
Cosmological super-resolution of the 21-cm signal

Simon Pochinda
Cavendish Laboratory
University of Cambridge
J. J. Thomson Avenue, Cambridge,
CB3 0HE, UK
sp2053@cam.ac.uk

Jiten Dhandha
Institute of Astronomy
University of Cambridge
Madingley Road, Cambridge,
CB3 0HA, UK

Anastasia Fialkov
Institute of Astronomy
University of Cambridge
Madingley Road, Cambridge,
CB3 0HA, UK

Eloy de Lera Acedo
Cavendish Laboratory
University of Cambridge
J. J. Thomson Avenue, Cambridge,
CB3 0HE, UK

Abstract

In this study, we train score-based diffusion models to super-resolve gigaparsec-scale cosmological simulations of the 21-cm signal. We examine the impact of network and training dataset size on model performance, demonstrating that a single simulation is sufficient for a model to learn the super-resolution task regardless of the initial conditions. Our best-performing model achieves pixelwise RMSE ~ 0.57 mK and dimensionless power spectrum residuals ranging from $10^{-2} - 10^{-1}$ mK² for 128^3 , 256^3 and 512^3 voxel simulation volumes at redshift 10. The super-resolution network ultimately allows us to utilize all spatial scales covered by the SKA1-Low instrument, and could in future be employed to help constrain the astrophysics of the early Universe.

1 Introduction

The Square Kilometre Array [5, 38, SKA] promises tomographic maps of the 21-cm line of neutral hydrogen along with detailed measurements of summary statistics [12, 31, 1, 27]. The SKA will consist of two phases SKA1 and SKA2 each covering low (50 – 350 MHz, $z \approx 3 - 27$) and mid (350 – 15400 MHz) frequencies. Here, we focus on the spatial scales probed by SKA1-Low, which will have frequency-dependent primary station beams covering a field-of-view of approximately ~ 1 cGpc at 100 MHz for $z \sim 10$ [23, Sec. 5.3]. Summary statistics of these large cosmological volumes, such as the 21-cm dimensionless power spectrum, will provide valuable information that can be used to constrain astrophysical models of the early Universe [e.g. 17, 15, 2, 29]. However, to fully utilize the information contained in the gigaparsec-scales covered by the SKA1-Low, it is necessary to simulate correspondingly large cosmic volumes at high resolution [21, 28]. This is complicated as producing astrophysical parameter constraints typically requires thousands of simulations to train reliable emulators for inference [2, 29] - which is not feasible, even with state-of-the-art semi-numerical codes.

The emergence of transformer-based architectures [45] has improved Convolutional Neural Networks (CNNs) substantially, where diffusion models [42, 18, 43] now beat out Generative Adversarial Networks [GANs, 16] in image synthesis [6]. Diffusion models have been employed for a range of tasks including text-to-image generation [e.g. Stable Diffusion, 35, 30], super-resolution [SR, 37, 25], image compression [48] among others. Diffusion models have also been adapted for various applications in astronomy such as galaxy image simulation [41], gravitational lens reconstruction

[19], sampling the initial conditions of the Universe [24], super-resolving cosmological density fields [36], and even conditional generation of 2D 21-cm differential brightness temperature images [49]. Based on their superior performance in image generation tasks and successful application in the aforementioned cosmological problems, we likewise adopt diffusion models in this study.

In this proof of concept work, we train score-based conditional diffusion models [43] for super-resolving large 3D cosmological high-resolution (HR) simulations of the 21-cm differential brightness temperature. The SR emulator is aided by conditional inputs consisting of a low-resolution (LR) 21-cm brightness temperature box, and simulation initial conditions. The generative model will help us overcome the computational constraints associated with large-scale HR simulations, which will allow us cover the spatial scales (k -modes) probed by the SKA1-Low.

2 Methods

2.1 Simulations of the 21-cm line

The simulations used in this SR study are semi-numerical simulations of the 21-cm differential brightness temperature generated by the code 21cmSPACE [46, 8–11, 3, 7, 32–34, 26, 13, 14, 39]. The dataset used in this work consists of 80 HR simulations of 768^3 cMpc³ volumes with 256^3 voxels each 3 cMpc across. While the simulations cover integer redshifts from $z = 6 - 49$, we focus on $z = 10$ and leave redshift conditional generation to a future study. The simulations are all run with the same astrophysical and cosmological parameters, but different initial conditions. These initial conditions consists of matter overdensity fields and baryon-dark matter relative velocity fields also spanning 768^3 cMpc³ with 256^3 voxels at 3 cMpc resolution. The conditional LR input is generated from the HR simulations by a trilinear downsampling to 12 cMpc voxels, followed by a trilinear upsampling to match the other input dimensions. The 80 simulations in the dataset took ~ 34000 CPU hours to complete (20 CPUs for ~ 20 hours for each simulation).

Additionally, we run a single large simulation to evaluate the model on the large spatial scales covered by SKA1-Low. This large simulation covers a 1536^3 cMpc³ volume with 512^3 voxels at 3 cMpc resolution from $z = 6 - 49$. In our analysis of the performance across spatial scales, we likewise choose to focus on $z = 10$. The large simulation took ~ 45000 CPU hours (128 CPUs for ~ 2 weeks runtime) to run.

2.2 Super-resolution network architecture and training procedure

We train score-based Variance Preserving Stochastic Differential Equation (VP SDE) diffusion models using the noise-weighted training objective as described in Song et al. [43]. We utilize the improved DDPM++ architecture for the denoising network, as implemented by Karras et al. [20], adapting it to handle 3D inputs. The models in this work use 4 residual blocks per resolution, resolution multipliers 1, 2, 4, 8, and 16 for the number of channels, and self-attention at the 8^3 and bottleneck resolution [45, 47]. For sampling, we employ the Euler-Maruyama algorithm with 100 steps for the reverse diffusion process. While we use the stochastic Euler-Maruyama sampler, the simulated data in this work does not include any inherent stochasticity. Thus, deterministic sampling methods like ODE-based solvers could also be explored for inference. However, 21cmSPACE can model stochastic processes which might benefit from the noise injected at every reverse diffusion step. Hence our choice of stochastic sampling is compatible with potential future extensions including stochastic processes in the simulations.

The training procedure for an epoch can be broken down into the following: A 256^3 HR simulation is drawn from the training dataset, along with the corresponding initial conditions. To accelerate training, each HR cube is split into 8×128^3 smaller cubes to train on two batches of 4 simulations. The conditional cubes are normalized by subtracting their mean and dividing by their standard deviation. As the statistics of the HR cube are unknown in a realistic scenario, the HR target is normalized to the mean and standard deviation of the downsampled LR cube. Input cubes are augmented by applying one of 24 unique rotations, increasing the effective size of our training dataset. The HR cube is attenuated according to the drawn noise level. The initial conditions and LR input are then concatenated to the noisy HR cube. A forward pass through the model is performed, followed by backpropagation to update the model weights.




Figure 1: *Left*: Models of varying network size ($\text{CBM} \in \{4, 8\}$) trained on 1.25%/35%/70% (1/28/56 models) of the dataset. *Right*: RMSE distribution of the model with $\text{CBM} = 8$ and trained on 1.25% of the dataset.

We train models using the Adam optimizer [22] with an initial learning rate of 10^{-3} . After the first 1000 epochs, the learning rate is reduced to 10^{-4} . Validation checks are conducted every 50 epochs after the model passes 1000 epochs of training. During validation, a 256^3 simulation and initial conditions are drawn from the validation set, and a sample from the model is generated. A checkpoint is saved if the pixelwise Root Mean Squared Error (RMSE) of the validation sample reaches a new minimum. The model is trained for ~ 12 hours on four NVIDIA A100-SXM-80GB GPUs.

3 Results

3.1 Effect of Network and Training Dataset Size

The dataset is split into three different training, validation, and test percentage configurations i.e.: (training, validation, test)% = (70, 20, 10)%, (35, 55, 10)%, and (1.25, 88.75, 10)%. In terms of number of simulations for training, validation and testing, this corresponds to (56, 16, 8) simulations, (28, 44, 8) simulations, and (1, 71, 8) simulations respectively. Additionally, we train two network configurations – a small network with a channel base multiplier (CBM) of 4, and a large network with a $\text{CBM} = 8$. With the two different network sizes and the three different dataset splits, a total of 6 models are trained. This setup allows us to examine how the size of the training dataset and the network architecture impact model performance. Each model is referred to according to the following convention: e.g. DM8.1 is the diffusion model (DM) with $\text{CBM} = 8$ and a rounded 1.25% data for training. For the post-training analysis we split the 10% (8 models) test data into 64×128^3 cubes, and rotate each simulation 24 unique times to effectively evaluate the network on 1536×128^3 boxes.

The left-hand side of Figure 1 shows the 95 percentile confidence interval (CI) for the pixelwise-RMSE across combinations of network size (CBM) and training dataset size. While the pixelwise-RMSE decreases with network size (from ~ 1.2 mK to ~ 0.6 mK), the overall noise level remains far below the error margins of current 21-cm experiments specialized in measuring the global 21-cm signal, such as REACH [4, 25 mK RMSE] and SARAS 3 [40, 213mK RMSE]. Models with $\text{CBM} = 4$ showed higher variance in the pixelwise-RMSE with the training dataset size, however all 95% CIs overlap significantly, showing no strong dependency on the size of the training dataset. For the larger networks ($\text{CBM} = 8$), the pixelwise-RMSE variance is smaller and more stable across training dataset size, suggesting the network performance is independent of the amount of training data. Hence, training a SR network for a set of astrophysical parameters is computationally feasible as only a single simulation is required to learn the high-redshift astrophysics regardless of the initial conditions. The right-hand side of Figure 1 shows the pixelwise-RMSE distribution for the 1536 test models sampled by DM8.1 along with 68% and 95% CIs. The majority of samples (99.15%) have $\text{RMSE} < 0.63$ mK while just a few outliers (0.85%) have $\text{RMSE} \approx 1.75$ mK.

While learning the super-resolution task from a single training simulation is a significant result, it is worth noting that each simulation contains millions of voxels. The physical size of the training




Figure 2: Row 1 shows slices from the 21-cm HR 512^3 , 256^3 , and 128^3 targets. Row 2 shows slices from sampled models at each scale using DM8.1. Each sample is independently generated from the model, rather than being zoomed-in sections of the largest cube. Row 3 presents histograms of the voxel distributions, and Row 4 compares the 21-cm dimensionless power spectrum of HR targets and SR predictions across different spatial scales.

simulations is large enough to neglect cosmic variance, which means the summary statistics of the simulations are entirely determined by the astrophysical and cosmological parameters. Since these parameters are identical for all simulations, this may explain why a single simulation is sufficient to train a super-resolution network.

3.2 Super-resolving cosmological volumes at different scales

Following the evaluation of model and training dataset sizes, we tested the best-performing model, DM8.1, across different spatial scales. For this analysis we used a large 1536^3 cMpc 3 simulation as described in Section 2.1. Figure 2 shows simulation and model slices and statistics for 512^3 , 256^3 , and 128^3 spatial scales. Specifically, the rows show target 21-cm HR simulation slices, sampled SR model slices, histograms of pixel probability distribution functions (PDF), and 21-cm dimensionless power spectra. It is worth emphasizing that the SR output in each column of Figure 2 is independently sampled from the model (i.e. the SR row is not simply crop-outs of the larger box). On the largest

scale DM8.1 achieved an $\text{RMSE}_{512^3} = 0.57_{-0.02}^{+0.02}$ mK on 24 rotations of the simulated box. To test intermediate scales, we further divided the large simulation into 8×256^3 cubes, on which DM8.1 achieved an $\text{RMSE}_{256^3} = 0.56_{-0.01}^{+0.03}$ mK. Finally, on the smallest scale of 128^3 (also previously shown in Figure 1), DM8.1 achieved an $\text{RMSE}_{128^3} = 0.56_{-0.03}^{+0.05}$ mK. DM8.1 also demonstrated residuals within the $10^{-2} - 10^{-1}$ mK² range for the dimensionless power spectrum, Δ_{21}^2 , which is below the expected noise level of SKA1-Low [21].

The pixel histograms in Figure 2 also reveal worse performance around the sharp peak at $T_{21} = 0$ mK, which corresponds to fully ionized regions where the 21-cm line is extinguished. In our training simulations the 21-cm signal is only observed in emission (i.e. $T_{21} \geq 0$). Consequently, one could explore adding a positivity constraint by learning a quantity where e.g. $\exp(T_{21}^*) = T_{21}$, which might reduce some of the observed discrepancies in the PDF in Figure 2. However, in reality and in the simulations, the 21-cm signal can assume both negative and positive values. The sign of the signal depends on the interplay between the initial conditions and the astrophysical processes as well as the redshift and spatial scales observed. As we ultimately aim to develop a methodology that generalizes across a range of redshifts, including those where T_{21} can also assume negative values, we do not adopt any positivity constraints as that would limit the applicability of our approach.

The SR network offers a significant reduction in computational resources taking just 4.7 hours to sample all 24 unique rotations of the 512^3 simulation cube on four NVIDIA A100-SXM-80GB GPUs.

4 Conclusions and Future Work

In this work, we trained score-based diffusion models to super-resolve gigaparsec-scale cosmological simulations of the 21-cm differential brightness temperature. We explored the impact of network architecture and training dataset size on the performance of trained SR networks. We demonstrated that the model performance in terms of the pixelwise-RMSE is independent of the training dataset size, i.e. a single simulation is enough to learn the data distribution regardless of the initial conditions. This is likely due to our training simulation volumes being large enough to neglect cosmic variance, such that the statistics are solely determined by the astrophysical and cosmological parameters, which are identical for all simulations. Moreover, we present sub-mK performance for the large models, which is well within the error margins of current 21-cm experiments [4, 40].

Finally, we examined our best-performing model across a range of spatial scales. We found $\text{RMSE} \sim 0.57$ mK across sampled 512^3 , 256^3 , and 128^3 simulation volumes, and in terms of the dimensionless power spectrum the residuals were of the order $10^{-2} - 10^{-1}$ mK², which is below the expected SKA1-Low noise level [21]. Additionally, our trained models successfully match the full probability distribution function (PDF) while achieving a low pixelwise-RMSE and low power spectrum residuals. While we believe jointly demonstrating a low RMSE, matching the power spectrum and the PDF, constitutes a challenging task in itself, and underscores the capabilities of the trained models, the model performance could also be evaluated based on higher-order statistics such as the bispectrum, which we will explore in future work.

In terms of future extensions to this study, training models with varying amounts of conditional information could also be explored. This could be attempted by systematically removing either of the initial conditions or the low-resolution input. In this proof-of-concept study, we provide the model with both the initial conditions and a low-resolution input. For this initial investigation, we considered the LR inputs essential, as they contain unique information about the astrophysical processes at play. These processes are independent from the initial conditions and are highly uncertain. Therefore, it would be extremely challenging - if not impossible - to reconstruct the 21-cm signal from the initial conditions without providing this extra information.

Finally, future extensions to speed up sampling, and add parameter-conditional generation should be explored. Faster sampling could be attempted using a methodology similar to Karras et al. [20], and parameter-conditional generation could be implemented using a latent diffusion model architecture [44, 35]. Parameter-conditioned super-resolution would allow us to use the largest spatial scales probed by the SKA1-Low to constrain astrophysical models, and infer the properties of the early Universe.

Acknowledgements

SP acknowledges the support from the Cambridge Commonwealth, European & International Trust in the form of a scholarship to undertake his degree and support this paper. JD acknowledges support from the Boustany Foundation and Cambridge Commonwealth Trust in the form of an Isaac Newton Studentship. AF is grateful for the support from the Royal Society through a University Research Fellowship. EdLA is supported by the Science and Technology Facilities Council (UK) through his Rutherford Fellowship.

This work used the DiRAC COSMA Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/P002293/1, ST/R002371/1 and ST/S002502/1, Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

This work was also supported by a seedcorn project which used the DiRAC Extreme Scaling service Tursa at the University of Edinburgh, managed by the Edinburgh Parallel Computing Centre on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The DiRAC service at Edinburgh was funded by BEIS, UKRI and STFC capital funding and STFC operations grants. DiRAC is part of the UKRI Digital Research Infrastructure

Finally, this work was also performed using resources provided by the Cambridge Service for Data Driven Discovery (CSD3) operated by the University of Cambridge Research Computing Service (www.csd3.cam.ac.uk), provided by Dell EMC and Intel using Tier-2 funding from the Engineering and Physical Sciences Research Council (capital grant EP/T022159/1), and DiRAC funding from the Science and Technology Facilities Council (www.dirac.ac.uk).

References

- [1] R. Barkana. The rise of the first stars: Supersonic streaming, radiative feedback, and 21-cm cosmology. *Phys. Rep.*, 645:1–59, July 2016. doi: 10.1016/j.physrep.2016.06.006.
- [2] H. T. J. Bevins, S. Heimersheim, I. Abril-Cabezas, A. Fialkov, E. de Lera Acedo, W. Handley, S. Singh, and R. Barkana. Joint analysis constraints on the physics of the first galaxies with low-frequency radio astronomy data. *Mon. Not. R. Astron. Soc.*, 527(1):813–827, Jan. 2024. doi: 10.1093/mnras/stad3194.
- [3] A. Cohen, A. Fialkov, and R. Barkana. The 21-cm BAO signature of enriched low-mass galaxies during cosmic reionization. *Mon. Not. R. Astron. Soc.*, 459(1):L90–L94, June 2016. doi: 10.1093/mnras/slw047.
- [4] E. de Lera Acedo, D. I. L. de Villiers, N. Razavi-Ghods, W. Handley, A. Fialkov, A. Magro, D. Anstey, H. T. J. Bevins, R. Chiello, J. Cumner, A. T. Josaitis, I. L. V. Roque, P. H. Sims, K. H. Scheutwinkel, P. Alexander, G. Bernardi, S. Carey, J. Cavillot, W. Croukamp, J. A. Ely, T. Gessey-Jones, Q. Gueuning, R. Hills, G. Kulkarni, R. Maiolino, P. D. Meerburg, S. Mittal, J. R. Pritchard, E. Puchwein, A. Saxena, E. Shen, O. Smirnov, M. Spinelli, and K. Zarb-Adami. The REACH radiometer for detecting the 21-cm hydrogen signal from redshift $z \approx 7.5$ –28. *Nature Astronomy*, 6:984–998, July 2022. doi: 10.1038/s41550-022-01709-9.
- [5] P. E. Dewdney, P. J. Hall, R. T. Schilizzi, and T. J. L. W. Lazio. The square kilometre array. *Proceedings of the IEEE*, 97(8):1482–1496, 2009. doi: 10.1109/JPROC.2009.2021005.
- [6] P. Dhariwal and A. Nichol. Diffusion models beat gans on image synthesis. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34 of *NIPS 2021*, pages 8780–8794. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/49ad23d1ec9fa4bd8d77d02681df5cfa-Paper.pdf.
- [7] A. Fialkov and R. Barkana. Signature of excess radio background in the 21-cm global signal and power spectrum. *Mon. Not. R. Astron. Soc.*, 486(2):1763–1773, June 2019. doi: 10.1093/mnras/stz873.

- [8] A. Fialkov, R. Barkana, D. Tseliakhovich, and C. M. Hirata. Impact of the relative motion between the dark matter and baryons on the first stars: semi-analytical modelling. *Mon. Not. R. Astron. Soc.*, 424(2):1335–1345, Aug. 2012. doi: 10.1111/j.1365-2966.2012.21318.x.
- [9] A. Fialkov, R. Barkana, E. Visbal, D. Tseliakhovich, and C. M. Hirata. The 21-cm signature of the first stars during the Lyman-Werner feedback era. *Mon. Not. R. Astron. Soc.*, 432(4): 2909–2916, July 2013. doi: 10.1093/mnras/stt650.
- [10] A. Fialkov, R. Barkana, A. Pinhas, and E. Visbal. Complete history of the observable 21 cm signal from the first stars during the pre-reionization era. *Mon. Not. R. Astron. Soc.*, 437(1): L36–L40, Jan. 2014. doi: 10.1093/mnrasl/slt135.
- [11] A. Fialkov, R. Barkana, and E. Visbal. The observable signature of late heating of the Universe during cosmic reionization. *Nature*, 506(7487):197–199, Feb. 2014. doi: 10.1038/nature12999.
- [12] S. R. Furlanetto, S. P. Oh, and F. H. Briggs. Cosmology at low frequencies: The 21 cm transition and the high-redshift Universe. *Phys. Rep.*, 433(4-6):181–301, Oct. 2006. doi: 10.1016/j.physrep.2006.08.002.
- [13] T. Gessey-Jones, N. S. Sartorio, A. Fialkov, G. M. Mirouh, M. Magg, R. G. Izzard, E. de Lera Acedo, W. J. Handley, and R. Barkana. Impact of the primordial stellar initial mass function on the 21-cm signal. *Mon. Not. R. Astron. Soc.*, 516(1):841–860, Oct. 2022. doi: 10.1093/mnras/stac2049.
- [14] T. Gessey-Jones, A. Fialkov, E. de Lera Acedo, W. J. Handley, and R. Barkana. Signatures of cosmic ray heating in 21-cm observables. *Mon. Not. R. Astron. Soc.*, 526(3):4262–4284, Dec. 2023. doi: 10.1093/mnras/stad3014.
- [15] T. Gessey-Jones, S. Pochinda, H. T. J. Bevins, A. Fialkov, W. J. Handley, E. de Lera Acedo, S. Singh, and R. Barkana. On the constraints on superconducting cosmic strings from 21-cm cosmology. *Mon. Not. R. Astron. Soc.*, 529(1):519–536, Mar. 2024. doi: 10.1093/mnras/stae512.
- [16] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27 of *NIPS 2014*. Curran Associates, Inc., 2014. URL https://proceedings.neurips.cc/paper_files/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf.
- [17] HERA Collaboration, Z. Abdurashidova, T. Adams, J. E. Aguirre, P. Alexander, Z. S. Ali, R. Baartman, Y. Balfour, R. Barkana, A. P. Beardsley, G. Bernardi, T. S. Billings, J. D. Bowman, R. F. Bradley, D. Breitman, P. Bull, J. Burba, S. Carey, C. L. Carilli, C. Cheng, S. Choudhuri, D. R. DeBoer, E. de Lera Acedo, M. Dexter, J. S. Dillon, J. Ely, A. Ewall-Wice, N. Fagnoni, A. Fialkov, R. Fritz, S. R. Furlanetto, K. Gale-Sides, H. Garsden, B. Glendenning, A. Gorce, D. Gorthi, B. Greig, J. Grobbelaar, Z. Halday, B. J. Hazelton, S. Heimersheim, J. N. Hewitt, J. Hickish, D. C. Jacobs, A. Julius, N. S. Kern, J. Kerrigan, P. Kittiwisit, S. A. Kohn, M. Kolopanis, A. Lanman, P. La Plante, D. Lewis, A. Liu, A. Loots, Y.-Z. Ma, D. H. E. MacMahon, L. Malan, K. Malgas, C. Malgas, M. Maree, B. Marero, Z. E. Martinot, L. McBride, A. Mesinger, J. Mirocha, M. Molewa, M. F. Morales, T. Mosiane, J. B. Muñoz, S. G. Murray, V. Nagpal, A. R. Neben, B. Nikolic, C. D. Nunhokee, H. Nuwegeld, A. R. Parsons, R. Pascua, N. Patra, S. Pieterse, Y. Qin, N. Razavi-Ghods, J. Robnett, K. Rosie, M. G. Santos, P. Sims, S. Singh, C. Smith, H. Swarts, J. Tan, N. Thyagarajan, M. J. Wilensky, P. K. G. Williams, P. van Wyngaarden, and H. Zheng. Improved Constraints on the 21 cm EoR Power Spectrum and the X-Ray Heating of the IGM with HERA Phase I Observations. *Astrophys. J.*, 945(2):124, Mar. 2023. doi: 10.3847/1538-4357/acaf50.
- [18] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33 of *NIPS 2020*, pages 6840–6851. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf.

- [19] K. Karchev, N. Anau Montel, A. Coogan, and C. Weniger. Strong-Lensing Source Reconstruction with Denoising Diffusion Restoration Models. In *36th Conference on Neural Information Processing Systems: Workshop on Machine Learning and the Physical Sciences*, 11 2022.
- [20] T. Karras, M. Aittala, T. Aila, and S. Laine. Elucidating the design space of diffusion-based generative models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 26565–26577. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/a98846e9d9cc01cfb87eb694d946ce6b-Paper-Conference.pdf.
- [21] H. D. Kaur, N. Gillet, and A. Mesinger. Minimum size of 21-cm simulations. *Monthly Notices of the Royal Astronomical Society*, 495(2):2354–2362, 05 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa1323. URL <https://doi.org/10.1093/mnras/staa1323>.
- [22] D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. *arXiv e-prints*, art. arXiv:1412.6980, Dec. 2014. doi: 10.48550/arXiv.1412.6980.
- [23] L. Koopmans, J. Pritchard, G. Mellema, J. Aguirre, K. Ahn, R. Barkana, I. van Bemmell, G. Bernardi, A. Bonaldi, F. Briggs, A. G. de Bruyn, T. C. Chang, E. Chapman, X. Chen, B. Courty, P. Dayal, A. Ferrara, A. Fialkov, F. Fiore, K. Ichiki, I. T. Illiev, S. Inoue, V. Jelic, M. Jones, J. Lazio, U. Maio, S. Majumdar, K. J. Mack, A. Mesinger, M. F. Morales, A. Parsons, U. Pen, M. Santos, R. Schneider, B. Semelin, R. S. de Souza, R. Subrahmanyam, T. Takeuchi, H. Vedantham, J. Wagg, R. Webster, S. Wyithe, K. K. Datta, and C. Trott. The Cosmic Dawn and Epoch of Reionisation with SKA. *PoS, AASKA14:001*, 2015. doi: 10.22323/1.215.0001.
- [24] R. Legin, M. Ho, P. Lemos, L. Perreault-Levasseur, S. Ho, Y. Hezaveh, and B. Wandelt. Posterior sampling of the initial conditions of the universe from non-linear large scale structures using score-based generative models. *Mon. Not. R. Astron. Soc.*, 527(1):L173–L178, Jan. 2024. doi: 10.1093/mnras/lsad152.
- [25] H. Li, Y. Yang, M. Chang, S. Chen, H. Feng, Z. Xu, Q. Li, and Y. Chen. Srdiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing*, 479:47–59, 2022. ISSN 0925-2312. doi: <https://doi.org/10.1016/j.neucom.2022.01.029>. URL <https://www.sciencedirect.com/science/article/pii/S0925231222000522>.
- [26] M. Magg, I. Reis, A. Fialkov, R. Barkana, R. S. Klessen, S. C. O. Glover, L.-H. Chen, T. Hartwig, and A. T. P. Schauer. Effect of the cosmological transition to metal-enriched star formation on the hydrogen 21-cm signal. *Mon. Not. R. Astron. Soc.*, 514(3):4433–4449, Aug. 2022. doi: 10.1093/mnras/stac1664.
- [27] A. Mesinger, editor. *The Cosmic 21-cm Revolution*. 2514-3433. IOP Publishing, 2019. ISBN 978-0-7503-2236-2. doi: 10.1088/2514-3433/ab4a73. URL <https://dx.doi.org/10.1088/2514-3433/ab4a73>.
- [28] O. S. D. O’Hara, F. Dulwich, E. de Lera Acedo, J. Dhandha, T. Gessey-Jones, D. Anstey, and A. Fialkov. Understanding spectral artefacts in SKA-LOW 21-cm cosmology experiments: the impact of cable reflections. *arXiv e-prints*, art. arXiv:2402.04008, Feb. 2024. doi: 10.48550/arXiv.2402.04008.
- [29] S. Pochinda, T. Gessey-Jones, H. T. J. Bevins, A. Fialkov, S. Heimersheim, I. Abril-Cabezas, E. de Lera Acedo, S. Singh, S. Sikder, and R. Barkana. Constraining the properties of Population III galaxies with multiwavelength observations. *Mon. Not. R. Astron. Soc.*, 531(1):1113–1132, June 2024. doi: 10.1093/mnras/stae1185.
- [30] D. Podell, Z. English, K. Lacey, A. Blattmann, T. Dockhorn, J. Müller, J. Penna, and R. Rombach. SDXL: Improving latent diffusion models for high-resolution image synthesis. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=di52zR8xgf>.
- [31] J. R. Pritchard and A. Loeb. 21 cm cosmology in the 21st century. *Reports on Progress in Physics*, 75(8):086901, Aug. 2012. doi: 10.1088/0034-4885/75/8/086901.

- [32] I. Reis, A. Fialkov, and R. Barkana. High-redshift radio galaxies: a potential new source of 21-cm fluctuations. *Mon. Not. R. Astron. Soc.*, 499(4):5993–6008, Dec. 2020. doi: 10.1093/mnras/staa3091.
- [33] I. Reis, A. Fialkov, and R. Barkana. The subtlety of Ly α photons: changing the expected range of the 21-cm signal. *Mon. Not. R. Astron. Soc.*, 506(4):5479–5493, Oct. 2021. doi: 10.1093/mnras/stab2089.
- [34] I. Reis, R. Barkana, and A. Fialkov. Shot noise and scatter in the star formation efficiency as a source of 21-cm fluctuations. *Mon. Not. R. Astron. Soc.*, 511(4):5265–5273, Apr. 2022. doi: 10.1093/mnras/stac411.
- [35] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, June 2022.
- [36] A. Rouhiainen, M. Münchmeyer, G. Shiu, M. Gira, and K. Lee. Superresolution emulation of large cosmological fields with a 3D conditional diffusion model. *Phys. Rev. D*, 109(12):123536, June 2024. doi: 10.1103/PhysRevD.109.123536.
- [37] C. Saharia, J. Ho, W. Chan, T. Salimans, D. J. Fleet, and M. Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4713–4726, 2023. doi: 10.1109/TPAMI.2022.3204461.
- [38] A. M. M. Scaife. Big telescope, big data: towards exascale with the Square Kilometre Array. *Philosophical Transactions of the Royal Society of London Series A*, 378(2166):20190060, Mar. 2020. doi: 10.1098/rsta.2019.0060.
- [39] S. Sikder, R. Barkana, A. Fialkov, and I. Reis. Strong 21-cm fluctuations and anisotropy due to the line-of-sight effect of radio galaxies at cosmic dawn. *Mon. Not. R. Astron. Soc.*, 527(4):10975–10985, Feb. 2024. doi: 10.1093/mnras/stad3847.
- [40] S. Singh, N. T. Jishnu, R. Subrahmanyam, N. Udaya Shankar, B. S. Girish, A. Raghunathan, R. Somashekar, K. S. Srivani, and M. Sathyanarayana Rao. On the detection of a cosmic dawn signal in the radio background. *Nature Astronomy*, 6:607–617, Feb. 2022. doi: 10.1038/s41550-022-01610-5.
- [41] M. J. Smith, J. E. Geach, R. A. Jackson, N. Arora, C. Stone, and S. Courteau. Realistic galaxy image simulation via score-based generative models. *Mon. Not. R. Astron. Soc.*, 511(2):1808–1818, Apr. 2022. doi: 10.1093/mnras/stac130.
- [42] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In F. Bach and D. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/sohl-dickstein15.html>.
- [43] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=PxtTIG12RRHS>.
- [44] A. Vahdat, K. Kreis, and J. Kautz. Score-based generative modeling in latent space. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 11287–11302. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/5dca4c6b9e244d24a30b4c45601d9720-Paper.pdf.
- [45] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30 of *NIPS 2017*. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.

- [46] E. Visbal, R. Barkana, A. Fialkov, D. Tseliakhovich, and C. M. Hirata. The signature of the first stars in atomic hydrogen at redshift 20. *Nature*, 487(7405):70–73, July 2012. doi: 10.1038/nature11177.
- [47] X. Wang, R. Girshick, A. Gupta, and K. He. Non-local neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [48] R. Yang and S. Mandt. Lossy image compression with conditional diffusion models. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36 of *NIPS 2023*, pages 64971–64995. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/ccf6d8b4a1fe9d9c8192f00c713872ea-Paper-Conference.pdf.
- [49] X. Zhao, Y.-S. Ting, K. Diao, and Y. Mao. Can diffusion model conditionally generate astrophysical images? *Mon. Not. R. Astron. Soc.*, 526(2):1699–1712, Dec. 2023. doi: 10.1093/mnras/stad2778.