HIGlow: Conditional Normalizing Flows for High-Fidelity HI Map Modeling

Roy Friedman Department of Computer Science The Hebrew University of Jerusalem Jerusalem, Israel roy.friedman@mail.huji.ac.il Sultan Hassan *^{†‡} Center for Computational Astrophysics Flatiron Institute New York, NY 10010 shassan@flatironinstitute.org

Abstract

Extracting the maximum amount of cosmological and astrophysical information from upcoming large-scale surveys remains a challenge. This includes evaluating the exact likelihood, parameter inference and generating new diverse synthetic examples of the incoming high-dimensional data sets. In this work, we propose the use of normalizing flows as a generative model of the neutral hydrogen (HI) maps from the CAMELS project. Normalizing flows have been very successful at parameter inference and generating new, realistic examples. Our model utilizes the spatial structure of the HI maps in order to faithfully follow the statistics of the data, allowing for high-fidelity sample generation and efficient parameter inference.

1 Introduction

As the fields of cosmology and astrophysics enter the era of big-data, next-generation large-scale surveys, such as the Square Kilometer Array [SKA, 1], the Hydrogen Epoch of Reionization Array [HERA, 2], the Low Frequency Array [LOFAR, 3], the Vera C. Rubin Observatory Legacy Survey of Space and Time [LSST, 4], Nancy Grace Roman Space Telescope [Roman, 5], Spectro-Photometer for the History of the Universe, Epoch of Reionization, and Ices Explorer [SPHEREx, 6] and Euclid [7], will enable imaging the neutral hydrogen (HI) in the early universe using unprecedented sensitivity and large fields of views. These surveys are expected to provide large amounts of high-dimensional data sets that require new tools, able to efficiently extract as much cosmological and astrophysical information as possible. A key challenge in the analysis of these new data sets lies in evaluating the exact likelihood of new samples in order to perform parameter inference that will enable translation of these growing observational efforts into tight theoretical constraints. Additionally, generating labeled synthetic large-scale HI maps, which faithfully capture the high-order statistics of the data, is equally important in order to facilitate accurate forecasting for upcoming surveys.

Traditional inference methods usually rely on the use of summary statistics, such as the sphericallyaveraged power spectrum, and are prone to loss of information due to data compression. However, a key lesson learnt from recent state-of-the-art convolutional neural networks (CNNs) is that performing inference at the field level is crucial to extracting the maximum amount of information [8, 9]. While there exist several families of generative models, such as generative adversarial networks [10] (GANs) and variational autoencoders [11] (VAEs), they are not able to provide full access to the exact likelihood of high-dimensional data. On the other hand, normalizing flows [12–16] are emerging as a powerful technique that provides access to high-dimensional likelihood and can generate high fidelity new samples.

Machine Learning and the Physical Sciences workshop, NeurIPS 2022.

^{*}Department of Astrophysical Sciences, Princeton University, Princeton, NJ, 08544

[†]Department of Physics & Astronomy, University of the Western Cape, Cape Town 7535

[‡]NHFP Hubble Fellow



Figure 1: (A) normalizing flow modules used in Glow, (B) Glow architecture and (C) module changes in order to make Glow conditional

Recently, a normalizing flow model called HIFlow [17] that is based on conditional masked autoregressive flows was introduced as a generative model for simulated HI maps. While successful, HIFlow does not utilize the spatial structure of these HI maps and as such is unable to capture the full statistics of the data. In this work we propose a new model, called *HIGlow*⁴, which is a conditional version of Glow [13], a powerful normalizing flow for modeling images. By using convolutional operations, our model is more suited to the spatial structure of the HI maps and is able to faithfully capture their high-frequency statistics. Additionally, using our conditional model we show how parameter inference is made possible using the posterior distribution of our model.

1.1 Simulation

We use simulations from the CAMELS project, recently introduced in Villaescusa-Navarro et al. [18]. CAMELS is a suite of thousands of simulations run with state-of-the-art cosmological hydrodynamic galaxy formation models, by varying two cosmological parameters (Ω_m and σ_8) and four other parameters that modify the strength of stellar feedback (ASN1, ASN2) and black hole feedback (AAGN1, AAGN2) relative to the original IllustrisTNG and Simba simulations. We focus our analysis on the Latin-hypercube (LH) set, which is a set of 1,000 simulations that explores this six-dimensional parameter space with uniform prior ranges.

Hassan et al. [17] showed that the distribution of the HI maps from the CAMELS simulations is largely dependent on the cosmological parameters Ω_m and σ_8 . As such, we chose to focus on a model that is conditional on only these two parameters. Nonetheless, our model can be extended in order to use more conditional parameters, if needed.

2 Methods

Normalizing flows are a flexible framework for modelling probability distributions. Intuitively, normalizing flows utilize the change of variable formula in order to define a parametric transformation of a simple distribution into a more flexible distribution, in order to model the distribution of data. See [16] for an in-depth discussion of normalizing flows and their applications.

⁴Code available at github.com/friedmanroy/HI-generation

2.1 Glow

Glow ([13]) is a normalizing flow model that utilizes 1×1 invertible convolutions in order to define invertible functions that take spatial information into account. The Glow architecture can be simplified into 3 modules which are stacked in order to create a single flow step, which are further stacked in order to get a flow block. These modules are:

Actnorm which scales and translates the different channels in the input.

Invertible 1×1 **convolutions** which act as a generalization of a permutation.

Affine coupling which is the workhorse of the Glow architecture and allows for invertible affine transformations of the input.

The Glow architecture further utilizes the spatial structure of the images by squeezing the spatial dimension into the channel dimension (see [19]), which allows for more efficient splitting in the affine coupling layers and the use of a multi-scale approach; Figure 1-A shows a schematic of the layers mentioned above while Figure 1-B shows a schematic of Glow's architecture. The implementations of Glow's affine coupling layers utilize CNNs, making it particularly useful for modelling images.

2.2 HIGlow

In conditional modelling, we are given a dataset of pairs $\mathcal{D} = \{(z_i, x_i)\}_{i=1}^N$ where the z_i s are the underlying parameters for the generation of x_i . In this case, the task is to model the conditional distribution $p_x(x_i|z_i)$. The Glow model, as described above, does not have the capability to model the required conditional probabilities.

In order to turn a normalizing flow model into a conditional model, it is enough to change the transformation into one that takes z, as well as the latent vector y, as an input: $x = f_{\theta}^{-1}(y, z)$. In practice, this can be accomplished by changing the modules described in section 2.1.

We follow the methods of [14, 15] to turn Glow into a conditional model, using variants of the actnorm and affine coupling modules. Given the parameters z, these new layer types are:

Affine injector, a variant of the actnorm layer, with the following transformation:

$$y = g_s(z) \circ x + g_t(z) \tag{1}$$

where $g_s(\cdot)$ and $g_t(\cdot)$ are neural networks. This transformation was placed instead of the actnorm in Glow.

Conditional affine coupling is the same as the original affine coupling, only z is added as an input to the affine transformation. This transformation was used instead of the standard affine couplings of Glow.

A schematic of both of the newly defined layers and how they compare to the standard Glow network can be seen in 1-C.

In order to insert conditional information into the affine coupling layer, as described above, a new constant channel was added to the affine coupling's input for every parameter. Additionally, the affine injector in Equation 1 is defined according to two functions: $g_s(z)$ and $g_t(z)$. In our implementation, both of these functions were implemented as multi-layer perceptrons (MLPs) with a single hidden layer from z to a scalar.

3 Results

For all experiments, we used HIGlow with 6 flow blocks, each containing 20 flow steps on 64×64 pixel HI maps. Both the conditional affine coupling as well as the affine injectors had 1 hidden layer of width 16. The model was trained for 1000 epochs on one GPU.

3.1 HI map generation

HIGlow can be used in order to generate new HI maps, conditional on the cosmological parameters Ω_m and σ_8 . As can be seen in Figure 2-A, samples generated from HIGlow are visually very similar to those from the CAMELS simulation, even for different parameter settings.



Figure 2: (A) Samples from CAMELS (top) next to samples generated from HIGlow (bottom) for different parameter settings. (B) Comparison of marginal power spectrum mean and standard deviation between CAMELS, HIGlow and HIFlow; (C) comparison of power spectra under different conditional parameters.



Figure 3: (A) Posterior distribution of HI maps with different parameters. (B) Estimated parameters using the mean of the posterior versus the true parameter values.

To ensure that the statistics of the generated HI maps follow the statistics of the CAMELS simulation data, we compare the power spectra of both methods. Figure 2-B shows the mean power spectrum and standard deviation from it as a function of spherical frequency of 500 maps unconditionally generated from HIGlow against 500 maps from CAMELS and 500 maps generated by HIFlow [18], a previous approach using autoregressive flows in order to model HI maps from CAMELS. HIGlow is able to correctly model the high-frequency information in the HI maps, unlike HIFlow. In Figure 2-C, the same curves can be seen for specific parameters values. In these HIGlow is able, again, to generate images with similar mean spectrum and standard deviation.

3.2 Posterior distribution and parameter inference

HIGlow can also be used in order to infer the underlying cosmological parameters, given a specific HI map. Bayes' law can be approximated in order to find the posterior distribution p(z|x):

$$p(z|x) = \frac{p(z)p(x|z)}{\int p(x,z)dz} \approx \frac{p(z)p(x|z)}{\sum_{i=1}^{N} p(x|\tilde{z}_i)p(\tilde{z}_i)}$$
(2)

where $\tilde{z}_1, \dots, \tilde{z}_N$ are N samples from p(z).

In CAMELS, the distribution over the cosmological parameters is uniform, so this is the setting we use in the following experiments, as it also follows for the test data. If a different prior over cosmological parameters is given, it can also easily be incorporated in this step. Figure 3-A shows examples of the posterior distributions of HI maps the model didn't see during training time. The posterior distribution in these plots is concentrated around the true parameter values and admits a measure of uncertainty, which may be useful in real experimental settings.

The posterior distribution can also be utilized in order to estimate the cosmological parameters given a new, unseen, HI map. To do so, we approximate the minimum mean-squared error (MMSE) estimator, given by $\hat{z} = \mathbb{E}_{z} [z|x]$, where \hat{z} is the estimated value. Assuming that HIGlow correctly captures the conditional distribution, the MMSE is assured to give the most accurate estimate (in expectation). Results for parameter inference can be seen in Figure 3; as can be seen, the predictions are typically close to the true values. This estimator achieves an RMSE of 0.058 for Ω_m and 0.065 for σ_8 .

4 Discussion

In this work, we have introduced a new generative model, HIGlow, capable of generating high-fidelity HI maps and allows for exact likelihood estimation of high-dimensional data points. HIGlow also allows for the generation of new data points conditional on cosmological parameters, with a similar power spectrum distribution to that of the CAMELS simulation suite. By utilizing the ability of this model to calculate the exact likelihood of maps conditional on a set of parameters, the posterior of the cosmological parameters conditioned on HI maps is made possible. This further allows for accurate parameter inference, which may prove helpful in upcoming large-scale cosmological surveys.

As HIGlow is an accurate model of the CAMELS simulation, it is our hope that HIGlow can be used to generate novel samples of HI maps under different cosmological parameters. Future work will focus on using HIGlow to forecasting for specific experiments such as CHIME [20], HIRAX [21] and SKA [1]. This will include a full treatment to all instrumental effects, such as angular resolution, thermal noise and foreground cleaning.

5 Broader Impact

In this work, a new generative model for HI maps in the early universe was introduced, allowing for faster generation of new HI maps as well as parameter inference. Hopefully, the methods introduced in this work will facilitate cosmological research. While in general machine learning can be misused, but we recognize no situation where the methodologies, models or implementations of this work can lead towards any negative societal impacts.

Acknowledgments and Disclosure of Funding

The authors would like to acknowledge support provided by the Simons Foundation as well as support for Program number HST-HF2-51507 provided by NASA through a grant from the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, incorporated, under NASA contract NAS5-26555.

References

 Garrelt Mellema, Léon V. E. Koopmans, Filipe A. Abdalla, Gianni Bernardi, Benedetta Ciardi, Soobash Daiboo, A. G. de Bruyn, Kanan K. Datta, Heino Falcke, Andrea Ferrara, Ilian T. Iliev, Fabio Iocco, Vibor Jelić, Hannes Jensen, Ronniy Joseph, Panos Labroupoulos, Avery Meiksin, Andrei Mesinger, André R. Offringa, V. N. Pandey, Jonathan R. Pritchard, Mario G. Santos, Dominik J. Schwarz, Benoit Semelin, Harish Vedantham, Sarod Yatawatta, and Saleem Zaroubi. Reionization and the Cosmic Dawn with the Square Kilometre Array. *Experimental Astronomy*, 36(1-2):235–318, August 2013. doi: 10.1007/s10686-013-9334-5.

- [2] David R DeBoer, Aaron R Parsons, James E Aguirre, Paul Alexander, Zaki S Ali, Adam P Beardsley, Gianni Bernardi, Judd D Bowman, Richard F Bradley, Chris L Carilli, et al. Hydrogen epoch of reionization array (hera). *Publications of the Astronomical Society of the Pacific*, 129(974):045001, 2017.
- [3] Michael P van Haarlem, Michael W Wise, AW Gunst, George Heald, John P McKean, Jason WT Hessels, A Ger de Bruyn, Ronald Nijboer, John Swinbank, Richard Fallows, et al. Lofar: The low-frequency array. *Astronomy & astrophysics*, 556:A2, 2013.
- [4] Željko Ivezić, Steven M Kahn, J Anthony Tyson, Bob Abel, Emily Acosta, Robyn Allsman, David Alonso, Yusra AlSayyad, Scott F Anderson, John Andrew, et al. Lsst: from science drivers to reference design and anticipated data products. *The Astrophysical Journal*, 873(2):111, 2019.
- [5] D Spergel, N Gehrels, C Baltay, D Bennett, J Breckinridge, M Donahue, A Dressler, BS Gaudi, T Greene, O Guyon, et al. Wide-field infrarred survey telescope-astrophysics focused telescope assets wfirst-afta 2015 report. arXiv preprint arXiv:1503.03757, 2015.
- [6] Olivier Doré, Jamie Bock, Matthew Ashby, Peter Capak, Asantha Cooray, Roland de Putter, Tim Eifler, Nicolas Flagey, Yan Gong, Salman Habib, et al. Cosmology with the spherex all-sky spectral survey. arXiv preprint arXiv:1412.4872, 2014.
- [7] Giuseppe D Racca, René Laureijs, Luca Stagnaro, Jean-Christophe Salvignol, José Lorenzo Alvarez, Gonzalo Saavedra Criado, Luis Gaspar Venancio, Alex Short, Paolo Strada, Tobias Bönke, et al. The euclid mission design. In *Space telescopes and instrumentation 2016: optical, infrared, and millimeter wave*, volume 9904, pages 235–257. SPIE, 2016.
- [8] Francisco Villaescusa-Navarro, Daniel Anglés-Alcázar, Shy Genel, David N Spergel, Yin Li, Benjamin Wandelt, Andrina Nicola, Leander Thiele, Sultan Hassan, Jose Manuel Zorrilla Matilla, et al. Multifield cosmology with artificial intelligence. arXiv preprint arXiv:2109.09747, 2021.
- [9] Sultan Hassan, Sambatra Andrianomena, and Caitlin Doughty. Constraining the astrophysics and cosmology from 21 cm tomography using deep learning with the ska. *Monthly Notices of the Royal Astronomical Society*, 494(4):5761–5774, 2020.
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144, 2020.
- [11] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114, 2013.
- [12] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- [13] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. Advances in Neural Information Processing Systems, 31, 2018.
- [14] You Lu and Bert Huang. Structured output learning with conditional generative flows. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5005–5012, 2020.
- [15] Andreas Lugmayr, Martin Danelljan, Luc Van Gool, and Radu Timofte. Srflow: Learning the superresolution space with normalizing flow. In *European Conference on Computer Vision*, pages 715–732. Springer, 2020.
- [16] George Papamakarios, Eric T Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. J. Mach. Learn. Res., 22(57): 1–64, 2021.
- [17] Sultan Hassan, Francisco Villaescusa-Navarro, Benjamin Wandelt, David N Spergel, Daniel Anglés-Alcázar, Shy Genel, Miles Cranmer, Greg L Bryan, Romeel Davé, Rachel S Somerville, et al. Hiflow: Generating diverse hi maps conditioned on cosmology using normalizing flow. arXiv preprint arXiv:2110.02983, 2021.

- [18] Francisco Villaescusa-Navarro, Daniel Anglés-Alcázar, Shy Genel, David N Spergel, Rachel S Somerville, Romeel Dave, Annalisa Pillepich, Lars Hernquist, Dylan Nelson, Paul Torrey, et al. The camels project: Cosmology and astrophysics with machine-learning simulations. *The Astrophysical Journal*, 915(1):71, 2021.
- [19] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- [20] Mandana Amiri, Kevin Bandura, Anja Boskovic, Tianyue Chen, Jean-François Cliche, Meiling Deng, Nolan Denman, Matt Dobbs, Mateus Fandino, Simon Foreman, et al. An overview of chime, the canadian hydrogen intensity mapping experiment. *The Astrophysical Journal Supplement Series*, 261(2):29, 2022.
- [21] LB Newburgh, K Bandura, MA Bucher, T-C Chang, HC Chiang, JF Cliche, R Davé, M Dobbs, C Clarkson, KM Ganga, et al. Hirax: a probe of dark energy and radio transients. In *Ground-based and Airborne Telescopes VI*, volume 9906, pages 2039–2049. SPIE, 2016.

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] See section 5
 - (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)? [N/A]
 - (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)? [N/A]
 - (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? [Yes] See section 1.1
 - (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? [No]
 - (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]