Super-Resolution Emulation of Large Cosmological Fields with a 3D Conditional Diffusion Model

Adam Rouhiainen
Department of Physics
University of Wisconsin-Madison
Madison, WI 53706
rouhiainen@wisc.edu

Michael Gira
Department of Electrical and Computer Engineering
University of Wisconsin-Madison
Madison, WI 53706

Moritz Münchmeyer
Department of Physics
University of Wisconsin-Madison
Madison, WI 53706

Kangwook Lee
Department of Electrical and Computer Engineering
University of Wisconsin-Madison
Madison, WI 53706

Gary Shiu
Department of Physics
University of Wisconsin-Madison
Madison, WI 53706

Abstract

High-resolution (HR) simulations of baryonic matter in cosmology often take millions of CPU hours. On the other hand, low resolution (LR) dark matter simulations of the same comological volume use minimal computing resources. In this paper we train a conditional diffusion model to upgrade LR dark matter simulations probabilistically to HR baryonic matter simulations. Our approach is based on the Palette diffusion model, which we generalize to 3 dimensions. Our super-resolution emulator is trained to perform outpainting, and we generate a simulation box with 8 times the volume of our Illustris TNG300 training data, constructed with over 9000 outpainting iterations.

1 Introduction

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

A novel development in this work is that we train the model to perform outpainting in 3 dimensions, allowing us to generate arbitrarily large simulation volumes, larger than the HR training data. This is

37th Conference on Neural Information Processing Systems (NeurIPS 2023).

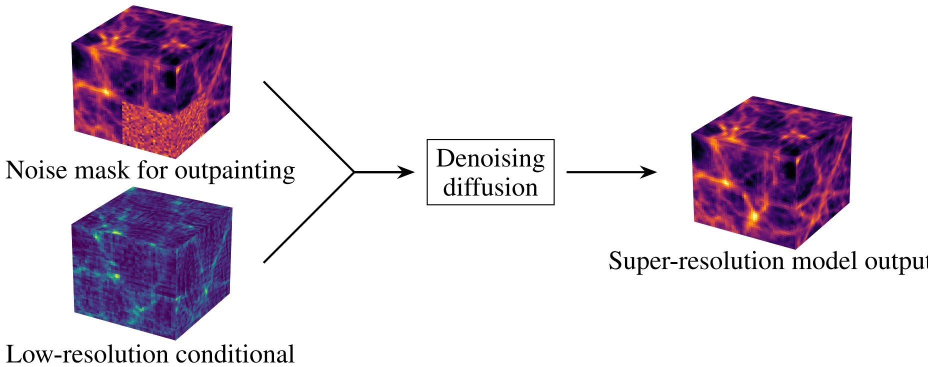

Figure 1: We train a diffusion model to do outpainting on cosmological fields. By iterating outpainting many times over, we construct large SR fields given an LR conditional to guide the large scale modes.

possible due to the locality of the underlying physics of structure formation. On large scales, the LR simulations with the same initial conditions give the same final result as the vastly more expensive HR simulation. We developed an iterative outpainting procedure where a large LR volume is upgraded to SR locally, patch by patch. To enforce consistency between patches, our SR diffusion model is conditioned both on the LR simulation and on neighboring SR samples.

Super-resolution of late time cosmological fields have been previously developed with GANs [7–10]. Further, the work [8] make large simulations by sewing together smaller SR volumes. However, they do not condition neighboring SR volumes on each other, which necessarily means that there will be boundary effects in the SR field. In our work, we explicitly condition the local SR volumes on their neighboring volumes, previously generated in the outpainting chain. Recently, [11] used a diffusion model to perform super-resolution on 2-dimensional dark matter fields. Their work is technically most similar to ours, with their approach including a Fourier filter on the large-scale structure data to boost the importance of learning the nonlinear scales.

2 Method - Super-resolution with a conditional diffusion model

A conditional denoising diffusion model is a probabilistic generative model that samples $\boldsymbol{y}$ from $p(\boldsymbol{y}|\boldsymbol{x})$, where the (SR) field $\boldsymbol{y}$ is conditional on a (LR) field $\boldsymbol{x}$. The model has a forward process for training, and a reverse process for generation. In the forward diffusion process, a field $\boldsymbol{y}_0$ iteratively gets added small amounts of noise with variances β_t, producing a final noisy field $\boldsymbol{y}_T$:

$$p(\boldsymbol{y}_T|\boldsymbol{y}_0) = \prod_{t=1}^{T} \mathcal{N}(\boldsymbol{y}_t|\sqrt{1-\beta_t}\,\boldsymbol{y}_{t-1}, \beta_t \mathbb{I}) = \mathcal{N}(\boldsymbol{y}_T|\sqrt{\gamma_T}\,\boldsymbol{y}_0, (1-\gamma_T)\mathbb{I}). \tag{1}$$

Here $\gamma_T \equiv \prod_{t=1}^{T}(1-\beta_t)$. The β_t are model hyperparameters chosen so that $\sqrt{\gamma_T}\,\boldsymbol{y}_0$ is small relative to $\boldsymbol{y}_T$, and thus we produce nearly true noise by the final step T.

The reverse process generates fields in the distribution of training data $\boldsymbol{y}_0$ from noise $\boldsymbol{y}_T$ by solving for $\boldsymbol{y}_{t-1}$ in terms of $\boldsymbol{y}_t$. To achieve this, a neural network $f_{\boldsymbol{\theta}}(\boldsymbol{x}, \boldsymbol{y}_t, \gamma_t)$ with parameters $\boldsymbol{\theta}$ is fit to the noise $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbb{I})$ by minimizing the loss function $\mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})}\mathbb{E}_{\boldsymbol{\epsilon},\gamma}\left\|f_{\boldsymbol{\theta}}(\boldsymbol{x}, \sqrt{\gamma}\,\boldsymbol{y}_0 + \sqrt{1-\gamma}\,\boldsymbol{\epsilon}, \gamma) - \boldsymbol{\epsilon}\right\|_2^2$. The reverse diffusion step to sample $\boldsymbol{y}_{t-1}$ from $P(\boldsymbol{y}_{t-1}|\boldsymbol{x}, \boldsymbol{y}_t)$ is (see [2] for details)

$$\boldsymbol{y}_{t-1} = \frac{1}{\sqrt{1-\beta_t}}\left(\boldsymbol{y}_t - \frac{\beta_t}{\sqrt{1-\gamma_t}} f_{\boldsymbol{\theta}}(\boldsymbol{x}, \boldsymbol{y}_t, \gamma_t)\right) + \mathcal{N}(0, \beta_t \mathbb{I}). \tag{2}$$

This reverse diffusion step is iterated to obtain a sample $\boldsymbol{y}_0$ from $\boldsymbol{y}_T$. Because $\boldsymbol{y}_T$ is a random noise field, we can obtain a variety of samples $\boldsymbol{y}_0$ without adjusting the network parameters $\boldsymbol{\theta}$.

2.1 Data

The training data is HR-LR pairs of 48^3 px cubes. The HR data to learn the baryonic physics of the large-scale structure comes from IllustrisTNG [12–16], a set of three gravo-magnetohydrodynamical simulations run on `Arepo` [17]. We use the TNG300 run, which simulated 2500^3 matter and 2500^3 dark matter particles in a 205 Mpc$/h \approx 300$ Mpc length cube. TNG300 computed over 10^7 time steps from redshift $z = 127$ to $z = 0$, taking 35 million core hours. We use the $z = 0.01$ snapshot. We use cloud-in-cell mass assignment to place the TNG300 baryons onto a 264^3 px cubic mesh.

We simulate dark matter fields to create the LR conditional training and LR test data, both containing the same physics from different initial seeds. Our conditional fields are LR compared to TNG300 as they are dark matter only simulations, contain fewer particles, and run across fewer time steps. For the LR *training* data, we use the same initial seed, box size, and cosmological parameters as TNG300, with 128^3 dark matter particles, computing ~ 2000 time steps from $z = 127$ to $z = 0.01$. For the LR *test* data, we run this simulation again with twice the length (8 times the volume) with a different seed. There is no HR truth for the LR test data.

Finally, the HR-LR training pairs are constructed by randomly cropping 16000 cubes of size 48^3 px $(17.1$ Mpc$/h)^3$ out of the full 264^3 px fields, with random rotations and mirrors.

2.2 Iterative outpainting to generate large fields

We generate large cosmological fields with iterative outpainting. The basic idea is that after an initial SR field is generated, every subsequent SR field is generated conditional on its surrounding SR field(s) already generated. Training the model to outpaint requires masking sub-volumes of the HR data in the HR-LR pairs, and thus masked sub-volumes at inference can be generated given surrounding SR data. An illustration is shown in Fig. 1, with further explanation in the Appendix or the full version of this paper [18].

2.3 Network details and training

We use a U-net [19] to learn the denoising function $f_{\boldsymbol{\theta}}(\boldsymbol{x}, \boldsymbol{y}_t, \gamma_t)$. Our U-net is based on the "guided diffusion" U-net used in [4], minimally modified for 3-dimensional data. We use Big-GAN residual blocks for two downsampling steps with the number of channels being 64, 128, and 256, and similarly two upsampling steps. We use 1 residual block per resolution. The model has 31.5 million parameters.

We train on the 48^3 px cube HR-LR pairs, with 2000 diffusion steps, and β_t linearly increasing from $\beta_0 = 10^{-6}$ to $\beta_T = 0.01$. During training, each HR field is given a random rectangular cuboid mask to learn outpainting. We use the Adam optimizer with a 10^{-4} learning rate and a 0.9999 EMA decay. With two A100s and batch size of 16 per GPU, the loss converges in 70 hours.

3 Results

We show visual results in Fig. 2 between the LR conditional test data and SR model outputs, on 528^3 px $(410$ Mpc$/h)^3$ cubes. Each SR field is constructed with 21^3 outpainting iterations in 24 px increments, taking about 120 hours to generate. Results are given with zero-mean overdensities.

We quantify our SR results for the large 528^3 px $(410$ Mpc$/h)^3$ cube with several summary statistics familiar to cosmology, comparing the LR test field, SR model output, and now include the HR TNG300 training field as a truth comparison. We show the one-point probability density function in Fig. 3 (top left); here the LR field has a large low-density spike where no particles are present in a voxel, and the SR field matches HR identically. The power spectrum $P(k)$ and bispectrum $B(k_1, k_2, k_3)$ are Fourier space two- and three-point correlation functions, defined by

$$P(k)\ \delta_{\mathrm{D}}(\boldsymbol{k} + \boldsymbol{k}') = \langle \delta(\boldsymbol{k}) \delta(\boldsymbol{k}') \rangle, \tag{3}$$

$$B(k_1, k_2, k_3)\ \delta_{\mathrm{D}}(\boldsymbol{k}_1 + \boldsymbol{k}_2 + \boldsymbol{k}_3) = \langle \delta(\boldsymbol{k}_1) \delta(\boldsymbol{k}_2) \delta(\boldsymbol{k}_3) \rangle. \tag{4}$$

Here δ_{D} is the Dirac delta function, and $\langle ... \rangle$ is an ensemble average. The power spectrum is plotted in Fig. 3 (top right), noting the 24 px outpainting scale $k_{24\text{ px}}$. Below $k_{24\text{ px}}$, the model loses information on the large length modes, but the SR field is still accurate with HR, especially in generating BAOs; therefore, the LR conditional is successfully guiding the large length scales of the SR outpaintings.

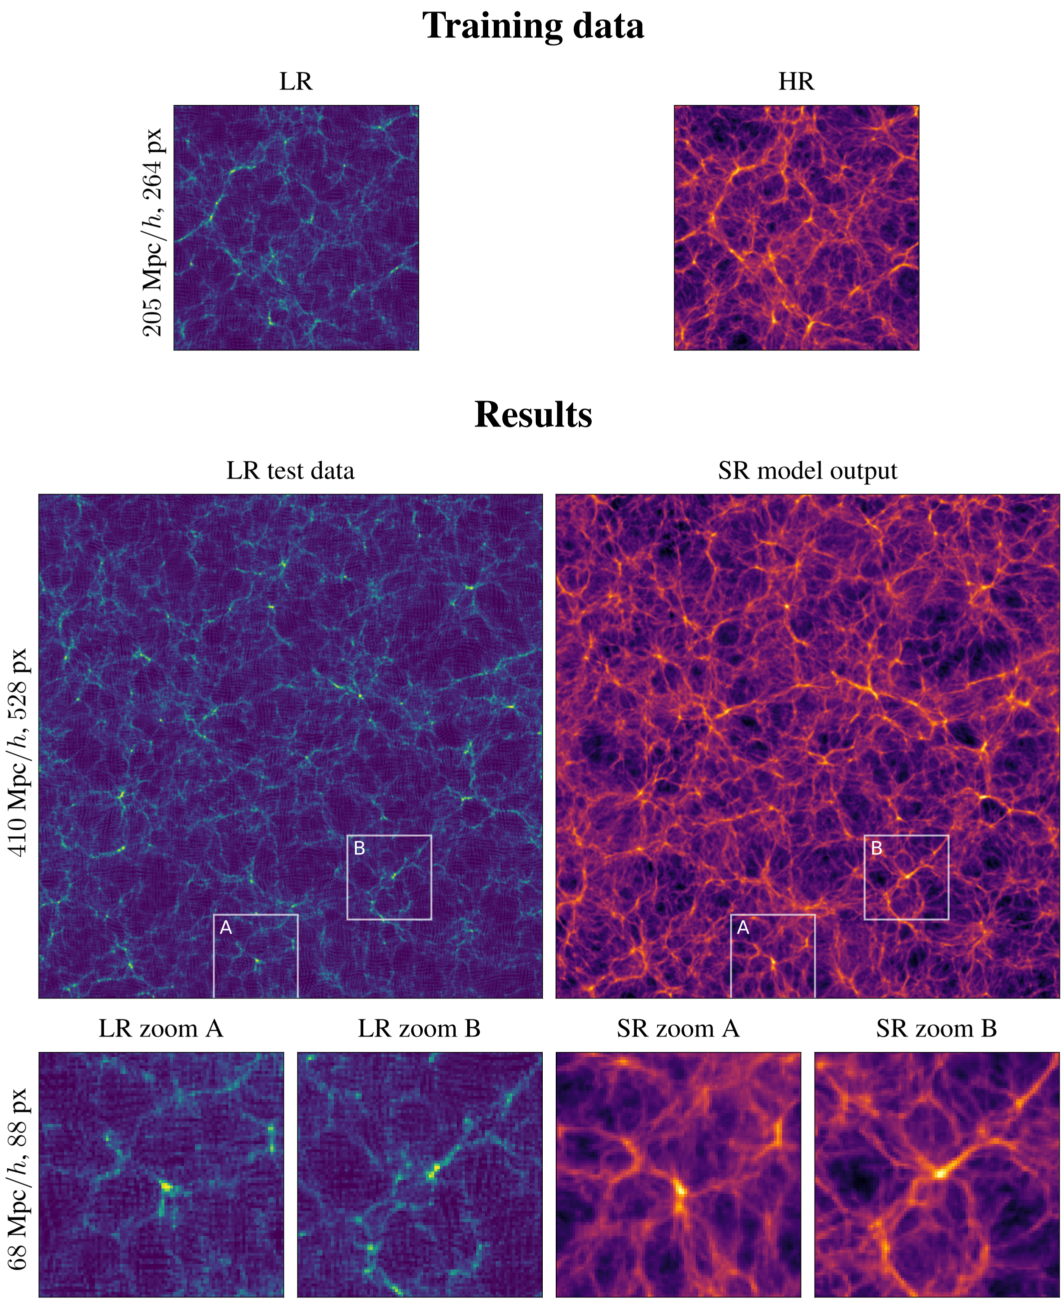

Figure 2: Matter density results for our super-resolution diffusion model in generating a volume larger than the entire training data volume. Images are 2-dimensional projections of depth 19 Mpc$/h$. (Top row) The training data comes from the single pair of boxes shown. The model trains on 48 px length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

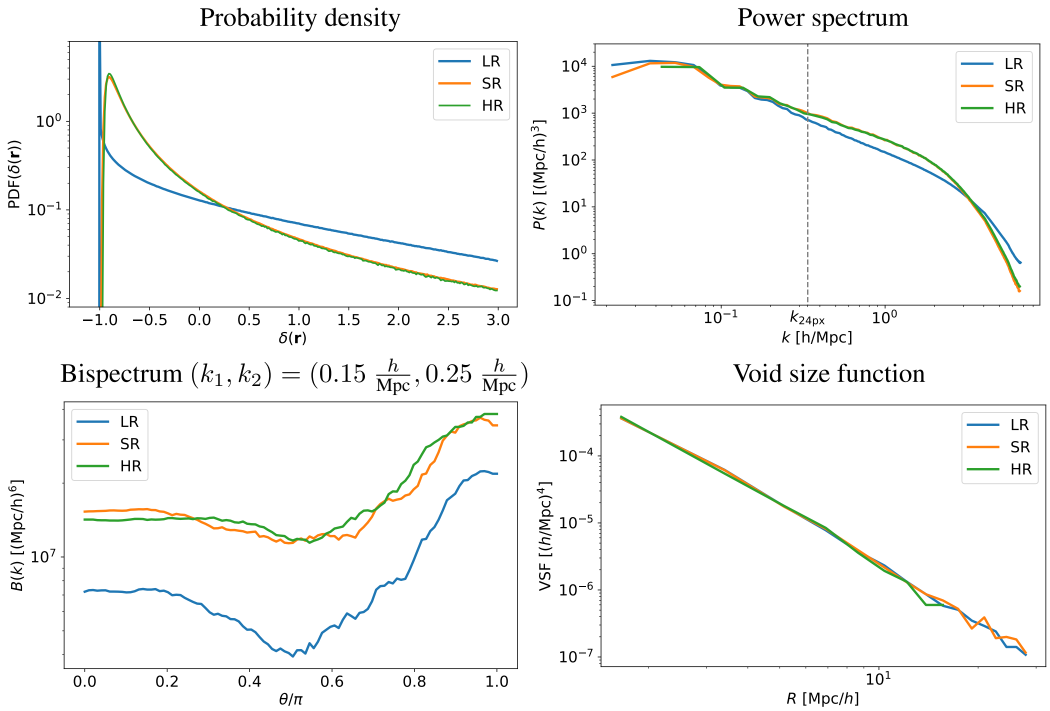

Figure 3: Summary statistics comparing the LR test field, SR model output, and HR training field. The SR and HR fields closely match, indicating the diffusion model has correctly learned the HR physics.

The bispectrum measures the non-Gaussian structure of the field. We show in Fig. 3 (bottom left) a projected bispectrum as a function of the angle θ between two vectors of length $k_1 = 0.15\ h/\mathrm{Mpc}$, $k_2 = 0.25\ h/\mathrm{Mpc}$ [20]. Finally, the void size function [21] calculated with `Pylians` [22] is shown in Fig. 3 (bottom right). For all the summary statistics calculated here, the SR model output closely matches the HR truth.

4 Conclusion

SR emulators are a promising tool to open the computational bottleneck of HR baryonic simulations in cosmology. In this work we evaluated the performance of a diffusion model on volumetric data, and made use of the property of locality in cosmological structure formation to develop a conditional outpainting scheme that can upgrade large LR volumes. The resulting SR volume matches the summary statistics of the training HR simulation closely. Leading up to this work, we found that training generative models in 3D is not always successful, as we intended to train normalizing flows such as Real NVP [23] and Glow [24] on the same task. While 2D results were promising, the flows were not accurate in 3D and we thus moved to the more expressive diffusion models.

Further research with 3-dimensional diffusion models would be greatly aided with a faster denoising algorithm. While this work uses a conditional denoising diffusion probabilistic model (DDPM), recent developments in diffusion research include the significantly faster denoising diffusion implicit model (DDIM) [25] and denoising probabilistic models solver (DPM-Solver) [26]. The DDIM algorithm and DPM-Solver advertise a factor of 10 to 100 reduction in the number of denoising steps with a trade-off of a small loss in accuracy. We tested an implementation of DDIM in preparing this work, but our results were subpar compared to DDPM. Nevertheless, it is likely that with more work the sample generation time can be reduced by a significant factor.

5 Acknowledgements

We thank Ying Fan for helpful discussion. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics under Award Numbers DE-SC-0023719 (G.S.) and DE-SC-0017647 (G.S.,M.M.). M.M is supported by NSF grant 2307109. Model training and inference used resources from the Data Science Institute at the University of Wisconsin-Madison.

References

[1] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 2015. PMLR. URL: `https://proceedings.mlr.press/v37/sohl-dickstein15.html`.

[2] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. 2020. `arXiv:2006.11239`.

[3] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. *Palette: Image-to-Image Diffusion Models*. Association for Computing Machinery, New York, NY, USA, 2022. `doi:10.1145/3528233.3530757`.

[4] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion Models Beat GANs on Image Synthesis. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL: `https://openreview.net/forum?id=AAWuCvzaVt`.

[5] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing Flows for Probabilistic Modeling and Inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021. URL: `http://jmlr.org/papers/v22/19-1028.html`.

[6] Zhifeng Kong and Kamalika Chaudhuri. The Expressive Power of a Class of Normalizing Flow Models, 2020. `arXiv:2006.00392`.

[7] Doogesh Kodi Ramanah, Tom Charnock, Francisco Villaescusa-Navarro, and Benjamin D Wandelt. Super-resolution emulator of cosmological simulations using deep physical models. *Monthly Notices of the Royal Astronomical Society*, 495(4):4227–4236, 2020. `doi:10.1093/mnras/staa1428`.

[8] Yin Li, Yueying Ni, Rupert A. C. Croft, Tiziana Di Matteo, Simeon Bird, and Yu Feng. Ai-assisted superresolution cosmological simulations. *Proceedings of the National Academy of Sciences*, 118(19):e2022038118, 2021. `doi:10.1073/pnas.2022038118`.

[9] Yueying Ni, Yin Li, Patrick Lachance, Rupert A. C. Croft, Tiziana Di Matteo, Simeon Bird, and Yu Feng. AI-assisted superresolution cosmological simulations – II. Halo substructures, velocities, and higher order statistics. *Monthly Notices of the Royal Astronomical Society*, 507(1):1021–1033, 2021. `doi:10.1093/mnras/stab2113`.

[10] Xiaowen Zhang, Patrick Lachance, Yueying Ni, Yin Li, Rupert A. C. Croft, Tiziana Di Matteo, Simeon Bird, and Yu Feng. AI-assisted super-resolution cosmological simulations III: Time evolution, 2023. `arXiv:2305.12222`.

[11] Andreas Schanz, Florian List, and Oliver Hahn. Stochastic Super-resolution of Cosmological Simulations with Denoising Diffusion Models, 2023. `arXiv:2310.06929`.

[12] Volker Springel, Rüdiger Pakmor, Annalisa Pillepich, Rainer Weinberger, Dylan Nelson, Lars Hernquist, Mark Vogelsberger, Shy Genel, Paul Torrey, Federico Marinacci, and Jill Naiman. First results from the IllustrisTNG simulations: matter and galaxy clustering. *Monthly Notices of the Royal Astronomical Society*, 475(1):676–698, 2017. `doi:10.1093/mnras/stx3304`.

[13] Annalisa Pillepich, Dylan Nelson, Lars Hernquist, Volker Springel, Rüdiger Pakmor, Paul Torrey, Rainer Weinberger, Shy Genel, Jill P Naiman, Federico Marinacci, and Mark Vogelsberger. First results from the IllustrisTNG simulations: the stellar mass content of groups and clusters of galaxies. *Monthly Notices of the Royal Astronomical Society*, 475(1):648–675, 2017. `doi:10.1093/mnras/stx3112`.

[14] Jill P Naiman, Annalisa Pillepich, Volker Springel, Enrico Ramirez-Ruiz, Paul Torrey, Mark Vogelsberger, Rüdiger Pakmor, Dylan Nelson, Federico Marinacci, Lars Hernquist, Rainer Weinberger, and Shy Genel. First results from the IllustrisTNG simulations: a tale of two elements – chemical evolution of magnesium and europium. *Monthly Notices of the Royal Astronomical Society*, 477(1):1206–1224, 2018. `doi:10.1093/mnras/sty618`.

[15] Federico Marinacci, Mark Vogelsberger, Rüdiger Pakmor, Paul Torrey, Volker Springel, Lars Hernquist, Dylan Nelson, Rainer Weinberger, Annalisa Pillepich, Jill Naiman, and Shy Genel. First results from the

IllustrisTNG simulations: radio haloes and magnetic fields. *Monthly Notices of the Royal Astronomical Society*, 480(4):5113–5139, 2018. doi:10.1093/mnras/sty2206.

[16] Dylan Nelson, Annalisa Pillepich, Volker Springel, Rainer Weinberger, Lars Hernquist, Rüdiger Pakmor, Shy Genel, Paul Torrey, Mark Vogelsberger, Guinevere Kauffmann, Federico Marinacci, and Jill Naiman. First results from the IllustrisTNG simulations: the galaxy colour bimodality. *Monthly Notices of the Royal Astronomical Society*, 475(1):624–647, 2017. doi:10.1093/mnras/stx3040.

[17] Volker Springel. E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh. *Monthly Notices of the Royal Astronomical Society*, 401(2):791–851, 2010. doi:10.1111/j.1365-2966.2009.15715.x.

[18] Adam Rouhiainen, Michael Gira, Moritz Münchmeyer, Kangwook Lee, and Gary Shiu. Super-Resolution Emulation of Large Cosmological Fields with a 3D Conditional Diffusion Model, 2023. arXiv:2311.05217.

[19] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing. doi:10.1007/978-3-319-24574-4_28.

[20] Elena Giusarma, Mauricio Reyes, Francisco Villaescusa-Navarro, Siyu He, Shirley Ho, and ChangHoon Hahn. Learning Neutrino Effects in Cosmology with Convolutional Neural Network. *The Astrophysical Journal*, 950(1):70, 2023. doi:10.3847/1538-4357/accd61.

[21] Arka Banerjee and Neal Dalal. Simulating nonlinear cosmological structure formation with massive neutrinos. *Journal of Cosmology and Astroparticle Physics*, 2016(11):015, 2016. doi:10.1088/1475-7516/2016/11/015.

[22] Francisco Villaescusa-Navarro. Pylians, 2020. URL: https://github.com/franciscovillaescusa/Pylians3.

[23] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using Real NVP. 2016. arXiv:1605.08803.

[24] Diederik P. Kingma and Prafulla Dhariwal. Glow: Generative Flow with Invertible 1x1 Convolutions. 2018. arXiv:1807.03039.

[25] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising Diffusion Implicit Models, 2022. arXiv:2010.02502.

[26] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. DPM-Solver: A Fast ODE Solver for Diffusion Probabilistic Model Sampling in Around 10 Steps, 2022. arXiv:2206.00927.

A Iterative outpainting illustration

This work uses massively iterative outpainting that can generate cosmological volumes much larger than the training data. Given a large LR field, an SR field is generated sequentially, patch by patch, with each new patch generated conditional on both the LR field and adjacent SR fields. The LR conditional guides the large length scales of the SR field, while the adjacent SR conditional ensures that newly generated SR fields are smooth and physically consistent with the larger SR volume. Training the model to outpaint requires masking sub-volumes of the HR data in the HR-LR pairs, as shown in Fig.1, and thus masked sub-volumes at inference can be generated given surrounding SR data. An illustration of the iterative outpainting is shown in Fig. 4.

Our iterative outpainting method is motivated by the nature of mode coupling in the large-scale structure. On large scales, for $k \lesssim 0.1\ h/\text{Mpc}$ today, the evolution of physical perturbations is linear, as modes evolve independently. On intermediate scales, $0.1\ h/\text{Mpc} \lesssim k \lesssim 0.5\ h/\text{Mpc}$ modes evolve nonlinearly but should be accurately captured by the LR dark matter N-body simulation. On smaller scales $k \gtrsim 0.5\ h/\text{Mpc}$ we want the diffusion model to model nonlinearity and baryonic feedback. This informs us about the minimum physical size required for the outpainting volumes, and the diffusion model we describe here has a fundamental mode of $k_{24\text{px}} = 0.34\ h/\text{Mpc}$. The diffusion model can modify the results of the LR simulation on physical scales smaller than this scale, and can thus take into account mode coupling on these scales. Mode coupling is included by the LR simulation on large scales, but cannot be modified by the diffusion model due to its outpainting window size, and we thus assume that the LR simulation is correct on these scales. In position space, the 24 px window corresponds to a physical length of $18.6\ \text{Mpc}/h$. This can be compared to the typical particle displacement of $5\ \text{Mpc}/h$, with an upper limit of $\sim 20\ \text{Mpc}/h$ [7]. Most of this displacement is already included in the LR simulation, so our window size is sufficient.

In detail, the iterative outpainting procedure works as follows. First, a 48^3 px SR cube is generated conditional on a 48^3 px LR cube; this first cube is shown in the top left of Fig. 4. We move in row-major order, outpainting 24 px at a time. The second SR volume generated is thus conditional on both its underlying 48^3 px LR cube, as well as the 24 px length right half of the first SR cube. After an entire plane is generated, the outpaintings move to the third dimension, with every subsequent plane conditional on the previously generated plane. The outpaintings continue in this way until the entire LR volume is generated to SR. We never break conditionality on previous adjacent SR regions, even after moving into the third dimension. Thus throughout the full SR volume, every locally outpainted volume is conditional on all adjacent previously generated volumes.

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

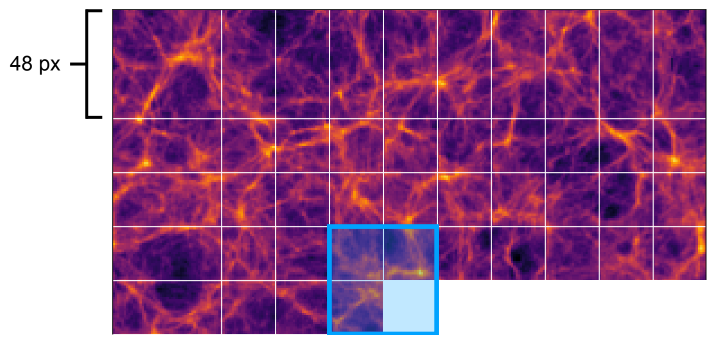

Figure 4: Iterative outpainting illustration, 2-dimensional projection. The blue cube contains the previously generated SR fields adjacent to the subsequent volume ready to be outpainted. After a new volume is generated, the blue cube then moves in row-major order for the next SR volume to be generated. After a plane of small volumes is generated to SR, the outpaintings continue in the third dimension (perpendicular to the page), and the next plane is generated, with each SR volume also conditional on the previous plane.