
Neural 3D Reconstruction of 21-cm Tomographic Data

Nashwan Sabti

William H. Miller III Department of Physics and Astronomy
Johns Hopkins University
Baltimore, MD 21218, USA
nash.sabti@gmail.com

Ram Reddy

Weinberg Institute for Theoretical Physics
The University of Texas at Austin
Austin, TX 78712, USA
ramreddy@utexas.edu

Julian B. Muñoz

Department of Astronomy
The University of Texas at Austin
Austin, TX 78712, USA
julianbmunoz@utexas.edu

Siddharth Mishra-Sharma


The NSF AI Institute for Artificial Intelligence and Fundamental Interactions
Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
Department of Physics, Harvard University, Cambridge, MA 02138, USA
smsharma@mit.edu

Taewook Youn

Laboratory for Elementary Particle Physics, Cornell University, Ithaca, NY 14853, USA
Department of Physics, Korea University Seoul 02841, Republic of Korea
taewook.youn@cornell.edu

Abstract

Analyses of the cosmic 21-cm signal are hampered by astrophysical foregrounds that are far stronger than the signal itself. These foregrounds, typically confined to a wedge-shaped region in Fourier space, often necessitate the removal of a vast majority of modes, thereby degrading the quality of the data anisotropically. To address this challenge, we introduce a novel deep generative model based on stochastic interpolants to reconstruct the 21-cm data lost to wedge filtering. Our method leverages the non-Gaussian nature of the 21-cm signal to effectively map wedge-filtered 3D lightcones to samples from the conditional distribution of wedge-recovered lightcones. We demonstrate how our method is able to restore spatial information effectively, considering both varying cosmological initial conditions and astrophysics, potentially offering new opportunities for 21-cm image analyses.

 Code, pre-trained models, and scripts for making plots in this paper can be found here.

1 Introduction

The high-redshift 21-cm signal serves as a promising tool for investigating the early Universe during the epochs of cosmic dawn and reionization [1–4]. Its detection, however, is significantly hindered by the presence of astrophysical foregrounds that are several orders of magnitude stronger than the

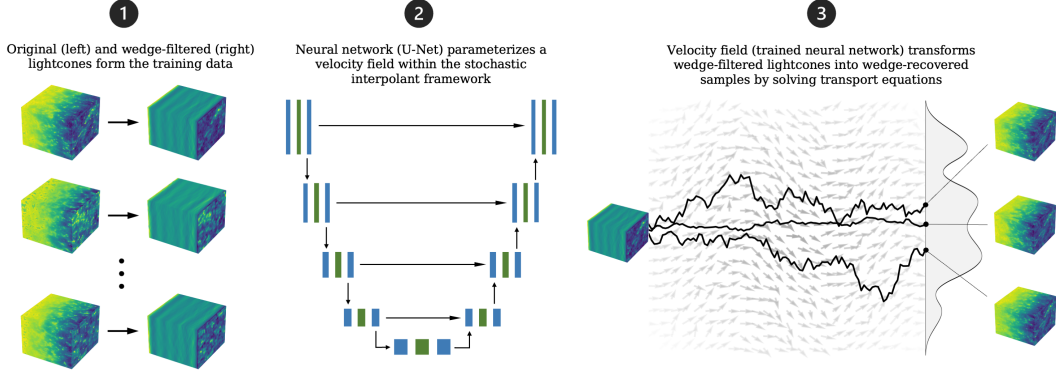


Figure 1: A schematic overview of our approach to develop a generative model aimed at reconstructing tomographic 21-cm maps at the field level from wedge-removed data. This method encompasses several stages, including data generation, neural-network training, and solving transport equations. **Part 1:** We generate 21-cm lightcones using simulations and remove Fourier modes that are contaminated by foregrounds, resulting in pairs of ‘original’ and wedge-filtered lightcones. **Part 2:** These data pairs are then used within the stochastic interpolant framework, where a neural network is trained to learn a time-dependent velocity field that can dynamically transport a wedge-filtered lightcone to a wedge-recovered lightcone. **Part 3:** A stochastic differential equation is solved, with the trained neural network acting as a drift function, to push an input wedge-filtered lightcone to a sample from the distribution of wedge-recovered lightcones conditioned on this input.

signal itself, see e.g. Refs. [5, 6]. These foregrounds interact with an interferometer’s frequency-dependent response to such emissions, introducing an anisotropic component in Fourier space that occupies a wedge-shaped region [7–14]. Consequently, only a small window remains where the 21-cm signal can be observed without substantial foreground contamination. While completely removing the foreground-dominated regions in Fourier space is always possible, this would clearly have a detrimental effect on the signal-to-noise of the information extracted with this data [15]. Moreover, if we are aiming for 3D tomographic imaging of the 21-cm signal, it is essential to retain as many Fourier modes as possible. This is because wedge filtering leads to a distortion of the original image beyond simple blurring, making an interpretation of what we see even more complicated [16].

To address the challenge posed by foreground contamination, a promising set of approaches based on machine learning has emerged [17–27]. One such approach leverages neural networks to identify correlations between Fourier modes inside and outside of the wedge region [23], which is largely possible due to the inherent non-Gaussian nature of the 21-cm signal [28–32]. Then, under certain conditions, it becomes feasible for the network to reconstruct a significant portion of the modes that were originally lost in the filtering process. A major advantage of this approach is that it starts with relatively uncontaminated 21-cm data and does not require precise noise modeling.

In this study, we build upon the work of Ref. [23] by exploring a generative method based on ‘stochastic interpolants’ [33, 34] that are able to continuously bridge two arbitrary probability densities. Within our context, this method makes it possible to connect the distribution of wedge-filtered 21-cm data to that of wedge-recovered data. The stochastic nature of this process allows for the same wedge-filtered 21-cm input to be transformed into a distribution of possible wedge-recovered lightcones. In this way, we can sample from this distribution and obtain realizations of wedge-recovered solutions — a generative model. We summarize the main differences between our approach and previous works: (1) Our method utilizes a generative model to obtain samples from the conditional distribution of wedge-recovered lightcones, (2) we do not consider binarized maps as in Ref. [23], but rather allow for the recovery of the full 21-cm brightness temperature, (3) cosmological initial conditions and a range of astrophysical parameters are varied in the data generation process, which results in the final distribution being marginalized over these effects, and (4) lightcones are considered instead of coeval boxes (where all voxels share the same redshift), as they more closely resemble real data and enable recovery over a range of redshifts. We provide an illustrative overview of our procedure in Fig. 1.

2 Methodology

2.1 Dataset and Forward Model

We make use of the code 21cmFAST [35, 36] to generate lightcones of the 21-cm brightness temperature. These lightcones have a volume of $(512)^3 \text{ Mpc}^3$ and feature a resolution of 4 Mpc along each axis, resulting in 128 voxels per axis. We post-process each box such that the line-of-sight axis has a redshift range $z \in [8.9, 11.3]$ and is centered at $z \approx 10$. We keep the default settings of 21cmFAST, modifying only cosmological initial conditions and a few key astrophysical parameters for each simulation. Specifically, we vary the following parameters within logflat priors whose limits are informed by the wedge-removed power-spectrum forecast done in Ref. [37]: The fraction of galactic gas in stars $\log_{10} f_{s,10} \in [-1.45, -1.15]$, the escape fraction of ionizing photons $\log_{10} f_{\text{esc},10} \in [-1.17, -0.84]$, and the specific X-ray luminosity $\log_{10} l_X \in [40.45, 40.55]$. We generate a total of 25,000 lightcones, each featuring unique cosmological initial conditions and astrophysical parameters as previously mentioned. We augment the data by rotating the boxes four times by 90° along the line-of-sight axis, obtaining a total of 100,000 lightcones. We then process the lightcones by removing modes dominated by intrinsic foregrounds and those within the foreground wedge. We employ the notation $k_{\parallel} \equiv k_z$ and $k_{\perp} \equiv (k_x^2 + k_y^2)^{1/2}$ to represent the Fourier modes along the line-of-sight and those perpendicular to it, respectively. To address intrinsic foregrounds, we remove all modes with $|k_{\parallel}| < 0.05 \text{ Mpc}^{-1}$ [10, 24]. We define the wedge region as $|k_{\parallel}| < |k_{\perp}| \tan \psi$, where ψ is the wedge angle, set equal to $\psi = 65^\circ$ for demonstrative purposes. Throughout the text, we will refer the removal of both types of foregrounds as ‘wedge filtering’. After nullifying the relevant modes, we apply the inverse Fourier transform to obtain the wedge-filtered lightcones back in image space.

2.2 Stochastic Interpolants

Our dataset consists of pairs of 21-cm lightcones and the corresponding wedge-filtered lightcones, $(x_{21}, x_{21}^{\text{wf}}) \sim \rho(x_{21}, x_{21}^{\text{wf}})$, defining our target and base quantities, respectively. The stochastic interpolant framework [38, 34] provides a way to dynamically transport a base measure (x_{21}^{wf}) onto a target (x_{21}) , allowing to sample from the conditional probability density $\rho(x_{21}|x_{21}^{\text{wf}})$. First, we construct the stochastic interpolant $x(t)$ that bridges two probability densities as follows [39]:

$$x(t) = \alpha(t)x_{21}^{\text{wf}} + \beta(t)x_{21} + \sigma(t)W(t), \quad (1)$$

where $t \in [0, 1]$ is a time variable, $\alpha(t) = \sigma(t) = 1 - t$, $\beta(t) = t^2$, and W is a Wiener process. Next, it can be shown that the velocity field corresponding to the interpolant in Eq. (1) is given by [39]:

$$v(t, x_{21}, x_{21}^{\text{wf}}) = \dot{\alpha}(t)x_{21}^{\text{wf}} + \dot{\beta}(t)x_{21} + \dot{\sigma}(t)W(t), \quad (2)$$

with the dot denoting derivation with respect to time t . We approximate this function with a neural network $\hat{v}(t, x(t))$ by minimizing the following objective with stochastic gradient descent [39]:

$$L[\hat{v}] = \int_0^1 dt \mathbb{E} \left[|\hat{v}(t, x(t)) - v(t, x_{21}, x_{21}^{\text{wf}})|^2 \right], \quad (3)$$

where the expectation value is taken over the data pairs $(x_{21}, x_{21}^{\text{wf}}) \sim \rho(x_{21}, x_{21}^{\text{wf}})$ and W . The time integral can be simply evaluated through Monte Carlo sampling. The learned velocity field will then act as a drift function in the stochastic differential equation (SDE) [40],

$$dX(t) = \hat{v}(t, X(t))dt + \sigma(t)dW(t), \quad (4)$$

whose solutions are such that $X(t=1) \sim \rho(x_{21}|x_{21}^{\text{wf}})$. We generate samples by discretizing the time interval $t \in [0, 1]$ into 500 steps along which we solve Eq. (4) using Heun’s second-order method.

2.3 Neural Network Architecture

The backbone of our architecture consists of convolutional neural networks, with the primary objective to represent the velocity field in Eq. (2). First, we take a data pair $(x_{21}, x_{21}^{\text{wf}})$ and assign a random time t sampled from a uniform distribution $\mathcal{U}(0, 1)$ to construct the stochastic interpolant $x(t)$ as in Eq. (1). These two quantities, t and $x(t)$, are the inputs to our model. We project the time variable




Figure 2: Reconstruction of a wedge-filtered lightcone over time. Top and bottom rows show a slice perpendicular to ($z \approx 9.1$) and along the line-of-sight axis, respectively. The leftmost column is a wedge-filtered input (x_{21}^{wf}) at time $t = 0$. Over time, this input gets transported to a sample from the conditional distribution $\rho(x_{21}|x_{21}^{wf})$ at $t = 1$, which can be compared to the ground truth in the sixth column. The last column shows the average of 2000 sampled lightcones.

t onto an embedding space of dimension 48 using sinusoidal functions [41], and then process it through a multi-layer perceptron consisting of two linear layers of hidden dimension 192 with SiLU activation functions. To aid in capturing temporal dependencies, the model is conditioned on this time variable by adding a linear projection of the embedded t to the output of a convolutional block as described below. Then, we put the stochastic interpolant through the main structure of our model, which is a 3D U-Net neural network [42]. The U-Net consists of 4 encoding blocks, followed by a bottom block, and finally 4 decoder blocks. Each of these blocks is made out of two pre-activation ‘convolutional blocks’, which is a sequence of an instance normalization layer, a ReLU activation, and a 3D convolutional layer. The time embedding is added to the output of the first of such two convolutional blocks. The encoding blocks are downsampled using a max pooling operation, while before each decoding block an upsampling procedure is performed using transposed convolutional layers. Skip connections between encoding and decoding blocks at each of the 4 levels of the U-Net are added to the output of the decoding blocks. Lastly, the output of the decoder part of the U-Net is processed with a final convolutional block.

2.4 Training and Evaluation

The data is randomly split into a 90%/10% train/test set and standardized using the mean and standard deviation of the training set. We train the model using the voxel-wise mean-squared-error (MSE) loss between the network output and the ground truth, as defined in Eq. (3). The model is optimized using the Adam optimizer [43], with an initial learning rate of 10^{-3} that is halved every two epochs. We use a batch size of 3 and train for a total of five epochs. The model’s performance is evaluated by comparing the MSE loss of wedge-recovered lightcones to the identity loss (MSE loss between the original and wedge-filtered data) and by comparing the voxel distribution of the wedge-recovered lightcones to that of the ground truth. Our base model has approximately 43M parameters and training takes roughly three days on an 80GB Nvidia A100 GPU.

3 Results

An example of a sample obtained by solving the SDE in Eq. (4) is presented in Fig. 2. The top and bottom rows show slices along the transverse and line-of-sight directions, respectively. The leftmost column displays the wedge-filtered image at $t = 0$. The subsequent columns illustrate the image’s evolution as it transforms into a wedge-recovered sample by $t = 1$. This can be compared to the ground-truth image in the sixth column. We repeat this process 2000 times and display the average of these samples in the last column. From this, we observe that large-scale correlations are reasonably well reconstructed, while small-scale features exhibit more randomness (that get averaged out in the last column). We further illustrate the reconstruction quality by plotting power spectra in Fig. 3. The left panel shows the 2D power spectrum for the reconstructed sample in Fig. 2 relative to the ground truth. Note that the Fourier modes to the right of the black dashed curve were initially zeroed out. Yet, we can recover the power within tens of percents in that region, even deep inside the

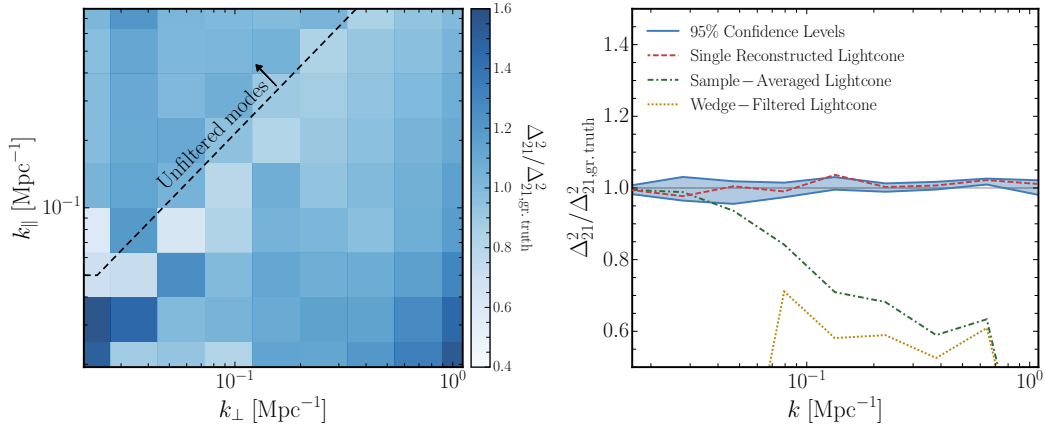


Figure 3: Recovered power spectra of the 21-cm signal. **Left:** 2D power spectrum in $(k_{\parallel}, k_{\perp})$ -space of the wedge-recovered sample in Fig. 2 at $t = 1$ relative to the ground truth. **Right:** Spherical power spectra. We generate 2000 wedge-recovered lightcone samples, compute their power spectra, and plot the 95% confidence levels in blue. The red, green, and yellow curves show the power spectrum of a single sample, the sample-averaged lightcone, and the wedge-filtered lightcone, respectively.

foreground wedge. The right panel of Fig. 3 shows the spherical power spectrum and highlights a key feature of this generative model. Namely, we can compute a power spectrum for each of the 2000 generated samples and obtain a distribution as a function of spherical wavenumber $k = (k_{\parallel}^2 + k_{\perp}^2)^{1/2}$. This then enables us to set confidence intervals on the retrieved power spectrum (blue solid curves). Our reconstructed power spectrum is within approximately 5% of the ground truth for all k values considered. This is in contrast to the sample-averaged lightcone case (green dash-dotted curve in this plot) and the wedge-filtered lightcone (yellow dotted curve).

4 Conclusions

We introduced a generative model designed to reconstruct 21-cm lightcones that were distorted due to foreground-wedge filtering. Our approach leverages the framework of stochastic interpolants, where a 3D neural network is trained to represent a velocity field that transports wedge-filtered lightcones onto wedge-recovered lightcones. This allowed us to sample from the distribution of wedge-recovered lightcones conditioned on an input wedge-filtered lightcone. The results presented here indicate that our model can recover a substantial portion of the lost data, both in image and Fourier space, producing reconstructed lightcones and power spectra that closely resemble the original.

References

- [1] S. Furlanetto, S. P. Oh, and F. Briggs, *Phys. Rept.* **433**, 181 (2006), arXiv:astro-ph/0608032 .
- [2] M. F. Morales and J. S. B. Wyithe, *Annual Review of Astronomy and Astrophysics* **48**, 127–171 (2010).
- [3] J. R. Pritchard and A. Loeb, *Rept. Prog. Phys.* **75**, 086901 (2012), arXiv:1109.6012 [astro-ph.CO] .
- [4] A. Liu and J. R. Shaw, *Publ. Astron. Soc. Pac.* **132**, 062001 (2020), arXiv:1907.08211 [astro-ph.IM] .
- [5] M. G. Santos, A. Cooray, and L. Knox, *Astrophys. J.* **625**, 575 (2005), arXiv:astro-ph/0408515 .
- [6] J. D. Bowman, M. F. Morales, and J. N. Hewitt, *Astrophys. J.* **695**, 183 (2009), arXiv:0807.3956 [astro-ph] .

- [7] M. F. Morales, B. Hazelton, I. Sullivan, and A. Beardsley, *The Astrophysical Journal* **752**, 137 (2012).
- [8] A. R. Parsons, J. C. Pober, J. E. Aguirre, C. L. Carilli, D. C. Jacobs, *et al.*, *The Astrophysical Journal* **756**, 165 (2012).
- [9] H. Vedantham, N. Udaya Shankar, and R. Subrahmanyan, *The Astrophysical Journal* **745**, 176 (2012).
- [10] J. C. Pober *et al.*, *Astrophys. J. Lett.* **768**, L36 (2013), arXiv:1301.7099 [astro-ph.CO] .
- [11] B. J. Hazelton, M. F. Morales, and I. S. Sullivan, *Astrophys. J.* **770**, 156 (2013), arXiv:1301.3126 [astro-ph.IM] .
- [12] N. Thyagarajan *et al.*, *Astrophys. J.* **776**, 6 (2013), arXiv:1308.0565 [astro-ph.CO] .
- [13] A. Liu, A. R. Parsons, and C. M. Trott, *Phys. Rev. D* **90**, 023018 (2014), arXiv:1404.2596 [astro-ph.CO] .
- [14] A. Liu, A. R. Parsons, and C. M. Trott, *Phys. Rev. D* **90**, 023019 (2014), arXiv:1404.4372 [astro-ph.CO] .
- [15] H.-J. Seo and C. M. Hirata, *Mon. Not. Roy. Astron. Soc.* **456**, 3142 (2016), arXiv:1508.06503 [astro-ph.CO] .
- [16] A. P. Beardsley, M. F. Morales, A. Lidz, M. Malloy, and P. M. Sutter, *Astrophys. J.* **800**, 128 (2015), arXiv:1410.5427 [astro-ph.CO] .
- [17] N. Gillet, A. Mesinger, B. Greig, A. Liu, and G. Ucci, *Mon. Not. Roy. Astron. Soc.* **484**, 282 (2019), arXiv:1805.02699 [astro-ph.CO] .
- [18] S. Hassan, S. Andrianomena, and C. Doughty, *Mon. Not. Roy. Astron. Soc.* **494**, 5761 (2020), arXiv:1907.07787 [astro-ph.CO] .
- [19] W. Li, H. Xu, Z. Ma, R. Zhu, D. Hu, *et al.*, *Mon. Not. Roy. Astron. Soc.* **485**, 2628 (2019), arXiv:1902.09278 [astro-ph.IM] .
- [20] T. L. Makinen, L. Lancaster, F. Villaescusa-Navarro, P. Melchior, S. Ho, *et al.*, *JCAP* **04**, 081 (2021), arXiv:2010.15843 [astro-ph.CO] .
- [21] P. Villanueva-Domingo and F. Villaescusa-Navarro, *Astrophys. J.* **907**, 44 (2021), arXiv:2006.14305 [astro-ph.CO] .
- [22] D. Prelogović, A. Mesinger, S. Murray, G. Fiameni, and N. Gillet, *Mon. Not. Roy. Astron. Soc.* **509**, 3852 (2021), arXiv:2107.00018 [astro-ph.CO] .
- [23] S. Gagnon-Hartman, Y. Cui, A. Liu, S. Ravanbakhsh, and J. Kennedy, *Mon. Not. Roy. Astron. Soc.* **504**, 4716 (2021), [Erratum: *Mon. Not. Roy. Astron. Soc.* 529, 2539–2542 (2024)], arXiv:2102.08382 [astro-ph.CO] .
- [24] J. Kennedy, J. C. Carr, S. Gagnon-Hartman, A. Liu, J. Mirocha, *et al.*, *Mon. Not. Roy. Astron. Soc.* **529**, 3684 (2024), arXiv:2308.09740 [astro-ph.CO] .
- [25] F. Shi, H. Chang, L. Zhang, H. Shan, J. Zhang, *et al.*, *Phys. Rev. D* **109**, 063509 (2024), arXiv:2310.06518 [astro-ph.CO] .
- [26] M. Bianco, S. K. Giri, D. Prelogović, T. Chen, F. G. Mertens, *et al.*, *Mon. Not. Roy. Astron. Soc.* **528**, 5212 (2024), arXiv:2304.02661 [astro-ph.IM] .
- [27] F. G. Mertens, J. Bobin, and I. P. Carucci, *Mon. Not. Roy. Astron. Soc.* **527**, 3517 (2023), arXiv:2307.13545 [astro-ph.CO] .
- [28] A. Gorce and J. R. Pritchard, *Mon. Not. Roy. Astron. Soc.* **489**, 1321 (2019), arXiv:1903.11402 [astro-ph.CO] .

- [29] H. Shimabukuro, S. Yoshiura, K. Takahashi, S. Yokoyama, and K. Ichiki, *Mon. Not. Roy. Astron. Soc.* **458**, 3003 (2016), arXiv:1507.01335 [astro-ph.CO] .
- [30] S. Majumdar, J. R. Pritchard, R. Mondal, C. A. Watkinson, S. Bharadwaj, *et al.*, *Mon. Not. Roy. Astron. Soc.* **476**, 4007 (2018), arXiv:1708.08458 [astro-ph.CO] .
- [31] C. A. Watkinson, S. K. Giri, H. E. Ross, K. L. Dixon, I. T. Iliev, *et al.*, *Mon. Not. Roy. Astron. Soc.* **482**, 2653 (2019), arXiv:1808.02372 [astro-ph.CO] .
- [32] A. Hutter, C. A. Watkinson, J. Seiler, P. Dayal, M. Sinha, *et al.*, *Mon. Not. Roy. Astron. Soc.* **492**, 653 (2020), arXiv:1907.04342 [astro-ph.CO] .
- [33] M. S. Alberg, N. M. Boffi, and E. Vanden-Eijnden, (2023), arXiv:2303.08797 [cs.LG] .
- [34] M. S. Alberg, M. Goldstein, N. M. Boffi, R. Ranganath, and E. Vanden-Eijnden, (2023), arXiv:2310.03725 [cs.LG] .
- [35] A. Mesinger, S. Furlanetto, and R. Cen, *Mon. Not. Roy. Astron. Soc.* **411**, 955 (2011), arXiv:1003.3878 [astro-ph.CO] .
- [36] S. G. Murray, B. Greig, A. Mesinger, J. B. Muñoz, Y. Qin, *et al.*, *J. Open Source Softw.* **5**, 2582 (2020), arXiv:2010.15121 [astro-ph.IM] .
- [37] C. A. Mason, J. B. Muñoz, B. Greig, A. Mesinger, and J. Park, (2022), 10.1093/mnras/stad2145, arXiv:2212.09797 [astro-ph.CO] .
- [38] M. S. Alberg and E. Vanden-Eijnden, (2023), arXiv:2209.15571 [cs.LG] .
- [39] Y. Chen, M. Goldstein, M. Hua, M. S. Alberg, N. M. Boffi, and E. Vanden-Eijnden, (2024), arXiv:2403.13724 [cs.LG] .
- [40] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, *et al.*, (2021), arXiv:2011.13456 [cs.LG] .
- [41] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” (2023), arXiv:1706.03762 [cs.CL] .
- [42] O. Ronneberger, P. Fischer, and T. Brox, (2015), arXiv:1505.04597 [cs.CV] .
- [43] D. P. Kingma and J. Ba, (2017), arXiv:1412.6980 [cs.LG] .