

Superresolution emulation of large cosmological fields with a 3D conditional diffusion model

Adam Rouhiainen,^{*} Moritz Münchmeyer, and Gary Shiu[✉]

Department of Physics, University of Wisconsin-Madison, Madison, Wisconsin 53706, USA

Michael Gira and Kangwook Lee[✉]

Department of Electrical and Computer Engineering, University of Wisconsin-Madison, Madison, Wisconsin 53706, USA



(Received 28 November 2023; accepted 23 May 2024; published 26 June 2024)

High-resolution (HR) simulations in cosmology, in particular when including baryons, can take millions of CPU hours. On the other hand, low-resolution (LR) dark matter simulations of the same cosmological volume use minimal computing resources. We develop a denoising diffusion superresolution emulator for large cosmological simulation volumes. Our approach is based on the image-to-image Palette diffusion model, which we modify to 3 dimensions. Our superresolution emulator is trained to perform outpainting, and can thus upgrade very large cosmological volumes from LR to HR using an iterative outpainting procedure. As an application, we generate a simulation box with 8 times the volume of the Illustris TNG300 training data, constructed with over 9000 outpainting iterations, and quantify its accuracy using various summary statistics.

DOI: [10.1103/PhysRevD.109.123536](https://doi.org/10.1103/PhysRevD.109.123536)

I. INTRODUCTION

Cosmological simulations of dark matter and baryonic matter are crucial for modern cosmology. They are required for theoretical studies, statistical method development, and parameter inference on real data. Unfortunately, high-resolution (HR) simulations with baryonic matter, even on moderate cosmological volumes, can require millions of CPU hours. On the other hand, low-resolution (LR) simulations of dark matter can easily be generated using much smaller computational resources. Our approach is to train a 3-dimensional diffusion model to upgrade LR simulations to superresolution (SR) simulations, which emulate the HR simulations at our targeted resolution. Because the small scales of the HR simulation are not contained in the LR simulation, an SR emulator should be probabilistic and able to generate many SR simulations consistent with the same LR simulation. We thus require a stochastic generative model, which learns probabilistic small-scale physics from HR training data. In this work, we choose a denoising diffusion probabilistic model [1–3], because they are currently the best performing models for image generation, generally outperforming GANs on various image tasks [4], and are much more expressive than normalizing flows in higher dimensions [5,6].

This work develops a stochastic treatment of superresolution emulation combined with a conditional outpainting

procedure that can make large volumes without boundary effects. Our outpainting method generates new SR volumes conditioned both on the LR simulation and on previously generated neighboring SR volumes, inspired by auto-regressive models such as PixelCNN [7]. By repeatedly outpainting smaller, local volumes many times over, we are able to generate simulation volumes larger than the entire HR training data volume. This is possible due to the locality of the underlying physics of structure formation, where on large scales, LR and HR simulations with the same initial conditions give the same final result. Therefore, we are able to emulate features with length scales larger than the outpainting window so long as such features are in the LR conditional. Details of our approach and its alternatives are discussed in Sec. IV A.

There have been various previous approaches to superresolution of cosmological data. In [8], the authors trained a Wasserstein GAN as a superresolution emulator at field level, in a similar setup to our work. However, their approach did not probabilistically model the SR conditioned on the LR simulation. Instead, they generate a deterministic SR solution, conditioned on the HR initial conditions field. This is an interesting and efficient approach but does not allow to explore the probabilistic space of consistent SR simulations given an LR simulation. In a different series of works [9–11], a GAN is used to generate SR simulations stochastically at particle level, increasing the particle count in the simulation boxes by several orders of magnitude. This is a powerful alternative

^{*}rouhiainen@wisc.edu

approach that is closer to N-body simulations, which naturally work with particles. We believe that both particle and field-level approaches can be useful. Cosmological survey analysis is usually done on a regular grid (for example, to take FFTs), so that particle positions are not necessarily required. On the other hand, the Lagrangian approach has the advantage that it can adaptively represent a physical field with a large dynamic range. Recently, [12] used a diffusion model to perform superresolution on 2-dimensional dark matter fields. Their work is technically most similar to ours, with their approach including a Fourier filter on the large-scale structure data to boost the importance of learning the nonlinear scales. Our work is volumetric and also includes an auto-regressive procedure to generate large volumes by conditioning sub-volumes on their neighbors.

Diffusion models have been used in astrophysics for other purposes than superresolution emulation. They have been used to generate astrophysical fields projected to 2 dimensions [13,14], and reconstruct strong gravitational lensing images [15]. Further, a 3-dimensional diffusion model was recently used to infer initial conditions from present day dark matter simulations, using data meshed onto 128^3 voxels [16].

This paper is organized as follows. We briefly review the theory behind denoising diffusion models in Sec. II. In Sec. III, we describe the HR and LR training data, and LR test set data, used in this project. In Sec. IV, we describe the diffusion model architecture with our iterative outpainting method. In Sec. V, we show our results for a generated field larger than the entire training data volume. In Sec. VI, we use the stochastic nature of diffusion models to generate a variety of SR fields conditional on a single LR field. In Appendix A, we demonstrate SR emulation of baryon acoustic oscillations, which have length scales larger than our model's outpainting window. In Appendix B, we test the robustness of the diffusion mode to slightly out-of-distribution data by varying the cosmological parameters of the LR field, and generating new SR fields.

II. DIFFUSION MODEL THEORETICAL BACKGROUND

The image-to-image denoising diffusion probabilistic model is a stochastic generative model that samples \mathbf{y} from $p(\mathbf{y}|\mathbf{x})$, where the image \mathbf{y} is conditional on a (usually low-resolution) image \mathbf{x} . Applications of image-to-image diffusion include colorization, denoising, JPEG resoration, inpainting, outpainting, and other types image reconstruction and superresolution. Here we briefly explain the theory behind this conditional diffusion model, with more details available in [2]. The equations here are valid in any number of spatial dimensions.

The diffusion model has a forward process for training, and a reverse process for generation. In the forward diffusion step, a field \mathbf{y}_t gets added noise from a field \mathbf{y}_{t-1} as

$$p(\mathbf{y}_t|\mathbf{y}_{t-1}) = \mathcal{N}(\mathbf{y}_t|\sqrt{1-\beta_t}\mathbf{y}_{t-1},\beta_t\mathbb{I}), \quad (1)$$

where β_t control the amount of noise reduction between steps. By iterating this forward diffusion process, the field \mathbf{y}_t at some step t can be written in terms of the original field \mathbf{y}_0 as

$$p(\mathbf{y}_t|\mathbf{y}_0) = \prod_{i=1}^t p(\mathbf{y}_i|\mathbf{y}_{i-1}) \quad (2)$$

$$= \mathcal{N}(\mathbf{y}_t|\sqrt{\gamma_t}\mathbf{y}_0, (1-\gamma_t)\mathbb{I}) \quad (3)$$

where $\gamma_t = \prod_{i=1}^t (1-\beta_i)$. A total of T steps are taken. The γ_t are chosen as model hyperparameters in such a way that by the end of the diffusion chain, $\sqrt{\gamma_T}\mathbf{y}_0$ is small relative to \mathbf{y}_T , and thus the noisy data has lost almost all resemblance to \mathbf{y}_0 .

The purpose of diffusion models is to generate fields in the distribution \mathbf{y}_0 without knowing the true value of \mathbf{y}_0 , and so we need an equation for \mathbf{y}_{t-1} in terms of \mathbf{y}_t without \mathbf{y}_0 . To achieve this, we fit a function $f_{\theta}(\mathbf{x}, \mathbf{y}_t, \gamma_t)$ to the noise $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbb{I})$ by minimizing the loss function

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y})} \mathbb{E}_{\epsilon, \gamma} \|f_{\theta}(\mathbf{x}, \sqrt{\gamma}\mathbf{y}_0 + \sqrt{1-\gamma}\epsilon, \gamma) - \epsilon\|_2^2 \quad (4)$$

where $\|\dots\|_2^2$ is the squared L_2 norm. In practice, the function $f_{\theta}(\mathbf{x}, \mathbf{y}_t, \gamma_t)$ is a neural network with parameters θ , and the details of our network are described in Sec. IV B.

There are multiple ways to solve for the reverse diffusion step [17,18], and we have a brief discussion of these various approaches in Sec. VII. In this work, we use the results of [2,3] to write the reverse diffusion step as

$$\mathbf{y}_{t-1} = \frac{1}{\sqrt{1-\beta_t}} \left(\mathbf{y}_t - \frac{\beta_t}{\sqrt{1-\gamma_t}} f_{\theta}(\mathbf{x}, \mathbf{y}_t, \gamma_t) \right) + \sqrt{\beta_t} \epsilon_t. \quad (5)$$

This reverse diffusion step is repeated to obtain a sample \mathbf{y}_0 from \mathbf{y}_T . Because the starting field \mathbf{y}_T is a random Gaussian field, we can obtain a variety of samples \mathbf{y}_0 without changing the network parameters θ .

III. DATA

Here we describe the HR TNG300 training data and its accompanying LR conditional training data. We also describe our the LR conditional test data, which has different initial conditions (phases), and a larger volume for Sec. V.

A. High-resolution data from TNG300

The HR data to learn the baryonic physics of the large-scale structure comes from IllustrisTNG [19–23], a set of three gravo-magnetohydrodynamical simulations run on AREPO [24]. We picked TNG300 because it covers a large enough volume to include both linear and nonlinear physics, and to demonstrate that a single HR simulation can be enough to train a model that generates larger volumes. The TNG300 run simulated 2500^3 baryon particles and 2500^3

dark matter particles in a $205 \text{ Mpc}/h \approx 300 \text{ Mpc}$ length cube, with cosmological parameters $\Omega_\Lambda = 0.6911$, $\Omega_M = 0.3089$, $\sigma_8 = 0.8159$, $n_s = 0.9667$, and $h = 0.6774$. The simulation has an initial field at redshift of $z = 127$ generated with NGenIC, using the Zel'dovich approximation. TNG300 computed over 10 million time steps down to $z = 0$, and we use the $z = 0.01$ snapshot. We use the cloud-in-cell mass assignment scheme to place the TNG300 baryons onto a 264^3 px cubic mesh. We will comment on increasing this resolution in Sec. VII. Our target field in the present work is the gas density (Particle type 0 in TNG). We chose this field because of our interest in applying the method to kinetic Sunyaev-Zeldovich science in the future (see Sec. VII), but for the purpose of method development the precise target of the superresolution task is not important. Although we only use one channel of data in this work, the gas particle density (baryons without stars and blackholes), we can in principle include more channels in our HR data to learn the temperature, electron density, etc. Generating more than one channel, possibly with an increased pixel resolution, is left to future work, with a discussion also in Sec. VII.

B. Low-resolution conditional data

Our goal is to quickly construct new HR matter fields from LR dark matter fields as a conditional for our diffusion model. With `Arepo`, we simulate dark matter fields to create the LR conditional training and LR test data. Our conditional fields are LR compared to TNG300 in three ways: they are dark matter only simulations with no baryonic physics, they contain far fewer particles, and they run across fewer time steps.

To simulate the LR training data, we use the same NGenIC initial seed, box size, and cosmological parameters as TNG300 described above, but now with 128^3 dark matter particles, computing about ~ 2000 time steps from $z = 127$ to $z = 0.01$. This simulation runs in about 10 CPU hours.

The HR-LR training pairs are constructed by randomly cropping 16000 cubes of size 48^3 px ($17.1 \text{ Mpc}/h$)³ out of the full 264^3 px fields, where 18% of the volume has been set aside as a validation set. Respecting the rotational symmetry of the superresolution operation, we also give each pair of cubes a random $\pi/2$ rotation and random chance of being mirrored. While the amount of training data is somewhat limited, we do not find evidence of over-training or memory of phases in our tests of sample diversity in Sec. VI. We also checked that our model did not perform better on the LR training data than on LR test data with various statistics.

For the LR test data, we run this simulation again, but now with a different NGenIC initial seed. We make two sets of LR test data. For the main result of this paper, presented in Sec. V, the LR field is a $410 \text{ Mpc}/h$ length box, having 8 times the volume of TNG300; therefore, we proportionally increase to 256^3 dark matter particles as to have the same particle density as the LR training data. For Sec. VI and Appendix B, we simulate LR fields at $205 \text{ Mpc}/h$. There is no HR truth for our LR test data, and so we will use summary statistics from the TNG300 training data as a truth comparison.

IV. CONDITIONAL DIFFUSION MODEL APPROACH

Our approach is based on the image-to-image Palette diffusion model [2,3], which is capable of performing outpainting to generate larger images. We modify the model to 3 dimensions and train it to be conditional on our 3-dimensional LR simulation.

A. Iterative outpainting to generate large fields

This work uses massively iterative outpainting that can generate cosmological volumes much larger than the training data. Given a large LR field, an SR field is

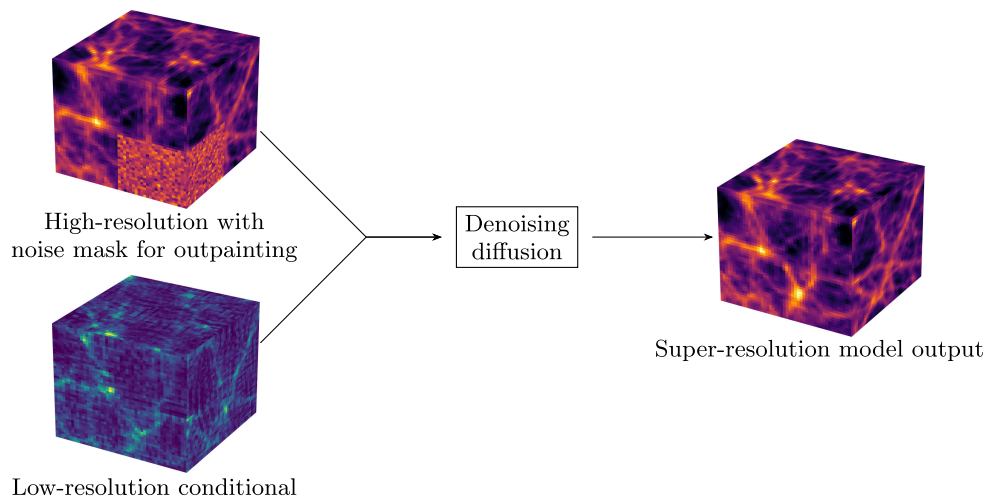


FIG. 1. We train a diffusion model to do outpainting on cosmological fields. By iterating this outpainting many times over, we are able to construct large SR fields given an LR conditional to guide the large scale modes.

generated sequentially, patch by patch, with each new patch generated conditional on both the LR field and adjacent SR fields. The LR conditional guides the large length scales of the SR field, while the adjacent SR conditional ensures that newly generated SR fields are smooth and physically consistent with the larger SR volume. Training the model to outpaint requires masking subvolumes of the HR data in the HR-LR pairs, as shown in Fig. 1, and thus masked subvolumes at inference can be generated given surrounding SR data. An illustration of the iterative outpainting is shown in Fig. 2.

Our iterative outpainting method is motivated by the nature of mode coupling in the large-scale structure. On large scales, for $k \lesssim 0.1 h/\text{Mpc}$ today, the evolution of physical perturbations is linear, as modes evolve independently. On intermediate scales, $0.1 h/\text{Mpc} \lesssim k \lesssim 0.5 h/\text{Mpc}$ modes evolve nonlinearly but should be accurately captured by the LR dark matter N-body simulation. On smaller scales $k \gtrsim 0.5 h/\text{Mpc}$ we want the diffusion model to model nonlinearity and baryonic feedback. This informs us about the minimum physical size required for the outpainting volumes, and the diffusion model we describe here has a fundamental mode of $k_{24\text{px}} = 0.34 h/\text{Mpc}$. The diffusion model can modify the results of the LR simulation on physical scales smaller than this scale, and can thus take into account mode coupling on these scales. Mode coupling is included by the LR simulation on large scales, but cannot be modified by the diffusion model due to its outpainting window size, and we thus assume that the LR simulation is correct on these scales. In position space, the 24 px window corresponds to a physical length of $18.6 \text{ Mpc}/h$. This can be compared to the typical particle displacement of $5 \text{ Mpc}/h$, with an upper limit of $\sim 20 \text{ Mpc}/h$ [8]. Most of this displacement is already included in the LR simulation, so our window size is sufficient.

It would also be possible to make large volumes without conditional patching by a different approach.¹ Since the diffusion model is convolution-based and therefore translationally invariant, by sharing the noise of adjacent volumes on the boundary, one could avoid boundary effects by construction. This is an elegant alternative that we will consider in future work. Our own approach is based on the autoregressive way of generating data, used very successfully e.g. in PixelCNN [7] and large language models [25], which generate output sequentially. The learning task in our approach is not the same as in the alternative approach, and both sample quality and sample diversity may be different, especially with limited training data. Our approach also has the benefit that it is possible to outpaint from high-resolution data which was not generated by a diffusion model. This would be useful for example if one

¹We thank the anonymous reviewer for suggesting this approach.




FIG. 2. Iterative outpainting illustration, 2-dimensional projection. The blue cube contains the previously generated SR fields adjacent to the subsequent volume ready to be outpainted. After a new volume is generated, the blue cube then moves in row-major order for the next SR volume to be generated. After a plane of small volumes is generated to SR, the outpaintings continue in the third dimension (perpendicular to the page), and the next plane is generated, with each SR volume also conditional on the previous plane.

has a high-resolution observation of smaller parts of the sky, but only a low-resolution observation of large scales.

In detail, the iterative outpainting procedure works as follows. First, a 48^3 px SR cube is generated conditional on a 48^3 px LR cube; this first cube is shown in the top left of Fig. 2. We move in row-major order, outpainting 24 px at a time. The second SR volume generated is thus conditional on both its underlying 48^3 px LR cube, as well as the 24 px length right half of the first SR cube. After an entire plane is generated, the outpaintings move to the third dimension, with every subsequent plane conditional on the previously generated plane. The outpaintings continue in this way until the entire LR volume is generated to SR. We never break conditionality on previous adjacent SR regions, even after moving into the third dimension. Thus throughout the full SR volume, every locally outpainted volume is conditional on all adjacent previously generated volumes.

In some small regions at the outpainting boundaries, a slight discontinuity develops in the SR model output. To remedy this, we apply a linear interpolation in the 2 px wide strip at the outpainting boundaries. This interpolation has negligible affect on the summary statistics, while making the results visually appear slightly more accurate.

A potential issue we have found is the generated data progressively slipping out of the distribution that the diffusion model was trained on. Suppose that at some early point i in the outpainting chain, the diffusion model generates an out-of-distribution $\mathbf{y}_i + \delta\mathbf{y}_i$ as a sample from $p(\mathbf{y}_i|\mathbf{x}, \mathbf{y}_{<i})$. The next sample would be generated from $p(\mathbf{y}_{i+1}|\mathbf{x}, \mathbf{y}_{<i}, \mathbf{y}_i + \delta\mathbf{y}_i)$, and because the model has not trained with any $\mathbf{y}_i + \delta\mathbf{y}_i$, the out-of-distribution problem may compound. This problem arose for training the model using apparently too few model parameters. In this case, the first SR cube in the outpainting chain was too Gaussian, albeit still resembling the conditional field. However, after

several more outpainting iterations, the model output quality collapsed, and appeared as very smooth blob. We overcame this problem by sufficiently increasing the number of model parameters (described below) as to never slip out of the model’s learned distribution.

B. Model details

Our model is a modified version of Palette based on code from [26]. We use a U-net [27] to learn the function $f_{\theta}(\mathbf{x}, \mathbf{y}_t, \gamma_t)$ that handles the denoising routine from \mathbf{y}_t to \mathbf{y}_{t-1} . The details of the U-net used here are based on the “guided diffusion” U-net described in [4], which we minimally modify to operate on 3-dimensional data. A diagram of the U-net is shown in Appendix C. The U-net takes as input 2 channels, \mathbf{y}_t and the conditional \mathbf{x} , and outputs 1 channel, a prediction for ϵ . It has two downsampling steps, with the number of channels being 64, 128, and 256, and similarly two upsampling steps. BigGAN residual blocks [28] are used for the downsampling and upsampling, alternating between standard residual blocks at each resolution. The middle section of our U-net has 2 residual blocks and a query-key-value attention layer [29]. We also experimented with the SR3 U-net introduced in [30], being the network used in the original Palette implementation, and we found similar model performance between these two U-nets.

Our residual blocks contain two $3 \times 3 \times 3$ convolutions, with embeddings to be dependent on the diffusion step number t . Prior to these convolutions is a group normalization [31] with sigmoid linear unit (SiLU) [32] activation

function. We also experimented with ReLU activation functions, with which we found similar model performance. We use 0.2 dropout before the second convolution in each residual block. There is additionally a $1 \times 1 \times 1$ convolution in residual skip connections when the number of channels is changed in the residual block. This model has 31.5 million learnable parameters.

C. Model training

We train the diffusion model on the LR-HR pairs of 48 px length fields described above. Due to the nature of matter clustering in cosmology, the distribution of values in both the LR and HR data is heavily skewed to the high-density tail, and so we preprocess the data by applying a natural logarithm. We want the diffusion model to take as input data values in the range $(-1, 1)$, so we additionally apply a sigmoid to the data. We should not allow the diffusion model to train on data with values well outside of $(-1, 1)$, for then the \mathbf{y}_T fields would not accurately be noisy but rather be spatially correlated with high density values in the truth \mathbf{y}_0 .

After each batch of HR-LR pairs is loaded as input to the model, each HR field is randomly assigned a mask, chosen from the 8 mask types shown in Fig. 3. We train with $T = 2000$ diffusion time steps with noise variances linearly increasing from $\beta_0 = 10^{-6}$ to $\beta_T = 10^{-2}$. We use the Adam optimizer with a learning rate of 10^{-4} , and a 0.9999 exponential moving average learning rate decay. We train on two A100 GPUs with a batch size of 16 per GPU,

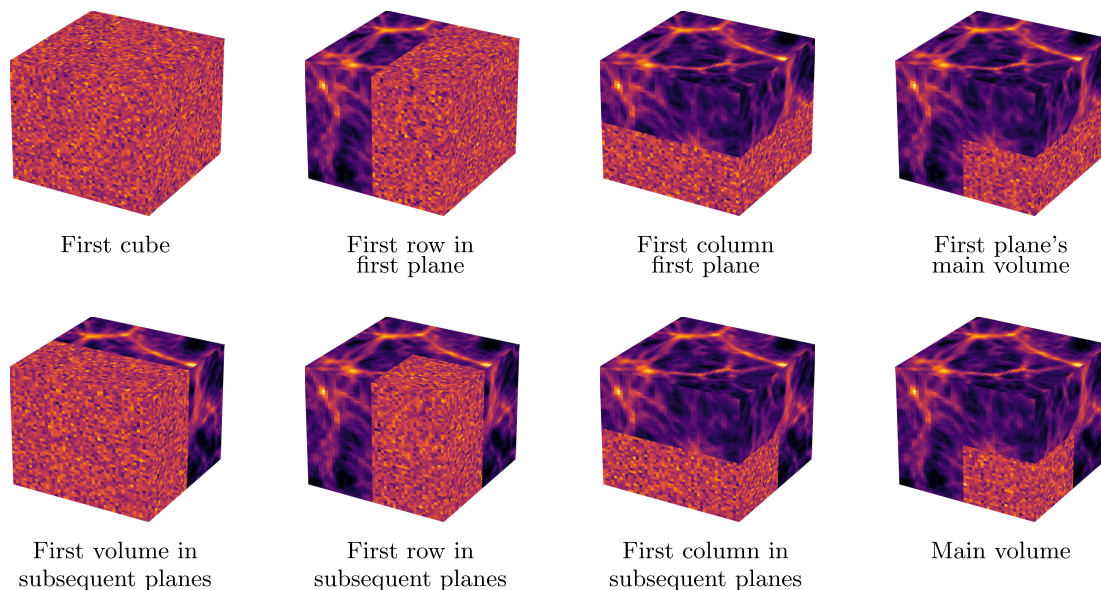


FIG. 3. There are 8 mask types (4 considering rotational symmetry) required to iteratively outpaint the full 3-dimensional volume. Each mask type is labeled as to where in the full volume it is used for outpainting. The top row of masks are used to outpaint the full first plane of cubes. The three “subsequent planes” masks on the bottom row are for outpainting the first row and column of every plane of cubes after the first plane. The vast majority of the outpaintings belong to the “main volume” mask (bottom right) containing noise in a single octant, used for volumes that are not located in the first plane of cubes, nor the first row or first column of any planes.

and the training loss Eq. (4) converges in 70 hours. The validation loss shows no sign of overtraining.

V. SUPERRESOLUTION RESULTS FOR A LARGE VOLUME

The main result of this work is the SR emulation of a 410 Mpc/h \approx 600 Mpc length cube, with 528^3 px, constructed with 21^3 outpainting iterations. As the training data came from the 205 Mpc length TNG300 cube meshed onto 264^3 px, our SR result thus has an increased volume of the entire training data by a factor of 8. We show our results in Fig. 4.

To generate our large SR cube, we used a linear noise schedule of $T = 1250$ steps, with $\beta_0 = 10^{-6}$ and $\beta_T = 1.5 \times 10^{-2}$. We found 1250 steps to be the lowest possible before losing any measurable amount of quality in the power spectrum; a discussion of potential further speed-ups is in Sec. VII. The time to generate the full cube of 21^3 outpainting volumes was 120 hours on a single A100 GPU.

We measure the accuracy of the diffusion model with several summary statistics familiar to cosmology. Results are computed on mean zero overdensities $\delta(\mathbf{r})$, its Fourier transform denoted by $\delta(\mathbf{k})$.

A. One-point probability distribution

We plot the one-point probability density function (PDF) in Fig. 5 (top left). The LR test data has many empty voxels, as it was created with 256^3 particles in a 528^3 px mesh. The diffusion model accurately generates SR data with the correct HR PDF, which is a smooth skew-Gaussian curve.

B. Power spectrum

The power spectrum $P(k)$ is a two-point correlation function in Fourier space, defined by

$$(2\pi)^3 P(k) \delta_{\text{D}}(\mathbf{k} + \mathbf{k}') = \langle \delta(\mathbf{k}) \delta(\mathbf{k}') \rangle. \quad (6)$$

Here δ_{D} is the Dirac delta function, and $\langle \dots \rangle$ is an ensemble average.

The power spectra comparing the LR, SR, and HR fields are plotted in Fig. 5 (top right). As the LR and SR fields have twice the box length of the HR field, they have additional small k Fourier modes around the $k \sim 0.2$ h/Mpc turning point. We see that below the outpainting mode of $k_{24\text{px}} = 0.34$ h/Mpc, the SR power spectrum is guided by the LR power spectrum, properly emulating the baryon acoustic oscillations (BAOs) in the LR conditional. We further investigate BAOs in Appendix A. Additionally, the SR field correctly begins to turn over for small k along with the LR conditional. At high k , the SR simulation correctly follows the HR power

spectrum, which is more suppressed for baryons than for dark matter, as physically expected. Our LR simulation dark matter power spectrum is not physically accurate at high k due to its low particle resolution, but the diffusion model nevertheless learns the correct baryonic power spectrum.

A possibly surprising fact is that in the smallest k power spectrum bin the SR power spectrum is somewhat suppressed with respect to the LR. In principle, the LR conditioner should dominate the power spectrum here. However, it appears that the mean density of the sequentially generated 24 px length SR volumes can slightly drift, introducing a modulation on the very largest scales. We expect that this problem could be resolved with a multi-scale conditioner, which will be discussed in Sec. VII. A simpler solution would be to Fourier filter the SR map and take the largest scales directly from the LR simulation.

C. Bispectrum

The power spectrum measures only the Gaussian structure, so we must also calculate the next higher moment to understand the non-Gaussian structure in our fields. The bispectrum $B(k_1, k_2, k_3)$ is a three-point correlation function defined by

$$(2\pi)^3 B(k_1, k_2, k_3) \delta_{\text{D}}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) = \langle \delta(\mathbf{k}_1) \delta(\mathbf{k}_2) \delta(\mathbf{k}_3) \rangle. \quad (7)$$


The bispectrum is a function of three variables, so for easier analysis, we show a lower dimensional projection of the bispectrum. We pick the first two vectors (k_1, k_2) to have a constant magnitude, and plot the bispectrum as a function of the angle θ between these two vectors.

The projected bispectrum is plotted in Fig. 5 (center row) for the two choices $(k_1, k_2) = (0.15$ h/Mpc, 0.25 h/Mpc) and $(k_1, k_2) = (0.6$ h/Mpc, 1.0 h/Mpc). These two pairs of (k_1, k_2) probe Fourier modes both below and above the $k_{24\text{px}}$ outpainting scale. The SR field has the same general bispectrum shape as HR for each of our two (k_1, k_2) choices, with higher accuracy for the modes above $k_{24\text{px}}$. It is likely that the slight mismatch of the bispectrum on scales around $k_{24\text{px}}$ is an imperfection of our trained outpainting, while at higher k the diffusion model perfectly accounts for mode-coupling. We will comment on potential improvements of the outpainting in Sec. VII. It could also suffice to increase the physical size of the outpainting window, which comes however at additional computational cost.

D. Void size function

Voids are an important feature of the highly non-Gaussian large-scale structure at late times. The SR field

Training data



Results




FIG. 4. Matter density results for our superresolution diffusion model in generating a volume larger than the entire training data volume. Images are 2-dimensional projections of depth 19 Mpc/h. Top row: the training data comes from the single pair of boxes shown. The model trains on 48 px length LR-HR pairs cut out of these boxes. Center left: LR conditional test data. Center right: SR model output generated with 21^3 outpainting iterations, having 8 times the volume of the entire training data. Bottom row: two zoom-ins of the LR and SR fields.

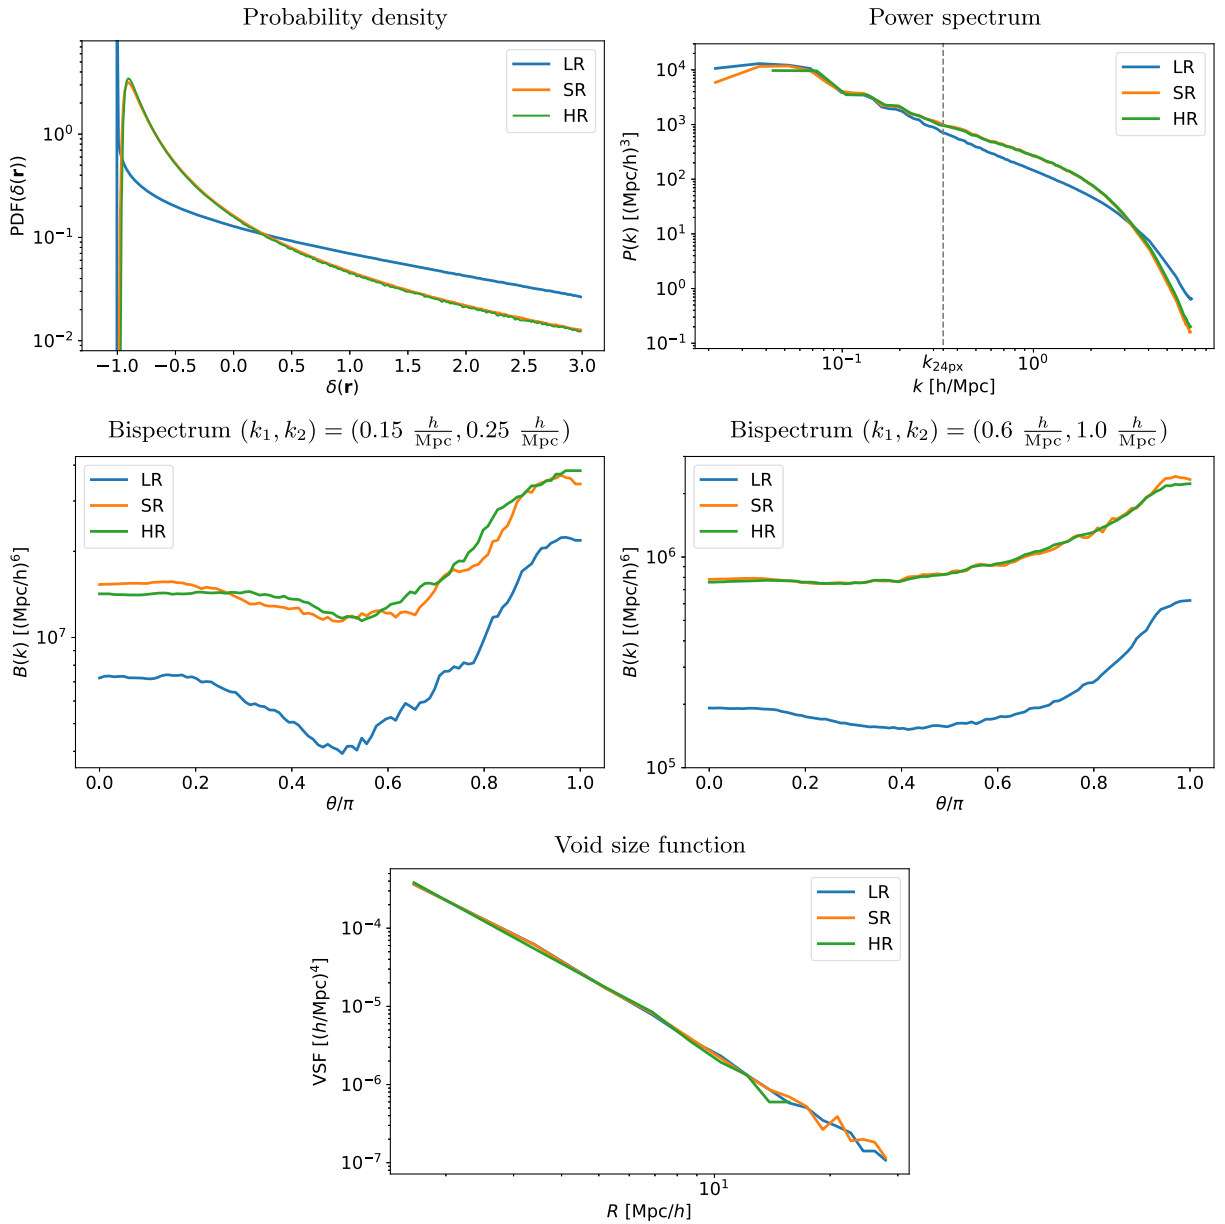


FIG. 5. Summary statistics comparing the 410 Mpc/h length LR test field and SR model output, and the 205 Mpc/h length HR training field as a truth comparison.

should retain the overall geometry of the LR conditioner, and so we expect the quantity and size of voids in the LR, SR, and HR fields to be similar. We calculate the void size function with PYLIANS [33], which uses the spherical overdensity void finder presented in [34]. The void finder smooths the field with a top-hat filter of radius R , and then finds underdense volumes of radius R below a given threshold.

The void size function is plotted in Fig. 5 (bottom) for a void threshold of 0.0 (compare to the PDF in Fig. 5). The HR void size function loses resolution at a radius of about $R = 15$ Mpc/h. The SR and LR void size functions match up to $R = 28$ Mpc/h, indicating that the SR field

respects the structure of the LR field for voids larger than can be accurately measured by the full TNG300 training volume.

VI. MODEL OUTPUT VARIETY FROM A SINGLE LOW-RESOLUTION FIELD

Due to the stochastic nature of denoising diffusion models, our SR emulator is able to sample from the manifold of SR solutions $p(y|x)$. In this section, we measure the variety of SR samples created from a single LR conditional field. Our method here may be compared to the detailed analysis in [12], where they generate multiple




FIG. 6. Visual results for a variety of SR fields from a single LR conditional field. Top row: LR and multiple SR realizations, 2-dimensional projects. Bottom row: High density features of the above SR fields in 3 dimensions. To better visually discriminate differences between fields, we plot two different SR fields in the same volume.

SR images conditional on a single LR image, as well as multiple SR images from multiple LR images.

We generate 25 SR fields of size 144^3 px (physical length $112 \text{ Mpc}/h$), each requiring 5^3 outpaintings, conditional on the same LR field. We use the same 1250 step noise schedule as in the previous section. Visual results are shown in Fig. 6 for the LR volume, and four samples of the SR results. It is difficult to discern differences between the SR samples side-by-side, and so we show also two SR fields in the same 3-dimensional plot. These point plots are constructed with the large density features of two SR fields, one field in red and the other in green. When plotted on top of each other, it becomes more apparent that each SR field has a different phase in their placement and shape of halos and filaments.

Summary statistics for the LR and 25 SR samples with 1σ variation between SR samples are shown in Fig. 7, with the HR training data as a truth comparison. The results are highly accurate between SR and HR for the PDF (top left), power spectrum (top right), and bispectrum with $(k_1, k_2) = (0.6 \text{ h}/\text{Mpc}, 1.0 \text{ h}/\text{Mpc})$ (bottom left). Both the power spectrum and the bispectrum in the small SR volumes show a sample variance due to the finite mode

number and the freedom of the diffusion model to modify nonlinear modes.

We further quantify the differences in the SR samples with their cross-correlations in Fourier space. For two fields $\delta_i(\mathbf{k})$ and $\delta_j(\mathbf{k})$, their cross-correlation coefficient is

$$r_{ij}(k) = \frac{P_{ij}(k)}{\sqrt{P_i(k)P_j(k)}}. \quad (8)$$

Here $P_{ij}(k)$ is the cross-power spectrum

$$(2\pi)^3 P_{ij}(k) \delta_D(\mathbf{k} + \mathbf{k}') = \langle \delta_i(\mathbf{k}) \delta_j(\mathbf{k}') \rangle. \quad (9)$$

We compute $r_{ij}(k)$ for all 300 different (i, j) pairs for our 25 SR fields, and plot the result in Fig. 7 (bottom right). The cross-correlation coefficient is close to 1 below the $k_{24\text{px}}$ outpainting scale as expected, as small k modes are guided by the single LR field. At nonlinear scales, the cross-correlation coefficient drops to nearly 0, signifying little correlation between SR samples. Additionally, there is almost no variance in the largest k -bins where the $r_{ij}(k)$ is nearly 0, and therefore very few SR samples are correlated at the smallest length scales.

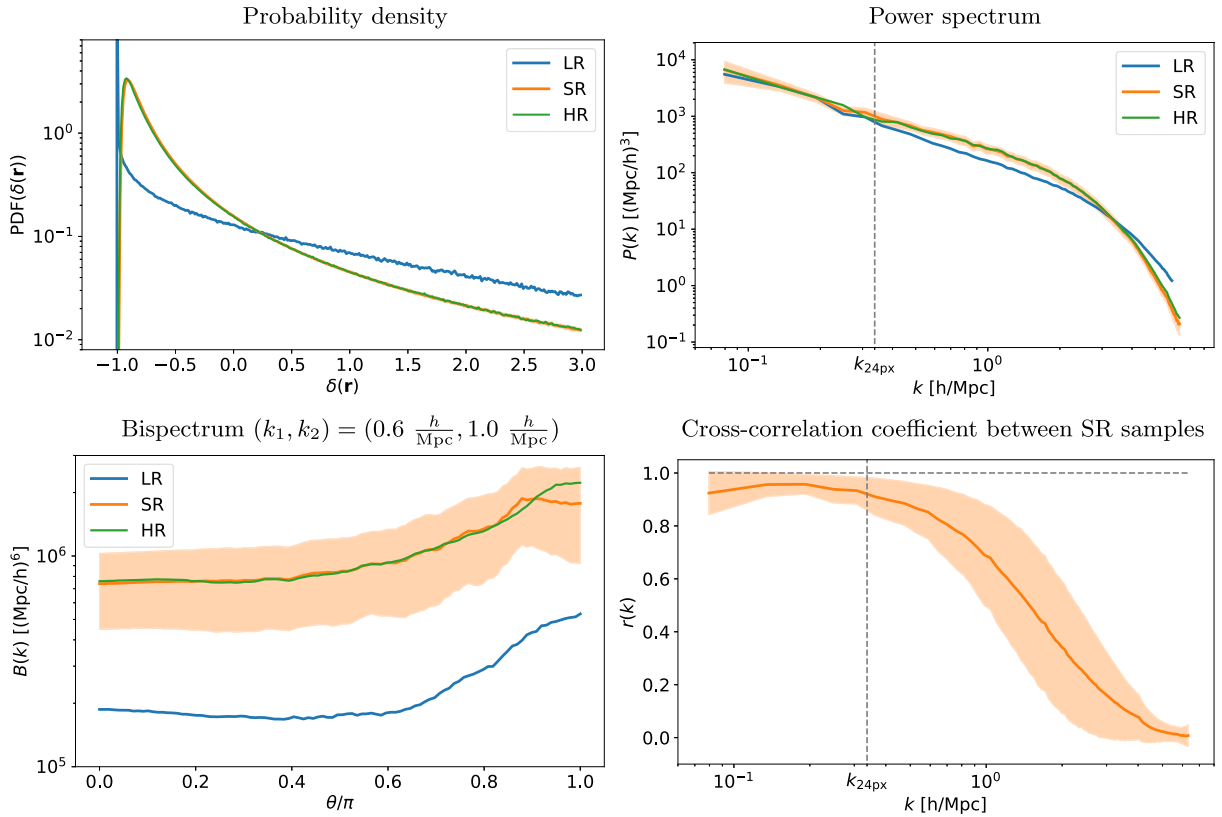


FIG. 7. PDF (top left), power spectrum (top right), and bispectrum (bottom left) for a variety of 25 SR fields generated from a single LR conditional field, with 1σ bands in the SR fields. The cross-correlation coefficient of Fourier modes (bottom right) is calculated between all pairs of the 25 SR fields.

VII. CONCLUSION

Super-resolution emulators are a promising tool to open the computational bottleneck of HR baryonic simulations in cosmology. In this work we evaluated the performance of a diffusion model on volumetric data, and developed a conditional outpainting scheme that can upgrade large LR volumes. We trained on data from the TNG300 simulation, and generated an SR field with 8 times the volume of the entire training data volume. Our diffusion model is capable of making accurate 3-dimensional SR emulation results, with the resulting SR volume matching the summary statistics of the training HR simulation closely. Our SR field is guided by the LR field at length scales larger than the outpainting scale. We also demonstrated the stochasticity of the diffusion model by generating a variety of SR fields conditional on a single LR field.

Leading up to this work, we found that training probabilistic generative models in 3 dimensions is not always successful, as we intended to train normalizing flows such as Real NVP [35] and Glow [36] on the same task. Normalizing flows have been recently used to learn the PDF of 2-dimensional projections of cosmological fields [37,38], and our 2-dimensional superresolution results were promising. However, the flows were not accurate in 3 spacial dimensions and we thus moved to the more

expressive diffusion models. On the other hand, diffusion models are far slower at inference than normalizing flows, and this makes it challenging to generate very large volumes or many volumes. Our 410 Mpc/h SR simulation took about 120 hours to generate with a single A100 GPU, and generation timescales linearly with the volume.

Further research with 3-dimensional diffusion models would be greatly aided with a faster denoising algorithm. While this work uses a conditional denoising diffusion probabilistic model (DDPM), recent developments in diffusion research include the significantly faster denoising diffusion implicit model (DDIM) [17] and denoising probabilistic models solver (DPM-Solver) [18]. The DDIM algorithm and DPM-Solver advertise a factor of 10 to 100 reduction in the number of denoising steps with a trade-off of a small loss in accuracy. We tested an implementation of DDIM in preparing this work, but our results were subpar compared to DDPM. Nevertheless, it is likely that with more work the sample generation time can be reduced by a significant factor.

It would be interesting to study different LR conditioners. Our LR conditioner is only 48 px in field of view, which could be increased, perhaps using a multiscale conditioner that includes up to the entire simulation. It is plausible that in this way the bispectrum at intermediate scales (see Fig. 5)

could be modeled even more precisely. This would also improve the modeling of SR squeezed limit N -point functions (beyond the response of the LR simulation). One may also wonder about the influence of the order of sample generation in the 3-dimensional volume. Our sample generation process breaks spatial homogeneity in principle, since we need to start generating fields somewhere in space. On the other hand, small-scale structure at some point in space is independent of small-scale structure far away from it, due to the locality of structure formation. An alternative approach which preserves spatial homogeneity exactly is discussed in Sec. IV A.

A physical application for our simulations is making simulations for $\text{CMB} \times \text{LSS}$ cross-correlation analyses [39]. An interesting case here is $k\text{SZ}$ tomography [40,41], which is sensitive to the cross-correlation of the galaxy density and the electron density $P_{ge}(k)$. As shown in [41], with SO and DESI, $P_{ge}(k)$ can be probed up to about $k \sim 8 \text{ Mpc}^{-1}$. This required resolution is near the Nyquist frequency of the pixelization used in this work. It would thus be important to scale up the resolution of our data for future studies, which TNG300 allows owing to its large 2500^3 particle count.

Cross-correlation analyses in general require several correlated fields, and the $k\text{SZ}$ tomography use case requires a two field output of electron density and galaxy density. Such fields are available in IllustrisTNG, and multiple fields could be generated by our diffusion model. The most apparent way to include multiple fields would be to add additional input and output channels to the diffusion model's U-net. While technically straightforward, this approach becomes computationally challenging, especially in 3 dimensions. Another possibility is to include a second machine learning model on top of the diffusion model output. For example, invertible mappings were recently developed [42] between different cosmological fields in the CAMELS simulations [43].

Our method could be applied to particle level super-resolution emulation with point cloud probabilistic diffusion models [44]. This would be closer to the superresolution emulator presented in the series of papers [9–11] or [12]. Point cloud diffusion models have been recently used to generate QCD jets with high precision [45,46]. For some analyses, a hybrid generator that can both make a continuous matter field and a point-cloud galaxy field would be useful.

Finally, it would be useful to be able to vary cosmological and astrophysical parameters. In Appendix B we explored transferring new cosmologies with different cosmological parameters from the LR conditional field to the SR field. Unfortunately, our diffusion model did not behave well with out-of-distribution LR fields. A more straightforward implementation of varied cosmological parameters requires multiple training simulations with different values of such parameters. For the IllustrisTNG simulations, this is not currently available, but the CAMELS project and its planned extensions can be useful. For our approach, we require

simulations that both include large linear and small nonlinear scales. In principle, such parameter dependence can be included in the diffusion model with techniques such as style transfer and Low-Rank Adaptation [47], the latter of which has been successful at concept tuning and concept fusion with text-to-image diffusion models [48]. With larger but still realistic GPU resources, it would be possible to generate a few large volume simulations with different astrophysical parameters for the same cosmology.

ACKNOWLEDGMENTS

We thank Ying Fan for helpful discussion. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics under Award Nos. DE-SC-0023719 (G. S.) and DE-SC-0017647 (G. S., M. M.). M. M. is supported by NSF Grant No. 2307109. Model training and inference used resources from the Data Science Institute at the University of Wisconsin-Madison. Support for this research was provided by the University of Wisconsin-Madison, Office of the Vice Chancellor for Research and Graduate Education with funding from the Wisconsin Alumni Research Foundation.

APPENDIX A: SR RESPONSE TO BAOS IN THE LR CONDITIONAL

Our diffusion model's SR output follows the basic structure of the LR power spectrum at large length scales, even though the SR field is constructed with many smaller conditional outputtings. We can demonstrate the response of our SR field to LR modes by examining BAOS. As we saw in Sec. V, the SR output was able to accurately emulate the BAOS around $k \sim 0.1 h/\text{Mpc}$, lower than our outputting scale of $k_{24\text{px}} = 0.34 h/\text{Mpc}$.

BAOs emulated from LR to SR

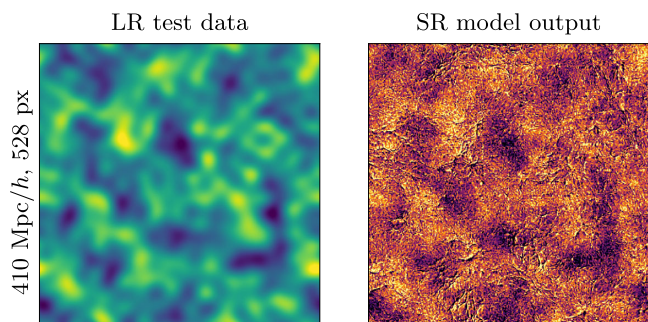


FIG. 8. Field-level visualization of BAOS in the LR and SR fields, $19 \text{ Mpc}/h$ depth 2-dimensional projections. We constructed two $410 \text{ Mpc}/h$ length LR boxes with the same initial seed, one with BAOS and one without, and we generated their respective SR fields. Shown are the LR and SR BAOS in position space, computed as the difference between fields with and without BAOS. Large-scale fluctuations are emulated from the LR conditional to the SR model output, even though the diffusion model only trained on $37 \text{ Mpc}/h$ length boxes.




FIG. 9. Power spectra for fields with and without BAOs, plotted as solid and dashed lines, respectively. When the LR conditional contains BAOs, they are emulated to SR, matching the HR training data. Conversely, an LR conditional without BAOs generates an SR field without BAOs. The SR model output emulates structure in the LR conditional at Fourier modes smaller than the $k_{24\text{px}}$ outpaiting scale.

We further illustrate this successful large-scale response of our model by comparing fields with and without BAOs. We constructed two 410 Mpc/h length LR boxes with the same initial seed, one with BAOs and one without. We have then generated their respective SR fields, and the one containing BAOs was presented in Sec V. We show field-level images of BAOs in Fig. 8, computed as the difference between the BAO and non-BAO fields. We see that large-scale fluctuations contained in the LR conditional field are emulated to the SR model output. (The SR fields additionally have high-frequency phase differences irrelevant to the BAOs, explained in Sec. VI.) In Fig. 9, we plot power spectra of our LR and SR fields comparing the existence of BAOs. The SR power spectra follow the same pattern as the LR power spectra below $k_{24\text{px}}$, whether there are BAOs or not. Large length structures are not learned from the HR field with our outpaiting method, and must accurately be represented in the LR field in order to be generated in the SR field.

APPENDIX B: VARYING COSMOLOGIES IN THE LR CONDITIONAL

In this Appendix, we generate new SR simulations with varied cosmological parameters in the LR conditional field as to examine the model’s robustness to out-of-distribution conditional data. We vary each of σ_8 and Ω_M by $\pm 5\%$ in the LR conditional from their fiducial values of $\sigma_8 = 0.8159$, $\Omega_M = 0.3089$, giving four experiments to test the sensitivity to out-of-distribution data. For our four varied cosmologies along with a fiducial cosmology, we simulate 205 Mpc/h length boxes of 128^3 dark matter particles to build the LR conditional, each with the same initial seed.

For each varied cosmology, we independently generate 8 SR fields, each of length 112 Mpc/h (144 px) consisting of 5^3 outpaiting iterations. The 8 fields are generated from

different volumes of the 205 Mpc/h conditional. By generating several SR fields, rather than a single large field, we get an estimate of the variance of our diffusion model output.

Here we use $T = 2000$ denoising steps (increased from 1250 previously) and use a noise schedule with $\beta_0 = 10^{-6}$, $\beta_T = 10^{-2}$. Increasing the number of steps while reducing the β_i between steps may possibly reduce errors in this out-of-distribution regime, but we did experiment with either noise schedule and obtained qualitatively similar results. We show power spectra ratios between the varied and fiducial SR fields for the four cosmologies in Fig. 10, along with the LR power spectra.

For σ_8^- , we expect the power spectrum to be decreased at nonlinear scales by a factor of about 0.95^2 , but our SR power spectrum only decreases slightly at these scales. There is a bit of success for σ_8^+ , as we can see that the model output has




FIG. 10. Power spectra for varying σ_8 and Ω_M by $\pm 5\%$ in the LR conditional field, shown as ratios over the fiducial power spectra with $\sigma_8 = 0.8159$, $\Omega_M = 0.3089$, for 8 independently generated volumes, the mean denoted by the bold line. In all four varied cosmologies, the SR field tends to have increased power at nonlinear scales. The model seems to be highly sensitive to out-of-distribution conditional fields.

nearly the proper 1.05^2 increase in power for $k < 1 \text{ h/Mpc}$. However, a major issue is that both σ_8^- and σ_8^+ significantly increase in power above $k = 1 \text{ h/Mpc}$, but we expect them to move in opposite directions. Considering Ω_M^- and Ω_M^+ , our SR model output has remained the same at nonlinear scales in the power spectrum. However, the model output again has a significant power increase above $k = 1 \text{ h/Mpc}$ for both Ω_M^- and Ω_M^+ .

In each of the four cosmologies, the SR model output did not accurately follow the expected power spectra curve compared to the fiducial cosmology, but rather the power at larger k tends to incorrectly increase significantly. We suspect this is a signature of pixel-wise noise not being properly denoised by the diffusion model. The diffusion model seems to be highly sensitive to out-of-distribution

conditional data and is not robust to emulating new cosmologies from the LR conditioner. To include cosmological parameter dependence, we will thus need training data that varies these parameters.

APPENDIX C: U-NET DIAGRAM

We show a diagram of the U-net used in this work in Fig. 11. This U-net is similar to that used in [4], built with a series of residual blocks, and downscaling and upscaling in the BigGAN residual blocks. To allow the U-net to be dependent on the diffusion step number t , each residual block has an embedding as an additional input. Each embedding goes through a linear layer and then gets added in the middle of the block, as shown.

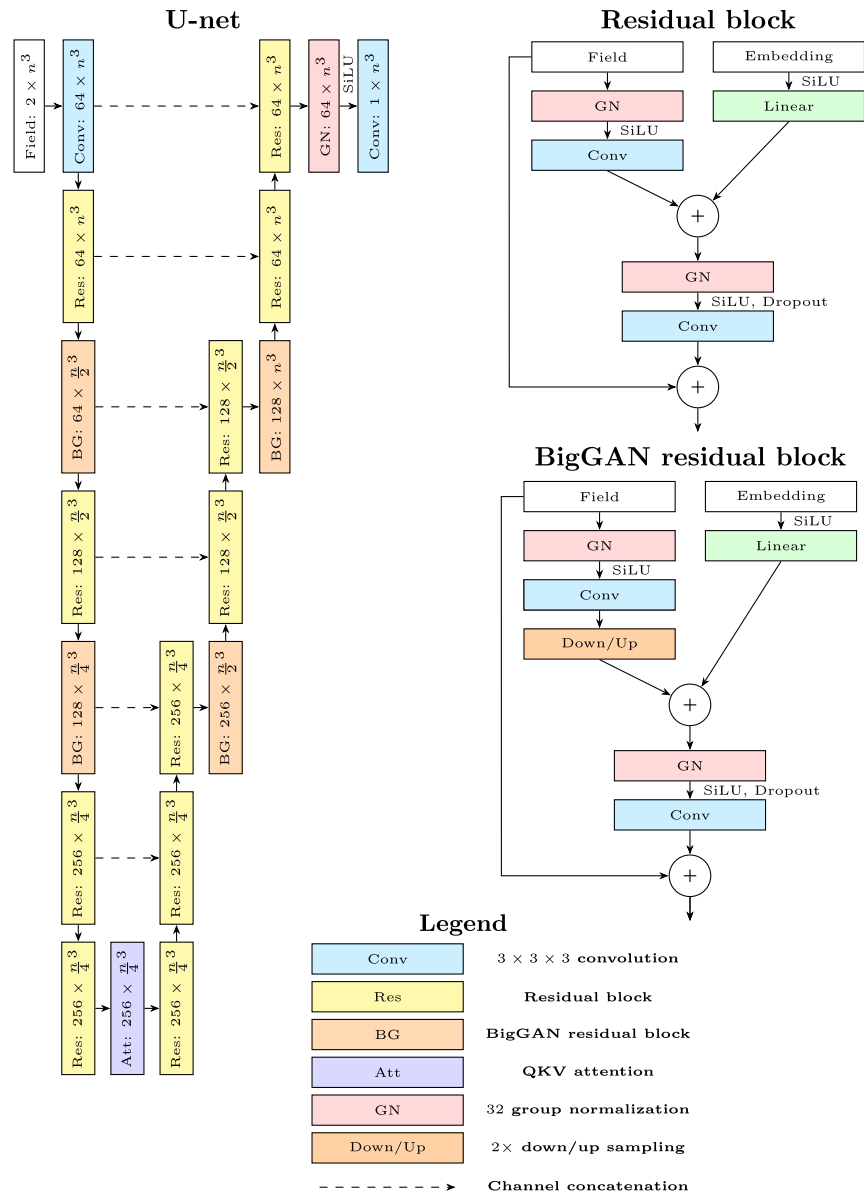


FIG. 11. The U-net used in this work (left) is made of residual blocks (top right), with BigGAN residual blocks (center right) used to downsample and upsample the fields. Shown after every layer name in the U-net is the layer’s output channels \times volume.

- [1] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, Deep unsupervised learning using nonequilibrium thermodynamics, in *Proceedings of the 32nd International Conference on Machine Learning*, Proceedings of Machine Learning Research Vol. 37, edited by F. Bach and D. Blei (PMLR, Lille, France, 2015), pp. 2256–2265.
- [2] J. Ho, A. Jain, and P. Abbeel, Denoising diffusion probabilistic models, [arXiv:2006.11239](https://arxiv.org/abs/2006.11239).
- [3] C. Saharia, W. Chan, H. Chang, C. Lee, J. Ho, T. Salimans, D. Fleet, and M. Norouzi, Palette: Image-to-image diffusion models, in *ACM SIGGRAPH 2022 Conference Proceedings* (Association for Computing Machinery, New York, NY, 2022), 10.1145/3528233.3530757.
- [4] P. Dhariwal and A. Nichol, Diffusion models beat GANs on image synthesis, in *Advances in Neural Information Processing Systems*, edited by M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J.W. Vaughan (Curran Associates, Inc., Vancouver, Canada, 2021), Vol. 34, pp. 8780–8794.
- [5] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, Normalizing flows for probabilistic modeling and inference, *J. Mach. Learn. Res.* **22**, 1 (2021).
- [6] Z. Kong and K. Chaudhuri, The expressive power of a class of normalizing flow models, [arXiv:2006.00392](https://arxiv.org/abs/2006.00392).
- [7] A. van den Oord, N. Kalchbrenner, L. Espeholt, k. kavukcuoglu, O. Vinyals, and A. Graves, Conditional image generation with PixelCNN decoders, in *Advances in Neural Information Processing Systems*, edited by D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (Curran Associates, Inc., Barcelona, Spain, 2016), Vol. 29.
- [8] D. Kodi Ramanah, T. Charnock, F. Villaescusa-Navarro, and B. D. Wandelt, Super-resolution emulator of cosmological simulations using deep physical models, *Mon. Not. R. Astron. Soc.* **495**, 4227 (2020).
- [9] Y. Li, Y. Ni, R. A. C. Croft, T. D. Matteo, S. Bird, and Y. Feng, AI-assisted superresolution cosmological simulations, *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2022038118 (2021).
- [10] Y. Ni, Y. Li, P. Lachance, R. A. C. Croft, T. Di Matteo, S. Bird, and Y. Feng, AI-assisted superresolution cosmological simulations—II. Halo substructures, velocities, and higher order statistics, *Mon. Not. R. Astron. Soc.* **507**, 1021 (2021).
- [11] X. Zhang, P. Lachance, Y. Ni, Y. Li, R. A. C. Croft, T. D. Matteo, S. Bird, and Y. Feng, AI-assisted superresolution cosmological simulations III: Time evolution, *Mon. Not. R. Astron. Soc.* **528**, 281 (2024).
- [12] A. Schanz, F. List, and O. Hahn, Stochastic super-resolution of cosmological simulations with denoising diffusion models, [arXiv:2310.06929](https://arxiv.org/abs/2310.06929).
- [13] N. Mudur and D. P. Finkbeiner, Can denoising diffusion probabilistic models generate realistic astrophysical fields?, in *Neural Information Processing Systems* (Curran Associates, Inc, New Orleans, LA, 2022), Vol. 36.
- [14] X. Zhao, Y.-S. Ting, K. Diao, and Y. Mao, Can diffusion model conditionally generate astrophysical images?, *Mon. Not. R. Astron. Soc.* **526**, 1699 (2023).
- [15] K. Karchev, N. Anau Montel, A. Coogan, and C. Weniger, Strong-lensing source reconstruction with denoising diffusion restoration models, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc, New Orleans, LA, 2022), Vol. 36.
- [16] R. Legin, M. Ho, P. Lemos, L. Perreault-Levasseur, S. Ho, Y. Hezaveh, and B. Wandelt, Posterior sampling of the initial conditions of the universe from non-linear large scale structures using score-based generative models, *Mon. Not. R. Astron. Soc.* **527**, L173 (2023).
- [17] J. Song, C. Meng, and S. Ermon, Denoising diffusion implicit models, [arXiv:2010.02502](https://arxiv.org/abs/2010.02502).
- [18] C. Lu, Y. Zhou, F. Bao, J. Chen, C. Li, and J. Zhu, DPM-solver: A fast ODE solver for diffusion probabilistic model sampling in around 10 steps, [arXiv:2206.00927](https://arxiv.org/abs/2206.00927).
- [19] V. Springel, R. Pakmor, A. Pillepich, R. Weinberger, D. Nelson, L. Hernquist, M. Vogelsberger, S. Genel, P. Torrey, F. Marinacci, and J. Naiman, First results from the IllustrisTNG simulations: Matter and galaxy clustering, *Mon. Not. R. Astron. Soc.* **475**, 676 (2017).
- [20] A. Pillepich, D. Nelson, L. Hernquist, V. Springel, R. Pakmor, P. Torrey, R. Weinberger, S. Genel, J. P. Naiman, F. Marinacci, and M. Vogelsberger, First results from the IllustrisTNG simulations: The stellar mass content of groups and clusters of galaxies, *Mon. Not. R. Astron. Soc.* **475**, 648 (2017).
- [21] J. P. Naiman, A. Pillepich, V. Springel, E. Ramirez-Ruiz, P. Torrey, M. Vogelsberger, R. Pakmor, D. Nelson, F. Marinacci, L. Hernquist, R. Weinberger, and S. Genel, First results from the IllustrisTNG simulations: A tale of two elements—chemical evolution of magnesium and europium, *Mon. Not. R. Astron. Soc.* **477**, 1206 (2018).
- [22] F. Marinacci, M. Vogelsberger, R. Pakmor, P. Torrey, V. Springel, L. Hernquist, D. Nelson, R. Weinberger, A. Pillepich, J. Naiman, and S. Genel, First results from the IllustrisTNG simulations: Radio haloes and magnetic fields, *Mon. Not. R. Astron. Soc.* **480**, 5113 (2018).
- [23] D. Nelson, A. Pillepich, V. Springel, R. Weinberger, L. Hernquist, R. Pakmor, S. Genel, P. Torrey, M. Vogelsberger, G. Kauffmann, F. Marinacci, and J. Naiman, First results from the IllustrisTNG simulations: The galaxy colour bimodality, *Mon. Not. R. Astron. Soc.* **475**, 624 (2017).
- [24] V. Springel, E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh, *Mon. Not. R. Astron. Soc.* **401**, 791 (2010).
- [25] T. Brown *et al.*, Language models are few-shot learners, in *Advances in Neural Information Processing Systems*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin (Curran Associates, Inc., Vancouver, Canada, 2020), Vol. 33, pp. 1877–1901.
- [26] L. Jiang, Palette: Image-to-Image Diffusion Models by Pytorch, <https://github.com/Janspiry/Palette-Image-to-Image-Diffusion-Models> (2022).
- [27] O. Ronneberger, P. Fischer, and T. Brox, U-Net: Convolutional networks for biomedical image segmentation, in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015*, edited by N. Navab, J. Hornegger, W.M. Wells, and A.F. Frangi (Springer International Publishing, Cham, 2015), pp. 234–241.
- [28] A. Brock, J. Donahue, and K. Simonyan, Large scale GAN training for high fidelity natural image synthesis, in *International Conference on Learning Representations* (2019), <https://openreview.net/forum?id=B1xsqj09Fm>.

- [29] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, Attention is all you need, [arXiv:1706.03762](https://arxiv.org/abs/1706.03762).
- [30] C. Saharia, J. Ho, W. Chan, T. Salimans, D. J. Fleet, and M. Norouzi, Image super-resolution via iterative refinement, *IEEE Trans. Pattern Anal. Mach. Intell.* **45**, 4713 (2023).
- [31] Y. Wu and K. He, Group normalization, *Int. J. Comput. Vis.* **128**, 742 (2020).
- [32] S. Elfving, E. Uchibe, and K. Doya, Sigmoid-weighted linear units for neural network function approximation in reinforcement learning, *Neural Netw.* **107**, 3 (2018), special issue on deep reinforcement learning.
- [33] F. Villaescusa-Navarro and Pylians, <https://github.com/franciscovillaescusa/Pylians3> (2020).
- [34] A. Banerjee and N. Dalal, Simulating nonlinear cosmological structure formation with massive neutrinos, *J. Cosmol. Astropart. Phys.* **11** (2016) 015.
- [35] L. Dinh, J. Sohl-Dickstein, and S. Bengio, Density estimation using Real NVP, [arXiv:1605.08803](https://arxiv.org/abs/1605.08803).
- [36] D. P. Kingma and P. Dhariwal, Glow: Generative flow with invertible 1×1 convolutions, [arXiv:1807.03039](https://arxiv.org/abs/1807.03039).
- [37] A. Rouhiainen, U. Giri, and M. Münchmeyer, Normalizing flows for random fields in cosmology, [arXiv:2105.12024](https://arxiv.org/abs/2105.12024).
- [38] A. Rouhiainen and M. Münchmeyer, De-noising non-Gaussian fields in cosmology with normalizing flows, [arXiv:2211.15161](https://arxiv.org/abs/2211.15161).
- [39] Y. Kvasiuk and M. Münchmeyer, An Auto-Differentiable Likelihood Pipeline for the Cross-Correlation of CMB and Large-Scale Structure due to the Kinetic Sunyaev-Zeldovich Effect, *Phys. Rev. D* **109**, 083515 (2024).
- [40] A.-S. Deutsch, E. Dimastrogiovanni, M. C. Johnson, M. Münchmeyer, and A. Terrana, Reconstruction of the remote dipole and quadrupole fields from the kinetic Sunyaev Zel'dovich and polarized Sunyaev Zel'dovich effects, *Phys. Rev. D* **98**, 123501 (2018).
- [41] K. M. Smith, M. S. Madhavacheril, M. Münchmeyer, S. Ferraro, U. Giri, and M. C. Johnson, KSZ tomography and the bispectrum, [arXiv:1810.13423](https://arxiv.org/abs/1810.13423).
- [42] S. Andrianomena, S. Hassan, and F. Villaescusa-Navarro, Invertible mapping between fields in CAMELS, [arXiv:2303.07473](https://arxiv.org/abs/2303.07473).
- [43] F. Villaescusa-Navarro *et al.*, The CAMELS project: Cosmology and astrophysics with machine-learning simulations, *Astrophys. J.* **915**, 71 (2021).
- [44] S. Luo and W. Hu, Diffusion probabilistic models for 3D point cloud generation, [arXiv:2103.01458](https://arxiv.org/abs/2103.01458).
- [45] V. Mikuni, B. Nachman, and M. Pettee, Fast point cloud generation with diffusion models in high energy physics, *Phys. Rev. D* **108**, 036025 (2023).
- [46] M. Leigh, D. Sengupta, G. Quétant, J. A. Raine, K. Zoch, and T. Golling, PC-JeDi: Diffusion for particle cloud generation in high energy physics, *SciPost Phys.* **16**, 018 (2024).
- [47] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, LoRA: Low-rank adaptation of large language models, in *International Conference on Learning Representations* (2022), <https://openreview.net/forum?id=nZeVKeeFYf9>.
- [48] Y. Gu, X. Wang, J. Z. Wu, Y. Shi, Y. Chen, Z. Fan, W. Xiao, R. Zhao, S. Chang, W. Wu, Y. Ge, Y. Shan, and M. Z. Shou, Mix-of-Show: Decentralized low-rank adaptation for multi-concept customization of diffusion models, [arXiv:2305.18292](https://arxiv.org/abs/2305.18292).