

Nonlinear 3D Cosmic Web Simulation with Heavy-Tailed Generative Adversarial Networks

Richard M. Feder*

*California Institute of Technology Division of Physics, Math, and Astronomy,
1200 East California Boulevard, Pasadena, CA 91125, USA*

Philippe Berger

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109, USA

George Stein

Berkeley Center for Cosmological Physics, University of California, Berkeley, CA 94720, USA

(Dated: May 26, 2020)

Fast and accurate simulations of the non-linear evolution of the cosmic density field are a major component of many cosmological analyses, but the computational time and storage required to run them can be exceedingly large. For this reason, we use generative adversarial networks (GANs) to learn a compressed representation of the 3D matter density field that is fast and easy to sample, and for the first time show that GANs are capable of generating samples at the level of accuracy of other conventional methods. Using sub-volumes from a suite of GADGET-2 N-body simulations, we demonstrate that a deep-convolutional GAN can generate samples that capture both large- and small-scale features of the matter density field, as validated through a variety of n -point statistics. The use of a data scaling that preserves high-density features and a heavy-tailed latent space prior allow us to obtain state of the art results for fast 3D cosmic web generation. In particular, the mean power spectra from generated samples agree to within 5% up to $k = 3$ and within 10% for $k \leq 5$ when compared with N-body simulations, and similar accuracy is obtained for a variety of bispectra. By modeling the latent space with a heavy-tailed prior rather than a standard Gaussian, we better capture sample variance in the high-density voxel PDF and reduce errors in power spectrum and bispectrum covariance on all scales. Furthermore, we show that a conditional GAN can smoothly interpolate between samples conditioned on redshift. Deep generative models, such as the ones described in this work, provide great promise as fast, low-memory, high-fidelity forward models of large-scale structure.

I. INTRODUCTION

Modeling the evolution of the large-scale structure of the Universe is a crucial part of many cosmological analyses. The initial linear growth of structures becomes highly non-linear at later times and smaller scales, such that analytical approaches are unable to accurately predict observations in those regimes. For this reason, numerical simulations [1] of a representative volume of the universe are required to bridge the gap between the early universe and present day observations.

N-body simulations, which numerically solve the Poisson equation, are a standard operating procedure for modeling large-scale structure. These methods are exact in the sense that they directly simulate the gravitational evolution of some set of collisionless dark matter “particles”, but the required mass and time resolution to make model predictions at scale is computationally expensive and requires significant data storage. Next-generation cosmological inferences will require thousands of simulations on the largest scales in order to constrain models like Λ CDM at the sub-percent level [2], a task that is currently intractable to all but the largest future com-

putational grants when using standard N-body methods. Another consideration is memory storage. At the moment, storing thousands of full scale N-body simulation snapshots is only feasible for petabyte-scale data storage facilities. Even with access to such resources, the analysis workflow may be cumbersome if one has to repeatedly transfer simulations between a cluster and local machine.

The need for fast and accurate simulations at a cheaper computational cost has led to the development of many “approximate” methods of cosmic structure simulation [3–11], which implement various physical approximations of gravity to decrease the simulation run-time and/or memory requirement while attempting to maintain a high degree of accuracy. These methods have shown to be very useful, but still require a non-negligible amount of computation time, and the accuracy of their results has room to be improved upon.

One may then wonder if moving past physical approximations of gravity, and instead allowing for arbitrary non-linear mappings from initial conditions (or some other prior distribution) to present day-observables, could provide more accurate simulations at a smaller computational cost. This application of machine learning (ML) techniques to aid in the efficient modeling of large scale structure has not gone un-investigated, and there has recently been great interest in the use of ML

* rfederst@caltech.edu

across the field¹. Image-to-image mapping techniques that transform initial conditions to the final structures [12–14] or augment dark matter-only simulation results with various physical effects not present in the simulation [15–20] have shown promise. However, these still require the generation of initial conditions, or the full simulation output, respectively. A fully generative model that does not require these intermediate steps is desirable.

In particular, generative adversarial networks (GANs) have garnered attention in recent years due to their ability to generate fast realizations of large scale structure. GANs are one class of deep generative modeling techniques that seeks to learn a probability distribution over an arbitrary data distribution, doing so through a two-player adversarial optimization scheme [21] (see Section II for more details).

To date, there have been a number of successful implementations of cosmological modeling tasks using GANs. They have been used to generate accurate weak lensing convergence maps as evaluated through the power spectrum and a variety of Minkowski functionals [22], to model 2D slices of the cosmic density field from 3D N-body simulations [23], and to generate 3D realizations of neutral hydrogen [24]. These methods learn a compressed representation of the data distribution and sample from the corresponding lower-dimensional latent space to produce new samples.

There are also a number of in-painting and super-resolution methods that do not involve the determination of a compressed data representation. One method that is particularly promising is [25], which uses a super-resolution GAN and deep physical model to map low resolution, 3D N-body simulations to their high resolution counterparts. This is similar in concept to [18] and [19], which used GANs to map dark matter density fields to corresponding halo number count maps and hydrodynamical quantities, respectively. [26] propose a super-resolution scheme for generating large-scale realizations of the matter density field hierarchically, treating the scalability problem separately from sample accuracy.

Three-dimensional N-body data sets are challenging to model because each sample realization contains many degrees of freedom. In addition, the voxel PDF for each realization has large tails, with matter densities spanning several orders of magnitude. Natural image data sets, a common focus of GAN applications, also contain tails in pixel space [27], however the application in cosmology is unique because the tails of the voxel PDF often have a disproportionate impact on the relevant summary statistics.

This exploratory work identifies a number of factors important for training GANs that, from a compressed representation, produce samples that are both accurate and preserve the sample variance of the data distribution. In a controlled setting, we demonstrate the need

for data scalings that preserve information in the high-density regime. Furthermore, we find that heavy-tailed distributions are more suitable priors for modeling the latent space, significantly improving the accuracy of our generated covariance and correlation matrices. Through these and other considerations, we obtain state of the art results for 3D cosmic web generation as reflected through the matter density power spectrum, a range of bispectra, and other cosmological summary statistics.

The report is structured as follows: in §2 we introduce the statistical and optimization framework of GANs. We describe the GADGET-2 data set and training procedure in §3 and outline the modeling choices specific to this work. The corresponding results for single redshift and redshift-interpolated models are presented in §4. Finally, we discuss our results in §5, concluding with a discussion of future directions and applications.

II. GENERATIVE ADVERSARIAL NETWORKS

Generative adversarial networks (GANs), originally introduced in [21], are a class of unsupervised models that learn an underlying probability distribution, which we denote $p_{data}(x)$, through adversarial training. In this paradigm, a generator G attempts to produce samples that appear real to a discriminator D . The discriminator is tasked with distinguishing between real training samples and synthetic ones.

The standard training process, shown schematically in Figure 1, goes as follows:

1. The input is either a random latent vector \mathbf{z} (typically drawn from a unit-normal distribution $\mathcal{N}(0, \mathbf{I})$), or some structured input.
2. The generator G , parameterized by a neural network, takes the latent vector \mathbf{z} and outputs a sample $G(\mathbf{z})$.
3. The discriminator D , also parameterized by a neural network, takes real samples \mathbf{x} and generated samples $G(\mathbf{z})$ as input, and outputs scores $D(\mathbf{x})$ and $D(G(\mathbf{z}))$, respectively. Each score represents the degree to which the discriminator believes a given sample comes from the data distribution $p_{data}(\mathbf{x})$. When the score is scaled to the range $[0, 1]$, it is sometimes interpreted as an implicit likelihood of the data given the discriminative model, $p(\mathbf{x}|D)$.
4. The discriminator predictions are compared to the ground truth labels, from which a loss $\mathcal{L}(D, G)$ is computed.
5. The loss is used to update the weights of G and D through backpropagation, flowing through the discriminator and then through the generator.
6. Steps 1-5 are repeated, looping through the training set over several epochs.

¹ Comprehensive list at github.com/georgestein/ml-in-cosmology

Formally, GAN training is a min-max problem, i.e. the generator and discriminator compete to minimize and maximize a loss function, respectively:

$$\min_G \max_D \mathcal{L}(G, D) = \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [\log(D(\mathbf{x}))] + \mathbb{E}_{\mathbf{z} \sim p_g(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))] \quad (1)$$

The first expectation on the right hand side updates the discriminator, while the second term updates both the discriminator and the generator. By training both networks simultaneously, the generator gradually learns to produce samples from the data distribution. As described in [21], the optimal generator and discriminator reach a Nash equilibrium, where neither network can make unilateral improvement over the other. That work also demonstrated that, for an optimal discriminator and certain choices of loss function, training the generator is equivalent to reducing the Jensen-Shannon divergence between the generated distribution $p_g(z)$ and the empirical distribution $p_{data}(x)$. The Jensen-Shannon divergence is the symmetrized version of the Kullback-Leibler divergence:

$$\mathbb{D}_{JS}(p_g(\mathbf{x})|p_{data}(\mathbf{x})) = \frac{1}{2}\mathbb{D}_{KL}(p_g||M) + \frac{1}{2}\mathbb{D}_{KL}(p_{data}||M) \quad (2)$$

where

$$M = \frac{p_{data}(\mathbf{x}) + p_g(\mathbf{z})}{2}. \quad (3)$$

A common assumption when using GANs for modeling is that while the observed data may be high-dimensional, the underlying data distribution exists on a lower-dimensional manifold. In this spirit, generator outputs $G(\mathbf{z})$ may be interpreted as samples drawn from some underlying manifold \mathcal{G} , for which $\dim \mathcal{G} = |\mathbf{z}|$ and $\text{Cov}(\mathcal{G}) = G$, where G is the neural network. In many applications, \mathbf{z} is chosen to have dimension $\mathcal{O}(100)$, though other techniques like in-painting and super-resolution may use higher-dimensional inputs. One advantage of learning a compressed representation of the data is that generating high-dimensional samples in the data space reduces to sampling from a simple, lower-dimensional latent space and performing a forward pass through the generator network G . This also means that sample generation is fast, since the only computation needed after a draw from the latent distribution is a forward pass through the network.

A. Conditional GANs

There are cases where one would like to incorporate additional information relevant to the data generation process. The conditional GAN (cGAN) is a natural extension of the GAN that addresses these cases [28]. The training problem is a modified version of the original ob-

jective:

$$\min_G \max_D \mathcal{L}(G, D) = \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [\log(D(\mathbf{x}|\mathbf{c}))] + \mathbb{E}_{\mathbf{z} \sim p_g(\mathbf{z})} [\log(1 - D(G(\mathbf{z}|\mathbf{c})))] \quad (4)$$

Here, \mathbf{c} denotes conditional information, which can be appended both to the latent vector of the generator and to the discriminator as an extra feature. A properly trained conditional generator is then able to interpolate within the conditional distribution $p(\mathbf{x}|\mathbf{c})$. If the marginal distribution $p(\mathbf{c})$ is known, then training the cGAN results in a model that can generate samples from the joint distribution $p(\mathbf{x}, \mathbf{c}) = p(\mathbf{c})p(\mathbf{x}|\mathbf{c})$. For example, in the case of the standard MNIST data set of handwritten digits [29], an obvious conditional distribution is $c \sim \text{Cat}(k)$, where $k \in \{0, 1, \dots, 9\}$. There is nothing that specifically constrains the form or dimensionality of the conditional information, so long as it can be passed through layers of the generator and discriminator.

The theoretical guarantees of GANs are compromised by a number of factors in practice. Because optimization is performed over network parameters and not over probability densities, convergence to the global optimum where $p_g = p_{data}$ may be complicated by critical points and/or inefficient training. Furthermore, unstable training dynamics can lead one of the networks to dominate the other such that gradients are suppressed and learning stops prematurely. When the generator dominates the discriminator, the pathology is referred to as “mode collapse”, in which the generator learns to model some small subset of the data distribution. Despite these practicalities, GANs with reasonable model capacity and sufficient training data can be trained successfully and produce high-fidelity samples from the underlying data distribution [30, 31].

III. DATA SET AND TRAINING PROCEDURE

A. N-body simulations with GADGET-2

While many types of cosmological data have non-linear properties, we choose to train a GAN to model the dark matter component of the “cosmic web” [32]. The cosmic web is a combination of baryonic and dark matter that evolves with the expansion of the universe and eventually clusters, forming the overdense seeds of galaxy and star formation. Due to properties of gravitational collapse in an expanding universe, the cosmic web is comprised of many coherent structures, such as one-dimensional filaments and two-dimensional sheets.

Our training set comes from a suite of thirty-two dark matter only simulations made with the software GADGET-2². Each simulation is initialized with a

² <https://wwwmpa.mpa-garching.mpg.de/gadget/>

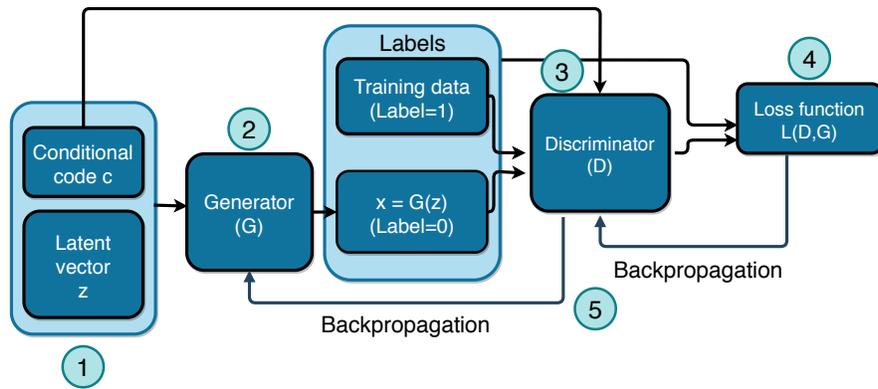


FIG. 1. Schematic for generative adversarial networks. As described in Section II, a latent vector is drawn from some prior (1) and used as input for the generator network (2) which outputs a sample in the data space. A batch of these samples, along with a batch of training set samples, are passed to the discriminator (3). The discriminator network outputs a set of scores representing how much it believes the samples come from the data distribution, and these scores are compared with ground truth labels to calculate a loss (4). This loss is then used to update the parameters of both the generator and discriminator network through backpropagation (5).

unique seed, using the cosmological parameters that match the best-fitting results of the WMAP+BOSS DR9 analysis [33]. The initial conditions are derived using second-order Lagrangian perturbation theory (2LPT) at a starting redshift of $z = 49$. Each simulation volume has 512^3 particles in a $(512 \text{ Mpc})^3$ volume that are evolved to the present day, with snapshots written to disk at $z = \{3.0, 2.0, 1.5, 1.0, 0.5, 0.25, 0.0\}$, after which they are placed on a 3D grid with 512 resolution elements per side. After gridding, the GADGET-2 density field is given in units of $10^{10} M_{\odot}/\text{cell}$, and we will refer to this normalized matter density field as $\tilde{\rho}(x)$.

A naive application of GANs to directly generate 512^3 -voxel samples would be challenging because of the high dimension of the data and the availability of training samples. Instead, we train on $(64 \text{ Mpc})^3$ sub-volumes, which gives us a training set with 16384 samples. We further exploit the isotropy of our samples to augment our data set through random rotations and flips, which increases the effective training set size by a factor of 6 (sides) \times 4 (90 deg. rotations) \times 2 (random flips) = 48.

B. Network architecture

We base our network architecture on [34], in which the generator upsamples a latent vector \mathbf{z} with 3D transpose convolutional layers and the discriminator downsamples with 3D convolutional layers. This is one example of a deep convolutional GAN (DCGAN, [35]). The parameters and activations used to configure the network are shown in Table I.

While a variety of large-scale, sophisticated neural network architectures have been presented in the machine learning literature, we found that (relatively) smaller networks had sufficient capacity to generate high-fidelity samples for our application and were also easier to train.

In particular, our results did not change significantly when the dimension of the latent space was increased, when the number of filters was increased, or when extra convolutional layers were added beyond the base architecture.

| | |
|------------------|------------------------------------|
| Filter sizes | {4, 4, 4, 4, 4} |
| n_{filter} | {256, 128, 64, 32, 1} [*] |
| Strides | {1, 2, 2, 2, 2} [*] |
| Padding: | {0, 0, 0, 0, 1} [*] |
| Layer Act. | Leaky ReLU (G), ReLU (D) |
| Final Act. | Tanh (G), Sigmoid (D) |
| Latent dimension | 100 |
| N_{params} | 4,393,920 |

TABLE I. Architecture specifications for each layer of the generator network (G) and discriminator network (D). Asterisk signifies that order of layers is reversed in the discriminator.

C. Data scaling

The evolved matter density field, $\tilde{\rho}(x)$, exhibits a dynamic range spanning several orders of magnitude. While not strictly prohibited, directly generating samples with values in this range would be difficult for a generator network. In fact, [35] argues that a bounded final activation allows the model to more quickly saturate the range of the training distribution. We use the transformation from [23],

$$c(x) = \frac{2\tilde{\rho}(x)}{\tilde{\rho}(x) + \kappa} - 1, \quad (5)$$

which scales our density fields to the range $[-1, 1]$, suitable for a generator network with a hyperbolic tangent (tanh) final activation.

In [23], the scaling parameter κ is set to 4. However, after performing a series of experiments, we choose to adopt a larger value, $\kappa = 45$, that better preserves information in the high-density regime and results in more accurate samples. Figure 2 illustrates how a small variation in scaled density can correspond to a large variation in the final matter density (see Appendix B for more detailed analysis). We experimented with the logarithmic scaling used in [24] as well as other piecewise scalings, but found none was able to capture the statistics of the matter density field as well as Equation (5).

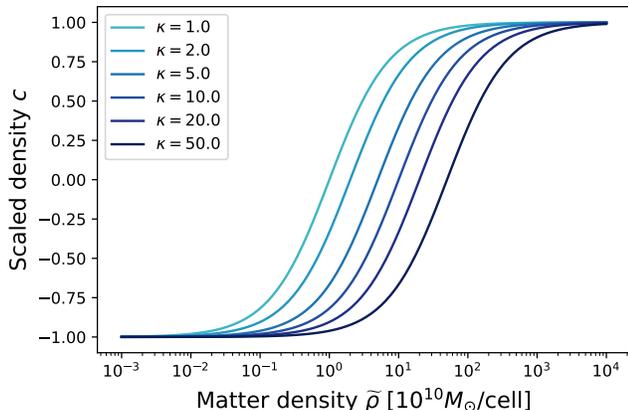


FIG. 2. Scaling from Equation (5) plotted for different values of κ . Scalings with small values of κ have smaller gradients at high density.

D. Modeling the latent space with a heavy-tailed prior

The dynamic range of voxel matter densities and of the matter density PDF make N-body simulations a unique data set. While many data sets have non-Gaussian tails in their pixel/voxel distributions, cosmological N-body data sets are particularly challenging to model because the least frequent, high-density voxels have the largest impact on the relevant summary statistics. For the data set considered in this work, the challenge is in capturing the mean voxel distribution function at high density while also capturing its sample variance.

The standard approach for sample generation with a GAN is to draw a noise vector \mathbf{z} from a Gaussian or uniform distribution. However, there is no strong motivation to use these as priors (beyond ease of sampling) if they do not reflect important features of the data. As shown in Section IV, it is only when a heavy-tailed prior is used that *both* accuracy and sample variance are captured by our model. To date, and to our knowledge, there are not many instances in the machine learning literature where heavy-tailed distributions are used to model the latent space. [36] develops non-parametric priors through on-the-shelf optimization tools in order to im-

prove latent space interpolation, though our work is less focused on semantic interpolation than common applications where natural image data sets are modeled. We choose a Student-t prior for the latent space and tune ν , the number of degrees of freedom in the distribution. The univariate Student-t distribution has probability density

$$P(z) = \frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\nu\pi}\Gamma(\frac{\nu}{2})} \left(1 + \frac{z^2}{\nu}\right)^{-\frac{\nu+1}{2}} \quad (6)$$

where $\Gamma(\cdot)$ is the Gamma function. When ν is smaller, the Student-t distribution has larger tails, and in the limit where $\nu \rightarrow \infty$, $P(z)$ converges to a Gaussian distribution. Through empirical tests, we find $\nu = 10$ performs best for our data set. We reiterate that our argument for modeling the latent space is not specific to the Student-t distribution. One might capture tail behavior with a Lorentzian distribution, a Cauchy distribution, or even a mixture of Gaussians.

E. Training Procedure

The generator and discriminator were trained for 150 epochs with learning rates of 1×10^{-4} and 5×10^{-5} , respectively, using the ADAM optimizer [37]. We chose to train with a batch size of 16, meaning gradient updates to the networks at each iteration are computed based on sixteen real and sixteen generated samples. Within each epoch, each training set mini-batch is chosen randomly, after which a set of random rotations and flips are applied to the samples. Following the training procedure from [34], the discriminator weights are only updated in a given batch iteration if its accuracy for that batch falls below 80%. This is one heuristic to prevent the discriminator from outperforming the generator, and in our case led to more stable training.

We also conducted several experiments with the Wasserstein GAN [38] with gradient penalty [39] scheme (WGAN-GP). The WGAN-GP trained stably for a range of architectures, optimizers and hyperparameters. While the resulting models were able to capture the mean summary statistics of the data as described in IV, they struggled to capture sample variance in the generator distribution. Because of this, we chose to use the standard GAN for the remainder of our experiments.

After some training, a GAN may be near its optimum, but if the learning rate is too large then the networks will be limited in their ability to further converge. With this in mind, we use a learning rate scheduler which periodically reduces the learning rate by a factor γ . We choose to decay the learning rates of both the generator and discriminator by 50% ($\gamma = 0.5$) after every 1000 gradient updates to the generator.

The losses of the generator and discriminator are plotted as a function of batch iteration in Figure 3. Despite small fluctuations which come from noisy estimates of the loss computed from mini-batches, training stabilizes

relatively quickly. The losses do not indicate when the models have converged, but help verify that the training dynamics of the GAN are controlled. As there are no asymptotic convergence guarantees when training GANs, one is at liberty to choose the model at any point in training where it best captures the desired summary statistics. To optimize our model selection, we save a version of the generator at the end of each epoch and determine for which model the mean voxel PDF and power spectra match the data best. As discussed in Appendix A, we find a correspondence between the scale width of the latent distribution and the mean power spectral amplitude of generated samples. After training, we use this to further calibrate the model to the training set by perturbing the scale width of the latent distribution, typically at the percent level.

Computations were done on the Texas Advanced Computing Cluster on a node with $4 \times$ Nvidia 1080-TI GPUs with 128 GB of RAM, and took roughly nine hours in wall clock time to train. The model and training were implemented using the PyTorch library³. We have made the code for this work publicly available on Github⁴.

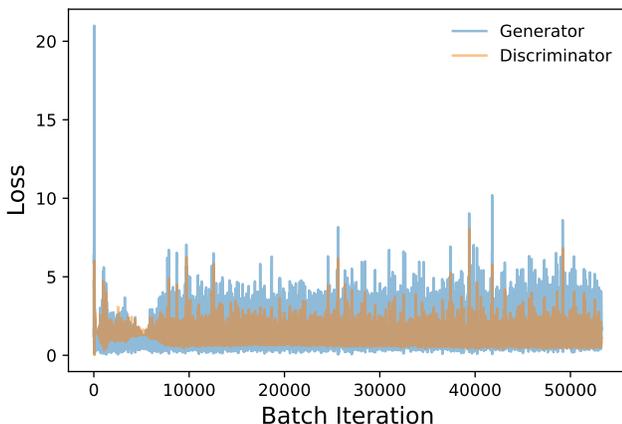


FIG. 3. Losses of the generator (blue) and discriminator (orange) as a function of batch iteration for a training run using a 3D DCGAN. The generator loss is computed from the second term in the objective from Equation (1), while the discriminator loss is the sum of the first and second terms. The losses do not indicate when the models have converged, but help verify that the training dynamics of the GAN are controlled.

IV. RESULTS

We first present the results of our model trained on the full N-body data set evolved to $z = 0$ (present day). In this section, the performance of our models is evaluated directly on the standard cosmological summary

statistics. These are more physically interpretable than discriminator-based metrics such as the Frechet Inception Distance (FID) or the Inception Score (IS) that are common in the machine learning literature. For $z = 0$, the model is trained on data that are scaled using (5) with parameter $\kappa = 45$. In our experiments, $\kappa = 45$ was the largest value where training was stable. Unless specified otherwise, our results use $N_s = 2000$ samples to evaluate each GAN model and the full training set for GADGET-2 ($N_s = 16384$).

Figure 4 shows visualizations of $(64 \text{ Mpc})^3$ volumes from GADGET-2 (top row) and from a trained GAN (bottom row). Samples from the generator network contain many of the filamentary structures known to populate the cosmic web. At $(64 \text{ Mpc})^3$, samples share similar types of features but also exhibit considerable variation by eye. The large variation on these scales is expected, as the homogeneity scale of the universe is typically accepted as $\gtrsim 100 \text{ h}^{-1} \text{ Mpc}$ [40, 41].

A. Voxel distribution function

Figure 5 shows voxel histograms for the scaled densities (top) and matter densities (bottom) of real and synthetic samples. Both are shown to highlight the fact that while the GAN trains on data scaled between $[1, -1]$, the re-scaling back to matter densities may produce unanticipated features. By using a larger scaling parameter for the data transformation, we prioritize information encoded in high-density voxels. This is reflected in the N-body and GAN-generated voxel PDFs, which agree well for the upper four orders of magnitude in density and diverge below $\tilde{\rho} \sim 10^{-1}$. The deviation at low density is not evident when looking at the scaled voxel PDFs alone, which appear consistent across the full range.

The voxel PDFs in Figure 5 also highlight the difference between models that use a Gaussian latent distribution and a heavy-tailed distribution. As alluded to in the previous section, initial tests produced models that generally underestimated the variance of the training set. Furthermore, this underestimated variance was present across several summary statistics. A similar phenomenon is observed and noted in [24]. If the variance of high-density voxels is underestimated, one should expect the variance of subsequent n -point statistics to also be underestimated. A number of potential remedies to this problem were attempted. Adding convolutional layers to the generator/discriminator did not yield noticeable improvements, nor did modifications of the size of the latent space, the number/size of filters in the networks, the scaling of the data or the training dynamics. However, one can see that the heavy-tailed GAN more faithfully captures variance in the high-density voxel PDF.

One validation for generative models of the matter density field is mass conservation. On the one hand, the question of mass conservation in generated simulations is ill-posed because there is no initial density field the

³ <https://pytorch.org/>

⁴ <https://github.com/RichardFeder/gan-work>

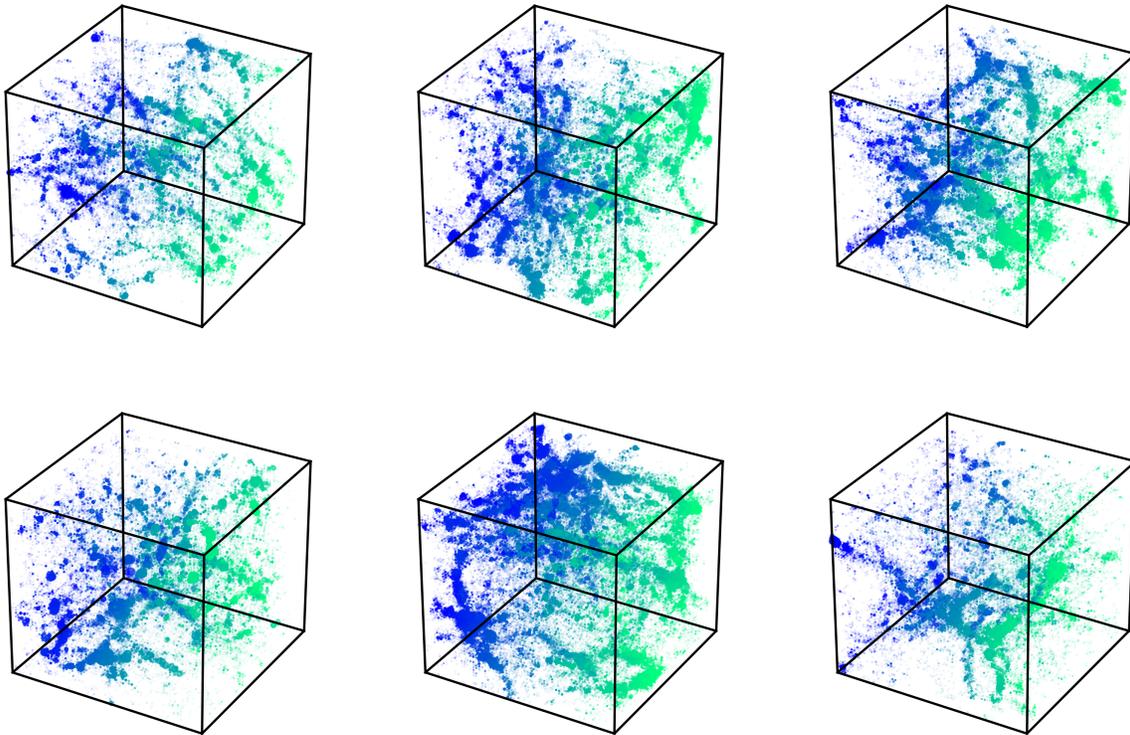


FIG. 4. 3D visualizations of GADGET sub-volumes (top row) and samples generated with a DCGAN (bottom row). Each sub-volume is visualized by plotting a number of points within the volume with size $R = (c(x) + 1)^5$, where $c(x)$ is the scaled density from (5) for a given location. This is done to highlight the relevant features of large scale structure. Color scale denotes position along the x -axis.

model uses as input. However, we can examine how the distribution of mean densities $P(\bar{\rho}_{\text{subvolume}})$ for an ensemble of generated sub-volumes compares to that of N-body simulations⁵. The average density of a sample is a weighted integral over the voxel PDF, so mismatches in the real/generated voxel PDFs should lead to changes in the distribution of average densities. Figure 6 compares estimates of $P(\bar{\rho}_{\text{subvolume}})$ from real and synthetic samples. Samples from the GAN trained with a Gaussian latent distribution have average densities $\sim 12\%$ higher compared to samples in the training set. This is consistent with the fact that on average, the high-density voxel PDF is higher than that of the training set. The heavy-tailed GAN does slightly better in this regard, with a median average density $\sim 5\%$ higher than that of the training set.

B. Power spectrum

Figure 7 shows a comparison of real and synthetic median power spectra, along with their 16th and 84th percentiles. The power spectrum is the Fourier transform of the two point correlation function $\xi(r)$:

$$P(|\vec{k}|) = \int \xi(\vec{r}) e^{-i\vec{k}\cdot\vec{r}} d^3\vec{r}. \quad (7)$$

where $\xi(|\vec{r}|) = \langle \delta(\vec{r}') \delta(\vec{r}' + \vec{r}) \rangle$, $k = 2\pi/\lambda$ is the fluctuation wavenumber and $\delta(\vec{r}) = \rho(\vec{r})/\bar{\rho} - 1$ is the density field contrast. We use the software package `nbodykit`⁶, which uses a Fast Fourier Transform (FFT) method to compute the power spectra of our samples. Also plotted are the mean absolute fractional deviations of generated power spectra with respect to GADGET-2 power spectra, i.e. $|(\bar{P}_{\text{gen}}(k)/\bar{P}_{\text{real}}(k)) - 1|$. Both GANs produce samples that are accurate well into the non-linear regime. However, the heavy-tailed GAN modeled produces more accurate power spectra, with significant improvement at

⁵ To clarify, $\bar{\rho}_{\text{subvolume}}$ denotes the average density of a given $(64 \text{ Mpc})^3$ sub-volume, and will be different from the average density of a full GADGET-2 simulation. For the training set, each $\bar{\rho}_{\text{subvolume}}$ is computed from an already normalized matter density field $\bar{\rho}(x)$.

⁶ <https://nbodykit.readthedocs.io/en/latest/index.html>

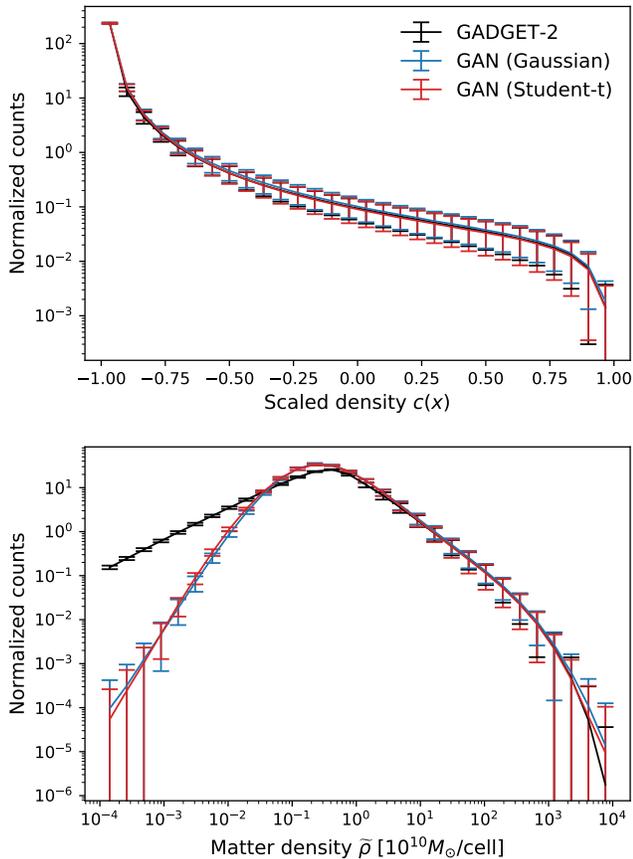


FIG. 5. Voxel probability distribution functions for scaled densities $[-1, 1)$ (top) and original matter densities (bottom) of GADGET-2 (black) and GAN-generated sub-volumes. Blue error bars correspond to a GAN model trained with a Gaussian latent space prior, while red error bars indicate results using a Student-t prior. Each set of error bars indicates the 1σ standard deviation within each bin, computed from 2000 samples.

large k . In particular, the mean power spectra from GADGET-2 and our heavy-tailed GAN agree within 5% for $k \leq 3$ and to $< 10\%$ up to $k = 5$. At these scales, baryonic effects have a significant impact on the matter power spectrum, and so we do not claim to capture the full matter power spectrum here. Nonetheless, our results on dark matter only simulations demonstrate the ability of our generative models to capture highly non-linear behavior on small scales. The ensemble of power spectra also capture super-sample variance from modes larger than 64 Mpc, since our training samples are sub-volumes from larger simulations.

1. Covariance and Correlation Matrices

Figures 9 and 10 show the covariance and correlation matrices for samples generated by Gaussian and heavy-tailed GANs, respectively, along with their devi-

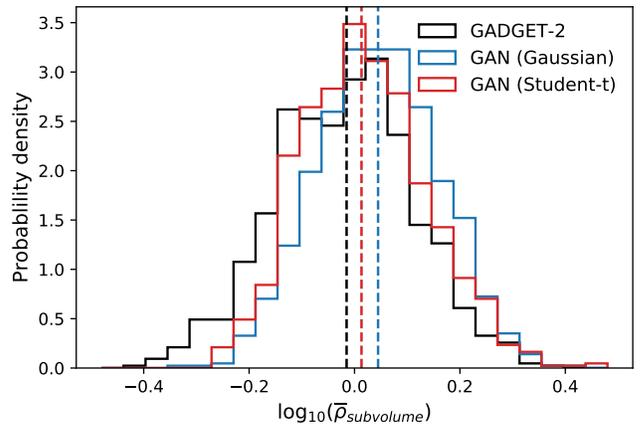


FIG. 6. Log-scaled average matter densities for GADGET-2 sub-volumes (black) and samples generated from GANs with Gaussian (blue) and Student-t (red) latent space priors. Dashed lines indicate the median average density of each distribution. Matter density PDFs are each calculated from 2000 samples.

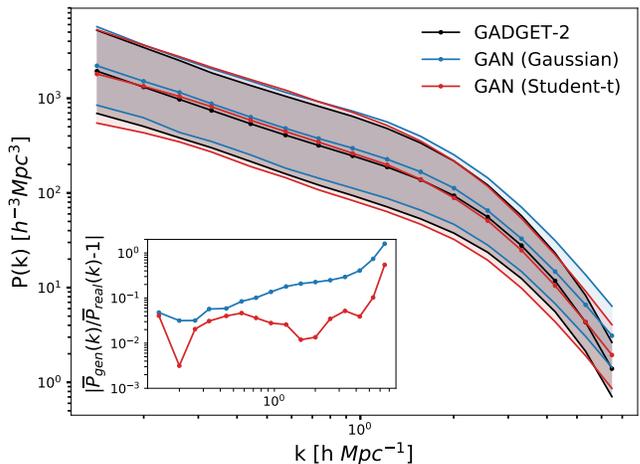


FIG. 7. Comparison of (log-scaled) power spectra between GADGET-2 sub-volumes at $z = 0$ (black) and samples drawn from GANs with Gaussian (blue) and Student-t (red) latent space priors. Shaded regions bound the 16th and 84th percentiles of the sample ensembles for each k bin. These percentiles are calculated from 2000 samples by each GAN model and from the full GADGET-2 data set. Bottom left corner: fractional deviations of mean power spectra between GAN and GADGET-2 samples.

ation from GADGET-2 N-body samples. We compute power spectrum covariance matrices with the sample covariance estimator

$$\hat{C}_{ij} = \frac{1}{N_s - 1} \sum_{n=1}^{N_s} [P_n(k_i) - \bar{P}(k_i)] [P_n(k_j) - \bar{P}(k_j)], \quad (8)$$

in which N_s is the number of samples, $P_n(k)$ is the dark matter power spectrum of sample n and $\bar{P}(k)$ is the mean

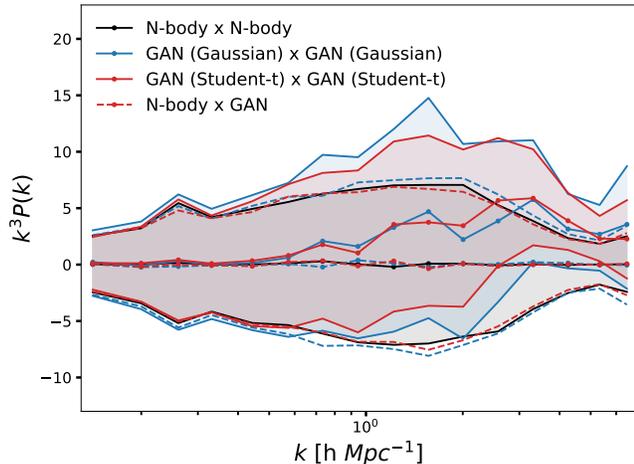


FIG. 8. Cross-power spectra for various combinations of GADGET-2 (black) and synthetic (blue, red) N-body sub-volumes. Dashed lines indicate cross power spectra between GAN generated samples and GADGET-2 sub-volumes. Each cross-spectrum combination is evaluated using 2000 sub-volume pairs. Shaded regions bound 16th and 84th percentiles of for each sample ensemble and within each k bin.

power spectrum across samples. The correlation matrix is further defined as

$$\hat{R}_{ij} = \frac{\hat{C}_{ij}}{\sqrt{\hat{C}_{ii}\hat{C}_{jj}}}. \quad (9)$$

The improvement from using a heavy-tailed prior is especially apparent when comparing covariance matrices. The covariance matrices computed from the heavy-tailed GAN are over 50% more accurate on all scales, agreeing to between 5% – 25% up to $k = 3$. The heavy-tailed GAN also improves the estimated correlation matrices on all scales, with notable improvement for correlations between large and small scales. At high k , covariance matrices from both GANs overestimate the true covariance, which is consistent with overestimates of the small-scale power spectrum seen in Figure 7. Both sets of synthetic correlation matrices follow the structure of the N-body correlation matrices. This is the first work to present power spectrum covariance matrices of GAN-generated 3D dark matter density fields.

2. Cross-power spectrum

Cross-power spectra help quantify the degree to which generated samples are correlated with one another and with the training set at different spatial scales. Figure 8 shows N-body \times N-body, N-body \times GAN and GAN \times GAN cross spectra, each evaluated from 2000 pairs of samples. Setting aside any large-scale effects that might come from evaluating sub-volumes of the same N-body

simulation, a random N-body sample should not (statistically) have any aligned structure with another random sample, i.e. the average cross spectrum should be zero. While the cross spectra between real and synthetic samples are consistent with zero, the non-zero GAN \times GAN cross-spectra observed for $k \geq 0.5$ indicate that our models produce some coherent structure across samples on these scales.

C. Bispectrum

A comparison of matter bispectra allows us to verify that our trained GANs learn to accurately model non-linear features of the cosmic web. The bispectrum is defined as the Fourier transform of the three-point correlation function:

$$(2\pi)^3 B(\vec{k}_1, \vec{k}_2, \vec{k}_3) \delta_D(\vec{k}_1 + \vec{k}_2 + \vec{k}_3) = \langle \delta(\vec{k}_1), \delta(\vec{k}_2), \delta(\vec{k}_3) \rangle \quad (10)$$

where δ_D is the Dirac delta. Figure 11 shows bispectra calculated from real and synthetic N-body samples, for a variety of large- and small-scale triangle configurations. We calculate sample bispectra using the software package `Pylians`⁷. The accuracy of the median generated bispectra is best at small scales and gets slightly worse for larger triangle configurations (in physical space), though notably these errors appear constant with respect to θ , the angle between wave-vectors. The angular dependence of each bispectrum configuration also appears to be captured by the 16th and 84th percentiles of the GAN sample ensembles. Overall, the heavy-tailed GAN yields a more accurate bispectrum distribution across samples, consistent with our earlier comparisons of matter power spectra and voxel density distributions.

D. Interpolation over redshift

In addition to single redshift sample generation, we consider GAN models that interpolate in a latent space augmented by conditional redshift information. In particular, we seek to establish that GANs can smoothly interpolate in redshift space for a given volume. For this task, we train on snapshots of twenty-four (512 Mpc)³ GADGET-2 volumes at redshifts $z = 3.0, 1.5, 0.5$ and 0.0, which gives us ~ 48000 training samples. While the scaling parameter $\kappa = 45$ worked well for N-body simulations at $z = 0$, such a scaling does not effectively capture information at higher redshifts, when the matter density field is closer to Gaussian. Indeed, attempts to train our conditional redshift model with $\kappa = 45$ were unsuccessful and often ended in divergent training dynamics. We instead use $\kappa = 4$, which allows for stable training, albeit at the cost of accuracy on small scales.

⁷ <https://github.com/franciscovillaescusa/Pylians>

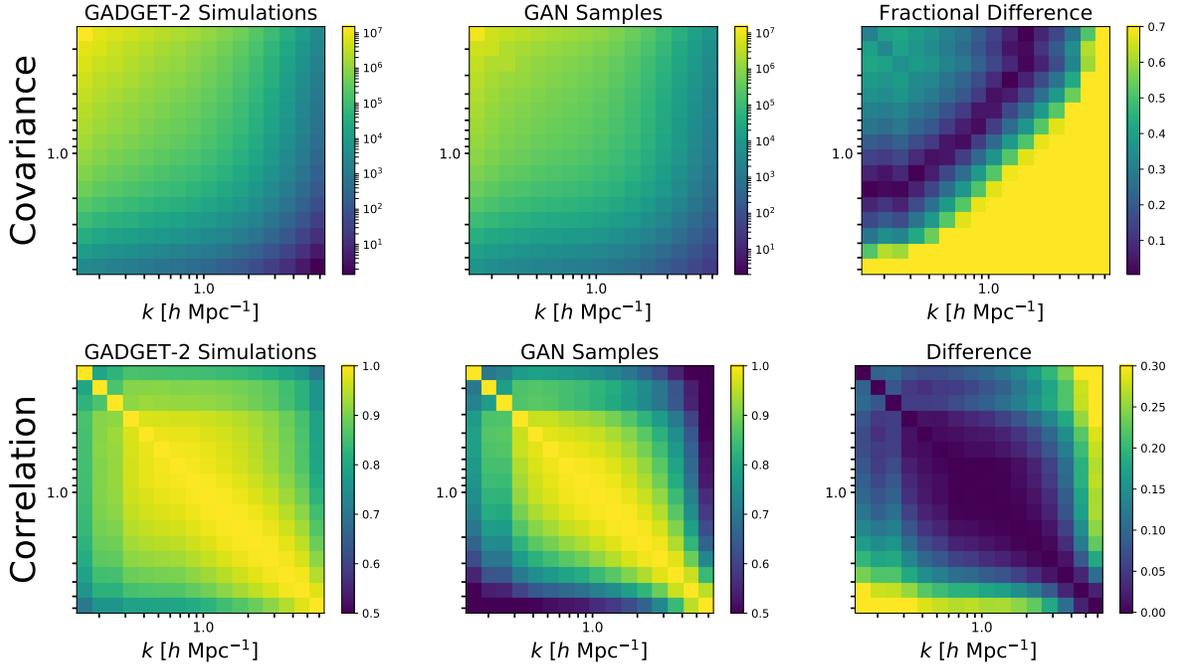


FIG. 9. **Top row:** Power spectrum covariance matrices of GADGET-2 N-body sub-volumes (left), synthetic samples drawn from a GAN with Gaussian latent space prior (middle), and their fractional difference (right). **Bottom row:** same as top row but for power spectrum correlation matrices and their difference. The covariance and correlation matrices from our GAN model are computed from 2000 synthetic samples, while those of GADGET-2 are computed using the full training set.

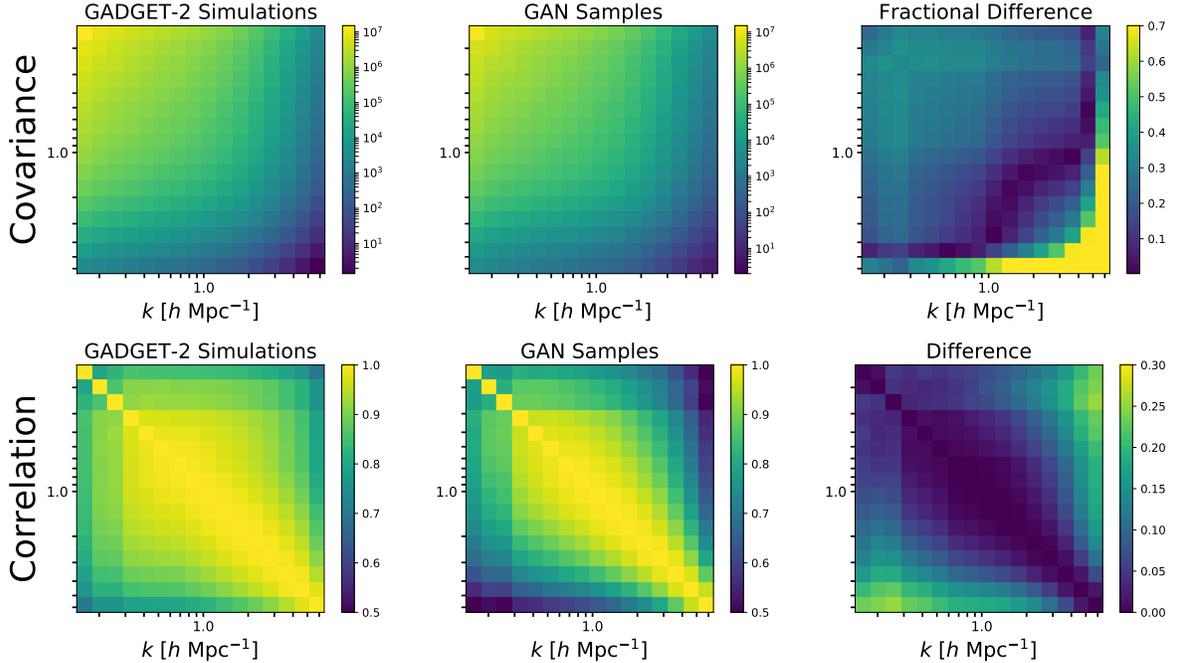


FIG. 10. The same as Figure 9, but for a GAN where the latent space prior is modeled with a Student-t distribution.

Figure 12 shows the extent to which our conditional GAN learns a non-linear mapping between redshift and the output density field. In the left panel, an initial power spectrum generated by the conditional GAN at $z = 3$ is

evolved linearly,

$$P(k, z) = \left(\frac{D(z)}{D(z=3)} \right)^2 P(k, z=3) \quad (11)$$

where $D(z)$ is the linear growth factor, while in the right

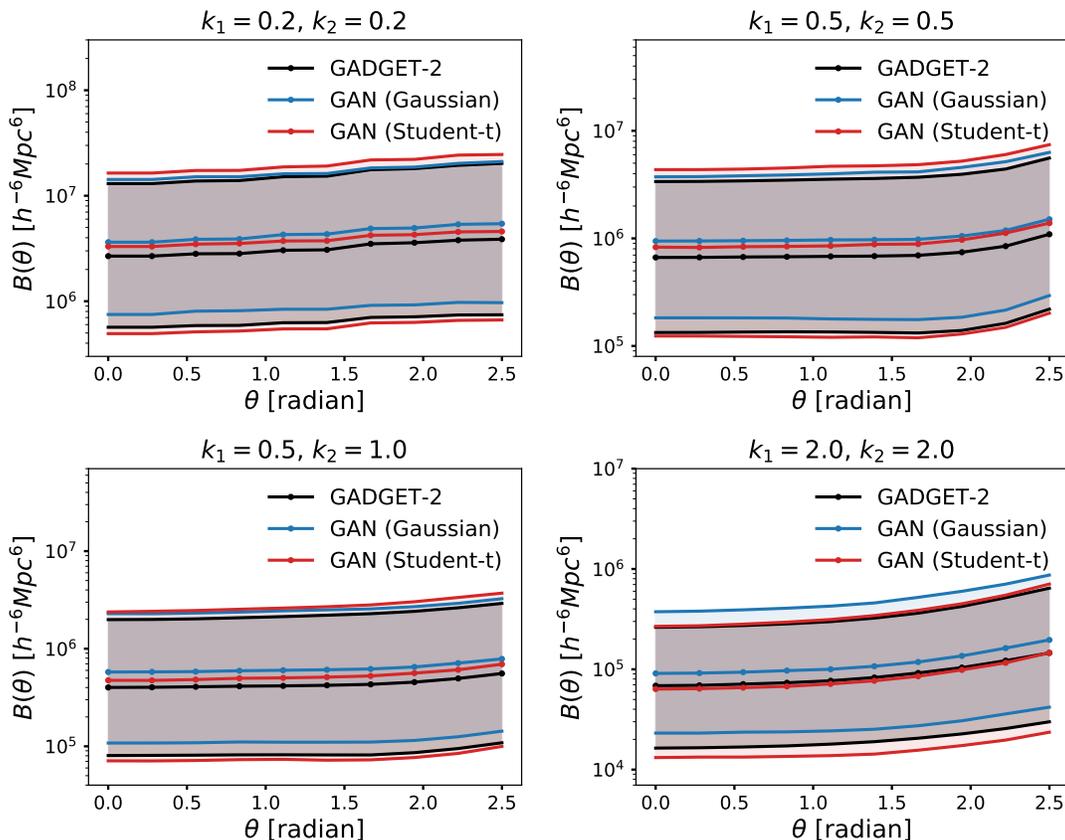


FIG. 11. Comparison of real (black) and synthetic (blue, red) matter bispectra for isosceles and other triangle configurations. Central lines indicate ensemble medians, while shaded regions bound the 16th and 84th percentiles for each set of bispectra. These percentiles are calculated from 2000 samples for each GAN model and from the full GADGET-2 data set.

panel the cGAN-interpolated matter power spectra are shown. We evaluate our model at both the aforementioned training redshifts and at intermediate redshifts $z = 2.0, 1.0$, and 0.25 . It is clear that while linear theory underestimates the growth of structure on small scales, the conditional GAN successfully captures the non-linear gravitational evolution. There is no obvious difference in cGAN performance evaluated at trained vs. interpolated redshifts, though interpolated power spectra at low redshift ($z < 1$) appear to be clustered around $z = 0.25$.

Figure 13 shows a sequence of cGAN-generated samples where the latent vector \mathbf{z} is fixed and the conditional redshift parameter is varied from $z = 2.8$ to the present day. Because we do not present consecutive redshift snapshots of the same volume to the discriminator and because we do not enforce mass conservation when varying the latent redshift parameter, we do not expect the cGAN to learn to simulate gravitational evolution per se, which would involve accurate modeling of the matter density displacement field. Nonetheless, our model learns a smooth mapping between redshift and the matter density field in real space – regions that are mildly clustered at high redshift appear to form coherent filaments as the conditional redshift is varied, with no obvious intermed-

iate discontinuities. This visual inspection helps establish the property that smooth paths in the latent space correspond to smooth paths in the data space, which reinforces the robustness of the compressed GAN representation.

V. DISCUSSION AND FUTURE WORK

In this work we have demonstrated that generative adversarial networks can learn a high-fidelity, compressed representation of the 3D dark matter density field. Once the generative model is trained, new samples can be generated in a fraction of a second – including the overhead from memory transfer between GPU and CPU, generating 512 samples (the volume equivalent of one N-body simulation) takes ~ 6 seconds. Furthermore, our trained generator only takes 17 Mb of storage, and storage of individual realizations amounts to saving latent vectors that can be passed through the generator when needed. We exploit the rotational invariance of our cosmological N-body simulations by using random rotations and flips to augment the training data set, which contributes to the stability of training. In fact, we are able to stably train our GANs with as few as 2048 samples (see Ap-

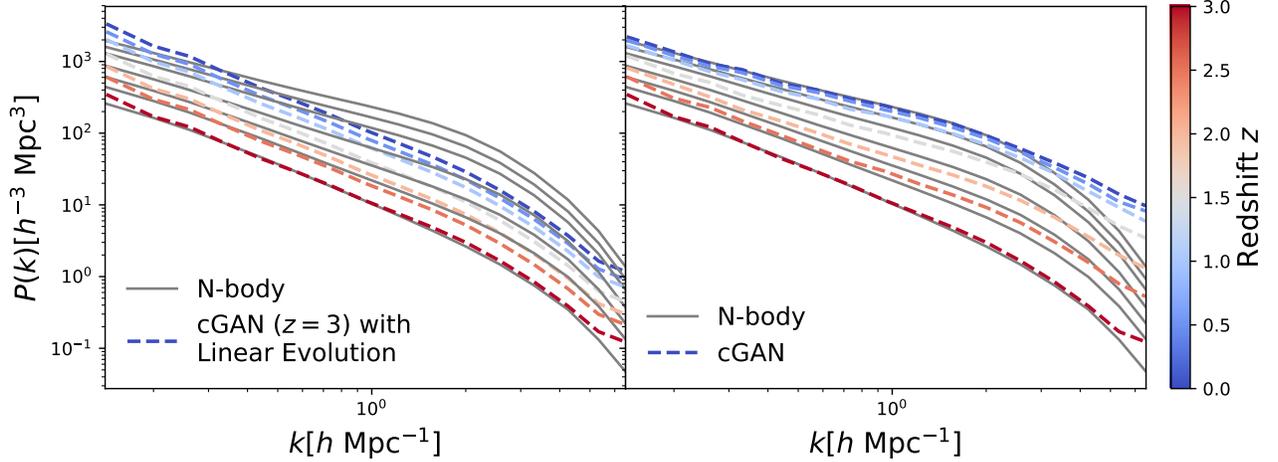


FIG. 12. **Left:** N-body power spectra (gray lines) compared with a median initial cGAN power spectrum at $z = 3$ that is evolved with linear theory (colored lines). **Right:** the same N-body power spectra, compared with cGAN-interpolated median power spectra. Power spectra in both panels are evaluated at redshifts $z = 3.0, 2.0, 1.5, 1.0, 0.5, 0.25$ and 0 , with color bar denoting the conditional input to the cGAN.

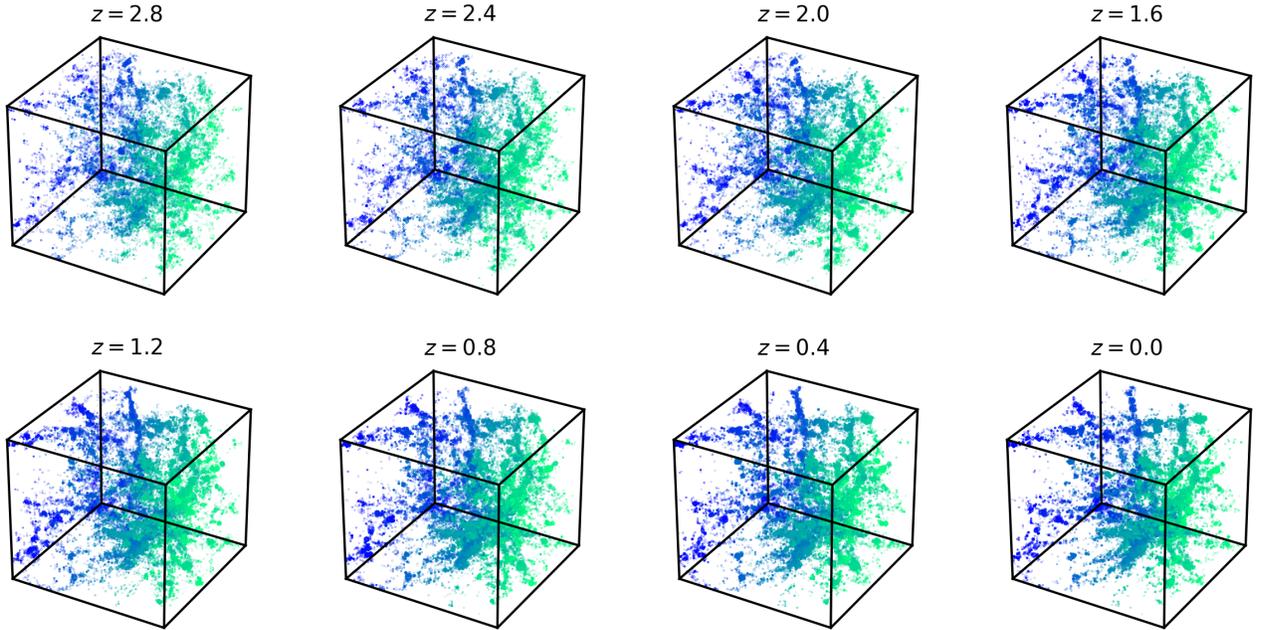


FIG. 13. 3D visualizations of cGAN-generated sub-volumes, with conditional redshift z linearly interpolated from $z = 2.8$ to $z = 0$. Sub-volumes are visualized using the same scheme as in Figure 4.

pendix C), equivalent to four $(512 \text{ Mpc})^3$ GADGET-2 simulations.

Synthetic samples were validated using one-, two- and three-point statistics that are accurate well into the non-linear regime where baryonic effects become important. This work presents the first estimates of power spectrum covariance matrices and bispectra for GANs sampling from a compressed representation of the data. While Figure 8 indicates some coherent structure in synthetic

samples on small scales, our models largely reproduce the power spectrum covariance structure of the matter density field when compared with N-body simulations.

This work also demonstrates that conditional GANs are capable of learning a smooth mapping between early and late times in the evolution of the matter density field, with fairly accurate interpolation as seen through the dark matter power spectrum. Future implementations might use a redshift-dependent data scaling that pre-

serves the information in linear fluctuations at high redshift and non-linear fluctuations at low redshift, though our attempts to do so were unsuccessful.

We also show that using a Student-t prior for the latent space results in a model that better captures sample variance across several summary statistics. While it is true that deep neural networks have high capacity, a given modeling task may be harder in the presence of distributional mismatch between the latent space and the data space. A similar point is argued in [42], which considers the role of dual data representations for classification. In that work, even when networks were given sufficient capacity to learn a dual transformation from position space to momentum space (where classification was straightforward) no trained model learned one. To our knowledge, this is the first application of heavy-tailed priors for modeling the latent space of a GAN. While we choose the Student-t distribution as our heavy-tailed prior, we remain agnostic to the “optimal” prior. Depending on the data set and purpose of the generative model, it may be possible to optimize the prior distribution non-parametrically as in [36]. We expect this modification of the standard GAN prior to be helpful for data sets with tails that impact the overall quality and diversity of the data distribution. While heavy-tailed priors are not necessarily the only remedy for issues addressed in this work, it was the only modification among several to the fiducial model and training procedure that led to models that were accurate and variance preserving. Furthermore, the change is very simple to implement – in our work, only one line of code needed to be changed.

Because our GANs learn a compressed representation of the data distribution, it may be possible to further decorrelate synthetic N-body ensembles by resampling the latent space of the generator after training. When one trains a GAN to model data with a lower-dimensional latent space, the underlying assumption is that a compressed representation of the data distribution exists on a manifold with dimension equal to that of the latent space. Mapping training samples to the posterior $P(\mathbf{z}|x_{data})$ is straightforward in models like variational autoencoders where the encoder network learns the mapping during training – for GANs this may be possible through a back propagation optimization scheme like in [43]. An estimate of the training set posterior could help determine which regions of the latent space are correlated. With an approximate posterior available, one might then train a normalizing flow to resample the latent space in a decorrelated manner, as has been done with variational autoencoders in [44].

A promising avenue in machine learning literature seeks to embed inductive biases into neural network models, such that they are guaranteed to respect the structures and symmetries of a system [45–47]. Such properties make generative models more robust and generalizable, and will be important for any applications of deep generative modeling to perform physical inference.

This work focuses on optimizing performance while us-

ing a fairly simple network architecture, but larger-scale applications may warrant more sophisticated schemes. If scaling up is done by stitching together sub-volumes, one needs to ensure continuity at the boundary of sub-volumes, and also ensure that spatial modes larger than the sub-volumes are properly modeled. One implementation uses super-resolution techniques and conditional neighbor information to generate large N-body volumes [26]. As described in Section IV, we experimented with simple extensions to our model, for example adding extra convolutional layers and expanding the latent space dimension, but did not see performance improvements. The issue of “going deeper” is known to be delicate in the GAN literature [31], requiring advanced regularization techniques that we do not attempt in this work, but it may be possible to optimize neural network architecture in a more detailed fashion to capture the proper correlation structures seen in N-body simulations. Once larger, high-fidelity volumes can be generated, it may be possible to further condition the generator on cosmological parameters, for example Ω_m and σ_8 .

ACKNOWLEDGMENTS

R.M.F is supported by the California Institute of Technology. P.B. was supported by Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Copyright 2020. All rights reserved. The authors would like to thank Tzu-Ching Chang, Olivier Doré, Michael Albergo, Jeremy Bernstein and Yun-Ting Cheng for useful discussions.

The authors acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing GPU cluster resources that have contributed to the research results reported within this paper. URL:<http://www.tacc.utexas.edu>

Appendix A: Exploring the connection between discriminative and cosmological bias

One challenge in using generative models to make predictions is that the statistics of generated samples are typically biased with respect to the data. For example, when generating cosmological volumes we never explicitly constrain the model to generate samples consistent up to the power spectrum, bispectrum, etc. Generative adversarial networks represent a framework in which an implicit likelihood is defined by a discriminator, which co-evolves with the generator – there is no tailored loss function based on summary statistics. Even if one were used, inefficient training and/or insufficient model capacity often prohibit neural networks from fully capturing a given target distribution. While the working assumption is that a trained generator will produce samples with a distribution equal to that of the true distribution, cer-

tain samples may have artifacts or excess power that make them outliers with respect to p_d , the data distribution. On the other hand, because sample generation with GANs is fast, we are well justified to explore bias reduction techniques that may degrade the nominal computational performance. Most bias reduction techniques in the machine learning literature cite improvement according to neural network based metrics such as the Frechet Inception Distance (FID) [48] and the Inception Score (IS) [49]. However, it is unclear whether bias reduction with respect to the discriminator corresponds to bias reduction for cosmological metrics.

One such bias reduction technique is discriminator rejection sampling (DRS) [50]. In standard rejection sampling, samples from a distribution X with density $f(X)$ are drawn using samples from another distribution Y with density $g(Y)$. This is done by computing $f(y)/Mg(y)$, where M is some finite bound on the likelihood ratio $f(X)/g(X)$. The idea behind DRS is that, rather than perform post-processing to correct individual samples, one can correct *ensemble based expectations* by training the discriminator network to learn the ratio $p_{data}(x)/p_g(x)$, where $p_{data}(x)$ and $p_g(x)$ are densities of the real and synthetic distributions, respectively. For a fixed generator, minimizing the loss of the discriminator defined by a sigmoid and trained with cross-entropy loss yields

$$D^*(x) = \frac{p_{data}(x)}{p_{data}(x) + p_g(x)} \quad (\text{A1})$$

where D^* is the optimal discriminator. One can also express the discriminator output in terms of the logit $\tilde{D}(x)$ (again for a sigmoid activation):

$$D^*(x) = \frac{1}{1 + e^{-\tilde{D}(x)}} = \frac{p_{data}(x)}{p_{data}(x) + p_g(x)} \quad (\text{A2})$$

Rearranging terms, the ratio $p_{data}(x)/p_g(x)$ can be expressed as

$$\frac{p_{data}(x)}{p_g(x)} = e^{\tilde{D}(x)}. \quad (\text{A3})$$

The likelihood ratio bound is set to $M = \max_x p_{data}(x)/p_g(x) = e^{\tilde{D}^*(x^*)}$. In practice we cannot compute the true maximum, so instead we estimate a lower bound using a large number of samples. Finally, the acceptance probability for rejection sampling can be written as $\alpha = e^{\tilde{D}^*(x) - \tilde{D}_M^*}$.

There are several caveats that degrade discriminator rejection sampling in practice. For one, the acceptance probability we estimate is only approximate to the true density ratio, since we cannot directly optimize over density functions. Furthermore, if the supports of p_d and p_g have a small intersection, then discriminator rejection sampling may yield vanishingly small acceptance probabilities. We counteract excessively small acceptance

probabilities by modifying the logit $F(x)$ of the final sigmoid activation that defines the acceptance probability:

$$\frac{1}{1 + e^{-F(x)}} = e^{\tilde{D}^*(x) - \tilde{D}_M^*}. \quad (\text{A4})$$

In the above equation, $F(x)$ can be computed from $\tilde{D}^*(x)$ and \tilde{D}_M^* , but in practice we include tunable parameters ϵ and γ :

$$\hat{F}(x) = \tilde{D}^*(x) - \tilde{D}_M^* - \log\left(1 - e^{\tilde{D}^*(x) - \tilde{D}_M^* - \epsilon}\right) - \gamma, \quad (\text{A5})$$

After training the generator, we train the discriminator for an additional five epochs (with the generator fixed) such that it minimizes its loss, which plateaus around 10^{-6} . Once this is done, we perform rejection sampling with hyperparameters $\epsilon = 0.01$ and $\gamma = 0.0$, which gives an acceptance rate of $\sim 3\%$. Larger values of γ lead to lower acceptance fractions, while more negative values yield the opposite. We calculate a lower bound for $\tilde{D}^*(x)$ from 2000 generated samples.

Figure 14 shows a comparison of synthetic power spectra from a heavy-tailed GAN, both with and without DRS. The choice of a relatively unbiased model serves as a null test for DRS, i.e. the accuracy of generated power spectra should remain the same or improve if DRS is effective. To the contrary, we find that the trained discriminator preferentially accepts samples with higher power on average, ironically biasing our model more. We experimented with different values of γ and found that when gamma was higher (i.e. when acceptance probabilities were lower), the accepted samples were more biased with respect to the mean power spectrum.

While DRS is evidently impractical for our purposes, it does provide a number of important insights to this work. For one, performing DRS demonstrates that our notion of cosmological bias (in one case, bias on the amplitude of the matter power spectrum) does not align with the more generic distributional ‘‘bias’’ as determined by the discriminator. DRS also gives us insight into the latent space of the generator. Figure 15 shows the ℓ_2 -norm distribution for a set of latent vectors drawn from the prior distribution, as well as latent vectors of samples accepted through DRS. One can see that DRS preferentially accepts samples with latent vectors with larger magnitudes on average, while the scale dependence of the re-sampled power spectra remains intact. This suggests a relationship between the width of the re-sampled latent distribution and the power amplitude of samples. While these results are specific to the models trained in this work, DRS provides an opportunity to reason about how bias of the data distribution in discriminator-based, implicit likelihood methods compare with more physical notions of bias.

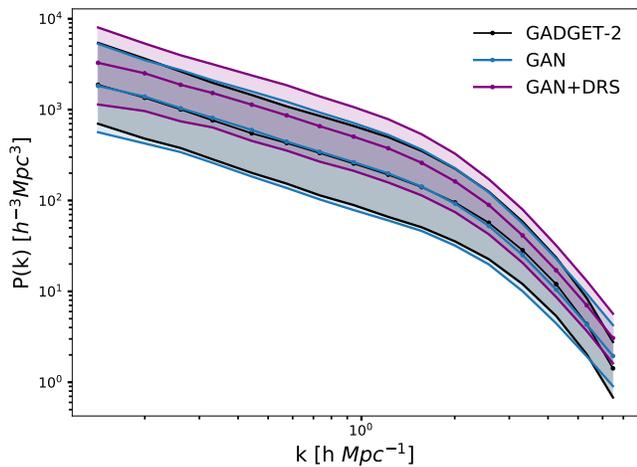


FIG. 14. Comparison of matter power spectra between GADGET-2 sub-volumes at $z = 0$ (black), samples drawn from a heavy-tailed GAN (blue) and samples obtained performing discriminator rejection sampling (DRS) on the heavy-tailed GAN (purple). Shaded regions bound the 16th and 84th percentiles of each ensemble and within each k bin. These percentiles are calculated based on 2000 samples for each model and from the full GADGET-2 data set.

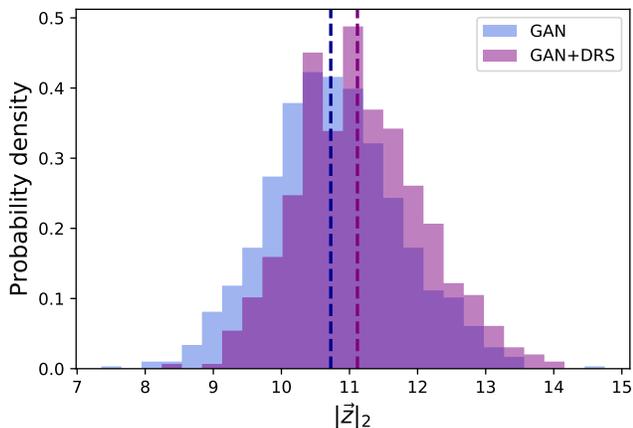


FIG. 15. Comparison of ℓ_2 -norm distribution for latent vectors sampled by a heavy-tailed GAN (blue) and resampled latent vectors obtained after performing discriminator rejection sampling (purple). Dashed vertical lines indicate the medians of the two distributions.

Appendix B: Scaling parameter

In Section III C, we advocate for the use of a data scaling that preserves high-density features that have the most significant impact on our summary statistics. We motivate this by observing that when the scaling parameter κ is small, the gradients in overdense regions (which can be very large in voxel space) get significantly suppressed in the scaled data. This is shown visually in Figure 2, which has (5) plotted for different choices of κ . Capturing gradients in the high-density regime ac-

curately with convolutions becomes more difficult, since small fluctuations in the final map may produce unrealistic values when scaled back to the original densities.

Another way to think about this problem is in terms of how a generic error gets propagated through the inverse data scaling. By computing the derivative of the inverse transformation, $d\rho/dc$, one can relate an error in c to error in ρ :

$$\delta\rho = \frac{2\kappa}{(1-c)^2} \delta c \quad (\text{B1})$$

For a fixed density ρ , the ratio of scaled values for different values of κ is

$$\frac{c_1}{c_2} = \left(\frac{\rho - \kappa_1}{\rho - \kappa_2} \right) \left(\frac{\rho + \kappa_2}{\rho + \kappa_1} \right) \quad (\text{B2})$$

For fixed ρ and fixed δc ,

$$\frac{\delta\rho_1(\kappa_1)}{\delta\rho_2(\kappa_2)} = \frac{\frac{2\kappa_1}{(1-c_1)^2} \delta c}{\frac{2\kappa_2}{(1-c_2)^2} \delta c} \quad (\text{B3})$$

$$= \frac{\kappa_1}{\kappa_2} \left(\frac{1-c_2}{1-c_1} \right)^2 \quad (\text{B4})$$

The last term in parentheses can be written as

$$\frac{1-c_2}{1-c_1} = \frac{\kappa_2}{\kappa_1} \left(\frac{\rho + \kappa_1}{\rho + \kappa_2} \right) \quad (\text{B5})$$

Finally,

$$\frac{\delta\rho_1(\kappa_1)}{\delta\rho_2(\kappa_2)} = \frac{\kappa_1}{\kappa_2} \left(\frac{\kappa_2}{\kappa_1} \right)^2 \left(\frac{\rho + \kappa_1}{\rho + \kappa_2} \right)^2 \quad (\text{B6})$$

$$= \left(\frac{\rho + \kappa_1}{\rho + \kappa_2} \right)^2 \frac{\kappa_2}{\kappa_1}. \quad (\text{B7})$$

As we are primarily concerned with the high-density regime, we can take the limit where $\rho \gg \kappa$ to conclude that $\delta\rho_1/\delta\rho_2 \approx \kappa_2/\kappa_1$. This calculation is ill-posed in the sense that we have not identified specific, quantifiable uncertainties of the GAN model. Nonetheless, it demonstrates in a fairly generic setting how data scalings affect models with finite precision in reference to some target distribution.

Figure 16 shows the effect of κ on the ability of our GANs to reproduce the power spectral distribution calculated from GADGET-2 N-body samples. While $\kappa = 45$ produces a model that outperforms others on nearly all scales, the improvement is most significant for $k \gtrsim 1$, where the power spectrum accuracy is over a factor of ten better than models trained with smaller values of κ . Beyond $\kappa = 45$, training was unstable and the generator was unable to learn a reliable representation of the data.

Appendix C: Training data set reduction effect on performance

As data-driven models, generative adversarial networks require a large amount of training data to learn

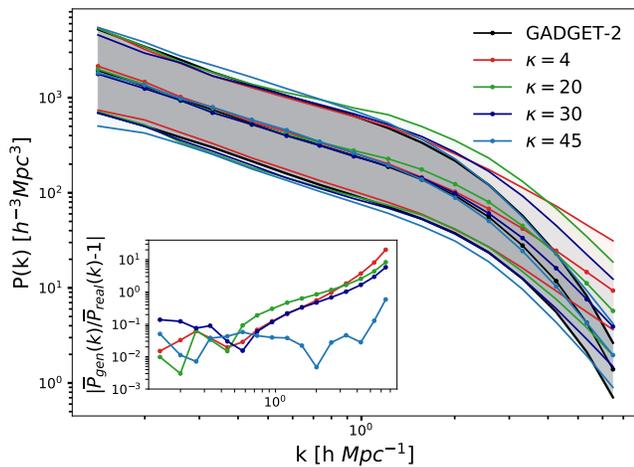


FIG. 16. Comparison of real and synthetic power spectra, varying the data scaling parameter κ from Equation (5) for each training run with a heavy-tailed GAN. Shaded regions bound the 16th and 84th percentiles of the ensemble within each k bin, which are calculated based on 2000 samples for each model and from the full GADGET-2 data set. Higher values of κ prioritize high-density features in the matter density field when the data are scaled to $[-1, 1]$.

the underlying data distribution. Much like network architectures and training hyperparameters, the size and diversity of the training data set are important factors in the stability of GAN training. For a non-standard training set like our own, it is unclear how many training samples are needed to successfully train our GAN models.

To understand the effect of training set size on performance, we train three configurations of the same model using one-half, one-fourth and one-eighth of the full training set. We fix all hyperparameters in the GAN across runs. While some runs on smaller training sets required restarting the training to achieve stable gradient updates in the first epoch, all models were able to train successfully. Figure 17 shows power spectra (top) and cross-spectra (bottom) generated from the trained models. The power spectra generated from all three models are accurate to within 10% up to $k = 2$, but none of them performs as well as with our full data set. There is no clear relationship between the training set size and power spectrum accuracy among our three configurations, however the model trained on the least amount of data ($n_{train} = 2048$) appears to have a larger cross-correlation on average between samples.

-
- [1] M. Vogelsberger, F. Marinacci, P. Torrey, and E. Puchwein, Cosmological simulations of galaxy formation, *Nature Reviews Physics* **2**, 42 (2020), arXiv:1909.07976 [astro-ph.GA].
- [2] L. Blot, P. S. Corasaniti, L. Amendola, and T. D. Kitching, Non-linear matter power spectrum covariance matrix errors and cosmological parameter uncertainties, *MNRAS* **458**, 4462 (2016), arXiv:1512.05383 [astro-ph.CO].
- [3] S. Tassev, M. Zaldarriaga, and D. J. Eisenstein, Solving large scale structure in ten easy steps with COLA, *J. Cosmology Astropart. Phys.* **2013**, 036 (2013), arXiv:1301.0322 [astro-ph.CO].
- [4] A. Izard, M. Crocce, and P. Fosalba, ICE-COLA: towards fast and accurate synthetic galaxy catalogues optimizing a quasi-N-body method, *MNRAS* **459**, 2327 (2016), arXiv:1509.04685 [astro-ph.CO].
- [5] F. Leclercq, B. Faure, G. Lavaux, B. D. Wandelt, A. H. Jaffe, A. F. Heavens, W. J. Percival, and C. Noûs, Perfectly parallel cosmological simulations using spatial comoving Lagrangian acceleration, arXiv e-prints, arXiv:2003.04925 (2020), arXiv:2003.04925 [astro-ph.CO].
- [6] Y. Feng, M.-Y. Chu, U. Seljak, and P. McDonald, FASTPM: a new scheme for fast simulations of dark matter and haloes, *MNRAS* **463**, 2273 (2016), arXiv:1603.00476 [astro-ph.CO].
- [7] P. Monaco, E. Sefusatti, S. Borgani, M. Crocce, P. Fosalba, R. K. Sheth, and T. Theuns, An accurate tool for the fast generation of dark matter halo catalogues, *MNRAS* **433**, 2389 (2013), arXiv:1305.1505 [astro-ph.CO].
- [8] M. Manera, R. Scoccimarro, W. J. Percival, L. Samushia, C. K. McBride, A. J. Ross, R. K. Sheth, M. White, B. A. Reid, A. G. Sánchez, R. de Putter, X. Xu, A. A. Berlind, J. Brinkmann, C. Maraston, B. Nichol, F. Montesano, N. Padmanabhan, R. A. Skibba, R. Tojeiro, and B. A. Weaver, The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: a large sample of mock galaxy catalogues, *MNRAS* **428**, 1036 (2013), arXiv:1203.6609 [astro-ph.CO].
- [9] F. S. Kitaura, G. Yepes, and F. Prada, Modelling baryon acoustic oscillations with perturbation theory and stochastic halo biasing., *MNRAS* **439**, L21 (2014), arXiv:1307.3285 [astro-ph.CO].
- [10] S. Avila, S. G. Murray, A. Knebe, C. Power, A. S. G. Robotham, and J. Garcia-Bellido, HALOGEN: a tool for fast generation of mock halo catalogues, *MNRAS* **450**, 1856 (2015), arXiv:1412.5228 [astro-ph.CO].
- [11] G. Stein, M. A. Alvarez, and J. R. Bond, The mass-Patch algorithm for fast generation of deep all-sky dark matter halo catalogues and its N-body validation, *MNRAS* **483**, 2236 (2019), arXiv:1810.07727 [astro-ph.CO].
- [12] P. Berger and G. Stein, A volumetric deep Convolutional Neural Network for simulation of mock dark matter halo catalogues, *MNRAS* **482**, 2861 (2019), arXiv:1805.04537 [astro-ph.CO].
- [13] S. He, Y. Li, Y. Feng, S. Ho, S. Ravanbakhsh, W. Chen, and B. Póczos, Learning to predict the cosmological structure formation, *Proceedings of the National Academy of Science* **116**, 13825 (2019), arXiv:1811.06533 [astro-ph.CO].
- [14] M. Bernardini, L. Mayer, D. Reed, and R. Feldmann, Predicting dark matter halo formation in N-body simulations with deep regression networks, arXiv e-prints, arXiv:1912.04299 (2019), arXiv:1912.04299 [astro-ph.CO].
- [15] C. Modi, Y. Feng, and U. Seljak, Cosmological recon-

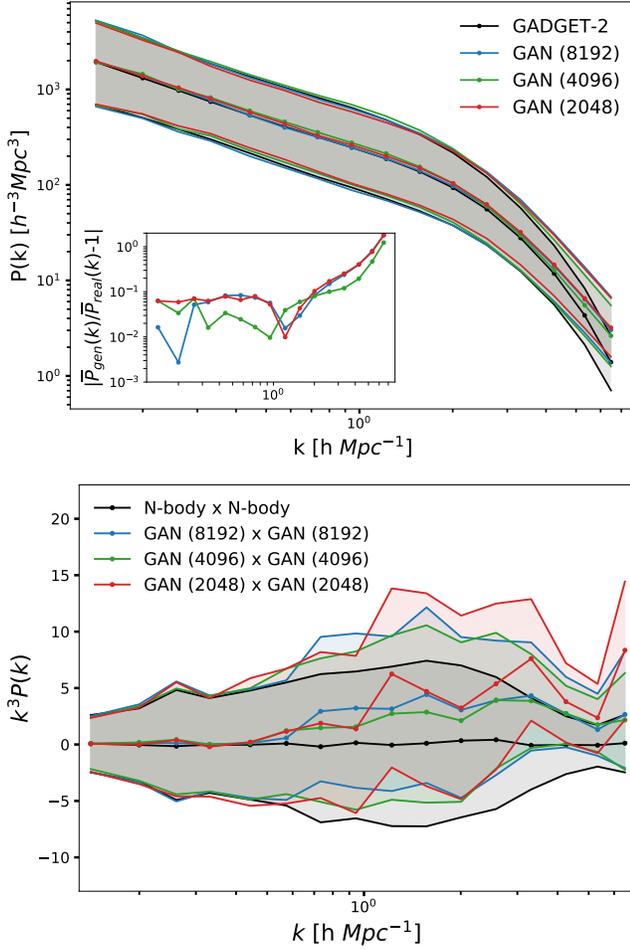


FIG. 17. Comparison of power spectra (top) and cross-spectra (bottom) from GADGET-2 and from GANs trained on one-half (8192 samples), one-fourth (4096 samples), and one-eighth (2048 samples) of the full training set. Shaded regions bound the 16th and 84th percentiles of the ensemble within each k bin, which are calculated based on 2000 samples for each model and from the full GADGET-2 data set.

struction from galaxy light: neural network based light-matter connection, *J. Cosmology Astropart. Phys.* **2018**, 028 (2018), arXiv:1805.02247 [astro-ph.CO].

- [16] E. Giusarma, M. Reyes Hurtado, F. Villaescusa-Navarro, S. He, S. Ho, and C. Hahn, Learning neutrino effects in Cosmology with Convolutional Neural Networks, arXiv e-prints , arXiv:1910.04255 (2019), arXiv:1910.04255 [astro-ph.CO].
- [17] F. List, I. Bhat, and G. F. Lewis, A black box for dark sector physics: predicting dark matter annihilation feedback with conditional GANs, *MNRAS* **490**, 3134 (2019), arXiv:1910.00291 [astro-ph.CO].
- [18] D. Kodi Ramanah, T. Charnock, and G. Lavaux, Painting halos from cosmic density fields of dark matter with physically motivated neural networks, *Phys. Rev. D* **100**, 043515 (2019), arXiv:1903.10524 [astro-ph.CO].
- [19] T. Tröster, C. Ferguson, J. Harnois-Déraps, and I. G. McCarthy, Painting with baryons: augmenting N-body simulations with gas using deep generative models, *MN-*

- RAS* **487**, L24 (2019), arXiv:1903.12173 [astro-ph.CO].
- [20] X. Zhang, Y. Wang, W. Zhang, Y. Sun, S. He, G. Contardo, F. Villaescusa-Navarro, and S. Ho, From Dark Matter to Galaxies with Convolutional Networks, arXiv e-prints , arXiv:1902.05965 (2019), arXiv:1902.05965 [astro-ph.CO].
- [21] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, Generative Adversarial Networks, arXiv e-prints , arXiv:1406.2661 (2014), arXiv:1406.2661 [stat.ML].
- [22] M. Mustafa, D. Bard, W. Bhimji, Z. Lukić, R. Al-Rfou, and J. M. Kratochvil, CosmoGAN: creating high-fidelity weak lensing convergence maps using Generative Adversarial Networks, *Computational Astrophysics and Cosmology* **6**, 1 (2019), arXiv:1706.02390 [astro-ph.IM].
- [23] A. C. Rodríguez, T. Kacprzak, A. Lucchi, A. Amara, R. Sgier, J. Fluri, T. Hofmann, and A. Réfrégier, Fast cosmic web simulations with generative adversarial networks, *Computational Astrophysics and Cosmology* **5**, 4 (2018), arXiv:1801.09070 [astro-ph.CO].
- [24] J. Zamudio-Fernandez, A. Okan, F. Villaescusa-Navarro, S. Bilaloglu, A. Derin Cengiz, S. He, L. Perreault Lévasseur, and S. Ho, HIGAN: Cosmic Neutral Hydrogen with Generative Adversarial Networks, arXiv e-prints , arXiv:1904.12846 (2019), arXiv:1904.12846 [astro-ph.CO].
- [25] D. Kodi Ramanah, T. Charnock, F. Villaescusa-Navarro, and B. D. Wandelt, Super-resolution emulator of cosmological simulations using deep physical models, arXiv e-prints , arXiv:2001.05519 (2020), arXiv:2001.05519 [astro-ph.CO].
- [26] N. Perraudin, A. Srivastava, A. Lucchi, T. Kacprzak, T. Hofmann, and A. Réfrégier, Cosmological N-body simulations: a challenge for scalable generative models, *Computational Astrophysics and Cosmology* **6**, 5 (2019), arXiv:1908.05519 [physics.comp-ph].
- [27] D. L. Ruderman, The statistics of natural images, *Network: Computation in Neural Systems* **5**, 517 (1994), <https://doi.org/10.1088/0954-898X.5.4.006>.
- [28] M. Mirza and S. Osindero, Conditional Generative Adversarial Nets, arXiv e-prints , arXiv:1411.1784 (2014), arXiv:1411.1784 [cs.LG].
- [29] Y. LeCun and C. Cortes, MNIST handwritten digit database, (2010).
- [30] T. Karras, S. Laine, and T. Aila, A Style-Based Generator Architecture for Generative Adversarial Networks, arXiv e-prints , arXiv:1812.04948 (2018), arXiv:1812.04948 [cs.NE].
- [31] A. Brock, J. Donahue, and K. Simonyan, Large Scale GAN Training for High Fidelity Natural Image Synthesis, arXiv e-prints , arXiv:1809.11096 (2018), arXiv:1809.11096 [cs.LG].
- [32] J. R. Bond, L. Kofman, and D. Pogosyan, How filaments of galaxies are woven into the cosmic web, *Nature* **380**, 603 (1996), arXiv:astro-ph/9512141 [astro-ph].
- [33] A. G. Sánchez, E. A. Kazin, F. Beutler, C.-H. Chuang, A. J. Cuesta, D. J. Eisenstein, M. Manera, F. Montesano, R. C. Nichol, N. Padmanabhan, W. Percival, F. Prada, A. J. Ross, D. J. Schlegel, J. Tinker, R. Tojeiro, D. H. Weinberg, X. Xu, J. Brinkmann, J. R. Brownstein, D. P. Schneider, and D. Thomas, The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological constraints from the full shape of the clustering wedges, *MNRAS* **433**, 1202

- (2013), arXiv:1303.4396 [astro-ph.CO].
- [34] J. Wu, C. Zhang, T. Xue, W. T. Freeman, and J. B. Tenenbaum, Learning a Probabilistic Latent Space of Object Shapes via 3D Generative-Adversarial Modeling, arXiv e-prints , arXiv:1610.07584 (2016), arXiv:1610.07584 [cs.CV].
- [35] A. Radford, L. Metz, and S. Chintala, Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks, arXiv e-prints , arXiv:1511.06434 (2015), arXiv:1511.06434 [cs.LG].
- [36] R. Singh, P. Turaga, S. Jayasuriya, R. Garg, and M. W. Braun, Non-Parametric Priors For Generative Adversarial Networks, arXiv e-prints , arXiv:1905.07061 (2019), arXiv:1905.07061 [cs.CV].
- [37] D. P. Kingma and J. Ba, Adam: A Method for Stochastic Optimization, arXiv e-prints , arXiv:1412.6980 (2014), arXiv:1412.6980 [cs.LG].
- [38] M. Arjovsky, S. Chintala, and L. Bottou, Wasserstein GAN, arXiv e-prints , arXiv:1701.07875 (2017), arXiv:1701.07875 [stat.ML].
- [39] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, Improved Training of Wasserstein GANs, arXiv e-prints , arXiv:1704.00028 (2017), arXiv:1704.00028 [cs.LG].
- [40] B. Pandey and S. Sarkar, Testing homogeneity in the Sloan Digital Sky Survey Data Release Twelve with Shannon entropy, MNRAS **454**, 2647 (2015), arXiv:1507.03124 [astro-ph.CO].
- [41] R. S. Gonçalves, G. C. Carvalho, J. Bengaly, C. A. P., J. C. Carvalho, A. Bernui, J. S. Alcaniz, and R. Maartens, Cosmic homogeneity: a spectroscopic and model-independent measurement, MNRAS **475**, L20 (2018), arXiv:1710.02496 [astro-ph.CO].
- [42] P. Betzler and S. Krippendorff, Connecting Dualities and Machine Learning, arXiv e-prints , arXiv:2002.05169 (2020), arXiv:2002.05169 [physics.comp-ph].
- [43] T. Schlegl, P. Seeböck, S. M. Waldstein, U. Schmidt-Erfurth, and G. Langs, Unsupervised Anomaly Detection with Generative Adversarial Networks to Guide Marker Discovery, arXiv e-prints , arXiv:1703.05921 (2017), arXiv:1703.05921 [cs.CV].
- [44] V. Böhm, F. Lanusse, and U. Seljak, Uncertainty Quantification with Generative Models, arXiv e-prints , arXiv:1910.10046 (2019), arXiv:1910.10046 [stat.ML].
- [45] C. Esteves, C. Allen-Blanchette, A. Makadia, and K. Daniilidis, Learning SO(3) Equivariant Representations with Spherical CNNs, arXiv e-prints , arXiv:1711.06721 (2017), arXiv:1711.06721 [cs.CV].
- [46] T. S. Cohen, M. Weiler, B. Kicanaoglu, and M. Welling, Gauge Equivariant Convolutional Networks and the Icosahedral CNN, arXiv e-prints , arXiv:1902.04615 (2019), arXiv:1902.04615 [cs.LG].
- [47] L. Falorsi, P. de Haan, T. R. Davidson, N. De Cao, M. Weiler, P. Forré, and T. S. Cohen, Explorations in Homeomorphic Variational Auto-Encoding, arXiv e-prints , arXiv:1807.04689 (2018), arXiv:1807.04689 [stat.ML].
- [48] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium, arXiv e-prints , arXiv:1706.08500 (2017), arXiv:1706.08500 [cs.LG].
- [49] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, Improved Techniques for Training GANs, arXiv e-prints , arXiv:1606.03498 (2016), arXiv:1606.03498 [cs.LG].
- [50] S. Azadi, C. Olsson, T. Darrell, I. Goodfellow, and A. Odena, Discriminator Rejection Sampling, arXiv e-prints , arXiv:1810.06758 (2018), arXiv:1810.06758 [stat.ML].