
HIDM: Emulating Large Scale HI Maps using Score-based Diffusion Models

Sultan Hassan *
New York University
Flatiron Institute
University of the Western Cape
sultan.hassan@nyu.edu

Sambatra Andrianomena
South African Radio Astronomy Observatory
Department of Physics & Astronomy
University of the Western Cape
andrianomena@gmail.com

Abstract

Efficiently analyzing maps from upcoming large-scale surveys requires gaining direct access to a high-dimensional likelihood and generating large-scale fields with high fidelity, which both represent major challenges. Using CAMELS simulations, we employ the state-of-the-art score-based diffusion models to simultaneously achieve both tasks. We show that our model, HIDM, is able to efficiently generate high fidelity large scale HI maps that are in a good agreement with the CAMELS’s power spectrum, probability distribution, and likelihood up to second moments. HIDM represents a step forward towards maximizing the scientific return of future large scale surveys.

1 Introduction

Imaging the neutral Hydrogen (HI) distribution – a promising tracer of dark matter and large scale structure – over cosmological scales is a primary target of several ongoing and upcoming surveys such as the Square Kilometer Array [SKA, Mellema et al., 2013], the Hydrogen Epoch of Reionization Array [HERA, DeBoer et al., 2017], the Low Frequency Array [LOFAR, van Haarlem et al., 2013] the Vera C. Rubin Observatory Legacy Survey of Space and Time [LSST, Ivezić et al., 2019], Nancy Grace Roman Space Telescope [Roman, Spergel et al., 2015], Spectro-Photometer for the History of the Universe, Epoch of Reionization, and Ices Explorer [SPHEREx, Doré et al., 2014], and Euclid [Racca et al., 2016]. A key challenge to analyzing outputs from these large scale surveys relies in the high-dimensionality, which poses several significant challenges including memory requirement and computational expense. For instance, traditional large scale simulations are computationally too expensive to enable field level inference via standard techniques such as Markov chain Monte Carlo (MCMC) samplers. To effectively analyze any large scale high-dimensional observed maps (including HI), access to the likelihood and ability to generate new synthetic examples conditioned on parameters are both required.

To achieve these goals, previous studies have used various techniques of generative models, such as variants of Generative Adversarial Networks [e.g. Zamudio-Fernandez et al., 2019, Villaescusa-Navarro et al., 2021a, Wadkar et al., 2021, Andrianomena et al., 2023], Variational Autoencoders [e.g. Lanusse et al., 2021, Cheng et al., 2021], and flow-based models [e.g. Hassan et al., 2022, Friedman and Hassan, 2022, Dai and Seljak, 2022]. However, most of these techniques fail to fully capture the high frequencies (i.e. small scale features), where non-linearity is higher, especially when applied on high resolution large scale maps. To maximize the scientific return of future surveys, an efficient and flexible deep learning model is needed to correctly model features over all scales. Currently, diffusion models achieve the state-of-the-art performance in image generation tasks². In addition,

*NASA Hubble Fellow

²<https://paperswithcode.com/method/diffusion>




Figure 1: Comparison between HIDM and CAMELS in terms of 64 random generated examples of large scale HI maps. Both simulations visually show indistinguishable diversity of the HI large scale fluctuations. These maps cover an area of $25 h^{-1} \text{ cMpc} \times 25 h^{-1} \text{ cMpc}$ with 64 pixels on a side, resulting in a resolution of $\sim 0.4 h^{-1} \text{ cMpc}$. The color scale in these maps show the column density range $\log_{10} N_{\text{HI}}/\text{cm}^{-2}=14\text{-}22$.

some previous works [Adam et al., 2022, Mudur and Finkbeiner, 2022] have shown that diffusion models are successful in generating posterior samples of realistic astrophysical fields. The general flow of these models is based on the idea of gradually injecting noise to convert the signal to a white noise, and then generating new examples by gradually removing back the noise [e.g. Song et al., 2020]. In addition and similar to the flow-based models, obtaining the likelihood is possible via a simple change of variables. Hence, developing diffusion models to generate large scale HI maps would be very useful to effectively analyze future surveys.

2 Training dataset

We use the publicly available HI maps from CAMELS Multifield Dataset [CMD, Villaescusa-Navarro et al., 2021b] an outcome of the CAMELS project [Villaescusa-Navarro et al., 2021a]. Through the variation of several cosmological and astrophysical parameters, the resulting maps from CAMELS simulations show a large diversity of HI fluctuations on cosmological scales as shown in the left panel of Figure 1. The HI maps considered in this study is from the IllustrisTNG simulations [Weinberger et al., 2017, Pillepich et al., 2018] at $z = 6$ with 64×64 pixels and size of $25 \text{ Mpc}/h$, where $h = 0.6711$.

3 Score-based diffusion models

We follow closely the implementation by Song et al. [2020] since it provides a somewhat faster sampling through solving a Stochastic Differential Equation (SDE) as well as access to the high dimensional likelihood. We here briefly describe the basic framework of diffusion models. As mentioned earlier, the diffusion process gradually corrupts data into random noise over time via the following SDE:

$$dx = f(x, t)dt + g(t)dw, \quad (1)$$

where $f(x, t)$ is the drift coefficient, $g(t)$ is diffusion coefficient, and dw is a standard Brownian motion. In this process, the distribution of data P_0 is gradually converted into a tractable prior




Figure 2: Comparison between HIDM (blue) and CAMELS (red) in terms of the mean (μ) and standard deviation (σ) of the power spectrum (right) and probability distribution (left) of HI large scale maps. Both simulations show similar statistics up to second moments.

distribution P_T over time. Due to the large nature of the diffusion processes, P_T would have no memory of the initial P_0 , and hence this provides essential tests for out-of-distribution detection [e.g. Graham et al., 2022]. The score function – the gradient of the prior $\nabla_x \log P_T$ – arises in the reverse process to generate samples via the following reverse-time SDE:

$$dx = [f(x, t) - g^2(t)\nabla_x \log P_t]dt + g(t)dw. \quad (2)$$

The score $\nabla_x \log P_t$ is estimated using a time-dependent score model (S) which is approximated by a standard U-Net architecture to minimize the following loss function \mathcal{L} :

$$\mathcal{L} = E_{x(0)} E_{x(t)} [||S(x(t), t) - \nabla_{x(t)} \log P_t(x(t)|x(0))||_2^2]. \quad (3)$$

In our case, we gradually perturb HI maps to the prior distribution assuming the following SDE: $dx = \lambda^t dw$, where $t \in [0, 1]$. This corresponds to a standard deviation (σ) of the marginal distribution $P_t(x(t)|x(0))$ of $\frac{1}{2 \log \lambda} (\lambda^{2t} - 1)$ that gradually converts HI maps to white noise. The resulting loss function in this case reads:

$$\mathcal{L} = \frac{1}{N} \sum_N (S(x + \sigma z, t) + z)^2, \quad (4)$$

where z are random samples from a normal distribution i.e. $\mathcal{N}(0, 1)$.

For reproducibility, our U-Net architecture consists of four convolutional blocks for down and up sampling with 3×3 64 filters in each layer that is conditioned on time (t). We use Adam optimizer with learning rate of 10^{-3} and a diffusion coefficient $\lambda = 25$ for a total of 500 training epochs (convergence achieved from first 100 epochs) with a mini batch size of 128. To generate new samples, one needs to solve the reverse-time SDE, which in our case reads:

$$dx = -\frac{1}{2} \lambda^{-2t} S(x, t) dt. \quad (5)$$

In particular, we start from a random white noise $P(T)$ and integrate this equation in the reverse direction from $t = T$ to 0 to generate new realization of HI maps $P(0)$. It is then straightforward to compute the exact likelihood with change of variable formula (similar to flow-based models) as follows:

$$\log P(0) = \log P(T) - \frac{1}{2} \int_0^T \frac{d\lambda^2}{dt} \text{div} S(x(t), t) dt, \quad (6)$$

where $\text{div} S$ is the divergence of the score function representing the trace of the Jacobian of the change of variable.

4 Results

Using the process explained in the previous section, we design HI map diffusion model (HIDM). Figure 1 shows 64 random samples from CAMELS (left) versus HIDM (right). This figure demonstrates the ability of HIDM to generate high quality synthetic examples of large scale HI maps with diversity that is indistinguishable from the target (CAMELS). To quantify the image generation quality by HIDM, we consider the probability distribution function (PDF, left) and power spectrum (P, right) as shown in Figure 2. We first select 1,000 random realizations from CAMELS and similarly generate 1,000 random examples from HIDM. We then compute the PDF over all the 1,000 realizations from each model, and measure the mean PDF ($\mu(\text{PDF})$) and standard deviation PDF ($\sigma(\text{PDF})$). We repeat the same procedure to compute the mean power ($\mu(\text{P})$) and standard deviation power ($\sigma(\text{P})$) out of the 1,000 random realizations. We show the corresponding statistics of CAMELS in red and HIDM in blue. It is evident that HIDM fully captures the PDF and power spectra up to the second moment (variance) over all 1,000 random realizations, albeit a minor difference in the large scale power. This shows the ability of HIDM to efficiently emulate the behaviour of CAMELS simulations. While a more comprehensive comparison in terms of higher order statistics might be needed to validate HIDM against CAMELS, it is already challenging for surveys to measure the average power spectrum due to the large observational systematics [e.g. Abdurashidova et al., 2022]. It is worth noting that sampling these 1,000 realizations from HIDM takes 20 seconds on a single GPU as compared to running 1,000 CAMELS simulations with each simulation run can take up to approximately ~ 48 hours on 24 CPU node. This is at least a factor of 10^6 speed gain as compared to CAMELS simulations. This speed again will enable performing inference at the field level in future surveys.

We additionally compute the likelihood over all 1,000 realizations and find $\log P(\text{CAMELS}) = 21.5 \pm 0.7$ bits/dim, whereas the $\log P(\text{HIDM}) = 20.9 \pm 0.4$ bits/dim. These likelihood values are consistent within $1-\sigma$ level. This shows that HIDM has successfully learned the likelihood of HI maps from the CAMELS simulations, and hence this model can efficiently perform implicit inference (i.e. simulation based inference) as soon as conditioning on cosmological and astrophysical parameters is implemented, which we leave to future works.

5 Conclusion

We have introduced HIDM, a generative model of high fidelity HI maps using score-based diffusion models trained on samples from the state-of-the-art CAMELS simulations. We have demonstrated the model ability to generate high quality synthetic examples of large scale HI maps with diversity that is indistinguishable from the target (see Figure 1). We have also shown, over 1,000 random samples, the model fully captures the CAMELS' PDF and power spectra, up to the second moment (variance, see Figure 2). By applying a simple change of variables, the model enables an exact likelihood computation. We have then shown that the model has successfully learned the target likelihood within $1-\sigma$ level.

In this initial attempt, we have verified the ability of score-based diffusion process to successfully model large scale HI maps and provide access to high dimensional likelihood. To maximize the use of this model, conditioning on the various cosmological and astrophysical parameters is required in order to enable field level inference. Conditioning on parameters can be easily implemented as an additional vector embedding next to the time information within the time-dependent U-Net architecture (the Score model). We leave developing the conditional score model on parameters to future works. The HI maps used in this study do not account for various instrumental effects such as foregrounds, thermal noise and angular resolution, which might reduce the applicability of HIDM to analyze maps from future surveys. We leave to future works adding all instrumental effects to perform accurate forecasting to future HI surveys. In addition, we plan to test the scalability of diffusion models as a function of resolution and dimension. The efficiency of HIDM is quite remarkable as compared to CAMELS simulations, achieving at least a factor of 10^6 speed gain, which will enable performing accurate inference at the field level, and hence maximizing the scientific return of future large scale surveys.

6 Broader impact

HIDM represents a step forward towards efficiently analyzing big data expected from large scale surveys. By accurately emulating large scale fields from simulations, HIDM could enable field-level inference, once conditioned on parameters, to minimize loss of complex information due to data compression down to summary statistics. The same analysis can be applied to large scale maps of different observables such as Carbon or Lyman- α emissions, which are the target for many of the upcoming intensity mapping surveys. These surveys will carry a wealth of information that will eventually constrain our understanding of galaxy evolution, formation and large scale structure, and hence the impact of developing emulators of large scale fields, such as HIDM, is wide enough to influence different communities in cosmology and astronomy.

Acknowledgments and Disclosure of Funding

SH acknowledges 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. SH also acknowledges support from the NYU Office of Postdoctoral Affairs and Simons Foundation. SA acknowledges financial support from the South African Radio Astronomy Observatory (SARAO).

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 Labropoulos, 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.
- 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, Carina Cheng, Eloy de Lera Acedo, Joshua S. Dillon, Aaron Ewall-Wice, Gcobisa Fadana, Nicolas Fagnoni, Randall Fritz, Steve R. Furlanetto, Brian Glendenning, Bradley Greig, Jasper Grobbelaar, Bryna J. Hazelton, Jacqueline N. Hewitt, Jack Hickish, Daniel C. Jacobs, Austin Julius, MacCalvin Kariseb, Saul A. Kohn, Telalo Lekalake, Adrian Liu, Anita Loots, David MacMahon, Lourence Malan, Cresshim Malgas, Matthys Maree, Zachary Martinot, Nathan Mathison, Eunice Matsetela, Andrei Mesinger, Miguel F. Morales, Abraham R. Neben, Nipanjana Patra, Samantha Pieterse, Jonathan C. Pober, Nima Razavi-Ghods, Jon Ringuette, James Robnett, Kathryn Rosie, Raddwine Sell, Craig Smith, Angelo Syce, Max Tegmark, Nithyanandan Thyagarajan, Peter K. G. Williams, and Haoxuan Zheng. Hydrogen Epoch of Reionization Array (HERA). *PASP*, 129(974):045001, April 2017. doi: 10.1088/1538-3873/129/974/045001.
- M. P. van Haarlem, M. W. Wise, A. W. Gunst, G. Heald, J. P. McKean, J. W. T. Hessels, A. G. de Bruyn, R. Nijboer, J. Swinbank, R. Fallows, M. Brentjens, A. Nelles, R. Beck, H. Falcke, R. Fender, J. Hörandel, L. V. E. Koopmans, G. Mann, G. Miley, H. Röttgering, B. W. Stappers, R. A. M. J. Wijers, S. Zaroubi, M. van den Akker, A. Alexov, J. Anderson, K. Anderson, A. van Ardenne, M. Arts, A. Asgekar, I. M. Avruch, F. Batejat, L. Bähren, M. E. Bell, M. R. Bell, I. van Bemmell, P. Bennema, M. J. Bentum, G. Bernardi, P. Best, L. Bîrzan, A. Bonafede, A. J. Boonstra, R. Braun, J. Bregman, F. Breitling, R. H. van de Brink, J. Broderick, P. C. Broekema, W. N. Brouw, M. Brüggen, H. R. Butcher, W. van Cappellen, B. Ciardi, T. Coenen, J. Conway, A. Coolen, A. Corstanje, S. Damstra, O. Davies, A. T. Deller, R. J. Dettmar, G. van Diepen, K. Dijkstra, P. Donker, A. Doorduyn, J. Dromer, M. Drost, A. van Duin, J. Eislöffel, J. van Enst, C. Ferrari, W. Frieswijk, H. Gankema, M. A. Garrett, F. de Gasperin, M. Gerbers, E. de Geus, J. M. Grießmeier, T. Grit, P. Gruppen, J. P. Hamaker, T. Hassall, M. Hoefl, H. A. Holties, A. Horneffer, A. van der Horst, A. van Houwelingen, A. Huijgen, M. Iacobelli, H. Intema, N. Jackson, V. Jelic, A. de Jong, E. Juette, D. Kant, A. Karastergiou, A. Koers, H. Kollen, V. I. Kondratiev, E. Kooistra, Y. Koopman, A. Koster, M. Kuniyoshi, M. Kramer, G. Kuper, P. Lambropoulos, C. Law, J. van Leeuwen, J. Lemaître, M. Loose, P. Maat, G. Macario, S. Markoff, J. Masters, R. A. McFadden,

D. McKay-Bukowski, H. Meijering, H. Meulman, M. Mevius, E. Middelberg, R. Millenaar, J. C. A. Miller-Jones, R. N. Mohan, J. D. Mol, J. Morawietz, R. Morganti, D. D. Mulcahy, E. Mulder, H. Munk, L. Nieuwenhuis, R. van Nieuwpoort, J. E. Noordam, M. Norden, A. Noutsos, A. R. Offringa, H. Olofsson, A. Omar, E. Orrú, R. Overeem, H. Paas, M. Pandey-Pommier, V. N. Pandey, R. Pizzo, A. Polatidis, D. Rafferty, S. Rawlings, W. Reich, J. P. de Reijer, J. Reitsma, G. A. Renting, P. Riemers, E. Rol, J. W. Romein, J. Roosjen, M. Ruiten, A. Scaife, K. van der Schaaf, B. Scheers, P. Schellart, A. Schoenmakers, G. Schoonderbeek, M. Serylak, A. Shulevski, J. Sluman, O. Smirnov, C. Sobey, H. Spreeuw, M. Steinmetz, C. G. M. Sterks, H. J. Stiepel, K. Stuurwold, M. Tagger, Y. Tang, C. Tasse, I. Thomas, S. Thoudam, M. C. Toribio, B. van der Tol, O. Usov, M. van Veelen, A. J. van der Veen, S. ter Veen, J. P. W. Verbiest, R. Vermeulen, N. Vermaas, C. Vocks, C. Vogt, M. de Vos, E. van der Wal, R. van Weeren, H. Weggemans, P. Weltevrede, S. White, S. J. Wijnholds, T. Wilhelmsson, O. Wucknitz, S. Yatawatta, P. Zarka, A. Zensus, and J. van Zwieten. LOFAR: The LOW-Frequency ARray. *AAP*, 556:A2, August 2013. doi: 10.1051/0004-6361/201220873.

Željko Ivezić, Steven M. Kahn, J. Anthony Tyson, Bob Abel, Emily Acosta, Robyn Allsman, David Alonso, Yusra AlSayyad, Scott F. Anderson, John Andrew, James Roger P. Angel, George Z. Angeli, Reza Ansari, Pierre Antilogus, Constanza Araujo, Robert Armstrong, Kirk T. Arndt, Pierre Astier, Éric Aubourg, Nicole Auza, Tim S. Axelrod, Deborah J. Bard, Jeff D. Barr, Aurelian Barrau, James G. Bartlett, Amanda E. Bauer, Brian J. Bauman, Sylvain Baumont, Ellen Bechtol, Keith Bechtol, Andrew C. Becker, Jacek Becla, Cristina Beldica, Steve Bellavia, Federica B. Bianco, Rahul Biswas, Guillaume Blanc, Jonathan Blazek, Roger D. Blandford, Josh S. Bloom, Joanne Bogart, Tim W. Bond, Michael T. Booth, Anders W. Borgland, Kirk Borne, James F. Bosch, Dominique Boutigny, Craig A. Brackett, Andrew Bradshaw, William Nielsen Brandt, Michael E. Brown, James S. Bullock, Patricia Burchat, David L. Burke, Gianpietro Cagnoli, Daniel Calabrese, Shawn Callahan, Alice L. Callen, Jeffrey L. Carlin, Erin L. Carlson, Srinivasan Chandrasekharan, Glenaver Charles-Emerson, Steve Chesley, Elliott C. Cheu, Hsin-Fang Chiang, James Chiang, Carol Chirino, Derek Chow, David R. Ciardi, Charles F. Claver, Johann Cohen-Tanugi, Joseph J. Cockrum, Rebecca Coles, Andrew J. Connolly, Kem H. Cook, Asantha Cooray, Kevin R. Covey, Chris Cribbs, Wei Cui, Roc Cutri, Philip N. Daly, Scott F. Daniel, Felipe Daruich, Guillaume Daubard, Greg Daues, William Dawson, Francisco Delgado, Alfred Dellapenna, Robert de Peyster, Miguel de Val-Borro, Seth W. Digel, Peter Doherty, Richard Dubois, Gregory P. Dubois-Felsmann, Josef Durech, Frossie Economou, Tim Eifler, Michael Eracleous, Benjamin L. Emmons, Angelo Fausti Neto, Henry Ferguson, Enrique Figueroa, Merlin Fisher-Levine, Warren Focke, Michael D. Foss, James Frank, Michael D. Freeman, Emmanuel Gangler, Eric Gawiser, John C. Geary, Perry Gee, Marla Geha, Charles J. B. Gessner, Robert R. Gibson, D. Kirk Gilmore, Thomas Glanzman, William Glick, Tatiana Goldina, Daniel A. Goldstein, Iain Goodenow, Melissa L. Graham, William J. Gressler, Philippe Gris, Leanne P. Guy, Augustin Guyonnet, Gunther Haller, Ron Harris, Patrick A. Hascall, Justine Haupt, Fabio Hernandez, Sven Herrmann, Edward Hileman, Joshua Hobbitt, John A. Hodgson, Craig Hogan, James D. Howard, Dajun Huang, Michael E. Huffer, Patrick Ingraham, Walter R. Innes, Suzanne H. Jacoby, Bhuvnesh Jain, Fabrice Jammes, M. James Jee, Tim Jenness, Garrett Jernigan, Darko Jevremović, Kenneth Johns, Anthony S. Johnson, Margaret W. G. Johnson, R. Lynne Jones, Claire Juramy-Gilles, Mario Jurić, Jason S. Kalirai, Nitya J. Kallivayalil, Bryce Kalmbach, Jeffrey P. Kantor, Pierre Karst, Mansi M. Kasliwal, Heather Kelly, Richard Kessler, Veronica Kinnison, David Kirkby, Lloyd Knox, Ivan V. Kotov, Victor L. Krabbendam, K. Simon Krughoff, Petr Kubánek, John Kuczewski, Shri Kulkarni, John Ku, Nadine R. Kurita, Craig S. Lage, Ron Lambert, Travis Lange, J. Brian Langton, Laurent Le Guillou, Deborah Levine, Ming Liang, Kian-Tat Lim, Chris J. Lintott, Kevin E. Long, Margaux Lopez, Paul J. Lotz, Robert H. Lupton, Nate B. Lust, Lauren A. MacArthur, Ashish Mahabal, Rachel Mandelbaum, Thomas W. Markiewicz, Darren S. Marsh, Philip J. Marshall, Stuart Marshall, Morgan May, Robert McKercher, Michelle McQueen, Joshua Meyers, Myriam Migliore, Michelle Miller, David J. Mills, Connor Miraval, Joachim Moeyens, Fred E. Moolekamp, David G. Monet, Marc Moniez, Serge Monkewitz, Christopher Montgomery, Christopher B. Morrison, Fritz Mueller, Gary P. Muller, Freddy Muñoz Arancibia, Douglas R. Neill, Scott P. Newbry, Jean-Yves Nief, Andrei Nomerotski, Martin Nordby, Paul O'Connor, John Oliver, Scot S. Olivier, Knut Olsen, William O'Mullane, Sandra Ortiz, Shawn Osier, Russell E. Owen, Reynald Pain, Paul E. Palecek, John K. Parejko, James B. Parsons, Nathan M. Pease, J. Matt Peterson, John R. Peterson, Donald L. Petravick, M. E. Libby Patrick, Cathy E. Petry, Francesco Pierfederici, Stephen Pietrowicz, Rob Pike, Philip A. Pinto, Raymond Plante, Stephen Plate, Joel P. Plutchak, Paul A. Price, Michael

- Prouza, Veljko Radeka, Jayadev Rajagopal, Andrew P. Rasmussen, Nicolas Regnault, Kevin A. Reil, David J. Reiss, Michael A. Reuter, Stephen T. Ridgway, Vincent J. Riot, Steve Ritz, Sean Robinson, William Roby, Aaron Roodman, Wayne Rosing, Cecille Roucelle, Matthew R. Rumore, Stefano Russo, Abhijit Saha, Benoit Sassolas, Terry L. Schalk, Pim Schellart, Rafe H. Schindler, Samuel Schmidt, Donald P. Schneider, Michael D. Schneider, William Schoening, German Schumacher, Megan E. Schwamb, Jacques Sebag, Brian Selvy, Glenn H. Sembroski, Lynn G. Seppala, Andrew Serio, Eduardo Serrano, Richard A. Shaw, Ian Shipsey, Jonathan Sick, Nicole Silvestri, Colin T. Slater, J. Allyn Smith, R. Chris Smith, Shahram Sobhani, Christine Soldahl, Lisa Storrie-Lombardi, Edward Stover, Michael A. Strauss, Rachel A. Street, Christopher W. Stubbs, Ian S. Sullivan, Donald Sweeney, John D. Swinbank, Alexander Szalay, Peter Takacs, Stephen A. Tether, Jon J. Thaler, John Gregg Thayer, Sandrine Thomas, Adam J. Thornton, Vaikunth Thukral, Jeffrey Tice, David E. Trilling, Max Turri, Richard Van Berg, Daniel Vanden Berk, Kurt Vetter, Françoise Virieux, Tomislav Vucina, William Wahl, Lucianne Walkowicz, Brian Walsh, Christopher W. Walter, Daniel L. Wang, Shin-Yawn Wang, Michael Warner, Oliver Wiecha, Beth Willman, Scott E. Winters, David Wittman, Sidney C. Wolff, W. Michael Wood-Vasey, Xiuqin Wu, Bo Xin, Peter Yoachim, and Hu Zhan. LSST: From Science Drivers to Reference Design and Anticipated Data Products. *ApJ*, 873(2):111, March 2019. doi: 10.3847/1538-4357/ab042c.
- D. Spergel, N. Gehrels, C. Baltay, D. Bennett, J. Breckinridge, M. Donahue, A. Dressler, B. S. Gaudi, T. Greene, O. Guyon, C. Hirata, J. Kalirai, N. J. Kasdin, B. Macintosh, W. Moos, S. Perlmutter, M. Postman, B. Rauscher, J. Rhodes, Y. Wang, D. Weinberg, D. Benford, M. Hudson, W. S. Jeong, Y. Mellier, W. Traub, T. Yamada, P. Capak, J. Colbert, D. Masters, M. Penny, D. Savransky, D. Stern, N. Zimmerman, R. Barry, L. Bartusek, K. Carpenter, E. Cheng, D. Content, F. Dekens, R. Demers, K. Grady, C. Jackson, G. Kuan, J. Kruk, M. Melton, B. Nemati, B. Parvin, I. Poberezhskiy, C. Peddie, J. Ruffa, J. K. Wallace, A. Whipple, E. Wollack, and F. Zhao. Wide-Field Infrared Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report. *arXiv e-prints*, art. arXiv:1503.03757, March 2015.
- Olivier Doré, Jamie Bock, Matthew Ashby, Peter Capak, Asantha Cooray, Roland de Putter, Tim Eifler, Nicolas Flagey, Yan Gong, Salman Habib, Katrin Heitmann, Chris Hirata, Woong-Seob Jeong, Raj Katti, Phil Korngut, Elisabeth Krause, Dae-Hee Lee, Daniel Masters, Phil Mausekopf, Gary Melnick, Bertrand Mennesson, Hien Nguyen, Karin Öberg, Anthony Pullen, Alvise Raccanelli, Roger Smith, Yong-Seon Song, Volker Toll, Steve Unwin, Tejaswi Venumadhav, Marco Viero, Mike Werner, and Mike Zemcov. Cosmology with the SPHEREX All-Sky Spectral Survey. *arXiv e-prints*, art. arXiv:1412.4872, December 2014.
- 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, Cyril Colombo, Adriano Calvi, Elena Maiorano, Osvaldo Piersanti, Sylvain Prezelus, Pierluigi Rosato, Jacques Pinel, Hans Rozemeijer, Valentina Lesna, Paolo Musi, Marco Sias, Alberto Anselmi, Vincent Cazaubiel, Ludovic Vaillon, Yannick Mellier, Jérôme Amiaux, Michel Berthé, Marc Sauvage, Ruyman Azzollini, Mark Cropper, Sabrina Pottinger, Knud Jahnke, Anne Ealet, Thierry Maciaszek, Fabio Pasian, Andrea Zaccari, Roberto Scaramella, John Hoar, Ralf Kohley, Roland Vavrek, Andreas Rudolph, and Micha Schmidt. The Euclid mission design. In Howard A. MacEwen, Giovanni G. Fazio, Makenzie Lystrup, Natalie Batalha, Nicholas Siegler, and Edward C. Tong, editors, *Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave*, volume 9904 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 99040O, July 2016. doi: 10.1117/12.2230762.
- Juan Zamudio-Fernandez, Atakan Okan, Francisco Villaescusa-Navarro, Seda Bilaloglu, Asena Derin Cengiz, Siyu He, Laurence Perreault Levasseur, and Shirley Ho. HIGAN: Cosmic Neutral Hydrogen with Generative Adversarial Networks. *arXiv e-prints*, art. arXiv:1904.12846, April 2019. doi: 10.48550/arXiv.1904.12846.
- Francisco Villaescusa-Navarro, Daniel Anglés-Alcázar, Shy Genel, David N. Spergel, Rachel S. Somerville, Romeel Dave, Annalisa Pillepich, Lars Hernquist, Dylan Nelson, Paul Torrey, Desika Narayanan, Yin Li, Oliver Philcox, Valentina La Torre, Ana Maria Delgado, Shirley Ho, Sultan Hassan, Blakesley Burkhart, Digvijay Wadekar, Nicholas Battaglia, Gabriella Contardo, and Greg L. Bryan. The CAMELS Project: Cosmology and Astrophysics with Machine-learning Simulations. *ApJ*, 915(1):71, July 2021a. doi: 10.3847/1538-4357/abf7ba.

- Digvijay Wadekar, Francisco Villaescusa-Navarro, Shirley Ho, and Laurence Perreault-Levasseur. HI-net: Generating Neutral Hydrogen from Dark Matter with Neural Networks. *ApJ*, 916(1):42, July 2021. doi: 10.3847/1538-4357/ac033a.
- Sambatra Andrianomena, Sultan Hassan, and Francisco Villaescusa-Navarro. Invertible mapping between fields in CAMELS. *arXiv e-prints*, art. arXiv:2303.07473, March 2023. doi: 10.48550/arXiv.2303.07473.
- François Lanusse, Rachel Mandelbaum, Siamak Ravanbakhsh, Chun-Liang Li, Peter Freeman, and Barnabás Póczos. Deep generative models for galaxy image simulations. *MNRAS*, 504(4): 5543–5555, July 2021. doi: 10.1093/mnras/stab1214.
- Ting-Yun Cheng, Marc Huertas-Company, Christopher J. Conselice, Alfonso Aragón-Salamanca, Brant E. Robertson, and Nesar Ramachandra. Beyond the hubble sequence - exploring galaxy morphology with unsupervised machine learning. *MNRAS*, 503(3):4446–4465, May 2021. doi: 10.1093/mnras/stab734.
- 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, Michael Eickenberg, Desika Narayanan, Shirley Ho, and Sambatra Andrianomena. HIFLOW: Generating Diverse HI Maps and Inferring Cosmology while Marginalizing over Astrophysics Using Normalizing Flows. *ApJ*, 937(2):83, October 2022. doi: 10.3847/1538-4357/ac8b09.
- Roy Friedman and Sultan Hassan. HIGlow: Conditional Normalizing Flows for High-Fidelity HI Map Modeling. *arXiv e-prints*, art. arXiv:2211.12724, November 2022. doi: 10.48550/arXiv.2211.12724.
- Biwei Dai and Uroš Seljak. Translation and rotation equivariant normalizing flow (TRENF) for optimal cosmological analysis. *MNRAS*, 516(2):2363–2373, October 2022. doi: 10.1093/mnras/stac2010.
- Alexandre Adam, Adam Coogan, Nikolay Malkin, Ronan Legin, Laurence Perreault-Levasseur, Yashar Hezaveh, and Yoshua Bengio. Posterior samples of source galaxies in strong gravitational lenses with score-based priors. *arXiv e-prints*, art. arXiv:2211.03812, November 2022. doi: 10.48550/arXiv.2211.03812.
- Nayantara Mudur and Douglas P. Finkbeiner. Can denoising diffusion probabilistic models generate realistic astrophysical fields? *arXiv e-prints*, art. arXiv:2211.12444, November 2022. doi: 10.48550/arXiv.2211.12444.
- Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations. *arXiv e-prints*, art. arXiv:2011.13456, November 2020. doi: 10.48550/arXiv.2011.13456.
- Francisco Villaescusa-Navarro, Shy Genel, Daniel Anglés-Alcázar, Leander Thiele, Romeel Dave, Desika Narayanan, Andrina Nicola, Yin Li, Pablo Villanueva-Domingo, Benjamin Wandelt, David N. Spergel, Rachel S. Somerville, Jose Manuel Zorrilla Matilla, Faizan G. Mohammad, Sultan Hassan, Helen Shao, Digvijay Wadekar, Michael Eickenberg, Kaze W. K. Wong, Gabriella Contardo, Yongseok Jo, Emily Moser, Erwin T. Lau, Luis Fernando Machado Poletti Valle, Lucia A. Perez, Daisuke Nagai, Nicholas Battaglia, and Mark Vogelsberger. The CAMELS Multifield Dataset: Learning the Universe’s Fundamental Parameters with Artificial Intelligence. *arXiv e-prints*, art. arXiv:2109.10915, September 2021b.
- Rainer Weinberger, Volker Springel, Lars Hernquist, Annalisa Pillepich, Federico Marinacci, Rüdiger Pakmor, Dylan Nelson, Shy Genel, Mark Vogelsberger, Jill Naiman, and Paul Torrey. Simulating galaxy formation with black hole driven thermal and kinetic feedback. *MNRAS*, 465(3):3291–3308, March 2017. doi: 10.1093/mnras/stw2944.
- Annalisa Pillepich, Volker Springel, Dylan Nelson, Shy Genel, Jill Naiman, Rüdiger Pakmor, Lars Hernquist, Paul Torrey, Mark Vogelsberger, Rainer Weinberger, and Federico Marinacci. Simulating galaxy formation with the IllustrisTNG model. *MNRAS*, 473(3):4077–4106, January 2018. doi: 10.1093/mnras/stx2656.

Mark S. Graham, Walter H. L. Pinaya, Petru-Daniel Tudosiu, Parashkev Nachev, Sebastien Ourselin, and M. Jorge Cardoso. Denoising diffusion models for out-of-distribution detection. *arXiv e-prints*, art. arXiv:2211.07740, November 2022. doi: 10.48550/arXiv.2211.07740.

Zara Abdurashidova, James E. Aguirre, Paul Alexander, Zaki S. Ali, Yanga Balfour, Adam P. Beardsley, Gianni Bernardi, Tashalee S. Billings, Judd D. Bowman, Richard F. Bradley, Philip Bull, Jacob Burba, Steve Carey, Chris L. Carilli, Carina Cheng, David R. DeBoer, Matt Dexter, Eloy de Lera Acedo, Taylor Dibblee-Barkman, Joshua S. Dillon, John Ely, Aaron Ewall-Wice, Nicolas Fagnoni, Randall Fritz, Steven R. Furlanetto, Kingsley Gale-Sides, Brian Glendenning, Deepthi Gorthi, Bradley Greig, Jasper Grobelaar, Ziyaad Halday, Bryna J. Hazelton, Jacqueline N. Hewitt, Jack Hickish, Daniel C. Jacobs, Austin Julius, Nicholas S. Kern, Joshua Kerrigan, Piyanat Kittiwisit, Saul A. Kohn, Matthew Kolopanis, Adam Lanman, Paul La Plante, Telalo Lekalake, David Lewis, Adrian Liu, David MacMahon, Lourence Malan, Cresshim Malgas, Matthys Maree, Zachary E. Martinot, Eunice Matsetela, Andrei Mesinger, Mathakane Molewa, Miguel F. Morales, Tshegofalang Mosiane, Steven G. Murray, Abraham R. Neben, Bojan Nikolic, Chuneeta D. Nunhokee, Aaron R. Parsons, Nipanjana Patra, Robert Pascua, Samantha Pieterse, Jonathan C. Pober, Nima Razavi-Ghods, Jon Ringuette, James Robnett, Kathryn Rosie, Peter Sims, Saurabh Singh, Craig Smith, Angelo Syce, Nithyanandan Thyagarajan, Peter K. G. Williams, Haoxuan Zheng, and HERA Collaboration. First Results from HERA Phase I: Upper Limits on the Epoch of Reionization 21 cm Power Spectrum. *ApJ*, 925(2):221, February 2022. doi: 10.3847/1538-4357/ac1c78.