
3D Reconstruction of Dark Matter Fields with Diffusion Models: Towards Application to Galaxy Surveys

Core Francisco Park^{*12} Nayantara Mudur¹² Carolina Cuesta-Lazaro²³ Yueying Ni² Victoria Ono⁴
Douglas P. Finkbeiner¹²

Abstract

Probabilistic diffusion models have shown great success in conditional image synthesis. In this work, we develop a high-resolution 3D diffusion model to reconstruct the dark matter density field from a galaxy distribution. We train a pixel space diffusion model at different resolutions on the CAMELS simulation and achieve good agreement in visual quality and summary statistics. However, we identify some challenges in scaling up the resolution. We then analyze the model’s ability to capture variations in simulation parameters and conclude that the model indeed captures the right change in the field when changing Ω_m . Next, we train our model on a more realistic dataset where the input conditioning consists of mass thresholded galaxy catalogs from CAMELS and find excellent adaptation of diffusion models to low galaxy density inputs. Finally, we show a preliminary application to a real galaxy catalog. Our results suggest that diffusion models are a powerful method to reconstruct the 3D dark matter field from galaxies.

1. Introduction

Λ CDM, the standard model of cosmology, postulates that 85% of the total matter of the universe consists of dark matter. Naturally, the nature of dark matter and the validity of the Λ CDM model are one of the most important questions in modern astrophysics. While dark matter is not directly observable yet by any means, cosmological N-body and hydrodynamic simulations allow us to calculate its hypothetical distribution given a model (Springel et al., 2005; Nelson et al., 2015; Garrison et al., 2019; Villaescusa-Navarro et al.,

2020). Some of these simulations also generate synthetic galaxy catalogs which can be connected to real galaxy surveys such as DESI (DESI Collaboration et al., 2016), Euclid (Laureijs et al., 2011), Roman (Spergel et al., 2015), or Rubin (LSST Science Collaboration et al., 2009). However, since these simulations are initialized with random initial conditions, their final states can only match our observable universe in a statistical sense, thus posing an inverse problem of finding the dark matter distribution compatible with our current universe.

The goal of reconstructing the underlying dark matter field in our universe has been addressed at different scales with different methods. At the largest scales, BORG (Jasche & Lavaux, 2019) uses a Bayesian method to reconstruct initial matter density fields. At smaller scales ~ 50 Mpc/h, (Hong et al., 2021) used deterministic convolutional neural networks and (Pfeifer et al., 2023) used constrained simulations to reconstruct the dark matter densities around the Milky Way.

Over the past decade, the machine learning (ML) community has developed probabilistic generative models such as variational autoencoders (Kingma & Welling, 2013), normalizing flows (Papamakarios et al., 2021) and diffusion models (Sohl-Dickstein et al., 2015), that model the data likelihood and enable sampling from the learned distribution. Diffusion generative models, in particular, (Ho et al., 2020; Kingma et al., 2021) have excelled at text-to-image synthesis tasks (Ramesh et al., 2021; Saharia et al., 2022; Rombach et al., 2022). Diffusion models involve learning a mapping between the standard normal distribution and the target distribution. In the noising direction, incremental amounts of noise are added to samples from the target distribution, constituting a discretization of a stochastic differential equation (SDE). The reverse SDE that transforms samples from the noise distribution to those from the target, can be derived by learning the score of the noised probability distributions (Anderson, 1982; Song et al., 2021). A neural network is used to parameterize the score model.

Recent works apply these advances in high dimensional inference to the dark matter reconstruction problem. Le-gin et al. (2024) used score based diffusion models to re-

^{*}Equal contribution ¹Department of Physics, Harvard University, MA, USA, 02138 ²Center for Astrophysics, MA, USA, 02138 ³Massachusetts Institute of Technology, MA, USA, 02138 ⁴Harvard College, MA, USA, 02138. Correspondence to: Core-Francisco Park <corefranciscopark@g.harvard.edu>.

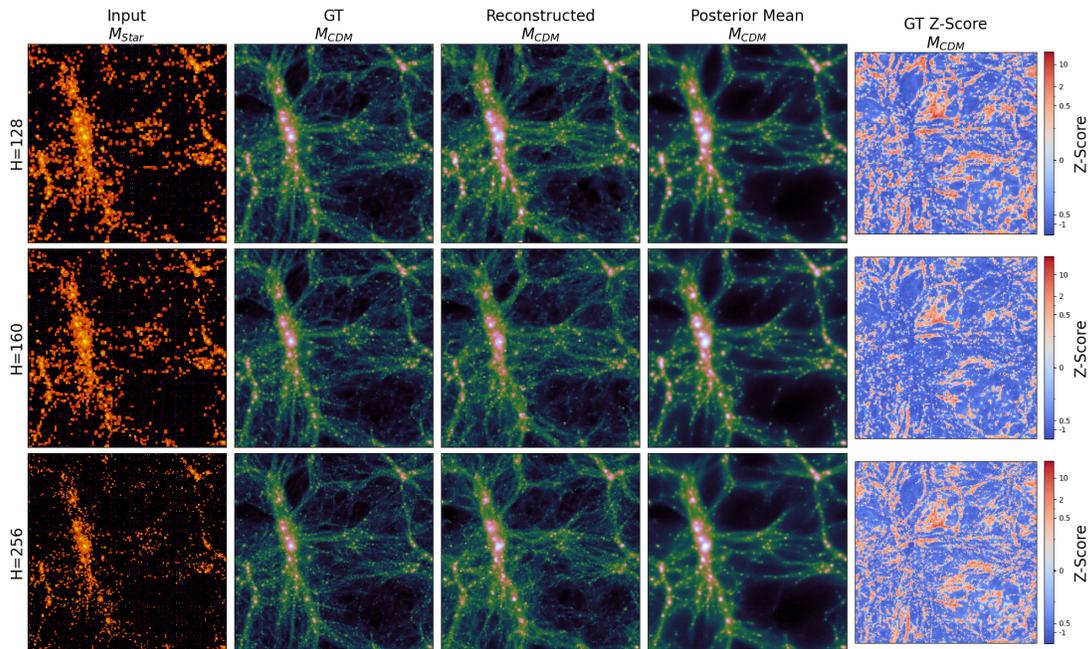


Figure 1. Generation Results of the Diffusion Model depending on Resolution All panels show a 2D projection integrating a $12.5 \text{ Mpc } h^{-1}$ thick slice of the density field. **First Column:** The input Stellar Mass map. **Second Column:** The ground truth(GT) dark matter density field. **Third Column:** A random sample from the reconstructed dark matter density field. **Fourth column:** The posterior mean over 128 generated samples of the dark matter density field. **Fifth column:** The signed Z-Score of the ground truth dark matter density field with respect to the posterior distribution from the diffusion model.

construct initial conditions given current dark matter fields while Park et al. (2023b); Ono et al. (2024) uses a conditional diffusion model to reconstruct the 2D dark matter density fields from stellar mass fields.

In this work, we scale spatially conditioned diffusion models to 3 dimensions to solve the inverse problem of generating the underlying dark matter distribution that resulted in an observed galaxy distribution. We focus on how the model performance scales when changing resolutions and sparsity of the data as well as how the model adapts to astrophysical parameter changes.

The paper is organized as follows. In Sec. 2, we introduce the datasets used in this work. In Sec. 3, we briefly summarize the models and metrics we use. In Sec. 4, we discuss our main results on resolution scaling (Sec. 4.1), cosmological parameter variation (Sec. 4.2), galaxy sparsity (Sec. 4.3) and application to real data (Sec. 4.4). Our main contributions are:

1. We train 3D diffusion models to solve the galaxy to dark matter inverse problem in cosmology.
2. We evaluate the 3D diffusion model’s statistical quality at different resolutions. (Sec. 4.1)

3. We evaluate the diffusion model’s ability to properly capture the effect of modulating simulation parameters, here Ω_m . (Sec. 4.2)
4. We demonstrate the diffusion model’s ability to generate detailed fields from sparse galaxy catalogs. (Sec. 4.3)
5. We apply our method to the galaxy catalog Cosmicflows-3 (Tully et al., 2016) and show preliminary results reconstructing the dark matter density around the Milky Way. (Sec. 4.4)

2. Data

CAMELS The CAMELS simulation suite(Villaescusa-Navarro et al., 2021a) is a set of state-of-the-art hydrodynamical simulations built for machine learning applications. The time evolution of 256^3 dark matter particles and 256^3 gas elements are tracked from redshift $z = 127$ to $z = 0$, with different astrophysical/cosmological parameters and initial conditions. In this work, we use the CAMELS Multifield Dataset(Villaescusa-Navarro et al., 2021b), which is a collection of 3D density fields generated from the simulations. Please refer to App. A.1 for more information.

Cosmicflows-3 Cosmicflows-3 (Tully et al., 2016) is a dataset compiling 17,669 galaxies’ sky coordinates and distances. Hong et al. (2021) used this dataset to reconstruct the local dark matter density deterministically. In this work, we select galaxies near the Milky Way by restricting the B-band magnitude to $M_B < -15$ and the galactic latitude to $|b| > 10^\circ$, resulting in 6,159 galaxies. Please refer to App. A.1 for more information.

3. Methods

3.1. Models

Variational Diffusion Models Variational Diffusion Models, introduced by Kingma et al. (2021) is a continuous time pixel space diffusion model where the noise schedule is parametrized by a parametric function $\gamma(t)$ and can be optimized jointly with the diffusion process. Variational Diffusion Models focus on approximating the variational lower bound(or ELBO) of the generative process with respect to the data. In the forward diffusion process, Gaussian noise is added to the data at each step until the data is indistinguishable to pure Gaussian noise while the backward process attempts to predict the noise added to generate data from noise. We use a conditional version of a U-Net (Ronneberger et al., 2015) like architecture as the noise prediction network and use a learned linear noise schedule (please refer to App.A.2 or Kingma et al. (2021) for further details). The noise prediction network is conditioned on the input stellar mass field as well as on the 6 astrophysical/cosmological parameters.

3.2. Summary Statistics

We evaluate the Power Spectrum (P(k)), the Probability Density Function (PDF) and the Reduced Wavelet Scattering Transform (Bruna & Mallat, 2012; Regalado-Saint Blancard et al., 2020) as summary statistics to evaluate our generations.

Power Spectrum The power spectrum is the angular averaged squared amplitude of the Fourier transform of a field. $P(k)$ can be thus be interpreted as an average power a field has at wavenumber k , corresponding to a spatial scale of $2\pi/k$. For a homogeneous and isotropic Gaussian random field, the power spectrum contains all the information describing the field. For these fields, odd numbered higher order spectra (e.g. bispectrum) vanish and even numbered ones (e.g. trispectra) can be expressed as a function of the power spectrum.

Probability Density Function The Probability Density Function(PDF), also called the one-point function in cosmology is simply the pixel values histogram of the field. We use a logarithmic density bins to compute the PDF.

Reduced Wavelet Scattering Transform The Wavelet Scattering Transform, introduced by Bruna & Mallat (2012), is a statistical probe designed to extract non-Gaussian properties the field. The first order coefficients are computed by convolving the field with a set of wavelets, extracting components around a certain wavevector, and integrating over the field. The second order coefficients are calculated by applying a non-linear transform (e.g. abs) on the convolved field from the first stage then convolving it by the set of wavelets again. The Reduced Wavelet Scattering Transform(RWST) (Regalado-Saint Blancard et al., 2020), efficiently reduces the number of coefficient by taking the angular average of these coefficients, which is a very efficient compression when the field is isotropic. We use 4 spatial scales and 4 angular scales resulting in 78 RWST coefficients. We use the implementation distributed in Park et al. (2023a).

4. Results

4.1. Performance at different resolutions

We explore the ability of a 3D pixel-space diffusion model to generate high resolution samples. We run experiments on the same density field with the same physical size but sampled at 128^3 , 160^3 and 256^3 resolutions. We draw $N=128$ dark matter fields corresponding to an input stellar mass field to validate our findings. Fig. 1 illustrates the input stellar mass field, the ground truth dark matter field, and the reconstructed dark matter field, all projected in 2 dimensions. One can visually check the quality of the generations. (We show additional generations in Fig. 10) The posterior mean is the mean over many diffusion samples, and one can easily identify which filaments are most robustly reproduced, and thus are more likely to exist. The Z-score is evaluated as $\frac{M_{CDM}^{GT} - \langle M_{CDM}^{Gen} \rangle}{Std(M_{CDM}^{Gen})}$, and shows which parts of the field deviate from the posterior mean more or less than the expected scatter given by the posterior standard deviation.

Since Park et al. (2023b) and Ono et al. (2024) have shown that diffusion models can generate 2D dark matter fields at 256^2 resolution with a statistically consistent P(k) and PDF, we project the generated 3D dark matter density fields to 2D by integrating over $12.5 \text{ Mpc } h^{-1}$ in one direction to make a comparison to these results. Fig. 2 shows these results. While the diffusion model can, in all three resolutions, generate fields with a power spectrum close to the ground truth, we see a significant lack of power at small scales (high k) for the 256^3 model. This effect is shown more drastically in the last panel of Fig. 2.

In Fig. 3 we show the PDF of the 2D projections. We find that despite the disagreement in the power spectrum, the 3D generations at 256^3 reproduce the PDF with good accuracy.

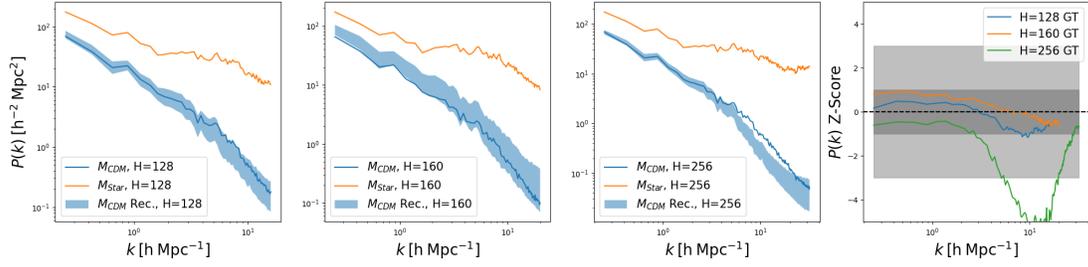


Figure 2. Power Spectrum of 2D Projections from 3D Generations The power spectrum range is calculated from $12.5 \text{ Mpc } h^{-1}$ 2D projections of 128 generated 3D dark matter density fields at resolutions $H = 128, 160$ and 256 in the first 3 columns. The last column depicts the z score of the ground truth dark matter power spectrum relative to the distribution obtained from the diffusion model for each resolution.

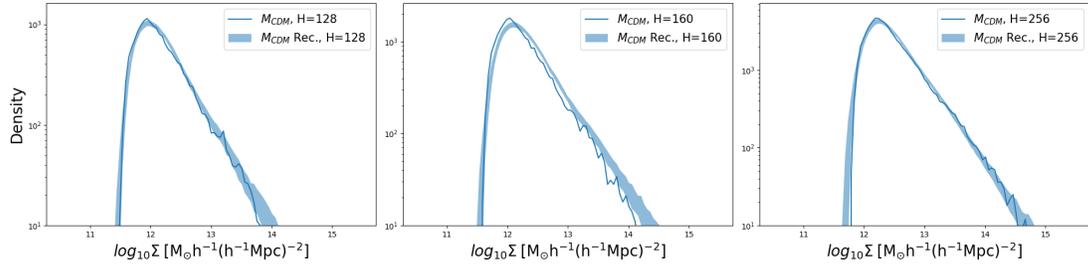


Figure 3. Probability Density Function of 2D Projections from the 3D Generation The probability density function range is calculated from $12.5 \text{ Mpc } h^{-1}$ 2D projections of 128 generated 3D dark matter density fields. The three panels respectively depicts the PDF at a resolution of $H=128, 160$ and 256 .

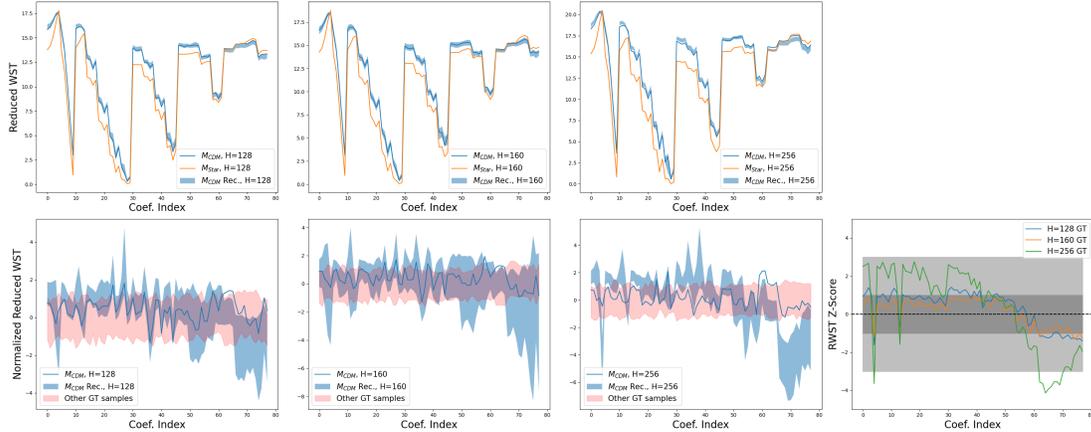


Figure 4. The Reduced Wavelet Scattering Transform Statistics of 2D Projections from the 3D Generation The RWST range is calculated from $12.5 \text{ Mpc } h^{-1}$ 2D projections of 128 generated 3D dark matter density fields. **Top Row:** We show, for resolutions $H = 128, 160$ and 256 , the raw RWST coefficients for the input stellar mass map, the ground truth dark matter map and the diffusion model generations. **Bottom Row:** We show in the first three panels the normalized RWST, where we normalize every coefficient by its mean and standard deviation over the whole CV set, for resolutions $H = 128, 160$ and 256 , and its variations within *different* ground truth fields. The last column depicts the z -score of the RWST of the ground truth field relative to the distribution from the diffusion model.

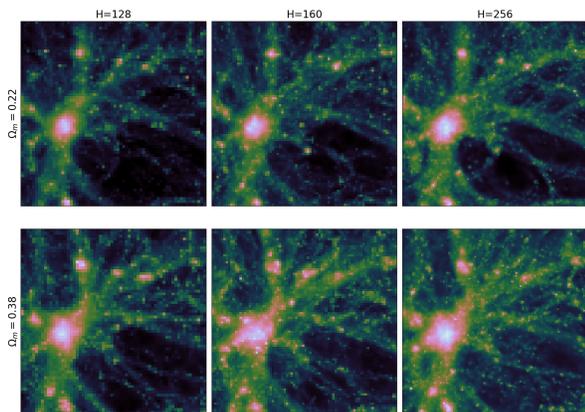


Figure 5. Generated Dark matter fields depending on Ω_m
Top Row: Generated Dark Matter sample with $\Omega_m = 0.22$
Bottom Row: Generated Dark Matter sample with $\Omega_m = 0.38$

We show additional results on the full 3D power spectrum and PDF in App. C.1. With these results, our conclusion remains that while 256^3 models approximate the PDF well, they tend to exhibit a lack of power at small scales.

Now we turn to a more complex statistic, the RWST, which can capture the non-Gaussian structure of the $z = 0$ dark matter density field. Fig. 4 shows the RWST coefficients of the ground truth dark matter map and the generated dark matter maps. We find an excellent agreement of nearly all coefficients at $H=128$ and 160 . Moreover, we find that the diffusion model correctly captures the per-initial condition variation of the RWST coefficients as seen in the second row of Fig. 4. However, as we have seen for the power spectrum, the agreement, and thus the Z-Score, for the 256^3 model is slightly degraded.

4.2. Astrophysical Conditioning

Another important aspect of a generative model applied to cosmology is its ability to capture variations in astrophysical/cosmological parameters. As one can see in Fig. 5, the variation of an underlying parameter (here Ω_m) changes the structure of the cosmic web and halo structures in our diffusion model, across all resolutions $H=[128, 160, 256]$. As intuitively expected a larger Ω_m results in a more massive halo, at least visually.

We verify this finding with summary statistics in Fig. 6. As we can see in the first two panels in Fig. 6, we find that the diffusion model is able to approximate the 3D power spectra ratios $\frac{P(k)_{\Omega_m=0.22}}{P(k)_{\Omega_m=0.3}}$ and $\frac{P(k)_{\Omega_m=0.38}}{P(k)_{\Omega_m=0.3}}$ accurately. Thus, the model is able to capture the modulation of the 2D power spectrum for fields corresponding to different parameters.

However, we find that similarly to other statistics, this agreement isn't as accurate with the 256^3 model. We show the 2D power spectrum results in App. C.2, Fig. 14. We find identical conclusions there.

4.3. Scaling to lower galaxy densities

Now, we attempt to analyze the adaptation of 3D diffusion models to a more realistic data. We process the halo catalogues from CAMELS (Villaescusa-Navarro et al., 2021a) and extract subhaloes as galaxies and apply a mass threshold of 10^7 , 10^8 , $10^9 M_\odot$. We generate galaxy overdensity fields from these catalogues and train a diffusion model for each density at 128^3 resolution. The results are illustrated in Fig. 7. We still find samples with good visual quality at very low galaxy densities. (More samples are shown in App. 11.) Since some constraining information from low mass galaxies are lost, we directly find that the smaller filaments in the posterior mean of the first row of Fig. 7 are lost in the posterior mean of the last row, as expected. Another effect of this change in the posterior distribution is clearly seen in the Z-Score. We find that as the input field gets sparser, the Z-Score map grows large valued regions, especially in the voids.

It is important to check whether sparse galaxy samples can cause a bias in the power spectrum. As we can see in Fig. 8, the generated distribution of the power spectrum agrees well with the ground truth. This can also be seen in the Z-Score in Fig. 8. We thus conclude that diffusion models are relatively robust to sparse conditioning samples and can generate consistent per-sample statistics while not over/under estimating the posterior mean.

4.4. Application to Cosmicflows-3

Finally, we apply our diffusion model to a real galaxy catalogue. We apply our model conditioned on synthetic galaxy catalogs with $M_{Gal} > 10^9 M_\odot$ on the Cosmicflows-3 (Tully et al., 2016) dataset. We bin the 6,159 selected galaxies (See App. A.1) into a grid with the same physical pixel resolution ($\text{Mpc } h^{-1} / n_{pix}$) as the 128^3 model then apply our model to generate 24 posterior dark matter density fields. One sample and the posterior mean is shown in Fig. 9, where we show 2D slices in x, y, z and a 3D projection. We find many similarities with the dark matter map in Hong et al. (2021), like the structure of the Virgo cluster or Ursa N/S. We also find that our *posterior* mean map looks very similar to Hong et al. (2021)'s map, which is exactly as expected for a deterministic CNN.

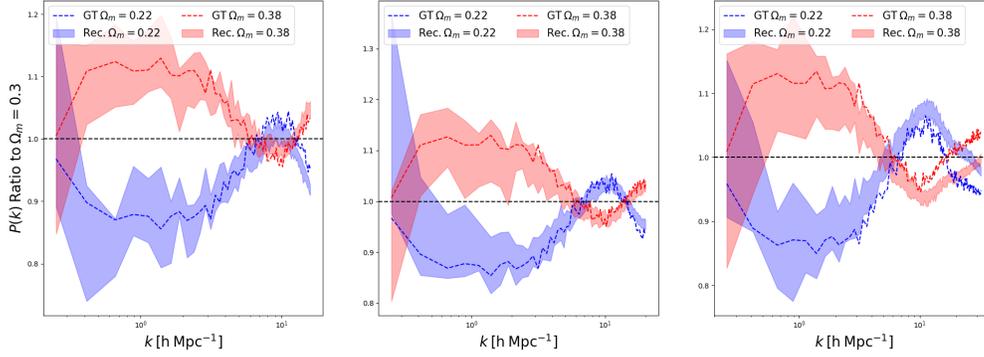


Figure 6. Ratio of the 3D Power Spectrum at a modified value of Ω_m relative to the fiducial value. We depict the generated power spectra range and the GT power spectrum in blue when $\Omega_m = 0.22$ and in red when $\Omega_m = 0.38$. The panels from left to right corresponds to resolution $H=128, 160$ and 256 .

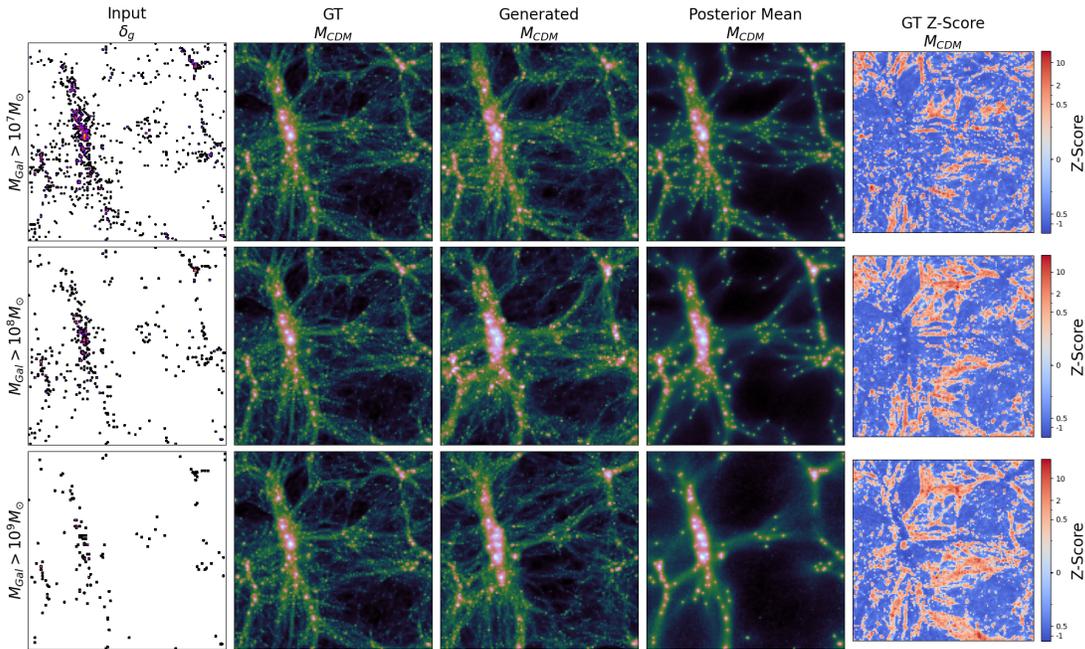


Figure 7. Generation Results of the Diffusion Model depending on the galaxy sparsity. **First Column:** The input Galaxy Overdensity map. **Second Column:** The ground truth(GT) dark matter density field. **Third Column:** A random sample from the reconstructed dark matter density field. **Fourth column:** The posterior mean over 128 generated samples of the dark matter density field. **Fifth column:** The signed Z-Score of the ground truth dark matter density field with respect to the posterior distribution of diffusion model.

5. Discussion

5.1. Limitations

We describe some limitations of our study:

3D Diffusion Models 3D diffusion models are memory/compute intensive and their generation speeds are some-

what limited on small number of GPUs. We were thus not able to do a proper hyperparameter grid search or inference time step adjustment.

Diffusion Loss Since our diffusion model’s generation quality keeps improving after the loss has saturated, it is rather ambiguous when to halt the training. In this study we chose

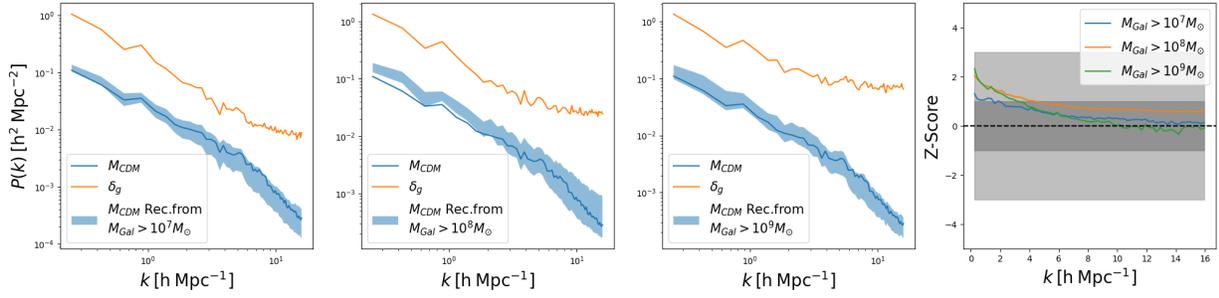


Figure 8. Power Spectrum of 2D Projections from 3D Generations from models conditioned on galaxy catalogs The power spectrum is calculated from $12.5 \text{ Mpc } h^{-1}$ 2D projections of 128 generated 3D dark matter density fields. The first three panels depict the power spectrum of the 2D projected fields when the model is conditioned on galaxies of $M_{Gal} > 10^7 M_{\odot}$, $M_{Gal} > 10^8 M_{\odot}$ and $M_{Gal} > 10^9 M_{\odot}$. The last panel shows the Z-Score of the GT power spectrum with respect to the generated distribution.

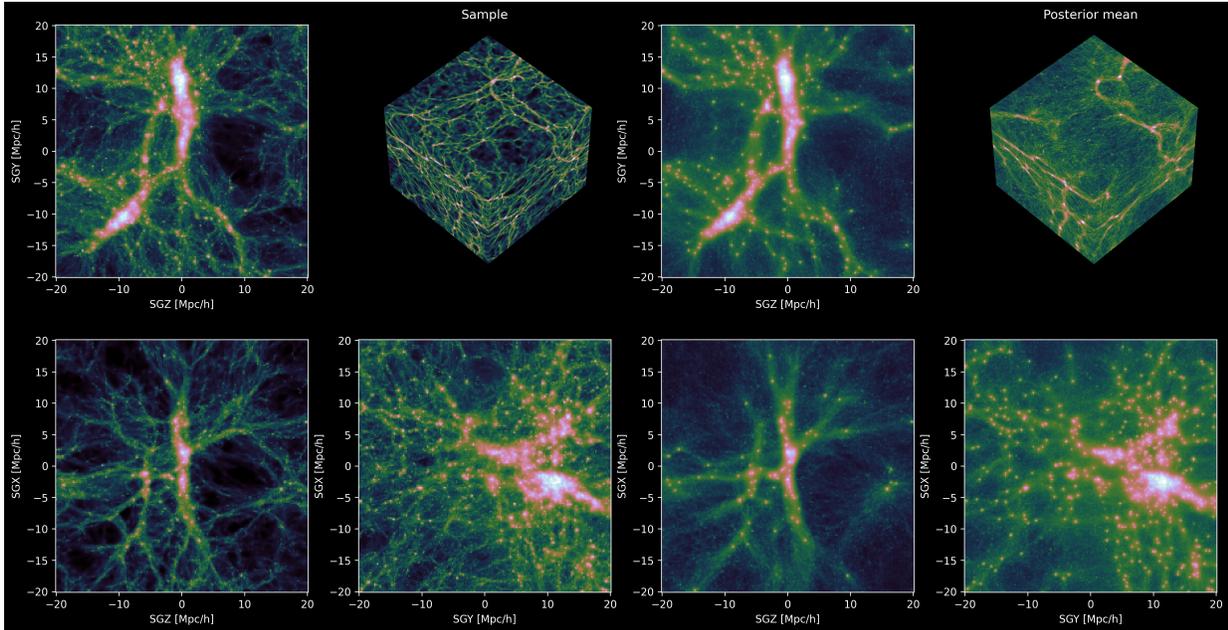


Figure 9. 3D Dark Matter Map Generated From Cosmicflows-3

Left: A single 3D sample from the diffusion model conditioned on $M_{Gal} > 10^9 M_{\odot}$ galaxy catalogs.

Right: The posterior mean 3D dark matter density field from the diffusion model conditioned on $M_{Gal} > 10^9 M_{\odot}$ galaxy catalogs.

a 320K threshold motivated by both the compute available and the generation quality, however there might be a more principled way of tackling this problem.

5.2. Future Studies

We describe some interesting future avenues:

Exploring more efficient 3D networks As our limitation above states, we expect that a more efficient 3D network will advance the progress of generative models in this domain.

Adding V_{pec} Hong et al. (2021) observed a great increase in quality when adding the peculiar velocity of galaxies in the model input. This, in principle, is also applicable to our method and would be an interesting future direction.

Towards Dark Matter Cross Correlation Studies One advantage to have a field level probabilistic representation of the dark matter density field is that one can design a cross correlation study with it. Since the diffusion model can draw from the posterior dark matter distribution, one might be able to determine which regions near the Milky Way could

be potentially good targets for dark matter search.

6. Conclusion

In this study, we explored the potential of 3D diffusion models to reconstruct the local dark matter density near the Milky Way. We identified that while the visual quality of the model is excellent at all resolutions, summary statistics reveals a degradation of quality at high resolutions. Nevertheless, the diffusion model still captures the correct modulation to the fields from variations in Ω_m and adapts extremely well to sparse galaxy catalogs. Motivated by the above, we showed a preliminary application to Cosmicflows-3 (Tully et al., 2016), a real galaxy catalog. Our results suggest that diffusion models are a powerful method to reconstruct the 3D dark matter density field from galaxies.

Acknowledgements

CFP greatly acknowledges the support of AstroAI and Cecilia Garraffo. CFP thanks the support of Aravinthan D.T. Samuel. The computations in this paper were run on the FASRC cluster supported by the FAS Division of Science Research Computing Group at Harvard University. CFP thanks Zechen Zhang for useful discussions.

References

Anderson, B. D. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.

Bruna, J. and Mallat, S. Invariant scattering convolution networks, 2012.

DESI Collaboration, Aghamousa, A., Aguilar, J., Ahlen, S., Alam, S., Allen, L. E., Allende Prieto, C., Annis, J., Bailey, S., Balland, C., Ballester, O., Baltay, C., Beaufore, L., Bebek, C., Beers, T. C., Bell, E. F., Bernal, J. L., Besuner, R., Beutler, F., Blake, C., Bleuler, H., Blomqvist, M., Blum, R., Bolton, A. S., Briceno, C., Brooks, D., Brownstein, J. R., Buckley-Geer, E., Burden, A., Burtin, E., Busca, N. G., Cahn, R. N., Cai, Y.-C., Cardiel-Sas, L., Carlberg, R. G., Carton, P.-H., Casas, R., Castander, F. J., Cervantes-Cota, J. L., Claybaugh, T. M., Close, M., Coker, C. T., Cole, S., Comparat, J., Cooper, A. P., Cousinou, M. C., Crocce, M., Cuby, J.-G., Cunningham, D. P., Davis, T. M., Dawson, K. S., de la Macorra, A., De Vicente, J., Delubac, T., Derwent, M., Dey, A., Dhungana, G., Ding, Z., Doel, P., Duan, Y. T., Ealet, A., Edelman, J., Eftekharzadeh, S., Eisenstein, D. J., Elliott, A., Escoffier, S., Evatt, M., Fagrellius, P., Fan, X., Fanning, K., Farahi, A., Farihi, J., Favole, G., Feng, Y., Fernandez, E., Findlay, J. R., Finkbeiner, D. P., Fitzpatrick, M. J., Flaughner, B., Flender, S., Font-Ribera,

A., Forero-Romero, J. E., Fosalba, P., Frenk, C. S., Fu-
magalli, M., Gaensicke, B. T., Gallo, G., Garcia-Bellido,
J., Gaztanaga, E., Pietro Gentile Fusillo, N., Gerard, T.,
Gershkovich, I., Giannantonio, T., Gillet, D., Gonzalez-
de-Rivera, G., Gonzalez-Perez, V., Gott, S., Graur, O.,
Gutierrez, G., Guy, J., Habib, S., Heetderks, H., Heet-
derks, I., Heitmann, K., Hellwing, W. A., Herrera, D. A.,
Ho, S., Holland, S., Honscheid, K., Huff, E., Hutchin-
son, T. A., Huterer, D., Hwang, H. S., Illa Laguna, J. M.,
Ishikawa, Y., Jacobs, D., Jeffrey, N., Jelinsky, P., Jen-
nings, E., Jiang, L., Jimenez, J., Johnson, J., Joyce, R.,
Jullo, E., Juneau, S., Kama, S., Karcher, A., Karkar, S.,
Kehoe, R., Kennamer, N., Kent, S., Kilbinger, M., Kim,
A. G., Kirkby, D., Kisner, T., Kitanidis, E., Kneib, J.-P.,
Koposov, S., Kovacs, E., Koyama, K., Kremin, A., Kron,
R., Kronig, L., Kueter-Young, A., Lacey, C. G., Lafever,
R., Lahav, O., Lambert, A., Lampton, M., Landriau, M.,
Lang, D., Lauer, T. R., Le Goff, J.-M., Le Guillou, L., Le
Van Suu, A., Lee, J. H., Lee, S.-J., Leitner, D., Lesser,
M., Levi, M. E., L’Huillier, B., Li, B., Liang, M., Lin, H.,
Linder, E., Loebman, S. R., Lukić, Z., Ma, J., MacCrann,
N., Magneville, C., Makarem, L., Manera, M., Manser,
C. J., Marshall, R., Martini, P., Massey, R., Matheson,
T., McCauley, J., McDonald, P., McGreer, I. D., Meisner,
A., Metcalfe, N., Miller, T. N., Miquel, R., Moustakas,
J., Myers, A., Naik, M., Newman, J. A., Nichol, R. C.,
Nicola, A., Nicolati da Costa, L., Nie, J., Niz, G., Nor-
berg, P., Nord, B., Norman, D., Nugent, P., O’Brien, T.,
Oh, M., Olsen, K. A. G., Padilla, C., Padmanabhan, H.,
Padmanabhan, N., Palanque-Delabrouille, N., Palmese,
A., Pappalardo, D., Pâris, I., Park, C., Patej, A., Peacock,
J. A., Peiris, H. V., Peng, X., Percival, W. J., Perruchot,
S., Pieri, M. M., Pogge, R., Pollack, J. E., Poppett, C.,
Prada, F., Prakash, A., Probst, R. G., Rabinowitz, D., Rai-
choor, A., Ree, C. H., Refregier, A., Regal, X., Reid, B.,
Reil, K., Rezaie, M., Rockosi, C. M., Roe, N., Ronayette,
S., Roodman, A., Ross, A. J., Ross, N. P., Rossi, G.,
Roza, E., Ruhlmann-Kleider, V., Rykoff, E. S., Sabiu, C.,
Samushia, L., Sanchez, E., Sanchez, J., Schlegel, D. J.,
Schneider, M., Schubnell, M., Secroun, A., Seljak, U.,
Seo, H.-J., Serrano, S., Shafieloo, A., Shan, H., Sharples,
R., Sholl, M. J., Shourt, W. V., Silber, J. H., Silva, D. R.,
Sirk, M. M., Slosar, A., Smith, A., Smoot, G. F., Som,
D., Song, Y.-S., Sprayberry, D., Staten, R., Stefanik, A.,
Tarle, G., Sien Tie, S., Tinker, J. L., Tojeiro, R., Valdes, F.,
Valenzuela, O., Valluri, M., Vargas-Magana, M., Verde,
L., Walker, A. R., Wang, J., Wang, Y., Weaver, B. A.,
Weaverdyck, C., Wechsler, R. H., Weinberg, D. H., White,
M., Yang, Q., Yeche, C., Zhang, T., Zhao, G.-B., Zheng,
Y., Zhou, X., Zhou, Z., Zhu, Y., Zou, H., and Zu, Y. The
DESI Experiment Part I: Science, Targeting, and Survey
Design. *arXiv e-prints*, art. arXiv:1611.00036, October
2016. doi: 10.48550/arXiv.1611.00036.

- Falcon, W. and The PyTorch Lightning team. PyTorch Lightning, March 2019. URL <https://github.com/Lightning-AI/lightning>.
- Garrison, L. H., Eisenstein, D. J., and Pinto, P. A. A high-fidelity realization of the euclid code comparison n-body simulation with abacus. *Monthly Notices of the Royal Astronomical Society*, 485(3):3370–3377, 2019.
- Hendrycks, D. and Gimpel, K. Gaussian error linear units (gelus), 2023.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Hong, S. E., Jeong, D., Hwang, H. S., and Kim, J. Revealing the local cosmic web from galaxies by deep learning. *The Astrophysical Journal*, 913(1):76, 2021.
- Jasche, J. and Lavaux, G. Physical bayesian modelling of the non-linear matter distribution: new insights into the nearby universe. *Astronomy & Astrophysics*, 625:A64, 2019.
- Kingma, D., Salimans, T., Poole, B., and Ho, J. Variational diffusion models. *Advances in neural information processing systems*, 34:21696–21707, 2021.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Laureijs, R., Amiaux, J., Arduini, S., Auguères, J. L., Brinchmann, J., Cole, R., Cropper, M., Dabin, C., Duvet, L., Ealet, A., Garilli, B., Gondoin, P., Guzzo, L., Hoar, J., Hoekstra, H., Holmes, R., Kitching, T., Maciaszek, T., Mellier, Y., Pasian, F., Percival, W., Rhodes, J., Saavedra Criado, G., Sauvage, M., Scaramella, R., Valenziano, L., Warren, S., Bender, R., Castander, F., Cimatti, A., Le Fèvre, O., Kurki-Suonio, H., Levi, M., Lilje, P., Meylan, G., Nichol, R., Pedersen, K., Popa, V., Rebolo Lopez, R., Rix, H. W., Rottgering, H., Zeilinger, W., Grupp, F., Hudelot, P., Massey, R., Meneghetti, M., Miller, L., Paltani, S., Paulin-Henriksson, S., Pires, S., Saxton, C., Schrabback, T., Seidel, G., Walsh, J., Aghanim, N., Amendola, L., Bartlett, J., Baccigalupi, C., Beaulieu, J. P., Benabed, K., Cuby, J. G., Elbaz, D., Fosalba, P., Gavazzi, G., Helmi, A., Hook, I., Irwin, M., Kneib, J. P., Kunz, M., Mannucci, F., Moscardini, L., Tao, C., Teyssier, R., Weller, J., Zamorani, G., Zapatero Osorio, M. R., Boulade, O., Foumond, J. J., Di Giorgio, A., Guttridge, P., James, A., Kemp, M., Martignac, J., Spencer, A., Walton, D., Blümchen, T., Bonoli, C., Bortoletto, F., Cerna, C., Corcione, L., Fabron, C., Jahnke, K., Ligorì, S., Madrid, F., Martin, L., Morgante, G., Pampalona, T., Prieto, E., Riva, M., Toledo, R., Trifoglio, M., Zerbi, F., Abdalla, F., Douspis, M., Grenet, C., Borgani, S., Bouwens, R., Courbin, F., Delouis, J. M., Dubath, P., Fontana, A., Frailis, M., Grazian, A., Koppenhöfer, J., Mansutti, O., Melchior, M., Mignoli, M., Mohr, J., Neissner, C., Noddle, K., Poncet, M., Scodreggio, M., Serano, S., Shane, N., Starck, J. L., Surace, C., Taylor, A., Verdoes-Kleijn, G., Vuerli, C., Williams, O. R., Zacchei, A., Altieri, B., Escudero Sanz, I., Kohley, R., Oosterbroek, T., Astier, P., Bacon, D., Bardelli, S., Baugh, C., Bellagamba, F., Benoist, C., Bianchi, D., Biviano, A., Branchini, E., Carbone, C., Cardone, V., Clements, D., Colombi, S., Conselice, C., Cresci, G., Deacon, N., Dunlop, J., Fedeli, C., Fontanot, F., Franzetti, P., Giocoli, C., Garcia-Bellido, J., Gow, J., Heavens, A., Hewett, P., Heymans, C., Holland, A., Huang, Z., Ilbert, O., Joachimi, B., Jennins, E., Kerins, E., Kiessling, A., Kirk, D., Kotak, R., Krause, O., Lahav, O., van Leeuwen, F., Lesgourgues, J., Lombardi, M., Magliocchetti, M., Maguire, K., Majerotto, E., Maoli, R., Marulli, F., Maurogordato, S., McCracken, H., McLure, R., Melchiorri, A., Merson, A., Moresco, M., Nonino, M., Norberg, P., Peacock, J., Pello, R., Penny, M., Pettorino, V., Di Porto, C., Pozzetti, L., Quercellini, C., Radovich, M., Rassat, A., Roche, N., Ronayette, S., Rossetti, E., Sartoris, B., Schneider, P., Semboloni, E., Serjeant, S., Simpson, F., Skordis, C., Smadja, G., Smartt, S., Spano, P., Spiro, S., Sullivan, M., Tilquin, A., Trotta, R., Verde, L., Wang, Y., Williger, G., Zhao, G., Zoubian, J., and Zucca, E. Euclid Definition Study Report. *arXiv e-prints*, art. arXiv:1110.3193, October 2011. doi: 10.48550/arXiv.1110.3193.
- Legin, R., Ho, M., Lemos, P., Perreault-Levasseur, L., Ho, S., Hezaveh, Y., and Wandelt, B. Posterior sampling of the initial conditions of the universe from non-linear large scale structures using score-based generative models. *Monthly Notices of the Royal Astronomical Society: Letters*, 527(1):L173–L178, 2024.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization, 2019.
- LSST Science Collaboration, Abell, P. A., Allison, J., Anderson, S. F., Andrew, J. R., Angel, J. R. P., Armus, L., Arnett, D., Asztalos, S. J., Axelrod, T. S., Bailey, S., Ballantyne, D. R., Bankert, J. R., Barkhouse, W. A., Barr, J. D., Barrientos, L. F., Barth, A. J., Bartlett, J. G., Becker, A. C., Becla, J., Beers, T. C., Bernstein, J. P., Biswas, R., Blanton, M. R., Bloom, J. S., Bochanski, J. J., Boeshaar, P., Borne, K. D., Bradac, M., Brandt, W. N., Bridge, C. R., Brown, M. E., Brunner, R. J., Bullock, J. S., Burgasser, A. J., Burge, J. H., Burke, D. L., Cargile, P. A., Chandrasekharan, S., Chartas, G., Chesley, S. R., Chu, Y.-H., Cinabro, D., Claire, M. W., Claver, C. F., Clowe, D., Connolly, A. J., Cook, K. H., Cooke, J., Cooray, A., Covey, K. R., Culliton, C. S., de Jong, R., de Vries, W. H., Debatista, V. P., Delgado, F., Dell’Antonio, I. P., Dhital, S., Di

- Stefano, R., Dickinson, M., Dilday, B., Djorgovski, S. G., Dobler, G., Donalek, C., Dubois-Felsmann, G., Durech, J., Eliasdottir, A., Eracleous, M., Eyer, L., Falco, E. E., Fan, X., Fassnacht, C. D., Ferguson, H. C., Fernandez, Y. R., Fields, B. D., Finkbeiner, D., Figueroa, E. E., Fox, D. B., Francke, H., Frank, J. S., Frieman, J., Fromenteau, S., Furqan, M., Galaz, G., Gal-Yam, A., Garnavich, P., Gawiser, E., Geary, J., Gee, P., Gibson, R. R., Gilmore, K., Grace, E. A., Green, R. F., Gressler, W. J., Grillmair, C. J., Habib, S., Haggerty, J. S., Hamuy, M., Harris, A. W., Hawley, S. L., Heavens, A. F., Hebb, L., Henry, T. J., Hileman, E., Hilton, E. J., Hoadley, K., Holberg, J. B., Holman, M. J., Howell, S. B., Infante, L., Ivezić, Z., Jacoby, S. H., Jain, B., Jedicke, Jee, M. J., Garrett Jernigan, J., Jha, S. W., Johnston, K. V., Jones, R. L., Juric, M., Kaasalainen, M., Styliani, Kafka, Kahn, S. M., Kaib, N. A., Kalirai, J., Kantor, J., Kasliwal, M. M., Keeton, C. R., Kessler, R., Knezevic, Z., Kowalski, A., Krabben-dam, V. L., Krughoff, K. S., Kulkarni, S., Kuhlman, S., Lacy, M., Lepine, S., Liang, M., Lien, A., Lira, P., Long, K. S., Lorenz, S., Lotz, J. M., Lupton, R. H., Lutz, J., Macri, L. M., Mahabal, A. A., Mandelbaum, R., Marshall, P., May, M., McGehee, P. M., Meadows, B. T., Meert, A., Milani, A., Miller, C. J., Miller, M., Mills, D., Minniti, D., Monet, D., Mukadam, A. S., Nakar, E., Neill, D. R., Newman, J. A., Nikolaev, S., Nordby, M., O'Connor, P., Oguri, M., Oliver, J., Olivier, S. S., Olsen, J. K., Olsen, K., Olszewski, E. W., Oluseyi, H., Padilla, N. D., Parker, A., Pepper, J., Peterson, J. R., Petry, C., Pinto, P. A., Pizagno, J. L., Popescu, B., Prsa, A., Radcka, V., Raddick, M. J., Rasmussen, A., Rau, A., Rho, J., Rhoads, J. E., Richards, G. T., Ridgway, S. T., Robertson, B. E., Roskar, R., Saha, A., Sarajedini, A., Scannapieco, E., Schalk, T., Schindler, R., Schmidt, S., Schmidt, S., Schneider, D. P., Schumacher, G., Scranton, R., Sebag, J., Seppala, L. G., Shemmer, O., Simon, J. D., Sivertz, M., Smith, H. A., Allyn Smith, J., Smith, N., Spitz, A. H., Stanford, A., Stassun, K. G., Strader, J., Strauss, M. A., Stubbs, C. W., Sweeney, D. W., Szalay, A., Szkody, P., Takada, M., Thorman, P., Trilling, D. E., Trimble, V., Tyson, A., Van Berg, R., Vanden Berk, D., VanderPlas, J., Verde, L., Vrsnak, B., Walkowicz, L. M., Wandelt, B. D., Wang, S., Wang, Y., Warner, M., Wechsler, R. H., West, A. A., Wiecha, O., Williams, B. F., Willman, B., Wittman, D., Wolff, S. C., Wood-Vasey, W. M., Wozniak, P., Young, P., Zentner, A., and Zhan, H. LSST Science Book, Version 2.0. *arXiv e-prints*, art. arXiv:0912.0201, December 2009. doi: 10.48550/arXiv.0912.0201.
- Nelson, D., Pillepich, A., Genel, S., Vogelsberger, M., Springel, V., Torrey, P., Rodriguez-Gomez, V., Sijacki, D., Snyder, G., Griffen, B., Marinacci, F., Blecha, L., Sales, L., Xu, D., and Hernquist, L. The illustris simulation: Public data release. *Astronomy and Computing*, 13:12–37, November 2015. ISSN 2213-1337. doi: 10.1016/j.ascom.2015.09.003. URL <http://dx.doi.org/10.1016/j.ascom.2015.09.003>.
- Ni, Y., Di Matteo, T., Bird, S., Croft, R., Feng, Y., Chen, N., Tremmel, M., DeGraf, C., and Li, Y. The astrid simulation: the evolution of supermassive black holes. *Monthly Notices of the Royal Astronomical Society*, 513(1):670–692, February 2022. ISSN 1365-2966. doi: 10.1093/mnras/stac351. URL <http://dx.doi.org/10.1093/mnras/stac351>.
- Ono, V., Park, C. F., Mudur, N., Ni, Y., Cuesta-Lazaro, C., and Villaescusa-Navarro, F. Debiasing with diffusion: Probabilistic reconstruction of dark matter fields from galaxies with camels, 2024.
- Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., and Lakshminarayanan, B. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021.
- Park, C. F., Allys, E., Villaescusa-Navarro, F., and Finkbeiner, D. Quantification of high-dimensional non-gaussianities and its implication to fisher analysis in cosmology. *The Astrophysical Journal*, 946(2):107, 2023a.
- Park, C. F., Ono, V., Mudur, N., Ni, Y., and Cuesta-Lazaro, C. Probabilistic reconstruction of dark matter fields from biased tracers using diffusion models, 2023b.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library, 2019.
- Pfeifer, S., Valade, A., Gottlöber, S., Hoffman, Y., Libeskind, N. I., and Hellwing, W. A. A local universe model for constrained simulations. *Monthly Notices of the Royal Astronomical Society*, 523(4):5985–5994, June 2023. ISSN 1365-2966. doi: 10.1093/mnras/stad1851. URL <http://dx.doi.org/10.1093/mnras/stad1851>.
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., and Sutskever, I. Zero-shot text-to-image generation, 2021.
- Regaldo-Saint Blancard, B., Levrier, F., Allys, E., Bellomi, E., and Boulanger, F. Statistical description of dust polarized emission from the diffuse interstellar medium: A rwst approach. *Astronomy & Astrophysics*, 642:A217, October 2020. ISSN 1432-0746. doi: 10.1051/0004-6361/202038044. URL <http://dx.doi.org/10.1051/0004-6361/202038044>.

- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. High-resolution image synthesis with latent diffusion models, 2022.
- Ronneberger, O., Fischer, P., and Brox, T. U-net: Convolutional networks for biomedical image segmentation, 2015.
- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S. K. S., Ayan, B. K., Mahdavi, S. S., Lopes, R. G., Salimans, T., Ho, J., Fleet, D. J., and Norouzi, M. Photorealistic text-to-image diffusion models with deep language understanding, 2022.
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pp. 2256–2265. PMLR, 2015.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=PxTIG12RRHS>.
- Spergel, D., Gehrels, N., Baltay, C., Bennett, D., Breckinridge, J., Donahue, M., Dressler, A., Gaudi, B. S., Greene, T., Guyon, O., Hirata, C., Kalirai, J., Kasdin, N. J., Macintosh, B., Moos, W., Perlmutter, S., Postman, M., Rauscher, B., Rhodes, J., Wang, Y., Weinberg, D., Benford, D., Hudson, M., Jeong, W. S., Mellier, Y., Traub, W., Yamada, T., Capak, P., Colbert, J., Masters, D., Penny, M., Savransky, D., Stern, D., Zimmerman, N., Barry, R., Bartusek, L., Carpenter, K., Cheng, E., Content, D., Dekens, F., Demers, R., Grady, K., Jackson, C., Kuan, G., Kruk, J., Melton, M., Nemati, B., Parvin, B., Poberezhskiy, I., Peddie, C., Ruffa, J., Wallace, J. K., Whipple, A., Wollack, E., and Zhao, F. Wide-Field Infrared Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report. *arXiv e-prints*, art. arXiv:1503.03757, March 2015. doi: 10.48550/arXiv.1503.03757.
- Springel, V., White, S. D., Jenkins, A., Frenk, C. S., Yoshida, N., Gao, L., Navarro, J., Thacker, R., Croton, D., Helly, J., et al. Simulations of the formation, evolution and clustering of galaxies and quasars. *nature*, 435(7042): 629–636, 2005.
- Tully, R. B., Courtois, H. M., and Sorce, J. G. Cosmicflows-3. *The Astronomical Journal*, 152(2):50, August 2016. ISSN 1538-3881. doi: 10.3847/0004-6256/152/2/50. URL <http://dx.doi.org/10.3847/0004-6256/152/2/50>.
- Villaescusa-Navarro, F., Hahn, C., Massara, E., Banerjee, A., Delgado, A. M., Ramanah, D. K., Charnock, T., Giusarma, E., Li, Y., Allys, E., et al. The quijote simulations. *The Astrophysical Journal Supplement Series*, 250 (1):2, 2020.
- Villaescusa-Navarro, F., Anglés-Alcázar, D., Genel, S., Spergel, D. N., Somerville, R. S., Dave, R., Pillepich, A., Hernquist, L., Nelson, D., Torrey, P., Narayanan, D., Li, Y., Philcox, O., La Torre, V., Maria Delgado, A., Ho, S., Hassan, S., Burkhardt, B., Wadekar, D., Battaglia, N., Contardo, G., and Bryan, G. L. The CAMELS Project: Cosmology and Astrophysics with Machine-learning Simulations. *Apj*, 915(1):71, July 2021a. doi: 10.3847/1538-4357/abf7ba.
- Villaescusa-Navarro, F., Genel, S., Angles-Alcazar, D., Thiele, L., Dave, R., Narayanan, D., Nicola, A., Li, Y., Villanueva-Domingo, P., Wandelt, B., Spergel, D. N., Somerville, R. S., Zorrilla Matilla, J. M., Mohammad, F. G., Hassan, S., Shao, H., Wadekar, D., Eickenberg, M., Wong, K. W. K., Contardo, G., Jo, Y., Moser, E., Lau, E. T., Machado Poletti Valle, L. F., Perez, L. A., Nagai, D., Battaglia, N., and Vogelsberger, M. The CAMELS Multifield Dataset: Learning the Universe’s Fundamental Parameters with Artificial Intelligence. *arXiv e-prints*, art. arXiv:2109.10915, September 2021b.
- Wu, Y. and He, K. Group normalization, 2018.

A. Data&Model details

A.1. Data

CAMELS We use the Camels Multifield Dataset (Villaescusa-Navarro et al., 2021b) from the CAMELS Simulation Suite (Villaescusa-Navarro et al., 2021a). In specific we use 3D density fields generated with PCS(Piecewise Cubic Spline) mass assignment from the Astrid suite(Ni et al., 2022) ran with the MP-Gadget code. These simulations follow the dynamics of 256^3 dark matter particles and 256^3 gas resolution elements in a comoving box with side $25 \text{ Mpc } h^{-1}$. We use the 256^3 density field generated at the current time, thus at a redshift of $z = 0$.

Splits We use the Latin Hypercube (LH) set as our training data, which consists of 1000 simulations with random initial conditions and random astrophysical/cosmological parameters. The range of the 6 parameters are: $0.1 \leq \Omega_m \leq 0.5$, $0.6 \leq \sigma_8 \leq 1.0$, $0.25 \leq (ASN1, AAGN1) \leq 4.00$ and $0.5 \leq (ASN2, AAGN2) \leq 2.0$. In Sec. 4.1, we use the Cosmic Variance (CV) set for evaluation, which consists of 27 simulations with the same astrophysical/cosmological parameters but different initial conditions. In Sec. 4.2, we use the One Parameter (1P) set, which consists of 61 simulations where the initial condition is fixed, and only one astrophysical parameter is varied in each simulation.

Galaxy Overdensity Field We use the Friends-of-Friends Subfind catalog from CAMELS to generate a galaxy overdensity field. We select all subhalos above a mass threshold of $10^{[7,8,9]} M_\odot$ for each of the fields in Sec. 4.3. We bin the galaxies into a grid of 256^3 then calculate the overdensity field as $\delta = \frac{f}{\langle f \rangle} - 1$ where f is the binned galaxy count and $\langle f \rangle$ is the average galaxy count over the field. We then downsample (see below) them to a 128^3 grid.

Normalization We normalize the M_{CDM} and M_{Star} fields by taking the logarithm of base 10 and standard scale the fields such that the whole LH set has a mean of zero and a standard deviation of unity. We add unity to the fields before the logarithm to avoid $-\infty$ values for pixels where the dark matter density or the stellar mass is null. For galaxy overdensity fields, we divide each overdensity field by its mean and add 2 to the overdensity field since the overdensity field can range down to -1 .

Augmentation We use flip, axis permutation (rotation + mirror) and periodic boundary translation to augment our training dataset. We did not use any augmentation to change the pixel values since we are interested in associating the pixel values to a real quantity.

Downsampling The original density fields are given in 256^3 , we use PyTorch’s(Paszke et al., 2019) trilinear interpolation function to down sample these grids to 128^3 and 160^3 .

Cosmicflows-3 Cosmicflows-3 (Tully et al., 2016) contains the sky coordinates and distances of 17,669 galaxies up to a maximum distance of 200 Mpc. In this work, we follow Hong et al. (2021) and use the cuts of $M_B < -15$ and $|b| > 10^\circ$ to select a volume limited catalog. The selected 6,159 galaxies are binned on a 512^3 grid and we calculate the overdensity field on this grid. Since our models are trained using periodic boundary conditions, we take the 256^3 central region of the 512^3 grid to pass into the network and again only take the 192^3 central region to produce the results in Fig. 9.

A.2. Model

Variational Diffusion Model We use variational diffusion models introduced by Kingma et al. (2021). We use an initial noise schedule of $\gamma(t) = 26.6 t - 13.3$, which corresponds to a initial SNR of $e^{13.3}$ and a terminal SNR of $e^{-13.3}$. We optimize the noise schedule jointly with the diffusion process by parametrizing gamma to be a learned linear function of t :

$$\gamma(t) = w * t + b \quad (1)$$

we initialize w to 26.6 and b to -13.3 . For further details, please refer to Kingma et al. (2021).

Network Architecture We use a conditional variant of the U-Net(Ronneberger et al., 2015). We use a network with [32,64,128,256] channels in each resolution block while we half the number of channels to [16,32,64,128] for the 256^3 network due to memory constraints. We use a GELU(Hendrycks & Gimpel, 2023) activationm Group Normalization(Wu & He, 2018) and residual connections in each block. The diffusion time is embedded with a 64 dimensional sinusoidal embedding which is then transformed into 256 dimensions via a MLP layer. The astrophysical parameters are also embedded into 64 dimensions via a 2 layer MLP with GELU activations. The spatial conditioning is simply concatenated to the U-Net input.

Training We train all models for $320K$ gradient steps with a batch size of 2. The batch size was limited by the GPU memory.

We use the AdamW(Loshchilov & Hutter, 2019) optimizer with a learning rate 3×10^{-4} and a weight decay of 1×10^{-5} . Although the loss saturates very quickly at $\sim 50K$ gradient steps, the generation quality keeps increasing.

Generation Since our model uses a continuous time formulation, the number of generation steps is simply a inference time hyperparameter. We choose $T = 250$ steps based on previous results in (Park et al., 2023b).

B. Additional generation results.

We show additional generation results without any cherry picking.

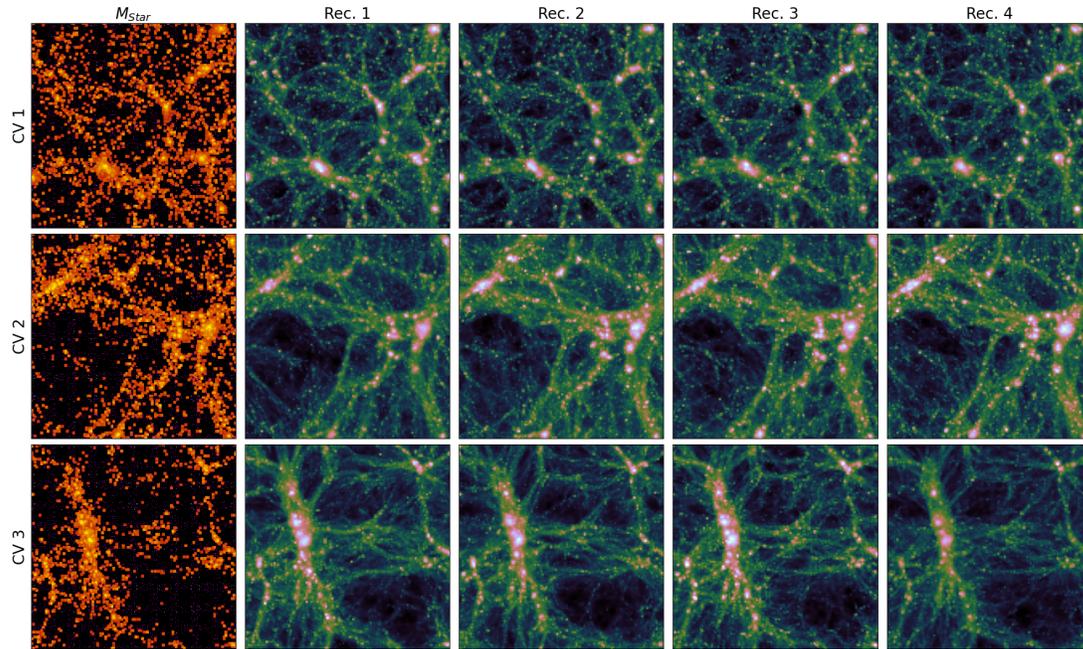


Figure 10. Multiple samples from the 128^3 variational diffusion model conditioned on the Stellar Mass Field. Each row is a different selection of cosmic variance, e.g. initial condition. The first column is the input stellar mass map and the 4 following columns are random samples from the diffusion model.

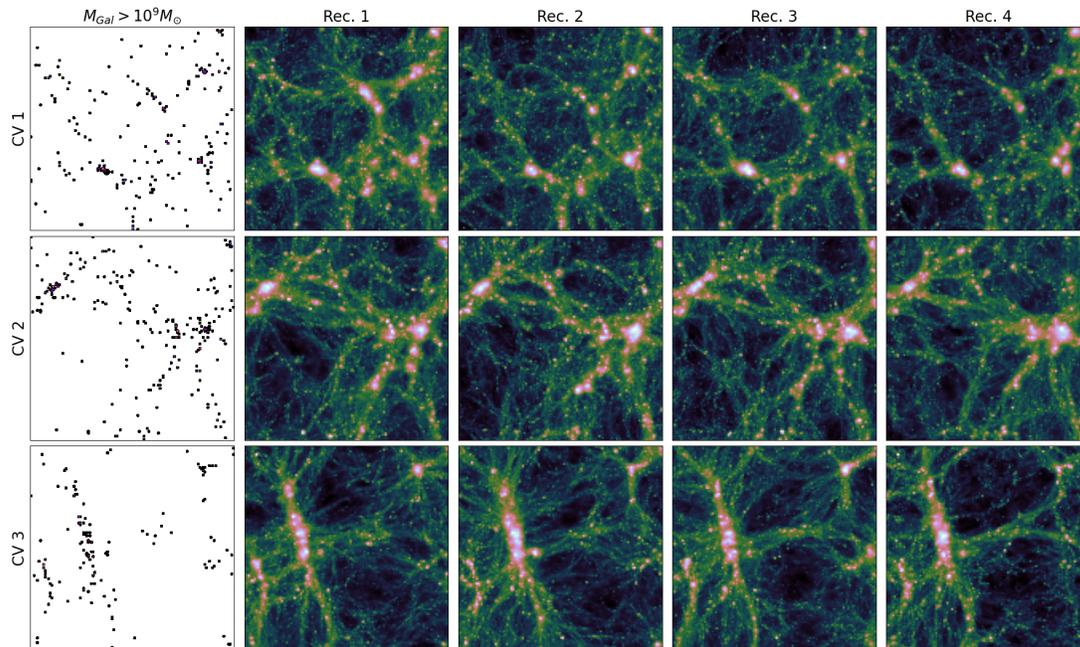


Figure 11. **Multiple samples from the 128^3 variational diffusion model conditioned on galaxy catalogs with $M_{gal} > 10^9 M_{\odot}$.** Each row is a different selection of cosmic variance, e.g. initial condition. The first column is the input galaxy overdensity map and the 4 following columns are random samples from the diffusion model. By carefully looking into void (low density) regions, one can notice that the variance is higher than in Fig. 10

C. Additional Analysis

Here, we show additional analysis not included in the main text.

C.1. Resolution Scaling

In Fig. 12, we show the 3D power spectrum of the generated fields and the ground truth field. We find similar conclusions as in the 2D case, just stronger in magnitude.

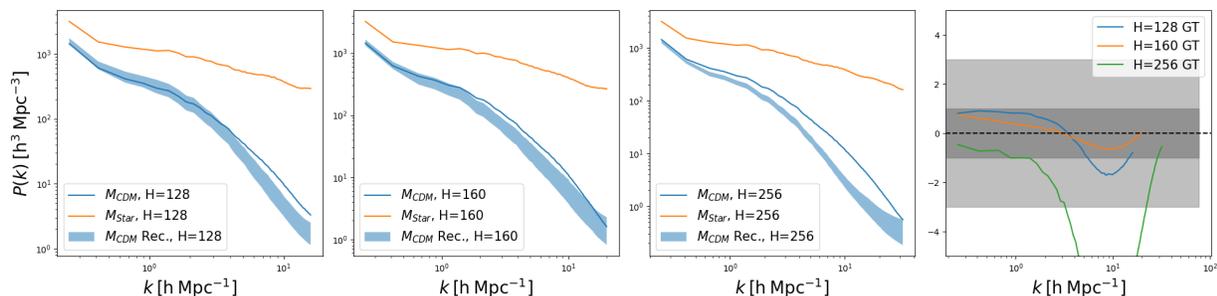


Figure 12. **The 3D Power Spectrum of the Generated fields.** The power spectra calculated from the whole 3D field. The first three panels respectively depicts the results at resolutions $H=128$, 160 and 256 . The last panel depicts the Z-Score of the ground truth P_k with respect to the generated distribution.

In Fig. 13, we show the 3D probability density field from the whole field. We find very good agreement on the PDF at 256^3

resolutions. Our conclusions remain consistent with Sec. 4.1.

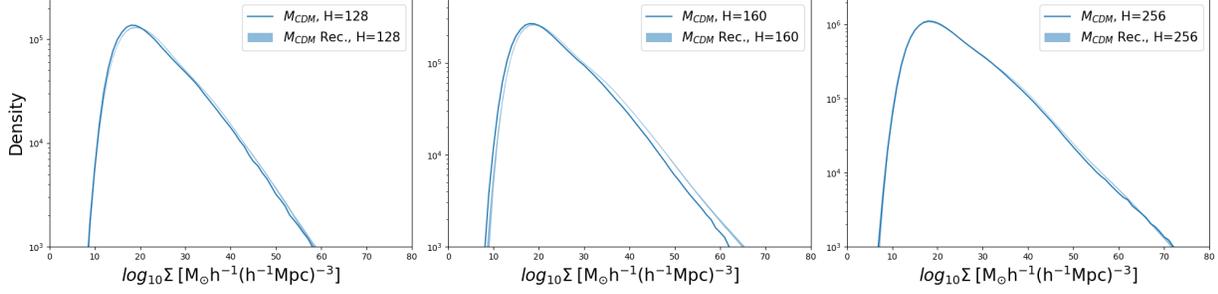


Figure 13. The 3D probability Density Function of the generated fields. The PDF is calculated on the whole 3D field. The three panels respectively depicts the results at resolutions $H=128$, 160 and 256 .

C.2. Astrophysical Conditioning

In Fig. 14, we show the ratio of the 2D projected power spectrum with $\Omega_m = 0.22$ to the one with $\Omega_m = 0.3$ (fiducial) as well as the ratio of the $\Omega_m = 0.38$ one to the $\Omega_m = 0.3$ one. We find the same conclusion as in Sec. 4.2.

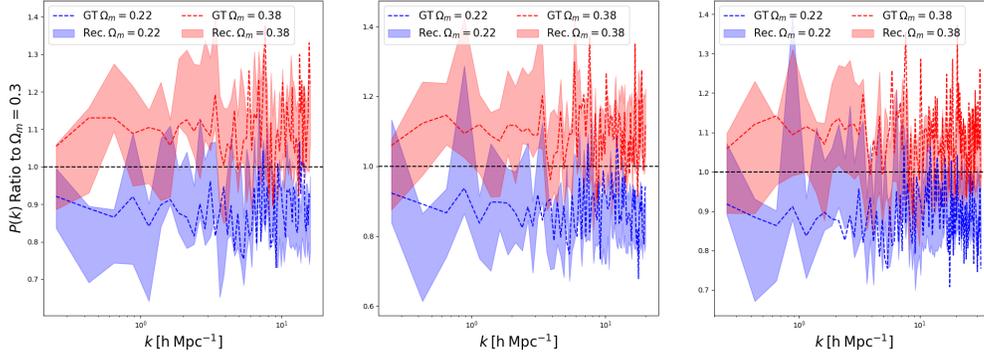


Figure 14. Ratio of the 2D Projected Power Spectrum at a modified value of Ω_m relative to the fiducial value. We depict in blue the generated power spectra range and the GT power spectrum when $\Omega_m = 0.22$ and in red when $\Omega_m = 0.38$. The panels from left to right corresponds to resolution $H=128$, 160 and 256 .

C.3. Sparsity Suite

The 2D PDF of the 3 diffusion model generations from the sparsity experiments are in Fig. 15.

The RWST statistics of the 3 diffusion model generations from the sparsity experiments are in Fig. 16. We find similar results to Fig. 4, the generations agree excellently to the GT field's RWST.

The 3D power spectrum of the 3 diffusion model generations from the sparsity experiments are in Fig. 17.

The 3D PDF of the 3 diffusion model generations from the sparsity experiments are in Fig. 18.

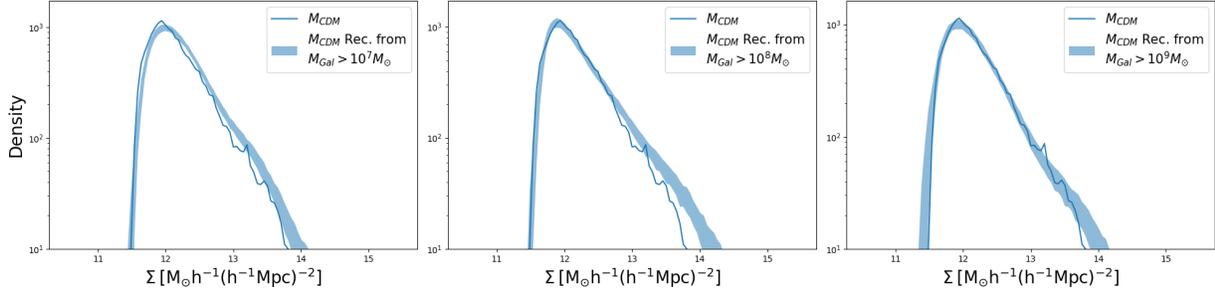


Figure 15. Probability Density Function of 2D Projections from the 3D Generation from models conditioned on galaxy catalogs The probability density function is calculated from $12.5 \text{ Mpc } h^{-1}$ 2D projections of 128 generated 3D dark matter density fields. The three panels depict the power spectrum of the 2D projected fields when the model is conditioned on galaxies of $M_{Gal} > 10^7 M_{\odot}$, $M_{Gal} > 10^8 M_{\odot}$ and $M_{Gal} > 10^9 M_{\odot}$.

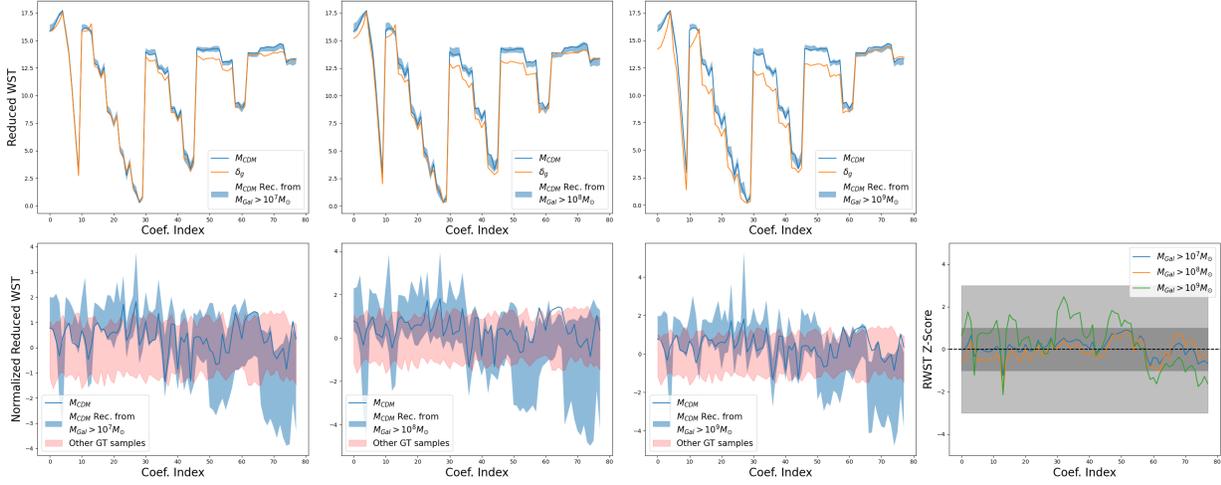


Figure 16. The RWST statistics of 2D Projections from the 3D Generation from models conditioned on galaxy catalogs The RWST is calculated from 12.5 h Mpc^{-1} 2D projections of 128 generated 3D dark matter density fields. The first three panels in each row depict the power spectrum of the 2D projected fields when the model is conditioned on galaxies of $M_{Gal} > 10^7 M_{\odot}$, $M_{Gal} > 10^8 M_{\odot}$ and $M_{Gal} > 10^9 M_{\odot}$.

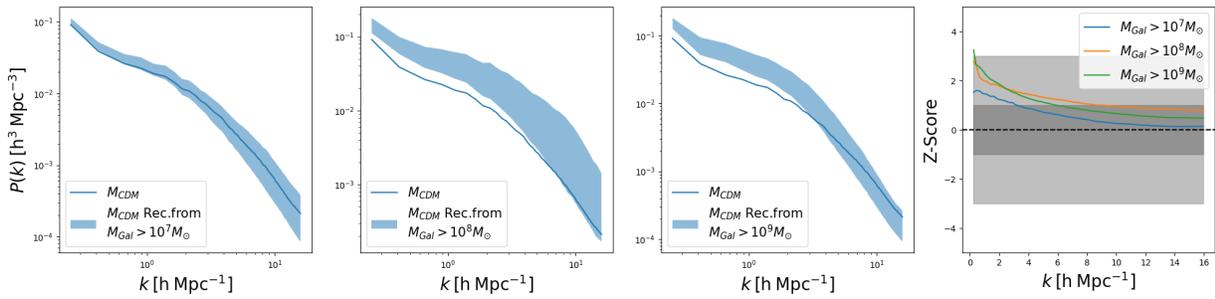


Figure 17. Power Spectrum of 3D Generations from models conditioned on galaxy catalogs The first three panels depict the power spectrum of the 2D projected fields when the model is conditioned on galaxies of $M_{Gal} > 10^7 M_{\odot}$, $M_{Gal} > 10^8 M_{\odot}$ and $M_{Gal} > 10^9 M_{\odot}$.

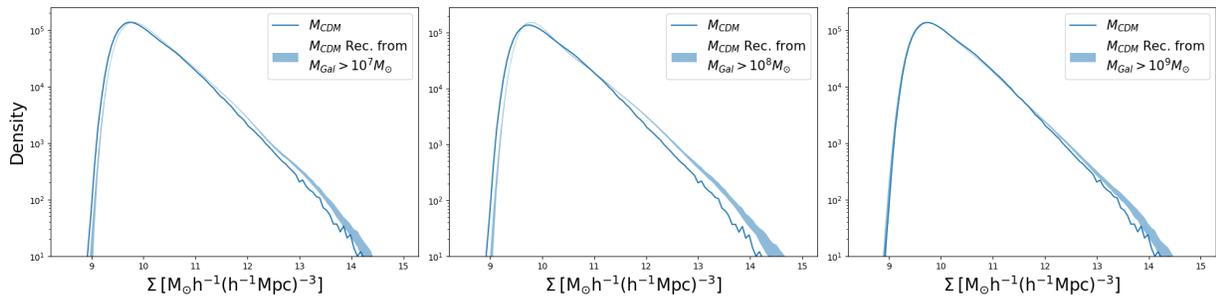


Figure 18. Probability Density Function of 3D generations from models conditioned on galaxy catalogs The three panels depict the power spectrum of the 2D projected fields when the model is conditioned on galaxies of $M_{Gal} > 10^7 M_{\odot}$, $M_{Gal} > 10^8 M_{\odot}$ and $M_{Gal} > 10^9 M_{\odot}$.

D. Computational Details

We implement our model in PyTorch([Paszke et al., 2019](#)) and PyTorchLightning([Falcon & The PyTorch Lightning team, 2019](#)). We train our models on a NVIDIA A100 80GB GPU.