



TARDIS. I. A Constrained Reconstruction Approach to Modeling the $z \sim 2.5$ Cosmic Web Probed by $\text{Ly}\alpha$ Forest Tomography

Benjamin Horowitz^{1,2,3} , Khee-Gan Lee^{1,3} , Martin White^{1,2,4} , Alex Krolewski⁴ , and Metin Ata³

¹Lawrence Berkeley National Lab, 1 Cyclotron Road, Berkeley, CA 94720, USA; bhorowitz@berkeley.edu

²Department of Physics, University of California at Berkeley, Berkeley, CA 94720, USA

³Kavli IPMU (WPI), UTIAS, The University of Tokyo, Kashiwa, Chiba 277-8583, Japan

⁴Department of Astronomy, University of California at Berkeley, Berkeley, CA 94720, USA

Received 2019 April 4; revised 2019 September 23; accepted 2019 October 11; published 2019 December 11

Abstract

Recent $\text{Ly}\alpha$ forest tomography measurements of the intergalactic medium (IGM) have revealed a wealth of cosmic structures at high redshift ($z \sim 2.5$). In this work, we present the Tomographic Absorption Reconstruction and Density Inference Scheme (TARDIS), a new chronocsmographic analysis tool for understanding the formation and evolution of these observed structures. We use maximum likelihood techniques with a fast nonlinear gravitational model to reconstruct the initial density field of the observed regions. We find that TARDIS allows accurate reconstruction of smaller-scale structures than standard Wiener-filtering techniques. Applying this technique to mock $\text{Ly}\alpha$ forest data sets that simulate ongoing and future surveys such as CLAMATO, Subaru PFS, or the ELTs, we are able to infer the underlying matter density field at observed redshift and classify the cosmic web structures. We find good agreement with the underlying truth in both the characteristic eigenvalues and eigenvectors of the pseudo-deformation tensor, with the eigenvalues inferred from 30 m class telescopes correlated at $r = 0.95$ relative to the truth. As an output of this method, we are able to further evolve the inferred structures to late time ($z = 0$) and also track the trajectories of coeval $z = 2.5$ galaxies to their $z = 0$ cosmic web environments.

Key words: cosmology: observations – galaxies: high-redshift – intergalactic medium – methods: numerical – techniques: spectroscopic

1. Introduction

A major goal of modern astrophysics is to understand how galaxies form and evolve from initial density fluctuations to the current day. Over the past few decades, it has become increasingly clear that the surrounding large-scale structures around galaxies play a critical role in their formation, morphology, and evolution (Dressler 1980; Kauffmann et al. 2004). There has also been new theoretical understanding of how these large-scale dark matter structures evolve, from both an analytical approach and numerical simulations (see Conselice 2014, for an overview). However, our understanding of the small-scale processes driving galaxy evolution remains poor, with many competing models (Conselice 2014; Naab & Ostriker 2017). Part of the challenge lies in the fact that most observations linking galaxy evolution and large-scale structure are at low redshifts, whereas most of the galaxy and star formation in the universe peaked at the so-called “Cosmic Noon” epoch at $z \sim 1.5\text{--}3$ (Madau & Dickinson 2014), which remains out of reach of most large-scale structure surveys.

There are many indications of the interconnected nature of cosmic structure and galactic evolution at high redshift. Numerous studies have found that low-redshift galaxies living in cluster environments have lower star formation rates and significantly older stellar ages than those in the field (Wake et al. 2005; Skibba et al. 2009). This indicates that these regions underwent significant star formation and quenching at high redshift ($z > 1.5$; Tran et al. 2010). This is further supported by simulation work showing that protoclusters produce roughly half of their stellar content at $2 < z < 4$ and are therefore an important contribution to the overall cosmic star formation rate (Chiang et al. 2017). Beyond protoclusters, there is evidence to suggest that star formation properties may

further depend on where the galaxy is first formed in the cluster or falls in along filamentary structure (Porter et al. 2008). Similarly, hydrodynamical simulations (Dubois et al. 2014) have suggested that the spin of galaxies may depend on the filament orientation, with simulated red and blue galaxies aligning perpendicular and parallel to the filament, respectively. Very few data are available tracing these cosmic structures at high redshift, but next-generation surveys will provide the depth over sufficient sky coverage to better constrain these astrophysical processes (Kartaltepe et al. 2019; Overzier & Kashikawa 2019).

Understanding these complex relationships between baryonic properties and dark matter in the context of the overall large-scale structure environment is not only useful in modeling galaxy formation but also crucial in exploiting galaxies as biased tracers of large-scale structure for cosmological constraints (Desjacques et al. 2018). The relationships between cosmic web structures and bias have been explored in the case of tidal shear bias (Baldauf et al. 2012) and, more recently, assembly bias (Ramakrishnan et al. 2019). Quantifying the sources of bias will be necessary when extending galaxy clustering surveys into the nonlinear regime where the particulars of the cosmic web may play a role (Alam et al. 2019) or in cosmic shear surveys where intrinsic alignments of galaxies will contribute substantial systematic uncertainty to precision cosmological measurements (Troxel & Ishak 2015).

So far, most studies of the cosmic web have used optically selected galaxies from spectroscopic redshift surveys as a tracer of the cosmic web. As a high number density (and therefore high spectroscopic sampling rate) is necessary for this sort of survey, this technique becomes increasingly expensive at higher redshift. The current state-of-the-art galaxy survey probing the high-redshift cosmic web is the VIPERS survey

(Guzzo et al. 2014) on the Very Large Telescope (VLT), which has obtained redshifts for 100,000 galaxies over 24 deg^2 as the largest-ever spectroscopic campaign on that facility. This enabled a cosmic web analysis in the redshift range $0.4 < z < 1.0$ (Malavasi et al. 2017), which suggested segregation of massive galaxies toward filaments already at this redshift. Over the next few years, new, massively multiplexed fiber spectrographs on 8 m class telescopes, such as VLT-MOONS (Cirasuolo et al. 2014) and Subaru Prime Focus Spectrograph (PFS; Takada et al. 2014), will allow such high sampling rate galaxy surveys to push to $z \sim 1.5$, but they would be prohibitively expensive at the “Cosmic Noon” epoch of $z \sim 2\text{--}3$.

In recent years, however, intergalactic medium (IGM) tomography (Pichon et al. 2001; Caucci et al. 2008; Lee et al. 2014; Stark et al. 2015) of the hydrogen Ly α forest has provided a complementary approach to mapping high-redshift large-scale structure. This technique uses dense configurations of closely spaced star-forming galaxies, in addition to quasars, as background sources to probe the three-dimensional (3D) structure of the optically thin IGM gas at $z > 2$ on scales of several comoving Mpc. The ongoing COSMOS Ly α Mapping and Tomographic Observations (CLAMATO) survey is the first observational program to implement IGM tomography and now has 240 sight lines covering an $\sim 600 \text{ arcmin}^2$ footprint within the COSMOS field, yielding a 3D tomographic map of the $2.05 < z < 2.55$ Ly α forest (Lee et al. 2018). A number of $z \sim 2.3$ cosmic structures have already been detected in the CLAMATO data, including protoclusters (Lee et al. 2016) and cosmic voids (Krolewski et al. 2018).

In the coming years, a number of next-generation spectroscopic surveys will radically increase the observational resources available for IGM tomography, including the Subaru PFS and Maunakea Spectroscopic Explorer (McConnachie et al. 2016). These telescopes will offer multiplex factors of several thousand over $\sim 1 \text{ deg}^2$ fields of view, allowing several times the volume of the current CLAMATO data to be observed within a single night. Meanwhile, with far sparser sight-line number density but significantly larger sky coverage, the Dark Energy Spectroscopic Instrument (DESI; Levi et al. 2013) could be another interesting platform for Ly α forest tomography to probe large-scale overdensities. Farther into the future, the 30 m class facilities, such as the Thirty Meter Telescope (TMT; Skidmore et al. 2015), Giant Magellan Telescope (GMT; Johns et al. 2012), and European Extremely Large Telescope (EELT; Evans et al. 2014), will have smaller fields of view but dramatically improved sensitivity for faint background sources at much greater sight-line densities that can probe spatial scales of $\sim 1 \text{ cMpc}$ and below. The need for accurate modeling of the formation and evolution of galaxies and galaxy clusters increases in order to maximize the science return of these facilities.

The current standard procedure for IGM tomography analysis is to create a Wiener-filtered absorption map from the observed Ly α absorption features (Pichon et al. 2001; Caucci et al. 2008; Lee et al. 2014). This absorption field can then be related to the underlying matter density through the fluctuating Gunn–Peterson approximation (FGPA). This Wiener filtering does not explicitly include information about the physical processes of the system and could, in an extreme case, lead to inferred matter distributions that cannot arise from gravitational evolution. In this work, we implement a different approach, finding the maximum a posteriori initial density field

that gives rise to the observed density field, often known as a “constrained realization.” This will constrain the transmitted flux⁵ field to those that are likely to arise from gravitational evolution, providing a more accurate reconstruction at $z = 2.5$. This epoch is particularly amenable to this technique, since the observed structures are only mildly nonlinear and have not yet undergone shell crossing. This will yield not only information on the underlying dark matter density field but also velocity information that allows us to deconvolve redshift- and real-space quantities (see Nusser & Haehnelt 1999 for a reconstruction method applied to 1D quasar Ly α forest sight lines and Pichon et al. 2001 for a full 3D convolution). This velocity information can also help inform the astrophysical processes occurring in the region; for example, combining the flux information, matter velocity information, and a galaxy catalog will provide insights into galaxy formation environmental dependence. In addition, since we have the $z = 2.5$ matter density and velocities, we are able to further evolve our field to $z = 0$ to infer the late-time fate of the observed structures.

Reconstructing the initial density field has additional advantages beyond possible improvements in late-time reconstruction. As there is currently no evidence for primordial non-Gaussianity (Planck Collaboration et al. 2018), the power spectrum of the initial density modes should provide a lossless statistic. The entire family of higher-order correlations (such as three-point functions, density peak counts, voids, topological measures, etc.) arises due to gravitational evolution of a density field described by a single power spectrum. In the case of galaxy large-scale structure surveys, there has already been work toward performing this optimal reconstruction (Seljak et al. 2017). As Ly α tomography builds up toward cosmological volumes, it would be worth exploring the application of the aforementioned techniques.

In this paper, we apply initial density reconstruction to mock observations of IGM tomography using the Tomographic Absorption Reconstruction and Density Inference Scheme (TARDIS). We overview the formalism in Section 2, describing the optimization scheme, forward model used, and measures of the cosmic web. In Section 3, we describe our mock data sets that simulate Ly α tomography observations. In Section 4 we describe our results, and we discuss next steps in Section 5.

2. Methodology

In order to implement our scheme to go from observed data to the system initial conditions, we need (a) a dynamic forward model (FastPM), (b) an absorption model (FGPA), (c) mapping from the field to data space (flux skewers), and (d) a noise model. In this section, we describe each component of our model.

2.1. Modeling

Here we summarize the optimization technique and standardize notation. For a more complete description, see Seljak (1998), Simon et al. (2009), Seljak et al. (2017), and Horowitz et al. (2019).

We measure N skewers of flux, assuming perfect identification of the continuum spectra each of length L , and stack those

⁵ It is a mild misnomer to refer to the Ly α transmission as a “flux,” but in this paper, we use both terms interchangeably.

into a full data vector, \mathbf{d} , of total dimension $N \times L$. This data vector will depend on the initial conditions we wish to estimate at a certain resolution M , s , the Ly α absorption model, and a noise term \mathbf{n} , which we choose to have the same dimension as the data, i.e.,

$$\mathbf{d} = \mathbf{R}(s) + \mathbf{n}, \quad (1)$$

where $\mathbf{R}: M^3 \rightarrow N \times L$ is the (nonlinear) response operator composed of a forward operator and a skewer-selector function. The Gaussian information is contained in covariance matrices, $\mathbf{S} = \langle \mathbf{s}\mathbf{s}^\dagger \rangle$ and $\mathbf{N} = \langle \mathbf{n}\mathbf{n}^\dagger \rangle$, for the estimated signal and noise components, which are assumed to be uncorrelated with each other, i.e., $\langle \mathbf{n}(\mathbf{R}(s))^\dagger \rangle = 0$. In this work, we are interested in maximizing the likelihood of some underlying signal given the data. The generic likelihood function can be written as

$$L(s|\mathbf{d}) = (2\pi)^{-(N+M)/2} \det(\mathbf{S}\mathbf{N})^{-1/2} \times \exp\left[-\frac{1}{2}\mathbf{s}^\dagger \mathbf{S}^{-1} \mathbf{s} + (\mathbf{d} - \mathbf{R}(s))^\dagger \mathbf{N}^{-1} (\mathbf{d} - \mathbf{R}(s))\right], \quad (2)$$

where we assume the signal covariance \mathbf{S} around some fiducial power spectra. The exponential in this likelihood can be interpreted as the sum of a prior term ($\mathbf{s}^\dagger \mathbf{S}^{-1} \mathbf{s}$) and a data-dependent term ($(\mathbf{d} - \mathbf{R}(s))^\dagger \mathbf{N}^{-1} (\mathbf{d} - \mathbf{R}(s))$), with the prefactor as a normalization term. Note that the minimum variance solution for the signal field can be found by minimizing,

$$\chi^2 = \mathbf{s}^\dagger \mathbf{S}^{-1} \mathbf{s} + (\mathbf{d} - \mathbf{R}(s))^\dagger \mathbf{N}^{-1} (\mathbf{d} - \mathbf{R}(s)), \quad (3)$$

with respect to s . Working in quadratic order around some fixed s_m , we have

$$\chi^2 = \chi_0^2 + 2\mathbf{g}(s - s_m) + (s - s_m)\mathbf{D}(s - s_m), \quad (4)$$

with gradient function

$$\mathbf{g} = \frac{1}{2} \frac{\partial \chi^2}{\partial s} = \mathbf{S}^{-1} s_m - \mathbf{R}'^\dagger(s_m) \mathbf{N}^{-1} (\mathbf{d} - \mathbf{R}(s_m)) \quad (5)$$

and curvature term⁶

$$\mathbf{D} = \frac{1}{2} \frac{\partial^2 \chi^2}{\partial s \partial s} = \mathbf{S}^{-1} + \mathbf{R}'^\dagger \mathbf{N}^{-1} \mathbf{R}'. \quad (6)$$

Calculation of the derivative term \mathbf{R}' requires calculation with respect to every initial mode. We use an automated differentiation framework in Appendix B of Feng et al. (2018) to calculate the Jacobian products of our evolution operator without running additional simulations. This avoids running additional involved simulations with respect to every mode, which would be prohibitively costly.

2.2. Optimization

As each iteration of the chain requires running a PM simulation, it is important to minimize computational time. While others have used Hamiltonian Markov Chain Monte Carlo algorithms to find fast reconstructions for galaxy surveys (see Jasche & Wandelt 2013; Wang et al. 2014, 2016a), in this work, we are instead finding the most likely map

reconstruction. We therefore use a limited-memory Broyden–Fletcher–Goldfarb–Shanno (LBFGS) algorithm (Press et al. 2002), a general technique for solving nonlinear optimization problems. Rather than sampling over the entire parameter space, LBFGS takes a quasi-Newtonian approach; i.e., it is similar to the standard Newton–Raphson method, but rather than calculating the inverse of the entire Hessian (a very large matrix for a density field on the scales of interest), it iteratively updates a pseudo-Hessian as the function is being optimized.

Quasi-Newtonian methods, like LBFGS, are only guaranteed to find extrema for convex optimization problems. For the case of large-scale structure, it was demonstrated that the posterior surface is multimodal at the smallest scales but not modes probed by next-generation large-scale structure surveys (Feng et al. 2018). This optimization technique was previously implemented for the case of cosmological shear measurements and cosmic microwave background reconstruction, finding fast numerical convergence even in very high dimensional parameter space (Horowitz et al. 2019), as well as in dark matter-only models (Seljak et al. 2017; Feng et al. 2018).

Our implementation is based on the `vmad` framework,⁷ an extension of the `abopt` framework used to perform similar reconstructions from late-time galaxy fields (Modi et al. 2018). This framework allows very fast reconstruction convergence; for the cases studied in this work, each reconstruction took approximately 5 CPU hr.

2.3. Response Function and Forward Model

Optimization over the initial density skewers requires defining a differential forward model that will allow us to define a χ^2 problem as in Equation (3) and gradient function as in Equation (5). This procedure is summarized in Figure 1.

2.3.1. Forward Evolution

Following the work of Feng et al. (2018), we first use Lagrangian perturbation theory (LPT) to evolve the initial conditions while the field is still almost entirely linear. We do this until $z = 100.0$, at which point we then use five steps of FastPM (Feng et al. 2016)⁸ to evolve until redshift $z = 2.5$.

There are fundamental limitations due to using a particle mesh framework with limited time steps and constraints imposed by the speed requirements for optimization. As discussed in Feng et al. (2016) and Dai et al. (2018), halos are not fully virialized when using these methods. This will not affect our ability to reconstruct structures on $>1 h^{-1} \text{Mpc}$ scales relevant for current and upcoming surveys. Similarly, we use a particle resolution of 128^3 for our reconstructions to allow fast optimization.

We use the $z = 2.5$ particle positions to generate a density field and infer the hydrogen Ly α optical depth using the FGPA, with $T = T_0(\rho/\bar{\rho})^{(\gamma-1)}$ with slope $\gamma = 1.6$ (Lee et al. 2015). Note that we calculate the optical depth first, which is then redshift space-distorted using the inferred velocity field. Then we compute the flux $F = \exp(-\tau)$ and select lines of sight matching the positions of the mock observations. The skewers are then smoothed with a $\sigma = 1.0 \text{Mpc/h}$ Gaussian filter to imitate spectrographic smoothing; this is a conservative estimate for upcoming surveys.

⁶ Note that we drop the \mathbf{R}'' term, as it fluctuates with mean zero and does not appreciably affect the optimization.

⁷ <https://github.com/rainwoodman/vmad>

⁸ <https://github.com/rainwoodman/fastpm>

Table 1
Mock Data Sets for Reconstructions

Name	N -body Code	LOS Separation (h^{-1} Mpc)	LOS Density (deg^{-2})	S/N_{\min} (\AA^{-1})	S/N_{\max} (\AA^{-1})	Description
T-TomODESI	TreePM	3.7	363	1.4	4.0	Dedicated survey with DESI spectrograph (4 m)
T-CLA/PFS	TreePM	2.4	863	1.4	10.0	Survey with 8–10 m class telescopes
T-30+T	TreePM	1.0	4970	2.8	10.0	Survey with 30 m class telescopes
F-CLA/PFS	FastPM	2.4	863	1.4	10.0	Same as T-CLA/PFS but using FastPM

Table 2
Cosmic Web Recovery at $z = 2.5$ (Eulerian Comparison)

Mock Data	Pearson Coefficients			Volume Overlap (%)			
	λ_1	λ_2	λ_3	Node	Filament	Sheet	Void
T-TomODESI	0.62	0.58	0.66	28	51	58	47
T-CLA/PFS	0.78	0.75	0.77	45	59	67	67
T-30+T	0.94	0.94	0.95	74	80	82	81

2.3.2. Overview of Forward Model

1. Initialize a Gaussian random field (the signal field).
2. Evolve field forward to $z = 2.5$ with FastPM.
3. Use FGPA to calculate a real-space Ly α optical depth.
4. Use the line-of-sight velocity field to shift the Ly α optical depth to redshift space.
5. Exponentiate the redshift-space optical depth field to get the transmitted flux field.
6. Select skewer sight lines from the redshift-space flux field.
7. Convolve skewers with Gaussian spectrograph smoothing.

3. Mock Data Sets

While the FastPM code provides a rapid convergence toward the underlying density field within the TARDIS framework, to rigorously test our reconstruction, we apply the formalism to mock data generated from well-characterized, large-volume, high-resolution N -body simulations. We therefore use a simulation volume run with TreePM (White 2002; White et al. 2010), which has been used for other work on Ly α forest tomography (Stark et al. 2015; Krolewski et al. 2018). This simulation uses 2560³ particles in a box with $256 h^{-1}$ Mpc along each dimension, with cosmological parameters $\Omega_m = 0.31$, $\Omega_b h^2 = 0.022$, $h = 0.677$, $n_s = 0.9611$, and $\sigma_8 = 0.83$. The initial conditions are generated using second-order LPT to $z_{ic} = 150$ and then further evolved using the TreePM code. The output was taken at $z = 2.5$ and 0 for comparison, and a $z = 2.5$ Ly α absorption field was generated using the FGPA with $T_0 = 2.0 \times 10^4$ and $\gamma = 1.6$.

We generated mock skewers from $(64 h^{-1} \text{Mpc})^3$ subvolumes of the TreePM simulation with different survey parameters to mimic various ongoing and upcoming IGM tomography surveys; these are summarized in Table 2. The most important survey parameter is the mean sight-line separation, or, equivalently, the areal density of background sources on the sky. This is typically set by the overall sensitivity of the telescope/instrument combination and desired integration time, but in this work, we simply quote the sight-line separation and minimal signal-to-noise ratio (S/N) for each survey; we refer the reader to Lee et al. (2014) for a more

detailed discussion with respect to observational strategy. The CLAMATO survey (Lee et al. 2018), which is currently ongoing with the Keck I telescope, achieves a mean separation of $2.4 h^{-1}$ Mpc between sight lines (albeit over a small footprint of 0.16 deg^2 at present). An IGM tomography program is currently being planned for the upcoming PFS (Sugai et al. 2015), which should achieve comparable spatial sampling to CLAMATO but over a much larger area ($\sim 15 \text{ deg}^2$). Further into the 2020s, 30 m class telescopes such as the TMT, ELT, and GMT will allow much greater sight-line densities by observing fainter background sources. While the exact parameters of future IGM tomography surveys on TMTs will depend on instruments that are largely still under early development, for now we assume a $1 h^{-1}$ Mpc sight-line separation. We also study a hypothetical dedicated IGM tomography program carried out with the DESI spectrograph, which is currently being installed on the 4 m Mayall telescope (Levi et al. 2013). Note that this is *not* the quasar Ly α forest survey currently being planned as part of the DESI cosmology program, which, at only $\sim 50\text{--}60 \text{ deg}^{-2}$, is far too sparse for cosmic web analysis. While the DESI instrument offers 5000 fibers over a 7.5 deg^2 field of view, we assume that 10% of the fibers will be dedicated to sky subtraction and a $1.7\times$ overhead factor in background sources will be targeted to maintain the specified sight-line density over a finite redshift range of $\delta z = 0.3$ (Lee et al. 2014). This implies a mean sight-line separation of $3.7 h^{-1}$ Mpc for a dedicated DESI tomography program.

For pixel noise, we assume Gaussian random noise that varies among different skewers but is constant along each skewer. To simulate a realistic distribution of skewer S/N , we follow the prescriptions in Stark et al. (2015) and Krolewski et al. (2018) and draw the individual skewers' S/N s from a power-law distribution with minimum value S/N_{\min} (i.e., $dn_{\text{los}}/dS/N \propto S/N^\alpha$) and spectral amplitude $\alpha = 2.7$. The S/N_{\min} is the same for both the DESI and CLAMATO/PFS mocks, since it reflects the actual minimal S/N in the real CLAMATO data, but for 30 m class telescopes, Lee et al. (2014) found that the S/N needs to be increased, as the tomographic reconstruction is no longer limited by the shot noise from finite skewer sampling. To be conservative, we also

Table 3
Cosmic Web Recovery at $z = 0$ (Lagrangian Comparison)

Mock Data	Pearson Coefficients			Volume Overlap (%)			
	λ_1	λ_2	λ_3	Node	Filament	Sheet	Void
T-TomODESI	0.58	0.40	0.34	20	42	54	31
T-CLA/PFS	0.70	0.54	0.47	41	50	54	37
T-30+T	0.82	0.67	0.54	48	55	62	46

impose a maximal S/N for all mock data sets (Lee et al. 2018), as specified in Table 1.

In addition to the random pixel noise, we add continuum error to account for the difficulty in identifying the intrinsic quasar or galaxy continuum. The ability to estimate the continuum is dependent on the S/N of the skewers, and we apply the fitted continuum error distribution of Krolewski et al. (2018) to our mock skewers. In particular, we take our observed flux to be

$$F_{\text{obs}} = \frac{F_{\text{sim}}}{1 + \delta_c}, \quad (7)$$

where δ_c is taken from an underlying Gaussian distribution with width σ_c depending on S/N along each skewer as

$$\sigma_c = \frac{0.205}{\text{S/N}} + 0.015, \quad (8)$$

where the constants are fitted from data from the CLAMATO field. While we add continuum errors to our mock spectra, we do not directly model continuum errors in TARDIS. This could be included as an off-diagonal term in the covariance matrix in future work.

In addition to the TreePM run, we have also generated mock skewers from FastPM using the exact same technique and parameters as in our forward model. This serves to isolate effects caused by known limitations of FastPM to resolve small-scale halo properties, as well as provide a tool for rapid consistency checks. These are applied toward the discussions regarding the code convergence in Appendix A and the method’s sensitivity to astrophysical assumptions in Appendix B.

4. Results

We apply the TARDIS method, described in Section 2, to the mock data set generated as described in Section 3. Broadly, we are interested in both how well we reconstruct cosmic structures at the observed redshift ($z = 2.5$) and the late-time ($z = 0$) fate of those structures. TARDIS solves for the initial density fluctuations within the volume, which one can then use to initialize a simulation using any cosmological N -body or hydrodynamical code to study the cosmic evolution of the large-scale structure realization. For convenience, however, in this paper, we continue to use FastPM to study the gravitational evolution of the TARDIS realizations at both $z = 2.5$ and 0. The $z = 2.5$ field is simply the best-fit TARDIS solution, whereas to get to $z = 0$, we evolve FastPM by another five steps. We then compare the resulting fields with the “truth” from the fiducial TreePM simulation volume.

Examples of reconstructed fields for initial density, $z = 2.5$ matter density and Ly α flux, and line-of-sight velocity and $z = 0$ matter density for T-CLA/PFS are shown in Figure 2. In comparison with the “true” fields, there is a strikingly good recovery of the overall filamentary backbone of the $z = 2.5$

matter density field, as well as the overall distribution of the velocity field. However, the TARDIS reconstruction appears to underestimate the overall amplitude of the density field, with less prominent density peaks in both the initial conditions and $z = 2.5$ matter density. As expected, the underestimated matter power propagates through to the evolved density field at $z = 0$, where the density peaks in the reconstruction are much less prominent than the true underlying density.

The underestimated matter amplitude appears to be a result of the reconstruction method and can be seen when we compare the reconstructed initial fluctuation power spectrum with that used to generate the “true” TreePM simulation volume (Figure 3). There is a shortfall in the recovered power in all of the mock reconstructions, especially on scales below the mean sight-line density of the mock data, but also on larger scales. This gets worse with the reduced sight-line density of the T-TomODESI reconstruction, while conversely, the improved sight-line sampling of the T-30+T mock allows a better job of recovering the true power spectrum, although there is still a shortfall at all scales. This is possibly due to the fact that the Ly α forest absorption blends and saturates in matter overdensities. In particular, at a fixed noise level, the Ly α forest features have a higher density resolution at lower absorption levels than at higher absorption levels due to the exponential FGPA mapping. For example, the optimization algorithm can distinguish between a 1σ and 2σ overdensity at higher significance than a 10σ and 11σ overdensity at a given flux noise level. While it might be possible to correct for this reduced power in the initial density fluctuations, this is a nontrivial process that we defer to an upcoming paper that will focus on modeling galaxy protoclusters within the TARDIS framework. It is also possible to adjust for this nonlinear noise bias at the power-spectrum level within a response formalism (Seljak et al. 2017; Horowitz et al. 2019).

Nevertheless, TARDIS appears to do a reasonable job in recovering the moderate-density cosmic web as seen in Figure 2. We thus focus on the large-scale cosmic web and compare the performance of TARDIS across cosmic time.

4.1. Classification of the Cosmic Web

For quantitative comparison of the large-scale structure recovery in TARDIS, we use the deformation tensor cosmic web classification of Krolewski et al. (2017) and described in Lee & White (2016), which was inspired by Bond et al. (1996), Hahn et al. (2007), and Forero-Romero et al. (2009). While there exist other cosmic web classification algorithms (see summary in Cautun et al. 2014), the deformation tensor approach has a strong physical interpretation within the Zel’dovich approximation (Zel’dovich 1970) and allows easy comparison to previous work in the context of Ly α forest tomography. However, in contrast to Lee & White (2016) and Krolewski et al. (2017), who measured the eigenvalues and eigenvectors of Wiener-filtered maps of the Ly α transmitted




Figure 1. Schematic illustration of our forward model (see Section 2.3.2). The underlying field we are optimizing for is the initial matter density field (left). The output of our forward model is the Ly α flux skewers probing the observational volume at the same positions as the data.

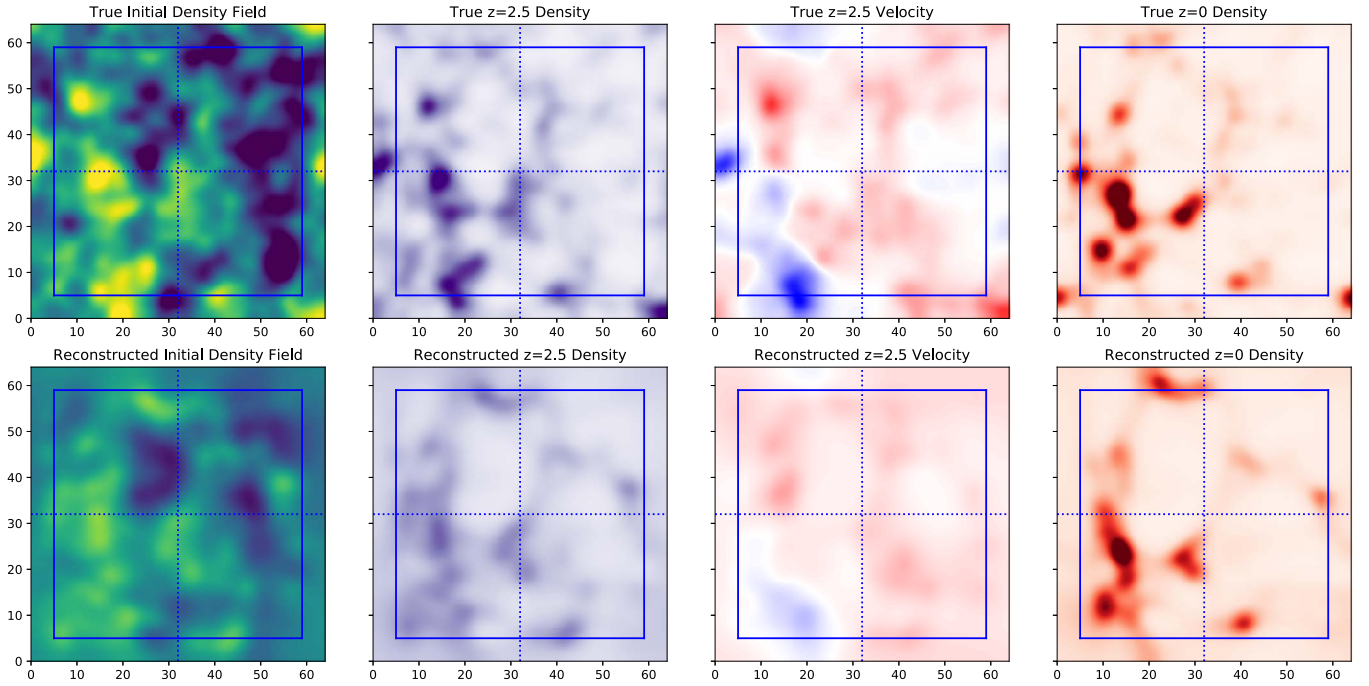


Figure 2. The reconstructions of various recovered quantities for the F-CLA/PFS mock data set, smoothed at $2 h^{-1}$ Mpc, are shown on the bottom row. The true corresponding fields from the FastPM simulation are shown on the top. In all panels, we project along a $5 h^{-1}$ Mpc slice. The region outside the solid blue box is masked in our analysis, while the dotted lines are merely to guide the eye. We find that the large-scale features are qualitatively captured well in the reconstructions.

flux, in this work, we directly measure the eigenvalues and eigenvectors of the dark matter fields reconstructed with TARDIS, which have first been smoothed with an $R = 2 h^{-1}$ Mpc Gaussian kernel.

The eigenvectors and eigenvalues of the deformation tensor relate directly to the flow of matter around that point in space; matter collapses along the axis of the eigenvector when the associated eigenvalue is positive and expands when it is negative. Points with three eigenvalues above some nonzero threshold value λ_{th} (as in Forero-Romero et al. 2009) are nodes (roughly corresponding to (proto)clusters), two values above

λ_{th} are filaments, one value above λ_{th} is a sheet, and zero values above λ_{th} are voids. The deformation tensor, D_{ij} , is defined as the Hessian of the gravitational potential, Φ , i.e.,

$$D_{ij} = \frac{\partial^2 \Phi}{\partial x_i \partial x_j}, \quad (9)$$

or, equivalently, in Fourier space in terms of the density field, δ_k , as

$$\tilde{D}_{ij} = \frac{k_i k_j}{k^2} \delta_k. \quad (10)$$

This tensor is then diagonalized to obtain the eigenvalues \hat{e}_1 , \hat{e}_2 , and \hat{e}_3 at each point on our spatial grid, ordered such that their corresponding eigenvalues are $\lambda_1 > \lambda_2 > \lambda_3$ (i.e., to demand that collapse first occurs along \hat{e}_1). Note that one could use the velocity field from the reconstruction itself to determine the flow at each point (e.g., Libeskind et al. 2013; Pahwa et al. 2016) instead of relying on the Zel’dovich approximation used in the classification here. We use the deformation tensor in order to stay consistent with past IGM tomography work (Lee & White 2016; Krolewski et al. 2017). Cosmic web directions for our reconstructed field are thus defined by the eigenvectors with associated eigenvalues used to classify the cosmic web.

We follow Lee & White (2016) and Krolewski et al. (2017) and define our threshold value λ_{th} for each simulated field such that the voids occupy 21% of the total volume at $z = 2.5$ and 27% at $z = 0$ (inspired by the redshift evolution in Cautun et al. 2014). The void fraction is somewhat arbitrary in the analysis as long as it is consistent between the mock reconstructions and true density field used for comparison.

4.2. Matter/Flux Density at $z \sim 2.5$

We compare the recovery of $z = 2.5$ Ly α flux to previously standard Wiener-filtering techniques. As we are assuming the FGPA, this reconstructed flux can be mapped directly to the density field. While past work on Wiener-filtered IGM tomographic maps (Cauci et al. 2008; Lee et al. 2018) smoothed the field on $1.4\times$ the mean sight-line spacing, for these comparisons, we smooth the respective matter fields with a $\sigma = 2 h^{-1}$ Mpc Gaussian kernel. The smaller smoothing scale is appropriate for our work because our method should be better able to infer nonlinear and semilinear structure between sight lines. For all plots, we treat the field in real space (without redshift-space distortions), since our optimization is over the initial real-space density field.

The reconstructed matter density fields from the various mock IGM tomography surveys (summarized in Table 2) are shown in the first row of Figure 4 in comparison with the true density field from the TreePM simulation. In all cases, they are smoothed with an $R = 2 h^{-1}$ Mpc Gaussian kernel. On large scales, the reconstructed density fields are well matched in terms of voids and sheets, but CLAMATO/PFS data miss out on some prominent filamentary structures and nodes as a consequence of the underestimated matter amplitude. The 30 + m telescopes, on the other hand, yield a matter density reconstruction with excellent fidelity over the entire volume.

We next calculate the characteristic eigenvalues of the deformation tensor, as described in Section 4.1, on the smoothed matter density fields. The scatter of the eigenvalues relative to the true underlying eigenvalues is plotted in Figure 6. This reflects how well we recover the amplitude of curvature of the matter density field along each cosmic web direction. The distribution of all three eigenvalues is unbiased relative to the truth, albeit with more scatter in the case of the sparser CLAMATO/PFS reconstruction. We quantify the agreement in terms of Pearson correlation coefficients, showing the scatter from a linear trend in Table 2. These show a strong correlation between the reconstructed and true eigenvalues, ranging from $r = [0.78, 0.75, 0.77]$ in recovering the three eigenvalues $[\lambda_1, \lambda_2, \lambda_3]$ for CLAMATO/PFS to the excellent




Figure 3. Top: power spectra of the reconstructed initial conditions for various experimental configurations, with the true initial conditions shown for comparison. Bottom: cross-correlation coefficient, $P_{RT}/\sqrt{P_{RR}P_{TT}}$, where P_{RT} is the cross-power between the true field and reconstructed field, P_{RR} is the reconstructed power spectrum, and P_{TT} is the true power spectrum. As the number of sight lines and spectral noise improve, power-spectrum reconstruction improves; however, there remains a residual noise bias for realistic experiments.

reconstruction of the 30 m class telescopes with correlation coefficients of $r = [0.94, 0.94, 0.95]$.

Next, we classify the each point within the density field as void, sheet, filament, or node, depending on how many of the eigenvalues are greater than the threshold value, $\lambda_i > \lambda_{\text{th}}$. In the true matter density field, we find that [22%, 50%, 25%, 3%] of the volume is occupied by voids, sheets, filaments, and nodes, respectively; by construction, the reconstructed matter fields show similar volume occupation fractions to within $\pm 2\%$. The volume overlap fractions between cosmic web classifications in the mock data reconstructions compared to the true matter field are listed in Table 2; these do not include a buffer region of $5 h^{-1}$ Mpc near the edge of the volume where we expect to be contaminated by boundary effects. For the CLAMATO/PFS mock reconstructions, the volume overlap fractions are $\sim 60\text{--}62\%$ for the sheets and voids, declining to 32% for the nodes. It is unsurprising that the nodes are more challenging to recover, since they occupy such a small fraction (3%) of the overall density field. These numbers are, on the surface, comparable to those found by Krolewski et al. (2017; their Table 1) for a similar CLAMATO-like mock data set but in fact somewhat better, since we are probing the matter field

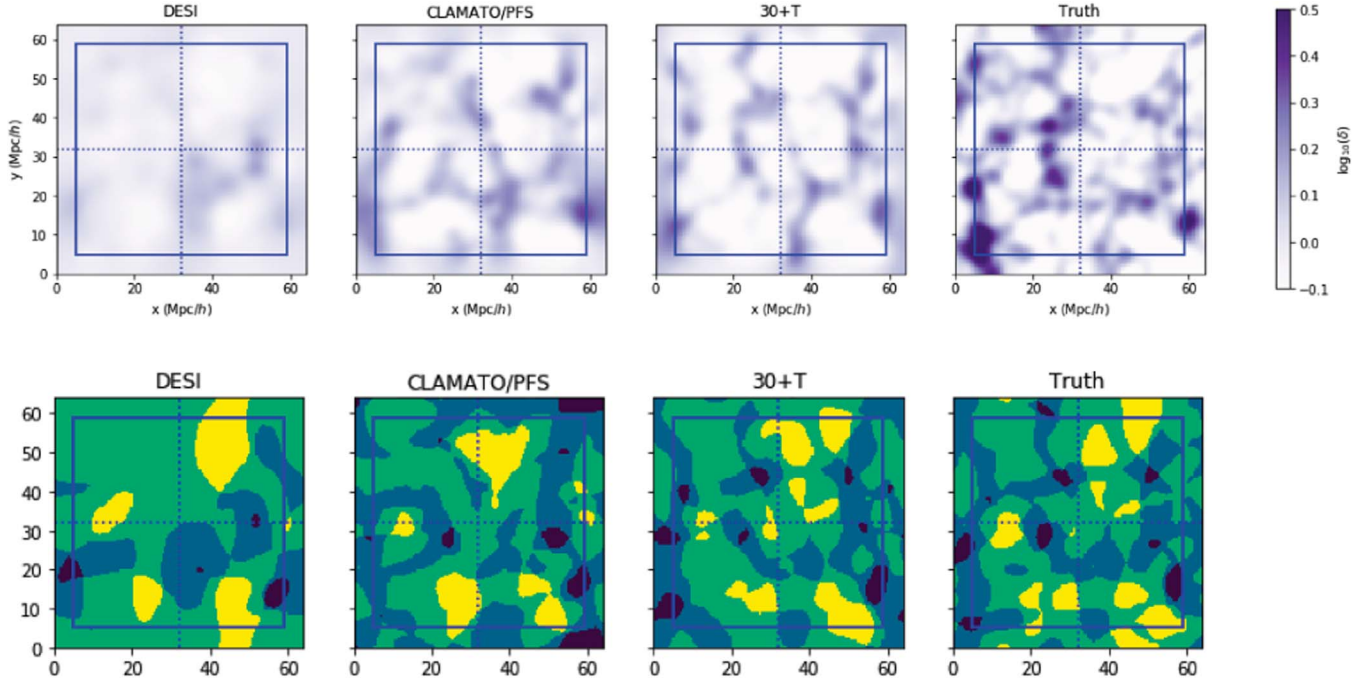


Figure 4. Comparison of the $z = 2.5$ reconstructed cosmic structures as classified by their eigenvalues from T-TomODESI, T-CLA/PFS, and T-30+T vs. the true $z = 2.5$ density field for an xy -slice. Fields have been smoothed by an $R = 2 h^{-1}$ Mpc Gaussian kernel. Top: matter density. Bottom: classification of cosmic structure. Dark blue indicates node, light blue indicates filament, green indicates sheet, and yellow indicates void. The region outside the solid blue box is masked in our analysis, while the dotted lines are to guide the eye. We find that our classification captures the visual appearance of the cosmic web well and that the recovered structure improves as the number of sight lines increases and noise decreases.




Figure 5. Probability density function showing the dot product of the eigenvectors from cosmic web reconstruction vs. the true cosmic web for various experimental configurations. Here $\cos \theta = 1.0$ indicates that the cosmic web structures are oriented the same way, while $\cos \theta = 0.0$ indicates perpendicular alignment. The horizontal dashed line indicates the expected distribution for randomly aligned structure. In T-30+T, the recovery of the cosmic web structure is nearly perfect, with only very slight misalignments, on average.

directly on $2 h^{-1}$ Mpc scales, whereas Krolewski et al. (2017) were evaluating the $\text{Ly}\alpha$ transmission field over coarser ($4 h^{-1}$ Mpc) scales in the equivalent case. This improvement is due to the fact that TARDIS incorporates the physics of gravitational evolution into its reconstructions, in contrast with Wiener filtering, which only assumes a correlation function. The 30 m class reconstruction, as expected, fares even better thanks to its finer sight-line sampling, with the voids, sheets, filaments, and nodes overlapping [81%, 82%, 80%, 74%] with the true matter density cosmic web.

To further illustrate the fidelity of the recovery, Figure 7 shows the confusion matrix, evaluated at all grid points in our volume, between the true cosmic web from the simulation and our reconstructions, finding good agreement. Overall,

we find that 80%, 60%, and 53% of the total observed volume is properly classified for T-30+T, T-CLA/PFS, and T-TomODESI, respectively. Allowing misclassification by a structurally adjacent type (i.e., void to sheet, sheet to void/filament, filament to sheet/node, and node to filament), the agreement goes up to 98%, 96%, and 95%, respectively. We also examine the eigenvector recovery by computing the dot product between the eigenvectors recovered from the reconstructions with those at the same Cartesian point in the true matter density field⁹; with a good recovery, the recovered eigenvectors would be well aligned with the true eigenvectors

⁹ These values only include structure in the observed region, excising an additional buffer of $2 h^{-1}$ Mpc near the survey boundary.

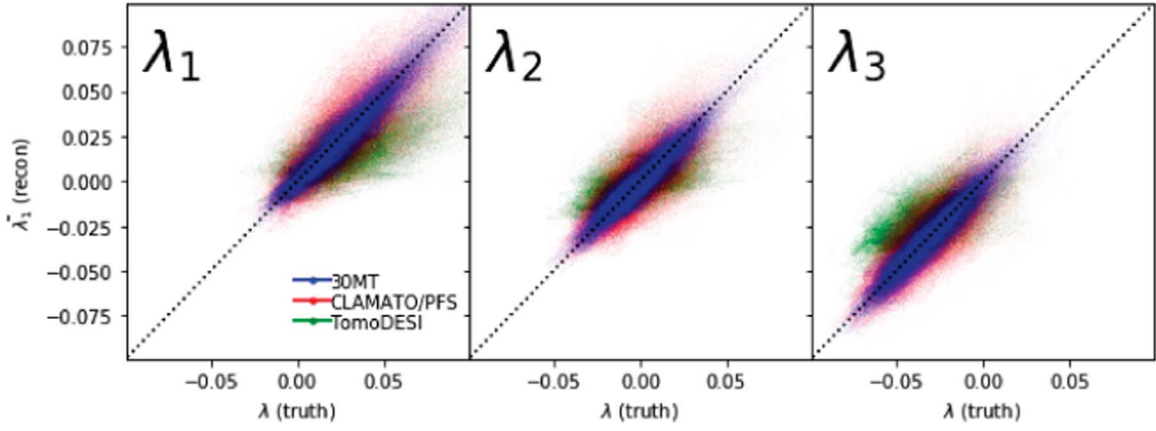


Figure 6. Point-by-point distribution of the eigenvalues inferred from the deformation tensor, smoothed by $2 h^{-1}$ Mpc. The magnitude of each eigenvalue indicates the magnitude of compression along the associated eigenvector. As sight lines increase and noise decreases, not only is there less scatter in the eigenvalues, but there is also less overall bias.




Figure 7. Confusion matrix for cosmic structures at $z = 2.5$ in real space showing the reconstructed fraction printed over each cell. For T-30+T, we correctly identify approximately 80% of the volume.

and lead to dot products of order unity. These are shown in Figure 5. For $[\hat{e}_1, \hat{e}_2, \hat{e}_3]$, we find average alignment cosine angles of $[0.80, 0.70, 0.80]$ for T-TomoDESI, $[0.87, 0.79, 0.80]$ for T-CLAMATO/PFS, and $[0.96, 0.92, 0.96]$ for T-30+T. This is again comparable to the results derived from Wiener-filtered flux maps in Krolewski et al. (2017) for the CLAMATO/PFS case but probing smaller scales.

4.3. Matter Density at $z = 0$

A main motivation for the TARDIS framework is inferring the late-time fate of structures and constituent galaxies found in regions observed by Ly α forest tomography. As an output of our model, we further evolve the particle field to $z = 0$ in order to study the reconstruction. We compare this evolved field with the TreePM “truth” at $z = 0$. The true underlying field contains cosmic structures with a mass fraction $[0.15, 0.49, 0.31, 0.05]$ and volume fraction $[0.02, 0.28, 0.48, 0.22]$ for [nodes, filaments, sheets, voids], respectively.

Eulerian (real) space provides a qualitative picture of the the structures reconstructed in this limit. In Figure 9 (top), we show the matter field and cosmic web reconstructed for different survey mock data. While they are qualitatively similar, as described in Section 4.2, the peaks of the $z = 2.5$ density field

are poorly reconstructed for realistic survey parameters. This results in significant drift of the Eulerian space structures and makes point-by-point comparisons difficult. This can be seen in Figure 9 (bottom), where the qualitative structure is quite similar, especially for 30+T, but the exact positions of nodes and filaments are slightly different relative to the true matter field. This leads to unsatisfactory cosmic web recovery when evaluated in the same way as $z = 2.5$.

However, the reconstructions’ cosmic web fidelity at $z = 0$ is a somewhat abstract concept, since the Eulerian matter density field is not accessible via any observations. Instead, we can evaluate the reconstructed field in Lagrange space, i.e., tracking the $z = 0$ environments sampled by test particles observed at $z = 2.5$. Since we expect galaxies to act roughly like test particles in the large-scale gravitational potential, this provides a direct connection to understanding the late-time fate of $z \approx 2.5$ galaxies observed in the same volume as the Ly α tomography data. We test this by the following. From the $z = 2.5$ density field reconstructed from the mock data reconstructions with TARDIS/FastPM, we select a set of test particles at Eulerian real-space positions $[x_{z25,i}, y_{z25,i}, z_{z25,i}]$, track them to their $z = 0$ Eulerian positions $[x_{z0,i}, y_{z0,i}, z_{z0,i}]$, and then evaluate their cosmic web eigenvalues and classifications (on the Eulerian real-space grid). From the TreePM “true”




Figure 8. Displacement fields from $z = 2.5$ to 0 for random matched particles between the TreePM truth and that reconstructed in the mock observed volume. The underlying $z = 0$ density field is also shown. TARDIS is able to well reconstruct the movement and $z = 0$ environment of test particles identified at $z = 2.5$.

matter density field at $z = 2.5$, we find matching test particles at the same Eulerian positions $[x_{z=2.5,i}, y_{z=2.5,i}, z_{z=2.5,i}]$ and again track them to their $z = 0$ positions and environments. This process is visualized in Figure 8, where we show the displacement vectors for particles from the reconstructions versus matched particles from the TreePM simulation evolved to $z = 0$.

The results of this exercise are shown in the $z = 0$ Lagrangian confusion matrix in Figure 10 and summarized in Table 3. For CLAMATO/PFS, we are able to successfully predict the $z = 0$ environment sampled by the test particles with $\sim 40\%$ – 50% fidelity, while this increases slightly to $\sim 50\%$ – 60% in the case of T-30+T. In both cases, $>90\%$ of the particles are predicted to lie within ± 1 of the correct cosmic web classification, with the exception of CLAMATO/PFS node particles that are misidentified as sheet particles in 15% of cases. Nonetheless, this demonstrates the remarkable ability of TARDIS to infer the $z = 0$ environment of galaxies observed at $z = 2.5$ across 10 Gyr of cosmic time.

5. Conclusion

We present the first use of initial density reconstruction on densely sampled Ly α forest data sets (often called “IGM tomography”) and show that by using this technique, we are able to accurately reconstruct large-scale properties within the survey volume over a range of scales. In particular, we are able to recover the characterization and orientation of the cosmic web at $z = 2.5$ in terms of the deformation eigenvalues and eigenvectors assuming mock data that reflect upcoming and future multiplexed spectroscopic instrumentation. In addition, we are able to recover the qualitative structure of the observed structures at late time, $z = 0$. We have also shown that the inferred flux maps from TARDIS are more accurate and have less variance than those from Wiener filtering. Excitingly, we argue that we would be able to predict the late-time environments of $z \approx 2.5$ galaxies that are coeval with our reconstructed IGM tomography volume. This provides a promising and direct route to studying galaxies and active galactic nuclei in the context of their surrounding cosmic web. For example, we would be able to identify the direct progenitors of $z = 0$ filament galaxies and study their $z = 2.5$ galaxy properties. While we are currently limited by noise levels and sight-line spacing, in future papers, we will explore ways to correct for underestimated fluctuation amplitude as a function of survey parameters.

While only explored indirectly (through $z = 0$ density reconstruction), a direct product of this technique is the particle velocity field at $z = 2.5$, which could have significant uses in informing astrophysical processes, as well as cosmological constraints. For example, it could allow accurate estimation of velocity dispersions in high-redshift protoclusters, which is currently uncertain due to challenges in disentangling galaxy peculiar motions from the large-scale Hubble expansion (Wang et al. 2016b; Cucciati et al. 2018; Topping et al. 2018). More generally, the velocity field reconstruction extends over the entire field and could be a useful addition beyond velocity fields from galaxy redshift-space distortions and kinetic Sunyaev–Zel’dovich effects (Sugiyama et al. 2017). While one might hope to use this reconstruction for constraining other exotic physics (such as using void velocity profiles to provide constraints on modified gravity (Falck et al. 2018) and neutrino mass (Massara et al. 2015)), the nature of our forward model will restrict the reconstructed maps to obey a Λ CDM cosmology. If alternative models were implemented efficiently into an N -body solver, their validity could be tested by comparing the best-fit likelihood values.

In this work, we have held the astrophysical and cosmological parameters constant. A more complete treatment would require varying these jointly with the underlying field; however, we view this as unnecessary at this point, since existing data cover a very limited volume with minimal cosmological constraining power. For next-generation surveys, which will greatly expand the footprint covered, it will be required to jointly vary these parameters as well. Within the FGPA approximation, the astrophysical parameters are not a significant limitation, since there are only two global parameters of interest (A_0 , γ), and our optimization scheme is fast enough that a naive Markov Chain Monte Carlo sampling would be sufficient to explore this parameter space. We explored the sensitivity of the reconstruction with respect to the absorption model in Appendix B.

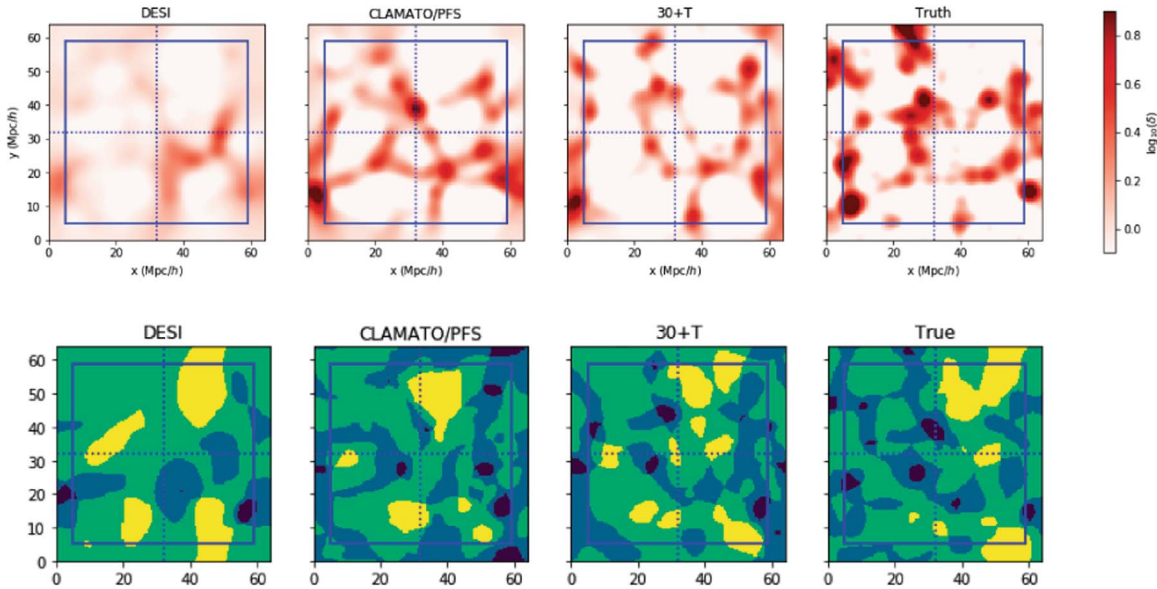


Figure 9. Comparison of the $z = 0$ inferred cosmic structure in Eulerian space from T-TomoDESI, T-CLA/PFS, and T-30+T vs. the true $z = 0$ density field. Fields have been smoothed at $2 h^{-1}$ Mpc. Top: matter density. Bottom: classification of cosmic structure. Dark blue indicates node, light blue indicates filament, green indicates sheet, and yellow indicates void. While the exact location of the structure is poorly constrained in real space, the overall structure is quite similar, especially with tight sight-line spacing.




Figure 10. Confusion matrix for cosmic web structures at $z = 0$ in Lagrange space (i.e., comparing particles with those matched in $z = 2.5$ positions) shown with the reconstructed fraction printed over each cell. While structure is not as well classified as at $z = 2.5$, classifications are approximately correct and tend toward morphologically similar environments. For comparison, the mass fractions residing in $z = 0$ nodes, filaments, sheets, and voids are [0.15, 0.49, 0.31, 0.05], respectively.

Our focus in this work is on reconstructing the moderate-density large-scale structure within the survey volume, and we demonstrated that we were able to recover qualitative structure over a range of scales. Going forward, it would be useful to study how well similar techniques would work on reconstructing halo-scale (i.e., $\leq 1 h^{-1}$ Mpc) structure, such as stacked halo and void profiles. However, going to this small-scale regime reconstruction will be limited by the specific astrophysical processes within the high-density regions where the FGPA will no longer hold. In particular, numerical hydrodynamic simulations have shown that there are significant deviations away from a simple temperature-density scaling relationship close to halos, in some cases even showing a turnover of the relationship (Sorini et al. 2018). It should be possible to extend the formalism proposed in this work and treat the variations from FGPA with some additional parameters to be fit for (or marginalized) in this limit, such as was done for galaxy surveys via a bias expansion

(Ata et al. 2015; Kitaura et al. 2016; Jasche & Lavaux 2019). One could also use grid-based approximation methods for baryonic effects (such as Dai et al. 2018) to provide a more precise formation formalism for halo substructure or a more accurate N -body-based approximation than FGPA (Sorini et al. 2016). It would be a natural extension to test this method on mock data generated from the NyX hydrodynamic simulations designed to accurately reproduce $\text{Ly}\alpha$ absorption physics (Almgren et al. 2013; Lukić et al. 2015). Other nontomographic techniques have shown great promise in detecting high-redshift clusters from $\text{Ly}\alpha$ observations (Cai et al. 2016), including a detection of a cluster at $z = 2.32$ (Cai et al. 2017), but these techniques probe scales of $\approx 10 h^{-1}$ Mpc.

On the other hand, additional work is needed to make this reconstruction technique useful for full-scale cosmological analysis. Directly extracting power-spectrum estimates from our reconstructed maps suffers from significant noise bias effects that would make them difficult to apply directly to

constrain cosmological parameters, as well as mode coupling effects due to the complexity of our forward model. Using a response formalism (as in Seljak et al. 2017; Feng et al. 2018; Horowitz et al. 2019) to estimate band powers would be straightforward and require $O(N)$ additional optimization runs to estimate N band powers. However, before using these reconstructions for cosmological analysis, additional considerations are necessary, such as incorporating light-cone effects (i.e., evolution) within the survey volume and including correlated errors within our model. While work in this direction is ongoing, upcoming and proposed Ly α tomography surveys will cover only a small sky fraction and are unlikely to be directly competitive with other cosmological surveys.

For future reconstruction efforts, the combination of galaxy surveys and Ly α tomographic mapping will be necessary in order to probe different redshift ranges with maximum efficiency. By including the galaxy density field in the reconstruction, we will be able to measure overdensities with higher precision than with IGM tomography alone. Furthermore, incorporating baryonic effects from hydrodynamical simulations can show how different components of the IGM trace the cosmic web at different redshifts (Martizzi et al. 2018). This will allow a joint understanding of the galaxy and IGM large-scale structure distribution and how they influence each other.

We appreciate helpful discussions with Uroš Seljak, Zarija Lukic, Chirag Modi, Teppei Okumura, Yu Feng, and David Spergel. B.H. is supported by the NSF Graduate Research Fellowship, award No. DGE 1106400, and JSPS via the GROW program. Kavli IPMU was established by the World Premier International Research Center Initiative (WPI), MEXT, Japan. This work was supported by JSPS KAKENHI grant No. JP18H05868.

This research used the resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under contract No. DEC02-05CH11231.

Appendix A Convergence

An important question with any optimization scheme is the convergence properties of the procedure. This is particularly important for nonlinear processes like structure evolution, where the likelihood surface is non-Gaussian and conceivably nonconvex. We divide the issue into two questions to explore in this appendix: how many iterations are necessary for us to be confident in our reconstruction technique, and how sensitive is the found solution to the initial optimization starting point? For both questions, we explore as a function of scale by looking at the reconstructed transfer function.

It has been shown that in the very low noise limit, the likelihood surface of possible initial conditions is multimodal, i.e., gravitational evolution is a noninjective map from initial conditions to late-time structure (Feng et al. 2018). However, this uncertainty is due to the shell-crossing degeneracy, which is only relevant for small-scale nonlinear structure not observed by even the optimistic configurations considered in this work. To study whether or not there is one “true” solution or whether there exist sufficiently different converged solutions, we perform the optimization analysis for the same mock catalog

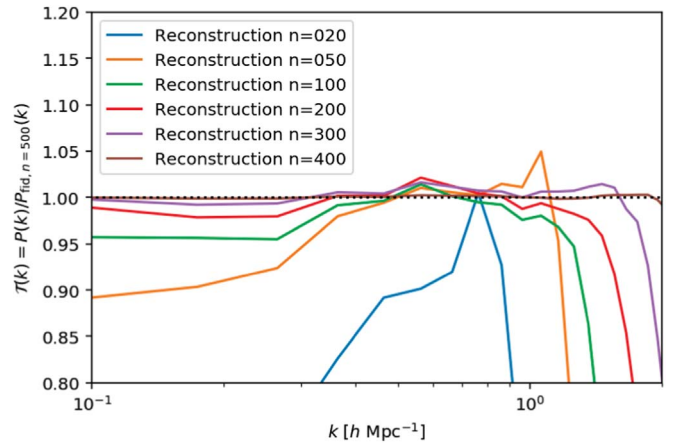


Figure 11. Transfer function with respect to a well-converged solution as a function of iteration number. As the iteration number progresses, smaller and smaller scales converge. In addition, there are larger modes on the order of the box size that are similarly slow to converge.

with different optimization starting points. In particular, we randomly choose a wide range of initial white-noise fields with variance spanning 3 orders of magnitude. We calculate their transfer functions after 100 iterations versus a fiducial “well-converged” solution that underwent 500 iteration steps. Up to the scales of interest for the structures studied in this work, $\approx 1 h^{-1} \text{Mpc}$, we find very good agreement between all different starting points. There are some differences of power on very large scales, reflective of the poor constraining power of modes of order the box size. The number density of modes per uniform bin scales as k^2 , resulting in significantly more weight placed on smaller modes, until the window function (depending on the smoothing scale and sight-line density) creates a sharp cutoff. If these larger modes are of significant interest, an adiabatic optimization scheme could be used wherein the optimization begins first on a smoothed version of the observed field, and then slowly, small-scale power is introduced back in by varying the smoothing scale as the optimization progresses (as done in Feng et al. 2018) or potentially directly using a multigrid preconditioner technique (Smith et al. 2007). Utilization of these techniques will likely be useful when extending this work for cosmological analysis.

The next important consideration is how long our scheme takes to be fully converged. We plot the transfer function as a function of the convergence step in Figure 11. The exact choice of cutoff depends on the scales of interest, but since we are fundamentally limited in the transverse direction by the line-of-sight density and the longitudinal direction by the spectrograph resolution, power above $k = 1.0 h/\text{Mpc}$ is mostly lost to the smoothing operations on our field. By $n = 100$, we find good agreement up to $k = 1.0 (h/\text{Mpc})$, and we use this criterion as an iteration limit in the main work.

Appendix B Sensitivity to Cosmology and Absorption Model

In the main body of this work, we have held cosmological and astrophysical parameters constant for the reconstructions. Here we briefly explore how wrong assumptions about the astrophysics or cosmology would bias our late-time density field.

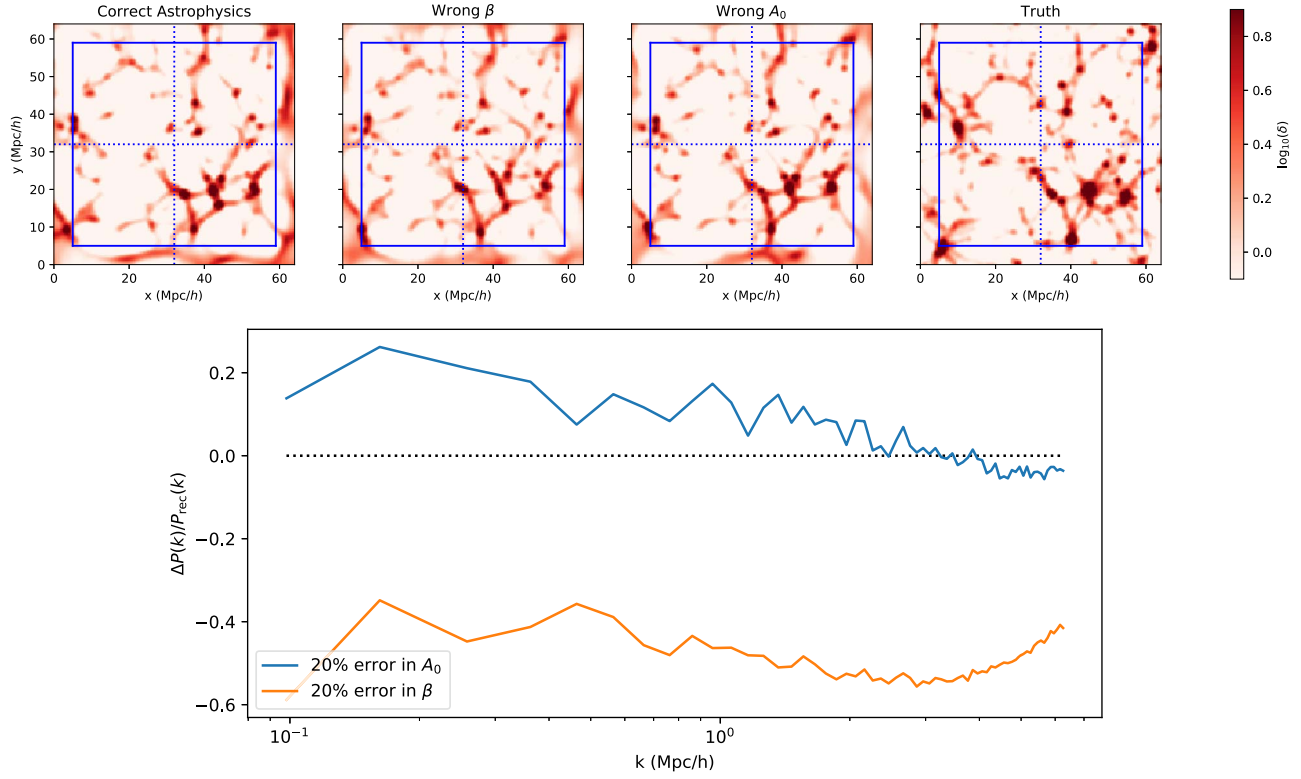


Figure 12. Effect of assuming the wrong astrophysical parameters on the $z = 0$ structure, both for a slice in real space (top) and the power spectra (bottom). Even under wrong astrophysical assumptions, we recover similar cosmic structures.

We use a different mock catalog, T-IDEAL, in order to examine the effects of varying the astrophysical parameters. This catalog has a constant S/N of 50 along each skewer, no continuum error, and a sight-line density twice that of T-30 + T. The idea of this superexperiment is to isolate the effects of the astrophysics from other potential sources of noise in the reconstruction. We perform our reconstructions assuming the “truth” astrophysics from our mock catalog, as well as the wrong the overall flux amplitude, $A_0 = \exp(-T_0)$, and the density scaling exponent, β .

We see the effects of wrong astrophysical assumptions in Figure 12. Even with rather radically different astrophysical assumptions, we find similar qualitative features in the late-time structure. On the power-spectrum level, we find that these wrong assumptions result primarily in a bias offset from the true power spectra. In practice, for surveys of the size studied in this work, it would be easily numerically tractable to sample over these parameters to perform the late-time reconstruction or, alternatively, to use Ly α tomography as a constraint on these parameters.

Appendix C Comparison to Wiener Filtering

A promising aspect of this initial density reconstruction technique is that the reconstructed $z \sim 2$ flux field should be strictly more accurate than that from direct Wiener filtering of the skewers. This is because direct Wiener filtering is a purely statistical process that does not take into account the physical evolution of the system under gravity, which further constrains the observed flux field. In this section, we review the Wiener-filtering technique that we compare our method against. For a

more general discussion of efficient Wiener filtering and associated optimal band-power construction, see Seljak (1998) and Horowitz et al. (2019). For a more thorough description in the context of the Ly α forest, see Stark et al. (2015).

As we are trying to reconstruct the optimal map given the data, we have to take into account the data–data covariance, C_{DD} ; the map–data covariance, C_{MD} ; and the overall map noise covariance, N_{ij} . The reconstructed map can then be expressed in terms of the observed flux, δ_F , as a standard Wiener filter by

$$\delta_F^{\text{rec}} = C_{MD} \cdot (C_{DD} + N)^{-1} \cdot \delta_F. \quad (11)$$

We approximate the covariance by assuming that $N_{ij} = n_i^2 \delta_{ij}$, where n_i is the pixel noise. This neglects the correlated error component of the continuum errors, but this is subdominant to the spectrograph noise and should not appreciably affect our reconstructed maps. The map–data and data–data covariances are therefore approximated as

$$C = \sigma_F^2 \exp \left[-\frac{\Delta x_{\perp}^2}{2l_{\perp}^2} - \frac{\Delta x_{\parallel}^2}{2l_{\parallel}^2} \right]. \quad (12)$$

In order to compare directly to the Wiener filter map, we use the inferred reconstructed flux map from TARDIS.

We apply the Wiener-filtering algorithm to the T-CLA/PFS mock catalog and compare along a number of slices to the TARDIS reconstruction. The results are shown in Figure 13. Overall, there is good agreement between all maps, with certain smaller-scale features better reconstructed in the TARDIS maps than the Wiener-filtered maps.

A well-known feature of the reconstructed maps is the presence of a bias caused by the presence of noise. We correct




Figure 13. Comparison of the true field, TARDIS reconstructed field, and Wiener-filtered field for the T-CLA/PFS mock. In the far left panels, we show the unsmoothed true flux field, with sight lines indicated as blue dots. The blue box indicates the boundaries of the survey, with the blue cross to help aid the eye in matching structures. We smooth the three rightmost column maps on $2 h^{-1}$ Mpc and project over a $5 h^{-1}$ Mpc slice. The recovered flux field is fairly similar between TARDIS and the Wiener filter.

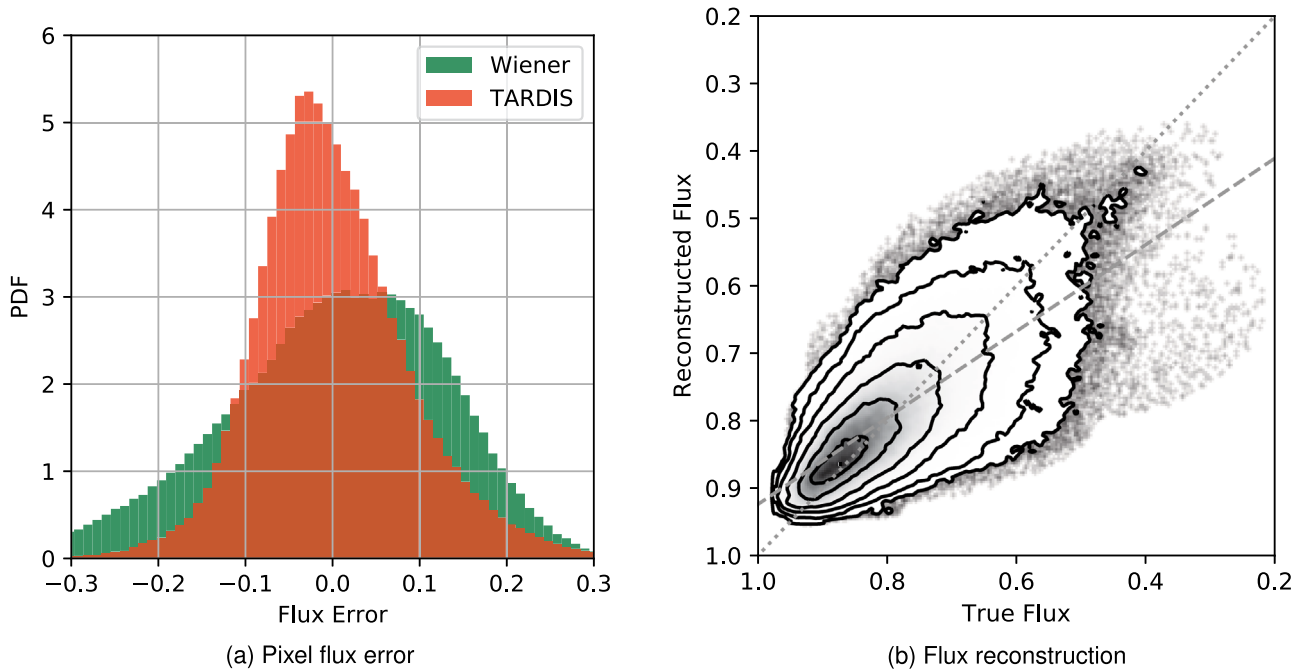


Figure 14. Comparing the flux reconstruction for the T-CIA/PFS mock catalog. For these comparisons, we have taken a central box that is $35 h^{-1}$ Mpc side length in order to mitigate potential boundary effects and smoothed the region with a $1.5 h^{-1}$ Mpc Gaussian. In this plot, we work in redshift space, unlike the other plots in the paper. (a) Comparison of the corrected fluxes for the Wiener filter map and TARDIS reconstruction vs. the true flux. (b) Scatter plot of the TARDIS reconstructed corrected flux vs. the true flux. Also shown is the linear fit of the uncorrected flux (dashed gray line), which was linearly transformed to the $x = y$ dotted line. If interpreted as a flux probability density function, each level surface indicates 0.5σ density. After this linear correction, the resulting TARDIS flux has no significant bias and mildly outperforms a linearly corrected Wiener-filtered map.

for this bias by a linear transformation calibrated from a separate simulated volume. The effect of this transformation is shown in Figure 14(b). We show the reconstructed flux error in Figure 14(a), showing that the TARDIS maps have a smaller flux error variance than the Wiener-filtered maps.

ORCID iDs

Benjamin Horowitz <https://orcid.org/0000-0001-7832-5372>
 Khee-Gan Lee <https://orcid.org/0000-0001-9299-5719>
 Martin White <https://orcid.org/0000-0001-9912-5070>
 Alex Krolewski <https://orcid.org/0000-0003-2183-7021>

References

- Alam, S., Zu, Y., Peacock, J. A., & Mandelbaum, R. 2019, *MNRAS*, **483**, 4501
 Almgren, A. S., Bell, J. B., Lijewski, M. J., Lukić, Z., & Van Andel, E. 2013, *ApJ*, **765**, 39
 Ata, M., Kitaura, F.-S., & Müller, V. 2015, *MNRAS*, **446**, 4250
 Baldauf, T., Seljak, U., Desjacques, V., & McDonald, P. 2012, *PhRvD*, **86**, 083540
 Bond, J. R., Kofman, L., & Pogosyan, D. 1996, *Natur*, **380**, 603
 Cai, Z., Fan, X., Bian, F., et al. 2017, *ApJ*, **839**, 131
 Cai, Z., Fan, X., Peirani, S., et al. 2016, *ApJ*, **833**, 135
 Cauci, S., Colombi, S., Pichon, C., et al. 2008, *MNRAS*, **386**, 211
 Cautun, M., van de Weygaert, R., Jones, B. J. T., & Frenk, C. S. 2014, *MNRAS*, **441**, 2923
 Chiang, Y.-K., Overzier, R. A., Gebhardt, K., & Henriques, B. 2017, *ApJL*, **844**, L23
 Cirasuolo, M., Afonso, J., Carollo, M., et al. 2014, *Proc. SPIE*, **9147**, 91470N
 Conselice, C. J. 2014, *ARA&A*, **52**, 291
 Cucciati, O., Lemaux, B. C., Zamorani, G., et al. 2018, *A&A*, **619**, 49
 Dai, B., Feng, Y., & Seljak, U. 2018, *JCAP*, **11**, 009
 Desjacques, V., Jeong, D., & Schmidt, F. 2018, *PhR*, **733**, 1
 Dressler, A. 1980, *ApJ*, **236**, 351
 Dubois, Y., Pichon, C., Welker, C., et al. 2014, *MNRAS*, **444**, 1453
 Evans, C. J., Puech, M., Barbu, B., et al. 2014, *Proc. SPIE*, **9147**, 914796
 Falck, B., Koyama, K., Zhao, G.-B., & Cautun, M. 2018, *MNRAS*, **475**, 3262
 Feng, Y., Chu, M.-Y., Seljak, U., & McDonald, P. 2016, *MNRAS*, **463**, 2273
 Feng, Y., Seljak, U., & Zaldarriaga, M. 2018, *JCAP*, **7**, 043
 Forero-Romero, J. E., Hoffman, Y., Gottlöber, S., Klypin, A., & Yepes, G. 2009, *MNRAS*, **396**, 1815
 Guzzo, L., Scodreggio, M., Garilli, B., et al. 2014, *A&A*, **566**, A108
 Hahn, O., Carollo, C. M., Porciani, C., & Dekel, A. 2007, *MNRAS*, **381**, 41
 Horowitz, B., Seljak, U., & Aslanyan, G. 2019, *JCAP*, **10**, 035
 Jasche, J., & Lavaux, G. 2019, *A&A*, **625**, 64
 Jasche, J., & Wandelt, B. D. 2013, *MNRAS*, **432**, 894
 Johns, M., McCarthy, P., Raybould, K., et al. 2012, *Proc. SPIE*, **8444**, 84441H
 Kartaltepe, J., Casey, C., Dickinson, M., et al. 2019, arXiv:1903.05026
 Kauffmann, G., White, S. D. M., Heckman, T. M., et al. 2004, *MNRAS*, **353**, 713
 Kitaura, F.-S., Ata, M., Angulo, R. E., et al. 2016, *MNRAS*, **457**, L113
 Krolewski, A., Lee, K.-G., Lukić, Z., & White, M. 2017, *ApJ*, **837**, 31
 Krolewski, A., Lee, K.-G., White, M., et al. 2018, *ApJ*, **861**, 60
 Lee, K.-G., Hennawi, J. F., Spergel, D. N., et al. 2015, *ApJ*, **799**, 196
 Lee, K.-G., Hennawi, J. F., White, M., et al. 2016, *ApJ*, **817**, 160
 Lee, K.-G., Hennawi, J. F., White, M., Croft, R. A. C., & Ozbek, M. 2014, *ApJ*, **788**, 49
 Lee, K.-G., Krolewski, A., White, M., et al. 2018, *ApJS*, **237**, 31
 Lee, K.-G., & White, M. 2016, *ApJ*, **831**, 181
 Levi, M., Bebek, C., Beers, T., et al. 2013, arXiv:1308.0847
 Libeskind, N. I., Hoffman, Y., Forero-Romero, J., et al. 2013, *MNRAS*, **428**, 2489
 Lukić, Z., Stark, C. W., Nugent, P., et al. 2015, *MNRAS*, **446**, 3697
 Madau, P., & Dickinson, M. 2014, *ARA&A*, **52**, 415
 Malavasi, N., Arnouts, S., Vibert, D., et al. 2017, *MNRAS*, **465**, 3817
 Martizzi, D., Vogelsberger, M., Artale, M. C., et al. 2018, arXiv:1810.01883v2
 Massara, E., Villaescusa-Navarro, F., Viel, M., & Sutter, P. M. 2015, *JCAP*, **11**, 018
 McConnachie, A., Babusiaux, C., Balogh, M., et al. 2016, arXiv:1606.00043
 Modi, C., Feng, Y., & Seljak, U. 2018, *JCAP*, **10**, 028
 Naab, T., & Ostriker, J. P. 2017, *ARA&A*, **55**, 59
 Nusser, A., & Haehnelt, M. 1999, *MNRAS*, **303**, 179
 Overzier, R., & Kashikawa, N. 2019, arXiv:1903.04980
 Pahwa, I., Libeskind, N. I., Tempel, E., et al. 2016, *MNRAS*, **457**, 695
 Pichon, C., Vergely, J. L., Rollinde, E., Colombi, S., & Petitjean, P. 2001, *MNRAS*, **326**, 597

- Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2018, arXiv:1807.06209
- Porter, S. C., Raychaudhury, S., Pimbblet, K. A., & Drinkwater, M. J. 2008, *MNRAS*, **388**, 1152
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 2002, *Numerical Recipes in C++: The Art of Scientific Computing* (Cambridge: Cambridge Univ. Press)
- Ramakrishnan, S., Paranjape, A., Hahn, O., & Sheth, R. K. 2019, *MNRAS*, **489**, 2977
- Seljak, U. 1998, *ApJ*, **503**, 492
- Seljak, U., Aslanyan, G., Feng, Y., & Modi, C. 2017, *JCAP*, **12**, 009
- Simon, P., Taylor, A. N., & Hartlap, J. 2009, *MNRAS*, **399**, 48
- Skibba, R. A., Bamford, S. P., Nichol, R. C., et al. 2009, *MNRAS*, **399**, 966
- Skidmore, W. & TMT International Science Development Teams, Science Advisory Committee, T. 2015, *RAA*, **15**, 1945
- Smith, K. M., Zahn, O., & Doré, O. 2007, *PhRvD*, **76**, 043510
- Sorini, D., Oñorbe, J., Hennawi, J. F., & Lukić, Z. 2018, *ApJ*, **859**, 125
- Sorini, D., Oñorbe, J., Lukić, Z., & Hennawi, J. F. 2016, *ApJ*, **827**, 97
- Stark, C. W., White, M., Lee, K.-G., & Hennawi, J. F. 2015, *MNRAS*, **453**, 311
- Sugai, H., Tamura, N., Karoji, H., et al. 2015, *JATIS*, **1**, 035001
- Sugiyama, N. S., Okumura, T., & Spergel, D. N. 2017, *JCAP*, **1**, 057
- Takada, M., Ellis, R. S., Chiba, M., et al. 2014, *PASJ*, **66**, R1
- Topping, M. W., Shapley, A. E., Steidel, C. C., Naoz, S., & Primack, J. R. 2018, *ApJ*, **852**, 134
- Tran, K.-V. H., Papovich, C., Saintonge, A., et al. 2010, *ApJL*, **719**, L126
- Troxel, M. A., & Ishak, M. 2015, *PhR*, **558**, 1
- Wake, D. A., Collins, C. A., Nichol, R. C., Jones, L. R., & Burke, D. J. 2005, *ApJ*, **627**, 186
- Wang, H., Mo, H. J., Yang, X., et al. 2016a, *ApJ*, **831**, 164
- Wang, H., Mo, H. J., Yang, X., Jing, Y. P., & Lin, W. P. 2014, *ApJ*, **794**, 94
- Wang, T., Elbaz, D., Daddi, E., et al. 2016b, *ApJ*, **828**, 56
- White, M. 2002, *ApJS*, **143**, 241
- White, M., Pope, A., Carlson, J., et al. 2010, *ApJ*, **713**, 383
- Zel'dovich, Y. B. 1970, *A&A*, **5**, 84