






# A Gigaparsec-scale Hydrodynamic Volume Reconstructed with Deep Learning

Cooper Jacobus<sup>1</sup> , Roger de Belsunce<sup>1,2</sup>, Solène Chabanier<sup>1</sup> , Peter Harrington<sup>1</sup>, J. D. Emberson<sup>3</sup> , Zarija Lukić<sup>1</sup>, and Salman Habib<sup>3</sup>

<sup>1</sup>Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

<sup>2</sup>Berkeley Center for Cosmological Physics, Department of Physics, University of California, Berkeley, CA 94720, USA

<sup>3</sup>Argonne National Laboratory, Lemont, IL 60439, USA

Received 2024 November 25; revised 2025 September 6; accepted 2025 September 8; published 2025 November 7

## Abstract

The next generation of spectroscopic surveys will map the large-scale structure of the Universe at high redshifts ( $2 \leq z \leq 5$ ) using millions of quasar spectra, enabling major advances in constraining both the standard cosmological model and its extensions. Robust cosmological analyses of these data sets require numerical simulations that both cover gigaparsec volumes and resolve features on  $\sim 10$  kpc scales and smaller. However, running such large-volume, high-resolution hydrodynamic simulations is computationally prohibitive. We present a generative deep learning model that enhances a low-resolution, gigaparsec-scale ( $960 h^{-1}$  Mpc) hydrodynamic simulation using a smaller ( $80 h^{-1}$  Mpc) high-resolution input hydrodynamic simulation as training data. The resulting enhanced simulation reproduces the line-of-sight power spectrum to within  $\sim 10\%$  and the three-dimensional power spectrum at the  $\sim 20\%$  level at intermediate to small scales ( $k \lesssim 2 h \text{ Mpc}^{-1}$ ). Our method shows strong promise for producing realistic simulations for cosmological analyses with current surveys such as the Dark Energy Spectroscopic Instrument and upcoming next-generation experiments, but further improvements are needed to accurately recover the large-scale modes. We publicly release the enhanced hydrodynamic simulation, along with a halo catalog from a companion  $N$ -body dark matter simulation to support the calibration of data analysis pipelines for these large-scale surveys.

*Unified Astronomy Thesaurus concepts:* [Cosmology \(343\)](#); [Intergalactic medium \(813\)](#); [Large-scale structure of the universe \(902\)](#); [Ly \$\alpha\$  forest \(980\)](#); [Convolutional neural networks \(1938\)](#)


## 1. Introduction

The currently observing Dark Energy Spectroscopic Instrument (DESI) and its planned successors DESI-II and Spec-S5 (D. J. Schlegel et al. 2022; R. Besuner et al. 2025) will observe beyond a million Ly $\alpha$  forest spectra requiring mocks that approximately span a volume of  $(10 h^{-1} \text{ Gpc})^3$  (J. Farr et al. 2020), enabling three-dimensional clustering measurements and baryon-acoustic-oscillation (BAO) studies at high redshift  $z \gtrsim 2$  with subpercent-level precision (DESI Collaboration et al. 2016; A. G. Adame et al. 2025a, 2025b). Extracting cosmological information from these surveys requires creating mock skies that simultaneously (i) cover the full survey footprint to capture large-scale modes and (ii) resolve small-scale features (i.e., gas fluctuations at the  $\sim 20 \text{ kpc } h^{-1}$  scale and below) influencing observed summary statistics such as the power spectrum (Z. Lukić et al. 2015; M. Walther et al. 2021; C. C. Doughty et al. 2023). Conventional Ly $\alpha$  hydrodynamic simulations commonly meet the second requirement but at prohibitive computational and memory costs for modeling large volumes.

The Ly $\alpha$  forest, absorption features in the spectra of distant quasars, traces density fluctuations in the neutral hydrogen along the line of sight. Observing and correlating a large number of sight lines traces the large-scale structure of our Universe down to  $h^{-1}$  Mpc scales and below (P. McDonald et al. 2000; P. McDonald 2003). Its one-dimensional power spectrum (P1D) already provides leading constraints on neutrino masses and dark matter models (R. A. C. Croft et al. 1999; P. McDonald et al. 2005;

U. Seljak et al. 2005; M. Viel et al. 2013; N. Palanque-Desabrouille et al. 2015; M. M. Ivanov et al. 2025). The two-point correlation function provides a unique measurement of the BAO at high redshift (A. G. Adame et al. 2025a). DESI is pushing this further by increasing the amount of observed spectra and by including information beyond the BAO feature either through compressed statistics such as redshift-space distortions and the Alcock–Paczynski effect (A. Cuceu et al. 2021) or by using the broadband shape (i.e., the full shape) of the correlations (F. Gerardi et al. 2022). Recent advances in measuring the three-dimensional power spectrum (P3D; see, e.g., R. de Belsunce et al. 2024) and in theoretical modeling using the effective field theory of large-scale structure (EFT; P. McDonald & A. Roy 2009; D. Baumann et al. 2012; J. J. M. Carrasco et al. 2012) extended to the Ly $\alpha$  forest (M. M. Ivanov 2024) allow for a consistent model spanning large to intermediate scales to jointly extract the BAO information as well as the shape of the power spectrum (or correlation function; A. Chudaykin & M. M. Ivanov 2025; R. de Belsunce et al. 2025a; B. Hadzhiyska et al. 2025). Realizing the full potential of this data set hinges on mock catalogs that reproduce both the P1D and P3D across the  $k \sim 0.001\text{--}10 h \text{ Mpc}^{-1}$  range.

State-of-the-art hydrodynamic solvers such as Nyx code (A. S. Almgren et al. 2013; J. Sexton et al. 2021), consistently capture nonlinear intergalactic medium (IGM) physics in the context of structure formation (R. Cen et al. 1994) but quickly become computationally prohibitively expensive.<sup>4</sup> Scaling them naively to survey volumes is infeasible: the numbers of grid cells and time steps grow as the scale is cubed, and even

 Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

<sup>4</sup> See also Illustris TNG (R. Weinberger et al. 2017; A. Pillepich et al. 2018), ACCEL<sup>2</sup> (S. Chabanier et al. 2024), CAMELS (F. Villaescusa-Navarro et al. 2021, 2023), and PRIYA (S. Bird et al. 2023).


storing a single snapshot of a  $1 \text{ Gpc } h^{-1}$  high-resolution run would require  $\sim 10^{15}$  grid elements, or  $>1 \text{ PB}$  of memory. Conversely, purely  $N$ -body or coarse-grid hydrodynamical runs require far less memory but smear out small-scale structure; see, e.g., ABACUSSUMMIT  $N$ -body simulations that paint the forest on top of the dark matter field (B. Hadzhiyska et al. 2023, 2025) adopting a fluctuating Gunn–Peterson approximation (FGPA) which has been calibrated on hydrodynamic simulations (R. A. C. Croft 1998; M. Qezlou et al. 2022). Further, the  $\text{Ly}\alpha$  Mass Association Scheme (S. Peirani et al. 2014, 2022), the iteratively matched statistics method (D. Sorini et al. 2016), Bias Assignment Method (F. Sinigaglia et al. 2022), and cosmic-web-dependent FGPA (F. Sinigaglia et al. 2024) have been used to connect observed flux to the underlying matter field. Recently, R. de Belsunce et al. (2025b) presented a perturbative technique to model the  $\text{Ly}\alpha$  forest directly at the field level.

Simulating the large-scale structure of the Universe while simultaneously capturing the relevant small-scale feature at high fidelity is challenging and requires accurate modeling of nonlinear processes in the IGM, such as small-scale gas flows, thermal feedback, and the interplay between gravity and hydrodynamics. To address this critical challenge, we build up on previous work from C. Jacobus et al. (2023) and demonstrate a promising path forward: combining a relatively inexpensive, coarse  $960 h^{-1} \text{ Mpc}$  hydrodynamic run with a generative deep learning model that we trained to reconstruct the small-scale baryonic physics. Our conditional generative-adversarial network (I. Goodfellow et al. 2014; M. Mirza & S. Osindero 2014) learns to reconstruct small-scale power and correct thermodynamic biases. The result is a  $960 \text{ Mpc } h^{-1} \text{ Ly}\alpha$  volume, shown in Figure 1, whose P1D agrees with a fully resolved reference simulation to better than 10%

across  $k = 0.1\text{--}5 h \text{ Mpc}^{-1}$  and whose flux PDF matches the reference to  $\sim 10\%$  over most of its dynamic range. The P3D of the reconstructed box agrees at the  $\sim 10\%$ – $20\%$  level in the range  $k \leq 2 h \text{ Mpc}^{-1}$ .

The reconstructed simulation we present here spans 3 orders of magnitude more volume than the input hydrodynamic  $\text{Ly}\alpha$  simulations while retaining effective resolution comparable to a simulation on a  $\sim 20 \text{ kpc } h^{-1}$  Eulerian grid. We quantify the performance of our model and the accuracy of our mock volume using the one-point probability density function (PDF), the P1D, and P3D and fit to the latter a perturbative model to measure the bias parameters. The reconstructed hydrodynamic volume is released to the community for survey analysis pipeline calibration and theoretical modeling. While our model is imperfect on large scales, these advances suggest a promising road map for developing survey-scale, high-fidelity mock skies that are required to interpret the unprecedented influx of data at high redshift from current and future surveys. In future work, we plan to address these imperfections using a hybrid approach that combines the current machine learning method with the perturbation theory.

This paper is organized as follows: We first discuss the simulations and data processing techniques we use to train our network in Section 2. We then briefly outline the design of our deep learning algorithm and the methods used to train it in Section 3. We introduce the theoretical model to validate the reconstruction in Section 4. Our results are presented in Section 5, where we discuss the accuracy of our algorithm; comment on some of the statistical features of our new, higher-fidelity, gigaparsec-scale volume; and explore how this volume can be employed for precision cosmology. We conclude and discuss future work in Section 6.



**Figure 1.** Visualization of the  $960 h^{-1}$  Mpc wide baryon density volume we have reconstructed using our machine learning model. The dashed box highlights a region  $80 h^{-1}$  Mpc wide, which is the size of our training data and comparable to existing simulations used for DESI (M. Walther et al. 2021, 2025) and other surveys.

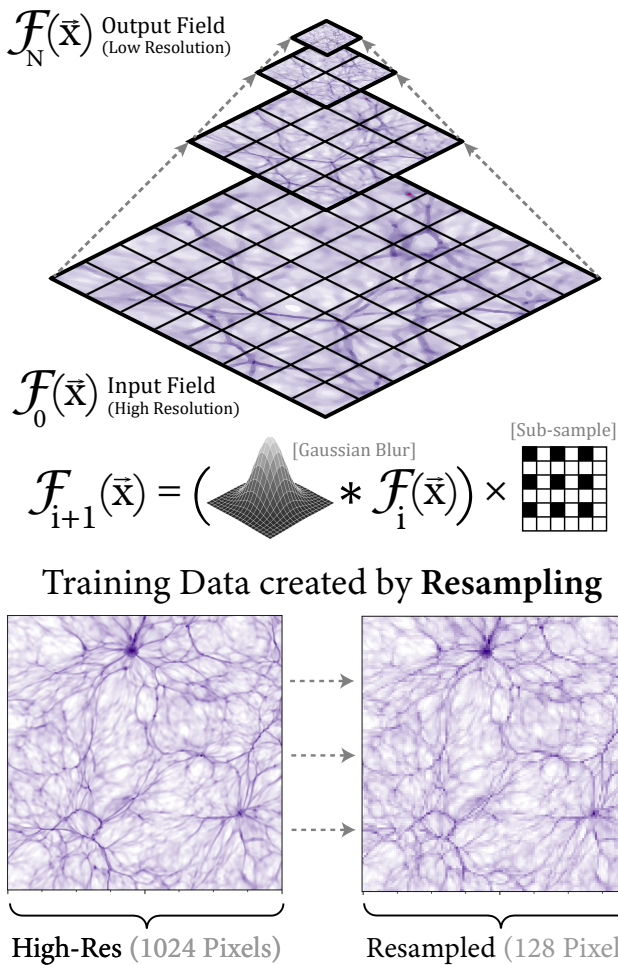
## 2. Training Data

We assemble our training and testing data from pairs of simulations run by the publicly available `NyX` code (A. S. Almgren et al. 2013; J. Sexton et al. 2021). The code follows the evolution of dark matter using collisionless self-gravitating Lagrangian particles that are coupled to baryonic ideal gas modeled on a uniform Cartesian grid. Although `NyX` is capable of employing adaptive mesh refinement (AMR), we do not make use of this functionality here as the  $\text{Ly}\alpha$  signal spans the vast low-density volumes and is fully absorbed close to halos where AMR would be relevant. To model the  $\text{Ly}\alpha$  forest, `NyX` models primordial gas composition of neutral and ionized hydrogen and helium, as well as free electrons. This type of simulation is very common in studies of the IGM and is used as a forward model in virtually any recent inference work using the  $\text{Ly}\alpha$  power spectrum (E. Boera et al. 2019; M. Walther et al. 2019; N. Palanque-Delabrouille et al. 2020; K. K. Rogers & H. V. Peiris 2021; M. Walther et al. 2021).

The simulations are initialized at  $z = 200$ , using the Zel’dovich approximation (Y. B. Zel’dovich 1970). Transfer functions were generated with the Boltzmann solver code CLASS (D. Blas et al. 2011). The cosmological parameters are set according to the Planck 2016 (Planck Collaboration et al. 2016) model:  $\Omega_b = 0.0487$ ,  $\Omega_m = 0.31$ ,  $H_0 = 67.5 \text{ km s}^{-1} \text{ Mpc}^{-1}$ ,  $n_s = 0.96$ , and

$\sigma_8 = 0.83$ . For our training data, we produce two pairs of simulations,  $80 h^{-1}$  Mpc on a side, where each pair has a high- and low-resolution run initialized from the same random initial conditions. The high-resolution simulations have  $4096^3$  elements (both dark matter particles and hydrodynamic grid cells) and are thus much more nearly converged and percent-level accurate, while the low-resolution simulations are much coarser with only  $512^3$  elements. All simulation runs have identical cosmology and UV background, and the two pairs of high- and low-resolution simulations differ only in the choice of initial condition. This enables the model to be trained on one pair of simulations and tested on another. We use the software suite `gimlet` (B. Friesen et al. 2016) to derive the optical depth and flux fields from our hydrodynamic fields and to calculate the relevant summary statistics from these fields.

In order to generate training data of high spectral fidelity but with a resolution equal to our inexpensive simulations, we use an iterative blurring and resampling method known as a Gaussian pyramid (P. Meer et al. 1987). We wish to reduce the data footprint of our high-resolution simulations by a factor of 512 while preserving the field’s spectral information. We achieve this by iteratively applying a Gaussian blur and subsampling our fields by a factor of 2, as depicted in Figure 2. We describe this method in more detail and demonstrate that it prevents aliasing in C. Jacobus et al. (2023).



**Figure 2.** Illustration of the process for creating training data target fields from high-resolution simulations. Fields are iteratively Gaussian blurred and subsampled until they match the grid size (and data footprint) of the low-resolution simulations. This method preserves most fine structures while dramatically reducing the memory necessary to store the density fields on our computer.

### 3. Machine Learning Model Design

To achieve the science goals described above, we must design a deep learning model to reconstruct physically realistic hydrodynamic fields from coarse-grid simulations while keeping the memory overall footprint constant. This presents

significant challenges due to sharp discontinuities between the input and target fields and inherent uncertainty in the mapping due to the sizable information content gap between resolutions. To address these challenges, we employ a conditional Generative Adversarial Network (GAN) framework (T. Karras et al. 2019, 2020) inspired by the fully convolutional TSIT model (L. Jiang et al. 2020). This architecture employs a multiscale encoder-generator structure, incorporating residual blocks to extract and process features at multiple spatial scales.

Our generator integrates multiscale Gaussian noise injection, inspired by StyleGAN (T. Karras et al. 2019), which treats noise injection scales as learnable parameters. This approach enhances stochasticity and better captures physically plausible variability across spatial scales (overcoming the limitations of single-point noise injection). This nondeterministic architecture enables the production of diverse yet realistic ensembles, facilitating the quantification of uncertainty. Point-wise variance across these ensembles of outputs strongly correlates with predictive error, as shown in C. Jacobus et al. (2023). This novel feature of our model naturally yields an uncertainty estimate that can be propagated throughout our analysis pipeline.

Since the baryon density and temperature fields in the simulations produced by  $N_{yx}$  contain structures spanning several orders of magnitude, their features (such as galaxy clusters and thermal shocks) are most effectively interpreted in logarithmic space. To accommodate this, we preprocess these hydrodynamic fields by log-normalizing them to fit within the range  $[-1, 1]$  before feeding them into our machine learning model.

To ensure the model learns to replicate the statistical properties of the data accurately, we employ a composite loss function composed of several terms. The primary loss component is the point-wise  $\mathcal{L}_1$  distance, which stabilizes training and ensures consistency of large-scale features. Additionally, a multiscale adversarial loss (T.-C. Wang et al. 2018; T. Park et al. 2019; L. Jiang et al. 2020) refines small-scale reconstruction using multiple patch-based discriminators. Given the crucial role of spectral accuracy in our application, we also incorporate a spectral loss term,  $\mathcal{L}_{fft}$ , which captures the differences in the truncated Fourier transforms between the predicted and target fields. For a more detailed description of our model architecture, training strategies, and inference methodology, see C. Jacobus et al. (2023).

#### 4. Theoretical Model

To validate the large-volume, low-resolution hydrodynamic simulation that we reconstructed using a generative deep learning model, we fit the resulting P3D using a theoretical model based on the EFT. The EFT formalism uses a perturbative expansion of the large-scale dynamics by only using the symmetries relevant to the tracer (P. McDonald & A. Roy 2009; D. Baumann et al. 2012; J. J. M. Carrasco et al. 2014; M. M. Ivanov 2022), which for the Ly $\alpha$  forest are the equivalence principle and rotational invariance around the line-of-sight direction  $\hat{z}$ , corresponding to the SO(2) group (P. McDonald et al. 2000; V. Desjacques et al. 2018; J. J. Givans & C. M. Hirata 2020; S.-F. Chen et al. 2021; M. M. Ivanov 2024; M. M. Ivanov et al. 2025; R. de Belsunce et al. 2025a). In the following, we briefly summarize the one-loop Ly $\alpha$  EFT model and refer the reader to V. Desjacques et al. (2018), M. M. Ivanov (2024), M. M. Ivanov et al. (2025), and R. de Belsunce et al. (2025a) for a more complete presentation.

In linear theory, the Ly $\alpha$  flux decrements are modeled using the bias expansion (P. McDonald et al. 2000; P. McDonald 2003)

$$\delta_F = b_1 \delta + b_\eta \eta, \quad (1)$$

with the dimensionless gradient of the peculiar velocity,  $\eta = -\partial_{\parallel v_{\parallel}}/\mathcal{H}$ , where  $\mathcal{H} = aH$  is the conformal Hubble parameter given by the scale factor  $a$  and the Hubble parameter  $H$ . The resulting tree-level power spectrum is connected to the linear matter power spectrum via

$$P^{\text{tree}}(k, \mu) = (b_1 + b_\eta f \mu^2)^2 P_{\text{lin}}(k), \quad (2)$$

where  $k$  is the Fourier wavenumber,  $\mu$  is the angle of  $k = \{k_{\parallel}, \mathbf{k}_{\perp}\}$  to the line of sight,  $\mu \equiv k_{\parallel}/k$ ,  $f$  is the (linear) growth rate, and the two linear bias parameters are  $b_1$  and  $b_\eta$ . All quantities are evaluated at the redshift of the simulation that we suppress in the following.

Schematically, the one-loop EFT power spectrum is based on four components:

$$P^{\text{theory}} = P^{\text{tree}} + P^{1\text{-loop}} + P^{\text{ct}} + P^{\text{st}}, \quad (3)$$

where  $P^{\text{tree}}$  is the infrared resummed linear theory power spectrum.<sup>5</sup> The one-loop Ly $\alpha$  power-spectrum component is given by

$$P^{1\text{-loop}}(k, \mu) = 2 \int_{\mathbf{q}} K_2^2(\mathbf{q}, \mathbf{k} - \mathbf{q}) P_{\text{lin}}(|\mathbf{k} - \mathbf{q}|) P_{\text{lin}}(q) \\ + 6K_1(\mathbf{k}) P_{\text{lin}}(k) \int_{\mathbf{q}} K_3(\mathbf{k}, -\mathbf{q}, \mathbf{q}) P_{\text{lin}}(q), \quad (4)$$

with higher-order redshift-space kernels,  $K_{2,3}$  given in Equation (3.19) in M. M. Ivanov (2024) with the shorthand notation  $\int_{\mathbf{q}} \equiv \int \frac{d^3q}{(2\pi)^3}$  to denote the three-dimensional integral over  $\mathbf{q}$ . The counterterms scale as  $k^2 P_{\text{lin}}(k)$ :

$$P^{\text{ct}}(k, \mu) = -2(c_0 + c_2 \mu^2 + c_4 \mu^4) K_1(\mathbf{k}) k^2 P_{\text{lin}}(k), \quad (5)$$

with the parameters  $c_{0,2,4}$  controlling the angular dependence and the stochastic contributions are a function of a constant shot noise piece and a scale- and angle-dependent term, accounting for small-scale physics:<sup>6</sup>

$$P^{\text{st}}(k, \mu) = P_{\text{shot}} + a_0 \frac{k^2}{k_{\text{NL}}^2} + a_2 \frac{k^2 \mu^2}{k_{\text{NL}}^2}. \quad (6)$$

Measurements of these parameters on different Ly $\alpha$  simulations have been presented in M. M. Ivanov (2024), B. Hadzhiyska et al. (2025), and R. de Belsunce et al. (2025a).

The one-loop Ly $\alpha$  power spectrum is described by the following set of nuisance parameters:

$$\{b_1, b_\eta, b_2, b_{\mathcal{G}_2}, b_{(KK)_{\parallel}}, b_{\Pi_{\parallel}^{[2]}}, b_{\delta\eta}, b_{\eta^2}\}, \quad (7)$$

in addition to the cubic EFT terms

$$\{b_{\Pi_{\parallel}^{[3]}}, b_{(K\Pi^{[2]})_{\parallel}}, b_{\delta\Pi_{\parallel}^{[2]}}, b_{\eta\Pi_{\parallel}^{[2]}}\}, \quad (8)$$

and the counter and stochastic terms

$$\{c_{0,2,4}, a_{0,2}, P_{\text{shot}}\}, \quad (9)$$

yielding a total of 18 parameters. We adopt the prior choices of Table 1 in R. de Belsunce et al. (2025a) and analytically marginalize over the parameters in Equations (8)–(9).

To fit the 3D power spectra, we follow the standard procedure in J. J. Givans et al. (2022) and sample the  $\chi^2$  function (jointly fitting all four angular and all  $k$  bins):

$$\chi^2 = \sum_i \frac{[P_i^{\text{data}} - P_i^{\text{model}}]^2}{2(P_i^{\text{data}})^2/N_i}. \quad (10)$$

The measured power spectrum is denoted by  $P_i^{\text{data}}$  in  $k$  and  $\mu$  bins with  $N_i$  Fourier modes per bin;  $P_i^{\text{model}}$  is the EFT theory prediction. The error bars are computed assuming a Gaussian covariance, given by  $P_i^{\text{data}} \sqrt{2/N_i}$ , which assumes the measured power along each of the axes to be independent. Note that we account for mode discreteness by evaluating the model vector at the power-spectrum-weighted mean  $k$  and  $\mu$  for each bin, computed via  $\bar{x} = \sum_i P_i^{\text{data}} x_i / \sum_i P_i^{\text{data}}$  for  $x \in \{k, \mu\}$ . For the present analysis, we do not vary cosmology and use the input cosmology of the simulation and use  $k_{\text{max}} = 1 \text{ h Mpc}^{-1}$  as the baseline maximum wavenumber for the large and  $k_{\text{max}} = 2 \text{ h Mpc}^{-1}$  for the smaller volumes.<sup>7</sup>

<sup>5</sup> This is computed using the Boltzmann solver CLASS-PT (D. Blas et al. 2011; A. Chudaykin et al. 2020).

<sup>6</sup> We emphasize that the EFT is only valid at larger scales than the nonlinear scale,  $k_{\text{NL}}$ , and should be treated as phenomenological beyond it. The scale is defined as  $\Delta^2 = k^3 P_{\text{lin}}(k, z)/(2\pi^2) \sim 1$ , i.e., where the dimensionless power spectrum reaches unity. For the current simulations, this corresponds to  $k_{\text{NL}} \approx 7 \text{ h Mpc}^{-1}$  at the redshift  $z = 3.0$ .

<sup>7</sup> For a discussion on the impact of scale cuts and biases induced by higher-order corrections, see Figure 3 in R. de Belsunce et al. (2025a).

## 5. Results

We present a new hydrodynamic simulation volume reconstructed using our deep learning method. The simulation is a cube  $960 h^{-1}$  Mpc wide represented by  $6144^3$  volumetric cells containing values of dark matter density, baryon density, temperature, velocity, and  $\text{Ly}\alpha$  optical depth. (For reference, the high- and low-resolution training volumes cover an  $80 h^{-1}$  Mpc wide cube with  $4096^3$  and  $512^3$  cells respectively; see Figure 1). Our simulation was run using `Nyx` at a resolution exactly matching that of our low-resolution training data and was likewise initialized at  $z = 200$ , with the exact same cosmology and UV background. We take a simulation output at a redshift of  $z = 3$  and then pass it through our deep learning model to improve the small-scale fidelity. The reduced simulation output in single precision requires 11 TB of computer memory for the grid representations of dark matter and baryon quantities, as well as the  $\text{Ly}\alpha$  optical depth. A high-resolution simulation of this volume would require  $>5.6$  PB.

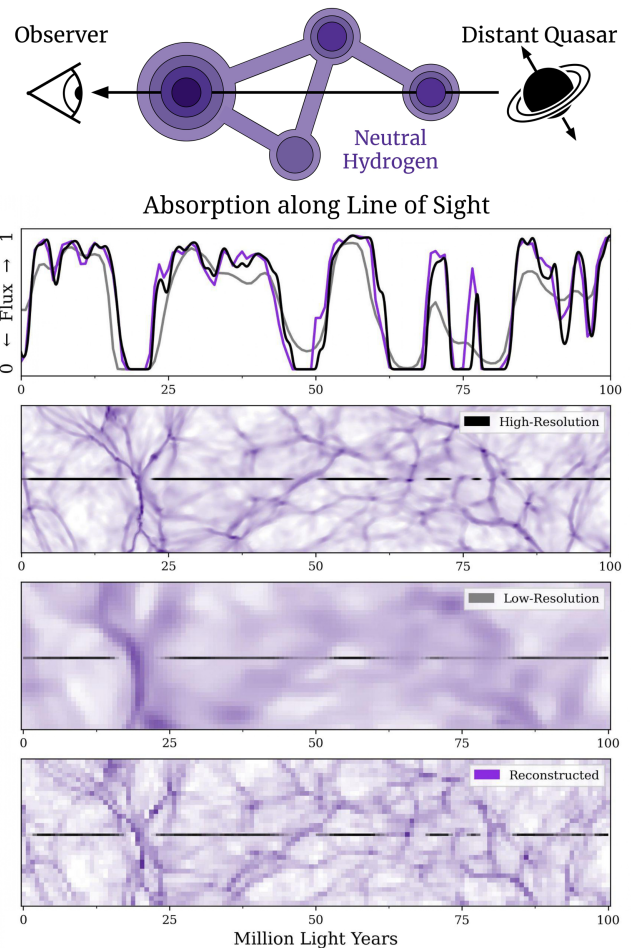
In the following subsections, we explore the accuracy of our Machine Learning-enhanced  $960 h^{-1}$  Mpc wide simulation by examining a few key summary statistics of the hydrodynamic fields and their corresponding  $\text{Ly}\alpha$  flux fields. We demonstrate the power of this larger volume to contain cosmological models by using it to fit an EFT model, introduced in Section 4.

### 5.1. Hydrodynamic Fields

The baryon density field of the  $960 h^{-1}$  Mpc wide volume we reconstructed using a generative deep learning model is visualized in Figure 1, the process of creating target and training data is visualized in Figure 2, and the improvement is visually demonstrated in Figure 3. We examine the statistical properties of the baryon density and temperature in Figure 4. We compare the probability distributions of our baryon density and temperature to those from a “true” high-resolution simulation of a much smaller volume. The output of our machine learning model presents a clear improvement over the low-resolution simulation volume; both distribution peaks have been noticeably corrected, and the distribution shapes have been adjusted to better match that of the high-resolution baseline.

While the machine learning model does capture the center of the distributions of the baryon density and temperature reasonably well, improving the low-resolution distributions greatly, it fails at accurately capturing the tails of these distributions. This is due to an underrepresentation of these regimes in the training data by a factor of  $10^5$ – $10^6$  and therefore a lack of incentive for the model to accurately reproduce them.

We show the 2D density–temperature distribution in the bottom panel of Figure 4. While the low-resolution distributions differ from the high-resolution case, they agree well with each other, indicating that the small volume captures the essential diversity of the large box. Because of our down-sampling scheme, much of the fine density–temperature structure is smoothed out, and the condensed, galaxy-forming tail is not well reproduced. However, this regime represents a very small fraction of the total volume. More importantly, our method significantly improves the recovery of the diffuse gas power-law slope,  $T = T_0(\rho/\langle\rho\rangle)^{\gamma-1}$ , reducing the error from



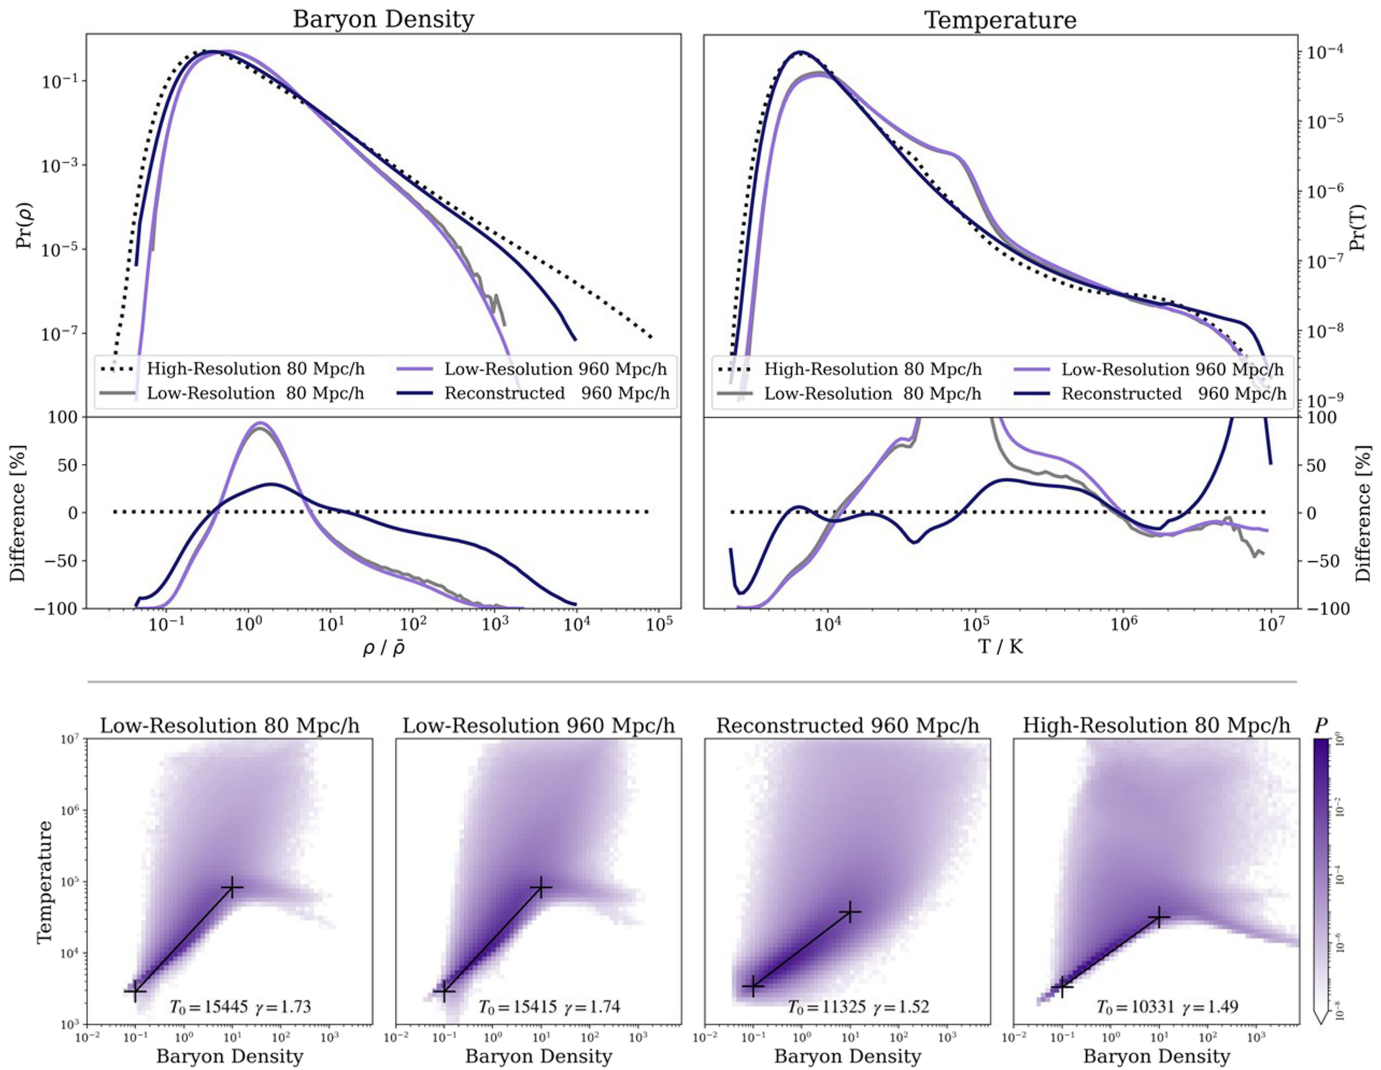
**Figure 3.** Illustration of the effects of resolution on  $\text{Ly}\alpha$  absorption features. We plot slices of baryon density  $\sim 30$  Mpc wide for high-resolutions, low-resolution, and reconstructed fields, visualized in log-space. Notice that absorption features are spectrally sharpest in the high-resolution simulation but that some spectral detail is recovered by our ML model.

$\sim 17\%$  to  $2\%$  as a fractional difference of the high-resolution slope,  $\gamma = 1.49$ . This improvement strongly affects the  $\text{Ly}\alpha$  signal and provides sharper constraints on the IGM’s ionization history.

### 5.2. Lyman- $\alpha$ Fields

$\text{Ly}\alpha$  absorption maps are calculated using the `gimlet` code (B. Friesen et al. 2016) along lines of sight or “skewers” that pass through the entire volume, parallel to one of the axes, simulating a random distribution of quasar spectra that probe our simulated fields, comparable to the data expected from spectroscopic surveys like DESI. The aggregate of all such lines of sight forms the  $\text{Ly}\alpha$  “field,” representing an absorption map of the entire volume, including redshift-space distortions along one of the axes. We calculate some key summary statistics of this field using `gimlet` and compare them to the target and input simulations to evaluate our method.

We find that our machine learning method reliably reconstructs key  $\text{Ly}\alpha$  forest statistics across a broad range of scales. Figure 5 compares the flux PDF (left) and the P1D (right) between the high-resolution target, the low-resolution baseline, and our ML-enhanced simulation. The ML model reproduces the P1D with high accuracy across all scales, while deviations in the PDF are more pronounced, particularly in the



**Figure 4.** (Top) The PDFs of the predicted density and temperature fields (dark blue lines) compared against the reference distributions from a target high-resolution simulation (dotted lines). (Bottom) The density–temperature phase distribution of the gas in our model prediction compared against the reference distributions from the low-resolution and high-resolution simulations. For each output, we plot the best-fit power-law relationship between density and temperature; our model presents an improvement in the slope of this power-law,  $\gamma$ , when compared against the low-resolution simulation but produces a much broader distribution.

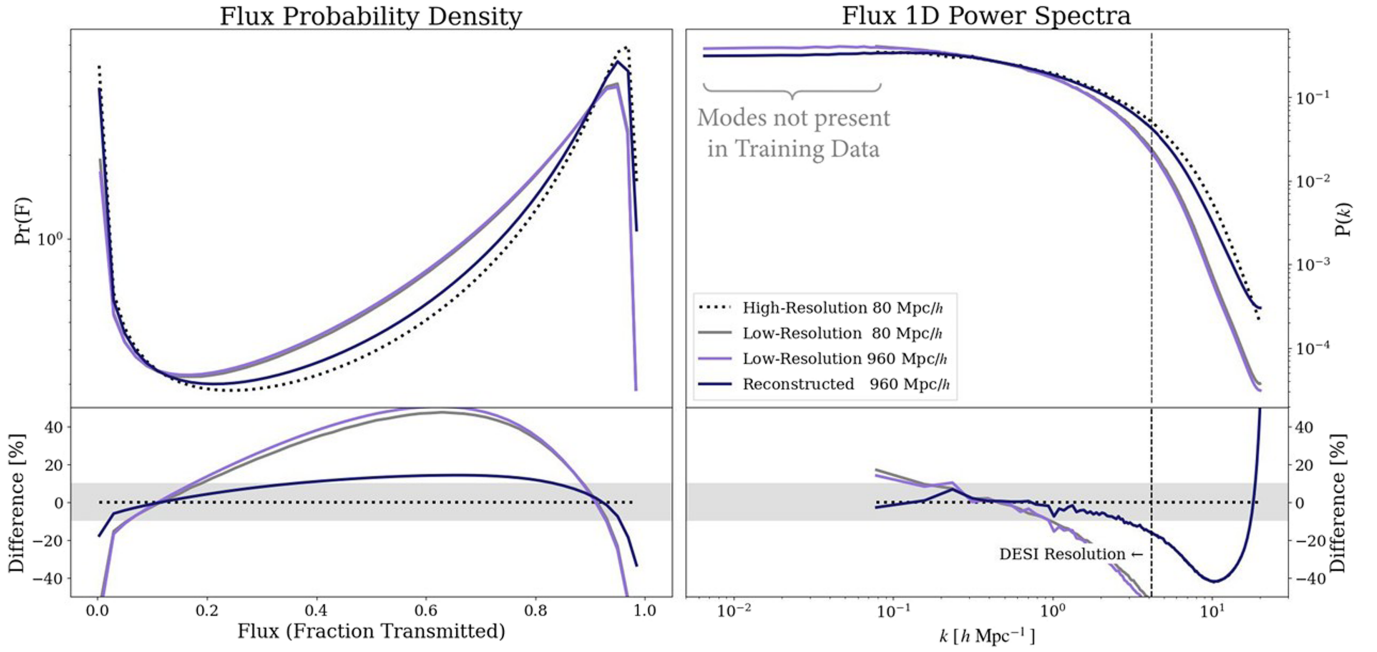
sparse areas of the distribution. This indicates that the method effectively captures the amplitude of small-scale flux fluctuations but is less accurate at recovering the detailed phase information needed to fully reproduce the PDF.

In particular, the flux PDF carries information about the nonlinear, halo-scale physics and reflects the thermal phase of the IGM, which cannot be accessed through power-spectrum measurements alone. The values of this metric for the input simulations (gray and purple lines in Figure 5 above) are off by as much as 50% when compared to the high-resolution target for intermediate flux values ( $F \sim 0.5$ ) due to the diffusivity of the low-resolution hydrodynamic fields, which would lead to significant errors if used for cosmological analysis. We show in Figure 5 that the PDF of our reconstructed Ly $\alpha$  field (black line) is within about 10% of the true (high-resolution) values. This offers a substantial improvement over the low-resolution input but suggests that the thermal phase of the IGM is still inaccurately captured.


Complementing the information captured by the flux PDF, the PID of the Ly $\alpha$  forest offers another key statistical measure that reflects the scale-dependent clustering of

transmitted flux along quasar sight lines, constraining both the properties of the IGM and critical cosmological parameters (P. McDonald et al. 2005; N. Palanque-Delabrouille et al. 2013). The PID is one of the most important observables that can be predicted by our cosmological models and is especially sensitive to the amplitude and shape of primordial density fluctuations and other large-scale physics. We plot the mean power spectrum  $\langle P(k) \rangle$  averaged over modes with magnitude  $k$  from all skewers, as done in Z. Lukić et al. (2015). We show comparisons between the low-resolution, ML-reconstructed, and “true” high-resolution 1D power spectra in Figure 5 above.

Our reconstructed Ly $\alpha$  PID is in agreement with the high-resolution test volume to within 10% for wavenumbers smaller than  $5 h \text{ Mpc}^{-1}$  and is in better agreement than the low-resolution simulation for all wavenumbers present in the test volume. At smaller scales, the remaining differences between our reconstructed PID and the high-resolution target likely arise from imperfect reconstruction of the nonlinear density–temperature relationships and thermal broadening effects that influence the detailed shape of absorption features. While our model captures the overall amplitude of small-scale power remarkably well, there



**Figure 5.** The probability density (PDF, left panel) and 1D power spectrum (PID, right panel) of Ly $\alpha$  flux in redshift space for four related simulation volumes: the high- and low-resolution Nyx simulations, the resized target, and our network output. The gray band in the residual highlights the 10% error range around the smaller,  $80 h^{-1}$  Mpc high-resolution simulation from our validation data set. The strong PID agreement demonstrates that small- and intermediate-scale fluctuations are well captured; in contrast, discrepancies in the PDF arise because the model’s limited correlation length prevents reconstruction of correct large-scale phases and spatial correlations.



**Figure 6.** The three-dimensional power spectrum (P3D) of the Ly $\alpha$  forest for four different simulations shown as a function of Fourier wavenumber,  $k$ , and the cosine of the angle with the line of sight,  $\mu = k_{\parallel}/|k|$ : (i) low-resolution  $960 h^{-1}$  Mpc box in purple; (ii) the reconstructed  $960 h^{-1}$  Mpc box in blue; (iii) the high-resolution input  $80 h^{-1}$  Mpc simulation in black dashed; and (iv) EFT fits to the ACCEL<sup>2</sup> simulation in pink with a box of length  $L = 160 h^{-1}$  Mpc (S. Chabanier et al. 2024; R. de Belsunce et al. 2025a). The panels show the 3D power spectrum divided into four separate  $\mu$  bins. The generative deep learning model is able to reconstruct the angular evolution for the  $960 h^{-1}$  Mpc box and an overall agreement at the  $\sim 20\%$  level.

are subtle differences in the thermal state of the IGM (see the difference in the tail of the phase distribution in Figure 4) that affect the fine structure and widths of some individual absorption lines, leading to residual discrepancies in statistics like the PID that are sensitive to these details. Fortunately, most of these discrepancies impact the power spectrum on scales smaller than what DESI can resolve.

While our machine learning model is only able to make adjustments to the field at small scales, by construction, it is worth noticing how the total power at very large scales (i.e., low wavenumbers) is affected. Our reconstruction produces a decrease of about 20% in the power with respect to the low-resolution gigaparsec-scale input simulation. This is most likely due to the insufficient correlation length (receptive field) of our ML model,

which cannot adequately capture very large-scale modes beyond the  $80 h^{-1}$  Mpc training volume size and which may slightly smooth out some of the very large-scale modes because it was impossible to represent them in the training data. We present some ideas for how to address this in future projects in Section 6.

To further quantify the performance of our model, we present the P3D of the Ly $\alpha$  flux in Figure 6. We present the 3D power spectrum as a function of Fourier wavenumber  $k$  and the angle between the wavevector and the line of sight, defined as  $\mu \equiv k_{\parallel}/|k|$ , illustrated in four bins with  $\Delta\mu = 0.25$ . Larger  $\mu$  values align closer to the line of sight. We compare four power spectra: (i) the low-resolution  $960 h^{-1}$  Mpc box in purple; (ii) the reconstructed  $960 h^{-1}$  Mpc box in blue; and (iii) the high-resolution input  $80 h^{-1}$  Mpc simulation in black dashed; for

comparison, we show (iv) EFT fits to the ACCEL<sup>2</sup> simulation in pink with a box of length  $L = 160 h^{-1} \text{Mpc}$  with  $N = 6144^3$  particles, yielding a physical resolution of  $25 h^{-1} \text{kpc}$  presented in R. de Belsunce et al. (2025a). We emphasize two aspects of this figure: First, the low-resolution simulation does not evolve as a function of  $\mu$  as expected from the high-resolution input simulation. For the reconstructed box, we see an evolution with an angle showing that the generative model is able to inject this information. Second, both high-resolution simulations (which share the same cosmology but use different hydrodynamic prescriptions) agree at the power-spectrum level, validating the used training simulation. We find that, for all four  $\mu$  bins, the power spectrum of the reconstructed simulation achieves a clear improvement, from an error of over 100% in the case of the low-resolution simulation down to around 20% with our method. Our ML method mostly recovers the  $\mu$  dependence, shifting the spectrum toward the target curve in all  $\mu$  bins and at all scales.

### 5.3. Inferring Parameters from the Simulations

In this section, we further quantify the performance of the generative model to reconstruct the low-resolution simulation and comment on their applicability for cosmological data analysis. We fit the power spectra shown in Figure 6 with the EFT theory model introduced in Section 4 and compare their bias parameters given in Equations (7)–(9). The key question we are interested in answering is how the reconstruction affects the large and small scales. EFT offers a pathway to explore this question as the linear bias parameters are (mostly) sensitive to the large-scale modes, while the higher-order bias parameters are dominated by small-scale modes.

In Figure 7, we show the marginalized 1D and 2D posteriors for the eight bias parameters for the input and reconstructed simulations used in this work. The low-resolution simulation is shown in purple, the reconstructed is shown in blue, and the high-resolution input is shown in gray. Following baseline expectations, the large volume of the coarse and reconstructed hydrodynamic simulation gives access to many quasi-linear modes yielding tight constraints on the bias parameters. For the higher-order bias parameters, these results are in qualitative agreement with the contour plots presented in R. de Belsunce et al. (2025a). The corresponding best-fit EFT spectra are shown in Figure 8 in the Appendix, with the mean best-fit values tabulated in Table 1 in the Appendix.

It is interesting to note that the higher-order bias parameters from the low- and high-resolution hydrodynamic simulations are fully consistent with each other. The reconstruction sharpens the marginalized posteriors, improving the consistency with the input high-resolution simulation. This is in contrast to the linear bias parameters, which are constrained mostly from large-scale modes. The marginalized posteriors on  $b_1$  shows tight constraints, with the reconstruction not improving the agreement with the input high-resolution simulation. The velocity bias  $b_\eta$  is in (slightly) better agreement with the input simulation. The discrepancy is also visible when assessing the quality of the fits in Figure 8 for the reconstructed volume. The main reason for the discrepancy in the linear bias parameters stems from the short correlation length ( $\sim 20 h^{-1} \text{Mpc}$ ) of our model. Thus, the reconstructed ML box does not have any large- to intermediate-scale information from the input simulation resulting in unrealistic (or even arbitrary) values of  $b_1$  and  $b_\eta$ . This maximum correlation length is determined by the structure of the

convolutional neural network and by the memory limitations of the GPUs used for inference. We emphasize that the reconstruction does improve the small-scale information of the coarse large-volume simulation as seen from the agreement of the PID and the higher-order bias parameters.

### 5.4. Halo Catalog

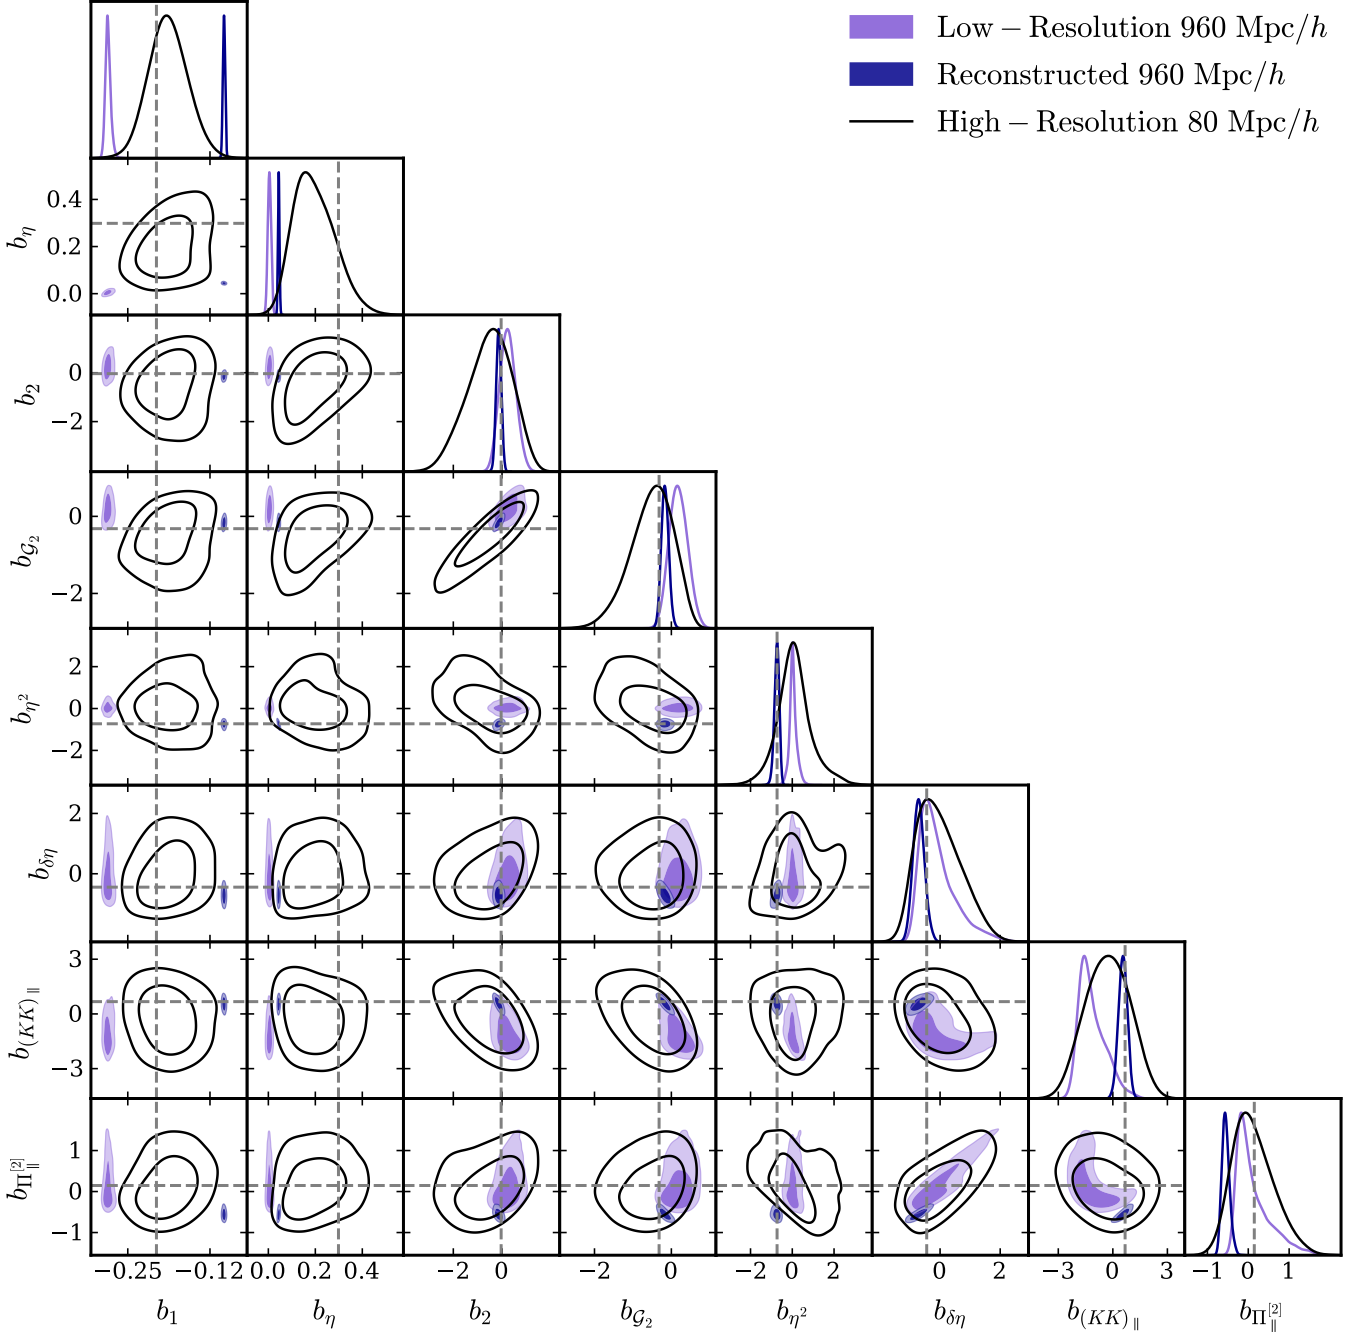
In addition to the hydrodynamic simulation discussed above, we performed a companion  $N$ -body simulation of the  $960 h^{-1} \text{Mpc}$  box using the Hardware/Hybrid Accelerated Cosmology Code (HACC; S. Habib et al. 2016). The initial noise field of the original N<sub>YX</sub> hydrodynamic simulation was used to generate positions and velocities of  $6144^3$  dark matter particles at a starting redshift of  $z = 200$ . The particles were subsequently evolved to  $z = 3$  using gravitational forces via typical  $N$ -body methods with a Plummer softening length of  $1.56 \text{kpc } h^{-1}$ . To produce a robust dark matter halo catalog from this  $n$ -body simulation, halos were identified using the friends-of-friends (FOF) method from the ARBORX library (A. Prokopenko et al. 2024), which is built into the parallel analysis framework of HACC. We employ a linking length of  $b = 0.168$  and assign the center of each FOF halo as the location of its most bound particle. Next, spherical overdensity (SO) halos are constructed around the center position of each FOF halo using an overdensity criterion of 200 times the critical density.

We make this halo catalog available alongside our hydrodynamic simulation so that it may be used to calibrate analysis techniques that combine Ly $\alpha$  fields with galaxy/quasar/halo data. The combination of these two data sets will allow cosmologists to explore the rich cross correlations between the Ly $\alpha$  forest and other cosmological tracers, developing new methodologies for extracting physical insights from existing and future surveys.

## 6. Conclusion

We present a generative deep learning model, introduced in Section 3, that enhances a low-resolution, gigaparsec-scale ( $L = 960 h^{-1} \text{Mpc}$ ) hydrodynamic simulation using a high-resolution ( $L = 80 h^{-1} \text{Mpc}$ ) input hydrodynamic simulation as training data, introduced in Section 2. The reconstructed simulation we present here is an essential step toward meeting the demands of current and next-generation spectroscopic surveys, which require both massive survey volumes and accurate small-scale modeling that have previously been out of reach with traditional simulation methods due to computational bottlenecks. By bridging the gap between computational affordability and physical fidelity, an approach like ours can enable the scientific community to create mock universes suitable for precision cosmological analysis, calibration of data pipelines, and testing of new statistical techniques, capabilities that are indispensable for extracting robust constraints from Ly $\alpha$  forest observations. Without such reconstructions, robust interpretation of Ly $\alpha$  forest data, especially in the context of large-scale structure, neutrino mass, and dark matter studies, would remain severely limited by either computational cost or insufficient resolution. Our method allows us to produce a mock universe over 1000 times larger than an input state-of-the-art simulation.

This volume was reconstructed from a relatively low-resolution simulation, using a deep learning algorithm that we



**Figure 7.** Drift plot of the quadratic bias parameters fit to three simulations: (i) the low-resolution  $960 h^{-1}$  Mpc (purple), (ii) the reconstructed  $960 h^{-1}$  Mpc box (blue); and (iii) the input high-resolution  $80 h^{-1}$  Mpc simulation (gray). We show the best-fit 1D and 2D marginalized posteriors using a maximum wavenumber of  $k_{\max} = 1 h \text{ Mpc}^{-1}$  ( $k_{\max} = 2 h \text{ Mpc}^{-1}$ ) for the fits of the large (small) boxes. We find promising agreement in the higher-order EFT bias parameters, but note that the linear bias parameters ( $b_1$  and  $b_\eta$ ) yield discrepant results at the  $2\sigma$ – $3\sigma$  level. The quality of the fits is shown in Figure 8 in the Appendix.

designed and optimized for this specific purpose. We have demonstrated that this translation method provides an accurate reconstruction over a wide range of scales down to  $\sim 0.3 h^{-1}$  Mpc. Our method provides a far less computationally expensive way to produce volumes like this one of large scales. Our network can be trained on a relatively small volume for some given cosmological and astrophysical parameters and then used to reconstruct an (arbitrarily) large volume.

Although our model presents a remarkable improvement in the fidelity and statistical accuracy of the reconstructed Ly $\alpha$

forest, particularly at the smaller scales, which cannot be accurately represented by low-resolution hydrodynamic simulations, there are still some notable shortcomings, particularly on large scales. As shown in Figures 7 and 8, the reconstructed Ly $\alpha$  forest does not accurately capture the large-scale modes. While increasing the size of the training set to include multiple high-resolution realizations would further reduce the uncertainty at intermediate scales (which are already very small as shown in the middle panel of Figure 8), the largest improvement will stem from increasing the model’s limited correlation length it can internally process.

Further, we also note some remaining shortcomings in our model’s hydrodynamic field reconstructions, particularly in and around the larger clusters and shocks. While these regions represent a very small percentage of the universe by volume and have only a marginal effect on the  $\text{Ly}\alpha$  statistics, they are of significant astrophysical interest and affect the hydrodynamic statistics shown in Figure 4. Because these regimes are especially dependent on subgrid dynamics, they are poorly captured by the input low-resolution simulation and present the greatest challenge to our machine learning model. However, our technique offers a significant improvement over others in capturing these shocks and accurately reproducing the statistics of the  $\text{Ly}\alpha$  forest.

We have shown that our reconstruction model is able to inject small-scale physics into larger-volume, low-resolution simulations. This is the first step toward constructing high-fidelity, large-volume mocks for surveys like DESI, WEAVE-QSO (M. M. Pieri et al. 2016), the Prime Focus Spectrograph (J. Greene et al. 2022), and 4MOST (R. S. de Jong et al. 2019). The most pressing shortcoming of the present model is that it cannot accurately reconstruct large-scale modes due to the short correlation length of the deep learning architecture. We leave the exploration of a hybrid approach (see, e.g., C. Modi & O. H. E. Philcox 2023) combining the present method for small-scale reconstruction and capturing (and constraining) the large scales using perturbation theory to future work. Recent progress in modeling the  $\text{Ly}\alpha$  forest directly at the field level (R. de Belsunce et al. 2025b) offers an interesting pathway for this hybrid approach.

### Acknowledgments

This work was supported by the Department of Energy, Laboratory Directed Research and Development program (PI: Zarija Lukić) at Lawrence Berkeley National Laboratory. Production of auxiliary simulations was supported by the DOE’s Office of Advanced Scientific Computing Research and Office of High Energy Physics through the Scientific Discovery through Advanced Computing (SciDAC) program. This research used 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. DEAC02-05CH11231. This research also utilized the resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357.

C.J. gratefully acknowledges scholarship support from the Goldwater Foundation and the NASA Astronaut Scholarship Foundation.

### Data Availability

The complete machine learning-enhanced simulation volume, including all hydrodynamic fields,  $\text{Ly}\alpha$  optical depth fields, and the associated dark matter halo catalog are publicly

available for access and download from the National Energy Research Scientific Computing Center (NERSC) “Science Gateways” portal:

<https://portal.nersc.gov/project/nyx/Gpc/>.

We invite members of the community to access these resources to experiment with this novel data product and conduct further analysis to enable future large-scale cosmological studies.


## Appendix Fit Details

In this section, we show the quality of the fits in Figure 8. It is interesting to note that the fits to the low- and high-resolution volumes yield very precise results. However, the perturbative model captures the small-scale features very well (mirroring the contour plot in Figure 7) but fails at capturing the large scales. This comparison transparently shows the regime where the generative model fails. The corresponding bias parameters for the fits are tabulated below in Table 1.

**Table 1**  
Mean Best-fit Values for the One-loop EFT Parameters Obtained from the Three Simulations

$b_{\mathcal{O}}$	Low-resolution	Reconstructed	High-resolution
$b_1$	$-0.2821^{+0.0046}_{-0.0033}$	$-0.0982 \pm 0.0018$	$-0.1997^{+0.0321}_{-0.0295}$
$b_{\eta}$	$0.0034 \pm 0.0071$	$0.0446^{+0.0032}_{-0.0032}$	$0.1672^{+0.1032}_{-0.0721}$
$b_2$	$0.0995^{+0.3467}_{-0.3114}$	$-0.1495 \pm 0.1046$	$-0.5521^{+0.8004}_{-1.0348}$
$b_{\mathcal{G}_2}$	$0.2469^{+0.2549}_{-0.2295}$	$-0.1666^{+0.0989}_{-0.0911}$	$-0.4863^{+0.4547}_{-0.6387}$
$b_{\eta^2}$	$0.0210^{+0.1745}_{-0.1352}$	$-0.7503^{+0.1028}_{-0.1286}$	$0.0707^{+0.8101}_{-0.6666}$
$b_{\delta\eta}$	$-0.1179^{+0.7233}_{-0.2658}$	$-0.7170^{+0.1977}_{-0.1793}$	$-0.6516^{+0.8986}_{-0.5490}$
$b_{(KK)\parallel}$	$-1.2576^{+0.9062}_{-0.4430}$	$0.5863 \pm 0.2429$	$0.4038 \pm 1.2596$
$b_{\Pi\parallel^{[2]}}$	$-0.1151^{+0.4858}_{-0.1490}$	$-0.5563^{+0.1044}_{-0.0768}$	$-0.3691^{+0.5898}_{-0.4020}$
$b_{\Pi\parallel^{[3]}}$	$1.0830 \pm 0.1186$	$1.8838 \pm 0.2614$	$1.6372 \pm 0.6625$
$b_{\delta\Pi\parallel^{[2]}}$	$-1.3370 \pm 0.2273$	$2.4556 \pm 0.4988$	$-0.6425 \pm 0.7162$
$b_{(K\Pi^{[2]})\parallel}$	$-1.2027 \pm 0.3184$	$-1.7333 \pm 0.7309$	$-0.6670 \pm 1.3111$
$b_{\eta\Pi\parallel^{[2]}}$	$0.3929 \pm 0.4059$	$-0.1171 \pm 0.8765$	$0.1663 \pm 1.3458$
$P_{\text{shot}}$	$0.2043 \pm 0.3899$	$1.9225 \pm 0.2716$	$0.5650 \pm 0.3600$
$a_0$	$0.1183 \pm 0.9992$	$-1.6936 \pm 0.6955$	$-0.3494 \pm 0.3167$
$a_2$	$0.5170 \pm 0.8914$	$-2.5761 \pm 0.9223$	$0.3657 \pm 0.4438$
$c_0$	$-0.0925 \pm 0.0922$	$-1.8439 \pm 0.1859$	$-0.2987 \pm 0.1509$
$c_2$	$0.0088 \pm 0.0344$	$2.8017 \pm 0.0673$	$0.2454 \pm 0.1074$
$c_4$	$-0.0451 \pm 0.0134$	$-1.9937 \pm 0.0244$	$-0.1394 \pm 0.0462$

**Note.** The corresponding contours are shown in Figure 7. The default fit is performed with  $k_{\text{max}} = 1 h \text{ Mpc}^{-1}$  ( $k_{\text{max}} = 2 h \text{ Mpc}^{-1}$ ) for the large (small) boxes. The counterterms are divided by  $(h \text{ Mpc}^{-1})^2$ , and we analytically marginalize over the parameters shown in the bottom part of the table. Note that their posteriors are recovered a posteriori from the chains.



**Figure 8.** Comparison of the quality of the EFT fits to the 3D power spectra of the three simulations. Left panel: low-resolution box; middle panel: reconstructed box; and right panel: high-resolution input box as a function of angular bin  $\mu$ . In each bottom panel, we show the residuals and guide the eye with a 2% error band shown in gray. We use a maximum wavenumber of  $k_{\max} = 1 \text{ h Mpc}^{-1}$  for the large simulations and  $k_{\max} = 2 \text{ h Mpc}^{-1}$  for the small volume.

## ORCID iDs

Cooper Jacobus <https://orcid.org/0000-0002-0498-2019>  
 Solène Chabanier <https://orcid.org/0000-0002-5692-5243>  
 J. D. Emberson <https://orcid.org/0000-0003-1406-0744>

## References

- Adame, A. G., Aguilar, J., Ahlen, S., et al. 2025a, *JCAP*, **01**, 124  
 Adame, A. G., Aguilar, J., Ahlen, S., et al. 2025b, *JCAP*, **02**, 021  
 Almgren, A. S., Bell, J. B., Lijewski, M. J., Lukić, Z., & Andel, E. V. 2013, *ApJ*, **765**, 39  
 Baumann, D., Nicolis, A., Senatore, L., & Zaldarriaga, M. 2012, *JCAP*, **07**, 051  
 Besuner, R., Dey, A., Drica-Wagner, A., et al. 2025, arXiv:2503.07923  
 Bird, S., Fernandez, M., Ho, M.-F., et al. 2023, *JCAP*, **10**, 037  
 Blas, D., Lesgourgues, J., & Tram, T. 2011, *JCAP*, **07**, 034  
 Boera, E., Becker, G. D., Bolton, J. S., & Nasir, F. 2019, *ApJ*, **872**, 101  
 Carrasco, J. J. M., Foreman, S., Green, D., & Senatore, L. 2014, *JCAP*, **07**, 057  
 Carrasco, J. J. M., Hertzberg, M. P., & Senatore, L. 2012, *JHEP*, **09**, 082  
 Cen, R., Miralda-Escudé, J., Ostriker, J. P., & Rauch, M. 1994, *ApJL*, **437**, L9  
 Chabanier, S., Ravoux, C., Lattille, L., et al. 2024, *MNRAS*, **534**, 2674  
 Chen, S.-F., Vlah, Z., & White, M. 2021, *JCAP*, **05**, 053  
 Chudaykin, A., & Ivanov, M. M. 2025, *PhRvD*, **111**, 083515  
 Chudaykin, A., Ivanov, M. M., Philcox, O. H. E., & Simonović, M. 2020, *PhRvD*, **102**, 063533  
 Croft, R. A. C. 1998, in 18th Texas Symposium on Relativistic Astrophysics and Cosmology, ed. A. V. Olinto, J. A. Frieman, & D. N. Schramm (River Edge, NJ: World Scientific), 664  
 Croft, R. A. C., Weinberg, D. H., Pettini, M., Hernquist, L., & Katz, N. 1999, *ApJ*, **520**, 1  
 Cuceu, A., Font-Ribera, A., Joachimi, B., & Nadathur, S. 2021, *MNRAS*, **506**, 5439  
 de Belsunce, R., Chen, S.-F., Ivanov, M. M., et al. 2025a, *PhRvD*, **111**, 063524  
 de Belsunce, R., Ivanov, M. M., Sullivan, J. M., Akitsu, K., & Chen, S.-F. 2025b, arXiv:2507.00284  
 de Belsunce, R., Philcox, O. H. E., Irsic, V., et al. 2024, *MNRAS*, **533**, 3756  
 de Jong, R. S., Agertz, O., Berbel, A. A., et al. 2019, *Msngr*, **175**, 3  
 DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016, arXiv:1611.00036  
 Desjacques, V., Jeong, D., & Schmidt, F. 2018, *JCAP*, **12**, 035  
 Doughty, C. C., Hennawi, J. F., Davies, F. B., Lukić, Z., & Oñorbe, J. 2023, *MNRAS*, **525**, 3790  
 Farr, J., Font-Ribera, A., du Mas des Bourboux, H., et al. 2020, *JCAP*, **03**, 068  
 Friesen, B., Almgren, A., & Lukić, Z. 2016, *ComAC*, **3**, 4  
 Gerardi, F., Cuceu, A., Font-Ribera, A., Joachimi, B., & Lemos, P. 2022, *MNRAS*, **518**, 2567  
 Givans, J. J., & Hirata, C. M. 2020, *PhRvD*, **102**, 023515  
 Givans, J. J., Font-Ribera, A., Slosar, A., et al. 2022, *JCAP*, **09**, 070  
 Goodfellow, I., Pouget-Abadie, J., Mirza, M., et al. 2014, in *Advanced in Neural Information Processing Systems*, 27, ed. Z. Ghahramani et al. (New York: Curran Associates), 2672  
 Greene, J., Bezanson, R., Ouchi, M., Silverman, J. & The PFS Galaxy Evolution Working Group 2022, arXiv:2206.14908  
 Habib, S., Pope, A., Finkel, H., et al. 2016, *NewA*, **42**, 49  
 Hadzhiyska, B., de Belsunce, R., Cuceu, A., et al. 2025, *MNRAS*, **540**, 1960  
 Hadzhiyska, B., Font-Ribera, A., Cuceu, A., et al. 2023, *MNRAS*, **524**, 1008  
 Ivanov, M. M. 2022, arXiv:2212.08488  
 Ivanov, M. M. 2024, *PhRvD*, **109**, 023507  
 Ivanov, M. M., Toomey, M. W., & Karaçaylı, N. G. 2025, *PhRvL*, **134**, 091001  
 Jacobus, C., Harrington, P., & Lukić, Z. 2023, *ApJ*, **958**, 21  
 Jiang, L., Zhang, C., Huang, M., et al. 2020, in *Computer Vision—ECCV 2020*, ed. A. Vedaldi et al. (Berlin: Springer), 206  
 Karras, T., Laine, S., & Aila, T. 2019, arXiv:1812.04948  
 Karras, T., Laine, S., Aittala, M., et al. 2020, arXiv:1912.04958  
 Lukić, Z., Stark, C. W., Nugent, P., et al. 2015, *MNRAS*, **446**, 3697  
 McDonald, P. 2003, *ApJ*, **585**, 34  
 McDonald, P., Miralda-Escudé, J., Rauch, M., et al. 2000, *ApJ*, **543**, 1  
 McDonald, P., & Roy, A. 2009, *JCAP*, **08**, 020  
 McDonald, P., Seljak, U., Cen, R., et al. 2005, *ApJ*, **635**, 761  
 Meer, P., Baugher, E. S., & Rosenfeld, A. 1987, *TPAMI*, **9**, 512  
 Mirza, M., & Osindero, S. 2014, arXiv:1411.1784  
 Modi, C., & Philcox, O. H. E. 2023, arXiv:2309.10270  
 Palanque-Delabrouille, N., Yèche, C., Schöneberg, N., et al. 2020, *JCAP*, **04**, 038  
 Palanque-Delabrouille, N., Yèche, C., Borde, A., et al. 2013, *A&A*, **559**, A85  
 Palanque-Delabrouille, N., Yèche, C., Baur, J., et al. 2015, *JCAP*, **11**, 011  
 Park, T., Liu, M.-Y., Wang, T.-C., & Zhu, J.-Y. 2019, in *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (Piscataway, NJ: IEEE)*, 2337  
 Peirani, S., Prunet, S., Colombi, S., et al. 2022, *MNRAS*, **514**, 3222  
 Peirani, S., Weinberg, D. H., Colombi, S., et al. 2014, *ApJ*, **784**, 11  
 Pieri, M. M., Bonoli, S., Chaves-Montero, J., et al. 2016, in *SF2A-2016: Proc. of the Annual Meeting of the French Society of Astronomy and Astrophysics*, ed. C. Reylé et al. (Paris: Société Française d'Astronomie et d'Astrophysique), 259  
 Pillepich, A., Springel, V., Nelson, D., et al. 2018, *MNRAS*, **473**, 4077  
 Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, *A&A*, **594**, A13  
 Prokopenko, A., Arndt, D., Lebrun-Grandié, D., et al. 2024, *IJHPCA*, **39**, 167  
 Qezlou, M., Newman, A. B., Rudie, G. C., & Bird, S. 2022, *ApJ*, **930**, 109  
 Rogers, K. K., & Peiris, H. V. 2021, *PhysRevLett*, **126**, 071302  
 Schlegel, D. J., Ferraro, S., Aldering, G., et al. 2022, arXiv:2209.03585  
 Seljak, U., Makarov, A., McDonald, P., et al. 2005, *PhRvD*, **71**, 103515  
 Sexton, J., Lukic, Z., Almgren, A., et al. 2021, *JOSS*, **6**, 3068  
 Sinigaglia, F., Kitaura, F.-S., Balaguera-Antolínez, A., et al. 2022, *ApJ*, **927**, 230  
 Sinigaglia, F., Kitaura, F. S., Nagamine, K., Oku, Y., & Balaguera-Antolínez, A. 2024, *A&A*, **682**, A21  
 Sorini, D., Oñorbe, J., Lukić, Z., & Hennawi, J. F. 2016, *ApJ*, **827**, 97  
 Viel, M., Becker, G. D., Bolton, J. S., & Haehnelt, M. G. 2013, *PhRvD*, **88**, 043502  
 Villaescusa-Navarro, F., Anglés-Alcázar, D., Genel, S., et al. 2021, *ApJ*, **915**, 71  
 Villaescusa-Navarro, F., Genel, S., Anglés-Alcázar, D., et al. 2023, *ApJS*, **265**, 54  
 Walther, M., Armengaud, E., Ravoux, C., et al. 2021, *JCAP*, **04**, 059  
 Walther, M., Oñorbe, J., Hennawi, J. F., & Lukić, Z. 2019, *ApJ*, **872**, 13  
 Walther, M., Schöneberg, N., Chabanier, S., et al. 2025, *JCAP*, **2025**, 099  
 Wang, T.-C., Liu, M.-Y., Zhu, J.-Y., et al. 2018, in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (Piscataway, NJ: IEEE)*, 8798  
 Weinberger, R., Springel, V., Hernquist, L., et al. 2017, *MNRAS*, **465**, 3291  
 Zel'dovich, Y. B. 1970, *A&A*, **5**, 84