

# UC Riverside

## UC Riverside Electronic Theses and Dissertations

### Title

Exploring Cosmology and Astrophysics With Cosmic Structure and Simulations: Cosmic Strings, the Lyman- $\alpha$  Forest, and Machine Learning

### Permalink

<https://escholarship.org/uc/item/1kd4m87z>

### Author

Fernandez, Martin

### Publication Date

2023

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA  
RIVERSIDE

Exploring Cosmology and Astrophysics With Cosmic Structure and Simulations:  
Cosmic Strings, the Lyman- $\alpha$  Forest, and Machine Learning

A Dissertation submitted in partial satisfaction  
of the requirements for the degree of

Doctor of Philosophy

in

Physics

by

Martin A. Fernandez

March 2023

Dissertation Committee:

Dr. Simeon Bird, Chairperson

Dr. George Becker

Dr. Anson D'Aloisio

Copyright by  
Martin A. Fernandez  
2023

The Dissertation of Martin A. Fernandez is approved:

---

---

---

Committee Chairperson

University of California, Riverside

## Acknowledgments

First and foremost, I thank my advisor Simeon Bird. Simeon gave me the freedom to explore a diversity of research topics, and in doing so develop a variety of technical skills that have become my greatest asset. Simeon also supported me throughout my graduate career personally, academically, and as a sort of computer wizard, who makes all errors disappear.

Several other faculty at UC Riverside deserve acknowledgment: Yanou Cui, for helping me bridge the gap between cosmology and physics; and George Becker & Anson D'Aloisio, for serving on both my qualifying and dissertation committees – two pretty important milestones!

Much appreciation is due my fellow Bird group members, in particular: Reza Monadi for his support, humor, and delicious cakes (all credit to Asieh); Phoebe Upton Sanderbeck for greatly improving my understanding of the astrophysics I'd missed while learning quantum mechanics instead; Ming-Feng Ho for many interesting discussions and insights on machine learning; and Mahdi Qezlou for always having interesting topics to discuss at group meetings.

There are several people from my time as an undergraduate who are especially deserving of my thanks: my first advisor, Kevin Covey, for introducing me to research, supporting my wandering science eye, and repeatedly reminding me that I really am the expert on my own research; Kristen Larson, from whom I first understood science as a technical *and* creative pursuit; and Regina Barber DeGraaff, who introduced me to STEM outreach, a very important part of my academic career.

## ABSTRACT OF THE DISSERTATION

Exploring Cosmology and Astrophysics With Cosmic Structure and Simulations:  
Cosmic Strings, the Lyman- $\alpha$  Forest, and Machine Learning

by

Martin A. Fernandez

Doctor of Philosophy, Graduate Program in Physics  
University of California, Riverside, March 2023  
Dr. Simeon Bird, Chairperson

In this dissertation, I present the work I have done using simulations to understand cosmology and astrophysics.

I developed a novel method for detecting the presence of cosmic strings, which are early universe topological defects, via their effect on cosmic filaments. The method was developed by inserting the effects of cosmic strings into N-body simulations, and determining what effects are most observable. I applied this method to filamentary structure from observations of galaxies. While there was no indication of the presence of cosmic strings in current observations, future prospects are promising.

I then turned my attention to smaller scales and lower redshift phenomena, determining how simulations of the Lyman- $\alpha$  forest are affected by including species specific transfer functions in the initialization of simulations. The effect on the Lyman- $\alpha$  forest flux power spectrum is particularly important, as it is widely used, and will be measured to increasingly high precision by future observations.

Continuing on that topic, I developed a machine learning approach to interpolating

Lyman- $\alpha$  forest summary statistics from simulations. This used Gaussian processes to construct an emulator for the Lyman- $\alpha$  forest flux power spectrum. Two suites of simulations were used to train the emulator: a small sample of high resolution simulations, and a large sample of low resolution simulations. This combination allowed an emulator that predicts high resolution simulation outputs, but costs a fraction of previously used methods. This is the multi-fidelity emulator model.

Finally, I present inference using a multi-fidelity emulator, trained on large volume ( $120\text{Mpc h}^{-1}$ )<sup>3</sup> and high resolution ( $1536^3$  and  $3072^2$  gas particles) simulations. The volume and resolution of these simulations allow a more robust modeling of the scales important to the Lyman- $\alpha$  forest. By combining the Lyman- $\alpha$  forest flux power spectrum with the mean temperature of the intergalactic medium, I constrain several cosmological and astrophysical parameters.

# Contents

<b>List of Figures</b>	<b>x</b>
<b>List of Tables</b>	<b>xii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Cosmic History . . . . .	1
1.2 Structure Formation . . . . .	5
1.3 The Lyman- $\alpha$ Forest . . . . .	7
1.4 Cosmological Simulations . . . . .	9
1.5 Chapter Outline . . . . .	12
<b>2 Cosmic Strings and the Cosmic Web</b>	<b>14</b>
2.1 Abstract . . . . .	14
2.2 Introduction . . . . .	15
2.3 Simulations . . . . .	19
2.3.1 Cosmic String Wakes . . . . .	20
2.3.2 Velocity . . . . .	23
2.4 Filament Identification . . . . .	25
2.4.1 SCMS+ . . . . .	26
2.4.2 DisPerSE . . . . .	29
2.5 Results . . . . .	32
2.5.1 Filament Population . . . . .	32
2.5.2 Spatial Distribution of Filaments . . . . .	33
2.5.3 SDSS Filament Catalog . . . . .	39
2.5.4 Roman Space Telescope Projection . . . . .	40
2.6 Box Size Scaling . . . . .	42
2.7 Conclusion . . . . .	42
<b>3 Transfer Functions in Simulations</b>	<b>45</b>
3.1 Abstract . . . . .	45
3.2 Introduction . . . . .	46
3.3 Simulations . . . . .	49



3.4	Methods & Results . . . . .	52
3.4.1	Curvature . . . . .	53
3.4.2	Doppler Width Methods . . . . .	55
3.4.3	Flux Power Spectrum . . . . .	62
3.4.4	Matter Power Spectrum . . . . .	66
3.5	Convergence . . . . .	68
3.6	Summary & Conclusions . . . . .	75
<b>4</b>	<b>Emulating the Lyman-<math>\alpha</math> Forest</b>	<b>78</b>
4.1	Abstract . . . . .	78
4.2	Introduction . . . . .	79
4.3	Simulations . . . . .	84
4.4	Emulators . . . . .	91
4.4.1	Gaussian process emulator . . . . .	92
4.4.2	Multi-Fidelity Emulation . . . . .	94
4.5	Results & Discussion . . . . .	100
4.5.1	Emulator Accuracy . . . . .	101
4.5.2	Emulator Runtime . . . . .	106
4.6	Non-linear Multi-Fidelity Emulator . . . . .	110
4.7	Conclusions . . . . .	112
<b>5</b>	<b>Inference via the Lyman-<math>\alpha</math> Forest</b>	<b>115</b>
5.1	Abstract . . . . .	115
5.2	Introduction . . . . .	116
5.3	Emulator Methods . . . . .	122
5.3.1	Gaussian Processes . . . . .	123
5.3.2	Single-Fidelity Emulator . . . . .	125
5.3.3	Multi-Fidelity Emulator . . . . .	127
5.4	Simulation Suite . . . . .	131
5.4.1	Cosmological & Astrophysical Parameters . . . . .	133
5.4.2	Sample Design . . . . .	136
5.4.3	Flux Power Spectrum . . . . .	138
5.4.4	IGM Mean Temperature . . . . .	139
5.5	Inference Scheme . . . . .	142
5.5.1	Flux Power Spectrum Data . . . . .	142
5.5.2	IGM Mean Temperature Data . . . . .	144
5.5.3	Likelihood . . . . .	146
5.6	Results . . . . .	148
5.6.1	Cosmology Results . . . . .	149
5.6.2	Astrophysics Results . . . . .	154
5.6.3	Maximum Posterior Predictions . . . . .	156
5.7	Conclusions . . . . .	160
5.8	Full Posteriors . . . . .	163
5.9	Inference Using Simulation Data . . . . .	168
5.10	Alternative Likelihood Results . . . . .	170

5.10.1	Parameter Priors . . . . .	170
5.10.2	BOSS DR9 Data . . . . .	170
5.10.3	Inflated BOSS Errors . . . . .	172
5.10.4	Mean Temperature Only Emulator . . . . .	174
5.10.5	Reduced Scale Range . . . . .	176
<b>6</b>	<b>Conclusions</b>	<b>179</b>
	<b>Bibliography</b>	<b>182</b>

# List of Figures

1.1	Dark matter and gas from a hydrodynamic simulations snapshot. . . . .	4
1.2	Examples of the observed and simulated Lyman- $\alpha$ forest spectra. . . . .	9
1.3	Simulation resolution example, temperature and density. . . . .	12
2.1	Setup for cosmic string & wake. . . . .	23
2.2	Visual wake versus no-wake comparison. . . . .	24
2.3	Average particle speed perpendicular to wake. . . . .	25
2.4	SCMS+ filaments. . . . .	27
2.5	SCMS+ filaments. . . . .	30
2.6	Filament number comparison (SCMS+, DisPerSE). . . . .	31
2.7	Filament number, increasing tension (SCMS+). . . . .	34
2.8	SCMS+ filament mass distributions. . . . .	34
2.9	Central filament overdensity, increasing tension. . . . .	36
2.10	Example of cosmic string detection method. . . . .	36
2.11	Filament slice overdensity (SCMS+). . . . .	37
2.12	Maximum number density excess by viewing angle. . . . .	38
2.13	Filament slice overdensity (SDSS). . . . .	39
2.14	Box size & resolution check (SCMS+). . . . .	43
3.1	Example of the curvature measurement. . . . .	54
3.2	Average curvature versus redshift . . . . .	56
3.3	Column density, Doppler width distribution at $z = 2.4$ . . . . .	58
3.4	Results for the Doppler width cutoff parameters. . . . .	59
3.5	Doppler width distributions at $z = 3$ . . . . .	61
3.6	Transfer function effect on flux power spectrum . . . . .	63
3.7	Flux power comparison with observations. . . . .	65
3.8	Baryon power versus CDM power. . . . .	67
3.9	Convergence for the curvature statistic. . . . .	69
3.10	Convergence for the Doppler width cutoff. . . . .	70
3.11	Convergence for the Doppler width distribution. . . . .	71
3.12	Convergence of the flux power spectrum. . . . .	72
3.13	Convergence with number of artificial spectra. . . . .	74

4.1	Simulation resolution and spectra comparison. . . . .	85
4.2	Example Lyman- $\alpha$ forest flux power spectra at two fidelities. . . . .	88
4.3	Simulation parameter limits and samples. . . . .	90
4.4	Latin hypercube sampling example. . . . .	91
4.5	Scale parameter for multi-fidelity emulator. . . . .	95
4.6	Multi-fidelity flowchart. . . . .	100
4.7	Emulator errors versus wavevector. . . . .	102
4.8	Small and large scale emulator error versus redshift. . . . .	103
4.9	Average multi-fidelity emulator errors, all scales and redshifts. . . . .	105
4.10	Emulator error versus number of training simulations. . . . .	107
4.11	Emulator error versus training sample computational runtime. . . . .	108
4.12	Non-linear multi-fidelity emulator performance. . . . .	110
5.1	Lyman- $\alpha$ forest flux power spectrum emulator errors . . . . .	127
5.2	IGM mean temperature emulator errors. . . . .	128
5.3	Simulation resolution example. . . . .	133
5.4	Sample parameters for LF and HF. . . . .	135
5.5	Simulated Lyman- $\alpha$ forest flux power spectra. . . . .	140
5.6	Simulated IGM mean temperature. . . . .	141
5.7	Leave-one-out errors compared to emulator errors. . . . .	145
5.8	Cosmology parameter posteriors. . . . .	150
5.9	Astrophysics parameter posteriors. . . . .	154
5.10	Observed Lyman- $\alpha$ forest flux power spectrum, and best fits. . . . .	157
5.11	Observed IGM mean temperature, and best fits. . . . .	159
5.12	Full parameter posteriors. . . . .	165
5.13	Post processing parameter posteriors. . . . .	167
5.14	Framework check using simulation as data. . . . .	169
5.15	Posteriors with parameter priors. . . . .	171
5.16	Previous observational data set. . . . .	173
5.17	Increasing observational uncertainties. . . . .	175
5.18	IGM mean temperature emulator posteriors. . . . .	177
5.19	Posteriors for scale restricted analysis. . . . .	178

# List of Tables

2.1	Cosmic String Simulation Parameters . . . . .	20
3.1	Transfer Function Simulations . . . . .	51
4.1	Multi-Fidelity Emulator Simulations . . . . .	87
5.1	Inference Simulations . . . . .	131
5.2	Maximum Posteriors for Astrophysics and Cosmology . . . . .	153
5.3	Derived Parameters . . . . .	153
5.4	Log-likelihoods for best predictions . . . . .	158
5.5	Maximum Posteriors for Post-processing Parameters . . . . .	166

# Chapter 1

## Introduction

### 1.1 Cosmic History

In order to discuss cosmic structure, a brief history of the Universe leading to the modern era is required. However, rather than list the various events in the history of the Universe chronologically, we start with what we observe now, and look back.

The Universe we observe today contains a plethora of different objects, including stars in galaxies, which house central supermassive black holes and are themselves housed in dark matter halos. These galaxies reside in the intergalactic medium (IGM), the low density region between galaxies that is mostly made up of ionized hydrogen and helium gas. Zooming out further, we observe galaxies clustering into what is often called the cosmic web (see Figure 1.1), with large nearly empty voids delineated by filamentary structures that intertwine to form high density knots. This large scale structure is embedded in a spacetime which is not only expanding, but accelerating in its expansion. Even with the briefest of summaries, questions abound.

It turns out the first event to consider, as we look back, relates to the IGM, and how the hydrogen and helium became ionized. While high redshift observed quasar absorption spectra have clear absorption troughs (called Gunn-Peterson troughs [1]), lower redshift spectra do not. This indicates that at lower redshift, the hydrogen is ionized and does not absorb the light from the quasar, while at higher redshift the hydrogen is neutral and does absorb the quasar light. Hence, the IGM was not always ionized, and it must have passed through a phase transition, where it went from neutral to ionized.

Between redshifts  $z \approx 3-4$ , the second electron of helium is ionized by high energy sources, such as quasars [2, 3, 4]. Before this, between redshifts  $z \approx 5.5 - 20$ , hydrogen is ionized, with the most likely sources being stars in galaxies [5]. These events, called helium and hydrogen reionization, also lead to heating, as energy is injected by the ionizing photons. This heating, combined with cooling from adiabatic expansion and atomic processes, are the primary influences on the thermal state of the IGM.

In order to ionize the IGM, sources of ionizing radiation are required. If reionization begins around  $z \approx 20$ , then the first stars and galaxies must begin forming at around the same time. This is preceded by what is often called the dark ages, the time before these first stars and galaxies formed, which lasts from  $z \approx 1100 - 20$ . During this time, the seeds of large scale structure continue to grow as dark matter continues to cluster via gravity. The baryonic matter also continues to cluster, enhanced by the existing dark matter structure.

However, before  $z \approx 1100$ , the Universe is so dense and hot that photons and baryons are in thermal equilibrium. Baryon overdensities do not grow during this time, though dark matter overdensities do, getting a jump start on structure formation. Within

the hot bath of photons and baryons, neutral atoms do not form, as the photons have enough energy and density to immediately ionize any that do form. When the Universe expands and cools enough for the photon energy to drop below the ionization threshold, neutral atoms start to form. This process is called recombination, and is the reason reionization has the re in front.

This is the point of no return for electromagnetic signals, as before this the universe is opaque to photons. Thus, we cannot observe an electromagnetic signal from this time, or earlier. However, a viable observable for this early stage in the history of the universe is primordial gravitational waves, which would travel unaffected by the thermal state of the universe.

Shortly after recombination, the photons decouple from the baryons, as they become increasingly neutral. This leads to the formation of the cosmic microwave background (CMB), or last scattering surface. The CMB has been measured to high precision by several surveys, the latest of which is Planck [6], and found to be well modeled by a black body spectrum at  $T = 2.726$  Kelvin at  $z = 0$ . Imprinted in the temperature fluctuations of the CMB is the pattern of overdensities and underdensities from acoustic waves traveling in the hot dense plasma prior to being frozen in by recombination. These waves set a scale for later structure formation, producing fluctuations in the density of baryons, which is called baryon acoustic oscillations (BAO).

Because the signal stops at this point, we also stop here, except to briefly outline a few more important stages in the history of the Universe. Looking back from recombination, the Universe becomes hotter and denser, leading to nuclei breaking into protons and



neutrons. The process by which nuclei form is called Big Bang Nucleosynthesis (BBN), and it sets the relative abundances of the main atomic constituents of the Universe.

The other important phase, even earlier, is inflation. This is the phase during which the Universe expands exponentially. The process of inflation produces a spectrum of primordial perturbations from quantum fluctuations in the inflation field. These perturbations are carried along with the inflation driven expansion, setting the stage for structure formation later.




Figure 1.1: Dark matter (left) and gas (right, colors are temperature) from a hydrodynamic simulations snapshot. The gas follows the structure of the dark matter, with high density regions hotter than low density regions. In this two-dimensional slice, the cosmic web filaments and knots are visible, and the voids can be inferred, but walls are not distinguishable.

## 1.2 Structure Formation

We begin the discussion of structure formation by laying out the mathematical description of the Universe as it expands in response to its constituent parts. This is the vaunted Friedmann equation:

$$H^2 = \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - c^2\frac{k}{a^2} + \frac{c^2}{3}\Lambda, \quad (1.1)$$

where  $a$  is the scale factor, related to the redshift by  $a = 1/(1+z)$ ,  $\rho$  is the energy density of the components of the universe (matter, radiation, neutrinos),  $k$  is the curvature of the universe (positive  $k$  for a closed universe, zero for flat, and negative for open), and  $\Lambda$  is the cosmological constant. These tell us about the Hubble parameter,  $H$ , which describes the expansion rate of the universe at a given time.

The density of any component can be described using the fluid equation, which conserves mass-energy,

$$\dot{\rho} = -3H(\rho + p/c^2) \quad (1.2)$$

Using this, the evolution of the energy density for each component can be determined. For radiation, the energy density goes as  $\rho_r \propto (1+z)^4$ , and for matter, the energy density goes as  $\rho_M \propto (1+z)^3$ . This means the energy density of the universe is dominated by radiation early on, but transitions to being dominated by matter as the radiation energy is diluted. Eventually, the matter energy density is also diluted to the point where the cosmological constant dominates, as this does not evolve with redshift.

It is common to rewrite the Friedmann equation using the evolution of each component, as well as critical density, which describes the density required for the Universe

to remain perfectly flat,  $\rho_{crit} = 3H^2/(8\pi G) = 2.775 \times 10^{11} h^2 \text{ M}_\odot \text{ Mpc}^{-3}$ . The density parameters can then be written as a fraction of the critical density,  $\Omega = \rho/\rho_{crit}$ . This leads to,

$$H(z) = H_0 \sqrt{\Omega_r(1+z)^4 + \Omega_M(1+z)^3 + \Omega_k(1+z)^2 + \Omega_\Lambda}. \quad (1.3)$$

Structures form via gravitation within the expanding universe described by the Hubble parameter. For dark matter, early density perturbations lead to overdensities which continue to collapse from the end of inflation onward, though only efficiently after matter domination begins. Perturbations in the density of baryons are washed away prior to recombination, due to the pressure of the tightly coupled photon-baryon fluid opposing the collapse due to gravity. However, baryon structures benefit from the early start of dark matter structure, collapsing more rapidly than they would in the absence of these dark matter potentials.

During the time when radiation dominates the energy density budget, which ends before recombination, dark matter perturbations grow logarithmically. When the matter energy density surpasses radiation, the dark matter perturbations have two modes, a growing mode that follows a power law, and a decaying mode. This is the onset of efficient dark matter structure formation, which allows baryons to collapse more rapidly, leading to the aforementioned cosmic web.

The density perturbations which dictate structure formation are described by a power spectrum,

$$P(k, t) = \frac{2\pi^2}{k^3} T^2(k, t) P(k, t = t_0), \quad (1.4)$$

where  $k$  is the wavevector,  $T(k, t)$  is the transfer function, and  $P(k, t = t_0)$  is the primordial

power spectrum. The transfer function encodes all the information on how the density perturbations evolve, thus it depends on the components, e.g. dark matter, baryons, radiation, neutrinos. The primordial power spectrum is most often parameterized as  $P(k, t = t_0) = A_s (k/k_s)^{n_s-1}$ . This form introduces two essential cosmological quantities, the primordial power amplitude  $A_s$ , and the primordial power index,  $n_s$ . The pivot scale,  $k_s$ , is somewhat arbitrary, but is most commonly  $k_s = 0.05 \text{ h/Mpc}$ .

One final useful measure of structure is the mass variance on a specific scale. The scale most often chosen is 8 comoving Mpc, as this is a typical scale for galaxy clusters. The variance is defined as

$$\sigma_8^2(t) = \frac{1}{2\pi^2} \int W_8^2(k) P(k, t) k^2 dk, \quad (1.5)$$

where  $W_8(k)$  is a window function picking out the appropriate scale, and  $\sigma_8$  is most often quoted for  $z = 0$ . Structure growth is linear at first, with the independent modes (scales) evolving independently. Structure growth becomes non-linear as overdensities compound, growing the gravitational potential. For a specific scale,  $R$ ,  $\sigma_R$  increases with time, and structure growth begins to become non-linear when  $\sigma_R$  is close to one.

### 1.3 The Lyman- $\alpha$ Forest

The Lyman- $\alpha$  forest is an observed phenomenon in the absorption spectra of distant quasars, which are supermassive black holes with extremely luminous accretion disks. The spectrum of quasars is close to uniform for a broad range of wavelengths, which makes them ideal for absorption line studies. What is observed is a series of narrow absorption dips in the quasar's spectrum, blueward of the rest frame Lyman- $\alpha$  wavelength, corresponding to

absorption by neutral hydrogen in the intergalactic medium. Specifically, in the rest frame of the neutral hydrogen, photons from the quasar at the Lyman- $\alpha$  wavelength,  $1215.67\text{\AA}$ , are absorbed. This is the first excited state of neutral hydrogen. After the neutral hydrogen is excited, it returns to its unexcited state and re-emits the photon, but now does so isotropically.

Because of the expansion of space, and thus the redshifting of the quasar spectrum, this absorption is repeated as shorter wavelength portions of the spectra become redshifted to the Lyman- $\alpha$  wavelength. When the photons intersect subsequent neutral hydrogen, a different part of the initial quasar spectrum is absorbed, dependent on the redshift. The observed wavelength of absorption is  $\lambda = 1215.67(1 + z)\text{\AA}$ , where  $z$  is the redshift of the absorbing hydrogen. The observable is the flux of photons reaching us from the initial quasar spectrum, where the flux is related to the optical depth by  $F = e^{-\tau}$ . As these features accumulate, the spectrum takes on the appearance of a forest of absorption features.

Figure 1.2 shows two examples of the Lyman- $\alpha$  forest, one from observations (top) and one from simulations (bottom). The observed Lyman- $\alpha$  forest contains other absorption features, from metal lines and Lyman- $\beta$  absorption. In this example, the spectra has been shifted into the rest frame wavelength of the quasar, and it can be seen that the forest begins at wavelengths shorter than the Lyman- $\alpha$  wavelength. The simulated example is the Lyman- $\alpha$  absorption along a skewer through the simulation, i.e. it is the calculated absorption of Lyman- $\alpha$  due to the simulated structure.

Because the Lyman- $\alpha$  forest probes linear and mildly non-linear scales, both astrophysical and cosmological effects need to be accounted for. The optimistic view of this

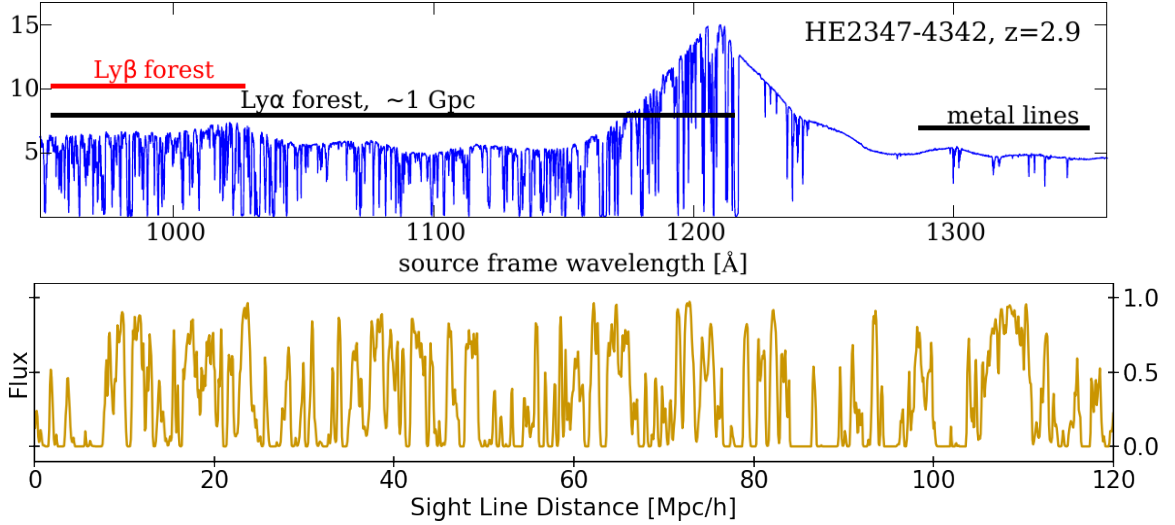


Figure 1.2: The top panel shows an example of the observed Lyman- $\alpha$  forest spectra from [7], and is in flux units of  $10^{-17} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ \AA}^{-1}$  versus wavelength. The observed forest begins at the Lyman- $\alpha$  transition at  $1215.67 \text{ \AA}$  and moves blueward. The bottom panel shows an example of the simulated Lyman- $\alpha$  forest spectra, with normalized flux versus sight line position (within the simulation). The simulated forest only contains the calculated Lyman- $\alpha$  absorption.

is that the Lyman- $\alpha$  forest can simultaneously constrain both, as long as both are modeled reasonably well. The extra cost is that simulations are required to fully exploit observations of the Lyman- $\alpha$  forest, as analytical models do not fully capture the range of scales relevant for the forest. Luckily, simulations have been able to make predictions for the Lyman- $\alpha$  forest for some time [8].

## 1.4 Cosmological Simulations

Simulations begin by initializing a distribution of particles, either as a grid or what is called a glass. The glass method evolves randomly placed particles with reversed gravity. The simulations used in this dissertation use a glass for the baryons, and a grid for the dark matter (then a second glass procedure for the combined baryons and dark matter,

to minimize CDM-baryon overlap). The particles are within a periodic boundary volume, which minimizes boundary effects.

The particles are displaced according to linear perturbation theory to produce an appropriate distribution at the starting redshift (often around  $z \sim 100$ ). The displacement comes from a Gaussian random field generated from the matter power spectrum at the starting redshift. The assumption of a Gaussian random field becomes less important as the initial redshift is increased, thus it is common to start simulations at high redshift. The particles are then evolved via the simulation code, allowing for non-linear structure growth.

Cosmological simulations can be run with only collisionless dark matter, or with both dark matter and gas. Modeling the dark matter is significantly simpler and less computationally costly. The main hurdle for simulations of dark matter, called N-body simulations, is to calculate gravitational forces efficiently, such that a large number of particles can be simulated. The best way to achieve that is to divide the gravitational calculations into short range and long range interactions.

For long range forces, i.e. for the gravitational effect from distant particles, a particle-mesh (PM) method is used. In the PM method, particles are interpolated onto a grid such that a density field can be approximated. The forces can then be calculated using the Poisson equation,

$$\nabla^2\phi = 4\pi G\rho(\mathbf{r}) \tag{1.6}$$

where  $\rho$  is the estimated density field, and  $\phi$  is the gravitational potential. This method is faster than calculating each individual gravitational interaction, but lacks the accuracy necessary for nearby particles, which will dictate the structure on small scales.

For short range forces, a tree algorithm is used. This method hierarchically groups particles based on distance, dividing the simulated space repeatedly until arriving at groups with a single particle. The number of force calculations can then be significantly reduced by walking along the tree, from 'root' to 'leaf.' Starting with the largest groups, a decision is made on whether a multipole approximation is sufficiently accurate. If the accuracy is sufficient, the calculation can be made for that group using a single multipole force. If the accuracy is not sufficient, the walk continues towards smaller groupings, and this is repeated, where the worst case (computationally speaking) is that individual particle interactions must be calculated. Overall, the method for gravity is called Tree-PM, the combination of the short range tree, and long range PM.

Gas, meaning the baryons, needs to be simulated differently than dark matter, requiring a hydrodynamical treatment. Gas is modeled using smoothed particle hydrodynamics (SPH). In SPH, tracer particles sample the underlying fluid, then a smoothed interpolation can be used to fully describe the fluid. The dynamics of the tracer particles come from a description of the gas using a Lagrangian, with kinetic energy and thermal potential energy. Using the Euler-Lagrange equations and thermodynamics, the equations of motion for the gas can be obtained, so long as a good estimate for the density is available. The SPH uses a kernel function to obtain this estimate.

Another requirement for the simulations used in this dissertation is appropriate heating and cooling of the gas, without which the IGM will not be properly modeled. Most often these processes are modeled using known cooling rates for the relevant species, with a uniform ionizing background that dictates the temperature and ionization state of the gas.



Finally, a few notes on simulation terminology. What is often referred to as the box size is the length of one side of the volume simulated, rather than the volume itself. This, as all positions, are in comoving coordinates. The particle load is the number of particles simulated, and is sometimes quoted as the cube root of the actual number. The mass resolution is determined from these two quantities and the relevant component density parameter, e.g. for gas mass resolution  $M_{gas} = \Omega_b \rho_c V / N$  for  $V$  box volume,  $N$  particle number, and critical density  $\rho_c$ . An example of the structure and temperature from a portion of simulations run at lower and higher resolution is shown in Figure 1.3.




Figure 1.3: Slice of the simulated gas (color coded by temperature) from a snapshot at  $z = 4$ , for two different resolution simulations. The bottom shows the lower resolution simulation, while the top shows the higher resolution simulation. In the higher resolution simulation more detailed structure can be seen (i.e. small scale structure).

## 1.5 Chapter Outline

In Chapter 2, the prospects for detecting cosmic strings (theoretical topological defects in spacetime) through their effect on filamentary structure is presented. The scales here are  $\sim 1 - 100$  Mpc, and the time frame is  $z \sim 30$  for the presence of the cosmic

strings, and  $z = 0 - 10$  for the observable effects of the string. This work was previously published [9], and is presented here with an additional figure (Figure 2.10), which serves only to illustrate the proposed detection method.

In Chapter 3, we turn our attention to small and intermediate scales, specifically the intergalactic medium (IGM) and the Lyman- $\alpha$  forest. We are now looking at scales of  $\sim 0.3 - 10$  Mpc, and the time frame is  $z = 2 - 6$ . This chapter details work done to quantify the effect of using separate transfer functions for baryons and dark matter in simulations. Again, this is a previously published work [10].

In Chapter 4, machine learning makes its first appearance. In this chapter, work exploring the use of a multi-fidelity Gaussian process emulator for the Lyman- $\alpha$  forest is presented. Specifically, this is a proof of concept and test of the effectiveness of this method of emulating the Lyman- $\alpha$  forest. This work was previously published [11], and is presented here with two additional figures (Figure 4.9 and Figure 4.4) and accompanying text, which summarizes the emulator accuracy, and provides an example of the Latin hypercube sampling.

Following on from the previous, Chapter 5, presents work using a multi-fidelity emulator to conduct cosmology and astrophysics parameter inference. Using large volume and high resolution simulations, along with observations of the Lyman- $\alpha$  forest flux power spectrum and IGM mean temperature, and a new likelihood code, constraints on various parameters are presented. This work is currently in review, with an expected publication date in 2023.

## Chapter 2

# Cosmic Strings and the Cosmic Web

### 2.1 Abstract

Cosmic strings are generically predicted in many extensions of the Standard Model of particle physics. We propose a new avenue for detecting cosmic strings through their effect on the filamentary structure in the cosmic web. Using cosmological simulations of the density wake from a cosmic string, we examine a variety of filament structure probes. We show that the largest effect of the cosmic string is an overdensity in the filament distribution around the string wake. The signal from the overdensity is stronger at higher redshift, and more robust with a wider field. We analyze the spatial distribution of filaments from a publicly available catalog of filaments built from SDSS galaxies. With existing data, we find no evidence for the presence of a cosmic string wake with string tension parameter  $G\mu$

above  $5 \times 10^{-6}$ . However, we project WFIRST will be able to detect a signal from such a wake at the 99% confidence level at redshift  $z = 2$ , with significantly higher confidence and the possibility of probing lower tensions ( $G\mu \sim 10^{-6}$ ), at  $z = 10$ . The sensitivity of this method is not competitive with constraints derived from the CMB. However, it provides an independent discovery channel at low redshift, which could be a smoking-gun in scenarios where the CMB bound can be weakened.

## 2.2 Introduction

Many well-motivated fundamental physics theories beyond the Standard Model predict the existence of cosmic strings, which are approximately one-dimensional stable relic objects. They may arise in super-string theory as fundamental objects [12, 13, 14, 15, 16] or as vortex-like configurations of quantum fields such as those originated from a  $U(1)$  symmetry breaking [17, 18]. A cosmic string network forms in the very early Universe, e.g., following a symmetry-breaking phase transition, and is expected to consist of stable horizon-length long strings together with sub-horizon loops that can decay away through gravitational radiation or matter emission [19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31]. The phenomenology of cosmic strings are characterized by their energy per unit length (tension)  $\mu$  that relates to the symmetry breaking scale  $\eta$  ( $\mu \sim \eta^2$ ) [32].

Cosmic strings have interested cosmologists and high energy theorists for decades, and can leave detectable signatures in a variety of observational windows. In the early days, strings were considered as potential large scale structure seeds [33, 34]. This role was later ruled out by CMB data. Nevertheless, cosmic strings may still imprint the CMB

as line discontinuities in the temperature map and the current constraint on the tension by Planck [35] is  $G\mu < 1.5 \times 10^{-7}$  at 95% confidence ( $G$ : Newton constant). A cosmic string background may produce detectable gravitational wave signatures that can reveal the expansion history of the early Universe as well as other beyond the Standard Model particle physics, and has thus received increasing attention in light of the recent LIGO detection [36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 31].

Cosmic strings can also be detected through lower redshift astrophysical observations, in particular through their distinct gravitational lensing effect which constrains  $G\mu < 2.3 \times 10^{-6}$  at 95% confidence [47]. Although these low  $z$  lensing observations typically give weaker constraints than those based on the CMB or stochastic gravitational wave background (with model dependence), they serve as important complementary probes. Meanwhile, recent work has brought up the possibility that the CMB and stochastic gravitational wave background limits may be alleviated or evaded in certain scenarios, e.g., if the cosmic strings form before or during early stage of inflation and re-enter back into the horizon only at late times [48, 46]. In such cases low  $z$  astrophysical signatures could be the smoking-gun for cosmic strings.

Early non-linearities are perhaps the most striking effect from cosmic strings in terms of structure formation. In particular, once inside the horizon, long strings straighten out at relativistic speeds [32]. The spacetime metric around these straight segments is conic, with a deficit angle of the size  $\sim 8\pi G\mu$  (i.e. one revolution around the string is less than  $2\pi$  radians) [49]. This causes particles moving relative to the string to be kicked towards the plane traced out by the string, with a magnitude proportional to the deficit angle. A wake

composed of in-falling particles is formed behind the string [50, 51] and the wake grows according to standard linear theory [52]. Such an effect can alter the large scale structure by introducing non-linearities earlier in the universe than would otherwise be possible.

Ultimately, these string induced non-linearities are overwhelmed by the growth of Gaussian fluctuations from inflation, erasing signatures of its existence at later times [53]. Because of this erasure of the early structure signal, the CMB and lensing effects have been the more robust ways in which to constrain the tension of cosmic strings.

In this work, we propose using the effect of the cosmic string on the filaments of the cosmic web as a new probe of cosmic strings. The large scale structure of the universe may be split into distinct components; “zero-dimensional” halos at the intersections of “one-dimensional” filaments, which outline “two-dimensional” walls containing “three-dimensional” voids. These components, listed in descending order of density, make up the cosmic web [54]. The large scale cosmic web has been observed through the distribution of galaxies [55, 56, 57, 58, 59, 60] and intensity maps of emission lines [61, 62, 63]. Future proposed and planned intensity mapping surveys [64, 65] and galaxy surveys such as those conducted by the square kilometre array [66], or dark energy spectroscopic instrument [67] as well as the EUCLID [68] and the Nancy Grace Roman Space Telescope (the Roman Space Telescope, formerly WFIRST) [69] satellites will extend our three-dimensional map of the cosmic web. Further boosting our interest in the large scale cosmic web is the recent development of a multitude of cosmic web identification algorithms [70], some of which have already identified filaments in existing galaxy surveys [71, 72].

We look at the effect that the passage of a long, straight cosmic string has on the filamentary structure in dark matter (DM) simulations. We examine filaments here for two reasons. First, while the wake is effectively a two dimensional feature, when projected onto the sky it would appear one dimensional. Second, filament finders are well developed and may be robustly applied to galaxy catalogues. It may be possible to detect the effects of a cosmic string in other cosmic web structures, such as walls and voids. We consider string tensions ranging from those consistent with Planck limits to a value similar to the gravitational lensing limit. This method is sensitive to higher redshift strings (when the string has passed through the simulation volume and the wake forms,  $z \sim 100$ ) and larger scales ( $\sim 10$  Mpc) than the lensing signal ( $z \lesssim 1$ ,  $\sim$ Mpc).

Refs. [73, 74, 75] showed that tensions of  $G\mu \sim 10^{-7}$  are detectable in cosmological simulations using the cosmic string wake. However, these methods require an accurate map of the three-dimensional dark matter distribution at  $z > 2$ , which is currently out of reach of foreseeable experiments. Another potential, albeit at present futuristic, method for detecting the wake from cosmic strings is 21 cm intensity mapping [76].

The effect of the physical shift and kick due to the string on the statistics of cosmic filaments has not been previously explored. While we show that the effects are generally too small to be competitive with limits from CMB constraints, the novel independent method we propose provides a valuable complementary probe, and would become competitive in scenarios when CMB bounds are weakened (for example, regrowing cosmic strings after inflationary dilution [46]).

In Section 2.3 we detail the suite of simulations that we will be using throughout, as well as the method we employ to add the effects of the cosmic string to the simulation. We also comment on the possibility of detection via the kinematics of the halos, as well the velocity kick feature of the wake implementation. In Section 2.4 we describe the two filament identification algorithms we use, comparing the filaments they identify qualitatively. The result of this section is a catalog of filaments that we analyze in Section 2.5, along with a set of publicly available filaments constructed from SDSS galaxies. We also demonstrate the projected sensitivities for the Roman space telescope in Section 2.5. In Section 2.6 we show that our results are independent of the box size and resolution of our simulations. Finally, in Section 2.7 we conclude, pointing to future prospects for cosmic string wake detection.

## 2.3 Simulations

We performed a suite of dark matter only cosmological simulations using MP-Gadget<sup>1</sup> [77]. MP-Gadget is a fork of Gadget-3 [78] modified for scalability. The initial power spectrum is generated via the Boltzmann code CLASS [79]. Radiation density is included in the background expansion rate, and the simulation box has periodic boundaries. The cosmology parameters are the current defaults of MP-Gadget and are consistent with the nine-year WMAP results [80] (our results are not sensitive to the exact values used). These, and the simulation settings used, can be seen in Table 2.1.

The main results of this work are from a set of five simulations (each with a different starting random seed, leading to distinct dark matter structures) with the parameters in

---

<sup>1</sup><https://github.com/sbird/MP-Gadget3>



Cosmology		Simulation	
$\Omega_0$	0.2814	$N$	$512^3$
$\Omega_\Lambda$	0.7186	$L$ [Mpc/h]	64
$\Omega_b$	0.0464	$\ell$ [Mpc/h]	250
$H_0$ [km/s/Mpc]	69.7	$z_0$	99
$\sigma_8$	0.810	$z_{cs}$	31
$n_s$	0.971		

Table 2.1: Cosmic String Simulation Parameters. Cosmology parameters are: total matter  $\Omega_0$ , dark energy  $\Omega_\Lambda$ , and baryon  $\Omega_b$  densities. Hubble constant  $H_0$ , density fluctuation  $\sigma_8$ , scalar spectral index  $n_s$ . Simulation parameters are: number of DM particles  $N$ , box length for main simulations  $L$ , box length for scaling simulation  $\ell$  (see Section 2.6), initial redshift  $z_0$ , wake insertion redshift  $z_{cs}$ .

Table 2.1 and a mass resolution of  $1.5 \times 10^8 M_\odot$ . The effects of the string wake were inserted at  $z = 31$ . The insertion redshift is chosen to be late enough that a string will have time to travel the box length since matter-radiation equality and early enough that non-linear structure has not yet formed in the wake. The insertion time does not affect our results as long as it is early enough that linear perturbation theory (Equations 2.2 & 2.3) is a good description. This will be until structure formation dominates the wake signal [53]. An additional simulation with a larger box size (250 Mpc) was run to ensure that our results were not strongly affected by box size or resolution (see Section 2.6).

### 2.3.1 Cosmic String Wakes

Cosmic strings create a deficit angle in the spacetime surrounding them. This means that as a cosmic string passes through matter, the trajectory of that matter is altered (for more details, see [49] or [32]). Specifically, the matter is “kicked” towards the worldsheet of the string with an additional velocity given by

$$\delta v = 4\pi G\mu\gamma u, \tag{2.1}$$

where  $G\mu$  is the dimensionless parameter defined by the tension of the cosmic string,  $u$  is the comoving speed of the cosmic string, and  $\gamma$  is the usual relativistic Lorentz factor introduced by transforming to the reference frame of the particle [32]. See Figure 2.1 for a cartoon of the cosmic string wake, deficit angle, and wake. We set the string speed as  $u = 0.6c$ , consistent with e.g., [37] (Eqs. 2.3 & 2.4 therein). The string speed only affects the magnitude of the velocity kick and displacement of particles on wake insertion. Therefore, the exact value is sub-dominant compared to the tension parameter, which in this study varies by a factor of 50, as compared with possible string speeds varying from the value used here on the order of a factor of  $\sim 1 - 5$  [32].

These in-falling particles create an overdensity along the worldsheet of the cosmic string. This overdensity, or cosmic string “wake”, is an early non-linear structure which perturbs the position and velocity of nearby particles. Instead of simulating the cosmic string directly, the wake is simulated following the method used in [75]. We simulate the wake once the string has passed by including the linear density perturbation induced by the string’s passage. For simulation purposes, a snapshot that is sufficiently late that the string will have passed through the box is chosen, and the particles are displaced and kicked at that time (by an amount corresponding to the prediction of linear theory between the time of the string’s passage and the time of wake insertion).

Perturbations due to the string wakes are inserted as

$$\delta x = \frac{3}{5} \delta v t_{eq} \frac{(z_{eq} + 1)^2}{(z + 1)} h, \quad (2.2)$$

$$\delta \dot{x} = \frac{2}{5} \delta v \frac{t_{eq}}{t} \frac{(z_{eq} + 1)^2}{(z + 1)^2}, \quad (2.3)$$

where  $t_{eq}, z_{eq}$  are the time and redshift of matter-radiation equality (and when the wake is

first formed),  $h$  is the Hubble factor, and  $t, z$  are the time and redshift at which the shift and kick are calculated (i.e. when the wake effect is inserted into the simulation). Note that our perturbations differ slightly from [75]: the factor of  $h$  in the shift is due to positions in MP-Gadget being in comoving kpc/h, one factor of  $z^{-1}$  in the velocity kick is due to MP-Gadget using physical peculiar velocities, and we have dropped the (good) approximation that  $z \approx z + 1 = 1/a$ .

In practice, the snapshot at which the wake's effects are being inserted is adjusted by first shifting the particle positions along the coordinate perpendicular to the plane of the string wake. The particles on either side are shifted towards the plane, then the velocity perturbation is added towards the plane of the wake. The simulation is then run to completion from the updated snapshot, which now includes the effects of the wake. Given that we expect  $\sim 10$  horizon-length strings per horizon [81, 82], we would not expect to see multiple string wakes in a region the size of our simulation boxes, therefore only one wake is inserted in each simulation.

The wake insertion is performed in the same way for each of the five distinct simulated cosmic structures. For each of these distinct structures, four different scenarios are simulated: one scenario where there is no wake, and the remaining scenarios where a wake is inserted with tensions  $G\mu = 10^{-7}, 10^{-6}$ , and  $5 \times 10^{-6}$ . An example can be seen in Figure 2.2, which shows the density of DM particles at two redshifts for the case with no wake (left) and with a  $G\mu = 4 \times 10^{-6}$  wake (right).




Figure 2.1: *Setup for cosmic string  $\mathcal{E}$  wake.* The cosmic string (bold, black) creates a deficit angle,  $8\pi G\mu$ , in the space around it. The string travels with speed  $u$ , leaving behind a wedge shaped overdensity called a wake with opening angle,  $8\pi G\mu\gamma$ , where  $\gamma = (1 - u^2)^{-1/2}$ .

### 2.3.2 Velocity

The effect of the cosmic string passage is a velocity kick towards the worldsheet of the string. This leads to an effect on the velocity distribution of the halos. At early times there is an excess velocity towards the wake. At some later point, depending on the string tension, the kicked halos cross the wake and the excess velocity is away from it. While this is a significant noticeable effect on the redshift space distribution of the halos, it depends strongly on the orientation between the observer and wake. Our simulations indicate that while even the lowest tension string wake would have a noticeable signal down to  $z = 2$ , the signal disappears when the angle between observer and the plane perpendicular to the wake is greater than  $\sim 10$  degrees at higher redshift ( $z \geq 5$ ) and  $\sim 3$  degrees at lower ( $z \leq 2$ ). A detection via this signal would rely on a highly fortuitous alignment.




Figure 2.2: *Visual wake versus no-wake comparison.* Projection of the particle density at  $z = 31$  (top) and  $z = 1$  (bottom). Left-hand panels are from a simulation with no wake, right-hand panels are from a simulation with a  $G\mu = 4 \times 10^{-6}$  wake inserted at  $z = 31$ .

These signals depend on the ratio of the average speed of particles towards the wake prior to wake insertion, and the velocity kick they receive. Figure 2.3 shows the average speed perpendicular to the wake against redshift. Also shown is the  $1\sigma$  region of these speeds (shaded). The velocity kicks from wake insertion are shown as dashed lines. The middle tension case traces out the particle speed fairly well, while the other two tensions trace out the edges of the  $1\sigma$  region. As might be expected, the observability of the effect

from the passage of a cosmic string seems to depend on how large a kick the particles receive compared to the intrinsic velocity dispersion from standard DM structure formation (i.e. a non-trivial number of halos change direction due to the wake).




Figure 2.3: *Average particle speed perpendicular to wake.* Average speed of DM particles in the direction perpendicular to the cosmic string wake. Dashed lines show the kick imparted on the particles (see Section 2.3.1) for the three different tensions used throughout.

## 2.4 Filament Identification

While the cosmic string wake is inserted via a position and velocity perturbation on the DM *particles*, we identify the filamentary structure from the resulting DM *halos* (these are constructed using a friends-of-friends algorithm in the simulation code). This is much more computationally reasonable, and agrees more closely with how filamentary structure is currently identified observationally (i.e. via galaxy surveys).

Many algorithms have been devised to identify cosmic web components within cosmological simulations or from galaxy catalogues (see [70] for a comparison of 12 algorithms). Here, we use two such algorithms; the subspace constrained mean shift (SCMS) algorithm [83, 84], and DisPerSE (Discrete Persistent Structures Extractor) [85, 86]. In both cases we identify only linear structures, the filaments of the cosmic web. By using two methods we mitigate the dependence of our results on the identification algorithm used, or options/parameters selected within each method. The two selected methods are at somewhat opposite ends in terms of complexity (with SCMS being fairly intuitive), and implementation (DisPerSE is a downloadable installation).

Below we briefly describe each algorithm and note our parameter/option choices in using them, as well as any necessary alterations made to conform to our simulation output (the most notable being the periodicity of our simulation box).

#### 2.4.1 SCMS+

SCMS is a gradient ascent method, which shifts tracer particles according to the Hessian (second derivative, or curvature) matrix of the density at each tracer position. The density is approximated with a Gaussian kernel density estimate (KDE) at each tracer position according to the population of dark matter halos in the simulation box. We follow the implementation as described in [84], which thresholds out halos in low density regions, then runs the SCMS algorithm to identify ridges in the KDE. Here, we use the DM halos as the starting positions of the tracers and set the smoothing length as 2 Mpc. Some modifications were implemented to correctly handle the periodicity of the simulation box.




Figure 2.4: *SCMS+ filaments*. (top left) converged SCMS tracers (black) outlining the spine of the filaments, halos-in-filaments (green), and halos too far from the spine (red). (top right) a split filament. (bottom right) a filament straddling the periodic boundary of the box. (bottom left) a “wiggly” filament.

A brief outline of the algorithm follows.

- (Thresholding) Tracers in low density regions are removed by calculating the KDE for all tracers and removing those with values lower than the mean KDE, as these are unlikely to reside in filaments ([70] figures 4 & 5). Particle separations account for the box periodicity, connecting particles using the smallest possible great circle arc.



- (SCMS) Using the subset left over from the thresholding step, the SCMS algorithm runs until a tracer is shifted by less than 1 kpc. For each iteration the Hessian matrix is calculated and the smallest two eigenvectors dictate which direction the halo moves to ascend the local ridge. Once a tracer is converged the next tracer is shifted until convergence, and so on (the KDE calculated here is from the DM halos, which are not shifted). To accommodate the periodic boundaries, the box (and the halos within) is centered on the current tracer position, and the separations are calculated between this tracer and tracer-centered halos. At the end of each step, if the tracer has shifted outside the box, it is moved back to the correct, periodic position by adding or subtracting the box size.

The output of the SCMS algorithm is a set of tracers converged on the ridges of the density field, and is not a set of filaments. To obtain separated filaments (and filament properties) we use both the tracers and the DM halos. First, the tracers are connected into segments if they are separated by less than a maximum distance, with a value that depends on the redshift. The segments are then connected iteratively using the same, redshift dependent maximum distance (from 250 kpc/h at  $z = 0$  to 1200 kpc/h at  $z = 10$ ) until a set of separated filaments is achieved. We call the entire method, from thresholding to filament separation, SCMS+. Some examples of filaments identified with this algorithm can be seen in Figure 2.4, which for a simulation with no wake inserted, shows the entire box in the top left corner (converged tracers in black, halos belonging to filaments in green, halos too far from filaments in red), then some zoomed in examples of filaments in the other panels.

The length of the filaments are determined by a smoothed stepping procedure, which sums the distances between neighboring tracers, which have been smoothed by 0.5 Mpc, starting from one end of each filament (a filament end is determined by identifying the tracer with the fewest directions occupied by another tracer). DM halos are then assigned to filaments based on proximity and a cutoff of 2 Mpc. Filaments that are shorter than the smoothing length (2 Mpc) or have fewer than 3 halos assigned to them are discarded. Finally, the mass of each filament is the sum of the halos assigned to that filament.

#### 2.4.2 DisPerSE

DisPerSE is a widely used cosmic web identification tool which extracts structure using the Morse-Smale complex of the input. The input is the Delaunay tessellation of the particle distribution. For details about DisPerSE, see the overview and tutorial at <http://www2.iap.fr/users/sousbie/web/html/indexd41d.html>. We implement the main program of DisPerSE, MSE, on the tessellated DM halo population from our simulations with a  $6\sigma$  persistence threshold. The output is converted using the skelconv program in DisPerSE, smoothed over 10 halos, and filaments are assembled if the angle between them is less than  $75^\circ$ .

DisPerSE can find filaments directly from DM particles<sup>3</sup>. To check the robustness of our results we subsample 5% of the DM particles from the simulations and run DisPerSE using the same settings, except for increasing the smoothing to 100 particles. We visually inspected the output and confirmed that approximately three-quarters of the filaments found

---

<sup>3</sup>In principle, so can SCMS, but it scales poorly to large tracer numbers.

from halos had approximately the same positions as the filaments found from particles. For ease of comparison with SCMS+ filaments, we use DM halos for the remainder of this work.

While the two filament identification algorithms employ fairly different methods, they both identify similar major filaments. DisPerSE detects substantially more filamentary structure as can be seen in Figure 2.5, which shows the filaments identified from a simulation




Figure 2.5: *SCMS+ filaments*. SCMS+ filaments (blue) and DisPerSE (brown) filaments, both using DM halos as the input. The top panel is at  $z = 10$ , while the bottom is at  $z = 0$ . While DisPerSE is clearly more sensitive, finding more filaments, the largest filaments are similar.




Figure 2.6: *Filament number comparison (SCMS+, DisPerSE)*. Number of filaments identified by the two methods as a function of redshift. Included are the filaments from the simulation without a wake, and a simulation with a  $G\mu = 5 \times 10^{-6}$  wake.

with no wake. Figure 2.6 compares the number of filaments identified by each method as the simulation evolves. Both algorithms detect the growth of structure up to the onset of dark energy domination and a subsequent decline in filament number. However, the two methods differ in filament number and redshift trend: DisPerSE initially has fewer filaments but much larger growth with redshift, which peaks at a higher level.

Both the number of filaments and the trend with redshift depend strongly on the values chosen for key parameters for each method; for DisPerSE the persistence threshold, and for SCMS+ the smoothing length. For example, reducing the persistence threshold of DisPerSE from  $6\sigma$  to  $2\sigma$  leads to a  $\sim 10$ -fold increase in filaments at all redshifts and moves the peak redshift to  $z = 2$ , preserving the late decline in filament number. Despite the differences in the filament populations these methods return, neither method shows a convincing signature of a cosmic string wake on structure formation. We also stress

that in the analysis presented, we are concerned with relative changes, comparing similar simulations where the only difference is the inclusion of a cosmic string wake.

For each cosmic string tension we have performed five different simulations using different initial realisations of cosmic structure. For each string tension, the filament catalogs from these five different structure realisations are combined, boosting the sample of available filaments. The final results are two catalogs (SCMS+, DisPerSE) with four filament populations each<sup>4</sup>. In Section 2.5 we analyze these populations for potential cosmic string signals, focusing on the SCMS+ catalog.

## 2.5 Results

Our goal is to determine if there are signals in the filament population indicating the previous passage of a cosmic string. With a catalog of filaments in hand, we compare the filament population and spatial distribution of filaments in simulations with and without cosmic string wakes.

### 2.5.1 Filament Population

Figure 2.7 shows the number of filaments for each of the four wake scenarios, plotted against redshift for the results from the SCMS+ filaments. Similar results are obtained with DisPerSE. The change in the number of filaments between the simulation without a wake and the three with wakes is small and inconsistent. It appears that the presence of the wake both separates and connects filament segments, in approximately equal amounts. This leads to little to no change in the number of filaments.

---

<sup>4</sup>(python) code used to produce these catalogs are available at <https://github.com/mafern/SCMSplus>.

The properties of the filament populations, for example the distribution of filament masses, are another potential signal. The procedure for determining the filament length and mass are outlined at the end of Section 2.4.1. As can be seen in Figure 2.8, which shows the distribution of filament masses for the SCMS+ filaments, the cosmic string wake has very little effect on the filament masses. The distribution of filament lengths is similarly devoid of a distinguishing signal. Visual inspection of the filament distribution reveals, however, that the cosmic string wake does alter the spatial distribution of filaments, enlarging some and disrupting others. Once averaged over the whole box, the overall distribution is unchanged. Filaments are re-ordered in a way which is indistinguishable from the random variance due to the realization of structure formation. One feature of structure resulting from the wake which distinguishes it from structure formation seeded from inflationary perturbations is that it has a preferred direction (i.e. towards the wake). This is the feature that we exploit in the following section.

### 2.5.2 Spatial Distribution of Filaments

Now we illustrate more sensitive observables for identifying filaments originated from string wakes. The effect of a cosmic string is to pull matter towards the worldsheet of its passage, motivating a look at the spatial distribution of filaments. In principle the spatial distribution of two-dimensional walls and three-dimensional voids will be affected similarly, however we remain focused on filaments here. We do this in two ways, the first being a comparison of the filament number density around the center of the wake to the overall filament number density. This is shown in Figure 2.9 for the SCMS+ catalog (we




Figure 2.7: *Filament number, increasing tension (SCMS+)*. Number of filaments identified by SCMS+ versus redshift. Even at very large tension, the effect on the number of filaments is small and noise dominated.

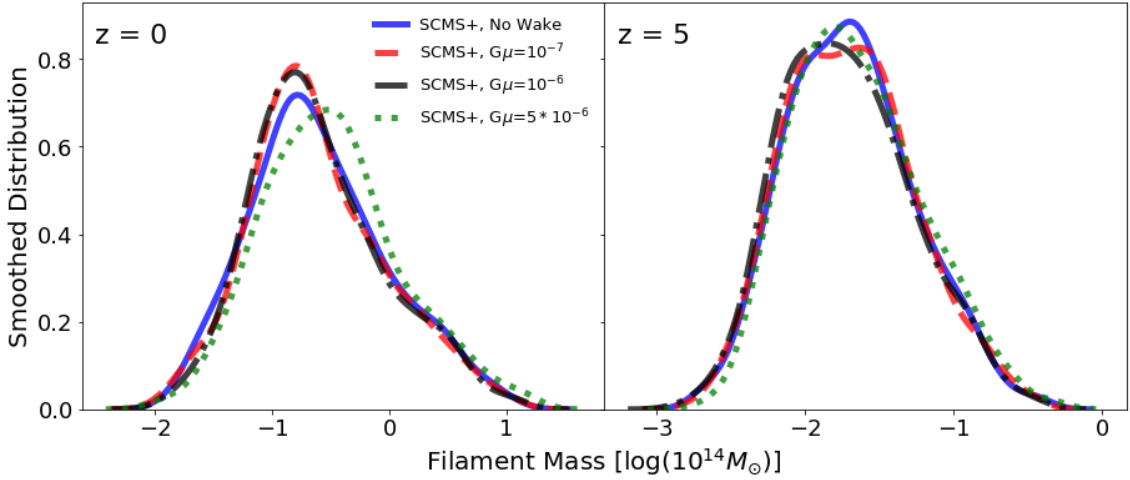


Figure 2.8: *SCMS+ filament mass distributions*. Smoothed, normalized distribution of filament masses at  $z = 0$  (left) and  $z = 5$  (right), for the three wake and no-wake cases. There are no significant changes in the masses of the filament population early or late in the simulation.

obtained similar results for the DisPerSE filaments). Any part of a filament within the central 10 Mpc/h is added to the number of filaments near the wake. The number density in this central region, surrounding the wake, is divided by the filament number density for the whole box. As a baseline the same calculation is repeated for simulations with no wake, comparing the central region to the rest of the box.

The smallest tension wake we include here ( $G\mu = 10^{-7}$ ) is not distinguishable from the simulation without a wake. The intermediate tension wake ( $G\mu = 10^{-6}$ ) shows signs at higher redshift ( $z > 5$ ) of a central overdensity in the number of filaments. The largest tension wake we include ( $G\mu = 5 \times 10^{-6}$ ) shows a clear overdensity in the central region at all redshifts, with the signal generally decreasing with time, with the exception of an uptick after redshift 2. This is likely due to the filaments coalescing into larger structures, which already have segments within the central region, see Figure 2.7). We do not see this trend in the lower tension and no-wake simulations because in these cases the filaments have not been significantly distorted/pulled towards the central wake.

In an observational setting, we would not know a priori where the wake is, and therefore the previous comparison would be of limited use. Rather than looking only at a slice of the box parallel to and around the wake center, we can look at slices centered from one side of the box to the other (see Figure 2.10). The number of filaments in each 10 Mpc/h slice can then be compared to the mean number of filaments in a slice, i.e. each slice filament number is divided by the mean slice filament number for the entire box. A value higher than one is overdense with respect to the mean slice. Comparing the overdensity in these slices highlights any clustering of filaments around a potential wake.






Figure 2.9: *Central filament overdensity, increasing tension.* Overdensity of filaments identified by SCMS+ versus redshift. The overdensity becomes more pronounced as the tension of the string is increased.




Figure 2.10: Proposed method for detecting cosmic strings by the overdensity of filaments. A priori, the location of the cosmic string wake would not be known (left panel), so searching for a filament overdensity would require scanning across the data set (right panel).

In Figure 2.11 the *slice overdensity* is apparent in the center of the box at all redshifts for the largest tension wake, while the smaller tension wakes show no robust signal at any redshift. The signal from the DisPerSE filaments is marginally smaller, but still shows that only the largest tension wake produces a noticeable overdensity. The underdensity at the edges of each panel is due to the finite size of the box. During wake insertion particles are kicked towards the center of the wake, away from the edges [75].

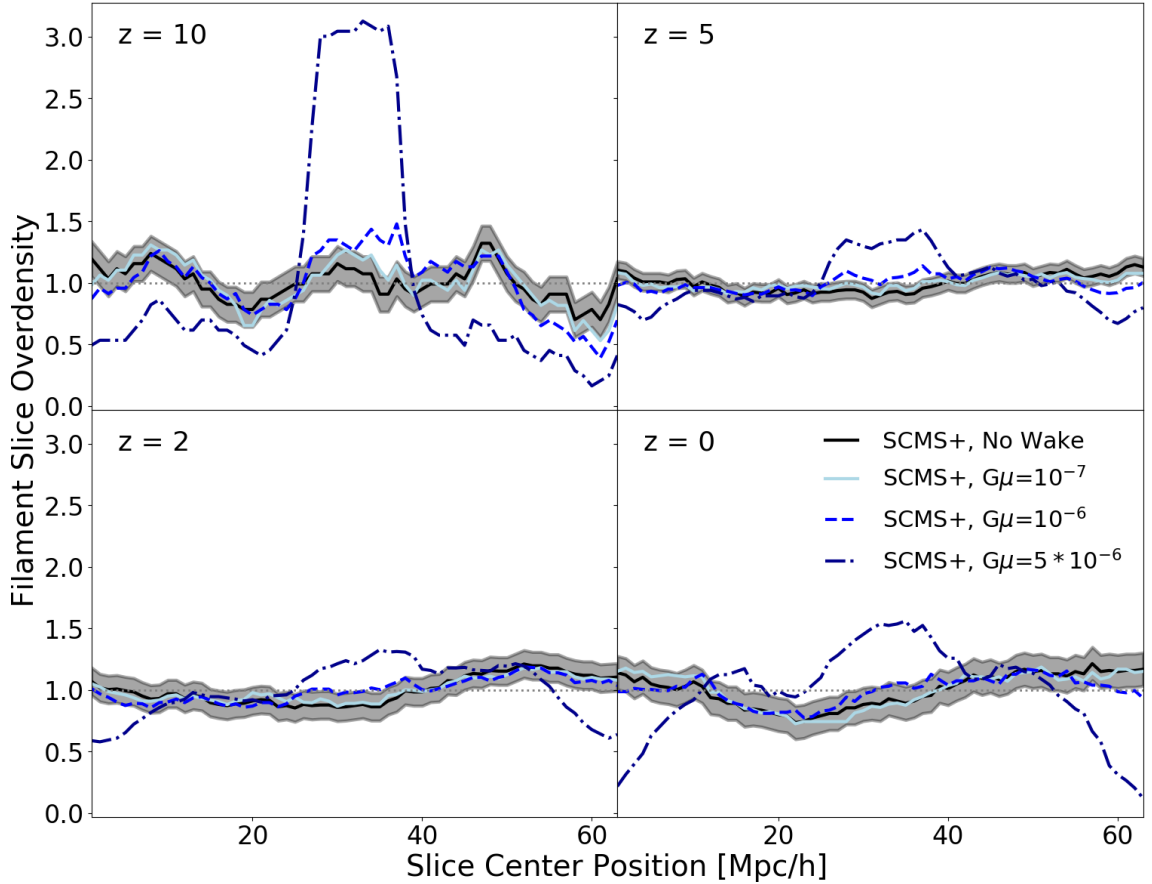


Figure 2.11: *Filament slice overdensity (SCMS+)*. Overdensity of filaments identified by the SCMS+ method in 10 Mpc/h slices versus the center of the slice. The grey shaded region shows the  $1\sigma$  area for the no wake case.

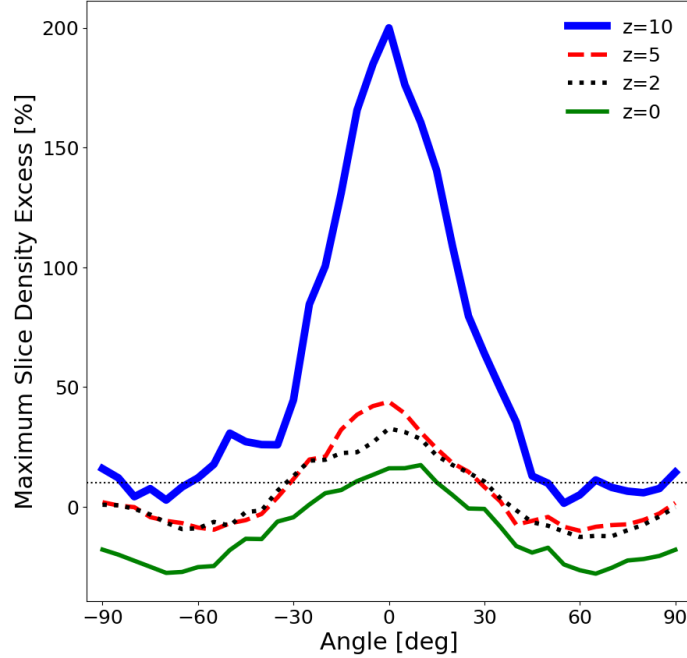


Figure 2.12: *Maximum number density excess by viewing angle.* The maximum percent difference between the number of filaments in a slice for the simulation with a  $G\mu = 5 \times 10^{-6}$  tension wake to the  $1\sigma$  region in the no-wake simulation (see grey shaded region in Figure 2.11). The dotted line marks a 10% excess.

The observability of this signal depends on the orientation of the wake with respect to the observer, i.e. how much of the wake the observer is looking through. By rotating our simulated filaments about an axis parallel to the plane of the wake we can simulate viewing the wake from edge-on (maximally visible) to face-on (minimally visible). After rotating the filaments, the box is re-formed to account for the periodic boundaries. Using the rotated filaments we can determine the viewing angles at which the signal persists (has at least one slice with a 10% density excess over the no wake  $1\sigma$  region). The average of the two possible rotations is shown in Figure 2.12, which indicates that for the highest tension string wake the signal in the SCMS+ filaments persists up to an angle of  $\sim 65$  degrees at  $z = 10$ ,  $\sim 30$  degrees at  $z = 5$ ,  $\sim 30$  degrees at  $z = 2$ , and  $\sim 15$  degrees  $z = 0$ . This is




Figure 2.13: *Filament slice overdensity (SDSS)*. Overdensity of SDSS filaments constructed using the filament catalog from [87]. Slices span 10 degrees in RA and are plotted over the RA slice center. Shown are slice overdensities using filaments at all redshifts (black), the  $1\sigma$  and  $2\sigma$  regions (dark and light shaded), and each  $\Delta z = 0.1$  bin from  $z = 0.005$  to  $z = 0.7$  (brown).

in contrast to the visibility of the overdensity of halos in the wake, which is robust to  $\sim 8$  degrees at higher redshift and  $\sim 4$  degrees at lower redshift.

### 2.5.3 SDSS Filament Catalog

We use the publicly available filament catalog from [87] to construct observed slice overdensities to compare to our simulated result (Figure 2.11). The main takeaway of this section is to show that this can be done in a relatively straight-forward way for any catalog of filaments.

The catalog from [87] identified filaments in a combined DR7/DR12 SDSS dataset using SCMS. The filament catalog is a set of points reported in right ascension (RA),

declination, and redshift in  $\Delta z = 0.005$  bins from  $z = 0.005$  to  $z = 0.7$ . The smoothing length used in the SCMS method for each redshift bin is also available. For filaments in each redshift bin, we convert the coordinates to Cartesian [88, 89], then separate the filaments using the method outlined in Section 2.4.1, using the reported smoothing lengths as a guide to the separation length. We then count the filaments in 10 degree slices of RA for each redshift bin.

Figure 2.13 shows the slice overdensities, grouped into  $\Delta z = 0.1$  bins (brown), as well as the overdensities for the entire set of filaments (black). The average of the whole filament set we treat as the no-wake case since cosmic string wakes will only affect a fraction of this volume. Also, shown are the  $1\sigma$  and  $2\sigma$  deviations from the full set of filaments. We compare the whole filament set to each individual slice in  $\Delta z = 0.1$  to determine whether an overdensity similar to what we see in Figure 2.11 appears in the observed filament population. Figure 2.13 shows that almost all slices fall within the  $1\sigma$  range around the total filament density using all SDSS galaxies. This suggests that a cosmic string wake formed with the largest tension,  $G\mu = 5 \times 10^{-6}$ , is unlikely except for a small region of parameter space where the wake is parallel to the line of sight.

#### 2.5.4 Roman Space Telescope Projection

Future surveys, such as the Roman Space Telescope [90] will provide high redshift ( $z > 2$ ) galaxy populations, which can be used to identify high redshift filament populations. For example, at  $z = 2$ , the Roman Space Telescope is expected to observe approximately 1000 galaxies per square degree per  $\Delta z$  over its  $\sim 2000$  square degree field of view [69], leading to  $\sim 2 \times 10^5$  galaxies (with  $\Delta z = 0.1$ ). Our simulations indicate a median of  $\sim 100$

halos per filament at that redshift, leading to  $\sim 2000$  filaments around  $z = 2$ . From Figure 2.11, it appears that there is a  $\sim 20\%$  enhancement (i.e. the maximum ratio between a simulation with a wake to the case without a wake) in the number of filaments around the cosmic string wake for the  $G\mu = 5 \times 10^{-6}$  case. Taking into account the angles at which the effect is still visible, and assuming all orientations are equally likely, the Roman Space Telescope should be able to detect this signal at the level of  $0.2/N_{fil}^{-1/2} \times 60/180 \sim 3\sigma$ , or 99% confidence.

We can repeat the same approximation at  $z = 10$ . At this high redshift the Roman Space Telescope, in the optimistic case, projects finding  $\sim 1000$  galaxies with the High-Latitude Survey [69]. Our simulations indicate at  $z = 10$  there are  $\sim 10$  galaxies per filament, so we could identify as many as  $\sim 100$  filaments at this redshift. For a tension of  $G\mu = 5 \times 10^{-6}$ , our simulations show an enhancement of a factor of  $\sim 3$ , with a signal that persists up to an observation angle of  $\sim 65$  degrees. This would be detectable by the Roman Space Telescope at  $z = 10$  at  $\sim 22\sigma$ . By contrast, a tension of  $G\mu = 10^{-6}$  shows an enhancement of  $\sim 0.1$  at  $\sim 25$  degrees. It would thus enhance the filament overdensity by only  $0.3\sigma$  and not be detectable.

Figure 2.11 shows that the signal at  $G\mu = 5 \times 10^{-6}$  increases at higher redshift, as the signal becomes more pronounced with respect to the filament overdensity induced by normal structure formation. However, for lower tensions with  $G\mu \leq 10^{-6}$  the increase in filament overdensity from the cosmic string wake is always less than the intrinsic variance that structure formation imparts in the spatial filament overdensity. Lower tensions thus produce a sharply reduced signal and so a floor in the cosmic string tension detectable with

filament finding methods. This floor is at  $G\mu \approx 10^{-6}$ , approximately where the relative magnitude of the cosmic string kick at string passage equals the mean particle velocity from structure formation (see Section 2.3.2).

## 2.6 Box Size Scaling

In this section we check that our results are not greatly affected by the box size or resolution. We compare one of the main simulations (64 Mpc) to a larger box size (250 Mpc), lower resolution simulation. Specifically, we look at the SCMS+ filaments from these two simulations, in the case where we have no wake and case where we have the largest tension wake used ( $G\mu = 5 \times 10^{-6}$ ). We look only at a single run of the simulation (as opposed to other figures, which are the combination of five runs with different structure seeds). The single simulation means greater variance in the signal, however it is clear from Figure 2.14 that while the increased box size has increased the number of filaments, it has not changed the lack of a signal between the no-wake and large-wake simulations (this is true for the DisPerSE filaments as well).

## 2.7 Conclusion

We have run a suite of simulations to demonstrate the effect of a cosmic string wake on structure formation, in particular the distribution of cosmic filaments. While simple measures, like the number of filaments, do not distinguish the presence of a cosmic




Figure 2.14: *Box size & resolution check (SCMS+)*. Fraction of maximum filaments versus redshift for simulations with different box sizes (64 & 250 Mpc), but the same structure seed and number of DM particles.

string wake in simulations, the spatial distribution of these filaments can. These potential signals are not competitive with constraints derived from the CMB, however they provide an independent constraint relevant at low redshift. They may thus be a viable channel for detection in future experiments or for constraining models that have a stronger signal at relatively low redshift ( $z \sim 5$ ). The most promising signals presented here are comparisons between the filament number density far from and near the cosmic string wake, especially across the plane perpendicular to the wake. Examples of this type of comparison are shown in Figures 2.9, 2.11. We found that the overdensity signal persists for angles between the wake and the observer of up to  $\sim 30$  degrees from  $z = 5$  down to  $z = 2$ .

The spatial distribution of filaments can be used to detect the imprint of a cosmic string wake at tensions of  $G\mu \gtrsim 5 \times 10^{-6}$  and redshifts  $z \gtrsim 2$ . Lower tensions of  $G\mu \lesssim 10^{-6}$



produce a smaller signal and are not detectable at any redshift from galaxy filaments. We found that filaments constructed from current galaxy catalogs (SDSS) show no sign of the overdensities we see in simulations in which a large string tension wake has been included. However, improved sensitivity is expected with new data from galaxy survey experiments in the coming years, such as the Roman Space Telescope [90] or EUCLID [91]. These surveys will easily detect filaments at  $z \sim 2$ , and present opportunities to probe cosmic strings with tensions between  $G\mu = 10^{-6}$  and  $G\mu = 5 \times 10^{-6}$ . If filament finders can be successfully used at higher redshifts, up to  $z \sim 10$ , very high significance detection may be achieved.

## Chapter 3

# Transfer Functions in Simulations

### 3.1 Abstract

Using a set of high resolution simulations, we quantify the effect of species specific initial transfer functions on probes of the IGM via the Lyman- $\alpha$  forest. We focus on redshifts 2 – 6, after H I reionization. We explore the effect of these initial conditions on measures of the thermal state of the low density IGM: the curvature, Doppler width cutoff, and Doppler width distribution. We also examine the matter and flux power spectrum, and potential consequences for constraints on warm dark matter models. We find that the curvature statistic is at most affected at the  $\approx 2\%$  level at  $z = 6$ . The Doppler width cutoff parameters are affected by  $\approx 5\%$  for the intercept, and  $\approx 8\%$  for the fit slope, though this is subdominant to sample variation. The Doppler width distribution shows a  $\approx 30\%$  effect at  $z = 3$ , however the distribution is not fully converged with simulation box size and resolution. The flux power spectrum is at most affected by  $\approx 5\%$  at high redshift and small scales. We discuss numerical convergence with simulation parameters.

## 3.2 Introduction

The intergalactic medium (IGM) occupies the space between galaxies and galaxy clusters, and houses the majority of baryonic matter in the universe. The major phase changes in the history of the IGM are fairly well understood, with recombination ( $z \sim 1100$ ) leading to the formation of a highly neutral IGM, and H I ( $z \sim 5.5 - 8$ ) [92, 93, 94, 95, 96] and He II ( $z \sim 3$ ) [2, 97, 98, 99, 100, 101] reionization events leading to the current, highly ionized IGM (for a review on the IGM, see [7]). The sources of the ionizing photons are thought to be stars in galaxies [5], and quasars [2, 3, 4] for H I and He II reionization, respectively.

During reionization, ionizing photons heat the IGM by tens of thousands of degrees. This heating, combined with cooling from adiabatic expansion and atomic processes, are the primary processes that influence the thermal state of the low density (1 – 100 times the cosmic mean density) IGM [102, 103, 104, 105, 106, 107]. The thermal energy of the IGM smooths and extends the distribution of the gas, which in turn affects structure formation. After each reionization event, the low density IGM cools asymptotically towards an equilibrium temperature [103, 108]. During this time the ionization state is well understood, as the neutral fraction is set by the equilibrium between photoionizations and recombinations. All of this makes the IGM, and especially the low density IGM, a valuable probe of the post-reionization universe ( $z < 6$ ) and the scales probed make it useful for both astrophysics and cosmology.

Conveniently, there are numerous observations probing intergalactic gas at  $2 < z < 6$ . Generally, these are observations of the Lyman- $\alpha$  forest, the series of absorption

features blueward of the rest-wavelength Lyman- $\alpha$  emission observed in quasar spectra [1]. A single forest spectrum is a one-dimensional map of the gaseous structure along that line of sight, making it a useful probe of structure formation. Knowledge of the large scale structure, either through the flux power spectrum or the inferred matter power spectrum, constrains warm dark matter (WDM) models [109, 110]. In addition to probing structure formation, the Lyman- $\alpha$  forest can be used to measure the thermal state of the IGM, leading to a set of measurements describing the thermal history of the IGM. Using the thermal and ionization history of the IGM, one can test models of the makeup and evolution of the ionizing background, and thus infer properties of the ionizing sources and sinks over time [96].

There are several ways in which Lyman- $\alpha$  forest spectra are processed to constrain cosmological models and the thermal state of intergalactic gas. Cosmological contexts generally make use of the flux power spectrum from a sample of Lyman- $\alpha$  forest spectra [111, 112, 113, 96]. The flux power is the Fourier transform of the flux over-density,  $\delta_F = F/\langle F \rangle - 1$ . The flux power spectrum is sensitive to cosmological parameters on large scales ( $k < 0.02$  s/km for velocity wavenumber  $k$ ), and constrains small scale smoothing at higher  $k$  [114]. For example, smoothing is enhanced in WDM models, leading to a reduction in power above some critical value of  $k$ , (dependent on the mass of the WDM particle). This makes the flux power spectrum a robust tool for constraining WDM models [110].

The spectral statistics used in determining the thermal state of the IGM are more varied. Common methods include statistics which encapsulate an entire forest spectrum [115, 116, 117, 118, 119, 120], as well as analyses which make use of absorption features

from spectra decomposed via Voigt profile fitting [121, 122, 104, 123, 124, 125]. The small scale flux power spectrum and the distribution of flux are also used to constrain the IGM thermal state [111, 126].

The Lyman- $\alpha$  forest probes scales on which non-linear structure growth is important, and so cosmological hydrodynamic simulations of the IGM are necessary to build a map between model parameters and observations. These simulations require two components: collisionless cold dark matter modelled using N-body techniques, and collisional baryons which include pressure forces. One common simplification is that, although baryons are evolved hydrodynamically, the initial conditions for both species are identical, using the transfer function for the total matter fluid [127].

Before recombination, baryons couple to radiation, suppressing their clustering on sub-horizon scales and reducing clustering relative to the dark matter. After recombination, baryons fall into the potential well of the cold dark matter and so the linear transfer functions differ by  $< 1\%$  at  $z = 0$ . The effect is larger at higher redshifts,  $z = 2 - 5$ , where the Lyman- $\alpha$  forest is a sensitive probe of the gas [128]. [129] showed that separate transfer functions can affect the one-dimensional Lyman- $\alpha$  forest flux power spectrum by  $5 - 10\%$  on scales  $0.001 - 0.01$  s/km in the redshift range  $z = 2 - 4$ .

The aim here is to determine whether species specific initial transfer functions have an appreciable effect on probes of the Lyman- $\alpha$  forest. We use the simulation technique developed in [129], which reproduces the theoretical offset between the dark matter and baryon power [130], to model separate initial transfer functions. Recently, [131] (see also [132, 133]) resolved this discrepancy by perturbing the particle masses, in agreement with

the results from [129]. We will examine the effect of these initial conditions on measures of the thermal state of the IGM; the curvature [119], Doppler width cutoff [121], and Doppler width distribution [126]. We also examine the effect on the matter and flux power spectrum, which could have consequences for warm dark matter models [134]. The simulations we use are higher resolution than in [129], allowing us to better probe smaller scales.

In Section 3.3 we outline the simulations and artificial spectra used throughout. In Section 3.4 we discuss the methods used to calculate each measure of the IGM, as well as the results of those calculations. Measures of the thermal history of the IGM, including the curvature, and the Doppler width cutoff and distribution, are covered in sections 3.4.1 & 3.4.2, respectively. The WDM relevant measures are examined in Sections 3.4.3 (flux power spectrum) and Section 3.4.4 (matter power spectrum). Section 3.5 discusses numerical convergence with box size, resolution, and number of artificial spectra used. In Section 3.6 we summarize and conclude.

We assume throughout a flat  $\Lambda$ CDM cosmology with  $\Omega_0 = \Omega_b + \Omega_{CDM} = 0.288$ ,  $\Omega_b = 0.0472$ ,  $h = 0.7$ ,  $n_s = 0.971$ , and  $\sigma_8 = 0.84$  (consistent with 9-year WMAP results [80]).

### 3.3 Simulations

Our set of hydrodynamical simulations were performed using the N-body and smoothed particle hydrodynamics (SPH) code MP-Gadget<sup>1</sup>, described in [135, 77]. MP-Gadget is a fork of Gadget-3, itself the descendent of Gadget-2 [78]. Initial conditions are generated with MP-GenIC, the initial conditions generator packaged with MP-Gadget.

---

<sup>1</sup><https://github.com/sbird/MP-Gadget3>

The initial power spectrum, and transfer functions are generated with the Boltzmann code CLASS [79].

Simulations using offset grids for both particle species (which is common in the literature) often introduce a spurious growing mode to the CDM-baryon difference. This can be avoided by using a glass to initialize the baryons [136, 129], or by an appropriate perturbation of the particle masses [132]. Two sets of simulations are used throughout this work, with initial conditions set using the baryon-glass method. Both sets of simulations use a glass to initialize the baryons and a grid to initialize the CDM. A glass procedure, with 14 time-steps, is then applied to the combined distribution to minimize CDM-baryon overlap, avoiding chance overdensities set by the initialization. The two sets of simulations then differ, with the first set using a single transfer function for both species, and the second set using separate, species specific transfer functions. Scale-dependent perturbations are included via first-order Lagrangian perturbation theory (during final preparation of this manuscript, [132, 131] proposed an alternative method based on second-order perturbation theory, which gives similar results). The phases of the Fourier modes are identical, leading to the same realization of cosmic structure on scales larger than the particle grid.

Gas is assumed to be in ionization equilibrium with a uniform ultraviolet background using the model of [137]<sup>2</sup>. [138] recently updated their UV background model and showed that simulations using uniform UV backgrounds do not accurately model the timing and photoheating associated with reionization. In our simulations reionization has completed by  $z = 6$  (the average neutral hydrogen fraction in low density regions of our simulations is less than 1%). Our results are generated in the redshift range  $2 < z < 6$ ,

---

<sup>2</sup>Specifically the 2011 update, <https://galaxies.northwestern.edu/uvb-fg09/>

Table 3.1: Transfer Function Simulations

Simulation	Box Volume	N	$M_{\text{gas}} (M_{\odot})$
Main	$(20 \text{ Mpc h}^{-1})^3$	$2 \times 1024^3$	$9.8 \times 10^4$
Low Res	$(20 \text{ Mpc h}^{-1})^3$	$2 \times 768^3$	$2.3 \times 10^5$
Small Box	$(10 \text{ Mpc h}^{-1})^3$	$2 \times 512^3$	$9.8 \times 10^4$

after hydrogen reionization. We do not implement He II reionization because the scale of our simulation box size is smaller than a typical He II bubble [139], leading to an effectively instantaneous reionization.

Star formation is implemented using the standard approach for Lyman- $\alpha$  forest analyses. Gas particles in the simulations are turned into stars using a simple density-based method; when they reach an overdensity  $\rho/\langle\rho\rangle > 1000$ , but remain at a temperature  $T < 10^5$ , they are turned into stars [140]. Our simulations do not include black hole or supernovae feedback.

The set of high-resolution simulations include our Main simulations, a simulation with lower gas mass resolution, and a simulation with a smaller box length of  $10 \text{ Mpc h}^{-1}$ . The latter two are used to check box size and gas mass resolution convergence. We discuss convergence for each result in the relevant results section, as well as in Section 3.5. All simulations start at  $z = 99$  and have periodic boundaries. Box volume, particle number, and gas particle mass resolution are reported in Table 3.1. The gas particle mass resolution is set so that the higher redshift Lyman- $\alpha$  forest is resolved [141].

Lyman- $\alpha$  absorption spectra are generated by sending random skewers through the simulation box using Fake Spectra Flux Extractor [142]<sup>3</sup>, described in [143]. Our analysis uses 5,000 randomly placed skewers, which are generated for each snapshot, leading to a

<sup>3</sup>[https://github.com/sbird/fake\\_spectra](https://github.com/sbird/fake_spectra)



large set of  $1 \text{ km s}^{-1}$  pixel width neutral hydrogen absorption spectra for redshifts in the range  $2 < z < 6$ . We discuss convergence of our results with number of sight lines in Section 3.5.

### 3.4 Methods & Results

In this section we examine the effect using species-specific initial conditions has on two commonly studied properties of the IGM, both of which use Lyman- $\alpha$  forest spectra. The first is the temperature-density relation of the low density IGM, which is generally parameterized as

$$T(\Delta) = T_0 \Delta^{\gamma-1}, \quad (3.1)$$

where  $\Delta$  is the matter overdensity,  $T_0$  is the temperature at mean density ( $\Delta = 1$ ), and  $\gamma-1$  is the power-law index [103, 108]. Throughout we focus on redshifts after H I reionization ( $\leq 6$ ), where adiabatic cooling and photoheating dominate the thermal state. This is the regime where the temperature-density relation parameterized above is generally valid (though it is not best described with a single temperature-density relation during He II reionization) [144, 145, 139]. We focus on three measures which probe the temperature-density relation of the IGM: the curvature (3.4.1), the Doppler width cutoff (3.4.2), and the Doppler width distribution (3.4.2).

The second property is the matter power spectrum of the IGM, which can constrain dark matter models, especially warm dark matter through its effect on structure formation. The matter power spectrum of the dim and diffuse IGM is not directly accessible. However, the flux power spectrum is a good proxy and allows constraints to be placed on the thermal

free-streaming of dark matter and thus a potential WDM particle mass. We examine the effect species specific initial conditions have on both the Lyman- $\alpha$  flux power spectrum (Section 3.4.3) and matter power spectrum (Section 3.4.4).

### 3.4.1 Curvature

The curvature statistic introduced in [119] has an approximately one-to-one relationship with the temperature of the IGM at an optimal overdensity. The temperature at the mean density can then be inferred using a temperature-density relationship slope calibrated from simulations. The curvature is essentially the second derivative, or curvature, of the flux. Specifically, it is given by  $\kappa \equiv F''/(1 + F'^2)^{3/2}$  and traces the ionized fraction of hydrogen. Higher temperature gas will show more thermal broadening in the absorption features of the spectra, while lower temperature gas will retain more small-scale spectral features. Because the curvature summarizes the entire spectrum it does not require decomposing spectra into individual absorbers, making it useful up to higher redshifts than the Doppler width methods (Section 3.4.2).

The simulated spectra are processed, and the curvature calculated following the general procedure in [119]:

1. Gaussian noise is added to the spectra such that the  $S/N \sim 20$ , then a cubic b-spline is fit to the flux iteratively. The initial break point spacing between the piecewise b-spline is set at  $50 \text{ km s}^{-1}$  and additional points are added to improve the fit until either a minimum resolution is reached ( $10 \text{ km s}^{-1}$ ) or the fit converges (the  $\chi^2$  value between spline and spectra changes by less than 3 between break point additions).




Figure 3.1: Example of the curvature measurement. Top: spectrum with noise (black) and the smoothed cubic spline fit to it (green). Middle: The region enclosed in the dotted box in the top panel, renormalized into  $10 \text{ Mpc h}^{-1}$  sections and rescaled such that the set of all these sections have the same mean flux across simulations. Bottom: the curvature of the middle panel. The single value reported for such a section is the mean absolute curvature value in regions where the renormalized and rescaled flux lies between  $0.1 < F < 0.9$ .

The resulting spline, an example of which can be seen in the top panel of Figure 3.1, is used in place of the spectrum for the rest of the analysis.

2. The spline is then renormalized by breaking it into  $10 \text{ Mpc/h}$  sections and dividing by the maximum value in that section. This normalizes the measure and avoids uncertainties due to continuum finding.

3. Each of these sections is then rescaled such that the mean flux of the entire set of sections is consistent with the model from [146], given by an effective optical depth,

$$\tau^{\text{eff}} = 0.0023(z + 1)^{3.65}.$$

The result of steps (ii) and (iii) are shown in the middle panel of Figure 3.1. Note that the values used in the processing outlined above (e.g. the S/N,  $\chi^2$  convergence value, etc.) are chosen either to agree with [119], to be reasonable in regards to observation, or simply to fit the artificial spectra well.

The curvature is then calculated, using only flux in the range  $0.1 < F < 0.9$ . The bottom panel of Figure 3.1 shows an example of the curvature, before restricting the flux range. For each section the mean absolute curvature is returned,  $\eta = \langle |\kappa| \rangle$ , and the average of  $\eta$  for each redshift is shown in Figure 3.2 (top), along with the fractional difference between the two (e.g.  $|x_1/x_2 - 1|$ ) in percentage (bottom). The squares (blue) show the results for the simulation which uses separate initial transfer functions, and the triangles (brown) show the simulation which uses the same transfer function. The difference between the two is remarkably small, peaking at  $< 2\%$  at  $z = 6$ .

For convergence testing, noise is not added to the spectra, though the spectra are still renormalized and rescaled. Convergence of the curvature is discussed in Section 3.5, and shown in Figure 3.9 for simulation parameters, and in Figure 3.13 (top panel) for convergence with number of sight lines used.

### 3.4.2 Doppler Width Methods

#### Doppler Width Cutoff

Another method used to determine the thermal state of the IGM, first introduced in [121], is fitting the lower cutoff in the Doppler width ( $b$ ) of spectral features as a function of their neutral hydrogen column density ( $N_{\text{HI}}$ ). The Doppler width characterizes the width




Figure 3.2: Results for the average curvature versus redshift. The curvature is robust up to higher redshift, hence the inclusion of redshifts up to  $z = 6$ . The agreement is extremely good, with a maximum difference of  $< 2\%$  at redshift  $z = 6$ .

of an absorption profile due to broadening from the motion of the particles constituting the absorber. Theoretically, an absorber has a minimum Doppler width, i.e. due entirely to thermal broadening with no additional effects such as broadening from a velocity gradient. The minimum will depend on the temperature of the absorber,  $b_{\text{therm}} = \sqrt{2k_b T/m}$ , where  $k_b$  is the Boltzmann constant,  $T$  the temperature, and  $m$  the proton mass. The temperature depends on the density, with higher density clouds having a higher temperature and thus a broader spectral profile. As the true density of an absorber is not observable, the column density of the absorber and its Doppler width are used as proxies. We then have a relation between the minimum Doppler width and the column density, where a higher column density

absorber has a higher temperature, and thus a larger minimum Doppler width. A lower cutoff, in the form of a power law, can then be fit to a set of Doppler width, column density measurements. Using simulations and observations of the  $N_{\text{HI}} - b$  cutoff parameters, the temperature-density relation is calibrated and the physical density inferred from the column densities [147].

This method requires decomposing spectra into features, with widths and amplitudes corresponding to the Doppler width and column density of the associated absorbers. We use Fake Spectra Flux Extractor [142] to decompose our artificial spectra, which are the optical depths ( $\tau$ ) along lines of sight through the simulations, into individual features by fitting the flux ( $F = e^{-\tau}$ ) using Voigt profiles. The Voigt profile width is the Doppler width,  $b$ , and the Voigt profile normalized amplitude is the neutral hydrogen column density,  $N_{\text{HI}}$ .

Perfect spectra and flawless profile fitting are not possible – Figure 3.3 shows a clear trend in the minimum Doppler width with column density, but there are points that lie below the visual cutoff. To best fit this minimum Doppler width cutoff, [121] developed an algorithm which fits the relation

$$\log_{10}(b) = \log_{10}(b_0) + (\Gamma - 1) \log_{10}(N_{\text{HI}}/N_{\text{HI},0}). \quad (3.2)$$

We follow most closely the algorithm choices used in [148], which includes an initial ( $\sigma$ ) rejection step to remove low points which in an observational setting are most likely metal-contaminated. Note that in that algorithm the  $\sigma$ -rejection step uses the RMS deviation while the cutoff fitting step uses the mean absolute deviation and a column density normalization of  $N_{\text{HI},0} = 10^{13.6}$ , in agreement with [121] and [148]. An example of the cutoff fit can be seen in Figure 3.3.




Figure 3.3: Examples of the distribution of column densities and Doppler widths at  $z = 2.4$  from the Main simulation with separate transfer functions. The green line is the best fit cutoff given by Equation 3.2 with  $b_0 = 17.4$ ,  $\Gamma - 1 = 0.153$ .

Figure 3.4 shows the effect separate transfer functions have on the cutoff parameters as a percent difference in the results between the Main simulations. Difficulties with fitting Voigt profiles at higher redshift due to fewer unsaturated features (a lower ionized fraction) lead to a more ambiguous cutoff in the distribution, so we only attempt to fit a cutoff up to  $z = 4$ . The difference in the two fit parameters brought about by the separate transfer functions is 2 – 5% for  $b_0$  and 2 – 8% for  $\Gamma - 1$ . This is similar to the level of convergence with resolution and spectral sample seen in Figure 3.10, indicating that the primary driver of any differences is most likely due to changes in the sample of spectral features used in the fit. Note that the effect is also not uniform with redshift, leading to an increase in some redshift bins, and a decrease in others.

Our results are consistent with [148] to within 5%, and [125] to within 10%, at a redshift of  $z = 2.4$  (where we overlap with both works). However, the sample used in fitting the lower cutoff has a large effect on the result. Specifically, the inclusion or omission of a few data points near the cutoff shifts the fit, while points away from the cutoff have little effect. This is further indicated by the continued variance in the fit parameters as the sample size is increased (see Figure 3.13, bottom two panels). The variance due to the sample size is on the same order ( $< 10\%$ ) as the difference we see between the two Main simulations in Figure 3.4. The effect of the sample used is likely the primary driver of the relatively weak convergence with simulation box size and resolution seen in Figure 3.10.




Figure 3.4: Results for the Doppler width fitted cutoff parameters. Shown is the percent absolute difference between the simulations with and without separate transfer functions. Top: the difference for the minimum width, i.e. the intercept for the fitted cutoff. At most the unscaled results differ by  $\approx 5\%$  at  $z = 4$  and the scaled results differ by  $\approx 2\%$ . Bottom: same as the top panel, but for the slope of the fitted cutoff. The agreement is good, peaking at  $\approx 8\%$  difference for both the scaled (at  $z = 3.2$ ) and unscaled results (at  $z = 2$ ).



Convergence of these results between the Main simulations and the Low Res and Small Box simulations are discussed in Section 3.5 and can be seen in Figure 3.10. Convergence with number of sight lines used is also discussed in the Section 3.5 and can be seen in Figure 3.13 (bottom two panels).

### **Doppler Width Distribution**

Recently, [149] and [126] have attempted to address the problems associated with fitting the lower cutoff of the  $N_{\text{HI}} - b$  distribution. Specifically, these works avoid fitting the lower cutoff entirely, which removes the strong dependence on just a few points near the cutoff (making poor use of most of the data), and potential systematics in the cutoff fitting algorithm.

Instead, these methods directly measure the traditional temperature-density parameters (mean temperature  $T_0$ , and power law index  $\gamma$ ) by comparing the observed  $N_{\text{HI}} - b$  distribution to simulations or simulation-derived emulators. [149] uses the full two-dimensional  $N_{\text{HI}} - b$  distribution and an emulator built on a grid of thermal parameters to estimate the observed  $T_0$  and  $\gamma$ . [126] bins the Doppler widths by column density, then uses the resulting set of one-dimensional distributions to infer the temperature-density parameters. Their simulations indicate that the lower column density bins are more sensitive to  $T_0$ , while the higher column density bins are sensitive to  $\gamma$ . The idea is to simultaneously fit the set of observed distributions with simulated distributions (with varying thermal parameters). An example of the Doppler width distributions is shown in the top panels of Figure 3.5, separated into the three column density bins highlighted in [126].




Figure 3.5: Top: Doppler width distributions at  $z = 3$  for the three column density bins highlighted in [126]. Shown are the distributions for the single transfer function simulation (blue) and species-specific transfer functions (brown) simulation. The distributions are from Voigt fitting 5,000 skewers for each simulation snapshot. Bottom: Residuals (as a percent) for the Doppler width distributions between our Main simulations. Our simulations differ by less than 1% in mean temperature ( $T_0$ ) and density index ( $\gamma$ ), measured from the  $z = 3$  snapshot.

To quantify the effect of using separate transfer functions, we find the maximum in the Doppler widths distribution residuals between our Main simulations across the bins shown in Figure 3.5. From the lower panels in that Figure, we see the (absolute) difference is at most  $\approx 30\%$ . The difference at other redshifts can be as much as  $\approx 70\%$  at  $z = 3.6 - 4$  ( $\approx 50\%$  for  $z = 2 - 2.4$ ), down to the  $\approx 30\%$  seen at and around  $z = 3$ . The transfer function affects the distribution at a similar level to the residuals [126] found in their fit to observations around  $z = 3$  (their figure 7). However, we may not be fully converged with simulation box size or resolution.

The  $\approx 30\%$  difference brought about by the transfer functions is similar to the difference we find between our Main simulations and the Small Box ( $\approx 50\%$ ) and Low Res

( $\approx 35\%$ ) simulations (see Figure 3.11). It is also similar to the level of convergence [126] found (their figure F1). Furthermore, the  $z = 3$  distribution is better converged than the distribution at other redshifts. The distribution from the Low Res simulations differ from the Main simulations by as much as 80% at lower redshifts ( $z = 2 - 2.4$ ) and as much as 100% at higher redshifts ( $z = 3.6 - 4$ ). The Box Size simulation distribution also differs by as much as 75% at lower redshifts, and as much as 150% at higher redshifts. It is not surprising that the higher redshift result is less well converged, since the lower ionized fraction at these redshifts leads to more highly saturated spectra, and thus Voigt fitting is more difficult. Regardless of these difficulties at higher redshift, the effect of the separate transfer functions on the distribution is on the same order as the effect from simulation box size and resolution at all redshifts explored here. This indicates that, in terms of the Doppler width distribution, our simulations are not fully converged (see also Section 3.5).

### 3.4.3 Flux Power Spectrum

Lyman- $\alpha$  forest spectra from the IGM can also be used to constrain cosmologies alternative to  $\Lambda$ CDM. For example, a warm dark matter particle suppresses structure relative to CDM on scales smaller than the WDM particle free-streaming scale [134]. Lyman- $\alpha$  forest spectra probe the scales relevant to the WDM model and can be used to estimate the clumping of matter (the matter power spectrum) through the observed flux distribution (the flux power spectrum) [140]. The flux power spectrum is  $P_F(k) = |L^{-1}\tilde{\delta}_F^2(k)|$ , where  $\tilde{\delta}_F^2(k)$  is the Fourier transform of the flux excess,  $\delta_F(k) = F(k)/\langle F(k) \rangle - 1$ , and  $L$  is the length of the sight lines in velocity space.




Figure 3.6: Ratio of flux power spectrum from the simulation with separate transfer functions to the simulation with a single transfer function. The red curve shows the result before rescaling the flux, while the black curve is the result after rescaling. The grey shaded region in each panel indicates where the flux power spectrum is converged with simulation box size and resolution to at least 10% (see Figure 3.12). On the scales probed here, the effect at  $z > 2$  tops out at  $\approx 5\%$ , while at  $z = 2$  the effect is  $< 1\%$ .

The effect of WDM on the flux power spectrum is to suppress high  $k$  ( $> 0.01$  s/km) power and marginally enhance low  $k$  ( $< 0.01$  s/km) power [150, 151]. The shape of the Lyman- $\alpha$  forest flux power spectrum can be used to measure the suppression scale, which directly constrains the WDM particle mass. The ever increasing number of observed quasar sight lines has meant that a statistically significant sample can be assembled to look at this effect [152, 153].

These constraints rely on accurate modeling of the flux power spectrum in simulations using a CDM or WDM model, coupled with observed Lyman- $\alpha$  forest spectra. Figure 3.6 shows the effect of using separate transfer functions for baryons and CDM on the Lyman- $\alpha$  forest flux power spectrum. Shown is the percent change in the power spectrum when using separate transfer functions instead of a single transfer function. The effect is generally strongest at the high end of the  $k$  range, with a decrease in power across all redshifts at  $k < 0.02$  for the rescaled result. The effect increases with redshift, however both the rescaled and unscaled results remain at  $\lesssim 5\%$  across the range of  $k$  our simulations reliably probe.

While observations extending to the largest wave numbers used here are not presently available, we can compare the flux power spectrum we obtain from our simulations to currently available data where they overlap. Figure 3.7 shows our flux power spectrum at  $z = 3$  and  $z = 4$  compared to data taken from [152]. Their estimate of the Lyman- $\alpha$  contribution to their total flux power, as well as their total flux power, are shown with the reported errors. Our results show a small underestimate of the power in comparison with them, but are roughly consistent. To estimate the level of agreement, we

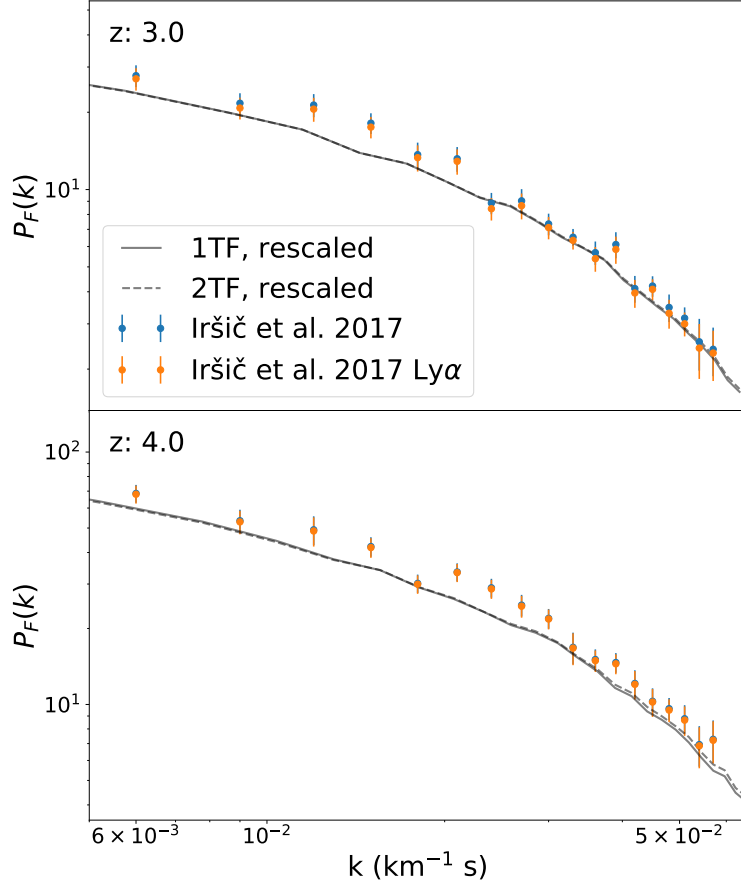


Figure 3.7: Comparison of the flux power spectra from the simulations presented here, to the data presented in [152]. In [152] they present the full flux, as well as an estimate on the contribution from only Lyman- $\alpha$ , which is comparable to our simulated results. For low  $k$ , the simulations appear to underpredict the power with the largest discrepancies at  $\approx 2.5\sigma$ , but most lie within  $1\sigma$ .

interpolate our flux power spectrum onto the wave numbers of the [152] data. We find that at most, our power spectrum differs by  $\sim 2.5\sigma$  from their data (where the deviation is their reported statistical and systematic errors, added in quadrature). For most data points, the difference is within  $1\sigma$ . In terms of this deviation, the difference between the single and separate transfer function cases is  $< 1\sigma$ .

### 3.4.4 Matter Power Spectrum

The total matter power spectrum is affected only at the 1 – 2% level at all scales probed with the simulations presented in this work ( $k = 0.6 - 100 \text{ h Mpc}^{-1}$ ). This is unsurprising, as the effect of the separate initial conditions is to reproduce the offset of the power between the baryons and dark matter, and not to change the total matter power spectrum.

Figure 3.8 shows the difference in the species specific matter power spectrum ratio (baryon power over CDM power). From this we see that the effect of the separate initial conditions is to decrease the power in the baryons, while retaining the behaviour at both higher redshift (for linear structure to dominate on a large range of scales) and at lower redshift (for baryons to collapse into non-linear structures at small scales). The offset is independent of  $k$ , and is consistent both with linear theory [154] and with [129].

We ran two additional simulations, both with the same volume and particle number as the Small Box simulation. The first used species-specific transfer functions, but initialized both the baryons and CDM using (offset) grids. Previous work [129] has shown that without additional adjustments [132, 131], this setup can lead to artificial growth in the CDM-baryon power difference. On the scales probed in this work, we find that the simulation using offset grids does a fairly good job in reproducing the power difference, agreeing with our separate transfer function Main simulation to within 2% on all the but the smallest scales ( $k > 7 \text{ h Mpc}^{-1}$ , where it stays within 5%). This indicates that the artificial growth in the power difference is sourced from larger scale velocity perturbations, and that the half-glass approach is unnecessary on the scales probed here.




Figure 3.8: Percent difference in the ratio of the species specific matter power spectrum for baryons over CDM. At  $z = 2$  and  $z = 6$  the simulations using a single transfer function predict more power in the baryons than the case with separate transfer functions. At  $z = 6$ , in the separate transfer function simulation we see higher power in the CDM, as expected (meanwhile the single transfer function simulation is consistent with  $P_b \approx P_{CDM}$ ). At  $z = 2$  we see the expected enhancement in baryon power on small scales.

The second additional simulation included WDM (a 5.3 keV WDM particle, consistent with [151, 155]). The change in the total matter power spectrum remained at the sub-percent level for all scales. The baryon-DM power difference agreed with the CDM simulations up to  $k \approx 3 - 4 \text{ h Mpc}^{-1}$ , with a relative drop in power ( $\sim 3\%$ ) at higher  $k$ , consistent with expectations (WDM suppressing small scale structure). This drop-off



was consistent between the simulations with and without species specific transfer functions, indicating that the effect of using separate transfer function is independent of the DM temperature.

### 3.5 Convergence

We check the convergence of our simulations with box size and mass resolution by running an additional four simulations. For both the separate and same transfer function cases we run a lower mass resolution simulation (Low Res) and a smaller box size simulation (Small Box). The simulation volume, particle number, and mass resolution can be seen in Table 3.1. The mass resolution used in the Main simulations agrees with [119], which previously showed convergence for the curvature at that resolution. [124] showed convergence for the  $N_{\text{HI}} - b$  cutoff parameters using the same set of simulations.

Figure 3.9 shows the fractional difference as a percentage between each of the Main simulations, and the two associated convergence simulations (called  $\delta$ ) for the curvature. The curvature calculated here uses spectra without added noise, obviating the need for a spline fit. Otherwise, the calculation is the same as that outlined in Section 3.4.1 (spectra are renormalized into 10 Mpc/h sections and the mean flux is rescaled). The curvature is well converged, with a difference of  $\lesssim 1\%$  for the Small Box simulations, and  $\sim 2\%$  for the Low Res simulations.

Figure 3.10 shows the convergence for the two fit parameters of the  $N_{\text{HI}} - b$  cutoff. As this method uses a population of features taken from each set of spectra, renormalizing between different volume simulations is not necessary. The logarithmic intercept,  $b_0$ , has




Figure 3.9: Convergence for the curvature. The convergence is quite good, staying within  $\sim 1\%$  for the smaller box simulations, and within  $\sim 2\%$  for the lower resolution simulations. This emphasizes the importance of the mass resolution in simulations aimed at probing the Lyman- $\alpha$  forest.

converged to an absolute difference of  $< 10\%$  at all redshifts. The logarithmic slope,  $\Gamma - 1$ , is less well converged, remaining within an absolute difference of  $< 20\%$  at all redshifts. However, the amount of data used to make the fit strongly affects the result. This can be seen in the bottom two panels of Figure 3.13, where the convergence with the number of sight lines used is shown. In contrast with the flux power spectrum and curvature, these fit parameters do not completely converge, instead exhibiting some variance all the way up to the inclusion of all 5,000 sight lines.

Figure 3.11 shows the convergence at  $z = 3$  for the Doppler width distribution. Specifically, Figure 3.11 shows the percent difference between the Doppler width distri-




Figure 3.10: Convergence for the Doppler width cutoff fit values. Top: the percent difference for the minimum width (cutoff intercept) fit values between the Main simulations and the Low Res and Small Box simulations. The convergence stays within 10%. Bottom: same as the top panel, but for the slope of the cutoff. The convergence is less good, staying within 20%.

butions for the Main and Small Box simulations, and between the Main and Low Res simulations. The top panels show this difference for simulations using species specific transfer functions, while the lower panels show this for the simulations using a single transfer function. The maximum difference across all three bins associated with the Small Box simulation is  $\approx 50\%$ , while the maximum difference associated with the Low Res simulations is  $\approx 35\%$ . This is a similar level of convergence to [126] (their Figure F1). However, as stated in Section 3.4.2, the level of convergence shown here is on the same order as the effect seen from using separate transfer functions ( $\sim 25\%$ ). We also note that redshifts around  $z = 3$

are better converged, with the Low Res and Small Box simulations differing by as much as  $\approx 80\%$  for redshifts between  $z = 2$  and  $z = 2.4$ , and at higher redshifts ( $z = 3.6 - 4$ ) they differ by as much as  $\approx 150\%$ . This indicates that for the Doppler width distribution, we are not fully converged with resolution or box size.




Figure 3.11: Convergence of the Doppler width distribution at a redshift of  $z = 3$ . Top: the percent difference between our Main simulation and the Small Box and Low Res simulations, for the simulations using species specific transfer function. Bottom: same as the top panel, but for the simulations using a single transfer function.

Figure 3.12 shows the convergence for the flux power spectrum at the two redshifts which span our analysis. We are well converged at low  $k$  ( $< 0.03 \text{ km}^{-1}\text{s}$ ) for all simulations and redshifts. At  $z = 6$  we are well converged ( $< 10\%$ ) for three cases, while the lower resolution simulation with separate transfer functions is not well converged beyond  $k \approx 0.05 \text{ km}^{-1}\text{s}$ . The shaded regions in Figure 3.6 which indicate the trusted  $k$  range for each redshift are based on this. At  $z = 2$  we are well converged across the range ( $k < 0.1 \text{ km}^{-1}\text{s}$ ) explored here.




Figure 3.12: Convergence of the flux power spectrum with simulation box volume and gas mass resolution. At early times ( $z = 6$ ) we are converged within 10% up to  $k \approx 0.05 \text{ km}^{-1}\text{s}$ . For larger  $k$ , the Low Res, separate transfer function simulation is less well converged. At late times ( $z = 2$ ) we are fully converged with resolution, however the convergence with box size is less good ( $\lesssim 10\%$ ).

The matter power spectrum is converged with gas mass resolution at all redshifts ( $2 < z < 6$ ) and scales ( $k = 0.6 - 100 \text{ h Mpc}^{-1}$ ) such that the difference between the higher and lower resolution simulations is  $\lesssim 5\%$ . Convergence with box volume is  $\lesssim 20\%$  in the range  $k = 10 - 100 \text{ h Mpc}^{-1}$ , and  $\approx 30\%$  from  $k = 0.6 - 10 \text{ h Mpc}^{-1}$ .

Finally, we check the convergence of the curvature, flux power spectrum, and  $N_{\text{HI}} - b$  cutoff fit parameters with the number of sight lines used in each of their calculation.

Figure 3.13 shows this convergence for the Main simulation using separate transfer functions. The convergence trends for all quantities are also seen in the simulations using a single transfer function. Each statistic is calculated using only the corresponding fraction of the 5,000 sight lines available (e.g. the 0.2 value for the  $N_{\text{HI}} - b$  cutoff fit parameters use only features from the first 1000 random sight lines). While the curvature is insensitive to the number of sight lines used (beyond  $\sim 50$  sight lines), the flux power spectrum depends strongly on the number. The convergence of the flux power spectrum shows the maximum difference between consecutive power spectra (the difference between the spectrum with  $x$  sight lines and the spectrum with  $x + 5$  sight lines). The  $N_{\text{HI}} - b$  cutoff fit parameters continue to fluctuate even with a large number of sight lines. The variance, even when using all of the sight lines in the  $N_{\text{HI}} - b$  cutoff fit parameters may explain the worse convergence with simulation size and resolution seen in Figure 3.10.




Figure 3.13: Convergence of curvature, flux power spectrum, and  $N_{\text{HI}} - b$  cutoff parameters with number of artificial spectra used. Shown here is the convergence for simulations with separate transfer functions. The flux power spectrum convergence shows the maximum change in the spectrum when additional sight lines are added. Note that the sight lines are randomly placed in the simulation box.

### 3.6 Summary & Conclusions

In this work we explored how switching from a single initial transfer function (for both baryons and CDM) to species specific transfer functions affects properties of the IGM. Using a set of high resolution simulations, we have quantified the effect this change has on probes of the IGM via the Lyman- $\alpha$  forest. The Main simulations presented here differ only in the transfer functions used; one uses a single transfer function, the other follows [129], adopting species specific transfer functions. Our work is motivated by simulations sometimes failing to match the theoretical offset between baryon and cold dark matter power, though we found that on the scales probed in this work ( $k = 0.6 - 10 \text{ h Mpc}^{-1}$ ), standard simulation methods did fairly well. Artificial spectra were extracted from snapshots in the  $2 < z < 6$  range and statistics relevant to the thermal history of the IGM and WDM models were calculated. Below we summarize the results of this work.

- The curvature is relatively unaffected by the species specific transfer functions, with a peak difference of  $< 2\%$  at  $z = 6$ . The effect of the transfer functions (on average, shifting  $\log_{10}\eta$  by  $\approx 0.01$ ) is about the  $1\sigma$  ( $2\sigma$ ) uncertainty observed in [120] ([119])
- The Doppler width cutoff fit parameters continue to vary with number of sight lines used, even when using a large number of sight lines (see Figure 3.13). This is likely due to the fitting method, which depends strongly on the few data points near the cutoff. The effect on these parameters is at most  $\approx 5\%$  for the fit intercept and  $\approx 8\%$  for the fit slope, however the variance due to sample size has a similar level of effect. The effect of the transfer functions seen here is at a similar level to the error [148], and generally less than the error [125] found from fitting to observations.



- The Doppler width distribution is affected by  $\approx 30\%$  around  $z = 3$ , which is similar to the difference between observation and best fit in [126]. However our simulations at  $z = 3$  are only converged to  $\approx 50\%$  with box size and  $\approx 35\%$  with resolution, the same order as the effect from the transfer functions.
- The flux power spectrum is affected more at high  $k$  ( $k > 0.05$  s/km) and redshift ( $z > 2$ ). However, the enhancement to the power is at most  $\approx 5\%$  for  $z > 2$  and  $\lesssim 1\%$  at  $z = 2$ . This level of effect is small compared to observational uncertainties, amounting to  $\approx 10\%$  of the uncertainty presented in [152] (and see Figure 3.7).

For measures of the thermal state of the IGM, the effect of separate transfer functions is either small ( $\sim 1\%$  for the curvature), or subdominant to either variance inherent to the method of calculation (Doppler width cutoff) or convergence (Doppler width distribution).

The flux power spectrum is relatively unaffected on the scales and times which are currently well observed, however our results indicate that it may become important on smaller scales or higher redshift. This may indicate that using separate transfer functions may be important for future observations and surveys. However, the effect is most pronounced at early times and on small scales, which constrain WDM most effectively. The importance of this effect will only increase as future measurements lead to a higher resolution flux power spectrum, constraining WDM models more stringently.

The future study of the IGM will be predicated on measuring absorption spectra at higher redshifts and at higher resolution using large optical and infrared telescopes in conjunction with broader surveys such as the James Webb Space Telescope [156]. Given

the ever increasing sample size and quality of IGM observations, it is paramount that simulations keep pace by improving their precision and modeling. The adjustment to the initial conditions in simulations explored here is one such improvement, but there are others which should be implemented as well, for example the modeling of He II reionization [139]. Improved simulations, in concert with future observations, will push the study of the IGM into the reionization epoch as it occurs, leading to a greater understanding of this relatively recent major phase transition, as well as the formation of the first galaxies and their subsequent evolution.

## Chapter 4

# Emulating the Lyman- $\alpha$ Forest

### 4.1 Abstract

In this work we extend our recently developed multi-fidelity emulation technique to the simulated Lyman- $\alpha$  forest flux power spectrum. Multi-fidelity emulation allows interpolation of simulation outputs between cosmological parameters using many cheap low-fidelity simulations and a few expensive high-fidelity simulations. Using a test suite of small box (30 Mpc/h) simulations, we show that multi-fidelity emulation is able to reproduce the Lyman- $\alpha$  forest flux power spectrum well, achieving an average accuracy when compared to a test suite of 0.8%. We further show that it has a substantially increased accuracy over single-fidelity emulators, constructed using either the high or low-fidelity simulations only. In particular, it allows the extension of an existing simulation suite to smaller scales and higher redshifts.

## 4.2 Introduction

The modern and future testbed for cosmology lies in small scales and non-linear structures. Cosmological analyses exploit observations of these scales to explore questions such as the nature of dark matter, the total neutrino mass, and the thermal history of the intergalactic medium (IGM). One of the most powerful probes of small scale structure is the Lyman- $\alpha$  forest, a series of absorption features in the spectrum of quasars [1, 157, 158, 159, 160, 92, 161, 162]. Numerical simulations are required to analyze these observations as they probe the non-linear regime. As these simulations are expensive, cosmologists build emulators [163, 164, 165], which interpolate a summary statistic (in this case the 1D Lyman- $\alpha$  forest flux power spectrum) between simulation outputs at different cosmological parameters.

A recent development in cosmology is the application of multi-fidelity emulators, which allow simulations with different particle loads, and thus costs, to be combined together [166]. Here we adapt the multi-fidelity emulation technique to the Lyman- $\alpha$  forest 1D flux power spectrum. Multi-fidelity emulation is especially useful in this context because the Lyman- $\alpha$  forest probes a range of redshifts, and is sensitive to smaller scales, which require higher resolution simulations, at higher redshifts [141]. The models developed here will allow a single emulator to target the wide range of scales probed by the Lyman- $\alpha$  forest, which would otherwise require a computationally infeasible number of very large simulations [167].

The Lyman- $\alpha$  forest is the result of overlapping neutral hydrogen absorption profiles in the spectra from distant luminous quasars, processed through the expansion of the

universe [1]. As light travels from the quasar, it passes through neutral hydrogen gas of varying densities. In the rest frame of those neutral hydrogen islands, light that has been redshifted close to the Lyman- $\alpha$  transition at  $1215.67\text{\AA}$  will be absorbed and the rest transmitted. This is repeated as the light continues to intersect more neutral hydrogen islands on its path towards us, the observers. The result is a quasar transmission spectra containing an overlapping field of absorption features that provides a proxy to the dark matter density along that sightline [168].

The densities probed by the Lyman- $\alpha$  forest, from redshift  $2 - 5$ , are  $\sim 1 - 100 \times$  the cosmological mean density. At these densities stellar winds and star formation effects are negligible, although black hole feedback is important [169, 170]. The densities along with the range of scales accessed has made the Lyman- $\alpha$  forest popular in cosmological studies, including: constraining the thermal history of the IGM and thus reionization [171, 124, 113, 96, 172, 173, 174], constraining cosmological parameters including the neutrino mass [140, 175, 176, 177, 178, 155, 179], and testing alternatives to cold dark matter [109, 150, 151, 155, 180, 181].

The Lyman- $\alpha$  forest 1D flux power spectrum is the most commonly used summary statistic for Lyman- $\alpha$  forest spectra. It probes small scale structure by measuring the two-point Fourier-space correlation between neutral hydrogen absorption within a sightline [168].

Current observational measurements of the Lyman- $\alpha$  forest flux power spectrum come from either a lower resolution, larger sample survey (SDSS, [182]), or various higher resolution, smaller sample surveys [152, 183, 184]. In [182], the flux power spectrum constructed from BOSS and eBOSS spectra accesses redshifts from  $z = 2.2 - 4.6$  (6% & 18%

average uncertainty, respectively) and scales from  $k \approx 0.001 - 0.02 \text{ km}^{-1} \text{ s}$  (6% & 14% average uncertainty, respectively). The small sample, higher resolution surveys generally access a similar redshift range, but shift both the largest and smallest scales to higher  $k$ . For example, in [183] (their conservative results), using spectra from multiple surveys (XQ-100, KODIAQ, and SQUAD) they access redshifts from  $z = 2 - 4.6$  (7% & 27% average uncertainty, respectively) and scales from  $k \approx 0.005 - 0.1 \text{ km}^{-1} \text{ s}$  (12% & 9% average uncertainty, respectively).

The Dark Energy Spectroscopic Instrument (DESI) will soon report its first year results. Ultimately it will increase the number of Lyman- $\alpha$  quasar spectra by a factor of four over SDSS. This corresponds to  $\sim 50$  quasars per square degree and a total of  $7 \times 10^5$  quasars over the 14,000 square degree survey footprint [67]. In addition, DESI is expected to measure the 1D flux power spectrum at smaller scales ( $k < 0.035 \text{ km}^{-1} \text{ s}$ ) and higher redshifts ( $z > 4.6$ ) than SDSS, achieving order of a few percent accuracy [185].

Extracting cosmological information from these observations will require simulations which follow the distribution of gas at relevant densities and on relevant scales. For the Lyman- $\alpha$  forest, box sizes of at least  $100 \text{ Mpc h}^{-1}$  and mean particle spacing of  $100/3072 \approx 0.03 \text{ Mpc h}^{-1}$  are necessary [167]. Earlier work has focused on methods which can reduce the cost of such simulations. [167] used a splicing technique to produce high resolution, large volume outputs from three sets of less computationally intensive simulations: low resolution, large volume; high resolution, small volume; and low resolution, small volume. [186] explored the use of Richardson extrapolation to enhance output resolution, in addition to testing the splicing technique.

Parameter inference tasks, such as a direct Markov Chain Monte Carlo analyses, require  $\sim 10^5 - 10^6$  model evaluations, indicating the number of simulations required by a naive approach. Even using techniques such as splicing, this is computationally infeasible. However, using a significantly reduced number of simulations ( $\sim 30$ ), an emulator can be constructed that effectively interpolates between this smaller set of simulations. In addition to the Lyman- $\alpha$  forest, emulators have been used extensively in cosmology for studying: the matter power spectrum [165, 187, 188, 189, 190, 191, 192], weak lensing [193, 194], the halo mass function [195, 196, 197], and the 21-cm signal [198, 199, 200, 201]. Still, the computational resources required to run  $\sim 30$  simulations with the requisite volume and resolution is highly restrictive, especially for the Lyman- $\alpha$  forest.

Here, we use modern machine learning techniques to alleviate the computational resource cost associated with constructing an emulator. Specifically, we are concerned with using machine learning to predict high resolution simulation outputs to a high degree of accuracy, while running a minimal number of high resolution simulations. One machine learning method that is suited to this task is a Gaussian Process (GP) emulator. Gaussian processes [202] are a means of interpolating between the simulation outputs, providing function prediction in a Bayesian framework. Essentially, a distribution of functions is learned through training on simulations, and the mean (best estimate) and variance (interpolation error) of the output can be returned for arbitrary simulation inputs. While other interpolation methods are possible, Gaussian processes have many benefits: the inherent quantification of prediction uncertainty, the option to incorporate prior knowledge, and the ability to interpolate within high-dimensional parameter space.

Previous uses of GP emulators for the Lyman- $\alpha$  forest have been shown to be effective at predicting summary statistics [77, 203, 204, 205, 206, 181]. [77] self-consistently showed that the predicted flux power from their GP emulator (trained with 21 simulations) agreed to within 1 – 2% of the corresponding simulation flux power spectrum. These GP emulators still require a substantial computational cost, as the full simulation suite must be run with sufficient volumes and resolutions for the Lyman- $\alpha$  forest.

Recently, [166] implemented a multi-fidelity GP emulator [207] for the matter power spectrum. Here, we combine and expand on the methods outlined in [77] and [166], to produce a multi-fidelity GP emulator for the Lyman- $\alpha$  forest flux power spectrum. In our multi-fidelity model, the training simulations are split into two fidelities; a large sample of low resolution simulations (low fidelity, LF), and a small subset of these simulations run at higher resolution (high fidelity, HF). Note that we use fidelity and resolution interchangeably throughout this work. Using these two training sets, the multi-fidelity emulator is trained to predict the 1D flux power spectrum that would be output by a *high* resolution simulation for arbitrary cosmological and astrophysical parameters.

A multi-fidelity emulator allows us to replace some of the HF simulations that would be needed in a single-fidelity emulator with LF simulations. This can dramatically reduce the computational cost of constructing an emulator, while retaining predictive power across parameter space. Using this method, emulators can be constructed that make use of the full range of scales and redshifts probed by Lyman- $\alpha$  forest observations. This enables analyses which can jointly constrain thermal, astrophysical, and cosmological parameters.



In our high resolution simulations, the Lyman- $\alpha$  forest flux power spectrum is converged to  $\approx 5\%$ . While the scales we probe in this work are resolved, the box size we use is smaller than required to analyze the full range of scales available in Lyman- $\alpha$  forest data, i.e. we cannot compute a likelihood function using all the real data without larger boxes. We therefore defer a full cosmological likelihood analysis to future work. Our goal is to quantitatively test the accuracy of the emulator output, as compared with the output from a set of testing simulations run at the same resolution as the HF training set, and demonstrate the validity and utility of the multi-fidelity technique. Specifically, we quantify the accuracy of the single- and multi-fidelity emulators with respect to true values from the testing simulations, thus determining how effective the multi-fidelity model is at producing high resolution outputs at minimal computational cost.

### 4.3 Simulations

Simulations were performed using MP-Gadget<sup>1</sup>, an N-body and smoothed particle hydrodynamics (SPH) code built on the solid base of Gadget-3 (last described in [78]). MP-Gadget has been substantially modified to include shared-memory parallelism using OpenMP, together with many other algorithmic improvements and new subgrid models, as described in [208, 209, 210]. The initial power spectrum and transfer functions are generated with the Boltzmann code CLASS [79]. Species-specific initial conditions are generated for baryons and dark matter [210, 10]. We include radiation in the cosmological background model and assume massless neutrinos.

---

<sup>1</sup><https://github.com/MP-Gadget/MP-Gadget>

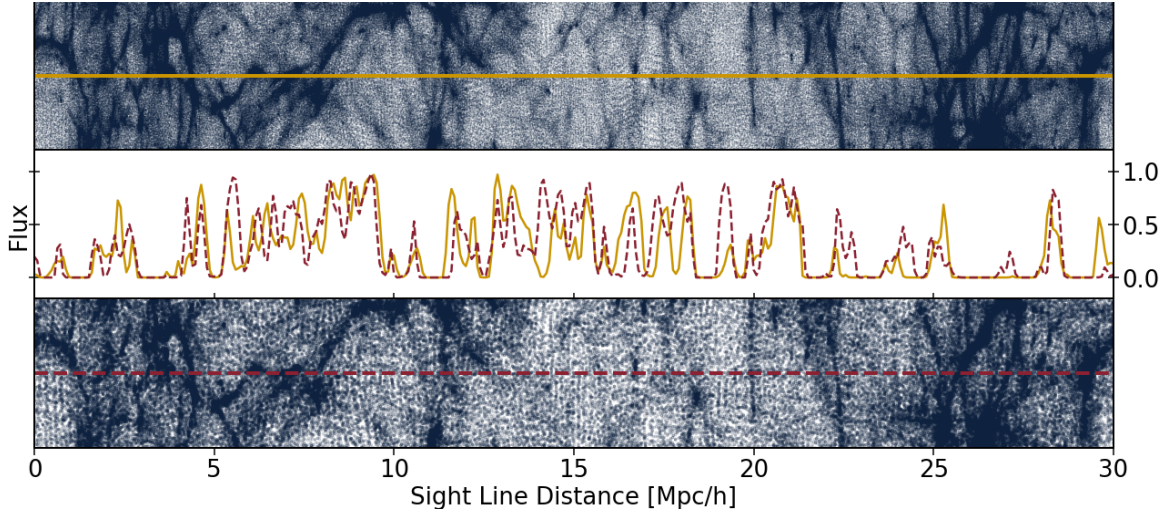


Figure 4.1: Example Lyman- $\alpha$  forest spectra and corresponding gas density from simulations at redshift 4. The top and bottom panel show the simulated gas surrounding the skewer which produced the spectra shown in the middle panel. Examples are shown for simulations run at high (top panel, yellow line) and low resolution (bottom panel, red line).

The physics models largely follow those used for the `ASTRID` simulation, which are described in [208]. Primary sources for these models, as well as changes from [208] are described below.

We use a cubic kernel for our density estimator (rather than a quintic kernel). For these simulations, the cubic kernel produced a neutral hydrogen column density distribution that was more consistent with observations for column densities between  $10^{20}$  and  $10^{22}$   $\text{cm}^{-2}$  [211]. We continue to use the pressure-entropy formulation of SPH. The smaller SPH kernel increases the noise within galaxies, but has minimal effect on the Lyman- $\alpha$  forest [212].

Star formation follows the model of [213], with our specific implementation as described in [214]. We lower the number of stars produced per gas particle from 4 (used in `ASTRID`) to 1 (as in `Illutris-TNG`), which speeds up the simulation without having an effect on the Lyman- $\alpha$  forest.

Black holes follow the model of [209]. We found that for the resolutions used here, the dynamic friction from gas led to a few black holes escaping from their dark matter halo, so we only use dynamic friction from dark matter and stars. The black hole feedback factor, which controls the fraction of luminosity that is converted to thermal energy, is an emulator parameter (BHF), with the associated parameter limits in Figure 4.3. The black hole feedback radius is fixed to  $3 \text{ kpc h}^{-1}$ , selected to be the average black hole feedback radius at the highest tested resolution when using a nearest neighbour distance. To accommodate a lower mass resolution than **ASTRID**, the minimum stellar mass needed in a halo to seed a black hole was increased to  $2 \times 10^8 M_\odot$ , and all black hole seeds start with a mass of  $5 \times 10^4 M_\odot$ .

Stellar winds are modeled following [215]. The decoupling distance for the winds is increased from  $20 \text{ kpc h}^{-1}$  to  $1 \text{ Mpc h}^{-1}$ , which allows the winds to recouple due to density changes rather than travel distance. The density threshold for wind recoupling is set to 10% of the star formation density threshold (which is 57.7 times the critical density). The minimum wind velocity is set to  $100 \text{ km/s}$ . Finally, metal return (gas enrichment) is disabled as it is not important for the Lyman- $\alpha$  forest and can be computationally expensive.

Gas is assumed to be in ionization equilibrium with a uniform ultraviolet background using the model of [138]. We boost the temperature of the gas to  $15000 \text{ K}$  the timestep after the gas is reionized, to model impulsive heating during hydrogen reionization from ionization fronts [107].

We implement He II reionization using the model of [139]. The input parameters for this model are: quasar mean bubble size and variance, redshifts for the start and completion

Table 4.1: Multi-Fidelity Emulator Simulations

Simulation	Box Volume	$N_{\text{part}}$	$M_{\text{gas}} (M_{\odot} h^{-1})$
LF	$(30 \text{ Mpc } h^{-1})^3$	$2 \times 256^3$	$[1.78, 2.37] \times 10^7$
HF	$(30 \text{ Mpc } h^{-1})^3$	$2 \times 512^3$	$[2.22, 2.96] \times 10^6$
Test	$(30 \text{ Mpc } h^{-1})^3$	$2 \times 512^3$	$[2.22, 2.96] \times 10^6$

of He II reionization  $z_i^{\text{He II}}, z_f^{\text{He II}}$ , and the quasar spectral index  $\alpha_q$  (which effectively scales the peak temperature during He II reionization). The quasar bubble size is reduced from the default of  $\sim 30$  Mpc, motivated by radiative transfer simulations [3], to 5 Mpc, due to our small box size.

Simulations are initialised at  $z = 99$  and finish at  $z = 2$ , and use periodic boundaries. Box volume, particle number, and gas particle mass resolution are reported in Table 4.1. The range given for the gas resolution is due to the varying value of  $h$  in our simulation suite. The gas particle mass resolution for our HF simulations does not meet the resolution that [141] recommend to resolve the forest at all redshifts of interest. However, the Lyman- $\alpha$  forest flux power spectrum from our HF simulations is converged to within  $\approx 5\%$  of a simulation that does meet the required resolution of [141]. We are interested in the performance of the multi-fidelity GP emulator in learning the mapping from low to high resolution, thus this slight lack of numerical convergence does not affect our results. Examples of the gas density (at  $z = 3.6$ ) for the two resolutions are shown in the top and bottom panels of Figure 4.1.

Lyman- $\alpha$  forest absorption spectra are generated using Fake Spectra Flux Extractor [142]<sup>2</sup>, described in [143]. We generate 32,000 (seeded) randomly placed skewers for

<sup>2</sup>[https://github.com/sbird/fake\\_spectra](https://github.com/sbird/fake_spectra)




Figure 4.2: Lyman- $\alpha$  forest flux power spectrum from LF and HF simulations. The top panel shows the flux power spectrum at redshift 3 from an LF simulation (blue), and its HF counterpart (yellow). The lower panels show the ratio of the flux power for these two resolutions, for all LF-HF simulation pairs, at  $z = 3.8, 4.6, \& 5.4$ . Dashed lines shown the highest  $k$  probed by BOSS/eBOSS [182], and the estimated reach for DESI [185].

each snapshot, from  $z = 5.4$  to  $z = 2.0$  in increments of  $\Delta z = 0.2$ . The pixel resolution is set to  $10 \text{ km s}^{-1}$ . An optical depth threshold of  $\tau < 10^6$  is set to eliminate damped Lyman- $\alpha$  systems. Example spectra from one LF and one HF simulation (at  $z = 4$ ) are shown in the

middle panel of Figure 4.1. Note that the LF simulation is not simply a smoothed version of the HF simulation, as the fine velocity structure of the gas moves the location of the absorption peaks.

These sets of neutral hydrogen absorption spectra are used to construct the Lyman- $\alpha$  forest flux power spectrum for each simulation, at each redshift. The flux power spectrum is defined as  $P_F(k) = |L^{-1}\tilde{\delta}_F^2(k)|$ , where  $\tilde{\delta}_F^2(k)$  is the Fourier transform of the flux excess,  $\delta_F(k) = F(k)/\langle F(k) \rangle - 1$ , and  $L$  is the length of the sightline. The reported flux power spectrum is averaged over all 32,000 spectra.

Figure 4.2 shows flux power spectra from a single LF simulation and its HF counterpart, and the ratio of these at several redshifts. While the exact difference between the LF and HF flux power spectra depends on simulation input parameters and redshift, in general the LF differs most from the HF at small scales. The enhanced power on large scales for the LF flux power spectra is consistent with [167], and is likely due to differences in heating and cooling during H I and He II reionization.

Figure 4.3 lists the input parameters that are varied across our suite of simulations, as well as their limits. Two parameters control the primordial power spectrum:  $n_s$  is the scalar spectral index (slope) and  $A_p$  is the amplitude (see [77] for more details). Three of the parameters relate to the He II reionization model:  $z_i^{\text{He II}}$  and  $z_f^{\text{He II}}$  are the redshifts for the start and end of He II reionization, and  $\alpha_q$  is the quasar spectral index. We vary the Hubble constant through  $h$ , and the total matter density through  $\Omega_M h^2$ . One parameter is varied for H I reionization:  $z^{\text{H I}}$  is the midpoint redshift of H I reionization. Finally, we vary one parameter for the black hole model:  $BHF$  is the black hole feedback factor,




Figure 4.3: Simulation parameter limits and samples. Parameters for the low resolution simulations (crosses) were determined by filling a Latin hypercube. See Figure 4.4 for good and bad examples of a Latin hypercube design. Initially, 30 low resolution samples were generated, then an additional 10 were added while maintaining the Latin hypercube method, hence the non-uniform spacing for the low resolution samples. The optimal subset of low resolution simulations is determined (see Section 4.4.2) and this subset is run at higher resolution (shown as red circles).

which controls the fraction of luminosity that is converted to thermal energy. Note that our simulations do not contain a kinetic feedback model. However, at  $z > 2$  it is expected that the thermal mode dominates. Also shown in Figure 4.3 are the LF and HF training samples. Note that the HF samples are a subset of the LF samples. The selection of the HF samples is described in Section 4.4.2.




Figure 4.4: Two examples of a Latin hypercube sample, where the example on the left poorly spans the parameter space, and the example on the right samples the parameter space well. The  $d$ -dimensional (in this case  $d = 2$ ) parameter space is divided into  $n$  sections (in this case  $n = 12$ ) so that only one sample is selected from each section. This can be done in  $n!$  different configurations (for  $n = 12$  this is  $\sim 4.8 \times 10^8$  configurations). To select the best configuration, random configurations can be generated and the one that maximizes the minimum distance between points is retained (often called maximin).

## 4.4 Emulators

In Section 4.4.1, we will briefly review emulation using a Gaussian process. In Section 4.4.2, we will review how the Gaussian process emulator can be extended to model simulations with different qualities using a multi-fidelity emulator, `MFEmuLator`. The multi-fidelity emulation technique of [207] will be reviewed in Section 4.4.2. Section 4.4.2 will discuss the differences in multi-fidelity emulator design between this work and [166]. Finally, Section 4.4.2 will outline how we select our training simulations in the parameter space.



#### 4.4.1 Gaussian process emulator

Gaussian process (GP) regression models [202] have been widely used to build cosmological emulators [163, 164, 165]. A GP provides closed-form expressions for predictions. In addition, a GP naturally comes with uncertainty quantification, which is useful when building an inference framework. In the context of emulation, a GP can be seen as a Bayesian prior for the simulation response. It is a prior because the emulator model is chosen to ensure smoothness and monotonicity features of the simulation response *before* data are collected [216].

Let  $\boldsymbol{\theta} \in \Theta \subseteq \mathbb{R}^d$  be the input cosmologies for the simulator, where  $d$  is the dimension of the parameters ( $d = 9$  for our emulator).  $f(\boldsymbol{\theta})$  is the corresponding output summary statistic. In this work, the summary statistic,  $f(\boldsymbol{\theta})$ , is the Lyman- $\alpha$  forest flux power spectrum. A GP regression model can be viewed as a prior on the response surface of our simulated Lyman- $\alpha$  forest flux power spectrum:

$$f(\boldsymbol{\theta}) \sim \mathcal{GP}(\mu(\boldsymbol{\theta}), k(\boldsymbol{\theta}, \boldsymbol{\theta}')), \quad (4.1)$$

where  $\mu(\boldsymbol{\theta}) = \mathbb{E}[f(\boldsymbol{\theta})]$  is the mean function, and  $k(\boldsymbol{\theta}, \boldsymbol{\theta}') = \text{Cov}[f(\boldsymbol{\theta}), f(\boldsymbol{\theta}')] is the covariance kernel function. In this work, we assume a zero mean function, and we used the same covariance function as [77], which will be defined later in this section.$

Suppose we run the simulations at  $n$  carefully chosen input cosmologies,  $\mathcal{D} = \{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_n\}$ , and we generate the corresponding Lyman- $\alpha$  forest flux power spectrum for each simulation,  $\mathbf{y} = \{f(\boldsymbol{\theta}_1), \dots, f(\boldsymbol{\theta}_n)\}$ . Conditioning on this training data, we can get the predictive distribution of  $f$  at a new input cosmology  $\boldsymbol{\theta}$  through the closed-form expression:

$$f(\boldsymbol{\theta}) | \mathbf{y}, \mathcal{D} \sim \mathcal{N}(\mu_n(\boldsymbol{\theta}), \sigma_n^2(\boldsymbol{\theta})), \quad (4.2)$$

where the mean and variance are:

$$\begin{aligned}\mu_n(\boldsymbol{\theta}) &= \mathbf{k}(\boldsymbol{\theta}, \mathcal{D})^\top \mathbf{K}(\mathcal{D})^{-1} \mathbf{y}; \\ \sigma_n^2(\boldsymbol{\theta}) &= k(\boldsymbol{\theta}, \boldsymbol{\theta}) - \mathbf{k}(\boldsymbol{\theta}, \mathcal{D})^\top \mathbf{K}(\mathcal{D})^{-1} \mathbf{k}(\boldsymbol{\theta}, \mathcal{D}).\end{aligned}\tag{4.3}$$

The vector  $\mathbf{k}(\boldsymbol{\theta}, \mathcal{D}) = [k(\boldsymbol{\theta}, \boldsymbol{\theta}_1), \dots, k(\boldsymbol{\theta}, \boldsymbol{\theta}_n)]$  represents the covariance between the new input cosmology,  $\boldsymbol{\theta}$ , and the training data. The matrix  $\mathbf{K}(\mathcal{D})$  is the covariance for the training data.

For the covariance kernel function, we choose the same kernel as in [77], which is a combination of a linear kernel and a radial basis kernel (RBF):

$$\begin{aligned}k(\boldsymbol{\theta}, \boldsymbol{\theta}'; \sigma_0, \mathbf{l}, \boldsymbol{\sigma}) &= k_{\text{RBF}}(\boldsymbol{\theta}, \boldsymbol{\theta}'; \sigma_0, \mathbf{l}) + k_{\text{LIN}}(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\sigma}) \\ &= \sigma_0^2 \exp\left(\sum_{i=1}^d -\frac{(\theta_i - \theta_i')^2}{2l_i^2}\right) + \sum_{i=1}^d \sigma_i^2 \theta_i \theta_i',\end{aligned}\tag{4.4}$$

where  $\sigma_0^2$  and  $\boldsymbol{\sigma}^2$  are the variance hyperparameters for the RBF kernel and the linear kernel, respectively.  $\mathbf{l}$  is the lengthscale parameter that controls the smoothness of the Gaussian process function. We applied Automatic Relevance Determination (ARD) for both linear and RBF kernels. That is, we assign one lengthscale  $l_i$  (variance  $\sigma_i$ ) hyperparameter for each input dimension  $i$  for the RBF and linear kernels. This allows the GP to dynamically learn the scale over which each input dimension varies, which corresponds to the degree of sensitivity of the flux power spectrum to the input parameter.

Although we do not explicitly write in the notation,  $f(\boldsymbol{\theta})$  is a single-valued output. Since our target summary statistic is a vector, we model each  $k$ -bin of the flux power spectrum with a separate GP. The primary reason is that the correlation between the low-fidelity and high-fidelity flux power spectrum changes depending on the scale. The multi-fidelity method can only capture this scale dependence if we model each scale separately.

### 4.4.2 Multi-Fidelity Emulation

We first introduce the Kennedy-O’Hagan model (KO model) [207] in Section 4.4.2. Section 4.4.2 describes the changes we have made to adapt the model from [166] to the Lyman- $\alpha$  forest. Finally, the strategy we employ for choosing parameters at which to generate high-fidelity training simulations is described in Section 4.4.2.

#### Kennedy O’Hagan Method

The KO model [207] was first introduced to model a sequence of computer codes with increasing fidelity. For simplicity, we assume there are only two fidelities: low-fidelity (LF) simulations with low resolution, and high-fidelity (HF) simulations with high resolution.

We define  $\{\mathbf{y}_{\text{LF}}, \mathbf{y}_{\text{HF}}\}$  as the Lyman- $\alpha$  forest flux power spectra in the training set.  $\mathbf{y}_{\text{LF}} = \{f_{\text{LF}}(\boldsymbol{\theta}_i^{\text{LF}})\}_{i=1}^{n_{\text{LF}}}$  and  $\mathbf{y}_{\text{HF}} = \{f_{\text{HF}}(\boldsymbol{\theta}_i^{\text{HF}})\}_{i=1}^{n_{\text{HF}}}$ , where  $n_{\text{LF}}$  and  $n_{\text{HF}}$  are the number of simulations in the low- and high-fidelity training sets. We use the KO method to model Lyman- $\alpha$  forest flux power spectra from different fidelities:

$$f_{\text{HF}}(\boldsymbol{\theta}) = \rho \cdot f_{\text{LF}}(\boldsymbol{\theta}) + \delta(\boldsymbol{\theta}), \quad (4.5)$$

where  $\rho$  is a trainable parameter describing a multiplicative correction between the low- and high-fidelity Lyman- $\alpha$  forest flux power spectra.  $\delta(\boldsymbol{\theta})$  is a GP independent of  $f_{\text{LF}}(\boldsymbol{\theta})$ , describing an additive correction between fidelities. In other words, Equation 4.5 assumes the high-fidelity Lyman- $\alpha$  forest flux power can be decomposed as the low-fidelity flux power multiplied by a correction parameter,  $\rho$ , and an additive bias function  $\delta(\boldsymbol{\theta})$ .

As mentioned in [166], the  $\rho$  parameter has to be scale-dependent (a function of  $k$ ) to model the well-known fact that small scales are less well resolved in smaller simulations. Here we use the same method as [166] and assume Equation 4.5 is a single-output GP model. We assign a KO model to each  $k$  bin of the data.<sup>3</sup> In this way, we can model  $\rho$  as a function of  $k$ , as shown in Figure 4.5.




Figure 4.5: The scale parameter,  $\rho$ , of the KO model (Equation 4.5) as a function of  $k$ . Different colors represent different redshifts. We separate the redshifts into four bins,  $2 \leq z < 3$ ,  $3 \leq z < 4$ ,  $4 \leq z < 5$ , and  $5 \leq z \leq 5.4$ . Within each redshift bin, we average over fixed  $k$  modes, which are linearly spaced between the maximum  $k$  and minimum  $k$  at the given redshift range. This emulator used 40 LF and 6 HF for training.

We also assign KO models for each redshift. As shown in Figure 4.5,  $\rho$  is a non-trivial function of both  $k$  and  $z$ , so we cannot simply use an emulator trained on one redshift to apply on another redshift.<sup>4</sup> We note that it is possible to assume a smooth function to

<sup>3</sup>We can easily get the same set of  $k$  bins for low- and high-fidelity by using the same spectral resolution for both simulations.

<sup>4</sup>See [204] for a Lyman- $\alpha$  forest emulator that uses a single GP for all redshifts, and achieves sub-percent

model  $\rho(k, z)$ . In practice, observational data are conditioned on a specific redshift, so training emulators on separate redshifts is sufficient for cosmology inference.

Figure 4.5 shows that  $\rho$  stays close to unity at large scales for most of the redshifts. At small scales, however, different redshifts require different values of  $\rho$ . At the middle redshifts ( $3 \leq z < 5$ ),  $\rho$  has a large positive deviation from unity. At the low redshifts ( $2 \leq z < 3$ ),  $\rho$  has a moderate deviation toward  $\rho > 1$ . The only exception is at the high redshifts ( $5 \leq z \leq 5.4$ ), which stays close to  $\rho = 1$  for all scales. This indicates that the correction to the Lyman- $\alpha$  forest flux power due to the resolution of the simulation varies with redshift, depending on the over-density probed by the forest.

### **Model differences from Ho et al. (2022)**

Here we highlight the ways in which we have adapted the model from [166], which emulated the non-linear matter power spectrum at  $z = 0 - 2$ , to the Lyman- $\alpha$  forest at  $z = 2 - 5.4$ . Since the redshift range is larger in this work, we employed a new strategy to select the optimal HF training set by averaging over the interpolation loss for all redshift bins. We will describe the strategy in detail in Section 4.4.2.

In [166], we outlined two multi-fidelity methods: a linear multi-fidelity emulator (the KO model, or the autoregression model (AR1)), and a non-linear multi-fidelity emulator (non-linear autoregressive GP, or NARGP, [217]). However, we found that NARGP requires more HF training simulations for the Lyman- $\alpha$  forest flux power than AR1, perhaps due to the wide range of redshifts used. We use the KO model for our main results, and describe the NARGP results in Section 4.6.

---

accuracy, albeit with some ambiguity between model parameters and redshift.

In this work, instead of emulating logarithm scaled powers, we adopted the mean-flux normalization strategy proposed in [77]. We normalize all flux power spectra in the training set by the median spectrum:

$$\begin{aligned}\mathbf{y}_{\text{LF}} &\leftarrow \frac{\mathbf{y}_{\text{LF}}}{\text{median}_i(\mathbf{y}_{\text{LF}})} - 1; \\ \mathbf{y}_{\text{HF}} &\leftarrow \frac{\mathbf{y}_{\text{HF}}}{\text{median}_i(\mathbf{y}_{\text{LF}})} - 1.\end{aligned}\tag{4.6}$$

The index  $i$  refers to one of the spectra in the training set,  $\mathbf{y}_{\text{LF}} = \{f_{\text{LF}}(\boldsymbol{\theta}_i^{\text{LF}})\}_{i=1}^{n_{\text{LF}}}$ . Equation 4.6 ensures the training sample distribution is close to having a zero mean, matching the prior of the GP emulator. We found that in practice this normalisation makes training the emulator substantially easier. Note that we normalize the HF training set using the same LF median spectrum. As the HF training set is small, the median spectrum estimate for HF is noisy, and so using it for normalization may introduce some unwanted training bias.

### Sampling strategy for high-fidelity simulations

The KO model approach can be seen as a Bayesian way to correct an emulator from low-fidelity to high-fidelity. Thus, if  $\mathbf{y}$  is the high-fidelity Lyman- $\alpha$  forest flux power spectrum used for training, and  $\boldsymbol{\theta}$  is the corresponding input parameters:

$$\begin{aligned}\mathbf{y} &= f_{\text{LF}}(\boldsymbol{\theta}) + (f_{\text{HF}}(\boldsymbol{\theta}) - f_{\text{LF}}(\boldsymbol{\theta})) \\ &= f_{\text{LF}}(\boldsymbol{\theta}) + \text{error}(\boldsymbol{\theta}).\end{aligned}\tag{4.7}$$

The emulation accuracy will be directly affected by how well an autoregressive construction can model  $\text{error}(\boldsymbol{\theta})$ . Usually, a large set of low-fidelity simulations are used as training data for  $f_{\text{LF}}(\boldsymbol{\theta})$  because they can be obtained cheaply. The quality of training data for  $\text{error}(\boldsymbol{\theta}) = f_{\text{HF}}(\boldsymbol{\theta}) - f_{\text{LF}}(\boldsymbol{\theta})$  thus relies on the choice of high-fidelity simulations.

In [166], we proposed an optimization strategy to select high-fidelity training simulations. A low-fidelity only emulator (LFEmu)<sup>5</sup> is trained on a subset of low-fidelity training simulations. The posterior means of the trained LFEmu are used to calculate the emulation errors from the remaining LF samples in the Latin Hypercube Sampling (LHS). By minimizing the emulation errors of LFEmu, we can grid search for the optimal set of cosmologies that best interpolates the parameter space using a small number of training simulations. Assuming LFEmu is correlated with HFEmu, we can use the selected optimal cosmologies as inputs for the HF training set. By ensuring the HF training set achieves a good interpolation, we mitigate emulation errors for the multi-fidelity emulator.

In practice, we employ a three-stage procedure for building a multi-fidelity emulator:

1. Prepare LF simulation suite.
2. Prepare HF simulation suite. This is done by using LFEmu to find the set of cosmologies that minimizes the interpolation loss.
3. Build MFEmulator. If the accuracy is not enough, go back to stage 1 or 2 to run more training simulations.

For stage (ii), to avoid wasting computational resources running more LF simulations, we directly use the LF simulation suite in stage (i) to build and validate the LFEmu. Thus, the cosmologies chosen for the HF set are a subset of the LF simulation LHS, which fulfills the nested training dataset design suggested in [207]. The benefit of using a nested data

---

<sup>5</sup>In a similar way, we call a high-fidelity only emulator, HFEmu.

structure,  $\theta_{\text{HF}} \subseteq \theta_{\text{LF}}$ , is that we can directly compute posterior means from the LF training set for cosmologies  $\theta_{\text{HF}}$ , without any interpolation in LF.

We note that it is possible to train a `MFEmulator` without using the LF simulations to optimize the HF points. However, if the selection of HF points are suboptimal (i.e., can barely interpolate in the prior volume), then the `MFEmulator` accuracy will be suboptimal. This is because the  $\text{error}(\theta)$  cannot be decomposed into an autoregressive structure easily.

To find the optimal HF training set across the full redshift range,  $z = 2 - 5.4$ , we train a `LFEmu` for each redshift and get the validation loss (we used mean squared errors). We sum up the validation loss for all redshifts and find a subset of cosmologies that minimizes the summed validation loss.

In Figure 4.6, we summarize the above described procedure in a flowchart. For a formal description, we refer to [166]. The only difference between the proposed procedure in Figure 4.6 and the procedure in [166] is that we optimize  $\theta_{\text{HF}}$  across redshifts  $z = 2 - 5.4$  in this work. Thus, we have an additional step to sum the `LFEmu` validation loss across redshifts.

Training an `LFEmu` on all possible low-fidelity subsets is computationally intensive. To reduce costs, we employed the greedy optimization strategy from [166]. We first explored all possible subsets for 3 design points within the LF LHS. For the optimal 4 design points, instead of exploring all possible subsets, we grew the subset one point at a time, fixing the previously chosen optimal 3 HF points. In the same line of thought, we grew the subset to 6 optimal design points for HF training cosmologies. Our final simulation suite of 40 LF and 6 HF samples, along with parameter limits, is shown in Figure 4.3.



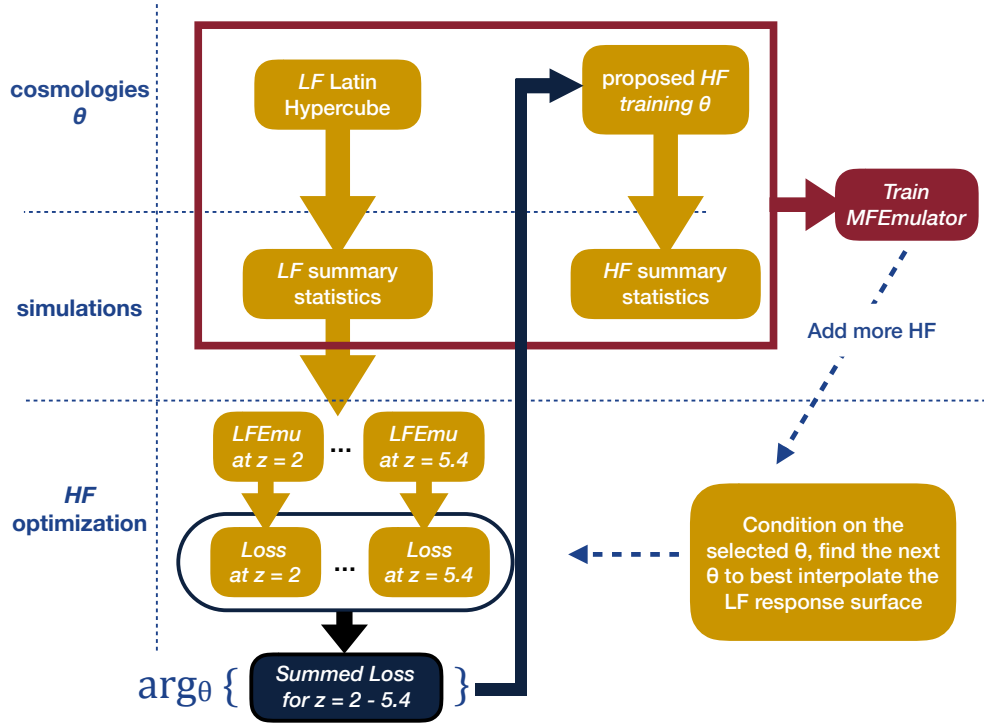


Figure 4.6: A flowchart for training a multi-fidelity emulator. We start with a low-fidelity (LF) set of simulations. We then select a subset of low-fidelity points to train an emulator, LFEmu, at each redshift. The validation loss for each LFEmu is computed using the rest of the LF simulations as a validation set. Finally, we sum up the validation losses at each redshift, and use the summed loss to propose a set of cosmologies  $\theta_{\text{HF}}$  which can best minimize the summed loss.

## 4.5 Results & Discussion

Using the flux power spectra from our LF and HF simulations, we train single-fidelity (one LF only, and one HF only) and multi-fidelity emulators. These trained emulators are used to predict the flux power spectrum output for a set of 10 simulation input parameters. We then compare these predictions to the corresponding testing simulations which were run at the same resolution as the HF simulations (see Table 4.1).

### 4.5.1 Emulator Accuracy

In the following, we only show results for the emulators that use the full available set of training simulations (40 LF and 6 HF). We have verified that using all available training simulations leads to the most accurate emulator. Section 4.5.2 shows how emulator accuracy degrades when a smaller subset of the available simulations is used.

Using the full set of available simulations, the mean prediction error for the multi-fidelity emulator is  $\langle |P_F^{\text{pred}}/P_F^{\text{true}} - 1| \rangle \approx 0.8\%$  (averaging across all scales, redshifts, and testing simulation outputs). For the LF single-fidelity emulator, the mean prediction error is  $\approx 4\%$ . This is not unexpected; there are real differences between the flux power spectra output by the low and high resolution simulations that are not being captured with this method. The 4% error may seem quantitatively quite good, considering the simpler methodology and reduced resource cost. However, there is no indication that the error could be reduced further with additional simulations (see Section 4.5.2).

For the HF single-fidelity emulator, the mean prediction error is  $\approx 3\%$ . This is likely limited by the sample size of the training set (6 simulations), leading to increased errors when making predictions for inputs that are far away from the training samples. It is important to note that the HF samples are selected to optimize the multi-fidelity emulator, rather than as an independent emulator (i.e. as a Latin hypercube sample). There is some indication that prior information from the LF training samples provides useful information about the best areas of parameter space to sample. To test this, we split our testing set (10 simulations, same resolution as the HF samples) into training and testing sets, then train all 210 combinations of 6 samples, and predict the outputs for the remaining 4 samples. The

error range from this exercise is 2.5 – 11.5% (5.5% mean error, 1.5% standard deviation). Though not a direct comparison, the 3% error we obtain from the HF single-fidelity emulator compares favorably with this, indicating that the HF samples selected are an improvement over using a Latin hypercube sampling scheme.




Figure 4.7: Comparing the prediction error as a function of (linearly binned) wavevector for multi- and single-fidelity emulators. This is the mean error across all redshifts and 10 test simulations. The shaded regions are the variance in the prediction error. Dashed lines show the highest  $k$  probed by BOSS/eBOSS [182], and the estimated reach for DESI [185].

Figure 4.7 shows the mean prediction error, averaged over all redshifts and 10 test simulation outputs, as a function of wavevector  $k$ . In this, and the following figures, the shaded region around the curves is the variance in prediction error ( $|P_F^{\text{pred}}/P_F^{\text{true}} - 1|$ ), to give a sense of how much the error varies beyond the mean. The multi-fidelity emulator outperforms the single-fidelity emulators at all scales, with an error between 0.5 – 1.5%. The LF (HF) single-fidelity emulator has error between 2 – 7% (2 – 6%).

Both single-fidelity emulators and the multi-fidelity emulator trend towards higher error for small scales. The LF emulator dips 1 – 2% around  $k \approx 0.02$  s/km. The dip occurs on scales at which the low resolution flux power spectra goes from overestimating to underestimating the high resolution power (see Figure 4.2, for  $z \leq 4.6$ ). The uptick in the multi-fidelity emulator error for the largest k-bins is also present in the HF emulator, indicating that there is a scarcity of large scale modes available in the emulator training.




Figure 4.8: Comparing the prediction error as a function of redshift and scale for multi- and single-fidelity emulators. This is mean error across all 10 test simulations. The blue shaded region shows the extent of helium reionization in our simulations (see parameter limits in Figure 4.3).

Figure 4.8 shows how the emulators perform as a function of redshift and scale. In the following, we define small scales as  $k > 0.02$  s/km and large scales as  $k \leq 0.02$  s/km (divided at the smallest scale accessible by BOSS/eBOSS data [182]). On large scales (Figure 4.8, left panel), the LF emulator error decreases with time until  $z \approx 4$ , then slowly increases. This trend is because, as can be seen in Figure 4.2 (left of the BOSS/eBOSS dashed line), the low resolution flux power comes into better agreement with the high resolution flux power as it nears  $z \approx 4$ . We also found that the LF simulations do not cool

as efficiently as the HF simulations after He II reionization, likely leading to the rise in error for  $z \leq 3.2$ .

On small scales (Figure 4.8, right panel), the LF emulator error is more variable. The dip in error seen around  $z \approx 4.6$  occurs as the low resolution flux power crosses from overestimating to underestimating the high resolution flux power (Figure 4.2 right of the BOSS/eBOSS dashed line). The subsequent rise in error is due to the loss of small-scale power and consequent under-estimation of the flux power spectrum in the low-fidelity simulations.

The trends in redshift and scale seen in the LF emulator performance are due not to interpolation error, but to the different numerical resolution of the two simulation fidelities, since this emulator is not predicting the flux power at the higher resolution. Some differences are connected to temperature differences between the LF and HF simulations. For low densities ( $\sim 1$  times the mean density), the LF simulations are colder than the HF simulations at high redshift, but come into better agreement leading up to He II reionization, after which they diverge from the HF simulations again. For higher densities ( $1 - 100$  times the mean density), the LF simulations are once again colder than the HF simulations at high redshift, but at lower redshift they are too hot (with a crossover at  $z \approx 4.6$ ). As higher redshifts probe lower densities, the error initially decreases with redshift, before rising again towards the lowest redshifts.

On both large and small scales the HF emulator errors are around 3.5%, dominated by sampling variance. During He II reionization the HF emulator has more variation in error, which is probably exacerbated by our small box sizes. The increased variation during He II

reionization further indicates that the primary source of error for the HF emulator is the sample size of the training set.




Figure 4.9: Errors for the multi-fidelity emulator, averaged over all test suite simulation predictions, as a function of redshift and scale. The error is largest at large scales and lower redshifts. This is likely a result of the volume of the simulations, and thus a lack of large scale modes for training. Errors also increase for small scales and intermediate redshifts, which is during He II reionization.

On small scales, the multi-fidelity emulator error is insensitive to redshift and small (0.9%). On large scales, the multi-fidelity emulator error slightly decreases until  $z = 4.2$ , then increases with the onset of He II reionization, before flattening again. The trend is also more variable during He II reionization, indicating that emulator finds it more difficult to learn the mapping during this process. However, the multi-fidelity emulator

still outperforms the single-fidelity emulators, with an error between  $\approx 0.4 - 1\%$ . The full picture of the multi-fidelity error can be seen in Figure 4.9, which shows the average (over test suite simulations) error (colors) as a function of the redshift and scale.

### 4.5.2 Emulator Runtime

While we have shown that the multi-fidelity emulator outperforms the single-fidelity emulators presented here, it still remains to show that it is more computationally cost efficient. We could, for example, add more training simulations to our single-fidelity HF emulator and get a similarly accurate high resolution emulator. However, the computational cost would increase significantly. By comparing the total emulator runtime to prediction error, we can determine the choice that balances computational cost and accuracy. In practice, the important question is to determine the computational cost at which a given emulation technique can achieve a desired accuracy. The computational cost of training the emulators is subdominant ( $\mathcal{O}(1)$  cpu-hours) to running the training simulations, so in the following we only consider the runtime for the simulation suites.

Figure 4.10 shows the mean prediction error (averaged over all redshifts and test outputs) as a function of the number of simulations used in the training set, for small and large scales (as defined in Section 4.5 and Figure 4.8). The solid lines show prediction errors for multi-fidelity emulators trained using 6 HF and a varying number of LF simulations. The small and large scale errors flatten out after  $\approx 30$  LF simulations are used in the training set. The LF simulations allow the emulator to determine how the flux power spectrum depends on the cosmological input parameters, and so this indicates that 30 LF




Figure 4.10: Emulator prediction error as a function of the number of simulations used in training the emulator. This is the mean error across all redshifts and 10 test simulation outputs. The prediction error is broken down into large ( $k < 0.02$  s/km) and small ( $k > 0.02$  s/km) scales, as in Figure 4.8. The solid lines show how the average error depends on the number of LF simulations, while the dashed lines show how the average error depends on the number of HF simulations once all LF simulations are included.

simulations are needed to explore our 9 parameter space. Other emulators range from using  $\approx 6$  simulations per parameter (e.g. [195, 166]) to 30 per parameter (e.g. [190]). The 3 – 4 simulations per parameter required here is unusually low, perhaps because the input parameters affect the flux power spectrum close to linearly in much of parameter space.

Dashed lines show prediction errors for multi-fidelity emulators trained using 40 LF and a varying number of HF simulations. Adding extra HF simulations to the training set has a larger impact than adding LF simulations. The addition of each HF simulation generally improves the emulator accuracy for small scales more than for large scales. This



is as expected, since the main purpose of the HF simulations is to learn the mapping from low to high resolution output, with small scales being more resolved in the HF simulations.

Figure 4.11 shows the emulator prediction errors as a function of the total runtime (cost of running the training simulations). All simulations were run on the Frontera super-computer at the Texas Advanced Computing Center. The cost is divided between the LF training simulations, which cost  $\approx 10$  node hours each, and the HF training simulations, which cost  $\approx 150$  node hours each.




Figure 4.11: Prediction error as a function of total training simulation computational cost. This is the mean error for 10 test simulations over all redshifts and scales. The solid line shows the prediction error when changing the number of HF simulations used in training the multi-fidelity emulator. Also shown are the single-fidelity emulators (red squares), and the multi-fidelity emulator trend with a varying number of LF training simulations (yellow, dashed).

The dashed trend shows the same emulators as Figure 4.10, but no longer divided by scale. Qualitatively it looks the same as both the large scale and small scale results from the previous figure. The error is flat after  $\approx 30$  LF training simulations, indicating that a similar accuracy can be achieved using the multi-fidelity emulator with  $\approx 30$  rather than 40 LF simulations. The most efficient 1% error emulator in this study is a multi-fidelity emulator using 30 LF and 5 HF training simulations (the cost for this was  $\approx 1050$  node hours). The most accurate emulator is the 40 LF, 6 HF multi-fidelity emulator, with error 0.8%, and cost  $\approx 1300$  node hours.

The dotted line (squares) shows the error and runtime for the single-fidelity emulators. Following from the 6 HF single-fidelity emulator result to the dashed (yellow) line, it can be seen that the addition of just a few LF training simulations quickly improves the accuracy. It can also be seen that in terms of computational cost, the multi-fidelity emulator is more efficient<sup>6</sup>. Note that the HF training simulations are not selected to optimize a single-fidelity emulator, but instead are selected to optimize the multi-fidelity emulator. They thus use prior information provided by the LF training simulations and so perform better than a naive Latin hypercube construction of a HF emulator using 6 training samples. Our multi-fidelity scheme is thus an even larger improvement on a single-fidelity model than Figure 4.11 suggests.

The solid line shows the error and runtime for the multi-fidelity emulator trained using 40 LF simulations, and 2–6 HF simulations. The point on the solid line corresponding to 40 LF, 2 HF has a similar cost, but slightly worse performance than the 5 HF single-

---

<sup>6</sup>At least for errors less than 4%, the approximate amount by which the LF simulations fail to be converged.

fidelity result. Adding a third HF training sample decreases the error more for the multi-fidelity emulator (error for 40 LF, 3 HF emulator) than it does for the single-fidelity emulator (error marked 6 HF). Adding a 6th HF simulation to the 40 LF, 5 HF multi-fidelity emulator produces a relatively small improvement in error, perhaps indicating that stochasticity in the simulations due to our relatively small box size, is beginning to dominate over interpolation error.

## 4.6 Non-linear Multi-Fidelity Emulator




Figure 4.12: Comparing the prediction error as a function of redshift and scale for linear and non-linear multi-fidelity emulators. This is mean error across all 10 test simulations. The shaded region shows the extent of helium reionization in our simulations (see parameter limits in Figure 4.3).

In the main text we have explored the effectiveness of a linear multi-fidelity emulator (the KO model, or AR1). In the linear model, the mapping from LF to HF is  $f_{\text{HF}}(\boldsymbol{\theta}) = \rho f_{\text{LF}}(\boldsymbol{\theta}) + \delta(\boldsymbol{\theta})$ , where  $f_{\text{HF}}$  and  $f_{\text{LF}}$  are the emulator predictions at those resolutions, and  $\rho$  is independent of the input parameters  $\boldsymbol{\theta}$ .

Here, we compare those results with the results using a non-linear multi-fidelity emulator (non-linear autoregressive GP, or NARGP). In the non-linear multi-fidelity model, proposed by [217], the mapping is a function of both the LF output and the input parameters. We model this as:

$$f_{\text{HF}}(\boldsymbol{\theta}) = \rho(\boldsymbol{\theta}, \tilde{f}_{\text{LF}}(\boldsymbol{\theta})) + \delta(\boldsymbol{\theta}),$$

such that  $\rho$  depends on both the input parameters and LF posterior output. The LF outputs, as is the case with the linear model, are median normalized such that the assumption on the Gaussian process of zero mean is more reasonable,  $\tilde{f}_{\text{LF}}(\boldsymbol{\theta}) = f_{\text{LF}}(\boldsymbol{\theta})/\mu_{\text{LF}} - 1$ .

Following [217],  $\rho$  is modelled as a Gaussian process with input from both input cosmologies for HF,  $\boldsymbol{\theta}_{\text{HF}}$ , and the output from LF,  $\tilde{f}_{\text{LF}}(\boldsymbol{\theta})$ . The NARGP construction results in a deep Gaussian process model [218]. We follow the approximation in [217] and replace  $\tilde{f}_{\text{LF}}(\boldsymbol{\theta})$  with its posterior distribution. Thus, the training reduces to training two regular GPs recursively.

In Figure 4.12 we show the prediction errors separated into small and large scales, as a function of redshift for both the linear and non-linear multi-fidelity emulators. They perform similarly, with the linear emulator being more accurate at low redshifts on all scales, and high redshifts for small scales. The difference between the average error for the linear and non-linear models (over all scales and redshifts) is 0.08%. This is in contrast to emulation of the matter power spectrum in [166], where the non-linear model outperformed the linear model.

It is worth noting that the non-linear model agrees closely with the linear model when using the full suite of training simulations, but lags behind the linear model when using

fewer HF training simulations. For example, the difference in the average error between linear and non-linear models using 40 LF and 3 HF is  $\approx 2\%$  (1.9% error for linear, 3.7% error for non-linear). When using 4 HF the difference is  $\approx 1\%$  (1.5%, 2.5%), and when using 5 HF the difference is  $\approx 0.4\%$  (0.9%, 1.3%). While differences in the effectiveness of the non-linear model may be due to the quantity being emulated (matter power versus flux power), one likely reason for the difference is the number of input parameters. In [166], five input parameters are used, while in this work we use nine. The non-linear model uses the posterior of the LF output, which requires Monte-Carlo sampling. It is possible that the additional dimensions degrade the performance of the Monte-Carlo integration, and thus the performance of the non-linear model. One other reason may be the larger number of hyperparameters that need to be optimized in the training.

## 4.7 Conclusions

We developed and tested a multi-fidelity emulator for the simulated Lyman- $\alpha$  forest flux power spectrum. Emulators address the growing computational demands of simulations, which must be run at increasingly high resolutions to allow analysis of the increasing quality and quantity of observational data. Here, we use a Gaussian process based emulator that addresses this demand by, in a Bayesian framework, training an interpolating function to predict the output (Lyman- $\alpha$  forest flux power spectrum) for a given input (simulation input parameters). Relatively few simulations are required to accurately predict across the span of input parameter space, making emulators especially useful for parameter inference problems.

The multi-fidelity framework allows a further reduction in computational cost by dividing the emulator training samples into multiple (in our case two) fidelities. The low-fidelity (low-resolution) training samples allow the emulator to learn how the the outputs depend on input parameters. The high-fidelity (high-resolution) training samples correct numerical errors in the low-fidelity emulator with a (parameter-dependent) mapping from low- to high-fidelity. Thus, the emulator can be trained with a large sample of low-fidelity training simulations and a small subset of high-fidelity training simulations.

Our training suite included 40 low resolution hydrodynamical simulations (30 Mpc/h simulation box length,  $256^3$  particles) and 6 high resolution hydrodynamical simulations (30 Mpc/h simulation box length,  $512^3$  particles). Using the Lyman- $\alpha$  forest flux power spectrum extracted from these simulations, we trained single- and multi-fidelity emulators to predict the high resolution flux power spectrum. Ten independent simulations were run to test the prediction accuracy of the trained emulators.

In summary, the multi-fidelity emulator:

- Modelled a redshift range  $5.4 - 2$  (in 18 redshift bins with  $\Delta z = 0.2$ ), on scales ranging from  $k = 1.4 \times 10^{-3}$  to  $k = 5.7 \times 10^{-2}$  s/km (in 25 bins).
- Achieved sub-1% error on most scales and redshifts when averaged over 10 test simulations.
- Reached an average error of 0.8%.
- Achieved 1% average error most cost efficiently using a training set with 30 low resolution and 5 high resolution simulations.

The low resolution single-fidelity emulator (4% average error) predicts the low resolution flux power, so it is limited by real differences between the output of the two resolutions. The high resolution single-fidelity emulator (3% average error) is limited by the small number of training samples. It is likely that the average error for the high resolution single-fidelity emulator could be improved to match the multi-fidelity emulator performance with the addition of more training simulations. However, the high resolution single-fidelity emulator quickly increases in computational cost with additional samples, and we expect it would thus be more expensive than our multi-fidelity emulator.

Some important caveats to our results are that the Lyman- $\alpha$  forest is converged at the  $\approx 5\%$  level in our high resolution simulations, and the box size is small. In a forthcoming work, a model that uses two different box sizes (rather than two different resolutions) to construct a multi-fidelity emulator will be developed and tested. While there is no direct evidence to suggest that changing the resolution or box size would significantly enhance or diminish the accuracy of the emulators presented here, it still remains to be tested on simulations with higher resolution and larger box sizes.

## Chapter 5

# Inference via the Lyman- $\alpha$ Forest

### 5.1 Abstract

In this work we expand on, and make use of, our recently developed Lyman- $\alpha$ forest multi-fidelity emulator for cosmological and astrophysical inference. The emulator is expanded to include the mean temperature of the intergalactic medium, in addition to the Lyman- $\alpha$ forest flux power spectrum. Our emulator is trained on large volume  $((120h/\text{Mpc})^3)$  and high particle load (48 simulations using  $1536^3$  & 3 simulations using  $3072^3$ ) simulations, and makes predictions with  $\approx 1\%$  error. The emulator is used within a Markov Chain Monte Carlo framework, with a new likelihood function, to constrain a set of nine parameters that are varied across the simulation suites. We find that the posteriors from our analysis depend strongly on the inclusion or omission of the lowest two redshift bins. When these redshifts are included, we find a preference for a low primordial power index,  $n_P = 0.879 \pm 0.012$ , and a primordial power amplitude,  $A_s = (1.90 \pm 0.13 \times 10^{-9})$ , consistent with Planck. When these redshifts are omitted, we find  $n_P = 0.927 \pm 0.017$ , con-



sistent with Planck, and a primordial power amplitude,  $A_s = (1.76 \pm 0.13 \times 10^{-9})$ , within  $2\sigma$  of Planck. The affect of the redshift range is also seen in the midpoint of reionization, which is  $z^{\text{H I}} = 6.85 \pm 0.28$  using the full redshift range, and  $z^{\text{H I}} = 7.87 \pm 0.25$  when using the reduced redshift range. The timing and duration of helium reionization is consistent between these two, running from  $z_i^{\text{He II}} \approx 4.0 - 2.8$ .

## 5.2 Introduction

Modern cosmological inference requires a constantly evolving set of tools, from improved modeling of cosmology and astrophysics, to new computational tools to keep pace with modern observations. In particular, the march of modern observations towards smaller scales and larger samples require more computationally intensive modeling, i.e. running larger volume and higher resolution simulations. The need for these costly simulations introduces its own problem; fewer of them can be run, reducing their usefulness in inference frameworks. Various methods have been developed to answer this last problem, many of which use machine learning methods.

The Lyman- $\alpha$  forest [1, 157, 158, 159, 160, 92, 161, 162] is an example of a cosmological observable that requires modern techniques to fully exploit. The Lyman- $\alpha$  forest is a series of spectral features, produced by overlapping neutral hydrogen absorption profiles in the spectra of distant luminous quasars as their light is processed through neutral hydrogen in the intergalactic medium (IGM) [1]. In the rest frame of this neutral hydrogen, light that has been redshifted close to the Lyman- $\alpha$  transition at  $1215.67\text{\AA}$  is absorbed. As the light traverses the IGM, it continues to intersect additional neutral hydrogen at lower redshift.

The result is a quasar transmission spectra containing an overlapping field of absorption features, which measures the neutral gas density along that sightline [168].

Specifically, the densities probed by the Lyman- $\alpha$ forest from redshift  $z = 2 - 5$ , are  $\sim 1 - 100 \times$  the cosmological mean density. For these redshifts and densities stellar winds and star formation effects are negligible, though feedback from black holes can be important [169, 170]. As such, many questions can be explored using the Lyman- $\alpha$ forest: the thermal and ionization history of the IGM, the nature of dark matter, the neutrino mass, and alternative cosmologies, to name a few. Some examples of works using the Lyman- $\alpha$ forest include: studies constraining the thermal history of the IGM and reionization [171, 124, 113, 96, 172, 173, 219], and constraining cosmological parameters including the neutrino mass [140, 175, 176, 177, 178, 155, 179], as well as studies testing alternatives to cold dark matter [109, 150, 151, 155, 180, 181].

There are an abundance of quasar spectra from both past and future large surveys, such as the Sloan Digital Sky Survey (SDSS) [182], and the Dark Energy Spectroscopic Instrument (DESI) [220]. DESI will increase the number of Lyman- $\alpha$ forest quasar spectra by a factor of four over SDSS [67]. DESI is also expected to measure the flux power spectrum at smaller scales ( $k \gtrsim 0.035 \text{ km}^{-1} \text{ s}$ ) and higher redshifts ( $z > 4.6$ ) than SDSS, achieving order of a few percent accuracy [185]. In addition to this, there are smaller sample size, higher resolution quasar surveys [152, 183, 184].

These observations bring with them interpretation challenges, requiring increasingly more sophisticated modeling, especially for observations like the Lyman- $\alpha$ forest, which probes scales requiring expensive numerical simulations. Simulations of the Lyman- $\alpha$ forest

must follow the distribution of gas on small scales, hence they must be well resolved hydrodynamics simulations. Specifically, simulations need to have a box side length of at least  $100 \text{ Mpc h}^{-1}$  and mean particle spacing of  $100/3072 \approx 0.03 \text{ Mpc h}^{-1}$  [167]. Because only a limited number of these expensive simulations are achievable, often researchers use methods to enhance the volume and resolution of their simulation outputs (splicing [167], Richardson extrapolation [186]), or use machine learning methods, such as emulators [163, 164, 165]. We will use an emulator, which provides a means of interpolating a summary statistic between simulation outputs for arbitrary cosmological and astrophysical input parameters. Specifically, we use a multi-fidelity emulator, which combines two different resolution training samples in a framework that allows predictions at high resolution.

Parameter inference tasks, such as Markov Chain Monte Carlo analysis, can require up to  $\sim 10^6$  model evaluations. This is prohibitive if each evaluation requires an expensive cosmological simulation, but easily achievable by drawing from the emulator based surrogate. An emulator is trained on a suite ( $\sim 30 - 50$ ) simulation outputs, and uses interpolation produce predictions for arbitrary input with negligible computational cost (once the emulator is constructed). Emulators have been used to study various cosmological probes: the matter power spectrum [165, 187, 188, 189, 190, 191, 192], weak lensing shear [193, 194], the halo mass function [195, 196, 197], and the 21-cm signal [198, 199, 200, 201].

Here, we use a Gaussian process (GP) based emulator following [77, 11]. Gaussian processes are flexible, fast and easy to train, producing good results even for our relatively small training data set. Gaussian processes [202] are a means of interpolating between the simulation outputs via a distribution of functions that is learned through training on

simulations. Draws from the ( $n$ -dimensional) learned distribution are potential predictions, with the mean serving as the actual prediction (best estimate). The prediction uncertainty comes from the variance of the learned distribution. The predictions can be made for arbitrary simulation inputs, within the parameter limits of the training set.

GP emulators have been used for Lyman- $\alpha$ forest summary statistics previously [77, 203, 204, 205, 206, 181]. In particular, [77] constructed a GP based emulator for the simulated Lyman- $\alpha$ forest flux power spectrum using 21 training simulations. This emulator used a single resolution for all of its training simulations. Thus, the emulator from [77] was restricted to volumes and resolutions for which  $\sim 20$  simulations could be generated.

One way to circumvent this problem is to use a multi-fidelity GP emulator model [207]. A multi-fidelity emulator replaces some of the more expensive simulations (high fidelity, HF) that would be needed in a single-fidelity emulator with cheaper simulations (low fidelity, LF). This model was used to emulate the matter power spectrum in [166], and more recently we used this model for the Lyman- $\alpha$ forest flux power spectrum [11]. In [11], the training simulations were split into two fidelities; a sample of 40 low resolution simulations (LF), and a subset of 6 of the LF simulations run at higher resolution (HF). The multi-fidelity emulator used in this work follows the same division, a small sample of HF simulations, supported by a large sample of much cheaper LF simulations.

The multi-fidelity model dramatically reduces the computational cost of constructing an emulator, while retaining predictive power across parameter space. The multi-fidelity model allows the construction of an emulator with eight times the mass resolution (and double the spatial resolution) of our lower fidelity simulations. In [11], the computation cost

was reduced by approximately two-thirds. In this work, the high resolution simulations cost  $\approx 16$  times more computational time than an equivalent simulation run at the low resolution. This means an emulator can be constructed that accesses the full range of scales and redshifts probed by Lyman- $\alpha$ forest observations. Using this emulator in an inference framework allows an analysis that jointly constrain thermal, astrophysical, and cosmological parameters.

Our aim in this work is to perform inference using the high resolution GP emulator constructed in [221] this and current observations of the Lyman- $\alpha$ forest. The emulator used here is similar to that described in [11]. The work here differs from the results of Ref. [11] in the volume and resolution of the simulations used. The volume has been increased by a factor of  $4^3 = 64$ , while the resolution for both the LF and HF sets have been improved by a factor of  $1.5^3 \approx 3.4$ . These puts the simulations in line with recommendations for simulating the Lyman- $\alpha$ forest, with a box side length of 120 Mpc/h, and  $1536^3$  and  $3072^3$  particles for the LF and HF set, respectively [167]. In a companion paper [221], the simulations used in this work are fully described, and the convergence of the Lyman- $\alpha$ forest statistics is tested.

Another major difference from [11] is that we build multi-fidelity GP emulators for two summary statistics of the Lyman- $\alpha$ forest: the Lyman- $\alpha$ forest flux power spectrum and the IGM mean temperature. The observations used in this work are from [182] for the flux power spectrum, and from [173] for the IGM mean temperature. Other works have used the Lyman- $\alpha$ forest to constrain cosmological and/or astrophysical parameters. This has included: constraining the thermal history of the IGM [113, 96, 110, 173, 219], probing warm dark matter models [109, 150, 155, 151, 222] and bounding the neutrino

mass [223, 155]. Using the larger scales measured by correlating absorption from different quasar sightlines has allowed detection of the baryon acoustic oscillations and constraints on the expansion of the universe [224, 225, 226, 227]. Our work has several novel features, including the use of full hydrodynamic simulations, the implementation of a physical model for hydrogen and helium reionization, and the use of a multi-fidelity emulator for sampling.

In summary, our method is: (1) construct a sample that spans the parameter space of interest, then run low fidelity simulations for each of these samples; (2) extract the summary statistics from these simulations, then train a single-fidelity emulator to determine the optimal simulations to run again for the high fidelity set; (3) using the summary statistics from the LF and HF sets, build a multi-fidelity emulator that accurately predicts these outputs for arbitrary inputs; (4) construct a likelihood that makes use of the emulator and observational data to determine the likelihood of the input parameters; (5) use this emulator and likelihood within an MCMC inference framework to constrain the input parameters.

We begin, in Section 5.3, by describing Gaussian processes, and the emulator methods used for both a single- and multi-fidelity emulator model. We then provide details on the simulation suite used for training the emulator in Section 3.3, including the parameters that are varied, the method of sampling those parameters to construct the final suite, as well as the method of calculating the summary statistics used (flux power and mean temperature). In Section 5.5 we outline and discuss the inference scheme, the observational data that is used and the construction of the likelihood. Section 4.5 details the results of running the likelihood in an MCMC framework, and we conclude in Section 4.7.

There are also appendices for extended and alternative results: the full set of

parameter posteriors in Appendix 5.8; a test where we use simulation input as the observation to see how well our likelihood recovers the truth in Appendix 5.9; results from an analysis that includes priors on some of the parameters in Appendix 5.10.1; results using an older observed flux power spectrum (DR9) in Appendix 5.10.2; results where we have inflated the observational errors in Appendix 5.10.3; results using only the mean temperature for inference in Appendix 5.10.4; and results using a reduced range for the scales in Appendices 5.10.5.

MCMC chains for all the results presented in this work along with files containing the training outputs used to construct the emulators<sup>1</sup>, as well as the code<sup>2</sup>, which includes the emulator, likelihood, and integration with the Cobaya MCMC package, are available publicly.

### 5.3 Emulator Methods

In this work, we perform parameter inference using a Markov Chain Monte Carlo (MCMC) method. Specifically we employ the Cobaya package [228, 229]. In order to perform parameter inference in this way, we require a much cheaper sampling scheme than running simulations for each step in the MCMC chain. This is the role of the emulator in the inference scheme.

We use a Gaussian process (GP) based emulator, which is discussed in Section 5.3.1. This allows accurate predictions of the simulated output, along with a measure of prediction uncertainty, without running the  $10^5 - 10^6$  samples required for an MCMC

---

<sup>1</sup><https://github.com/mafern/InferenceLyaData>

<sup>2</sup>[https://github.com/sbird/lya\\_emulator\\_full](https://github.com/sbird/lya_emulator_full)

analysis to converge. The methodology of the emulator is largely the same as in [11]. In addition to a brief discussion of both single- and multi-fidelity emulator models, any differences from [11] are discussed in Section 5.3.2 and Section 5.3.3. The two emulator models make use of our low fidelity (LF) and high fidelity (HF) training samples, which come from the low and high resolution simulation suites discussed in Section 3.3.

### 5.3.1 Gaussian Processes

Gaussian processes [202] are a trainable method for providing function prediction in a Bayesian framework. In this case, the GP is used to interpolate simulation outputs as a function of the simulation input parameters that are varied in the training samples (see Figure 5.4). In a Gaussian process, a distribution of functions is learned through the training samples (simulation input/output pairs), and the mean of the learned distribution serves as the prediction, while the variance of the distribution quantifies the prediction uncertainty. This allows a prediction to be cheaply produced for arbitrary simulation inputs (within the parameter space of the training samples – GPs work poorly for extrapolation). We use the GPy [230] package to implement Gaussian processes within our emulator model.

If  $P_F(\boldsymbol{\theta})$  is the simulated Lyman- $\alpha$  forest flux power spectrum (the following also applies to the IGM mean temperature,  $T(\boldsymbol{\theta})$ ) as a function of  $\boldsymbol{\theta}$ , the input parameters to the simulation, then we are modeling this output as draws from a learned distribution,

$$P_F(\boldsymbol{\theta}) \sim GP(\mu(\boldsymbol{\theta}), k(\boldsymbol{\theta}, \boldsymbol{\theta}')), \quad (5.1)$$

where  $\mu(\boldsymbol{\theta})$  and  $k(\boldsymbol{\theta}, \boldsymbol{\theta}')$  are the mean and covariance function for the distribution, respectively. As in [11], we rescale the training samples by their median (the LF median for both



single- and multi-fidelity emulators) so that we may assume a zero mean function during the GP training. The predictions are then multiplied by this rescaling factor to produce sensible predictions. With a zero mean function, it only remains for the covariance function to be trained, which amounts to selecting a covariance kernel and training the free parameters of that kernel. Here, as in [11], we use the sum of a radial basis function (RBF) kernel and a linear kernel to model our outputs.

The total kernel (RBF + linear kernels) is given by:

$$k_{\text{RBF}}(\boldsymbol{\theta}, \boldsymbol{\theta}'; \sigma_0, \boldsymbol{l}) + k_{\text{LIN}}(\boldsymbol{\theta}, \boldsymbol{\theta}'; \boldsymbol{\sigma}) = \sigma_0^2 \exp\left(\sum_{i=1}^d -\frac{(\theta_i - \theta_i')^2}{2l_i^2}\right) + \sum_{i=1}^d \sigma_i^2 \theta_i \theta_i', \quad (5.2)$$

where  $d$  is the dimensionality of the input parameters. The hyperparameters that are learned from the training samples are: variances for the RBF ( $\sigma_0^2$ ) and linear ( $\boldsymbol{\sigma}^2$ ) kernels, and the lengthscale controlling the smoothness of the RBF kernel,  $\boldsymbol{l}$ . The RBF lengthscale and linear variance hyperparameters are shown as vectors here, indicating that an independent value is assigned for each dimension  $d$  of the input for each of these hyperparameters. This is called Automatic Relevance Determination (ARD), and we use this in all cases for the LF Lyman- $\alpha$  forest flux power spectrum. Including ARD allows the GP to train a separate hyperparameter for each dimension of the input, such that dimensions that have a more sensitive effect on the prediction can be given a smaller scale, e.g. the effect of  $A_p$  on the mean temperature is negligible, while  $\alpha_q$  has a large effect. For the IGM mean temperature GP, we only use ARD for the RBF kernel. Unlike for the flux power predictions, including ARD for the linear variance hyperparameter had no effect on the accuracy of the mean temperature predictions.

The prediction and predicted uncertainty for a Gaussian process come from the

mean and variance of the learned distribution, and can be written as the following closed-form expressions (for a full derivation, see [202]):

$$\begin{aligned}\mu(\boldsymbol{\theta}) &= \mathbf{K}(\boldsymbol{\theta}, \mathcal{D})^\top \mathbf{K}(\mathcal{D})^{-1} \mathbf{P}_F(\mathcal{D}); \\ \sigma^2(\boldsymbol{\theta}) &= k(\boldsymbol{\theta}, \boldsymbol{\theta}) - \mathbf{K}(\boldsymbol{\theta}, \mathcal{D})^\top \mathbf{K}(\mathcal{D})^{-1} \mathbf{K}(\boldsymbol{\theta}, \mathcal{D}).\end{aligned}\tag{5.3}$$

where  $\boldsymbol{\theta}$  are the new inputs for which a prediction is desired,  $\mathcal{D}$  is the vector of training inputs,  $\mathbf{P}_F(\mathcal{D})$  ( $\mathbf{T}(\mathcal{D})$  for the mean temperature) is the vector of training outputs,  $\mathbf{K}(\mathcal{D})$  is the training data covariance matrix, and  $\mathbf{K}(\boldsymbol{\theta}, \mathcal{D})$  is the vector of covariances between the training data and new inputs.

In the following, we discuss the single- and multi-fidelity emulator models we use for our inference framework. The code used to generate training samples, train and test single- and multi-fidelity emulators, and perform inference via MCMC is publicly available at: [https://github.com/sbird/lya\\_emulator\\_full](https://github.com/sbird/lya_emulator_full).

### 5.3.2 Single-Fidelity Emulator

The single-fidelity emulator used in this work follows the prescription outlined in Section 5.3.1. Training samples extracted from the suite of LF simulations (see Section 3.3) are used to train the GP (both for the flux power spectrum and the mean temperature). For the mean temperature, a single GP is sufficient, as there is a single value per redshift. For the flux power, a GP is trained for each redshift, but across the full  $k$  range. This differs from [11], which trained a GP for every  $k$  bin for every redshift. In the case of a single-fidelity emulator, the accuracy was not substantially improved by training GPs for each  $k$  bin. All results using the single-fidelity emulator use only the LF simulations for training, thus the predictions are for the lower resolution flux power spectrum and mean

temperature. With only three HF samples, a single-fidelity emulator for the HF outputs is dominated by prediction errors.

To measure the accuracy of the single-fidelity emulator predictions, we employ a method called leave-one-out. In turn, each training sample is left out of the training set, the emulator is then trained with this reduced set, and a prediction is made for the omitted sample. The relative error between the prediction and true output from the omitted sample is then retained and this is repeated for each training sample. Due to the mean flux rescaling, which produces ten samples for each simulation, the samples that are omitted and then predicted are all those associated with each simulation, rather than each sample. The distribution of relative errors from this calculation are shown for the flux power spectrum in Figure 5.1, and for the mean temperature in Figure 5.2. These figures show the errors for a single-fidelity emulator that is predicting LF outputs (grey), and the errors for a multi-fidelity emulator that is predicting HF outputs (yellow, discussed in Section 5.3.3).

In the left panels of Figures 5.1 and 5.2, the absolute error normalized by the prediction uncertainty is shown. For the distributions with large enough sample sizes, these appear to be normally distributed, as expected. The right panels show the absolute relative error, a good indicator of the accuracy of the emulator predictions. The flux power single-fidelity emulator has an average error of 0.7%, with a maximum of 8%. The mean temperature single-fidelity emulator has an average of approximately 1%, with a small number of errors at, or slightly above 10%.

This is consistent with [77], which constructed an emulator for the Lyman- $\alpha$  forest flux power spectrum from 21 simulations and found an error of 1 – 2%. In [77], the




Figure 5.1: Leave-one-out errors for the Lyman- $\alpha$  forest flux power spectrum emulator. The left panel shows the absolute error in units of prediction uncertainty. The right panel shows the relative errors, an estimate of how well the emulator predicts the flux power spectrum. The distributions show the prediction error for each redshift of each flux power sample included in the final training sample. The grey distribution shows the emulator error for a single-fidelity emulator, predicting the flux power for the LF simulations. The yellow distribution shows the errors for a multi-fidelity emulator, predicting the flux power for the HF simulations.

simulations are smaller volume and lower resolution, cover a smaller redshift range than that used here, and use different physics models (simplified star formation, no helium reionization model). Thus, the rough agreement for emulator accuracy is encouraging, as the emulator framework works well for fairly different sets of training samples.

### 5.3.3 Multi-Fidelity Emulator

For the single-fidelity emulator, the training samples are all from the LF suite, thus the predictions correspond to the lower resolution simulations which produced the LF training samples. Rather than run a similar number of significantly more expensive simulations at a higher resolution, we employ the multi-fidelity model which requires only

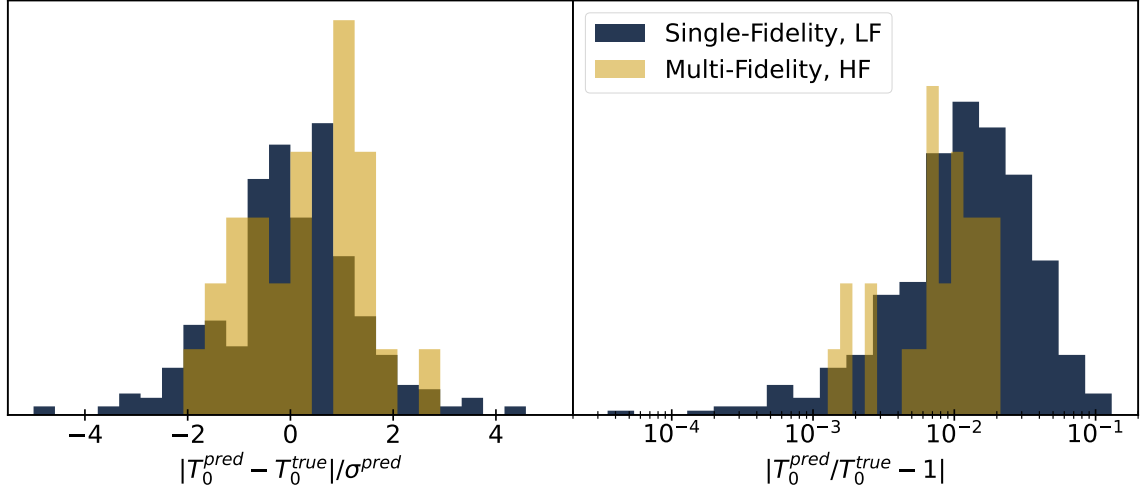


Figure 5.2: Same as Figure 5.1, but for the IGM mean temperature emulator. The distributions show the prediction error for each redshift of each simulation included in the final training sample.

a few HF training samples to make predictions of the higher resolution outputs.

Previously, we developed and tested multi-fidelity emulators for the Lyman- $\alpha$  forest flux power spectrum [11]. Essentially, the multi-fidelity model allows simulations with different particle loads, and thus costs, to be combined together. Thus, some of the HF simulations that would be needed in an HF single-fidelity emulator are replaced with LF simulations. This allows construction of an emulator that makes full use of the range of scales and redshifts probed by Lyman- $\alpha$  forest observations, at a significantly reduced computational cost (e.g. each HF simulation requires approximately sixteen times more node-hours to run than an LF simulation).

The multi-fidelity model on which our implementation is based is from [207]. Our two fidelities correspond to simulations run at two different resolutions (see Table 4.1): lower resolution simulations for the low fidelity (LF) samples, and higher resolution simulations for the high fidelity (HF) samples. The multi-fidelity prediction for the HF outputs (shown here

for the Lyman- $\alpha$ forest flux power spectrum, but equally valid for the mean temperature) is given by:

$$P_F^{\text{HF}}(\boldsymbol{\theta}) = \rho \cdot P_F^{\text{LF}}(\boldsymbol{\theta}) + \delta(\boldsymbol{\theta}), \quad (5.4)$$

where  $\rho$  is a parameter, and  $\delta(\boldsymbol{\theta})$  is a GP (independent of the LF output), both of which are optimized using the training samples. Since  $\rho$  is a multiplicative correction, and  $\delta(\boldsymbol{\theta})$  is an additive correction, this is often called a linear multi-fidelity model. We implement our multi-fidelity models using the package Emukit [231].

As in the single-fidelity emulator, the LF and HF inputs are normalized by the LF median (the HF sample is too small to obtain a reliable median). In our multi-fidelity model, ARD is used for the LF kernels, but not the HF kernels, due to the smaller number of HF samples. Once again a single GP per redshift is used for the Lyman- $\alpha$ forest flux power spectrum, while a single GP for the full redshift range is used for the mean temperature.

In the multi-fidelity emulator constructed in [11], the Lyman- $\alpha$ forest flux power spectrum was modeled with a separate GP for each  $k$  bin, and achieved sub-1% accuracy using 30 LF and 5 HF simulations. In that work, the emulator accuracy was worst for  $k > 0.02$  s/km and  $k < 0.002$  s/km. The simulations used in this work have eight times the volume, with an improvement in the mass resolution of more than three times, as compared to the simulations from [11]. The improved volume and resolution applies to both the LF and HF suites used for training in this work. In particular, the LF simulations are much more nearly converged on their own. Due to these improvements, we expect our Lyman- $\alpha$ forest flux power emulator to be only marginally affected by the difference in HF training sample size (we use 3 HF in this work) and the choice to use a single GP per redshift bin.

Additionally, though the emulator presented here supports scales smaller than  $k = 0.02$ , the observations we use only extend to  $k = 0.02$ . The reason we opt for a single GP per redshift bin is purely economical. Using a GP for each  $k$  bin leads to very large memory usage for the emulator, and substantially slows down the training and prediction steps in the inference framework.

Figures 5.1 and 5.2 show the leave-one-out errors for the multi-fidelity emulator (yellow). The error is approximately 1% for the multi-fidelity emulator predictions, similar to the single-fidelity emulator. Note that because there are only three HF simulations, the multi-fidelity emulator error distributions have much smaller sample sizes, and have been scaled for clarity. These errors are consistent with [11], though in that work the errors were calculated using a suite of test samples, separate from the training samples, rather than the leave-one-out method used here.

We quantify this by returning to the emulator from [11], and calculating the leave-one-out errors for that suite of training simulations. The errors using a test suite of simulations were on average  $\approx 0.8\%$ . The leave-one-out errors are on average  $\approx 0.8\%$ . The very marginal difference here indicates that the errors for the emulator used here, were they calculated using a test suite of simulations, would still be  $\approx 1\%$ . Regardless, the emulator accuracy indicates that the emulators are performing well, with potential errors significantly smaller than the uncertainty associated with the observational data used in our inference scheme (for the data set we use [182], the uncertainty is on average  $\approx 7\%$ ).

Simulation	Box Volume	$N_{\text{gas}}$	$M_{\text{gas}} (M_{\odot} \text{ h}^{-1})$
LF	$(120 \text{ Mpc h}^{-1})^3$	$1536^3$	$[5.29, 6.98] \times 10^6$
HF	$(120 \text{ Mpc h}^{-1})^3$	$3072^3$	$[6.73, 7.97] \times 10^5$

Table 5.1: Low-Fidelity (LF) and High-Fidelity (HF) simulation suite details.  $N_{\text{gas}}$  is the number of gas particles simulated,  $M_{\text{gas}}$  is the resulting mass resolution of those particles.

## 5.4 Simulation Suite

Simulations were performed using MP-Gadget<sup>3</sup>, an N-body and smoothed particle hydrodynamics (SPH) code. MP-Gadget is based on Gadget-3 (last described in [78]), but has been modified to include shared-memory parallelism using OpenMP, and various algorithmic improvements and new subgrid models [208, 209, 210]. The initial power spectrum and transfer functions are generated with the Boltzmann code CLASS [79]. Species-specific initial conditions are generated separately for baryons and dark matter [210, 10].

Simulations are initialised at  $z = 99$  and finish at  $z = 2.2$ . Box volume (simulations use periodic boundaries), number of gas particles, and gas particle mass resolution are reported in Table 5.1. The range given for the gas mass resolution is due to the varying value of  $h$  in our simulation suite ( $\Omega_b h^2$  is fixed, at a value of 0.0224). In [141], the recommended gas mass resolution for the Lyman- $\alpha$  forest is  $2 \times 10^5 M_{\odot} \text{ h}^{-1}$  at  $z = 5$  and  $1.61 \times 10^6 M_{\odot} \text{ h}^{-1}$  at  $z = 2$ . Our HF simulations have better gas mass resolution than the recommendation for lower redshifts, and slightly worse resolution than recommended at higher redshifts. However, we include a temperature boost [107] due to impulsive heating during hydrogen reionization from ionization fronts. This increase the thermal free-streaming length at high redshift, which reduces the resolution requirements. Also, we currently restrict our analysis to  $z = 2.2$  to  $z = 4.6$ , so the recommendation at  $z = 5$  is more strict than required in this

<sup>3</sup><https://github.com/MP-Gadget/MP-Gadget>



work.

In [167], simulations with box side length of at least  $100 \text{ Mpc } h^{-1}$ , and a corresponding gas mass resolution of  $4.3 \times 10^5 M_{\odot}/h$  were recommended for the Lyman- $\alpha$  forest. Our HF simulations are larger volume than the above recommendation, with slightly lower gas mass resolution. In [167], they found that simulations run with a similar resolution to our HF samples were no worse than 5% converged. When the reduced resolution requirement due to the aforementioned temperature boost is considered, the simulations used in this work simultaneously achieve the volume and resolution recommended in [167], without the need for splicing.

Simulation physics models are described in the companion paper to this work [221]. The models are the same as in [11], which largely follow those used for the **ASTRID** simulation, which are described in more detail in [208].

One note on the black holes, which follow the model of [209]. The black hole feedback factor,  $\epsilon_{AGN}$ , which controls the fraction of luminosity that is converted to thermal energy, is varied across the simulation suite. However, since the constraints on this parameter are relatively weak, we add a Gaussian prior on the parameter, which dominates the posterior for  $\epsilon_{AGN}$ . Instead, we apply a Gaussian prior centered on the middle of the parameter limits for  $\epsilon_{AGN}$ , and do not attempt to constrain the black hole feedback.

An example of the gas density and temperature (colors) at  $z = 4$  are shown in the top (high resolution) and bottom (low resolution) panels of Figure 5.3. The middle panel of Figure 5.3 shows the Lyman- $\alpha$  forest spectra for the skewers passing through the middle of the top and bottom panels.




Figure 5.3: Example Lyman- $\alpha$  forest spectra and corresponding gas density and temperature (colors) from an LF and HF simulation at redshift  $z = 4$ . The top and bottom panel show the simulated gas surrounding the skewer which produced the spectra shown in the middle panel. Examples are shown for simulations run at high (top panel, yellow line) and low resolution (bottom panel, red line).

#### 5.4.1 Cosmological & Astrophysical Parameters

Figure 5.4 shows the parameters that are varied across our suite of simulations, as well as their limits. The cosmology parameters include the scalar spectral index (slope),  $n_p$ , and amplitude,  $A_p$  [77]. Also varied are the Hubble constant through  $h$ , and the total matter density through  $\Omega_M h^2$ , though,  $h$  is poorly constrained by the Lyman- $\alpha$  forest. The upper limit for  $h$  is slightly lower than the most recent value measured by the local distance ladder (see Figure 2 in [232]), which has increased since we planned our emulator.

The astrophysics parameters include three to control the He II reionization model [139]:  $z_i^{\text{He II}}$  and  $z_f^{\text{He II}}$  are the redshifts for the start and end of He II reionization, and  $\alpha_q$  is the quasar spectral index (which scales the peak temperature during He II reionization).  $z^{\text{H I}}$  is the midpoint redshift of H I reionization. Finally,  $\epsilon_{AGN}$  is the black hole feedback

factor.

The He II reionization model also takes in the mean and variance of the He III bubble size, which we set to 20 and 0 Mpc, respectively. This differs from that found in [3] (approximately 35 Mpc mean, 10 Mpc variance). However, the effect from changes to the mean bubble size or variance (these are partially degenerate with each other) is degenerate with the other He II reionization parameters (namely  $\alpha_q$ ). We checked the effect of setting the variance to zero by running simulations (30 Mpc/h box side length,  $256^3$  particles) with a mean quasar bubble size of 10 Mpc with variances of 5 Mpc and zero. Setting the bubble size variance to zero decreased the temperature of the IGM by  $\approx 1000$  K at redshifts  $z < 3$ , which has the same affect as raising  $\alpha_q$ .

We have expanded the range covered by the  $\alpha_q$  parameter after calculating posterior constraints, to ensure that our emulator covers the full range of the posterior.

There are two further parameters that are varied during the post-processing of the spectra extracted from the simulations. When calculating the Lyman- $\alpha$  forest flux power spectrum from the set of simulated spectra, the mean flux is rescaled based on [146], which found that the effective optical depth scales with redshift as  $\tau_{\text{H I}}^{\text{eff}} = (0.0023 \pm 0.0007) \times (1 + z)^{3.65 \pm 0.21}$ . The mean flux rescaling implemented here generalizes this to  $\tau_{\text{H I}}^{\text{eff}} = 0.0023 \times \tau_0 \times (1 + z)^{3.65 + d\tau_0}$ . Each simulated flux power spectrum is post-processed into ten flux power spectra corresponding to unique values for the slope ( $d\tau_0$ , values between  $-0.4$  and  $0.25$ ) and amplitude ( $\tau_0$ , values between  $0.75$  and  $1.25$ ) in the above expression. Therefore, the final set of flux power spectra is ten times the number of simulations, or 480 LF, and 30 HF simulated flux power spectra.




Figure 5.4: Parameter limits for each of the nine parameters varied in the simulation suites. Samples for both the LF (grey crosses) and HF (red circles) are also shown. The blue boxes indicate two LF simulations which are included in the flux power spectrum emulator, but not in the mean temperature emulator (due to the loss of the simulation outputs before the mean temperatures were extracted).

### 5.4.2 Sample Design

The LF simulation parameters were selected using either a Latin hypercube sampling scheme, or using Bayesian optimization. The 48 LF simulation parameters were selected in the following manner: an initial Latin hypercube sample of 30 simulations was generated and run, then Bayesian optimization (see the description below of how we implemented this) was used to select 2 further simulations, after which a second Latin hypercube sample of 8 was run, followed by 2 more Bayesian optimization samples, and finally 6 samples from a Latin hypercube to extend the range for the  $\alpha_q$  parameter. The final set of LF samples is shown in Figure 5.4, indicated by the grey crosses.

Figure 5.4 also shows two LF simulations, indicated by blue boxes, which are used in the flux power emulator, but not in the IGM mean temperature emulator. The particle data for these simulations was lost before the mean temperature was extracted from them.

We found that the Bayesian optimization samples were no more effective at improving the emulator accuracy than the Latin hypercube samples, thus we generated only four samples using the Bayesian optimization method. The four Bayesian optimization samples were selected using a method similar to [203]. An emulator trained using the current sample of LF simulations is used to calculate an acquisition function for arbitrary input parameters. This acquisition function balances an exploitation term (preferring regions of high likelihood) with an exploration term (preferring regions with high uncertainty). The likelihood used to compute the exploitation is the same used for our inference, which is described in Section 5.5.3. The uncertainty for the exploration term is the emulator prediction uncertainty.

The parameters that maximize the acquisition function are selected for the next sample, and the process can then be iterated, either by running the sample and adding it to the LF set, or by using the emulator prediction for the new sample (this is called the batch mode for Bayesian optimization). In addition to the lack of improvement over Latin hypercube samples, we found that unless the exploration term was significantly down-weighted, our implementation of Bayesian optimization always preferred certain parameters to be at the edge of their range, where the uncertainty was highest. Effectively this means that the Bayesian optimization samples we use here were run in pure-exploitation mode, primarily selected as the parameters with the highest likelihood.

The red circles shown in Figure 5.4 are the HF samples. These are specially chosen LF samples that we ran at higher resolution ( $3072^3$  particles) to construct the multi-fidelity emulator. The HF samples were chosen by training a single-fidelity emulator using all combinations of two LF samples (excluding the aforementioned simulations without mean temperature data) to predict the remaining LF samples for both the flux power and mean temperature [166]. The two LF simulations which produced the most accurate predictions (defined as minimizing the loss,  $(pred - true)^2 / true^2$ ) for the combination of the flux power and mean temperature were then run at the higher resolution. These first two HF samples were selected prior to the LF simulations that were run to expand the  $\alpha_q$  parameter. The final HF sample was chosen using the same method, but conditioned on the previously selected two samples. Each remaining LF sample was added to the set of two previously selected samples, an emulator was trained using each of these three sample combinations, then the loss for each combination was calculated. Using both the loss from this method,

and the value of  $A_p$  for each sample (among the varied parameters,  $A_p$  has by far the strongest effect on the amount of computational time the simulation requires to complete), we selected a third sample to run at HF that had low loss and low  $A_p$ . This final HF sample was chosen from the full set of LF samples (after the  $\alpha_q$  expansion simulations were added).

### 5.4.3 Flux Power Spectrum

The Lyman- $\alpha$  forest flux power spectrum is calculated using Lyman- $\alpha$  forest absorption spectra generated using Fake Spectra Flux Extractor [142]<sup>4</sup>, described in [143]. For each snapshot (from  $z = 4.6$  to  $z = 2.2$  in increments of  $\Delta z = 0.2$ ) of each simulation, 32,000 randomly placed skewers are generated, along which the Lyman- $\alpha$  absorption spectra is calculated. The random placement is seeded, such that the same skewers (orientation and location) are used for every simulation. The pixel resolution for these spectra is set to  $10 \text{ km s}^{-1}$ . Example spectra from one LF and one HF simulation (at  $z = 4$ ) are shown in the middle panel of Figure 5.3.

Before using these spectra to calculate the flux power, we set an optical depth threshold of  $\tau < 10^6$ , to eliminate damped Lyman- $\alpha$  systems (DLAs). This removes spectra that contain any optical depth value greater than the threshold. In order to determine the appropriate optical depth threshold for these simulations, we fit a sample of Lyman- $\alpha$  spectra with Voigt profiles and returned the maximum optical depth and column density for each. A column density of  $\approx 10^{20} \text{ cm}^{-2}$  corresponded to the aforementioned value,  $\tau \approx 10^6$ .

This filter removes approximately 2% of the spectra at  $z = 2.2$ , and 4% at  $z = 4$

---

<sup>4</sup>[https://github.com/sbird/fake\\_spectra](https://github.com/sbird/fake_spectra)

(obviating the need for a resampling scheme to keep a sufficient number of spectra). This is roughly consistent with [233] (see their Table 1, small and large DLAs), which ranged between  $\approx 3\%$  and  $\approx 13\%$  for  $2 < z < 4.43$ . For a lower threshold associated with Lyman limit systems ( $\tau_{max} < 10^3$ ), the filter removes 10% at  $z = 2.2$  and 40% at  $z = 4$ . This is also consistent with [233]. We use the DLA level threshold and apply an LLS/DLA correction in the likelihood (described in Section 5.5.3) to handle LLS and any over/under-subtraction of DLAs.

The filtered Lyman- $\alpha$ absorption spectra are used to construct the Lyman- $\alpha$ forest flux power spectrum, defined as  $P_F(k) = |L^{-1}\tilde{\delta}_F^2(k)|$ , where  $\tilde{\delta}_F^2(k)$  is the Fourier transform of the flux excess,  $\delta_F(k) = F(k)/\langle F(k) \rangle - 1$ , and  $L$  is the length of the sightline. The flux power spectra used here are averaged over all spectra that remain from the initial set of 32,000, for each redshift of each simulation.

Figure 5.5 shows flux power spectra from all the LF (blue, solid) and HF (gold, dashed) simulations at several redshifts. Shown below each redshift panel in red are the ratios of the flux power spectra for simulations that were run at both resolutions. Note that there are ten times more flux power spectra shown here than simulations, due the previously discussed post-processing method of mean flux rescaling.

#### 5.4.4 IGM Mean Temperature

The IGM mean temperatures are extracted directly from the simulation snapshots. First, the temperature and density for all the gas particles in the simulation are retrieved, then all particles that are within 5% of the critical density are retained. The median






Figure 5.5: Flux power spectra from all LF (blue, solid) and HF (gold, dashed) simulations at  $z = 2.2$ ,  $z = 3.0$ ,  $z = 3.8$ , and  $z = 4.6$ . In each panel below the flux power, the ratio of the flux power spectra for simulations run at both resolutions is shown,  $LF/HF$ .




Figure 5.6: IGM mean temperatures from all LF (blue, solid) and HF (gold, dashed) simulations for the full redshift range. The bottom panel shows the ratio of the mean temperatures from simulations that were run at both resolutions,  $LF/HF$ .

temperature of these retained particles is used as the IGM mean temperature.

Figure 5.6 shows the mean temperatures for all the LF (blue, solid) and HF (gold, dashed) simulations for the full redshift range used in this work. The bottom panel shows the ratio of mean temperatures from the simulations that were run at both resolutions.

All of the Lyman- $\alpha$  forest flux power spectra, IGM mean temperatures, trained emulators, as well as select MCMC chains are available at: <https://github.com/mafern/InferenceLyaData>. The Lyman- $\alpha$  forest spectra, extracted from the simulation and used to calculate the flux power, are available upon reasonable request.

## 5.5 Inference Scheme

Now that we have a  $\approx 1\%$  accurate emulator for prediction of the simulated Lyman- $\alpha$  forest flux power and IGM mean temperature, we can conduct inference by comparing these to their corresponding observations. To accomplish that, we use the code Cobaya [229, 228] to conduct inference via Markov Chain Monte Carlo (MCMC) methods. The overall inference scheme is:

1. Train a single- or multi-fidelity GP emulator on our LF and HF simulation training samples for the Lyman- $\alpha$  forest flux power spectrum, and IGM mean temperature.
2. Use the emulator to predict the two observables for arbitrary input parameters (see Section 5.4.1 for the set of parameters that are varied).
3. Calculate a likelihood comparing these predictions to their observational counterparts.
4. Use the calculated likelihoods for an MCMC analysis.

In Section 5.5.1, the observational Lyman- $\alpha$  forest flux power spectrum data [182] is discussed, then the IGM mean temperature data [173] is discussed in Section 5.5.2. Details of the likelihood calculation used in the MCMC sampling is discussed in Section 5.5.3.

### 5.5.1 Flux Power Spectrum Data

We use the observed Lyman- $\alpha$  forest flux power spectrum from [182], which is based on the Baryon Oscillation Spectroscopic Survey (BOSS) and extended-Boss (eBOSS) quasar samples [234, 235]. In [182], the BOSS/eBOSS quasar samples are refined to remove spectra that have not been visually inspected, and to remove spectra with broad absorption

lines. Sky lines and damped Lyman- $\alpha$  absorbers (DLAs) are masked. The sample of Lyman- $\alpha$  forests from the set of remaining quasar spectra is then further refined based on cuts to the spectral resolution, signal-to-noise ratio, number of masked pixels, and forest length. Along with a windowing function to correct for the instrument response, an estimate for the noise and metal power (estimated using longer wavelength segments of the quasar spectra) are then subtracted to obtain the final Lyman- $\alpha$  forest flux power spectrum. Note that while DLAs are masked in the pipeline of [182], we employ a correction for damped systems in our likelihood (see Section 5.5.3).

The redshifts and scales covered by these observations set the redshift range and scales we use in our emulator, namely  $z = 2.2 - 4.6$  (redshift bin size of  $\Delta z = 0.2$ ), and  $k \approx 0.001 - 0.02 \text{ km}^{-1} \text{ s}$  (over 35 bins, spacing of  $5.42 \times 10^{-4} \text{ km}^{-1} \text{ s}$ ). The average uncertainty in the Lyman- $\alpha$  forest flux power from [182] ranges from  $\approx 6\%$  at low redshifts or large scales, to  $\approx 16\%$  for high redshifts or small scales.

There are higher resolution, smaller sample survey based Lyman- $\alpha$  forest flux power measurements [152, 183, 184]. These Lyman- $\alpha$  forest flux power measurements generally access a similar redshift range to [182], but shift both the largest and smallest scales to higher  $k$ . For example, [183] used spectra from multiple surveys (XQ-100, KODIAQ, and SQUAD) to measure the Lyman- $\alpha$  forest flux power at redshifts  $z = 2 - 4.6$  and scales  $k \approx 0.005 - 0.1 \text{ km}^{-1} \text{ s}$  (with mostly higher uncertainties than [182]). Combining these higher resolution observations with the larger sample size observations will allow inference based on a broader range of scales than either data set allows separately. For now, we focus on the data set from [182], and leave the higher resolution data sets for future work.

### 5.5.2 IGM Mean Temperature Data

We use the mean IGM temperatures from [173], which are derived from simulation modelling of high resolution quasar spectra from the KODIAQ survey [236]. In [173], their co-added quasar spectra are manually checked, and the sample is refined to remove spectra that do not contain Lyman- $\alpha$ forest, do contain DLAs or sub-DLAs, have large gaps in the spectra, or do not meet a signal-to-noise ratio cut. A cut is also made to remove regions of the spectra that may lie near to a quasar proximity zone. Metal lines are removed by fitting Voigt profiles to the spectra and removing features with Doppler widths  $b \leq 8 \text{ km s}^{-1}$ , which are assumed to be metal lines.

To derive temperatures from the observed quasar spectra, [173] calculated several summary statistics and compared them to those derived from spectra drawn from simulations. The simulations they used were similar in resolution to our HF suite (gas mass resolution of  $\sim 10^5 M_{\odot}$ ), though much smaller in volume (10 Mpc/h box side length).

The redshifts covered by these observed IGM mean temperatures is  $z = 2.0 - 3.8$  (in  $\Delta z = 0.2$  bins), though we do not use the mean temperature from the  $z = 2.0$  bin, for consistency with the available Lyman- $\alpha$ forest flux power data. The average uncertainty for this data set is  $\approx 10\%$ , once again larger than the prediction uncertainties from our emulator. Note that [173] provide mean temperatures derived from four different statistics: the Lyman- $\alpha$ forest flux power spectrum, curvature, wavelet decomposition, and Doppler width distribution. We use the Lyman- $\alpha$ forest flux power spectrum derived temperatures in the main body of this work, but show results using these other data sets in Appendix 5.10.4.

Because these observed mean temperatures are themselves derived using a suite

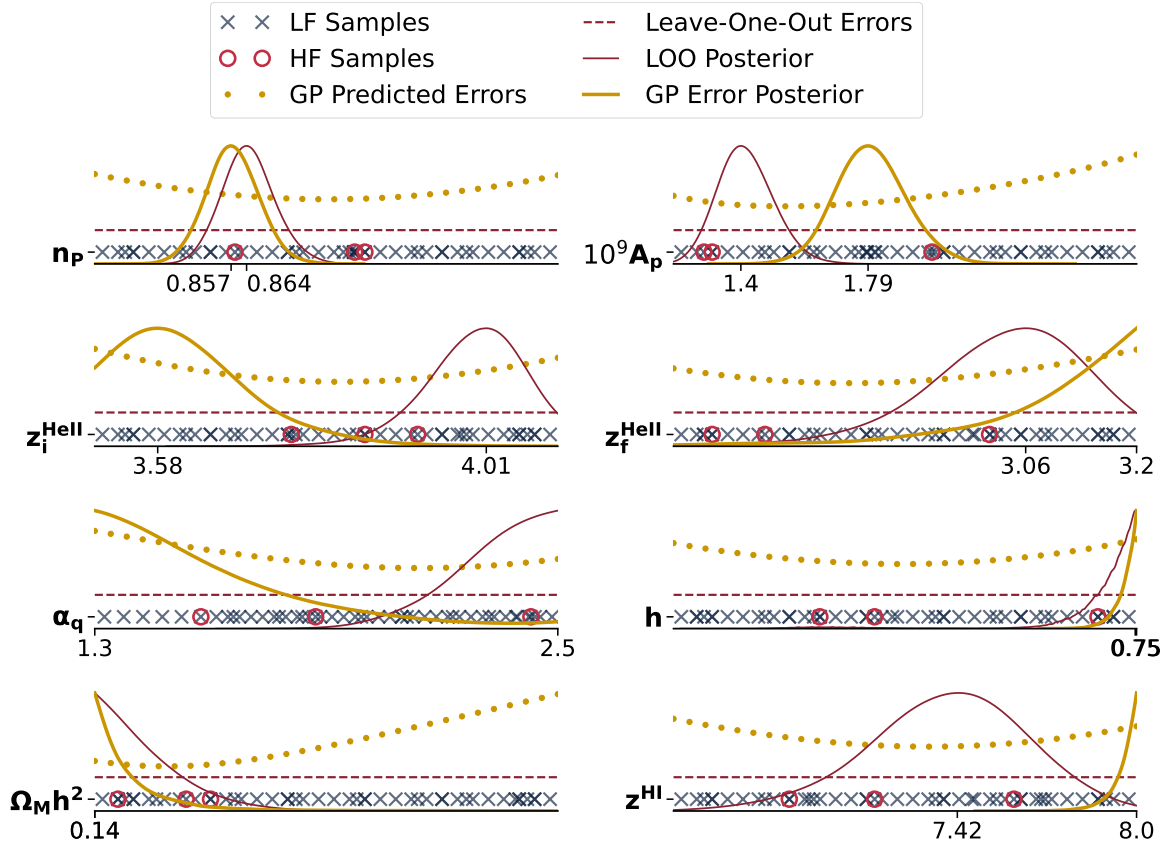


Figure 5.7: Emulator error and leave-one-out errors across parameter space. For eight of the input parameters, the training samples (grey crosses for LF, red circles for HF), GP emulator errors (yellow dots), leave-one-out errors (red dashed) are shown. The flux power only chains shown were run using the GP emulator error in the likelihood (thick, yellow) and using the leave-one-out errors in the likelihood (thin, red). Maximum posterior parameters are labeled for each chain.

of simulations, it would be ideal to remove the middle step in future work, i.e. calculate mean temperatures using observed spectra and our own set of simulations. We leave this for future work.

### 5.5.3 Likelihood

The likelihood we use for inference (i.e. within the MCMC method), is a log normal summed over all redshifts, e.g. for the flux power:

$$\log \mathcal{L} = -\frac{1}{2} \sum_{z=2.2}^{z=4.6} \left( \mathbf{P}_F^{\text{diff}} \cdot \mathbf{K}^{-1} \cdot \mathbf{P}_F^{\text{diff}} + \det(\mathbf{K}) \right) \quad (5.5)$$

where  $\mathbf{P}_F^{\text{diff}} = \mathbf{P}_F^{\text{sim}} - \mathbf{P}_F^{\text{obs}}$  is the vector difference between the simulation prediction and the observation. The covariance matrix,  $\mathbf{K}$ , includes emulator errors and uncertainty from the observations. The calculation for the IGM mean temperature likelihoods are similar, though simpler due to being single valued per redshift.

There are two potential emulator errors that can be used in Equation 5.5. The first is the error prediction from the GP, which is the estimated uncertainty on the prediction. The second is the absolute leave-one-out errors from Section 5.3, which are the actual errors in predictions. Figure 5.7 compares these two options, showing the training samples (grey crosses for LF, red circles for HF), GP emulator errors (yellow dots), leave-one-out errors (red dashed) Figure 5.7 also shows flux power only chains run using the GP emulator error (calculated by fixing all but one parameter to their median value, and varying the unfixed parameter) in the likelihood (thick, yellow) and using the leave-one-out errors (averaged over samples) in the likelihood (thin, red).

Both errors are redshift and scale dependent, but the GP error is parameter dependent as well. In Figure 5.7, the GP error (yellow dots) is lowest where the HF samples (red circles) are located, and high at the edges of parameter space. The chain using the GP errors (thick, yellow) prefers the edges for many of the parameters, where the error is largest (except in the case of  $\Omega_M h^2$ , which is partially degenerate with  $h$ ). Because the GP errors

are parameter dependent, they unphysically drive the chain towards the edges, without improving the fit to observations. The GP appears to overestimate the magnitude of the error (see left panel of Figure 5.1) in certain regions, which is not supported by the results of the leave-one-out errors. This overestimation means the GP error is non-negligible as compared to the observational errors for those same regions, possibly due to the high dimensionality of the parameter space. We opt to use the leave-one-out errors in our analysis, to avoid this effect.

For the Lyman- $\alpha$  forest flux power spectrum, a data correction for potential DLA contamination and Si III correlated absorption is applied within the likelihood. The Si III correction follows [162], and adds one additional free parameter ( $f_{\text{SiIII}}$ ) to the inference. The DLA correction is from [233]. We simplify that correction by adding only two additional free parameters to the inference, one parameter covering sub-DLAs and Lyman limit systems (LLS) called  $\alpha_{\text{lls}}$ , and one parameter covering DLAs, called  $\alpha_{\text{dla}}$ . In [233], there are four parameters, with sub-DLAs separate from LLSs, and DLAs divided into two categories. For each of the parameters, a redshift and scale dependent correction is applied, where a positive (negative) value for the parameter implies that our simulation has underestimated (overestimated) the number of absorbers in that category. We set the two DLA parameters equal to each other, as well as the sub-DLA and LLS parameters. This eliminates a strong negative correlation between these two sets of parameters, wherein one parameter can be made large and positive if the other is made large and negative. Neighbouring column density bins have similar shapes to their DLA correction, and the flux power spectrum data set is unable to distinguish between them. A linear combination of the two can be chosen



which is close to zero, causing the degeneracy. Thus, the free parameters used here cover column densities between  $1.6 \times 10^{17} - 10^{20} \text{ cm}^{-2}$  for  $\alpha_{\text{lls}}$ , and  $10^{21} - \infty \text{ cm}^{-2}$  for  $\alpha_{\text{dla}}$ .

These data corrections are applied directly to the emulator predictions for the Lyman- $\alpha$ forest flux power, before calculating the difference between the predictions and observations. See Appendix 5.8 for posteriors for these parameters. There are no data corrections or priors specific to the IGM mean temperature.

For MCMC chains run using both the Lyman- $\alpha$ forest flux power and IGM mean temperature likelihoods, the two likelihoods are calculated separately and added. Because the flux power spectrum has larger magnitude likelihoods (due to the extra dimension of information,  $k$  and additional redshift bins), the mean temperature likelihood is scaled before being added. The scaling is determined from chains run with each likelihood separately, such that the mean likelihood is approximately the same between the two sources (the factor is  $\sim 10$ ).

In Appendix 5.9, we test the likelihood by performing inference with one of the high resolution simulation outputs as data (replacing the observed flux power and mean temperature). The likelihood recovers the correct parameters well, indicating that the likelihood is behaving as expected.

## 5.6 Results

Using the inference scheme outlined in Section 5.5, which makes use of the emulators detailed in Section 5.3 (themselves trained on the simulations from Section 3.3), we can now constrain the cosmological and astrophysical parameters listed in Section 5.4.1. To do

this, we make use of the Cobaya package [229, 228, 237, 238] to run MCMC chains using the observational data discussed in Sections 5.5.1 & 5.5.2, and our simulation based emulator. The MCMC sampler uses the Metropolis method discussed in [237], and uses a Gaussian + exponential proposal distribution that dynamically learns the proposal covariance.

Priors in the form of parameter limits, as well as initial proposal scales for each parameter are passed to Cobaya. The former constrains the region explored by the sampler, while the latter is an initial guess for the width of the proposal distribution for each parameter, which is dynamically updated as the chain runs. Convergence is determined using the Gelman-Rubin statistic,  $R$ , also detailed in [237]. The chains presented here were all run until a convergence of  $R - 1 < 0.01$ , with results plotted for those chains for samples at  $R - 1 < 1$ .

First, we discuss the constraints we obtain on the cosmological parameters: the two primordial power spectrum parameters,  $n_P$  and  $A_p$ ; the Hubble constant  $h$ ; and the matter density  $\Omega_M h^2$ . We then separately discuss the constraints on the astrophysical parameters: the three parameters defining the He II reionization model,  $z_i^{\text{He II}}$ ,  $z_f^{\text{He II}}$ , and  $\alpha_q$ ; and the midpoint of H I reionization,  $z^{\text{H I}}$ . Discussion of the full posteriors, which include cosmology, astrophysics, mean flux, and data correction parameters, is in Appendix 5.8. We defer this to the appendix as the correlations between each of these sets of parameters is weak.

### 5.6.1 Cosmology Results

We begin with the results for the cosmology parameters (but note that all parameters are from the same MCMC analyses): the primordial matter power spectrum amplitude

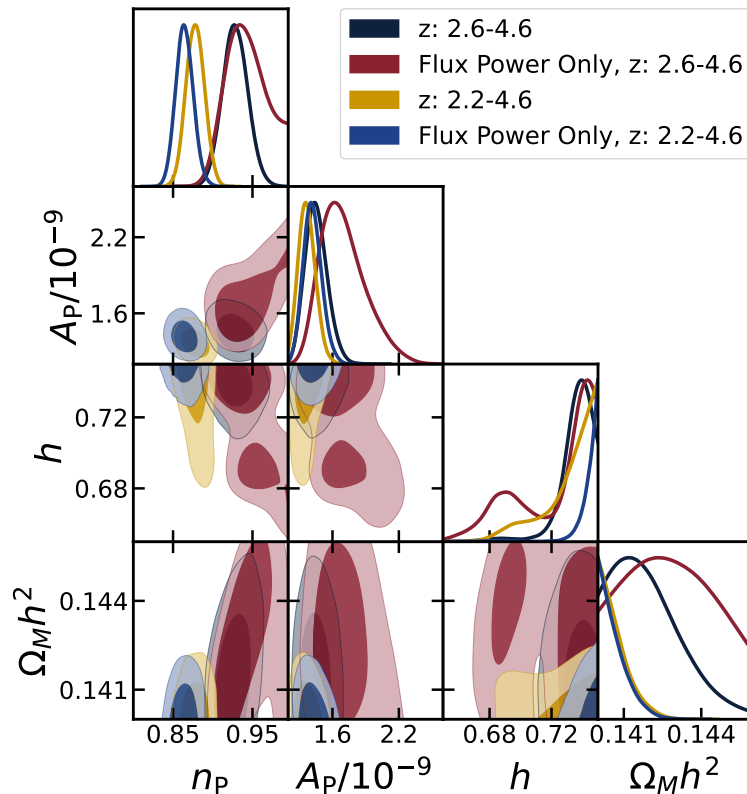


Figure 5.8: Posteriors for the cosmology parameters,  $n_P$ ,  $A_p$ ,  $h$ , and  $\Omega_M h^2$ . Results are from four MCMC chains, all using the multi-fidelity emulator. There are two pairs of chains, one set (black, red) uses a limited redshift range ( $z = 2.6 - 4.6$ ), while the other (yellow, blue) uses the full redshift range ( $z = 2.2 - 4.6$ ). Within each pair there is a chain that only uses the flux power spectrum (red, blue) and a chain that uses both the flux power and mean temperature (black, yellow). Our main result is the chain using the limited redshift range, and both the flux power and mean temperature (black).

$A_p$  and index  $n_P$ , the Hubble constant through  $h$ , and the matter density through  $\Omega_M h^2$ . Figure 5.8 shows the results for these parameters from four different chains. All chains use the multi-fidelity emulator and have no priors. Two pairs of chains are shown with different redshift ranges ( $z = 2.2 - 4.6$ ,  $z = 2.6 - 4.6$ ), each with and without the mean temperature included in the likelihood.

First, the primordial power spectrum index,  $n_P$ . The preferred value for  $n_P$  is consistent between chains using both flux power and mean temperature, and the chains using only the flux power. However, the chains differ significantly when the redshift range is changed. For the chain run using the full redshift range and mean temperature (yellow), the preferred value is  $n_P = 0.879 \pm 0.012$ . This is lower than other results, including from Planck [239], which found a value of  $n_P = 0.965 \pm 0.004$ . The full redshift range result is  $\approx 7\sigma$  different from Planck. When the lowest two redshift bins are removed from the analysis (black),  $n_P = 0.927 \pm 0.017$ . This is  $\approx 2\sigma$  different from Planck. This result is also more consistent with previous constraints using the same Lyman- $\alpha$  forest flux power spectrum observations. In [155] and [182],  $n_P$  was constrained to  $n_P = 0.953 \pm 0.007$  and  $n_P = 0.955 \pm 0.005$ , respectively. Our results are  $5 - 6\sigma$  different when using the full redshift range, and  $\approx 1.5\sigma$  different when using the reduced redshift range.

There are several possibilities for the difference in  $n_P$  between analyses that do and do not include the  $z = 2.2$  and  $z = 2.4$  redshift bins. This difference could point to unknown issues within simulations at these low redshifts, or could be a sign that more work needs to be done to mimic the observational flux power. The observations themselves could be the culprit, for example, misestimation of the continuum at lower redshifts, where the

resolution becomes important, i.e. where the systematics dominate the statistics. There is some indication (Figure 5.10 and Appendix 5.10.3) that the observational errors may be underestimated. If errors in the simulations or observations cannot explain this difference, it may be a sign of running of the spectral index.

The primordial power spectrum amplitude,  $A_p$  is consistent between three of the four chains, with the flux power only, reduced redshift result (red) shifted to higher values.  $A_p$  corresponds to the more traditional parameter,  $A_s$ , via  $A_s = (0.4/2\pi)^{n_P-1} A_p$ , and derived values for  $A_s$  are included in Table 5.3. The amplitude varies between  $A_p = (1.36 - 1.44) \times 10^{-9}$ , with the outlier at  $A_p = 1.62 \times 10^{-9}$ , with corresponding values of  $A_s$  between  $A_s = (1.76 - 1.94) \times 10^{-9}$ . Planck [239] found a value of  $A_s = (2.101_{-0.034}^{+0.031}) \times 10^{-9}$ . Our results are all within  $\approx 2\sigma$  of Planck.

The results for  $h$  are consistent with previous studies using the Lyman- $\alpha$ forest, which found that the Lyman- $\alpha$ forest is only weakly dependent on  $h$  [240, 155], thus we do not constrain  $h$ . For these chains,  $h$  moves towards the high edge of our limits. This is likely a result of the emulator errors being largest at the edges of parameter space, driving the likelihood when there is not a clear preference from the data. Chains were run with priors from both Planck [239] and SHOES [241], but these only constrained  $h$ , and had only marginal effects on other parameters. A more in depth discussion of the effect of using parameter priors on the posterior is in Appendix 5.10.1.

$\Omega_M h^2$  is weakly constrained when the reduced redshift range is used (black, red), and shifts to the lower edge when the full redshift range is used (yellow, blue). For the reduced redshift range, we find  $\Omega_M h^2 = 0.141 \pm 0.001$  for the combined analysis, and

Parameter	$z : 2.6 - 4.6$	$z : 2.2 - 4.6$	$z : 2.6 - 4.6$ (Flux Power Only)
$n_P$	$0.927 \pm 0.017$	$0.879 \pm 0.012$	$0.934 \pm 0.03$
$A_p/10^{-9}$	$1.44 \pm 0.1$	$1.36 \pm 0.08$	$1.62 \pm 0.22$
$h^\dagger$	$0.740 \pm 0.01$	$0.750 \pm 0.03$	$0.740 \pm 0.06$
$\Omega_M h^2$	$0.141 \pm 0.001$	$0.14 \pm 0.001$	$0.142 \pm 0.002$
$z_i^{\text{He II}}$	$4.00 \pm 0.04$	$3.95 \pm 0.05$	$4.10 \pm 0.44$
$z_f^{\text{He II}}$	$2.78 \pm 0.04$	$2.82 \pm 0.03$	$2.6 \pm 0.26$
$\alpha_q$	$1.46 \pm 0.06$	$1.58 \pm 0.06$	$2.50 \pm 0.27$
$z^{\text{H I}}$	$7.87 \pm 0.25$	$6.85 \pm 0.28$	$8.00 \pm 0.39$
$\epsilon_{AGN}$	$0.035 \pm 0.01$	$0.03 \pm 0.009$	$0.039 \pm 0.013$

Table 5.2: Maximum posterior values and  $1\sigma$  uncertainties for the parameters sampled in this analysis.  $\dagger$  We do not constrain  $h$  well in this analysis, so the uncertainty is significantly underestimated.

Parameter	$z : 2.6 - 4.6$	$z : 2.2 - 4.6$	$z : 2.6 - 4.6$ (Flux Power Only)
$A_s/10^{-9}$	$1.76 \pm 0.15$	$1.90 \pm 0.13$	$1.94 \pm 0.31$
$\sigma_8$	$0.825 \pm 0.072$	$0.838 \pm 0.067$	$0.873 \pm 0.159$

Table 5.3: Constraints on derived parameters, with estimates on their  $1\sigma$  uncertainties.

$\Omega_M h^2 = 0.142 \pm 0.002$  when only the flux power is used. The corresponding value from Planck is  $\Omega_M h^2 = 0.1424 \pm 0.001$ , which places both of these results within one sigma of Planck.

The highest likelihood cosmological parameters for each chain are compiled, along with the astrophysical parameters, in Table 5.2. In Table 5.3 derived parameters are listed for  $A_s$  and  $\sigma_8$ .

The values for  $A_s$  have already been discussed. Values for  $\sigma_8$  are calculated using the package Code for Anisotropies in the Microwave Background (CAMB) [242]. Errors are the root squared sum of the fractional errors for the parameters used to calculate  $\sigma_8$ ,

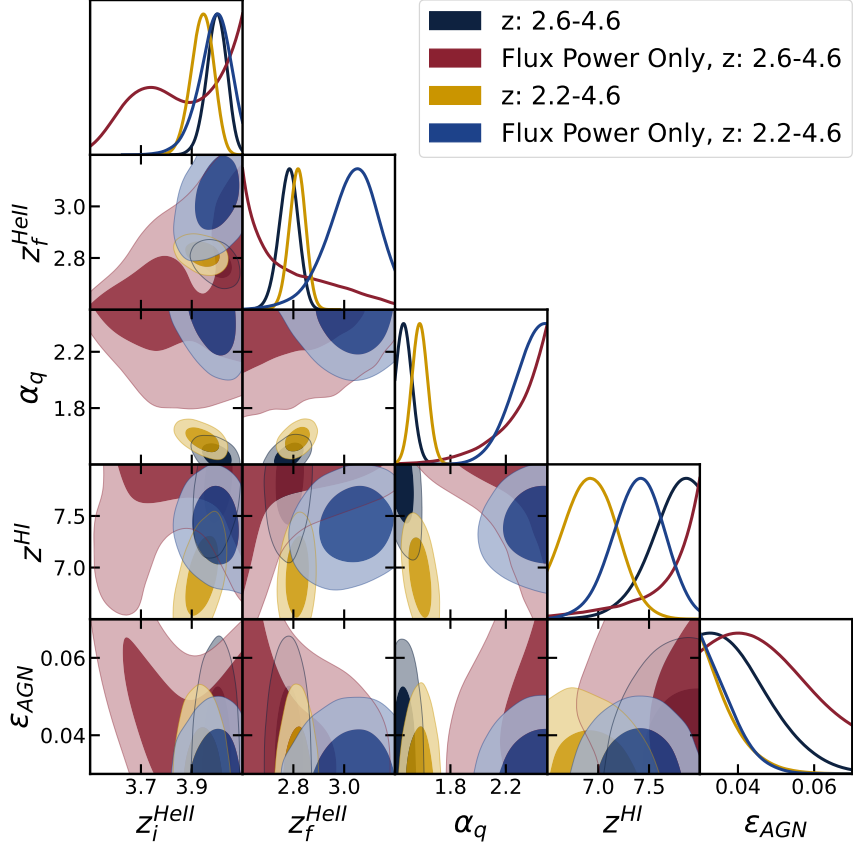


Figure 5.9: Posteriors for the astrophysics parameters,  $z_i^{\text{He II}}$ ,  $z_f^{\text{He II}}$ ,  $\alpha_q$ ,  $z^{\text{H I}}$ , and  $\epsilon_{\text{AGN}}$ . Results are from the same chains described in Section 5.6.1.

times the best value (thus they are likely exaggerated). For the chains run using both data sets,  $\sigma_8 = 0.825 - 0.838$  with uncertainties of  $\approx 0.07$ , well within  $0.5\sigma$  of the Planck result,  $\sigma_8 = 0.811 \pm 0.006$ . When only the flux power is used,  $\sigma_8 = 0.873 \pm 0.159$ , which is within one sigma of Planck, mostly due to the large uncertainty.

## 5.6.2 Astrophysics Results

We now move on to the astrophysical parameters: the starting and ending redshift of He II reionization,  $z_i^{\text{He II}}$  and  $z_f^{\text{He II}}$ , the quasar spectral index,  $\alpha_q$ , the midpoint of H I

reionization,  $z^{\text{H I}}$ , and the black hole feedback factor,  $\epsilon_{AGN}$ . Figure 5.9 shows the results for these parameters from the same set of chains as the cosmology results.

First, the redshifts for He II reionization. The initial redshift is consistent between three of the four chains, with the chain using only the flux power and the reduced redshift range being only mildly constrained. For the three consistent chains, the start of helium reionization is  $z_i^{\text{He II}} = 3.95 - 4.00$ , with one sigma uncertainty of  $\approx 0.05$ . The final redshift is consistent between the chains using both the mean temperature and flux power, but varies significantly when the mean temperature is not included. For the chains using both data sets, the end of helium reionization is  $z_f^{\text{He II}} = 2.78 - 2.82$ , with one sigma uncertainty of  $\approx 0.04$ . Once again, the chain without the mean temperature and with the reduced redshift range is not well constrained. Other than this, the other three chains are consistent with previous studies that have found that He II reionization ends at  $\sim z = 3$  [243, 244, 245].

The flux power only chains struggle to constrain  $\alpha_q$ , with values shifting towards the high edge, signifying that the error dominates for that parameter. Both chains using the mean temperature and flux power prefer lower  $\alpha_q$  than were initially included in the training samples, which led us to add more simulations at low  $\alpha_q$ , as described in Section 5.4.2. Lower  $\alpha_q$  corresponds to more heating from quasars during He II reionization, so this parameter is correlated with both the initial and final redshift of He II reionization. If He II reionization starts earlier or ends later, the IGM requires more heating from quasars to match the observations, while the opposite is true for late starting, or early ending He II reionization. The preferred value we find here is  $\alpha_q = 1.46 - 1.58$ , with one sigma uncertainty of  $\approx 0.06$ .

The midpoint of hydrogen reionization varies the most across this set of chains.



Note, the redshift range explored here ( $z = 2.2 - 4.6$ ) is well after the completion of hydrogen reionization, even in models where it ends late. The bit of information we do have is the effect of hydrogen reionization as the IGM comes back into equilibrium, i.e. the trend in the flux power and temperature at  $z = 4.6$ . This seems to be enough to place reasonably constraints on the midpoint of reionization, however the constraints are not robust to changes in the analysis.

The chains that include the mean temperature, but have different redshift ranges (omitting or including  $z = 2.2$  and  $z = 2.4$ ), sit at opposite ends of the parameter range. The chain using the full redshift range prefers  $z^{\text{H I}} \approx 6.85 \pm 0.28$ . This is consistent with some astrophysical probes, which indicate that the midpoint should be closer to  $z = 7$  [246]. The chain using the reduced redshift range prefers  $z^{\text{H I}} = 7.87 \pm 0.25$ . This is consistent with Planck [239], which found a midpoint of  $z^{\text{H I}} = 7.68 \pm 0.79$ .

These chains make it clear that the inclusion of the mean temperature in our likelihood framework has significantly improved the astrophysics constraints, indicating that the mean temperature of the IGM carries valuable information. The highest likelihood astrophysical parameters for each chain are compiled, along with the cosmological parameters, in Table 5.2. Full posteriors (including correlations between the astrophysical and cosmological parameters) are in Appendix 5.8.

### 5.6.3 Maximum Posterior Predictions

We can now compare the maximum posterior emulator predictions for the Lyman- $\alpha$  forest flux power spectrum and IGM mean temperatures to the observations. We begin

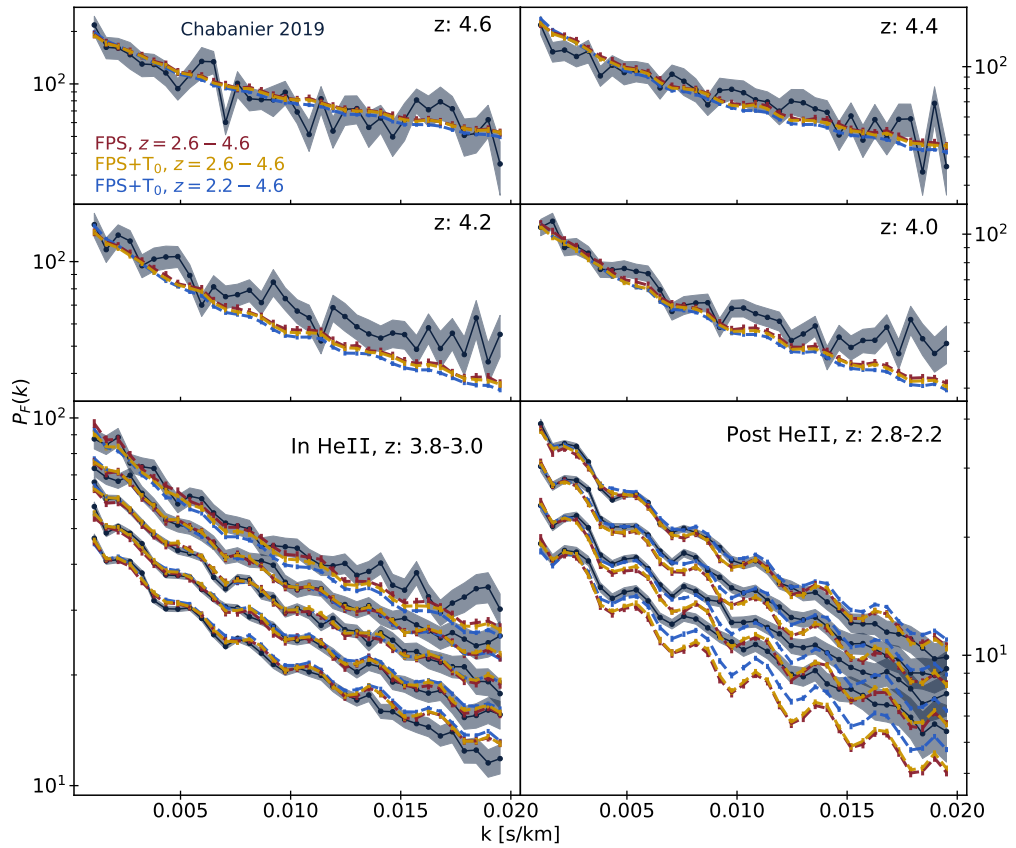


Figure 5.10: Observed Lyman- $\alpha$  forest flux power spectrum [182], from  $z = 4.6$  down to  $z = 2.2$  (black lines and circles, with shading corresponding to one sigma uncertainty). Also shown are three predictions for the Lyman- $\alpha$  forest flux power spectrum from our multi-fidelity emulator corresponding to the maximum posterior input parameters compiled in Table 5.2. The negative of the log-likelihood for these fits is compiled in Table 5.4.

Redshift	4.6	4.4	4.2	4.0	3.8	3.6	3.4	3.2	3.0	2.8	2.6	2.4	2.2
$z : 2.6 - 4.6$	24	19	25	30	16	8	15	17	25	26	25	47	75
FPS Only	24	19	23	28	17	8	17	20	28	27	27	48	76
$z : 2.2 - 4.6$	22	23	30	37	23	9	14	19	34	29	18	19	27

Table 5.4: Negative of the log-likelihood ( $\chi^2$ ) fit values for the flux power spectrum (smaller values correspond with a better fit). The third row is for the best fit parameters from a flux power only chain, using  $z = 2.6 - 4.6$ .

with the flux power.

Figure 5.10 shows the Lyman- $\alpha$  forest flux power spectrum from [182], along with their estimated one sigma uncertainty (black). Also shown in Figure 5.10 are predictions from our multi-fidelity emulator based on the maximum posterior input parameters from MCMC analysis with only the Lyman- $\alpha$  forest flux power emulator (reduced redshift range, red), and MCMC analysis using both the mean temperature and flux power emulators (reduced redshift range, yellow, and full redshift range, blue). The correlation between Lyman- $\alpha$  and Si III absorption can be seen at lower redshifts in the form of regular oscillations in the power spectrum (in Section 5.5.3 we describe the correction we make for Si III).

The fits are generally good, though from  $z = 4.2$  to  $z = 3.8$ , the prediction is consistently lower than the observations. For the chains run without the two lowest redshift bins, the associated fits for those redshifts are poor, as expected. For the chain using the full redshift range, the fit also degrades for  $z = 2.2$  on small scales, indicating that the Si III correction could use improvement (one possibility is a scale dependent correction). The negative of the log-likelihood for each of these fits, divided into redshifts, is compiled in Table 5.4. The best-fit flux power spectrum is not significantly affected by the inclusion of the  $T_0$  data in the likelihood, as is expected if the  $T_0$  data mainly breaks degeneracies.




Figure 5.11: IGM mean temperatures from [173] (black lines and circles, with shading corresponding to one sigma uncertainty). Specifically their temperatures derived from the flux power spectrum calculated using high resolution Lyman- $\alpha$ -forest spectra. Also shown are predictions for the mean temperature from our multi-fidelity emulator corresponding to the the same three maximum posterior input parameters used in Figure 5.10.

We note that even when chains are run for each redshift separately (and predictions made for these independently), the emulator framework is unable to fit all of the redshifts to the same level.

Figure 5.11 shows the IGM mean temperature from [173]. Specifically, the mean temperature derived from their calculated Lyman- $\alpha$ -forest flux power spectrum is shown (black). Also shown in Figure 5.11 are predictions from our multi-fidelity emulator based on the maximum posterior input parameters from the same chains used in Figure 5.10. Due to the larger uncertainty in the observed mean temperature, both predictions are consistent with the observations. In Figure 5.11 we include the predictions for  $z > 3.8$ , outside the

range of the observations, only to highlight the higher redshift trends indicated by the different emulator predictions. Both the flux power and mean temperature emulators affect the highest likelihood mean temperature prediction.

## 5.7 Conclusions

Using a simulation based machine learning emulator for two Lyman- $\alpha$  forest summary statistics, along with their observed counterparts, we performed inference on a set of cosmological and astrophysical parameters. The summary statistics we use are the Lyman- $\alpha$  forest flux power spectrum and the mean temperature of the intergalactic medium. The parameters we constrain describe hydrogen and helium reionization and the primordial power spectrum.

The emulator uses a multi-fidelity method to produce predictions at high resolution from Gaussian processes trained on two suites of hydrodynamical simulations. These training simulations are split into low fidelity (lower resolution) and high fidelity (higher resolution) suites, run with volumes of  $(120h/\text{Mpc})^3$  and gas particle loads of  $1536^3$  and  $3072^3$ , respectively. The emulator allows us to produce cheap,  $\approx 1\%$  accurate predictions for the simulated summary statistics for arbitrary parameter inputs, using the 48 low fidelity and 3 high fidelity training simulations.

The inference scheme includes a likelihood function within an MCMC framework (using the package Cobaya [229, 228]). Our likelihood includes corrections to the Lyman- $\alpha$  forest flux power spectrum from correlated Si III absorption and from the presence of damped Lyman- $\alpha$  systems. A redshift dependent mean flux rescaling is also applied. All

together, there are an additional 5 post processing parameters associated with the Lyman- $\alpha$  forest flux power spectrum, all of which are included in the MCMC sampling.

In summary, the analysis performed here supports:

- helium reionization beginning at  $z_i^{\text{He II}} = 4.00 \pm 0.04$  and running to  $z_i^{\text{He II}} = 2.78 \pm 0.04$ , with a low quasar spectral index, which corresponds to more heating during this process;
- an early midpoint of reionization, at  $z^{\text{H I}} = 7.87 \pm 0.25$ , consistent Planck [239], though the Planck measurement has large uncertainty,  $z_{\text{Planck}}^{\text{H I}} = 7.68 \pm 0.79$ ;
- a value for the primordial power index,  $n_P = 0.927 \pm 0.017$ , within  $2\sigma$  of Planck;
- and a primordial power amplitude  $A_p = (1.44 \times 10^{-9})$ , which along with the value of  $n_P$ , translates to  $A_s = (1.76 \pm 0.13 \times 10^{-9})$ , within  $2\sigma$  of Planck;
- Our analysis does not constrain  $\epsilon_{\text{AGN}}$  (the black hole feedback factor) or  $h$ , and only weakly constrains  $\Omega_M h^2$ .

The above results are from an analysis that used the redshift range  $z = 2.6 - 4.6$ . When including the full redshift range,  $z = 2.2 - 4.6$ , several parameters have significant shifts. The midpoint of hydrogen reionization shifts later,  $z^{\text{H I}} = 6.85 \pm 0.28$ , which is consistent with some measures [246], and, due to the large uncertainty of the Planck measurement, is also consistent with Planck. The value for the primordial power index becomes very low,  $n_P = 0.879 \pm 0.012$ , as much as  $7\sigma$  different from past measurements [239, 155, 182]. The resulting primordial power amplitude is  $A_s = (1.90 \pm 0.13 \times 10^{-9})$ , within  $1\sigma$  of Planck.

The low value for  $n_P$ , when using the full redshift range, could be due to the improved volume and resolution of the simulations used in our analysis. These avoid under-sampling of large scales and smoothing of small scales, which may give a better estimate of the true Lyman- $\alpha$ -forest flux power at lower redshifts. Other potential explanations include errors in the continuum fitting for the observations, or underestimation of the uncertainty in the observed flux power (Appendix 5.10.3). The instrumental resolution for the lowest redshifts means systematics dominate the error budget, which may bias our results. Should none of these potential causes be the true culprit, then the result using the full redshift range may be evidence for running of the primordial power index.

The Lyman- $\alpha$ -forest probes scales ( $\sim 1 - 100 \times$  the cosmological mean density) and redshifts ( $z = 2 - 5$ ) that make it a powerful tool for cosmological and astrophysical enquiry. The abundance of observations of the Lyman- $\alpha$ -forest add to this strength, and future surveys will only make analyses, like the one presented here, more robust. For example, the Dark Energy Spectroscopic Instrument (DESI) [220] is expected to accurately measure the Lyman- $\alpha$ -flux power spectrum at small scales,  $k \approx 0.035 \text{ km}^{-1} \text{ s}$ , and high redshifts,  $z > 4.6$  [185].

In this work we combined the Lyman- $\alpha$ -forest flux power spectrum with the mean temperature of the intergalactic medium. The inclusion of this second stream of information improved many of the constraints on the sampled parameters, especially the parameters related to helium and hydrogen reionization. One of the strengths of our framework is that it can be expanded and improved in several ways: adding other summary statistics (drawn from the training simulations, and from observations), expanding the parameter

space, including more training simulations to improve emulator accuracy, and incorporating future observations.

This may include emulating both the matter power spectrum and flux power spectrum, which would allow a measurement of the bias between these two statistics. Also, small sample size, but high resolution Lyman- $\alpha$ forest spectra could be incorporated to extend the analysis to smaller scales, which the current simulations would support. In a forthcoming work, we use the same framework and constraints presented here, along with the degeneracy between the primordial power amplitude and sum of neutrino masses [247], to constrain the neutrino mass via the Lyman- $\alpha$ forest.

## 5.8 Full Posteriors

Here, we present the full posteriors, including the correlations between the cosmology and astrophysics parameter sets, and the results for the post-processing parameters, which include parameters for data corrections (Si III and DLA), and for mean flux rescaling. The correlations between parameters in each of these sets is weak, as can be seen in Figure 5.12. For the chains run using both the mean temperature and flux power, the main correlations are: a slightly negative correlation between  $n_P$  and  $A_p$ ; a negative correlation between  $z_i^{\text{He II}}$  and both  $\alpha_q$  and  $z_f^{\text{He II}}$ ; a positive correlation between  $z_i^{\text{He II}}$  and  $z^{\text{H I}}$ ; a negative correlation between  $\alpha_q$  and  $z^{\text{H I}}$ ; and a positive correlation between  $\alpha_q$  and  $z_f^{\text{He II}}$ .

The only correlation between astrophysics and cosmology is a positive correlation between  $A_p$  and  $\alpha_q$ . Since the Lyman- $\alpha$ forest flux power spectrum is not strongly affected by  $\alpha_q$ , this correlation is most likely from the IGM mean temperature likelihood. This



may indicate that a larger  $\alpha_q$  (which corresponds to less heating) is appropriate when the primordial power spectrum amplitude is larger, which leads to more structure.

The  $z_i^{\text{He II}}-z_f^{\text{He II}}-\alpha_q$  correlations indicate that the model prefers either: a late start to helium reionization with less heating and an abbreviated reionization duration; or an early start to reionization with more heating, where the helium is mostly all ionized and the IGM temperature curves starts turning over before  $z_f^{\text{He II}}$  (flattens and begins cooling because the model has nothing left to ionize/heat). [173] also found that models for the thermal history support both pictures; a late and rapid reionization, or an extended reionization with the peak temperature occurring earlier.

The correlations between  $z^{\text{H I}}$  and both  $\alpha_q$  and  $z_i^{\text{He II}}$  are natural within the model. An earlier midpoint of reionization means that He II reionization should also start earlier, such that the temperature boost from H I reionization does dissipate too far to match observations. Likewise, a later midpoint allows for a later start to He II reionization. This is also indicated by the correlation with  $\alpha_q$ , which indicates that a later midpoint requires less heating to match observations, thus a larger  $\alpha_q$ .

The correlation between  $n_P$  and  $A_p$  is unsurprising, as they are both inputs to the primordial power spectrum, and to fit some modes, a raise works as well as a tilt.

Results for the data correction and mean flux scaling parameters are shown in Figure 5.13. The flux power only results are included here to highlight the difference made by the inclusion of the mean temperature likelihood. The post-processing parameters are well constrained and mostly consistent across the chains using the mean temperature and flux power, but vary when the redshift range is changed.

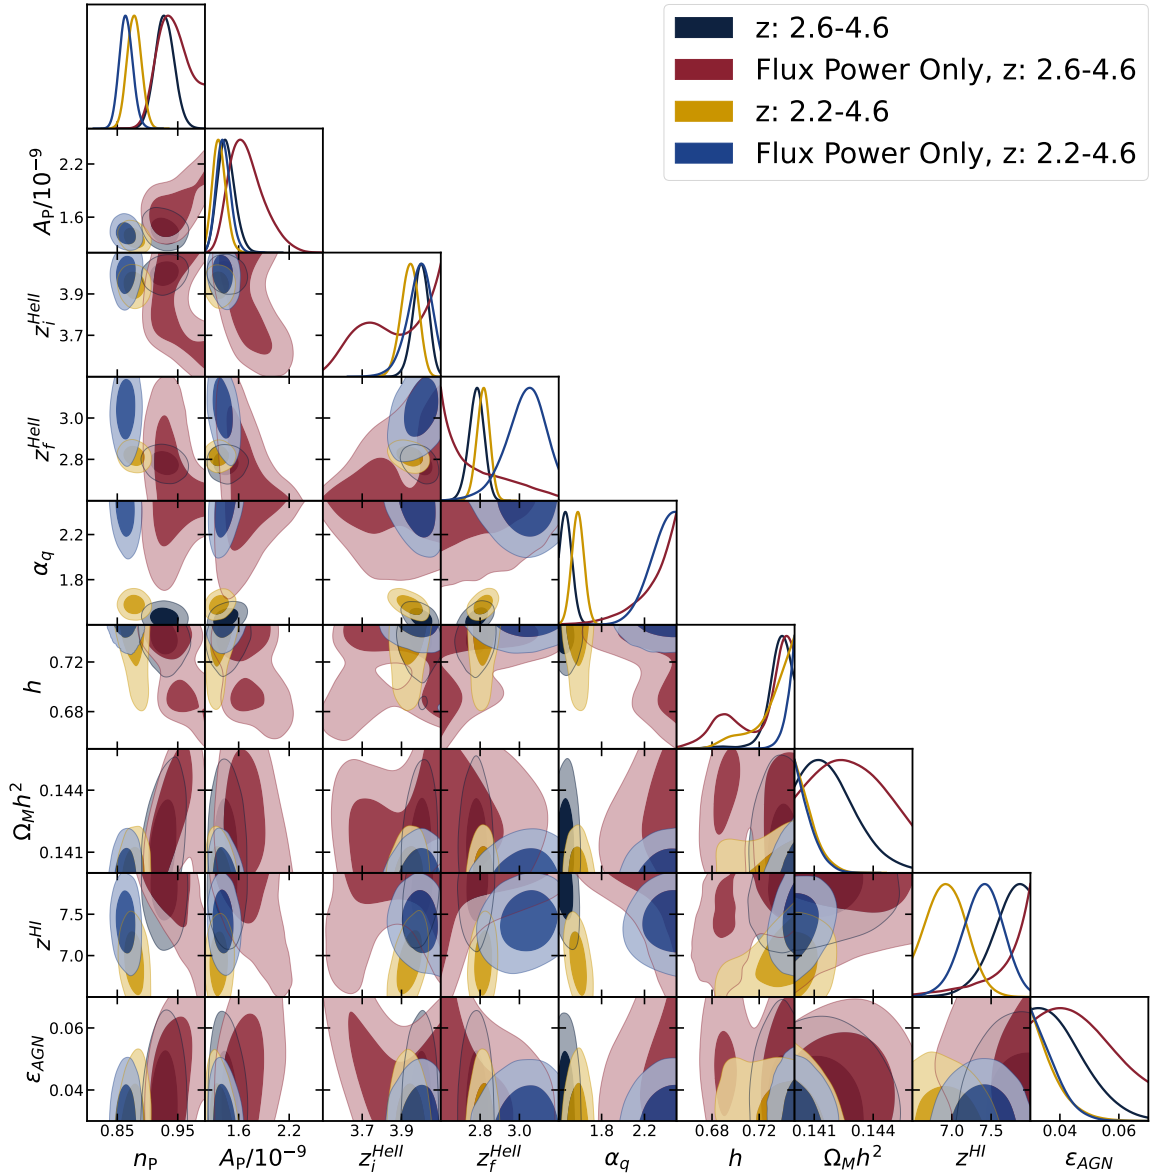


Figure 5.12: Posteriors for the full set of simulation parameters, both cosmological and astrophysical. Shown are the same chains from the figures in Section 4.5, thus the correlations between the cosmology and astrophysics parameters are the only new information here.

Parameter	$z : 2.6 - 4.6$	$z : 2.2 - 4.6$	$z : 2.6 - 4.6$ (Flux Power Only)
$d\tau_0$	$0.0 \pm 0.04$	$-0.28 \pm 0.03$	$-0.01 \pm 0.06$
$\tau_0$	$1.09 \pm 0.02$	$1.17 \pm 0.01$	$1.08 \pm 0.03$
$\alpha_{lls}$	$0.13 \pm 0.03$	$0.02 \pm 0.02$	$0.107 \pm 0.036$
$\alpha_{dla}$	$-0.009 \pm 0.005$	$-0.006 \pm 0.003$	$-0.001 \pm 0.009$
fSiIII	$0.0098 \pm 0.0005$	$0.0087 \pm 0.0004$	$0.0096 \pm 0.0005$

Table 5.5: Maximum posterior values and  $1\sigma$  uncertainties for the post processing parameters sampled in this analysis. The parameter limits for these are:  $d\tau_0 : [-0.4, 0.25]$ ,  $\tau_0 : [0.75, 1.25]$ ,  $\alpha_{lls} : [-1, 1]$ ,  $\alpha_{dla} : [-0.3, 0.3]$ , and fSiIII:  $[-0.03, 0.03]$ .

The Si III correlated absorption parameter,  $fSiIII$ , is well constrained, and small. The effect of this correction can be clearly seen in the oscillations of the flux power spectrum in Figure 5.10. The inclusion of the full redshift range shifts the preferred value to slightly smaller values.

The DLA correction parameters,  $\alpha_{lls}$  and  $\alpha_{dla}$ , are well constrained. There are small differences between the four chains for  $\alpha_{dla}$ , all of which prefer a value near zero. For the DLA correction  $\alpha_{dla}$ , this is not surprising, as DLAs were removed from both the observed and simulated Lyman- $\alpha$ -forest spectra used to construct the flux power spectrum. There are larger differences between the chains for  $\alpha_{lls}$ . When the full redshift range is used,  $\alpha_{lls}$  is consistent with zero, and when the reduced redshift range is used,  $\alpha_{lls}$  is positive. A positive value for  $\alpha_{lls}$  indicates that some subtraction due to LLS and sub-DLAs is required. A positive value therefore makes sense, as we did not attempt to remove LLS and sub-DLAs from the simulated flux power.

The mean flux rescaling parameters,  $\tau_0$  and  $d\tau_0$ , are well constrained and consistent when the redshift range is held fixed. For chains run with the full redshift range, the

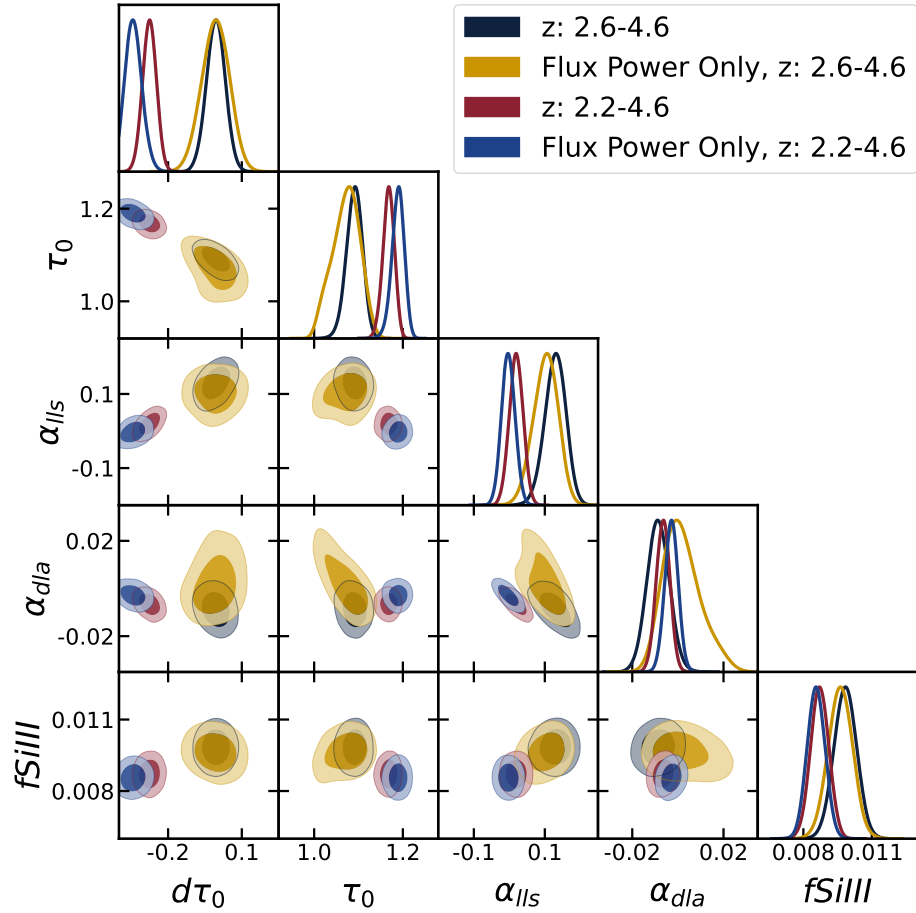


Figure 5.13: Posteriors for the post-processing parameters, which include the mean flux rescaling parameters,  $\tau_0$ ,  $d\tau_0$ , the DLA correction parameters,  $\alpha_{lls}$ ,  $\alpha_{dla}$ , and the Si III correlated absorption correction parameter,  $fSiIII$ . The same chains are shown here as in the main results.

amplitude,  $\tau_0 \approx 1.2$  with an index,  $d\tau_0 \approx -0.3$ . When the reduced redshift range is used, the amplitude,  $\tau_0 \approx 1.1$ , with an index consistent with zero. As a reminder, these parameters modify the mean flux relation from [146], so a value of  $\tau_0 = 1$  and  $d\tau_0 = 0$  would correspond to agreement with that model.

Given the maximum posterior values from the chains using both data sets, the mean flux relation would be  $\tau_{\text{H I}}^{\text{eff}} = 0.0025 \times (1 + z)^{3.65}$ . This is consistent with the model from which we started [146], as well as the mean fluxes from [248].

We tested whether a second mean flux rescaling slope would improve the fit to the observed Lyman- $\alpha$  forest flux power (Figure 5.10), especially at lower redshifts and smaller scales. To do this, we added a second mean flux slope to the MCMC sampled parameters, and assigned each on to a specific redshift range (we tested this using a redshift pivot of  $z = 3$  and  $z = 3.6$ ). Posterior constraints on the other parameters from a chain run using the second mean flux slope were unaffected and the fit was not improved.

## 5.9 Inference Using Simulation Data

In this section we test our inference framework by using high resolution simulation outputs in place of the observational data, to see how well we recover the correct parameters. Figure 5.14 shows the results of this, with the dashed red lines indicating the correct parameters. We successfully recover all of the parameters within one sigma. From this result, the importance of some of the parameters can be seen. For example, the broad posterior of  $\Omega_M h^2$  and  $\epsilon_{\text{AGN}}$  indicate that these parameters do not strongly affect the flux power or mean temperature.




Figure 5.14: Posteriors when using one of the high resolution simulation outputs as data, i.e. replacing observations with a known truth. The true parameter values for the simulation that is used for the flux power and mean temperature are indicated by the red dashed lines.

## 5.10 Alternative Likelihood Results

### 5.10.1 Parameter Priors

In this section we compare posteriors that use priors for black hole feedback,  $\epsilon_{AGN}$ , and the matter density,  $\Omega_M h^2$ , to our main results, which did not include parameter priors. From Planck [239], we have a prior for  $\Omega_M h^2 = 0.1424 \pm 0.001$ . The prior for the black hole feedback factor,  $\epsilon_{AGN}$ , is included to effectively remove this parameter from the inference, as it is not well constrained and does not have a strong effect on the flux power or mean temperature (the prior is  $\epsilon_{AGN} = 0.05 \pm 0.01$ ). All priors are Gaussian and implemented in the likelihood.

The chains are once again divided into redshift ranges, with one pair using  $z = 2.2 - 4.6$  and the other using  $z = 2.6 - 4.6$ . Figure 5.15 shows the results when using priors, as well as our main results without priors, for comparison. When we include the Planck prior on  $\Omega_M h^2$  and the prior on  $\epsilon_{AGN}$ , all the other parameters are changed only marginally. This is encouraging, as our results are robust to the inclusion of priors, and the correlations between most sets of parameters are weak. The only parameter that does change is  $h$ , which for the full redshift range chain begins to develop a second mode in the posterior near  $h = 0.69$ . This is likely due to the degeneracy between  $h$  and  $\Omega_M h^2$ .

### 5.10.2 BOSS DR9 Data

In this section we compare posteriors obtained using a previous observational data set, specifically the flux power spectrum from [112], which is based on BOSS DR9 quasar spectra. Figure 5.16 shows the posteriors for chains run with the newer data set (DR14,

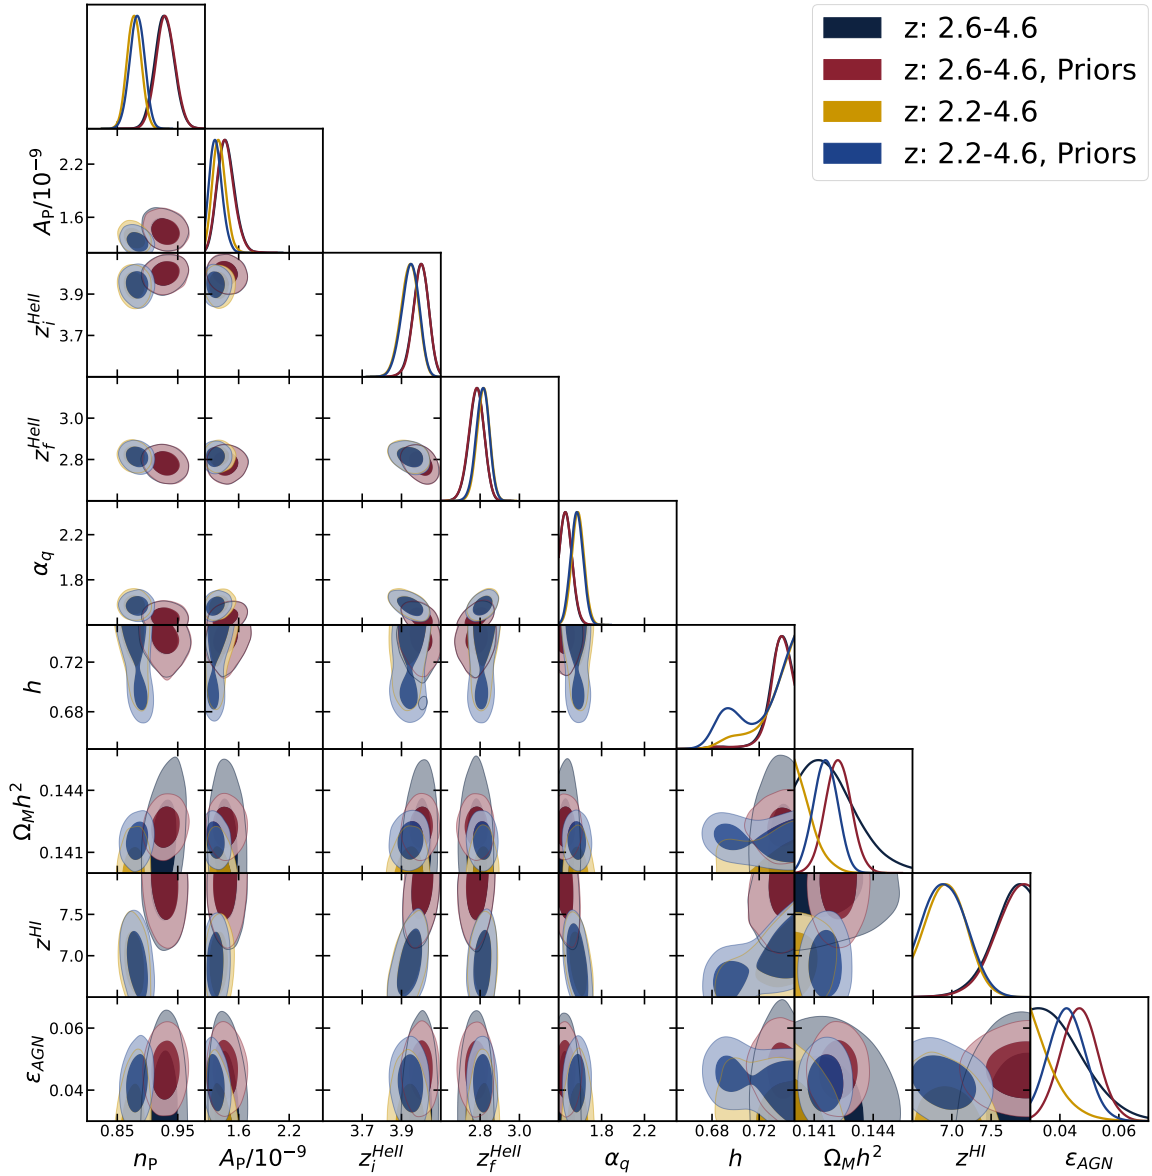


Figure 5.15: Posteriors for chains run using priors on  $\epsilon_{AGN}$  and  $\Omega_M h^2$  (red, blue), as compared with their counterparts run without priors (black, yellow).



black and red), as well as chains run with the previous data set (DR9, yellow and blue). We only show parameters that shifted due to the change in the observational data set.

For both redshift ranges, the primordial power spectrum index,  $n_P$ , shifts to lower values when DR9 is used. Conversely, the amplitude,  $A_p$ , is raised when switching to DR9. Through a degeneracy between the primordial power and mean flux parameters,  $d\tau_0$  shifts in the opposite direction to  $n_P$ , as does  $\tau_0$  with respect to  $A_p$ . The reason for these changes, as well as the changes in the other parameters, likely comes from both the improved uncertainty in the newer data set (due to larger samples and updated processing), and from the inclusion of the  $z = 4.6$  bin, which was not available in the DR9 data set.

The shift from the omission of the lowest redshift bins is mostly consistent between the two observational data sets. This may indicate that the root of the discrepancy between the different redshift range results is either in the instrumentation, or in the simulations.

### 5.10.3 Inflated BOSS Errors

In this section we compare posteriors obtained when the observational uncertainty is increased. Specifically, we increase the uncertainty by a factor of  $\sqrt{2}$  and 2. These results, along with a chain run using the default uncertainty, are shown in Figure 5.17. As the uncertainty is increased from the default (black), to  $\sqrt{2}$  times the default (red), and finally 2 times the default (yellow), the value for  $n_P$  shifts to higher values. Many of the other parameters are only marginally affected (note that we include the most affected parameters in the figure), or simply the width of the posterior increases. The other exception to this is  $d\tau_0$ , which is degenerate with  $n_P$ .




Figure 5.16: Posteriors for a chain run using observations from a previous data release, DR9 (red), compared to our main result, using DR14 (black).

This may indicate that the uncertainties on the observed flux power are underestimated. We ran a similar set of chains wherein we increased the emulator errors (up to 10 times the default), but this only marginally affect the parameters.

#### 5.10.4 Mean Temperature Only Emulator

In this section, we present results from chains run using only the mean temperature likelihood. Shown in Figure 5.18 are four chains, each using one of the observational mean temperatures, which are derived from different Lyman- $\alpha$  forest summary statistics: the flux power spectrum (blue), using the Doppler width distribution (BPDF, yellow), using the curvature statistic (red), and using a wavelet decomposition (black).

Most of the cosmology parameters are unconstrained by the mean temperature. The one exception is a slight preference for high  $A_p$ , though this is most likely due to the emulator error dominating the information. The three He II reionization parameters are well constrained by the mean temperature, as is expected. Because the mean temperature traces the history of He II reionization, the duration and magnitude of heating are naturally constrained. There is very little difference between the different mean temperature observations, with the BPDF derived temperature differing the most, specifically indicating a later start to He II reionization, with less heating.

The midpoint of H I reionization is also somewhat constrained, preferring the lowest values, thus a late midpoint. This may be due to errors, as was the likeliest case with  $n_p$ . However, there is some information in the mean temperature, as the IGM cools from the completion of H I reionization down towards our highest redshift at  $z = 4.6$ . The




Figure 5.17: Posteriors for a chain run using observations with inflated errors: the default errors (black), root two times the default (red), and two times the default (yellow). These use only the flux power, and include the full redshift range,  $z = 2.2 - 4.6$ .

mean temperature also has a strong effect on the midpoint when combined with the flux power, another indication that the mean temperature provides information on the midpoint redshift.

### 5.10.5 Reduced Scale Range

In this section, we present results from chains run using a reduced range of scales. Both chains are run using the multi-fidelity emulator, the full redshift range, with both the mean temperature and flux power likelihoods. One chain is run with the three largest scales (three smallest  $k$ ) omitted. Note that this is roughly the scales where systematics in the continuum dominate over statistical error. Another chain is run leaving out the five smallest scales (five largest  $k$ ). The results of these chains are shown in Figure 5.19 (red and yellow), along with a chain using the full range of scales (black).

Several parameters are relatively stable between these two chains:  $A_p$ , the three He II reionization parameters, and  $\epsilon_{AGN}$ . The value for  $n_P$  is slightly higher when the largest scales are omitted, highlighting the importance of the large scales, and thus the volume used in our simulations. The midpoint of reionization redshift,  $z^{\text{H I}}$ , is not well constrained when the large scales are omitted, likely as a result of the higher value of  $n_P$ . When the large scales are omitted  $\Omega_M h^2$  shifts off of the edge slightly, and this is accompanied by a secondary mode strengthening in  $h$ . This may indicate that the value for  $h$  and  $\Omega_M h^2$  are partially driven by overfitting of small scales, at the expense of large scales. These results indicate that the increased volume of the simulations used in this work have helped the analysis, by better resolving the largest scales.




Figure 5.18: Posteriors for chains run with only the mean temperature likelihood. Shown are chains using each of the four observational measurements of the mean temperature: using the flux power spectrum (blue), using the Doppler width distribution (BPDF, yellow), using the curvature statistic (red), and using a wavelet decomposition (black). The main results of this work use the flux power derived mean temperatures.




Figure 5.19: Posteriors for chains run with reduced ranges for the scales used in the flux power likelihood. Shown are a chain without the three largest scales (red), a chain without the five smallest scales (yellow), and a chain using all scales (black), for comparison.

## Chapter 6

# Conclusions

*“in all of our education, whether it’s in institutions or not, in homes or streets or wherever, whether it’s scholarly or whether it’s experiential, there is a kind of progression. We move from data to information to knowledge to wisdom. . . . And how quickly we can forget that wisdom without any data, is just a hunch.”*

– The Source of Self-Regard, by Toni Morrison

In the work presented in this dissertation, small and large scales have been explored in an attempt to understand both cosmology and astrophysics. A common theme has been the use of simulations, which are especially important for small scales, as these follow non-linear structure formation, which is not well modeled analytically (though semi-analytic models exist). In the final two works presented, machine learning became a key tool, and it will continue to be, not just in cosmology and astrophysics, but in all areas of science.

In Chapter 2, a method for detecting theoretical cosmic strings was presented. While the method required cosmic strings with tensions larger than some current bounds



indicate, there are reasonable models which circumvent those bounds, and the novel detection method would serve as a independent signal. Future work in this area could include developing methods to add cosmic string loops to simulations, to determine how much they affect structure formation, and in particular early massive galaxy formation.

In Chapter 3, an updated method of initializing simulations was tested, specifically the use of distinct transfer functions for baryons and dark matter. The effect of this change on the simulated Lyman- $\alpha$  forest was quantified. Though the changes to the summary statistics measured was modest, they were sufficient to be important in future work, where observations will increase in size and quality.

In Chapter 4, a method for emulating the simulated Lyman- $\alpha$  forest flux power spectrum was presented. This method used a machine learning method called Gaussian processes to construct a multi-fidelity emulator from a suite of training simulations at two resolutions. The multi-fidelity emulator worked well, achieving sub-1% error predictions for the higher resolution output, at approximately one-third the computational cost.

In Chapter 5, cosmological and astrophysical parameter inference was presented, making use of the aforementioned multi-fidelity emulator model. In this case, the training simulations were much larger volume, and higher resolution. This allowed the inference framework to make use of all the available observed scales. Using a new likelihood, which incorporates observations and simulation based emulator predictions for the Lyman- $\alpha$  forest flux power spectrum and mean temperature of the intergalactic medium, constraints are found for several parameters.

There are several areas of possible future work related to machine learning and the

Lyman- $\alpha$  forest. The most immediate is using the multi-fidelity emulator inference framework, along with a known degeneracy between the mass of neutrinos and the primordial power amplitude, to place constraints on the neutrino mass. The current inference framework would also support small sample, higher resolution observations of the Lyman- $\alpha$  forest, which could be combined with the currently used observations to cover a wider range of scales. Another use for the current inference framework is to constrain the bias between matter power and flux power by emulating the matter power spectrum in addition to the flux power spectrum.

There are several potential improvements to the machine learning methods. Development of a model that more economically trains Gaussian processes for each scale in the flux power spectrum would allow even more accurate predictions. Current and future work will expand the multi-fidelity model to more than two fidelities, e.g. improving volume in addition to resolution.

An interesting avenue of current and future work is the use of new summary statistics. One such statistic is the scattering transform, which may be a more information rich summary of the Lyman- $\alpha$  forest than current commonly used statistics.

I end with a quote, which summarizes many academic journeys, including my own:

*“We have to continually be jumping off cliffs and developing our wings on the way down.”*

– If This Isn’t Nice What Is?, by Kurt Vonnegut

# Bibliography

- [1] James E. Gunn and Bruce A. Peterson. On the Density of Neutral Hydrogen in Intergalactic Space. *ApJ*, 142:1633–1636, November 1965.
- [2] Piero Madau, Francesco Haardt, and Martin J. Rees. Radiative Transfer in a Clumpy Universe. III. The Nature of Cosmological Ionizing Sources. *ApJ*, 514(2):648–659, Apr 1999.
- [3] Matthew McQuinn, Adam Lidz, Matias Zaldarriaga, Lars Hernquist, Philip F. Hopkins, Suvendra Dutta, and Claude-André Faucher-Giguère. He II Reionization and its Effect on the Intergalactic Medium. *ApJ*, 694(2):842–866, April 2009.
- [4] Francesco Haardt and Piero Madau. Radiative Transfer in a Clumpy Universe. IV. New Synthesis Models of the Cosmic UV/X-Ray Background. *ApJ*, 746(2):125, Feb 2012.
- [5] R. J. Bouwens, R. Smit, I. Labbé, M. Franx, J. Caruana, P. Oesch, M. Stefanon, and N. Rasappu. The Lyman-Continuum Photon Production Efficiency  $\xi_{ion}$  of  $z \sim 4 - 5$  Galaxies from IRAC-based  $H\alpha$  Measurements: Implications for the Escape Fraction and Cosmic Reionization. *ApJ*, 831(2):176, November 2016.
- [6] Planck Collaboration, N. Aghanim, Y. Akrami, F. Arroja, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak, R. Battye, K. Benabed, J. P. Bernard, M. Bersanelli, P. Bielewicz, J. J. Bock, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J. F. Cardoso, J. Carron, B. Casaponsa, A. Challinor, H. C. Chiang, L. P. L. Colombo, C. Combet, D. Contreras, B. P. Crill, F. Cuttaia, P. de Bernardis, G. de Zotti, J. Delabrouille, J. M. Delouis, F. X. Désert, E. Di Valentino, C. Dickinson, J. M. Diego, S. Donzelli, O. Doré, M. Douspis, A. Ducout, X. Dupac, G. Efstathiou, F. Elsner, T. A. Enßlin, H. K. Eriksen, E. Falgarone, Y. Fantaye, J. Fergusson, R. Fernandez-Cobos, F. Finelli, F. Forastieri, M. Frailis, E. Franceschi, A. Frolov, S. Galeotta, S. Galli, K. Ganga, R. T. Génova-Santos, M. Gerbino, T. Ghosh, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gruppuso, J. E. Gudmundsson, J. Hamann, W. Handley, F. K. Hansen, G. Helou, D. Herranz, S. R. Hildebrandt, E. Hivon, Z. Huang, A. H. Jaffe, W. C. Jones, A. Karacki,

- E. Keihänen, R. Keskitalo, K. Kiiveri, J. Kim, T. S. Kisner, L. Knox, N. Krachmalnicoff, M. Kunz, H. Kurki-Suonio, G. Lagache, J. M. Lamarre, M. Langer, A. Lasenby, M. Lattanzi, C. R. Lawrence, M. Le Jeune, J. P. Leahy, J. Lesgourgues, F. Levrier, A. Lewis, M. Liguori, P. B. Lilje, M. Lilley, V. Lindholm, M. López-Cañiego, P. M. Lubin, Y. Z. Ma, J. F. Macías-Pérez, G. Maggio, D. Maino, N. Mandolesi, A. Mangilli, A. Marcos-Caballero, M. Maris, P. G. Martin, M. Martinelli, E. Martínez-González, S. Matarrese, N. Mauri, J. D. McEwen, P. D. Meerburg, P. R. Meinhold, A. Melchiorri, A. Mennella, M. Migliaccio, M. Millea, S. Mitra, M. A. Miville-Deschênes, D. Molinari, A. Moneti, L. Montier, G. Morgante, A. Moss, S. Mottet, M. Münchmeyer, P. Natoli, H. U. Nørgaard-Nielsen, C. A. Oxborrow, L. Pagano, D. Paoletti, B. Partridge, G. Patanchon, T. J. Pearson, M. Peel, H. V. Peiris, F. Perrotta, V. Pettorino, F. Piacentini, L. Polastri, G. Polenta, J. L. Puget, J. P. Rachen, M. Reinecke, M. Remazeilles, C. Renault, A. Renzi, G. Rocha, C. Rosset, G. Roudier, J. A. Rubiño-Martín, B. Ruiz-Granados, L. Salvati, M. Sandri, M. Savelainen, D. Scott, E. P. S. Shellard, M. Shiraishi, C. Sirignano, G. Sirri, L. D. Spencer, R. Sunyaev, A. S. Suur-Uski, J. A. Tauber, D. Tavagnacco, M. Tenti, L. Terenzi, L. Toffolatti, M. Tomasi, T. Trombetti, J. Valiviita, B. Van Tent, L. Vibert, P. Vielva, F. Villa, N. Vittorio, B. D. Wandelt, I. K. Wehus, M. White, S. D. M. White, A. Zacchei, and A. Zonca. Planck 2018 results. I. Overview and the cosmological legacy of Planck. *A&A*, 641:A1, September 2020.
- [7] Matthew McQuinn. The Evolution of the Intergalactic Medium. *ARA&A*, 54:313–362, September 2016.
- [8] Lars Hernquist, Neal Katz, David H. Weinberg, and Jordi Miralda-Escudé. The Lyman-Alpha Forest in the Cold Dark Matter Model. *ApJL*, 457:L51, February 1996.
- [9] M. A. Fernandez, Simeon Bird, and Yanou Cui. Cosmic filaments from cosmic strings. *Phys. Rev. D*, 102(4):043509, August 2020.
- [10] M. A. Fernandez, Simeon Bird, and Phoebe Upton Sanderbeck. Effect of separate initial conditions on the lyman- $\alpha$  forest in simulations. *MNRAS*, 503(2):1668–1679, May 2021.
- [11] M. A. Fernandez, Ming-Feng Ho, and Simeon Bird. A multifidelity emulator for the Lyman- $\alpha$  forest flux power spectrum. *MNRAS*, 517(3):3200–3211, December 2022.
- [12] Edmund J. Copeland, Robert C. Myers, and Joseph Polchinski. Cosmic F and D strings. *JHEP*, 06:013, 2004.
- [13] Gia Dvali and Alexander Vilenkin. Formation and evolution of cosmic D strings. *JCAP*, 0403:010, 2004.
- [14] Joseph Polchinski. Introduction to cosmic F- and D-strings. In *String theory: From gauge interactions to cosmology. Proceedings, NATO Advanced Study Institute, Cargèse, France, June 7-19, 2004*, pages 229–253, 2004.

- [15] Mark G. Jackson, Nicholas T. Jones, and Joseph Polchinski. Collisions of cosmic F and D-strings. *JHEP*, 10:013, 2005.
- [16] S. H. Henry Tye, Ira Wasserman, and Mark Wyman. Scaling of multi-tension cosmic superstring networks. *Phys. Rev.*, D71:103508, 2005. [Erratum: *Phys. Rev.*D71,129906(2005)].
- [17] Holger Bech Nielsen and P. Olesen. Vortex Line Models for Dual Strings. *Nucl. Phys.*, B61:45–61, 1973. [,302(1973)].
- [18] T. W. B. Kibble. Topology of Cosmic Domains and Strings. *J. Phys.*, A9:1387–1398, 1976.
- [19] A. Vilenkin. Gravitational radiation from cosmic strings. *Phys. Lett.*, 107B:47–50, 1981.
- [20] Tanmay Vachaspati and Alexander Vilenkin. Gravitational Radiation from Cosmic Strings. *Phys. Rev.*, D31:3052, 1985.
- [21] Neil Turok. Grand Unified Strings and Galaxy Formation. *Nucl. Phys.*, B242:520–541, 1984.
- [22] C. J. Burden. Gravitational Radiation From a Particular Class of Cosmic Strings. *Phys. Lett.*, 164B:277–281, 1985.
- [23] Ken D. Olum and J. J. Blanco-Pillado. Radiation from cosmic string standing waves. *Phys. Rev. Lett.*, 84:4288–4291, 2000.
- [24] J. N. Moore, E. P. S. Shellard, and C. J. A. P. Martins. On the evolution of Abelian-Higgs string networks. *Phys. Rev.*, D65:023503, 2002.
- [25] Daiju Matsunami, Levon Pogosian, Ayush Saurabh, and Tanmay Vachaspati. Decay of Cosmic String Loops Due to Particle Radiation. *Phys. Rev. Lett.*, 122(20):201301, 2019.
- [26] Graham Vincent, Nuno D. Antunes, and Mark Hindmarsh. Numerical simulations of string networks in the Abelian Higgs model. *Phys. Rev. Lett.*, 80:2277–2280, 1998.
- [27] Neil Bevis, Mark Hindmarsh, Martin Kunz, and Jon Urrestilla. CMB power spectrum contribution from cosmic strings using field-evolution simulations of the Abelian Higgs model. *Phys. Rev.*, D75:065015, 2007.
- [28] Yanou Cui and David E. Morrissey. Non-Thermal Dark Matter from Cosmic Strings. *Phys. Rev.*, D79:083532, 2009.
- [29] Daniel G. Figueroa, Mark Hindmarsh, and Jon Urrestilla. Exact Scale-Invariant Background of Gravitational Waves from Cosmic Defects. *Phys. Rev. Lett.*, 110(10):101302, 2013.

- [30] T. W. B. Kibble and Neil Turok. Self-intersection of cosmic strings. *Physics Letters B*, 116(2-3):141–143, 10 1982.
- [31] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, C. Affeldt, M. Afrough, B. Agarwal, M. Agathos, K. Agatsuma, N. Aggarwal, O. D. Aguiar, L. Aiello, A. Ain, P. Ajith, B. Allen, G. Allen, A. Allocca, P. A. Altin, A. Amato, A. Ananyeva, S. B. Anderson, W. G. Anderson, S. Antier, S. Appert, K. Arai, M. C. Araya, J. S. Areeda, N. Arnaud, K. G. Arun, S. Ascenzi, G. Ashton, M. Ast, S. M. Aston, P. Astone, P. Aufmuth, C. Aulbert, K. AultONeal, A. Avila-Alvarez, S. Babak, P. Bacon, M. K. M. Bader, S. Bae, P. T. Baker, F. Baldaccini, G. Ballardín, S. W. Ballmer, S. Banagiri, J. C. Barayoga, S. E. Barclay, B. C. Barish, D. Barker, F. Barone, B. Barr, L. Barsotti, M. Barsuglia, D. Barta, J. Bartlett, I. Bartos, R. Bassiri, A. Basti, J. C. Batch, C. Baune, M. Bawaj, M. Bazzan, B. Bécsy, C. Beer, M. Bejger, I. Belahcene, A. S. Bell, B. K. Berger, G. Bergmann, C. P. L. Berry, D. Bersanetti, A. Bertolini, J. Betzwieser, S. Bhagwat, R. Bhandare, I. A. Bilenko, G. Billingsley, C. R. Billman, J. Birch, R. Birney, O. Birnholtz, S. Biscans, A. Bisht, M. Bitossi, C. Biwer, M. A. Bizouard, J. K. Blackburn, J. Blackman, C. D. Blair, D. G. Blair, R. M. Blair, S. Bloemen, O. Bock, N. Bode, M. Boer, G. Bogaert, A. Bohe, 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, J. E. Broida, A. F. Brooks, D. A. Brown, D. D. Brown, N. M. Brown, S. Brunett, 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. A. Callister, E. Calloni, J. B. Camp, M. Canepa, P. Canizares, K. C. Cannon, H. Cao, J. Cao, C. D. Capano, E. Capocasa, F. Carbognani, S. Caride, M. F. Carney, J. Casanueva Diaz, C. Casentini, S. Caudill, M. Cavaglia, F. Cavalier, R. Cavalieri, G. Cella, C. B. Cepeda, L. Cerboni Baiardi, G. Cerretani, E. Cesarini, S. J. Chamberlin, M. Chan, S. Chao, P. Charlton, E. Chassande-Mottin, D. Chatterjee, B. D. Cheeseboro, H. Y. Chen, Y. Chen, H. P. Cheng, A. Chincarini, A. Chiummo, T. Chmiel, H. S. Cho, M. Cho, J. H. Chow, N. Christensen, Q. Chu, A. J. K. Chua, S. Chua, A. K. W. Chung, S. Chung, G. Ciani, R. Ciolfi, C. E. Cirelli, A. Cirone, F. Clara, J. A. Clark, F. Cleva, C. Cocchieri, E. Coccia, P. F. Cohadon, A. Colla, C. G. Collette, L. R. Cominsky, M. Constanancio, L. Conti, S. J. Cooper, P. Corban, T. R. Corbitt, K. R. Corley, N. Cornish, A. Corsi, S. Cortese, C. A. Costa, M. W. Coughlin, S. B. Coughlin, J. P. Coulon, S. T. Countryman, P. Couvares, P. B. Covas, E. E. Cowan, D. M. Coward, M. J. Cowart, D. C. Coyne, R. Coyne, J. D. E. Creighton, T. D. Creighton, J. Cripe, S. G. Crowder, T. J. Cullen, A. Cumming, L. Cunningham, E. Cuoco, T. Dal Canton, S. L. Danilishin, S. D’Antonio, K. Danzmann, A. Dasgupta, C. F. Da Silva Costa, V. Dattilo, I. Dave, M. Davier, D. Davis, E. J. Daw, B. Day, S. De, D. DeBra, J. Degallaix, M. De Laurentis, S. Deléglise, W. Del Pozzo, T. Denker, T. Dent, V. Dergachev, R. De Rosa, R. T. DeRosa, R. DeSalvo, J. Devenson, R. C. Devine, S. Dhurandhar, M. C. Díaz, L. Di Fiore, M. Di Giovanni, T. Di Girolamo, A. Di Lieto, S. Di Pace, I. Di Palma, F. Di Renzo, Z. Doctor, V. Dolique, F. Donovan, K. L. Dooley,

S. Doravari, I. Dorrington, R. Douglas, M. Dovale Álvarez, T. P. Downes, M. Drago, R. W. P. Drever, J. C. Driggers, Z. Du, M. Ducrot, J. Duncan, S. E. Dwyer, T. B. Edo, M. C. Edwards, A. Effler, H. B. Eggenstein, P. Ehrens, J. Eichholz, S. S. Eikenberry, R. A. Eisenstein, R. C. Essick, Z. B. Etienne, T. Etzel, M. Evans, T. M. Evans, M. Factourovich, V. Fafone, H. Fair, S. Fairhurst, X. Fan, S. Farinon, B. Farr, W. M. Farr, E. J. Fauchon-Jones, M. Favata, M. Fays, H. Fehrmann, J. Feicht, M. M. Fejer, A. Fernandez-Galiana, I. Ferrante, E. C. Ferreira, F. Ferrini, F. Fidecaro, I. Fiori, D. Fiorucci, R. P. Fisher, M. Fitz-Axen, R. Flaminio, M. Fletcher, H. Fong, P. W. F. Forsyth, S. S. Forsyth, J. D. Fournier, S. Frasca, F. Frasconi, Z. Frei, A. Freise, R. Frey, V. Frey, E. M. Fries, P. Fritschel, V. V. Frolov, P. Fulda, M. Fyffe, H. Gabbard, M. Gabel, B. U. Gadre, S. M. Gaebel, J. R. Gair, L. Gammaitoni, M. R. Ganija, S. G. Gaonkar, F. Garufi, S. Gaudio, G. Gaur, V. Gayathri, N. Gehrels, G. Gemme, E. Genin, A. Gennai, D. George, J. George, L. Gergely, V. Germain, S. Ghonge, Abhirup Ghosh, Archisman Ghosh, S. Ghosh, J. A. Giaime, K. D. Giardino, A. Giazotto, K. Gill, L. Glover, E. Goetz, R. Goetz, S. Gomes, G. González, J. M. Gonzalez Castro, A. Gopakumar, M. L. Gorodetsky, S. E. Gossan, M. Gosselin, R. Gouaty, A. Grado, C. Graef, M. Granata, A. Grant, S. Gras, C. Gray, G. Greco, A. C. Green, P. Groot, H. Grote, S. Grunewald, P. Gruning, G. M. Guidi, X. Guo, A. Gupta, M. K. Gupta, K. E. Gushwa, E. K. Gustafson, R. Gustafson, B. R. Hall, E. D. Hall, G. Hammond, M. Haney, M. M. Hanke, J. Hanks, C. Hanna, M. D. Hannam, O. A. Hannuksela, J. Hanson, T. Hardwick, J. Harms, G. M. Harry, I. W. Harry, M. J. Hart, C. J. Haster, K. Haughian, J. Healy, A. Heidmann, M. C. Heintze, H. Heitmann, P. Hello, G. Hemming, M. Hendry, I. S. Heng, J. Hennig, J. Henry, A. W. Heptonstall, M. Heurs, S. Hild, D. Hoak, D. Hofman, K. Holt, D. E. Holz, P. Hopkins, C. Horst, J. Hough, E. A. Houston, E. J. Howell, Y. M. Hu, E. A. Huerta, D. Huet, B. Hughey, S. Husa, S. H. Huttner, T. Huynh-Dinh, N. Indik, D. R. Ingram, R. Inta, G. Intini, H. N. Isa, J. M. Isac, M. Isi, B. R. Iyer, K. Izumi, T. Jacqmin, K. Jani, P. Jaranowski, S. Jawahar, F. Jiménez-Forteza, W. W. Johnson, D. I. Jones, R. Jones, R. J. G. Jonker, L. Ju, J. Junker, C. V. Kalaghatgi, V. Kalogera, S. Kandhasamy, G. Kang, J. B. Kanner, S. Karki, K. S. Karvinen, M. Kasprzack, M. Katolik, E. Katsavounidis, W. Katzman, S. Kaufer, K. Kawabe, F. Kéfélian, D. Keitel, A. J. Kembball, R. Kennedy, C. Kent, J. S. Key, F. Y. Khalili, I. Khan, S. Khan, Z. Khan, E. A. Khazanov, N. Kijbunchoo, Chunglee Kim, J. C. Kim, W. Kim, W. S. Kim, Y. M. Kim, S. J. Kimbrell, E. J. King, P. J. King, R. Kirchhoff, J. S. Kissel, L. Kleybolte, S. Klimenko, P. Koch, S. M. Koehlenbeck, S. Koley, V. Kondrashov, A. Kontos, M. Korobko, W. Z. Korth, I. Kowalska, D. B. Kozak, C. Krämer, V. Kringel, B. Krishnan, A. Królak, G. Kuehn, P. Kumar, R. Kumar, S. Kumar, L. Kuo, A. Kutynia, S. Kwang, B. D. Lackey, K. H. Lai, M. Landry, R. N. Lang, J. Lange, B. Lantz, R. K. Lanza, A. Lartaux-Vollard, P. D. Lasky, M. Laxen, A. Lazzarini, C. Lazzaro, P. Leaci, S. Leavey, C. H. Lee, H. K. Lee, H. M. Lee, H. W. Lee, K. Lee, J. Lehmann, A. Lenon, M. Leonardi, N. Leroy, N. Letendre, Y. Levin, T. G. F. Li, A. Libson, T. B. Littenberg, J. Liu, R. K. L. Lo, N. A. Lockerbie, L. T. London, J. E. Lord, M. Lorenzini, V. Lorette, M. Lormand, G. Losurdo, J. D. Lough, C. O. Lousto, G. Lovelace, H. Lück, D. Lumaca, A. P. Lundgren, R. Lynch, Y. Ma, S. Macfoy,

B. Machenschalk, M. MacInnis, D. M. Macleod, I. Magaña Hernandez, F. Magaña-Sand oval, L. Magaña Zertuche, R. M. Magee, E. Majorana, I. Maksimovic, N. Man, V. Mandic, V. Mangano, G. L. Mansell, M. Manske, M. Mantovani, F. Marchesoni, F. Marion, S. Márka, Z. Márka, C. Markakis, A. S. Markosyan, E. Maros, F. Martelli, L. Martellini, I. W. Martin, D. V. Martynov, K. Mason, A. Masserot, T. J. Massinger, M. Masso-Reid, S. Mastrogiovanni, A. Matas, F. Matichard, L. Matone, N. Mavalvala, N. Mazumder, R. McCarthy, D. E. McClelland, S. McCormick, L. McCuller, S. C. McGuire, G. McIntyre, J. McIver, D. J. McManus, T. McRae, S. T. McWilliams, D. Meacher, G. D. Meadors, J. Meidam, E. Mejuto-Villa, A. Melatos, G. Mendell, R. A. Mercer, E. L. Merilh, M. Merzougui, S. Meshkov, C. Messenger, C. Messick, R. Metzдорff, P. M. Meyers, F. Mezzani, H. Miao, C. Michel, H. Middleton, E. E. Mikhailov, L. Milano, A. L. Miller, A. Miller, B. B. Miller, J. Miller, M. Millhouse, O. Minazzoli, Y. Minenkov, J. Ming, C. Mishra, S. Mitra, V. P. Mitrofanov, G. Mitsu-makher, 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, B. Mours, C. M. Mow-Lowry, G. Mueller, A. W. Muir, Arunava Mukherjee, D. Mukherjee, S. Mukherjee, N. Mukund, A. Mullavey, J. Munch, E. A. M. Muniz, P. G. Murray, K. Napier, I. Nardecchia, L. Naticchioni, R. K. Nayak, G. Nelemans, T. J. N. Nelson, M. Neri, M. Nery, A. Neunzert, J. M. Newport, G. Newton, K. K. Y. Ng, T. T. Nguyen, D. Nichols, A. B. Nielsen, S. Nissanke, A. Nitz, A. Noack, F. Nocera, D. Nolting, M. E. N. Normandin, L. K. Nuttall, J. Oberling, E. Ochsner, E. Oelker, G. H. Ogin, J. J. Oh, S. H. Oh, F. Ohme, M. Oliver, P. Oppermann, Richard J. Oram, B. O'Reilly, R. Ormiston, L. F. Ortega, R. O'Shaughnessy, D. J. Ottaway, H. Overmier, B. J. Owen, A. E. Pace, J. Page, M. A. Page, A. Pai, S. A. Pai, J. R. Palamos, O. Palashov, C. Palomba, A. Pal-Singh, H. Pan, B. Pang, P. T. H. Pang, 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, B. L. Pearlstone, M. Pedraza, R. Pedurand, L. Pekowsky, A. Pele, S. Penn, C. J. Perez, A. Perreca, L. M. Perri, H. P. Pfeiffer, M. Phelps, O. J. Piccinni, M. Pichot, F. Piergiovanni, V. Pierro, G. Pillant, L. Pinard, I. M. Pinto, M. Pitkin, R. Poggiani, P. Popolizio, E. K. Porter, A. Post, J. Powell, J. Prasad, J. W. W. Pratt, V. Predoi, T. Prestegard, M. Prijatelj, M. Principe, S. Privitera, R. Prix, G. A. Prodi, L. G. Prokhorov, O. Puncken, M. Punturo, P. Puppo, M. Pürerer, H. Qi, J. Qin, S. Qiu, V. Quetschke, E. A. Quintero, R. Quitzow-James, F. J. Raab, D. S. Rabeling, H. Radkins, P. Raffai, S. Raja, C. Rajan, M. Rakhmanov, K. E. Ramirez, P. Rapagnani, V. Raymond, M. Razzano, J. Read, T. Regimbau, L. Rei, S. Reid, D. H. Reitze, H. Rew, S. D. Reyes, F. Ricci, P. M. Ricker, S. Rieger, K. Riles, M. Rizzo, N. A. Robertson, R. Robie, F. Robinet, A. Rocchi, L. Rolland, J. G. Rollins, V. J. Roma, J. D. Romano, R. Romano, C. L. Romel, J. H. Romie, D. Rosińska, M. P. Ross, S. Rowan, A. Rüdiger, P. Ruggi, K. Ryan, S. Sachdev, T. Sadecki, L. Sadeghian, M. Sakellariadou, L. Salconi, M. Saleem, F. Salemi, A. Samajdar, L. Sammut, L. M. Sampson, E. J. Sanchez, V. Sandberg, B. Sandeen, J. R. Sanders, B. Sassolas, P. R. Saulson, O. Sauter, R. L. Savage, A. Sawadsky, P. Schale, J. Scheuer, E. Schmidt, J. Schmidt, P. Schmidt, R. Schnabel, R. M. S. Schofield, A. Schönbeck, E. Schreiber, D. Schuette, B. W.



Schulte, B. F. Schutz, S. G. Schwalbe, J. Scott, S. M. Scott, E. Seidel, D. Sellers, A. S. Sengupta, D. Sentenac, V. Sequino, A. Sergeev, D. A. Shaddock, T. J. Shaffer, A. A. Shah, M. S. Shahriar, L. Shao, B. Shapiro, P. Shawhan, A. Sheperd, D. H. Shoemaker, D. M. Shoemaker, K. Siellez, X. Siemens, M. Sieniawska, D. Sigg, A. D. Silva, A. Singer, L. P. Singer, A. Singh, R. Singh, A. Singhal, A. M. Sintes, B. J. J. Slagmolen, B. Smith, J. R. Smith, R. J. E. Smith, E. J. Son, J. A. Sonnenberg, B. Sorazu, F. Sorrentino, T. Souradeep, A. P. Spencer, A. K. Srivastava, A. Staley, D. A. Steer, M. Steinke, J. Steinlechner, S. Steinlechner, D. Steinmeyer, B. C. Stephens, R. Stone, K. A. Strain, G. Stratta, S. E. Strigin, R. Sturani, A. L. Stuver, T. Z. Summerscales, L. Sun, S. Sunil, P. J. Sutton, B. L. Swinkels, M. J. Szczepańczyk, M. Tacca, D. Talukder, D. B. Tanner, M. Tápai, A. Taracchini, J. A. Taylor, R. Taylor, T. Theeg, E. G. Thomas, M. Thomas, P. Thomas, K. A. Thorne, K. S. Thorne, E. Thrane, S. Tiwari, V. Tiwari, K. V. Tokmakov, K. Toland, M. Tonelli, Z. Tornasi, C. I. Torrie, D. Töyrä, F. Travasso, G. Traylor, D. Trifirò, J. Trinastic, M. C. Tringali, L. Trozzo, K. W. Tsang, M. Tse, R. Tso, D. Tuyenbayev, K. Ueno, 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. Vand er-Hyde, L. van der Schaaf, J. V. van Heijningen, A. A. van Veggel, M. Vardaro, V. Varma, S. Vass, M. Vasúth, A. Vecchio, G. Vedovato, J. Veitch, P. J. Veitch, K. Venkateswara, G. Venugopalan, D. Verkindt, F. Vetrano, A. Viceré, A. D. Viets, S. Vinciguerra, D. J. Vine, J. Y. Vinet, S. Vitale, T. Vo, H. Vocca, C. Vorvick, D. V. Voss, W. D. Vousden, S. P. Vyatchanin, A. R. Wade, L. E. Wade, M. Wade, R. Walet, M. Walker, L. Wallace, S. Walsh, G. Wang, H. Wang, J. Z. Wang, M. Wang, Y. F. Wang, Y. Wang, R. L. Ward, J. Warner, M. Was, J. Watchi, B. Weaver, L. W. Wei, M. Weinert, A. J. Weinstein, R. Weiss, L. Wen, E. K. Wessel, P. Weßels, T. Westphal, K. Wette, J. T. Whelan, B. F. Whiting, C. Whittle, D. Williams, R. D. Williams, A. R. Williamson, J. L. Willis, B. Willke, M. H. Wimmer, W. Winkler, C. C. Wipf, H. Wittel, G. Woan, J. Woehler, J. Wofford, K. W. K. Wong, J. Worden, J. L. Wright, D. S. Wu, G. Wu, W. Yam, H. Yamamoto, C. C. Yancey, M. J. Yap, Hang Yu, Haocun Yu, M. Yvert, A. Zadrožny, M. Zanolin, T. Zelenova, J. P. Zendri, M. Zevin, L. Zhang, M. Zhang, T. Zhang, Y. H. Zhang, C. Zhao, M. Zhou, Z. Zhou, S. J. Zhu, X. J. Zhu, M. E. Zucker, J. Zweizig, LIGO Scientific Collaboration, and Virgo Collaboration. Constraints on cosmic strings using data from the first Advanced LIGO observing run. *Phys. Rev. D*, 97(10):102002, 5 2018.

- [32] A. Vilenkin and E. P. S. Shellard. *Cosmic Strings and Other Topological Defects*. Cambridge University Press, 2000.
- [33] Alexander Vilenkin. Cosmic strings. *Phys. Rev. D*, 24(8):2082–2089, 10 1981.
- [34] A. Vilenkin. Cosmological Density Fluctuations Produced by Vacuum Strings. *Phys. Rev. Letters*, 46(17):1169–1172, 4 1981.
- [35] Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, N. Bartolo, E. Battaner, R. Battye, K. Benabed, A. Benoît,

A. Benoit-Lévy, J. P. Bernard, M. Bersanelli, P. Bielewicz, J. Bobin, J. J. Bock, A. Bonaldi, L. Bonavera, J. R. Bond, J. Borrill, F. R. Bouchet, M. Bridges, M. Bucher, C. Burigana, R. C. Butler, J. F. Cardoso, A. Catalano, A. Challinor, A. Chamballu, L. Y. Chiang, H. C. Chiang, 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, L. Danese, R. D. Davies, R. J. Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, J. M. Delouis, F. X. Désert, J. M. Diego, H. Dole, S. Donzelli, O. Doré, M. Douspis, A. Ducout, J. Dunkley, X. Dupac, G. Efstathiou, T. A. Enßlin, H. K. Eriksen, J. Fergusson, F. Finelli, O. Forni, M. Frailis, E. Franceschi, S. Galeotta, K. Ganga, 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. Hildebrandt, E. Hivon, M. Hobson, W. A. Holmes, A. Hornstrup, W. Hovest, K. M. Huffenberger, T. R. Jaffe, A. H. Jaffe, W. C. Jones, M. Juvela, E. Keihänen, R. Kesitalo, T. S. Kisner, 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. Lesgourgues, M. Liguori, P. B. Lilje, M. Linden-Vørnle, M. López-Cañiego, P. M. Lubin, J. F. Macías-Pérez, B. Maffei, D. Maino, N. Mandolesi, M. Maris, D. J. Marshall, P. G. Martin, E. Martínez-González, S. Masi, S. Matarrese, F. Matthai, P. Mazzotta, J. D. McEwen, A. Melchiorri, 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, 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, F. Pasian, G. Patanchon, H. V. Peiris, 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, C. R  th, R. Rebolo, M. Remazeilles, C. Renault, S. Ricciardi, T. Riller, C. Ringeval, I. Ristorcelli, G. Rocha, C. Rosset, G. Roudier, M. Rowan-Robinson, 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, 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, L. Valenziano, J. Valiviita, B. Van Tent, J. Varis, P. Vielva, F. Villa, N. Vittorio, L. A. Wade, B. D. Wandelt, D. Yvon, A. Zacchei, and A. Zonca. Planck 2013 results. XXV. Searches for cosmic strings and other topological defects. *AAP*, 571:A25, 11 2014.

- [36] Yanou Cui, Marek Lewicki, David E. Morrissey, and James D. Wells. Cosmic Archaeology with Gravitational Waves from Cosmic Strings. *Phys. Rev.*, D97(12):123505, 2018.
- [37] Yanou Cui, Marek Lewicki, David E. Morrissey, and James D. Wells. Probing the pre-BBN universe with gravitational waves from cosmic strings. *JHEP*, 01:081, 2019.
- [38] B. P. Abbott et al. Constraints on cosmic strings using data from the first Advanced LIGO observing run. *Phys. Rev.*, D97(10):102002, 2018.

- [39] Pierre Auclair et al. Probing the gravitational wave background from cosmic strings with LISA. *arXiv*, 2019.
- [40] Chia-Feng Chang and Yanou Cui. Stochastic Gravitational Wave Background from Global Cosmic Strings. *arXiv*, 2019.
- [41] Yann Gouttenoire, Graldine Servant, and Peera Simakachorn. BSM with Cosmic Strings: Heavy, up to EeV mass, Unstable Particles. *arXiv*, 2019.
- [42] Yann Gouttenoire, Graldine Servant, and Peera Simakachorn. Beyond the Standard Models with Cosmic Strings. *arXiv*, 2019.
- [43] Wilfried Buchmuller, Valerie Domcke, Hitoshi Murayama, and Kai Schmitz. Probing the scale of grand unification with gravitational waves. *arXiv*, 2019.
- [44] Simone Blasi, Vedran Brdar, and Kai Schmitz. Fingerprint of Low-Scale Leptogenesis in the Primordial Gravitational-Wave Spectrum. *arXiv e-prints*, 4 2020.
- [45] Jeff A. Dror, Takashi Hiramatsu, Kazunori Kohri, Hitoshi Murayama, and Graham White. Testing the Seesaw Mechanism and Leptogenesis with Gravitational Waves. *Phys. Rev. Lett.*, 124(4):041804, 2020.
- [46] Yanou Cui, Marek Lewicki, and David E. Morrissey. Gravitational Wave Bursts as Harbingers of Cosmic Strings Diluted by Inflation. *arXiv*, 2019.
- [47] Eric Morganson, Phil Marshall, Tommaso Treu, Tim Schrabback, and Roger D. Blandford. Direct observation of cosmic strings via their strong gravitational lensing effect - II. Results from the HST/ACS image archive. *MNRAS*, 406(4):2452–2472, 8 2010.
- [48] G. S. F. Guedes, P. P. Avelino, and L. Sousa. Signature of inflation in the stochastic gravitational wave background generated by cosmic string networks. *Phys. Rev.*, D98(12):123505, 2018.
- [49] Robert H. Brandenberger. Cosmic strings and the large-scale structure of the Universe. *Physica Scripta Volume T*, 36:114–126, 1 1991.
- [50] J. Silk and A. Vilenkin. Cosmic Strings and Galaxy Formation. *Phys. Rev. Letters*, 53(17):1700–1703, 10 1984.
- [51] Nathalie Deruelle and Bernard Linet. A cosmic string shock wave. *Classical and Quantum Gravity*, 5(1):55–58, 1 1988.
- [52] Y. B. Zel'Dovich. Reprint of 1970A&A.....5...84Z. Gravitational instability: an approximate theory for large density perturbations. *Astronomy and Astrophysics*, 500:13–18, 3 1970.
- [53] Disrael Camargo Neves da Cunha, Robert H. Brandenberger, and Oscar F. Hernández. Disruption of cosmic string wakes by Gaussian fluctuations. *Phys. Rev. D*, 93(12):123501, 6 2016.

- [54] J. Richard Bond, Lev Kofman, and Dmitry Pogosyan. How filaments of galaxies are woven into the cosmic web. *Nature (London)*, 380(6575):603–606, 4 1996.
- [55] Matthew Colless, Gavin Dalton, Steve Maddox, Will Sutherland, Peder Norberg, Shaun Cole, Joss Bland-Hawthorn, Terry Bridges, Russell Cannon, Chris Collins, Warrick Couch, Nicholas Cross, Kathryn Deeley, Roberto De Propris, Simon P. Driver, George Efstathiou, Richard S. Ellis, Carlos S. Frenk, Karl Glazebrook, Carole Jackson, Ofer Lahav, Ian Lewis, Stuart Lumsden, Darren Madgwick, John A. Peacock, Bruce A. Peterson, Ian Price, Mark Seaborne, and Keith Taylor. The 2dF Galaxy Redshift Survey: spectra and redshifts. *MNRAS*, 328(4):1039–1063, 12 2001.
- [56] D. Heath Jones, Mike A. Read, Will Saunders, Matthew Colless, Tom Jarrett, Quentin A. Parker, Anthony P. Fairall, Thomas Mauch, Elaine M. Sadler, Fred G. Watson, Donna Burton, Lachlan A. Campbell, Paul Cass, Scott M. Croom, John Dawe, Kristin Fiegert, Leela Frankcombe, Malcolm Hartley, John Huchra, Dionne James, Emma Kirby, Ofer Lahav, John Lucey, Gary A. Mamon, Lesa Moore, Bruce A. Peterson, Sayuri Prior, Dominique Proust, Ken Russell, Vicky Safouris, Ken-Ichi Wakamatsu, Eduard Westra, and Mary Williams. The 6dF Galaxy Survey: final redshift release (DR3) and southern large-scale structures. *MNRAS*, 399(2):683–698, 10 2009.
- [57] Daniel J. Eisenstein, David H. Weinberg, Eric Agol, Hiroaki Aihara, Carlos Allende Prieto, Scott F. Anderson, James A. Arns, Éric Aubourg, Stephen Bailey, Eduardo Balbinot, Robert Barkhouser, Timothy C. Beers, Andreas A. Berlind, Steven J. Bickerton, Dmitry Bizyaev, Michael R. Blanton, John J. Bochanski, Adam S. Bolton, Casey T. Bosman, Jo Bovy, W. N. Brandt, Ben Breslauer, Howard J. Brewington, J. Brinkmann, Peter J. Brown, Joel R. Brownstein, Dan Burger, Nicolas G. Busca, Heather Campbell, Phillip A. Cargile, William C. Carithers, Joleen K. Carlberg, Michael A. Carr, Liang Chang, Yanmei Chen, Cristina Chiappini, Johan Comparat, Natalia Connolly, Marina Cortes, Rupert A. C. Croft, Katia Cunha, Luiz N. da Costa, James R. A. Davenport, Kyle Dawson, Nathan De Lee, Gustavo F. Porto de Mello, Fernando de Simoni, Janice Dean, Saurav Dhital, Anne Ealet, Garrett L. Ebelke, Edward M. Edmondson, Jacob M. Eiting, Stephanie Escoffier, Massimiliano Esposito, Michael L. Evans, Xiaohui Fan, Bruno Femenía Castellá, Leticia Dutra Ferreira, Greg Fitzgerald, Scott W. Fleming, Andreu Font-Ribera, Eric B. Ford, Peter M. Frinchaboy, Ana Elia García Pérez, B. Scott Gaudi, Jian Ge, Luan Ghezzi, Bruce A. Gillespie, G. Gilmore, Léo Girardi, J. Richard Gott, Andrew Gould, Eva K. Grebel, James E. Gunn, Jean-Christophe Hamilton, Paul Harding, David W. Harris, Suzanne L. Hawley, Frederick R. Hearty, Joseph F. Hennawi, Jonay I. González Hernández, Shirley Ho, David W. Hogg, Jon A. Holtzman, Klaus Honscheid, Naohisa Inada, Inese I. Ivans, Linhua Jiang, Peng Jiang, Jennifer A. Johnson, Cathy Jordan, Wendell P. Jordan, Guinevere Kauffmann, Eyal Kazin, David Kirkby, Mark A. Klaene, G. R. Knapp, Jean-Paul Kneib, C. S. Kochanek, Lars Koesterke, Juna A. Kollmeier, Richard G. Kron, Hubert Lampeitl, Dustin Lang, James E. Lawler, Jean-Marc Le Goff, Brian L. Lee, Young Sun Lee, Jarron M. Leisenring, Yen-Ting Lin, Jian Liu, Daniel C. Long, Craig P. Loomis, Sara Lucatello, Britt Lundgren, Robert H.

Lupton, Bo Ma, Zhibo Ma, Nicholas MacDonald, Claude Mack, Suvrath Mahadevan, Marcio A. G. Maia, Steven R. Majewski, Martin Makler, Elena Malanushenko, Viktor Malanushenko, Rachel Mandelbaum, Claudia Maraston, Daniel Margala, Paul Mase-man, Karen L. Masters, Cameron K. McBride, Patrick McDonald, Ian D. McGreer, Richard G. McMahon, Olga Mena Requero, Brice Ménard, Jordi Miralda-Escudé, Heather L. Morrison, Fergal Mullally, Demetri Muna, Hitoshi Murayama, Adam D. Myers, Tracy Naugle, Angelo Fausti Neto, Duy Cuong Nguyen, Robert C. Nichol, David L. Nidever, Robert W. O’Connell, Ricardo L. C. Ogando, Matthew D. Olmstead, Daniel J. Oravetz, Nikhil Padmanabhan, Martin Paegert, Nathalie Palanque-Delabrouille, Kaike Pan, Parul Pandey, John K. Parejko, Isabelle Pâris, Paulo Pellegrini, Joshua Pepper, Will J. Percival, Patrick Petitjean, Robert Pfaffenberger, Janine Pforr, Stefanie Phleps, Christophe Pichon, Matthew M. Pieri, Francisco Prada, Adrian M. Price-Whelan, M. Jordan Raddick, Beatriz H. F. Ramos, I. Neill Reid, Celine Reyle, James Rich, Gordon T. Richards, George H. Rieke, Marcia J. Rieke, Hans-Walter Rix, Annie C. Robin, Helio J. Rocha-Pinto, Constance M. Rockosi, Natalie A. Roe, Emmanuel Rollinde, Ashley J. Ross, Nicholas P. Ross, Bruno Rossetto, Ariel G. Sánchez, Basilio Santiago, Conor Sayres, Ricardo Schiavon, David J. Schlegel, Katharine J. Schlesinger, Sarah J. Schmidt, Donald P. Schneider, Kris Sellgren, Alaina Shelden, Erin Sheldon, Matthew Shetrone, Yiping Shu, John D. Silverman, Jennifer Simmerer, Audrey E. Simmons, Thirupathi Sivarani, M. F. Skrutskie, Anže Slosar, Stephen Smee, Verne V. Smith, Stephanie A. Snedden, Keivan G. Stassun, Oliver Steele, Matthias Steinmetz, Mark H. Stockett, Todd Stollberg, Michael A. Strauss, Alexander S. Szalay, Masayuki Tanaka, Aniruddha R. Thakar, Daniel Thomas, Jeremy L. Tinker, Benjamin M. Tofflemire, Rita Tojeiro, Christy A. Tremonti, Mariana Vargas Magaña, Licia Verde, Nicole P. Vogt, David A. Wake, Xiaoke Wan, Ji Wang, Benjamin A. Weaver, Martin White, Simon D. M. White, John C. Wilson, John P. Wisniewski, W. Michael Wood-Vasey, Brian Yanny, Naoki Yasuda, Christophe Yèche, Donald G. York, Erick Young, Gail Zasowski, Idit Zehavi, and Bo Zhao. SDSS-III: Massive Spectroscopic Surveys of the Distant Universe, the Milky Way, and Extra-Solar Planetary Systems. *The Astronomical Journal*, 142(3):72, 9 2011.

- [58] Jeffrey A. Newman, Michael C. Cooper, Marc Davis, S. M. Faber, Alison L. Coil, Puragra Guhathakurta, David C. Koo, Andrew C. Phillips, Charlie Conroy, Aaron A. Dutton, Douglas P. Finkbeiner, Brian F. Gerke, David J. Rosario, Benjamin J. Weiner, C. N. A. Willmer, Renbin Yan, Justin J. Harker, Susan A. Kassin, N. P. Konidaris, Kamson Lai, Darren S. Madgwick, K. G. Noeske, Gregory D. Wirth, A. J. Connolly, N. Kaiser, Evan N. Kirby, Brian C. Lemaux, Lihwai Lin, Jennifer M. Lotz, G. A. Luppino, C. Marinoni, Daniel J. Matthews, Anne Metevier, and Ricardo P. Schiavon. The DEEP2 Galaxy Redshift Survey: Design, Observations, Data Reduction, and Redshifts. *ApJS*, 208(1):5, 9 2013.
- [59] J. Liske, I. K. Baldry, S. P. Driver, R. J. Tuffs, M. Alpaslan, E. Andrae, S. Brough, M. E. Cluver, M. W. Grootes, M. L. P. Gunawardhana, L. S. Kelvin, J. Loveday, A. S. G. Robotham, E. N. Taylor, S. P. Bamford, J. Bland-Hawthorn, M. J. I. Brown,

- M. J. Drinkwater, A. M. Hopkins, M. J. Meyer, P. Norberg, J. A. Peacock, N. K. Agius, S. K. Andrews, A. E. Bauer, J. H. Y. Ching, M. Colless, C. J. Conselice, S. M. Croom, L. J. M. Davies, R. De Propris, L. Dunne, E. M. Eardley, S. Ellis, C. Foster, C. S. Frenk, B. Häußler, B. W. Holwerda, C. Howlett, H. Ibarra, M. J. Jarvis, D. H. Jones, P. R. Kafle, C. G. Lacey, R. Lange, M. A. Lara-López, Á. R. López-Sánchez, S. Maddox, B. F. Madore, T. McNaught-Roberts, A. J. Moffett, R. C. Nichol, M. S. Owers, D. Palamara, S. J. Penny, S. Phillipps, K. A. Pimblet, C. C. Popescu, M. Prescott, R. Proctor, E. M. Sadler, A. E. Sansom, M. Seibert, R. Sharp, W. Sutherland, J. A. Vázquez-Mata, E. van Kampen, S. M. Wilkins, R. Williams, and A. H. Wright. Galaxy And Mass Assembly (GAMA): end of survey report and data release 2. *MNRAS*, 452(2):2087–2126, 9 2015.
- [60] M. Scodeggio, L. Guzzo, B. Garilli, B. R. Granett, M. Bolzonella, S. de la Torre, U. Abbas, C. Adami, S. Arnouts, D. Bottini, A. Cappi, J. Coupon, O. Cucciati, I. Davidzon, P. Franzetti, A. Fritz, A. Iovino, J. Krywult, V. Le Brun, O. Le Fèvre, D. Maccagni, K. Małek, A. Marchetti, F. Marulli, M. Polletta, A. Pollo, L. A. M. Tasca, R. Tojeiro, D. Vergani, A. Zanichelli, J. Bel, E. Branchini, G. De Lucia, O. Ilbert, H. J. McCracken, T. Moutard, J. A. Peacock, G. Zamorani, A. Burden, M. Fumana, E. Jullo, C. Marinoni, Y. Mellier, L. Moscardini, and W. J. Percival. The VIMOS Public Extragalactic Redshift Survey (VIPERS). Full spectroscopic data and auxiliary information release (PDR-2). *Astronomy and Astrophysics*, 609:A84, 1 2018.
- [61] Tzu-Ching Chang, Ue-Li Pen, Kevin Bandura, and Jeffrey B. Peterson. An intensity map of hydrogen 21-cm emission at redshift  $z \sim 0.8$ . *Nature (London)*, 466(7305):463–465, 7 2010.
- [62] C. J. Anderson, N. J. Luciw, Y. C. Li, C. Y. Kuo, J. Yadav, K. W. Masui, T. C. Chang, X. Chen, N. Oppermann, Y. W. Liao, U. L. Pen, D. C. Price, L. Staveley-Smith, E. R. Switzer, P. T. Timbie, and L. Wolz. Low-amplitude clustering in low-redshift 21-cm intensity maps cross-correlated with 2dF galaxy densities. *MNRAS*, 476(3):3382–3392, 5 2018.
- [63] Ely Kovetz, Patrick C. Breyse, Adam Lidz, Jamie Bock, Charles M. Bradford, Tzu-Ching Chang, Simon Foreman, Hamsa Padmanabhan, Anthony Pullen, Dominik Riechers, Marta B. Silva, and Eric Switzer. Astrophysics and Cosmology with Line-Intensity Mapping. *BAAS*, 51(3):101, 5 2019.
- [64] M. Santos, P. Bull, D. Alonso, S. Camera, P. Ferreira, G. Bernardi, R. Maartens, M. Viel, F. Villaescusa-Navarro, F. B. Abdalla, M. Jarvis, R. B. Metcalf, A. Pourtsidou, and L. Wolz. Cosmology from a SKA HI intensity mapping survey. In *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, page 19, 4 2015.
- [65] Emanuele Castorina, Simon Foreman, Dionysios Karagiannis, Adrian Liu, Kiyoshi W. Masui, Pieter D. Meerburg, Laura B. Newburgh, Paul O’Connor, Andrej Obuljen, Hamsa Padmanabhan, J. Richard Shaw, Anže Slosar, Paul Stankus, Peter T. Tim-

- bie, Benjamin Wallisch, and Martin White. Packed Ultra-wideband Mapping Array (PUMA): Astro2020 RFI Response. *arXiv e-prints*, page arXiv:2002.05072, 2 2020.
- [66] Square Kilometre Array Cosmology Science Working Group, David J. Bacon, Richard A. Battye, Philip Bull, Stefano Camera, Pedro G. Ferreira, Ian Harrison, David Parkinson, Alkistis Pourtsidou, Mario G. Santos, Laura Wolz, Filipe Abdalla, Yashar Akrami, David Alonso, Sambatra Andrianomena, Mario Ballardini, Jose Luis Bernal, Daniele Bertacca, Carlos A. P. Bengaly, Anna Bonaldi, Camille Bonvin, Michael L. Brown, Emma Chapman, Song Chen, Xuelei Chen, Steven Cunningham, Tamara M. Davis, Clive Dickinson, Jose Fonseca, Keith Grainge, Stuart Harper, Matt J. Jarvis, Roy Maartens, Natasha Maddox, Hamsa Padmanabhan, Jonathan R. Pritchard, Alvise Raccanelli, Marzia Rivi, Sambit Roychowdhury, Martin Sahlen, Dominik J. Schwarz, Thilo M. Siewert, Matteo Viel, Francisco Villaescusa-Navarro, Yidong Xu, Daisuke Yamauchi, and Joe Zuntz. Cosmology with Phase 1 of the Square Kilometre Array; Red Book 2018: Technical specifications and performance forecasts. *arXiv e-prints*, page arXiv:1811.02743, 11 2018.
- [67] DESI Collaboration, Amir Aghamousa, Jessica Aguilar, Steve Ahlen, Shadab Alam, Lori E. Allen, Carlos Allende Prieto, James Annis, Stephen Bailey, Christophe Balleland, Otger Ballester, Charles Baltay, Lucas Beaufore, Chris Bebek, Timothy C. Beers, Eric F. Bell, José Luis Bernal, Robert Besuner, Florian Beutler, Chris Blake, Hannes Bleuler, Michael Blomqvist, Robert Blum, Adam S. Bolton, Cesar Briceno, David Brooks, Joel R. Brownstein, Elizabeth Buckley-Geer, Angela Burden, Etienne Burtin, Nicolas G. Busca, Robert N. Cahn, Yan-Chuan Cai, Laia Cardiel-Sas, Raymond G. Carlberg, Pierre-Henri Carton, Ricard Casas, Francisco J. Castander, Jorge L. Cervantes-Cota, Todd M. Claybaugh, Madeline Close, Carl T. Coker, Shaun Cole, Johan Comparat, Andrew P. Cooper, M. C. Cousinou, Martin Crocce, Jean-Gabriel Cuby, Daniel P. Cunningham, Tamara M. Davis, Kyle S. Dawson, Axel de la Macorra, Juan De Vicente, Timothée Delubac, Mark Derwent, Arjun Dey, Govinda Dhungana, Zhejie Ding, Peter Doel, Yutong T. Duan, Anne Ealet, Jerry Edelstein, Sarah Eftekharzadeh, Daniel J. Eisenstein, Ann Elliott, Stéphanie Escoffier, Matthew Evatt, Parker Fagrelus, Xiaohui Fan, Kevin Fanning, Arya Farahi, Jay Farihi, Ginevra Favole, Yu Feng, Enrique Fernandez, Joseph R. Findlay, Douglas P. Finkbeiner, Michael J. Fitzpatrick, Brenna Flaugher, Samuel Flender, Andreu Font-Ribera, Jaime E. Forero-Romero, Pablo Fosalba, Carlos S. Frenk, Michele Fumagalli, Boris T. Gaensicke, Giuseppe Gallo, Juan Garcia-Bellido, Enrique Gaztanaga, Nicola Pietro Gentile Fusillo, Terry Gerard, Irena Gershkovich, Tommaso Giannantonio, Denis Gillet, Guillermo Gonzalez-de-Rivera, Violeta Gonzalez-Perez, Shelby Gott, Or Graur, Gaston Gutierrez, Julien Guy, Salman Habib, Henry Heetderks, Ian Heetderks, Katrin Heitmann, Wojciech A. Hellwing, David A. Herrera, Shirley Ho, Stephen Holland, Klaus Honscheid, Eric Huff, Timothy A. Hutchinson, Dragan Huterer, Ho Seong Hwang, Joseph Maria Illa Laguna, Yuzo Ishikawa, Dianna Jacobs, Niall Jeffrey, Patrick Jelinsky, Elise Jennings, Linhua Jiang, Jorge Jimenez, Jennifer Johnson, Richard Joyce, Eric Jullo, Stéphanie Juneau, Sami Kama, Armin Karcher, Sonia Karkar, Robert Kehoe, Noble Kennamer, Stephen Kent, Martin Kilbinger,

Alex G. Kim, David Kirkby, Theodore Kisner, Ellie Kitanidis, Jean-Paul Kneib, Sergey Koposov, Eve Kovacs, Kazuya Koyama, Anthony Kremin, Richard Kron, Luzius Kronig, Andrea Kueter-Young, Cedric G. Lacey, Robin Lafever, Ofer Lahav, Andrew Lambert, Michael Lampton, Martin Landriau, Dustin Lang, Tod R. Lauer, Jean-Marc Le Goff, Laurent Le Guillou, Auguste Le Van Suu, Jae Hyeon Lee, Su-Jeong Lee, Daniela Leitner, Michael Lesser, Michael E. Levi, Benjamin L’Huillier, Baojiu Li, Ming Liang, Huan Lin, Eric Linder, Sarah R. Loebman, Zarija Lukić, Jun Ma, Niall MacCrann, Christophe Magneville, Laleh Makarem, Marc Manera, Christopher J. Manser, Robert Marshall, Paul Martini, Richard Massey, Thomas Matheson, Jeremy McCauley, Patrick McDonald, Ian D. McGreer, Aaron Meisner, Nigel Metcalfe, Timothy N. Miller, Ramon Miquel, John Moustakas, Adam Myers, Milind Naik, Jeffrey A. Newman, Robert C. Nichol, Andrina Nicola, Luiz Nicolati da Costa, Jundan Nie, Gustavo Niz, Peder Norberg, Brian Nord, Dara Norman, Peter Nugent, Thomas O’Brien, Minji Oh, Knut A. G. Olsen, Cristobal Padilla, Hamsa Padmanabhan, Nikhil Padmanabhan, Nathalie Palanque-Delabrouille, Antonella Palmese, Daniel Pappalardo, Isabelle Pâris, Changbom Park, Anna Patej, John A. Peacock, Hiranya V. Peiris, Xiyang Peng, Will J. Percival, Sandrine Perruchot, Matthew M. Pieri, Richard Pogge, Jennifer E. Pollack, Claire Poppett, Francisco Prada, Abhishek Prakash, Ronald G. Probst, David Rabinowitz, Anand Raichoor, Chang Hee Ree, Alexandre Refregier, Xavier Regal, Beth Reid, Kevin Reil, Mehdi Rezaie, Constance M. Rockosi, Natalie Roe, Samuel Ronayette, Aaron Roodman, Ashley J. Ross, Nicholas P. Ross, Graziano Rossi, Eduardo Roza, Vanina Ruhlmann-Kleider, Eli S. Rykoff, Cristiano Sabiu, Lado Samushia, Eusebio Sanchez, Javier Sanchez, David J. Schlegel, Michael Schneider, Michael Schubnell, Aurélia Secroun, Uros Seljak, Hee-Jong Seo, Santiago Serrano, Arman Shafieloo, Huanyuan Shan, Ray Sharples, Michael J. Sholl, William V. Shourt, Joseph H. Silber, David R. Silva, Martin M. Sirk, Anze Slosar, Alex Smith, George F. Smoot, Debopam Som, Yong-Seon Song, David Sprayberry, Ryan Staten, Andy Stefanik, Gregory Tarle, Suk Sien Tie, Jeremy L. Tinker, Rita Tojeiro, Francisco Valdes, Octavio Valenzuela, Monica Valluri, Mariana Vargas-Magana, Licia Verde, Alistair R. Walker, Jiali Wang, Yuting Wang, Benjamin A. Weaver, Curtis Weaverdyck, Risa H. Wechsler, David H. Weinberg, Martin White, Qian Yang, Christophe Yèche, Tianmeng Zhang, Gong-Bo Zhao, Yi Zheng, Xu Zhou, Zhimin Zhou, Yaling Zhu, Hu Zou, and Ying Zu. The DESI Experiment Part I: Science, Targeting, and Survey Design. *arXiv e-prints*, page arXiv:1611.00036, October 2016.

- [68] Euclid Collaboration, A. Blanchard, S. Camera, C. Carbone, V. F. Cardone, S. Casas, S. Ilić, M. Kilbinger, T. Kitching, M. Kunz, F. Lacasa, E. Linder, E. Majerotto, K. Markovič, M. Martinelli, V. Pettorino, A. Pourtsidou, Z. Sakr, A. G. Sánchez, D. Sapone, I. Tutusaus, S. Yahia-Cherif, V. Yankelevich, S. Andreon, H. Aussel, A. Balaguera-Antolínez, M. Baldi, S. Bardelli, R. Bender, A. Biviano, D. Bonino, A. Boucaud, E. Bozzo, E. Branchini, S. Brau-Nogue, M. Brescia, J. Brinchmann, C. Burigana, R. Cabanac, V. Capobianco, A. Cappi, J. Carretero, C. S. Carvalho, R. Casas, F. J. Castander, M. Castellano, S. Cavuoti, A. Cimatti, R. Cledassou, C. Colodro-Conde, G. Congedo, C. J. Conselice, L. Conversi, Y. Copin, L. Corcione, J. Coupon, H. M. Courtois, M. Cropper, A. Da Silva, S. de la Torre, D. Di Ferdinand



- o, F. Dubath, F. Ducret, C. A. J. Duncan, X. Dupac, S. Dusini, G. Fabbian, M. Fabricius, S. Farrens, P. Fosalba, S. Fotopoulou, N. Fourmanoit, M. Frailis, E. Franceschi, P. Franzetti, M. Fumana, S. Galeotta, W. Gillard, B. Gillis, C. Giocoli, P. Gómez-Alvarez, J. Graciá-Carpio, F. Grupp, L. Guzzo, H. Hoekstra, F. Hormuth, H. Israel, K. Jahnke, E. Keihanen, S. Kermiche, C. C. Kirkpatrick, R. Kohley, B. Kubik, H. Kurki-Suonio, S. Ligori, P. B. Lilje, I. Lloro, D. Maino, E. Maiorano, O. Marggraf, N. Martinet, F. Marulli, R. Massey, E. Medinaceli, S. Mei, Y. Mellier, B. Metcalf, J. J. Metge, G. Meylan, M. Moresco, L. Moscardini, E. Munari, R. C. Nichol, S. Niemi, A. A. Nucita, C. Padilla, S. Paltani, F. Pasian, W. J. Percival, S. Pires, G. Polenta, M. Poncet, L. Pozzetti, G. D. Racca, F. Raison, A. Renzi, J. Rhodes, E. Romelli, M. Roncarelli, E. Rossetti, R. Saglia, P. Schneider, V. Scottez, A. Secroun, G. Sirri, L. Stanco, J. L. Starck, F. Sureau, P. Tallada-Crespí, D. Tavagnacco, A. N. Taylor, M. Tenti, I. Tereno, R. Toledo-Moreo, F. Torradeflot, L. Valenziano, T. Vassallo, G. A. Verdoes Kleijn, M. Viel, Y. Wang, A. Zacchei, J. Zoubian, and E. Zucca. Euclid preparation: VII. Forecast validation for Euclid cosmological probes. *arXiv e-prints*, page arXiv:1910.09273, 10 2019.
- [69] D. Spergel, N. Gehrels, C. Baltay, D. Bennett, J. Brekinridge, M. Donahue, A. Dressler, B. S. Gaudi, T. Greene, O. Guyon, C. Hirata, J. Kalirai, N. J. Kasdin, B. Macintosh, W. Moos, S. Perlmutter, M. Postman, B. Rauscher, J. Rhodes, Y. Wang, D. Weinberg, D. Benford, M. Hudson, W. S. Jeong, Y. Mellier, W. Traub, T. Yamada, P. Capak, J. Colbert, D. Masters, M. Penny, D. Savransky, D. Stern, N. Zimmerman, R. Barry, L. Bartusek, K. Carpenter, E. Cheng, D. Content, F. Dekens, R. Demers, K. Grady, C. Jackson, G. Kuan, J. Kruk, M. Melton, B. Nemat, B. Parvin, I. Poberezhskiy, C. Peddie, J. Ruffa, J. K. Wallace, A. Whipple, E. Wollack, and F. Zhao. Wide-Field Infrared Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report. *arXiv e-prints*, page arXiv:1503.03757, 3 2015.
- [70] Noam I. Libeskind, Rien van de Weygaert, Marius Cautun, Bridget Falck, Elmo Tempel, Tom Abel, Mehmet Alpaslan, Miguel A. Aragón-Calvo, Jaime E. Forero-Romero, Roberto Gonzalez, Stefan Gottlöber, Oliver Hahn, Wojciech A. Hellwing, Yehuda Hoffman, Bernard J. T. Jones, Francisco Kitaura, Alexander Knebe, Serena Manti, Mark Neyrinck, Sebastián E. Nuza, Nelson Padilla, Erwin Platen, Nesar Ramachandra, Aaron Robotham, Enn Saar, Sergei Shandarin, Matthias Steinmetz, Radu S. Stoica, Thierry Sousbie, and Gustavo Yepes. Tracing the cosmic web. *MNRAS*, 473(1):1195–1217, 1 2018.
- [71] E. Tempel, R. S. Stoica, V. J. Martínez, L. J. Liivamägi, G. Castellan, and E. Saar. Detecting filamentary pattern in the cosmic web: a catalogue of filaments for the SDSS. *MNRAS*, 438(4):3465–3482, 3 2014.
- [72] Nicola Malavasi, Nabila Aghanim, Marian Douspis, Hideki Tanimura, and Victor Bonjean. Characterising filaments in the SDSS volume from the galaxy distribution. *arXiv e-prints*, page arXiv:2002.01486, 2 2020.

- [73] Samuel Laliberte, Robert Brandenberger, and Disrael Camargo Neves da Cunha. Cosmic String Wake Detection using 3D Ridgelet Transformations. *arXiv e-prints*, page arXiv:1807.09820, 7 2018.
- [74] Disrael Camargo Neves da Cunha. Signature of a Cosmic String Wake at  $z = 3$ . *arXiv e-prints*, page arXiv:1810.07737, 10 2018.
- [75] Disrael Cunha, Joachim Harnois-Deraps, Robert Brandenberger, Adam Amara, and Alexandre Refregier. Dark Matter Distribution Induced by a Cosmic String Wake in the Nonlinear Regime. *arXiv e-prints*, page arXiv:1804.00083, 3 2018.
- [76] Robert H. Brandenberger, Rebecca J. Danos, Oscar F. Hernández, and Gilbert P. Holder. The 21 cm signature of cosmic string wakes. *JCAP*, 2010(12):028, 12 2010.
- [77] Simeon Bird, Keir K. Rogers, Hiranya V. Peiris, Licia Verde, Andreu Font-Ribera, and Andrew Pontzen. An emulator for the Lyman- $\alpha$  forest. *JCAP*, 2019(2):050, February 2019.
- [78] Volker Springel. The cosmological simulation code GADGET-2. *MNRAS*, 364(4):1105–1134, December 2005.
- [79] Julien Lesgourgues. The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview. *arXiv e-prints*, page arXiv:1104.2932, April 2011.
- [80] G. Hinshaw, D. Larson, E. Komatsu, D. N. Spergel, C. L. Bennett, J. Dunkley, M. R. Nolta, M. Halpern, R. S. Hill, N. Odegard, L. Page, K. M. Smith, J. L. Weiland, B. Gold, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, E. Wollack, and E. L. Wright. Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results. *ApJS*, 208(2):19, October 2013.
- [81] Tanmay Vachaspati, Levon Pogosian, and Danièle A. Steer. Cosmic strings. *Scholarpedia*, 10(2):31682, 1 2015.
- [82] Jose J. Blanco-Pillado, Ken D. Olum, and Benjamin Shlaer. Large parallel cosmic string simulations: New results on loop production. *Phys. Rev. D*, 83(8):083514, 4 2011.
- [83] Umut Ozertem and Deniz Erdogmus. Locally defined principal curves and surfaces. *J. Mach. Learn. Res.*, 12(1249):1249–1286, 7 2011.
- [84] Yen-Chi Chen, Shirley Ho, Peter E. Freeman, Christopher R. Genovese, and Larry Wasserman. Cosmic web reconstruction through density ridges: method and algorithm. *MNRAS*, 454(1):1140–1156, 11 2015.
- [85] T. Sousbie. The persistent cosmic web and its filamentary structure - I. Theory and implementation. *MNRAS*, 414(1):350–383, 6 2011.
- [86] T. Sousbie, C. Pichon, and H. Kawahara. The persistent cosmic web and its filamentary structure - II. Illustrations. *MNRAS*, 414(1):384–403, 6 2011.

- [87] Yen-Chi Chen, Shirley Ho, Jon Brinkmann, Peter E. Freeman, Christopher R. Genovese, Donald P. Schneider, and Larry Wasserman. Cosmic web reconstruction through density ridges: catalogue. *MNRAS*, 461(4):3896–3909, 10 2016.
- [88] Astropy Collaboration, T. P. Robitaille, E. J. Tollerud, P. Greenfield, M. Droettboom, E. Bray, T. Aldcroft, M. Davis, A. Ginsburg, A. M. Price-Whelan, W. E. Kerzendorf, A. Conley, N. Crighton, K. Barbary, D. Muna, H. Ferguson, F. Grollier, M. M. Parikh, P. H. Nair, H. M. Unther, C. Deil, J. Woillez, S. Conseil, R. Kramer, J. E. H. Turner, L. Singer, R. Fox, B. A. Weaver, V. Zabalza, Z. I. Edwards, K. Azalee Bostroem, D. J. Burke, A. R. Casey, S. M. Crawford, N. Dencheva, J. Ely, T. Jenness, K. Labrie, P. L. Lim, F. Pierfederici, A. Pontzen, A. Ptak, B. Refsdal, M. Servillat, and O. Streicher. Astropy: A community Python package for astronomy. *Astronomy and Astrophysics*, 558:A33, 10 2013.
- [89] A. M. Price-Whelan, B. M. Sipőcz, H. M. Günther, P. L. Lim, S. M. Crawford, S. Conseil, D. L. Shupe, M. W. Craig, N. Dencheva, A. Ginsburg, J. T. VanderPlas, L. D. Bradley, D. Pérez-Suárez, M. de Val-Borro, (Primary Paper Contributors, T. L. Aldcroft, K. L. Cruz, T. P. Robitaille, E. J. Tollerud, (Astropy Coordination Committee, C. Ardelean, T. Babej, Y. P. Bach, M. Bachetti, A. V. Bakanov, S. P. Bamford, G. Barentsen, P. Barmby, A. Baumbach, K. L. Berry, F. Biscani, M. Boquien, K. A. Bostroem, L. G. Bouma, G. B. Brammer, E. M. Bray, H. Breytenbach, H. Buddelmeijer, D. J. Burke, G. Calderone, J. L. Cano Rodríguez, M. Cara, J. V. M. Cardoso, S. Cheedella, Y. Copin, L. Corrales, D. Crichton, D. D’Avella, C. Deil, É. Depagne, J. P. Dietrich, A. Donath, M. Droettboom, N. Earl, T. Erben, S. Fabbro, L. A. Ferreira, T. Finethy, R. T. Fox, L. H. Garrison, S. L. J. Gibbons, D. A. Goldstein, R. Gommers, J. P. Greco, P. Greenfield, A. M. Groener, F. Grollier, A. Hagen, P. Hirst, D. Homeier, A. J. Horton, G. Hosseinzadeh, L. Hu, J. S. Hunkeler, Ž. Ivezić, A. Jain, T. Jenness, G. Kanarek, S. Kendrew, N. S. Kern, W. E. Kerzendorf, A. Khvalko, J. King, D. Kirkby, A. M. Kulkarni, A. Kumar, A. Lee, D. Lenz, S. P. Littlefair, Z. Ma, D. M. Macleod, M. Mastroiello, C. McCully, S. Montagnac, B. M. Morris, M. Mueller, S. J. Mumford, D. Muna, N. A. Murphy, S. Nelson, G. H. Nguyen, J. P. Ninan, M. Nöthe, S. Ogaz, S. Oh, J. K. Parejko, N. Parley, S. Pascual, R. Patil, A. A. Patil, A. L. Plunkett, J. X. Prochaska, T. Rastogi, V. Reddy Janga, J. Sabater, P. Sakurikar, M. Seifert, L. E. Sherbert, H. Sherwood-Taylor, A. Y. Shih, J. Sick, M. T. Silbiger, S. Singanamalla, L. P. Singer, P. H. Sladen, K. A. Sooley, S. Sornarajah, O. Streicher, P. Teuben, S. W. Thomas, G. R. Tremblay, J. E. H. Turner, V. Terrón, M. H. van Kerkwijk, A. de la Vega, L. L. Watkins, B. A. Weaver, J. B. Whitmore, J. Woillez, V. Zabalza, and (Astropy Contributors. The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package. *The Astronomical Journal*, 156:123, 9 2018.
- [90] 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. Mor-tonson, 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*, page arXiv:1305.5425, 5 2013.
- [91] Luca Amendola, Stephen Appleby, Anastasios Avgoustidis, David Bacon, Tessa Baker, Marco Baldi, Nicola Bartolo, Alain Blanchard, Camille Bonvin, Stefano Borgani, Enzo Branchini, Clare Burrage, Stefano Camera, Carmelita Carbone, Luciano Casarini, Mark Cropper, Claudia de Rham, Jörg P. Dietrich, Cinzia Di Porto, Ruth Durrer, Anne Ealet, Pedro G. Ferreira, Fabio Finelli, Juan García-Bellido, Tommaso Giannantonio, Luigi Guzzo, Alan Heavens, Lavinia Heisenberg, Catherine Heymans, Henk Hoekstra, Lukas Hollenstein, Rory Holmes, Zhiqi Hwang, Knud Jahnke, Thomas D. Kitching, Tomi Koivisto, Martin Kunz, Giuseppe La Vacca, Eric Linder, Marisa March, Valerio Marra, Carlos Martins, Elisabetta Majerotto, Dida Markovic, David Marsh, Federico Marulli, Richard Massey, Yannick Mellier, Francesco Montanari, David F. Mota, Nelson J. Nunes, Will Percival, Valeria Pettorino, Cristiano Porciani, Claudia Quercellini, Justin Read, Massimiliano Rinaldi, Domenico Sapone, Ignacy Sawicki, Roberto Scaramella, Constantinos Skordis, Fergus Simpson, Andy Taylor, Shaun Thomas, Roberto Trotta, Licia Verde, Filippo Vernizzi, Adrian Vollmer, Yun Wang, Jochen Weller, and Tom Zlosnik. Cosmology and fundamental physics with the Euclid satellite. *Living Reviews in Relativity*, 21(1):2, 4 2018.
- [92] Xiaohui Fan, Michael A. Strauss, Robert H. Becker, Richard L. White, James E. Gunn, Gillian R. Knapp, Gordon T. Richards, Donald P. Schneider, J. Brinkmann, and Masataka Fukugita. Constraining the Evolution of the Ionizing Background and the Epoch of Reionization with  $z \sim 6$  Quasars. II. A Sample of 19 Quasars. *AJ*, 132(1):117–136, July 2006.
- [93] B. E. Robertson, R. S. Ellis, J. S. Dunlop, R. J. McLure, and D. P. Stark. Early star-forming galaxies and the reionization of the Universe. *Nature*, 468:49–55, November 2010.
- [94] E. Komatsu et al. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation. *ApJS*, 192:18, February 2011.
- [95] Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak, R. Battye, K. Benabed, J. P. Bernard, M. Bersanelli, P. Bielewicz, J. J. Bock, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J. F. Cardoso, J. Carron, A. Challinor, H. C. Chiang, J. Chluba, L. P. L. Colombo, C. Combet, D. Contreras, B. P. Crill, F. Cuttaia, P. de Bernardis, G. de Zotti, J. Delabrouille, J. M. Delouis, E. Di Valentino, J. M. Diego, O. Doré, M. Douspis, A. Ducout, X. Dupac, S. Dusini, G. Efstathiou, F. Elsner, T. A. Enßlin, H. K. Eriksen, Y. Fantaye, M. Farhang, J. Fergusson, R. Fernandez-Cobos, F. Finelli, F. Forastieri, M. Frailis, A. A. Fraisse, E. Franceschi, A. Frolov, S. Galeotta, S. Galli, K. Ganga, R. T. Génova-Santos, M. Gerbino, T. Ghosh, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gruppuso, J. E. Gudmundsson, J. Hamann, W. Handley,

- F. K. Hansen, D. Herranz, S. R. Hildebrandt, E. Hivon, Z. Huang, A. H. Jaffe, W. C. Jones, A. Karakci, E. Keihänen, R. Keskitalo, K. Kiiveri, J. Kim, T. S. Kisner, L. Knox, N. Krachmalnicoff, M. Kunz, H. Kurki-Suonio, G. Lagache, J. M. Lamarre, A. Lasenby, M. Lattanzi, C. R. Lawrence, M. Le Jeune, P. Lemos, J. Lesgourgues, F. Levrier, A. Lewis, M. Liguori, P. B. Lilje, M. Lilley, V. Lindholm, M. López-Cañiego, P. M. Lubin, Y. Z. Ma, J. F. Macías-Pérez, G. Maggio, D. Maino, N. Mandolesi, A. Mangilli, A. Marcos-Caballero, M. Maris, P. G. Martin, M. Martinelli, E. Martínez-González, S. Matarrese, N. Mauri, J. D. McEwen, P. R. Meinhold, A. Melchiorri, A. Mennella, M. Migliaccio, M. Millea, S. Mitra, M. A. Miville-Deschênes, D. Molinari, L. Montier, G. Morgante, A. Moss, P. Natoli, H. U. Nørgaard-Nielsen, L. Pagano, D. Paoletti, B. Partridge, G. Patanchon, H. V. Peiris, F. Perrotta, V. Pettorino, F. Piacentini, L. Polastri, G. Polenta, J. L. Puget, J. P. Rachen, M. Reinecke, M. Remazeilles, A. Renzi, G. Rocha, C. Rosset, G. Roudier, J. A. Rubiño-Martín, B. Ruiz-Granados, L. Salvati, M. Sandri, M. Savelainen, D. Scott, E. P. S. Shellard, C. Sirignano, G. Sirri, L. D. Spencer, R. Sunyaev, A. S. Suur-Uski, J. A. Tauber, D. Tavagnacco, M. Tenti, L. Toffolatti, M. Tomasi, T. Trombetti, L. Valenziano, J. Valiviita, B. Van Tent, L. Vibert, P. Vielva, F. Villa, N. Vittorio, B. D. Wandelt, I. K. Wehus, M. White, S. D. M. White, A. Zacchei, and A. Zonca. Planck 2018 results. VI. Cosmological parameters. *arXiv e-prints*, page arXiv:1807.06209, July 2018.
- [96] Elisa Boera, George D. Becker, James S. Bolton, and Fahad Nasir. Revealing Reionization with the Thermal History of the Intergalactic Medium: New Constraints from the Ly $\alpha$  Flux Power Spectrum. *ApJ*, 872(1):101, February 2019.
- [97] Jordi Miralda-Escudé, Martin Haehnelt, and Martin J. Rees. Reionization of the Inhomogeneous Universe. *ApJ*, 530(1):1–16, Feb 2000.
- [98] J. Stuart B. Wyithe and Abraham Loeb. Reionization of Hydrogen and Helium by Early Stars and Quasars. *ApJ*, 586(2):693–708, April 2003.
- [99] Steven R. Furlanetto and S. Peng Oh. The History and Morphology of Helium Reionization. *ApJ*, 681(1):1–17, July 2008.
- [100] Michael Shull, Kevin France, Charles Danforth, Britton Smith, and Jason Tumlinson. Hubble/COS Observations of the Quasar HE 2347-4342: Probing the Epoch of He II Patchy Reionization at Redshifts  $z = 2.4$ - $2.9$ . *arXiv e-prints*, page arXiv:1008.2957, August 2010.
- [101] Gábor Worseck, J. Xavier Prochaska, Joseph F. Hennawi, and Matthew McQuinn. Early and Extended Helium Reionization over More Than 600 Million Years of Cosmic Time. *ApJ*, 825(2):144, July 2016.
- [102] J. Miralda-Escudé and M. J. Rees. Reionization and thermal evolution of a photoionized intergalactic medium. *MNRAS*, 266:343–352, January 1994.
- [103] Lam Hui and Nickolay Y. Gnedin. Equation of state of the photoionized intergalactic medium. *MNRAS*, 292(1):27–42, November 1997.

- [104] J. Schaye, T. Theuns, M. Rauch, G. Efstathiou, and W. L. W. Sargent. The thermal history of the intergalactic medium\*. *MNRAS*, 318:817–826, November 2000.
- [105] L. Hui and Z. Haiman. The Thermal Memory of Reionization History. *ApJ*, 596:9–18, October 2003.
- [106] Phoebe R. Upton Sanderbeck, Anson D’Aloisio, and Matthew J. McQuinn. Models of the thermal evolution of the intergalactic medium after reionization. *MNRAS*, 460(2):1885–1897, August 2016.
- [107] Anson D’Aloisio, Matthew McQuinn, Oliver Maupin, Frederick B. Davies, Hy Trac, Spencer Fuller, and Phoebe R. Upton Sanderbeck. Heating of the Intergalactic Medium by Hydrogen Reionization. *ApJ*, 874(2):154, April 2019.
- [108] Matthew McQuinn and Phoebe R. Upton Sanderbeck. On the intergalactic temperature-density relation. *MNRAS*, 456(1):47–54, February 2016.
- [109] Matteo Viel, Julien Lesgourgues, Martin G. Haehnelt, Sabino Matarrese, and Antonio Riotto. Constraining warm dark matter candidates including sterile neutrinos and light gravitinos with WMAP and the Lyman- $\alpha$  forest. *Phys. Rev. D*, 71(6):063534, March 2005.
- [110] Michael Walther, Jose Oñorbe, Joseph F. Hennawi, and Zarija Lukić. New Constraints on IGM Thermal Evolution from the Ly $\alpha$  Forest Power Spectrum. *ApJ*, 872(1):13, February 2019.
- [111] Matias Zaldarriaga, Lam Hui, and Max Tegmark. Constraints from the Ly $\alpha$  Forest Power Spectrum. *ApJ*, 557(2):519–526, August 2001.
- [112] Nathalie Palanque-Delabrouille, Christophe Yèche, Arnaud Borde, Jean-Marc Le Goff, Graziano Rossi, Matteo Viel, Éric Aubourg, Stephen Bailey, Julian Bautista, Michael Blomqvist, Adam Bolton, James S. Bolton, Nicolás G. Busca, Bill Carithers, Rupert A. C. Croft, Kyle S. Dawson, Timothée Delubac, Andreu Font-Ribera, Shirley Ho, David Kirkby, Khee-Gan Lee, Daniel Margala, Jordi Miralda-Escudé, Demitri Muna, Adam D. Myers, Pasquier Noterdaeme, Isabelle Pâris, Patrick Petitjean, Matthew M. Pieri, James Rich, Emmanuel Rollinde, Nicholas P. Ross, David J. Schlegel, Donald P. Schneider, Anže Slosar, and David H. Weinberg. The one-dimensional Ly $\alpha$  forest power spectrum from BOSS. *A&A*, 559:A85, November 2013.
- [113] Fahad Nasir, James S. Bolton, and George D. Becker. Inferring the IGM thermal history during reionization with the Lyman  $\alpha$  forest power spectrum at redshift  $z \simeq 5$ . *MNRAS*, 463(3):2335–2347, December 2016.
- [114] Girish Kulkarni, Joseph F. Hennawi, Jose Oñorbe, Alberto Rorai, and Volker Springel. Characterizing the Pressure Smoothing Scale of the Intergalactic Medium. *ApJ*, 812(1):30, October 2015.
- [115] T. Theuns and S. Zaroubi. A wavelet analysis of the spectra of quasi-stellar objects. *MNRAS*, 317:989–995, October 2000.

- [116] T. Theuns, S. Zaroubi, T.-S. Kim, P. Tzanavaris, and R. F. Carswell. Temperature fluctuations in the intergalactic medium. *MNRAS*, 332:367–382, May 2002.
- [117] M. Zaldarriaga. Searching for Fluctuations in the Intergalactic Medium Temperature Using the Ly $\alpha$  Forest. *ApJ*, 564:153–161, January 2002.
- [118] A. Lidz, C.-A. Faucher-Giguère, A. Dall’Aglio, M. McQuinn, C. Fechner, M. Zaldarriaga, L. Hernquist, and S. Dutta. A Measurement of Small-scale Structure in the  $2.2 \leq z \leq 4.2$  Ly $\alpha$  Forest. *ApJ*, 718:199–230, July 2010.
- [119] George D. Becker, James S. Bolton, Martin G. Haehnelt, and Wallace L. W. Sargent. Detection of extended He II reionization in the temperature evolution of the intergalactic medium. *MNRAS*, 410(2):1096–1112, January 2011.
- [120] E. Boera, M. T. Murphy, G. D. Becker, and J. S. Bolton. The thermal history of the intergalactic medium down to redshift  $z = 1.5$ : a new curvature measurement. *MNRAS*, 441:1916–1933, July 2014.
- [121] Joop Schaye, Tom Theuns, Anthony Leonard, and George Efstathiou. Measuring the equation of state of the intergalactic medium. *MNRAS*, 310(1):57–70, November 1999.
- [122] Massimo Ricotti, Nickolay Y. Gnedin, and J. Michael Shull. The Evolution of the Effective Equation of State of the Intergalactic Medium. *ApJ*, 534(1):41–56, May 2000.
- [123] Patrick McDonald, Jordi Miralda-Escudé, Michael Rauch, Wallace L. W. Sargent, Tom A. Barlow, and Renyue Cen. A Measurement of the Temperature-Density Relation in the Intergalactic Medium Using a New Ly $\alpha$  Absorption-Line Fitting Method. *ApJ*, 562(1):52–75, November 2001.
- [124] James S. Bolton, George D. Becker, Martin G. Haehnelt, and Matteo Viel. A consistent determination of the temperature of the intergalactic medium at redshift  $z = 2.4$ . *MNRAS*, 438(3):2499–2507, March 2014.
- [125] Hector Hiss, Michael Walther, Joseph F. Hennawi, José Oñorbe, John M. O’Meara, Alberto Rorai, and Zarija Lukić. A New Measurement of the Temperature-density Relation of the IGM from Voigt Profile Fitting. *ApJ*, 865(1):42, September 2018.
- [126] Prakash Gaikwad, Raghunathan Srianand, Martin G. Haehnelt, and Tirthankar Roy Choudhury. A consistent and robust measurement of the thermal state of the IGM at  $2 \leq z \leq 4$  from a large sample of Ly $\alpha$  forest spectra: Evidence for late and rapid HeII reionization. *arXiv e-prints*, page arXiv:2009.00016, August 2020.
- [127] J. D. Emberson, Nicholas Frontiere, Salman Habib, Katrin Heitmann, Patricia Larsen, Hal Finkel, and Adrian Pope. The Borg Cube Simulation: Cosmological Hydrodynamics with CRK-SPH. *ApJ*, 877(2):85, June 2019.
- [128] S. Naoz and R. Barkana. Growth of linear perturbations before the era of the first galaxies. *MNRAS*, 362(3):1047–1053, September 2005.

- [129] Simeon Bird, Yu Feng, Christian Pedersen, and Andreu Font-Ribera. More accurate simulations with separate initial conditions for baryons and dark matter. *arXiv e-prints*, page arXiv:2002.00015, January 2020.
- [130] Raul E. Angulo, Oliver Hahn, and Tom Abel. How closely do baryons follow dark matter on large scales? *MNRAS*, 434(2):1756–1764, September 2013.
- [131] Cornelius Rampf, Cora Uhlemann, and Oliver Hahn. Cosmological perturbations for two cold fluids in  $\Lambda$ CDM. *arXiv e-prints*, page arXiv:2008.09123, August 2020.
- [132] Oliver Hahn, Cornelius Rampf, and Cora Uhlemann. Higher-order initial conditions for mixed baryon-CDM simulations. *arXiv e-prints*, page arXiv:2008.09124, August 2020.
- [133] Michaël Michaux, Oliver Hahn, Cornelius Rampf, and Raul E. Angulo. Accurate initial conditions for cosmological N-body simulations: Minimizing truncation and discreteness errors. *arXiv e-prints*, page arXiv:2008.09588, August 2020.
- [134] Vijay K. Narayanan, David N. Spergel, Romeel Davé, and Chung-Pei Ma. Constraints on the Mass of Warm Dark Matter Particles and the Shape of the Linear Power Spectrum from the Ly $\alpha$  Forest. *ApJ*, 543(2):L103–L106, November 2000.
- [135] Simeon Bird, Yacine Ali-Haïmoud, Yu Feng, and Jia Liu. An efficient and accurate hybrid method for simulating non-linear neutrino structure. *MNRAS*, 481(2):1486–1500, December 2018.
- [136] Naoki Yoshida, Naoshi Sugiyama, and Lars Hernquist. The evolution of baryon density fluctuations in multicomponent cosmological simulations. *MNRAS*, 344(2):481–491, September 2003.
- [137] Claude-André Faucher-Giguère, Adam Lidz, Matias Zaldarriaga, and Lars Hernquist. A New Calculation of the Ionizing Background Spectrum and the Effects of He II Reionization. *ApJ*, 703(2):1416–1443, October 2009.
- [138] Claude-André Faucher-Giguère. A cosmic UV/X-ray background model update. *MNRAS*, 493(2):1614–1632, April 2020.
- [139] Phoebe Upton Sanderbeck and Simeon Bird. Inhomogeneous He II reionization in hydrodynamic simulations. *MNRAS*, 496(4):4372–4382, August 2020.
- [140] Matteo Viel, Martin G. Haehnelt, and Volker Springel. Inferring the dark matter power spectrum from the Lyman  $\alpha$  forest in high-resolution QSO absorption spectra. *MNRAS*, 354(3):684–694, November 2004.
- [141] James S. Bolton and George D. Becker. Resolving the high redshift Ly $\alpha$  forest in smoothed particle hydrodynamics simulations. *MNRAS*, 398(1):L26–L30, September 2009.
- [142] Simeon Bird. FSFE: Fake Spectra Flux Extractor, October 2017.



- [143] Simeon Bird, Martin Haehnelt, Marcel Neeleman, Shy Genel, Mark Vogelsberger, and Lars Hernquist. Reproducing the kinematics of damped Lyman  $\alpha$  systems. *MNRAS*, 447(2):1834–1846, February 2015.
- [144] Hy Trac, Renyue Cen, and Abraham Loeb. Imprint of Inhomogeneous Hydrogen Reionization on the Temperature Distribution of the Intergalactic Medium. *ApJ*, 689(2):L81, December 2008.
- [145] Steven R. Furlanetto and S. Peng Oh. The Temperature-Density Relation of the Intergalactic Medium after Hydrogen Reionization. *ApJ*, 701(1):94–104, August 2009.
- [146] T. S. Kim, J. S. Bolton, M. Viel, M. G. Haehnelt, and R. F. Carswell. An improved measurement of the flux distribution of the Ly $\alpha$  forest in QSO absorption spectra: the effect of continuum fitting, metal contamination and noise properties. *MNRAS*, 382(4):1657–1674, December 2007.
- [147] Joop Schaye. Model-independent Insights into the Nature of the Ly $\alpha$  Forest and the Distribution of Matter in the Universe. *ApJ*, 559(2):507–515, October 2001.
- [148] Gwen C. Rudie, Charles C. Steidel, and Max Pettini. The Temperature-Density Relation in the Intergalactic Medium at Redshift  $z = 2.4$ . *ApJ*, 757(2):L30, October 2012.
- [149] Hector Hiss, Michael Walther, Jose Oñorbe, and Joseph F. Hennawi. A Novel Statistical Method for Measuring the Temperature-Density Relation in the IGM Using the  $b - N_{HI}$  Distribution of Absorbers in the Ly $\alpha$  Forest. *ApJ*, 876(1):71, May 2019.
- [150] Matteo Viel, George D. Becker, James S. Bolton, and Martin G. Haehnelt. Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman- $\alpha$  forest data. *Phys. Rev. D*, 88(4):043502, August 2013.
- [151] Vid Iršič, Matteo Viel, Martin G. Haehnelt, James S. Bolton, Stefano Cristiani, George D. Becker, Valentina D’Odorico, Guido Cupani, Tae-Sun Kim, Trystyn A. M. Berg, Sebastian López, Sara Ellison, Lise Christensen, Kelly D. Denney, and Gábor Worseck. New constraints on the free-streaming of warm dark matter from intermediate and small scale Lyman- $\alpha$  forest data. *Phys. Rev. D*, 96(2):023522, July 2017.
- [152] Vid Iršič, Matteo Viel, Trystyn A. M. Berg, Valentina D’Odorico, Martin G. Haehnelt, Stefano Cristiani, Guido Cupani, Tae-Sun Kim, Sebastian López, Sara Ellison, George D. Becker, Lise Christensen, Kelly D. Denney, Gábor Worseck, and James S. Bolton. The Lyman  $\alpha$  forest power spectrum from the XQ-100 Legacy Survey. *MNRAS*, 466(4):4332–4345, April 2017.
- [153] Michael Walther, Joseph F. Hennawi, Hector Hiss, Jose Oñorbe, Khee-Gan Lee, Alberto Rorai, and John O’Meara. A New Precision Measurement of the Small-scale Line-of-sight Power Spectrum of the Ly $\alpha$  Forest. *ApJ*, 852(1):22, January 2018.
- [154] Ryan M. O’Leary and Matthew McQuinn. The Formation of the First Cosmic Structures and the Physics of the  $z \sim 20$  Universe. *ApJ*, 760(1):4, November 2012.

- [155] Nathalie Palanque-Delabrouille, Christophe Yèche, Nils Schöneberg, Julien Lesgourgues, Michael Walther, Solène Chabanier, and Eric Armengaud. Hints, neutrino bounds, and WDM constraints from SDSS DR14 Lyman- $\alpha$  and Planck full-survey data. *JCAP*, 2020(4):038, April 2020.
- [156] George Becker, Anson D’Aloisio, Frederick B. Davies, Joseph F. Hennawi, and Robert A. Simcoe. Studying the Reionization Epoch with QSO Absorption Lines. *BAAS*, 51(3):440, May 2019.
- [157] Tom Theuns, Anthony Leonard, George Efstathiou, F. R. Pearce, and P. A. Thomas. P $\hat{3}$ M-SPH simulations of the Ly $\alpha$  forest. *MNRAS*, 301(2):478–502, December 1998.
- [158] Patrick McDonald, Jordi Miralda-Escudé, Michael Rauch, Wallace L. W. Sargent, Tom A. Barlow, Renyue Cen, and Jeremiah P. Ostriker. The Observed Probability Distribution Function, Power Spectrum, and Correlation Function of the Transmitted Flux in the Ly $\alpha$  Forest. *ApJ*, 543(1):1–23, November 2000.
- [159] Lam Hui, Scott Burles, Uroš Seljak, Robert E. Rutledge, Eugene Magnier, and David Tytler. On Estimating the QSO Transmission Power Spectrum. *ApJ*, 552(1):15–35, May 2001.
- [160] M. Viel, S. Matarrese, H. J. Mo, M. G. Haehnelt, and Tom Theuns. Probing the intergalactic medium with the Ly $\alpha$  forest along multiple lines of sight to distant QSOs. *MNRAS*, 329(4):848–862, February 2002.
- [161] Matteo Viel and Martin G. Haehnelt. Cosmological and astrophysical parameters from the Sloan Digital Sky Survey flux power spectrum and hydrodynamical simulations of the Lyman  $\alpha$  forest. *MNRAS*, 365(1):231–244, January 2006.
- [162] Patrick McDonald, Uroš Seljak, Scott Burles, David J. Schlegel, David H. Weinberg, Renyue Cen, David Shih, Joop Schaye, Donald P. Schneider, Neta A. Bahcall, John W. Briggs, J. Brinkmann, Robert J. Brunner, Masataka Fukugita, James E. Gunn, Željko Ivezić, Stephen Kent, Robert H. Lupton, and Daniel E. Vanden Berk. The Ly $\alpha$  Forest Power Spectrum from the Sloan Digital Sky Survey. *ApJS*, 163(1):80–109, March 2006.
- [163] Katrin Heitmann, David Higdon, Charles Nakhleh, and Salman Habib. Cosmic Calibration. *Astrophys. J. Let.*, 646(1):L1–L4, July 2006.
- [164] Salman Habib, Katrin Heitmann, David Higdon, Charles Nakhleh, and Brian Williams. Cosmic calibration: Constraints from the matter power spectrum and the cosmic microwave background. *Phys. Rev. D*, 76(8):083503, October 2007.
- [165] Katrin Heitmann, David Higdon, Martin White, Salman Habib, Brian J. Williams, Earl Lawrence, and Christian Wagner. The Coyote Universe. II. Cosmological Models and Precision Emulation of the Nonlinear Matter Power Spectrum. *ApJ*, 705(1):156–174, November 2009.

- [166] Ming-Feng Ho, Simeon Bird, and Christian R. Shelton. Multifidelity emulation for the matter power spectrum using Gaussian processes. *MNRAS*, 509(2):2551–2565, January 2022.
- [167] Arnaud Borde, Nathalie Palanque-Delabrouille, Graziano Rossi, Matteo Viel, James S. Bolton, Christophe Yèche, Jean-Marc LeGoff, and Jim Rich. New approach for precise computation of Lyman- $\alpha$  forest power spectrum with hydrodynamical simulations. *JCAP*, 2014(7):005, July 2014.
- [168] Rupert A. C. Croft, David H. Weinberg, Neal Katz, and Lars Hernquist. Recovery of the Power Spectrum of Mass Fluctuations from Observations of the Ly $\alpha$  Forest. *ApJ*, 495(1):44–62, March 1998.
- [169] Matteo Viel, Joop Schaye, and C. M. Booth. The impact of feedback from galaxy formation on the Lyman  $\alpha$  transmitted flux. *MNRAS*, 429(2):1734–1746, February 2013.
- [170] Solène Chabanier, Frédéric Bournaud, Yohan Dubois, Nathalie Palanque-Delabrouille, Christophe Yèche, Eric Armengaud, Sébastien Peirani, and Ricarda Beckmann. The impact of AGN feedback on the 1D power spectra from the Ly  $\alpha$  forest using the Horizon-AGN suite of simulations. *MNRAS*, 495(2):1825–1840, June 2020.
- [171] J. S. Bolton, M. Viel, T. S. Kim, M. G. Haehnelt, and R. F. Carswell. Possible evidence for an inverted temperature-density relation in the intergalactic medium from the flux distribution of the Ly $\alpha$  forest. *MNRAS*, 386(2):1131–1144, May 2008.
- [172] Xiaohan Wu, Matthew McQuinn, Rahul Kannan, Anson D’Aloisio, Simeon Bird, Federico Marinacci, Romeel Davé, and Lars Hernquist. Imprints of temperature fluctuations on the  $z \sim 5$  Lyman- $\alpha$  forest: a view from radiation-hydrodynamic simulations of reionization. *MNRAS*, 490(3):3177–3195, December 2019.
- [173] Prakash Gaikwad, Raghunathan Srianand, Martin G. Haehnelt, and Tirthankar Roy Choudhury. A consistent and robust measurement of the thermal state of the IGM at  $2 \leq z \leq 4$  from a large sample of Ly  $\alpha$  forest spectra: evidence for late and rapid He II reionization. *MNRAS*, 506(3):4389–4412, September 2021.
- [174] Bruno Villasenor, Brant Robertson, Piero Madau, and Evan Schneider. Inferring the Thermal History of the Intergalactic Medium from the Properties of the Hydrogen and Helium Lyman-alpha Forest. *arXiv e-prints*, page arXiv:2111.00019, October 2021.
- [175] Patrick McDonald, Uroš Seljak, Renyue Cen, David Shih, David H. Weinberg, Scott Burles, Donald P. Schneider, David J. Schlegel, Neta A. Bahcall, John W. Briggs, J. Brinkmann, Masataka Fukugita, Željko Ivezić, Stephen Kent, and Daniel E. Vanden Berk. The Linear Theory Power Spectrum from the Ly $\alpha$  Forest in the Sloan Digital Sky Survey. *ApJ*, 635(2):761–783, December 2005.
- [176] Matteo Viel, Martin G. Haehnelt, and Antony Lewis. The Lyman  $\alpha$  forest and WMAP year three. *MNRAS*, 370(1):L51–L55, July 2006.

- [177] Uroš Seljak, Alexey Makarov, Patrick McDonald, Scott F. Anderson, Neta A. Bahcall, J. Brinkmann, Scott Burles, Renyue Cen, Mamoru Doi, James E. Gunn, Željko Ivezić, Stephen Kent, Jon Loveday, Robert H. Lupton, Jeffrey A. Munn, Robert C. Nichol, Jeremiah P. Ostriker, David J. Schlegel, Donald P. Schneider, Max Tegmark, Daniel E. Berk, David H. Weinberg, and Donald G. York. Cosmological parameter analysis including SDSS Ly $\alpha$  forest and galaxy bias: Constraints on the primordial spectrum of fluctuations, neutrino mass, and dark energy. *Phys. Rev. D*, 71(10):103515, May 2005.
- [178] Uros Seljak, Anze Slosar, and Patrick McDonald. Cosmological parameters from combining the Lyman- $\alpha$  forest with CMB, galaxy clustering and SN constraints. *JCAP*, 2006(10):014, October 2006.
- [179] Mathias Garny, Thomas Konstandin, Laura Sagunski, and Matteo Viel. Neutrino mass bounds from confronting an effective model with BOSS Lyman- $\alpha$  data. *JCAP*, 2021(3):049, March 2021.
- [180] Antonella Garzilli, Andrii Magalich, Oleg Ruchayskiy, and Alexey Boyarsky. How to constrain warm dark matter with the Lyman- $\alpha$  forest. *MNRAS*, 502(2):2356–2363, April 2021.
- [181] Keir K. Rogers and Hiranya V. Peiris. Strong Bound on Canonical Ultralight Axion Dark Matter from the Lyman-Alpha Forest. *Phys. Rev. L*, 126(7):071302, February 2021.
- [182] Solène Chabanier, Nathalie Palanque-Delabrouille, Christophe Yèche, Jean-Marc Le Goff, Eric Armengaud, Julian Bautista, Michael Blomqvist, Nicolas Busca, Kyle Dawson, Thomas Etourneau, Andreu Font-Ribera, Youngbae Lee, Hélion du Mas des Bourboux, Matthew Pieri, James Rich, Graziano Rossi, Donald Schneider, and Anže Slosar. The one-dimensional power spectrum from the SDSS DR14 Ly $\alpha$  forests. *JCAP*, 2019(7):017, July 2019.
- [183] Naim Göksel Karaçaylı, Nikhil Padmanabhan, Andreu Font-Ribera, Vid Iršič, Michael Walther, David Brooks, Enrique Gaztañaga, Robert Kehoe, Michael Levi, Pierros Ntelis, Nathalie Palanque-Delabrouille, and Gregory Tarlé. Optimal 1D Ly  $\alpha$  forest power spectrum estimation - II. KODIAQ, SQUAD, and XQ-100. *MNRAS*, 509(2):2842–2855, January 2022.
- [184] Aaron Day, David Tytler, and Bharat Kambalur. Power spectrum of the flux in the Lyman-alpha forest from high-resolution spectra of 87 QSOs. *MNRAS*, 489(2):2536–2554, October 2019.
- [185] Monica Valluri, Solene Chabanier, Vid Irsic, Eric Armengaud, Michael Walther, Connie Rockosi, Miguel A. Sanchez-Conde, Leandro Beraldo e Silva, Andrew P. Cooper, Elise Darragh-Ford, Kyle Dawson, Alis J. Deason, Simone Ferraro, Jaime E. Forero-Romero, Antonella Garzilli, Ting Li, Zarija Lukic, Christopher J. Manser, Nathalie Palanque-Delabrouille, Corentin Ravoux, Ting Tan, Wenting Wang, Risa Weschler,

- Andreia Carrillo, Arjun Dey, Sergey E. Kogosov, Yao-Yuan Mao, Paulo Montero-Camacho, Ekta Patel, Graziano Rossi, L. Arturo Urena-Lopez, and Octavio Valenzuela. Snowmass2021 Cosmic Frontier White Paper: Prospects for obtaining Dark Matter Constraints with DESI. *arXiv e-prints*, page arXiv:2203.07491, March 2022.
- [186] Zarija Lukić, Casey W. Stark, Peter Nugent, Martin White, Avery A. Meiksin, and Ann Almgren. The Lyman  $\alpha$  forest in optically thin hydrodynamical simulations. *MNRAS*, 446(4):3697–3724, February 2015.
- [187] Katrin Heitmann, Earl Lawrence, Juliana Kwan, Salman Habib, and David Higdon. The Coyote Universe Extended: Precision Emulation of the Matter Power Spectrum. *ApJ*, 780(1):111, January 2014.
- [188] Earl Lawrence, Katrin Heitmann, Juliana Kwan, Amol Upadhye, Derek Bingham, Salman Habib, David Higdon, Adrian Pope, Hal Finkel, and Nicholas Frontiere. The Mira-Titan Universe. II. Matter Power Spectrum Emulation. *Astrophys. J. Let.*, 847(1):50, September 2017.
- [189] Benjamin Giblin, Matteo Cataneo, Ben Moews, and Catherine Heymans. On the road to per cent accuracy - II. Calibration of the non-linear matter power spectrum for arbitrary cosmologies. *MNRAS*, 490(4):4826–4840, December 2019.
- [190] Euclid Collaboration, M. Knabenhans, J. Stadel, D. Potter, J. Dakin, S. Hannestad, T. Tram, S. Marelli, A. Schneider, R. Teyssier, P. Fosalba, S. Andreon, N. Auricchio, C. Baccigalupi, A. Balaguera-Antolínez, M. Baldi, S. Bardelli, P. Battaglia, R. Bender, A. Biviano, C. Bodendorf, E. Bozzo, E. Branchini, M. Brescia, C. Burigana, R. Cabanac, S. Camera, V. Capobianco, A. Cappi, C. Carbone, J. Carretero, C. S. Carvalho, R. Casas, S. Casas, M. Castellano, G. Castignani, S. Cavuoti, R. Cledassou, C. Colodro-Conde, G. Congedo, C. J. Conselice, L. Conversi, Y. Copin, L. Corcione, J. Coupon, H. M. Courtois, A. Da Silva, S. de la Torre, D. Di Ferdinando, C. A. J. Duncan, X. Dupac, G. Fabbian, S. Farrens, P. G. Ferreira, F. Finelli, M. Frailis, E. Franceschi, S. Galeotta, B. Garilli, C. Giocoli, G. Gozaliasl, J. Graciá-Carpio, F. Grupp, L. Guzzo, W. Holmes, F. Hormuth, H. Israel, K. Jahnke, E. Keihänen, S. Kermiche, C. C. Kirkpatrick, B. Kubik, M. Kunz, H. Kurki-Suonio, S. Ligi, P. B. Lilje, I. Lloro, D. Maino, O. Marggraf, K. Markovic, N. Martinet, F. Marulli, R. Massey, N. Mauri, S. Maurogordato, E. Medinaceli, M. Meneghetti, B. Metcalf, G. Meylan, M. Moresco, B. Morin, L. Moscardini, E. Munari, C. Neissner, S. M. Niemi, C. Padilla, S. Paltani, F. Pasian, L. Patrizii, V. Pettorino, S. Pires, G. Polenta, M. Poncet, F. Raison, A. Renzi, J. Rhodes, G. Riccio, E. Romelli, M. Roncarelli, R. Saglia, A. G. Sánchez, D. Sapone, P. Schneider, V. Scottez, A. Secroun, S. Serrano, C. Sirignano, G. Sirri, L. Stanco, F. Sureau, P. Tallada Crespí, A. N. Taylor, M. Tenti, I. Tereno, R. Toledo-Moreo, F. Torradeflot, L. Valenziano, J. Valiviita, T. Vassallo, M. Viel, Y. Wang, N. Welikala, L. Whittaker, A. Zacchei, and E. Zucca. Euclid preparation: IX. EuclidEmulator2 - power spectrum emulation with massive neutrinos and self-consistent dark energy perturbations. *MNRAS*, 505(2):2840–2869, August 2021.

- [191] Giovanni Aricò, Raul E. Angulo, Sergio Contreras, Lurdes Ondaro-Mallea, Marcos Pellejero-Ibañez, and Matteo Zennaro. The BACCO simulation project: a baryonification emulator with neural networks. *MNRAS*, 506(3):4070–4082, September 2021.
- [192] Sambit K. Giri and Aurel Schneider. Emulation of baryonic effects on the matter power spectrum and constraints from galaxy cluster data. *JCAP*, 2021(12):046, December 2021.
- [193] J. Harnois-Déraps, B. Giblin, and B. Joachimi. Cosmic shear covariance matrix in  $\Lambda$ CDM: Cosmology matters. *Astron. Astrophys.*, 631:A160, November 2019.
- [194] Christopher T. Davies, Marius Cautun, Benjamin Giblin, Baojiu Li, Joachim Harnois-Déraps, and Yan-Chuan Cai. Constraining cosmology with weak lensing voids. *MNRAS*, 507(2):2267–2282, October 2021.
- [195] Thomas McClintock, Eduardo Rozo, Matthew R. Becker, Joseph DeRose, Yao-Yuan Mao, Sean McLaughlin, Jeremy L. Tinker, Risa H. Wechsler, and Zhongxu Zhai. The Aemulus Project. II. Emulating the Halo Mass Function. *ApJ*, 872(1):53, February 2019.
- [196] Takahiro Nishimichi, Masahiro Takada, Ryuichi Takahashi, Ken Osato, Masato Shirasaki, Taira Oogi, Hironao Miyatake, Masamune Oguri, Ryoma Murata, Yosuke Kobayashi, and Naoki Yoshida. Dark Quest. I. Fast and Accurate Emulation of Halo Clustering Statistics and Its Application to Galaxy Clustering. *ApJ*, 884(1):29, October 2019.
- [197] Sebastian Bocquet, Katrin Heitmann, Salman Habib, Earl Lawrence, Thomas Uram, Nicholas Frontiere, Adrian Pope, and Hal Finkel. The Mira-Titan Universe. III. Emulation of the Halo Mass Function. *ApJ*, 901(1):5, September 2020.
- [198] Nicholas S. Kern, Adrian Liu, Aaron R. Parsons, Andrei Mesinger, and Bradley Greig. Emulating Simulations of Cosmic Dawn for 21 cm Power Spectrum Constraints on Cosmology, Reionization, and X-Ray Heating. *ApJ*, 848(1):23, October 2017.
- [199] Aviad Cohen, Anastasia Fialkov, Rennan Barkana, and Raul A. Monsalve. Emulating the global 21-cm signal from Cosmic Dawn and Reionization. *MNRAS*, 495(4):4845–4859, July 2020.
- [200] H. T. J. Bevins, W. J. Handley, A. Fialkov, E. de Lera Acedo, and K. Javid. GLOB-ALEMU: a novel and robust approach for emulating the sky-averaged 21-cm signal from the cosmic dawn and epoch of reionization. *MNRAS*, 508(2):2923–2936, December 2021.
- [201] Christian H. Bye, Stephen K. N. Portillo, and Anastasia Fialkov. 21cmVAE: A Very Accurate Emulator of the 21 cm Global Signal. *ApJ*, 930(1):79, May 2022.
- [202] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.

- [203] Keir K. Rogers, Hiranya V. Peiris, Andrew Pontzen, Simeon Bird, Licia Verde, and Andreu Font-Ribera. Bayesian emulator optimisation for cosmology: application to the Lyman-alpha forest. *JCAP*, 2019(2):031, February 2019.
- [204] Christian Pedersen, Andreu Font-Ribera, Keir K. Rogers, Patrick McDonald, Hiranya V. Peiris, Andrew Pontzen, and Anže Slosar. An emulator for the Lyman- $\alpha$  forest in beyond- $\Lambda$ CDM cosmologies. *JCAP*, 2021(5):033, May 2021.
- [205] Michael Walther, Eric Armengaud, Corentin Ravoux, Nathalie Palanque-Delabrouille, Christophe Yèche, and Zarija Lukić. Simulating intergalactic gas for DESI-like small scale Lyman $\alpha$  forest observations. *JCAP*, 2021(4):059, April 2021.
- [206] Keir K. Rogers and Hiranya V. Peiris. General framework for cosmological dark matter bounds using N -body simulations. *Phys. Rev. D*, 103(4):043526, February 2021.
- [207] MC Kennedy and A O’Hagan. Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1):1–13, 03 2000.
- [208] Simeon Bird, Yueying Ni, Tiziana Di Matteo, Rupert Croft, Yu Feng, and Nianyi Chen. The ASTRID simulation: galaxy formation and reionization. *MNRAS*, 512(3):3703–3716, May 2022.
- [209] Yueying Ni, Tiziana Di Matteo, Simeon Bird, Rupert Croft, Yu Feng, Nianyi Chen, Michael Tremmel, Colin DeGraf, and Yin Li. The ASTRID simulation: the evolution of supermassive black holes. *MNRAS*, 513(1):670–692, June 2022.
- [210] Simeon Bird, Yu Feng, Christian Pedersen, and Andreu Font-Ribera. More accurate simulations with separate initial conditions for baryons and dark matter. *JCAP*, 2020(6):002, June 2020.
- [211] Ming-Feng Ho, Simeon Bird, and Roman Garnett. Damped Lyman- $\alpha$  absorbers from Sloan digital sky survey DR16Q with Gaussian processes. *MNRAS*, 507(1):704–719, October 2021.
- [212] Simeon Bird, Mark Vogelsberger, Debora Sijacki, Matias Zaldarriaga, Volker Springel, and Lars Hernquist. Moving-mesh cosmology: properties of neutral hydrogen in absorption. *MNRAS*, 429(4):3341–3352, March 2013.
- [213] Volker Springel and Lars Hernquist. Cosmological smoothed particle hydrodynamics simulations: a hybrid multiphase model for star formation. *MNRAS*, 339(2):289–311, February 2003.
- [214] Yu Feng, Tiziana Di-Matteo, Rupert A. Croft, Simeon Bird, Nicholas Battaglia, and Stephen Wilkins. The BlueTides simulation: first galaxies and reionization. *MNRAS*, 455(3):2778–2791, January 2016.
- [215] Takashi Okamoto, Carlos S. Frenk, Adrian Jenkins, and Tom Theuns. The properties of satellite galaxies in simulations of galaxy formation. *MNRAS*, 406(1):208–222, July 2010.

- [216] Thomas J. Santner, Brian J. Williams, and William I. Notz. *The Design and Analysis of Computer Experiments*. Springer series in statistics. Springer, 2003.
- [217] P. Perdikaris, M. Raissi, A. Damianou, N. D. Lawrence, and G. E. Karniadakis. Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling. *Proceedings of the Royal Society of London Series A*, 473(2198):20160751, February 2017.
- [218] Andreas Damianou and Neil D. Lawrence. Deep Gaussian processes. In Carlos M. Carvalho and Pradeep Ravikumar, editors, *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, volume 31 of *Proceedings of Machine Learning Research*, pages 207–215, Scottsdale, Arizona, USA, 29 Apr–01 May 2013. PMLR.
- [219] Bruno Villasenor, Brant Robertson, Piero Madau, and Evan Schneider. Inferring the Thermal History of the Intergalactic Medium from the Properties of the Hydrogen and Helium Ly $\alpha$  Forest. *ApJ*, 933(1):59, July 2022.
- [220] B. Abareshi, J. Aguilar, S. Ahlen, Shadab Alam, David M. Alexander, R. Alfarsy, L. Allen, C. Allende Prieto, O. Alves, J. Ameel, E. Armengaud, J. Asorey, Alejandro Aviles, S. Bailey, A. Balaguera-Antolínez, O. Ballester, C. Baltay, A. Bault, S. F. Beltran, B. Benavides, S. BenZvi, A. Berti, R. Besuner, Florian Beutler, D. Bianchi, C. Blake, P. Blanc, R. Blum, A. Bolton, S. Bose, D. Bramall, S. Brieden, A. Brodzeller, D. Brooks, C. Brownell, E. Buckley-Geer, R. N. Cahn, Z. Cai, R. Canning, R. Cappasso, A. Carnero Rosell, P. Carton, R. Casas, F. J. Castander, J. L. Cervantes-Cota, S. Chabanier, E. Chaussidon, C. Chuang, C. Circosta, S. Cole, A. P. Cooper, L. da Costa, M. C. Cousinou, A. Cuceu, T. M. Davis, K. Dawson, R. de la Cruz-Noriega, A. de la Macorra, A. de Mattia, J. Della Costa, P. Demmer, M. Derwent, A. Dey, B. Dey, G. Dhungana, Z. Ding, C. Dobson, P. Doel, J. Donald-McCann, J. Donaldson, K. Douglass, Y. Duan, P. Dunlop, J. Edelman, S. Eftekharzadeh, D. J. Eisenstein, M. Enriquez-Vargas, S. Escoffier, M. Evatt, P. Fagrellius, X. Fan, K. Fanning, V. A. Fawcett, S. Ferraro, J. Ereza, B. Flaugher, A. Font-Ribera, J. E. Forero-Romero, C. S. Frenk, S. Fromenteau, B. T. Gänsicke, C. Garcia-Quintero, L. Garrison, E. Gaztañaga, F. Gerardi, H. Gil-Marín, S. Gontcho A. Gontcho, Alma X. Gonzalez-Morales, G. Gonzalez-de-Rivera, V. Gonzalez-Perez, C. Gordon, O. Graur, D. Green, C. Grove, D. Gruen, G. Gutierrez, J. Guy, C. Hahn, S. Harris, D. Herrera, Hiram K. Herrera-Alcantar, K. Honscheid, C. Howlett, D. Huterer, V. Iršič, M. Ishak, P. Jelinsky, L. Jiang, J. Jimenez, Y. P. Jing, R. Joyce, E. Jullo, S. Juneau, N. G. Karaçaylı, M. Karamanis, A. Karcher, T. Karim, R. Kehoe, S. Kent, D. Kirkby, T. Kisner, F. Kitaura, S. E. Koposov, A. Kovács, A. Kremin, Alex Krolewski, B. L’Huillier, O. Lahav, A. Lambert, C. Lamman, Ting-Wen Lan, M. Landriau, S. Lane, D. Lang, J. U. Lange, J. Lasker, L. Le Guillou, A. Leauthaud, A. Le Van Suu, Michael E. Levi, T. S. Li, C. Magneville, M. Manera, Christopher J. Manser, B. Marshall, Paul Martini, W. McCollam, P. McDonald, Aaron M. Meisner, J. Mena-Fernández, J. Meneses-Rizo, M. Mezcua, T. Miller, R. Miquel, P. Montero-Camacho, J. Moon, J. Moustakas, E. Mueller, Andrea Muñoz-Gutiérrez, Adam D. Myers, S. Nadathur, J. Najita,



- L. Napolitano, E. Neilsen, Jeffrey A. Newman, J. D. Nie, Y. Ning, G. Niz, P. Norberg, Hernán E. Noriega, T. O’Brien, A. Obuljen, N. Palanque-Delabrouille, A. Palmese, P. Zhiwei, D. Pappalardo, X. Peng, W. J. Percival, S. Perruchot, R. Pogge, C. Poppett, A. Porredon, F. Prada, J. Prochaska, R. Pucha, A. Pérez-Fernández, I. Pérez-Ràfols, D. Rabinowitz, A. Raichoor, S. Ramirez-Solano, César Ramírez-Pérez, C. Ravoux, K. Reil, M. Rezaie, A. Rocher, C. Rockosi, N. A. Roe, A. Roodman, A. J. Ross, G. Rossi, R. Ruggeri, V. Ruhlmann-Kleider, C. G. Sabiu, S. Safonova, K. Said, A. Saintonge, Javier Salas Catonga, L. Samushia, E. Sanchez, C. Saulder, E. Schaan, E. Schlafly, D. Schlegel, J. Schmoll, D. Scholte, M. Schubnell, A. Secroun, H. Seo, S. Serrano, Ray M. Sharples, Michael J. Sholl, Joseph Harry Silber, D. R. Silva, M. Sirk, M. Siudek, A. Smith, D. Sprayberry, R. Staten, B. Stupak, T. Tan, Gregory Tarlé, Suk Sien Tie, R. Tojeiro, L. A. Ureña-López, F. Valdes, O. Valenzuela, M. Valluri, M. Vargas-Magaña, L. Verde, M. Walther, B. Wang, M. S. Wang, B. A. Weaver, C. Weaverdyck, R. Wechsler, Michael J. Wilson, J. Yang, Y. Yu, S. Yuan, Christophe Yèche, H. Zhang, K. Zhang, Cheng Zhao, Rongpu Zhou, Zhimin Zhou, H. Zou, J. Zou, S. Zou, and Y. Zu. Overview of the Instrumentation for the Dark Energy Spectroscopic Instrument. *AJ*, 164(5):207, November 2022.
- [221] Simeon Bird and M. A. Fernandez. Simulation Suite Paper Placeholder. *JCAP*, xxx(x):xxxx–xxxx, March 2023.
- [222] Bruno Villasenor, Brant Robertson, Piero Madau, and Evan Schneider. New Constraints on Warm Dark Matter from the Lyman- $\alpha$  Forest Power Spectrum. *arXiv e-prints*, page arXiv:2209.14220, September 2022.
- [223] Christophe Yèche, Nathalie Palanque-Delabrouille, Julien Baur, and Hélion du Mas des Bourboux. Constraints on neutrino masses from Lyman-alpha forest power spectrum with BOSS and XQ-100. *JCAP*, 2017(6):047, June 2017.
- [224] Anže Slosar, Andreu Font-Ribera, Matthew M. Pieri, James Rich, Jean-Marc Le Goff, Éric Aubourg, Jon Brinkmann, Nicolas Busca, Bill Carithers, Romain Charlassier, Marina Cortês, Rupert Croft, Kyle S. Dawson, Daniel Eisenstein, Jean-Christophe Hamilton, Shirley Ho, Khee-Gan Lee, Robert Lupton, Patrick McDonald, Bumbarija Medolin, Demitri Muna, Jordi Miralda-Escudé, Adam D. Myers, Robert C. Nichol, Nathalie Palanque-Delabrouille, Isabelle Pâris, Patrick Petitjean, Yodovina Piškur, Emmanuel Rollinde, Nicholas P. Ross, David J. Schlegel, Donald P. Schneider, Erin Sheldon, Benjamin A. Weaver, David H. Weinberg, Christophe Yèche, and Donald G. York. The Lyman- $\alpha$  forest in three dimensions: measurements of large scale flux correlations from BOSS 1st-year data. *JCAP*, 2011(9):001, September 2011.
- [225] Anže Slosar, Vid Iršič, David Kirkby, Stephen Bailey, Nicolás G. Busca, Timothée Delubac, James Rich, Éric Aubourg, Julian E. Bautista, Vaishali Bhardwaj, Michael Blomqvist, Adam S. Bolton, Jo Bovy, Joel Brownstein, Bill Carithers, Rupert A. C. Croft, Kyle S. Dawson, Andreu Font-Ribera, J. M. Le Goff, Shirley Ho, Klaus Honscheid, Khee-Gan Lee, Daniel Margala, Patrick McDonald, Bumbarija Medolin, Jordi Miralda-Escudé, Adam D. Myers, Robert C. Nichol, Pasquier Noterdaeme, Nathalie

- Palanque-Delabrouille, Isabelle Pâris, Patrick Petitjean, Matthew M. Pieri, Yodovina Piškur, Natalie A. Roe, Nicholas P. Ross, Graziano Rossi, David J. Schlegel, Donald P. Schneider, Nao Suzuki, Erin S. Sheldon, Uroš Seljak, Matteo Viel, David H. Weinberg, and Christophe Yèche. Measurement of baryon acoustic oscillations in the Lyman- $\alpha$  forest fluctuations in BOSS data release 9. *JCAP*, 2013(4):026, April 2013.
- [226] Hélión du Mas des Bourboux, James Rich, Andreu Font-Ribera, Victoria de Sainte Agathe, James Farr, Thomas Etourneau, Jean-Marc Le Goff, Andrei Cuceu, Christophe Balland, Julian E. Bautista, Michael Blomqvist, Jonathan Brinkmann, Joel R. Brownstein, Solène Chabanier, Edmond Chaussidon, Kyle Dawson, Alma X. González-Morales, Julien Guy, Brad W. Lyke, Axel de la Macorra, Eva-Maria Mueller, Adam D. Myers, Christian Nitschelm, Andrea Muñoz Gutiérrez, Nathalie Palanque-Delabrouille, James Parker, Will J. Percival, Ignasi Pérez-Ràfols, Patrick Petitjean, Matthew M. Pieri, Corentin Ravoux, Graziano Rossi, Donald P. Schneider, Hee-Jong Seo, Anže Slosar, Julianna Stermer, M. Vivek, Christophe Yèche, and Samantha Youles. The Completed SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations with Ly $\alpha$  Forests. *ApJ*, 901(2):153, October 2020.
- [227] Andrei Cuceu, Andreu Font-Ribera, Seshadri Nadathur, Benjamin Joachimi, and Paul Martini. New constraints on the expansion rate at redshift 2.3 from the Lyman- $\alpha$  forest. *arXiv e-prints*, page arXiv:2209.13942, September 2022.
- [228] Jesús Torrado and Antony Lewis. Cobaya: Bayesian analysis in cosmology. *Astrophysics Source Code Library*, record ascl:1910.019, October 2019.
- [229] Jesús Torrado and Antony Lewis. Cobaya: code for Bayesian analysis of hierarchical physical models. *JCAP*, 2021(5):057, May 2021.
- [230] GPpy. GPpy: A gaussian process framework in python. <http://github.com/SheffieldML/GPy>, since 2012.
- [231] Andrei Paleyes, Mark Pullin, Maren Mahsereci, Cliff McCollum, Neil D. Lawrence, and Javier Gonzalez. Emulation of physical processes with Emukit. *arXiv e-prints*, page arXiv:2110.13293, October 2021.
- [232] Elcio Abdalla, Guillermo Franco Abellán, Amin Aboubrahim, Adriano Agnello, Özgür Akarsu, Yashar Akrami, George Alestas, Daniel Aloni, Luca Amendola, Luis A. Anchordoqui, Richard I. Anderson, Nikki Arendse, Marika Asgari, Mario Ballardini, Vernon Barger, Spyros Basilakos, Ronaldo C. Batista, Elia S. Battistelli, Richard Battye, Micol Benetti, David Benisty, Asher Berlin, Paolo de Bernardis, Emanuele Berti, Bohdan Bilenko, Simon Birrer, John P. Blakeslee, Kimberly K. Boddy, Clecio R. Bom, Alexander Bonilla, Nicola Borghi, François R. Bouchet, Matteo Braglia, Thomas Buchert, Elizabeth Buckley-Geer, Erminia Calabrese, Robert R. Caldwell, David Camarena, Salvatore Capozziello, Stefano Casertano, Geoff C. F. Chen, Jens Chluba, Angela Chen, Hsin-Yu Chen, Anton Chudaykin, Michele Cicoli, Craig J. Copi, Fred Courbin, Francis-Yan Cyr-Racine, Božena Czerny, Maria Dainotti, Guido

D’Amico, Anne-Christine Davis, Javier de Cruz Pérez, Jaume de Haro, Jacques Delabrouille, Peter B. Denton, Suhail Dhawan, Keith R. Dienes, Eleonora Di Valentino, Pu Du, Dominique Eckert, Celia Escamilla-Rivera, Agnès Ferté, Fabio Finelli, Pablo Fosalba, Wendy L. Freedman, Noemi Frusciante, Enrique Gaztañaga, William Giarè, Elena Giusarma, Adrià Gómez-Valent, Will Handley, Ian Harrison, Luke Hart, Dhiraj Kumar Hazra, Alan Heavens, Asta Heinesen, Hendrik Hildebrandt, J. Colin Hill, Natalie B. Hogg, Daniel E. Holz, Deanna C. Hooper, Nikoo Hosseininejad, Dragan Huterer, Mustapha Ishak, Mikhail M. Ivanov, Andrew H. Jaffe, In Sung Jang, Karsten Jedamzik, Raul Jimenez, Melissa Joseph, Shahab Joudaki, Marc Kamionkowski, Tanvi Karwal, Lavrentios Kazantzidis, Ryan E. Keeley, Michael Klasen, Eiichiro Komatsu, Léon V. E. Koopmans, Suresh Kumar, Luca Lamagna, Ruth Lazkoz, Chung-Chi Lee, Julien Lesgourgues, Jackson Levi Said, Tiffany R. Lewis, Benjamin L’Huillier, Matteo Lucca, Roy Maartens, Lucas M. Macri, Danny Marfatia, Valerio Marra, Carlos J. A. P. Martins, Silvia Masi, Sabino Matarrese, Arindam Mazumdar, Alessandro Melchiorri, Olga Mena, Laura Mersini-Houghton, James Mertens, Dinko Milaković, Yuto Minami, Vivian Miranda, Cristian Moreno-Pulido, Michele Moresco, David F. Mota, Emil Mottola, Simone Mozzon, Jessica Muir, Ankan Mukherjee, Suvodip Mukherjee, Pavel Naselsky, Pran Nath, Savvas Nesseris, Florian Niedermann, Alessio Notari, Rafael C. Nunes, Eoin Ó Colgáin, Kayla A. Owens, Emre Özülker, Francesco Pace, Andronikos Paliathanasis, Antonella Palmese, Supriya Pan, Daniela Paoletti, Santiago E. Perez Bergliaffa, Leandros Perivolaropoulos, Dominic W. Pesce, Valeria Pettorino, Oliver H. E. Philcox, Levon Pogolian, Vivian Poulin, Gaspard Poulot, Marco Raveri, Mark J. Reid, Fabrizio Renzi, Adam G. Riess, Vivian I. Sabla, Paolo Salucci, Vincenzo Salzano, Emmanuel N. Saridakis, Bangalore S. Sathyaprakash, Martin Schmaltz, Nils Schöneberg, Dan Scolnic, Anjan A. Sen, Neelima Sehgal, Arman Shafieloo, M. M. Sheikh-Jabbari, Joseph Silk, Alessandra Silvestri, Foteini Skara, Martin S. Sloth, Marcelle Soares-Santos, Joan Solà Peracaula, Yu-Yang Songsheng, Jorge F. Soriano, Denitsa Staicova, Glenn D. Starkman, István Szapudi, Elsa M. Teixeira, Brooks Thomas, Tommaso Treu, Emery Trott, Carsten van de Bruck, J. Alberto Vazquez, Licia Verde, Luca Visinelli, Deng Wang, Jian-Min Wang, Shao-Jiang Wang, Richard Watkins, Scott Watson, John K. Webb, Neal Weiner, Amanda Weltman, Samuel J. Witte, Radosław Wojtak, Anil Kumar Yadav, Weiqiang Yang, Gong-Bo Zhao, and Miguel Zumalacárregui. Cosmology intertwined: A review of the particle physics, astrophysics, and cosmology associated with the cosmological tensions and anomalies. *Journal of High Energy Astrophysics*, 34:49–211, June 2022.

- [233] Keir K. Rogers, Simeon Bird, Hiranya V. Peiris, Andrew Pontzen, Andreu Font-Ribera, and Boris Leistedt. Simulating the effect of high column density absorbers on the one-dimensional Lyman  $\alpha$  forest flux power spectrum. *MNRAS*, 474(3):3032–3042, March 2018.
- [234] Kyle S. Dawson, David J. Schlegel, Christopher P. Ahn, Scott F. Anderson, Éric Aubourg, Stephen Bailey, Robert H. Barkhouser, Julian E. Bautista, Alessandra Beifiori, Andreas A. Berlind, Vaishali Bhardwaj, Dmitry Bizyaev, Cullen H. Blake,

Michael R. Blanton, Michael Blomqvist, Adam S. Bolton, Arnaud Borde, Jo Bovy, W. N. Brandt, Howard Brewington, Jon Brinkmann, Peter J. Brown, Joel R. Brownstein, Kevin Bundy, N. G. Busca, William Carithers, Aurelio R. Carnero, Michael A. Carr, Yanmei Chen, Johan Comparat, Natalia Connolly, Frances Cope, Rupert A. C. Croft, Antonio J. Cuesta, Luiz N. da Costa, James R. A. Davenport, Timothée Delubac, Roland de Putter, Saurav Dhital, Anne Ealet, Garrett L. Ebelke, Daniel J. Eisenstein, S. Escoffier, Xiaohui Fan, N. Filiz Ak, Hayley Finley, Andreu Font-Ribera, R. Génova-Santos, James E. Gunn, Hong Guo, Daryl Haggard, Patrick B. Hall, Jean-Christophe Hamilton, Ben Harris, David W. Harris, Shirley Ho, David W. Hogg, Diana Holder, Klaus Honscheid, Joe Huehnerhoff, Beatrice Jordan, Wendell P. Jordan, Guinevere Kauffmann, Eyal A. Kazin, David Kirkby, Mark A. Klaene, Jean-Paul Kneib, Jean-Marc Le Goff, Khee-Gan Lee, Daniel C. Long, Craig P. Loomis, Britt Lundgren, Robert H. Lupton, Marcio A. G. Maia, Martin Makler, Elena Malanushenko, Viktor Malanushenko, Rachel Mandelbaum, Marc Manera, Claudia Maraston, Daniel Margala, Karen L. Masters, Cameron K. McBride, Patrick McDonald, Ian D. McGreer, Richard G. McMahon, Olga Mena, Jordi Miralda-Escudé, Antonio D. Montero-Dorta, Francesco Montesano, Demitri Muna, Adam D. Myers, Tracy Naugle, Robert C. Nichol, Pasquier Noterdaeme, Sebastián E. Nuza, Matthew D. Olmstead, Audrey Oravetz, Daniel J. Oravetz, Russell Owen, Nikhil Padmanabhan, Nathalie Palanque-Delabrouille, Kaike Pan, John K. Parejko, Isabelle Pâris, Will J. Percival, Ismael Pérez-Fournon, Ignasi Pérez-Ràfols, Patrick Petitjean, Robert Pfaffenberger, Janine Pforr, Matthew M. Pieri, Francisco Prada, Adrian M. Price-Whelan, M. Jordan Raddick, Rafael Rebolo, James Rich, Gordon T. Richards, Constance M. Rockosi, Natalie A. Roe, Ashley J. Ross, Nicholas P. Ross, Graziano Rossi, J. A. Rubiño-Martin, Lado Samushia, Ariel G. Sánchez, Conor Sayres, Sarah J. Schmidt, Donald P. Schneider, C. G. Scóccola, Hee-Jong Seo, Alaina Shelden, Erin Sheldon, Yue Shen, Yiping Shu, Anže Slosar, Stephen A. Smee, Stephanie A. Snedden, Fritz Stauffer, Oliver Steele, Michael A. Strauss, Alina Streblyanska, Nao Suzuki, Molly E. C. Swanson, Tomer Tal, Masayuki Tanaka, Daniel Thomas, Jeremy L. Tinker, Rita Tojeiro, Christy A. Tremonti, M. Vargas Magaña, Licia Verde, Matteo Viel, David A. Wake, Mike Watson, Benjamin A. Weaver, David H. Weinberg, Benjamin J. Weiner, Andrew A. West, Martin White, W. M. Wood-Vasey, Christophe Yèche, Idit Zehavi, Gong-Bo Zhao, and Zheng Zheng. The Baryon Oscillation Spectroscopic Survey of SDSS-III. *AJ*, 145(1):10, January 2013.

- [235] Kyle S. Dawson, Jean-Paul Kneib, Will J. Percival, Shadab Alam, Franco D. Albareti, Scott F. Anderson, Eric Armengaud, Éric Aubourg, Stephen Bailey, Julian E. Bautista, Andreas A. Berlind, Matthew A. Bershad, Florian Beutler, Dmitry Bizyaev, Michael R. Blanton, Michael Blomqvist, Adam S. Bolton, Jo Bovy, W. N. Brandt, Jon Brinkmann, Joel R. Brownstein, Etienne Burtin, N. G. Busca, Zheng Cai, Chia-Hsun Chuang, Nicolas Clerc, Johan Comparat, Frances Cope, Rupert A. C. Croft, Irene Cruz-Gonzalez, Luiz N. da Costa, Marie-Claude Cousinou, Jeremy Darling, Axel de la Macorra, Sylvain de la Torre, Timothée Delubac, Hélión du Mas des Bourboux, Tom Dwelly, Anne Ealet, Daniel J. Eisenstein, Michael Eracleous, S. Escoffier, Xiaohui Fan, Alexis Finoguenov, Andreu Font-Ribera, Peter Frinchaboy,

Patrick Gaulme, Antonis Georgakakis, Paul Green, Hong Guo, Julien Guy, Shirley Ho, Diana Holder, Joe Huehnerhoff, Timothy Hutchinson, Yipeng Jing, Eric Jullo, Vikrant Kamble, Karen Kinemuchi, David Kirkby, Francisco-Shu Kitaura, Mark A. Klaene, Russ R. Laher, Dustin Lang, Pierre Laurent, Jean-Marc Le Goff, Cheng Li, Yu Liang, Marcos Lima, Qiufan Lin, Weipeng Lin, Yen-Ting Lin, Daniel C. Long, Britt Lundgren, Nicholas MacDonald, Marcio Antonio Geimba Maia, Elena Malanushenko, Viktor Malanushenko, Vivek Mariappan, Cameron K. McBride, Ian D. McGreer, Brice Ménard, Andrea Merloni, Andres Meza, Antonio D. Montero-Dorta, Demitri Muna, Adam D. Myers, Kirpal Nandra, Tracy Naugle, Jeffrey A. Newman, Pasquier Noterdaeme, Peter Nugent, Ricardo Ogando, Matthew D. Olmstead, Audrey Oravetz, Daniel J. Oravetz, Nikhil Padmanabhan, Nathalie Palanque-Delabrouille, Kaike Pan, John K. Parejko, Isabelle Pâris, John A. Peacock, Patrick Petitjean, Matthew M. Pieri, Alice Pisani, Francisco Prada, Abhishek Prakash, Anand Raichoor, Beth Reid, James Rich, Jethro Ridl, Sergio Rodriguez-Torres, Aurelio Carnero Rosell, Ashley J. Ross, Graziano Rossi, John Ruan, Mara Salvato, Conor Sayres, Donald P. Schneider, David J. Schlegel, Uros Seljak, Hee-Jong Seo, Branimir Sesar, Sarah Shandera, Yiping Shu, Anže Slosar, Flavia Sobreira, Alina Streblyanska, Nao Suzuki, Donna Taylor, Charling Tao, Jeremy L. Tinker, Rita Tojeiro, Mariana Vargas-Magaña, Yuting Wang, Benjamin A. Weaver, David H. Weinberg, Martin White, W. M. Wood-Vasey, Christophe Yeche, Zhongxu Zhai, Cheng Zhao, Gong-bo Zhao, Zheng Zheng, Guangtun Ben Zhu, and Hu Zou. The SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Overview and Early Data. *AJ*, 151(2):44, February 2016.

- [236] J. M. O’Meara, N. Lehner, J. C. Howk, J. X. Prochaska, A. J. Fox, M. S. Peeples, J. Tumlinson, and B. W. O’Shea. The Second Data Release of the KODIAQ Survey. *AJ*, 154(3):114, September 2017.
- [237] Antony Lewis. Efficient sampling of fast and slow cosmological parameters. *Phys. Rev. D*, 87(10):103529, May 2013.
- [238] Antony Lewis and Sarah Bridle. Cosmological parameters from CMB and other data: A Monte Carlo approach. *Phys. Rev. D*, 66(10):103511, November 2002.
- [239] Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak, R. Battye, K. Benabed, J. P. Bernard, M. Bersanelli, P. Bielewicz, J. J. Bock, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J. F. Cardoso, J. Carron, A. Challinor, H. C. Chiang, J. Chluba, L. P. L. Colombo, C. Combet, D. Contreras, B. P. Crill, F. Cuttaia, P. de Bernardis, G. de Zotti, J. Delabrouille, J. M. Delouis, E. Di Valentino, J. M. Diego, O. Doré, M. Douspis, A. Ducout, X. Dupac, S. Dusini, G. Efstathiou, F. Elsner, T. A. Enßlin, H. K. Eriksen, Y. Fantaye, M. Farhang, J. Fergusson, R. Fernandez-Cobos, F. Finelli, F. Forastieri, M. Frailis, A. A. Fraisse, E. Franceschi, A. Frolov, S. Galeotta, S. Galli, K. Ganga, R. T. Génova-Santos, M. Gerbino, T. Ghosh, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gruppuso, J. E. Gudmundsson, J. Hamann, W. Handley, F. K. Hansen, D. Herranz, S. R. Hildebrandt, E. Hivon, Z. Huang, A. H. Jaffe,

- W. C. Jones, A. Karakci, E. Keihänen, R. Keskitalo, K. Kiiveri, J. Kim, T. S. Kisner, L. Knox, N. Krachmalnicoff, M. Kunz, H. Kurki-Suonio, G. Lagache, J. M. Lamarre, A. Lasenby, M. Lattanzi, C. R. Lawrence, M. Le Jeune, P. Lemos, J. Lesgourgues, F. Levrier, A. Lewis, M. Liguori, P. B. Lilje, M. Lilley, V. Lindholm, M. López-Cañego, P. M. Lubin, Y. Z. Ma, J. F. Macías-Pérez, G. Maggio, D. Maino, N. Mandolesi, A. Mangilli, A. Marcos-Caballero, M. Maris, P. G. Martin, M. Martinelli, E. Martínez-González, S. Matarrese, N. Mauri, J. D. McEwen, P. R. Meinhold, A. Melchiorri, A. Mennella, M. Migliaccio, M. Millea, S. Mitra, M. A. Miville-Deschênes, D. Molinari, L. Montier, G. Morgante, A. Moss, P. Natoli, H. U. Nørgaard-Nielsen, L. Pagano, D. Paoletti, B. Partridge, G. Patanchon, H. V. Peiris, F. Perrotta, V. Pettorino, F. Piacentini, L. Polastri, G. Polenta, J. L. Puget, J. P. Rachen, M. Reinecke, M. Remazeilles, A. Renzi, G. Rocha, C. Rosset, G. Roudier, J. A. Rubiño-Martín, B. Ruiz-Granados, L. Salvati, M. Sandri, M. Savelainen, D. Scott, E. P. S. Shellard, C. Sirignano, G. Sirri, L. D. Spencer, R. Sunyaev, A. S. Suur-Uski, J. A. Tauber, D. Tavagnacco, M. Tenti, L. Toffolatti, M. Tomasi, T. Trombetti, L. Valenziano, J. Valiviita, B. Van Tent, L. Vibert, P. Vielva, F. Villa, N. Vittorio, B. D. Wandelt, I. K. Wehus, M. White, S. D. M. White, A. Zacchei, and A. Zonca. Planck 2018 results. VI. Cosmological parameters. *Astron. Astrophys.*, 641:A6, September 2020.
- [240] Nathalie Palanque-Delabrouille, Christophe Yèche, Julien Baur, Christophe Magueville, Graziano Rossi, Julien Lesgourgues, Arnaud Borde, Etienne Burtin, Jean-Marc LeGoff, James Rich, Matteo Viel, and David Weinberg. Neutrino masses and cosmology with Lyman-alpha forest power spectrum. *JCAP*, 2015(11):011–011, November 2015.
- [241] Adam G. Riess, Wenlong Yuan, Lucas M. Macri, Dan Scolnic, Dillon Brout, Stefano Casertano, David O. Jones, Yukei Murakami, Gagandeep S. Anand, Louise Breuval, Thomas G. Brink, Alexei V. Filippenko, Samantha Hoffmann, Saurabh W. Jha, W. D’arcy Kenworthy, John Mackenty, Benjamin E. Stahl, and WeiKang Zheng. A Comprehensive Measurement of the Local Value of the Hubble Constant with 1 km s<sup>-1</sup> Mpc<sup>-1</sup> Uncertainty from the Hubble Space Telescope and the SH0ES Team. *Astrophys. J. Let.*, 934(1):L7, July 2022.
- [242] Antony Lewis and Anthony Challinor. CAMB: Code for Anisotropies in the Microwave Background. Astrophysics Source Code Library, record ascl:1102.026, February 2011.
- [243] Matthew McQuinn. The Implications of Gunn-Peterson Troughs in the He II Ly $\alpha$  Forest. *ApJL*, 704(2):L89–L92, October 2009.
- [244] Gábor Worseck, J. Xavier Prochaska, Matthew McQuinn, Aldo Dall’Aglio, Cora Fechner, Joseph F. Hennawi, Dieter Reimers, Philipp Richter, and Lutz Wisotzki. The End of Helium Reionization at  $z \sim 2.7$  Inferred from Cosmic Variance in HST/COS He II Ly $\alpha$  Absorption Spectra. *ApJL*, 733(2):L24, June 2011.
- [245] Gábor Worseck, Frederick B. Davies, Joseph F. Hennawi, and J. Xavier Prochaska. The Evolution of the He II-ionizing Background at Redshifts 2.3  $\leq z \leq$  3.8 Inferred from

- a Statistical Sample of 24 HST/COS He II Ly $\alpha$  Absorption Spectra. *ApJ*, 875(2):111, April 2019.
- [246] Alexa M. Morales, Charlotte A. Mason, Sean Bruton, Max Gronke, Francesco Haardt, and Claudia Scarlata. The Evolution of the Lyman-alpha Luminosity Function during Reionization. *ApJ*, 919(2):120, October 2021.
- [247] Christian Pedersen, Andreu Font-Ribera, Thomas D. Kitching, Patrick McDonald, Simeon Bird, Anže Slosar, Keir K. Rogers, and Andrew Pontzen. Massive neutrinos and degeneracies in Lyman-alpha forest simulations. *JCAP*, 2020(4):025, April 2020.
- [248] George D. Becker, Paul C. Hewett, Gábor Worseck, and J. Xavier Prochaska. A refined measurement of the mean transmitted flux in the Ly $\alpha$  forest over  $2 < z < 5$  using composite quasar spectra. *MNRAS*, 430(3):2067–2081, April 2013.
- [249] Huanqing Chen, Rupert Croft, and Nickolay Y. Gnedin. Reconstructing Large-scale Temperature Profiles around  $z \sim 6$  Quasars. *arXiv e-prints*, page arXiv:2208.13787, August 2022.
- [250] Christian Pedersen, Andreu Font-Ribera, and Nickolay Y. Gnedin. Compressing the cosmological information in one-dimensional correlations of the Lyman- $\alpha$  forest. *arXiv e-prints*, page arXiv:2209.09895, September 2022.
- [251] Euclid Collaboration, Mischa Knabenhans, Joachim Stadel, Stefano Marelli, Doug Potter, Romain Teyssier, Laurent Legrand, Aurel Schneider, Bruno Sudret, Linda Blot, Saeeda Awan, Carlo Burigana, Carla Sofia Carvalho, Hannu Kurki-Suonio, and Gabriele Sirri. Euclid preparation: II. The EUCLIDEMULATOR - a tool to compute the cosmology dependence of the nonlinear matter power spectrum. *MNRAS*, 484(4):5509–5529, April 2019.
- [252] Justin Alsing, Tom Charnock, Stephen Feeney, and Benjamin Wandelt. Fast likelihood-free cosmology with neural density estimators and active learning. *MNRAS*, 488(3):4440–4458, September 2019.
- [253] Richard Neveux, Etienne Burtin, Vanina Ruhlmann-Kleider, Arnaud de Mattia, Agne Semenaite, Kyle S. Dawson, Axel de la Macorra, Will J. Percival, Graziano Rossi, Donald P. Schneider, and Gong-Bo Zhao. Combined full shape analysis of BOSS galaxies and eBOSS quasars using an iterative emulator. *arXiv e-prints*, page arXiv:2201.04679, January 2022.
- [254] Supranta S. Boruah, Tim Eifler, Vivian Miranda, and Sai Krishanth P. M. Accelerating cosmological inference with Gaussian processes and neural networks – an application to LSST Y1 weak lensing and galaxy clustering. *arXiv e-prints*, page arXiv:2203.06124, March 2022.
- [255] Sihan Yuan, Lehman H. Garrison, Daniel J. Eisenstein, and Risa H. Wechsler. Stringent  $\sigma_8$  constraints from small-scale galaxy clustering using a hybrid MCMC+emulator framework. *arXiv e-prints*, page arXiv:2203.11963, March 2022.

- [256] Yi Ji, Simon Mak, Derek Soeder, J-F Paquet, and Steffen A. Bass. A graphical multi-fidelity Gaussian process model, with application to emulation of expensive computer simulations. *arXiv e-prints*, page arXiv:2108.00306, July 2021.
- [257] Keir K. Rogers and Hiranya V. Peiris. Strong Bound on Canonical Ultralight Axion Dark Matter from the Lyman-Alpha Forest. *Phys. Rev. L*, 126(7):071302, February 2021.
- [258] Zhongxu Zhai, Jeremy L. Tinker, Matthew R. Becker, Joseph DeRose, Yao-Yuan Mao, Thomas McClintock, Sean McLaughlin, Eduardo Rozo, and Risa H. Wechsler. The Aemulus Project. III. Emulation of the Galaxy Correlation Function. *ApJ*, 874(1):95, March 2019.
- [259] Christian Arnold, Baojiu Li, Benjamin Giblin, Joachim Harnois-Déraps, and Yan-Chuan Cai. FORGE – the f(R) gravity cosmic emulator project I: Introduction and matter power spectrum emulator. *arXiv e-prints*, page arXiv:2109.04984, September 2021.
- [260] Timothy Wing Hei Yiu, Janis Fluri, and Tomasz Kacprzak. A tomographic spherical mass map emulator of the KiDS-1000 survey using conditional generative adversarial networks. *arXiv e-prints*, page arXiv:2112.12741, December 2021.
- [261] Neerav Kaushal, Francisco Villaescusa-Navarro, Elena Giusarma, Yin Li, Conner Hawry, and Mauricio Reyes. NECOLA: Toward a Universal Field-level Cosmological Emulator. *ApJ*, 930(2):115, May 2022.
- [262] Jie Wang. An Intuitive Tutorial to Gaussian Processes Regression. *arXiv e-prints*, page arXiv:2009.10862, September 2020.
- [263] Yen-Chi Chen, Shirley Ho, Ananth Tenneti, Rachel Mand elbaum, Rupert Croft, Tiziana DiMatteo, Peter E. Freeman, Christopher R. Genovese, and Larry Wasserman. Investigating galaxy-filament alignments in hydrodynamic simulations using density ridges. *MNRAS*, 454(3):3341–3350, 12 2015.
- [264] T. W. B. Kibble. Topology of cosmic domains and strings. *Journal of Physics A Mathematical General*, 9(8):1387–1398, 8 1976.
- [265] Tanmay Vachaspati and Alexander Vilenkin. Gravitational radiation from cosmic strings. *Phys. Rev. D*, 31:3052–3058, 6 1985.
- [266] M. B. Hindmarsh and T. W. B. Kibble. Cosmic strings. *Reports on Progress in Physics*, 58(5):477–562, 5 1995.
- [267] M. Alice Gasparini, Phil Marshall, Tommaso Treu, Eric Morganson, and Florian Dubath. Direct observation of cosmic strings via their strong gravitational lensing effect - I. Predictions for high-resolution imaging surveys. *MNRAS*, 385(4):1959–1964, 4 2008.



- [268] Robert H. Brandenberger. Topological Defects and Structure Formation. *International Journal of Modern Physics A*, 9(13):2117–2189, 1 1994.
- [269] T. Vachaspati. Cosmic strings and the large-scale structure of the Universe. *Phys. Rev. Letters*, 57(13):1655–1657, 9 1986.
- [270] Edmund J. Copeland and T. W. B. Kibble. Cosmic strings and superstrings. *Proceedings of the Royal Society of London Series A*, 466(2115):623–657, 3 2010.
- [271] Joseph Polchinski. Introduction to Cosmic F- and D-Strings. *arXiv e-prints*, pages hep-th/0412244, 12 2004.