Closing the resolution gap in Lyman alpha simulations with deep learning

Cooper Jacobus
Department of Astronomy
University of California, Berkeley
Berkeley, CA 94720
cjacobus@berkeley.edu

Peter Harrington
Lawrence Berkeley National Laboratory
Berkeley, CA 94720, USA
pharrington@lbl.gov

Zarija Lukić
Lawrence Berkeley National Laboratory
Berkeley, CA 94720, USA
zarija@lbl.gov

Abstract

In recent years, super-resolution and related approaches powered by deep neural networks have emerged as a compelling option to accelerate computationally expensive cosmological simulations, which require modeling complex multi-physics systems in large spatial volumes. However, training such models in a physically consistent way is not always feasible or well-defined, as the data volume output by a super-resolution model may be too large, and the spatiotemporal dynamics of the simulation as well as the statistics of key observables like Lyman alpha (Lyα) flux are very sensitive to changes in resolution. In this work we address both challenges simultaneously, training neural networks to synthesize Lyα and other hydrodynamic fields with correct statistics on the relevant length scales but represented on the coarse grid of the input simulations. Effectively, our method is capable of 8x super-resolving a coarse simulation in-place without increasing memory footprint, using just a single pair of simulations for training. With chunked inference, we are able to apply the model to simulations of arbitrary size after training, and demonstrate this capability on a very large volume simulation spanning 600 Mpc/h.

1 Introduction

The Lyman alpha (Lyα) forest is a primary probe of cosmological structure at redshifts $2 < z < 6$ with length scales ranging from sub-megaparsec (Mpc) to tens of Mpc, complementing other probes by extending the reach to significantly smaller scales. This allows extracting information on neutrino masses, dark matter properties, inflation and reionization in the early universe [18; 15; 4; 25; 21; 8; 2; 22]. Answering these fundamental cosmological questions necessitates not only precise observations, but also the creation of simulated model universes which help constrain parameters of a cosmological model. To compare numerical simulations with the growing body of increasingly accurate measurements of the Lyα forest, it is critical that the simulated models are of precision comparable to that of the data, but doing so requires state-of-the-art supercomputing resources. More specifically, to resolve the Lyα forest at percent-level accuracy, the spatial resolution required is $\Delta x \sim 20$ kpc/h [13; 23] which implies simulated volumes of $L_{\rm box} \sim 80 - 160$ Mpc/h with 4096^3-8192^3 resolution elements.

Machine Learning and the Physical Sciences workshop, NeurIPS 2022.

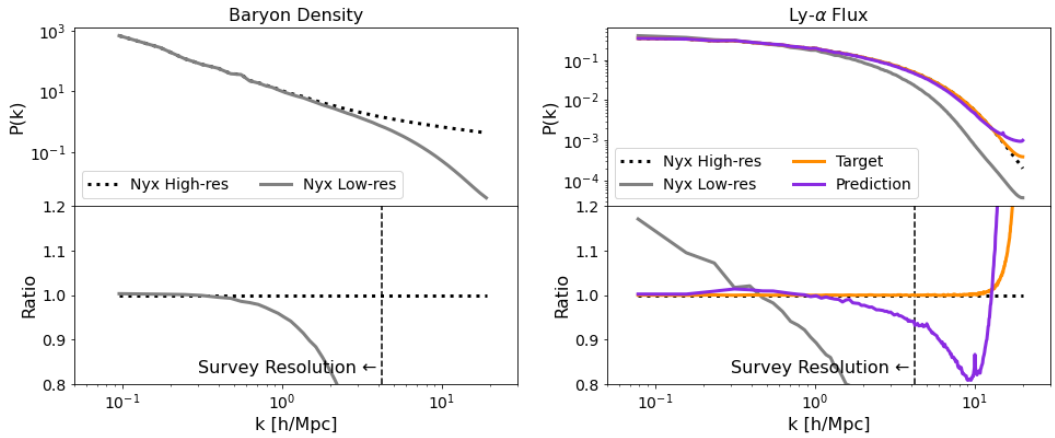

Figure 1: The Power Spectrum of baryon density (left) and Lyα flux (right) from our validation simulation. The discrepancy at high frequencies in baryon density (and other hydrodynamic quantities) at low resolution (grey) leads to errors on all frequencies in Lyα flux. Our model (purple) produces an estimate of flux from the low-resolution simulation which avoids these errors, by training on a target (orange) subsampled from the high-resolution simulation. The ratios are taken with respect to the high-resolution simulation. The vertical dashed line indicates a rough cutoff for DESI resolution.

At such a box size, the simulated volumes are still a factor of $\sim$20 smaller on a side than the volume which will be observed by the ongoing DESI survey [5]. Because of the low-density of quasars in previous large-volume spectroscopic surveys, a 3D Lyman-α power spectrum (P3D) measurement – measuring both large ($\sim$ 100 Mpc) and small-scale ($\sim$ 1Mpc) flux fluctuations at the same time – has never been performed. However, the P3D is a cosmic gold mine that can strongly improve many cosmological constraints, and the large increase of close quasar lines of sight in DESI will allow for such a measurement for the first time. To interpret the observational measurement, the gap in spatial scales between simulated and observed volumes is a key challenge which needs to be addressed. An attractive potential solution to this problem is to ease the computation burden of full-physics, full-resolution simulations by leveraging deep neural networks for super-resolution or similar coarse-to-fine mappings, an idea that has received widespread interest in the general cosmological community [6; 7; 20; 19; 11; 10; 14; 3].

Super-resolution and the more general idea of mapping from coarse, computationally cheap inputs to the desired fine-grained outputs is appealing, but meets several challenges when applied to multi-physics cosmological simulations [6]. On extreme computational scales, the cosmological volumes of interest may be simply too large to apply standard super-resolution, as our model output may overwhelm storage resources [11]. Compounding the problem, defining the training targets themselves may be a non-trivial and resolution-dependent task. For example, in the case of Lyα flux considered here, the operation to compute optical depth from hydrodynamic fields such as baryon density, temperature, and local velocity (an analytical integral done in simulation post-processing), does not commute with downsampling or upsampling operators, and this strongly degrades the shape of standard summary statistics like power spectra (see Figure 1).

In this work, we address the aforementioned concerns related to data volumes, computational constraints, and resolution-dependence, and present a deep neural network capable of generating Lyα with high accuracy at the relevant scales from coarse inputs. We achieve this by training an adversarial U-Net on carefully sampled data, and arrive at a model which is capable of effectively performing 8x super-resolution in-place with relatively little training data. Once trained, our model can then be applied to very large volume simulations and produce mock Lyα skies with power spectra accurate to a few percent, even at the length scale limit of modern surveys such as DESI.

2 Simulation Data and Resolution Dependence

We use simulations from the cosmological hydrodynamics code `Nyx` [1] as the basis for our investigation. `Nyx` is a state-of-the-art, adaptive mesh, N-body and gas dynamics code for large-scale

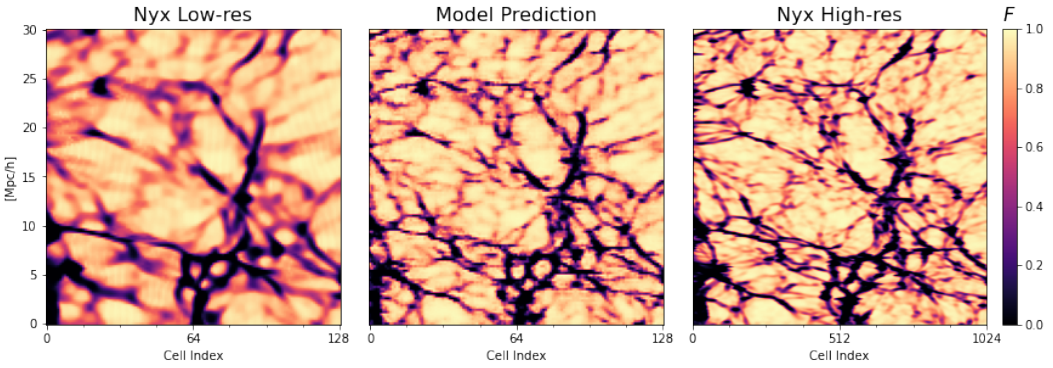

Figure 2: Comparison of simulated and predicted Lyman-α Flux, F (the fraction of light that passes through a region). The left and right panels illustrate how the scales and features in simulated Lyα are heavily dependent on the resolution used. Our model prediction (middle panel) is much closer to the high-resolution result, despite only requiring the low-resolution hydrodynamic fields as input, and making Lyα predictions only on the reduced grid of the low-resolution box.

cosmological simulations, modeling temporal evolution of the universe via a system of discrete dark matter particles gravitationally coupled to an inviscid ideal fluid in a co-moving, expanding box.

The core hydrodynamic outputs of a `Nyx` simulation are baryon density, temperature, and velocity. These quantities can then be post-processed, accounting for ultraviolet (UV) background and atomic physics, to yield an optical depth τ indicating the local opacity to a certain wavelength of light, in our case that of the Lyman α absorption line in neutral hydrogen. From τ one can attain a Lyα flux map $F = \exp(-\tau)$ representing the fraction of light that passes through given points in the IGM. We visualize example flux maps from low and high resolution `Nyx` simulations in Figure 2, rendering F in 2D slices with lines-of-sight along the horizontal direction.

The grid resolution per box side in `Nyx` is determined by the number N^3 of dark matter particles, and important Lyα summary statistics like the power spectrum $P(k)$ are extremely sensitive to the chosen effective resolution [13]. Unfortunately, unlike raw hydrodynamic or N-body outputs, Lyα flux accuracy is severely degraded even at the lowest wavenumbers for under-resolved simulations when compared to the reference high-resolution run, as shown in the right panel of Figure 1. As high-resolution simulations are computationally infeasible and generate overwhelming data volumes, we aim to perform effective 8x super-resolution by correcting Lyα flux F from a coarse simulation $(N = 512, L = 80\,\mathrm{Mpc}/h)$ towards that of a corresponding high resolution run $(N = 4096, L = 80\,\mathrm{Mpc}/h)$, while still representing F on the coarse grid to minimize the disk footprint of outputs.

To train our model, we feed low-resolution hydrodynamic fields (baryon density, temperature, and velocity) from a coarse simulation as input and predict the corresponding subsampled Lyα flux F from the high-resolution simulation. We use one pair of coarse and fine simulations for training, and another pair for validation, and all data is sampled from snapshots of the simulations at redshift $z = 3$. Each pair is matched by sharing identical initial conditions, up to the resolution limit in Fourier space.

To preserve the statistics of the high-resolution field when downsampling Lyα flux for use as a training target, we must use direct subsampling to guarantee that the power spectrum will be unaffected at low k (see the orange line in Figure 1) and that the probability density will be representative. This comes at the expense of aliasing at the 1-4 voxel scales, and is also not entirely rotation invariant, as there is no cell in the larger field that lies exactly in the center of a cell in the smaller field. However, all alternative methods for downsampling perturb the Lyα power spectrum too much, making them inadequate for our use-case.

3 Model Architecture and Training

Since our formulation of the coarse-to-fine task does not require actually adding resolution elements in the spatial dimensions, we choose the U-Net architecture [16] for its simplicity and strong performance in cosmological applications [6; 19]. Our U-Net employs residual convolution blocks

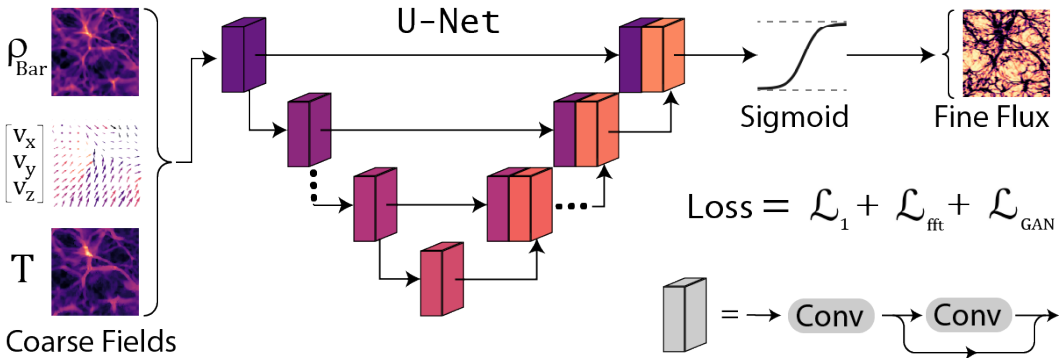

Figure 3: Our model architecture is a modified 3D U-Net with extra internal residual connections, taking the coarse hydrodynamic fields as input and predicting a corrected Lyα flux field as output.

between each down-sampling and up-sampling step, which improves the accuracy of the power spectrum and the reproduction of fine features. To constrain the output field to $F \in (0, 1)$ we use a sigmoid activation at the final layer, which performed better than ReLU in our tests.

We train our model with a weighted combination of L_1 loss and adversarial loss $\mathsf{L}_{\mathrm{GAN}}$, which is given by a multi-scale patch-based discriminator [24]. We also improve the accuracy of the power spectra further by enforcing an additional Fourier-based loss $\mathsf{L}_{\mathrm{fft}}$, which measures the difference between the Fourier transforms of the prediction and target, truncated at $k = 10$ h/Mpc. For training, we normalize our input physical fields following [6], and leave F in its natural $(0, 1)$ range. We list hyperparameters in Appendix A.

Since our model is fully convolutional, we can train on small crops (128^3 voxels) and apply on larger ones (e.g., 512^3 voxels) at inference time, depending on GPU memory limits. For full-box inference on very large simulations, we partition the volume into chunks (with some padding to minimize edge effects) and run inference in parallel, then reassemble the chunks into a contiguous array.

4 Results

After training, we evaluate our model on a separate volume of the same dimensions. In Figure 2 we plot 2D slices of the Lyα flux from our model, compared against the reference high-resolution simulation. We also plot the coarse Lyα flux given by analytical computation of τ directly from the low-resolution hydrodynamic fields, which is what one would see in the absence of our model. Our model is able to greatly improve over the coarse baseline, capturing both large and small-scale features well and closely matching the spatial patterns of the high-resolution run. The visual discrepancy between the model output and high-resolution target is only noticeable at the very finest scales.

For a quantitative assessment of our model, we focus on the Lyα power spectrum $P(k)$, arguably the most important summary statistic for Lyα flux analysis, and take the $P(k)$ in the reference high-resolution simulation as the ground truth. As shown in Figure 1 the power spectrum of our estimated field is significantly better than that of the low-resolution result, particularly over scales which will be observable in surveys like DESI ($k < 4.2$ h/Mpc). In particular, our estimated $P(k)$ achieves $<\sim 1\%$ difference for $k < 2$ h/Mpc, and $<\sim 5\%$ difference for $k < 5$ h/Mpc. This is far better than the low-resolution result which has $> 10\%$ difference for most of this range, even deviating at low k as well as high k. We observe larger divergence at scales beyond the DESI limit, as expected from using the subsampled high-resolution Lyα as a training target. However at such scales the observational uncertainty is too large to make proper use of high-resolution simulations in the first place, so we do not consider this a limitation of the model.

5 Conclusion

In this work we present a novel approach to estimating Lyα flux from coarse simulations, taking care to curate and sample a training set from paired low- and high-resolution `Nyx` simulations. We demonstrate that our approach allows accurate reconstruction of the statistics of Lyα flux, closely

matching a full-resolution simulation across the relevant scales, without adding any additional resolution elements. We thus effectively perform 8x super-resolution, but in a physically consistent and computationally practical manner, such that our model can be easily applied to large-scale cosmological volumes. We provide an illustrative demonstration of this exciting capability in Appendix B. In general, our method could be used to close the major gap in spatial scales between current simulation capabilities and upcoming observational products.

Broader Impact

We are not aware of any negative social impacts of the work presented here. While generative adversarial networks similar to the techniques used in this work can be used outside the domain of cosmology for malicious cases such as deepfakes, our architecture and adversarial loss terms consist of known techniques already present in the machine learning community and the novelty of our work is the deployment of these techniques in a practical and careful manner on cosmological modeling tasks. A potential positive impact of our method is the reduction in energy consumption due to computational savings avoiding the expensive full-resolution simulations.

Acknowledgments and Disclosure of Funding

This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.

References

[1] A. S. Almgren, J. B. Bell, M. J. Lijewski, Z. Lukić, and E. V. Andel. Nyx: A massively parallel amr code for computational cosmology. *The Astrophysical Journal*, 765(1):39, feb 2013. doi: 10.1088/0004-637x/765/1/39. URL https://doi.org/10.1088/0004-637x/765/1/39.

[2] E. Armengaud, N. Palanque-Delabrouille, C. Yèche, D. J. E. Marsh, and J. Baur. Constraining the mass of light bosonic dark matter using SDSS Lyman- forest. *Monthly Notices of the Royal Astronomical Society*, 471(4):4606–4614, 07 2017. ISSN 0035-8711. doi: 10.1093/mnras/stx1870. URL https://doi.org/10.1093/mnras/stx1870.

[3] M. Bernardini, R. Feldmann, D. Anglés-Alcázar, M. Boylan-Kolchin, J. Bullock, L. Mayer, and J. Stadel. From EMBER to FIRE: predicting high resolution baryon fields from dark matter simulations with deep learning. *Monthly Notices of the Royal Astronomical Society*, 509(1):1323–1341, 10 2021. ISSN 0035-8711. doi: 10.1093/mnras/stab3088. URL https://doi.org/10.1093/mnras/stab3088.

[4] S. Chabanier, N. Palanque-Delabrouille, C. Yèche, J.-M. Le Goff, E. Armengaud, J. Bautista, M. Blomqvist, N. Busca, K. Dawson, T. Etourneau, A. Font-Ribera, Y. Lee, H. du Mas des Bourboux, M. Pieri, J. Rich, G. Rossi, D. Schneider, and A. Slosar. The one-dimensional power spectrum from the SDSS DR14 Ly forests. *J. Cosmology Astropart. Phys.* , 2019(7):017, July 2019. doi: 10.1088/1475-7516/2019/07/017.

[5] A. e. a. DESI Collaboration. The DESI Experiment Part I: Science,Targeting, and Survey Design. *arXiv e-prints*, art. arXiv:1611.00036, Oct. 2016.

[6] P. Harrington, M. Mustafa, M. Dornfest, B. Horowitz, and Z. Lukić. Fast, high-fidelity ly forests with convolutional neural networks. *The Astrophysical Journal*, 929(2):160, apr 2022. doi: 10.3847/1538-4357/ac5faa. URL https://doi.org/10.3847/1538-4357/ac5faa.

[7] B. Horowitz, M. Dornfest, Z. Lukić, and P. Harrington. Hyphy: Deep generative conditional posterior mapping of hydrodynamical physics. *arXiv preprint arXiv:2106.12675*, 2021.

[8] V. Iršič, M. Viel, M. G. Haehnelt, J. S. Bolton, S. Cristiani, G. D. Becker, V. D'Odorico, G. Cupani, T.-S. Kim, T. A. M. Berg, S. López, S. Ellison, L. Christensen, K. D. Denney, and G. Worseck. New constraints on the free-streaming of warm dark matter from intermediate and small scale lyman- forest data. *Phys. Rev. D*, 96:023522, Jul 2017. doi: 10.1103/PhysRevD.96.023522. URL https://link.aps.org/doi/10.1103/PhysRevD.96.023522.

[9] L. Jiang, C. Zhang, M. Huang, C. Liu, J. Shi, and C. C. Loy. TSIT: A Simple and Versatile Framework for Image-to-Image Translation. In A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, editors, *Computer Vision – ECCV 2020*, pages 206–222, Cham, 2020. Springer International Publishing. ISBN 978-3-030-58580-8.

[10] D. Kodi Ramanah, T. Charnock, F. Villaescusa-Navarro, and B. D. Wandelt. Super-resolution emulator of cosmological simulations using deep physical models. *Monthly Notices of the Royal Astronomical Society*, 495(4):4227–4236, 05 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa1428. URL https://doi.org/10.1093/mnras/staa1428.

[11] Y. Li, Y. Ni, R. A. C. Croft, T. D. Matteo, S. Bird, and Y. Feng. Ai-assisted superresolution cosmological simulations. *Proceedings of the National Academy of Sciences*, 118(19):e2022038118, 2021. doi: 10.1073/pnas.2022038118. URL https://www.pnas.org/doi/abs/10.1073/pnas.2022038118.

[12] J. H. Lim and J. C. Ye. Geometric gan. *arXiv preprint arXiv:1705.02894*, 2017.

[13] Z. Lukić, C. W. Stark, P. Nugent, M. White, A. A. Meiksin, and A. Almgren. The Lyman forest in optically thin hydrodynamical simulations. *Mon. Not. Royal Astro. Soc.* , 446(4):3697–3724, Feb. 2015. doi: 10.1093/mnras/stu2377.

[14] Y. Ni, Y. Li, P. Lachance, R. A. C. Croft, T. Di Matteo, S. Bird, and Y. Feng. AI-assisted superresolution cosmological simulations – II. Halo substructures, velocities, and higher order statistics. *Monthly Notices of the Royal Astronomical Society*, 507(1):1021–1033, 07 2021. ISSN 0035-8711. doi: 10.1093/mnras/stab2113. URL https://doi.org/10.1093/mnras/stab2113.

[15] N. Palanque-Delabrouille, C. Yèche, J. Baur, C. Magneville, G. Rossi, J. Lesgourgues, A. Borde, E. Burtin, J.-M. LeGoff, J. Rich, M. Viel, and D. Weinberg. Neutrino masses and cosmology with lyman-alpha forest power spectrum. *Journal of Cosmology and Astroparticle Physics*, 2015(11):011–011, nov 2015. doi: 10.1088/1475-7516/2015/11/011. URL https://doi.org/10.1088/1475-7516/2015/11/011.

[16] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

[17] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29, 2016.

[18] U. Seljak, A. Slosar, and P. McDonald. Cosmological parameters from combining the lyman- forest with CMB, galaxy clustering and SN constraints. *Journal of Cosmology and Astroparticle Physics*, 2006 (10):014–014, oct 2006. doi: 10.1088/1475-7516/2006/10/014. URL https://doi.org/10.1088/1475-7516/2006/10/014.

[19] L. Thiele, F. Villaescusa-Navarro, D. N. Spergel, D. Nelson, and A. Pillepich. Teaching Neural Networks to Generate Fast Sunyaev-Zel'dovich Maps. *The Astrophysical Journal*, 902(2):129, Oct. 2020. doi: 10.3847/1538-4357/abb80f.

[20] T. Tröster, C. Ferguson, J. Harnois-Déraps, and I. G. McCarthy. Painting with baryons: augmenting N-body simulations with gas using deep generative models. *Monthly Notices of the Royal Astronomical Society*, 487(1):L24–L29, July 2019. doi: 10.1093/mnrasl/slz075.

[21] M. Viel, G. D. Becker, J. S. Bolton, and M. G. Haehnelt. Warm dark matter as a solution to the small scale crisis: New constraints from high redshift lyman- forest data. *Phys. Rev. D*, 88:043502, Aug 2013. doi: 10.1103/PhysRevD.88.043502. URL https://link.aps.org/doi/10.1103/PhysRevD.88.043502.

[22] M. Walther, J. Oñorbe, J. F. Hennawi, and Z. Lukić. New constraints on IGM thermal evolution from the ly forest power spectrum. *The Astrophysical Journal*, 872(1):13, feb 2019. doi: 10.3847/1538-4357/aafad1. URL https://doi.org/10.3847/1538-4357/aafad1.

[23] M. Walther, E. Armengaud, C. Ravoux, N. Palanque-Delabrouille, C. Yèche, and Z. Lukić. Simulating intergalactic gas for DESI-like small scale Lyman forest observations. *J. Cosmology Astropart. Phys.* , 2021(4):059, Apr. 2021. doi: 10.1088/1475-7516/2021/04/059.

[24] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018.

[25] C. Yèche, N. Palanque-Delabrouille, J. Baur, and H. du Mas des Bourboux. Constraints on neutrino masses from lyman-alpha forest power spectrum with BOSS and XQ-100. *Journal of Cosmology and Astroparticle Physics*, 2017(06):047–047, jun 2017. doi: 10.1088/1475-7516/2017/06/047. URL https://doi.org/10.1088/1475-7516/2017/06/047.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] **Code & data release is WIP at the moment but we plan to share as soon as we can.**
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] **We did not have time for this before the submission deadline due to computational limitations. This will appear in future publications.**
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Appendices

A Additional model details

Our multi-scale patch-based discriminator is based on the implementation from TSIT [9], a state-of-the-art image synthesis model. The adversarial loss term L_{GAN} is given by a weighted combination of hinge [12] loss and feature-matching (Feat) loss [17] to improve training stability for the U-Net. We list network hyperparameters in Table 1.

Table 1: Hyperparameters for our netowrks.

Hyperparameter	Value
Global batch size	16
Training crop size	128^3
Optimizer	Adam($\beta_1 = 0.5$, $\beta_2 = 0.999$)
Learning rate	2E-4
λ_{L1}	1500
λ_{fft}	500
λ_{Feat}	100
Number of U-Net layers	4 levels, 2 layers per level
Multi-scale discriminators	2
Discriminator layers	4

Our model takes around 25 hours to train on 16 A100 GPUs.

B Large-scale inference

We demonstrate the preliminary use of our model for synthesizing Lyα on large-scale cosmological volumes by applying the trained network to a very large box of size 600 Mpc/h. We do not run a dedicated 600 Mpc/h simulation for this demonstration, but rather re-purpose an existing simulation from an unrelated investigation. Unfortunately this means there is a mismatch in resolution with respect to our training and validation data; namely, the coarse inputs we use for training and validation have an effective resolution of $\sim$156 kpc/h while the large volume has an effective resolution of $\sim$98 kpc/h. Consequently, to match resolution we must downsample the data from the large box by a factor of $\sim$0.63, which we do so in chunks due to computational limitations, and this process introduces minor artifacts.

Nevertheless we show a visualization of our model output in Figure 4, which renders a 2D slice similar to those in Figure 2 but covering a much greater spatial extent. The data and computational savings enabled by our model in this context are enormous – we estimate running a high-resolution simulation at that scale would require $\sim$1 billion CPU hours and the outputs would consume $\sim$700 TBs of disk space. By correcting the Lyα flux on the coarse grid, we are able to reduce the disk footprint of the output Lyα flux to 200GB, several orders of magnitude lower. The physical volume used to train our model is a factor of $\sim$422 smaller than the volume we demonstrate inference on here.

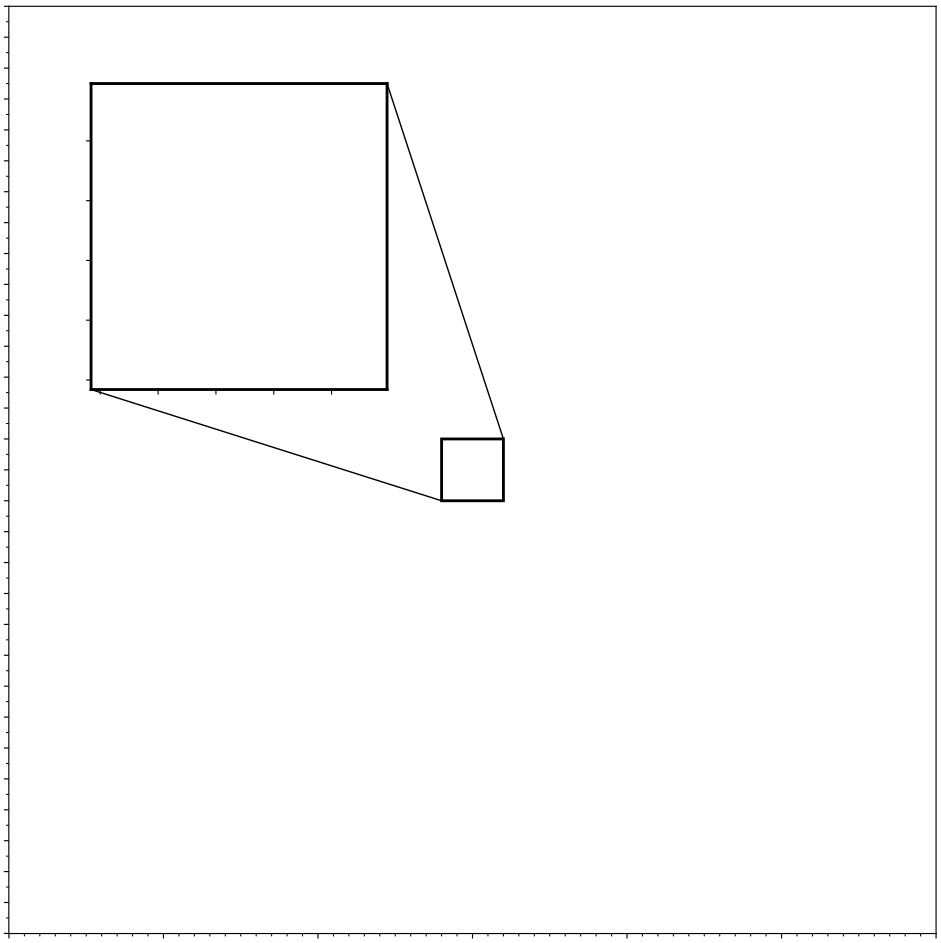

Figure 4: Visualization of the Lyα prediction from our model when running inference on a large-scale volume with side-length 600 Mpc/h. The inset shows a zoom on a region roughly 40 Mpc/h in size, which is nearly the size of our training volume.