

RECEIVED: August 16, 2024
REVISED: December 21, 2024
ACCEPTED: February 26, 2025
PUBLISHED: March 28, 2025

Field-level emulation of cosmic structure formation with cosmology and redshift dependence

Drew Jamieson , ^a Yin Li , ^b Francisco Villaescusa-Navarro , ^c Shirley Ho  and David N. Spergel , ^{c,d}

^a Max-Planck-Institut für Astrophysik,
Karl-Schwarzschild-Straße 1, 85748 Garching, Germany

^b Department of Mathematics and Theory, Peng Cheng Laboratory,
Shenzhen, Guangdong 518066, China

^c Center for Computational Astrophysics, Flatiron Institute,
162 5th Avenue, New York, NY 10010, U.S.A.

^d Simons Foundation,
160 5th Avenue, New York, NY 10010, U.S.A.

E-mail: jamieson@mpa-garching.mpg.de, eelregit@gmail.com,
fvillaescusa@flatironinstitute.org, sho@flatironinstitute.org,
dspergel@flatironinstitute.org

ABSTRACT: We present a field-level emulator for large-scale structure, capturing the cosmology dependence and the time evolution of cosmic structure formation. The emulator maps linear displacement fields to their corresponding nonlinear displacements from N-body simulations at specific redshifts. Designed as a neural network, the emulator incorporates style parameters that encode dependencies on Ω_m and the linear growth factor $D(z)$ at redshift z . We train our model on the six-dimensional N-body phase space, predicting particle velocities as the time derivative of the model's displacement outputs. This innovation results in significant improvements in training efficiency and model accuracy. Tested on diverse cosmologies and redshifts not seen during training, the emulator achieves percent-level accuracy on scales of $k \sim 1 \text{ Mpc}^{-1} h$ at $z = 0$, with improved performance at higher redshifts. We compare predicted structure formation histories with N-body simulations via merger trees, finding consistent merger event sequences and statistical properties.

KEYWORDS: cosmic web, cosmological simulations, galaxy clustering, Machine learning

ARXIV EPRINT: [2408.07699](https://arxiv.org/abs/2408.07699)

Contents

1	Introduction	1
2	Model	2
3	Training	5
4	Results	7
4.1	Phase space	7
4.2	Merger trees	10
5	Conclusion	13

1 Introduction

As cosmological data analysis pushes to smaller scales, utilizes higher-order statistics [1–4], and implements field-level analyses [5–11] and simulations-based inference schemes [12–17], there is a growing demand for computationally efficient and accurate methods of predicting the outcomes of nonlinear cosmic structure formation. Traditional analysis methods of N -point statistics beyond the power spectrum and bispectrum require large mock datasets for accurate covariance estimation. The same is true for alternative summary statistics such as k -nearest neighbours [18–20] or scattering wavelets [21–25]. Simulations-based inference methods and field-level analyses require generating many accurate realizations of the late-time density field for constraining model parameters and initial conditions reconstruction. As next-generation galaxy survey data become available from DESI [26, 27], Euclid [28], LSST at the Vera C. Rubin Observatory [29], SPHEREx [30], and Subaru Prime Focus Spectrograph [31], obtaining the best possible constraints on cosmological parameters and initial conditions will require accelerated, highly accurate predictions for survey observables.

The most accurate method to predict the late-time, nonlinear matter field come from N -body simulations. These are too computationally expensive to use directly for inference. Fast, particle-mesh simulations [32–36] significantly reduce the computational cost of full N -body simulations at the expense of small-scale inaccuracy. Improving small scales in particle mesh algorithms requires both increasing the mesh resolution and decreasing the time step size, which quickly becomes too computationally expensive for inference. Achieving small-scale accuracy at low computational cost seem to require sacrificing the interpretability of simulations that directly integrate the equations of motion in favor of surrogate models, such as convolutional neural networks, that are not readily interpretable but can achieve highly accurate results. This is the strategy that we employ.

Recent advancements in machine-learning methods for modelling nonlinear cosmological structure formation have considerably improved computational efficiency and small-scale accuracy. These include error corrections for fast particle-mesh-based simulations [37], interpolation between two simulation snapshots [38], generative methods for super-resolution

simulations [39–42] along with their redshift dependence [43], normalizing flows [44, 45], and mapping between standard Λ CDM and modified gravity cosmologies [46]. The first neural network emulating the N-body mapping at the field level, including cosmology dependence, was demonstrated in [47]. This built upon previous work emulating particle-mesh simulations [48, 49]. The N-body neural network model was tested on a set of initial conditions, such as spherical collapse and pairs of isolated coupled modes, demonstrating its ability to accurately predict the physics of cosmic structure formation far outside of its training data [50]. This field-level emulator has been implemented in hybrid Lagrangian bias modelling of the galaxy field [51] and Hamiltonian Monte Carlo Bayesian inference of initial conditions [52].

In this work, we extend the field-level N-body emulator by adding redshift dependence and training on simulation snapshots at multiple redshifts. Due to the time dependence and autodifferentiability of this new model, we can efficiently obtain the N-body particle velocities as the time derivative of the output particle displacements. Since we can evaluate these velocities on the fly during training, we can define a loss function that depends on both the particle positions and velocities, training on the six-dimensional N-body phase space. Enforcing the physical constraint that velocities must equal the displacement time derivatives improves training efficiency and increases the model’s accuracy, especially for the velocity field. Our model is an example of physics-informed neural networks (PINNs) [53, 54]. We impose a prior on the relationship between the time-dependence of the particle coordinates and the particle velocities explicitly through our model design and, as we will describe later, through our choice of loss function.

The emulator is extremely fast, capable of predicting the nonlinear displacement and velocity fields for 128^3 particles in half a second on a single GPU. It is parallelized and scales well with multi-GPU processing, capable of generating realizations for an arbitrarily large box size. The model achieves percent-level accuracy down to scales of $k \sim 1 \text{ Mpc}^{-1} h$.

This paper is structured as follows. In section 2 we present the model architecture. In section 3 we describe the training data and loss function. Section 4 contains the results of our model tested on simulations that were not in its training set. We conclude in section 5.

2 Model

We describe N-body particles by their coordinates, \mathbf{x} , in a periodic simulation box. Each particle is associated with a site \mathbf{q} on a regular cubic lattice, so its position at redshift z is defined as,

$$\mathbf{x}(\mathbf{q}, z) = \mathbf{q} + \boldsymbol{\Psi}(\mathbf{q}, z). \quad (2.1)$$

Here $\boldsymbol{\Psi}$ is the displacement field, and \mathbf{q} are the Lagrangian coordinates of the particle. In the linear, Zeldovich approximation (ZA), the displacement field evolves as,

$$\boldsymbol{\Psi}_{\text{ZA}}(\mathbf{q}, z) = \frac{D(z)}{D(z_i)} \boldsymbol{\Psi}(\mathbf{q}, z_i), \quad (2.2)$$

where $D(z)$ is the linear growth factor, and z_i is a redshift chosen early enough that the displacement field is well-described by linear theory. As the nonlinearity of gravitational

clustering becomes important at late times, this linear approximation becomes inaccurate and nonperturbative methods of simulating cosmic structure formation, such as N-body simulations, become necessary (see [55] for a review).

The full phase space of the N-body system includes the particle velocities,

$$\dot{\Psi}(\mathbf{q}, z) = D(z)f(z)H(z) \frac{d\Psi(\mathbf{q}, z)}{dD(z)}. \quad (2.3)$$

Here $f(z) = d \log D(z) / d \log(a)$ is the linear growth rate, a is the scale factor, and $H(z)$ is the Hubble rate. The particle velocities are needed to model redshift-space distortions in galaxy surveys. We designed our field-level emulator to predict the nonlinear particle displacements and velocities at any redshift ranging from $z = 3$ –0 based on the ZA displacement field at the target redshift.

The field-level emulator is a convolutional neural network augmented with style parameters. We implement and train the model in PyTorch [56], using the `map2map` (<https://github.com/eelregit/map2map>) code base. The model and its parameters can be found at https://github.com/dsjamieson/m2m_nbody_zdep.

The neural network implements a U-Net/V-Net design, shown in figure 1. The input has three channels corresponding to the Cartesian components of the ZA displacements at the desired redshift, arranged in a 3D grid. The input passes through four ResNet $3 \times 3 \times 3$ convolutions. The first of these convolutional operations transforms the 3 input channels to 64 internal channels. The subsequent convolutions map from 64 to 64 internal channels. After the four convolutional operations, a copy of the outcome is stored for use on the upsampling side of the network before the result is downsampled with a $2 \times 2 \times 2$ convolutional kernel. The model repeats this sequence of operations two times, except these intermediate levels of the network use two instead of four convolutions before downsampling. After the fourth and deepest level of the network, we do not store a copy of the outcome. The model then upsamples the data with a $2 \times 2 \times 2$ convolutional filter. The stored outcome from the previous level is then appended to the data, doubling the number of internal channels to 128 before convolving twice and upsampling again. The first of these convolutions maps the 128 channels back down to 64. The model repeats this once before the final level of the network, which performs four convolutions. The final operation maps the internal 64 channels to a three-channel final result. The model then adds the input to the output. This addition means that the output represents the Cartesian components of the difference between the nonlinear and linear displacements. See [47] for additional details.

This architecture has a finite receptive field corresponding to 48 grid sites on either side of a given focal cell. We compute the receptive field's size based on the sequence of convolutions and downsampling/upsampling operations. Predicting a single particle's displacement,¹ a box of size 97^3 , centred on the focal particle would need to be passed through the network, which corresponds to a Lagrangian volume of comoving size $(189.45 \text{ Mpc } h^{-1})^3$. Since the network is missing information for all cells in the padding region, it cannot accurately predict the displacements for these particles and we must discard this padded region from the output.

¹The downsampling and upsampling operations will not process data that is odd in dimension, so the smallest size box that the model can process is for 2^3 particles, or 98^3 with the padding.

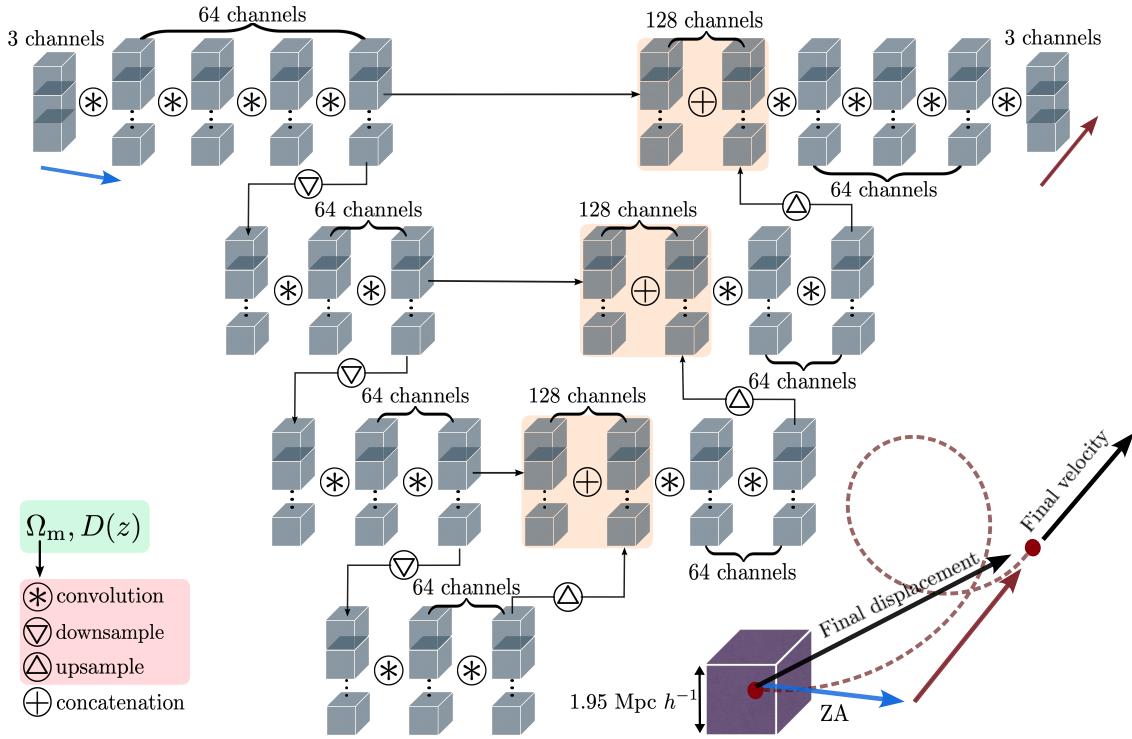


Figure 1. Here, we depict the architecture of the field-level emulator. In the bottom-right corner, we depict a single particle’s Lagrangian cell. Its initial ZA displacement is the blue vector. The 3D vector field of ZA displacements is input for the model. The red dashed curve shows the particle’s trajectory. The field-level emulator predicts the red vector, the residual between the linear and nonlinear displacements. The time derivative of the model’s displacement output is the final velocity vector. All convolutional operations, including the downsampling and upsampling ones, are modulated by parameters that depend on Ω_m and $D(z)$, allowing the emulator to capture cosmology and redshift dependence.

The network can process a box of even size as long as it is padded by 48 cells on all sides. The dimensions of the input and output are $3 \times (96 + N)^3$, which we crop down to the desired size of $3 \times N^3$ by discarding the padding. Since the simulations in our training and testing datasets are run in periodic boxes, the padding must be done periodically. There is an advantage to the limited receptive field of the network: it preserves the ZA field on large scales where linear theory is accurate.²

The network described so far can train on data from simulation snapshots at a single redshift for a fixed cosmology. To expand its capabilities, allowing the network to learn the N-

²One caveat here is that the neural network could induce long-range correlations if it added a global displacement, rotation, or compression to the box. A global translation would violate translation invariance, and a global rotation would violate isotropy. Due to our data augmentation strategy, randomly transforming the crops by symmetries of the cube, this cannot happen for the trained model. Such effects would also increase the loss function in a way that would be easy for the network to optimize away during training, so this is not such a significant caveat. One might also worry that the model would miss the separate universe couplings between small-scale modes and modes much larger than the receptive field. Since the form of the separate universe response is known analytically [57], one could correct for these in post-processing after running the emulator.

body mapping as a function of Ω_m and redshift, we augment it to include *style parameters* [58]. Before any of the convolutions including downsampling/upsampling operations are performed, the values of Ω_m and $D(z)$ for the snapshot are passed to the network and mapped to internal arrays matching the dimensions of the convolutional kernels. These parameters then modulate the network weights. In effect, this elevates the single neural network to a two-parameter family of neutral networks, continuously and autodifferentiably parameterized by the style inputs Ω_m and $D(z)$. We use $D(z)$ to quantify time dependence because it makes the particle velocities easier to obtain from eq. (2.3). The other Λ CDM cosmological parameters (σ_8 , n_s , Ω_b , and h), affect only the input, through the shape of the matter transfer functions. The physics of gravitational clustering does not depend on these parameters, so they would not be informative to the neural network as it learns the N-body mapping and should not be included as input for the style parameters.

3 Training

The style parameters and network parameters are trained simultaneously by randomly sampling snapshots from a suite of simulations with varied cosmological parameters and a fixed set of snapshot redshifts. We use the Quijote Latin hypercube simulations [59], which were run with 512^3 particles in a box of side length $1 \text{ Gpc } h^{-1}$. The Lagrangian space resolution for all of these simulations is $1.95 \text{ Mpc } h^{-1}$. Since the emulator operates in Lagrangian space, it can only reliably make predictions for this specific resolution. There are 2000 simulations in the simulations dataset. Each simulation has a unique set of the five Λ CDM cosmological parameters $\Omega_m \in [0.1, 0.5]$, $\Omega_b \in [0.015, 0.035]$, $h \in [0.5, 0.9]$, $n_s \in [0.85, 1.05]$, and $\sigma_8 \in [0.6, 1.0]$ sampled on a 5D latin hypercube. Each simulation also has unique initial conditions. Snapshots are stored at redshifts $z \in \{0, 0.5, 1, 2, 3\}$.

We partition the 2000 simulations into three sets: 1874 for training, 122 for validation, and 4 for testing. The previous work [47] had roughly equal numbers of testing and validation simulations. However, since the model errors depend only on the values of Ω_m and σ_8 , we take the most extreme values among these two parameters from the previous testing set and add one simulation from the Quijote fiducial cosmology (see [59]) for our testing data. In total, there are 5 test simulations. We added all other simulations to the training data. To encourage isotropy, we use data augmentation, randomly transforming the input and target data by the symmetries of a cube (see [47] for details).

The most challenging part of the mapping to learn is when the nonlinearity is strongest, at redshift $z = 0$. We already have a pretrained model for this redshift, so it is most efficient to initialize the new model to be nearly identical to the old model and have it learn the comparatively easier mapping at earlier redshifts. We accomplish this by initializing all network parameters with those from the pretrained, fixed-redshift model. Then we initialize the new style parameters randomly. We first trained only the style parameters for 10 epochs, holding the network parameters fixed. We then train all parameters for another 10 epochs.

Including redshift dependence has several benefits. Since the model is automatically differentiable with respect to its style inputs, the particle velocities can be obtained as the time derivatives of their displacements. Taking the forward mode derivative, or Jacobian-vector-product (JVP) [60], with respect to the style input $D(z)$, we use eq. (2.3) to obtain the

model prediction for the velocities. The fixed-redshift model [47] needed two separate neural networks: one for displacements and one for velocities. Now, we only need the displacement model, which eliminates the training time for the second model. This is already a significant advantage, but there is another benefit from adding redshift dependence. We can evaluate the velocities along with the displacements during training. This allows us to include terms involving the N-body particle velocities in the loss function. The model is then trained on the full N-body phase space. The network learns the particle trajectories as parametric curves depending on $D(z)$. The network simultaneously learns that the tangent vectors obtained by taking the time derivatives of these trajectories should match the particle velocities. This imposes a stringent physical constraint, requiring the particle velocities to equal the displacement time derivatives, which greatly improves training efficiency and model accuracy. This can also be viewed as multi-task learning [61].

The loss function we use contains four terms. The first is the mean squared error (MSE) for the particle displacements,

$$L_{\Psi} = \frac{1}{N} \sum_{i=1}^N |\Psi_{i,\text{SNN}} - \Psi_{i,\text{SIM}}|^2, \quad (3.1)$$

where $\Psi_{i,\text{SNN}}$ is the emulator displacement prediction for the i^{th} particle and $\Psi_{i,\text{SIM}}$ is the true N-body displacement. Here SIM refers to the simulations and SNN refers to the styled neural network. The sum is over the N particles in the cropped subbox. We also construct the Eulerian density fields, $\delta_{i,\text{SNN}}$ and $\delta_{i,\text{SIM}}$, by distributing the particles to a 3D grid using the trilinear, cloud-in-cell scheme. The second term in the loss is the MSE of the Eulerian densities,

$$L_{\delta} = \frac{1}{N} \sum_{i=1}^N |\delta_{i,\text{SNN}} - \delta_{i,\text{SIM}}|^2. \quad (3.2)$$

Here the sum is over grid sites, but the number of these is equal to the number of particles in the batch for our choice of Eulerian grid resolution. For the third loss term, we use eq. 2.3 and compute the MSE for the particle velocities

$$L_{\dot{\Psi}} = \frac{1}{N} \sum_{i=1}^N |\dot{\Psi}_{i,\text{SNN}} - \dot{\Psi}_{i,\text{SIM}}|^2. \quad (3.3)$$

The final term corresponds to the MSE loss of the Eulerian momentum field. We distribute the particle velocities to an Eulerian vector field mesh with the same resolution as the initial Lagrangian grid and compute,

$$L_{\mathbf{p}} = \frac{1}{N} \sum_{i=1}^N |\mathbf{p}_{i,\text{SNN}} - \mathbf{p}_{i,\text{SIM}}|^2, \quad (3.4)$$

where \mathbf{p} is the Eulerian momentum per particle mass in the grid cells. The final loss function for a snapshot at redshift z is

$$L = \frac{1}{1+z} \log(L_{\Psi}^6 L_{\delta}^2 L_{\dot{\Psi}}^3 L_{\mathbf{p}}^3). \quad (3.5)$$

The exponents are somewhat arbitrary, but they have the effect of weighting the terms of the loss function differently. We have not optimized these exponents, but this choice achieves good training efficiency. The prefactor in the loss function down-weights high-redshift snapshots, where the N-body mapping is easier for the neural network to learn. During the initial phase of training, where we fix the network parameters from the previous, fixed-redshift model and train only the style parameters, we omit this factor.

We could also choose an Eulerian grid of a higher resolution in order to improve performance on small scales. However, this tends to result in tradeoffs, degrading large-scale accuracy in favour of the smaller scales that dominate the information content in the fields. It may be possible to implement weighting schemes to tailor the loss function to preserve accuracy on larger scales, but we leave this to future work.

4 Results

4.1 Phase space

We evaluate the accuracy of the emulator by constructing the Eulerian density and momentum auto power spectra and cross power spectra for the emulator output and the N-body simulation ground truth. For the density fields $\delta_{\text{SNN}}(\mathbf{x})$ and $\delta_{\text{SIM}}(\mathbf{x})$ we distribute the particles to a 512^3 mesh using the CIC interpolation scheme and estimate the power spectra,

$$\langle \delta_I(\mathbf{k})\delta_J(\mathbf{k}') \rangle = (2\pi)^3 \delta_{\text{D}}^{(3)}(\mathbf{k} + \mathbf{k}') P_{\delta,IJ}(k), \quad (4.1)$$

where I and J can be either SNN or SIM.

For the momentum fields $\mathbf{p}_{\text{SNN}}(\mathbf{x})$ and $\mathbf{p}_{\text{SIM}}(\mathbf{x})$, we distribute the particles weighted by their velocity components to three grids (one for each velocity component) using the CIC scheme and estimator the power spectra

$$\langle \mathbf{p}_I(\mathbf{k}) \cdot \mathbf{p}_J(\mathbf{k}') \rangle = (2\pi)^3 \delta_{\text{D}}^{(3)}(\mathbf{k} + \mathbf{k}') P_{p,IJ}(k). \quad (4.2)$$

Note that this involves the dot product between the modes of the two vector fields.

We define the stochasticities,

$$s = 1 - \frac{(P_{\text{SNN} \times \text{SIM}})^2}{P_{\text{SNN}} P_{\text{SIM}}}, \quad (4.3)$$

where P_{SNN} and P_{SIM} are the auto power spectra for the emulator and the simulation respectively and $P_{\text{SNN} \times \text{SIM}}$ is their cross power spectrum. We also define the fractional transfer function error,

$$\frac{\Delta T}{T} = \sqrt{\frac{P_{\text{SNN}}}{P_{\text{SIM}}}} - 1. \quad (4.4)$$

For individual modes, the stochasticities quantify the error in phases and the transfer function errors quantify the error in the mode amplitudes. Binning the power spectrum, however, mixes contributions from phase and amplitude errors.

The emulator never encounters any of the 5 test simulations during training. The cosmological parameters for these simulations are listed in table 1. We reran these simulations

Ω_m	Ω_b	h	n_s	σ_8
0.3175	0.04900	0.6711	0.9624	0.834
0.1663	0.04783	0.6173	1.1467	0.6461
0.1289	0.06325	0.7293	1.1537	0.9489
0.4599	0.04055	0.7287	0.8505	0.7011
0.4423	0.03533	0.8267	1.0009	0.9151

Table 1. This table lists the cosmological parameters of the five simulations used to test our emulator. The first row corresponds to the Quijote fiducial cosmology, while the other four correspond to extreme corners of the (Ω_m, σ_8) subspace of the 5-parameter Qujote Latin hypercube.

and output a denser sample of redshifts, $z \in \{0.0, 0.25, 0.5, 0.75, 1.0, 1.5, 2.0, 2.5, 3.0\}$ to test the model performance when interpolating between the five fixed redshifts in the training data.

In figure 2 we show the errors as a function of scale with the color of each curve indicating redshift. The Eulerian density errors originate solely from errors in the particle displacements. Results from the original model from [47] are shown as the dashed red curves. At redshift $z = 0$, the stochasticities of the new time-dependent model are comparable to those of the original model, and the transfer function errors are generally improved over the original model. For reference, in the fiducial cosmology at redshift $z = 0$, stochasticities between the ZA and N-body density fields are 1% at $k = 0.093 h \text{ Mpc}^{-1}$ and 10% at $k = 0.30 h \text{ Mpc}^{-1}$.

The momentum errors arise from a combination of the particle velocities and displacements, which determine how the velocity components are distributed onto the Eulerian grid. The redshift-dependent model achieves significantly reduced errors compared with the original model at $z = 0$. This demonstrates the efficacy and constraining power of imposing the physical connection between displacement time dependence and velocity and training directly on the full dynamical N-body phase space. Stochasticities between the ZA and N-body density fields in the $z = 0$ fiducial cosmology are 1% at $k = 0.044 h \text{ Mpc}^{-1}$ and 10% at $k = 0.18 h \text{ Mpc}^{-1}$.

The emulator errors depend systematically on cosmology as well as redshift. These dependencies have the same underlying cause: the level of nonlinear clustering. The model performs worse at low redshift because the degree of nonlinear clustering is higher at late times. The model also performs worst on cosmologies with high σ_8 and low Ω_m , as was found in [47]. These cosmologies have the highest degree of nonlinearity, and correspondingly, the highest abundances of halos. More specifically, when Ω_m is high, the error has only weak dependence on σ_8 , as seen from the first and fourth columns in figure 2. Conversely, the second and fifth columns of figure 2 demonstrate that when Ω_m is low the error depends strongly on σ_8 . Similarly, the first and second columns of figure 2 show that when σ_8 is low the error depends very weakly on Ω_m . The two right-most columns of figure 2 show that high values of σ_8 lead to very strong dependence of the error on Ω_m .

The error depends monotonically and correlates positively with σ_8 . The amplitude of small-scale power, and thus the formation of halos is controlled by σ_8 . The highly chaotic virialized motion inside of halos makes it the most difficult regime of the N-body map for the emulator to learn, so increasing σ_8 tends to degrade the emulator's accuracy. The

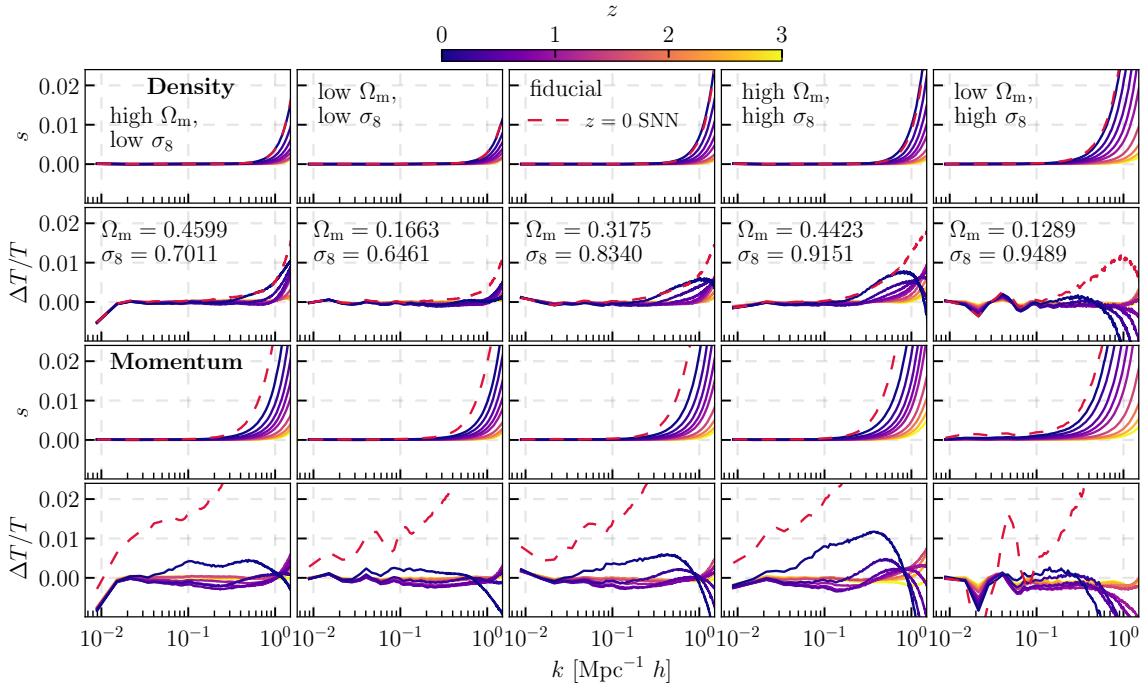


Figure 2. This figure shows the SNN emulator’s power spectrum errors. The top row shows the cross-correlation error, or stochasticities, for the density field. The second row shows the amplitude errors, or relative transfer function errors, for the densities. The third and fourth rows show the same for the momentum field. The colour of each curve indicates its redshift, according to the colour bar at the top. The columns show the errors for specific cosmologies, corresponding to extreme corners of the Latin hypercube, and the fiducial cosmology at the center of the hypercube. The red dashed line shows the results from the original, redshift-independent model [47], which could only map to $z = 0$ and had two independent neural networks for the displacements and velocities.

small-scale errors scale as k^2 . Since we are dealing with displacements, mass is automatically conserved and there can be no k^0 contribution to the error. The next leading term should scale as k^2 if the network predictions are isotropic, so this error scaling confirms that the data augmentation used to enforce isotropy is effective.

Ref. [51] found that parameters used to model small-scale stochasticity and redshift space distortions can absorb much of the excess error for high σ_8 cosmologies. This significantly mitigates the systematic cosmology-dependence of the error. Since there is a large degree of uncertainty over the values of these parameters, and they will be treated as nuisance parameters and marginalized over, the σ_8 dependence of the errors may not be a significant issue for inference. This will be tested more rigorously in future work.

Another notable feature of the errors is the large-scale bias appearing in the first column of figure 2. Here, the transfer function for the most extreme, high Ω_m , low σ_8 testing simulation has a systematically low transfer function in the first power spectrum bin. Since this cosmology is near the corner of the Latin hypercube where performance is expected to be worst, and since there is a low number of nearby samples in the training dataset here, this indicates the model’s limitations when extrapolating in cosmological parameter space. This large-scale bias is absent for cosmologies further away from the edge of the Latin hypercube.

In the right-most column of figure 2, we also see oscillatory errors due to the emulator not perfectly predicting the BAO amplitude. These features of the emulator errors are subpercent level, and would likely improve with more training data.

Errors in higher-order statistics of the original model were explored in ref. [47]. Field-level error diagnostics were also considered in ref. [51]. The redshift-dependent model performs comparably to these for real space density statistics at $z = 0$, and significantly better for both redshift space and real-space statistics at higher redshifts, so we do not present these higher-order statistics errors here.

The errors increase smoothly and monotonically with decreasing redshift. This demonstrates that the emulator interpolates effectively between the small number of fixed-redshift snapshots in its training data. The model is not overfitting to match the training redshifts, otherwise we would see oscillatory features for errors at intermediate times. A reason that could explain the lack of overfitting for such a small sample size of redshifts is that redshift is not strongly connected to gravitational clustering as a time coordinate. Cosmic time, for example, varies with cosmology at fixed redshift. The network is exposed to a wide sample of cosmic times, due to the range of cosmological parameters in its training data. The network can infer from this data how to accurately interpolate between redshifts at fixed cosmology.

4.2 Merger trees

A more stringent test of the field-level emulator’s accuracy is to analyze the sequence of structure formation that it predicts and check whether this matches the N-body simulation. Structure formation occurs in a bottom-up sequence, where small objects collapse first and then merge to form larger objects, which continue merging and forming increasingly massive objects.

To test our emulator’s predictions for formation histories, we reran the fiducial cosmology simulation, outputting 64 snapshots between redshifts $z = 3$ and $z = 0$. We then ran the Gadget4 [62] halo finder, subhalo finder SUBFIND [63], and merger-tree construction algorithm [64]. Figure 3 displays four merger trees for final halos occurring in the emulator prediction and the N-body simulation. In each panel, the vertical axis is redshift, and the three horizontal axes are the x , y , and z Cartesian coordinates with the final halo at the origin. The N-body results are in the top three panels of each subplot, and the emulator results appear in the bottom three panels. Overall, the merger sequences from the emulator and simulation qualitatively agree. The halos have comparable environments, satellites, and substructures throughout their formation histories.

A closer inspection of figure 3 uncovers some differences. The emulator will occasionally merge a satellite that persists at late times in the simulation or miss the merger of a satellite. In the bottom left, the emulator misses a late-time satellite not resolved by the halo finder. In the bottom right, the emulator has merged a satellite into the main branch halo that did not merge in the simulation. In the top left, the emulator has a large satellite with a trajectory in the x -direction that does not follow the simulation. The emulator does not perfectly predict all small-scale structure formation details, but it faithfully reproduces the most salient features of the merger histories on \sim Mpc scales. A schematic version of these merger trees that more clearly demonstrates the sequences of events is depicted in figure 4.

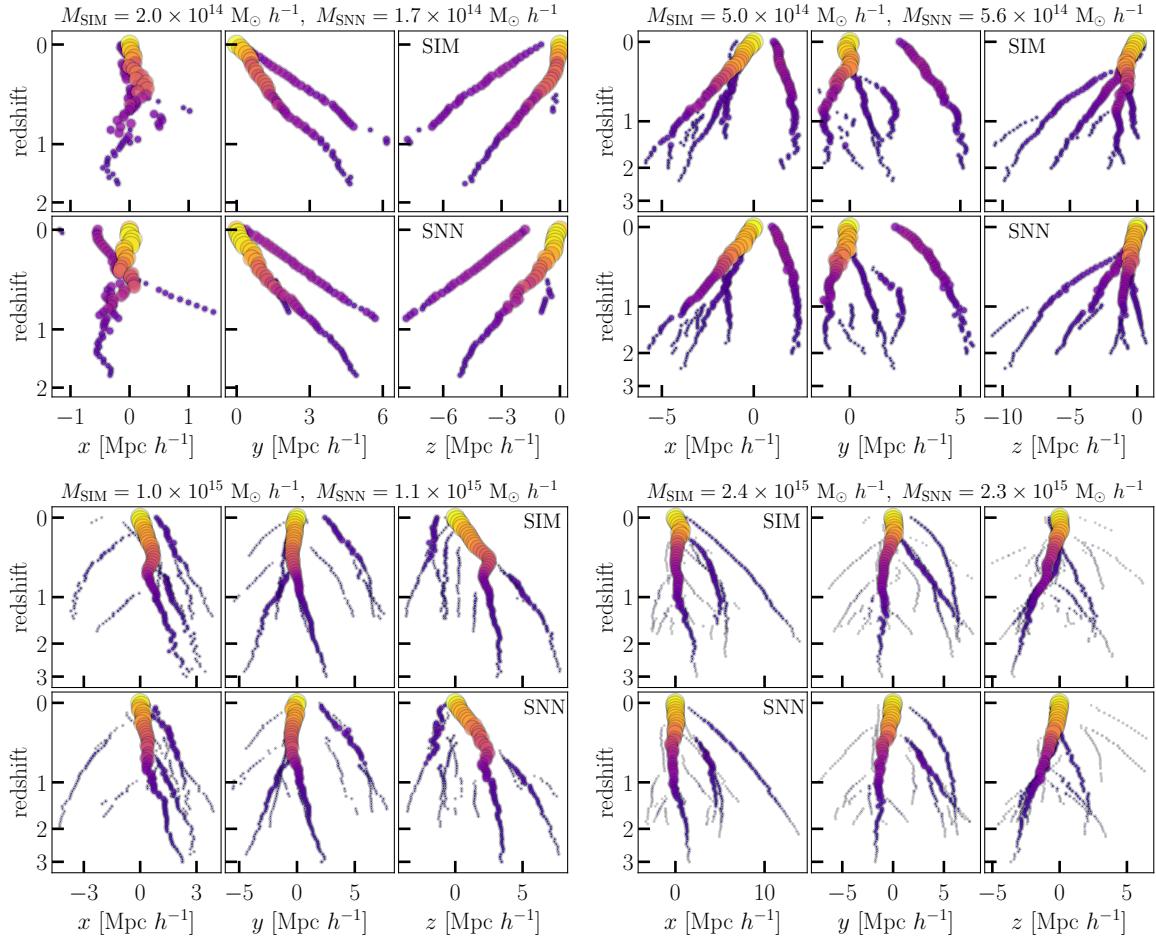


Figure 3. Merger trees for four halos of different masses. Each of the four sets of plots shows the positions of halos along the x , y , and z directions on the horizontal axes as a function of redshift on the vertical axis. The upper panels show the merger trees from the N-body simulations and the lower panels show the merger trees from the SNN emulator. The main branch and satellites are included. The colour and size of each circle indicate the mass of the object relative to the final halo at redshift zero. The size and color of the points indicate the masses of the halos, but the absolute size is arbitrary, which is why the minor progenitors do not overlap with the major ones at redshift $z = 0$.

Following [43], we also test the statistical properties among the whole set of merger trees. Specifically, we quantify the statistics of mass accretion rates and progenitor formation times for main branch halos in each merger tree (see [65] for a detailed discussion of merger tree statistics, including those we use here). Note that a comparison to the results of [43] is not possible since their resolution is significantly higher, making the substructures they resolve much lower in mass. Also, their model is generative, whereas ours is deterministic.

Between two sequential snapshots, the main branch halo in the merger tree will undergo a mass change from M_i in the earlier snapshot at scale factor a_i , to M_{i+1} in the later snapshot at scale factor a_{i+1} . We define the fractional rate of mass change between these snapshots,

$$\frac{M'}{M} \equiv \frac{\log(M_{i+1}/M_i)}{\log(a_{i+1}/a_i)}. \quad (4.5)$$

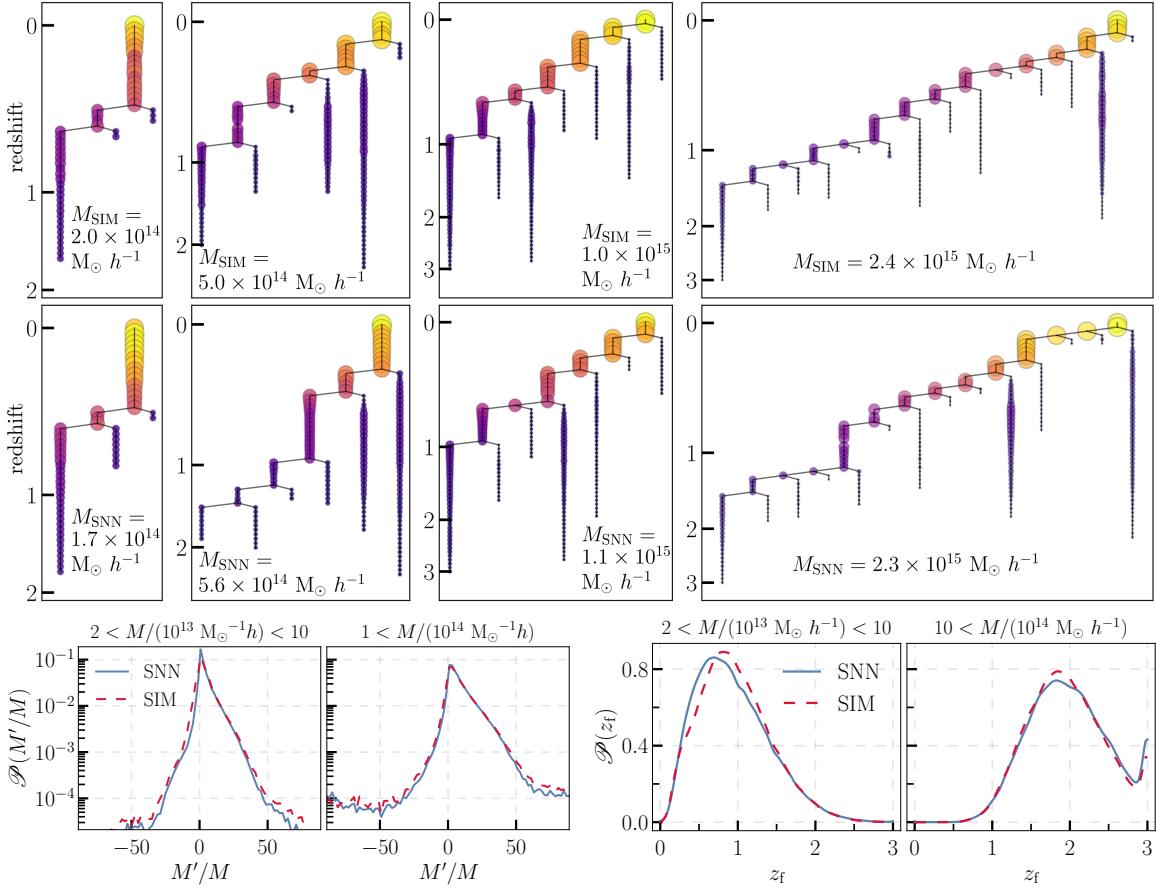


Figure 4. The top two rows show the merger trees for four halos of different masses. The top row shows the simulations merger trees and the second row shows the SNN emulator merger trees for the same halos. The colour and size of each circle indicate the halo’s mass relative to the final halo at redshift zero. The bottom row shows the probability densities for the logarithmic rate of mass accretion (left) and the earliest redshift along the main branch (right), for two different halo mass bins.

We then estimate the probability density function (PDF) among all main-branch halos in each merger tree between every sequential pair of snapshots. These are displayed in the bottom-left of figure 4 for two different mass bins. The emulator results agree well with the simulation results in the high-mass bin. The plot is on a log scale, so the emulator PDF for the low-mass bin is more sharply peaked and has smaller dispersion compared with the simulation. This discrepancy is unsurprising. The emulator makes errors on \sim Mpc scales similar to the size of a halo. Thus, the emulator predicts fuzzier halos than the simulation, with a more diffuse density profile. This tends to smooth out the mass accretion, leading to more incremental accretion events and fewer major mergers with dramatic changes in mass. This effect is much more pronounced for low-mass halos.

Since the halo profiles from the emulator differ systematically from those of the N-body simulation, statistics of mass-selected populations will also differ. This was investigated at redshift $z = 0$ in a previous paper for halos in real space and redshift space [47]. The

catalogues of collapsed objects from the emulator match those from the simulation well, it is only the properties assigned to these objects using spherical overdensity that differ significantly. One could therefore use abundance matching or modify the halo property assignment methods (see for example [66–68]) to improve the accuracy of the halo catalogues.

We also define the earliest redshift that a main branch halo appears as z_f . The lower-right panels of figure 4 display the PDFs for z_f in two mass bins, which agree between the emulator and the simulation. Since higher-mass halos form from a sequence of merger events of smaller-mass halos, they typically have earlier z_f compared with lower-mass halos. This is reflected in both the simulation and the emulator. For the low-mass bin at around redshift $z = 1$, the simulation halos appear slightly earlier compared with those in the emulator, which can also be understood in terms of the emulator’s more diffuse halos. This diffuseness means that the halo finder resolves these late-time, low-mass emulator objects at slightly later times than in the simulation. The peak near redshift $z = 3$ for the high-mass bin the bottom-right panel of figure 4 is due to the fact that we have no snapshot earlier than $z = 3$, so all halos that form at earlier redshifts are placed in this bin.

These merger tree statistics will be sensitive to the choice of redshift sampling. However, our sampling choice is already much higher than that of the training data. The consistency between the emulator and simulation merger sequences confirms that the emulator accurately models structure formation histories on Mpc scales. This agreement suggests that the emulator can be used alongside semi-analytic models (SAMs) mapping from the initial conditions to galaxies and incorporating hydrodynamic effects of Baryonic physics (see e.g. refs. [69–71] and references therein). This is outside the scope of the current work, but we plan to investigate this in the future.

5 Conclusion

In this work, we presented results and tests of a new field-level emulator for large-scale structure formation. The emulator is a neural network with style parameters that capture cosmology and redshift dependence. The model maps the ZA Lagrangian displacement field at redshift z to the nonlinear displacement field at the same redshift. The model predicts particle velocities through the JVP, forward-mode automatic differentiation of the predicted displacements with respect to time. We trained the model on the six-dimensional phase space of N-body evolution.

The emulator achieves percent-level accuracy on scales of $k \sim \text{Mpc } h^{-1}$ in the amplitudes and phases of the density field and momentum field at redshift $z = 0$, as shown in figure 2. The emulator performs significantly better at higher redshifts when the nonlinearity is less severe. Our new time-dependent model demonstrates significant improvement over the previous fixed-redshift model [47] due to the physical constraint imposed during training, that the velocities correspond to time derivatives of the displacements.

We also tested the model’s consistency in reproducing the history of structure formation through dark matter merger trees. We found consistency between the emulator predictions and the simulations in terms of halo environments and merger sequences. The emulator smooths out the mass accretion somewhat compared to the simulations, which is most noticeable for low-mass halos.

Using mixed-precision model evaluations, the emulator can process a box of 128^3 particles (an input of 224^3 with padding) in under half a second on GPU. The model is parallelized for multi-GPU processing and scales well with box size. The accuracy and speed of our model make it particularly well-suited to field-level inference. The model can also efficiently generate large, high-quality mock datasets for simulations-based inference schemes and high-order N-point statistics analyses. Our field-level emulator provides a power tool for achieving robust and optimal constraints from next-generation cosmological survey data.

Acknowledgments

We thank Raul Angulo, Ludvig Doeser, Marcos Pellejero Ibáñez, Jens Jasche, Eiichiro Komatsu, and Volker Springel for their helpful discussions. We thank the Flatiron Scientific Computing Core for providing the computational resources used in this work. The Flatiron Institute is supported by the Simons Foundation.

References

- [1] D. Gualdi and L. Verde, *Integrated trispectrum detection from BOSS DR12 NGC CMASS*, *JCAP* **09** (2022) 050 [[arXiv:2201.06932](https://arxiv.org/abs/2201.06932)] [[INSPIRE](#)].
- [2] O.H.E. Philcox, J. Hou and Z. Slepian, *A First Detection of the Connected 4-Point Correlation Function of Galaxies Using the BOSS CMASS Sample*, [arXiv:2108.01670](https://arxiv.org/abs/2108.01670) [[INSPIRE](#)].
- [3] J. Hou, Z. Slepian and R.N. Cahn, *Measurement of parity-odd modes in the large-scale 4-point correlation function of Sloan Digital Sky Survey Baryon Oscillation Spectroscopic Survey twelfth data release CMASS and LOWZ galaxies*, *Mon. Not. Roy. Astron. Soc.* **522** (2023) 5701 [[arXiv:2206.03625](https://arxiv.org/abs/2206.03625)] [[INSPIRE](#)].
- [4] O.H.E. Philcox, *Probing parity violation with the four-point correlation function of BOSS galaxies*, *Phys. Rev. D* **106** (2022) 063501 [[arXiv:2206.04227](https://arxiv.org/abs/2206.04227)] [[INSPIRE](#)].
- [5] SIMBIG collaboration, *Field-level simulation-based inference of galaxy clustering with convolutional neural networks*, *Phys. Rev. D* **109** (2024) 083536 [[arXiv:2310.15256](https://arxiv.org/abs/2310.15256)] [[INSPIRE](#)].
- [6] C.H. Hahn et al., *SimBIG: mock challenge for a forward modeling approach to galaxy clustering*, *JCAP* **04** (2023) 010 [[arXiv:2211.00660](https://arxiv.org/abs/2211.00660)] [[INSPIRE](#)].
- [7] S. Stopyra et al., *Towards accurate field-level inference of massive cosmic structures*, *Mon. Not. Roy. Astron. Soc.* **527** (2023) 1244 [[arXiv:2304.09193](https://arxiv.org/abs/2304.09193)] [[INSPIRE](#)].
- [8] A. Andrews, J. Jasche, G. Lavaux and F. Schmidt, *Bayesian field-level inference of primordial non-Gaussianity using next-generation galaxy surveys*, *Mon. Not. Roy. Astron. Soc.* **520** (2023) 5746 [[arXiv:2203.08838](https://arxiv.org/abs/2203.08838)] [[INSPIRE](#)].
- [9] J. Stadler, F. Schmidt and M. Reinecke, *Cosmology inference at the field level from biased tracers in redshift-space*, *JCAP* **10** (2023) 069 [[arXiv:2303.09876](https://arxiv.org/abs/2303.09876)] [[INSPIRE](#)].
- [10] N.-M. Nguyen et al., *How Much Information Can Be Extracted from Galaxy Clustering at the Field Level?*, *Phys. Rev. Lett.* **133** (2024) 221006 [[arXiv:2403.03220](https://arxiv.org/abs/2403.03220)] [[INSPIRE](#)].
- [11] I. Babić, F. Schmidt and B. Tucci, *Straightening the Ruler: Field-Level Inference of the BAO Scale with LEFTfield*, [arXiv:2407.01524](https://arxiv.org/abs/2407.01524) [[INSPIRE](#)].

- [12] J. Alsing, T. Charnock, S. Feeney and B. Wandelt, *Fast likelihood-free cosmology with neural density estimators and active learning*, *Mon. Not. Roy. Astron. Soc.* **488** (2019) 4440 [[arXiv:1903.00007](https://arxiv.org/abs/1903.00007)] [[INSPIRE](#)].
- [13] K. Cranmer, J. Brehmer and G. Louppe, *The frontier of simulation-based inference*, *Proc. Nat. Acad. Sci.* **117** (2020) 30055 [[arXiv:1911.01429](https://arxiv.org/abs/1911.01429)] [[INSPIRE](#)].
- [14] E. Massara et al., *SimBIG: Cosmological Constraints using Simulation-Based Inference of Galaxy Clustering with Marked Power Spectra*, [arXiv:2404.04228](https://arxiv.org/abs/2404.04228) [[INSPIRE](#)].
- [15] C. Modi et al., *Sensitivity analysis of simulation-based inference for galaxy clustering*, *Mon. Not. Roy. Astron. Soc.* **536** (2024) 254 [[arXiv:2309.15071](https://arxiv.org/abs/2309.15071)] [[INSPIRE](#)].
- [16] B. Tucci and F. Schmidt, *EFTofLSS meets simulation-based inference: σ_8 from biased tracers*, *JCAP* **05** (2024) 063 [[arXiv:2310.03741](https://arxiv.org/abs/2310.03741)] [[INSPIRE](#)].
- [17] K. Lin, M. von Wietersheim-Kramsta, B. Joachimi and S. Feeney, *A simulation-based inference pipeline for cosmic shear with the Kilo-Degree Survey*, *Mon. Not. Roy. Astron. Soc.* **524** (2023) 6167 [[arXiv:2212.04521](https://arxiv.org/abs/2212.04521)] [[INSPIRE](#)].
- [18] S. Yuan, T. Abel and R.H. Wechsler, *Robust cosmological inference from non-linear scales with k -th nearest neighbour statistics*, *Mon. Not. Roy. Astron. Soc.* **527** (2023) 1993 [[arXiv:2310.04501](https://arxiv.org/abs/2310.04501)] [[INSPIRE](#)].
- [19] W.R. Coulton, T. Abel and A. Banerjee, *Small-scale signatures of primordial non-Gaussianity in k -nearest neighbour cumulative distribution functions*, *Mon. Not. Roy. Astron. Soc.* **534** (2024) 1621 [[arXiv:2309.15151](https://arxiv.org/abs/2309.15151)] [[INSPIRE](#)].
- [20] Y. Wang, A. Banerjee and T. Abel, *Detection of spatial clustering in the 1000 richest SDSS DR8 redMaPPer clusters with nearest neighbor distributions*, *Mon. Not. Roy. Astron. Soc.* **514** (2022) 3828 [[arXiv:2112.04502](https://arxiv.org/abs/2112.04502)] [[INSPIRE](#)].
- [21] S. Cheng, Y.-S. Ting, B. Ménard and J. Bruna, *A new approach to observational cosmology using the scattering transform*, *Mon. Not. Roy. Astron. Soc.* **499** (2020) 5902 [[arXiv:2006.08561](https://arxiv.org/abs/2006.08561)] [[INSPIRE](#)].
- [22] G. Valogiannis and C. Dvorkin, *Towards an optimal estimation of cosmological parameters with the wavelet scattering transform*, *Phys. Rev. D* **105** (2022) 103534 [[arXiv:2108.07821](https://arxiv.org/abs/2108.07821)] [[INSPIRE](#)].
- [23] M. Eickenberg et al., *Wavelet Moments for Cosmological Parameter Estimation*, [arXiv:2204.07646](https://arxiv.org/abs/2204.07646) [[INSPIRE](#)].
- [24] G. Valogiannis and C. Dvorkin, *Going beyond the galaxy power spectrum: An analysis of BOSS data with wavelet scattering transforms*, *Phys. Rev. D* **106** (2022) 103509 [[arXiv:2204.13717](https://arxiv.org/abs/2204.13717)] [[INSPIRE](#)].
- [25] SIMBIG collaboration, *Galaxy clustering analysis with SimBIG and the wavelet scattering transform*, *Phys. Rev. D* **109** (2024) 083535 [[arXiv:2310.15250](https://arxiv.org/abs/2310.15250)] [[INSPIRE](#)].
- [26] DESI collaboration, *The DESI Experiment Part I: Science, Targeting, and Survey Design*, [arXiv:1611.00036](https://arxiv.org/abs/1611.00036) [[INSPIRE](#)].
- [27] DESI collaboration, *DESI 2024 VI: cosmological constraints from the measurements of baryon acoustic oscillations*, *JCAP* **02** (2025) 021 [[arXiv:2404.03002](https://arxiv.org/abs/2404.03002)] [[INSPIRE](#)].
- [28] L. Amendola et al., *Cosmology and fundamental physics with the Euclid satellite*, *Living Rev. Rel.* **21** (2018) 2 [[arXiv:1606.00180](https://arxiv.org/abs/1606.00180)] [[INSPIRE](#)].

[29] LSST collaboration, *LSST: from Science Drivers to Reference Design and Anticipated Data Products*, *Astrophys. J.* **873** (2019) 111 [[arXiv:0805.2366](https://arxiv.org/abs/0805.2366)] [[INSPIRE](#)].

[30] SPHEREx collaboration, *Cosmology with the SPHEREx All-Sky Spectral Survey*, [arXiv:1412.4872](https://arxiv.org/abs/1412.4872) [[INSPIRE](#)].

[31] PFS TEAM collaboration, *Extragalactic science, cosmology, and Galactic archaeology with the Subaru Prime Focus Spectrograph*, *Publ. Astron. Soc. Jap.* **66** (2014) R1 [[arXiv:1206.0737](https://arxiv.org/abs/1206.0737)] [[INSPIRE](#)].

[32] Y. Feng, M.-Y. Chu, U. Seljak and P. McDonald, *FastPM: a new scheme for fast simulations of dark matter and haloes*, *Mon. Not. Roy. Astron. Soc.* **463** (2016) 2273 [[arXiv:1603.00476](https://arxiv.org/abs/1603.00476)] [[INSPIRE](#)].

[33] S. Tashev, M. Zaldarriaga and D. Eisenstein, *Solving Large Scale Structure in Ten Easy Steps with COLA*, *JCAP* **06** (2013) 036 [[arXiv:1301.0322](https://arxiv.org/abs/1301.0322)] [[INSPIRE](#)].

[34] C. Howlett, M. Manera and W.J. Percival, *L-PICOLA: A parallel code for fast dark matter simulation*, *Astron. Comput.* **12** (2015) 109 [[arXiv:1506.03737](https://arxiv.org/abs/1506.03737)] [[INSPIRE](#)].

[35] C. Modi, F. Lanusse and U. Seljak, *FlowPM: Distributed TensorFlow implementation of the FastPM cosmological N-body solver*, *Astron. Comput.* **37** (2021) 100505 [[arXiv:2010.11847](https://arxiv.org/abs/2010.11847)] [[INSPIRE](#)].

[36] Y. Li et al., *Differentiable Cosmological Simulation with the Adjoint Method*, *Astrophys. J. Suppl.* **270** (2024) 36 [[arXiv:2211.09815](https://arxiv.org/abs/2211.09815)] [[INSPIRE](#)].

[37] N. Kaushal et al., *NECOLA: Toward a Universal Field-level Cosmological Emulator*, *Astrophys. J.* **930** (2022) 115 [[arXiv:2111.02441](https://arxiv.org/abs/2111.02441)] [[INSPIRE](#)].

[38] C. Chen et al., *Learning the Evolution of the Universe in N-body Simulations*, in the proceedings of the 34th Conference on Neural Information Processing Systems, Online Conference, Canada (2020) [[arXiv:2012.05472](https://arxiv.org/abs/2012.05472)] [[INSPIRE](#)].

[39] D.K. Ramanah, T. Charnock, F. Villaescusa-Navarro and B.D. Wandelt, *Super-resolution emulator of cosmological simulations using deep physical models*, *Mon. Not. Roy. Astron. Soc.* **495** (2020) 4227 [[arXiv:2001.05519](https://arxiv.org/abs/2001.05519)] [[INSPIRE](#)].

[40] Y. Li et al., *AI-assisted super-resolution cosmological simulations*, [arXiv:2010.06608](https://arxiv.org/abs/2010.06608) [[DOI:10.1073/pnas.2022038118](https://doi.org/10.1073/pnas.2022038118)] [[INSPIRE](#)].

[41] Y. Ni et al., *AI-assisted superresolution cosmological simulations — II. Halo substructures, velocities, and higher order statistics*, *Mon. Not. Roy. Astron. Soc.* **507** (2021) 1021 [[arXiv:2105.01016](https://arxiv.org/abs/2105.01016)] [[INSPIRE](#)].

[42] D. Schaurecker et al., *Super-resolving Dark Matter Halos using Generative Deep Learning*, [arXiv:2111.06393](https://arxiv.org/abs/2111.06393) [[INSPIRE](#)].

[43] X. Zhang et al., *AI-assisted super-resolution cosmological simulations III: time evolution*, *Mon. Not. Roy. Astron. Soc.* **528** (2024) 281 [[arXiv:2305.12222](https://arxiv.org/abs/2305.12222)] [[INSPIRE](#)].

[44] B. Dai and U. Seljak, *Translation and rotation equivariant normalizing flow (TRENF) for optimal cosmological analysis*, *Mon. Not. Roy. Astron. Soc.* **516** (2022) 2363 [[arXiv:2202.05282](https://arxiv.org/abs/2202.05282)] [[INSPIRE](#)].

[45] B. Dai and U. Seljak, *Multiscale Flow for robust and optimal cosmological analysis*, *Proc. Nat. Acad. Sci.* **121** (2024) e2309624121 [[arXiv:2306.04689](https://arxiv.org/abs/2306.04689)] [[INSPIRE](#)].

[46] D. Saadeh, K. Koyama and X. Morice-Atkinson, *A field-level emulator for modified gravity*, *Mon. Not. Roy. Astron. Soc.* **537** (2025) 448 [[arXiv:2406.03374](https://arxiv.org/abs/2406.03374)] [[INSPIRE](#)].

[47] D. Jamieson et al., *Field-level Neural Network Emulator for Cosmological N-body Simulations*, *Astrophys. J.* **952** (2023) 145 [[arXiv:2206.04594](https://arxiv.org/abs/2206.04594)] [[INSPIRE](#)].

[48] S. He et al., *Learning to Predict the Cosmological Structure Formation*, *Proc. Nat. Acad. Sci.* **116** (2019) 13825 [[arXiv:1811.06533](https://arxiv.org/abs/1811.06533)] [[INSPIRE](#)].

[49] R. Alves de Oliveira et al., *Fast and Accurate Non-Linear Predictions of Universes with Deep Learning*, in the proceedings of the 34th Conference on Neural Information Processing Systems, Online Conference, Canada (2020) [[arXiv:2012.00240](https://arxiv.org/abs/2012.00240)] [[INSPIRE](#)].

[50] D. Jamieson et al., *Simple lessons from complex learning: what a neural network model learns about cosmic structure formation*, *PNAS Nexus* **2** (2023) gac250 [[arXiv:2206.04573](https://arxiv.org/abs/2206.04573)] [[INSPIRE](#)].

[51] M.P. Ibanez, R.E. Angulo, D. Jamieson and Y. Li, *Hybrid bias and displacement emulators for field-level modelling of galaxy clustering in real and redshift space*, *Mon. Not. Roy. Astron. Soc.* **529** (2024) 89 [[arXiv:2307.09134](https://arxiv.org/abs/2307.09134)] [[INSPIRE](#)].

[52] L. Doeser et al., *Bayesian inference of initial conditions from non-linear cosmic structures using field-level emulators*, *Mon. Not. Roy. Astron. Soc.* **535** (2024) 1258 [[arXiv:2312.09271](https://arxiv.org/abs/2312.09271)] [[INSPIRE](#)].

[53] M. Raissi, P. Perdikaris and G.E. Karniadakis, *Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations*, *J. Comput. Phys.* **378** (2019) 686 [[arXiv:1711.10561](https://arxiv.org/abs/1711.10561)] [[INSPIRE](#)].

[54] Z. Hao et al., *Physics-Informed Machine Learning: A Survey on Problems, Methods and Applications*, [arXiv:2211.08064](https://arxiv.org/abs/2211.08064).

[55] R.E. Angulo and O. Hahn, *Large-scale dark matter simulations*, [arXiv:2112.05165](https://arxiv.org/abs/2112.05165) [[DOI:10.1007/s41115-021-00013-z](https://doi.org/10.1007/s41115-021-00013-z)] [[INSPIRE](#)].

[56] A. Paszke et al., *PyTorch: An Imperative Style, High-Performance Deep Learning Library*, in *Advances in Neural Information Processing Systems 32*, H. Wallach et al. eds., Curran Associates, Inc., (2019), pg. 8026.

[57] T. Baldauf, U. Seljak, L. Senatore and M. Zaldarriaga, *Galaxy Bias and non-Linear Structure Formation in General Relativity*, *JCAP* **10** (2011) 031 [[arXiv:1106.5507](https://arxiv.org/abs/1106.5507)] [[INSPIRE](#)].

[58] T. Karras et al., *Analyzing and Improving the Image Quality of StyleGAN*, [arXiv:1912.04958](https://arxiv.org/abs/1912.04958).

[59] F. Villaescusa-Navarro et al., *The Quijote simulations*, *Astrophys. J. Suppl.* **250** (2020) 2 [[arXiv:1909.05273](https://arxiv.org/abs/1909.05273)] [[INSPIRE](#)].

[60] A. Griewank and A. Walther, *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*, Second Edition, Society for Industrial and Applied Mathematics (2008) [[DOI:10.1137/1.9780898717761](https://doi.org/10.1137/1.9780898717761)].

[61] Y. Zhang and Q. Yang, *A Survey on Multi-Task Learning*, *IEEE Trans. Knowl. Data Eng.* **34** (2021) 5586.

[62] V. Springel, R. Pakmor, O. Zier and M. Reinecke, *Simulating cosmic structure formation with the gadget-4 code*, *Mon. Not. Roy. Astron. Soc.* **506** (2021) 2871 [[arXiv:2010.03567](https://arxiv.org/abs/2010.03567)] [[INSPIRE](#)].

[63] V. Springel, S.D.M. White, G. Tormen and G. Kauffmann, *Populating a cluster of galaxies. 1. Results at $z = 0$* , *Mon. Not. Roy. Astron. Soc.* **328** (2001) 726 [[astro-ph/0012055](https://arxiv.org/abs/astro-ph/0012055)] [[INSPIRE](#)].

[64] V. Springel et al., *Simulating the joint evolution of quasars, galaxies and their large-scale distribution*, *Nature* **435** (2005) 629 [[astro-ph/0504097](https://arxiv.org/abs/astro-ph/0504097)] [[INSPIRE](#)].

[65] C. Srisawat et al., *Sussing Merger Trees: The Merger Trees Comparison Project*, *Mon. Not. Roy. Astron. Soc.* **436** (2013) 150 [[arXiv:1307.3577](https://arxiv.org/abs/1307.3577)] [[INSPIRE](#)].

- [66] A. Izard, M. Crocce and P. Fosalba, *ICE-COLA: Towards fast and accurate synthetic galaxy catalogues optimizing a quasi N-body method*, *Mon. Not. Roy. Astron. Soc.* **459** (2016) 2327 [[arXiv:1509.04685](https://arxiv.org/abs/1509.04685)] [[INSPIRE](#)].
- [67] B. Dai, Y. Feng, U. Seljak and S. Singh, *High mass and halo resolution from fast low resolution simulations*, *JCAP* **04** (2020) 002 [[arXiv:1908.05276](https://arxiv.org/abs/1908.05276)] [[INSPIRE](#)].
- [68] Y. Wu, H. Guo and V. Springel, *Improving the accuracy of halo mass based statistics for fast approximate N-body simulations*, *Mon. Not. Roy. Astron. Soc.* **531** (2024) 4944.
- [69] C.M. Baugh, *A primer on hierarchical galaxy formation: The semi-analytical approach*, *Rept. Prog. Phys.* **69** (2006) 3101 [[astro-ph/0610031](https://arxiv.org/abs/astro-ph/0610031)] [[INSPIRE](#)].
- [70] R.S. Somerville and R. Davé, *Physical Models of Galaxy Formation in a Cosmological Framework*, *Ann. Rev. Astron. Astrophys.* **53** (2015) 51 [[arXiv:1412.2712](https://arxiv.org/abs/1412.2712)] [[INSPIRE](#)].
- [71] M. Barrera et al., *The MillenniumTNG Project: semi-analytic galaxy formation models on the past lightcone*, *Mon. Not. Roy. Astron. Soc.* **525** (2023) 6312 [[arXiv:2210.10419](https://arxiv.org/abs/2210.10419)] [[INSPIRE](#)].