A hierarchical field-level inference approach to reconstruction from sparse Lyman- α forest data

Natalia Porqueres¹, Oliver Hahn², Jens Jasche³, and Guilhem Lavaux⁴

¹ Imperial Centre for Inference and Cosmology, Imperial College London, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, UK

e-mail: n.porqueres@imperial.ac.uk

² Laboratoire Lagrange, Université Côte d’Azur, Observatoire de la Côte d’Azur, CNRS, Bvd de l’Observatoire, CS 34229, 06304 Nice, France

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

⁴ CNRS & Sorbonne Université, UMR7095, Institut d’Astrophysique de Paris, 75014 Paris, France

Received 23 May 2020 / Accepted 14 August 2020

ABSTRACT

We address the problem of inferring the three-dimensional matter distribution from a sparse set of one-dimensional quasar absorption spectra of the Lyman- α forest. Using a Bayesian forward modelling approach, we focus on extending the dynamical model to a fully self-consistent hierarchical field-level prediction of redshift-space quasar absorption sightlines. Our field-level approach rests on a recently developed semiclassical analogue to Lagrangian perturbation theory (LPT), which improves over noise problems and interpolation requirements of LPT. It furthermore allows for a manifestly conservative mapping of the optical depth to redshift space. In addition, this new dynamical model naturally introduces a coarse-graining scale, which we exploited to accelerate the Markov chain Monte-Carlo (MCMC) sampler using simulated annealing. By gradually reducing the effective temperature of the forward model, we were able to allow it to first converge on large spatial scales before the sampler became sensitive to the increasingly larger space of smaller scales. We demonstrate the advantages, in terms of speed and noise properties, of this field-level approach over using LPT as a forward model, and, using mock data, we validated its performance to reconstruct three-dimensional primordial perturbations and matter distribution from sparse quasar sightlines.

Key words. large-scale structure of Universe – dark matter – methods: statistical – methods: data analysis

1. Introduction

A fundamental task in cosmology consists of relating the structures we see in the late Universe with the primordial density fluctuations when the Universe was still close to homogeneous. While the primordial density fluctuations are well-described by a Gaussian distribution (Planck Collaboration IX 2020), gravitational collapse in the cosmological context leads to intricate structures with a large density contrast over the age of the Universe. Non-linear dynamics produce a matter density field with a highly non-Gaussian complex statistical structure, which makes the analysis of the late-time Universe very challenging. A detailed modelling of the cosmic matter distribution would require describing the high-order statistics corresponding to the filamentary structure of the cosmic web. At present, a closed-form description of the non-linear density field in terms of a high-dimensional multivariate probability distribution does not exist. Although there are approximations to reproduce the statistical behaviour of the dark matter density field (e.g. log-normal distribution or multivariate Gaussians, Lahav et al. 1994; Zaroubi et al. 1999; Kitaura & Enßlin 2008; Kitaura et al. 2009; Jasche & Kitaura 2010), they only parameterise the one- and two-point statistics and fail to reproduce more complex structures such as filaments (Baugh et al. 1995; Peacock & Dodds 1996; Smith et al. 2003).

Here, we instead follow the forward modelling approach to relate initial conditions and observables. It consists of a data

model that describes how the continuous three-dimensional field of initial matter fluctuations affects a set of predicted observables, which are then compared to data. In this work, we specifically focus on the observed flux in quasar absorption spectra as our observable. We are then interested in the inverse problem: given a set of quasar spectra, we want to infer the underlying matter distribution and the corresponding primordial fluctuations. In this context, the data model should describe everything that may happen between the initial fluctuations and the observation of the spectra, which includes the time-evolution of the matter density and the cosmic structure formation model, as well as sparse sampling of the observable. In this way, every data point is used, rather than relying on summary statistics that do not capture all the information and whose distributions are not well known.

Structure formation in a cold dark matter (Λ CDM) cosmology proceeds through the collapse of baryons and dark matter that can be well approximated as cold in comparison to the velocities induced by gravity. The dynamics and growth of cosmic structures are then described by Lagrangian perturbation theory (LPT; Zel’dovich 1970; Bouchet et al. 1992) well, which directly describes the motion of fluid elements. LPT is, however, only valid before the crossing of fluid trajectories and, therefore, restricted to large scales or early times. Due to its simplicity, especially when truncated at first or second order, many forward modelling approaches in cosmology rely on LPT to describe the dynamics of (cold collisionless) matter (see,

e.g., Jasche & Wandelt 2013; Kitaura 2013; Wang et al. 2013; Bos et al. 2019; Ata et al. 2020). Since it is no longer valid at and after shell-crossing, LPT has to be employed in a way that shell-crossed scales are filtered out prior to employing it (see, e.g., Sahni & Coles 1995). However, determining the scale of shell-crossing is approximate and usually challenging.

Another downside of LPT is that it predicts fluid trajectories, while in many cases one is instead interested in density or velocity fields at fixed spatial, that is Eulerian, coordinates. Numerically, an Eulerian density field can only be obtained by interpolating the fluid elements back to an Eulerian grid. This can be achieved using a particle injection scheme such as cloud-in-cell deposit (CIC; Hockney & Eastwood 1981), which is impacted by particle sampling noise. This can also be achieved by interpolating from a tessellation of the distribution function (e.g., Abel et al. 2012), which is computationally more expensive than CIC. Using Eulerian perturbation theory is not an option since it requires going to very high orders to achieve comparable accuracy to LPT (e.g., Bouchet 1996).

An alternative exists in “field-based” approaches that directly operate by predicting the Eulerian density based on the notion of particle trajectories. The semiclassical approach to describe cold collisionless dynamics presented in Uhlemann et al. (2019), but see also Short & Coles (2006a,b), provides such an alternative to LPT. This approach, which we name as propagator perturbation theory (PPT) here, translates LPT into an action and then uses a propagator to evolve a wave function, which encodes the cosmological perturbations. From the evolved wave function, the Eulerian density and velocity fields are readily obtained. The PPT approach introduces an additional free parameter, an effective $\hbar$, which acts as a natural smoothing, or coarse-graining scale. We note that PPT is fundamentally different from Schrödinger-Poisson (SP) analogues as effective models (cf. Widrow & Kaiser 1993; Uhlemann et al. 2014; Kopp et al. 2017; Garmy et al. 2020; Eberhardt et al. 2020) of cold Vlasov-Poisson (VP) dynamics (VP underlies all CDM non-linear cosmological structure formation; Peebles 1980). PPT is not a fully non-linear model like SP but a perturbative analogue to LPT, more similar in spirit to the Burgers approach of Matarrese & Mohayaee (2002). However, PPT gives easier access to phase space statistics than LPT by absorbing the “sum-over-streams” into a propagator (cf. Uhlemann et al. 2019).

Before shell-crossing, the first order of the PPT approach provides results equivalent to the first-order LPT in the limit of vanishing $\hbar$. At shell-crossing, while LPT leads to infinite densities, the PPT density remains finite and, after shell-crossing, the PPT density presents interference patterns in multi-stream regions. These interference patterns, therefore, provide a natural way to detect the shell-crossing scale. These oscillations naturally encode stream-averaged velocity fields (and higher moments) that are notoriously expensive to obtain for cold Vlasov dynamics (cf. Pueblas & Scoccimarro 2009; Hahn et al. 2015; Buehlmann & Hahn 2019).

By predicting Eulerian fields, PPT overcomes the particle sampling problem of LPT where particles cluster in the high-density regions and the under-densities are affected by high levels of shot noise, as illustrated in Fig. 1. This is especially relevant for the analysis of Lyman- α (Ly- α) forest observations since these data are particularly sensitive to under-dense regions in the matter distribution (cf. Peirani et al. 2014; Sorini et al. 2016).

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; Slosar et al.

2011; Busca et al. 2013; Palanque-Delabrouille et al. 2015; Rossi et al. 2015; Nasir et al. 2016; Yèche et al. 2017; Rossi 2017; Bautista et al. 2017; Boera et al. 2019; Blomqvist et al. 2019; Maitra 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). While various approaches to perform three-dimensional density reconstructions have been proposed in the literature, they are based on Wiener filter techniques (Ozbek et al. 2016; Stark et al. 2015a; Ravoux et al. 2020; Newman et al. 2020) or they assume the density amplitudes to be log-normally distributed (Kitaura et al. 2012; Gallerani et al. 2011). These approaches fail to reproduce the high-order statistics of the filamentary matter distribution. To reproduce the high-order statistics, Porqueres et al. (2019a) and Horowitz et al. (2019) recently used a large-scale optimisation approach to fit a gravitational structure growth model to simulated Ly- α data.

In this work, we employed for the first time PPT in a Bayesian forward model to infer the dark matter density from the Ly- α forest. In particular, we use the extension of the BORG framework (Jasche & Kitaura 2010; Jasche & Wandelt 2013; Lavaux et al. 2019) to the analysis of the Ly- α forest presented in Porqueres et al. (2019a), combined with a redshift-space optical depth field obtained from our extension of PPT presented here.

Our inference framework consists of a Gaussian prior on the primordial matter fluctuations, a physical model of structure formation to evolve the density field (in this case, the PPT), and a likelihood based on the fluctuating Gunn-Peterson approximation (FGPA, Gunn & Peterson 1965). To extract the large-scale structure information from the data, the BORG framework employs a Markov chain Monte-Carlo (MCMC) sampler. MCMC methods typically require a warm-up phase before they reach the target distribution and acquire a stationary state. This warm-up phase can be computationally expensive.

To accelerate the warm-up phase, we exploit the fact that PPT comes with a built-in tuneable scale parameter, $\hbar$, which controls an effective phase space resolution. One can think of this as an effective temperature so that a large $\hbar$ corresponds to a high temperature. The effective temperature controls to which features particle trajectories will be able to respond. In this work, we show that the computational costs of the warm-up phase can be reduced by performing a simulated annealing with the PPT model and taking advantage of the lower complexity of coarser scales. Such a procedure has been explored in the field of image processing (Gidas 1989; Alexander et al. 2003) and consists of walking down a hierarchy of scales from coarsest to finest resolution. At one level of resolution, we can focus on a particular scale: coarsest scales are frozen from the lowest resolution, and smallest scales are still evolving and will continue fluctuating in higher resolutions. Decreasing the effective $\hbar$ over the course of the chain thus corresponds to annealing and allows the trajectories to respond to increasingly finer structures. For high- $\hbar$ modelling, low spatial resolution can be used, allowing for further speed-up. By consistently changing the resolution and $\hbar$, we can then perform a simulated annealing that substantially reduces the computational cost of the warm-up phase of the MCMC sampler.

The paper is organised as follows. Section 2 provides a brief description of the PPT model and its extension to include redshift space distortions. Section 3 gives an overview of our Bayesian inference framework, BORG, as required for this work. In Sect. 4, we described the simulated data employed in testing and validating the method. The simulated annealing is described

Fig. 1. Density field obtained with LPT (*left*) and corresponding signal-to-noise due to the particle distribution (*right*). In the LPT model, particles cluster at high-densities, poorly sampling the low-density regimes from which the Ly- α forest arises.

in Sect. 5, showing that this strategy reduces the computational cost of the warm-up phase of the Markov sampler. The inference results from Ly- α forest data in redshift space are presented in Sect. 6.1. Finally, Sect. 7 summarises the results.

2. Propagator perturbation theory for the Ly- α forest

In this section, we briefly describe the PPT model as relevant for this work and present its extension to redshift space. For a more detailed description, its derivation and proof that it has LPT as its classical limit, we kindly refer the reader to Uhlemann et al. (2019).

2.1. Background

The PPT model relies on semiclassical dynamics to evolve the dark matter density using a propagator. In the Zel’dovich approximation, a fluid element moves in a time D_+ from its initial (Lagrangian) coordinate $\mathbf{q}$ to its final (Eulerian) coordinate $\mathbf{x}$ on a straight line, where D_+ is the linear theory growth factor. The classical action of this motion is therefore simply

$$S_0(\mathbf{x}, \mathbf{q}; a) = \frac{1}{2} \frac{(\mathbf{x} - \mathbf{q})^2}{D_+(a)}. \quad (1)$$

This action can be promoted to a (free) propagator K_0 using the Dirac-Feynman (Dirac 1933; Feynman 1948) trick

$$K_0(\mathbf{x}, \mathbf{q}; a) = (2\pi i \hbar D_+(a))^{-3/2} \exp \left[\frac{i}{\hbar} S_0(\mathbf{x}, \mathbf{q}; a) \right]. \quad (2)$$

This propagator can be used to compute the transition amplitude from an initial state represented by a wave function ψ_0 to the final state

$$\psi(\mathbf{x}, a) = \int d^3q K_0(\mathbf{x}, \mathbf{q}; a) \psi_0(\mathbf{q}). \quad (3)$$

Herein, $\hbar$ has no physical meaning, but instead is a free parameter that controls an effective smoothing scale, which we exploited to our advantage later. Uhlemann et al. (2019) have shown that this approach converges rigorously to the Zel’dovich approximation, and can be upgraded to second order LPT by adding a time-independent potential to the action. Here we only considered the free propagator however.

2.2. Implementation

We now describe our specific implementation of PPT and how late-time density and velocity fields can be obtained. Since the Zel’dovich approximation has pure growing-mode solutions only, the initial state has only one degree of freedom, the “back-scaled”¹ gravitational potential ϕ^{ic} which is given at a fictitious initial time when $D_+ = 0$ and represents a homogeneous state of the Universe before structure formation begins. The corresponding wave function has to have only phase perturbations and is given by

$$\psi_0(\mathbf{q}) := \exp \left[-\frac{i}{\hbar} \phi^{\text{ic}}(\mathbf{q}) \right]. \quad (4)$$

We note that the density associated with the initial wave function $\rho_0 := \psi_0 \bar{\psi}_0 = 1$ corresponds to the uniform mean density (an overline denotes a complex conjugate). We often refer also to the primordial density fluctuations, by which we mean the field $\delta^{\text{ic}} := \nabla^2 \phi^{\text{ic}}$, corresponding to the linear theory total matter density with the growth factor scaled out.

The evolved state ψ is obtained by computing the propagation of the field ψ_0 through Eq. (3), which mathematically corresponds to a convolution integral. This is numerically most conveniently carried out as a multiplication in Fourier space. Using the discrete Fourier transform (DFT), and exploiting circular convolution on a periodic domain, we find the full expression relating $\phi^{\text{ic}}(\mathbf{q})$ and ψ as

$$\psi(\mathbf{x}, a) = \text{DFT}^{-1} \left[\exp \left(-i \hbar \frac{k^2}{2} D_+(a) \right) \text{DFT} \left[\exp \left(-\frac{i}{\hbar} \phi^{\text{ic}}(\mathbf{q}) \right) \right] \right], \quad (5)$$

where $\mathbf{q}$ and $\mathbf{x}$ are discrete on a regular three-dimensional grid.

By construction, this wave function encodes all the phase-space information, that is, the full cumulant hierarchy (cf. Uhlemann 2018). In our case we are interested in the normalised density $\rho := 1 + \delta$, where δ is the fractional overdensity, and the peculiar momentum field $\mathbf{j} := (1 + \delta)\mathbf{v}$, where $\mathbf{v}$ is the peculiar velocity. These are given in terms of the propagated wave function $\psi = \psi(\mathbf{x}, a)$ as

$$\rho = \psi \bar{\psi} \quad \text{and} \quad \mathbf{j} = \frac{i \hbar}{2} (\psi \nabla \bar{\psi} - \bar{\psi} \nabla \psi). \quad (6)$$

This density agrees with a smoothed version of the Zel’dovich approximation before shell-crossing, where $\hbar$ controls the smoothing² (Short & Coles 2006a,b; Uhlemann et al. 2019). After shell-crossing, the Zel’dovich approximation is of course no longer valid since it does not account for secondary infall, and collapsed structures simply disperse again. More importantly, shell-crossing is accompanied by the formation of caustics, regions of infinite density, and multi-stream flow (cf. Arnold et al. 1982; Hidding et al. 2014). While the density in the classical approach becomes infinite or multi-valued, in PPT ρ remains finite and develops interference patterns in multi-stream regions, all regulated by the finite $\hbar$.

¹ Back-scaling means that we think of the potential at the initial time ϕ^{ic} and the linear potential ϕ at the target time a_{target} related by the factor $\phi^{\text{ic}} = \phi \times \frac{a_{\text{target}}}{D_+(a_{\text{target}})} \lim_{a \rightarrow 0} \frac{D_+(a)}{a}$.

² Following from Nyquist-Shannon, the smallest possible $\hbar$ fulfills the condition $|\Delta \phi| / \hbar < \pi$, where $\Delta \phi$ is the difference of the gravitational potential in neighbouring voxels.

2.3. Redshift-space distortions

Cosmological observations take place in redshift space rather than in comoving physical space. For all practical purposes, we can make the approximation of a distant observer, which implies that the redshift space distortion can be chosen to coincide with a Cartesian axis. Specifically, a particle is not observed to be at its Eulerian position $\mathbf{x}$, as we discussed in Sect. 2.1, but instead at its redshift space position $\mathbf{s}$, because of deviations from pure Hubble expansion (peculiar velocities). In LPT, this is given by

$$\mathbf{s} := \mathbf{x} + f(a) (\boldsymbol{\Psi} \cdot \hat{\mathbf{e}}_{\text{LOS}}) \hat{\mathbf{e}}_{\text{LOS}}, \quad (7)$$

where $\hat{\mathbf{e}}_{\text{LOS}}$ is a unit vector pointing along the line-of-sight (which we shall without loss of generality assume to be along the z -axis), $\boldsymbol{\Psi}$ is the displacement field between Lagrangian and Eulerian coordinates, $\boldsymbol{\Psi} := \mathbf{x} - \mathbf{q}$, and $f = d \log D_+ / d \log a$. This is quite obviously simply a velocity dependent displacement, and it can therefore be trivially included in an additional propagator from Eulerian to redshift space, given as

$$K_{\text{RSD}}(\mathbf{s}, \mathbf{x}; a) = N \exp \left[\frac{i}{\hbar} \frac{1}{2} \frac{((\mathbf{s} - \mathbf{x}) \cdot \hat{\mathbf{e}}_{\text{LOS}})^2}{f(a) D_+(a)} \right], \quad (8)$$

with N a normalisation that has to be suitably chosen. Effectively, at leading order PPT, the propagators can be trivially combined into a single propagator from Lagrangian space to redshift space, which in Fourier space takes the form

$$\widehat{K}(\mathbf{k}; a) := \exp \left[-\frac{i\hbar}{2} \left(k^2 + f(a) (\mathbf{k} \cdot \hat{\mathbf{e}}_{\text{LOS}})^2 \right) D_+(a) \right]. \quad (9)$$

While we did not use the next-to-leading order (NLO) version of PPT (cf. Uhlemann et al. 2019) here, the propagation to redshift space can also be applied at NLO, by performing the propagation to redshift space after carrying out the “kick-drift-kick” endpoint approximation to the path integral (their Eq. (D4)).

2.4. Modelling of the Ly- α -forest

We have explained above how PPT can be used to predict a quasi-linear density field $\rho = \psi \bar{\psi}$, consistent with the Zel’dovich approximation, from a wave function ψ propagated forwards to time a . In order to model the absorption of photons from the quasar, we employed the fluctuating Gunn-Peterson approximation (Gunn & Peterson 1965). The fractional transmitted flux is given by

$$F = e^{-\tau}, \quad (10)$$

where τ is the optical depth. In Eulerian space, the optical depth field reads

$$\tau(\mathbf{x}) := A \rho(\mathbf{x})^\beta, \quad (11)$$

where A and β are heuristic parameters, which are given by the physical state of the intergalactic medium. In a next and final step, we want to map this optical depth to redshift space and compute the transmitted quasar flux. In order to achieve this, we construct a new wave function that transports the optical depth, that is, we re-scale the amplitude such that

$$\chi_0(\mathbf{x}) := \sqrt{A} \rho^{\frac{\beta-1}{2}}(\mathbf{x}) \psi(\mathbf{x}). \quad (12)$$

This new wave function obeys $\chi_0 \bar{\chi}_0 = \tau$, while the phase information (i.e. the velocity) of the evolved wave function ψ is untouched. We note that this is essentially just a direct

application of Madelung’s interpretation of the wave function (Madelung 1927). This means we can exploit that the amplitude of the wave function has a conserved current. This is a crucial property since the optical depth τ is conserved under the mapping between physical and redshift space (e.g. Seljak 2012). We can therefore use the RSD propagator in Eq. (8) to propagate the χ field to redshift space, manifestly conserving τ , by evaluating

$$\chi(\mathbf{s}) := \int d^3x K_{\text{RSD}}(\mathbf{s}, \mathbf{x}; a) \chi_0(\mathbf{x}). \quad (13)$$

We note that this is essentially a “non-linear velocity RSD” (cf. Cieplak & Slosar 2016, their Eq. (4.3), but at a field level and using quasilinear Zel’dovich velocities at the order we are considering here). In a final step, we can obtain the three-dimensional quasar flux field F in redshift space by evaluating

$$F(\mathbf{s}) = \exp[-\chi \bar{\chi}]. \quad (14)$$

In Fig. 2, we show a comparison of the density field obtained with PPT (right panels) and the Zel’dovich approximation (left panels) in comoving physical space (top panels), and in redshift space (bottom panels), showing that PPT and LPT provide the same structures in real and redshift space. In the right panels, we show the quasar flux field F in physical and in redshift space.

3. The BORG framework for Ly- α forest

As mentioned above, we implemented PPT as a forward model in the BORG framework to infer the three-dimensional matter distribution underlying Ly- α forest data. In this section, we provide a summary of the algorithm. A more detailed description of the BORG framework can be found in Jasche & Wandelt (2013), Jasche et al. (2015), Lavaux & Jasche (2016), Jasche & Lavaux (2019) and, more specifically, the extension of BORG to the Ly- α forest analysis is described in Porqueres et al. (2019a).

The BORG framework is a Bayesian inference method 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 and structure formation models to observations. By using non-linear structure growth models, the BORG algorithm can exploit the full statistical power of high-order statistics of the matter distribution imprinted by gravitational clustering. This dynamical model links the primordial density fluctuations to the present large-scale structures. Therefore, the forward modelling approach allows translating the problem of inferring non-linear matter density fields into the inference of the spatial distribution of the primordial density fluctuations, which are well described by Gaussian statistics (Planck Collaboration IX 2020). The BORG algorithm, therefore, infers the initial matter fluctuations, the dark matter distribution and its dynamical properties from observations.

While the BORG framework incorporates several dynamical models based on Lagrangian perturbation theory and particle-mesh models, in this work, we include the PPT. Besides the advantage to work directly with fields, PPT allows for a reduction of the computational costs of the algorithm (see Sect. 5).

We tested the inference with PPT by applying the BORG framework to the analysis of simulated Ly- α forest data. The field-based approach of the PPT provides an advantage when analysing Ly- α forest observations since these data arise from low-density regimes, which are impacted by particle sampling noise in the standard LPT. To model the Ly- α forest, we used a Gaussian likelihood based on the fluctuating Gunn-Peterson

Fig. 2. Density fields (*left and middle column*) and quasar flux field in the FGP approximation (*right column*) in physical space (*top panels*) and in redshift space (*bottom panels*), with the line-of-sight direction upwards. *Leftmost panels:* use the Zel'dovich approximation and CIC deposit, the others use the PPT formalism. In all cases we used 256^3 resolution elements, the box size (and extent of each image) used for this comparison is $256 h^{-1}$ Mpc at $z = 2.5$. The thickness of the projected slice is $2 h^{-1}$ Mpc and white pixels in the CIC panels indicate zero particles deposited.

approximation,

$$P(\delta^{\text{ic}}, \delta^{\text{f}}|F) = \prod_{n,x} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{((F_n)_x - \exp(-\tau))^2}{2\sigma^2}\right], \quad (15)$$

where n labels the lines of sight and x runs over the pixels along a line of sight. The hierarchical representation of the algorithm is illustrated in Fig. 3.

At its core, the BORG framework employs MCMC techniques. This method allows inference of the full posterior distribution from which we can quantify the uncertainties in our results. However, the inference of the density field typically involves $\mathcal{O}(10^7)$ free parameters, corresponding to the discretised volume elements of the observed domain. To explore efficiently this high-dimensional parameter-space, the BORG framework uses a Hamiltonian Monte Carlo (HMC) method, which exploits the information in the gradients and adapts to the geometry of the problem. We need, therefore, the gradient of the dynamical forward model. The PPT gradient is derived in Appendix A. More details about the HMC and its implementation are described in Jasche & Kitaura (2010) and Jasche & Wandelt (2013).

4. The data

To test the inference framework, we generated artificial mock observations emulating the properties of present Ly- α forest sur-

veys such as the CLAMATO survey (Stark et al. 2015b; Lee et al. 2018) and LATIS (Newman et al. 2020). In this section, we describe the properties of the artificial data.

Mock data are constructed by first generating Gaussian initial conditions on a cubic Cartesian grid of side length of $128 h^{-1}$ Mpc with a resolution of $1 h^{-1}$ Mpc. To generate primordial Gaussian density fluctuations we used a cosmological matter power-spectrum including the baryonic wiggles calculated according to the prescription provided by Eisenstein & Hu (1998, 1999). We further assumed a standard Λ CDM cosmology with the following set of parameters: $\Omega_{\text{m}} = 0.31$, $\Omega_{\Lambda} = 0.69$, $\Omega_{\text{b}} = 0.022$, $h = 0.6777$, $\sigma_8 = 0.83$, $n_{\text{s}} = 0.9611$ (Planck Collaboration XIII 2016). Here $H_0 = 100 h \text{ km s}^{-1} \text{ Mpc}^{-1}$.

To generate realisations of the non-linear density field, we evolve the Gaussian primordial fluctuations via the forward model (PPT or LPT, respectively). A three-dimensional quasar flux field is generated by applying the FGPA model in Eqs. (10) and (11), 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, randomly distributed with a mean separation of $8 h^{-1}$ Mpc. The separation between lines of sight is the most important parameter of Ly- α surveys. Present surveys like CLAMATO (Lee et al. 2018) and LATIS (Newman et al. 2020)

Fig. 3. Hierarchical representation of the BORG inference framework for the analysis of Ly- α forest data. Primordial fluctuations δ_{ic} encoded in a set of Fourier modes at $z \approx 1000$ are obtained from the prior $P(\delta_{ic}|\Omega)$, where Ω represents the cosmological parameters. These initial conditions are evolved to $z = 2.5$ using PPT, which provides the optical depth τ_{PPT} and the evolved density, δ_{PPT} . The optical depth is then used to generate quasar spectra based on the fluctuating Gunn-Peterson approximation (FGPA). F_{obs} indicates the data. Purple boxes indicate deterministic transition while green boxes are probability distributions. Iterating this procedure results in a Monte Carlo Markov chain that explores the joint posterior distribution of the three-dimensional matter distribution underlying Ly- α forest observations.

achieve an average separation of $2.4 h^{-1}$ Mpc. Finally, we added Gaussian pixel-noise to the flux with $\sigma = 0.03$. This σ results in a signal-to-noise of $S/N = 2$, which corresponds to the majority of lines of sight in the CLAMATO survey.

5. Simulated annealing

In this section, we describe how we use a simulated annealing strategy to accelerate the warm-up phase of the MCMC sampler. Specifically, PPT has a built-in coarse-graining scale through the effective $\hbar$ which one can think of as an effective energy (or temperature) scale. By gradually reducing $\hbar$, one can, therefore, implement a simulated annealing procedure for the forward model. We compare the computational costs of the warm-up phase to the standard LPT approach. In this comparison, we use the PPT in physical (not redshift) space, as described in Sect. 2.2.

5.1. Annealing strategy

As discussed above, our inference method employs an MCMC 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 initialised 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 the stationary state, we may start recording samples to perform statistical analyses of the inference problem. The initial warm-up phase of the Markov sampler can be costly since it typically requires a high number of samples. In this work, we introduced simu-

lated annealing to reduce the computational cost of the warm-up phase.

The idea behind the simulated annealing is to take advantage of the lower complexity at larger scales and work down a hierarchy from the coarsest to the finest scales of the density field. For that, we start sampling only the large scales of the density field and, once these are converged, we map the density into a finer resolution and sample higher k -modes. This process reduces the computational cost of the warm-up phase in two ways. First, we only need to iterate enough to allow the relatively local structure to converge since the larger structures already converged at coarser levels. Secondly, the number of modes to be sampled in the coarser resolution is smaller, allowing rapid sampling.

When we increase the resolution, a key question is how to anneal in a way that the features of interest are represented in the current level and mapped to the next finer resolution. To answer this question, we make use of insights from re-normalisation theory, as it has been previously done in the field of image processing (see, e.g. Gidas 1989; Alexander et al. 2003). The insight from re-normalisation theory is that the effective temperature for a given feature is scale-dependent. This temperature can be translated into the smoothness of the density field: at some intermediate resolution, coarse scales are converged (“frozen”), and finer scales are still evolving (“hot”). This means that we can concentrate on the intermediate scales and ignore the changes in the smaller scales: we focus on the coarsest not-converged scales. In the PPT model, this is possible by changing the $\hbar$ parameter that controls the effective phase space resolution. We can think of $\hbar$ as an effective temperature, with high $\hbar$ corresponding to high temperature (coarser resolution). Gradually decreasing $\hbar$ corresponds to allowing the algorithm to respond to increasingly finer structures. For high $\hbar$, therefore, we can sample at low resolution since the algorithm only responds to large scales. We then perform a simulated annealing by consistently changing the resolution and $\hbar$. This is illustrated in Fig. 4.

In this work, we rely on heuristic rules to determine the changes of $\hbar$ at each level. We start with a high $\hbar$ and a coarse resolution of $N = 32^3$ voxels. Once the density field is converged, we open new modes by increasing the number of voxels to $N = 64^3$. The algorithm can now respond to these new modes, which are not yet converged. For this reason, we initially keep the same $\hbar$ and reduce it after few iterations, when the evolution of the small scales starts to saturate. Reducing $\hbar$ results in a sharper density and, therefore, the method becomes sensitive to smaller scales. We repeat these steps every time we change the resolution of the density field.

5.2. The warm-up phase of the Markov chain with annealing

In this Bayesian approach, we keep the cosmology fixed, and specify a prior on the initial power spectrum. However, the power spectrum of the inferred matter distribution is conditioned by the data, and we can use the posterior $P(k)$ as a diagnostic for the effectiveness of the inference since the power spectrum of the simulation differs from the prior. To monitor the initial warm-up phase of the Markov sampler, we follow a similar approach to our previous works (Jasche & Wandelt 2013; Jasche & Lavaux 2017, 2019; Ramanah et al. 2019; Porqueres et al. 2019a,b): we initialised the Markov chain with an over-dispersed state and traced the systematic drift of inferred quantities towards their preferred regions in the parameter space. Specifically, we initialised the Markov chain with a random Gaussian initial density field scaled by a factor 10^{-3} and monitored the drift

Fig. 4. Annealing of the density field. The density field is inferred hierarchically, starting with the largest scales and, once these are converged, opening new modes. These panels correspond to $N = 32^3$ voxels (left), $N = 64^3$ (middle) and $N = 128^3$ (right). The h parameter is decreased over the course of the chain, allowing the algorithm to respond to finer structures. *From left to right: h decreased from 0.15 to 0.09.*

Fig. 5. Burn-in of the posterior initial matter power spectra. *Left panel:* corresponds to PPT with annealing, and the *right panel* corresponds to standard LPT. The colour scale shows the evolution of the matter power spectrum with the number of samples. The dashed lines indicate the underlying power spectrum and the 1- and 2- σ cosmic variance limits. The Markov chain is initialised 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, recovering the true matter power spectrum at the end of the warm-up phase. Monitoring this drift allows us to identify when the Markov chain approaches a stationary distribution and provides unbiased estimates of the target distribution. The annealing with PPT reduces significantly the number of samples required in the warm-up phase, moving the chain faster to the target distribution. This is achieved by first sampling the coarser scales and gradually allowing the algorithm to respond to increasingly finer scales.

of corresponding posterior power-spectra during the warm-up phase. Figure 5 presents the results of this exercise for the standard LPT and the annealing with PPT. 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. While both forward models correctly recover the fiducial power spectrum, the simulated annealing with the PPT speeds up the burn-in phase, reducing the number of samples needed to reach the target distribution. Figure 6 shows the evolution of the amplitude of

the different Fourier modes in the posterior power spectrum with the number of MCMC samples. While the amplitudes of the $P(k)$ start to evolve in the first 50 samples for the PPT, the modes in the LPT take more than 150 samples to start evolving significantly. These results show that the annealing with PPT allows moving the chain faster towards the high probability regions of the parameter space, reducing the computational cost of the warm-up phase.

To test for residual correlations between different Fourier modes, we estimated the covariance matrix of power-spectrum

Fig. 6. Amplitudes of the posterior primordial matter power-spectrum at different Fourier modes traced during the warm-up phase of the MCMC sampler for the LPT (*upper panel*) and PPT with annealing (*lower panel*). As can be seen, initially, modes perform a coherent drift towards the high probability region in posterior distribution and start oscillating around their fiducial values once the Markov chain has reached a stationary state. The fiducial values are reached faster with the PPT annealing, reducing the computational cost of the warm-up phase.

amplitudes from our ensemble of Markov samples. Figure 7 shows that the covariance matrix for PPT with annealing has a clear diagonal structure, equivalent to the covariance for the standard LPT. This confirms that the annealing does not introduce spurious correlations between scales.

5.3. Correlation length

By design, subsequent samples in Markov chains are correlated. The statistical efficiency of an MCMC algorithm is determined by the effective 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 density amplitude at different locations of the box. For the amplitude at a given voxel, θ , the auto-correlation for samples with a given lag in the chain can be estimated as

$$C_n(\theta) = \frac{1}{N-n} \sum_{i=0}^{N-n} \frac{(\theta^i - \langle \theta \rangle)(\theta^{i+n} - \langle \theta \rangle)}{\text{Var}(\theta)} \quad (16)$$

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 presents the number of transitions required to generate one more effectively independent sample.

Figure 8 presents the results of this test for the standard LPT and the annealing with PPT. As can be seen, the correlation is generally lower for the PPT. This confirms that reducing the computational costs of the warm-up phase by annealing does not come at the expense of introducing longer correlations in the chain. Therefore, the annealing with the PPT results in a net speed-up of the Markov sampler to reach the target distribution.

Without annealing, the PPT shows an equivalent warm-up phase and correlation length to LPT (350 samples). However, the PPT is still an advantage over LPT since it provides a more accurate description of the density at low-density regimes (see Sect. 5.4).

5.4. Comparison to LPT

In this section, we compare the large-scale structures obtained with LPT and PPT. More specifically, this section focuses on comparing the profiles of cosmic structures in the density fields. For this, we evolved a set of initial conditions with both forward models (PPT and LPT) and compared the profiles of individual voids and clusters.

To compare the density profiles of a cluster (or a void), we randomly chose a local maximum (or minimum) in the final density field. We, then, determined the density profiles in spherical concentric shells. Figure 9 shows the density profiles for a cluster and a void obtained with PPT and LPT. Both models provide the same profiles, indicating that the PPT and LPT describe equivalent cosmic structures. While the standard deviation of the cluster profile is similar for both methods, the void profile shows a larger uncertainty region for LPT. This larger uncertainty in the void is due to the particle sampling noise introduced by the LPT. Since most of the particles cluster in high-density regions, voids are impacted by higher uncertainty in the LPT. The field-level approach of the PPT overcomes this problem, showing a lower standard deviation in the void profile.

6. Inference results

In this section, we present the results of applying our algorithm to Ly- α forest data in redshift space. We show that our method infers unbiased density fields and corresponding power-spectra at all scales considered in this work. We also perform a posterior predictive test for quasar spectra, showing that the inferred quantities can explain the data within the noise uncertainty.

6.1. Inferred density fields

As discussed above, our method uses a forward modelling approach, fitting a physical dynamical model to Ly- α forest data. This provides the full posterior distribution, from which we draw samples of the initial matter fluctuations and the non-linear spatial matter distribution at $z = 2.5$. In this section, the dynamical model is the PPT in redshift space, described in Sect. 2.4. Since the optical depth is conserved under the mapping between physical and redshift space, our inference in redshift space focuses on the optical depth field, while the density field is obtained in real space.

Fig. 7. Estimated correlation matrix of power spectrum amplitudes with the mean value, normalised using the variance of amplitudes of the power spectrum modes, for PPT (*left panel*) and LPT (*right panel*). We computed the correlation matrix from 600 samples after the warm-up phase. The low off-diagonal terms indicate that the annealing method does not introduce any erroneous mode coupling.

Figure 10 shows slices through the true fields, and the ensemble mean and variances of inferred three-dimensional fields, computed from 600 samples. A first visual comparison between ground truth and the inferred ensemble mean final density and optical depth fields shows that the algorithm correctly recovered the large-scale structure from Ly- α forest data. We note that the optical depth field is in redshift space while the final density field is in physical space. The lower right panel of Fig. 10 shows the corresponding standard deviations of the amplitudes, which is estimated from the samples in the Markov chain. The estimated density standard deviation correlates with the inferred density field. The same is true for the optical depth field. This is expected for a non-linear data model, which couples signal and noise. Higher uncertainty regions correspond to over-densities since the absorption saturates at high density, making the signal weaker. Once the light absorption is saturated, the data only provide information on a minimally lowest density threshold required to explain the observations. The line saturation effectively removes constraints from data above some lower threshold, nullifying the impact of higher density amplitudes in the likelihood. This leads the algorithm to use solely the prior to fill the missing pieces in high-density regions. The inference is thus not impacted at all by this observational limitation. In future development of the method, tighter constraints of the density amplitude at high densities could be achieved by modelling the absorption line profile since the line saturation introduces broadening of the profile.

While the standard deviations of the fields at $z = 2.5$ present a structure that correlates with the density, the standard deviation of the initial conditions is Gaussian noise, as shown in the upper panels of Fig. 10. This indicates that our forward model correctly propagates the information between the initial and final density field.

As discussed above, the mean separation between lines of sight is the most relevant parameter in Ly- α forest surveys. A particular challenge is to recover the density field in between one-dimensional lines of sight. To test the performance of our algorithm, we computed the Pearson correlation coefficients between the true density field and several density samples in our Markov chain. Figure 11 shows the correlation coefficients for a slice of width $1 h^{-1}$ Mpc, indicating the value of the Pearson

coefficient at different locations across this slice. This permits us to track the correlations on and in-between lines of sight (indicated with dashed lines). At the position of the lines of sight, the correlation is typically $>80\%$. For a better understanding of the fluctuations in the correlation, we indicated the position of lines of sight in the neighbouring slices (dotted lines). Most of the dotted lines correspond to a peak in the correlation, which indicates that the algorithm can interpolate the information between lines of sight. Also, there are regions of $15 h^{-1}$ Mpc without neighbouring lines of sight where the correlation is still $>70\%$ (see the regions centred at 36 and $100 h^{-1}$ Mpc), indicating that the method can recover the cosmic large-scale structure in the unobserved regions between observed lines of sight. We note that this plot is not comparable to Fig. 11 in Porqueres et al. (2019a) since the signal-to-noise of the mock data, the distribution of lines of sight and the width of the slice are different. A more detailed analysis of the impact of the mean line of sight separation on the accuracy of the results will be included in future work.

As a posterior test, we estimate the mean and variance of posterior power-spectra measured from the ensemble of Markov samples. The result is shown in Fig. 12. Although our method does not sample the different modes of the power spectrum, this is not enforced on the density samples. This means that, if the data requires it, the primordial power spectrum can be overwritten. Reconstructing three-dimensional density fields from one-dimensional Ly- α data is technically challenging. Previous approaches of inferring the density field from the Ly- α forest failed at recovering the correct power-spectrum amplitudes. 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 log-normal prior for the evolved density field. This approach inferred correct power-spectrum amplitudes at large scales $k < 0.1 h^{-1}$ Mpc but obtained erroneous excess power at smaller scales. Horowitz et al. (2019) used an optimisation approach to fit a dynamical forward model to the data, but the method obtains power-spectra that severely underestimate the power of density amplitudes. Typically, deviations from the fiducial power spectrum indicate the breakdown of the assumptions in the data model or the inference method. Figure 12 shows that our method recovers the fiducial power spectrum within the

Fig. 8. Autocorrelation of the density amplitudes as a function of the sample lag in the Markov chain for LPT (*upper panel*) and PPT (*lower panel*). This autocorrelation is estimated from 2000 samples after the warm-up phase. The correlation length of the sampler can be estimated by determining the point when correlations drop below 0.1 for the first time. The annealing with the PPT model does not introduce longer correlations in the density sampler.

$1\text{-}\sigma$ cosmic variance uncertainty at scales considered in this work. This demonstrates that our method is capable of inferring the matter distributions with the correct power spectrum from noisy Ly- α data in redshift space.

6.2. Posterior predictive tests

Posterior predictions allow testing of whether the inferred density fields provide accurate explanations for the data (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. Figure 13 shows the result of this test for one line of sight in redshift space, showing that the posterior predicted quasar spectrum recovers the data input within the observational 1σ uncertainty region. The 1σ region corresponds to the standard deviation of the noise added in this line of sight. This demonstrates that the method correctly locates absorber positions and corresponding amplitudes of the underlying densities and, therefore, the inferred quantities can explain the data at the level of the noise.

Fig. 9. Comparison of density profiles obtained with PPT and LPT. *Upper panel:* cluster density profile obtained by evolving the same initial conditions with PPT and LPT. *Lower panel:* test for a void profile. This demonstrates that PPT and LPT provide equivalent cosmic structures. The shaded regions indicate the uncertainty region of the profiles, corresponding to the standard deviation of 50 realisations. While the PPT and LPT show similar uncertainty regions for the cluster profile, the LPT has a larger standard deviation than the PPT in the void profile. This larger uncertainty is due to the poor sampling of the voids in the LPT model since most of the particles cluster in high-density regions.

6.3. Velocity fields from the PPT

The dynamical model in our algorithm allows us to naturally infer the velocity field since it derives from the initial perturbations. This velocity information can provide significant information on the formation of structures since it allows discrimination between peculiar velocities and the Hubble flow.

Figure 14 shows a slice through the line-of-sight component of the velocity field from the ground truth and the mean and standard deviation estimated from 200 samples. A visual comparison between the true and mean velocity fields shows that the algorithm recovers the true velocity from Ly- α forest data. This method, therefore, provides velocity fields constrained by the data at $z > 2$, where this information is challenging to obtain otherwise.

Figure 15 shows a zoom-in on the inferred mean density field and the corresponding velocity components (v_x, v_y) obtained with PPT. PPT provides more accurate peculiar velocities than

Fig. 10. Slices through ground truth initial (left upper panel), true evolved density field (left middle panel), true optical depth field (left lower panel), inferred ensemble mean initial (middle upper panel), ensemble mean evolved (middle-lower panel) density field and ensemble mean optical depth (lower middle panel) computed from 600 MCMC samples. The density fields are in physical space, obtained with the PPT as indicated in Eq. (6). The optical depth field is in redshift space, corresponding to $\tau = \chi \bar{\chi}$ with χ from Eq. (13). Comparison between these panels shows that the method recovers the structure of the true density fields with high accuracy. Right panels: standard deviations of inferred amplitudes of initial (upper right panel), final density fields (middle right panel) and optical depth (lower right panel). We note that the standard deviation of the density σ_{δ^f} but the mean density is plotted as $\log_{10}(2 + \langle \delta^f \rangle)$. We note that the uncertainty of δ^f and of the optical depth present a structure that correlates with the corresponding field. In contrast, the standard deviation of the initial conditions are homogeneous and show no correlation with the initial density field, indicating that the dynamical model correctly propagates the information between the primordial matter fluctuations and the final density and absorption fields.

the LPT in voids and filaments. In the LPT approach, one obtains the momentum field since the velocities are associated with massive particles. This means that we need to divide by the density field to obtain the peculiar velocities $\mathbf{v} = \mathbf{j}/\rho$. This requires smoothing of the density field to avoid empty regions with $\rho = 0$ and, in filaments, we need to average over enough particles. PPT overcomes these problems by operating on the field.

7. Summary and discussion

While LPT provides a good description of the dynamics of cold dark matter before shell-crossing, it requires to interpolate the fluid elements to an Eulerian field to obtain density and velocity fields in physical space. The discrete fluid elements further introduce a particle sampling noise in the fields. Since most of the particles cluster in the over-densities, voids and under-dense

regions are more impacted by the sampling noise. These under-dense regions are especially relevant when analysing Ly- α forest observations since these data mostly arise from sheets and voids.

A recent alternative to LPT is a field-based approach that directly predicts the Eulerian density field. The PPT, presented in Uhlemann et al. (2019), provides such an alternative to LPT, overcoming particle sampling noise, and giving easy access to the full Boltzmann hierarchy. PPT uses a propagator to evolve a wave function that encodes the density and velocity (and higher moment) information.

In this work we employ, for the first time, PPT in a Bayesian forward model to infer primordial fluctuations from sparse redshift-space quasar flux spectra, connected by the dynamical model and the fluctuating Gunn-Peterson approximation (FGPA). This framework is based on a Gaussian prior for the primordial fluctuations and a likelihood based on the quasar flux

Fig. 11. Pearson coefficient of the true density field and 300 density samples in our Markov chain. The red line corresponds to the mean of the correlation, and the shaded region indicates the standard deviation. The correlation is computed for a slice of width $1 h^{-1} \text{ Mpc}$ (the x -axis indicate the position across this slice). The dashed lines indicate the position of the lines of sight in the slice, and the dotted lines indicate the position of lines of sight in neighbouring slices. The Pearson coefficient is >0.7 at most of the locations in this slice, including regions where there are no neighbouring lines of sight. This indicates that the algorithm can interpolate the information between lines of sight and correctly recover the structures in unobserved regions.

Fig. 12. Mean posterior matter power-spectrum. The mean and the standard deviation of the initial matter power spectrum have been computed from 300 density samples of the Markov chain obtained after the warm-up phase. The standard deviation is plotted, but it is too small to be visible, 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 fiducial power-spectrum amplitudes within the 1- σ cosmic variance uncertainty limit throughout the entire range of Fourier modes considered in this work.

Fig. 13. Posterior predictive flux for a spectrum with noise $\sigma = 0.03$. The posterior predicted flux (orange line) is computed from the ensemble mean optical depth field in redshift space. The blue shaded region indicates the 1- σ region, corresponding to the standard deviation of the noise in this line of sight. This test checks 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 within the noise uncertainty.

field for the Ly- α forest. We fixed the cosmology. To explore the parameter space, our method employs MCMC techniques.

Furthermore, the PPT approach introduces a free parameter, $\hbar$, that acts as a natural smoothing (or temperature) scale. Allowing $\hbar$ to evolve over time allows performing a simulated annealing, thereby selecting the features and scales that the algorithm is sensitive to. Taking advantage of the lower complexity of coarser scales, we decrease $\hbar$ over the course of the chain, allowing the algorithm to respond to increasingly finer structures. By comparing to the standard LPT, we have shown that the PPT annealing reduces the computational cost of the warm-up phase of the MCMC sampler. With our implementation serving as a proof-of-concept, we find that multi-scale techniques are able to accelerate MCMC burn-in significantly. More sophisticated algorithms might be possible in the future that exploit this aspect further.

Since cosmological observations take place in redshift space, we have derived the RSD within the PPT formalism. We showed that RSDs can be easily included via an additional propagator between physical Eulerian space and redshift space. Since the optical depth is conserved under the mapping between physical and redshift space, the RSD propagator can be applied to a wave function that directly encodes the optical depth instead of the density. Based on PPT, we are therefore able to provide a forward model mapping primordial fluctuations to quasar flux in redshift space at the field level.

We have tested our inference method in redshift space with simulated data. These tests showed that our method recovers the underlying initial and final density field (in physical space) and the optical depth field (in redshift space). Our method based on PPT is able to correctly propagate the information between the initial matter fluctuations and the density and absorption field at $z = 2.5$.

From the dynamical forward model, we can easily derive peculiar velocity fields constrained by the data. To obtain peculiar velocities with the standard LPT, one needs to divide the momentum field by the density. This requires to previously

Fig. 14. Slices through ground truth (*left panel*) and mean (*middle panel*) velocity field in the direction of the line of sight, estimated from 200 samples. Comparison between these panels shows that the method recovers the true velocity field. *Right panel*: standard deviation.

Fig. 15. Zoom-in on the density field. The vector field shows the velocities derived from PPT, showing matter flowing out of the void and falling into the gravitational potential of the cluster. Our method provides consistent velocity and density fields that can be used to study structure formation. In particular, the velocity derived from the PPT provides more accurate estimates than the standard LPT in voids and filaments.

smooth the density field to avoid empty regions and, in filaments, one needs to average over a sufficient number of particles to obtain the correct velocity in the filament. These problems are naturally overcome by PPT since it operates at the field level. Therefore, PPT provides a better estimate of peculiar velocities in filaments and voids. This clearly demonstrates the advantage of field-based over fluid-element based dynamical³ models.

Acknowledgements. NP and OH thank Alan Heavens, Cornelius Rampf and Cora Uhlemann for discussions and comments on the draft. NP acknowledges funding from STFC through Imperial College Astrophysics Consolidated Grant ST/5000372/1. OH acknowledges funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant Agreement No. 679145, project “COSMO-SIMS”). GL acknowledges financial support from the ILP LABEX, under reference ANR-10-LABX-63, which is financed by French state funds managed by the ANR within the programme “Investissements d’Avenir” under reference ANR-11-IDEX-0004-02. GL also acknowledges financial support from the ANR BIG4, under reference ANR-16-CE23-0002. This work was carried out within the Aquila Consortium.

References

- Abel, T., Hahn, O., & Kaehler, R. 2012, *MNRAS*, **427**, 61
- Alexander, S. K., Fieguth, P., & Vrsay, E. R. 2003, in *Energy Minimization Methods in Computer Vision and Pattern Recognition*, eds. A. Rangarajan, M. Figueiredo, & J. Zerubia (Berlin, Heidelberg: Springer), 194
- Arnold, V. I., Shandarin, S. F., & Zeldovich, I. B. 1982, *Geophys. Astrophys. Fluid Dyn.*, **20**, 111
- Ata, M., Kitaura, F. S., Lee, K. G., et al. 2020, ArXiv e-prints [arXiv:2004.11027]
- Baugh, C. M., Gaztanaga, E., & Efstathiou, G. 1995, *MNRAS*, **274**, 1049
- Bautista, J. E., Busca, N. G., Guy, J., et al. 2017, *A&A*, **603**, A12
- Bird, S., Peiris, H. V., Viel, M., & Verde, L. 2011, *MNRAS*, **413**, 1717
- Blomqvist, M., du Mas des Bourboux, H., Busca, N. G., et al. 2019, *A&A*, **629**, A86
- Boera, E., Becker, G. D., Bolton, J. S., & Nasir, F. 2019, *ApJ*, **872**, 101
- Bos, E. G. P., Kitaura, F.-S., & van de Weygaert, R. 2019, *MNRAS*, **488**, 2573
- Bouchet, F. R. 1996, in *Dark Matter in the Universe*, eds. S. Bonometto, J. R. Primack, & A. Provenzale, 565
- Bouchet, F. R., Juszkiewicz, R., Colombi, S., & Pellat, R. 1992, *ApJ*, **394**, L5
- Buehlmann, M., & Hahn, O. 2019, *MNRAS*, **487**, 228
- Busca, N. G., Delubac, T., Rich, J., et al. 2013, *A&A*, **552**, A96
- Cieplak, A. M., & Slosar, A. 2016, *JCAP*, **2016**, 016
- 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
- Dirac, P. A. M. 1933, *Phys. Z. SowjUn.*, **3**, 64
- Eberhardt, A., Banerjee, A., Kopp, M., & Abel, T. 2020, *Phys. Rev. D*, **101**, 043011
- Eisenstein, D. J., & Hu, W. 1998, *ApJ*, **496**, 605
- Eisenstein, D. J., & Hu, W. 1999, *ApJ*, **511**, 5
- Feynman, R. P. 1948, *Rev. Mod. Phys.*, **20**, 367
- Gallerani, S., Kitaura, F. S., & Ferrara, A. 2011, *MNRAS*, **413**, L6
- Garny, M., Konstandin, T., & Rubira, H. 2020, *JCAP*, **2020**, 003
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. 2004, *Bayesian Data Analysis*, 2nd edn. (Chapman and Hall/CRC)
- Gidas, B. 1989, *IEEE Trans. Pattern Anal. Mach. Intell.*, **11**, 164
- Gunn, J. E., & Peterson, B. A. 1965, *ApJ*, **142**, 1633
- Hahn, O., Angulo, R. E., & Abel, T. 2015, *MNRAS*, **454**, 3920
- He, S., Alam, S., Ferraro, S., Chen, Y.-C., & Ho, S. 2018, *Nat. Astron.*, **2**, 401
- Hidding, J., Shandarin, S. F., & van de Weygaert, R. 2014, *MNRAS*, **437**, 3442
- Hockney, R. W., & Eastwood, J. W. 1981, *Computer Simulation Using Particles* (New York: McGraw-Hill)
- Horowitz, B., Lee, K.-G., White, M., Krolewski, A., & Ata, M. 2019, *ApJ*, **887**, 61
- 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. 2013, *MNRAS*, **429**, L84
- Kitaura, F. S., & Enßlin, T. A. 2008, *MNRAS*, **389**, 497
- Kitaura, F. S., Jasche, J., Li, C., et al. 2009, *MNRAS*, **400**, 183
- Kitaura, F.-S., Gallerani, S., & Ferrara, A. 2012, *MNRAS*, **420**, 61
- Kopp, M., Vattis, K., & Skordis, C. 2017, *Phys. Rev. D*, **96**, 123532
- Lahav, O., Fisher, K. B., Hoffman, Y., Scharf, C. A., & Zaroubi, S. 1994, *ApJ*, **423**, L93
- Lavaux, G., & Jasche, J. 2016, *MNRAS*, **455**, 3169

³ <https://aquila-consortium.org>

- Lavaux, G., Jasche, J., & Leclercq, F. 2019, *MNRAS*, submitted [arXiv:1909.06396]
- Lee, K.-G., Krolewski, A., White, M., et al. 2018, *ApJS*, 237, 31
- Madelung, E. 1927, *Z. Phys.*, 40, 322
- Maitra, S., Srianand, R., Petitjean, P., et al. 2019, *MNRAS*, 490, 3633
- Matarrese, S., & Mohayaee, R. 2002, *MNRAS*, 329, 37
- Nasir, F., Bolton, J. S., & Becker, G. D. 2016, *MNRAS*, 463, 2335
- Newman, A. B., Rudie, G. C., Blanc, G. A., et al. 2020, *ApJ*, 891, 147
- 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
- Peacock, J. A., & Dodds, S. J. 1996, *MNRAS*, 280, L19
- Peebles, P. J. E. 1980, *The Large-scale Structure of the Universe* (Princeton: Princeton University Press)
- Peirani, S., Weinberg, D. H., Colombi, S., et al. 2014, *ApJ*, 784, 11
- Planck Collaboration XIII. 2016, *A&A*, 594, A13
- Planck Collaboration IX. 2020, *A&A*, 641, A9
- Porqueres, N., Jasche, J., Lavaux, G., & Enßlin, T. 2019a, *A&A*, 630, A151
- Porqueres, N., Kodi Ramanah, D., Jasche, J., & Lavaux, G. 2019b, *A&A*, 624, A115
- Pueblas, S., & Scoccimarro, R. 2009, *Phys. Rev. D*, 80, 043504
- Ramanah, D. K., Lavaux, G., Jasche, J., & Wandelt, B. D. 2019, *A&A*, 621, A69
- Ravoux, C., Armengaud, E., Walther, M., et al. 2020, *J. Cosmol. Astropart. Phys.*, 7, 010
- Rossi, G. 2017, *ApJS*, 233, 12
- Rossi, G., Yèche, C., Palanque-Delabrouille, N., & Lesgourgues, J. 2015, *Phys. Rev. D*, 92, 063505
- Sahni, V., & Coles, P. 1995, *Phys. Rep.*, 262, 1
- Seljak, U. 2012, *JCAP*, 2012, 004
- Seljak, U., Slosar, A., & McDonald, P. 2006, *JCAP*, 2006, 014
- Short, C. J., & Coles, P. 2006a, *JCAP*, 2006, 012
- Short, C. J., & Coles, P. 2006b, *JCAP*, 2006, 016
- Slosar, A., Font-Ribera, A., Pieri, M. M., et al. 2011, *JCAP*, 2011, 001
- Smith, R. E., Peacock, J. A., Jenkins, A., et al. 2003, *MNRAS*, 341, 1311
- Sorini, D., Oñorbe, J., Lukić, Z., & Hennawi, J. F. 2016, *ApJ*, 827, 97
- 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
- Uhlemann, C. 2018, *JCAP*, 2018, 030
- Uhlemann, C., Kopp, M., & Haugg, T. 2014, *Phys. Rev. D*, 90, 023517
- Uhlemann, C., Rampf, C., Gosenca, M., & Hahn, O. 2019, *Phys. Rev. D*, 99, 083524
- Viel, M., Haehnelt, M. G., & Lewis, A. 2006, *MNRAS*, 370, L51
- Wang, H., Mo, H. J., Yang, X., & van den Bosch, F. C. 2013, *ApJ*, 772, 63
- Widrow, L. M., & Kaiser, N. 1993, *ApJ*, 416, L71
- Yèche, C., Palanque-Delabrouille, N., Baur, J., & du Mas des Bourboux, H. 2017, *JCAP*, 2017, 047
- Zaroubi, S., Hoffman, Y., & Dekel, A. 1999, *ApJ*, 520, 413
- Zel'dovich, Y. B. 1970, *A&A*, 5, 84

Appendix A: Adjoint gradient of the PPT

The inference of the density field requires inferring the amplitudes of the primordial density 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. To explore this high-dimensional parameter space efficiently, the BORG framework employs a Hamiltonian Monte Carlo (HMC) method, which adapts to the geometry of the problem by using the information in the gradients. Therefore, this algorithm requires the derivatives of the forward model. In this section, we derive the gradient of the PPT.

More specifically, the HMC relies on the availability of a gradient of the posterior distribution. Therefore, we need to compute the gradient of the log-likelihood with respect to the gravitational potential ϕ^{ic} .

A.1. Gradient of PPT in real space

First, we derive the gradient corresponding to the PPT described in Sects. 2.1 and 2.2. The gradient of PPT in redshift space is derived in the following section.

We want the gradient of the likelihood with respect to the gravitational potential ϕ_p^{ic} ,

$$\frac{\partial \log \mathcal{L}}{\partial \phi_p^{\text{ic}}} = \sum_l \frac{\partial \log \mathcal{L}}{\partial \delta_l^f} \frac{\partial \delta_l^f}{\partial \phi_p^{\text{ic}}} \quad (\text{A.1})$$

where $\log \mathcal{L}$ is the log-likelihood function and δ_l^f indicates the final density field at voxel l . In the Gunn-Peterson approximation, the derivative of the data model is

$$\frac{\partial \log \mathcal{L}}{\partial \delta_l^f} = \sum_n \frac{(F_n)_l - \exp[-A(1 + \delta_l^f)_l]^\beta}{\sigma^2} \times A\beta(1 + \delta_l^f)_l^{\beta-1} \exp[-A(1 + \delta_l^f)_l]^\beta \quad (\text{A.2})$$

where n runs over the different lines of sight. However, in this section we focus on the gradient of the dynamical forward model, independently of the data model. Therefore, we not specify the derivative of the likelihood with respect to the final density.

The final density field is given by $\delta_l^f = \psi_l \bar{\psi}_l$. Therefore, Eq. (A.1) reads

$$\frac{\partial \log \mathcal{L}}{\partial \phi_p^{\text{ic}}} = \sum_l \frac{\partial \log \mathcal{L}}{\partial \delta_l^f} \left[\frac{\partial \psi_l}{\partial \phi_p^{\text{ic}}} \bar{\psi}_l + \psi_l \frac{\partial \bar{\psi}_l}{\partial \phi_p^{\text{ic}}} \right] \quad (\text{A.3})$$

The derivative of ψ with respect to the initial gravitational potential is

$$\frac{\partial \psi_l}{\partial \phi_p^{\text{ic}}} = \sum_m M_{lm} \hat{K}_m M'_{mp} \left(\frac{-i}{\hbar} \right) e^{-\frac{i}{\hbar} \phi_p^{\text{ic}}} \quad (\text{A.4})$$

and

$$\frac{\partial \bar{\psi}_l}{\partial \phi_p^{\text{ic}}} = \sum_m M_{lm} \bar{\hat{K}}_m M'_{mp} \left(\frac{i}{\hbar} \right) e^{\frac{i}{\hbar} \phi_p^{\text{ic}}} \quad (\text{A.5})$$

where M indicate Fourier transforms and $\hat{K}_m = \exp[-i\frac{1}{2}\hbar k_m^2 D_+]$ is the free propagator in Fourier space.

Finally, the gradient reads

$$\frac{\partial \log \mathcal{L}}{\partial \phi_p^{\text{ic}}} = \frac{i}{\hbar} \sum_l \frac{\partial \log \mathcal{L}}{\partial \delta_l^f} \times \sum_m M_{lm} \left[\bar{\hat{K}}_m M'_{mp} e^{\frac{i}{\hbar} \phi_p^{\text{ic}}} - \hat{K}_m M'_{mp} e^{-\frac{i}{\hbar} \phi_p^{\text{ic}}} \right] \quad (\text{A.6})$$

A.2. Gradient PPT with RSD

Here we derived the gradient of the PPT in redshift for the Ly- α forest, described in Sect. 2.4. As before, the HMC requires the gradient of the likelihood with respect to the gravitational potential ϕ^{ic} .

$$\frac{\partial \log \mathcal{L}}{\partial \phi_p^{\text{ic}}} = \sum_l \frac{\partial \log \mathcal{L}}{\partial \tau_l} \frac{\partial \tau_l}{\partial \phi_p^{\text{ic}}} \quad (\text{A.7})$$

where τ_l is the optical depth at voxel l . In this case, the derivative of the data model is

$$\frac{\partial \log \mathcal{L}}{\partial \tau_l} = - \sum_n \frac{(F_n)_l - \exp(-\tau_l)}{\sigma^2} \exp(-\tau_l) \quad (\text{A.8})$$

where n runs over the lines of sight.

The optical depth is given by $\tau = \chi \bar{\chi}$. Therefore, Eq. (A.7) is

$$\frac{\partial \log \mathcal{L}}{\partial \phi_p^{\text{ic}}} = \sum_l \frac{\partial \log \mathcal{L}}{\partial \tau_l} \left[\frac{\partial \chi_l}{\partial \phi_p^{\text{ic}}} \bar{\chi}_l + \chi_l \frac{\partial \bar{\chi}_l}{\partial \phi_p^{\text{ic}}} \right] \quad (\text{A.9})$$

Developing the first term, we obtain

$$\sum_l \frac{\partial \log \mathcal{L}}{\partial \tau_l} \frac{\partial \chi_l}{\partial \phi_p^{\text{ic}}} \bar{\chi}_l = \sum_l \frac{\partial \log \mathcal{L}}{\partial \tau_l} \bar{\chi}_l \times \sum_m M_{lm} (\hat{K}_{\text{RSD}})_m \sum_n M'_{mn} \frac{\partial \chi_0}{\partial \phi^{\text{ic}}} \quad (\text{A.10})$$

where M indicate a Fourier transform and $\hat{K}_{\text{RSD}}$ is the propagator in Eq. (9).

Introducing $\chi_0 = \sqrt{A} (\psi \bar{\psi})^{\frac{\beta-1}{2}} \psi$ in the previous equation, we get

$$\begin{aligned} \sum_l \frac{\partial \log \mathcal{L}}{\partial \tau_l} \frac{\partial \chi_l}{\partial \phi_p^{\text{ic}}} \bar{\chi}_l &= \sqrt{A} \sum_l \frac{\partial \log \mathcal{L}}{\partial \tau_l} \bar{\chi}_l \\ &\times \sum_m M_{lm} (\hat{K}_{\text{RSD}})_m \\ &\times \sum_n M'_{mn} \left[(\psi \bar{\psi})_n^{\frac{\beta-1}{2}} \frac{\partial \psi_n}{\partial \phi^{\text{ic}}} \right. \\ &\left. + \frac{\beta-1}{2} (\psi \bar{\psi})_n^{\frac{\beta-3}{2}} \left(\frac{\partial \psi_n}{\partial \phi_p^{\text{ic}}} \bar{\psi}_n + \psi_n \frac{\partial \bar{\psi}_n}{\partial \phi_n^{\text{ic}}} \right) \psi_n \right] \quad (\text{A.11}) \end{aligned}$$

Similarly, we can obtain the second term of Eq. (A.9) and combine them, obtaining

$$\begin{aligned} \sum_l \frac{\partial \log \mathcal{L}}{\partial \tau_l} \frac{\partial \chi_l}{\partial \phi_p^{\text{ic}}} \bar{\chi}_l &= \sqrt{A} \sum_l \frac{\partial \log \mathcal{L}}{\partial \tau_l} \bar{\chi}_l \\ &\times \sum_m M_{lm} \left\{ (\hat{K}_{\text{RSD}})_m \sum_n M'_{mn} \frac{\beta+1}{2} \rho_n^{\frac{\beta-1}{2}} \right. \\ &+ (\bar{\hat{K}}_{\text{RSD}})_m \sum_n M'_{mn} \frac{\beta-1}{2} \rho_n^{\frac{\beta-3}{2}} (\bar{\psi})_n^2 \left. \frac{\partial \psi_n}{\partial \phi_p^{\text{ic}}} \right. \\ &+ \left[(\hat{\bar{K}}_{\text{RSD}})_m \sum_n M'_{mn} \frac{\beta+1}{2} \rho_n^{\frac{\beta-1}{2}} \right. \\ &\left. + (\hat{K}_{\text{RSD}})_m \sum_n M'_{mn} \frac{\beta-1}{2} \rho_n^{\frac{\beta-3}{2}} (\psi)_n^2 \right] \frac{\partial \bar{\psi}_n}{\partial \phi_p^{\text{ic}}} \left. \right\} \quad (\text{A.12}) \end{aligned}$$

with the derivatives from Eqs. (A.4) and (A.5).