can undo changes in the background expansion induced by the presence of massive neutrinos at late times. This was visible in the luminosity distance, where for $z < 2$ the posteriors for the ΛCDM and νwCDM chains agreed.

This effect had two important consequences. First, much larger neutrino masses can be accommodated within the data: the upper bound of M_ν is as high as 0.38 eV (95% C.L., 3 PCs) or 0.55 eV (95% C.L., 5 PCs), when including weak lensing data. Second, at higher redshifts, where dark energy is a subdominant component in the universe and we do not have supporting data, there are large ($\sim 1\sigma$) deviations from ΛCDM in the background, since such large neutrino masses are not generally allowed by the combination of CMB and BAO data in ΛCDM (which is why analyses using combinations of these datasets find stringent upper bounds $M_\nu < 0.12$ eV). Furthermore, these large changes to the expansion history of the universe are also visible in the matter power spectrum at $z \approx 0$, where there is a large amplitude increase on large scales in νwCDM with respect to ΛCDM and, consequently, on the shear power spectrum, since it is a measure of the integrated large-scale structure along the line-of-sight.

Since BAO measurements can probe amplitude shifts in distances, our results can be seen as a compelling case to pursue the BAO signal with high-z tracers of the underlying baryonic density field [265, 266]. High-redshift Type IA supernovae can also be critical, since they can constrain changes in the shape of the luminosity distance induced by the transition between the low- and high-z behavior. These measurements are within reach in the near future. For instance, recently there was a first detection of the BAO with the Lyman-α forest, at $z = 2.4$ [267]. Furthermore, experiments like eBOSS [268] and DESI [269] will provide distance measurements at these redshifts in the very near future.

The large amplitude offset in the shear power spectrum also makes it an exciting candidate to falsify ΛCDM, particularly since future surveys like the Vera Rubin Observatory, Euclid, and WFIRST will reach very large scales, down to $\ell_{min} \approx 30$ (although this might not be enough to mitigate systematics that could hinder the use of this observable to falsify ΛCDM). As an aside, having more precise weak lensing information will be interesting due to the S_8 tension between *Planck* and other weak lensing surveys, since

to date there is no satisfying mechanism to solve it (although plenty of exotic models have been proposed); several weak lensing measurements [19, 236, 270–272] seem to have values of S_8 that are slightly lower than those of *Planck*.

Previous works have shown that allowing neutrino mass to vary when inferring cosmological parameters from weak lensing datasets lowers S_8: for instance, Ref. [236] showed that their best-fit value of S_8 was lowered with respect to ΛCDM by 0.5σ when allowing neutrino mass to vary. Adding neutrinos to *Planck* data has a similar effect, meaning that simply extending the base ΛCDM model by allowing $M_\nu \neq 0.06$ eV does not solve the tension. Here we have seen that our PC-built equation of state further reduces S_8, and would have this effect if we considered a CMB likelihood and a WL likelihood independently, meaning that the additional freedom given to the dark energy component does not solve the tension either.

Note that, conversely, previous works that used nonparametric equations of state for dark energy were optimistic about the prospect of an evolving dark energy to solve cosmological tensions: Ref. [260] found that they could mitigate the tensions in H_0 (between local H_0 and *Planck*) and Ω_m (between BOSS and *Planck*) with such an evolving dark energy model. Furthermore, Ref. [273] claimed that extending ΛCDM to include neutrinos could solve the S_8 tension; however, this paper did not consider *separately* the effect of adding neutrinos to datasets that favor high S_8 (e.g. CMB) and datasets that favor low S_8 (e.g. WL), which is an important distinction that we have made here.

Our goal for this chapter was to construct a model-independent $w(z)$ and see how the constraints on M_ν compared not only to ΛCDM but also to other works that have considered M_ν in the context of specific parametrizations for $w(z)$. As we have already mentioned, Ref. [241] found $M_\nu \leq 0.19$ eV in their phantom, two-parameter dark energy cosmology. However, Ref. [258] showed, by choosing specific parametrizations for $w(z)$ and projecting them onto the PC basis, that the first three PCs contain most of the information, which means that two-parameter models could be neglecting relevant information, leading to constraints on M_ν that are artificially tight. Our results corroborate this conclusion: we find that the first three PCs can be constrained with better than

10% precision. The crucial consequence of this is that with only 3 PCs we have already opened up the allowed parameter space of neutrinos considerably ($M_\nu \lesssim 0.4$). This puts into question the claim that current data prefers $M_\nu \ll 0.3$ eV, since we have shown that this depends quite dramatically on our assumption about the behavior of dark energy.

However, this is a double-edged sword: while having too few PCs might yield artificially tight constraints, very high PCs cannot be probed well by the SNe data. Since most of the information is contained in the first few PCs, the higher PCs are essentially given free rein to alter the background cosmology, consequently opening up the parameter space of neutrinos even more. As we point out throughout the text, the discrepancies between ΛCDM and $\nu\Lambda$CDM are significantly larger with 5 PCs than with 3 PCs. Looking at Figure 2 it is apparent that as we add more PCs to parametrize $w(z)$ the constraint on M_ν is loosened. This means it is possible for theories of $w(z)$ to fit the SNe data but increase M_ν posteriors considerably. One must therefore be cautious, and we leave results with 5 PCs as an extreme example.

Ultimately, there is no theoretical reason to favor the CPL (or similar) parametrizations over others, and models yet to be investigated might rely on higher PCs to be distinguished from ΛCDM [274], which is why pursing a model-independent approach to parametrizing $w(z)$ (and being cautious when building a basis so as to not have many unconstrained basis vectors) is an attractive alternative.

Unlike parametric forms for $w(z)$, using PCs allows us to remain agnostic with respect to what the alternative to the cosmological constant may be, since at present there is no strong theoretical or observational support for any particular exotic dark energy scenario. It is worth noting that the main caveat of the PCA method lies in the fact that it does not assess by itself the physical plausibility of the $w(z)$ shapes that are marginalized over.

Finally, it is important to keep in mind that, as we mentioned in Section 8.1, PCs oscillate quite drastically at ultra-low redshifts (outside the region of SNe support), but they do so by construction: they are unphased by physical assumptions, and are only driven by the data. Although many dark energy models are very smooth, there is no *a priori* reason to believe that dark energy models with a low-redshift oscillatory $w(z)$ are

unphysical (and in fact some models do predict such behavior [259]). If one wanted to use the PCA method but had a strong reason to rule out the low-redshift oscillations allowed by the data, one could impose a prior that would punish such behavior, thus limiting the behavior of the PCs outside the redshift range supported by the data. For example, one could impose that the PCs only have non-zero weight at $z > z_{\rm min}$.

In our implementation, where we have not imposed any prior at low-redshift to remain agnostic, the low-redshift oscillations degrade our ability to use local H_0 to constrain M_ν. We are therefore not claiming that every model of dark energy would allow $M_\nu \sim 0.5$ eV, but merely that it is still possible to generate models that could make such extreme masses compatible with the data.

Chapter 9

Flow-Based Likelihoods for Non-Gaussian Inference

There are three key ingredients when doing inference: data, a model, and a likelihood function (in the case of Bayesian inference, there is also a prior). The quality of the inferred parameters of course hinges on the quality of the data, how closely the model approximates the process that gave rise to the data, and how accurately the likelihood maps the probability of observing the data given the model. Throughout this work, we are going to be concerned with the last of these.

In data analysis across many disciplines the likelihood used for inference is often assumed to be Gaussian. Gaussian likelihoods are attractive and widespread because they are well understood and inference boils down to obtaining a covariance matrix for the data. However, in reality, it is not generally true that the underlying likelihood of a dataset is Gaussian.

There are several general points one can keep in mind when considering the applicability of a Gaussian likelihood function. One is the central limit theorem (CLT), which states that the difference between a sample mean and a true mean, normalized by the standard error, tends toward the standard normal distribution as the number of samples tends to infinity. This is generally invoked in favor of using Gaussian likelihoods, even when datasets are known to be non-Gaussian. For example, for data in Fourier space

(such as power spectra), many modes contribute at high wave numbers and thus the CLT approximately applies. Conversely, at low wave numbers few modes contribute and the CLT is not applicable. Another important point is that, in works where the covariance matrix of the data is an estimated quantity, meaning it is not known *a priori*, it becomes a stochastic object with some uncertainty. Thus, to obtain the likelihood of the data given the estimated covariance matrix, the likelihood function has to be marginalized over the true (unknown) covariance, conditioned on the estimated one. If the original likelihood is assumed to be Gaussian, this marginalization step in fact leads to a multivariate t-distribution [275]. Finally, it is also important to consider that systematic effects might introduce correlations in the data that are not Gaussian.

There are also field-specific factors that can inform the choice of likelihood function. Although the method we develop in this work is applicable for inference with any dataset, here we apply it to a cosmological one. There are several puzzling tensions in different cosmological datasets that have so far stood the test of time [19, 20, 276] that would be interesting to reconsider under a new, more accurate likelihood. For cosmological data, we can use knowledge about the physical processes that give rise to an observable to understand its Gaussianity. For instance, while some cosmological fields such as the cosmic microwave background (CMB) *are* essentially Gaussian, those that follow from nonlinear gravitational collapse − such as distributions of galaxies and cosmic shear − are highly non-Gaussian. It is therefore likely that the underlying likelihoods of such fields are themselves non-Gaussian: a nonlinear function of a Gaussian random variable is not Gaussian distributed.

Recently, several works have studied the impact of using Gaussian likelihoods to infer parameters from non-Gaussian cosmological data, showing that it can bias the posteriors of inferred cosmological parameters and underestimate uncertainties [275, 277–279]. The reason behind this is quite intuitive: the use of Gaussian likelihoods to model non-Gaussian processes becomes a source of systematic error.

One promising approach to avoid relying on a Gaussian likelihood is to use what are

called *data-driven likelihoods* (DDL): likelihoods that are learned directly from the data.[1] Nevertheless, there is not yet a unified framework to deal with non-Gaussian likelihoods. Instead, data-driven approaches rely heavily on trial and error, and tend to not generalize well in different contexts [277]. In fact, we will show that depending on the type of non-Gaussianity in the data some DDL methods might not be appropriate and thus offer little improvement upon a multivariate normal (MVN) likelihood.

Here, we use a *flow-based generative model* [280] to achieve this task. Generative models are models from which new samples can be drawn when trained, by for example learning the underlying true likelihood that gave rise to the data. Flow-based models, in particular, consist of applying a series of transformations (a flow) to a simple probability density that is easy to sample from (e.g., a univariate Gaussian) to transform it into a (theoretically) arbitrarily complex probability density function [281]. If the transformations are known, an exact form for the resulting likelihood can be obtained and sampled from. By parametrizing the transformations with neural networks, these models can be made very expressive (in this context, "expressive" refers to a model that can capture complex features in the data accurately). Flow-based models have been used to enrich variational posteriors or enhance other generative models (e.g., [281–285]) and for density estimation, to learn complicated probability distributions (e.g., governing images of human faces) [286–291]. The driving interest behind these models has been their capacity to learn complicated probability distributions, but understanding how faithfully they reconstruct the true data distribution is an open area of research [292, 293].

We seek to understand the quality of the learned likelihood function with the goal of establishing *flow-based likelihoods* (FBLs) as an extremely flexible and adaptable method to obtain data-driven likelihoods that are more accurate than other methods to deal with

[1] We detail two points with respect to this terminology. We use the term DDL to refer to likelihoods that are more flexible than a multivariate Gaussian (or other ubiquitous, simple parametric likelihoods) and do not necessarily fix the functional form at the outset. Note that, while for a MVN likelihood the covariance matrix can be estimated from data too, the functional form of the likelihood is otherwise fixed. The second point is what is meant by "data". We only have one universe, so in cosmology we often have to rely on mock (simulated) realizations of the universe to assess the fit of a model to some data. Because the method presented in this chapter is applied in the context of cosmology, we will technically be using *mock data*-driven likelihoods. But because our method is equally applicable in any field, one could imagine disciplines where experiments are repeatable and thus the likelihoods are indeed learned from real data.

known (or suspected) non-Gaussianities in datasets. FBLs have appeared in the machine learning and cosmology/astrophysics literature in other settings, such as in likelihood-free inference (e.g., DELFI) [294–298] and simulation-based inference [21, 299–301]. We demonstrate that FBLs can match the underlying likelihood of the data very well, even in high-dimensional spaces, by comparing the non-Gaussian features of the generated samples and the data. Furthermore, we show that this is not trivial: under certain conditions, the samples generated can be indistinguishable from the original data and yet the likelihood can be significantly biased and imprecise. We apply the FBL to mock weak lensing power spectra, which we show have strong non-Gaussianities. This observable is representative of many observables in cosmology for which, given a set of cosmological model parameters, we can generate sophisticated mock data with forward simulations (including instrumental effects, selection biases, and other systematics) but we cannot write down a tractable likelihood. Along the way we create a thorough pipeline, incorporating some elements from previous works [277, 278], which can be applied to any dataset to analyze the presence (or absence) of non-Gaussianities. With this pipeline we are also able to gain some insight into why some data-driven methods work for some datasets and not for others. Seeing that FBLs are capable of overcoming barriers that other data-driven methods can face depending on the type of non-Gaussianity in the data further emphasizes the advantage of using FBLs for inference. *We emphasize that no aspect of the NG pipeline nor FBLs are specific to cosmology and therefore these methods can be applied in any context.*

In Section 9.1 we briefly present two data-driven likelihood methods, Gaussian mixture models (GMMs) and independent component analysis (ICA), which we use to compare to our FBL. We also outline the basic principles behind normalizing flows and describe FFJORD [280], a flow-based generative model that uses ordinary differential equations to evolve the initial probability density. In Section 9.2 we present the exhaustive tests we carry out to characterize non-Gaussianity in the data. In Section 9.3 we introduce the observable to which we apply our flow-based likelihood (FBL), the weak lensing convergence power spectrum, and finally we show our results applying the FBL to this

data in Section 9.4. We discuss our results and conclude in Sections 9.5 and 9.6. All the code used in this work will be made available upon publication.

Throughout the remainder of the chapter, we will use the terms Gaussian, multivariate Gaussian, and multivariate Normal interchangeably. Furthermore, we will abbreviate "non-Gaussian" (and its derivatives) as NG when convenient.

9.1 Data-Driven Likelihoods

A probability density function (PDF) can be estimated by drawing sufficient samples from it. This is the key idea behind data-driven likelihoods. With access to sufficient (mock) data, it is not necessary to impose a restrictive functional form for the likelihood *a priori*; instead, we can think of independent (mock) catalogs as independent draws from the underlying true likelihood function, and can use this fact to estimate the likelihood directly from the samples. The benefit of this method is that, if an appropriate method for density estimation is used, the reconstructed likelihood can take into account any non-Gaussian features in the data.[2]

We refer to likelihoods learned from the data as data-driven likelihoods (DDL). In this work, we implement three different DDLs. The first two, used as our baseline, were chosen to serve as a direct comparison of the work in Ref. [277], which studied NG in large-scale structure data. These are Gaussian mixture models and independent component analysis. We briefly summarize key aspects of each of these in the section below, but due to their ubiquity in the literature we provide references for more detailed explanations. The third method is what we refer to as flow-based likelihoods (FBL): we propose using the likelihood learned by flow-based generative models as a DDL. We discuss this in detail below.

[2]We are using the terms PDF and likelihood interchangeably, although they are different objects. We explain this in Appendix G.1.

9.1.1 Gaussian Mixture Model

As its name indicates, a Gaussian mixture model (GMM) is simply a convex combination (mixture) of multiple Gaussians with unknown means and covariances, each with a weight that determines their contribution to the PDF. In a multidimensional setting, the final PDF for a vector $\mathbf{x}$ in a GMM with K components can be written as

$$\hat{p}_{\text{GMM}}(\mathbf{x}) = \sum_{i=1}^{K} \phi_i \mathcal{N}(\mathbf{x}|\mu_i, \Sigma_i), \tag{9.1}$$

where μ_i, Σ_i and ϕ_i are the mean, covariance, and weight of the ith Gaussian in the mixture, respectively, and $\mathcal{N}$ is the multivariate normal PDF. The weights are normalized such that $\sum_{i=1}^{K} \phi_i = 1$. The number of parameters in these models is given by

$$K\left(d + \frac{1}{2}d(d+1)\right) + K, \tag{9.2}$$

where d is the dimension of the data vector $\mathbf{x}$. This is because for each of the K components we have to learn a d-dimensional mean vector and a covariance matrix with $\frac{1}{2}d(d+1)$ degrees of freedom (since it is positive semidefinite)[3]. We also have to learn K weights, one for each component. Notice that for $K = 1$ we recover a standard MVN likelihood.

We use the `scikit-learn` [302] implementation of GMMs, which uses expectation maximization [303] to estimate the model parameters, and the Bayesian information criterion (BIC) to decide how many components to include in the mixture. This method of model selection considers the maximum likelihood of a model while penalizing model complexity. We refer the reader to Ref. [304] for a thorough explanation of GMMs and the expectation maximization procedure.

[3]The number of parameters can also be decreased by putting additional constraints on the covariance matrix, such as having different components share a covariance matrix or forcing it to be diagonal, but we will use full covariances throughout this chapter.

9.1.2 Independent Component Analysis

Independent Component Analysis (ICA) is typically used to isolate linear mixtures of independent sources. For some observed data vector $\mathbf{x}$, this amounts to the assumption

$$\mathbf{x} = \mathbf{A}\mathbf{s}, \tag{9.3}$$

where $\mathbf{A}$ is an unknown matrix, called the mixing matrix, that mixes the sources $\mathbf{s}$. The goal of ICA is to solve for $\mathbf{A}$, which is done by actually solving for its inverse, $\mathbf{A^{-1}} = \mathbf{W}$,

$$\hat{\mathbf{s}} = \mathbf{W}\mathbf{x}, \tag{9.4}$$

such that $\hat{\mathbf{s}} \approx \mathbf{s}$.

ICA solves for $\mathbf{W}$ by breaking it up into three different linear operations. The first one is decorrelating the data (i.e. projecting the data onto the principal components), an operation usually referred to as principal components analysis. The data is then normalized, and finally a rotation matrix is solved for such that the statistical independence of the sources is maximized. We refer the readers to Ref. [305] for details on how this maximization is carried out. We use the `scikit-learn` implementation of the ICA algorithm.

For our purposes, what matters is that ICA provides a way of turning the d-dimensional likelihood for $\mathbf{x}$ into d one-dimensional likelihoods by converting $\mathbf{x}$ into d independent components[4]:

$$\hat{\mathbf{s}} \equiv \mathbf{x}_{\rm ICA} = \mathbf{W}\mathbf{x} = \{\mathbf{x}_{1,\rm ICA}, ..., \mathbf{x}_{N,\rm ICA}\}, \tag{9.5}$$

where N is the number of ICA components. In this work, we set $N = d$, although in principle ICA can also be used for dimensionality reduction.

The one-dimensional probability density of each component, $\hat{p}_n$, can then be estimated using a Kernel Density Estimator (KDE). We use a Gaussian kernel, with a standard

[4]Recall that statistical independence requires that neither second-order nor higher-order correlations exist.

deviation[5] given by Scott's bandwidth [306]. Finally, the likelihood for $\mathbf{x}$ is a factorial distribution, a product of the N independent PDFs:

$$\hat{p}_{\text{ICA}}(\mathbf{x}) = \prod_{n=1}^{N} \hat{p}_n(\mathbf{x}). \tag{9.6}$$

9.1.3 Flow-Based Likelihood

Flow-based Generative Models

For a given data vector $\mathbf{x}$, the goal of generative models is to learn the distribution that $\mathbf{x}$ is drawn from: $\mathbf{x} \sim p^*(\mathbf{x})$, where the asterisk denotes the true underlying distribution. Once an estimate of p^* is obtained, new samples of $\mathbf{x}$ can be generated, thus the name of this class of models.

In flow-based methods, the generative process starts by drawing a sample $\mathbf{z}$ from a (simple) probability distribution that has a tractable PDF and is easy to sample from (such as a univariate Gaussian),

$$\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z}). \tag{9.7}$$

This prior distribution then undergoes a "flow", which means that it is transformed repeatedly while conserving its probability. The optimization process relies on finding a series of transformations such that the resulting distribution approaches $p^*(\mathbf{x})$.

Consider a series of transformations that are bijective, so that the relationship between $\mathbf{z}$ and $\mathbf{x}$ can be summarized as

$$\mathbf{x} \equiv \mathbf{h}_0 \underset{g_K}{\overset{f_1}{\longleftrightarrow}} \mathbf{h_1} \underset{g_{K-1}}{\overset{f_2}{\longleftrightarrow}} ... \underset{g_2}{\overset{f_{K-1}}{\longleftrightarrow}} \mathbf{h_{K-1}} \underset{g_1}{\overset{f_K}{\longleftrightarrow}} \mathbf{h_K} \equiv \mathbf{z}, \tag{9.8}$$

such that $g = g_1 \circ g_2 \circ ... g_K$, $f = f_1 \circ f_2 \circ ... f_K$ and $\mathbf{z} = f_\theta(\mathbf{x}) = g_\theta^{-1}(\mathbf{x})$. If the transformation $f_\theta(\mathbf{x})$ is learned from the data, then the invertibility criterion ensures that, after drawing a sample of $\mathbf{z}$, f_θ can be inverted to generate new samples of $\mathbf{x}$. Furthermore, notice that enforcing the transformations to be bijective enforces volume preservation (and unlike

[5] More generally, kernels are controlled by a parameter called the bandwidth. It decides how much to smooth each data point, and therefore controls the bias-variance trade-off. For the case of a Gaussian kernel, the bandwidth is the same as the standard deviation.

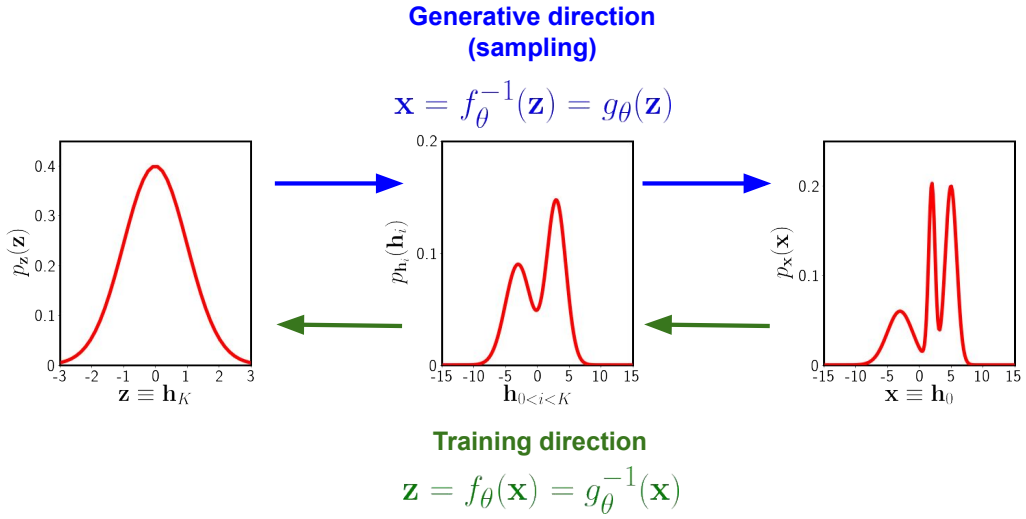

Figure 9.1: Schematic depiction of a flow-based generative model in one dimension. The left-hand panel shows a standard Normal, which is chosen as the base distribution $p_\mathbf{z}(\mathbf{z} \equiv \mathbf{h}_K)$, following the notation in Eq. (9.8). The right-hand panel shows the target distribution, i.e. the distribution of the data $p_\mathbf{x}(\mathbf{x} \equiv \mathbf{h}_0)$, which is visibly more complex than the prior. The middle panel shows the distribution at an intermediate point in the flow, $\mathbf{h}_{0<i<K}$. The green arrows reflect the direction in which training takes place: the model is fed the data $\mathbf{x}$, which then undergoes the flow into the latent space. The blue arrows reflect the generative, or sampling, direction, whereby a sample is drawn from the prior distribution and undergoes the inverse flow to generate a data sample.

other generative models such as variational autoencoders [145], the dimension of the latent space and the data is the same). In theory, if the transformations are expressive enough, $p_\mathbf{z}$ can be turned into an arbitrarily complex distribution. To this end, the transformations tend to involve (invertible) neural networks. Figure 9.1 depicts how flow-based generative models work in a simplified one-dimensional setting.

Ultimately, Eq. (9.8) is simply a concatenated sequence of changes of variables, so the log-density of the final distribution can be written as a sum of the log-PDF of the original distribution plus the sum of the log-determinant of the Jacobian matrix of each transformation:

$$\log p_\mathbf{x}(\mathbf{x}) = \log p_\mathbf{z}(\mathbf{z}) + \log \left| \det \left(\frac{d\mathbf{z}}{d\mathbf{x}} \right) \right| \tag{9.9}$$

$$= \log p_\mathbf{z}(\mathbf{z}) + \sum_{i=1}^{K} \log \left| \det \left(\frac{d\mathbf{h}_i}{d\mathbf{h}_{i-1}} \right) \right|. \tag{9.10}$$

The loss function used for training is then simply the negative log-likelihood over the

entire training set $\mathcal{D}$:

$$\mathcal{L} = -\frac{1}{|\mathcal{D}|} \sum_{\mathbf{x} \in \mathcal{D}} \log p_\theta(\mathbf{x}), \tag{9.11}$$

and the model can be trained using stochastic gradient descent.

There are two main catches to flow-based models: one, the transformations must be easily invertible, and two, they must have an easy-to-compute Jacobian determinant, since computing the determinant of a matrix scales as n^3 for an $n \times n$ matrix.

Models in the literature have tackled these issues in different ways: the form of the transformation can be restricted such that determinant identities can be exploited [281]; the models can be made autoregressive, such that the Jacobian is lower triangular [307]; or dimensions can be split up and affine transformations used such that the Jacobian is easy to compute [287–289]. All these methods have their own advantages and drawbacks, but they all have one feature in common: they sacrifice model expressivity to improve the speed of the Jacobian determinant computation.

FFJORD: Free-Form Jacobian of Reversible Dynamics

Due to the limited expressivity of the aforementioned flow-based models, we opt for FFJORD [280],[6] which replaces the transformation function with an integral of continuous-time dynamics, giving rise to continuous normalizing flows (CNFs) [308]. We summarize key details of FFJORD in this section but refer the reader to Refs. [280, 308] for additional details on the derivation and optimization procedure.

For FFJORD, the transformation from prior to data is seen as an evolution in time. Let us define the observable data as $\mathbf{z}(t_1)$ and the original sample drawn from the prior as $\mathbf{z}_0 = \mathbf{z}(t_0)$. If we then define the time evolution of $\mathbf{z}$ as an ordinary differential equation (ODE)

$$\frac{\partial \mathbf{z}(t)}{\partial t} = f(\mathbf{z}(t), t; \theta), \tag{9.12}$$

(where f can be a neural network and θ its parameters) we can obtain $\mathbf{z}(t_1)$ by solving the ODE subject to the initial condition $\mathbf{z}_0 = \mathbf{z}(t_0)$. The change in log-density is given

[6] `https://github.com/rtqichen/ffjord`

by the instantaneous change of variables formula:

$$\frac{\partial \log p(\mathbf{z}(t))}{\partial t} = -\mathrm{Tr}\left(\frac{\partial f}{\partial \mathbf{z}(t)}\right), \tag{9.13}$$

and thus the total change in log-density can be obtained by integrating across time:

$$\log p(\mathbf{z}(t_1)) = \log p(\mathbf{z}(t_0)) - \int_{t_0}^{t_1} \mathrm{Tr}\left(\frac{\partial f}{\partial \mathbf{z}(t)}\right) dt. \tag{9.14}$$

Optimizing Eq. (9.14) is non-trivial, and requires a continuous-time analog to back-propagation. It can nevertheless be combined with gradient-based optimization methods to fit the parameters θ. In general, computing the trace of this transformation scales as $\mathcal{O}(n^2)$, but by using an unbiased stochastic estimator of Eq. (9.14) Ref. [280] decreases [7] the complexity further to $\mathcal{O}(n)$. This makes FFJORD scalable without having to constrain the Jacobian, yielding a very expressive model.

A key point we want to make is that the focus in works that employ these model is generally on their capacity as generators, or their abstract capacity to increase the complexity of distributions (e.g. as variational posteriors). The log-likelihood is used as the training objective, and the model quality is for example judged by the quality of the samples produced (often in a qualitative fashion). Here, together with quantitatively discussing sample quality, we zero-in on the quality of the likelihood itself (showing that sample likelihood is not necessarily indicative of likelihood quality, Appendix G.2), with the goal of establishing FBLs as very accurate and versatile DDLs that can themselves be used for inference.

9.2 Measuring Non-Gaussianity

Quantifying the non-Gaussianity of a dataset is the first step in understanding the applicability or shortcoming of applying a Gaussian likelihood for inference. We propose carrying out three different tests to quantify the NG of a given dataset, which rely on diagnosing deviations from the null hypothesis that the underlying likelihood is a MVN,

[7] It is shown that the variance of the log-likelihood induced by the estimator is less than 10^{-4}.

and act at different "resolutions":

1. t-statistic of skewness and excess kurtosis of each bin, which quantifies the NG of bins independently;

2. transcovariance matrix [278], which considers the NG of all pairs of bins;

3. the Kullback-Leibler (KL) divergence [309] of all the data with respect to a MVN distribution.

Notice that the latter two tests are sensitive to all higher-order correlations in the data. The different scopes of these three tests culminate in a very holistic indicator of NG when combined. As we will show in subsequent sections, the pipeline's tiered scope is capable of shedding light into why some DDLs work in some settings and not others, with the important consequence that methods that have applied DDLs in the past may have failed to capture non-Gaussianities adequately.

9.2.1 t-statistic of Skewness and Excess Kurtosis

By having many mocks of a given observable, we can obtain an estimate of the distribution for each bin (e.g. the power spectrum at a specific multipole number). We then calculate the t-statistic of skewness and excess kurtosis of each to diagnose a deviation away from Gaussianity (under the Gaussian assumption, the null hypothesis is zero skewness and excess kurtosis). Henceforth, for conciseness we refer to the excess kurtosis simply as kurtosis.

Recall that the t-statistic is basically a measure of how many standard deviations away from the null hypothesis a given measure is. For a given parameter β,

$$t = \frac{\hat{\beta} - \beta_{\text{null}}}{\text{SE}(\hat{\beta})}, \tag{9.15}$$

where $\hat{\beta}$ is the estimated value of β, β_{null} is the value of β under the null hypothesis, and $\text{SE}(\hat{\beta})$ is the standard error of $\hat{\beta}$. t is thus a dimensionless quantity that measures

the deviation of the estimated parameter away from the null hypothesis in units of the standard error.

For our purposes, β is going to be either the skewness or the kurtosis of a sample, and $\beta_{\text{null}} = 0$ under the null hypothesis that the samples are drawn from a Gaussian distribution. The variance of the skewness of a random sample of size n from a Normal distribution is

$$\widehat{\text{Var}}_{\text{skew}} = \frac{6n(n-1)}{(n-2)(n+1)(n+3)}, \tag{9.16}$$

and the variance of the kurtosis

$$\widehat{\text{Var}}_{\text{kurt}} = \frac{24n(n-1)^2}{(n-3)(n-2)(n+3)(n+5)}. \tag{9.17}$$

Taking $\text{SE} = \sqrt{\widehat{\text{Var}}}$, we can thus obtain the t-statistic for the skewness and kurtosis of each bin. One thing to notice is that the t-statistic is an extensive quantity: it depends on the number of data points in a bin (or, equivalently, the number of mocks).

9.2.2 Pairwise Non-Gaussianity of Data Points

Following Ref. [278], we use the basic observation that a sum of two independent Gaussian random variables should itself be a Gaussian random variable – if $x_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$ and $x_j \sim \mathcal{N}(\mu_j, \sigma_j^2)$, then $x_i + x_j \sim \mathcal{N}(\mu_i + \mu_j, \sigma_i^2 + \sigma_j^2)$, where $\mathcal{N}(\mu, \sigma^2)$ is the normal distribution with mean μ and standard deviation σ[8] – to test the pairwise non-Gaussianity of data bins.

To perform this test, Ref. [278] considers the pairwise sum of all the bins in a given observable. Consider an ensemble of N realizations of a d-dimensional observable $\mathbf{x}$. Denoting the d elements of the ith data vector as x_i^u, where $i \in [1, N]$ and $u \in [1, d]$, for

[8]Ref. [278] also tests two other conditions, regarding the product and quotient of normal random variables. The expected distributions are a superposition of χ^2 random variables and the Cauchy distribution, respectively. These are both very sharply peaked distributions and we found that estimating their density with a limited number of mock samples was unreliable, so we only show results for the sum.

each pair x_i^u, x_i^v with $u \neq v$ it is straightforward to obtain the sum

$$s_i^{u,v} = x_i^u + x_i^v. \tag{9.18}$$

For each pair of bins (u, v), there are N samples of the distribution $s_i^{u,v}$.

The N samples are grouped into b bins of a histogram $\mathcal{H}_b$. Under the Gaussian assumption, if $N \to \infty$ and $b \to \infty$ the histogram will tend to a Gaussian distribution. The deviation from the estimate of the density of $s_i^{u,v}$ and the expected normal distribution can be calculated using the mean integrated square error (MISE):

$$\frac{1}{b}\sum_{a=1}^{b}[\mathcal{H}_a(s_i^{u,v}) - \mathcal{N}]^2 \equiv S_{u,v}^{+}. \tag{9.19}$$

Each pair of bins can be pre-processed by mean-subtracting and whitening the 2×2 covariance matrix, to destroy all the Gaussian (second-order) correlations in the data, such that any remaining correlations are necessarily of non-Gaussian origin (in practice, this means diagonalizing the covariance matrix; we also normalize the variance of each dimension such that the final covariance matrix is the identity matrix). After these two steps each bin should be a draw from a standard univariate normal $\mathcal{N}(0, 1)$ if they were originally truly Gaussian, and consequently $s_i^{u,v} \sim \mathcal{N}(0, 2)$.

By finding the MISE for each pair of points, Ref. [278] builds a *transcovariance* matrix S^+: while covariance matrices measure Gaussian correlations between pairs of parameters/bins, the transcovariance matrix measures non-Gaussian correlations. The total contamination of the uth data point is then simply the sum over a column of the matrix:

$$\epsilon_u^{+} = \sum_{v \neq u} S_{u,v}^{+}. \tag{9.20}$$

We carry out this procedure with one importance difference: instead of using a histogram of data points, we use a kernel density estimator instead. We prefer this methodology because while the density estimate of a histogram is strongly dependent on the

number of bins, and thus applying Eq. (9.19) requires finding the value of b that minimizes $S^{+}_{u,v}$, the KDE is insensitive to this (although it does have other tuning parameters, as discussed in Section **??**). We therefore modify Eq. (9.19):

$$\frac{1}{b}\sum_{a=1}^{b}[\mathcal{K}_a(s_i^{u,v}) - \mathcal{N}(0,2)]^2 \equiv S^{+}_{u,v}, \tag{9.21}$$

where $\mathcal{K}(\cdot)$ denotes the KDE, and here b is simply the number of discrete values at which we estimate the KDE and the normal distribution. We use a Gaussian kernel with standard deviation given by Scott's bandwidth. Just like for the NG test above, we note that ϵ^{+} is extensive, since it relies on summing over the columns of a matrix whose dimension depends on the number of data bins.

9.2.3 Nonparametric Kullback-Leibler (KL) divergence

Following Ref. [277], we use a nonparametric estimator of the KL divergence to quantify the non-Gaussianity in a dataset. The KL divergence is a well-known measure of the (dis)similarity between two PDFs p and q:

$$D_{n,m}(p||q) = \int p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})} d\mathbf{x}. \tag{9.22}$$

For cases in which p and q are unknown and we instead just have ensembles of draws from unknown distributions, Ref. [310] derived an unbiased estimator of the KL divergence that essentially relies on estimating the probability density using k-nearest neighbors (kNN). Consider two densities p and q, defined on $\mathcal{R}^d$, and independent and identically distributed (i.i.d.) d-dimensional samples $\{X_1, X_2, ...X_n\}$ and $\{Y_1, Y_2, ...Y_m\}$ drawn from each, respectively. Letting $\rho_k(i)$ be the Euclidean distance between X_i and its kNN in $\{X_j\}_{j\neq i}$, and $\nu_k(i)$ the distance from X_i to its kNN in $\{Y_j\}$, the KL divergence estimator can be written as:

$$\hat{D}_{n,m}(p||q) = \frac{d}{n}\sum_{i=1}^{n} \log \frac{\nu_k(i)}{\rho_k(i)} + \log \frac{m}{n-1}. \tag{9.23}$$

Ref. [277] proposed estimating non-Gaussianity in an ensemble of mocks by comparing

the nonparametric KL divergence between samples of a mock observable and samples drawn from a MVN with mean and covariance taken from the mocks. If the true likelihood that gave rise to the data were Gaussian, then the KL divergence between these two sets of samples would vanish, while deviations away from zero would indicate the presence of NG in the data.

Despite the theoretical appeal of this test, we find that for the number of data samples and bins we use, it is not quite as robust as the other two. This is due to the curse of dimensionality: kNN-based algorithms struggle in high-dimensional spaces due to the fast increase of volume with increasing dimensions, which makes the data sparse.

In Appendix G.4 we show the variability of the KL divergence estimate when comparing two MVN distributions for different random seeds. We find that two sample distributions drawn from the same likelihood can have as little as $\sim 20\%$ overlap between them. So when showing results for this test, we will keep this lower bound in mind to judge whether a deviation can be due to NG or simply due to random error. Note that, despite these limitations, we include this test because we find that for data with significant NG (like the one used in this chapter) the KL divergence estimate is much larger than the random scatter.

9.3 Mock Weak Lensing Convergence Power Spectra

9.3.1 Background

In the standard cosmological picture, infinitesimal initial fluctuations in the matter density field evolved through gravitational collapse in a highly non-linear fashion to yield all the structure we have in the universe today. Mapping the distribution of matter on cosmic scales is non-trivial because $\sim 85\%$ of all the matter is dark: it does not interact with light and therefore we cannot observe it directly. One approach to do so is to use weak gravitational lensing (WL): as photons from far-away background sources (such as galaxies) travel towards us, the cosmic web itself acts as a gravitational lens, distorting their paths and thus distorting the shape of the sources.

Galaxies that are nearby will be lensed by similar structure so they will have correlated shapes. WL galaxy surveys look at millions of galaxies [311] to find statistical correlations in their shapes. The large number of galaxies is needed both due to the intrinsic shape noise of galaxies as well as the weakly correlated nature of the underlying signal. These surveys construct shear maps, which can be used to reconstruct the projected distribution of matter in the universe between us and the sources. The resulting surface mass density maps are referred to as convergence maps. By providing a direct view into the distribution of dark matter across cosmic times, convergence maps can constrain cosmological parameters, the halo mass function, and can be cross-correlated with images at other frequencies to learn about halo bias and dependence of astrophysical processes on the dark matter density.

9.3.2 Mock Data

In this work, we focus on the weak lensing convergence power spectrum. Like many other observables in cosmology that are the product of the highly non-linear process of structure formation on small scales, writing down a tractable (and correct) likelihood function for weak lensing observables is not trivial. Instead, we have to rely on simulations: given a set of cosmological parameters that we seek to constrain, we can run complex forward simulations to generate mock data, which we can then use to infer cosmological parameters from real data.

In Appendix G.6 we detail how we obtain the mock convergence maps and power spectra used in this work. We summarize some key details here. We obtain 75,000 mock maps by running four N-body simulations with different initial seeds and using `LensTools` [312] to generate convergence maps from them. Each map has a sky coverage of 3.5 $\times$ 3.5 deg and 1024 $\times$ 1024 pixels. We set the density of matter $\Omega_m = 0.3$, the density of dark energy $\Omega_\Lambda = 0.7$, the density of baryons $\Omega_b = 0.046$, the variance of matter fluctuations $\sigma_8 = 0.8$, the scalar spectral index $n_\mathrm{s} = 1$, and the Hubble constant $H_0 = 72$ km/s/Mpc. We obtain the convergence power spectrum in 34 different multipole bins uniformly distributed in log space for $\ell = [100, 5000]$ (past $\ell = 5000$ the calculated

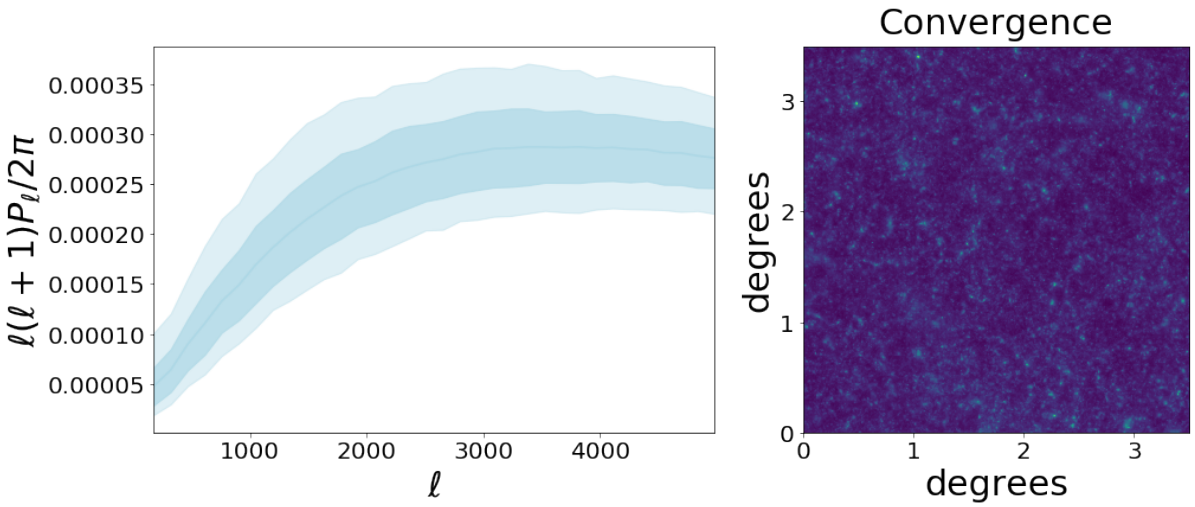

Figure 9.2: *Left*: mock convergence power spectra as a function of multipole number ℓ. *Right*: mock convergence map with a sky coverage of 3.5×3.5 deg. and 1024×1024 pixels.

power spectrum deviates from theory significantly [313]). Figure 9.2 displays an example of a mock convergence map (right) and the 68% CL and 95% CLs of the convergence power spectra in blue (left).

We point out that for the results presented in this chapter we do not include observational effects in the maps such as noise or filtering. The reason for this is twofold. First, as we will show below, the NG is strongest at the lowest multipoles, where even a pessimistic amount of noise level would not Gaussianize the data. Furthermore, while adding noise could Gaussianize high-ℓ bins, the objective of this work is to analyze the capacity of different DDLs to capture non-Gaussianities, and thus how strong the non-Gaussianities have to be in order to be adequately picked up, so we opt against introducing noise. In follow-up work where we study the impact of using FBLs on parameter inference, we take into account observational effects that would actually be present in the data.

9.3.3 Weak Lensing Likelihood Non-Gaussianity

If a field is decomposed into spherical harmonics, with coefficients $a_{\ell,m}$, then an angular power spectrum bin at a given ℓ is given by:

$$P_\ell = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} |a_{\ell,m}|^2. \tag{9.24}$$

The transformation into $a_{\ell,m}$ space is linear, so Gaussian fields will have Gaussian-distributed $a_{\ell,m}$s. Because each bin is a sum over quadratic combinations of Gaussian-distributed variables, power spectrum bins are governed by Gamma distributions. The applicability of a Gaussian likelihood therefore decreases with decreasing multipole number (the sum is over less modes as ℓ decreases), which is precisely why CMB analyses only apply a Gaussian likelihood at high multipoles [314]. In WL analyses this distinction is not commonly made, and there are two additional effects that are important. The underlying field is itself non-Gaussian, meaning that the $a_{\ell,m}$s are not necessarily Gaussian-distributed. Furthermore, the fact that galaxies are discrete tracers (as opposed to a smooth random field like the CMB) also increases the skewness [315]. Finally, recall that, compared to the CMB, WL observables have the added complexity of requiring computationally-expensive forward-modelling.

Ref. [278] showed that there are significant non-Gaussian correlations in the cosmic shear correlation functions (the real-space analogue of the power spectrum) for the weak lensing survey CFHTLenS [316]. Furthermore, they showed that because the one-point correlation statistics are skewed, WL datasets are likely to lead to a systematically low lensing amplitude. Since the WL amplitude increases with Ω_m and σ_8, the authors suggest that this could (partially) explain the discrepancy in the value of $S_8 = \sigma_8\sqrt{\Omega_m}$ between WL and CMB surveys (the so-called S_8 tension [19]). Finally, they show that the non-Gaussianities become more relevant on larger angular scales, meaning that this issue will be more relevant for upcoming wide-angle surveys such as *Euclid* and the Vera Rubin Observatory. This can be understood as a break of the CLT on these scales. These results were confirmed and expanded in Ref. [315].

For our suite of NG tests (Section 9.2), we use the power spectra from an ensemble of 2,048 maps instead of the full set. The reason why we use this seemingly arbitrary number of mocks is to have a direct point of comparison to the caliber of non-Gaussianity discussed in the context of galaxy power spectra in Ref. [277]: there, the authors only have access to 2,048 mocks for their observable. As we discussed above, our measures of non-Gaussianity are extensive (both in terms of number of bins and in terms of number of

mocks), so an "apples to apples" comparison between two observables requires the same number of mocks and bins. We relegate the details of this comparison to Appendix G.5, but in short, we find that the WL power spectrum is significantly more non-Gaussian than the galaxy power spectrum at the scales considered, and its potential to have a larger effect on biasing inferred parameters makes it a more exciting target of data-driven likelihood methods.

Finally, we pre-process the mock power spectra by subtracting off the mean and whitening them using the Cholesky decomposition of the precision matrix: $\Sigma^{-1} = \mathbf{L}\mathbf{L}^{\mathrm{T}}$, where Σ is the covariance matrix and the superscript T denotes the transpose. The data is whitened by applying the linear transformation $\mathbf{L}$ on the mean-subtracted mock data matrix. It can subsequently be unwhitened by applying the inverse transformation.

The results of the NG tests are shown in Figure 9.3. The strong non-Gaussian signatures are apparent in all three tests. The top row shows the absolute value of the t-statistic for the distribution of each bin as red crosses. The dark (light) gray shaded region is the 1σ (2σ) confidence level (CL) obtained from 2,048 Gaussian realizations drawn with mean and covariance extracted from the mock power spectra. As expected, the Gaussian CL matches a t value of 1 (2). Clearly, the strongest NG correspond to the largest scales (lowest ℓ), but notice that for nearly all bins the t-statistic is significantly larger than 2σ, and even 3σ. In Appendix G.6.2 we show the individual distributions for each bin, together with each KDE fit, which make the statistically-significant skewness and kurtosis visible for many of them.

The middle row shows the S^+ matrix for the mock data (left) and for Gaussian samples drawn with mean and covariance given by the mock data (center). The stark difference between both panels is easily visible by eye. The sum of each column in S^+, ϵ^+, is shown on the right. The red crosses represent the mock data while the gray circles correspond to the Gaussian samples. There is a non-negligible gap between both, and in fact the mean of the red crosses (shown as a horizontal red bar) is ~ 8 times greater than the mean for the Gaussian mocks.

Finally, the histograms in the bottom row reflect the nonparametric KL divergence

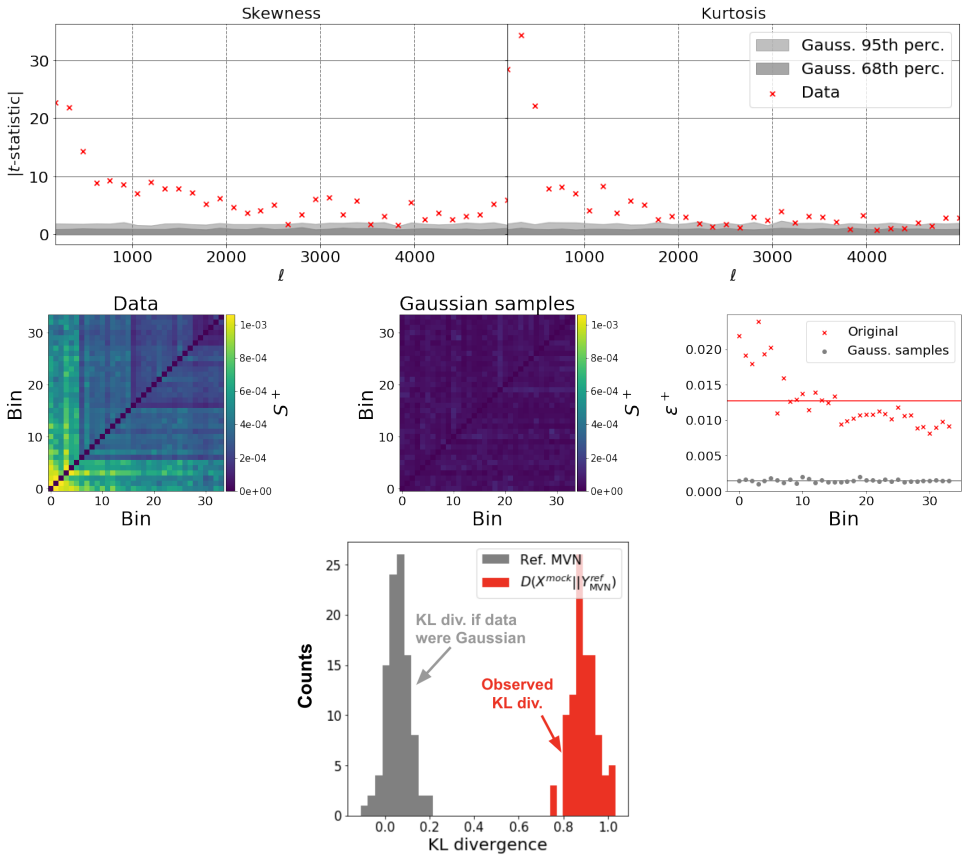

Figure 9.3: *Top*: absolute value of the t-statistic of skewness (left) and kurtosis (right) of individual bins for 2,048 weak lensing mock power spectra in red. The gray contours correspond to averaging 100 different sets of 2,048 samples drawn from a multivariate normal likelihood with the same mean and covariance as the mock data. As expected, they correspond to t-statistic values of 1 (dark gray) and 2 (light gray). The fact that most of the red crosses lie above the Gaussian contours reflects the per-bin NG in the data. *Middle*: the S^+ matrix for the mock data (left) and equivalent Gaussian samples (middle). The sum over columns of each matrix, $\epsilon_u^+ \equiv \sum S_{u,v}^+$, is shown on the right as red crosses and gray circles, respectively. The red line is a factor of 8 higher than the gray line. Comparing the two S^+ matrices, there is structure in the data that is not present in the Gaussian samples. This, and the vertical offset between the red and gray crosses on the rightmost panel, reflect the pairwise NG in the data. *Bottom*: nonparametric KL divergence estimate between the mock data and their Gaussian counterparts (red), and the Gaussian samples with themselves (gray). The fact that the gray histogram is not perfectly centered at zero is due to the slight variability of the KL estimator in 34 dimensions, given the number of mocks considered (Appendix G.4). If the data were truly Gaussian, we would expect the red histogram to lie on (or near, see Appendix G.4) the gray histogram; the large horizontal offset reflects the NG in the distribution as a whole.

test with the number of nearest neighbors set to $k = 10$. Each histogrammed data point represents a KL divergence test between the 2,048 mock power spectra and 2,048 Gaussian power spectra drawn from an analogous MVN. We repeated this procedure 100 different times, with the red histogram showing the estimated KL divergence distribution. The gray histogram is the KL divergence of an ensemble of Gaussian mocks with another ensemble of Gaussian mocks, and therefore serves as a reference for the expected nonparametric KL divergence. The large horizontal offset between these two distribution reflects the fact that the true likelihood for the power spectra is not a multivariate Gaussian.

9.4 Learning Flow-Based Likelihoods (FBL)

9.4.1 Network architecture and training procedure

Before training, we apply the same pre-processing steps as the ones mentioned in Section 9.3: we subtract off the mean and whiten the data using the Cholesky decomposition of the precision matrix. We do so because our tests on toy Gaussian data suggest that training benefits from standardizing the variance of each bin (see Appendix G.2), both in terms of speed and in terms of the quality of the final likelihood fit. Furthermore, by destroying Gaussian correlations in the data, the network can focus on picking up non-Gaussian signatures, and by subtracting off the mean all bins are equally important to the network.

In terms of the network architecture, we stack a single continuous normalizing flow with a hidden layer of dimension $d = 64$. We use exponential linear unit (ELU) activation functions, given by:

$$\mathrm{ELU}(\mathbf{x}) = \max(0, \mathbf{x}) + \min(0, \alpha * (\exp(\mathbf{x}) - 1)), \tag{9.25}$$

with $\alpha = 1$ and $*$ denoting element-wise multiplication. The network has 13,449 parameters.

Our mock dataset consists of 75,000 samples. We reserve 10% for testing, 10% for validation, and the remaining 80% for training. We train the network with batches of 500 samples using the Adam optimizer [164] and a learning rate of 0.001. We do not introduce any regularization and thus rely on the validation loss to gauge overfitting. The results shown in the remainder of the chapter correspond to training for 75 epochs.

We have checked that our results are robust to different activation functions, network architectures, and learning rate. Furthermore, we find that as the network is trained on $\gtrsim 20,000$ samples, it is able to learn the likelihood without overfitting (unsurprising, since as a rule of thumb the number of training samples has to be at least greater than the number of parameters in the model).

Before training on real data, we analyzed the fidelity of FBLs on toy Gaussian data. We detail our results in Appendix G.2, but summarize some of the relevant findings for training here. By studying the learned likelihood in Gaussian problems with singular and non-singular covariance matrices, we noticed that although the sample quality was excellent in both, the recovered likelihood becomes significantly biased and imprecise as the determinant of the covariance approaches zero. Conversely, for a full-rank covariance the likelihood was recovered perfectly to within sampling error. Therefore, we tried whitening the data before training and found that indeed training is much faster and well-behaved when we do so.

9.4.2 Results

In Figure 9.4 we show the log-likelihood of the mock samples under a MVN likelihood versus that of the FBL (top), as well as the progression of the test loss as a function of iteration number (bottom). Note that we do not expect the FBL to match the likelihood under a MVN (quite the opposite, given the level of NG detected in Section 9.3), we simply show it to see how the likelihood values are distorted with respect to the Gaussian likelihood commonly used for inference. Interestingly, we can actually see that after a full epoch FFJORD has learned the multivariate Normal likelihood: the likelihood values of the test set under the FBL (blue) are in perfect agreement with the MVN likelihood

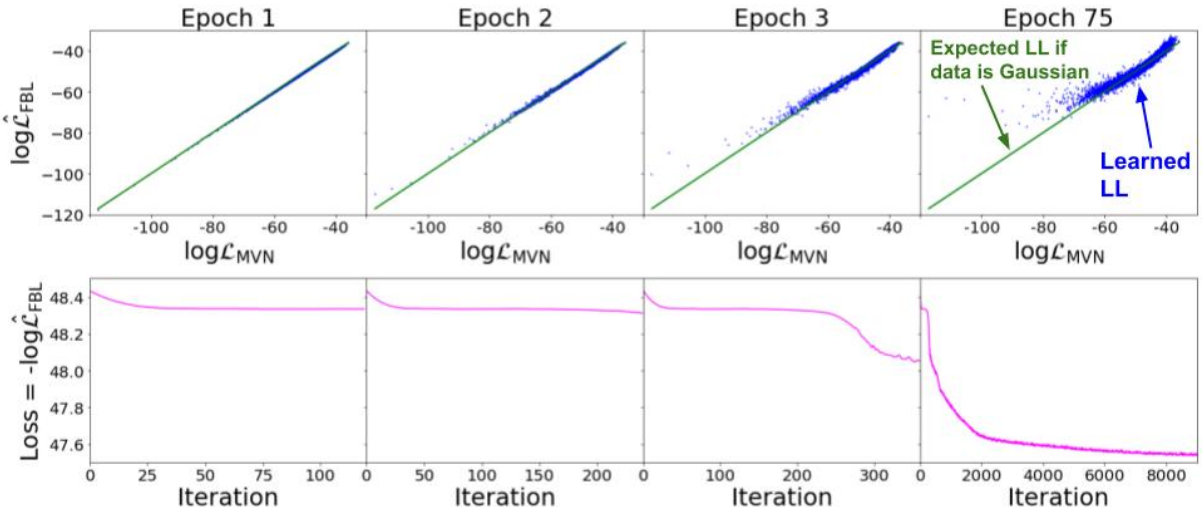

Figure 9.4: *Top*: In blue, log-likelihood of the test set samples under a multivariate normal likelihood versus the flow-based likelihood for various epochs. The green line is the MVN log-likelihood against itself, shown to make deviations from Gaussianity in the FBL more obvious. *Bottom*: test loss as a function of batch iteration number. Note that there are 120 iterations per epoch.

(green line), and the loss plateaus due to having found a local minimum with the MVN likelihood. With further training, however, the network is able to pick up on the non-Gaussianities and we can see the loss starts decreasing while simultaneously the deviation away from the MVN likelihood becomes stronger in the test set.

We compare the FBL to two other DDLs: ICA with 34 components (i.e. no dimensionality reduction) and a Gaussian mixture model with two components (chosen by minimizing the Bayesian information criterion; from here onwards, we refer to this model as GMM2). Note that we fit these two likelihoods with the full set of weak lensing mocks as well, to ensure that the comparison to the FBL was fair. In Appendix G.3 we show how our results are affected by fitting the DDLs with many less mocks, to mimick the data-limited regime in which some previous works that have applied DDLs in cosmology have found themselves in. In short, we find that the DDLs severely underestimate the non-Gaussianities in the data, concluding that the claimed parameter shifts works that have applied DDLs in this regime have observed are unlikely to truly incorporate the full extent of the impact that non-Gaussinities can have, if modeled correctly.

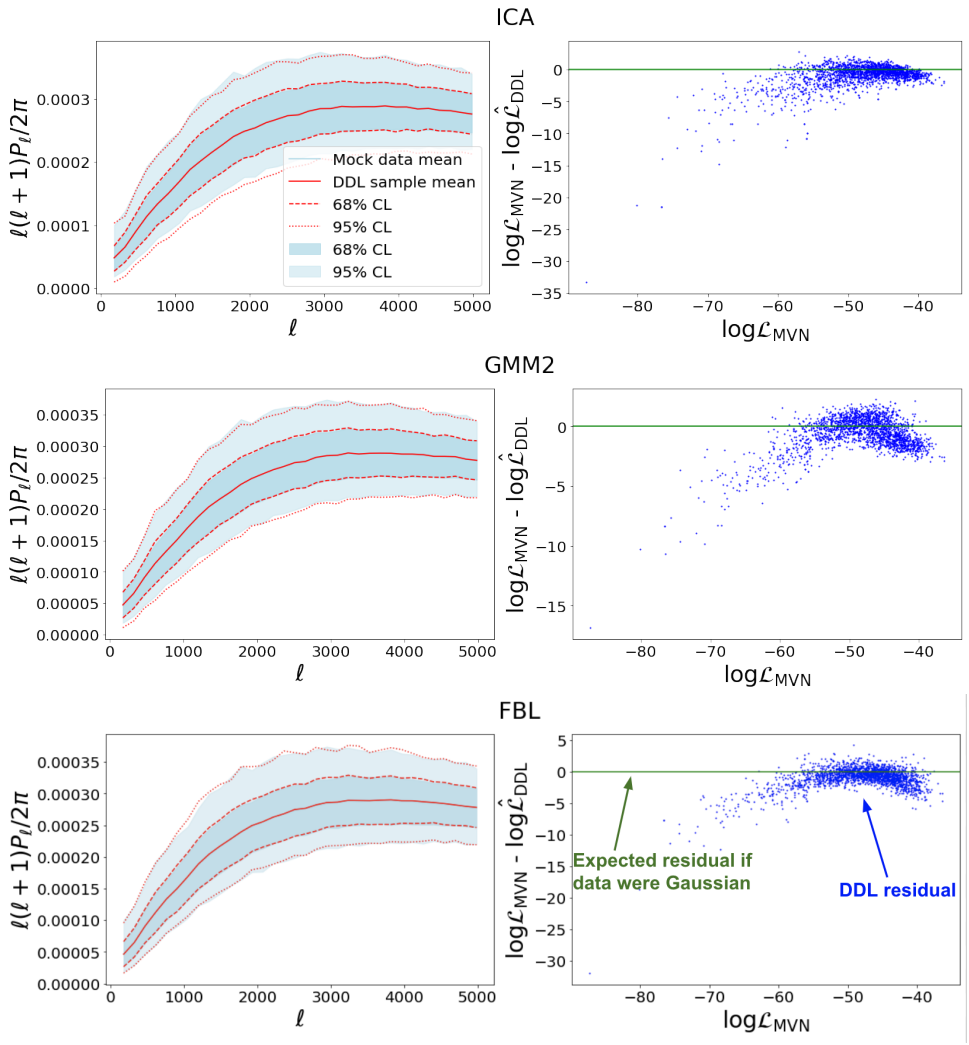

Figure 9.5: Each row corresponds to a different data-driven likelihood. From top to bottom: ICA, GMM2, and FBL. *Left*: mock weak lensing convergence power spectrum 68% and 95% CLs (light blue) together with the CLs obtained from sampling the DDLs (red), as a function of multipole number ℓ. *Right*: residual between the log-likelihood of each test set sample under a MVN likelihood and under the DDL. While the samples drawn from all three DDLs appear to match the data when looking at the mean, 1 and 2σ intervals, the likelihood is significantly different.

In Figure 9.5, each row corresponds to a different DDL likelihood: ICA, GMM2, and FBL from top to bottom. The left column shows the true mean, 68% and 95% CLs of the mock data in blue, and the same quantities obtained from sampling the DDL and unwhitening in red. The 1σ contours of all likelihoods show excellent agreement with the data, while the 2σ boundaries show small discrepancies, especially for ICA.

The right-hand panels show the residual between the MVN likelihood and the DDL for the test set samples. The non-zero residual for most test set samples indicates that the DDLs are picking up on some NG features in the data. All three DDLs show a similar

moon-shaped residual, although it is more apparent for GMM2 and FBL than for ICA. This shows that the MVN likelihood is not capturing the tail-ends of the data distribution correctly, which could have important implications in inference pipelines.

By generating samples from all three likelihoods we can also carry out the three non-Gaussianity tests detailed in Section 9.2. In Figure 9.6, the red contours in the first three columns correspond to the 68% and 95% CLs of the skewness, kurtosis, and ϵ^+, from left to right, from drawing sets of 2,048 samples from each likelihood 100 different times. They can be contrasted to the ones obtained from 100 different draws of 2,048 power spectra from the full set of mocks, which are shown in blue, and to the ones that correspond to draws from a MVN, shown in gray. Furthermore, the fourth column shows the KL divergence between the mock data and a MVN likelihood (red); between the mock data and the DDL (green); and between a MVN and itself (gray).

ICA and GMM2 each seem to have different strengths. ICA is able to pick up on the strong skewness and kurtosis of the first 2-4 bins, since they are the ones that display the strongest non-Gaussianities. On the other hand, the GMM2 approximates the true ϵ^+ while failing to capture most of the skewness and kurtosis. Interestingly, by looking at the right-hand panels of these two DDLs, we can see that capturing the pairwise NGs is much more strongly correlated to improving the KL divergence of the distribution as a whole. The shortcomings of GMM2 in terms of the skewness and kurtosis are overcome by the FBL, which shows excellent agreement with the data except the skewness at the highest multipoles, and the KL divergence estimate between the FBL and the data (green) has near-perfect alignment with the reference distribution (gray).

9.5 Discussion

We have shown that, for our mock weak lensing data, neither the GMM nor the ICA likelihoods fully succeed in capturing the NG in the data, while the FBL does so extremely well. By considering the underlying principles behind each of the three DDLs we can try to understand their performance when applied to the task of capturing non-Gaussianities.

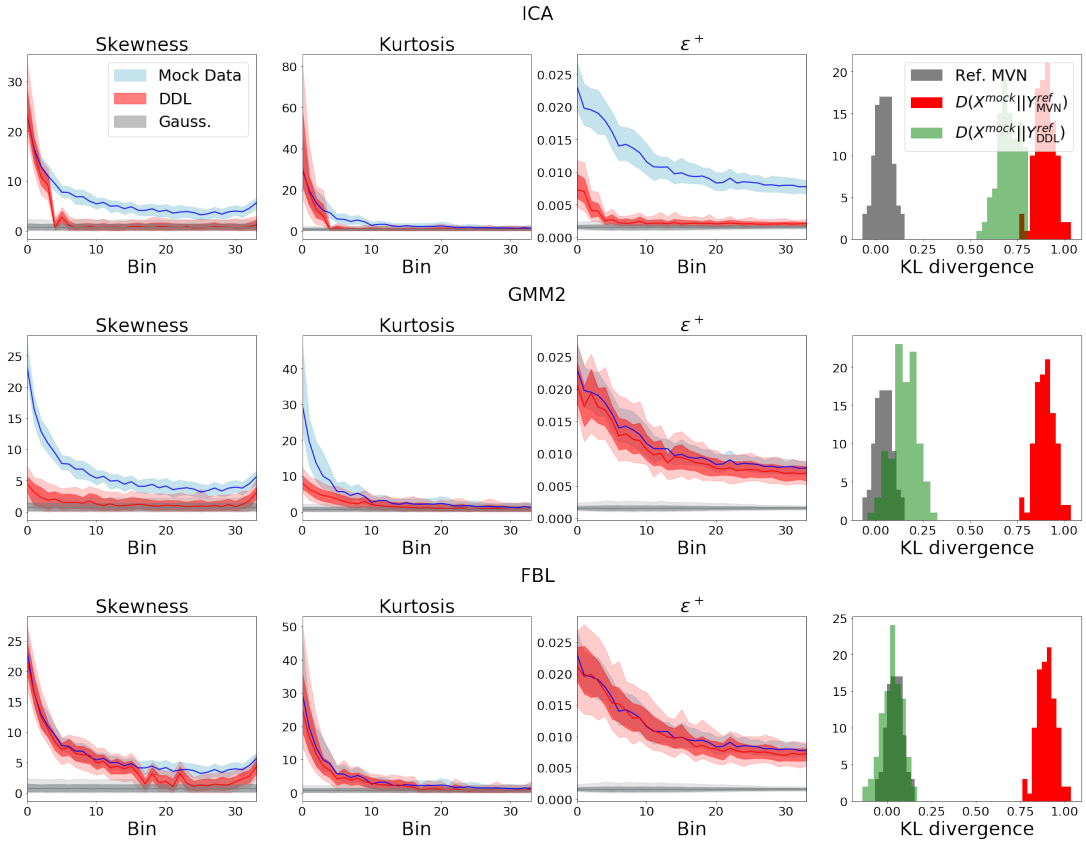

Figure 9.6: Each row corresponds to a different data-driven likelihood. From top to bottom: ICA, GMM2, and FBL. From left to right: absolute value of the t-statistic of the skewness, kurtosis, and ϵ^+ of each bin for 100 different sets of 2,048 mock WL power spectra (blue), DDL samples (red), and Gaussian samples (gray). The vertical offset between the blue and gray shaded regions reflects the high NG in the data. Overlap between the red and blue contours indicates that the DDL has captured the NG. The rightmost panel shows the KL divergence between a Gaussian with itself (gray), between the data and the MVN likelihood (red), and between the data and the DDL likelihood (green). The horizontal offset between the red and gray histograms reflects the fact that the true likelihood is not MVN. If the DDL has captured the true likelihood correctly, we expect the green histogram to lie on (or near, see Appendix G.4) the gray histogram.

The fact that ICA is predicated upon maximizing the non-Gaussianity of the components means that it is able to reproduce the distributions of independent bins, even when these have significant skewness or kurtosis. However, because it also requires independence between dimensions, it destroys the non-Gaussian correlations between datapoints that the S^+ test is sensitive to. On the other hand, GMMs cannot account for strong distortions away from Gaussianity of individual bins. The FFJORD-based FBLs are built from transformations with unconstrained Jacobians, and are thus given much more freedom of expressivity than the other DDLs considered in this work. Nevertheless, although the FBL outperforms the ICA and GMM likelihoods in terms of the three NG tests carried out, we find that there is a minimum threshold for the per-bin NG below which the FBL struggles to distinguish from a Gaussian (as can be seen for the skewness at high mutipoles).

The strengths and weaknesses of each of the three DDLs offers strong evidence that data volume is not the only factor that will determine the success or failure of a DDL. Ultimately, having some understanding of the type of NG present in the data is crucial to select the right model and not underestimate the impact of NG when inferring parameters from the data. Using our multi-resolution NG tests, which focus on increasingly coarser levels of non-Gaussianity in the data, is beneficial to faithfully diagnose the quality of the DDLs: in isolation, they could mislead one into having false confidence in the learned likelihood (like the bin-wise non-Gaussianity for ICA or ϵ^+ for GMM2), but taken together they succeed in identifying shortcomings in each of these models. The fact that the FBLs succeed in capturing the different types of NG diagnosed through the three tests suggests that FBLs are likely to be widely applicable across datasets and domains much more readily than ICA and GMM, which can require a trial-and-error procedure [277].

The impact of using FBLs for parameter inference, instead of MVN or another DDL, is left for follow-up work, but looking at the deviations from the Gaussian expectation can give us some insight. By comparing the evolution of the loss during training to the likelihood of the samples, we can see how the loss is progressively minimized as the likelihood of the test samples deviates from the Gaussian expectation. The trained model

shows that the FBL boosts the lowest likelihood values the most, although it also boosts the highest ones. This suggests that, at the very least, misuse of MVN likelihoods could be underestimating uncertainties in inferred parameters.

9.6 Conclusion

In this chapter, we have investigated the use of data-driven likelihoods to capture non-Gaussianities in the data. In particular, we have suggested exploiting flow-based machine learning models. These models are interesting because the loss function used for optimization is the negative log-likelihood of the data itself. We focus on the quality of this optimized likelihood and its capacity to pick up non-Gaussian signatures in the training data, with the goal of using it for inference. We refer to it as flow-based likelihood, or FBL.

We applied the FBL to a significantly non-Gaussian mock cosmological dataset: the weak lensing convergence power spectrum. We built on the work of Refs. [277, 278] to design a suite of tests that seek to capture different non-Gaussian features in the data − from a bin at a time to the observable as a whole − and used it to gauge to what extent three different data-driven likelihood methods succeed in capturing the non-Gaussianities in the data they are fit on. We used two DDLs used in Ref. [277] for different cosmological large-scale structure observables, ICA and GMM, and contrasted them to our proposed FBL. An interesting point to keep in mind is that the non-Gaussianities exhibited by the mock weak lensing data are much stronger than the ones of the mock galaxy power spectra used in that work, meaning that while they found that using DDLs lead to small posterior shifts compared to a Gaussian likelihood, the shift could be greater for a dataset such as this one (see also Appendices G.3, G.4, and G.5).

We found that the FBL captured the underlying likelihood much better than the other two DDLs we considered: neither GMM nor ICA fully succeeded in capturing a vast portion of the non-Gaussianities. Through our three NG tests we were able to gain some insight into the applicability of each of these three DDLs, showing that the

NG structure in the data can determine whether a given DDL is appropriate or not. Seeing the strong impact of pairwise non-Gaussian correlations in WL data, and the shortcomings of ICA in addressing them, is particularly interesting since works such as Ref. [313] used an ICA dimensionality reduction before performing inference on weak lensing data, and concluded that the impact of a non-Gaussian likelihood was small.

Unlike GMM and ICA, which require some restrictions in order to fit the data (e.g., independence between components in ICA), the FBLs used in this work are the product of transformations with unconstrained Jacobians, which allows them to be much more expressive. Not only do they succeed in capturing the NG in the data, but this freedom makes it likely for them to be widely applicable across datasets with different types of NG. Furthermore, the flexibility of the FBL when fitting different types of NG could prove beneficial not only to prevent having to follow a trial-and-error procedure to find appropriate DDLs for different observables, but also to avoid choosing a wrong one altogether.

One final consideration when weighing what DDL to apply is also the quantity of (mock) data available for training. While all data-driven methods are data-intensive, out of the three methods we consider, ICA requires significantly less parameters than GMMs, and GMMs than FBLs. Thus, depending on the computational expense required to generate mock data, employing a FBL could be prohibitive. In such a setting, looking at the type of NG features in the data through an approach like the one we suggest could aid the selection of an adequate DDL that is more restrictive than an FBL but requires less data.

When placing this work in the context of other research that has studied the impact of non-Gaussianity in WL data in particular, we want to point out that data is sometimes "Gaussianized", by for instance combining bins) before applying Gaussian and/or non-Gaussian likelihoods. Such works have concluded that NGs do not shift posteriors considerably (e.g., [296, 317]). However, this process can destroy potentially-useful information, and we therefore advocate not Gaussianizing the data and instead opting for more accurate non-Gaussian likelihoods, such as FBLs. This plight is addressed in the context

of the thermal Sunyaev-Zel'dovich one-point PDF in Ref. [318], which discusses the unfortunate need to Gaussianize clearly non-Gaussian data due to not having access to an adequate non-Gaussian likelihood, with the consequence of weakening the constraining power of the PDF. We emphasize that even highly complex and data-driven approaches can (inadvertently) Gaussianize the data in their pipelines (e.g., [319]), and thus advocate for scrutinizing data processing steps to ensure that non-Gaussian information isn't erased.

This work has required the use of simulated data to fit the three DDLs. The use of mocks in canonical large scale structure analyses is widespread, usually to estimate covariance matrices. As we have previously emphasized, adding theoretical and observational systematics can be easier with mocks than it is to incorporate directly into a likelihood. Mocks used for inference are extensively validated against the data they are needed for, but there is always the possibility that the way in which these effects are incorporated into mocks is insufficient or incorrect. Just like this could lead to an incorrect covariance being extracted from mocks, it can lead to DDLs that miss certain effects that are pertinent to the data. Understanding whether this deficit can be more important for a machine learning-based DDL such as FBLs, given that domain adaptation is nontrivial, than it is for other DDLs is left to future work.

We emphasize that this method is certainly not restricted to cosmological data. Inference in any domain can be improved by relaxing assumptions about the Gaussianity of the data. In particular, there is nothing about FBLs or the NG pipeline in this work that is specific to cosmology, and can thus easily be applied in any domain. Nevertheless, there are intriguing tensions in cosmological data that have caused much interest in the community [19, 20, 276] and, thus, make it an exciting target of FBLs. By relying on the Gaussian approximation for inference, we might be biasing inferred parameters or being falsely confident in them. While in the past this approximation may have been sufficient, as data precision increases considerably we enter a regime in which errors induced by using incorrect likelihoods can be significant.

Final Remarks

The work presented in this dissertation was done with the goal of getting a step closer to answering a single question: *is the ΛCDM model, and by extension the primary theories it is based on, the correct model of the universe?* We have presented a series of statistical techniques that can push the boundaries of the ΛCDM model and test it in novel ways. Our work has spanned a variety of fields, and a breadth of theory, simulations, and applications.

The first part of the dissertation focused on one of the largest enigmas that the ΛCDM model is predicated upon (and gives it 75% of its acronym): cold dark matter. Because there is widespread agreement between the CDM paradigm and observations on cosmological scales, we set our sights on an observable that can probe much smaller ($\sim$ kpc) scales: strong lens images. This observable can test our current understanding of CDM structure formation in a so-far untested regime. On these scales exotic dark matter physics can lead to interesting phenomenology that is observably different from CDM predictions, meaning that we can learn something fundamental about the particle nature of dark matter by adequately probing them. We presented novel ways of mapping the distribution of dark matter on these scales that offer a variety of advantages compared to traditional methods: sensitivity to lower masses (smaller scales), much more model-independence, and a significant speed-up.

In the second part of the dissertation we addressed the more practical concern of how to correctly determine the cosmological parameters (and consequently assess the adequacy of the ΛCDM model) from data. Without ensuring that this step is done correctly, we can't be certain whether apparent deviations from ΛCDM mean we need new theories to explain the universe we observe today. We explored a data-driven approach to incorporate our ignorance about elements of the theories underlying ΛCDM into how the model is parametrized (and ultimately into parameter constraints). We did so with the equation of state of dark energy, and other works have done so for reionization and inflation [320, 321]. We also explored the issue of likelihood Gaussianity, an assumption that is made in the vast majority of cosmological analyses. We developed a method to learn any likelihood from (simulated) data instead of having to impose it, and analyzed how

common approaches used to move away from a Gaussian likelihood might inadvertently be less effective than expected. At the time of writing this dissertation, we are working on applying our FBL to weak lensing data to get a definitive answer on whether using Gaussian likelihoods has biased constraints on Ω_m and σ_8 and artificially given rise to the S_8 tension.[9]

It is interesting to reflect on the reach of our work, and how we envision it will fit into the broader field in the years to come. The years encompassed in this work, from 2016 to 2020, have seen important shifts in the interests and focus of the field as a whole. To put this into perspective, in 2016 the H_0 tension between the CMB and local measurements was only at the level of $\sim 2\sigma$ [322], while today it is a $> 5\sigma$ tension [21], with papers discussing it being published nearly on a daily basis. Seminal discoveries have included the first direct detection of gravitational waves by LIGO [323], and the first image of a black hole by the Event Horizon Telescope [324–329].

Strong lensing has moved from being a probe of dark matter physics on sub-galactic scales used by a small community to a very exciting and growing subfield, arousing interest from astrophysicists, particle physicists, and cosmologists alike. Since our initial paper on the convergence power spectrum [4], Ref. [200] used a semi-analytic galaxy formation model to compute the power spectrum and found broad agreement with the predictions of our work, and Ref. [120] did an in-depth study into how source and lens properties can affect the observability of the substructure power spectrum in strong lenses. Key members of the lensing community and collaborations that traditionally rely on other methods have recognized its potential, with Ref. [232] having analyzed a strongly-lensed image and put an upper bound on the amplitude of the power spectrum, and most recently the TDCOSMO collaboration including the power spectrum extracted from image residuals in their analysis [121]. Concurrently, there have been several ML methods put forth to accelerate and enhance the extraction of small-scale dark matter information from strong lens images that fall into the statistical detection category [153, 154, 330].

So far, in this dissertation we have strongly emphasized the benefits of statistical de-

[9]Recall $S_8 \equiv \sigma_8 \sqrt{\Omega_m/0.3}$.

tection techniques, in terms of harnessing the additional constraining power of a huge but unresolvable population of very-low-mass halos. However, this discussion was implicitly from the point of view of constraining dark matter. There are other settings in which direct detection methods can have an advantage, if obtaining a lens model including perturbers is a precursor to another science objective: statistical methods have the downside of not being easy to visualize or incorporate into a high-fidelity lens model. In these settings, we believe that our direct detection-like ML implementations, which can serve a dual purpose − enhancing traditional analyses or replacing them altogether − will make them a particularly attractive model to incorporate into the analysis of real images.

One such setting is time-delay cosmography, which has now established itself as a powerful and independent way of measuring H_0 from multiply-lensed quasars [21, 210, 211].[10] So far the measurement of H_0 from time delays seems to agree with local measurements, thus widening the chasm between low-redshift and high-redshift determinations of H_0. Time delays aren't independent of the perturber problem we've been concerned with, since accurate lens modeling (including substructure) is a necessary ingredient before estimating H_0 (the projected gravitational potential of the lens enters the time delay equation). In a small sample of 6 lenses Ref. [121] found that unaccounted for structure can inflate error bars on $H_=$ between $0.7 - 2.4\%$ in individual systems, with the highest source redshift $z_{s,\max} = 2.355$. They establish that, intuitively, the error scales with the square root of the LOS volume (divided by the longest time delay), meaning that for higher-redshift sources the additional error could be considerably larger. Their results mean that the true error from time delays is somewhat higher than the claimed H_0 precision of 2.4% for the joint sample of 6 lenses analyzed by the HOLiCOW collaboration [21] (5 of them are common to both analyses). Clearly it will be crucial to remove this contribution to the uncertainty to shrink error bars to the sub-percent level. We are excited by the prospect of our work on quantifying and identifying invisible structures in strong lens images aiding H_0 estimates from time delays and helping in the quest to pinpoint the source of the H_0 tension.

[10] It can also constrain other cosmological parameters, such as the curvature and the equation of state of dark energy, but with much less precision than other cosmological observables.

As a final point on this topic, we emphasize that constraints extracted from strong lens images with respect to dark matter and H_0, although not quite at the level of other measurements at present, are based on $\mathcal{O}(10)$ images. In the next several years, the number of high-quality images at our disposal is going to increase by several orders of magnitude, so we are optimistic that, hand-in-hand with improved ways of understanding the contribution of perturbers, we'll soon be able to place state-of-the-art constraints on both the nature of dark matter and the expansion rate of the universe with strong lens images.

These last few years have also coincided with the machine learning boom, which has trickled into many scientific disciplines. The original polarization towards machine learning methods in the field − on one side reticence toward these "black box" models, and on the other boundless optimism about their applicability and capability − has given way to a fruitful middle ground, with excitement about their capacity for data processing and analysis compensated by thorough and ingenious ways of improving their interpretability and thus our comfort in using them for science. We foresee a hybrid approach between ML and canonical modeling techniques becoming increasingly popular in the near future.

Our flow-based likelihood approach is one such example, where the "black box" is used to learn a likelihood, which is extensively validated, and is not the final product of the analysis but rather a tool that can subsequently be used to perform a traditional likelihood analysis. The huge advantage of this approach is that, although it incorporates an ML element, the data analysis pipeline maintains a fully Bayesian treatment of the data, thus quelling a common complaint of inferring parameters with ML models in cosmology (the adequacy of error estimates from (Bayesian) machine learning models remains an open problem).

As the quality and volume of our data increases, and our error bars shrink, pursuing such hybrid approaches is more relevant than ever before. As we emphasized throughout the text, evaluating − and wherever necessary, changing − canonical modeling choices and assumptions is going to become crucial. We are excited by the prospect of re-

analyzing data that have given rise to some of the most important cosmological tensions with FBLs. Whether we find that using accurate non-Gaussian likelihoods does or does not significantly shift parameters with respect to a Gaussian likelihood can increase our confidence in our current model of the universe or offer decisive evidence to catalyze the move towards a new model.

Finally, looking to the future impact of ML in the cosmology community more broadly, we believe we are going to see important shifts away from having to rely on summary statistics towards more "map-level" inference: instead of having to reduce observations to a low-dimensional vector (for example a 3D distribution of millions of galaxies into a power spectrum), erasing a lot of valuable information in the process, we are going to see a rise in methods that can perform inference directly on the full-dimensional dataset. Cosmologists often have to rely on two-point functions (and sometimes three-point functions), so this would be a marked improvement, leading to significantly tighter constraints on cosmological parameters. Map-level comparisons to theory can also make deviations from ΛCDM more obvious (for instance if they are hidden as higher-order correlations).

We believe several factors are going to allows for this shift. The first of these are huge improvements in the complexity, detail, and availability of forward-modelling techniques for a wide range of observables. The second is the speed-up provided by machine learning in terms of forward-model emulators and in terms of data analysis. The final key element is the development of novel data-driven likelihood methods, such as our flow-based likelihoods, or likelihood-free methods such as DELFI, since traditional parametric likelihoods will be inadequate for these purposes.

Appendix A

Direct Detection of Dark Matter Substructure in Strong Lens Images with Convolutional Neural Networks

A.1 Subhalo population characteristics

Here we show relevant features of the subhalo populations in the images that were simulated following the procedure in Section 3.1, and are in the training/validation/test sets we used.

Single subhalo:

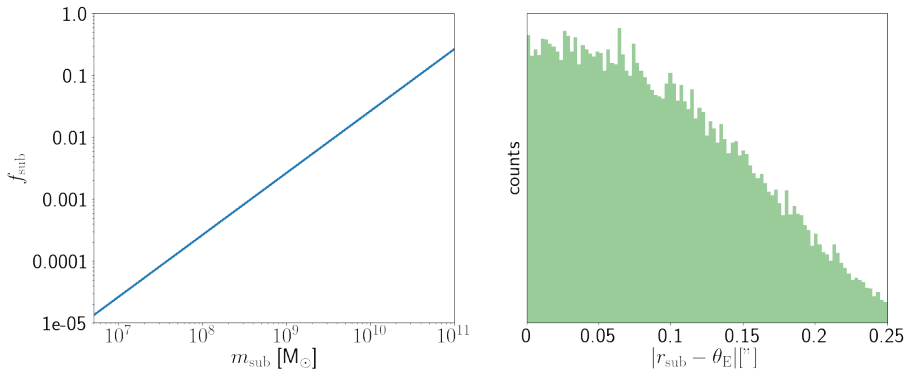

Figure A.1: Images with a single subhalo *Left*: mass fraction in substructure $f_{\rm sub}$ as a function of the mass of the subhalo in the image. *Right*: absolute distance from the subhalo's position $r_{\rm sub}$ to the Einstein ring $\theta_{\rm E}$ in arcseconds.

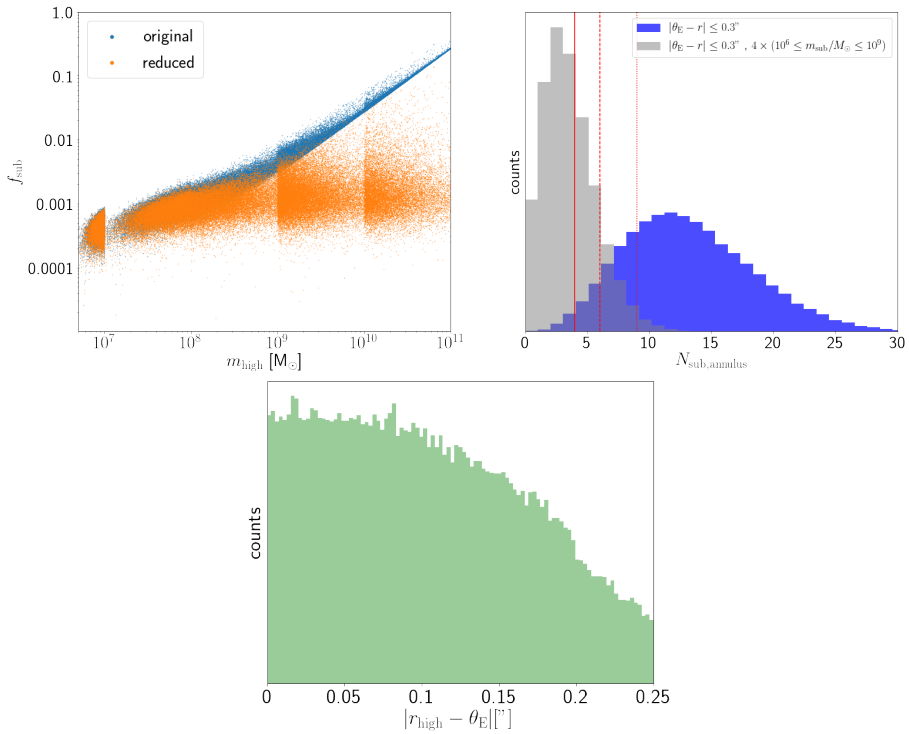

Figure A.2: Images with $N_{\rm sub}$ constrained *Top left*: in blue, the mass fraction in substructure $f_{\rm sub}$ as a function of the highest subhalo mass in the image $m_{\rm high}$. The "reduced" markers correspond to having removed the most massive subhalo. *Top right*: in blue, the total number of subhalos in an annulus 0.6″wide centered on the Einstein ring. In gray, the number of subhalos with masses between $4 \times 10^6 - 4 \times 10^9$ $M_\odot$ in this same area. The solid, dashed, and dotted red lines correspond to the 50th, 16th/84th and 2.5th/97.5th percentiles, respectively. *Bottom*: absolute distance from the position of the most massive subhalo $r_{\rm high}$ to the Einstein ring $\theta_{\rm E}$, in units of arcseconds.

Figure A.1 shows, on the left, the correspondence between the subhalo mass and the mass fraction in substructure, and on the right the absolute distance between the subhalo's position and the Einstein ring. Since these images have a single subhalo, the correspondence between $f_{\rm sub}$ and $m_{\rm sub}$ is trivially one-to-one. Also, notice that the distribution of the subhalo's position $r_{\rm sub}$ is not uniform from 0″- 0.25″because of the additional constraint on the minimum intensity at the subhalo position.

$N_{\rm sub}$-*bound:*

Figure A.2 shows different features of the subhalo populations over the entire training/validation/test sets images, where the number of subhalos per image was constrained. The top left panel shows the fraction of mass in substructure as a function of the highest-mass subhalo in the image. Most images have a fraction of mass in substructure well below 1% (in particular all the images with $m_{\rm high} < 10^9$ $M_\odot$) and, as expected, $f_{\rm sub}$ increases considerably with increasing $m_{\rm high}$. The orange markers show the value of $f_{\rm sub}$

if the most massive subhalo in the image is removed. This shows that for the majority of the images with $m_{\rm high} \gtrsim 10^9$ M$_\odot$ the value of $f_{\rm sub}$ is driven by the single most massive halo, although there is some scatter. In Section 3.2 we discuss the effect (or lack thereof) of the images with high $f_{\rm sub}$ due to more than a single subhalo.

The top right panel shows the number of subhalos within an annulus of width 0.6″centered on the Einstein ring together with a histogram of the number of subhalos in the annulus only with masses between $4\times10^6-4\times10^9$ M$_\odot$, to have a direct point of comparison to the numbers cited in [108], where it was found that, in the context of CDM, for $f_{\rm sub}=0.3\%$ there should be $\sim 7\pm1$ subhalos with masses between $4\times10^6-4\times10^9$ M$_\odot$ in an annulus of this width. We can see that the number of subhalos (and $f_{\rm sub}$) lie comfortably in the lower end of the expectations within CDM. Finally, the bottom panel shows the absolute distance between the position of the most massive subhalo and the Einstein radius. Again it is apparent that the distribution is not perfectly uniform due to the fact that the most massive subhalo has to lie at a point on the image with non-negligible intensity.

$f_{\rm sub}$-bound:

Figure A.3 is analogous to Figure A.2 but for the dataset where the mass fraction in substructure per image was constrained to be $f_{\rm sub}=1\pm0.05\%$. The top left panel shows that, up to $f_{\rm sub}M_{\rm host}$, the mass fraction in substructure is $1\pm0.05\%$. Past this point, as explained in Section 3.1.1, the single most massive subhalo already saturates this bound so instead, for a given $m_{\rm high} > f_{\rm sub}M_{\rm host}$, we generate a population of subhalos that does obey the bound on $f_{\rm sub}$ but has $m'_{\rm high} < m_{\rm high}$, and then add $m_{\rm high}$ to the image. The orange markers show the "reduced" value of $f_{\rm sub}$, where the most massive subhalo isn't taken into account. It can be seen that it obeys the constraint of $1\pm0.05\%$.

The top right panel again shows the number of subhalos in an annulus that is 0.6″wide centered on the Einstein ring in blue, and in gray the number of subhalos with masses between $4\times10^6-4\times10^9$ M$_\odot$. It has a mean of 21 subhalos, which is expected since these images have a $\sim 3\times$ greater value of $f_{\rm sub}$ than those in Ref. [108]. The bottom panel shows, as before, the absolute distance from the position of the most massive subhalo to the Einstein ring.

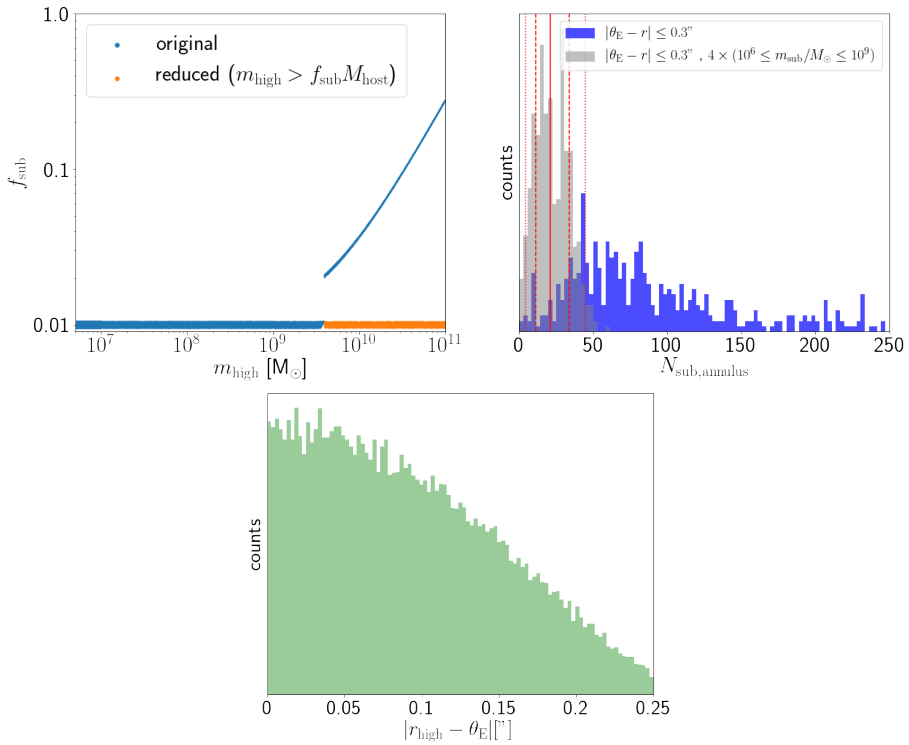

Figure A.3: Images with $f_{\rm sub}$ constrained *Top left*: in blue, the mass fraction in substructure $f_{\rm sub}$ as a function of the highest subhalo mass in the image $m_{\rm high}$. The "reduced" markers correspond to having removed the most massive subhalo for images with $m_{\rm high} > f_{\rm sub} M_{\rm host}$ to show that for these images $f_{\rm sub}$ is still bound to $1 \pm 0.05\%$ when the most massive subhalo isn't taken into account. *Top right*: in blue, the total number of subhalos in an annulus $0.6''$wide centered on the Einstein ring. In gray, the number of subhalos with masses between $4 \times 10^6 - 4 \times 10^9$ M$_\odot$ in this same area. The solid, dashed, and dotted red lines correspond to the 50th, 16th/84th and 2.5th/97.5th percentiles, respectively. *Bottom*: absolute distance from the position of the most massive subhalo $r_{\rm high}$ to the Einstein ring $\theta_{\rm E}$, in units of arcseconds.

Appendix B

Extracting the Subhalo Mass Function from Strong Lens Images using Image Segmentation

B.1 Single subhalo pixel predictions

In the main text, we examine the per-pixel predictions on the images of the test set in Fig. 4.7. There, each panel corresponded to the true label of the pixel and the x-axis represented the predicted class for the pixel. Here we flip the information around. Each panel in Fig. B.1 corresponds to pixels that are predicted to be part of the class indicated by the title. The x-axis then shows the true label and the panels are again normalized to unity. The blue and orange lines are for images with no noise and 1% noise, respectively. The solid lines do not use a probability threshold for the pixel prediction, while the dotted lines enforce that a pixel will not be predicted to a subhalo class unless the probability is at least 50%.

It is now clear that the class that is predicted is very likely to be correct. For the images with no noise, every panel has the class with highest probability correct for the prediction. The probability threshold reduces the number of pixels predicted to belong to the light subhalos. This slightly increases the fraction of pixels coming for the correct

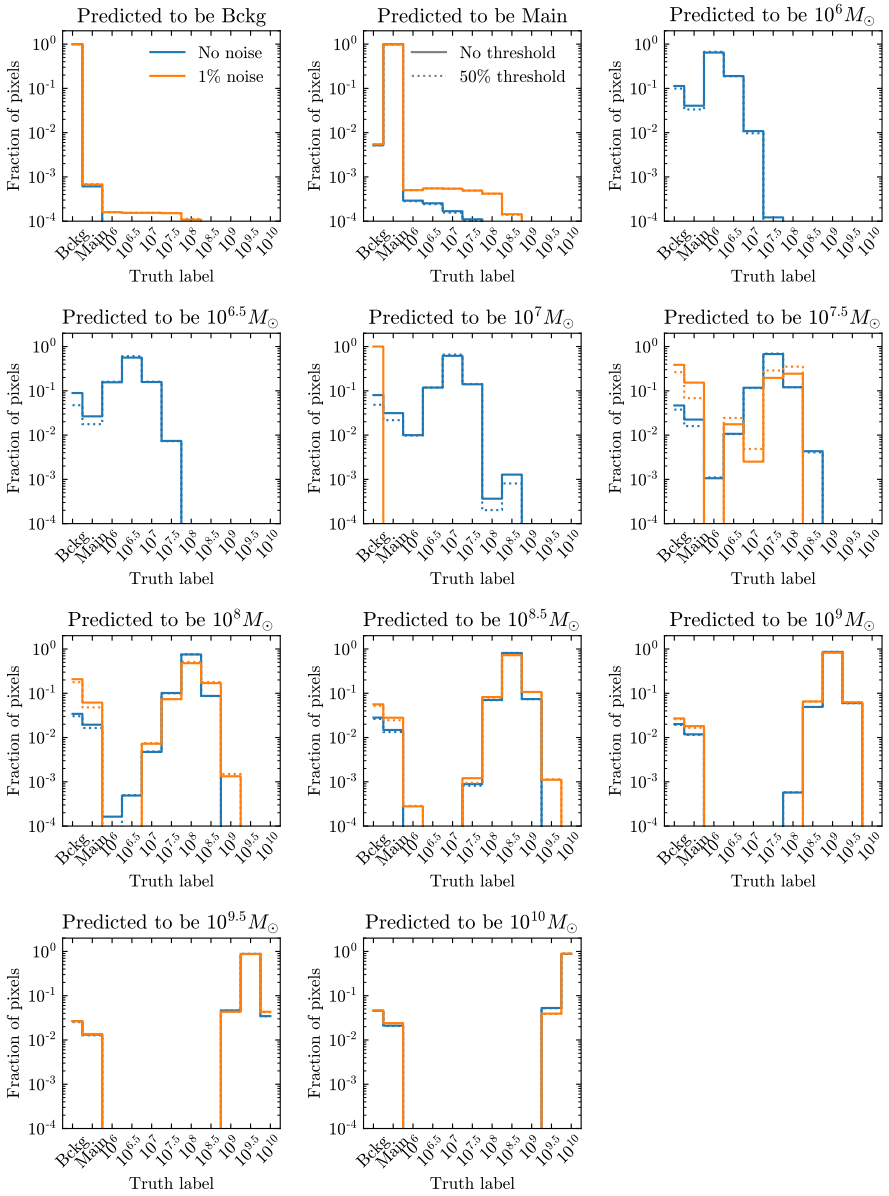

Figure B.1: Each panel corresponds to pixels which are predicted of the indicated class. The x-axis denotes the class that the pixels belong to at truth-level. Each panel is normalized to unity. The results for images without (with 1%) noise are shown in the blue (orange) lines. The solid lines denote when the predicted class has no probability threshold. The dotted lines require that the predicted probability is at least 50% before assigning a pixel to a subhalo class.

class while reducing the predictions from pixels which were supposed to belong to the background or main lens.

With noise, the network essentially does not predict pixels to belong to subhalos lighter than $m \leq 10^7$ $M_\odot$. While it predicts some pixels in the $10^{7.5}$ $M_\odot$ bin, these are not very accurate, although the probability threshold does help. Once the subhalos are as heavy as 10^8 $M_\odot$, the predicted class usually comes from pixels which at truth level were that same class.

B.2 Confusion matrix

In Fig. B.2, we show the confusion matrix for the individual pixels for all the images in the test set. The columns correspond to the true class (target) of a given pixel, while the rows show what the model classifies it as (predicted). We normalize the columns such that they sum to unity: with this choice, the values in each column show the probability that a pixel with a given true label was predicted to be in each of the eleven possible classes. The left panel corresponds to a model trained and tested on images with no noise, while the right panel corresponds to a model trained and tested on images with Gaussian noise with a standard deviation of 1% of the mean brightness of the pixels near the Einstein ring. All images (with or without noise) are convolved with a Gaussian point spread function (PSF) with a full-width half-max of $0.07''$, which roughly corresponds to the PSF size of the Hubble Space Telescope.

We can see from the left panel of Fig. B.2 that the matrix is mostly diagonal, implying remarkable accuracy overall. The two dark squares in the lower-left corner show that background and main lens pixels are almost always classified correctly.

For the subhalo classes, we see that there is always a non-zero probability of pixels getting assigned to the two adjacent classes by the model. When the network misclassifies a pixel, it is often still locating a subhalo but getting a slightly higher or lower mass estimate. This is not surprising since the subhalo masses can lie anywhere within their class and can therefore lie near the boundaries of adjacent classes. It is remarkable that

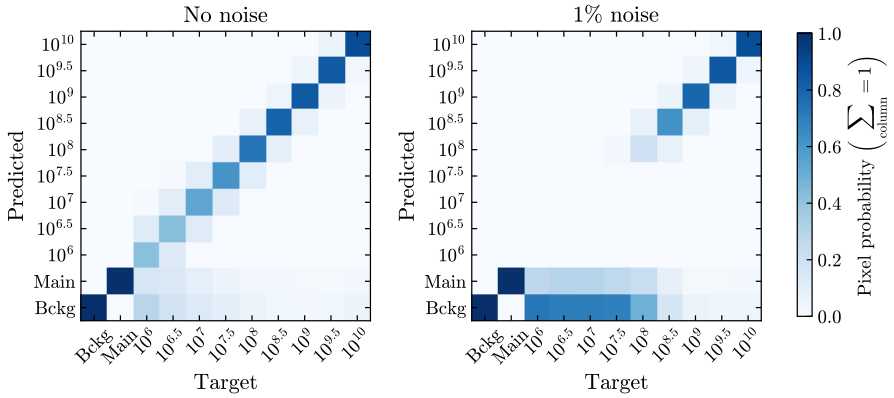

Figure B.2: Confusion matrix for the pixels of the 10^4 test images. Each image contains 80×80 pixels which are then placed in the matrix according to their truth target label and the label predicted by the network. The matrix is then normalized so the sum of the columns is unity. Most of the data lies along the diagonal, showing good accuracy. The network makes some errors by correctly locating the subhalo, but predicting the neighboring mass bin. When including noise, subhalos with $m \lesssim 10^8$ $M_\odot$ are likely to be missed altogether and get predicted as background or the main lens.

even for extremely low-mass subhalos $\lesssim 10^7$ $M_\odot$, the highest probability lies along the diagonal, although there is also a non-negligible probability of getting classified as main or background, as expected. Impressively, the network is able to identify these light subhalos despite their perturbations to the brightness being less than a 0.1% effect.

While the method in principle has an exceptionally good sensitivity, it invariably degrades with noise. When including noise at the level of 1% of the mean image brightness, the lower mass reach of the model is around 10^8 $M_\odot$. This 1% noise is found to be a realistic approximations for sources brighter than magnitude 20 (App. B.4). Above this mass, the matrix is again nearly diagonal. Below this mass, nearly all of the pixels are predicted as main or background because the effect induced by the substructure is less than the noise. One important conclusion we draw is that, despite losing sensitivity, the network does not add spurious subhalos randomly across the image when noise is added.

B.3 Subhalo predictions with probability threshold

Table B.1: The number of subhalos predicted to each class. In order for a pixel to get assigned to a subhalo class, the probability must pass a 50% threshold. The subhalo counts are then computed by examining the pixels which at truth level correspond to the substructure. The subhalo is assigned to the class with the most predictions from these pixels. For each row, there are 1000 images with a single subhalo in the indicated mass bin. The results for the network on images without noise and with noise are indicated by the number without and with brackets.

True Class	Not Detected	10^6 $M_\odot$	$10^{6.5}$ $M_\odot$	10^7 $M_\odot$	$10^{7.5}$ $M_\odot$	10^8 $M_\odot$	$10^{8.5}$ $M_\odot$	10^9 $M_\odot$	$10^{9.5}$ $M_\odot$	10^{10} $M_\odot$
10^6 $M_\odot$	410 [1000]	463 [0]	117 [0]	9 [0]	1 [0]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]
$10^{6.5}$ $M_\odot$	300 [1000]	135 [0]	455 [0]	100 [0]	10 [0]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]
10^7 $M_\odot$	178 [996]	5 [0]	128 [0]	577 [0]	107 [1]	5 [3]	0 [0]	0 [0]	0 [0]	0 [0]
$10^{7.5}$ $M_\odot$	103 [963]	0 [0]	6 [0]	132 [0]	653 [4]	105 [31]	1 [2]	0 [0]	0 [0]	0 [0]
10^8 $M_\odot$	33 [620]	0 [0]	0 [0]	0 [0]	115 [2]	778 [299]	73 [79]	1 [0]	0 [0]	0 [0]
$10^{8.5}$ $M_\odot$	19 [140]	0 [0]	0 [0]	0 [0]	4 [0]	90 [70]	840 [716]	47 [74]	0 [0]	0 [0]
10^9 $M_\odot$	2 [13]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]	78 [81]	872 [841]	48 [65]	0 [0]
$10^{9.5}$ $M_\odot$	1 [4]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]	62 [52]	879 [892]	58 [52]
10^{10} $M_\odot$	2 [2]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]	0 [0]	36 [42]	962 [956]

Table B.1 shows the subhalo detection accuracy when using a probability threshold. The results are quite similar to those without the probability threshold in Tab. 4.1. The probability threshold defaults individual pixels to the background class unless the probability for a subhalo class is larger than 50%. In general, this causes a few of the subhalos which were detected without the threshold to now be not detected. However, it also causes a few to jump to an adjacent mass bin. These subhalos have are nearly equally predicted to belong to two different classes, which happens when the mass is near the edge of a mass bin. If a pixel around the edge is now assigned to the background, the secondary class may now have more predicted pixels, resulting in a changed mass.

B.4 How realistic is 1% white noise?

Throughout this work, we have shown how the U-Net is able to detect dark matter subhalos in strong lens images without noise and when 1% white noise is included. Framing the noise as a fraction of the mean image brightness allowed us the easily see when the network would lose sensitivity to subhalos of a certain mass (Sec. 4.3). However, it is reasonable to ask how realistic this is. For instance, the gravitational imaging method searches for subhalos in pixels with a signal-to-noise ratio (S/N) of at least 3. We point out that using 1% noise does not imply that all of the pixels have S/N of 100. Instead, there is wide range, with some pixels having very strong signal compared to the noise, and other pixels that are swamped by noise.

To get a better sense of this, we simulated images including a more realistic noise similar to that of the *Hubble Space Telescope*. In particular, this requires us to set the brightness of the sky, the brightness of the source, the threshold brightness, and the exposure time. We generated 100 images using an exposure of 5400 seconds, assuming the background sky has a magnitude of 22, and that the threshold is 25.9 magnitude. For each image, we compute the signal-to-noise for each pixel and generate the cumulative distribution, starting from the highest ratio. In effect, this counts the number of pixels which have a signal-to-noise greater than a given number. These distributions are shown

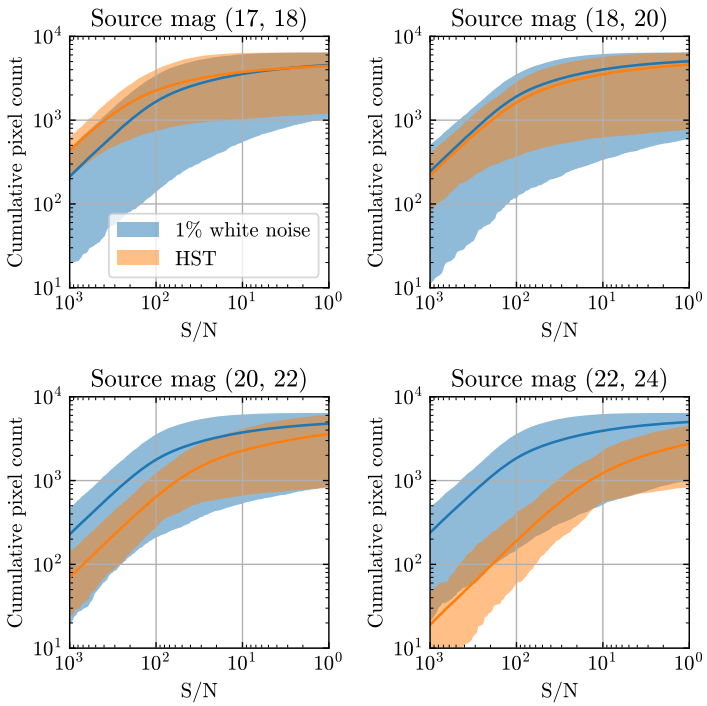

Figure B.3: The number of pixels in an 80×80 pixel strong lens image that have signal-to-noise (S/N) larger than the number on the x-axis. The blue denotes using Gaussian white noise with a standard deviation 1% of the mean image brightness. The orange is the result of using an HST-like simulation. The bands show the range over 100 images. For bright sources (lower magnitude), the HST-like noise results in larger S/N than 1% Gaussian white noise. In dimmer sources, the 1% Gaussian white noise is an optimistic assumption.

in Fig. B.3, where the blue regions show our simple white noise assumption and the orange displays the HST-like noise. The bands span the range observed for the 100 images.

In the upper left panel, the source has a magnitude between 17 and 18, which is much brighter than the sky. This implies that the Poisson shot noise for the HST-like scenario dominates over the white noise. The orange band is near the top of or above the blue-band; there are more pixels with large S/N for HST-like noise than if we use 1% white noise, making substructure easier to detect. The upper right panel shows that our approximation is very similar to HST for sources with magnitudes between 18 and 20. The source light in the bottom left panel is now comparable with the background sky. The blue band is above the orange band, indicating that our noise approximation is optimistic for such sources. Finally, the lower right panel has very dim sources. In these images, the noise from the sky dominates and is much larger than our 1% estimate.

As a proof-of-concept, we used 1% Gaussian white noise, which yields similar noise profiles as HST for sources brighter than magnitude 20. This reduced a number of

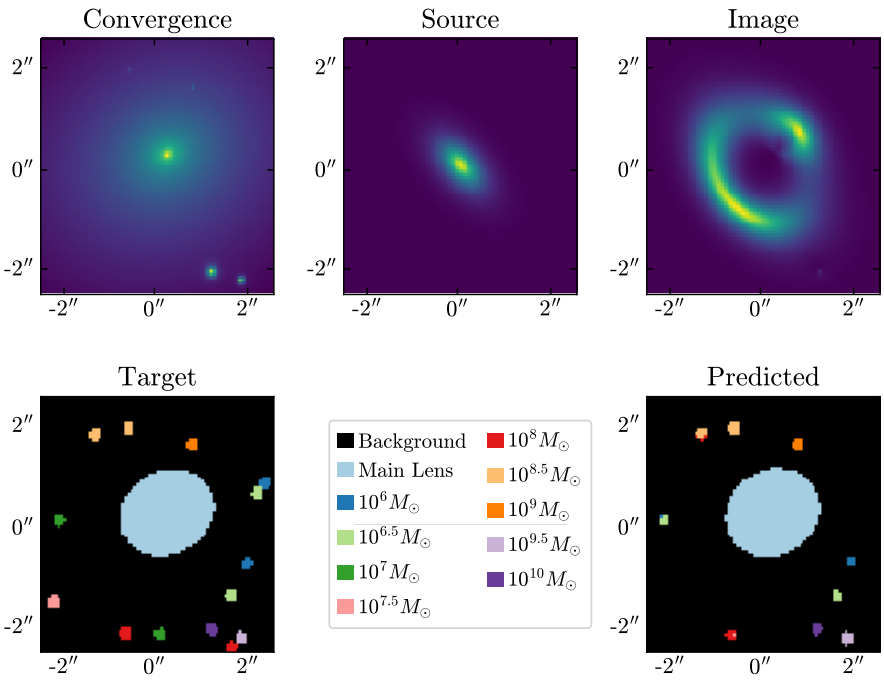

Figure B.4: Example of finding subhalos with image segmentation. The simulated convergence field of lens halos (upper-left panel) is used to define target labels (lower-left). We define our subhalo labels to be a circle with radius 2 pixels centered on the subhalo. Even though this is not representative of relative convergences of the different masses, it was found to lead to more stable results. A light source (upper-middle) is placed behind the lens, resulting in an observed image (upper-right) after ray-tracing the light through the gravitational potential of the lens. A neural network takes in only the observed image and predicts the label for each pixel (lower-right).

variables in our simulation, such as the source brightness. To apply image segmentation to real strong lens images, one should include all of these considerations.

B.5 Domain adaptation

Machine learning models work by discovering correlations in training data. It is challenging to get them to work when new data is outside the realm of what they have been trained on, and a whole subfield is dedicated to the problem of adapting to new domains.

Both in the training and in the testing presented, each image had either no substructure or exactly one subhalo. Due to the steep slope of the subhalo mass function, we expect that many subhalos should be present in strong lensing images [143]. We therefore assess whether the network is capable of generalizing to other lensing situations after this training regiment.

In Fig. B.4 we show an example with no noise, where we included many subhalos

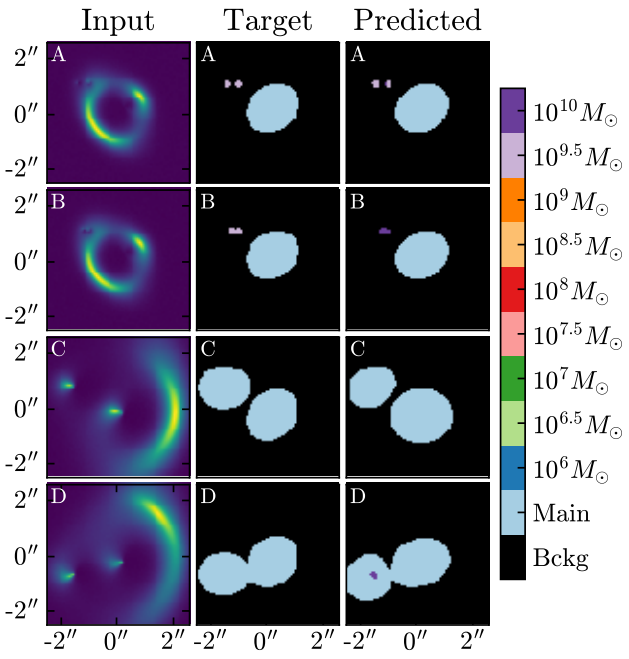

Figure B.5: The left column displays the strongly lensed images used as input to the neural network. The middle column contains the pixel-by-pixel labels we want the network to reproduce. The right column shows the network predictions.

both near the Einstein ring and closer to the edge. Near the edge, the network is able to detect many of the heavier subhalos, but does not capture the lighter ones. The fact that training on a single subhalo generalizes to many subhalos is remarkable. This opens the possibility of using the results of the network to infer the subhalo mass function using many fewer images than if we could only detect a single subhalo at a time, which was precisely the object of Chapter 4.

In panels A and B of Fig. B.5, we examine the effect of two subhalos being close or overlapping. We choose two subhalos in the $10^{9.5}$ $M_\odot$ bin because they are easy to identify by eye. The source light and main galaxy lens are kept constant throughout these images to see only the effect due to subhalos, and we include 1% Gaussian noise. In A, the two subhalos are far enough away from each other that the network is able to resolve them separately. In panels B, the two subhalos are close enough that the true pixels are touching each other. The network does not correctly identify two individual subhalos in this case, but it does classify the pixels as belonging to a single, higher mass subhalo, 10^{10} $M_\odot$. An animation of the subhalo traversing the image and its effect on the network output can be found at this link.

Finally, in panels C and D of Fig. B.5 a second large halo is included in the image with a mass on the same order of the main lens, meaning that there are two main lens halos rather than a single one. The combined lenses result in much larger distortions of the light, as shown in the input images in the left columns. The Einstein ring in the training images was always around $1''$ in radius, while in these images it is closer to $2''$. This is why the network predicts the central lens as around 30% too large. While the network is capable of identifying the presence and location of a second main lens, its shape is not captured well. In fact, we can see in panel D, how it adds a heavy subhalo to the center of the second lens. An animation showing the different positions of the second lens, and its effect on pixel classification, is available at this link.

Appendix C

Power Spectrum of Dark Matter Substructure in Strong Gravitational Lenses

C.1 Deriving $\mathcal{P}_{\rm t}$

Let us consider a subhalo population that is uniformly distributed. Starting with Eq. (5.41) and setting $g(h) \equiv r_{\rm t} - r_{\rm t,0}\left(\frac{m}{m_0}\right)^{1/3}\left(\frac{\sqrt{r^2+h^2}}{r_{\rm 3D,0}}\right)^{\nu}$, we obtain

$$\begin{aligned}\mathcal{P}_{\rm t}(r_{\rm t}|m,r) &= \frac{1}{Z}\int dh\, \mathcal{P}_{3D}(r_{3D})\, \delta(g(h)) \\ &= \frac{1}{Z}\frac{1}{2AR_{\rm max}}\int dh\, \delta(g(h)) \\ &= \frac{1}{Z}\frac{1}{2AR_{\rm max}}\frac{2}{|g'(h_i)|} \\ &= \frac{1}{R_{\rm max}|g'(h_i)|},\end{aligned} \tag{C.1}$$

where h_i is the solution of $g(h_i) = 0$, and where we used $Z = 1/A$. Then,

$$\begin{aligned}|g'(h_i)| = \nu\frac{r_{\rm t}}{r_{\rm 3D,0}^2}\sqrt{r_{\rm 3D,0}^2\left[\left(\frac{m_0}{m}\right)^{1/3}\frac{r_{\rm t}}{r_{\rm t,0}}\right]^{2/\nu} - r^2} \\ \times\left[\left(\frac{m_0}{m}\right)^{1/3}\frac{r_{\rm t}}{r_{\rm t,0}}\right]^{-2/\nu}.\end{aligned} \tag{C.2}$$

where $0 \leq r \lesssim b$. Letting $x^2 = r_{\rm 3D,0}^2 \left[\left(\frac{m_0}{m}\right)^{1/3} \frac{r_{\rm t}}{r_{\rm t,0}} \right]^{2/\nu}$, we can do the following expansion:

$$\sqrt{x^2 - r^2} = x\sqrt{1 - \frac{r^2}{x^2} + ...} \approx x, \tag{C.3}$$

where we have used the fact that $x^2 \gg r^2$. In reality this equality does not hold perfectly: when subhalos are at 3D halo-centric distances close to (or below) the Einstein radius, their tidal radius can be such that x^2 is comparable to (or less than) r^2. However, we can take advantage of the fact that the volume in which $r_{\rm t}$ takes on such small values makes up only $\sim 1\%$ of the entire line-of-sight volume within the host, so the number of subhalos with these tidal radii will make up a minute portion of the entire subhalo population after projection onto the lens plane.

Then,

$$|g'(h_i)| = \nu \frac{r_{\rm t}}{r_{\rm 3D,0}} \left[\left(\frac{m_0}{m}\right)^{1/3} \frac{r_{\rm t}}{r_{\rm t,0}} \right]^{-1/\nu} \tag{C.4}$$

and plugging this into Eq. (C.1),

$$\mathcal{P}_{\rm t}(r_{\rm t}|m) = \frac{1}{\nu R_{\rm max}} \frac{r_{\rm 3D,0}}{r_{\rm t}} \left[\left(\frac{m_0}{m}\right)^{1/3} \frac{r_{\rm t}}{r_{\rm t,0}} \right]^{1/\nu}. \tag{C.5}$$

In fact $\mathcal{P}_{\rm t}$ is unchanged in a case where $\mathcal{P}_{\rm r}$ has some radial dependence. Using as an example $\mathcal{P}_{\rm r}(r) = (1/2\pi b)(1/r)$, we obtain $\mathcal{P}_{3D}(r) = (1/4\pi b R_{\rm max})(1/r)$. Then,

$$\begin{aligned} \mathcal{P}_{\rm t}(r_{\rm t}|m,r) &= \frac{1}{Z} \frac{1}{4\pi b R_{\rm max}} \int dh \, \frac{1}{r} \, \delta(g(h)) \\ &= \frac{2\pi b r}{4\pi b R_{\rm max} r} \frac{2}{|g'(h_i)|} \\ &= \frac{1}{R_{\rm max}|g'(h_i)|}. \end{aligned} \tag{C.6}$$

C.2 SIDM convergence profile

To normalize Eq. (5.60) (i.e., determine $m_{\rm b}$), we simply integrate the profile out to infinity, which gives us Eq. (5.61). To obtain the convergence profile, we calculate the projection integral

$$\kappa_{\rm tBurk}(r) = \frac{1}{\Sigma_{\rm crit}} \int_{-\infty}^{\infty} \rho_{\rm tBurk}(\sqrt{h^2 + r^2})\, dh, \tag{C.7}$$

where r is the 2D radial coordinate on the lens plane, h is the line-of-sight coordinate, and thus $R = \sqrt{h^2 + r^2}$. We can in fact simplify this expression by doing a slight change of variables. We can rewrite Eq. (5.60) as

$$\rho_{\rm tBurk}(y) = \frac{m_{\rm b}}{4\pi r_{\rm s}^3} \frac{1}{(p+y)(p^2+y^2)} \left(\frac{\tau^2}{y^2+\tau^2}\right), \tag{C.8}$$

where $y = R/r_{\rm s}$ and $\tau = r_{\rm t}/r_{\rm s}$. Then, with $l = h/r_{\rm s}$ and $x = r/r_{\rm s}$,

$$\begin{aligned}
\kappa_{\rm tBurk}(x) &= \frac{r_{\rm s}}{\Sigma_{\rm crit}} \int_{-\infty}^{\infty} \rho(\sqrt{l^2 + x^2})\, dl \\
&= \frac{m_{\rm b}}{2\pi\Sigma_{\rm crit} r_{\rm s}^2}\, \tau^2 \Bigg\{ \pi \Bigg(\frac{2p\sqrt{\frac{1}{\tau^2+x^2}}}{p^4 - \tau^4} - \frac{\sqrt{\frac{1}{x^2-p^2}}}{p(\tau^2+p^2)} \\
&\quad - \frac{\sqrt{\frac{1}{x^2+p^2}}}{p^3 - p\tau^2} \Bigg) + \frac{2\arctan\left[\frac{p}{\sqrt{x^2-p^2}}\right]}{\sqrt{x^2-p^2}(p^3+p\tau^2)} - \\
&\quad \frac{2\tanh^{-1}\left[\frac{p}{\sqrt{p^2+x^2}}\right]}{\sqrt{x^2+p^2}(p^3-p\tau^2)} + \frac{4\tau\tanh^{-1}\left[\frac{\tau}{\sqrt{x^2+\tau^2}}\right]}{\sqrt{x^2+\tau^2}(p^4-\tau^4)} \Bigg\}.
\end{aligned} \tag{C.9}$$

C.3 Table of Constants and Variables

Table C.1

Constant or Variable	Value	Description
M_{lens}	1.8×10^{12} $M_\odot$	Lens mass
R_{max}	409.6 kpc	Maximum radius of the lens
b	6.3 kpc	Einstein radius of the lens
Σ_{crit}	3×10^9 $M_\odot/\text{kpc}^2$	Critical surface mass density
m_{high}	10^8 $M_\odot$	Upper bound for the subhalo mass
m_{low}	10^5 $M_\odot$	Lower bound for the subhalo mass
m_*	2.52×10^7 $M_\odot$	
β	-1.9	Subhalo mass function slope
r_{s}		Subhalo scale radius
$r_{\text{s,0}}$	0.1 kpc	
r_{t}		Subhalo tidal radius
$r_{\text{t,0}}$	1 kpc	
r_{3D}		3D halocentric distance to a subhalo
$r_{\text{3D,0}}$	100 kpc	
m_0	10^6 $M_\odot$	

Appendix D

Gravitational Lensing and the Power Spectrum of Dark Matter Substructure: Insights from the ETHOS N-body Simulations

D.1 Truncated convergence profiles

The truncated NFW profile (tNFW) [160] is given by

$$\rho_{\rm tNFW}(r) = \frac{m_{\rm NFW}}{4\pi r(r+r_{\rm s})^2}\left(\frac{r_{\rm t}^2}{r^2+r_{\rm t}^2}\right), \tag{D.1}$$

where $r_{\rm s}$ is the scale radius, $r_{\rm t}$ is the tidal radius and $m_{\rm NFW}$ is defined below. Integrating this profile along the line of sight and diving by the critical density for lensing $\Sigma_{\rm crit}$, we obtain the tNFW convergence profile:

$$\kappa_{\rm tNFW}(x) = \frac{m_{\rm NFW}}{\Sigma_{\rm crit} r_{\rm s}^2} \frac{\tau^2}{2\pi(\tau^2+1)^2} \Bigg[\frac{\tau^2+1}{x^2-1}(1-F(x)) + 2F(x) - \frac{\pi}{\sqrt{\tau^2+x^2}} + \frac{\tau^2-1}{\tau\sqrt{\tau^2+x^2}} L(x) \Bigg], \tag{D.2}$$

where

$$x = \frac{r}{r_{\rm s}}, \quad \tau = \frac{r_{\rm t}}{r_{\rm s}}, \tag{D.3}$$

$$F(x) = \frac{\cos^{-1}(1/x)}{\sqrt{x^2-1}}, \tag{D.4}$$

$$L(x) = \ln\left(\frac{x}{\sqrt{\tau^2+x^2}+\tau} \right), \tag{D.5}$$

and

$$m = \frac{m_{\rm NFW}\tau^2}{(\tau^2+1)^2} \left[(\tau^2-1)\ln(\tau) + \tau\pi - (\tau^2+1) \right]. \tag{D.6}$$

The truncated Burkert profile [4, 183] is given by:

$$\rho_{\rm tBurk}(r) = \frac{m_{\rm b}}{4\pi(r+p\,r_{\rm s})(r^2+p^2 r_{\rm s}^2)} \left(\frac{r_{\rm t}^2}{r^2+r_{\rm t}^2} \right), \tag{D.7}$$

where $r_{\rm b}$ is the core radius, and the scale mass $m_{\rm b}$ is the mass within the core. Here we set $r_{\rm b} = p\,r_{\rm s}$, where p is a constant that represents the size of the core as a fraction of the scale radius. The convergence field of this density profile is then [4]:

$$\kappa_{\rm tBurk}(x) = \frac{m_{\rm b}}{8\pi\Sigma_{\rm crit} r_{\rm s}^2}\,\tau^2\Bigg\{\pi\Bigg(\frac{2p\sqrt{\frac{1}{\tau^2+x^2}}}{p^4-\tau^4} - \frac{\sqrt{\frac{1}{x^2-p^2}}}{p(\tau^2+p^2)} - \frac{\sqrt{\frac{1}{x^2+p^2}}}{p^3-p\tau^2}\Bigg) + \frac{2\arctan\left[\frac{p}{\sqrt{x^2-p^2}}\right]}{\sqrt{x^2-p^2}(p^3+p\tau^2)} - \frac{2\tanh^{-1}\left[\frac{p}{\sqrt{p^2+x^2}}\right]}{\sqrt{x^2+p^2}(p^3-p\tau^2)} + \frac{4\tau\tanh^{-1}\left[\frac{\tau}{\sqrt{x^2+\tau^2}}\right]}{\sqrt{x^2+\tau^2}(p^4-\tau^4)}\Bigg\}, \tag{D.8}$$

where again x and τ are defined as in Eq. (D.3).

D.2 Features of the convergence maps

The tables in this Section display some quantities of interest extracted from the two simulations. The main value quoted corresponds to the median across 90 projections for a box size with $L = 100$ kpc, while the errors correspond to the 90% confidence interval. The exception to this is the first entry in each table, $N_{\rm sub}(L = 300$ kpc), since this is a quantity extracted from the original simulations before doing any projections. The Einstein radius is fixed to $1''$. With our cosmology, this corresponds to a physical size $R_{\rm E} = 6.18$ kpc at $z = 0.5$[1] and $R_{\rm E} = 8.10$ kpc at $z = 1$.

Notice that in the "All subhalos" column the $m_{\rm eff} \equiv \langle m^2\rangle/\langle m\rangle$ entry displays very large upper bounds. This is due to the fact that when there are very few subhalos with masses $> 10^9$ $M_\odot$ they only rarely get projected into the region of interest. Even a one order of magnitude difference in the maximum subhalo mass can change $m_{\rm eff}$ considerably, yielding a very large upper bound. This is reflected in the large 90% confidence interval in the red lines of Figure 6.3.

[1]We will be using this value in the $z = 0$ tables as well for the reasons outlined in 6.4.1.

D.2.1 CDM

Table D.1: Relevant properties of the CDM subhalo population at $z = 0$

	$m_{\rm high} = 10^8$ M$_\odot$	$m_{\rm high} = 10^9$ M$_\odot$	All subhalos
$N_{\rm sub}(L_{\rm box} = 300$ kpc$)$	9810	10007	10031
$N_{\rm sub}(L_{\rm box} = 100$ kpc$)$	1004^{+116}_{-97}	1026^{+148}_{-116}	1024^{+104}_{-109}
$N_{\rm sub}(R_{\rm E})$	12^{+8}_{-6}	13^{+8}_{-5}	13^{+7}_{-6}
$\bar{\kappa}_{\rm sub}$	$\left(3.28^{+0.54}_{-0.37}\right) \times 10^{-4}$	$\left(5.06^{+1.19}_{-0.97}\right) \times 10^{-4}$	$\left(1.22^{+0.94}_{-0.51}\right) \times 10^{-3}$
$\langle m \rangle$ [M$_\odot$]	$\left(7.75^{+0.41}_{-0.57}\right) \times 10^6$	$\left(1.20^{+0.18}_{-0.20}\right) \times 10^7$	$\left(2.99^{+2.34}_{-1.33}\right) \times 10^7$
$m_{\rm eff} \equiv \langle m^2 \rangle / \langle m \rangle$ [M$_\odot$]	$\left(2.72^{+0.17}_{-0.27}\right) \times 10^7$	$\left(1.64^{+0.83}_{-0.83}\right) \times 10^8$	$\left(7.79^{+20.1}_{-7.30}\right) \times 10^9$
$r_{\rm t,max}$ [kpc]	$8.17^{+11.67}_{-2.70}$	$9.51^{+5.53}_{-3.33}$	$10.96^{+26.50}_{-2.88}$
$r_{\rm s,min}$ [kpc]	$0.03^{+0.00}_{-0.01}$	$0.03^{+0.00}_{-0.01}$	$0.03^{+0.00}_{-0.01}$

Table D.2: Relevant properties of the CDM subhalo population at $z = 0.5$

	$m_{\rm high} = 10^8$ M$_\odot$	$m_{\rm high} = 10^9$ M$_\odot$	All subhalos
$N_{\rm sub}(L_{\rm box} = 300$ kpc$)$	6516	6651	6669
$N_{\rm sub}(L_{\rm box} = 100$ kpc$)$	587^{+70}_{-57}	584^{+103}_{-59}	596^{+88}_{-59}
$N_{\rm sub}(R_{\rm E})$	9^{+3}_{-4}	9^{+5}_{-5}	9^{+5}_{-5}
$\bar{\kappa}_{\rm sub}$	$\left(1.90^{+0.16}_{-0.22}\right) \times 10^{-4}$	$\left(2.93^{+0.72}_{-0.58}\right) \times 10^{-4}$	$\left(4.18^{+7.87}_{-1.55}\right) \times 10^{-4}$
$\langle m \rangle$ [M$_\odot$]	$\left(7.67^{+0.48}_{-0.58}\right) \times 10^6$	$\left(1.20^{+0.23}_{-0.19}\right) \times 10^7$	$\left(1.66^{+3.04}_{-0.58}\right) \times 10^7$
$m_{\rm eff} \equiv \langle m^2 \rangle / \langle m \rangle$ [M$_\odot$]	$\left(2.53^{+0.45}_{-0.37}\right) \times 10^7$	$\left(1.09^{+1.04}_{-0.44}\right) \times 10^8$	$\left(5.95^{+128.00}_{-4.99}\right) \times 10^8$
$r_{\rm t,max}$ [kpc]	$9.62^{+4.14}_{-3.95}$	$10.90^{+6.38}_{-2.94}$	$12.19^{+6.12}_{-2.87}$
$r_{\rm s,min}$ [kpc]	$0.06^{+0.02}_{-0.02}$	$0.07^{+0.02}_{-0.02}$	$0.06^{+0.02}_{-0.02}$

Table D.3: Relevant properties of the CDM subhalo population at $z = 1$

	$m_{\rm high} = 10^8$ M$_\odot$	$m_{\rm high} = 10^9$ M$_\odot$	All subhalos
$N_{\rm sub}(L_{\rm box} = 300$ kpc$)$	4694	4783	4798
$N_{\rm sub}(L_{\rm box} = 100$ kpc$)$	446^{+80}_{-49}	459^{+75}_{-46}	453^{+90}_{-54}
$N_{\rm sub}(R_{\rm E})$	11^{+4}_{-6}	11^{+4}_{-6}	11^{+5}_{-6}
$\bar{\kappa}_{\rm sub}$	$\left(7.93^{+1.71}_{-1.26}\right) \times 10^{-5}$	$\left(1.52^{+0.53}_{-0.52}\right) \times 10^{-4}$	$\left(1.74^{+2.63}_{-0.64}\right) \times 10^{-4}$
$\langle m \rangle$ [M$_\odot$]	$\left(7.57^{+0.66}_{-0.58}\right) \times 10^{6}$	$\left(1.43^{+0.36}_{-0.36}\right) \times 10^{7}$	$\left(1.66^{+1.91}_{-0.51}\right) \times 10^{7}$
$m_{\rm eff} \equiv \langle m^2 \rangle / \langle m \rangle$ [M$_\odot$]	$\left(2.60^{+0.33}_{-0.38}\right) \times 10^{7}$	$\left(2.05^{+1.14}_{-0.95}\right) \times 10^{8}$	$\left(2.82^{+44.1}_{-1.64}\right) \times 10^{8}$
$r_{\rm t,max}$ [kpc]	$10.12^{+7.08}_{-3.28}$	$14.27^{+6.15}_{-5.47}$	$16.00^{+10.70}_{-5.72}$
$r_{\rm s,min}$ [kpc]	$0.08^{+0.03}_{-0.01}$	$0.08^{+0.03}_{-0.02}$	$0.08^{+0.03}_{-0.02}$

D.2.2 ETHOS4

Table D.4: Relevant properties of the ETHOS4 subhalo population at $z = 0$

	$m_{\rm high} = 10^8$ M$_\odot$	$m_{\rm high} = 10^9$ M$_\odot$	All subhalos
$N_{\rm sub}(L_{\rm box} = 300$ kpc$)$	821	898	918
$N_{\rm sub}(L_{\rm box} = 100$ kpc$)$	93^{+28}_{-15}	100^{+35}_{-15}	101^{+47}_{-14}
$N_{\rm sub}(R_{\rm E})$	1^{+2}_{-1}	2^{+2}_{-2}	2.00^{+2}_{-2}
$\bar{\kappa}_{\rm sub}$	$\left(4.06^{+1.33}_{-1.08}\right) \times 10^{-5}$	$\left(1.28^{+0.49}_{-0.59}\right) \times 10^{-4}$	$\left(7.36^{+12.80}_{-5.65}\right) \times 10^{-4}$
$\langle m \rangle$ [M$_\odot$]	$\left(1.05^{+0.15}_{-0.19}\right) \times 10^{7}$	$\left(2.92^{+1.77}_{-1.11}\right) \times 10^{7}$	$\left(1.89^{+2.58}_{-1.47}\right) \times 10^{8}$
$m_{\rm eff} \equiv \langle m^2 \rangle / \langle m \rangle$ [M$_\odot$]	$\left(3.07^{+0.84}_{-0.90}\right) \times 10^{7}$	$\left(3.63^{+1.79}_{-2.40}\right) \times 10^{8}$	$\left(1.08^{+2.35}_{-1.02}\right) \times 10^{10}$
$r_{\rm t,max}$ [kpc]	$10.86^{+18.13}_{-3.54}$	$13.85^{+12.21}_{-4.80}$	$20.28^{+45.90}_{-6.99}$
$r_{\rm s,min}$ [kpc]	$0.06^{+0.00}_{-0.02}$	$0.06^{+0.00}_{-0.02}$	$0.06^{+0.00}_{-0.02}$

Table D.5: Relevant properties of the ETHOS4 subhalo population at $z = 0.5$

	$m_{\rm high} = 10^8$ M$_\odot$	$m_{\rm high} = 10^9$ M$_\odot$	All subhalos
$N_{\rm sub}(L_{\rm box} = 300$ kpc$)$	579	629	642
$N_{\rm sub}(L_{\rm box} = 100$ kpc$)$	57^{+19}_{-10}	60^{+19}_{-8}	62^{+17}_{-8}
$N_{\rm sub}(R_{\rm E})$	$0.50^{+1.5}_{-0.5}$	$0.5^{+2}_{-0.5}$	$0.5^{+2}_{-0.5}$
$\bar{\kappa}_{\rm sub}$	$\left(3.43^{+1.38}_{-1.02}\right) \times 10^{-5}$	$\left(7.17^{+2.71}_{-3.30}\right) \times 10^{-5}$	$\left(1.91^{+8.38}_{-0.77}\right) \times 10^{-4}$
$\langle m \rangle$ [M$_\odot$]	$\left(1.45^{+0.38}_{-0.32}\right) \times 10^7$	$\left(2.77^{+1.16}_{-1.01}\right) \times 10^7$	$\left(8.06^{+37.30}_{-3.25}\right) \times 10^7$
$m_{\rm eff} \equiv \langle m^2 \rangle / \langle m \rangle$ [M$_\odot$]	$\left(3.75^{+0.70}_{-0.76}\right) \times 10^7$	$\left(1.47^{+3.26}_{-0.88}\right) \times 10^8$	$\left(1.31^{+17.30}_{-0.36}\right) \times 10^9$
$r_{\rm t,max}$ [kpc]	$13.33^{+59.01}_{-5.75}$	$20.22^{+34.73}_{-10.95}$	$26.57^{+25.22}_{-14.72}$
$r_{\rm s,min}$ [kpc]	$0.11^{+0.02}_{-0.04}$	$0.11^{+0.02}_{-0.04}$	$0.11^{+0.02}_{-0.04}$

Table D.6: Relevant properties of the ETHOS4 subhalo population at $z = 1$

	$m_{\rm high} = 10^8$ M$_\odot$	$m_{\rm high} = 10^9$ M$_\odot$	All subhalos
$N_{\rm sub}(L_{\rm box} = 300$ kpc$)$	462	511	521
$N_{\rm sub}(L_{\rm box} = 100$ kpc$)$	56^{+14}_{-14}	62^{+12}_{-14}	65^{+11}_{-20}
$N_{\rm sub}(R_{\rm E})$	1^{+3}_{-1}	1^{+3}_{-1}	1^{+3}_{-1}
$\bar{\kappa}_{\rm sub}$	$\left(1.17^{+0.55}_{-0.44}\right) \times 10^{-5}$	$\left(4.68^{+2.68}_{-2.74}\right) \times 10^{-5}$	$\left(9.58^{+11.30}_{-6.86}\right) \times 10^{-5}$
$\langle m \rangle$ [M$_\odot$]	$\left(9.33^{+2.87}_{-2.75}\right) \times 10^6$	$\left(3.33^{+2.28}_{-1.59}\right) \times 10^7$	$\left(6.26^{+10.10}_{-3.81}\right) \times 10^7$
$m_{\rm eff} \equiv \langle m^2 \rangle / \langle m \rangle$ [M$_\odot$]	$\left(3.32^{+1.11}_{-1.75}\right) \times 10^7$	$\left(2.71^{+2.41}_{-1.70}\right) \times 10^8$	$\left(7.84^{+49.50}_{-5.85}\right) \times 10^8$
$r_{\rm t,max}$ [kpc]	$14.17^{+21.39}_{-5.24}$	$22.73^{+22.16}_{-11.53}$	$23.11^{+21.78}_{-9.12}$
$r_{\rm s,min}$ [kpc]	$0.12^{+0.04}_{-0.04}$	$0.12^{+0.04}_{-0.04}$	$0.12^{+0.04}_{-0.02}$

Appendix E

Quantifying the Line-of-Sight Halo Contribution to the Dark Matter Convergence Power Spectrum from Strong Gravitational Lenses

E.1 Scaling Relations for tNFW Projected Density

In this Appendix we derive the scaling laws for the surface density that we used in §7.1.2. Let us start with Eq. (7.33), which is trivial because scaling the surface density everywhere by a constant amount scales the total mass by the same amount. Eq. (7.34) can be derived by first denoting

$$\Sigma(\vec{r};\, m', r_{\rm s}', \tau) = \Sigma(\eta\vec{r};\, m, r_{\rm s}, \tau). \tag{E.1}$$

The new mass m' can then be found by integrating over the whole 2D plane:

$$m' = \int d^2\vec{r}\, \Sigma(\vec{r};\, m', r'_{\rm s}, \tau') \tag{E.2}$$

$$= \frac{1}{\eta^2} \int d^2(\eta\vec{r})\, \Sigma(\eta\vec{r};\, m, r_{\rm s}, \tau) \tag{E.3}$$

$$= \frac{m}{\eta^2}. \tag{E.4}$$

The new scale radius can be found by setting $|\vec{r}| = r'_{\rm s}$ in Eq. (E.1), which gives the relation

$$r'_{\rm s} = \frac{r_{\rm s}}{\eta}. \tag{E.5}$$

E.2 2D Fourier Transform of a Projection

In this Appendix we derive a useful relationship between the Fourier transform of the dimensionless tNFW density profile ϕ (defined in Eq. 7.47) and its projection.

Let us write the effective convergence as

$$\kappa_{i,\rm eff}(s) = \frac{1}{\Sigma_{{\rm crit},l}} \frac{m_{{\rm eff},i}}{r^2_{s,{\rm eff},i}} \Theta\left(\frac{s}{r_{s,{\rm eff},i}}\,;\, \tau_i\right), \tag{E.6}$$

where Θ is defined in terms of ϕ as

$$\Theta(t\,;\,\tau) \equiv \int dw\, \phi(\sqrt{t^2+w^2}\,;\,\tau) \tag{E.7}$$

and its Fourier transform is

$$\tilde{\kappa}_{i,\rm eff}(k) = \frac{m_{{\rm eff},i}}{\Sigma_{{\rm crit},l}} \tilde{\Theta}\left(r_{s,{\rm eff},i} k\,;\, \tau_i\right). \tag{E.8}$$

If we write the Fourier transform of Θ explicitly, we can relate it to the Fourier transform

of ϕ. Let us start with

$$\begin{aligned}\tilde{\Theta}(k\,;\,\tau) &= \int d^2\vec{t}\,\exp[-i\vec{k}\cdot\vec{t}]\Theta(t\,;\,\tau)\\ &= \int d^2\vec{t}\,dw\,\exp[-i\vec{k}\cdot\vec{t}]\phi(\sqrt{t^2+w^2}\,;\,\tau).\end{aligned} \tag{E.9}$$

We can combine $(\vec{t}, w)$ into a 3D vector $\vec{r}$. We can also replace $\vec{k}\cdot\vec{t}$ with $\vec{k}\cdot\vec{r}$ since $\vec{k}$ is perpendicular to the z-axis. We then obtain

$$\begin{aligned}\tilde{\Theta}(k\,;\,\tau) &= \int d^3\vec{r}\,\exp[-i\vec{k}\cdot\vec{r}]\phi(r\,;\,\tau)\\ &= \tilde{\phi}(k\,;\,\tau).\end{aligned} \tag{E.10}$$

Therefore, we can write the Fourier transform of the effective convergence of each interloper i as

$$\tilde{\kappa}_{i,\mathrm{eff}}(k) = \frac{m_{\mathrm{eff},i}}{\Sigma_{\mathrm{crit},l}}\tilde{\phi}\left(r_{s,\mathrm{eff},i}k\,;\,\tau_i\right). \tag{E.11}$$

E.3 The Area of the Double Cone

The radius of the double cone (Fig. 7.3) increases linearly from 0 to its maximum at the lens with comoving distance from the observer χ, and it decreases linearly from its maximum at the lens to 0 with the comoving distance from the lens $\chi - \chi_l$. With this in mind, we can write the follow relation for the radius of the double cone $R(\chi)$:

$$\frac{R(\chi)}{R(\chi_l)} = \begin{cases} a_1(\chi + b_1) & \chi \leq \chi_l \\ \\ a_2(\chi + b_2) & \chi > \chi_l, \end{cases} \tag{E.12}$$

where a_1, a_2, b_1, and b_2 are the linearity constants which we will obtain from the boundary values. We know that

$$\frac{R(0)}{R(\chi_l)} = 0 \quad \Rightarrow \quad b_1 = 0 \tag{E.13}$$

$$\frac{R(\chi_l)}{R(\chi_l)} = 1 \quad \Rightarrow \quad a_{1,2}(\chi_l + b_{1,2}) = 1 \tag{E.14}$$

$$\frac{R(\chi_s)}{R(\chi_l)} = 0 \quad \Rightarrow \quad a_2(\chi_s - b_2) = 0. \tag{E.15}$$

Solving for the constants we get

$$\frac{R(\chi)}{R(\chi_l)} = \begin{cases} \dfrac{\chi}{\chi_l} & \chi \leq \chi_l \\ \dfrac{\chi_s - \chi}{\chi_s - \chi_l} & \chi > \chi_l. \end{cases} \tag{E.16}$$

Then, the normalized area is

$$\frac{S(\chi)}{A/a^2(\chi_l)} = \begin{cases} \left(\dfrac{\chi}{\chi_l}\right)^2 & \chi \leq \chi_l \\ \left(\dfrac{\chi_s - \chi}{\chi_s - \chi_l}\right)^2 & \chi > \chi_l \end{cases} \tag{E.17}$$

$$= \frac{\chi^2}{\chi_l^2} g^2(\chi) \tag{E.18}$$

where A is the physical area of the main lens.

E.4 Number density of subhalos

Here we derive Eq. (7.72), which gives the number of subhalos per area as a function of mass and lens redshift. In particular, the expression is shown in terms of the fraction of halo mass in substructure, f_{sub,z_l}, which is commonly used in the literature as a proxy for the subhalo mass function normalization.

At any redshift, the total mass in subhalos within the mass range $[10^5 - 10^8]$ $\mathrm{M}_\odot$ can

be written as

$$M_{\rm sub} = A\bar{\kappa}_{\rm sub}\Sigma_{\rm crit}, \tag{E.19}$$

where A is the area of the lens, and $\bar{\kappa}_{\rm sub}$ is the average convergence of the subhalos. We assume a moderately elliptical isothermal lens, so that the convergence of the main lens near the Einstein radius is $\kappa_l = 0.5$, and roughly 40% of matter within the Einstein radius comes from stars [224]. Therefore, we can write

$$f_{{\rm sub},z_l} = \frac{2}{0.6}\bar{\kappa}_{\rm sub}, \tag{E.20}$$

which we can substitute into Eq. (E.19) and get

$$M_{\rm sub} = \frac{0.6}{2}Af_{{\rm sub},z_l}\Sigma_{{\rm crit},z_l}. \tag{E.21}$$

Our goal is to express the normalization of $n_{\rm sub}$ evaluated at $z_l = 0.5$, denoted F, in terms of $f_{\rm sub,0.5}$ evaluated at that same redshift. So we write

$$n_{\rm sub}(m, z_l = 0.5) = Fm^{\beta}. \tag{E.22}$$

The total mass in subhalos can be calculated from $n_{\rm sub}$:

$$\begin{aligned} M_{\rm sub} &= \int dA \int_{m_{\rm low}}^{m_{\rm high}} dm\, m\, n_{\rm sub}(m, 0.5) \\ &= AF \int_{m_{\rm low}}^{m_{\rm high}} dm\, m^{\beta+1} \\ &= AF\frac{m_{\rm high}^{2+\beta} - m_{\rm low}^{2+\beta}}{2+\beta}. \end{aligned} \tag{E.23}$$

Combined with Eq. (E.21), we get

$$F = \Sigma_{\rm crit,0.5}\frac{0.6f_{\rm sub,0.5}}{2}\frac{2+\beta}{m_{\rm high}^{2+\beta} - m_{\rm low}^{2+\beta}}. \tag{E.24}$$

Plugging this into Eq. (E.22) we obtain the final expression for the number density

of subhalos:

$$n_{\rm sub}(m, z_l = 0.5) = \frac{0.3\Sigma_{\rm crit,0.5} f_{\rm sub,0.5}(2+\beta)}{(m_{\rm high}^{2+\beta} - m_{\rm low}^{2+\beta})} m^{\beta}. \tag{E.25}$$

The redshift dependence of the halo number density per comoving area is found to be $(1+z_l)^{1/2}$ [218]. The number density per physical area will then depend on lens redshift as $(1+z_l)^{5/2}$ where the extra factor of 2 comes from the scale factor. We now have the final expression for the subhalo mass function:

$$n_{\rm sub}(m, z_l) = n_{\rm sub}(m, z_l = 0.5) \frac{(1+z_l)^{5/2}}{(1+0.5)^{5/2}}. \tag{E.26}$$

E.5 $k \to 0$ Limit of the Perturber Power Spectrum

For interlopers, we first calculate the Fourier transform of the profile ϕ (defined in Eq. 7.47) in the $k \to 0$ limit. Letting $\zeta \equiv (D_l r_s / g(\chi) D_\chi) k$, we can write

$$\begin{aligned} \lim_{\zeta \to 0} \tilde{\phi}(\zeta\,;\,\tau) &= \int_0^\infty 4\pi\xi^2\, d\xi \, \lim_{\zeta\to 0} \left[\frac{\sin(\zeta\xi)}{\zeta\xi}\right] \phi(\xi\,;\,\tau) \\ &= \int_0^\infty 4\pi\xi^2\, d\xi \, \phi(\xi\,;\,\tau) = 1. \end{aligned} \tag{E.27}$$

This means that in the $k \to 0$ limit, Eq. (7.57) gives

$$\begin{aligned} P_{\rm I,0} \equiv P_{\rm I}(k \to 0) = &\left(\frac{4\pi G}{c^2}\right)^2 D_l^2 \int_0^{\chi_s} d\chi \, \frac{W_{\rm I}^2(\chi)}{g^2(\chi)\chi^2} \\ &\times \int_{m_{\rm low}}^{m_{\rm high}} dm\, n(m,\chi)\, m^2, \end{aligned} \tag{E.28}$$

where $W_{\rm I}(\chi)$ and $g(\chi)$ depend on z_l and z_s. The intrinsic halo parameters completely drop out because $\int d^2\vec{q}\,\mathcal{P}(\vec{q}\,|\,m,\chi) = 1$.

From Eq. (E.28), there are two limits where the interloper power spectrum (in Fig. 7.4) goes to zero. In one limit, $z_l \to 0$, the power goes to zero due to the factor of D_l^2. In the other limit, $z_l \to z_s$, the power goes to zero because $f(\chi) \to 0$ for all χ, and W_I in the integrand contains a factor of $f(\chi)$.

For subhalos, we take the $k \to 0$ limit by applying Eq. (E.27) to Eq. (7.71), which gives

$$P_{\mathrm{S},0} = \frac{1}{\Sigma_{\mathrm{crit}}^2} \int_{m_{\mathrm{low}}}^{m_{\mathrm{high}}} dm\, m^2 n_{\mathrm{sub}}(m). \tag{E.29}$$

E.6 Curl and Divergence Components

Helmholtz Decomposition in 2D: We can express a 2D vector function $\vec{\alpha}(\vec{x})$ in a volume V from its divergence $\nabla \cdot \vec{\alpha}$, curl $\nabla \times \vec{\alpha}$, and its values on the boundary ∂V. We use the fact that $\frac{1}{2\pi} \ln |\vec{x} - \vec{x}'|$ is the Green's function for the Laplacian in 2D, i.e. we can write

$$\delta_{\mathrm{2D}}(\vec{x} - \vec{x}') = \frac{1}{2\pi} \nabla^2 \ln |\vec{x} - \vec{x}'|, \tag{E.30}$$

where δ_{2D} is the 2D delta function and ∇^2 is the Laplace operator that acts on $\vec{x}$.

We have

$$\vec{\alpha}(\vec{x}) = \int_V dV'\, \vec{\alpha}(\vec{x}') \delta(\vec{x} - \vec{x}') \tag{E.31}$$

$$= \frac{1}{2\pi} \nabla^2 \int_V dV'\, \vec{\alpha}(\vec{x}') \ln |\vec{x} - \vec{x}'|. \tag{E.32}$$

Using the identities

$$\nabla^2 \vec{q} = \nabla(\nabla \cdot \vec{q}) - \nabla \times (\nabla \times \vec{q}) \tag{E.33}$$

$$\vec{q} \cdot \nabla \phi = -\phi(\nabla \cdot \vec{q}) + \nabla \cdot (\phi \vec{q}) \tag{E.34}$$

$$\vec{q} \times \nabla \phi = \phi(\nabla \times \vec{q}) - \nabla \times (\phi \vec{q}), \tag{E.35}$$

we obtain

$$\vec{\alpha}(\vec{x}) = \frac{1}{2\pi}\nabla\left(\nabla\cdot\int_V dV'\,\vec{\alpha}(\vec{x}')\ln|\vec{x}'-\vec{x}|\right) - \frac{1}{2\pi}\nabla\times\left(\nabla\times\int_V dV'\,\vec{\alpha}(\vec{x}')\ln|\vec{x}'-\vec{x}|\right) \tag{E.36}$$

$$\vec{\alpha}(\vec{x}) = -\frac{1}{2\pi}\nabla\left(\int_V dV'\,\vec{\alpha}(\vec{x}')\cdot\nabla'\ln|\vec{x}'-\vec{x}|\right) - \frac{1}{2\pi}\nabla\times\left(\int_V dV'\,\vec{\alpha}(\vec{x}')\times\nabla'\ln|\vec{x}'-\vec{x}|\right) \tag{E.37}$$

$$\begin{aligned}\vec{\alpha}(\vec{x}) = {} & \frac{1}{2\pi}\nabla\left(\int_V dV'\,\ln|\vec{x}-\vec{x}'|\nabla'\cdot\vec{\alpha}(\vec{x}')\right) \\ & - \frac{1}{2\pi}\nabla\left(\int_V dV'\,\nabla'\cdot[\ln|\vec{x}-\vec{x}'|\vec{\alpha}(\vec{x}')]\right) \\ & + \frac{1}{2\pi}\nabla\times\left(\int_V dV'\,\ln|\vec{x}-\vec{x}'|\nabla'\times\vec{\alpha}(\vec{x}')\right) \\ & - \frac{1}{2\pi}\nabla\times\left(\int_V dV'\,\nabla'\times[\ln|\vec{x}-\vec{x}'|\vec{\alpha}(\vec{x}')]\right),\end{aligned} \tag{E.38}$$

where ∇' acts on $\vec{x}'$. Now we use the divergence theorem to write

$$\begin{aligned}\vec{\alpha}(\vec{x}) = {} & \frac{1}{2\pi}\nabla\left(\int_V dV'\,\ln|\vec{x}-\vec{x}'|\nabla'\cdot\vec{\alpha}(\vec{x}')\right) \\ & - \frac{1}{2\pi}\nabla\left(\oint_{\partial V} dS'\,\hat{n}'\cdot[\ln|\vec{x}-\vec{x}'|\vec{\alpha}(\vec{x}')]\right) \\ & - \frac{1}{2\pi}\nabla\times\left(\int_V dV'\,\ln|\vec{x}-\vec{x}'|\nabla'\times\vec{\alpha}(\vec{x}')\right) \\ & + \frac{1}{2\pi}\nabla\times\left(\oint_{\partial V} dS'\,\hat{n}'\times[\ln|\vec{x}-\vec{x}'|\vec{\alpha}(\vec{x}')]\right),\end{aligned} \tag{E.39}$$

where $\hat{n}'$ is the unit vector normal to the boundary ∂V. If $\vec{\alpha}$ vanishes faster than $1/|\vec{x}'|\ln|\vec{x}'|$, the boundary terms vanish as we make V infinitely large. This allows us to

write

$$\begin{aligned}
\vec{\alpha}(\vec{x}) &= \frac{1}{2\pi}\nabla\left(\int_V dV'\,\ln|\vec{x}-\vec{x}'|\nabla'\cdot\vec{\alpha}(\vec{x}')\right) \\
&\quad -\frac{1}{2\pi}\nabla\times\left(\int_V dV'\,\ln|\vec{x}-\vec{x}'|\nabla'\times\vec{\alpha}(\vec{x}')\right) \qquad \text{(E.40)} \\
&= \frac{1}{\pi}\int_V dV'\,\{\nabla\cdot\ln|\vec{x}-\vec{x}'|\}\left[\frac{1}{2}\nabla'\cdot\vec{\alpha}(\vec{x}')\right] \\
&\quad -\frac{1}{\pi}\int_V dV'\,\{\nabla\times\ln|\vec{x}-\vec{x}'|\}\left[\frac{1}{2}\nabla'\times\vec{\alpha}(\vec{x}')\right] \qquad \text{(E.41)} \\
&= \frac{1}{\pi}\int_V dV'\left\{\frac{\vec{x}-\vec{x}'}{|\vec{x}-\vec{x}'|^2}\right\}\underbrace{\left[\frac{1}{2}\nabla'\cdot\vec{\alpha}(\vec{x}')\right]}_{\kappa_{\text{eff}}} \\
&\quad -\frac{1}{\pi}\int_V dV'\left\{\hat{x}_3\times\frac{\vec{x}-\vec{x}'}{|\vec{x}-\vec{x}'|^2}\right\}\underbrace{\left[\frac{1}{2}\nabla'\times\vec{\alpha}(\vec{x}')\right]}_{\kappa_{\text{curl}}}, \qquad \text{(E.42)}
\end{aligned}$$

where $\hat{x}_3 \equiv \hat{x}_1 \times \hat{x}_2$ is the unit vector that is perpendicular to the 2D plane on which $\vec{\alpha}$ lives. $\hat{x}_1$ and $\hat{x}_2$ are the unit vectors of the 2D plane.

Five-Point Stencil: For a function $f : \mathbb{R} \to \mathbb{R}$, the first derivative can be approximated by [331],

$$f'(x) = \text{5pt}[f](x\,;\,h) + \frac{h^4}{30}f^{(5)}(x) + \mathcal{O}(h^5), \qquad \text{(E.43)}$$

where

$$\begin{aligned}
\text{5pt}[f](x\,;\,h) \equiv\ & \frac{2}{3}\left[f(x+h) - f(x-h)\right] \\
& -\frac{1}{12}\left[f(x+2h) - f(x-2h)\right]. \qquad \text{(E.44)}
\end{aligned}$$

The error scales with the fourth power of the discrete interval size h and the 5th derivative of the function. Defining $\vec{x} \equiv (x_1, x_2)$, $\alpha_1(x) \equiv \hat{x}_1 \cdot \vec{\alpha}(x, x_2)$, and $\alpha_2(x) \equiv \hat{x}_2 \cdot \vec{\alpha}(x_1, x)$, the divergence of $\vec{\alpha}$ can be calculated as

$$\nabla\cdot\vec{\alpha} \cong \text{5pt}[\alpha_1](x_1\,;\,h) + \text{5pt}[\alpha_2](x_2\,;\,h), \qquad \text{(E.45)}$$

where the leading error term is

$$\frac{h^4}{30}\left(\alpha_1^{(5)}(x_1) + \alpha_2^{(5)}(x_2)\right). \tag{E.46}$$

Numerical Artifacts: As with the divergence, we calculate the curl using the five-point stencil. Defining $\vec{x} \equiv (x_1, x_2)$, $\alpha_1(x) \equiv \hat{x}_1 \cdot \vec{\alpha}(x_1, x)$, and $\alpha_2(x) \equiv \hat{x}_2 \cdot \vec{\alpha}(x, x_2)$, the curl of $\vec{\alpha}$ is calculated as

$$\nabla \times \vec{\alpha} \cong 5\text{pt}[\alpha_1](x_2\,;\,h) - 5\text{pt}[\alpha_2](x_1\,;\,h), \tag{E.47}$$

where the leading error term is

$$\frac{h^4}{30}\left(\alpha_1^{(5)}(x_2) - \alpha_2^{(5)}(x_1)\right). \tag{E.48}$$

To study how this numerical effect appears, we simulate a lensing system with a main lens and subhalos (i.e. no interlopers). From Eq. (7.77), we know that we should have $\kappa_{\text{curl}} = 0$. Therefore, any non-zero value we get after calculating the curl using Eq. (E.47) will be a numerical artifact. In Fig. E.6.1, we see that this numerical effect is only present at the centers of subhalos, as well as the center of the main lens, where the central cusp has a large 5th derivative, which increases the error. Nevertheless, it is more that 2 orders of magnitude smaller than the curl that we calculate in Fig. 7.7.

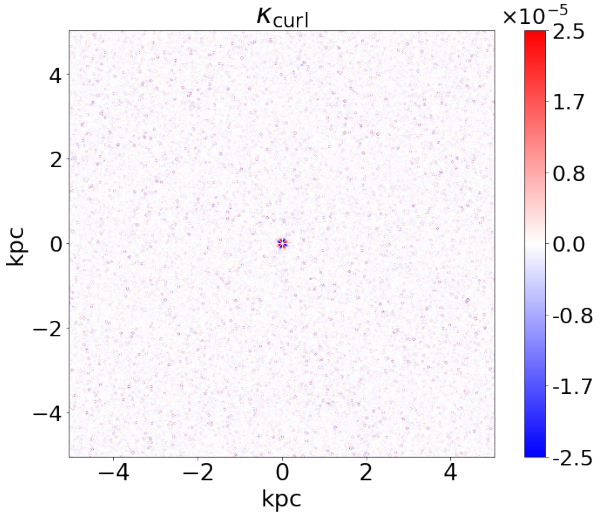

Figure E.6.1: κ_{curl} of a simulated lensing system described in Section 7.2.1 without the interlopers. Without interlopers there is nothing that can source a curl component. This figure shows the numerical error in calculating the curl using discrete pixel values.

Appendix F

Observable Predictions for Massive-Neutrino Cosmologies with Model-Independent Dark Energy

F.1 Results for the *Reduced* dataset

We show the 1σ confidence levels for the posteriors of the luminosity distance and the Hubble expansion rate (Figure F.1.1) and the matter power spectrum at $z = 0$ (Figure F.1.2) for the *Reduced* dataset, in analogy with Figures 8.4 and 8.5 in the main text for the *All* dataset. Note that we do not show the shear power spectrum since the *Reduced* dataset does not have weak lensing data.

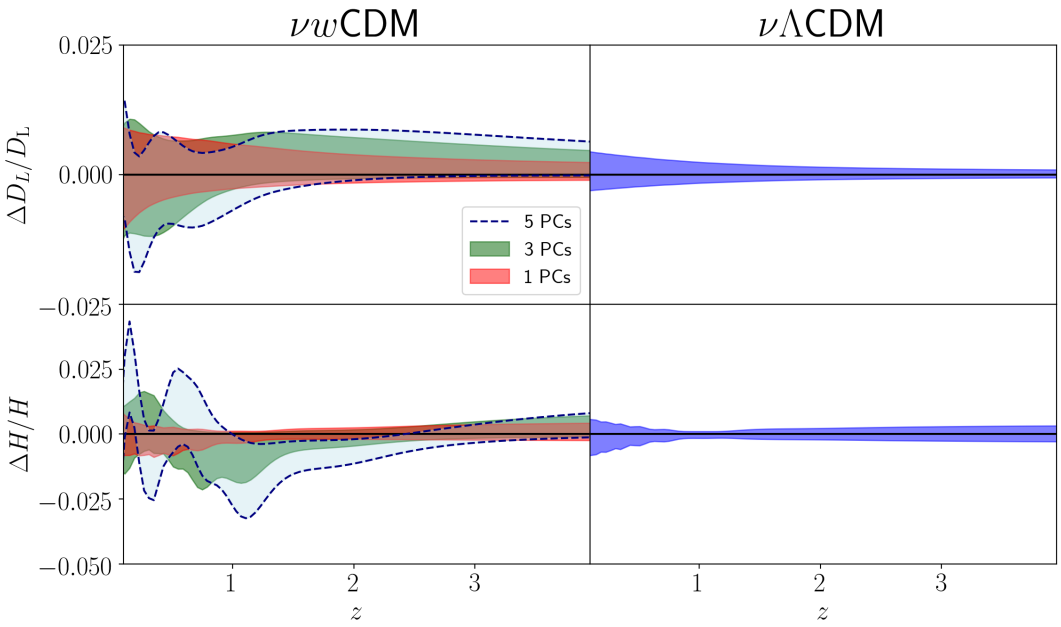

Figure F.1.1: Fractional difference for the luminosity distance $D_L(z)$ (top row) and the Hubble expansion rate as a function of redshift $H(z)$ (bottom row) for νwCDM (left) and $\nu\Lambda$CDM (right), with respect to ΛCDM, for the *Reduced* dataset. Bands correspond to 1σ confidence levels.

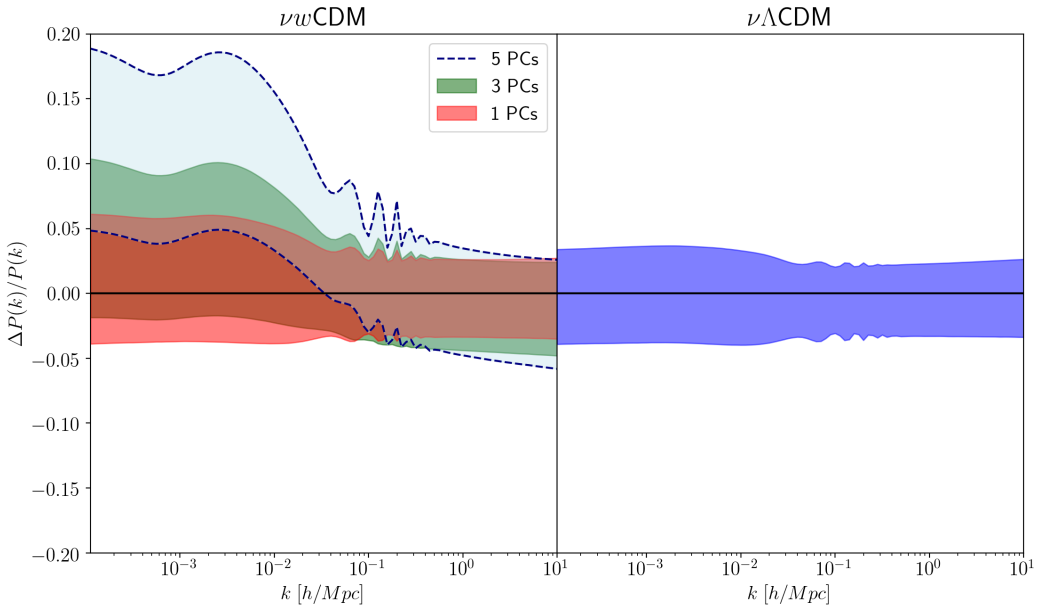

Figure F.1.2: Fractional difference for the matter power spectrum at $z = 0$ $P(k)$ for νwCDM (left) and $\nu\Lambda$CDM (right) with respect to ΛCDM, for the *Reduced* dataset. Bands correspond to 1σ confidence levels.

Appendix G

Flow-Based Likelihoods for Non-Gaussian Inference

G.1 Likelihood versus Sampling density

Throughout this chapter, we have referred to the sampling density learned from the data, $\hat{p}_{\rm DDL}(\mathbf{x})$, as a learned likelihood, although these are distinct entities. In this Appendix we explain the reason behind our use of this terminology, and reconcile the difference between these two functions.

For some data $\mathbf{x}$, the likelihood is a function of the model parameters $\boldsymbol{\theta}$: $\mathcal{L}_{\mathbf{x}}(\boldsymbol{\theta})$. The likelihood is not a probability density function, and as such does not have the same restrictions (e.g., having to integrate to unity in the case of a continuous distribution). Conversely, for a choice of parameters $\boldsymbol{\theta}$ the probability density function is in fact normalized to unity when integrated over the data: $\int d\mathbf{x} p_{\boldsymbol{\theta}}(\mathbf{x}) = 1$. Although they have the exact same functional form (up to constant factors), they represent different things: the likelihood is a function of the model parameters with the data fixed, while the PDF is a function of the data with the model parameters fixed.

In this chapter, for the DDL methods discussed we used N mocks for a given value of the cosmological model parameters $\boldsymbol{\theta}$. This is analogous to likelihood analyses that use the mocks for a given cosmology to obtain the covariance matrix and then employ a

MVN likelihood:

$$p(\mathbf{x}|\boldsymbol{\theta}) \propto \exp\left[(\mathbf{x}-\boldsymbol{\mu}(\boldsymbol{\theta}))^T\widehat{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu}(\boldsymbol{\theta}))\right], \tag{G.1}$$

where we have put a hat on the precision matrix to emphasize that it is an estimated quantity. This means that Σ is evaluated at some fiducial cosmology and assumed to be cosmology-independent. In principle, Σ does vary with cosmology, and this dependence can have a significant impact on inferred parameters [332–334].

Note that the dependence on the cosmological parameters enters the likelihood through the estimate of the mean $\boldsymbol{\mu}(\boldsymbol{\theta})$. We can rewrite Eq. (G.1) as

$$p(\mathbf{x}|\boldsymbol{\theta}) \propto \exp\left[\Delta(\boldsymbol{\theta})^T\widehat{\Sigma}^{-1}\Delta(\boldsymbol{\theta})\right], \tag{G.2}$$

to emphasize that the parameter dependence is encoded in the difference between the data and the model $\Delta(\boldsymbol{\theta})$. To find the parameters that provide the best fit for the data, it is necessary to have a model to obtain $\boldsymbol{\mu}$ for any point in the allowed parameter space.

In this chapter, our estimates from mock data are analogous to the step of estimating the covariance matrix in the MVN example above in the sense that all the mocks correspond to a single value of the cosmological parameters. Unlike in the MVN case, what we estimate from the mocks is the full sampling density $\hat{p}_{\rm DDL}(\mathbf{x})$. However, we interchangeably use the term DDL to refer to this function because it can in fact be used as a likelihood [277]. All we require is an array $\Delta(\boldsymbol{\theta})$ and we can apply $\hat{p}_{\rm DDL}$ as a likelihood, which will now depend on the cosmological parameters: $\hat{p}_{\rm DDL}(\boldsymbol{x}|\boldsymbol{\theta})$. Thus, the terminology data-driven likelihoods in this chapter foreshadows the possibility of using them as true likelihoods for inference. Although the focus here was showing their ability to capture NG in the data, using them as actual likelihoods is the ultimate goal of developing this method, and is done in forthcoming work.

We emphasize that, just like in the case for a MVN likelihood, where using a fixed-cosmology covariance matrix is an approximation made for convenience (otherwise every point in the parameter space of $\boldsymbol{\theta}$ requires a whole suite of mock data), the impact of

using a DDL estimated with mock data from a single cosmology has to be studied. The cosmology dependence will also be addressed in forthcoming work.

G.2 Toy Problems

To study the quality of the likelihoods provided by a flow-based model such as FFJORD, we first analyze the results with data whose likelihood is known and we can sample from. The advantage of these toy problems is that (1) we can compare the likelihood given by the model and the true likelihood and (2) by being able to sample from the likelihood, we can obtain a virtually infinite number of unique training samples.

G.2.1 In two dimensions

Our first stepping stone is simple two-dimensional data. The benefit in two dimensions is that we can visualize the reconstructed PDF. We train FFJORD on a dataset whose PDF is comprised of 8 equal, symmetric Gaussians arranged around a circle. We use three stacked continuous normalizing flows with hidden dimensions $d = 256$.

The results are shown in Figure G.2.1. The top left panel is the true PDF, while the middle panel is the one obtained after training FFJORD. In Ref. [280], the authors also showed these two panels. What interests us is the difference between them and how this translates into log-likelihood values. The top right panel shows the residual, which is on the order of 10%. The middle row shows the same but for the marginalized probability distributions along the x and y axes. Finally, the bottom row shows the fractional log-likelihood residual between the true log-likelihood $\mathcal{L}$ and the model log-likelihood $\hat{\mathcal{L}}$. We can see that the scatter is around 10-15%.

G.2.2 In higher dimensions

The next step is gauging the flow-based likelihood obtained in higher dimensions, on the order of the mock observables discussed in the main text. We train FFJORD on samples

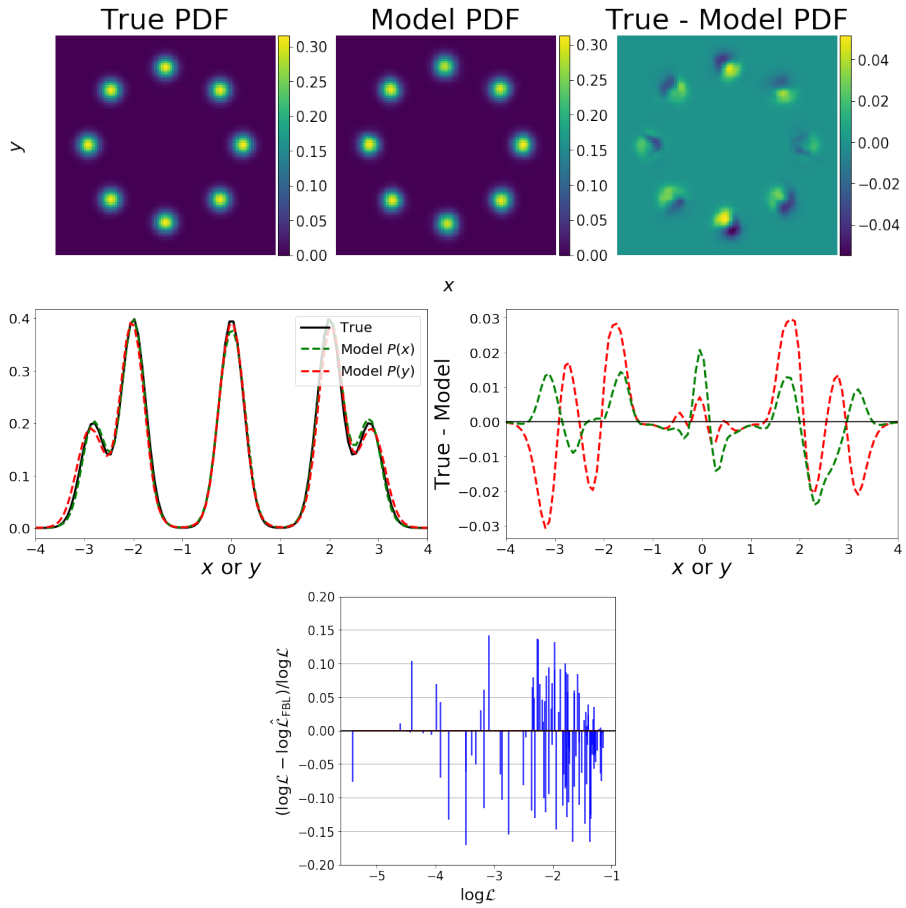

Figure G.2.1: *Top*: true (left) and flow-reconstructed PDF (middle), together with their residual (right). *Middle*: true and reconstructed marginal PDF (left) along x (red) or y (green), together with their residual (right). *Bottom*: fractional difference between the true log-likelihood $\mathcal{L}$ and the FBL $\hat{\mathcal{L}}$ as a function of the true log-likelihood.

drawn from a 30-dimensional Gaussian distribution with mean

$$\mu = 10\sin(x) \tag{G.3}$$

and a covariance matrix obtained from a squared exponential kernel K:

$$K(x, x') = \sigma^2 \exp\left(-\frac{(x - x')^2}{2L^2}\right), \tag{G.4}$$

where σ is the standard deviation and L the correlation lengthscale. We set $\sigma = 1$ and $L = \sqrt{8}$. Note that this choice of mean and covariance was completely arbitrary.

Something that we were interested in was how the learning was inhibited (or not) by using an approximately singular vs. non-singular covariance matrix. As given above, the squared exponential kernel using bin numbers as values for x and x' yields a singular covariance. By adding "noise" to the diagonal (adding the identity matrix), we can turn it into a full rank matrix.

The results training on data drawn from a non-singular covariance matrix are shown in Figure G.2.2. The two rows of Panel (a) show the difference between the reconstructed FBL covariance (precision) matrix and the true one, as well as between the sample MVN covariance (precision) and the true one. It can be seen that the error in the FBL matrices is of the same order than that of the sampled MVN. The left figure in Panel (b) shows the mean, 68% and 95% CLs from 2,000 samples drawn from a MVN with mean and covariance given by Eqs. (G.3) and (G.4) in blue, and the mean, 68% and 95% CLs obtained from sampling the learned FBL in red. The right one shows the difference between the true and sample FBL mean (dashed red) and true and sample MVN mean (blue), which shows that the error in the FBL's mean is on the same order of magnitude as the sampling error.

Finally, Panel (c) shows what the values of the log-likelihood are under a MVN for the test data with different combinations of μ and Σ, in blue. The green line common to all the subpanels is the true log-likelihood, and the red points are the log-likelihood given by the trained FBL for the same data. Clearly, the small scatter about the green

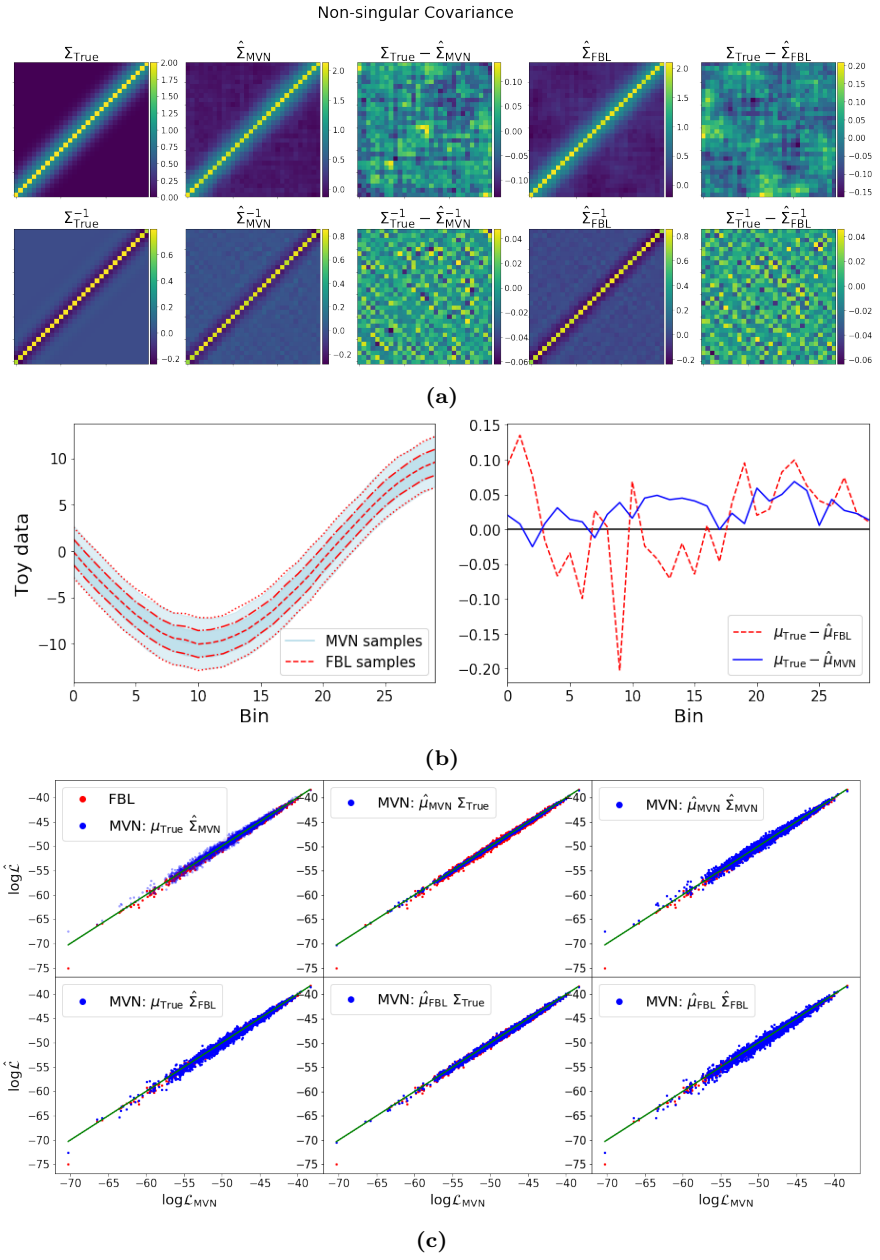

Figure G.2.2: Results when training FFJORD on Gaussian 30-dimensional data with mean and covariance given by Eqs. (G.3) and (G.4), and noise added to the covariance to make it full rank. (a): The true covariance Σ_{True} and precision matrix $\Sigma^{-1}_{\text{True}}$, together with the MVN sample covariance $\hat{\Sigma}_{\text{MVN}}$ and precision $\hat{\Sigma}^{-1}_{\text{MVN}}$, and FBL sample covariance $\hat{\Sigma}_{\text{FBL}}$ and precision $\hat{\Sigma}^{-1}_{\text{FBL}}$, obtained from 2,000 samples. It can be seen that the FFJORD-reconstructed matrices are accurate to within sampling error. (b): On the left, the mean, 68% and 95% CLs for the toy data in blue, and for 2,000 FBL samples in red. On the right, the difference between the true mean and the FBL sample mean (red) and between the true mean and MVN sample mean (blue). The FBL mean, 1 and 2σ contours match the data very well; the mean residual is of the same order of magnitude as the error due to sampling. (c): the green line corresponds to the true log-likelihood of the test data. In blue, the log-likelihoood of the test data under a MVN for different combinations of the mean and covariance in blue. In red, the log-likelihood of the test data under the FBL. Agreement between the red scatter points and the green line would reflect that the FBL has learned the likelihood perfectly. The small scatter around the green line can be attributed to sampling error, as it is also present when using a sample mean and covariance with a MVN likelihood.

line visible in the FBL values is of the same order than that due simply to sampling error when estimating the covariance matrix from a finite number of samples.

Figure G.2.3 shows the analogous results but for the case where no constant is added to the diagonal terms of the covariance matrix and it is thus singular. Everything else is the same as in the toy problem above. The interesting thing to notice is that, although the sample quality is still excellent, the quality of the likelihood is significantly worse. In Panel (c) we see that the FBL has a much larger scatter as well as a constant offset with respect to the true likelihood. Furthermore, the MVN likelihood is much more sensitive to tiny deviations away form the true mean, as can be seen in the middle-bottom subpanel: even though the difference between the FBL sample mean and the true mean is very small, the likelihood values are abysmal. This is a general feature of MVN likelihoods with (nearly) singular covariances. Comparing the true and reconstructed precision matrices also reflects the fact that the model has not been able to learn the likelihood correctly.

The dissonance between sample quality and likelihood quality is very interesting, and emphasizes the value in scrutinizing the likelihoods that flow-based models are learning, as we do in this work. The point is that simply sampling from the FBL and looking at the distribution of generated samples does not guarantee that the likelihood learned by the model is actually correct.

G.3 Data-driven Likelihoods in Data-Limited Regimes

As their name suggests, data-driven likelihoods can only truly thrive with plentiful data. Ultimately, these methods are trying to estimate densities (often in very high-dimensional spaces). We illustrate the potential shortcomings of reaching conclusions by using DDLs fit on a limited number of mocks in Figure G.3.1. It shows the same results as in Figure 9.6 for the ICA likelihood, but fit on 2,048 convergence power spectrum mocks instead of 75,000. Comparing these two figures it is clear that $2,048$ mocks are not enough for the DDL to capture the skewness and kurtosis in the first few bins. It is therefore possible that works that have operated in such data-limited regimes (e.g. [277, 313] have

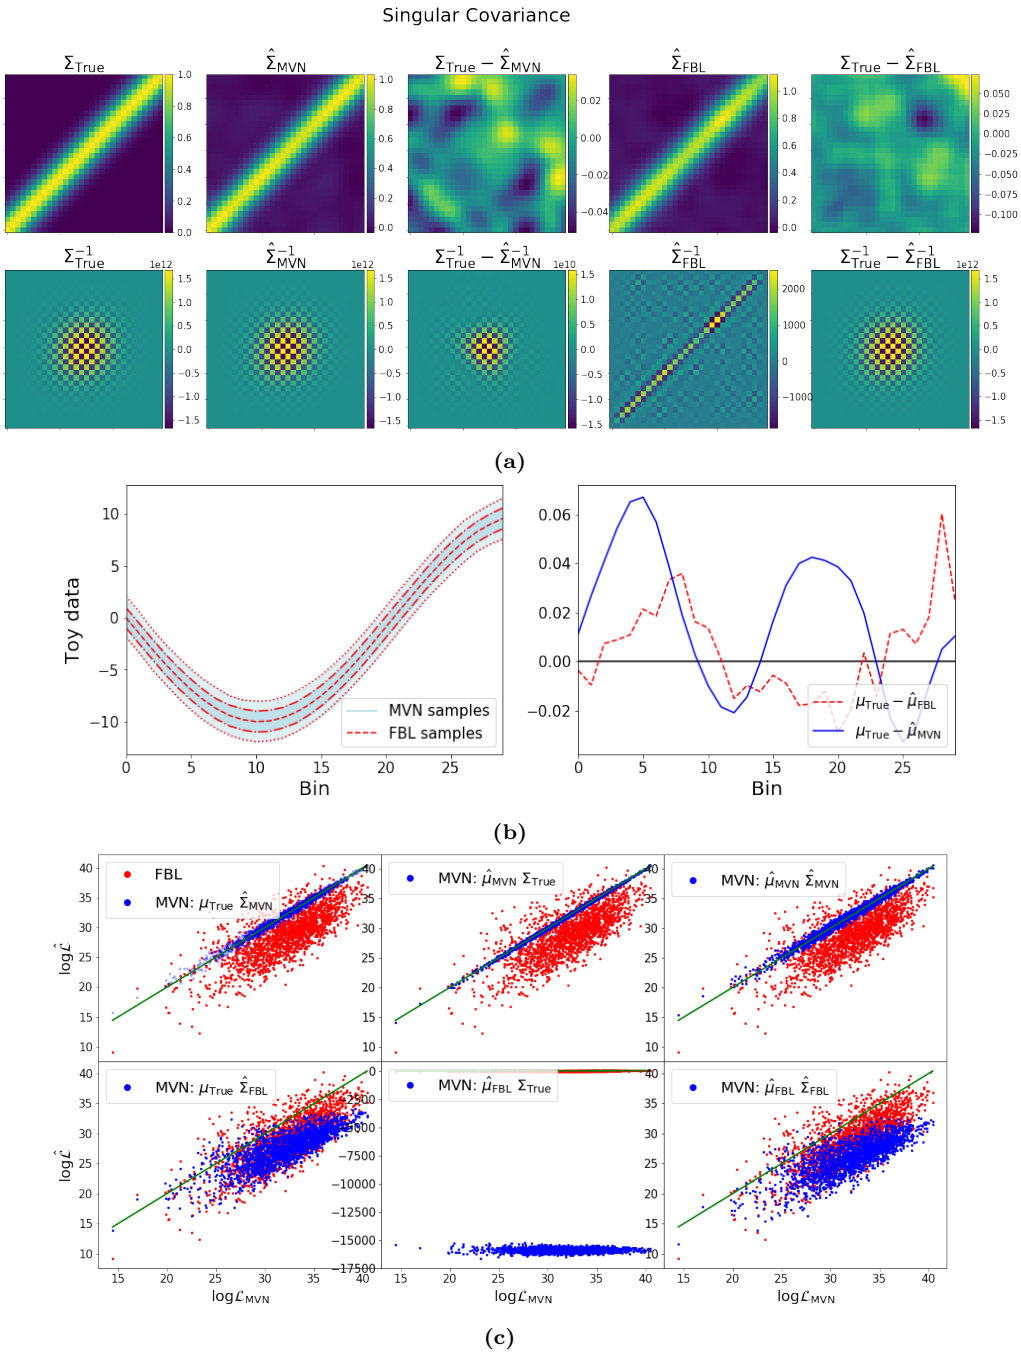

Figure G.2.3: Results when training FFJORD on Gaussian 30-dimensional data with mean and covariance given by Eqs. (G.3) and (G.4). (a): The true covariance Σ_{True} and precision matrix $\Sigma^{-1}_{\text{True}}$, together with the MVN sample covariance $\hat{\Sigma}_{\text{MVN}}$ and precision $\hat{\Sigma}^{-1}_{\text{MVN}}$, and FBL sample covariance $\hat{\Sigma}_{\text{FBL}}$ and precision $\hat{\Sigma}^{-1}_{\text{FBL}}$, obtained from 2,000 samples. (b): On the left, the mean, 68% and 95% CLs for the toy data in blue, and for 2,000 FBL samples in red. On the right, the difference between the true mean and the FBL sample mean (red) and between the true mean and MVN sample mean (blue). (c): the green line corresponds to the true log-likelihood of the test data. In blue, the log-likelihoood of the test data under a MVN for different combinations of the mean and covariance in blue. In red, the log-likelihood of the test data under the FBL.

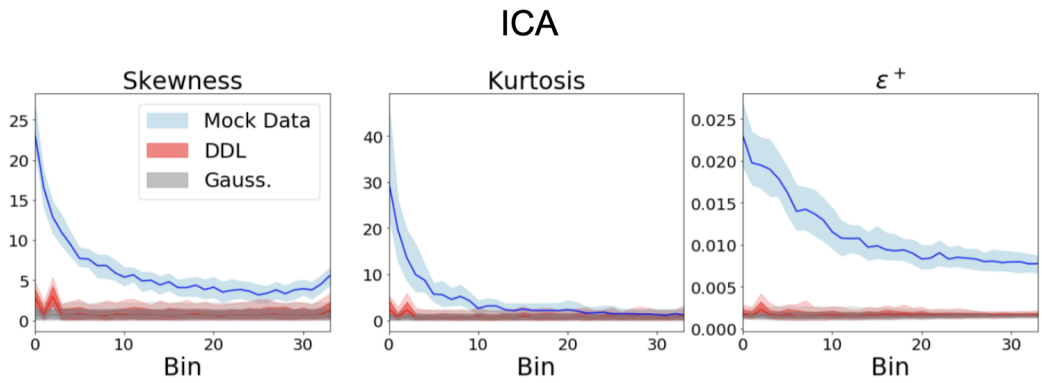

Figure G.3.1: Same as the first three panels for the ICA likelihood in Figure 9.6, but fit on 2,048 mocks instead of the full 75,000. From left to right: absolute value of the t-statistic of the skewness, kurtosis, and ϵ^+ of each bin for 100 different sets of 2,048 mock WL power spectra (blue), ICA samples (red), and Gaussian samples (gray). A vertical offset between the blue and gray contours reflects the NG in the data. Lack of overlap between the red and blue contours indicates that the ICA likelihood has not succeeded in capturing the NG.

underestimated the effect that NGs can have on inferred parameters.

G.4 Robustness of the nonparametric KL Divergence Estimator

In this appendix we test the robustness of the nonparametric KL Divergence estimator introduced in Section 9.2. To do so, we estimate the same reference KL divergence histogram as the one shown in gray in Figure G.5.1 for two different Gaussian datasets drawn from the same likelihood. Figure G.4.1 shows the results using different random seeds. It can be seen that the degree of overlap can vary quite a bit, even from samples drawn from the same underlying distribution. The leftmost panel, with an overlapping area equal to 0.19, is the lowest we found. Therefore, while the large horizontal offset seen between the reference and data histograms in Figure 9.3 and Figure 9.6 still seems statistically significant, the one seen for GMM2 in Figure 9.6 does not. This is likely due to the curse of dimensionality, since the KL estimator relies on a kNN algorithm, and further emphasizes the importance of including more fine-grained measures of NG that are more robust with limited numbers of samples when quantifying non-Gaussianity.

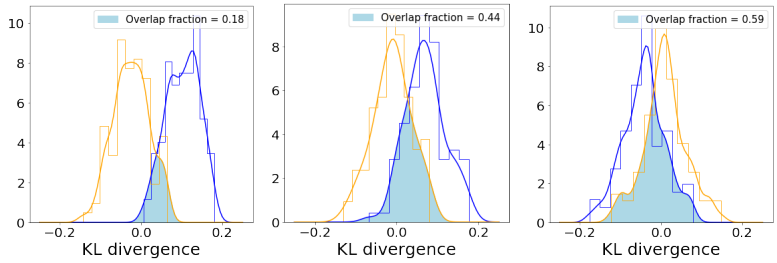

Figure G.4.1: KL divergence estimate (Eq. 9.23) of two ensembles of draws with respect to the same reference distribution, all drawn from a 34-dimensional multivariate Normal likelihood with mean and covariance from the weak lensing convergence power spectra mocks. Each panel corresponds to a different random seed.

G.5 Galaxy Power Spectrum Non-Gaussianity

Ref. [277] studied non-Gaussianities in two large-scale structure observables, one of them being the galaxy power spectrum. They used the MultiDark-PATCHY mock catalogs [335], which were built to match the BOSS Data Release 12 observations. In particular, they looked at the 2,048 mocks for the North Galactic Cap (NGC) in the redshift bin $0.2 < z < 0.5$ and obtained mock power spectra using `NBodyKit` [202]. They used the nonparametric KL divergence test described in Section 9.2 to establish the likelihood non-Gaussianity.

They then sought to build a data-driven likelihood that would incorporate the NG. They showed that the estimated KL divergence was unchanged with a GMM likelihood[1], but nearly vanished with an ICA likelihood. They used this likelihood to perform importance sampling on an MCMC chain that had been previously analyzed with a Gaussian likelihood [193] (essentially re-weighting the points by the ratio of the ICA likeliood to the MVN likelihood) and found small shifts ($< 0.5\sigma$) in relevant cosmological parameters.

We produced mock galaxy power spectra from the MultiDark-PATCHY catalogs using the same procedure as Ref. [277], so we refer the reader to them for details. Ultimately we obtain 2,048 mock power spectra in 37 bins, encompassing the power spectrum monopole, quadrupole, and hexadecapole. For the remainder of this section, however, we truncate

[1]When we try fitting a GMM to the same mock dataset, we find that the BIC is minimized for a single component. The fact that the BIC increases monotonically with more components is due to the number of mocks being too small to fit the large number of parameters without overfitting. For example, with our WL mocks we find that with 2,048 samples the BIC behaves similarly, while with more mocks it finds the minimum at $K = 2$.

the power spectrum at 34 bins, so that the NG tests (Section 9.2) are directly comparable to the ones in Section 9.3 for the WL power spectra (recall that the tests are extensive). Note that this means that our measures of non-Gaussianity are conservative compared to those in Ref. [277]. Just like we did for the WL mocks, we mean-subtract and whiten the data before running it through the NG tests.

Figure G.5.1 is analogous to the one shown for the WL mock data in the main text showing the non-Gaussianity results from our three tests. In addition, Figure G.5.2 shows the individual galaxy power spectrum bin distributions. Comparing the values of the t-statistics here and the ones for the WL mocks reveals that the skewness and kurtosis are much more pronounced in the latter. The same can be said when comparing the mock and Gaussian S^+ matrices, as well as the vertical offset between the mock and Gaussian ϵ^+.

Keeping in mind the discussion in the main text on the applicability of different DDLs depending on the type of NG in the data, we can see that the fact that ϵ^+ is virtually indistinguishable between the data and the Gaussian samples, while some individual bins do exhibit non-negligible skewness and kurtosis, could explain why the ICA likelihood worked better in this setting than the GMM did (though the small number of mocks might also be a contributing factor).

The top two panels of Figure G.5.1 also show the 68% and 95% CLs from 2,048 samples drawn from the ICA likelihood. It can be seen that, while slightly higher than the Gaussian ones, few of the data points with strong NG fall within the 2σ boundary of the ICA samples. This corroborates our results on WL data and Appendix G.4, which showed that even if the KL divergence is small between a catalog of data and a data-driven likelihood, the DDL may not be accurately capturing the NG in the data. We thus conclude that it is likely that the non-Gaussian signatures in galaxy power spectrum data could actually have a larger impact than was found in Ref. [277], but the limited number of mocks makes this a difficult task for a DDL to solve adequately.

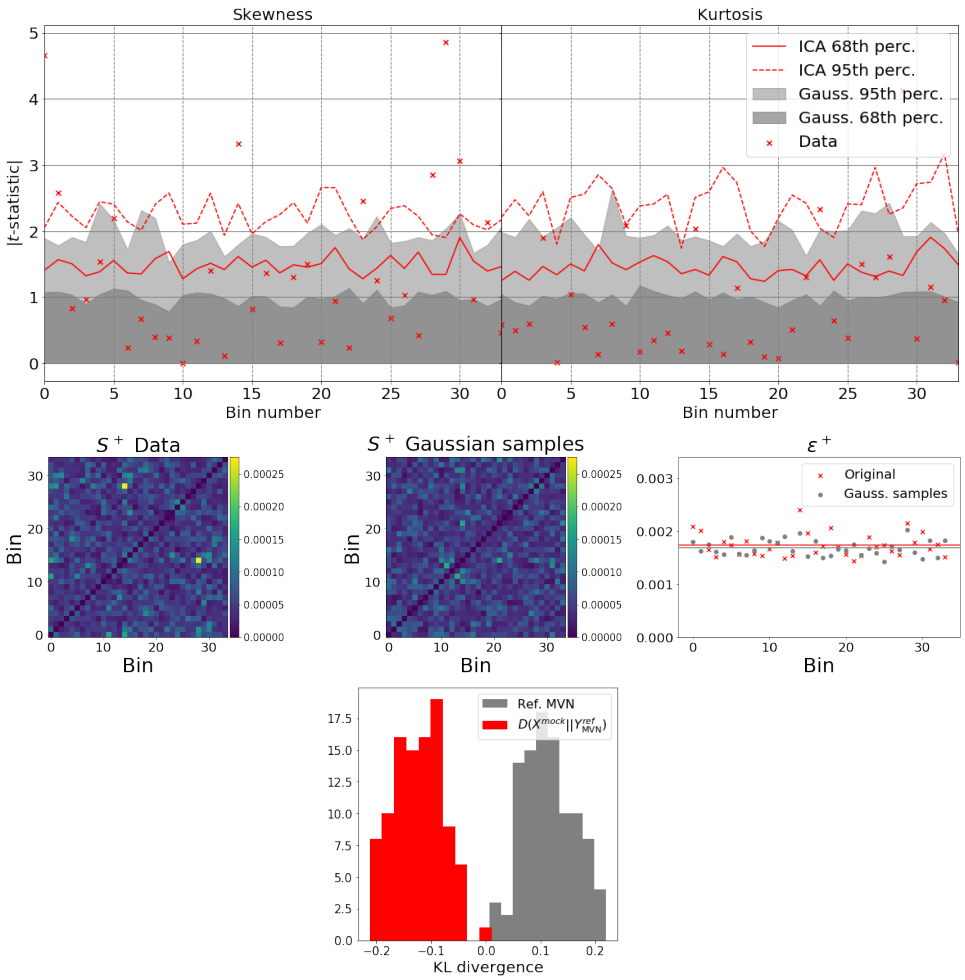

Figure G.5.1: Same as Figure 9.3 but for the BOSS DR12 mock galaxy power spectra, as a function of bin number. *Top*: the red crosses are the absolute value of the t-statistic of skewness (left) and kurtosis (right) of individual bins for 2,048 mock galaxy power spectra. The gray contours correspond to the one (dark gray) and two sigma (light gray) confidence level when averaging 100 different sets of 2,048 samples drawn from a multivariate normal likelihood with the same mean and covariance as the mock data. As expected, they correspond to t-statistic values of 1 (dark gray) and 2 (light gray). In addition, the red solid and dashed lines show the 68% and 95% CLs for 2,048 draws from the ICA likelihood fit on the data. Notice that ICA fails to capture the strong deviations from Gaussianity in the skewness. *Middle*: the S^+ matrix for the mock data (left) and equivalent Gaussian samples (middle). The sum over columns of each matrix, $\epsilon^+_u \equiv \sum S^+_{u,v}$, is shown on the right as red crosses and gray circles, respectively. *Bottom*: nonparametric KL divergence estimate between the mock data and their Gaussian counterparts (red), and the Gaussian samples with themselves (gray). The fact that the gray histogram is not perfectly centered at zero is due to the slight variability of the KL estimator in 34 dimensions , given the number of mocks considered (Appendix G.4).

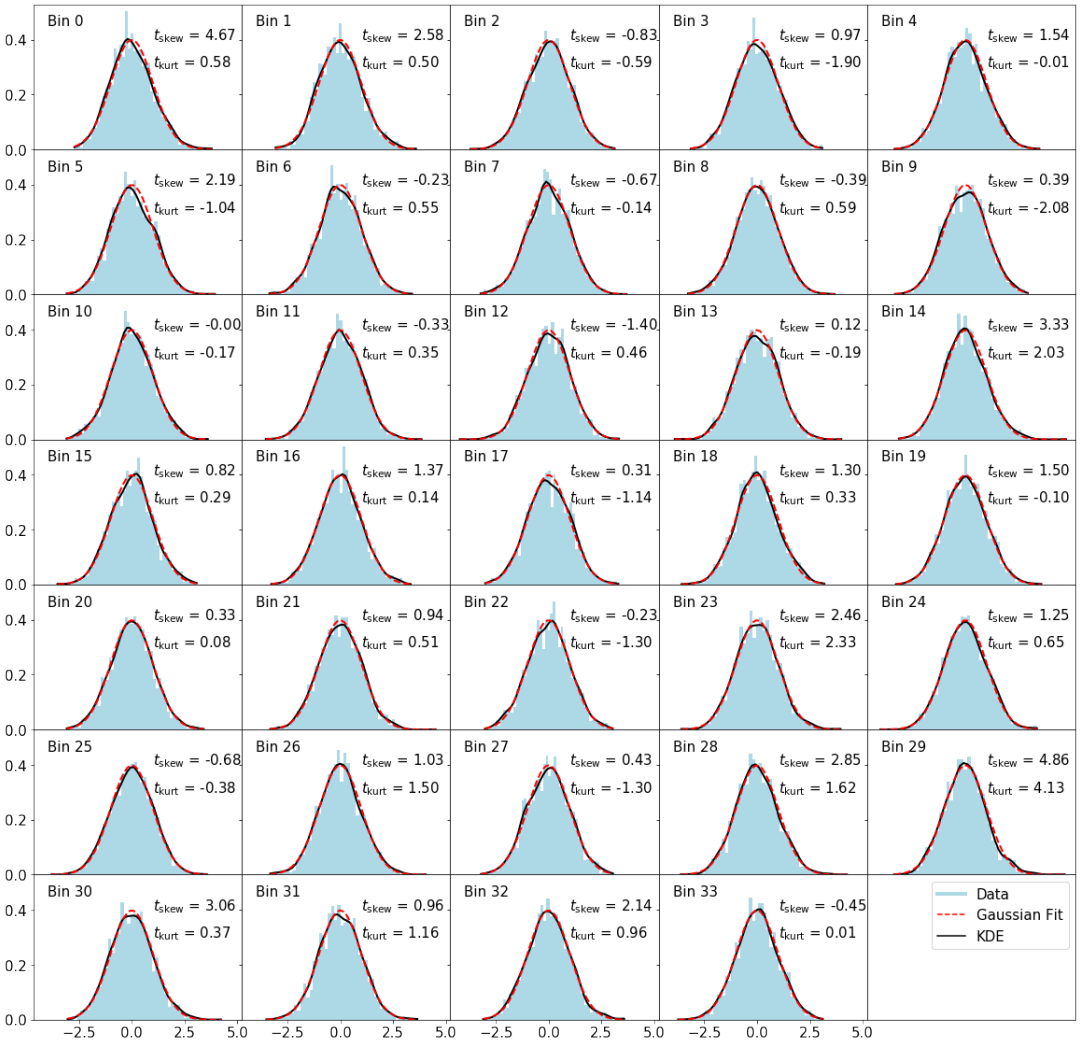

Figure G.5.2: Distribution for each of the 34 bins in the ensemble of 2,048 mock galaxy power spectra for the BOSS DR12 North Galactic Cap in the redshift bin $0.2 < z < 0.5$. The histogrammed values are shown in light blue, while the KDE is shown in black. The dashed red line is a Gaussian fit to each bin.

G.6 Mock Weak Lensing Convergence Maps

G.6.1 Generating the mocks

Generating mock WL maps involves running full N-body simulations, which consist of evolving millions of particles under the force of gravity to simulate the formation of structure from an early time, when the matter field was nearly homogeneous, to late times, where it is highly clustered. As a first step, we generated a primordial power spectrum for a given choice of cosmological parameters[2] with the Boltzmann solver `CAMB` [254]. From the power spectrum we used `N-GenIC` [336] to generate the initial conditions for the particles in the simulation box, and finally `Gadget2` [337] to evolve the particles. We used a box with a comoving length of 240 Mpc/h on each side and 512^3 particles. The particles were initialized at redshift $z = 100$, evolved until $z = 0$, and snapshots were saved at 60 different redshifts between $z = 3$ and $z = 0$. For reasons described below, we ran four different N-body simulations, all with the same underlying cosmological parameters but with different seeds for the initial density and velocity perturbations.

To generate the convergence maps we used the software package `LensTools` [312], which implements a multi-lens-plane algorithm for ray-tracing. This algorithm approximates the three-dimensional distribution of matter $\delta(\boldsymbol{x}, z)$ (obtained from an N-body simulation) between the source redshift z_s and us as a series of discrete two-dimensional planes perpendicular to the line of sight, with thickness $\triangle$ and surface mass density σ:

$$\sigma(\boldsymbol{x}, z) = \frac{3H_0^2\Omega_m\chi(z)}{2c^2a(z)} \int_\triangle d\chi' \delta(\boldsymbol{x}, z(\chi')), \tag{G.5}$$

where χ is the comoving distance, $a = 1/(1+z)$ is the scale factor, Ω_m the matter density, H_0 the Hubble constant and c the speed of light. Since the surface density and the gravitational potential ϕ are related via the two-dimensional Poisson equation,

$$\nabla^2_{\boldsymbol{x}}\phi(\boldsymbol{x}, z) = 2\sigma(\boldsymbol{x}, z), \tag{G.6}$$

[2]We set $\Omega_m = 0.3$, $\Omega_\Lambda = 0.7$, $\Omega_\mathrm{b} = 0.046$, $\sigma_8 = 0.8$, $n_s = 1$ and $H_0 = 72$ km/s/Mpc.

the gravitational potential can be solved for. Then, the angular photon trajectory $\boldsymbol{\beta}$ can be calculated using the geodesic equation. Finally, the shear $\boldsymbol{\gamma}$ and the convergence κ can be obtained, since they are elements in the Jacobian of the angular trajectory of a photon as a function of its initial position $\boldsymbol{\theta}$:

$$\frac{\partial \boldsymbol{\beta}}{\partial \boldsymbol{\theta}} = \begin{pmatrix} 1 - \kappa + \boldsymbol{\gamma}_1 & -\boldsymbol{\gamma}_2 \\ -\boldsymbol{\gamma}_2 & 1 - \kappa - \boldsymbol{\gamma}_1 \end{pmatrix}. \tag{G.7}$$

To make the two-dimensional density planes we cut each snapshot at three points (55, 167, 278) Mpc and project slabs (along all three axes) that are 111 Mpc thick around each cut point. The planes are generated at a resolution of 4096×4096 pixels to make sure that small-scale information is preserved. To build the convergence maps we place a source at $z_{\rm s} = 1$ and 30 planes between $z_{\rm s}$ and $z = 0$, where a plane at a given redshift is randomly chosen from the planes made from each of the four different N-body simulations at that redshift. This is done to generate random realizations of the convergence field that are statistically independent. We created 1024×1024 pixel convergence maps that have a sky coverage of 3.5×3.5 deg. The left panel of Figure 9.2 shows an example of a simulated convergence map. We obtained the convergence power spectra through `LensTools` as well. We ultimately obtain 75,000 mock power spectra in 34 bins, uniformly distributed in log space for $\ell = [100, 5000]$, where the lower limit is set by the size of the map and the upper limit by the fact that at higher multipoles the numerical power spectra were shown to deviate from theory significantly [313]. The 68% CL and 95% CLs of our mock observable can be seen in the right panel of Figure 9.2.

G.6.2 Convergence Power Spectrum Distribution per Bin

In Section 9.3 we showed the absolute value of the t-statistic for each bin in the ensemble of 2,048 mock convergence power spectra. In Figure G.6.1 we show the actual distribution for each bin as blue histograms. We also show a Gaussian fit to the histogram in dashed red, and a KDE of the distribution in black. For many of the bins, the deviation away from zero skewness and kurtosis can be seen by eye by comparing the Gaussian fit to

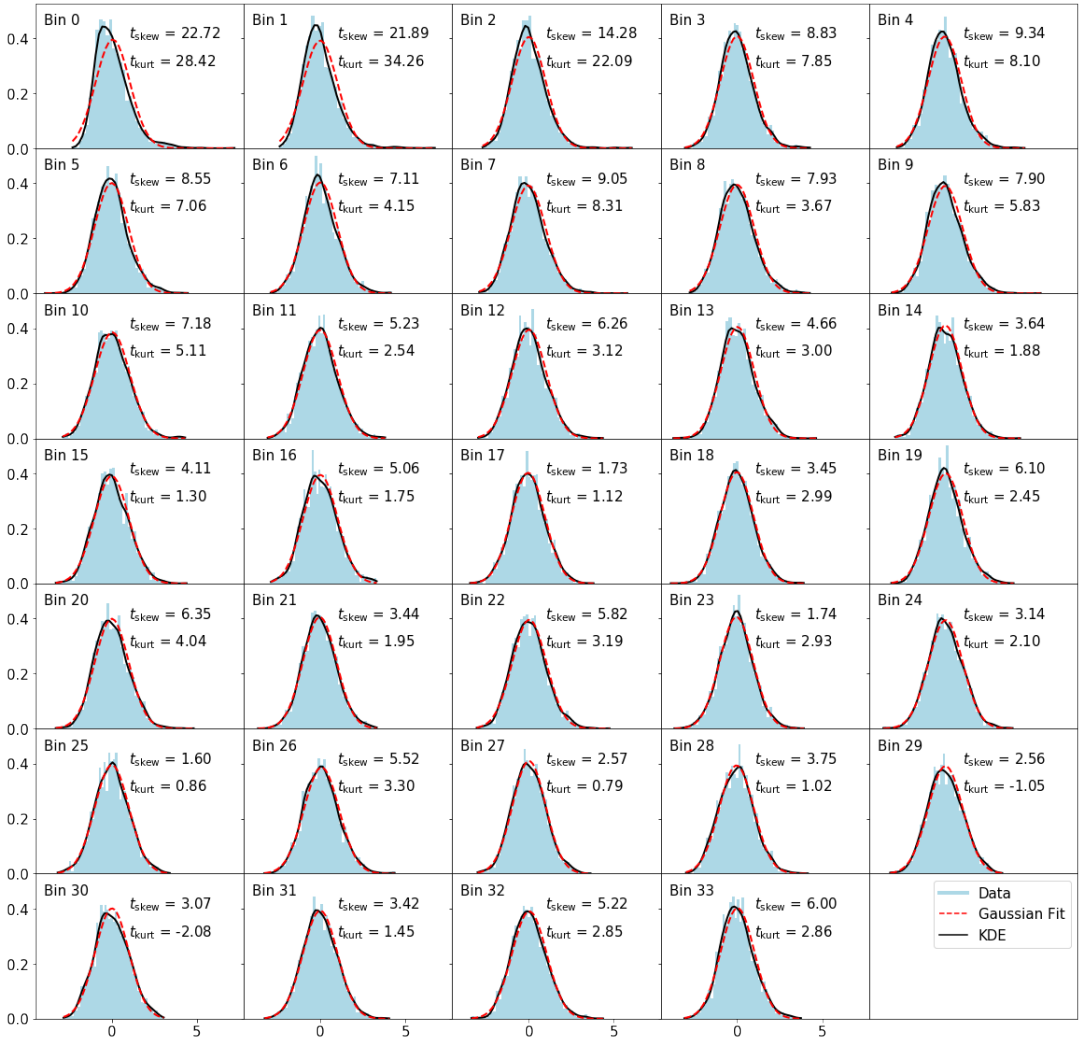

Figure G.6.1: Distribution for each of the 34 bins in an ensemble of 2,048 mock convergence power spectra. The histogrammed values are shown in light blue, while the KDE is shown in black. The dashed red line is a Gaussian fit to each bin.

either the KDE or the histogram.

References

[1] Ana Diaz Rivero and Cora Dvorkin. Direct Detection of Dark Matter Substructure in Strong Lens Images with Convolutional Neural Networks. *Phys. Rev. D*, 101(2):023515, 2020.

[2] Bryan Ostdiek, Ana Diaz Rivero, and Cora Dvorkin. Detecting subhalos in strong gravitational lens images with image segmentation, 2020.

[3] Bryan Ostdiek, Ana Diaz Rivero, and Cora Dvorkin. Extracting the Subhalo Mass Function from Strong Lens Images with Image Segmentation. 9 2020.

[4] Ana Diaz Rivero, Francis-Yan Cyr-Racine, and Cora Dvorkin. Power spectrum of dark matter substructure in strong gravitational lenses. *Phys. Rev.*, D97(2):023001, 2018.

[5] Ana Díaz Rivero, Cora Dvorkin, Francis-Yan Cyr-Racine, Jesús Zavala, and Mark Vogelsberger. Gravitational Lensing and the Power Spectrum of Dark Matter Substructure: Insights from the ETHOS N-body Simulations. *Phys. Rev.*, D98(10):103517, 2018.

[6] Atinc Çağan Şengül, Arthur Tsang, Ana Diaz Rivero, Cora Dvorkin, Hong-Ming Zhu, and Uroš Seljak. Quantifying the Line-of-Sight Halo Contribution to the Dark Matter Convergence Power Spectrum from Strong Gravitational Lenses. 6 2020.

[7] Ana Diaz Rivero, V. Miranda, and Cora Dvorkin. Observable Predictions for Massive-Neutrino Cosmologies with Model-Independent Dark Energy. *Phys. Rev. D*, 100(6):063504, 2019.

[8] Ana Diaz Rivero and Cora Dvorkin. Flow-Based Likelihoods for Non-Gaussian Inference. 7 2020.

[9] V. F. Mukhanov and G. V. Chibisov. Quantum fluctuations and a nonsingular universe. *Soviet Journal of Experimental and Theoretical Physics Letters*, 33:532, May 1981.

[10] D. N. Spergel, L. Verde, H. V. Peiris, E. Komatsu, M. R. Nolta, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters. ApJS, 148(1):175–194, September 2003.

[11] N. Aghanim et al. Planck 2018 results. VI. Cosmological parameters. 2018.

[12] Adam G. Riess, Alexei V. Filippenko, Peter Challis, Alejandro Clocchiatti, Alan Diercks, Peter M. Garnavich, Ron L. Gilliland, Craig J. Hogan, Saurabh Jha, Robert P. Kirshner, and et al. Observational evidence from supernovae for an accelerating universe and a cosmological constant. *The Astronomical Journal*, 116(3):1009–1038, Sep 1998.

[13] S. Perlmutter, G. Aldering, M. Della Valle, S. Deustua, R. S. Ellis, S. Fabbro, A. Fruchter, G. Goldhaber, D. E. Groom, I. M. Hook, and et al. Discovery of a supernova explosion at half the age of the universe. *Nature*, 391(6662):51–54, Jan 1998.

[14] D. M. Scolnic et al. The Complete Light-curve Sample of Spectroscopically Confirmed SNe Ia from Pan-STARRS1 and Cosmological Constraints from the Combined Pantheon Sample. *Astrophys. J.*, 859(2):101, 2018.

[15] Daniel J. Eisenstein, Idit Zehavi, David W. Hogg, Roman Scoccimarro, Michael R. Blanton, Robert C. Nichol, Ryan Scranton, Hee-Jong Seo, Max Tegmark, Zheng Zheng, and et al. Detection of the baryon acoustic peak in the large-scale correlation function of sdss luminous red galaxies. *The Astrophysical Journal*, 633(2):560–574, Nov 2005.

[16] S. Alam, M. Ata, S. Bailey, F. Beutler, D. Bizyaev, J. A. Blazek, A. S. Bolton, J. R. Brownstein, A. Burden, C.-H. Chuang, J. Comparat, A. J. Cuesta, K. S. Dawson, D. J. Eisenstein, S. Escoffier, H. Gil-Marín, J. N. Grieb, N. Hand, S. Ho, K. Kinemuchi, D. Kirkby, F. Kitaura, E. Malanushenko, V. Malanushenko, C. Maraston, C. K. McBride, R. C. Nichol, M. D. Olmstead, D. Oravetz, N. Padmanabhan, N. Palanque-Delabrouille, K. Pan, M. Pellejero-Ibanez, W. J. Percival, P. Petitjean, F. Prada, A. M. Price-Whelan, B. A. Reid, S. A. Rodríguez-Torres, N. A. Roe, A. J. Ross, N. P. Ross, G. Rossi, J. A. Rubiño-Martín, S. Saito, S. Salazar-Albornoz, L. Samushia, A. G. Sánchez, S. Satpathy, D. J. Schlegel, D. P. Schneider, C. G. Scóccola, H.-J. Seo, E. S. Sheldon, A. Simmons, A. Slosar, M. A. Strauss, M. E. C. Swanson, D. Thomas, J. L. Tinker, R. Tojeiro, M. V. Magaña, J. A. Vazquez, L. Verde, D. A. Wake, Y. Wang, D. H. Weinberg, M. White, W. M. Wood-Vasey, C. Yèche, I. Zehavi, Z. Zhai, and G.-B. Zhao. The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample. *Mon. Not. R. Astron. Soc.* , 470:2617–2652, September 2017.

[17] Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barand ela, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, R. Barrena, J. G. Bartlett, E. Battaner, R. Battye, K. Benabed, A. Benoît, A. Benoit-Lévy, J. P. Bernard, M. Bersanelli, P. Bielewicz, I. Bikmaev, A. Blanchard, J. Bobin, J. J. Bock, H. Böhringer, A. Bonaldi, J. R. Bond, J. Borrill, F. R. Bouchet, H. Bourdin, M. Bridges, M. L. Brown, M. Bucher, R. Burenin, C. Burigana, R. C. Butler, J. F. Cardoso, P. Carvalho, A. Catalano, A. Challinor, A. Chamballu, R. R. Chary, L. Y. Chiang, H. C. Chiang, G. Chon, P. R. Christensen, S. Church, D. L. Clements, S. Colombi, L. P. L. Colombo, F. Couchot, A. Coulais, B. P. Crill, A. Curto, F. Cuttaia, A. Da Silva, H. Dahle, L. Danese, R. D. Davies, R. J. Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, J. M. Delouis, J. Démoclès, F. X. Désert, C. Dickinson, J. M. Diego, K. Dolag,

H. Dole, S. Donzelli, O. Doré, M. Douspis, X. Dupac, G. Efstathiou, T. A. Enßlin, H. K. Eriksen, F. Finelli, I. Flores-Cacho, O. Forni, M. Frailis, E. Franceschi, S. Fromenteau, S. Galeotta, K. Ganga, R. T. Génova-Santos, M. Giard, G. Giardino, Y. Giraud-Héraud, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gregorio, A. Gruppuso, F. K. Hansen, D. Hanson, D. Harrison, S. Henrot-Versillé, C. Hernández-Monteagudo, D. Herranz, S. R. Hildebrand t, E. Hivon, M. Hobson, W. A. Holmes, A. Hornstrup, W. Hovest, K. M. Huffenberger, G. Hurier, T. R. Jaffe, A. H. Jaffe, W. C. Jones, M. Juvela, E. Keihänen, R. Keskitalo, I. Khamitov, T. S. Kisner, R. Kneissl, J. Knoche, L. Knox, M. Kunz, H. Kurki-Suonio, G. Lagache, A. Lähteenmäki, J. M. Lamarre, A. Lasenby, R. J. Laureijs, C. R. Lawrence, J. P. Leahy, R. Leonardi, J. León-Tavares, J. Lesgourgues, A. Liddle, M. Liguori, P. B. Lilje, M. Linden-Vørnle, M. López-Caniego, P. M. Lubin, J. F. Macías-Pérez, B. Maffei, D. Maino, N. Mand olesi, A. Marcos-Caballero, M. Maris, D. J. Marshall, P. G. Martin, E. Martínez-González, S. Masi, S. Matarrese, F. Matthai, P. Mazzotta, P. R. Meinhold, A. Melchiorri, J. B. Melin, L. Mendes, A. Mennella, M. Migliaccio, S. Mitra, M. A. Miville-Deschênes, A. Moneti, L. Montier, G. Morgante, D. Mortlock, A. Moss, D. Munshi, P. Naselsky, F. Nati, P. Natoli, C. B. Netterfield, H. U. Nørgaard-Nielsen, F. Noviello, D. Novikov, I. Novikov, S. Osborne, C. A. Oxborrow, F. Paci, L. Pagano, F. Pajot, D. Paoletti, B. Partridge, F. Pasian, G. Patanchon, O. Perdereau, L. Perotto, F. Perrotta, F. Piacentini, M. Piat, E. Pierpaoli, D. Pietrobon, S. Plaszczynski, E. Pointecouteau, G. Polenta, N. Ponthieu, L. Popa, T. Poutanen, G. W. Pratt, G. Prézeau, S. Prunet, J. L. Puget, J. P. Rachen, R. Rebolo, M. Reinecke, M. Remazeilles, C. Renault, S. Ricciardi, T. Riller, I. Ristorcelli, G. Rocha, M. Roman, C. Rosset, G. Roudier, M. Rowan-Robinson, J. A. Rubiño-Martín, B. Rusholme, M. Sandri, D. Santos, G. Savini, D. Scott, M. D. Seiffert, E. P. S. Shellard, L. D. Spencer, J. L. Starck, V. Stolyarov, R. Stompor, R. Sudiwala, R. Sunyaev, F. Sureau, D. Sutton, A. S. Suur-Uski, J. F. Sygnet, J. A. Tauber, D. Tavagnacco, L. Terenzi, L. Toffolatti, M. Tomasi, M. Tristram, M. Tucci, J. Tuovinen, M. Türler, G. Umana, L. Valenziano, J. Valiviita, B. Van Tent, P. Vielva, F. Villa, N. Vittorio, L. A. Wade, B. D. Wandelt, J. Weller, M. White, S. D. M. White, D. Yvon, A. Zacchei, and A. Zonca. Planck 2013 results. XX. Cosmology from Sunyaev-Zeldovich cluster counts. A&A, 571:A20, November 2014.

[18] Marian Douspis, Laura Salvati, and Nabila Aghanim. On the tension between large scale structures and cosmic microwave background, 2019.

[19] M. A. Troxel et al. Dark Energy Survey Year 1 results: Cosmological constraints from cosmic shear. *Phys. Rev.*, D98(4):043528, 2018.

[20] L. Verde, T. Treu, and A. G. Riess. Tensions between the Early and the Late Universe. 2019.

[21] Kenneth C Wong, Sherry H Suyu, Geoff C-F Chen, Cristian E Rusu, Martin Millon, Dominique Sluse, Vivien Bonvin, Christopher D Fassnacht, Stefan Taubenberger, Matthew W Auger, and et al. H0licow xiii. a 2.4% measurement of h0 from lensed quasars: 5.3 tension between early and late-universe probes. *Monthly Notices of the Royal Astronomical Society*, Jun 2020.

[22] Adam G. Riess et al. New Hubble Space Telescope Discoveries of Type Ia Su-

pernovae at z¿=1: Narrowing Constraints on the Early Behavior of Dark Energy. *Astrophys. J.*, 659:98–121, 2007.

[23] Pierre Astier et al. The Supernova legacy survey: Measurement of omega(m), omega(lambda) and W from the first year data set. *Astron. Astrophys.*, 447:31–48, 2006.

[24] E. Komatsu, K. M. Smith, J. Dunkley, C. L. Bennett, B. Gold, G. Hinshaw, N. Jarosik, D. Larson, M. R. Nolta, L. Page, D. N. Spergel, M. Halpern, R. S. Hill, A. Kogut, M. Limon, S. S. Meyer, N. Odegard, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation. ApJS, 192:18, February 2011.

[25] Daniel J. Eisenstein et al. Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies. *Astrophys. J.*, 633:560–574, 2005.

[26] Y. Akrami et al. Planck 2018 results. I. Overview and the cosmological legacy of Planck. 2018.

[27] M. Vogelsberger, S. Genel, V. Springel, P. Torrey, D. Sijacki, D. Xu, G. Snyder, D. Nelson, and L. Hernquist. Introducing the Illustris Project: simulating the coevolution of dark and visible matter in the Universe. *Mon. Not. R. Astron. Soc.* , 444:1518–1547, October 2014.

[28] M. Vogelsberger, S. Genel, V. Springel, P. Torrey, D. Sijacki, D. Xu, G. Snyder, S. Bird, D. Nelson, and L. Hernquist. Properties of galaxies reproduced by a hydrodynamic simulation. Nature, 509:177–182, May 2014.

[29] V. Springel, R. Pakmor, A. Pillepich, R. Weinberger, D. Nelson, L. Hernquist, M. Vogelsberger, S. Genel, P. Torrey, F. Marinacci, and J. Naiman. First results from the IllustrisTNG simulations: matter and galaxy clustering. *Mon. Not. R. Astron. Soc.* , 475:676–698, March 2018.

[30] R. G. Carlberg. Star Stream Folding by Dark Galactic Sub-Halos. *Astrophys. J.*, 705:L223–L226, 2009.

[31] Raymond G. Carlberg. Dark Matter Sub-Halo Counts via Star Stream Crossings. *Astrophys. J.*, 748:20, 2012.

[32] R. G. Carlberg. The Dynamics of Star Stream Gaps. *Astrophys. J.*, 775:90, 2013.

[33] W. H. W. Ngan and R. G. Carlberg. Using Gaps in N-body Tidal Streams to Probe Missing Satellites. ApJ, 788:181, June 2014.

[34] J. Bovy, D. Erkal, and J. L. Sanders. Linear perturbation theory for tidal streams and the small-scale CDM power spectrum. *Mon. Not. R. Astron. Soc.* , 466:628–668, April 2017.

[35] Denis Erkal, Vasily Belokurov, Jo Bovy, and Jason L. Sanders. The number and size of subhalo-induced gaps in stellar streams. *Mon. Not. R. Astron. Soc.* , 463(1):102–119, Nov 2016.

[36] J. L. Sanders, J. Bovy, and D. Erkal. Dynamics of stream-subhalo interactions. *Mon. Not. R. Astron. Soc.* , 457:3817–3835, April 2016.

[37] Denis Erkal and Vasily Belokurov. Forensics of subhalo–stream encounters: the three phases of gap growth. *Mon. Not. Roy. Astron. Soc.*, 450(1):1136–1149, 2015.

[38] Denis Erkal and Vasily Belokurov. Properties of Dark Subhaloes from Gaps in Tidal Streams. *Mon. Not. Roy. Astron. Soc.*, 454(4):3542–3558, 2015.

[39] Malte Buschmann, Joachim Kopp, Benjamin R. Safdi, and Chih-Liang Wu. Stellar Wakes from Dark Matter Subhalos. *Phys. Rev. Lett.*, 120(21):211101, 2018.

[40] Nilanjan Banik, Gianfranco Bertone, Jo Bovy, and Nassim Bozorgnia. Probing the nature of dark matter particles with stellar streams. 2018.

[41] Ken Van Tilburg, Anna-Maria Taki, and Neal Weiner. Halometry from Astrometry. 2018.

[42] J. R. Bond and A. S. Szalay. The Collisionless Damping of Density Fluctuations in an Expanding Universe. ApJ, 274:443–468, 1983.

[43] Julianne J. Dalcanton and Craig J. Hogan. Halo cores and phase space densities: Observational constraints on dark matter physics and structure formation. *Astrophys. J.*, 561:35–45, 2001.

[44] Paul Bode, Jeremiah P. Ostriker, and Neil Turok. Halo formation in warm dark matter models. *Astrophys. J.*, 556:93–107, 2001.

[45] D. Boyanovsky, H. J. de Vega, and N. Sanchez. The dark matter transfer function: free streaming, particle statistics and memory of gravitational clustering. *Phys. Rev. D*, 78:063546, 2008.

[46] Daniel Boyanovsky and Jun Wu. Small scale aspects of warm dark matter : power spectra and acoustic oscillations. *Phys. Rev. D*, 83:043524, 2011.

[47] E. D. Carlson, M. E. Machacek, and L. J. Hall. Self-interacting dark matter. ApJ, 398:43–52, October 1992.

[48] Celine Boehm, Alain Riazuelo, Steen H. Hansen, and Richard Schaeffer. Interacting dark matter disguised as warm dark matter. *Phys. Rev. D*, 66:083505, 2002.

[49] Lotty Ackerman, Matthew R. Buckley, Sean M. Carroll, and Marc Kamionkowski. Dark Matter and Dark Radiation. *Phys. Rev. D*, 79:023519, 2009.

[50] Jonathan L. Feng, Manoj Kaplinghat, Huitzu Tu, and Hai-Bo Yu. Hidden Charged Dark Matter. *J. Cosmol. Astropart. Phys.*, 07:004, 2009.

[51] David E. Kaplan, Gordan Z. Krnjaic, Keith R. Rehermann, and Christopher M. Wells. Atomic Dark Matter. JCAP, 05:021, 2010.

[52] Laura G. van den Aarssen, Torsten Bringmann, and Christoph Pfrommer. Is dark matter with long-range interactions a solution to all small-scale problems of Λ CDM cosmology? *Phys. Rev. Lett.*, 109:231301, 2012.

[53] Francis-Yan Cyr-Racine and Kris Sigurdson. The cosmology of atomic dark matter. *Phys. Rev. D*, 87:103515, 2013.

[54] Xiaoyong Chu and Basudeb Dasgupta. Dark Radiation Alleviates Problems with Dark Matter Halos. *Phys. Rev. Lett.*, 113(16):161301, 2014.

[55] Manuel A. Buen-Abad, Gustavo Marques-Tavares, and Martin Schmaltz. Non-Abelian dark matter and dark radiation. *Phys. Rev.*, D92(2):023531, 2015.

[56] Zackaria Chacko, Yanou Cui, Sungwoo Hong, Takemichi Okui, and Yuhsin Tsai. Partially Acoustic Dark Matter, Interacting Dark Radiation, and Large Scale Structure. *J. High Energy Phys.*, 12:108, 2016.

[57] P. Ko and Yong Tang. Light dark photon and fermionic dark radiation for the Hubble constant and the structure formation. *Phys. Lett.*, B762:462–466, 2016.

[58] Matthew R. Buckley, Jesus Zavala, Francis-Yan Cyr-Racine, Kris Sigurdson, and Mark Vogelsberger. Scattering, Damping, and Acoustic Oscillations: Simulating the Structure of Dark Matter Halos with Relativistic Force Carriers. *Phys. Rev. D*, 90(4):043524, 2014.

[59] J. A. Schewtschenko, R. J. Wilkinson, C. M. Baugh, C. Boehm, and S. Pascoli. Dark matter-radiation interactions: the impact on dark matter haloes. *Mon. Not. R. Astron. Soc.* , 449(4):3587–3596, 2015.

[60] Mark Vogelsberger, Jesus Zavala, Francis-Yan Cyr-Racine, Christoph Pfrommer, Torsten Bringmann, and Kris Sigurdson. ETHOS – an effective theory of structure formation: dark matter physics as a possible explanation of the small-scale CDM problems. *Mon. Not. Roy. Astron. Soc.*, 460(2):1399–1416, 2016.

[61] David N. Spergel and Paul J. Steinhardt. Observational evidence for selfinteracting cold dark matter. *Phys. Rev. Lett.*, 84:3760–3763, 2000.

[62] N. Yoshida, V. Springel, S. D. M. White, and G. Tormen. Weakly Self-interacting Dark Matter and the Structure of Dark Halos. ApJ, 544:L87–L90, December 2000.

[63] R. Davé, D. N. Spergel, P. J. Steinhardt, and B. D. Wandelt. Halo Properties in Cosmological Simulations of Self-interacting Cold Dark Matter. ApJ, 547:574–589, February 2001.

[64] P. Colín, V. Avila-Reese, O. Valenzuela, and C. Firmani. Structure and Subhalo Population of Halos in a Self-interacting Dark Matter Cosmology. ApJ, 581:777–793, December 2002.

[65] M. Vogelsberger, J. Zavala, and A. Loeb. Subhaloes in self-interacting galactic dark matter haloes. *Mon. Not. R. Astron. Soc.* , 423:3740–3752, July 2012.

[66] Miguel Rocha, Annika H. G. Peter, James S. Bullock, Manoj Kaplinghat, Shea Garrison-Kimmel, Jose Onorbe, and Leonidas A. Moustakas. Cosmological Simulations with Self-Interacting Dark Matter I: Constant Density Cores and Substructure. *Mon. Not. Roy. Astron. Soc.*, 430:81–104, 2013.

[67] Annika H. G. Peter, Miguel Rocha, James S. Bullock, and Manoj Kaplinghat. Cosmological Simulations with Self-Interacting Dark Matter II: Halo Shapes vs. Observations. *Mon. Not. R. Astron. Soc.*, 430:105, 2013.

[68] J. Zavala, M. Vogelsberger, and M. G. Walker. Constraining self-interacting dark matter with the Milky Way's dwarf spheroidals. *Mon. Not. R. Astron. Soc.* , 431:L20–L24, April 2013.

[69] Manoj Kaplinghat, Ryan E. Keeley, Tim Linden, and Hai-Bo Yu. Tying Dark Matter to Baryons with Self-interactions. *Phys. Rev. Lett.*, 113:021302, 2014.

[70] Manoj Kaplinghat, Sean Tulin, and Hai-Bo Yu. Dark Matter Halos as Particle Colliders: A Unified Solution to Small-Scale Structure Puzzles from Dwarfs to Clusters. *Phys. Rev. Lett.*, 116:041302, 2016.

[71] Julio F. Navarro, Carlos S. Frenk, and Simon D. M. White. A Universal density profile from hierarchical clustering. *Astrophys. J.*, 490:493–508, 1997.

[72] Wayne Hu, Rennan Barkana, and Andrei Gruzinov. Cold and fuzzy dark matter. *Phys. Rev. Lett.*, 85:1158–1161, 2000.

[73] Lam Hui, Jeremiah P. Ostriker, Scott Tremaine, and Edward Witten. Ultralight scalars as cosmological dark matter. *Phys. Rev.*, D95(4):043541, 2017.

[74] Xiaolong Du, Christoph Behrens, and Jens C. Niemeyer. Substructure of fuzzy dark matter haloes. *Mon. Not. R. Astron. Soc.*, 465(1):941–951, 2017.

[75] Philip Mocz, Mark Vogelsberger, Victor Robles, Jesus Zavala, Michael Boylan-Kolchin, and Lars Hernquist. Galaxy Formation with BECDM: I. Turbulence and relaxation of idealised haloes. *Mon. Not. R. Astron. Soc.*, 471:4559, 2017.

[76] A. M. Brooks and A. Zolotov. Why Baryons Matter: The Kinematics of Dwarf Spheroidal Satellites. ApJ, 786:87, May 2014.

[77] A. M. Brooks, M. Kuhlen, A. Zolotov, and D. Hooper. A Baryonic Solution to the Missing Satellites Problem. ApJ, 765:22, March 2013.

[78] K. S. Arraki, A. Klypin, S. More, and S. Trujillo-Gomez. Effects of baryon removal on the structure of dwarf spheroidal galaxies. *Mon. Not. R. Astron. Soc.* , 438:1466–1482, February 2014.

[79] Jose Oñorbe, Michael Boylan-Kolchin, James S. Bullock, Philip F. Hopkins, Dušan Kerěs, Claude-André Faucher-Giguère, Eliot Quataert, and Norman Murray. Forged in FIRE: cusps, cores, and baryons in low-mass dwarf galaxies. *Mon. Not. R. Astron. Soc.*, 454(2):2092–2106, 2015.

[80] A. R. Wetzel, P. F. Hopkins, J.-h. Kim, C.-A. Faucher-Giguère, D. Kereš, and E. Quataert. Reconciling Dwarf Galaxies with ΛCDM Cosmology: Simulating a Realistic Population of Satellites around a Milky Way-mass Galaxy. ApJ, 827:L23, August 2016.

[81] T. Sawala, C. S. Frenk, A. Fattahi, J. F. Navarro, R. G. Bower, R. A. Crain, C. Dalla

Vecchia, M. Furlong, J. C. Helly, A. Jenkins, K. A. Oman, M. Schaller, J. Schaye, T. Theuns, J. Trayford, and S. D. M. White. The APOSTLE simulations: solutions to the Local Group's cosmic puzzles. *Mon. Not. R. Astron. Soc.* , 457:1931–1943, April 2016.

[82] P. Creasey, O. Sameie, L. V. Sales, H.-B. Yu, M. Vogelsberger, and J. Zavala. Spreading out and staying sharp - creating diverse rotation curves via baryonic and self-interaction effects. *Mon. Not. R. Astron. Soc.* , 468:2283–2295, June 2017.

[83] T. Sawala, P. Pihajoki, P. H. Johansson, C. S. Frenk, J. F. Navarro, K. A. Oman, and S. D. M. White. Shaken and stirred: the Milky Way's dark substructures. *Mon. Not. R. Astron. Soc.* , 467:4383–4400, June 2017.

[84] S. Garrison-Kimmel, A. R. Wetzel, J. S. Bullock, P. F. Hopkins, M. Boylan-Kolchin, C.-A. Faucher-Giguere, D. Keres, E. Quataert, R. E. Sanderson, A. S. Graus, and T. Kelley. Not so lumpy after all: modeling the depletion of dark matter subhalos by Milky Way-like galaxies. *Mon. Not. R. Astron. Soc.*, 471:1709, January 2017.

[85] Mark Vogelsberger, Federico Marinacci, Paul Torrey, and Ewald Puchwein. Cosmological Simulations of Galaxy Formation. *Nature Rev. Phys.*, 2(1):42–66, 2020.

[86] Alex Fitts et al. FIRE in the Field: Simulating the Threshold of Galaxy Formation. *Mon. Not. R. Astron. Soc.*, 471:3547, 2017.

[87] J. I. Read, G. Iorio, O. Agertz, and F. Fraternali. The stellar mass-halo mass relation of isolated field dwarfs: a critical test of ΛCDM at the edge of galaxy formation. *Mon. Not. R. Astron. Soc.* , 467:2019–2038, May 2017.

[88] Gregory A. Dooley, Annika H. G. Peter, Tianyi Yang, Beth Willman, Brendan F. Griffen, and Anna Frebel. An observer's guide to the (Local Group) dwarf galaxies: predictions for their own dwarf satellite populations. *Mon. Not. Roy. Astron. Soc.*, 471(4):4894–4909, 2017.

[89] R. G. Carlberg. Modeling GD-1 Gaps in a Milky Way Potential. ApJ, 820(1):45, Mar 2016.

[90] J. Bovy. Detecting the Disruption of Dark-Matter Halos with Stellar Streams. *Phys. Rev. Lett.*, 116(12):121301, March 2016.

[91] Robert Feldmann and Douglas Spolyar. Detecting Dark Matter Substructures around the Milky Way with Gaia. *Mon. Not. Roy. Astron. Soc.*, 446:1000–1012, 2015.

[92] Shu-de Mao and Peter Schneider. Evidence for substructure in lens galaxies? *Mon. Not. Roy. Astron. Soc.*, 295:587–594, 1998.

[93] R. Benton Metcalf and Piero Madau. Compound Gravitational Lensing as a Probe of Dark Matter Substructure within Galaxy Halos. *Astrophys. J.*, 563(1):9–20, dec 2001.

[94] M. Chiba. Probing Dark Matter Substructure in Lens Galaxies. ApJ, 565:17–23, January 2002.

[95] R. B. Metcalf and H. Zhao. Flux Ratios as a Probe of Dark Substructures in Quadruple-Image Gravitational Lenses. ApJ, 567:L5–L8, March 2002.

[96] Neal Dalal and Christopher S. Kochanek. Strong lensing constraints on small-scale linear power. 2002.

[97] Charles R. Keeton, B. Scott Gaudi, and A. O. Petters. Identifying lenses with small-scale structure. i. cusp lenses. *Astrophys. J.*, 598:138–161, 2003.

[98] R. B. Metcalf and A. Amara. Small-scale structures of dark matter and flux anomalies in quasar gravitational lenses. *Mon. Not. R. Astron. Soc.* , 419:3414–3425, February 2012.

[99] N. Dalal and C. S. Kochanek. Direct detection of CDM substructure. *Astrophys. J.*, 572:25–33, 2002.

[100] R. Fadely and C. R. Keeton. Substructure in the lens HE 0435-1223. *Mon. Not. R. Astron. Soc.* , 419:936–951, January 2012.

[101] A. M. Nierenberg, T. Treu, S. A. Wright, C. D. Fassnacht, and M. W. Auger. Detection of substructure with adaptive optics integral field spectroscopy of the gravitational lens B1422+231. *Mon. Not. Roy. Astron. Soc.*, 442(3):2434–2445, 2014.

[102] A. M. Nierenberg, T. Treu, G. Brammer, A. H. G. Peter, C. D. Fassnacht, C. R. Keeton, C. S. Kochanek, K. B. Schmidt, D. Sluse, and S. A. Wright. Probing dark matter substructure in the gravitational lens HE0435-1223 with the WFC3 grism. *Mon. Not. R. Astron. Soc.*, 471(2224), 2017.

[103] L. V. E. Koopmans. Gravitational-mass imaging of CDM substructure. *Mon. Not. Roy. Astron. Soc.*, 363:1136, 2005.

[104] S. Vegetti and L. V. E. Koopmans. Bayesian Strong Gravitational-Lens Modelling on Adaptive Grids: Objective Detection of Mass Substructure in Galaxies. *Mon. Not. Roy. Astron. Soc.*, 392:945, 2009.

[105] S. Vegetti, O. Czoske, and L. V. E. Koopmans. Quantifying dwarf satellites through gravitational imaging: the case of SDSSJ120602.09+514229.5. *Mon. Not. R. Astron. Soc.* , 407:225–231, September 2010.

[106] S. Vegetti, L. V. E. Koopmans, A. Bolton, T. Treu, and R. Gavazzi. Detection of a dark substructure through gravitational imaging. *Mon. Not. R. Astron. Soc.* , 408(4):1969–1981, Nov 2010.

[107] S. Vegetti, D. J. Lagattuta, J. P. McKean, M. W. Auger, C. D. Fassnacht, and L. V. E. Koopmans. Gravitational detection of a low-mass dark satellite galaxy at cosmological distance. Nature, 481(7381):341–343, Jan 2012.

[108] S. Vegetti and L. V. E. Koopmans. Statistics of mass substructure from strong gravitational lensing: quantifying the mass fraction and mass function. *Mon. Not. R. Astron. Soc.* , 400:1583–1592, December 2009.

[109] S. Vegetti, L. V E Koopmans, M. W. Auger, T. Treu, and A. S. Bolton. Inference of the cold dark matter substructure mass function at z = 0.2 using strong gravitational lenses. *Mon. Not. R. Astron. Soc.*, 442(3):2017–2035, may 2014.

[110] Leonidas A. Moustakas and R. Benton Metcalf. Detecting dark matter substructure spectroscopically in strong gravitational lenses. *Mon. Not. Roy. Astron. Soc.*, 339:607, 2003.

[111] Yashar Hezaveh, Neal Dalal, Gilbert Holder, Michael Kuhlen, Daniel Marrone, Norman Murray, and Joaquin Vieira. Dark Matter Substructure Detection Using Spatially Resolved Spectroscopy of Lensed Dusty Galaxies. ApJ, 767(1):9, Apr 2013.

[112] Yashar D. Hezaveh et al. Detection of lensing substructure using ALMA observations of the dusty galaxy SDP.81. *Astrophys. J.*, 823(1):37, 2016.

[113] Francis Yan Cyr-Racine, Leonidas A. Moustakas, Charles Keeton, Kris Sigurdson, and Daniel A. Gilman. Dark census: Statistically detecting the satellite populations of distant galaxies. *Physical Review D*, 94(4), 8 2016.

[114] S. Birrer, A. Amara, and A. Refregier. Lensing substructure quantification in RXJ1131-1231: a 2 keV lower bound on dark matter thermal relic mass. JCAP, 05:037, May 2017.

[115] Brendon J. Brewer, David Huijser, and Geraint F. Lewis. Trans-Dimensional Bayesian Inference for Gravitational Lens Substructures. *Mon. Not. R. Astron. Soc.*, 455(2):1819–1829, aug 2015.

[116] Tansu Daylan, Francis-Yan Cyr-Racine, Ana Diaz Rivero, Cora Dvorkin, and Douglas P. Finkbeiner. Probing the small-scale structure in strongly lensed systems via transdimensional inference. *Astrophys. J.*, 854(2):141, 2018.

[117] Yashar Hezaveh, Neal Dalal, Gilbert Holder, Theodore Kisner, Michael Kuhlen, and Laurence Perreault Levasseur. Measuring the power spectrum of dark matter substructure using strong gravitational lensing. *JCAP*, 1611(11):048, 2016.

[118] S. Chatterjee and L. V. E. Koopmans. The inner mass power spectrum of galaxies using strong gravitational lensing: beyond linear approximation. *Mon. Not. R. Astron. Soc.* , 474:1762–1772, February 2018.

[119] E. Ritondale, S. Vegetti, G. Despali, M. W. Auger, L. V. E. Koopmans, and J. P. McKean. Low-mass halo perturbations in strong gravitational lenses at redshift z $\sim$ 0.5 are consistent with CDM. *Mon. Not. Roy. Astron. Soc.*, 485(2):2179–2193, 2019.

[120] Francis-Yan Cyr-Racine, Charles R. Keeton, and Leonidas A. Moustakas. Beyond subhalos: Probing the collective effect of the Universe's small-scale structure with gravitational lensing. *Phys. Rev.*, D100(2):023013, 2019.

[121] Daniel Gilman, Simon Birrer, and Tommaso Treu. Tdcosmo iii: Dark matter substructure meets dark energy – the effects of (sub)halos on strong-lensing measurements of h_0, 2020.

[122] Masamune Oguri and Philip J. Marshall. Gravitationally lensed quasars and supernovae in future wide-field optical imaging surveys. *Mon. Not. R. Astron. Soc.* , 405(4):2579–2593, Jul 2010.

[123] R. S. Pawase, F. Courbin, C. Faure, R. Kokotanekova, and G. Meylan. A 7 deg^2 survey for galaxy-scale gravitational lenses with the HST imaging archive. *Mon. Not. R. Astron. Soc.* , 439(4):3392–3404, Apr 2014.

[124] Thomas E. Collett. THE POPULATION OF GALAXY–GALAXY STRONG LENSES IN FORTHCOMING OPTICAL IMAGING SURVEYS. *The Astrophysical Journal*, 811(1):20, sep 2015.

[125] Yann Lecun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pages 2278–2324, 1998.

[126] Yann LeCun and Corinna Cortes. MNIST handwritten digit database. 2010.

[127] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.

[128] C. Schaefer, M. Geiger, T. Kuntzer, and J. P. Kneib. Deep convolutional neural networks as strong gravitational lens detectors. A&A, 611:A2, Mar 2018.

[129] Andrew Davies, Stephen Serjeant, and Jane M. Bromley. Using convolutional neural networks to identify gravitational lenses in astronomical images. *Mon. Not. R. Astron. Soc.* , 487(4):5263–5271, Aug 2019.

[130] Yashar D. Hezaveh, Laurence Perreault Levasseur, and Philip J. Marshall. Fast Automated Analysis of Strong Gravitational Lenses with Convolutional Neural Networks. *Nature*, 548:555–557, 2017.

[131] Laurence Perreault Levasseur, Yashar D. Hezaveh, and Risa H. Wechsler. Uncertainties in Parameters Estimated with Neural Networks: Application to Strong Gravitational Lensing. *Astrophys. J.*, 850(1):L7, 2017.

[132] Warren R. Morningstar, Yashar D. Hezaveh, Laurence Perreault Levasseur, Roger D. Blandford, Philip J. Marshall, Patrick Putzky, and Risa H. Wechsler. Analyzing interferometric observations of strong gravitational lenses with recurrent and convolutional neural networks. 2018.

[133] Simon Birrer and Adam Amara. Lenstronomy: multi-purpose gravitational lens modelling software package. 2018.

[134] R. Kormann, P. Schneider, and M. Bartelmann. Isothermal elliptical gravitational lens models. A&A, 284:285–299, April 1994.

[135] Piero Madau, Jürg Diemand, and Michael Kuhlen. Dark matter subhalos and the dwarf satellites of the milky way. *The Astrophysical Journal*, 679(2):1260–1271, jun 2008.

[136] D. D. Xu, Dominique Sluse, Liang Gao, Jie Wang, Carlos Frenk, Shude Mao,

and Peter Schneider. How well can cold-dark-matter substructures account for the observed lensing flux-ratio anomalies? 2013.

[137] Jurg Diemand, Michael Kuhlen, and Piero Madau. Dark matter substructure and gamma-ray annihilation in the Milky Way halo. *Astrophys. J.*, 657:262–270, 2007.

[138] Jurg Diemand, Michael Kuhlen, and Piero Madau. Formation and evolution of galaxy dark matter halos and their substructure. *Astrophys. J.*, 667:859–877, 2007.

[139] J. Diemand, M. Kuhlen, P. Madau, M. Zemp, B. Moore, D. Potter, and J. Stadel. Clumps and streams in the local dark matter distribution. *Nature*, 454:735–738, 2008.

[140] S. Vegetti, G. Despali, M. R. Lovell, and W. Enzi. Constraining sterile neutrino cosmologies with strong gravitational lensing observations at redshift z $\sim$ 0.2. *Mon. Not. Roy. Astron. Soc.*, 481(3):3661–3669, 2018.

[141] Julio F. Navarro, Carlos S. Frenk, and Simon D. M. White. The Structure of Cold Dark Matter Halos. ApJ, 462:563, May 1996.

[142] P. A. R. Ade et al. Planck 2015 results. XIII. Cosmological parameters. *Astron. Astrophys.*, 594:A13, 2016.

[143] Volker Springel, Jie Wang, Mark Vogelsberger, Aaron Ludlow, Adrian Jenkins, Amina Helmi, Julio F. Navarro, Carlos S. Frenk, and Simon D. M. White. The Aquarius Project: the subhalos of galactic halos. *Mon. Not. Roy. Astron. Soc.*, 391:1685–1711, 2008.

[144] A. S. Fruchter and R. N. Hook. Drizzle: A Method for the Linear Reconstruction of Undersampled Images. PASP, 114(792):144–152, Feb 2002.

[145] Diederik P Kingma and Max Welling. Auto-Encoding Variational Bayes. *arXiv e-prints*, page arXiv:1312.6114, Dec 2013.

[146] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017.

[147] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. volume 9 of *Proceedings of Machine Learning Research*, pages 249–256, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. JMLR Workshop and Conference Proceedings.

[148] Quinn E. Minor, Manoj Kaplinghat, and Nan Li. A Robust Mass Estimator for Dark Matter Subhalo Perturbations in Strong Gravitational Lenses. ApJ, 845(2):118, Aug 2017.

[149] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012.

[150] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going Deeper with Convolutions. *arXiv e-prints*, page arXiv:1409.4842, Sep 2014.

[151] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. *arXiv e-prints*, page arXiv:1512.03385, Dec 2015.

[152] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely Connected Convolutional Networks. *arXiv e-prints*, page arXiv:1608.06993, Aug 2016.

[153] Johann Brehmer, Siddharth Mishra-Sharma, Joeri Hermans, Gilles Louppe, and Kyle Cranmer. Mining for Dark Matter Substructure: Inferring subhalo population properties from strong lenses with machine learning. 2019.

[154] Stephon Alexander, Sergei Gleyzer, Evan McDonough, Michael W. Toomey, and Emanuele Usai. Deep Learning the Morphology of Dark Matter Substructure. 2019.

[155] Ran Li, Carlos S. Frenk, Shaun Cole, Qiao Wang, and Liang Gao. Projection effects in the strong lensing study of subhaloes. *Mon. Not. Roy. Astron. Soc.*, 468(2):1426–1432, 2017.

[156] Giulia Despali, Simona Vegetti, Simon D. M. White, Carlo Giocoli, and Frank C. van den Bosch. Modelling the line-of-sight contribution in substructure lensing. *Mon. Not. Roy. Astron. Soc.*, 475(4):5424–5442, 2018.

[157] Warren R. Morningstar, Laurence Perreault Levasseur, Yashar D. Hezaveh, Roger Blandford, Phil Marshall, Patrick Putzky, Thomas D. Rueter, Risa Wechsler, and Max Welling. Data-Driven Reconstruction of Gravitationally Lensed Galaxies using Recurrent Inference Machines. 2019.

[158] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. *arXiv e-prints*, page arXiv:1505.04597, May 2015.

[159] P. A. R. Ade et al. Planck 2015 results. XIII. Cosmological parameters. *Astron. Astrophys.*, 594:A13, 2016.

[160] Edward A. Baltz, Phil Marshall, and Masamune Oguri. Analytic models of plausible gravitational lens potentials. JCAP, 01:015, 2009.

[161] Zhong-Qiu Zhao, Peng Zheng, Shou-tao Xu, and Xindong Wu. Object Detection with Deep Learning: A Review. *arXiv e-prints*, page arXiv:1807.05511, July 2018.

[162] Shervin Minaee, Yuri Boykov, Fatih Porikli, Antonio Plaza, Nasser Kehtarnavaz, and Demetri Terzopoulos. Image Segmentation Using Deep Learning: A Survey. *arXiv e-prints*, page arXiv:2001.05566, January 2020.

[163] Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *arXiv e-prints*, page arXiv:1502.03167, February 2015.

[164] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv e-prints*, page arXiv:1412.6980, December 2014.

[165] J.A. Nelder and R. Mead. A Simplex Method for Function Minimization. *Comput. J.*, 7:308–313, 1965.

[166] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, CJ Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake Vand erPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1. 0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.

[167] LSST Science Collaboration, P. A. Abell, J. Allison, S. F. Anderson, J. R. Andrew, J. R. P. Angel, L. Armus, D. Arnett, S. J. Asztalos, T. S. Axelrod, and et al. LSST Science Book, Version 2.0. *ArXiv e-prints*, December 2009.

[168] A. Refregier, A. Amara, T. D. Kitching, A. Rassat, R. Scaramella, and J. Weller. Euclid Imaging Consortium Science Book. *arXiv e-prints*, page arXiv:1001.0061, January 2010.

[169] Curtis McCully, Charles R. Keeton, Kenneth C. Wong, and Ann I. Zabludoff. Quantifying Environmental and Line-of-Sight Effects in Models of Strong Gravitational Lens Systems. *Astrophys. J.*, 836(1):141, 2017.

[170] Asantha Cooray and Ravi K. Sheth. Halo models of large scale structure. *Phys. Rep.*, 372:1–129, 2002.

[171] Gerard Gilmore, Mark I. Wilkinson, Rosemary F. G. Wyse, Jan T. Kleyna, Andreas Koch, N. Wyn Evans, and Eva K. Grebel. The Observed properties of Dark Matter on small spatial scales. *Astrophys. J.*, 663:948–959, 2007.

[172] M. G. Walker and J. Peñarrubia. A Method for Measuring (Slopes of) the Mass Profiles of Dwarf Spheroidal Galaxies. ApJ, 742:20, November 2011.

[173] Nicola Cristiano Amorisco, Jesus Zavala, and Thomas J. L. de Boer. Dark matter cores in the Fornax and Sculptor dwarf galaxies: joining halo assembly and detailed star formation histories. *Astrophys. J.*, 782:L39, 2014.

[174] Davide Fiacconi, Piero Madau, Doug Potter, and Joachim Stadel. Cold Dark Matter Substructures in Early-Type Galaxy Halos. *Astrophys. J.*, 824:144, 2016.

[175] Ravi K. Sheth and Bhuvnesh Jain. Substructure and the halo model of large scale structure. *Mon. Not. R. Astron. Soc.*, 345:529, 2003.

[176] S. Refsdal and R. Stabell. Gravitational micro-lensing for large sources. A&A, 250:62–66, October 1991.

[177] B. Neindorf. A probability theoretical access to extragalactic microlensing. A&A,

404:83–92, June 2003.

[178] C. R. Keeton. Gravitational lensing with stochastic substructure: Effects of the clump mass function and spatial distribution. *ArXiv e-prints*, August 2009.

[179] Simona Vegetti and Mark Vogelsberger. On the density profile of dark matter substructure in gravitational lens galaxies. *Mon. Not. Roy. Astron. Soc.*, 442(4):3598–3603, 2014.

[180] A. R. Duffy, J. Schaye, S. T. Kay, and C. Dalla Vecchia. Dark matter halo concentrations in the Wilkinson Microwave Anisotropy Probe year 5 cosmology. *Mon. Not. R. Astron. Soc.* , 390:L64–L68, October 2008.

[181] Robert T. Chamberlain, Neal Dalal, Andrew Hearin, and Paul Ricker. Probing Satellite Quenching With Galaxy Clustering. *Mon. Not. R. Astron. Soc.*, 451(2):1496–1501, 2015.

[182] Yuedong Fang, Joseph Clampitt, Neal Dalal, Bhuvnesh Jain, Eduardo Rozo, John Moustakas, and Eli Rykoff. Tidal stripping as a test of satellite quenching in redMaPPer clusters. *Mon. Not. R. Astron. Soc.*, 463(2):1907–1915, 2016.

[183] A. Burkert. The Structure of dark matter halos in dwarf galaxies. *Astrophys. J.*, 447:L25, 1995. [Astrophys. J.447,L25(1995)].

[184] Charles R. Keeton. A Catalog of Mass Models for Gravitational Lensing. 2001.

[185] N. Murray. Star Formation Efficiencies and Lifetimes of Giant Molecular Clouds in the Milky Way. ApJ, 729:133, March 2011.

[186] R. G. Gratton, E. Carretta, and A. Bragaglia. Multiple populations in globular clusters. Lessons learned from the Milky Way globular clusters. A&A Rev., 20:50, February 2012.

[187] Qiuhan He, Ran Li, Sungsoon Lim, Carlos S. Frenk, Shaun Cole, Eric W. Peng, and Qiao Wang. Globular clusters vs dark matter haloes in strong lensing observations. 2017.

[188] D. D. Xu, S. Mao, J. Wang, V. Springel, L. Gao, S. D. M. White, C. S. Frenk, A. Jenkins, G. Li, and J. F. Navarro. Effects of Dark Matter Substructures on Gravitational Lensing: Results from the Aquarius Simulations. *Mon. Not. R. Astron. Soc.*, 398:1235, 2009.

[189] D. D. Xu, Shude Mao, Andrew Cooper, Liang Gao, Carlos Frenk, Raul Angulo, and John Helly. On the Effects of Line-of-Sight Structures on Lensing Flux-ratio Anomalies in a LCDM Universe. *Mon. Not. R. Astron. Soc.*, 421:2553, 2012.

[190] Yao-Yuan Mao, Marc Williamson, and Risa H. Wechsler. The Dependence of Subhalo Abundance on Halo Concentration. *Astrophys. J.*, 810(1):21, 2015.

[191] A. Sonnenfeld, R. Gavazzi, S. H. Suyu, T. Treu, and P. J. Marshall. The SL2S Galaxy-scale Lens Sample. III. Lens Models, Surface Photometry, and Stellar Masses for the Final Sample. ApJ, 777:97, November 2013.

[192] D. Xu, V. Springel, D. Sluse, P. Schneider, A. Sonnenfeld, D. Nelson, M. Vogelsberger, and L. Hernquist. The inner structure of early-type galaxies in the Illustris simulation. *Mon. Not. R. Astron. Soc.*, 469:1824–1848, August 2017.

[193] Florian Beutler, Hee-Jong Seo, Shun Saito, Chia-Hsun Chuang, Antonio J. Cuesta, Daniel J. Eisenstein, Hector Gil-Marin, Jan Niklas Grieb, Nick Hand, Francisco-Shu Kitaura, Chirag Modi, Robert C. Nichol, Matthew D. Olmstead, Will J. Percival, Francisco Prada, Ariel G. Sanchez, Sergio Rodriguez-Torres, Ashley J. Ross, Nicholas P. Ross, Donald P. Schneider, Jeremy Tinker, Rita Tojeiro, and Mariana Vargas-Magana. The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: anisotropic galaxy clustering in Fourier space. *Monthly Notices of the Royal Astronomical Society*, 466(2):2242–2260, 12 2016.

[194] Charles R. Keeton. Analytic cross sections for substructure lensing. *Astrophys. J.*, 584:664–674, 2003.

[195] Simon Birrer, Cyril Welschen, Adam Amara, and Alexandre Refregier. Line-of-sight effects in strong lensing: Putting theory into practice. *JCAP*, 1704(04):049, 2017.

[196] Francis-Yan Cyr-Racine, Kris Sigurdson, Jesus Zavala, Torsten Bringmann, Mark Vogelsberger, and Christoph Pfrommer. ETHOS—an effective theory of structure formation: From dark particle physics to the matter distribution of the Universe. *Phys. Rev.*, D93(12):123527, 2016.

[197] M. R. Lovell, J. Zavala, M. Vogelsberger, X. Shen, F.-Y. Cyr-Racine, C. Pfrommer, K. Sigurdson, M. Boylan-Kolchin, and A. Pillepich. ETHOS - an effective theory of structure formation: predictions for the high-redshift Universe - abundance of galaxies and reionization. *Mon. Not. R. Astron. Soc.* , 477:2886–2899, July 2018.

[198] Volker Springel, Simon D. M. White, Giuseppe Tormen, and Guinevere Kauffmann. Populating a cluster of galaxies. 1. Results at z = 0. *Mon. Not. Roy. Astron. Soc.*, 328:726, 2001.

[199] S. Vegetti, L. V. E. Koopmans, A. Bolton, T. Treu, and R. Gavazzi. Detection of a dark substructure through gravitational imaging. *Mon. Not. R. Astron. Soc.* , 408:1969–1981, November 2010.

[200] Sean Brennan, Andrew J. Benson, Francis-Yan Cyr-Racine, Charles R. Keeton, Leonidas A. Moustakas, and Anthony R. Pullen. Quantifying the power spectrum of small-scale structure in semi-analytic galaxies. *Mon. Not. Roy. Astron. Soc.*, 488(4):5085–5092, 2019.

[201] S. Garrison-Kimmel, M. Boylan-Kolchin, J. S. Bullock, and E. N. Kirby. Too big to fail in the Local Group. *Mon. Not. R. Astron. Soc.* , 444:222–236, October 2014.

[202] Nick Hand, Yu Feng, Florian Beutler, Yin Li, Chirag Modi, Uroš Seljak, and Zachary Slepian. nbodykit: An Open-source, Massively Parallel Toolkit for Large-scale Structure. AJ, 156(4):160, October 2018.

[203] Mark Vogelsberger, Jesús Zavala, Katelin Schutz, and Tracy R. Slatyer. Evaporat-

ing the Milky Way halo and its satellites with inelastic self-interacting dark matter. 2018.

[204] Adam S. Bolton, Scott Burles, Léon V. E. Koopmans, Tommaso Treu, Raphaël Gavazzi, Leonidas A. Moustakas, Randall Wayth, and David J. Schlegel. The Sloan Lens ACS Survey. V. The Full ACS Strong Lens Sample. *The Astrophysical Journal*, 682(2):964–984, may 2008.

[205] Anson D'Aloisio and Priyamvada Natarajan. Cosmography with cluster strong lenses: the influence of substructure and line-of-sight haloes. *Mon. Not. R. Astron. Soc.* , 411(3):1628–1640, March 2011.

[206] Daniel Gilman, Simon Birrer, Tommaso Treu, Anna Nierenberg, and Andrew Benson. Probing dark matter structure down to 10^7 solar masses: flux ratio statistics in gravitational lenses with line-of-sight haloes. *Mon. Not. Roy. Astron. Soc.*, 487(4):5721–5738, 2019.

[207] Daniel Gilman, Simon Birrer, Anna Nierenberg, Tommaso Treu, Xiaolong Du, and Andrew Benson. Warm dark matter chills out: constraints on the halo mass function and the free-streaming length of dark matter with 8 quadruple-image strong gravitational lenses. *Mon. Not. Roy. Astron. Soc.*, 491(4):6077–6101, 2020.

[208] Daniel Gilman, Xiaolong Du, Andrew Benson, Simon Birrer, Anna Nierenberg, and Tommaso Treu. Constraints on the mass-concentration relation of cold dark matter halos with 11 strong gravitational lenses. *Mon. Not. Roy. Astron. Soc.*, 492(1):L12–L16, 2020.

[209] Jen-Wei Hsueh, Wolfgang Enzi, Simona Vegetti, Matt Auger, Christopher D. Fassnacht, Giulia Despali, Leon V. E. Koopmans, and John P. McKean. SHARP – VII. New constraints on the dark matter free-streaming properties and substructure abundance from gravitationally lensed quasars. *Mon. Not. Roy. Astron. Soc.*, 492(2):3047–3059, 2020.

[210] Cristian E. Rusu et al. H0LiCOW XII. Lens mass model of WFI2033-4723 and blind measurement of its time-delay distance and H_0. 2019.

[211] Geoff C. F. Chen et al. A SHARP view of H0LiCOW: H_0 from three time-delay gravitational lens systems with adaptive optics imaging. *Mon. Not. Roy. Astron. Soc.*, 490(2):1743–1773, 2019.

[212] Peter Schneider, Jürgen Ehlers, and Emilio E. Falco. *Gravitational Lenses*. 1992.

[213] Simon Birrer, Cyril Welschen, Adam Amara, and Alexandre Refregier. Line-of-sight effects in strong lensing: Putting theory into practice. *JCAP*, 1704(04):049, 2017.

[214] D. N. Limber. The Analysis of Counts of the Extragalactic Nebulae in Terms of a Fluctuating Density Field. ApJ, 117:134, January 1953.

[215] Scott Dodelson. *Modern Cosmology*. Academic Press, Elsevier Science, 2003.

[216] Ravi K. Sheth, H. J. Mo, and Giuseppe Tormen. Ellipsoidal collapse and an im-

proved model for the number and spatial distribution of dark matter haloes. 1999.

[217] Darren S. Reed, Richard Bower, Carlos S. Frenk, Adrian Jenkins, and Tom Theuns. The halo mass function from the dark ages through the present day. *Mon. Not. R. Astron. Soc.* , 374(1):2–15, January 2007.

[218] Carlo Giocoli, Giuseppe Tormen, Ravi K. Sheth, and Frank C. van den Bosch. The substructure hierarchy in dark matter haloes. *Mon. Not. R. Astron. Soc.* , 404(1):502–517, May 2010.

[219] Nobuhiro Okabe, Toshifumi Futamase, Masaru Kajisawa, and Risa Kuroshima. Subaru Weak-lensing Survey of Dark Matter Subhalos in the Coma Cluster: Subhalo Mass Function and Statistical Properties. ApJ, 784(2):90, April 2014.

[220] G. De Lucia, G. Kauffmann, V. Springel, S. D. M. White, B. Lanzoni, F. Stoehr, G. Tormen, and N. Yoshida. Substructures in cold dark matter haloes. *Mon. Not. R. Astron. Soc.* , 348(1):333–344, February 2004.

[221] Michael Boylan-Kolchin, Volker Springel, Simon D. M. White, and Adrian Jenkins. There's no place like home? Statistics of Milky Way-mass dark matter haloes. *Mon. Not. R. Astron. Soc.* , 406(2):896–912, August 2010.

[222] L. Gao, J. F. Navarro, C. S. Frenk, A. Jenkins, V. Springel, and S. D. M. White. The Phoenix Project: the dark side of rich Galaxy clusters. *Mon. Not. R. Astron. Soc.* , 425(3):2169–2186, September 2012.

[223] Hao-Yi Wu, Oliver Hahn, Risa H. Wechsler, Peter S. Behroozi, and Yao-Yuan Mao. Rhapsody. II. Subhalo Properties and the Impact of Tidal Stripping From a Statistical Sample of Cluster-size Halos. ApJ, 767(1):23, April 2013.

[224] M. W. Auger, T. Treu, A. S. Bolton, R. Gavazzi, L. V. E. Koopmans, P. J. Marshall, K. Bundy, and L. A. Moustakas. The Sloan Lens ACS Survey. IX. Colors, Lensing, and Stellar Masses of Early-Type Galaxies. ApJ, 705(2):1099–1115, November 2009.

[225] William H. Press and Paul Schechter. Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation. ApJ, 187:425–438, February 1974.

[226] Michael S. Warren, Kevork Abazajian, Daniel E. Holz, and Luís Teodoro. Precision Determination of the Mass Function of Dark Matter Halos. ApJ, 646(2):881–885, August 2006.

[227] Jeremy Tinker, Andrey V. Kravtsov, Anatoly Klypin, Kevork Abazajian, Michael Warren, Gustavo Yepes, Stefan Gottlöber, and Daniel E. Holz. Toward a Halo Mass Function for Precision Cosmology: The Limits of Universality. ApJ, 688(2):709–728, December 2008.

[228] Tiago Castro, Valerio Marra, and Miguel Quartin. Constraining the halo mass function with observations. *Mon. Not. R. Astron. Soc.* , 463(2):1666–1677, December 2016.

[229] Joel R. Brownstein, Adam S. Bolton, David J. Schlegel, Daniel J. Eisenstein, Christopher S. Kochanek, Natalia Connolly, Claudia Maraston, Parul Pandey,

Stella Seitz, David A. Wake, W. Michael Wood-Vasey, Jon Brinkmann, Donald P. Schneider, and Benjamin A. Weaver. The BOSS Emission-Line Lens Survey (BELLS). I. A Large Spectroscopically Selected Sample of Lens Galaxies at Redshift ~0.5. ApJ, 744(1):41, January 2012.

[230] M. W. Auger, T. Treu, A. S. Bolton, R. Gavazzi, L. V. E. Koopmans, P. J. Marshall, L. A. Moustakas, and S. Burles. The Sloan Lens ACS Survey. X. Stellar, Dynamical, and Total Mass Correlations of Massive Early-type Galaxies. ApJ, 724(1):511–525, November 2010.

[231] Andrew Robertson, Graham P. Smith, Richard Massey, Vincent Eke, Mathilde Jauzac, Matteo Bianconi, and Dan Ryczanowski. What does strong gravitational lensing? The mass and redshift distribution of high-magnification lenses. *arXiv e-prints*, page arXiv:2002.01479, February 2020.

[232] D. Bayer, S. Chatterjee, L. V. E. Koopmans, S. Vegetti, J. P. McKean, T. Treu, and C. D. Fassnacht. Observational constraints on the sub-galactic matter-power spectrum from galaxy-galaxy strong gravitational lensing. 2018.

[233] Sreedevi Varma, Malcolm Fairbairn, and Julio Figueroa. Dark Matter Subhalos, Strong Lensing and Machine Learning. 5 2020.

[234] S. Perlmutter et al. Measurements of Omega and Lambda from 42 high redshift supernovae. *Astrophys. J.*, 517:565–586, 1999.

[235] S. Alam, M. Ata, S. Bailey, F. Beutler, D. Bizyaev, J. A. Blazek, A. S. Bolton, J. R. Brownstein, A. Burden, C.-H. Chuang, J. Comparat, A. J. Cuesta, K. S. Dawson, D. J. Eisenstein, S. Escoffier, H. Gil-Marín, J. N. Grieb, N. Hand, S. Ho, K. Kinemuchi, D. Kirkby, F. Kitaura, E. Malanushenko, V. Malanushenko, C. Maraston, C. K. McBride, R. C. Nichol, M. D. Olmstead, D. Oravetz, N. Padmanabhan, N. Palanque-Delabrouille, K. Pan, M. Pellejero-Ibanez, W. J. Percival, P. Petitjean, F. Prada, A. M. Price-Whelan, B. A. Reid, S. A. Rodríguez-Torres, N. A. Roe, A. J. Ross, N. P. Ross, G. Rossi, J. A. Rubiño-Martín, S. Saito, S. Salazar-Albornoz, L. Samushia, A. G. Sánchez, S. Satpathy, D. J. Schlegel, D. P. Schneider, C. G. Scóccola, H.-J. Seo, E. S. Sheldon, A. Simmons, A. Slosar, M. A. Strauss, M. E. C. Swanson, D. Thomas, J. L. Tinker, R. Tojeiro, M. V. Magaña, J. A. Vazquez, L. Verde, D. A. Wake, Y. Wang, D. H. Weinberg, M. White, W. M. Wood-Vasey, C. Yèche, I. Zehavi, Z. Zhai, and G.-B. Zhao. The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample. *Mon. Not. R. Astron. Soc.* , 470:2617–2652, September 2017.

[236] Chiaki Hikage et al. Cosmology from cosmic shear power spectra with Subaru Hyper Suprime-Cam first-year data. *Submitted to: Publ. Astron. Soc. Jap.*, 2018.

[237] P. F. de Salas, D. V. Forero, C. A. Ternes, M. Tortola, and J. W. F. Valle. Status of neutrino oscillations 2018: 3σ hint for normal mass ordering and improved CP sensitivity. *Phys. Lett.*, B782:633–640, 2018.

[238] Michael J. Mortonson, Wayne Hu, and Dragan Huterer. Falsifying Paradigms for Cosmic Acceleration. *Phys. Rev.*, D79:023004, 2009.

[239] Eduardo J. Ruiz and Dragan Huterer. Testing the dark energy consistency with geometry and growth. *Phys. Rev.*, D91:063009, 2015.

[240] V. Miranda and Cora Dvorkin. Model-Independent Predictions for Smooth Cosmic Acceleration Scenarios. *Phys. Rev.*, D98(4):043537, 2018.

[241] Sunny Vagnozzi, Suhail Dhawan, Martina Gerbino, Katherine Freese, Ariel Goobar, and Olga Mena. Constraints on the sum of the neutrino masses in dynamical dark energy models with $w(z) \geq -1$ are tighter than those obtained in ΛCDM. 2018.

[242] Shouvik Roy Choudhury and Abhishek Naskar. Strong Bounds on Sum of Neutrino Masses in a 12 Parameter Extended Scenario with Non-Phantom Dynamical Dark Energy ($w(z) \geq -1$) with CPL parameterization. 2018.

[243] Shouvik Roy Choudhury and Sandhya Choubey. Updated Bounds on Sum of Neutrino Masses in Various Cosmological Scenarios. *JCAP*, 1809(09):017, 2018.

[244] H. Gil-Marín, W. J. Percival, A. J. Cuesta, J. R. Brownstein, C.-H. Chuang, S. Ho, F.-S. Kitaura, C. Maraston, F. Prada, S. Rodríguez-Torres, A. J. Ross, D. J. Schlegel, D. P. Schneider, D. Thomas, J. L. Tinker, R. Tojeiro, M. Vargas Magaña, and G.-B. Zhao. The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: BAO measurement from the LOS-dependent power spectrum of DR12 BOSS galaxies. *Mon. Not. R. Astron. Soc.* , 460:4210–4219, August 2016.

[245] F. Beutler, C. Blake, M. Colless, D. H. Jones, L. Staveley-Smith, L. Campbell, Q. Parker, W. Saunders, and F. Watson. The 6dF Galaxy Survey: baryon acoustic oscillations and the local Hubble constant. *Mon. Not. R. Astron. Soc.* , 416:3017–3032, October 2011.

[246]

[247] N. Aghanim et al. Planck 2015 results. XI. CMB power spectra, likelihoods, and robustness of parameters. 2015.

[248] P.A.R. Ade et al. Planck 2015 results. XV. Gravitational lensing. 2015.

[249] T. M. C. Abbott et al. Dark Energy Survey year 1 results: Cosmological constraints from galaxy clustering and weak lensing. *Phys. Rev.*, D98(4):043526, 2018.

[250] E. Krause et al. Dark Energy Survey Year 1 Results: Multi-Probe Methodology and Simulated Likelihood Analyses. *Submitted to: Phys. Rev. D*, 2017.

[251] Ryuichi Takahashi, Masanori Sato, Takahiro Nishimichi, Atsushi Taruya, and Masamune Oguri. Revising the Halofit Model for the Nonlinear Matter Power Spectrum. *Astrophys. J.*, 761:152, 2012.

[252] L. Casarini, S. A. Bonometto, E. Tessarotto, and P. S. Corasaniti. Extending the Coyote emulator to dark energy models with standard w_0-w_a parametrization of the equation of state. *JCAP*, 1608(08):008, 2016.

[253] Antony Lewis. Efficient sampling of fast and slow cosmological parameters. *Phys.*

Rev., D87:103529, 2013.

[254] Antony Lewis and Sarah Bridle. Cosmological parameters from CMB and other data: A Monte Carlo approach. Phys. Rev. D, 66:103511, 2002.

[255] Antony Lewis, Anthony Challinor, and Anthony Lasenby. Efficient computation of CMB anisotropies in closed FRW models. *Astrophys. J.*, 538:473–476, 2000.

[256] Cullan Howlett, Antony Lewis, Alex Hall, and Anthony Challinor. CMB power spectrum parameter degeneracies in the era of precision cosmology. *JCAP*, 1204:027, 2012.

[257] R. Hounsell, D. Scolnic, R. J. Foley, R. Kessler, V. Miranda, A. Avelino, R. C. Bohlin, A. V. Filippenko, J. Frieman, S. W. Jha, P. L. Kelly, R. P. Kirshner, K. Mandel, A. Rest, A. G. Riess, S. A. Rodney, and L. Strolger. Simulations of the WFIRST Supernova Survey and Forecasts of Cosmological Constraints. ApJ, 867:23, Nov 2018.

[258] Roland de Putter and Eric V. Linder. Being PC: Principal Components and Dark Energy. 2008.

[259] Michael Mortonson, Wayne Hu, and Dragan Huterer. Hiding dark energy transitions at low redshift. Phys. Rev. D, 80:067301, September 2009.

[260] Gong-Bo Zhao et al. Dynamical dark energy in light of the latest observations. *Nat. Astron.*, 1(9):627–632, 2017.

[261] D. Spergel, N. Gehrels, J. Breckinridge, M. Donahue, A. Dressler, B. S. Gaudi, T. Greene, O. Guyon, C. Hirata, J. Kalirai, N. J. Kasdin, W. Moos, S. Perlmutter, M. Postman, B. Rauscher, J. Rhodes, Y. Wang, D. Weinberg, J. Centrella, W. Traub, C. Baltay, J. Colbert, D. Bennett, A. Kiessling, B. Macintosh, J. Merten, M. Mortonson, M. Penny, E. Rozo, D. Savransky, K. Stapelfeldt, Y. Zu, C. Baker, E. Cheng, D. Content, J. Dooley, M. Foote, R. Goullioud, K. Grady, C. Jackson, J. Kruk, M. Levine, M. Melton, C. Peddie, J. Ruffa, and S. Shaklan. WFIRST-2.4: What Every Astronomer Should Know. *ArXiv e-prints*, May 2013.

[262] Emmanuel Schaan, Elisabeth Krause, Tim Eifler, Olivier Doré, Hironao Miyatake, Jason Rhodes, and David N. Spergel. Looking through the same lens: Shear calibration for LSST, Euclid, and WFIRST with stage 4 CMB lensing. *Phys. Rev.*, D95(12):123512, 2017.

[263] Elisabeth Krause and Tim Eifler. cosmolike − cosmological likelihood analyses for photometric galaxy surveys. *Mon. Not. Roy. Astron. Soc.*, 470(2):2100–2112, 2017.

[264] Luca Amendola et al. Cosmology and fundamental physics with the Euclid satellite. *Living Rev. Rel.*, 16:6, 2013.

[265] N. G. Busca, T. Delubac, J. Rich, S. Bailey, A. Font-Ribera, D. Kirkby, J.-M. Le Goff, M. M. Pieri, A. Slosar, É. Aubourg, J. E. Bautista, D. Bizyaev, M. Blomqvist, A. S. Bolton, J. Bovy, H. Brewington, A. Borde, J. Brinkmann, B. Carithers, R. A. C. Croft, K. S. Dawson, G. Ebelke, D. J. Eisenstein, J.-C. Hamilton, S. Ho, D. W. Hogg, K. Honscheid, K.-G. Lee, B. Lundgren, E. Malanushenko,

V. Malanushenko, D. Margala, C. Maraston, K. Mehta, J. Miralda-Escudé, A. D. Myers, R. C. Nichol, P. Noterdaeme, M. D. Olmstead, D. Oravetz, N. Palanque-Delabrouille, K. Pan, I. Pâris, W. J. Percival, P. Petitjean, N. A. Roe, E. Rollinde, N. P. Ross, G. Rossi, D. J. Schlegel, D. P. Schneider, A. Shelden, E. S. Sheldon, A. Simmons, S. Snedden, J. L. Tinker, M. Viel, B. A. Weaver, D. H. Weinberg, M. White, C. Yèche, and D. G. York. Baryon acoustic oscillations in the Lyα forest of BOSS quasars. A&A, 552:A96, April 2013.

[266] J. M. Le Goff, C. Magneville, E. Rollinde, S. Peirani, P. Petitjean, C. Pichon, J. Rich, C. Yeche, E. Aubourg, N. Busca, R. Charlassier, T. Delubac, J. C. Hamilton, N. Palanque Delabrouille, I. Pâris, and M. Vargas Magaña. Simulations of BAO reconstruction with a quasar Ly-α survey. A&A, 534:A135, October 2011.

[267] Anse Slosar et al. Measurement of Baryon Acoustic Oscillations in the Lyman-alpha Forest Fluctuations in BOSS Data Release 9. *JCAP*, 1304:026, 2013.

[268] K. S. Dawson, J.-P. Kneib, W. J. Percival, S. Alam, F. D. Albareti, S. F. Anderson, E. Armengaud, É. Aubourg, S. Bailey, J. E. Bautista, A. A. Berlind, M. A. Bershady, F. Beutler, D. Bizyaev, M. R. Blanton, M. Blomqvist, A. S. Bolton, J. Bovy, W. N. Brandt, J. Brinkmann, J. R. Brownstein, E. Burtin, N. G. Busca, Z. Cai, C.-H. Chuang, N. Clerc, J. Comparat, F. Cope, R. A. C. Croft, I. Cruz-Gonzalez, L. N. da Costa, M.-C. Cousinou, J. Darling, A. de la Macorra, S. de la Torre, T. Delubac, H. du Mas des Bourboux, T. Dwelly, A. Ealet, D. J. Eisenstein, M. Eracleous, S. Escoffier, X. Fan, A. Finoguenov, A. Font-Ribera, P. Frinchaboy, P. Gaulme, A. Georgakakis, P. Green, H. Guo, J. Guy, S. Ho, D. Holder, J. Huehnerhoff, T. Hutchinson, Y. Jing, E. Jullo, V. Kamble, K. Kinemuchi, D. Kirkby, F.-S. Kitaura, M. A. Klaene, R. R. Laher, D. Lang, P. Laurent, J.-M. Le Goff, C. Li, Y. Liang, M. Lima, Q. Lin, W. Lin, Y.-T. Lin, D. C. Long, B. Lundgren, N. MacDonald, M. A. Geimba Maia, E. Malanushenko, V. Malanushenko, V. Mariappan, C. K. McBride, I. D. McGreer, B. Ménard, A. Merloni, A. Meza, A. D. Montero-Dorta, D. Muna, A. D. Myers, K. Nandra, T. Naugle, J. A. Newman, P. Noterdaeme, P. Nugent, R. Ogando, M. D. Olmstead, A. Oravetz, D. J. Oravetz, N. Padmanabhan, N. Palanque-Delabrouille, K. Pan, J. K. Parejko, I. Pâris, J. A. Peacock, P. Petitjean, M. M. Pieri, A. Pisani, F. Prada, A. Prakash, A. Raichoor, B. Reid, J. Rich, J. Ridl, S. Rodriguez-Torres, A. Carnero Rosell, A. J. Ross, G. Rossi, J. Ruan, M. Salvato, C. Sayres, D. P. Schneider, D. J. Schlegel, U. Seljak, H.-J. Seo, B. Sesar, S. Shandera, Y. Shu, A. Slosar, F. Sobreira, A. Streblyanska, N. Suzuki, D. Taylor, C. Tao, J. L. Tinker, R. Tojeiro, M. Vargas-Magaña, Y. Wang, B. A. Weaver, D. H. Weinberg, M. White, W. M. Wood-Vasey, C. Yeche, Z. Zhai, C. Zhao, G.-b. Zhao, Z. Zheng, G. Ben Zhu, and H. Zou. The SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Overview and Early Data. AJ, 151:44, February 2016.

[269] Amir Aghamousa et al. The DESI Experiment Part I: Science, Targeting, and Survey Design. 2016.

[270] H. Hildebrandt, M. Viola, C. Heymans, S. Joudaki, K. Kuijken, C. Blake, T. Erben, B. Joachimi, D. Klaes, L. Miller, C. B. Morrison, R. Nakajima, G. Verdoes Kleijn, A. Amon, A. Choi, G. Covone, J. T. A. de Jong, A. Dvornik, I. Fenech Conti, A. Grado, J. Harnois-Déraps, R. Herbonnet, H. Hoekstra, F. Köhlinger,

J. McFarland, A. Mead, J. Merten, N. Napolitano, J. A. Peacock, M. Radovich, P. Schneider, P. Simon, E. A. Valentijn, J. L. van den Busch, E. van Uitert, and L. Van Waerbeke. KiDS-450: cosmological parameter constraints from tomographic weak gravitational lensing. *Mon. Not. R. Astron. Soc.* , 465:1454–1498, February 2017.

[271] F. Köhlinger et al. KiDS-450: The tomographic weak lensing power spectrum and constraints on cosmological parameters. *Mon. Not. Roy. Astron. Soc.*, 471(4):4412–4435, 2017.

[272] S. Joudaki, C. Blake, C. Heymans, A. Choi, J. Harnois-Deraps, H. Hildebrandt, B. Joachimi, A. Johnson, A. Mead, D. Parkinson, M. Viola, and L. van Waerbeke. CFHTLenS revisited: assessing concordance with Planck including astrophysical systematics. *Mon. Not. R. Astron. Soc.* , 465:2033–2052, February 2017.

[273] Vivian Poulin, Kimberly K. Boddy, Simeon Bird, and Marc Kamionkowski. Implications of an extended dark energy cosmology with massive neutrinos for cosmological tensions. *Phys. Rev.*, D97(12):123504, 2018.

[274] M. J. Mortonson, D. Huterer, and W. Hu. Figures of merit for present and future dark energy probes. Phys. Rev. D, 82(6):063004, September 2010.

[275] Elena Sellentin and Alan F. Heavens. Parameter inference with estimated covariance matrices. *Mon. Not. Roy. Astron. Soc.*, 456(1):L132–L136, 2016.

[276] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, and et al. Planck intermediate results. *Astronomy & Astrophysics*, 607:A95, Nov 2017.

[277] ChangHoon Hahn, Florian Beutler, Manodeep Sinha, Andreas Berlind, Shirley Ho, and David W. Hogg. Likelihood Non-Gaussianity in Large-Scale Structure Analyses. *Mon. Not. Roy. Astron. Soc.*, 485(2):2956–2969, 2019.

[278] Elena Sellentin and Alan F. Heavens. On the insufficiency of arbitrarily precise covariance matrices: non-Gaussian weak lensing likelihoods. *Mon. Not. Roy. Astron. Soc.*, 473(2):2355–2363, 2018.

[279] Elena Sellentin and Jean-Luc Starck. Debiasing inference with approximate covariance matrices and other unidentified biases. JCAP, 2019(8):021, August 2019.

[280] Will Grathwohl, Ricky T. Q. Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. FFJORD: free-form continuous dynamics for scalable reversible generative models. *CoRR*, abs/1810.01367, 2018.

[281] Danilo Jimenez Rezende and Shakir Mohamed. Variational Inference with Normalizing Flows. *arXiv e-prints*, page arXiv:1505.05770, May 2015.

[282] Héctor J. Hortúa, Luigi Malago, and Riccardo Volpi. Constraining the Reionization History using Bayesian Normalizing Flows. *arXiv e-prints*, page arXiv:2005.07694, May 2020.

[283] Vanessa BÃ¶hm and UroÅ¡ Seljak. Probabilistic auto-encoder, 2020.

[284] Johann Brehmer and Kyle Cranmer. Flows for simultaneous manifold learning and density estimation, 2020.

[285] Biwei Dai and Uros Seljak. Sliced iterative generator, 2020.

[286] George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation, 2018.

[287] Laurent Dinh, David Krueger, and Yoshua Bengio. NICE: Non-linear Independent Components Estimation. *arXiv e-prints*, page arXiv:1410.8516, Oct 2014.

[288] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using Real NVP. *arXiv e-prints*, page arXiv:1605.08803, May 2016.

[289] Diederik P. Kingma and Prafulla Dhariwal. Glow: Generative Flow with Invertible 1x1 Convolutions. *arXiv e-prints*, page arXiv:1807.03039, Jul 2018.

[290] He Jia and Uroš Seljak. Normalizing Constant Estimation with Gaussianized Bridge Sampling. *arXiv e-prints*, page arXiv:1912.06073, December 2019.

[291] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows, 2019.

[292] Eric Nalisnick, Akihiro Matsukawa, Yee Whye Teh, Dilan Gorur, and Balaji Lakshminarayanan. Do deep generative models know what they don't know?, 2019.

[293] Vanessa BÃ¶hm, FranÃ§ois Lanusse, and UroÅ¡ Seljak. Uncertainty quantification with generative models, 2019.

[294] George Papamakarios, David Sterratt, and Iain Murray. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. volume 89 of *Proceedings of Machine Learning Research*, pages 837–848. PMLR, 16–18 Apr 2019.

[295] Justin Alsing, Benjamin Wandelt, and Stephen Feeney. Massive optimal data compression and density estimation for scalable, likelihood-free inference in cosmology. *Monthly Notices of the Royal Astronomical Society*, 477(3):2874 –– 2885, Mar 2018.

[296] Peter L. Taylor, Thomas D. Kitching, Justing Alsing, Benjamin D. Wandelt, Stephen M. Feeney, and Jason D. McEwen. Cosmic shear: Inference from forward models. *Physical Review D*, 100(2), Jul 2019.

[297] Conor Durkan, George Papamakarios, and Iain Murray. Sequential neural methods for likelihood-free inference, 2018.

[298] Jan-Matthis Lueckmann, Giacomo Bassetto, Theofanis Karaletsos, and Jakob H. Macke. Likelihood-free inference with emulator networks. volume 96 of *Proceedings of Machine Learning Research*, pages 32–53. PMLR, 02 Dec 2019.

[299] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 2020.

[300] Alvaro Tejero-Cantero, Jan Boelts, Michael Deistler, Jan-Matthis Lueckmann,

Conor Durkan, Pedro J. GonÃ§alves, David S. Greenberg, and Jakob H. Macke. sbi: A toolkit for simulation-based inference. *Journal of Open Source Software*, 5(52):2505, 2020.

[301] Miles D. Cranmer, Richard Galvez, Lauren Anderson, David N. Spergel, and Shirley Ho. Modeling the gaia color-magnitude diagram with bayesian neural flows to constrain distance estimates, 2019.

[302] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[303] T. K. Moon. The expectation-maximization algorithm. *IEEE Signal Processing Magazine*, 13(6):47–60, 1996.

[304] Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006.

[305] Jonathon Shlens. A Tutorial on Independent Component Analysis. *arXiv e-prints*, page arXiv:1404.2986, April 2014.

[306] D. W. Scott. *Multivariate Density Estimation*. 1992.

[307] Diederik P. Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improving Variational Inference with Inverse Autoregressive Flow. *arXiv e-prints*, page arXiv:1606.04934, June 2016.

[308] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural Ordinary Differential Equations. *arXiv e-prints*, page arXiv:1806.07366, June 2018.

[309] S. Kullback and R. A. Leibler. On information and sufficiency. *Ann. Math. Statist.*, 22(1):79–86, 03 1951.

[310] Q. Wang, S. R. Kulkarni, and S. Verdu. Divergence estimation for multidimensional densities via k-nearest-neighbor distances. *IEEE Transactions on Information Theory*, 55(5):2392–2405, 2009.

[311] Malin Velander, Edo van Uitert, Henk Hoekstra, Jean Coupon, Thomas Erben, Catherine Heymans, Hendrik Hildebrandt, Thomas D. Kitching, Yannick Mellier, Lance Miller, Ludovic Van Waerbeke, Christopher Bonnett, Liping Fu, Stefania Giodini, Michael J. Hudson, Konrad Kuijken, Barnaby Rowe, Tim Schrabback, and Elisabetta Semboloni. CFHTLenS: the relation between galaxy dark matter haloes and baryons from weak gravitational lensing. *Monthly Notices of the Royal Astronomical Society*, 437(3):2111–2136, 11 2013.

[312] A. Petri. Mocking the weak lensing universe: The LensTools Python computing package. *Astronomy and Computing*, 17:73–79, October 2016.

[313] Arushi Gupta, José Manuel Zorrilla Matilla, Daniel Hsu, and Zoltán Haiman. Non-Gaussian information from weak lensing data via deep learning. Phys. Rev. D,

97(10):103515, May 2018.

[314] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, and et al. Planck 2018 results. v. cmb power spectra and likelihoods. *Astronomy & Astrophysics*, 641:A5, Sep 2020.

[315] Elena Sellentin, Catherine Heymans, and Joachim Harnois-Déraps. The skewed weak lensing likelihood: why biases arise, despite data and theory being sound. *Mon. Not. R. Astron. Soc.* , 477(4):4879–4895, July 2018.

[316] Catherine Heymans, Ludovic Van Waerbeke, Lance Miller, Thomas Erben, Hendrik Hildebrandt, Henk Hoekstra, Thomas D. Kitching, Yannick Mellier, Patrick Simon, Christopher Bonnett, Jean Coupon, Liping Fu, Joachim Harnois Déraps, Michael J. Hudson, Martin Kilbinger, Koenraad Kuijken, Barnaby Rowe, Tim Schrabback, Elisabetta Semboloni, Edo van Uitert, Sanaz Vafaei, and Malin Velander. CFHTLenS: the Canada-France-Hawaii Telescope Lensing Survey. *Mon. Not. R. Astron. Soc.* , 427(1):146–166, November 2012.

[317] Chien-Hao Lin, Joachim Harnois-DÃ©raps, Tim Eifler, Taylor Pospisil, Rachel Mandelbaum, Ann B. Lee, and Sukhdeep Singh. Non-gaussianity in the weak lensing correlation function likelihood - implications for cosmological parameter biases, 2019.

[318] J. Colin Hill, Blake D. Sherwin, Kendrick M. Smith, Graeme E. Addison, Nick Battaglia, Elia S. Battistelli, J. Richard Bond, Erminia Calabrese, Mark J. Devlin, Joanna Dunkley, Rolando Dunner, Thomas Essinger-Hileman, Megan B. Gralla, Amir Hajian, Matthew Hasselfield, Adam D. Hincks, Renee Hlozek, John P. Hughes, Arthur Kosowsky, Thibaut Louis, Danica Marsden, Kavilan Moodley, Michael D. Niemack, Lyman A. Page, Bruce Partridge, Benjamin Schmitt, Neelima Sehgal, Jonathan L. Sievers, David N. Spergel, Suzanne T. Staggs, Daniel S. Swetz, Robert Thornton, Hy Trac, and Edward J. Wollack. The atacama cosmology telescope: A measurement of the thermal sunyaev-zel'dovich one-point pdf, 2014.

[319] Justin Alsing, Tom Charnock, Stephen Feeney, and Benjamin Wandelt. Fast likelihood-free cosmology with neural density estimators and active learning. *Monthly Notices of the Royal Astronomical Society*, Jul 2019.

[320] Georges Obied, Cora Dvorkin, Chen Heinrich, Wayne Hu, and Vinicius Miranda. Inflationary features and shifts in cosmological parameters from planck 2015 data. *Physical Review D*, 96(8), Oct 2017.

[321] Georges Obied, Cora Dvorkin, Chen Heinrich, Wayne Hu, and V. Miranda. Inflationary versus reionization features from planck 2015 data. *Physical Review D*, 98(4), Aug 2018.

[322] Adam G. Riess, Lucas M. Macri, Samantha L. Hoffmann, Dan Scolnic, Stefano Casertano, Alexei V. Filippenko, Brad E. Tucker, Mark J. Reid, David O. Jones, Jeffrey M. Silverman, and et al. A 2.4% determination of the local value of the hubble constant. *The Astrophysical Journal*, 826(1):56, Jul 2016.

[323] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ack-

ley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, C. Affeldt, M. Agathos, K. Agatsuma, N. Aggarwal, O. D. Aguiar, L. Aiello, A. Ain, P. Ajith, B. Allen, A. Allocca, P. A. Altin, S. B. Anderson, W. G. Anderson, K. Arai, M. A. Arain, M. C. Araya, C. C. Arceneaux, J. S. Areeda, N. Arnaud, K. G. Arun, S. Ascenzi, G. Ashton, M. Ast, S. M. Aston, P. Astone, P. Aufmuth, C. Aulbert, S. Babak, P. Bacon, M. K. M. Bader, P. T. Baker, F. Baldaccini, G. Ballardin, S. W. Ballmer, J. C. Barayoga, S. E. Barclay, B. C. Barish, D. Barker, F. Barone, B. Barr, L. Barsotti, M. Barsuglia, D. Barta, J. Bartlett, M. A. Barton, I. Bartos, R. Bassiri, A. Basti, J. C. Batch, C. Baune, V. Bavigadda, M. Bazzan, B. Behnke, M. Bejger, C. Belczynski, A. S. Bell, C. J. Bell, B. K. Berger, J. Bergman, G. Bergmann, C. P. L. Berry, D. Bersanetti, A. Bertolini, J. Betzwieser, S. Bhagwat, R. Bhandare, I. A. Bilenko, G. Billingsley, J. Birch, R. Birney, O. Birnholtz, S. Biscans, A. Bisht, M. Bitossi, C. Biwer, M. A. Bizouard, J. K. Blackburn, C. D. Blair, D. G. Blair, R. M. Blair, S. Bloemen, O. Bock, T. P. Bodiya, M. Boer, G. Bogaert, C. Bogan, A. Bohe, P. Bojtos, C. Bond, F. Bondu, R. Bonnand, B. A. Boom, R. Bork, V. Boschi, S. Bose, Y. Bouffanais, A. Bozzi, C. Bradaschia, P. R. Brady, V. B. Braginsky, M. Branchesi, J. E. Brau, T. Briant, A. Brillet, M. Brinkmann, V. Brisson, P. Brockill, A. F. Brooks, D. A. Brown, D. D. Brown, N. M. Brown, C. C. Buchanan, A. Buikema, T. Bulik, H. J. Bulten, A. Buonanno, D. Buskulic, C. Buy, R. L. Byer, M. Cabero, L. Cadonati, G. Cagnoli, C. Cahillane, J. Calderón Bustillo, T. Callister, E. Calloni, J. B. Camp, K. C. Cannon, J. Cao, C. D. Capano, E. Capocasa, F. Carbognani, S. Caride, J. Casanueva Diaz, C. Casentini, S. Caudill, M. Cavaglià, F. Cavalier, R. Cavalieri, G. Cella, C. B. Cepeda, L. Cerboni Baiardi, G. Cerretani, E. Cesarini, R. Chakraborty, T. Chalermsongsak, S. J. Chamberlin, M. Chan, S. Chao, P. Charlton, E. Chassande-Mottin, H. Y. Chen, Y. Chen, C. Cheng, A. Chincarini, A. Chiummo, H. S. Cho, M. Cho, J. H. Chow, N. Christensen, Q. Chu, S. Chua, S. Chung, G. Ciani, F. Clara, J. A. Clark, F. Cleva, E. Coccia, P.-F. Cohadon, A. Colla, C. G. Collette, L. Cominsky, M. Constancio, A. Conte, L. Conti, D. Cook, T. R. Corbitt, N. Cornish, A. Corsi, S. Cortese, C. A. Costa, M. W. Coughlin, S. B. Coughlin, J.-P. Coulon, S. T. Countryman, P. Couvares, E. E. Cowan, D. M. Coward, M. J. Cowart, D. C. Coyne, R. Coyne, K. Craig, J. D. E. Creighton, T. D. Creighton, J. Cripe, S. G. Crowder, A. M. Cruise, A. Cumming, L. Cunningham, E. Cuoco, T. Dal Canton, S. L. Danilishin, S. D'Antonio, K. Danzmann, N. S. Darman, C. F. Da Silva Costa, V. Dattilo, I. Dave, H. P. Daveloza, M. Davier, G. S. Davies, E. J. Daw, R. Day, S. De, D. DeBra, G. Debreczeni, J. Degallaix, M. De Laurentis, S. Deléglise, W. Del Pozzo, T. Denker, T. Dent, H. Dereli, V. Dergachev, R. T. DeRosa, R. De Rosa, R. DeSalvo, S. Dhurandhar, M. C. Díaz, L. Di Fiore, M. Di Giovanni, A. Di Lieto, S. Di Pace, I. Di Palma, A. Di Virgilio, G. Dojcinoski, V. Dolique, F. Donovan, K. L. Dooley, S. Doravari, R. Douglas, T. P. Downes, M. Drago, R. W. P. Drever, J. C. Driggers, Z. Du, M. Ducrot, S. E. Dwyer, T. B. Edo, M. C. Edwards, A. Effler, H.-B. Eggenstein, P. Ehrens, J. Eichholz, S. S. Eikenberry, W. Engels, R. C. Essick, T. Etzel, M. Evans, T. M. Evans, R. Everett, M. Factourovich, V. Fafone, H. Fair, S. Fairhurst, X. Fan, Q. Fang, S. Farinon, B. Farr, W. M. Farr, M. Favata, M. Fays, H. Fehrmann, M. M. Fejer, D. Feldbaum, I. Ferrante, E. C. Ferreira, F. Ferrini, F. Fidecaro, L. S. Finn, I. Fiori, D. Fiorucci, R. P. Fisher, R. Flaminio, M. Fletcher, H. Fong, J.-D. Fournier, S. Franco, S. Frasca, F. Frasconi, M. Frede, Z. Frei, A. Freise, R. Frey, V. Frey, T. T. Fricke, P. Fritschel,

V. V. Frolov, P. Fulda, M. Fyffe, H. A. G. Gabbard, J. R. Gair, L. Gammaitoni, S. G. Gaonkar, F. Garufi, A. Gatto, G. Gaur, N. Gehrels, G. Gemme, B. Gendre, E. Genin, A. Gennai, J. George, L. Gergely, V. Germain, Abhirup Ghosh, Archisman Ghosh, S. Ghosh, J. A. Giaime, K. D. Giardina, A. Giazotto, K. Gill, A. Glaefke, J. R. Gleason, E. Goetz, R. Goetz, L. Gondan, G. González, J. M. Gonzalez Castro, A. Gopakumar, N. A. Gordon, M. L. Gorodetsky, S. E. Gossan, M. Gosselin, R. Gouaty, C. Graef, P. B. Graff, M. Granata, A. Grant, S. Gras, C. Gray, G. Greco, A. C. Green, R. J. S. Greenhalgh, P. Groot, H. Grote, S. Grunewald, G. M. Guidi, X. Guo, A. Gupta, M. K. Gupta, K. E. Gushwa, E. K. Gustafson, R. Gustafson, J. J. Hacker, B. R. Hall, E. D. Hall, G. Hammond, M. Haney, M. M. Hanke, J. Hanks, C. Hanna, M. D. Hannam, J. Hanson, T. Hardwick, J. Harms, G. M. Harry, I. W. Harry, M. J. Hart, M. T. Hartman, C.-J. Haster, K. Haughian, J. Healy, J. Heefner, A. Heidmann, M. C. Heintze, G. Heinzel, H. Heitmann, P. Hello, G. Hemming, M. Hendry, I. S. Heng, J. Hennig, A. W. Heptonstall, M. Heurs, S. Hild, D. Hoak, K. A. Hodge, D. Hofman, S. E. Hollitt, K. Holt, D. E. Holz, P. Hopkins, D. J. Hosken, J. Hough, E. A. Houston, E. J. Howell, Y. M. Hu, S. Huang, E. A. Huerta, D. Huet, B. Hughey, S. Husa, S. H. Huttner, T. Huynh-Dinh, A. Idrisy, N. Indik, D. R. Ingram, R. Inta, H. N. Isa, J.-M. Isac, M. Isi, G. Islas, T. Isogai, B. R. Iyer, K. Izumi, M. B. Jacobson, T. Jacqmin, H. Jang, K. Jani, P. Jaranowski, S. Jawahar, F. Jiménez-Forteza, W. W. Johnson, N. K. Johnson-McDaniel, D. I. Jones, R. Jones, R. J. G. Jonker, L. Ju, K. Haris, C. V. Kalaghatgi, V. Kalogera, S. Kandhasamy, G. Kang, J. B. Kanner, S. Karki, M. Kasprzack, E. Katsavounidis, W. Katzman, S. Kaufer, T. Kaur, K. Kawabe, F. Kawazoe, F. Kéfélian, M. S. Kehl, D. Keitel, D. B. Kelley, W. Kells, R. Kennedy, D. G. Keppel, J. S. Key, A. Khalaidovski, F. Y. Khalili, I. Khan, S. Khan, Z. Khan, E. A. Khazanov, N. Kijbunchoo, C. Kim, J. Kim, K. Kim, Nam-Gyu Kim, Namjun Kim, Y.-M. Kim, E. J. King, P. J. King, D. L. Kinzel, J. S. Kissel, L. Kleybolte, S. Klimenko, S. M. Koehlenbeck, K. Kokeyama, S. Koley, V. Kondrashov, A. Kontos, S. Koranda, M. Korobko, W. Z. Korth, I. Kowalska, D. B. Kozak, V. Kringel, B. Krishnan, A. Królak, C. Krueger, G. Kuehn, P. Kumar, R. Kumar, L. Kuo, A. Kutynia, P. Kwee, B. D. Lackey, M. Landry, J. Lange, B. Lantz, P. D. Lasky, A. Lazzarini, C. Lazzaro, P. Leaci, S. Leavey, E. O. Lebigot, C. H. Lee, H. K. Lee, H. M. Lee, K. Lee, A. Lenon, M. Leonardi, J. R. Leong, N. Leroy, N. Letendre, Y. Levin, B. M. Levine, T. G. F. Li, A. Libson, T. B. Littenberg, N. A. Lockerbie, J. Logue, A. L. Lombardi, L. T. London, J. E. Lord, M. Lorenzini, V. Loriette, M. Lormand, G. Losurdo, J. D. Lough, C. O. Lousto, G. Lovelace, H. Lück, A. P. Lundgren, J. Luo, R. Lynch, Y. Ma, T. MacDonald, B. Machenschalk, M. MacInnis, D. M. Macleod, F. Magaña Sandoval, R. M. Magee, M. Mageswaran, E. Majorana, I. Maksimovic, V. Malvezzi, N. Man, I. Mandel, V. Mandic, V. Mangano, G. L. Mansell, M. Manske, M. Mantovani, F. Marchesoni, F. Marion, S. Márka, Z. Márka, A. S. Markosyan, E. Maros, F. Martelli, L. Martellini, I. W. Martin, R. M. Martin, D. V. Martynov, J. N. Marx, K. Mason, A. Masserot, T. J. Massinger, M. Masso-Reid, F. Matichard, L. Matone, N. Mavalvala, N. Mazumder, G. Mazzolo, R. McCarthy, D. E. McClelland, S. McCormick, S. C. McGuire, G. McIntyre, J. McIver, D. J. McManus, S. T. McWilliams, D. Meacher, G. D. Meadors, J. Meidam, A. Melatos, G. Mendell, D. Mendoza-Gandara, R. A. Mercer, E. Merilh, M. Merzougui, S. Meshkov, C. Messenger, C. Messick, P. M. Meyers, F. Mezzani, H. Miao, C. Michel, H. Middleton, E. E. Mikhailov, L. Milano, J. Miller, M. Mill-

house, Y. Minenkov, J. Ming, S. Mirshekari, C. Mishra, S. Mitra, V. P. Mitrofanov, G. Mitselmakher, R. Mittleman, A. Moggi, M. Mohan, S. R. P. Mohapatra, M. Montani, B. C. Moore, C. J. Moore, D. Moraru, G. Moreno, S. R. Morriss, K. Mossavi, B. Mours, C. M. Mow-Lowry, C. L. Mueller, G. Mueller, A. W. Muir, Arunava Mukherjee, D. Mukherjee, S. Mukherjee, N. Mukund, A. Mullavey, J. Munch, D. J. Murphy, P. G. Murray, A. Mytidis, I. Nardecchia, L. Naticchioni, R. K. Nayak, V. Necula, K. Nedkova, G. Nelemans, M. Neri, A. Neunzert, G. Newton, T. T. Nguyen, A. B. Nielsen, S. Nissanke, A. Nitz, F. Nocera, D. Nolting, M. E. N. Normandin, L. K. Nuttall, J. Oberling, E. Ochsner, J. O'Dell, E. Oelker, G. H. Ogin, J. J. Oh, S. H. Oh, F. Ohme, M. Oliver, P. Oppermann, Richard J. Oram, B. O'Reilly, R. O'Shaughnessy, C. D. Ott, D. J. Ottaway, R. S. Ottens, H. Overmier, B. J. Owen, A. Pai, S. A. Pai, J. R. Palamos, O. Palashov, C. Palomba, A. Pal-Singh, H. Pan, Y. Pan, C. Pankow, F. Pannarale, B. C. Pant, F. Paoletti, A. Paoli, M. A. Papa, H. R. Paris, W. Parker, D. Pascucci, A. Pasqualetti, R. Passaquieti, D. Passuello, B. Patricelli, Z. Patrick, B. L. Pearlstone, M. Pedraza, R. Pedurand, L. Pekowsky, A. Pele, S. Penn, A. Perreca, H. P. Pfeiffer, M. Phelps, O. Piccinni, M. Pichot, M. Pickenpack, F. Piergiovanni, V. Pierro, G. Pillant, L. Pinard, I. M. Pinto, M. Pitkin, J. H. Poeld, R. Poggiani, P. Popolizio, A. Post, J. Powell, J. Prasad, V. Predoi, S. S. Premachandra, T. Prestegard, L. R. Price, M. Prijatelj, M. Principe, S. Privitera, R. Prix, G. A. Prodi, L. Prokhorov, O. Puncken, M. Punturo, P. Puppo, M. Pürrer, H. Qi, J. Qin, V. Quetschke, E. A. Quintero, R. Quitzow-James, F. J. Raab, D. S. Rabeling, H. Radkins, P. Raffai, S. Raja, M. Rakhmanov, C. R. Ramet, P. Rapagnani, V. Raymond, M. Razzano, V. Re, J. Read, C. M. Reed, T. Regimbau, L. Rei, S. Reid, D. H. Reitze, H. Rew, S. D. Reyes, F. Ricci, K. Riles, N. A. Robertson, R. Robie, F. Robinet, A. Rocchi, L. Rolland, J. G. Rollins, V. J. Roma, J. D. Romano, R. Romano, G. Romanov, J. H. Romie, D. Rosińska, S. Rowan, A. Rüdiger, P. Ruggi, K. Ryan, S. Sachdev, T. Sadecki, L. Sadeghian, L. Salconi, M. Saleem, F. Salemi, A. Samajdar, L. Sammut, L. M. Sampson, E. J. Sanchez, V. Sandberg, B. Sandeen, G. H. Sanders, J. R. Sanders, B. Sassolas, B. S. Sathyaprakash, P. R. Saulson, O. Sauter, R. L. Savage, A. Sawadsky, P. Schale, R. Schilling, J. Schmidt, P. Schmidt, R. Schnabel, R. M. S. Schofield, A. Schönbeck, E. Schreiber, D. Schuette, B. F. Schutz, J. Scott, S. M. Scott, D. Sellers, A. S. Sengupta, D. Sentenac, V. Sequino, A. Sergeev, G. Serna, Y. Setyawati, A. Sevigny, D. A. Shaddock, T. Shaffer, S. Shah, M. S. Shahriar, M. Shaltev, Z. Shao, B. Shapiro, P. Shawhan, A. Sheperd, D. H. Shoemaker, D. M. Shoemaker, K. Siellez, X. Siemens, D. Sigg, A. D. Silva, D. Simakov, A. Singer, L. P. Singer, A. Singh, R. Singh, A. Singhal, A. M. Sintes, B. J. J. Slagmolen, J. R. Smith, M. R. Smith, N. D. Smith, R. J. E. Smith, E. J. Son, B. Sorazu, F. Sorrentino, T. Souradeep, A. K. Srivastava, A. Staley, M. Steinke, J. Steinlechner, S. Steinlechner, D. Steinmeyer, B. C. Stephens, S. P. Stevenson, R. Stone, K. A. Strain, N. Straniero, G. Stratta, N. A. Strauss, S. Strigin, R. Sturani, A. L. Stuver, T. Z. Summerscales, L. Sun, P. J. Sutton, B. L. Swinkels, M. J. Szczepańczyk, M. Tacca, D. Talukder, D. B. Tanner, M. Tápai, S. P. Tarabrin, A. Taracchini, R. Taylor, T. Theeg, M. P. Thirugnanasambandam, E. G. Thomas, M. Thomas, P. Thomas, K. A. Thorne, K. S. Thorne, E. Thrane, S. Tiwari, V. Tiwari, K. V. Tokmakov, C. Tomlinson, M. Tonelli, C. V. Torres, C. I. Torrie, D. Töyrä, F. Travasso, G. Traylor, D. Trifirò, M. C. Tringali, L. Trozzo, M. Tse, M. Turconi, D. Tuyenbayev, D. Ugolini, C. S. Unnikrishnan, A. L. Urban, S. A. Usman, H. Vahlbruch,

G. Vajente, G. Valdes, M. Vallisneri, N. van Bakel, M. van Beuzekom, J. F. J. van den Brand, C. Van Den Broeck, D. C. Vander-Hyde, L. van der Schaaf, J. V. van Heijningen, A. A. van Veggel, M. Vardaro, S. Vass, M. Vasúth, R. Vaulin, A. Vecchio, G. Vedovato, J. Veitch, P. J. Veitch, K. Venkateswara, D. Verkindt, F. Vetrano, A. Viceré, S. Vinciguerra, D. J. Vine, J.-Y. Vinet, S. Vitale, T. Vo, H. Vocca, C. Vorvick, D. Voss, W. D. Vousden, S. P. Vyatchanin, A. R. Wade, L. E. Wade, M. Wade, S. J. Waldman, M. Walker, L. Wallace, S. Walsh, G. Wang, H. Wang, M. Wang, X. Wang, Y. Wang, H. Ward, R. L. Ward, J. Warner, M. Was, B. Weaver, L.-W. Wei, M. Weinert, A. J. Weinstein, R. Weiss, T. Welborn, L. Wen, P. Weßels, T. Westphal, K. Wette, J. T. Whelan, S. E. Whitcomb, D. J. White, B. F. Whiting, K. Wiesner, C. Wilkinson, P. A. Willems, L. Williams, R. D. Williams, A. R. Williamson, J. L. Willis, B. Willke, M. H. Wimmer, L. Winkelmann, W. Winkler, C. C. Wipf, A. G. Wiseman, H. Wittel, G. Woan, J. Worden, J. L. Wright, G. Wu, J. Yablon, I. Yakushin, W. Yam, H. Yamamoto, C. C. Yancey, M. J. Yap, H. Yu, M. Yvert, A. Zadrożny, L. Zangrando, M. Zanolin, J.-P. Zendri, M. Zevin, F. Zhang, L. Zhang, M. Zhang, Y. Zhang, C. Zhao, M. Zhou, Z. Zhou, X. J. Zhu, M. E. Zucker, S. E. Zuraw, and J. Zweizig. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116:061102, Feb 2016.

[324] The Event Horizon Telescope Collaboration. First m87 event horizon telescope results. i. the shadow of the supermassive black hole, 2019.

[325] The Event Horizon Telescope Collaboration. First m87 event horizon telescope results. ii. array and instrumentation, 2019.

[326] The Event Horizon Telescope Collaboration. First m87 event horizon telescope results. iii. data processing and calibration, 2019.

[327] The Event Horizon Telescope Collaboration. First m87 event horizon telescope results. iv. imaging the central supermassive black hole, 2019.

[328] The Event Horizon Telescope Collaboration. First m87 event horizon telescope results. v. physical origin of the asymmetric ring, 2019.

[329] The Event Horizon Telescope Collaboration. First m87 event horizon telescope results. vi. the shadow and mass of the central black hole, 2019.

[330] Sreedevi Varma, Malcolm Fairbairn, and Julio Figueroa. Dark Matter Subhalos, Strong Lensing and Machine Learning. *arXiv e-prints*, page arXiv:2005.05353, May 2020.

[331] T. Sauer. *Numerical Analysis*. Pearson custom library. Pearson, 2013.

[332] Eifler, T., Schneider, P., and Hartlap, J. Dependence of cosmic shear covariances on cosmology - impact on parameter estimation. *A&A*, 502(3):721–731, 2009.

[333] Christopher B Morrison and Michael D Schneider. On estimating cosmology-dependent covariance matrices. *Journal of Cosmology and Astroparticle Physics*, 2013(11):009–009, nov 2013.

[334] Martin White and Nikhil Padmanabhan. Including parameter dependence in the

data and covariance for cosmological inference. *Journal of Cosmology and Astroparticle Physics*, 2015(12):058–058, dec 2015.

[335] Francisco-Shu Kitaura, Sergio Rodriguez-Torres, Chia-Hsun Chuang, Cheng Zhao, Francisco Prada, Hector Gil-Marin, Hong Guo, Gustavo Yepes, Anatoly Klypin, Claudia G. Scaccola, Jeremy Tinker, Cameron McBride, Beth Reid, Ariel G. Sanchez, Salvador Salazar-Albornoz, Jan Niklas Grieb, Mariana Vargas-Magana, Antonio J. Cuesta, Mark Neyrinck, Florian Beutler, Johan Comparat, Will J. Percival, and Ashley Ross. The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: mock galaxy catalogues for the BOSS Final Data Release. *Monthly Notices of the Royal Astronomical Society*, 456(4):4156–4173, 01 2016.

[336] V. Springel. N-GenIC: Cosmological structure initial conditions. Astrophysics Source Code Library, February 2015.

[337] V. Springel. The cosmological simulation code GADGET-2. *Mon. Not. R. Astron. Soc.* , 364:1105–1134, December 2005.