Qeios

1

Research Article

vGAN: A Generative Adversarial Emulator for Cosmic Web with Neutrinos

Neerav Kaushal¹, Elena Giusarma², Mauricio Reyes-Hurtado²

1. Institute of Computing and Cybersystems, Michigan Technological University, United States; 2. Department of Physics, Michigan Technological University, United States

Understanding the impact of neutrino masses on the evolution of Universe is a crucial aspect of modern cosmology. Due to their large free streaming lengths, neutrinos significantly influence the formation of cosmic structures at non-linear scales. To maximize the information yield from current and future galaxy surveys, it is essential to generate precise theoretical predictions of structure formation. One approach to achieve this is by running large sets of cosmological numerical simulations, which is a computationally intensive process. In this study, we propose a deep learningbased generative adversarial network (GAN) model to emulate the Universe for a variety of neutrino masses. Our model called ν GAN (for neutrino GAN) is able to generate 2D cosmic webs of the Universe for a number of neutrino masses ranging from 0.0 eV to 0.4 eV. The generated maps exhibit statistical independence, lack correlations with training data, and very closely resemble the distribution of matter in true maps. We assess the accuracy of our results both visually and through key statistics used in cosmology and computer vision analyses. Our results indicate that samples generated by ν GAN are accurate within a 5% error on power spectrum between k = 0.01 to $k = 0.5 h {
m Mpc}^{-1}$. Although this accuracy covers the mildly non-linear scales, consistent with other works and observations, achieving higher accuracy at fully non-linear scales requires more sophisticated models, such as diffusion models. Nevertheless, our work opens up new avenues for building emulators to generate fast and massive neutrino simulations, potentially revolutionizing cosmological predictions and analyses. This work serves as a proof-of-concept, paving the way for future extensions with higherresolution 3D data and advanced generative models.

Corresponding author: Neerav Kaushal, kaushal@mtu.edu

1. Introduction

Neutrinos are among the most abundant particles in the Universe, with number densities only slightly lower than those of photons. In the early Universe, neutrinos were relativistic, contributing to the radiation energy density during that epoch. Unlike photons, neutrinos have now become non-relativistic and possess rest mass, thereby contributing to the total matter density of the Universe. This indicates that relic neutrinos can have significant effects on cosmological observables, particularly influencing the background evolution, the spectra of matter perturbations, and the anisotropies in the Cosmic Microwave Background (CMB).

Next-generation large-scale structure (LSS) surveys, including Dark Energy Spectroscopic Instrument (DESI) [1], eROSITA [2], Euclid [3], Nancy Grace Roman Space Telescope (WFIRST) [4], and Rubin Observatory's Legacy Survey of Space and Time (LSST)^[5], will map vast cosmological volumes with high precision, producing datasets that span a wide range of redshifts and cosmic epochs. These surveys are expected to significantly improve constraints on cosmological parameters, including the sum of neutrino masses, thereby enhancing our understanding of the fundamental role neutrinos play in the evolution of the Universe. Recent analyses from DESI and the Atacama Cosmology Telescope (ACT) ^{[6][7]} have already refined cosmological constraints on $\sum m_{\nu}$. Current upper bounds from joint analyses of CMB and LSS data lie in the range $\sum m_{\nu} < 0.06-0.08$ eV at the 95% confidence level, assuming a minimal Λ CDM+ $\sum m_{\nu}$ framework.

Extracting the full scientific potential of these upcoming surveys requires theoretical predictions of comparable precision. In particular, accurate comparison between observations and theory demands rigorous modeling of the spatial distribution of matter and luminous tracers, especially in scenarios involving massive neutrinos. Analytic tools such as perturbation theory ^[8], can provide reliable predictions on quasi-linear scales, [e.g. ^{[9][10][11][12][13][14][15][16][17][18][19][20][21][22][23][24][25][26][27][28]], but most of the cosmological information resides in the fully non-linear regime, which requires more sophisticated approaches. In the absence of a comprehensive analytical model, numerical simulations with massive neutrinos are crucial for studying these non-linear scales and validating observations. In particular, N-body simulations are among the most crucial tools for evolving cosmological matter fluctuations under gravity alone, allowing for direct comparisons with theoretical predictions. These simulations are instrumental in generating mock galaxy catalogs, computing covariance matrices, and optimizing observational strategies. Recent developments in N-body simulations incorporating massive}

neutrinos have significantly advanced our understanding of their impact on clustering at fully non-linear scales in real space^{[29][30][31][32][33]}, the clustering and abundance of halos and cosmic voids^{[34][35][36]}, and matter clustering in real space^{[37][38]}. However, a significant drawback of cosmological N-body simulations is their high computational cost. A single simulation demands extensive computational resources and can take days or even weeks to complete. This computational bottleneck restricts the amount of information that can be extracted from observational data and tested against theoretical models. Consequently, there is a pressing need for faster methods to generate cosmological simulations that maintain both accuracy and reliability, thereby accelerating the process and enhancing our ability to extract and analyze cosmological information.

Over the past decade, deep learning has emerged as a powerful tool for emulating high-resolution numerical simulations and accelerating cosmological predictions. Convolutional neural networks (CNNs) have been widely used for parameter inference, weak lensing map generation, and super-resolution of N-body simulations [39][40][41][42][43][44][45][46][47][48][49][50][51][52]]. CNNs have also been successfully applied to model the impact of massive neutrinos on structure formation ^{[53][54]}, demonstrating their effectiveness in extracting cosmological information from complex datasets. These models can efficiently learn to emulate the evolution of cosmic structures from high-fidelity simulations, significantly reducing computational costs.

More recently, generative models such as Generative Adversarial Networks (GANs) [40][55][56][57][58][59], normalizing flows^{[60][61]}, and diffusion models ^{[62][63][64][52]} have demonstrated remarkable success in generating synthetic cosmological data that closely match high-fidelity simulations. While CNNs and generative models have been extensively applied to a range of cosmological tasks, GAN-based emulators specifically tailored for massive neutrino cosmologies remain relatively unexplored. Our work aims to fill this gap by developing a GAN framework conditioned on neutrino mass, enabling rapid generation of cosmic web realizations across different neutrino mass scenarios. This approach builds on the success of generative models for cosmic structure emulation while addressing the unique challenges posed by massive neutrino effects on non-linear structure formation.

In particular, we employ deep GANs to generate 2D cosmic web realizations of the Universe, conditioned on a range of neutrino masses. We assume a scenario of three degenerate neutrinos, where the electron, muon, and tau neutrino species share the same mass. The choice of GANs is motivated by their ability to effectively learn the complex probability distributions underlying the data, enabling the generation of new, random, statistically independent, and identically distributed samples after training on N-body simulations. These generated samples are uncorrelated with the training examples. Our model, named ν GAN, is specifically conditioned on neutrino masses, allowing it to produce dark matter cosmic webs for any given neutrino mass after training. This approach substantially reduces the computational burden associated with generating variable-mass neutrino simulations using traditional methods, as numerous new samples can be generated within seconds.

To validate our results, we employ various summary statistics, including the power spectrum, transfer function, pixel intensity histograms, peak statistics, and structural similarity tests.

This study serves as a proof of concept, demonstrating that generative models, such as GANs, are a feasible and effective approach to emulate cosmological simulations with massive neutrinos. Building on this foundation, future work will extend the framework to more advanced generative architectures, including diffusion models and normalizing flows, which promise improved accuracy, better coverage of the data distribution, and greater fidelity on non-linear scales.

This paper is organized as follows. Section 2 briefly introduces and discusses conditional GAN and the data we used for training our GAN model. Section 3 details the training process, model architecture, and hyperparameters while section 4 presents the quantitative results. Finally, we draw our conclusions in section 5.

2. Methods

2.1. Conditional GAN

A conditional Generative Adversarial Network or CGAN ^[65] is an extension of the traditional GAN framework, designed to generate data samples with specific characteristics or attributes. In a standard GAN, two neural networks, a generator and a discriminator, are employed. The generator produces synthetic data that mimics real data, while the discriminator distinguishes between real and generated data. These networks compete during training: the generator improves its ability to create realistic data, and the discriminator enhances its skill in differentiating real from fake data. Both networks undergo simultaneous training, and if the optimization process is executed effectively, the generator will acquire the ability to generate data that matches the data distribution. In a CGAN, an additional conditional element, often in the form of extra input data or labels, is introduced. This conditional information is used to guide the generation process, enabling the model to generate data samples with specific attributes or characteristics.

In this study, we condition both the generator (G) and the discriminator (D) on a parameter denoted as y, representing the neutrino mass. By doing so, we ensure that the generated data samples are specifically tailored to different neutrino masses.

The algorithm works as follows:

- 1. The input consists of a random noise vector, denoted as z, which can originate from various distributions such as the Gaussian distribution (typically sampled from a unit-normal distribution, N(0,1)), a uniform distribution, or other structured inputs.
- 2. The generator G, which is parameterized by a neural network, takes the latent vector z and an additional random variable y (serving as a conditioning factor), and produces the output G(z, y).
- 3. The discriminator D, also parameterized by a neural network, takes real data samples x and the synthetic samples generated by G, named G(z, y). It then provides scores for both, denoted as D(x) and D(G(z, y)) respectively. These scores represents the discriminator's confidence in whether the samples are real, i.e. originating from the actual data distribution $p_{data}(x)$. When scaled to the range [0, 1], this score can also be loosely interpreted as an implicit likelihood of the data given D, i.e., p(x|D).
- 4. The predictions made by D are compared to the actual, true labels, and a loss is computed, represented as L(D,G).
- 5. This loss is backpropagated through D and then through G, to update the weights and biases of both networks.
- 6. These steps, from 1 to 5, are iteratively repeated over multiple epochs while processing the entire dataset.

2.2. Data

In the following, we present a detailed description of the various steps involved in the generation and preprocessing of data:

1. We run N-body simulations using the COLA (COmoving Lagrangian Acceleration) approximation ^{[66][67]} with the MG-PICOLA code 1 ^[68]. COLA is an approximate simulation method designed to efficiently evolve large-scale structure by combining Lagrangian perturbation theory with N-body dynamics, achieving faster computation while preserving large-scale accuracy. We used MG-PICOLA instead of full N-body simulations because it provides a computationally efficient

compromise, retaining accuracy on large and mildly non-linear scales at a fraction of the computational cost. The simulations track the evolution of 1024^3 cold dark matter (CDM) particles in the presence of neutrinos with masses 0.0, 0.1, and 0.4 eV from a redshift of z = 9 to z = 0 in 50 timesteps in a simulation volume of size 500 Mpc/h. The cosmological parameters used for these simulations are $\Omega_M = 0.3175$, $\Omega_B = 0.0490$, $n_s = 0.9624$, $\sigma_8 = 0.8340$, H = 67.11, and $\sum m_{\nu} = (0.0, 0.1, 0.4)$ eV. We generate two realizations for each neutrino mass, resulting in a total of six realizations. Each simulation provides the 3D spatial coordinates of 1024^3 CDM particles.

- 2. Following Rodríguez et al.^[40], for each realization, we take this 3D cube of particle positions and divide the positions along the x-axis into 1000 equal segments. We then extract 2D slices of particle positions in the y-z plane and select 500 non-consecutive 2D slices. We repeat this process along the y and z axes, obtaining 1500 2D slices per realization, and a total of 9000 slices for the entire dataset (i.e., six realizations).
- 3. These 2D slices are then pixelized into 256×256 images, where each pixel value represents the number of particles within that pixel. This process effectively transforms the data into 2D grayscale images of size 256×256 . The images are then smoothed using a Gaussian filter with a standard deviation of 1, resulting in floating-point pixel values spanning a wide range of magnitudes, rather than discrete integers. As part of data augmentation, each image is randomly flipped horizontally, vertically, or both, each with a 25% probability; the remaining 25% of the time, the image is left unaltered.
- 4. The data is scaled to [-1,1], following $\frac{[40]}{}$, using the same scaling parameter a = 4 as in that work. This scaling improves model performance and ensures compatibility with the generator's final tanh activation. The original data (ρ) and the scaled data ($\rho(x)$) are related by the following transformation:

$$\rho(x) = \frac{2x}{x+a} - 1. \tag{1}$$

This transformation is analogous to the logarithmic function. As the cosmic web of the Universe spans a dynamic range of magnitudes, from nearly empty cosmic voids to super-massive galaxy clusters, this transformation enhances the contrast of filaments, galaxy sheets, and dark matter halos. The parameter a in equation [eqn:scaling] controls the median value of the images. The choice a = 4 balances compression of extreme densities while retaining sufficient contrast in intermediate regions, ensuring the transformed data are centered around zero and compatible with the generator's tanh activation, promoting stable learning.

3. Implementation

Learning the optimal parameters for the discriminator and generator networks, and effectively optimizing a GAN, are challenging tasks. To address these challenges, we employed a CGAN based on the Wasserstein GAN (WGAN) framework ^[69].

Unlike the discriminator in a conventional GAN, which outputs a probability value indicating the confidence in the realness of the sample, the discriminator in a WGAN assigns a score based on the Wasserstein distance (also known as the Earth Mover's distance) between the real and fake distributions. This distance measures the divergence between two probability distributions. Consequently, the discriminator in a WGAN is referred to as a critic.

Our WGAN model, named ν GAN, consists of a generator and a critic. The generator includes one linear layer and six transposed convolutional layers. The neutrino mass $\sum m_{\nu}$ is concatenated to the latent vector z of size 200 at the beginning and processed by a linear layer. This is followed by upsampling through six transposed convolutional layers with filter sizes of 5 and 3, and strides of 2 and 1, respectively. Each upsampling operation is followed by batch normalization and a ReLU activation, except for the final layer, which uses a tanh activation without batch normalization. The architecture of the generator is detailed in Table 1.

Layer	Operations	Filter	Dimension
z			${ m bs} imes 200$
h0	linear + identity		$\mathrm{bs} imes 512 imes 16 imes 16$
h1	deconv + BatchNorm + ReLU	5 imes 5	$\mathrm{bs} imes 256 imes 32 imes 32$
h2	deconv + BatchNorm + ReLU	5 imes 5	$\mathrm{bs} imes 128 imes 64 imes 64$
h3	deconv + BatchNorm + ReLU	3 imes 3	$\mathrm{bs} imes 128 imes 64 imes 64$
h4	deconv + BatchNorm + ReLU	5 imes 5	$\mathrm{bs} imes 64 imes 128 imes 128$
h5	deconv + BatchNorm + ReLU	3 imes 3	$\mathrm{bs} imes 64 imes 128 imes 128$
h6	deconv + Tanh	5 imes 5	$\mathrm{bs} imes 1 imes 256 imes 256$

Table 1. Generator architecture of our ν GAN. bs refers to the batch size.

The discriminator in our model is composed of four convolutional layers followed by one linear layer. Each convolutional layer uses a filter of size 5 and a stride of 2, doubling the number of channels and halving the feature size with each operation. These layers are followed by batch normalization and a leaky ReLU activation function with a parameter of 0.2. After the convolutional operations, the data is flattened, and the neutrino mass is concatenated to it. Finally, a linear layer is applied to produce the desired output. The detailed architecture of the discriminator is presented in Table 2.

Layer	Operations	Filter	Dimension
X			bs imes 1 imes 256 imes 256
h0	conv + BatchNorm + Leaky ReLU	5 imes 5	bs imes 64 imes 128 imes 128
h1	conv + BatchNorm + Leaky ReLU	5 imes 5	bs imes 128 imes 64 imes 64
h2	conv + BatchNorm + Leaky ReLU	5 imes 5	bs imes 256 imes 32 imes 32
h3	conv + BatchNorm + Leaky ReLU	5 imes 5	bs imes 512 imes 16 imes 16
h4	linear + identity		bs imes 1

 Table 2. Discriminator architecture of our µGAN. bs refers to the batch size.

The networks were trained until convergence, defined by a stable distance between the generated and real images, was achieved. The hyperparameters used for training are detailed in Table 3.

Hyperparameter	Value
Learning rate (G, D)	$(10^{-5},10^{-5})$
Batch size	16
Latent vector size	200
Latent vector distribution	Standard Normal
Optimizer	ADAM
Gradient Penalty	1000
eta_1,eta_2	0.5, 0.999
Augmentation	True
Epochs	300

Table 3. Hyperparameters used for model training.

4. Results

In this section, we evaluate the performance of our model through both qualitative and quantitative diagnostics. We quantitatively assess the results using a range of summary statistics: power spectrum and transfer function, pixel intensity histogram, pixel peak histogram, and the Multi-Scale Structural Similarity Index to evaluate mode collapse. We also present a visual comparison between the synthetic and real images.

4.1. Power Spectrum and Transfer Function

Figure 1 (top panels) presents the average 2D power spectra computed over 500 realizations, comparing MG-PICOLA simulations (black curve) with the ν GAN model (red curve) for different neutrino masses. The power spectrum analysis provides a quantitative assessment of how well ν GAN captures the statistical properties of large-scale structure formation across various scales. The results demonstrate that the density distribution from ν GAN closely matches the N-body simulations in both amplitude and shape across nearly all scales, particularly at large scales where the agreement is strongest.



Figure 1. Power spectrum and transfer function comparison. The top panel shows the average 2D power spectra of the N-body images (black curves) and the ones generated by *v*GAN (red curves) for various neutrino masses. The difference in power spectra is very small (within 5%) at linear and mildly non-linear scales.

To better quantify the agreement between the two datasets, the bottom panels of Figure 1 display the transfer function, which compares the relative clustering amplitude of ν GAN-generated maps and N-body simulation as a function of the wavenumber. 2 The transfer function should ideally be equal to one if the emulator perfectly reproduces the simulation data. Between k = 0.01 and k = 0.65, ν GAN achieves an accuracy within 5%. However, for larger k ($k > 0.7 h Mpc^{-1}$), deviations become more pronounced. The transfer function shows increasing fluctuations and variance at these smaller scales, indicating the model's limited ability to capture non-linear clustering due to the lower resolution of the MG-PICOLA training data.

Nevertheless, these preliminary results demonstrate that the ν GAN model outperforms the previous work of Giusarma et al.^[70] on mildly non-linear scales ($0.3 < k < 0.65 h \text{Mpc}^{-1}$), achieving an improved accuracy of 5% compared to the 7% reported in our previous study, despite being trained on lower-resolution MG-PICOLA simulations with 2D data. Given that our earlier work used high-resolution 3D simulations ^[9], this suggests that incorporating high-resolution 3D data in future work could further enhance the model's accuracy and reliability.

We note that the power spectrum and transfer function results are averaged over 500 realizations, providing a robust statistical estimate and reducing sample variance.

A key advantage of this approach is that, once trained, the model can generate new cosmic web realizations for different neutrino masses within seconds on a modern Graphics Processing Unit (GPU). This represents a substantial improvement over traditional N-body simulations, which require several hundred to thousands of core-hours per neutrino mass case. As a result, our method provides a gain of several orders of magnitude in computational efficiency, enabling rapid and scalable cosmological emulation.

4.2. Visual Comparison

The top two panels display 10 cosmic web images from N-body simulations, while the bottom two panels show images generated by ν GAN. Each bright spot in the images represents the average number of dark matter particles or the density contrast at that pixel location. It is important to note that the pixel values are scaled to the range of [-1,1].

Figure 2 presents 10 original images of the cosmic web from standard N-body simulations (top) and 10 synthesized samples generated by ν GAN without neutrinos (bottom) over a scale of 500 Mpc/h. It is evident that ν GAN effectively captures the prominent visual features of the data. The structures of filaments and halos are well reproduced, demonstrating the capability of GANs to accurately replicate the cosmic web. In particular, the cosmic web structures produced by the our model are visually indistinguishable from the originals.

To validate the parameterization of our model on neutrino masses, we generate images using the same latent vector space but with different neutrino masses. Figure 3 displays images from N-body simulations (top) and ν GAN (bottom) using various neutrino masses. The latent vector in ν GAN and the random seed in the simulations are kept constant to facilitate a direct visual comparison of the images under the same fixed conditions. Figure 3 appears visually very similar across different neutrino masses, reflecting that the impact of neutrino mass on small-scale clustering is too subtle to be visually discernible in 2D cosmic web images, even though the model encodes this dependence in its statistical structure.

N-body simulation samples



Figure 2. The top panel displays the cosmic webs from our simulations, while the bottom panel shows those generated by ν GAN. In the top images, the random seed was fixed during the simulations, and in the bottom images, the latent vector was fixed. As a result, the images in the top row appear similar to each other, and the same consistency is observed in the bottom row.

4.3. Pixel Intensity Histogram

We analyze pixel intensity histograms as complementary diagnostics for assessing ν GAN's ability to replicate cosmic structure. Specifically, we generate mass map histograms, which represent the distribution of mass densities in cosmic web images. Since pixel values in these images correspond to mass density, mass map histograms are also referred to as pixel intensity histograms $\frac{[71][72]}{1}$. By comparing these histograms for ν GAN-generated images and N-body simulations, we assess how

accurately our model reproduces the true mass distribution of the Universe. Figure 4 illustrates the distribution of mass map pixels (N_{pixels}) for both N-body and ν GAN-generated maps. The top panel shows the mass histograms, while the bottom panel presents the fractional difference between the predictions and the actual values. The comparison reveals a generally good agreement between the truth (black curve) and the ν GAN prediction (red curve). However, for higher pixel intensity values (between 0.25 and 0.8), ν GAN predicts a slightly lower number of pixels than the actual N-body samples. The bottom panel indicates that the largest deviations occur at these higher intensity values, with discrepancies of approximately 20% – 30%. While these deviations are moderate, their potential impact on cosmological parameter inference should be assessed in future work to determine whether additional corrections or improved training are necessary for precision analyses. Similar trends have been observed in previous works of ^[55] and ^[73].



Figure 4. A comparison of pixel intensity histogram of the samples generated from the N-body simulations and our _{*L*}GAN model. The curves are averaged over 500 samples. The major difference is at lower pixel intensity values.

4.4. Pixel Peak Histogram

While the power spectrum provides a complete statistical description of a Gaussian random field, the cosmic web exhibits significant non-Gaussian features that require higher-order statistics for a more detailed characterization. Traditional higher-order statistics, such as the bispectrum and three-point correlation function, provide valuable insights into non-linear structure formation but are computationally expensive to evaluate ^{[74][75][76][77][78][79]}. To address these challenges, an alternative approach known as "peak statistics" is widely used to extract non-Gaussian features from cosmic web data. Peak statistics have been particularly effective in analyzing weak lensing convergence maps and large-scale structure distributions ^{[72][80]}.

In this context, a peak is defined as a pixel in the mass map that has a greater intensity than all of its surrounding pixels. Specifically, a pixel is classified as a peak if its value exceeds that of all 24 neighboring pixels within a 5×5 region. By identifying and quantifying these peaks, we gain additional insight into the clustering of matter, as peaks often correspond to high-density regions such as galaxy clusters.

Figure 5 presents the distribution of peaks (denoted as N_{peaks}) in the mass map, capturing the local maxima across the dataset. To ensure statistical robustness, this process is applied to 500 independent simulated images, from which we compute the median histogram of peak counts along with the 16th and 84th percentile intervals, capturing the spread of peak values.

The top panel of Figure 5 displays the histograms of peak counts, while the bottom panel illustrates the fractional difference in peak distributions between the ν GAN-generated images and the N-body simulation results. The overall close agreement between the two datasets demonstrates that ν GAN accurately reproduces the number and distribution of peaks. However, small deviations at the highest peak values suggest that the model may slightly underestimate the frequency of the most massive structures, potentially indicating limitations in capturing the extreme non-linear regime of structure formation.



Figure 5. A comparison of pixel peaks. The solid lines represent the median histogram derived from 500 samples generated by *v*GAN and N-body simulations. The shaded areas around each line indicate the 16th and 84th percentiles of the distribution. Note that pixel values have been scaled to the range [-1, 1].

4.5. Multi-Scale Structural Similarity Index (MS-SSIM)

A common challenge with GANs is that, if poorly trained, the generator may repeatedly produce a limited subset of the data distribution, failing to capture the full data variance. This phenomenon, known as **mode collapse**, results in outputs that represent only a narrow range of the target data3. Our ν GAN avoids mode collapse by using a larger latent vector (z) of size 200 from a standard gaussian distribution. Evidence of this can be seen in Figure 2, where each generated image is distinctly different, indicating that the model captures a wide range of data features.

To further assess mode collapse, we employ the Multi-scale Structural Similarity Index (MS-SSIM), a widely used metric in image analysis. MS-SSIM helps detect mode collapse, especially relevant in cosmology, where summary statistics may still match true data distributions even when mode collapse is present, potentially concealing the issue.

The MS-SSIM score between two images ranges from 0 to 1, where 1 indicates identical images and 0 represents completely different images. As the mass maps are stochastic and only statistically similar, we are not interested in similarity between specific maps. Following Perraudin et al.^[56], we calculate the MS-SSIM between a large ensemble of 2,000 images from N-body simulations and generated by ν GAN. This approach allows us to evaluate overall similarity across an image set rather than focusing on pairwise comparisons. Table 4 reports MS-SSIM scores between N-body and ν GAN-generated maps for different neutrino masses. The scores range from 0.41 to 0.58, indicating moderate structural similarity between the generated and real samples. This level of similarity reflects that ν GAN does not replicate individual maps pixel-by-pixel (as expected for stochastic realizations), but rather captures the overall statistical structure of the cosmic web. Importantly, the MS-SSIM values remain sufficiently low to avoid mode collapse, confirming that the model generates diverse samples across the data distribution.

Neutrino mass (in eV)	MS-SSIM score
0.0	0.58
0.1	0.48
0.4	0.41

Table 4. MS-SSIM scores for *v*GAN for each neutrino mass

5. Conclusions

In this work, we demonstrated how GANs can be employed to model the dark matter cosmic web of the Universe, incorporating the effects of neutrinos. We developed a deep Wasserstein GAN, termed ν GAN, which conditions on neutrino masses to generate statistically independent and uncorrelated samples of the 2D cosmic web. By using a larger latent dimension, ν GAN effectively avoids mode collapse, as confirmed by the multi-scale structural similarity (MS-SSIM) scores between generated and real samples. Our model converges efficiently, producing samples that are visually indistinguishable from N-body simulation data.

To validate ν GAN's accuracy, we evaluated various cosmological and computer vision statistics, including the power spectrum, transfer function, pixel intensity histogram, peak statistics, and MS-SSIM scores, all of which show strong agreement between ν GAN-generated samples and N-body simulations up to mildly non-linear scales. This indicates that ν GAN produces genuinely novel data, rather than simply replicating the training set, as observed in previous studies^{[40][55][56][65][73][81]}.

A key advantage of this approach is that, once trained, the model can generate new cosmic web realizations for different neutrino masses within seconds on a modern Graphics Processing Unit (GPU). This represents a substantial improvement over traditional N-body simulations, which require several hundred to thousands of core-hours per neutrino mass case. As a result, our method provides a gain of several orders of magnitude in computational efficiency, enabling rapid and scalable cosmological emulation. Given the data demands of upcoming large-scale cosmological surveys, efficient methods like ν GAN will be crucial for generating theoretical predictions, facilitating advanced analyses using deep learning^[82] and sophisticated statistical methods^[83].

This work represents a proof-of-concept toward developing rapid, differentiable emulators of the neutrino cosmic web, demonstrating the feasibility and potential of GAN-based approaches for accelerating cosmological simulation pipelines. Beyond technical validation, the cosmological implications of our results warrant discussion. The ~5% accuracy achieved in the power spectrum on linear and mildly nonlinear scales suggests that ν GAN can already serve as a fast emulator for summary statistics relevant to upcoming surveys such as DESI, Euclid, and LSST, especially for large-scale analyses targeting BAO and growth rates. However, the observed deviations of ~20–30% in peak statistics at high intensities indicate that caution is required when using ν GAN outputs to infer cosmological parameters sensitive to small-scale nonlinearities, such as σ_8 or $\sum m_{\nu}$. Future work should quantify how these discrepancies propagate into parameter uncertainties via forward modeling pipelines or likelihood analyses. Nevertheless, this proof-of-concept highlights the potential of generative models to accelerate cosmological inference while retaining acceptable accuracy for many applications.

6. Future Works

While our results demonstrate that ν GAN effectively captures the large-scale statistical properties of the cosmic web, several avenues remain for improvement. One key limitation of the current model is its restriction to 2D projections of the cosmic web. Expanding the framework to generate full 3D density fields is a natural next step, enabling more direct comparisons with state-of-the-art numerical simulations and observational datasets. This will require adapting 3D generative architectures and optimizing their scalability for high-dimensional data.

Additionally, while GANs provide fast sample generation, they suffer from inherent drawbacks such as mode collapse and difficulties in learning the true probability distribution of the data. To address these issues, our ongoing research explores the integration of normalizing flows ^[84] and diffusion models ^[85]. Normalizing flows explicitly learn the data probability distribution and provide exact likelihood estimates, making them well-suited for cosmological emulation ^{[60][86][87]}. Diffusion models, on the other hand, have demonstrated superior performance in high-resolution image generation and improved mode coverage compared to ν GAN ^{[62][64][52]}. Implementing diffusion-based emulators for cosmology could further enhance the accuracy and reliability of deep generative methods in structure formation modeling.

Future work will focus on improving accuracy at non-linear scales by incorporating high-resolution 3D simulations and leveraging diffusion-based models, thereby extending the applicability of generative emulators to precision cosmology. While our approach shows promising results, it represents an initial step toward developing more sophisticated generative models, such as diffusion models and normalizing flows, which have the potential to further improve accuracy at non-linear scales.

Furthermore, given that the current approach is trained on low-resolution MG-PCOLA simulations, we plan to incorporate high-resolution N-body simulations such as QUIJOTE ^[9] to improve small-scale accuracy and reduce discrepancies in peak statistics. Extending ν GAN to generate simulations across a broader range of cosmological parameters, including baryonic effects, will further increase its applicability to next-generation surveys.

Acknowledgments

We thank Francisco Villaescusa–Navarro for valuable discussions. NK, EG, and MR acknowledge the IT department at Michigan Technological University for their assistance in managing the computing cluster.

Footnotes

¹<u>https://github.com/HAWinther/MG-PICOLA-PUBLIC</u>

² The transfer function T(k) is defined as the square root of the ratio of the power spectra of the density field predicted by the ν GAN and the target respectively, $T(k) = \sqrt{P_{\text{pred}}(k)/P_{\text{target}}(k)}$

³ For example, a GAN trained on a dataset of digits from 0 to 9 might only produce a few digits (e.g., 2, 4, and 9) if it suffers from mode collapse.

References

- 1. [^]Aghamousa, Amir et al. The DESI Experiment Part I: Science, Targeting, and Survey Design. 2016. Available from: https://arxiv.org/abs/1611.00036
- 2. [^]A. Merloni, P. Predehl, W. Becker, H. Böhringer, T. Boller, et al. (2012). eROSITA Science Book: Mapping the S tructure of the Energetic Universe. arXiv e-prints. :arXiv:1209.3114. doi:10.48550/arXiv.1209.3114
- 3. [△]Luca Amendola, Stephen Appleby, Anastasios Avgoustidis, David Bacon, Tessa Baker, et al. (2018). Cosmolo gy and fundamental physics with the Euclid satellite. Living Reviews in Relativity. 21(1):2. doi:10.1007/s41114 -017-0010-3
- 4. [^]D. Spergel, N. Gehrels, J. Breckinridge, M. Donahue, A. Dressler, et al. (2013). WFIRST-2.4: What Every Astron omer Should Know. arXiv e-prints. :arXiv:1305.5425. doi:10.48550/arXiv.1305.5425
- 5. [^]Mandelbaum R, Eifler T, Hložek R, et al. (2018). "The LSST Dark Energy Science Collaboration (DESC) Science Requirements Document." arXiv e-prints. arXiv:1809.01669. https://arxiv.org/abs/1809.01669.
- 6. [△]M. Abdul Karim, et al. DESI DR2 Results II: Measurements of Baryon Acoustic Oscillations and Cosmologic al Constraints. 2025. Available from: https://arxiv.org/abs/2503.14738
- 7. [^]Calabrese, Erminia and others. The Atacama Cosmology Telescope: DR6 Constraints on Extended Cosmol ogical Models. 2025. Available from: https://arxiv.org/abs/2503.14454
- 8. [^]F. Bernardeau, S. Colombi, E. Gaztañaga, R. Scoccimarro. (2002). Large-scale structure of the Universe and cosmological perturbation theory. 367(1-3):1–248. doi:10.1016/S0370-1573(02)00135-7
- 9. ^{a, b, c}Francisco Villaescusa-Navarro, ChangHoon Hahn, Elena Massara, Arka Banerjee, Ana Maria Delgado, et al. (2020). The Quijote Simulations. 250(1):2. doi:10.3847/1538-4365/ab9d82
- 10. ^ΔLado Samushia, Zachary Slepian, Francisco Villaescusa-Navarro. (2021). Information Content of Higher-O rder Galaxy Correlation Functions. arXiv e-prints. :arXiv:2102.01696. Available from: https://arxiv.org/abs/21 02.01696
- 11. [△]Davide Gualdi, Hector Gil-Marin, Licia Verde. (2021). Joint analysis of anisotropic power spectrum, bispectr um and trispectrum: application to N-body simulations. arXiv e-prints. :arXiv:2104.03976. Available from: ht tps://arxiv.org/abs/2104.03976

- 12. [△]Joseph Kuruvilla, Nabila Aghanim. (2021). Information content in mean pairwise velocity and mean relati ve velocity between pairs in a triplet. arXiv e-prints. :arXiv:2102.06709. Available from: https://arxiv.org/abs/ 2102.06709
- 13. [△]Adrian E. Bayer, Francisco Villaescusa-Navarro, Elena Massara, Jia Liu, David N. Spergel, et al. Detecting ne utrino mass by combining matter clustering, halos, and voids. 2021. Available from: https://arxiv.org/abs/21 02.05049
- 14. [^]Arka Banerjee, Emanuele Castorina, Francisco Villaescusa-Navarro, Travis Court, Matteo Viel. (2020). Wei ghing neutrinos with the halo environment. 2020(6):032. doi:10.1088/1475-7516/2020/06/032
- 15. [△]ChangHoon Hahn, Francisco Villaescusa-Navarro, Emanuele Castorina, Roman Scoccimarro. (2020). Cons training Mv with the bispectrum. Part I. Breaking parameter degeneracies. 2020(3):040. doi:10.1088/1475-75 16/2020/03/040
- 16. [△]Cora Uhlemann, Oliver Friedrich, Francisco Villaescusa-Navarro, Arka Banerjee, Sandrine Codis. (2020). Fi sher for complements: extracting cosmology and neutrino mass from the counts-in-cells PDF. 495(4):4006– 4027. doi:10.1093/mnras/staa1155
- 17. [△]Oliver Friedrich, Cora Uhlemann, Francisco Villaescusa-Navarro, Tobias Baldauf, Marc Manera, et al. (202
 0). Primordial non-Gaussianity without tails how to measure fNL with the bulk of the density PDF. 498(1):
 464–483. doi:10.1093/mnras/staa2160
- 18. [△]Elena Massara, Francisco Villaescusa-Navarro, Shirley Ho, Neal Dalal, David N. Spergel. (2021). Using the Marked Power Spectrum to Detect the Signature of Neutrinos in Large-Scale Structure. 126(1):011301. doi:10. 1103/PhysRevLett.126.011301
- 19. [△]Ji-Ping Dai, Licia Verde, Jun-Qing Xia. (2020). What can we learn by combining the skew spectrum and the power spectrum? 2020(8):007. doi:10.1088/1475-7516/2020/08/007
- 20. [^]E. Allys, T. Marchand, J.-F. Cardoso, F. Villaescusa-Navarro, S. Ho, et al. (2020). New interpretable statistics f or large-scale structure analysis and generation. 102(10):103506. doi:10.1103/PhysRevD.102.103506
- 21. [△]Arka Banerjee, Tom Abel. (2021). Nearest neighbour distributions: New statistical measures for cosmologic al clustering. 500(4):5479–5499. doi:10.1093/mnras/staa3604
- 22. [△]Arka Banerjee, Tom Abel. (2021). Cosmological cross-correlations and nearest neighbour distributions. 50
 4(2):2911–2923. doi:10.1093/mnras/stab961
- 23. [△]Davide Gualdi, Sergi Novell, Héctor Gil-Marín, Licia Verde. (2021). Matter trispectrum: theoretical modellin g and comparison to N-body simulations. 2021(1):015. doi:10.1088/1475-7516/2021/01/015

- 24. [△]Utkarsh Giri, Kendrick M. Smith. (2020). Exploring KSZ velocity reconstruction with N-body simulations a nd the halo model. arXiv e-prints. :arXiv:2010.07193. Available from: https://arxiv.org/abs/2010.07193
- 25. [△]Lucia F. de la Bella, Nicolas Tessore, Sarah Bridle. (2020). The unequal-time matter power spectrum: impa ct on weak lensing observables. arXiv e-prints. :arXiv:2011.06185. Available from: https://arxiv.org/abs/2011.0 6185
- 26. [△]ChangHoon Hahn, Francisco Villaescusa-Navarro. (2021). Constraining Mv with the bispectrum. Part II. T he information content of the galaxy bispectrum monopole. 2021(4):029. doi:10.1088/1475-7516/2021/04/02 9
- 27. [△]Georgios Valogiannis, Cora Dvorkin. (2021). Towards an Optimal Estimation of Cosmological Parameters with the Wavelet Scattering Transform. arXiv e-prints. :arXiv:2108.07821. Available from: https://arxiv.org/a bs/2108.07821
- 28. ^ΔJoseph Kuruvilla. (2021). Cosmology with the kinetic Sunyaev-Zeldovich effect: Independent of the optical depth and σ8. arXiv e-prints. :arXiv:2109.13938. Available from: https://arxiv.org/abs/2109.13938
- ^ASimeon Bird, Matteo Viel, Martin G. Haehnelt. (2012). Massive neutrinos and the non-linear matter power spectrum. Monthly Notices of the Royal Astronomical Society. 420(3):2551–2561. doi:10.1111/j.1365-2966.201 1.20222.x
- 30. [△]Julien Lesgourgues, Sergio Pastor. (2012). Neutrino mass from Cosmology. arXiv e-prints. :arXiv:1212.6154.
 Available from: https://arxiv.org/abs/1212.6154
- 31. [△]F. Villaescusa-Navarro, M. Vogelsberger, M. Viel, A. Loeb. (2013). Neutrino signatures on the high-transmiss ion regions of the Lyman α forest. 431:3670–3677. doi:10.1093/mnras/stt452
- 32. [△]F. Villaescusa-Navarro, F. Marulli, M. Viel, E. Branchini, E. Castorina, et al. (2014). Cosmology with massive neutrinos I: towards a realistic modeling of the relation between matter, haloes and galaxies. 3:011. doi:10.10 88/1475-7516/2014/03/011
- 33. [△]Marco Peloso, Massimo Pietroni, Matteo Viel, Francisco Villaescusa-Navarro. The effect of massive neutrin os on the BAO peak. arXiv 2015. doi:10.48550/ARXIV.1505.07477
- 34. [△]E. Castorina, E. Sefusatti, R. K. Sheth, F. Villaescusa-Navarro, M. Viel. (2014). Cosmology with massive neut rinos II: on the universality of the halo mass function and bias. 2:049. doi:10.1088/1475-7516/2014/02/049
- 35. [^]E. Castorina, C. Carbone, J. Bel, E. Sefusatti, K. Dolag. (2015). DEMNUni: the clustering of large-scale structu res in the presence of massive neutrinos. 7:043. doi:10.1088/1475-7516/2015/07/043
- 36. [△]Massara E, Villaescusa-Navarro F, Viel M, Sutter PM (2015). "Voids in massive neutrino cosmologies." JCAP.
 11:018. doi: 10.1088/1475-7516/2015/11/018.

- 37. [△]Francisco Villaescusa-Navarro, Arka Banerjee, Neal Dalal, Emanuele Castorina, Roman Scoccimarro, et al.
 (2018). The imprint of neutrinos on clustering in redshift space. The Astrophysical Journal. 861(1):53. doi:10.3
 847/1538-4357/aac6bf
- 38. [△]J. Bel, A. Pezzotta, C. Carbone, E. Sefusatti, L. Guzzo. (2019). Accurate fitting functions for peculiar velocity s pectra in standard and massive-neutrino cosmologies. Astronomy & Astrophysics. 622:A109. doi:10.1051/00 04-6361/201834513
- ^ANayantara Mudur, Carolina Cuesta-Lazaro, Douglas P. Finkbeiner. (2023). Cosmological Field Emulation a nd Parameter Inference with Diffusion Models. arXiv e-prints. :arXiv:2312.07534. doi:10.48550/arXiv.2312.07 534
- 40. ^{a, b, c, d, e}Andres C. Rodríguez, Tomasz Kacprzak, Aurelien Lucchi, Adam Amara, Raphaël Sgier, et al. (2018). Fast cosmic web simulations with generative adversarial networks. Computational Astrophysics and Cosm ology. 5(1):4. doi:10.1186/s40668-018-0026-4
- 41. [△]Siyu He, Yin Li, Yu Feng, Shirley Ho, Siamak Ravanbakhsh, et al. (2019). Learning to predict the cosmologic al structure formation. Proceedings of the National Academy of Science. 116(28):13825–13832. doi:10.1073/p nas.1821458116
- 42. [△]Renan Alves de Oliveira, Yin Li, Francisco Villaescusa-Navarro, Shirley Ho, David N. Spergel. (2020). Fast an d Accurate Non-Linear Predictions of Universes with Deep Learning. arXiv e-prints. :arXiv:2012.00240. Avail able from: https://arxiv.org/abs/2012.00240
- 43. [△]Neerav Kaushal, Francisco Villaescusa-Navarro, Elena Giusarma, Yin Li, Conner Hawry, et al. (2022). NEC
 OLA: Toward a Universal Field-level Cosmological Emulator. 930(2):115. doi:10.3847/1538-4357/ac5c4a
- 44. [△]Victoria Ono, Core Francisco Park, Nayantara Mudur, Yueying Ni, Carolina Cuesta-Lazaro, et al. (2024). De biasing with Diffusion: Probabilistic Reconstruction of Dark Matter Fields from Galaxies with CAMELS. 970 (2):174. doi:10.3847/1538-4357/ad5957
- 45. [△]Core Francisco Park, Victoria Ono, Nayantara Mudur, Yueying Ni, Carolina Cuesta-Lazaro. (2023). Probabil istic reconstruction of Dark Matter fields from biased tracers using diffusion models. arXiv e-prints. :arXiv:2 311.08558. doi:10.48550/arXiv.2311.08558
- 46. [△]Yin Li, Yueying Ni, Rupert A. C. Croft, Tiziana Di Matteo, Simeon Bird, et al. (2021). AI-assisted superresolut ion cosmological simulations. Proceedings of the National Academy of Sciences. 118(19). doi:10.1073/pnas.20 22038118
- 47. [^]Yueying Ni, Yin Li, Patrick Lachance, Rupert A. C. Croft, Tiziana Di Matteo, et al. (2021). AI-assisted superre solution cosmological simulations II. Halo substructures, velocities, and higher order statistics. Monthly N

otices of the Royal Astronomical Society. 507(1):1021–1033. doi:10.1093/mnras/stab2113

- 48. [^]Xiaowen Zhang, Patrick Lachance, Yueying Ni, Yin Li, Rupert A. C. Croft, et al. (2024). AI-assisted super-res olution cosmological simulations III: time evolution. 528(1):281–293. doi:10.1093/mnras/stad3940
- 49. [^]Xiaowen Zhang, Patrick Lachance, Ankita Dasgupta, Rupert A. C. Croft, Tiziana Di Matteo, et al. (2024). AI -assisted super-resolution cosmological simulations IV: An emulator for deterministic realizations. arXiv eprints. :arXiv:2408.09051. doi:10.48550/arXiv.2408.09051
- 50. [^]Andreas Schanz, Florian List, Oliver Hahn. (2023). Stochastic Super-resolution of Cosmological Simulation s with Denoising Diffusion Models. arXiv e-prints. :arXiv:2310.06929. doi:10.48550/arXiv.2310.06929
- 51. [△]Adam Rouhiainen, Moritz Münchmeyer, Gary Shiu, Michael Gira, Kangwook Lee. (2024). Superresolution emulation of large cosmological fields with a 3D conditional diffusion model. 109(12):123536. doi:10.1103/Ph ysRevD.109.123536
- 52. ^{a, b, c}Mudur N, Finkbeiner DP (2022). "Can denoising diffusion probabilistic models generate realistic astrop hysical fields?" In: 36th Conference on Neural Information Processing Systems: Workshop on Machine Lear ning and the Physical Sciences. https://arxiv.org/abs/2211.12444.
- 53. [△]Elena Giusarma, Mauricio Reyes, Francisco Villaescusa-Navarro, Siyu He, Shirley Ho. (2019). Predicting cos mological massive neutrino simulation with convolutional neural networks. In: 33th conference on Neural I nformation Processing Systems. Available from: https://arxiv.org/abs/1910.04255v2
- 54. ^ΔElena Giusarma, Mauricio Reyes, Francisco Villaescusa-Navarro, Siyu He, Shirley Ho, et al. (2023). Learnin g neutrino effects in cosmology with convolutional neural network. The Astrophysical Journal. 950(1):70. do i:10.3847/1538-4357/accd61
- 55. ^{a, b, c}Andrius Tamosiunas, Hans A. Winther, Kazuya Koyama, David J. Bacon, Robert C. Nichol, et al. (2021). I nvestigating cosmological GAN emulators using latent space interpolation. 506(2):3049–3067. doi:10.1093/ mnras/stab1879
- 56. ^{a, b, c}Nathanaël Perraudin, Sandro Marcon, Aurelien Lucchi, Tomasz Kacprzak. (2020). Emulation of cosmol ogical mass maps with conditional generative adversarial networks. arXiv e-prints. :arXiv:2004.08139. Avai lable from: https://arxiv.org/abs/2004.08139
- 57. [△]Sambatra Andrianomena, Francisco Villaescusa-Navarro, Sultan Hassan. (2022). Emulating cosmological multifields with generative adversarial networks. In: 36th Conference on Neural Information Processing Sys tems: Workshop on Machine Learning and the Physical Sciences. Available from: https://arxiv.org/abs/2211. 05000

- 58. [△]Prabh Bhambra, Benjamin Joachimi, Ofer Lahav, Davide Piras. psi-GAN: A power-spectrum-informed gen erative adversarial network for the emulation of large-scale structure maps across cosmologies and redshif ts. Monthly Notices of the Royal Astronomical Society.: Oxford University Press (OUP) 2024. pp. 3138–3157. d oi:10.1093/mnras/stae2810
- 59. [△]Davide Piras, Benjamin Joachimi, Francisco Villaescusa-Navarro. Fast and realistic large-scale structure fr om machine-learning-augmented random field simulations. Monthly Notices of the Royal Astronomical So ciety.: Oxford University Press (OUP) 2023. pp. 668–683. doi:10.1093/mnras/stad052
- 60. ^{a, b}Roy Friedman, Sultan Hassan. (2022). HIGlow: Conditional Normalizing Flows for High-Fidelity HI Map Modeling. arXiv e-prints. :arXiv:2211.12724. doi:10.48550/arXiv.2211.12724
- 61. [△]Christopher C. Lovell, Stephen M. Wilkins, Peter A. Thomas, Matthieu Schaller, Carlton M. Baugh, et al. (20
 21). A machine learning approach to mapping baryons on to dark matter haloes using the EAGLE and C-EA
 GLE simulations. doi:10.1093/mnras/stab3221
- 62. ^{a, b}Nayantara Mudur, Carolina Cuesta-Lázaro, Douglas P. Finkbeiner. (2023). Cosmological field emulation and parameter inference with diffusion models. ArXiv. abs/2312.07534. Available from: https://api.semantics cholar.org/CorpusID:266174439
- 63. [△]Nayantara Mudur, Carolina Cuesta-Lazaro, Douglas P. Finkbeiner. (2024). Diffusion-HMC: Parameter infer ence with diffusion-model-driven hamiltonian monte carlo. The Astrophysical Journal. 978(1):64. doi:10.384 7/1538-4357/ad8bc3
- 64. ^{a, b}Tanner Sether, Elena Giusarma, Mauricio Reyes-Hurtado. (2024). Probabilistic Galaxy Field Generation with Diffusion Models. In: 38th conference on Neural Information Processing Systems. Available from: http s://arxiv.org/abs/2412.05131
- 65. ^{a, b}Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, et al. Generative adve rsarial networks. arXiv 2014. doi:10.48550/ARXIV.1406.2661
- 66. [^]Svetlin Tassev, Matias Zaldarriaga, Daniel J. Eisenstein. (2013). Solving large scale structure in ten easy ste ps with COLA. Journal of Cosmology and Astroparticle Physics. 2013(06):036–036. doi:10.1088/1475-7516/20 13/06/036
- 67. [^]Svetlin Tassev, Daniel J. Eisenstein, Benjamin D. Wandelt, Matias Zaldarriaga. sCOLA: The n-body COLA m ethod extended to the spatial domain. 2015. Available from: https://arxiv.org/abs/1502.07751
- 68. [^]Bill S. Wright, Hans A. Winther, Kazuya Koyama. (2017). COLA with massive neutrinos. 2017(10):054. doi:1 0.1088/1475-7516/2017/10/054

- 69. [≜]Martin Arjovsky, Soumith Chintala, Léon Bottou. Wasserstein GAN. arXiv 2017. doi:10.48550/ARXIV.1701.07 875
- 70. [△]Elena Giusarma, Mauricio Reyes, Francisco Villaescusa-Navarro, Siyu He, Shirley Ho, et al. (2023). Learnin g Neutrino Effects in Cosmology with Convolutional Neural Network. 950(1):70. doi:10.3847/1538-4357/accd
 61
- 71. [△]M. Gatti, C. Chang, O. Friedrich, B. Jain, D. Bacon, et al. (2020). Dark Energy Survey Year 3 results: cosmolog y with moments of weak lensing mass maps - validation on simulations. 498(3):4060–4087. doi:10.1093/m nras/staa2680
- 72. ^{a, b}T. Kacprzak, D. Kirk, O. Friedrich, A. Amara, A. Refregier, et al. (2016). Cosmology constraints from shear p eak statistics in Dark Energy Survey Science Verification data. 463(4):3653–3673. doi:10.1093/mnras/stw207
- 73. ^{a, b}Mustafa Mustafa, Deborah Bard, Wahid Bhimji, Zarija Lukić, Rami Al-Rfou, et al. (2019). CosmoGAN: cre ating high-fidelity weak lensing convergence maps using Generative Adversarial Networks. Computational Astrophysics and Cosmology. 6(1):1. doi:10.1186/s40668-019-0029-9
- 74. [△]Zachary Slepian, Daniel J. Eisenstein. (2016). Accelerating the two-point and three-point galaxy correlatio n functions using Fourier transforms. 455(1):L31–L35. doi:10.1093/mnrasl/slv133
- 75. [▲]Zachary Slepian, Daniel J. Eisenstein. (2015). Computing the three-point correlation function of galaxies in O(N2) time. 454(4):4142–4158. doi:10.1093/mnras/stv2119
- 76. [^]Zachary Slepian, Daniel J. Eisenstein. (2018). A practical computational method for the anisotropic redshift -space three-point correlation function. 478(2):1468–1483. doi:10.1093/mnras/sty1063
- 77. [△]Oliver H. E. Philcox, Zachary Slepian, Jiamin Hou, Craig Warner, Robert N. Cahn, et al. (2022). ENCORE: an
 O (Nq2) estimator for galaxy N-point correlation functions. 509(2):2457–2481. doi:10.1093/mnras/stab3025
- 78. [△]Oliver H. E. Philcox, Zachary Slepian. (2022). Efficient computation of N-point correlation functions in D di mensions. Proceedings of the National Academy of Science. 119(33):e2111366119. doi:10.1073/pnas.2111366119
- 79. [△]Lado Samushia, Zachary Slepian, Francisco Villaescusa-Navarro. (2021). Information content of higher or der galaxy correlation functions. 505(1):628–641. doi:10.1093/mnras/stab1199
- 80. [△]Nicolas Martinet, Peter Schneider, Hendrik Hildebrandt, HuanYuan Shan, Marika Asgari, et al. (2017). KiDS -450: Cosmological constraints from weak-lensing peak statistics II: Inference from shear peaks using n-bo dy simulations. Monthly Notices of the Royal Astronomical Society. 474(1):712–730. doi:10.1093/mnras/stx27 93

- 81. [^]Richard M. Feder, Philippe Berger, George Stein. (2020). Nonlinear 3D cosmic web simulation with heavy-t ailed generative adversarial networks. Physical Review D. 102(10). doi:10.1103/physrevd.102.103504
- 82. [^]Jorit Schmelzle, Aurelien Lucchi, Tomasz Kacprzak, Adam Amara, Raphael Sgier, et al. Cosmological mode l discrimination with deep learning. arXiv 2017. doi:10.48550/ARXIV.1707.05167
- 83. [△]Andrea Petri, Zoltán Haiman, Lam Hui, Morgan May, Jan M. Kratochvil. (2013). Cosmology with minkows ki functionals and moments of the weak lensing convergence field. Phys Rev D. 88:123002. doi:10.1103/Phys RevD.88.123002
- 84. [△]George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, Balaji Lakshminarayan an. (2019). Normalizing Flows for Probabilistic Modeling and Inference. arXiv e-prints. :arXiv:1912.02762. do i:10.48550/arXiv.1912.02762
- 85. [^]Jonathan Ho, Ajay Jain, Pieter Abbeel. (2020). Denoising Diffusion Probabilistic Models. arXiv e-prints. :arX iv:2006.11239. doi:10.48550/arXiv.2006.11239
- 86. [^]Sultan Hassan, Sambatra Andrianomena, Caitlin Doughty. (2020). Constraining the astrophysics and cos mology from 21 cm tomography using deep learning with the SKA. 494(4):5761–5774. doi:10.1093/mnras/sta a1151
- 87. [△]Christopher C. Lovell, Sultan Hassan, Francisco Villaescusa-Navarro, Shy Genel, ChangHoon Hahn, et al. (2
 023). A Hierarchy of Normalizing Flows for Modelling the Galaxy-Halo Relationship. In: Machine learning f or astrophysics. p. 21. doi:10.48550/arXiv.2307.06967

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.