SF2A 2025
A. Siebert, K. Baillié, M. Béthermin, F. Cantalloube, E. Josselin, N. Lagarde, J. Malzac, J. Richard, L. Selliez, O. Venot (eds)

PROBING COSMIC WEB STRUCTURES WITH GENERATIVE MODELS

M. Noor[1], T. Bonnaire[1,2], A. Decelle[3] and N. Aghanim[1]

Abstract. The cosmic web, consisting of an intricate network of voids, filaments, walls, and clusters, contains key information about cosmological parameters and the evolution of large-scale structures in the universe. In this context, very large ensembles of cosmological simulations are required to analyse upcoming galaxy survey data. To circumvent the associated computational cost, we aim at constructing a generative diffusion model to emulate the cosmic web while preserving its statistical properties. In this context, we explore how self-attention maps, which are learnt neural network features embedded within the diffusion model, can offer a new way to probe the underlying structure of the cosmic web. In addition to emulating high-fidelity simulations, our approach could also provide a quantitative assessment of the performance of generative models.

Keywords: Cosmic web, Large-scale structure, Cosmological simulations, Diffusion models, Self-attention maps

1 Introduction

The distribution of matter in the universe arises from the anisotropic gravitational collapse of primordial fluctuations. The resulting large-scale structure is characterised by a filamentary web-like pattern known as the cosmic web (Bond et al. 1996), composed primarily of four main components: nodes, filaments, walls and voids. This structure, revealed in large galaxy surveys (Geller & Huchra 1989; Tegmark et al. 2004; Colless et al. 2003; Huchra et al. 2005) and reproduced in numerical simulations (Doroshkevich et al. 1980; Klypin & Shandarin 1983; Jenkins et al. 1998; Springel et al. 2005), encodes key information about cosmological parameters and structure formation in the linear and nonlinear regimes.

Extracting this information from large galaxy surveys like Euclid[*] requires vast ensembles of numerical simulations, motivating the use of generative models to bypass the computational cost. To this end, we present a generative diffusion model that emulates the cosmic web with high visual and statistical fidelity.

Generative diffusion models (Sohl-Dickstein et al. 2015; Ho et al. 2020) often incorporate self-attention maps (Vaswani et al. 2023), which are learnt neural network features capturing important semantic information across multiple scales. To better understand how diffusion models learn the matter distribution in cosmological simulations, we perform a quantitative analysis of their attention maps, evaluating how effectively they capture different cosmic web environments. We adopt the T-Web algorithm (Hahn et al. 2007) for classifying dark matter density fields into distinct components, acting as a reference, which we use to assess the structural information captured by attention maps.

2 Diffusion model for cosmic web emulation

We utilise diffusion models to emulate the matter distribution in the Quijote suite (Villaescusa-Navarro et al. 2020), a large set of N-body dark matter simulations spanning thousands of cosmological models across a broad

[1] Université Paris-Saclay, CNRS, Institut d'Astrophysique Spatiale, 91405, Orsay, France

[2] Laboratoire de Physique de l'École normale supérieure, ENS, Université PSL, CNRS, Sorbonne Université, Université Paris Cité, F-75005 Paris, France

[3] Escuela Técnica Superior de Ingenieros Industriales, Universidad Politécnica de Madrid, Calle de José Gutiérrez Abascal 2, Madrid 28006, Spain

[*] https://www.esa.int/Science_Exploration/Space_Science/Euclid

range of parameters that govern structure formation. These models learn the underlying data distribution by means of a denoising process that stochastically transports random noise toward the empirical data distribution. This denoising process is learnt by a neural network, allowing the model to transform random noise into high-fidelity samples of dark matter density fields after training. In our work, we focus on a subset of $100,000$ 32×32 log-transformed density fields extracted from the Quijote simulations, taken from the fiducial cosmology set produced using the Planck Collaboration et al. (2020) cosmological parameters. We incorporate self-attention layers at multiple spatial resolutions throughout the neural network architecture, allowing the model to capture cosmic web features across a wide range of scales.

Figure 1 (left) shows a visual comparison between the density fields from a Quijote test dataset and the trained diffusion model. Evidently, the filamentary structure of the cosmic web is captured very well. Quantitatively, we compute the mean matter power spectrum over the test dataset and generated samples. This reveals a very good agreement between their amplitudes and shapes, as seen in the right plot of Fig. 1.

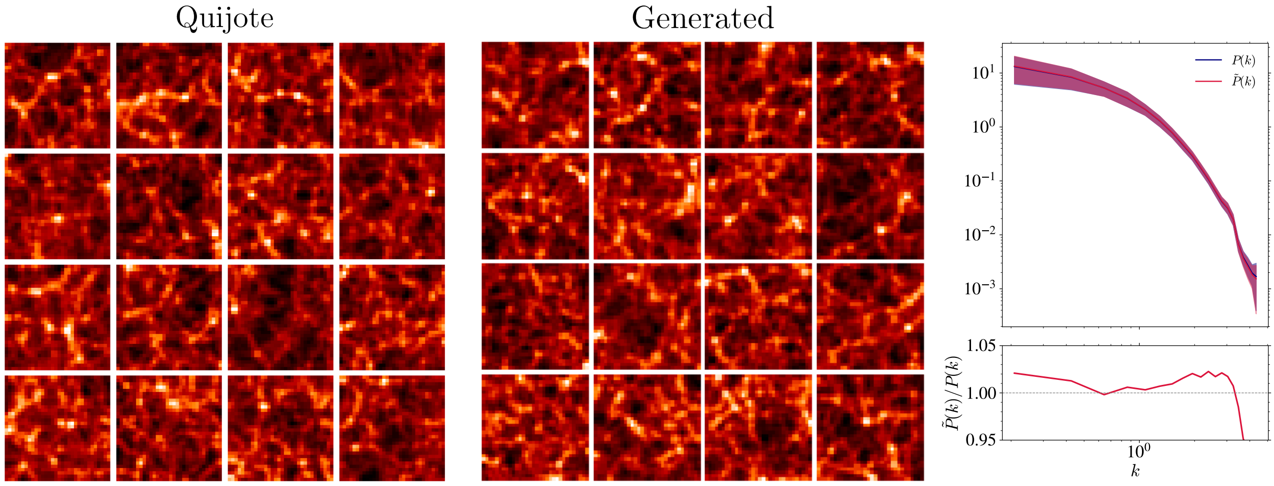

Fig. 1. Left: 16 32×32 samples of both the true (left) and generated log density fields (right). **Right:** Mean matter power spectra of the true (blue) and generated density fields (red). Shaded regions indicate the 1σ dispersion around the spectra of generated samples. The lower plot displays the ratio of mean amplitudes, with the mean power spectrum of the true density fields from Quijote as reference. The mean spectra of the generated density fields are reproduced by the trained diffusion model to within $\approx 3\%$ at all scales.

3 Attention maps for probing cosmic web structures

Self-attention maps in diffusion models quantify how strongly each spatial region in an input field is correlated with every other region. Each spatial location in the input has its own attention map, where the size of the receptive field of the map is defined by its resolution. Maps at higher resolutions capture fine-grained, ie. small-scale structures, while lower resolution maps emphasise larger scale patterns. This multiscale representation of structures should thus offer a new way to probe the main distinct cosmic web environments.

To evaluate what attention maps actually learn, we adopt the T-Web algorithm to segment input density fields from Quijote into their cosmic web environments. The T-Web takes a smoothed density field as input, computes the Hessian of the gravitational potential, also known as the tidal tensor, and evaluates its eigenvalues. Each grid cell is then classified based on the number of eigenvalues above a chosen threshold. Grid cells with zero, one, two or three eigenvalues above the threshold are labelled as voids, walls, filaments and nodes, respectively. In this work, we use this segmentation as ground truth to assess whether attention maps at different resolutions are sensitive to different components of the cosmic web.

Figure 2 illustrates one such example, where we present an input density field, its corresponding T-Web segmentation, and two attention maps extracted from different attention layers of the trained model. The first attention map is strongly activated by overdense structures, capturing the filamentary network along with prominent node regions, while large void regions show little signal, indicating that overdensities are well recovered. In contrast, the second attention map highlights underdense, void-like regions, showing strong signal in these low-density regions and minimal response to overdensities. These differences show that certain layers of the model learn distinct cosmic web patterns across multiple scales.

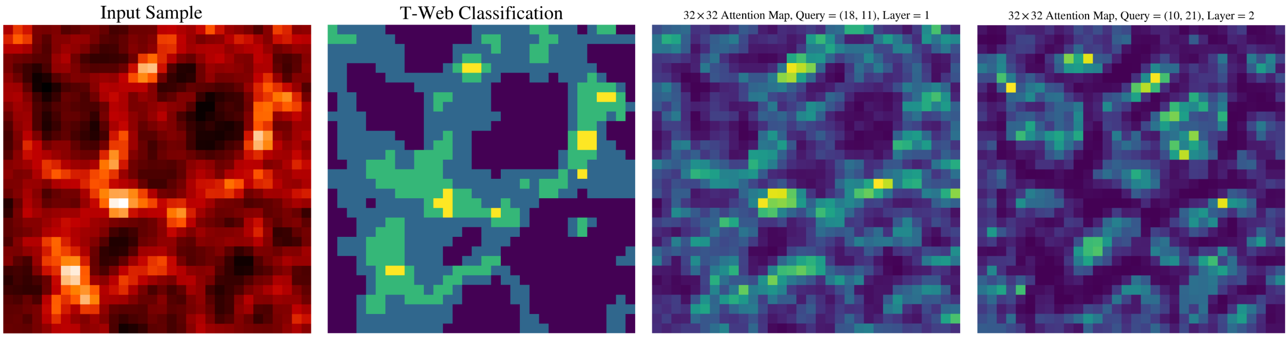

Fig. 2. From left to right: An input log density field from the QUIJOTE suite, T-Web segmentation mask showing the classification of the four cosmic web components (voids, walls, filaments and nodes, from darkest to brightest), and two attention maps of resolution 32×32 from different layers of the trained diffusion model. The first attention map responds to overdense regions, capturing filaments and nodes, while the second map is more sensitive to underdense, void-like regions.

Beyond visual inspection, we use statistical estimators to quantitatively assess the ability of attention maps to capture cosmic web environments. One such tool is the cross-power spectrum, which measures correlations between two spatial fields in Fourier space, averaged over all modes with the same wavenumber k:

$$P_{\delta A}(k) = \langle \tilde{\delta}(\mathbf{k}) \tilde{A}^*(\mathbf{k}) \rangle_k, \tag{3.1}$$

where δ denotes the log density field, and A is the attention map. Since each pixel in an attention map corresponds to a region of the density field (with size depending on the map resolution), and we know which cosmic web environment that region belongs to from the T-Web classification scheme (nodes, filaments, walls and voids), we can compute the mean cross-power spectrum between the density field and the subset of attention maps corresponding to each environment. This environment-dependent approach allows us to investigate how attention maps associated with different structures respond to the density field across multiple scales.

We compute the mean cross-power spectra between input density fields and their associated attention maps of resolution 32×32 across multiple layers of the trained model to quantify correlations at different scales. Across layers, significant correlations appear at large spatial scales, with weaker signals at smaller ones. While some layers exhibit clear and interpretable results, others are less informative. To highlight some representative examples, we focus on layers 1 and 2 in Fig. 3. In layer 1, attention maps associated with node and filament positions in the density field capture strong positive correlations, while those associated with void positions exhibit significant anti-correlations. Walls generally show weaker responses, which is consistent with the difficulty of identifying them even with dynamical classifiers such as T-Web. In layer 2, attention maps capture negative correlations across all environments, demonstrating sensitivity to underdense regions. These results confirm, both visually and quantitatively, that attention maps from different layers of the trained model respond to distinct structural patterns of the cosmic web.

4 Conclusions

We have shown that generative diffusion models can successfully learn the distribution of dark matter density fields from the QUIJOTE simulations, demonstrating high fidelity in capturing the spatial structure of the cosmic web. This indicates that the model captures not only two-point statistics such as the matter power spectrum, but also higher-order correlations inherent to the non-Gaussian structure of the cosmic web.

To better understand how diffusion models capture cosmic web structures, we perform a quantitative analysis of self-attention maps in our trained model. We find that attention maps from different layers are sensitive to distinct features, where some respond predominantly to overdense structures such as filaments and nodes, while others emphasise underdense regions such as voids. These results, supported by both visual inspection and statistical analyses, show that the network crafts multiscale features containing physically meaningful information about the cosmic web. Our findings demonstrate that not only can diffusion models emulate the

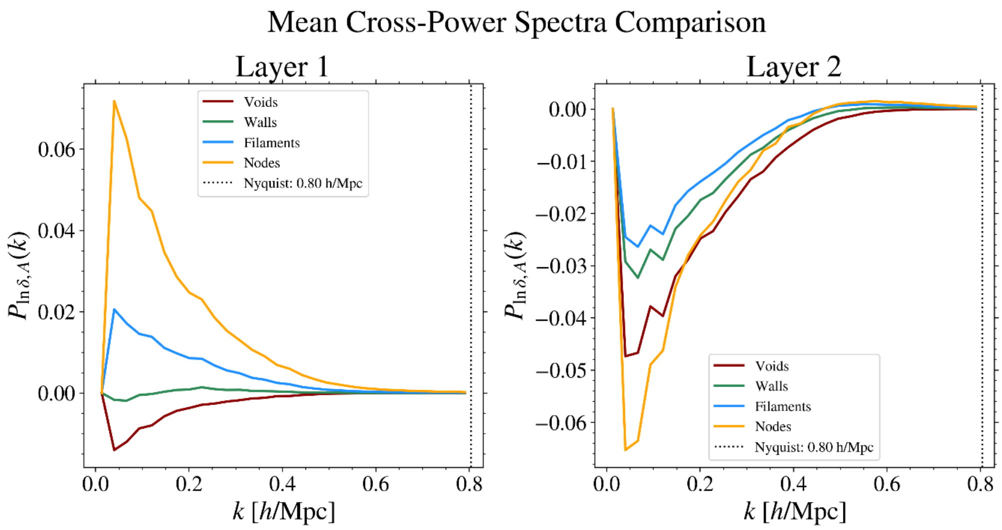

Fig. 3. Cross-power spectra between input log density fields and their associated 32×32 attention maps for layers 1 (left) and 2 (right) of the trained diffusion model. The spectra are computed separately for maps corresponding to different cosmic web environments, as classified by T-Web. Layer 1 displays strong positive correlations for nodes and filaments, and anti-correlations for voids. However, layer 2 exhibits an overall negative correlation across all component cases.

cosmic web with high statistical fidelity, but also that their self-attention maps encode important structural information, potentially offering a novel way to probe cosmic web structures.

MN thanks CNES and the DIM-IdF Origines for funding his PhD grant.

References

Bond, J. R., Kofman, L., & Pogosyan, D. 1996, Nature, 380, 603–606

Colless, M., Peterson, B. A., Jackson, C., et al. 2003, The 2dF Galaxy Redshift Survey: Final Data Release

Doroshkevich, A. G., Kotok, E. V., Poliudov, A. N., et al. 1980, MNRAS, 192, 321

Geller, M. J. & Huchra, J. P. 1989, Science, 246, 897

Hahn, O., Porciani, C., Carollo, C. M., & Dekel, A. 2007, Monthly Notices of the Royal Astronomical Society, 375, 489–499

Ho, J., Jain, A., & Abbeel, P. 2020, Denoising Diffusion Probabilistic Models

Huchra, J., Jarrett, T., Skrutskie, M., et al. 2005, in Astronomical Society of the Pacific Conference Series, Vol. 329, Nearby Large-Scale Structures and the Zone of Avoidance, ed. A. P. Fairall & P. A. Woudt, 135

Jenkins, A., Frenk, C. S., Pearce, F. R., et al. 1998, The Astrophysical Journal, 499, 20–40

Klypin, A. A. & Shandarin, S. F. 1983, MNRAS, 204, 891

Planck Collaboration, Aghanim, N., Akrami, Y., Ashdown, M., et al. 2020, Astronomy & Astrophysics, 641, A6

Sohl-Dickstein, J., Weiss, E. A., Maheswaranathan, N., & Ganguli, S. 2015, Deep Unsupervised Learning using Nonequilibrium Thermodynamics

Springel, V., White, S. D. M., Jenkins, A., et al. 2005, Nature, 435, 629

Tegmark, M., Strauss, M. A., Blanton, M. R., et al. 2004, Physical Review D, 69

Vaswani, A., Shazeer, N., Parmar, N., et al. 2023, Attention Is All You Need

Villaescusa-Navarro, F., Hahn, C., Massara, E., et al. 2020, The Astrophysical Journal Supplement Series, 250, 2