Inferring high-redshift large-scale structure dynamics from the Lyman- α forest

Natalia Porqueres^{1,2}, Jens Jasche³, Guilhem Lavaux^{4,5}, and Torsten Enßlin^{1,6}

¹ Max-Planck-Institute für Astrophysik (MPA), Karl-Schwarzschild-Strasse 1, 85741 Garching, Germany
e-mail: natalia@mpa-garching.mpg.de

² Excellence Cluster Universe, Technische Universität München, Boltzmannstrasse 2, 85748 Garching, Germany

³ The Oskar Klein Centre, Department of Physics, Stockholm University, Albanova University Center, 106 91 Stockholm, Sweden

⁴ Institut Lagrange de Paris (ILP), Sorbonne Universités, 98bis Boulevard Arago, 75014 Paris, France

⁵ Institut d'Astrophysique de Paris (IAP), UMR 7095, CNRS – UPMC Université Paris 6, Sorbonne Universités, 98bis Boulevard Arago, 75014 Paris, France

⁶ Ludwig-Maximilians-Universität München, Geschwister-Scholl-Platz 1, 80539 München, Germany

Received 5 July 2019 / Accepted 13 September 2019

ABSTRACT

One of the major science goals over the coming decade is to test fundamental physics with probes of the cosmic large-scale structure out to high redshift. Here we present a fully Bayesian approach to infer the three-dimensional cosmic matter distribution and its dynamics at $z > 2$ from observations of the Lyman- α forest. We demonstrate that the method recovers the unbiased mass distribution and the correct matter power spectrum at all scales. Our method infers the three-dimensional density field from a set of one-dimensional spectra, interpolating the information between the lines of sight. We show that our algorithm provides unbiased mass profiles of clusters, becoming an alternative for estimating cluster masses complementary to weak lensing or X-ray observations. The algorithm employs a Hamiltonian Monte Carlo method to generate realizations of initial and evolved density fields and the three-dimensional large-scale flow, revealing the cosmic dynamics at high redshift. The method correctly handles multi-modal parameter distributions, which allow constraining the physics of the intergalactic medium with high accuracy. We performed several tests using realistic simulated quasar spectra to test and validate our method. Our results show that detailed and physically plausible inference of three-dimensional large-scale structures at high redshift has become feasible.

Key words. large-scale structure of Universe – dark matter

1. Introduction

Currently, cosmology is at a crossroad. While the standard model of cosmology, the Λ cold dark matter (Λ CDM), fits the bulk of cosmological observations to extraordinary accuracy, some tensions between the model and observations seem to persist and increase (Planck Collaboration XIII 2016; Planck Collaboration IX 2019). Amongst them are the H_0 and σ_8 tensions (e.g. Planck Collaboration XXIV 2016; Riess et al. 2016, 2018; Köhlinger et al. 2017; Abbott et al. 2018; Rusu et al. 2019). Additionally, Lyman- α auto-correlation and cross-correlations with quasar data reported a 2.3σ tension with the flat Λ CDM prediction obtained from *Planck* observations (see e.g. Delubac et al. 2015; du Mas des Bourboux et al. 2017).

The resolution of these tensions may encompass systematic effects, but may also be the first signs of new physics indicated by novel cosmological data of increasing quality. New observations and better control on systematic effects in data analyses are inevitable to gain new insights into the physical processes driving the evolution of the universe. For this reason, in this work, we present a novel and statistically rigorous approach to extract cosmologically relevant and significant information from high-redshift Lyman- α forest observations tracing the dynamic evolution of cosmic structures.

Detailed analyses of the spatial distribution of matter and growth of cosmic structures can provide significant information

to discriminate between models of homogeneous dark energy and modifications of gravity, to determine neutrino masses and their mass hierarchy, and to investigate the dynamical clustering behaviour of warm or cold dark matter (e.g. Colombi et al. 1996; Huterer et al. 2015; Basilakos & Nesseris 2016; Frenk & White 2012; Boyle & Komatsu 2018; Mishra-Sharma et al. 2018).

To harvest this information, cosmology now turns to analyse the inhomogeneous matter distribution with next-generation galaxy surveys such as the Large Synoptic Survey Telescope (LSST) and the *Euclid* satellite mission probing the galaxy distribution out to redshifts $z \sim 3$ and beyond (LSST Science Collaboration 2009; Refregier et al. 2010; Laureijs et al. 2011; Alam et al. 2016).

While most of this research focuses on studying the cosmic large-scale structure with galaxy clustering and weak lensing observations, the Lyman- α (Ly- α) forest has the potential to provide important complementary information. Firstly, the Ly- α forest probes the matter distribution at higher redshift with higher spatial resolution than can be achieved with galaxy sampling rates. By probing scales down to 1 Mpc, the Ly- α forest is sensitive to neutrino masses and dark matter models (Viel et al. 2013; Rossi 2014, 2017; Palanque-Delabrouille et al. 2015; Rossi et al. 2015; Yèche et al. 2017). Secondly, while galaxy clustering probes high-density regions, the Ly- α forest is particularly sensitive to under-densities in the matter distribution (Peirani et al. 2014; Sorini et al. 2016). In addition, the Ly- α at

$z > 2$ is redshifted to an optical band for which the atmosphere is transparent. The Ly- α emission becomes then one of the more powerful probes at redshifts that are otherwise challenging to access with ground-based galaxy surveys (Lee 2016).

Specifically, auto-correlations of the Ly- α forest along the line of sight have been used to infer cosmological parameters which agree by and large with CMB observations (e.g. Seljak et al. 2006; Viel et al. 2006; Slosar et al. 2011; Busca et al. 2013; Bautista et al. 2017; Blomqvist et al. 2019). Ly- α forest data have also been used to constrain neutrino masses (Palanque-Delabrouille et al. 2015; Rossi et al. 2015; Yèche et al. 2017), test warm dark matter models (Viel et al. 2013), and study the thermal history of the intergalactic medium (IGM; Nasir et al. 2016; Boera et al. 2019).

In line with these promises, a large number of surveys mapping the Ly- α forest at higher redshifts has been launched. The eBOSS observations in the SDSS DR14 are expected to provide the spectra of 435 000 quasars over 7500 deg^2 and re-observe the lines of sight from the previous data release with low signal-to-noise ($S/N < 3$, Myers et al. 2015). DESI (Dark Energy Spectroscopic Instrument, Levi et al. 2013; DESI Collaboration 2016) will map the Ly- α forest over a larger fraction of the sky ($14\,000 \text{ deg}^2$), targeting 50 high-redshift quasars per deg^2 . Increasing the density of tracers, the ongoing COSMOS Lyman Alpha Mapping and Tomography survey (CLAMATO, Lee et al. 2018) maps the Ly- α forest over a small part of the sky (600 arcmin^2) with a separation between lines of sight of $2.4 h^{-1} \text{ Mpc}$, 20 times smaller than the BOSS survey (Eisenstein et al. 2011; Dawson et al. 2013) and smaller than the expected separation in DESI ($3.7 h^{-1} \text{ Mpc}$, Horowitz et al. 2019). Future surveys like Maunakea Spectroscopic Explorer (MSE, McConnachie et al. 2016) and LSST (LSST Science Collaboration 2009) will increase the amount of available Ly- α forest observations by a factor of 10.

At present, major analyses of the Ly- α forest focus only on the analysis of the matter power spectrum (e.g. Croft et al. 1998; Seljak et al. 2006; Viel et al. 2006; Bird et al. 2011; Palanque-Delabrouille et al. 2015; Rossi et al. 2015; Nasir et al. 2016; Yèche et al. 2017; Boera et al. 2019). However, these approaches ignore significant amounts of information contained in the higher-order statistics of the matter density field as generated by non-linear gravitational dynamics in the late time universe (He et al. 2018).

Capturing the full information content of the cosmic large-scale structure requires a field-based approach to infer the entire three-dimensional cosmic large-scale structure from observations. This poses a particular challenge for the analyses of Ly- α forest observations, which provide sparse inherently one-dimensional information along the lines of sight. Various approaches to perform three-dimensional density reconstructions from one-dimensional Ly- α forests have been proposed in the literature (e.g. Kitaura et al. 2012; Cisewski et al. 2014; Stark et al. 2015a; Ozbek et al. 2016; Horowitz et al. 2019). Gallerani et al. (2011) and Kitaura et al. (2012) proposed a Gibbs sampling scheme to jointly infer density and velocity fields and corresponding power-spectra. However, these approaches assume matter density amplitudes to be log-normally distributed. The log-normal distribution reproduces one- and two-point statistics but fails to reproduce higher-order statistics associated with the filamentary dark matter distribution. In an attempt to extrapolate information from one-dimensional quasar spectra into the three-dimensional volume, Cisewski et al. (2014) applied a local polynomial smoothing method. Ozbek et al. (2016) and Stark et al. (2015a) employed a Wiener filtering approach to reconstruct the three-dimensional density field between lines of sight of Ly- α forest data. In

order to reproduce higher-order statistics, Horowitz et al. (2019) recently used a large-scale optimization approach to fit a gravitational structure growth model to Ly- α data, showing that this approach allows recovering the more filamentary structure of the cosmic web. Although the approach improves over linear and isotropic Wiener filtering approaches, it shows systematic deviations of reconstructed matter power-spectra and underestimates density amplitudes at scales corresponding to the mean separation between lines of sight (Horowitz et al. 2019).

To go beyond previous approaches, in this work, we present a fully Bayesian and statistically rigorous approach to perform dynamical matter clustering analyses with high-redshift Ly- α forest data while accounting for all uncertainties inherent to the observations. The aim of this work is to provide a novel and fully Bayesian approach to infer physically plausible three-dimensional density and velocity fields from Ly- α forest data. Our approach builds upon the algorithm for Bayesian Origin Reconstruction from Galaxies (BORG, Jasche & Wandelt 2013; Jasche & Lavaux 2019), which employs physical models of structure formation and sophisticated Markov chain Monte Carlo (MCMC) techniques to optimally extract large-scale structure information from data and quantify corresponding uncertainties.

To make such inferences feasible, we developed a likelihood based on the fluctuating Gunn-Peterson approximation (FGPA, Gunn & Peterson 1965) and jointly constrained the astrophysical properties of the IGM. We tested and validated our approach with simulated data emulating the CLAMATO survey, showing that the algorithm recovers the unbiased dark-matter field.

The paper is organized as follows. Section 4 provides a brief overview of our Bayesian inference framework, BORG, as required for this work. In Sect. 2, we discuss the physics of the Ly- α forest underlying the data model described in Sect. 3. Section 5 describes the generation of simulated data employed to test and validate the algorithm. The performance of the algorithm is tested in Sect. 6 and the results are shown in Sect. 7, where we also discuss possible applications of the method to scientifically exploit the Ly- α forest. Finally, Sect. 9 summarizes the results and discusses further extensions of the algorithm.

2. The physics of the Lyman- α forest

Developing a Bayesian framework to infer the matter distribution from the Ly- α forest requires to incorporate the corresponding light absorption physics into the data model.

When traversing the universe, photons emitted from quasars are scattered by neutral hydrogen (HI) gas residing inside cosmic large-scale structures. Scattering of quasar light results in observed spectra covered with absorption features referred to as the Ly- α forest. While high-density regions obscure the light and attenuate quasar fluxes, under-dense regions are almost transparent. Therefore, the signal comes from the under-dense regions. Due to the cosmological redshift, different HI regions absorb photons from different wavelengths in the quasar spectrum (see e.g. Mo et al. 2010). Consequently, every HI absorber leaves an absorption feature to the spectrum, permitting to trace the distance and density of HI regions along the line of sight. The Ly- α forest, therefore, provides a formidable probe of the cosmic large-scale structure along the observers past light cone.

Since the Ly- α forest generated at $z > 2$ is redshifted to the optical atmospheric window, it can be observed by ground-based telescopes, becoming one of the more relevant probes at redshifts that are otherwise challenging to access with galaxy surveys. Observations of the Ly- α forest at lower redshift, from $z = 1.5$ down to $z = 0$ (Davé et al. 1999; Davé 2001; Williger et al. 2010),

require UV space-based spectrographs as Faint Object Spectrograph in the *Hubble* Space Telescope (Bahcall et al. 1993, 1996; Weymann et al. 1998).

In contrast to galaxy surveys, requiring assumptions on galaxy formation to model galaxy biases, modelling the Ly- α forest does not involve complicated models. At $z > 2$, HI gas, producing the Ly- α forest, is in radiative equilibrium of photoionization due to ultraviolet (UV) background and adiabatic cooling due to the cosmic expansion (Hui & Gnedin 1997). This yields a tight relation between the temperature of the IGM and the dark matter density ρ

$$T = T_0 \left(\frac{\rho}{\bar{\rho}} \right)^{\gamma-1}, \quad (1)$$

where T_0 and γ are constants that depend on the re-ionization history and the spectral shape of the UV background sources (Hui & Gnedin 1997).

In thermal equilibrium, HI number density can be expressed as (Hui & Gnedin 1997)

$$n_{\text{HI}}(\mathbf{x}) = \frac{\alpha(T(\mathbf{x}))}{\Gamma} \left[n_0(1+z)^3 (1+\delta(\mathbf{x})) \right]^2 \propto (1+\delta(\mathbf{x}))^\beta, \quad (2)$$

where $\mathbf{x}$ is the comoving position, $\alpha(T) \propto T^{-0.7}$ is the radiative recombination rate of HI, Γ is the photoionization rate of neutral hydrogen, $n_0(1+z)^3$ is the mean baryon number density at redshift z and δ is the dark matter density contrast. On scales larger than Jeans' scale (100 kpc), thermal pressure in the gas is negligible and the IGM traces the dark matter distribution with high accuracy (Peirani et al. 2014).

Combining Eqs. (1) and (2) leads to the fluctuating Gunn-Peterson approximation (FGPA, Gunn & Peterson 1965) for the transmitted flux:

$$F(z, \hat{\mathbf{x}}) = \exp \left[-A(1+\delta(\mathbf{x}))^\beta \right], \quad (3)$$

with z the considered absorption redshift, $\mathbf{x}$ the corresponding comoving distance, $\hat{\mathbf{x}}$ the associated unit vector, $\beta = 2 - 0.7(\gamma - 1)$ and $A \propto (1+z)^6 T_0^{-0.7} \Gamma^{-1}$. The connection between the redshift and the cartesian coordinate of the regular grid $\mathbf{x}$ is given by the projector along the line of sight (see Appendix A). We note that the astrophysical properties of HI gas are encoded in the parameters A and β of the FGPA model. Although this work does not include a treatment of redshift space distortions (RSD) or thermal broadening of the absorption lines, we intend to extend the current formalism to include it in future works.

3. A data model for the Ly- α forest

To incorporate Ly- α forest observations into the Bayesian framework, we now built a data model using the light absorption physics described in Sect. 2. Specifically, we assume the FGPA transmission model and Gaussian pixel noise for measured quasar spectra. A corresponding likelihood distribution can then be expressed as:

$$P(\delta^f | F) = \prod_{n,x} \frac{1}{\sqrt{2\pi}\sigma} \times \exp \left[-\frac{\left((F_n)_x - \exp \left[-A(1+\delta_x^f)^\beta \right] \right)^2}{2\sigma^2} \right], \quad (4)$$

where n labels respective quasar spectra, x indexes different volume elements intersected by the n th quasar line of sight and δ^f is the non-linear final density contrast evaluated at $z = 2.5$.

Fig. 1. Flow chart depicting the iterative block sampling approach of the BORG inference framework. As can be seen, the BORG algorithm first infers three-dimensional initial and evolved density fields from quasar spectra using assumed astrophysical parameters of the FGPA model. Then, using realizations of the evolved density field, the algorithm updates these astrophysical parameters and a new iteration of the process begins. Iterating this procedure results in a valid Markov Monte Carlo chain that correctly explores the joint posterior distribution of the three-dimensional matter distribution and astrophysical properties underlying Ly- α forest observations.

This likelihood is then implemented into the large-scale structure sampler of the BORG framework. The corresponding physical forward modelling approach proceeds as follows. Using realizations of the three-dimensional field of primordial fluctuations, the dynamical structure formation model evaluates non-linear realizations of the dark matter distribution at $z = 2.5$. Using these matter field realizations and the FGPA data model, BORG predicts quasar spectra that are compared to actual observations via the Gaussian likelihood of Eq. (4). More details on the sampling procedure are given in Appendix B. To efficiently implement this non-linear and non-Gaussian forward modelling approach into an MCMC framework, we employed a Hamiltonian Monte Carlo (HMC) method, as described in Appendix C.

In this work, we focus on the inference of the spatial dark matter distribution and its dynamics. For this reason, we treat the parameters A and β mostly as nuisance parameters. Following similar approaches described in previous works (Jasche & Wandelt 2013; Jasche & Lavaux 2017, 2019; Kodi Ramanah et al. 2019), we use efficient slice sampling techniques to sample these parameters. We note that these parameters can carry valuable information on the astrophysical properties of the IGM. As illustrated in Sect. 7.4, our inference can also be used to make statements on the astrophysics of the IGM.

The entire iterative sampling approach is depicted in Fig. 1. As can be seen, the method iteratively infers three-dimensional matter density fields and parameters of the FGPA model, resulting in a valid MCMC approach to jointly explore density fields and astrophysical parameters.

4. Method

As mentioned above, this work extends our previously developed BORG algorithm in order to analyse the spatial matter distribution underlying Ly- α forest observations. Here we provide only a brief summary of the algorithm and corresponding concepts, but interested readers can find more detailed descriptions in our previous works (Jasche & Wandelt 2013; Jasche et al. 2015; Lavaux & Jasche 2016; Jasche & Lavaux 2017, 2019).

The BORG algorithm is a large-scale Bayesian inference framework aiming at inferring the non-linear spatial dark matter distribution and its dynamics from cosmological data sets. The underlying idea is to fit full dynamical gravitational structure formation models to observations tracing the spatial distribution of matter. Using non-linear structure growth models, the BORG algorithm can exploit the full statistical power of higher-order statistics imprinted to the matter distribution by gravitational clustering. In fitting dynamical structure growth models to data, the task of inferring non-linear matter density fields turns into a statistical initial conditions problem aiming at inferring the spatial distribution of primordial matter fluctuations from which present structures formed via gravitational structure growth. As such the BORG algorithm naturally links primordial initial conditions to late time observations and permits to infer dynamical properties and the structure formation history from present galaxy observations (Jasche & Wandelt 2013; Jasche et al. 2015; Lavaux & Jasche 2016; Jasche & Lavaux 2017, 2019).

The BORG algorithm employs a large-scale structure posterior distribution based on the well-developed prior understanding of almost Gaussian primordial density fluctuations and gravitational structure formation to predict physically plausible realizations of present matter density fields. More specifically BORG encodes a Gaussian prior for the initial density contrast at an initial cosmic scale factor of $a = 10^{-3}$. Initial and evolved density fields are linked by deterministic gravitational evolution mediated by various physics models of structure growth. Specifically, BORG incorporates physical models based on Lagrangian Perturbation Theory (LPT) but also fully non-linear particle mesh (PM) models.

Our Bayesian inference approach has been previously shown to perform accurate dynamic mass estimation and provide mass measurements that agree well with complimentary standard weak lensing and X-ray observations (Jasche & Lavaux 2019). More recently, we have shown, that BORG can exploit the geometric shapes of the cosmic large-scale structure to significantly improve constraints on cosmological parameters via the Alcock-Paczynski test (Kodi Ramanah et al. 2019). Additional projects, conducted with the BORG algorithm, can be found at our web-page of the Aquila consortium¹.

In this work, we rely on an LPT approach to structure formation since the Ly- α forest mostly arises from under-dense regions that can be conveniently modelled by perturbation theory (Peirani et al. 2014; Sorini et al. 2016). The dynamical model permits to recover the non-linear higher-order statistics associated with the filamentary matter distribution of the cosmic large-scale structure. Our approach also immediately provides detailed inferences of large-scale velocity fields and structure growth information at high redshifts as is demonstrated in the remainder of this work.

A feature of particular relevance to this work is that BORG employs a modular statistical programming engine that executes a statistically rigorous Markov chain to marginalize out any nuisance parameters associated to the data model, such as unknown biases or systematic effects. This statistical programming engine permits us to easily and straightforwardly implement complex data models. In particular, in this work, we perform joint analyses of the cosmological large-scale structure and unknown astrophysical properties of the IGM medium described by the parameters of the FGPA model. The corresponding iterative block sampling inference approach is illustrated in Fig. 1.

5. Generating artificial Ly- α forest observations

To test our Ly- α forest inference framework, we generate artificial mock observations emulating the CLAMATO survey (Stark et al. 2015b; Lee et al. 2018).

Mock data are constructed by first generating Gaussian initial conditions on a cubic Cartesian grid of side length of $256 h^{-1}$ Mpc with a resolution of $1 h^{-1}$ Mpc. To generate primordial Gaussian density fluctuations we use a cosmological matter power-spectrum including the Baryonic wiggles calculated according to the prescription provided by Eisenstein & Hu (1998, 1999). We further assume a standard Λ CDM cosmology with the following set of parameters: $\Omega_m = 0.31$, $\Omega_\Lambda = 0.69$, $\Omega_b = 0.022$, $h = 0.6777$, $\sigma_8 = 0.83$, $n_s = 0.9611$ (Planck Collaboration XIII 2016). We assumed $H = 100 h \text{ km s}^{-1} \text{ Mpc}^{-1}$.

To generate realizations of the non-linear density field, we evolve these Gaussian initial conditions via Lagrangian Perturbation Theory. This involves simulating displacements for 512^3 particles in the LPT simulation. Final density fields are constructed by estimating densities via the cloud-in-cell scheme from simulated particles on a Cartesian equidistant grid with 256^3 volume elements. We further apply a Gaussian smoothing kernel of $\sigma = 0.5 h^{-1}$ Mpc to the fields to simulate the difference between dark matter and gas density fields (see e.g. Peirani et al. 2014; Stark et al. 2015b). A three dimensional quasar flux field is generated by applying the FGPA model, given in Eq. (3), to the final density field and assuming constant parameters $A = 0.35$ and $\beta = 1.56$ at $z = 2.5$, corresponding to the values in Stark et al. (2015b).

From this three-dimensional quasar flux field, we generate individually observed skewers by tracing lines of sight through the volume. Specifically, we generate a total of 1024 lines of sight parallel to the z -axis of the box, distributed in a regular grid of 32×32 lines of sight spaced regularly with a separation of $8 h^{-1}$ Mpc. The separation between lines of sight is the most limiting factor of Ly- α surveys. Although the CLAMATO survey achieves a separation of $2.4 h^{-1}$ Mpc, the lines of sight in this work are separated by $8 h^{-1}$ Mpc. This gives us the opportunity to explore the potential of our method to reconstruct the three-dimensional density field between lines of sight. Finally, we added Gaussian pixel-noise to the flux with $\sigma = 0.03$.

Although tests are performed with the same noise distribution in all lines of sight, Sect. 8 presents an analysis with a different signal-to-noise ratio (S/N) for each line of sight.

6. Testing sampler performance

In this section, we present a series of tests to evaluate the performance of our method. In particular, we focus on the convergence (Sect. 6.1) and efficiency (Sect. 6.2) to infer the underlying dark matter density field. We provide tests of the posterior power-spectra and perform posterior predictive tests for quasar spectra to check that the inferred model agrees with the data.

6.1. The warm-up phase of the sampler

In the large sample limit, any properly set up Markov chain is guaranteed to approach a stationary distribution that provides an unbiased estimate of the target distribution. While Markov chains are typically started from a place remote from the target distribution after a finite amount of transition steps, the chain acquires a stationary state. Once the chain is in a stationary state, we may start recording samples to perform statistical analyses of the inference problem. To test when the Markov sampler has

¹ <https://aquila-consortium.org>

Fig. 2. Burn-in of the posterior matter power spectra from the inferred initial conditions from a BORG analysis. The colour scale shows the evolution of the matter power spectrum with the number of samples: the Markov chain is initialized with a Gaussian initial density field scaled by a factor 10^{-3} and the amplitudes of the power spectrum systematically drift towards the fiducial values. Monitoring this drift allows us to identify when the Markov chain approaches a stationary distribution and provides unbiased estimates of the target distribution. The dashed lines indicate the underlying power spectrum and the 1- and 2- σ cosmic variance limits. At the end of the warm-up phase, the algorithm recovers the true matter power spectrum in all range of Fourier modes.

passed its initial warm-up phase, we follow a similar approach as described in previous works (Jasche & Wandelt 2013; Jasche & Lavaux 2017, 2019; Kodi Ramanah et al. 2019; Porqueres et al. 2019), by initializing the Markov chain with an over-dispersed state and trace the systematic drift of inferred quantities towards their preferred regions in parameter space. Specifically, we initialized the Markov chain with a random Gaussian cosmological density field scaled by a factor 10^{-3} and monitored the drift of corresponding posterior power-spectra during the initial warm-up phase. The results of this exercise are presented in Fig. 2. As can be seen, successive measurements of the posterior power-spectrum during the initial warm-up phase show a systematic drift of power-spectrum amplitudes towards their fiducial values. The sampler, therefore, correctly recovers the power of the initial density field and moves the chain towards regions of high probability in the parameter space. By the end of the warm-up phase, the sampler has found an unbiased representation of the matter distribution at all Fourier modes considered in this work. Starting the sampler from an over-dispersed state, therefore, provides us with an important diagnostics to test the validity of the sampling algorithm.

6.2. Statistical efficiency of sampler

By design, subsequent samples in Markov chains are correlated. The statistical efficiency of an MCMC algorithm is determined by the number of independent samples that can be drawn from

a chain of a given length. To estimate the statistical efficiency of the sampler, we estimate the correlation length of the astrophysical parameters A and β along the Markov chain. For a parameter θ the auto-correlation for samples with a given lag in the chain can be estimated as

$$C_n(\theta) = \frac{1}{N-n} \sum_{i=0}^{N-n} \frac{(\theta^i - \langle \theta \rangle)(\theta^{i+n} - \langle \theta \rangle)}{\text{Var}(\theta)} \quad (5)$$

where n is the lag in MCMC samples, $\langle \theta \rangle$ is the mean and $\text{Var}(\theta)$ is the variance. We typically determine the correlation length by estimating the lag n_C at which the auto-correlation C_n dropped below 0.1. The number n_C therefore present the number of transitions required to produce one independent sample. The results of this test are presented in Fig. 6. As can be seen, correlation lengths for parameters A and β amount to 640 and 830 samples, respectively. To significantly improve statistical efficiency, in the future, we use similar strategies to obtain a faster mixing of the chain as described in Kodi Ramanah et al. (2019).

6.3. Posterior predictions for quasar spectra

To test whether inferred density fields provide accurate explanations for the observations, we perform a simple posterior predictive test (see e.g. Gelman et al. 2004). Generally, posterior predictive tests provide good diagnostics about the adequacy of data models in explaining observations and identifying possible systematic problems with the inference.

Here we use density fields and the astrophysical parameters A and β , inferred by BORG, to predict expected quasar fluxes. If posterior predictions agree with actual observations within the uncertainty bounds, then the data model can be considered to be sufficient to analyse the data. In contrast, any misfits or systematic deviations would indicate a problem of the data model or the inference process.

Specifically, we applied the FGPA model given in Eq. (3) to inferred density fields:

$$(F_{\text{pp}})_x = \frac{1}{N} \sum_{i=0}^N \exp \left[-A_i (1 + (\delta_i^f)_x)^{\beta_i} \right] \quad (6)$$

where F_{pp} is the posterior-predicted flux, i labels the samples and x labels the volume elements.

The result of this test is presented in Fig. 3. As can be seen, the posterior predicted quasar spectrum nicely traces the data input within the observational 1σ uncertainty region. This demonstrates that the method correctly locates absorber positions and corresponding amplitudes of the underlying density field.

While Fig. 3 shows results only for a single line of sight, more generally, we also explored the flux errors for all lines of sight. The corresponding distribution of flux errors for posterior predictions is presented in Fig. 4. The distribution of flux errors corresponds to the distribution of pixel-noise in the spectra, demonstrating that our method is close to the theoretical optimum. This indicates that our method exhibits good control of the data model, including the handling of nuisance parameters and uncertainties.

Since many studies of the Ly- α forest are based on the power spectrum of the flux, we tested the flux power spectrum of the posterior-predicted fluxes. The results of this test are presented in Fig. 5, showing that the power spectrum of the recovered spectra matches the true flux power spectrum very accurately. We note that we computed the power spectrum of the flux from the noiseless data as the posterior-predicted fluxes do not contain noise.

Fig. 3. *Top panel:* posterior predictive flux for a spectrum with noise $\sigma = 0.03$. The posterior predicted flux is computed from the ensemble mean density field, shown in the lower panel (orange line). The blue shaded region indicates the $1-\sigma$ region, corresponding to the standard deviation of the noise in this line of sight. These tests check whether the data model can accurately account for the observations. Any significant mismatch would immediately indicate a breakdown of the applicability of the data model or error of the inference framework. Our method recovers the transmitted flux fraction correctly, confirming that the data model can accurately account for the observations. *Bottom panel:* comparison of the inferred ensemble mean density field along the line of sight to the ground truth. It can be seen that high-density regions yield a suppression of transmitted flux, while under-dense regions transmit the quasar signal, in agreement with the FGPA model. The density field has been smoothed with a Gaussian smoothing kernel of $\sigma = 0.5 h^{-1}$ Mpc to simulate the difference between dark matter and gas density fields (Peirani et al. 2014). This smoothing kernel is applied to the dark matter density field before generating the mock data.

6.4. Isotropy of the velocity field

One-dimensional Ly- α forest data introduce a particular geometry of the data into the inference problem. In particular, data are only available along one-dimensional lines of sight. In this work, we are interested in recovering the three-dimensional density field from Ly- α forest data by fitting a dynamical model. This model describes the formation of structures by displacing matter from its initial conditions to their final Eulerian positions. As such, we have found, that the large-scale velocity field is particularly sensitive to systematic effects introduced by the data, in particular, the survey geometry. Since the distribution of lines of sight defines a preferred axis, we thus test the isotropy of the velocity field to ensure that the algorithm correctly recov-

ered a physically plausible three-dimensional velocity and corresponding density fields. The results of this test are shown in Fig. 7, showing that there is no preferred component of the three-dimensional velocity field. All velocity components have a similar distribution indicating that the algorithm correctly accounted and corrected for the geometry of the survey. This result is further supported by the recovery of an isotropic primordial power-spectrum as described in Sect. 7.2.

7. Analysing the LSS in Ly- α forest data

In this section, we present the results of applying our algorithm to simulated Ly- α forest data. We show that our method infers physically plausible and unbiased density fields and

Fig. 4. Histogram of the error in the fractional transmitted flux, which is computed as the difference between posterior-predicted fluxes and input spectra. The distribution of flux error matches the distribution of pixel noise in the data, indicating that the method is close to the theoretical optimum. This distribution of error flux can be compared to the one obtained in previous works, demonstrating that our method recovers the fluxes with significantly higher accuracy.

Fig. 5. One-dimensional power spectrum of the flux. The orange line corresponds to the mean flux power spectrum computed from posterior-predicted spectra. The blue-shaded region corresponds to the standard deviation between the different lines of sight. The blue line indicates the flux power spectrum for the mock data. This test shows that the algorithm recovers the flux spectrum inside the $1\text{-}\sigma$ uncertainty limit.

corresponding power-spectra at all scales considered in this work. We further demonstrate that the algorithm accurately extrapolates information into unobserved regions between lines of sight. We further illustrate the performance of the algorithm by recovering mass and velocity profiles of clusters and voids.

7.1. Inference of matter density fields at high-redshift

As discussed above, the BORG algorithm performs a full-scale Bayesian analysis of the cosmological large-scale structure in Ly- α forest data. This is achieved by using the physical forward modelling approach to fit a dynamical model of structure

Fig. 6. Autocorrelation of the parameters as a function of the sample lag in the Markov chain. The correlation length of the sampler can be estimated by determining the point when correlations drop below 0.1 for the first time. These parameters have a correlation length of 830 samples for β and 640 for A .

Fig. 7. Distribution of the three Cartesian components of the velocity. The three distributions have the same mean and variance, indicating that the density field is isotropic as expected.

growth to the data. As a result, we simultaneously obtain inferences of the primordial field of fluctuations from which structures formed, the non-linear spatial matter distribution at a redshift of $z=2.5$ and corresponding velocity fields. More specifically, the BORG algorithm provides us with an ensemble of realizations of these fields drawn from the corresponding large-scale structure posterior distribution discussed in Appendix B.3. This ensemble of realizations from the Markov chain enables us to estimate any desired statistical summary and quantify corresponding uncertainties.

As an example, in Fig. 8, we illustrate ensemble mean and variances for primordial and final density fields. Specifically, Fig. 8 shows slices through the true density, and the ensemble mean and variances of inferred three-dimensional density fields, computed from 12600 samples. A first visual comparison between ground truth and the inferred ensemble mean final density fields in the lower panels of Fig. 8 illustrates that the algorithm correctly recovered the three-dimensional large-scale

Fig. 8. Slices through ground truth initial (left upper panel), final density field (left lower panel), inferred ensemble mean initial (middle upper panel) and ensemble mean final (middle-lower panel) density field computed from 12600 MCMC samples. Comparison between these panels shows that the method recovers the structure of the true density field with high accuracy. We note that the algorithm infers the correct amplitudes of the density field. *Right panels:* standard deviations of inferred amplitudes of initial (upper right panel) and final density fields (lower right panel). The standard deviation of the final density field shows lower values at the position of the lines of sight. We note that the uncertainty of δ^{f} presents a structure that correlates with the density field. Particularly, the variance is higher in high-density regions due to the saturation of the absorbed flux. In contrast, the standard deviation of the initial conditions are homogeneous and show no correlation with the initial density field due to the propagation of information from the final to the initial density field via the dynamical model.

structure from Ly- α forest data. On first sight, both fields are visually almost indistinguishable, indicating the high quality of the inference. The lower right panel of Fig. 8 shows the corresponding standard deviations for density amplitudes for respective volume elements as estimated from the Markov chain. It can be seen that the estimated density standard deviations correlate with the inferred density field. This is expected for a non-linear data model, which couples signal and noise. In particular, one can observe that uncertainties are lowest for under-dense regions where Ly- α forest data provide high signal-to-noise, as quasar light is simply transmitted by structures. In contrast, one can observe the highest uncertainties for over-dense structures. Since over-dense structures absorb quasar light, this decreases the signal in these regions: the absorption is saturated at high densities. Even more, if structures are sufficiently dense, then they absorb all quasar light and the absorption becomes saturated. Once light absorption is saturated, the data provide no further information about the actual density amplitude of the absorber and data only provide information on a minimally lowest density threshold required to explain the observations. Figure 8 therefore, clearly demonstrates that our method correctly accounts for and quantifies uncertainties inherent to Ly- α forest observations. It is interesting to remark, that while observations of galaxy clustering are most informative in high-density regions, the Ly- α forest is particularly informative for under-dense regions and therefore provides complementary information to galaxy surveys. Figure 8 demonstrates that, in addition to inferring the correct filamentary structure, our method clearly infers the correct amplitudes

of the density field. Our results can also be compared to Fig. 2 in Horowitz et al. (2019), which presents a systematic underestimation of the density amplitude.

7.2. Analysing posterior power-spectra

As demonstrated in the previous section, the algorithm provides accurate reconstructions of final density fields. To provide more quantitative tests and to estimate whether inferred density fields are physically plausible, we perform analyses of posterior power-spectra measured from inferred density fields.

Reconstructing three-dimensional density fields from one-dimensional Ly- α data is technically challenging. For example, Kitaura et al. (2012) used a Gibbs sampling approach to sample the large- and small-scales of the density field separately with a lognormal prior for the evolved density field. This approach inferred correct power-spectrum amplitudes at large scales $k < 0.1 h \text{Mpc}^{-1}$ but obtained erroneous excess power at smaller scales, which was attributed to the inadequacy of the log-normal approximation or complex correlations in the data (Kitaura et al. 2012). Horowitz et al. (2019) used an optimization approach to fit a dynamical forward model to the data, but the method obtains power-spectra that severely underestimate the power of density amplitudes.

As we aim to optimally extract cosmologically relevant information from Ly- α forest data, we are particularly interested in recovering physically plausible density fields from observations. Erroneous power in inferred power-spectra is often a sign

Fig. 9. Mean posterior matter power-spectrum. The mean and standard deviation of the matter power spectrum have been computed from 8060 density samples of the Markov chain obtained after the warm-up phase shown in Fig. 2. Although the standard deviation is plotted, it is too small to be seen, showing the stability of the posterior power-spectrum. The dashed line indicates the underlying power spectrum and the 1- and 2- σ cosmic variance limit. The algorithm recovers the power-spectrum amplitudes within the 1- σ cosmic variance uncertainty limit throughout the entire range of Fourier modes considered in this work.

of untreated or unknown survey systematics and indicates a break down of the assumptions underlying the data model or the inference approach. In this work, we use a sophisticated MCMC framework to marginalize out nuisance parameters and quantify uncertainties inherent to the observations. As already indicated in Sect. 6.1, our approach is able to identify the correct power distribution of density amplitudes from observations. To further quantify this result, we estimate the mean and variance of posterior power-spectra measured from the ensemble of Markov samples. The results are presented in Fig. 9. As can be seen, our method recovers the correct fiducial cosmological power-spectrum within the 1- σ cosmic variance uncertainty at all Fourier modes considered in this work. The unbiased recovery of the power-spectrum clearly indicates that the method correctly accounted for uncertainties and systematic effects of the data.

Systematic effects, such as survey geometries, often introduce spurious correlations between power-spectrum modes, when not accounted for properly. This is of particular relevance for Ly- α data, which introduces a particular survey geometry through the set of one-dimensional lines of sight. To test for residual correlations between Fourier modes, we estimated the covariance matrix of power-spectrum amplitudes from our ensemble of Markov samples. As shown in Fig. 10 this covariance matrix shows a clear diagonal structure with the expected correlations at $k \sim 0.5\text{--}0.7 h \text{Mpc}^{-1}$ due to the lines of sight grid. This test shows that our method correctly accounted for survey geometries and other systematic effects that could introduce erroneous mode coupling.

Fig. 10. Estimated correlation matrix of power spectrum amplitudes with the mean value, normalized using the variance of amplitudes of the power spectrum modes. The low off-diagonal contributions are a clear indication that our method correctly accounted and corrected for otherwise erroneous mode coupling, typically introduced by survey systematic effects and uncertainties. The expected correlations correspond to the lines of sight grid.

In summary, these tests demonstrate that our method is capable of inferring physically plausible matter density fields with correct power distribution from noisy Ly- α data.

7.3. Recovering information between lines of sight

The previous section describes that inferred density fields are physically plausible realizations of the matter field underlying the Ly- α observations. Here we want to quantify to what accuracy the method recovers the underlying ground truth density field. A particular challenge is to correctly recover the density field in between one-dimensional lines of sight. As illustrated above, our method is capable of recovering the three-dimensional density field from a set of one-dimensional quasar spectra using the Bayesian physical forward modelling approach. In particular, we determine the performance of our algorithm to recover the density field from data by estimating the Pearson correlation coefficient (see Appendix E) between the inferred and ground truth density fields. We estimate the Pearson correlation coefficient as a function of the position on the x -axis, which permits us to track the correlations on and in-between observed lines of sight. Figure 11 shows the Pearson coefficient averaged over 300 samples of the chain together with its 1- σ credibility interval. It can be clearly seen that inferred density fields correlate with the ground truth density typically to more than 98% along the lines of sight. But even in-between lines of sight the correlation is on average larger than 90%. These high Pearson correlations demonstrate the capability of our method to recover the cosmic large-scale structure in unobserved regions between observed lines of sight.

7.4. Cluster and void profiles

Above we showed that the BORG algorithm recovers physically plausible density fields from Ly- α data that agree with the expected statistical behaviour of a Λ CDM density field. Here we want to test if the algorithm also correctly recovers the properties

Fig. 11. Pearson coefficient of the inferred and true density field as a function of the position across lines of sight. The dashed lines indicate the positions of the lines of sight. The Pearson coefficient is >0.9 at any location in the density field. This, combined with Fig. 8 showing that the algorithm recovers the correct amplitudes, indicates that the algorithm can interpolate the information between lines of sight and correctly recover the structures in unobserved regions.

of individual cosmic structures at their respective locations in the three-dimensional volume. To illustrate this fact, we study the properties of inferred clusters and voids. In particular, we show that the algorithm correctly recovers mass density and velocity profiles of cosmic structures.

To test the properties of an inferred cluster, we randomly chose a peak in the final density field. Then we determined the mass and velocity profiles in spherical shells around the identified centre of the cluster. In particular, we estimated the cumulative radial mass profiles as:

$$M(<r) = \frac{1}{N} \sum_{i=0}^N m_p K_{\text{part}}^i(<r) \quad (7)$$

where r is the radial distance from the cluster centre, i labels the Markov samples, N is the total number of samples and $K_{\text{part}}^i(<r)$ is the number of simulation particles of Markov sample i inside a sphere of radius r around the cluster centre. Finally, $m_p = 287 \times 10^6 M_\odot$ is the particle mass of the simulation, which is obtained as

$$m_p = \frac{\bar{\rho} V_{\text{box}}}{N_{\text{part}}^{\text{total}}} \quad (8)$$

where $\bar{\rho}$ is the mean cosmic density, V_{box} is the volume of the inferred density field and $N_{\text{part}}^{\text{total}}$ is the total number of particles used in the simulation. Figure 12 shows the cluster mass profile measured from the inferred and true density fields. The inferred mass profile shown here is the average over 60 samples. Figure 12 shows that our method provides unbiased mass estimates of cluster mass profiles, becoming an alternative to measuring cluster masses, complementary to weak lensing or X-ray observations.

The method also provides inferred density and velocity profiles of voids. Figure 12 shows the density and velocity profiles measured from the inferred and true density fields. The void is defined spherically and the velocity and density profiles are

spherically averaged:

$$\rho(r) = \frac{1}{N} \sum_{i=0}^N \frac{3m_p K_{\text{part}}^i(<r)}{4\pi r^3} \quad (9)$$

$$v(r) = \frac{1}{N} \sum_{i=0}^N \frac{1}{K_{\text{part}}^i(<r)} \sum_p^{K_{\text{part}}^i(<r)} v_p \quad (10)$$

where i runs over the samples, N is the total number of samples, r is the radius of the sphere centred in the volume element with the lowest density, m_p is the particle mass described in Eq. (8), $K_{\text{part}}^i(<r)$ is the number of particles inside the sphere and v_p are the velocities of the particles provided by the dynamical model. The profiles shown in Fig. 12 are obtained from averaging over 60 samples. Voids constitute the dominant volume fraction of the Universe. Since the effect of matter in voids is mitigated, they are ideal to study the diffuse components of the Universe such as dark energy (Granett et al. 2008; Biswas et al. 2010; Lavaux & Wandelt 2012; Bos et al. 2012) and gravity (Li 2011; Clampitt et al. 2013; Spolyar et al. 2013; Hamaus et al. 2016). Voids are also interesting tools to constrain the neutrino mass (Massara et al. 2015; Kreisch et al. 2019; Sahlén 2019; Schuster et al. 2019); the neutrino free-streaming length falls within the range of typical void size. Therefore, the void profiles obtained with this framework provide an alternative to constrain the neutrino mass.

7.5. Velocity field and structure formation

The dynamical model employed in BORG allows us to naturally infer the velocity field since it derives from the initial perturbations. The velocity information allows discriminating between peculiar velocities and the Hubble flow. The combination of velocity and density fields can provide significant information on the formation of structures and galaxies.

Our method provides velocity fields at $z = 2.5$, where this kind of data is usually hard to obtain. Figure 13 shows the velocity along the line of sight. The line-of-sight velocity provides a tool to measure the kinetic Sunyaev–Zeldovich effect (Sunyaev & Zeldovich 1972, 1980; Vishniac 1987) by cross-correlating the velocity field with the CMB and study the expansion of the Universe.

While the hierarchical structure formation model has been tested in the nearby Universe (Bond et al. 1982; Blumenthal et al. 1984; Davis et al. 1985), our method provides a framework to test it at high redshift. The method provides consistent velocity and density fields, as shown in Fig. 14, which shows a zoom-in on the inferred mean density field and the corresponding velocity components (v_x, v_y). Therefore, we can test the hierarchical structure formation model at high redshift by combining the density and velocity information and testing the predictions of the method with observations. Additionally, these density and velocity fields can provide some insights into galaxy formation and evolution by studying the effect of large-scale structures in galaxy populations (Porqueres et al. 2018). Therefore, this method for the Ly- α forest allows extending investigations of the nature of galaxies or AGN to the high-redshift Universe.

7.6. The astrophysical parameters

Current efforts of the science community aim at constraining the astrophysics of the IGM (see e.g. Lee 2012; Rorai et al. 2017; Eilers et al. 2017) since the thermal history of

Fig. 12. *Top:* cluster mass (*left*) and radial velocity (*right*) profiles from the inferred ensemble mean density field and the simulation. These profiles are computed by averaging over the mass and velocity profiles measured on 60 different density samples. The shaded regions show the standard deviation between these samples. The structured is traversed by 4 lines of sight with a Gaussian noise distribution of $\sigma = 0.03$. The algorithm recovers the correct mass profile from the simulation. This demonstrates that the method provides unbiased mass estimates, correcting for the astrophysical bias of the IGM. *Bottom:* spherically averaged density (*left*) and velocity (*right*) profiles of a void, obtained by averaging the void profile measured in 60 density samples. The standard deviation between the samples is shown as the shaded region. This void is traversed by 6 lines of sight with a Gaussian noise distribution of $\sigma = 0.03$. These plots show that the method recovers the underlying density and velocity profiles of the simulation.

the IGM holds the key to understanding hydrogen reionization at $z > 6$ (Lee 2012). Particularly, analysis of the Ly- α forest attempt to measure the parameters T_0 and γ . The parameter T_0 is related to the spectral shapes of the ultraviolet background. Therefore, it contains information on the intensity of the ionizing background, which is relevant to understand the nature of the first luminous sources and their impact on the subsequent generation of galaxies (Becker et al. 2015). The parameter γ defines the temperature-density relation indicating whether over-dense regions are hotter than the under-dense regions ($\gamma > 1$) or vice versa ($\gamma < 1$). While Becker et al. (2007), Bolton et al. (2008), Viel et al. (2009), Rorai et al. (2017) found evidence for $\gamma < 1$, Calura et al. (2012), Garzilli et al. (2012) found $\gamma \approx 1$. Other approaches (Schaye et al. 2000; Theuns & Zaroubi 2000; McDonald et al. 2001; Lidz et al. 2010; Rudie et al. 2012; Bolton et al. 2014) found $\gamma > 1$, which is more in agreement with the theoretically predicted $\gamma \approx 1.6$ for a post-reionization IGM (Hui & Gnedin 1997). However, the value of γ is still an open debate

since the different values are obtained from different kind of data, affected by systematic uncertainties in a different way (Eilers et al. 2017).

The algorithm presented in this work infers the astrophysical parameters of the FGPA, which are directly related to T_0 and γ . Figure 15 shows the posterior distribution and correlations of these parameters and the noise variance σ , obtained from 6800 samples. This shows that the method recovers the true value of the posterior distribution of the FGPA parameters, corresponding to $T_0 = (206.5 \pm 0.1) 10^3 \Gamma^{-1/0.7} \text{ K}$ and $\gamma = 1.614 \pm 0.004$, where the uncertainty is derived from the standard deviation of A and β . While the more accurate measurements of γ have an uncertainty above 7% (Rudie et al. 2012; Eilers et al. 2017), this test shows that our method can constrain γ with an uncertainty below 1%. The reason for this tight constraint of γ is probably the fact that higher-order statistics of the density field can break parameter degeneracy (Schmidt et al. 2019). Although constraining γ from real Ly- α forest data require to model the continuum flux of the

Fig. 13. Projection of the velocity along lines of sight, which are parallel to the x -axis of the plot.

Fig. 14. Zoom-in on the density field. The vector field shows the velocities on top of the inferred density field, showing matter flowing out of the void and falling into the gravitational potential of the cluster. Therefore, the method provides consistent velocity and density fields.

quasar, our method holds the promise to shed new light on the temperature-density debate.

Note in Fig. 15 that parameters A and β do not show a strong correlation with the noise variance σ , which is jointly sampled. It is also worth noticing that β presents a bimodal distribution. Optimization methods (see e.g. Horowitz et al. 2019) can only explore local extremes of the distribution, becoming sub-optimal to explore multi-modal distributions. However, our MCMC approach can characterize the bimodal distribution and provide the corresponding uncertainties, which are necessary to interpret the data correctly. Therefore, our algorithm provides a framework to constrain the astrophysics of the IGM jointly with the density field. To the knowledge of the authors, this work presents the first approach that attempts to jointly quantify the astrophysical parameters and the 3d cosmic structure.

8. Study with low signal-to-noise data

The previous tests and analysis were performed using simulated data with a noise distribution that was the same for all lines of sight. However, in real data, different lines of sight present different signal-to-noise ratio (S/N), which depend on the integration time of the spectrum. For this reason, this section presents an analysis with mock data with different S/N for each spectrum. The distribution of lines of sight is the same as described in Sect. 5 but the Gaussian pixel-noise is characterized by a different S/N in each line of sight. Therefore, the S/N is constant along one line of sight but varies among lines of sight, following a power-law as described in Krolewski et al. (2018) and Horowitz et al. (2019).

To simulate realistic pixel-noise, we draw S/N values for each line of sight according to

$$\frac{dn_{\text{los}}}{dS/N} \propto S/N^\alpha \quad (11)$$

with $\alpha = 2.7$ (Horowitz et al. 2019). The distribution of S/N is shown in Fig. 17, which can be compared to Fig. 2 in Krolewski et al. (2018). Mimicking the CLAMATO survey, we limit the values of possible S/N to $S/N_{\text{min}} = 1.4$ and $S/N_{\text{max}} = 10$ (Horowitz et al. 2019). The Gaussian pixel-noise is then drawn according to the S/N of each line of sight. The individual σ of the noise distribution for each line of sight is sampled jointly with the astrophysical parameters A and β . Figure 16 shows that the method can recover the individual noise distribution for different spectra.

Figure 18 shows a posterior predicted test for a low S/N spectrum ($S/N = 2$), illustrating the performance of the method in a low S/N regime. This figure shows that the algorithm can predict the input flux and recover the matter density along the line of sight in the $1-\sigma$ uncertainty limits.

Figure 19 shows the probability distribution of the difference between the true density field and the ensemble mean density, demonstrating that the algorithm recovers the unbiased matter density. This test, therefore, shows that our method can be successfully applied to low S/N data, providing the unbiased dark matter density field, jointly inferred with the astrophysical parameters and the characterization of the noise distribution.

9. Summary and discussion

Observations of the late time cosmic large-scale structure can provide information on fundamental physics driving the dynamic evolution of our Universe only if we manage to accurately account for non-linear data models and systematic effects on the data, and if we connect theory to observations. A detailed physical understanding of our universe, therefore, requires linking the current increase of high-quality cosmological data with the development of novel analysis methods capable of making the most of such observations. Traditional analyses of cosmic structures are limited to exploiting only one- and two-point statistics, but ignore significant information encoded in higher-order statistics associated with the filamentary features of the late time cosmic matter distribution. To fully account for the information content of the cosmic large-scale structure, we need to follow a field-based approach to extract the entire three-dimensional matter distribution from observations. In this work we were particularly interested in studying the cosmic matter distributions at redshifts $z > 2$ with one-dimensional Lyman- α forest observations.

Fig. 15. Corner plot for the unknown parameters A and β of the FGPA model and the standard deviation of the Gaussian noise σ . As indicated in the plot, different panels show different marginal posterior distributions for respective parameters. Black solid lines in uni-variate marginal distributions indicate the true fiducial parameter values. It can be seen that the method correctly infers the underlying parameters and quantifies corresponding uncertainties. It is interesting to remark, that the astrophysical parameters A and β show no a posteriori correlations with the noise standard deviation σ .

Fig. 16. Posterior distribution of the noise standard deviations for three lines of sight. The method samples the individual σ of each line of sight, recovering the input value (solid black line).

Fig. 17. Distribution of signal-to-noise (S/N) per pixel in different lines of sight. This distribution of S/N follows a power-law, mimicking the CLAMATO survey (Krolewski et al. 2018). The S/N is constant along each line of sight but varies among different lines of sight.

To extract the full three-dimensional information content from the Ly- α forest, we have proposed a Bayesian physical forward modelling approach to fit a numerical model of gravitational structure formation to data. Building upon the previously presented BORG algorithm, our fully probabilistic framework infers the three-dimensional matter density field and its dynamics from the Lyman- α forest. The method is a considerable improvement over previous approaches (e.g. Kitaura et al. 2012; Horowitz et al. 2019) as it infers unbiased dark matter fields from the Ly- α forest for the first time. While the previous approaches failed to recover the matter-power spectrum at all range of Fourier modes. In particular, the approach followed by Horowitz et al. (2019) produces an underestimation of the power spectrum, corresponding to too low amplitudes of the inferred density field. We have shown that our method recovers the unbiased dark matter distribution and physically plausible density and velocity fields from Lyman- α data.

To make such inferences feasible, we used a likelihood function based on the fluctuating Gunn-Peterson approximation (FGPA). In addition to modelling the dynamical evolution and

Fig. 18. *Top panel:* posterior predictive flux for a spectrum with $S/N = 2$. The posterior predicted flux is computed from the ensemble mean density field, shown in the lower panel (orange line). The blue shaded region indicates the $1-\sigma$ region, corresponding to the standard deviation of the noise in this line of sight. Our method recovers the transmitted flux fraction correctly, confirming that the data model can accurately account for the observations also in the low S/N regime. *Bottom panel:* comparison of the inferred ensemble mean density field along the line of sight to the ground truth. The density field has been smoothed with a Gaussian smoothing kernel of $\sigma = 0.5 h^{-1} \text{ Mpc}$ to simulate the difference between dark matter and gas density fields (Peirani et al. 2014). This smoothing kernel is applied to the dark matter density field before generating the mock data.

spatial distribution of cosmic matter, our approach also accounts for systematic effects arising from the astrophysical properties of the IGM by self-consistently marginalizing out corresponding nuisance parameters of the FGPA model. We have shown that this approach can provide tight constraints on the astrophysical parameters T_0 and γ , which provide significant information about the reionization history of the Universe and the early luminous sources.

Our hierarchical Bayesian framework encodes non-linear dynamical models. This work is based on the Lagrangian Perturbation Theory (LPT). As demonstrated by hydrodynamical simulations, the Ly- α forest arises from regions where the density of matter is within a factor ten of the cosmic mean density, and thus is still close to the linear regime of gravitational collapse (Peirani et al. 2014; Sorini et al. 2016).

We tested our approach using realistic mock observations emulating the CLAMATO survey to infer the dark matter density and the velocity field at redshift $z = 2.5$ with a resolution of $1 h^{-1} \text{ Mpc}$. These tests demonstrate that our method recovers unbiased reconstructions of the non-linear spatial matter dis-

tribution and its power-spectrum from observations of the Ly- α forest. The inferred mean density field presents a high correlation with the true density at any location. This indicates that our method can interpolate the information between lines of sight. While previous approaches tried to address this challenge with polynomial smoothing (Cisewski et al. 2014) or Wiener filtering (Ozbek et al. 2016; Stark et al. 2015a), the dynamical model in our method interpolates the information by accounting for the high-order statistics of the density field.

The dynamical model naturally infers velocity fields jointly with the density field. Therefore, our method provides velocity fields at $z > 2$, where this kind of data is usually hard to obtain. From the velocity fields, we can derive the velocities along the lines of sight or radial velocities. These velocity fields can be cross-correlated with the CMB to detect the kinetic Sunyaev–Zeldovich effect. Additionally, the consistent velocity and density fields provide a framework to test hierarchical structure formation and galaxy formation models at high redshift.

We have shown that the method provides accurate mass and velocity profiles for cosmic structures, such as voids and

Fig. 19. Probability distribution of the difference between the true density field and the ensemble mean density. The algorithm correctly recovers the unbiased density field from the Ly- α forest data with low S/N and individual noise distribution for each spectra.

clusters. Therefore, this method provides an alternative to measuring cluster masses, complementary to X-ray and lensing measurements. The method also provides the velocity and density profile of voids. Since the non-linear effects are mitigated in voids, they are sensitive to the diffuse components of the Universe such as dark energy. Therefore, the velocity and density profiles of voids can be used to discriminate between homogeneous dark energy and modifications of gravity.

Since the Ly- α probes the matter distribution down to few megaparsecs, it is sensitive to neutrino masses. This high resolution and the void profiles provided by our algorithm can constrain neutrino masses. Additionally, this high resolution of the Ly- α forest allows us to infer the density field at the 1 Mpc scale. By applying new techniques compatible with [Kodi Ramanah et al. \(2019\)](#), this high-resolution density field can provide tight constraints of the cosmological parameters by studying structure geometry.

Summarizing, our method clearly demonstrates the feasibility of detailed and physically plausible inferences of three-dimensional large-scale structures at high redshift from the Ly- α forest observations. The proposed approach, therefore, opens a new window to study cosmology and structure formation at high² redshifts³.

Acknowledgements. This work has made use of the C2PAP cluster at the Leibniz-Rechenzentrum. NP is supported by the DFG cluster of excellence “Origin and Structure of the Universe”. GL acknowledges financial support from the ILP LABEX (under reference ANR-10-LABX-63) and from “Programme National de Cosmologie et Galaxies” (PNCG) of CNRS/INSU, France, and the ANR BIG4, under reference ANR-16-CE23-0002. This work was carried out within the Aquila Consortium.

References

- Abbott, T. M. C., Abdalla, F. B., Alarcon, A., et al. 2018, *Phys. Rev. D*, **98**, 043526
- Alam, S., Ho, S., & Silvestri, A. 2016, *MNRAS*, **456**, 3743
- Bahcall, J. N., Bergeron, J., Boksenberg, A., et al. 1993, *ApJS*, **87**, 1
- Bahcall, J. N., Bergeron, J., Boksenberg, A., et al. 1996, *ApJ*, **457**, 19
- Basilakos, S., & Nesseris, S. 2016, *Phys. Rev. D*, **94**, 123525

² www.universe-cluster.de

³ <https://aquila-consortium.org>

- Bautista, J. E., Busca, N. G., Guy, J., et al. 2017, *A&A*, **603**, A12
- Becker, G. D., Rauch, M., & Sargent, W. L. W. 2007, *ApJ*, **662**, 72
- Becker, G. D., Bolton, J. S., & Lidz, A. 2015, *PASA*, **32**, e045
- Bird, S., Peiris, H. V., Viel, M., & Verde, L. 2011, *MNRAS*, **413**, 1717
- Biswas, R., Alizadeh, E., & Wandelt, B. D. 2010, *Phys. Rev. D*, **82**, 023002
- Blomqvist, M., du Mas des Bourboux, H., Busca, N. G., et al. 2019, *A&A*, **629**, A86
- Blumenthal, G. R., Faber, S. M., Primack, J. R., & Rees, M. J. 1984, *Nature*, **311**, 517
- Boera, E., Becker, G. D., Bolton, J. S., & Nasir, F. 2019, *ApJ*, **872**, 101
- Bolton, J. S., Viel, M., Kim, T. S., Haehnelt, M. G., & Carswell, R. F. 2008, *MNRAS*, **386**, 1131
- Bolton, J. S., Becker, G. D., Haehnelt, M. G., & Viel, M. 2014, *MNRAS*, **438**, 2499
- Bond, J. R., Szalay, A. S., & Turner, M. S. 1982, *Phys. Rev. Lett.*, **48**, 1636
- Bos, E. G. P., van de Weygaert, R., Dolag, K., & Pettorino, V. 2012, *MNRAS*, **426**, 440
- Boyle, A., & Komatsu, E. 2018, *JCAP*, **2018**, 035
- Busca, N. G., Delubac, T., Rich, J., et al. 2013, *A&A*, **552**, A96
- Calura, F., Tescari, E., D’Odorico, V., et al. 2012, *MNRAS*, **422**, 3019
- Cisewski, J., Croft, R. A. C., Freeman, P. E., et al. 2014, *MNRAS*, **440**, 2599
- Clampitt, J., Cai, Y.-C., & Li, B. 2013, *MNRAS*, **431**, 749
- Colombi, S., Dodelson, S., & Widrow, L. M. 1996, *ApJ*, **458**, 1
- Croft, R. A. C., Weinberg, D. H., Katz, N., & Hernquist, L. 1998, in *Large Scale Structure: Tracks and Traces*, eds. V. Mueller, S. Gottloeber, J. P. Muecket, & J. Wambsganss, 69
- Davé, R. 2001, in *Astrophysical Ages and Times Scales*, eds. T. von Hippel, C. Simpson, & N. Manset, *ASP Conf. Ser.*, **245**, 561
- Davé, R., Hernquist, L., Katz, N., & Weinberg, D. H. 1999, *ApJ*, **511**, 521
- Davis, M., Efstathiou, G., Frenk, C. S., & White, S. D. M. 1985, *ApJ*, **292**, 371
- Dawson, K. S., Schlegel, D. J., Ahn, C. P., et al. 2013, *AJ*, **145**, 10
- Delubac, T., Bautista, J. E., Busca, N. G., et al. 2015, *A&A*, **574**, A59
- DESI Collaboration 2016, ArXiv e-prints [arXiv:1611.00036]
- Duane, S., Kennedy, A. D., Pendleton, B. J., & Roweth, D. 1987, *Phys. Lett. B*, **195**, 216
- du Mas des Bourboux, H., Le Goff, J. M., Blomqvist, M., et al. 2017, *A&A*, **608**, A130
- Eilers, A.-C., Hennawi, J. F., & Lee, K.-G. 2017, *ApJ*, **844**, 136
- Eisenstein, D. J., & Hu, W. 1998, *ApJ*, **496**, 605
- Eisenstein, D. J., & Hu, W. 1999, *ApJ*, **511**, 5
- Eisenstein, D. J., Weinberg, D. H., Agol, E., et al. 2011, *ApJ*, **142**, 72
- Frenk, C. S., & White, S. D. M. 2012, *Ann. Phys.*, **524**, 507
- Gallerani, S., Kitaura, F. S., & Ferrara, A. 2011, *MNRAS*, **413**, L6
- Garzilli, A., Bolton, J. S., Kim, T. S., Leach, S., & Viel, M. 2012, *MNRAS*, **424**, 1723
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. 2004, *Bayesian Data Analysis*, 2nd edn. (Chapman and Hall/CRC)
- Granett, B. R., Neyrinck, M. C., & Szapudi, I. 2008, *ApJ*, **683**, L99
- Gunn, J. E., & Peterson, B. A. 1965, *ApJ*, **142**, 1633
- Hamaus, N., Pisani, A., Sutter, P. M., et al. 2016, *Phys. Rev. Lett.*, **117**, 091302
- He, S., Alam, S., Ferraro, S., Chen, Y.-C., & Ho, S. 2018, *Nat. Astron.*, **2**, 401
- Horowitz, B., Lee, K. G., White, M., Krolewski, A., & Ata, M. 2019, ArXiv e-prints [arXiv:1903.09049]
- Hui, L., & Gnedin, N. Y. 1997, *MNRAS*, **292**, 27
- Huterer, D., Kirkby, D., Bean, R., et al. 2015, *Astropart. Phys.*, **63**, 23
- Jasche, J., & Kitaura, F. S. 2010, *MNRAS*, **407**, 29
- Jasche, J., & Lavaux, G. 2017, *A&A*, **606**, A37
- Jasche, J., & Lavaux, G. 2019, *A&A*, **625**, A64
- Jasche, J., & Wandelt, B. D. 2013, *MNRAS*, **432**, 894
- Jasche, J., Leclercq, F., & Wandelt, B. D. 2015, *JCAP*, **1**, 036
- Kitaura, F.-S., Gallerani, S., & Ferrara, A. 2012, *MNRAS*, **420**, 61
- Kodi Ramanah, D., Lavaux, G., Jasche, J., & Wandelt, B. D. 2019, *A&A*, **621**, A69
- Köhlinger, F., Viola, M., Joachimi, B., et al. 2017, *MNRAS*, **471**, 4412
- Kreisch, C. D., Pisani, A., Carbone, C., et al. 2019, *MNRAS*, **488**, 4413
- Krolewski, A., Lee, K.-G., White, M., et al. 2018, *ApJ*, **861**, 60
- Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, ArXiv e-prints [arXiv:1110.3193]
- Lavaux, G., & Jasche, J. 2016, *MNRAS*, **455**, 3169
- Lavaux, G., & Wandelt, B. D. 2012, *ApJ*, **754**, 109
- Lee, K.-G. 2012, *ApJ*, **753**, 136
- Lee, K. G. 2016, in *The Zeldovich Universe: Genesis and Growth of the Cosmic Web*, eds. R. van de Weygaert, S. Shandarin, E. Saar, & J. Einasto, *IAU Symp.*, **308**, 360
- Lee, K.-G., Krolewski, A., White, M., et al. 2018, *ApJS*, **237**, 2
- Levi, M., Bebek, C., Beers, T., et al. 2013, ArXiv e-prints [arXiv:1308.0847]
- Li, B. 2011, *MNRAS*, **411**, 2615
- Lidz, A., Faucher-Giguère, C.-A., Dall’Aglio, A., et al. 2010, *ApJ*, **718**, 199

- LSST Science Collaboration (Abell, P. A., et al.) 2009, ArXiv e-prints [arXiv:[0912.0201](#)]
- Massara, E., Villaescusa-Navarro, F., Viel, M., & Sutter, P. M. 2015, *JCAP*, **2015**, 018
- McConnachie, A., Babusiaux, C., Balogh, M., et al. 2016, ArXiv e-prints [arXiv:[1606.00043](#)]
- McDonald, P., Miralda-Escudé, J., Rauch, M., et al. 2001, *ApJ*, **562**, 52
- Mishra-Sharma, S., Alonso, D., & Dunkley, J. 2018, *Phys. Rev. D*, **97**, 123544
- Mo, H., van den Bosch, F. C., & White, S. 2010, *Galaxy Formation and Evolution* (Cambridge, UK: Cambridge University Press)
- Myers, A. D., Palanque-Delabrouille, N., Prakash, A., et al. 2015, *ApJS*, **221**, 27
- Nasir, F., Bolton, J. S., & Becker, G. D. 2016, *MNRAS*, **463**, 2335
- Ozbek, M., Croft, R. A. C., & Khandai, N. 2016, *MNRAS*, **456**, 3610
- Palanque-Delabrouille, N., Yèche, C., Baur, J., et al. 2015, *JCAP*, **2015**, 011
- Peirani, S., Weinberg, D. H., Colombi, S., et al. 2014, *ApJ*, **784**, 11
- Planck Collaboration XIII. 2016, *A&A*, **594**, A13
- Planck Collaboration XXIV. 2016, *A&A*, **594**, A24
- Planck Collaboration IX. 2019, A&A, submitted [arXiv:[1905.05697](#)]
- Porqueres, N., Jasche, J., Enßlin, T. A., & Lavaux, G. 2018, *A&A*, **612**, A31
- Porqueres, N., Kodi Ramanah, D., Jasche, J., & Lavaux, G. 2019, *A&A*, **624**, A115
- Refregier, A., Amara, A., Kitching, T. D., et al. 2010, ArXiv e-prints [arXiv:[1001.0061](#)]
- Riess, A. G., Macri, L. M., Hoffmann, S. L., et al. 2016, *ApJ*, **826**, 56
- Riess, A. G., Casertano, S., Yuan, W., et al. 2018, *ApJ*, **861**, 126
- Rorai, A., Becker, G. D., Haehnelt, M. G., et al. 2017, *MNRAS*, **466**, 2690
- Rossi, G. 2014, ArXiv e-prints [arXiv:[1406.5411](#)]
- Rossi, G. 2017, *ApJS*, **233**, 12
- Rossi, G., Yèche, C., Palanque-Delabrouille, N., & Lesgourgues, J. 2015, *Phys. Rev. D*, **92**, 063505
- Rudie, G. C., Steidel, C. C., & Pettini, M. 2012, *ApJ*, **757**, L30
- Rusu, C. E., Wong, K. C., Bonvin, V., et al. 2019, MNRAS, submitted [arXiv:[1905.09338](#)]
- Sahlén, M. 2019, *Phys. Rev. D*, **99**, 063525
- Schaye, J., Theuns, T., Rauch, M., Efstathiou, G., & Sargent, W. L. W. 2000, *MNRAS*, **318**, 817
- Schmidt, F., Elsner, F., Jasche, J., Nguyen, N. M., & Lavaux, G. 2019, *JCAP*, **2019**, 042
- Schuster, N., Hamaus, N., Pisani, A., et al. 2019, ArXiv e-prints [arXiv:[1905.00436](#)]
- Seljak, U., Slosar, A., & McDonald, P. 2006, *JCAP*, **2006**, 014
- Slosar, A., Font-Ribera, A., Pieri, M. M., et al. 2011, *JCAP*, **2011**, 001
- Sorini, D., Oñorbe, J., Lukić, Z., & Hennawi, J. F. 2016, *ApJ*, **827**, 97
- Spolyar, D., Sahlén, M., & Silk, J. 2013, *Phys. Rev. Lett.*, **111**, 241103
- Stark, C. W., White, M., Lee, K.-G., & Hennawi, J. F. 2015a, *MNRAS*, **453**, 311
- Stark, C. W., Font-Ribera, A., White, M., & Lee, K.-G. 2015b, *MNRAS*, **453**, 4311
- Sunyaev, R. A., & Zeldovich, Y. B. 1972, *Comm. Astrophys. Space Phys.*, **4**, 173
- Sunyaev, R. A., & Zeldovich, I. B. 1980, *MNRAS*, **190**, 413
- Theuns, T., & Zaroubi, S. 2000, *MNRAS*, **317**, 989
- Viel, M., Haehnelt, M. G., & Lewis, A. 2006, *MNRAS*, **370**, L51
- Viel, M., Bolton, J. S., & Haehnelt, M. G. 2009, *MNRAS*, **399**, L39
- Viel, M., Becker, G. D., Bolton, J. S., & Haehnelt, M. G. 2013, *Phys. Rev. D*, **88**, 043502
- Vishniac, E. T. 1987, *ApJ*, **322**, 597
- Weymann, R. J., Jannuzi, B. T., Lu, L., et al. 1998, *ApJ*, **506**, 1
- Williger, G. M., Carswell, R. F., Weymann, R. J., et al. 2010, *MNRAS*, **405**, 1736
- Yèche, C., Palanque-Delabrouille, N., Baur, J., & du Mas des Bourboux, H. 2017, *JCAP*, **2017**, 047

Appendix A: Projector along the line of sight

The Ly- α forest is measured along lines of sight to background sources. The volume elements intersected by the line of sight are identified from the wavelengths in the spectrum: first, the redshift is computed from the wavelength

$$z = \frac{\lambda - 1}{\lambda_0} \quad (\text{A.1})$$

where $\lambda_0 = 1216 \text{ \AA}$ corresponding to the wavelength of the Ly- α transition. The redshift is then used to compute the comoving distance from tabulated values of $d_{\text{com}}(z)$ and finally, this is transformed into the regular grid coordinates as

$$x = d_{\text{com}}(z) \cos(\xi) \cos(\alpha), \quad (\text{A.2})$$

$$y = d_{\text{com}}(z) \cos(\xi) \sin(\alpha), \quad (\text{A.3})$$

$$z = d_{\text{com}}(z) \sin(\xi), \quad (\text{A.4})$$

with ξ being the declination and α being the right ascension of the background source.

Appendix B: Method

In this section, we present a more detailed description of our algorithm. We describe the data model, likelihood, and prior. We briefly discuss how the high-dimensional joint posterior distribution can be explored via a multiple-block sampling scheme.

B.1. Prior

BORG includes a physical density prior involving a model for structure formation. Therefore, it translates the inference of the density field into the inference of the initial conditions, which obey Gaussian statistics. The link between the initial and evolved density fields is given by the deterministic structure formation model. This makes the prior for the final density contrast highly non-Gaussian and non-linear.

In Fourier space, the prior for the initial density contrast δ^{ic} is a multivariate Gaussian with zero mean and diagonal covariance matrix $\hat{S}$ corresponding to the initial cosmological power spectrum

$$P(\hat{\delta}^{\text{ic}}|\hat{S}) = \frac{1}{\sqrt{|2\pi\hat{S}|}} \exp\left(-\frac{1}{2} \sum_{k,k'} \hat{\delta}_k^{\text{ic}} \hat{S}_{kk'}^{-1} \hat{\delta}_{k'}^{\text{ic}}\right) \quad (\text{B.1})$$

where the hat denotes Fourier-space. The elements in the matrix $\hat{S}$ are fixed parameters that characterize the variance of the initial density field. This covariance matrix is diagonal in Fourier space due to the statistical homogeneity of the initial density field. The diagonal coefficients are related to the initial power spectrum $P(k)$ as

$$\hat{S}_{kk}^{-1} = \sqrt{\frac{P(k)}{(2\pi)^{3/2}}}. \quad (\text{B.2})$$

The power spectrum is chosen to follow the prescription of (Eisenstein & Hu 1998, 1999), including BAO wiggles.

The prior suggests physically reasonable density fields from the space of all possible states. However, the prior does not limit the space of initial conditions that is explored to match the data. If unlikely events are required to explain the data, the method explores prior regions that are unlikely.

The next step is to evolve the initial density field to the final density. We obtain the prior for the final density contrast δ^{f} at a scale factor a by using conditional probabilities:

$$P(\delta^{\text{f}}) = \int P(\delta^{\text{f}}, \delta^{\text{ic}}) d\delta^{\text{ic}} \quad (\text{B.3})$$

$$= \int P(\delta^{\text{f}}|\delta^{\text{ic}}) P(\delta^{\text{ic}}) d\delta^{\text{ic}}. \quad (\text{B.4})$$

Since gravity is deterministic, the conditional probability is a Dirac delta of the structure formation model $M(a, \delta^{\text{ic}})$:

$$P(\delta^{\text{f}}|\delta^{\text{ic}}) = \prod_x \delta^D(\delta_x^{\text{f}} - M(a, \delta^{\text{ic}})_x). \quad (\text{B.5})$$

Once the gravitational model is defined, a prior distribution for the evolved density field can be obtained in two steps: first, a realization of the initial conditions is drawn from the Gaussian prior. It is then evolved forward in time with the structure formation model. This procedure provides samples from the joint prior distribution of the initial and evolved density fields:

$$P(\delta^{\text{f}}, \delta^{\text{ic}}) = P(\delta^{\text{ic}}) \prod_x \delta^D(\delta_x^{\text{f}} - M(a, \delta^{\text{ic}})_x). \quad (\text{B.6})$$

where x labels the volume elements. The evolved density field is smoothed using a Gaussian kernel with a smoothing length of $\sigma = 0.5 h^{-1} \text{ Mpc}$.

B.2. Likelihood

We developed a likelihood based on the Fluctuating Gunn-Peterson Approximation (FGPA, Gunn & Peterson 1965):

$$F_x = \exp\left[-A(1 + \delta_x)^\beta\right] \quad (\text{B.7})$$

where the index x labels the volume elements, F is the transmitted flux fraction, δ is the density contrast and A and β are astrophysical parameters that are related to the neutral hydrogen.

It is common to assume that the noise in the data is Gaussian (Bird et al. 2011; Cisewski et al. 2014; Lee et al. 2018; Ozbek et al. 2016; Horowitz et al. 2019). Labelling the lines of sight with an index n , the likelihood reads

$$P(\delta^{\text{ic}}, \delta^{\text{f}}|F) = \prod_{n,x} \frac{1}{\sqrt{2\pi\sigma^2}} \times \exp\left[-\frac{\left((F_n)_x - \exp[-A(1 + \delta_x)^\beta]\right)^2}{2\sigma^2}\right] \quad (\text{B.8})$$

where n labels the lines of sight, x runs over the volume elements intersected by the n th line of sight and F is the observed flux. The constants A and β are astrophysical parameters that depend on the physics of the neutral hydrogen that produces the absorption lines. Although these parameters are interesting on their own because they can provide information about the IGM, they are not interesting for our cosmological analysis and we marginalize over them.

B.3. Posterior

The posterior distribution of the density field can be written as

$$P(\delta^{\text{f}}|F) = \frac{P(F|\delta^{\text{f}})P(\delta^{\text{f}})}{P(F)}, \quad (\text{B.9})$$

where $P(F)$ is the normalization. In terms of the initial density field, the posterior can be written as

$$P(\delta^f, \delta^{ic}|F, S) = \frac{P(F|\delta^f, \delta^{ic})P(\delta^f, \delta^{ic}|S)}{P(F)} \quad (\text{B.10})$$

$$= P(F|\delta^f) \frac{P(\delta^{ic}|S)}{P(F)} \prod_x \delta^D[\delta_x^f - M(a, \delta^{ic})_x]$$

where x labels the volume elements.

By marginalizing over the final density field, we obtain the posterior that links the primordial density field with the observed flux in the Ly- α forest:

$$P(\delta^{ic}|F, S) = P(F|M(a, \delta^{ic})) \frac{P(\delta^{ic}|S)}{P(F)} \quad (\text{B.11})$$

$$= P(F|\delta^{ic}) \frac{P(\delta^{ic}|S)}{P(F)} \quad (\text{B.12})$$

B.4. Modules

We made use of the statistical modular framework of BORG (Jasche & Wandelt 2013). This approach allows us to model complex data models as a Bayesian hierarchical problem.

Since it is not efficient to sample from the high-dimensional posterior (B.12), BORG makes use of the Metropolis-Hastings theorem which allows us to break the high-dimensional sampling into the sampling of the density field δ^{ic} and the rest of the inferred quantities θ_i separately:

$$(\delta^{ic})^{s+1} \sim P(\delta^{ic}|\theta^s, F) \quad (\text{B.13})$$

$$\theta^{s+1} \sim P(\theta^s | (\delta^{ic})^s, F) \quad (\text{B.14})$$

where s indicates the sample number, l labels the volume elements and θ are the astrophysical parameters of the IGM, which are sampled using a slice sampling scheme. We obtained samples of the high-dimensional posterior distribution using a Markov chain Monte Carlo in this block sampling scheme.

Appendix C: Hamiltonian Monte Carlo sampling

The inference of the density field requires inferring the amplitudes of the primordial density field at different volume elements of a regular grid, commonly between 128^3 and 256^3 volume elements. This implies 10^6 to 10^7 free parameters. Algorithms that perform a random walk, like the Metropolis-Hastings, fail in this high-dimensional parameter space. However, the Hamiltonian Monte Carlo (HMC) method provides an efficient sampling method in high-dimensional non-linear parameter spaces since it uses the information in the gradients to avoid random walks. In this section, we briefly review the HMC framework used for this work. A more detailed review, as well as its application to the large-scale structure analysis, is presented in Jasche & Kitaura (2010) and Jasche & Wandelt (2013).

C.1. Algorithm

Suppose that we want to generate samples from a probability distribution $P(\mathbf{x})$. Then the HMC scheme interprets the negative logarithm of the distribution as a potential $\psi(\mathbf{x}) = -\ln(P(\mathbf{x}))$ (Duane et al. 1987).

We construct a Hamiltonian with this potential $\psi(x)$ describing the dynamics of the phase space. In order to add a kinetic term, we introduce a set of momenta p_i and a mass matrix M as

$$H(\mathbf{x}, \mathbf{p}) = \sum_i \sum_j \frac{1}{2} p_i M_{ij}^{-1} p_j + \psi(x_i). \quad (\text{C.1})$$

The mass matrix M is positive definite and affects the performance of the sampler. These masses characterize the inertia of the parameters when moving through the parameter space.

Then we can obtain the probability distribution as

$$e^{-H} = P(\mathbf{x}) \exp\left(\sum_i \sum_j \frac{1}{2} p_i M_{ij}^{-1} p_j\right). \quad (\text{C.2})$$

This ensures that the joint distribution can be split into a Gaussian distribution for the momenta p_i and the distribution $P(\mathbf{x})$ to generate the samples. Therefore, the sets $\mathbf{p}$ and $\mathbf{x}$ are independent and the auxiliary momenta can be marginalized over.

In order to generate samples from the joint distribution, we first need to draw a set of momenta from the distribution defined by the kinematic energy, which is a Gaussian distribution with covariance $\mathbf{M}$. The system evolves deterministically from an initial point in the high-dimensional parameter space, for a pseudo-time τ , following the Hamilton equations:

$$\frac{dx_i}{dt} = \frac{\partial H(\mathbf{x}, \mathbf{p})}{\partial p_i} \quad (\text{C.3})$$

$$\frac{dp_i}{dt} = -\frac{\partial H(\mathbf{x}, \mathbf{p})}{\partial x_i} = -\frac{\partial \psi(\mathbf{x})}{\partial x_i}. \quad (\text{C.4})$$

The new position $(\mathbf{x}', \mathbf{p}')$ in the phase space is found by integrating the Hamilton equations. Then this new point is accepted or rejected according to the original Metropolis-Hastings acceptance rule:

$$P_A = \min\{1, \exp[-H(\mathbf{x}', \mathbf{p}') - H(\mathbf{x}, \mathbf{p})]\}. \quad (\text{C.5})$$

The Hamiltonian equations guarantees a unity acceptance rate since they ensure energy conservation. However, numerical inaccuracies in the integration scheme result in smaller acceptance rates. After accepting a new sample, the algorithm discards the momenta and restarts the sampling scheme by drawing a new set of momenta, which ensures the exploration of the system. Summarizing, the HMC sampling scheme is done in two steps: the first step performs a Gibbs sampling to get a set of Gaussian distributed momenta and the second step computes the deterministic dynamical trajectory on the posterior distribution.

Appendix D: Hamilton equations for the LSS

The HMC exploits the gradient of the logarithmic posterior distribution to optimally explore the parameter space. Therefore, the algorithm also requires the derivatives of the posterior distribution.

We can derive the potential in the Hamilton equations from Eqs. (B.1) and (B.9)

$$\psi = -\ln[P(\delta^{ic}|F)] \quad (\text{D.1})$$

$$= \psi_{\text{likelihood}}(\delta^{ic}) + \psi_{\text{prior}}(\delta^{ic}), \quad (\text{D.2})$$

where F is the data and

$$\psi_{\text{prior}}(\delta^{\text{ic}}) = \frac{1}{2} \sum_{xy} \delta_x^{\text{ic}} S_{xy}^{-1} \delta_y^{\text{ic}}, \quad (\text{D.3})$$

$$\psi_{\text{likelihood}}(\delta^{\text{ic}}) = \sum_n \sum_x \frac{\left((F_n)_x - \exp[A(1 + M(a, \delta^{\text{ic}})_x)^\beta] \right)^2}{2\sigma^2} + \frac{1}{2N} \sum_n \ln(2\pi\sigma^2)^N \quad (\text{D.4})$$

where n runs over the lines of sight, x and y label the volume elements and N is the number of voxels of the n th line of sight.

The forces are computed by differentiating with respect to δ^{ic}

$$\frac{\partial \psi(\delta^{\text{ic}})}{\partial \delta_x^{\text{ic}}} = \frac{\partial \psi_{\text{prior}}(\delta^{\text{ic}})}{\partial \delta_x^{\text{ic}}} + \frac{\partial \psi_{\text{likelihood}}(\delta^{\text{ic}})}{\partial \delta_x^{\text{ic}}}. \quad (\text{D.5})$$

The prior gradient is

$$\frac{\partial \psi_{\text{prior}}(\delta^{\text{ic}})}{\partial \delta_x^{\text{ic}}} = \sum_y S_{xy}^{-1} \delta_y^{\text{ic}} \quad (\text{D.6})$$

and the corresponding likelihood term

$$\begin{aligned} \frac{\partial \psi_{\text{likelihood}}(\delta^{\text{ic}})}{\partial \delta_x^{\text{ic}}} &= \sum_n \frac{(F_n)_x - \exp[-A(1 + M(a, \delta^{\text{ic}})_x)^\beta]}{\sigma^2} \\ &\quad \times A\beta(1 + M(a, \delta^{\text{ic}})_x)^{\beta-1} \\ &\quad \times \exp\left[-A(1 + M(a, \delta^{\text{ic}})_x)^\beta\right] \times \frac{\partial M(a, \delta^{\text{ic}})_x}{\partial \delta_x^{\text{ic}}}. \end{aligned} \quad (\text{D.7})$$

The parameters of the likelihood A , β and the noise variance σ are sampled using a slice sampler. The equations of motion are integrated using a leapfrog integration scheme [Duane et al. \(1987\)](#), which is exactly reversible and symplectic, maintaining the detailed balance of the Markov chain.

Appendix E: Pearson's correlation coefficient

The Pearson's correlation coefficient is a measurement of the linear correlation between two variables. It has a value between 1 and -1 , with 1 indicating linear positive correlation and -1 anti-correlation. The Pearson's coefficient is defined as

$$r_{xy} = \frac{\sum_i (x_i - \langle x \rangle)(y_i - \langle y \rangle)}{\left[\sum_i (x_i - \langle x \rangle)^2 \right]^{1/2} \left[\sum_i (y_i - \langle y \rangle)^2 \right]^{1/2}} \quad (\text{E.1})$$

where i labels the samples and $\langle x \rangle$ indicates the sample mean of x .