# A Fast Generative Framework for High-dimensional Posterior Sampling: Application to CMB Delensing

Mohammad-Hadi Sotoudeh<sup>123</sup> Pablo Lemos<sup>123</sup> Laurence Perreault-Levasseur<sup>1234</sup>

#### Abstract

We introduce a deep generative framework for high-dimensional Bayesian inference that enables efficient posterior sampling. As telescopes and simulations rapidly expand the volume and resolution of astrophysical data, fast simulation-based inference methods are increasingly needed to extract scientific insights. While diffusion-based approaches offer high-quality generative capabilities, they are hindered by slow sampling speeds. Our method performs posterior sampling an order of magnitude faster than a diffusion baseline. Applied to the problem of CMB delensing, it successfully recovers the unlensed CMB power spectrum from simulated observations. The model also remains robust to shifts in cosmological parameters, demonstrating its potential for out-of-distribution generalization and application to observational cosmological data.

# 1. Introduction

The volume and complexity of astrophysical data are rising significantly due to modern instruments and simulations (Bellm et al., 2019; Ntampaka et al., 2021). Space missions and ground-based telescopes such as the James Webb Space Telescope (Gardner et al., 2006), Euclid (Racca et al., 2016), Square Kilometre Array (Dewdney et al., 2009), LSST (Ivezić et al., 2019), Simons Observatory (Galitzki et al., 2024), and CMB-S4 (Abazajian et al., 2016) will not only bring about exponential growth to the size of the observed data but will also increase the resolution and data acquisition rate to unprecedented levels.

ML4Astro 2025, Vancouver, CA. Copyright 2025 by the author(s).

Bayesian inference provides a principled framework for estimating physical parameters from observations. However, the scale and dimensionality of modern datasets pose substantial challenges for modeling the posterior distribution. For likelihood-based models, inference can become intractable because the likelihood function often breaks down as its simplifying assumptions fail at higher data fidelity. Even moving to likelihood-free (or implicit-likelihood) approaches, which bypass the need for an explicit likelihood by relying on forward simulations, fails to fully address key bottlenecks. Specifically, the computational cost of posterior sampling often grows prohibitively with data complexity.

In recent years, machine learning (ML) methods have become increasingly popular for addressing challenges in Bayesian inference. These techniques have driven significant advancements in simulation-based inference (SBI) approaches (Cranmer et al., 2020; Durkan et al., 2020). However, several challenges remain in realizing the full potential of these methods for scientific discovery. Key obstacles include obtaining reliable uncertainty estimates, generalizing to real data, developing training schemes that work across problems with minimal fine-tuning, and enabling fast inference.

This paper introduces a deep generative framework for fast, high-dimensional Bayesian inference. Our approach builds on the Hierarchical Probabilistic U-Net (HPU-Net) architecture (Kohl et al., 2019), which enables learning structured probabilistic representations using the variational autoencoder (VAE) inference framework (Kingma & Welling, 2022). The VAE-based design allows for fast sampling. We demonstrate the framework's practical utility on the task of CMB delensing — removing the effects of weak gravitational lensing from the observed cosmic microwave background — and show that the model rapidly and reliably recovers the primordial<sup>1</sup> CMB power spectrum with reasonable uncertainty estimates.

<sup>&</sup>lt;sup>1</sup>Department of Physics, Université de Montréal, Montréal, Canada <sup>2</sup>Mila - Quebec AI Institute, Montréal, Canada <sup>3</sup>Ciela, Montreal Institute for Astrophysics and Machine Learning, Montréal, Canada <sup>4</sup>Center for Computational Astrophysics, Flatiron Institute, New York, USA. Correspondence to: Mohammad-Hadi Sotoudeh <hadi.sotoudeh@umontreal.ca>.

<sup>&</sup>lt;sup>1</sup>Here, *primordial* refers to the CMB signal in the absence of gravitational lensing.

# 2. Methods

## 2.1. Architecture

Our inference framework draws samples from the highdimensional posterior distribution  $p(\mathbf{x}|\mathbf{y})$  of parameters  $\mathbf{x}$ given the observed data  $\mathbf{y}$ .

To make the problem more tractable, we separate the task of learning the posterior mean from modeling the dispersion of samples around the mean. This decomposition enables using lighter-weight probabilistic networks, which are typically more expensive to train than deterministic ones. This is achieved through two separately trained networks:

- 1. The *Mean Network* learns the posterior mean,  $\bar{\mathbf{x}}(\mathbf{y}) = \mathbb{E}_{\mathbf{x}} \left[ p(\mathbf{x}|\mathbf{y}) \right].$
- 2. The *Dispersion Network* models the variability of posterior samples around the mean. It generates samples from the distribution of deviations, defined as the difference between a posterior sample and the mean,  $\delta := \mathbf{x} \bar{\mathbf{x}}$ .

The Mean Network follows the U-Net architecture (Ronneberger et al., 2015) and produces deterministic outputs. The Dispersion Network can be any generative model; we use the Hierarchical Probabilistic U-Net (HPU-Net) (Kohl et al., 2019), which introduces stochasticity by adding sampling layers to the expanding path of the U-Net.

#### 2.2. Training Objective

#### MEAN NETWORK

The Mean Network minimizes the Mean Squared Error (MSE) between its output  $\hat{\mathbf{x}}$  and the target x:

$$\ell_{\text{MeanNet}} = \frac{1}{2} \left\| \hat{\mathbf{x}} - \mathbf{x} \right\|_2^2 \,. \tag{1}$$

According to Adler & Öktem (2018), the optimal solution of training a deterministic neural network (i.e., one without sampling layers) using MSE loss is the posterior mean. Hence, if properly trained, the Mean Network is guaranteed to learn the posterior mean.

#### DISPERSION NETWORK

The Dispersion Network models the joint distribution of deviations  $\boldsymbol{\delta}$  and some latent variables  $\mathbf{z}$ , conditioned on observed data  $\mathbf{y}$ :  $p(\boldsymbol{\delta}, \mathbf{z}|\mathbf{y}) = p(\mathbf{z}|\mathbf{y}) p(\boldsymbol{\delta}|\mathbf{z}, \mathbf{y})$ . It is trained similarly to a VAE, where the network is tasked with maxi-

mizing the evidence lower bound (ELBO), i.e., minimizing

$$\ell_{\text{DispersionNet}} = \mathbb{E}_{z \sim q(\mathbf{z}|\boldsymbol{\delta}, \mathbf{y})} \left[ -\ln p(\boldsymbol{\delta}|\mathbf{z}, \mathbf{y}) \right] \\ + D_{\text{KL}} \left[ q(\mathbf{z}|\boldsymbol{\delta}, \mathbf{y}) \parallel p(\mathbf{z}|\mathbf{y}) \right] \quad (2)$$
$$= \ell_{\text{rec}} + \ell_{\text{KL}} .$$

The first term is the reconstruction loss, which measures how well the generated data matches the expected output. The second term is the Kullback–Leibler (KL) divergence between the variational posterior  $q(\mathbf{z}|\boldsymbol{\delta}, \mathbf{y})$  and the prior  $p(\mathbf{z}|\mathbf{y})$ . For further details, see Kohl et al. (2019).

We employ a diagonal Gaussian negative log-likelihood reconstruction loss,

$$\ell_{\rm rec} = \frac{1}{2} \left( \ln |\hat{\boldsymbol{\Sigma}}| + (\boldsymbol{\delta} - \hat{\boldsymbol{\mu}})^\top \hat{\boldsymbol{\Sigma}}^{-1} (\boldsymbol{\delta} - \hat{\boldsymbol{\mu}}) \right) = \frac{1}{2} \sum_{i} \left[ \ln(\hat{\sigma}_i^2) + \frac{(\delta_i - \hat{\mu}_i)^2}{\hat{\sigma}_i^2} \right],$$
(3)

where  $\hat{\mu}$  and  $\hat{\Sigma}$  (or  $\hat{\mu}_i$  and  $\hat{\sigma}_i^2$ ) are estimated from multiple samples generated by the Dispersion Network, and  $\delta = \mathbf{x} - \hat{\mathbf{x}}$  is the deviation of the ground-truth target from MeanNet's output. The sum is over all output pixels *i*.

This choice of reconstruction loss offers two advantages: (i) For networks with *leaky paths*<sup>2</sup>, maximizing the ELBO can lead to the KL vanishing problem (Bowman et al., 2016; Fu et al., 2019), causing the network to ignore the latent variables and collapse to a deterministic solution. Maximizing the ELBO under the proposed reconstruction loss does not suffer from this problem; (ii) In VAE-like frameworks, *directly optimizing* the variance  $\hat{\sigma}_i^2$  can cause it to collapse to zero (Rezende & Viola, 2018), resulting in a deterministic model. By estimating the variance from posterior samples, we avoid this collapse while including a loss term that provides a useful signal for calibrating the model's uncertainty estimates. A key component of this design is the VAE-based approach, which allows fast sampling and makes training feasible. In contrast, implementing this strategy with a diffusion model would be significantly more expensive.

#### 2.3. Baseline

For some experiments, we compare performance against a baseline in which the Dispersion Network is replaced with a conditional diffusion model, trained using the denoising diffusion probabilistic model (DDPM) framework of Ho et al. (2020).

Further details on the architecture, training procedure, and baseline can be found in Appendix B.

 $<sup>^{2}</sup>$ A *leaky path* is a route in the network's computation graph that allows information to flow from the ground truth to the output without passing through the latent space - for example, skip connections in the U-Net architecture.



*Figure 1.* GRF rotation results: (a) Comparing empirical and theoretical moments for a test example. The pixel-wise empirical means and standard deviations are estimated using 1000 posterior samples. (b) TARP coverage test results.

#### **3. Experiments and Results**

#### 3.1. Problem 1: Rotating Gaussian Random Fields (GRFs)

This experiment evaluates our framework on a problem with an analytically tractable posterior: a linear inverse problem, where the data y and parameters x are related through  $y = \mathbf{R}\mathbf{x} + \mathbf{n}$ . Here, **R** is a transformation matrix, and **n** is a random noise vector. If the parameter prior  $p(\mathbf{x})$ and the noise distribution  $p(\mathbf{n})$  are both Gaussian, then the posterior distribution  $p(\mathbf{x}|\mathbf{y})$  is also Gaussian (Kaipio & Somersalo, 2005), with the following parameters:

$$\boldsymbol{\mu}_{\text{post}} = (\mathbf{M} + \mathbf{M}^{\top})^{-1} \mathbf{d} \qquad \boldsymbol{\Sigma}_{\text{post}}^{-1} = \mathbf{M}, \\ \mathbf{d} := 2 \left( \boldsymbol{\Sigma}_{\text{pri}}^{-1} \boldsymbol{\mu}_{\text{pri}} + \mathbf{R}^{\top} \boldsymbol{\Sigma}_{n}^{-1} \mathbf{y} \right) \qquad \mathbf{M} := \left( \boldsymbol{\Sigma}_{\text{pri}}^{-1} + \mathbf{R}^{\top} \boldsymbol{\Sigma}_{n}^{-1} \mathbf{R} \right)$$
(4)

where the *pri*, *post*, and *n* subscripts indicate the prior, posterior, and noise distributions, respectively. We set **R** to be a 90° rotation. To evaluate performance on this task, we compare  $\mu_{post}$  and  $\Sigma_{post}$  with the empirical mean and covariance estimated from the model's predictions (Figure 1a). We also employ the Test of Accuracy with Random Points (TARP) (Lemos et al., 2023) to assess the calibration of the credible regions produced using the model (Figure 1b).

#### 3.2. Problem 2: CMB Delensing

CMB lensing occurs when the cosmic microwave background is distorted by weak gravitational lensing from intervening cosmic structures between the last scattering surface and Earth (Lewis & Challinor, 2006; Hanson et al., 2010). Correcting for this effect is crucial for obtaining unbiased estimates of cosmological parameters, studying the distribution of matter in the universe, and removing the lensinginduced B-mode pattern, which can be confused with signals from primordial gravitational waves. Recent CMB lensing results and analyses are presented in Aghanim et al. (2020); Pan et al. (2023); Qu et al. (2024). Lately, deep learning has also been applied to CMB delensing. Specifically, Caldeira et al. (2019); Li et al. (2022); Yan et al. (2023a;b) used convolutional neural networks to produce point estimates of unlensed CMB maps. More recently, Flöss et al. (2024) proposed a diffusion-based approach.

In our delensing experiments, we use the Code for Anisotropies in the Microwave Background (CAMB) (Lewis & Bridle, 2002) to generate simulated CMB observations, with and without lensing effects. Specifically, we use the lensed and unlensed TT power spectra from CAMB to generate realizations of CMB maps corresponding to each spectrum. The inference framework is tasked with predicting the difference between the lensed and unlensed maps, given the lensed map as input. To evaluate performance on this task, we compare the power spectrum of the delensed CMB maps with the target (unlensed) power spectrum (Figure 2a). We also use TARP to assess the model's uncertainty calibration (Figure 2b). Finally, to test robustness, we evaluate the model on out-of-distribution (OOD) CMB maps generated using different cosmological parameters (Figure 2c). To construct an OOD dataset, we retain the same noise realizations as in the in-distribution dataset; however, we vary the matter density parameter  $\Omega_m$ by a multiple of the Planck satellite measurement uncertainty  $\sigma_{\Omega_m} = 0.0073$ , which affects the TT power spectra generated by CAMB. Using this procedure, we generate several OOD datasets with different  $\Omega_m$  values.

Further details on the data generation procedure can be found in Appendix A.

## 4. Discussion and Conclusion

Our framework closely approximates the posterior distribution, with slightly conservative credible intervals. In the GRF rotation experiment (Figure 1), the empirical and theoretical moments align closely, and the model's uncertainty estimates are well calibrated. For CMB delensing (Figure 2a), the unlensed spectrum consistently falls within the uncertainty regions computed from the power spectrum of

A Fast Generative Framework for High-dimensional Posterior Sampling



Figure 2. CMB delensing results: (a) Comparison of the lensed (input), unlensed (target), and delensed (obtained from model outputs) TT power spectra for six test examples. The delensed spectrum is represented by mean and uncertainty regions computed from 1000 posterior samples. Relative differences  $\Delta D_{\ell}/D_{\ell}$  are computed with respect to the mean delensed spectrum. (b) TARP coverage test results. (c) Out-of-distribution performance for different  $\Omega_m$  values, with all other cosmological parameters fixed. Each column corresponds to an  $\Omega_m$  value, which differs from the training (base) value by a factor of the Planck satellite measurement error  $\sigma_{\Omega_m}$ . Each row corresponds to a unique noise realization used to generate the lensed and unlensed maps.

posterior samples. Furthermore, the model generalizes well to out-of-distribution (OOD) inputs (Figure 2c): unless the matter density parameter  $\Omega_m$  is significantly perturbed from training values, the predicted uncertainty bands continue to contain the true unlensed spectrum. This indicates potential for application to other OOD scenarios, most importantly for generalizing from simulations to observations.

Table 1. Comparison of the time required to generate 50 posterior samples for a batch of 4 test examples on an NVIDIA H100 GPU. Uncertainties are estimated from 10 runs. The numbers next to the diffusion baselines indicate the number of denoising steps.

Model	Sampling Time (s)
Our Framework	$0.31\pm0.05$
Diffusion Baseline-100	$12.85\pm0.09$
Diffusion Baseline-1000	$125.6\pm0.4$

Furthermore, our inference framework provides uncertainty estimates, offering an advantage over deep learning approaches to CMB delensing that produce only point estimates. In addition, leveraging a VAE-based architecture enables fast inference (Bond-Taylor et al., 2022). As shown in Table 1, our model is at least 40 times faster than diffusionbased baselines. Beyond their slow sampling, conditional diffusion models are also known to produce overconfident uncertainty estimates (Tachella & Pereyra, 2023; Bourdin et al., 2024). By comparison, our model's uncertainty estimates are slightly conservative. Lastly, the proposed reconstruction loss prevents KL vanishing and enables implicit uncertainty calibration, ensuring latent variables remain informative throughout training and preventing collapse to a deterministic solution.

Future directions include improving uncertainty calibration, scaling the framework to higher-dimensional problems, applying it to observational data, and conducting more extensive comparisons against baselines. To improve calibration, one could explore modified loss functions that encourage predicting extreme values, thereby reducing the conservativeness of credible intervals. Scaling to higher dimensions would better highlight the advantages of this approach over diffusion models, particularly as the sampling speed gap becomes more pronounced. Applying the framework to observational data includes both training directly on real observations and investigating how models trained on simulations generalize — or transfer — to real-world settings.

In this paper, we presented a framework for highdimensional posterior sampling based on the Hierarchical Probabilistic U-Net. This method is fast, robust, and grounded in a sound theoretical foundation. It is well-suited for simulation-based inference scenarios where direct likelihood modeling is infeasible, enabling the characterization of high-dimensional posteriors and uncertainty estimation in physical measurements using neural networks.

## Software and Data

The code for generating data and training models will be made publicly available in conjunction with the journal publication.

## Acknowledgements

This work was carried out as part of the Simons Collaboration on *Learning the Universe*, which supported this research. It was enabled in part by computational resources provided by Calcul Québec, Compute Canada, and the Digital Research Alliance of Canada. MHS was supported by an MSc Excellence Scholarship from IVADO. We thank the reviewers for their valuable comments. MHS gratefully acknowledges Steven Dillmann's support during the preparation of this work for the ML4Astro workshop.

## **Impact Statement**

This paper presents work whose goal is to advance the field of Machine Learning and its applications in Astrophysics. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

#### References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mane, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viegas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. Tensorflow: Largescale machine learning on heterogeneous distributed systems, 2016.
- Abazajian, K. N., Adshead, P., Ahmed, Z., Allen, S. W., Alonso, D., Arnold, K. S., Baccigalupi, C., Bartlett, J. G., Battaglia, N., Benson, B. A., Bischoff, C. A., Borrill, J., Buza, V., Calabrese, E., Caldwell, R., Carlstrom, J. E., Chang, C. L., Crawford, T. M., Cyr-Racine, F.-Y., Bernardis, F. D., de Haan, T., di Serego Alighieri, S., Dunkley, J., Dvorkin, C., Errard, J., Fabbian, G., Feeney, S., Ferraro, S., Filippini, J. P., Flauger, R., Fuller, G. M., Gluscevic, V., Green, D., Grin, D., Grohs, E., Henning, J. W., Hill, J. C., Hlozek, R., Holder, G., Holzapfel, W., Hu, W., Huffenberger, K. M., Keskitalo, R., Knox, L., Kosowsky, A., Kovac, J., Kovetz, E. D., Kuo, C.-L., Kusaka, A., Jeune, M. L., Lee, A. T., Lilley, M., Loverde, M., Madhavacheril, M. S., Mantz, A., Marsh, D. J. E., McMahon, J., Meerburg, P. D., Meyers, J., Miller, A. D., Munoz, J. B., Nguyen, H. N., Niemack, M. D.,

Peloso, M., Peloton, J., Pogosian, L., Pryke, C., Raveri, M., Reichardt, C. L., Rocha, G., Rotti, A., Schaan, E., Schmittfull, M. M., Scott, D., Sehgal, N., Shandera, S., Sherwin, B. D., Smith, T. L., Sorbo, L., Starkman, G. D., Story, K. T., van Engelen, A., Vieira, J. D., Watson, S., Whitehorn, N., and Wu, W. L. K. Cmb-s4 science book, first edition, 2016.

- Adler, J. and Öktem, O. Deep bayesian inversion, 2018. URL https://arxiv.org/abs/1811.05910.
- Aghanim, N., Akrami, Y., Ashdown, M., Aumont, J., Baccigalupi, C., Ballardini, M., Banday, A. J., Barreiro, R. B., Bartolo, N., Basak, S., Benabed, K., Bernard, J.-P., Bersanelli, M., Bielewicz, P., Bock, J. J., Bond, J. R., Borrill, J., Bouchet, F. R., Boulanger, F., Bucher, M., Burigana, C., Calabrese, E., Cardoso, J.-F., Carron, J., Challinor, A., Chiang, H. C., Colombo, L. P. L., Combet, C., Crill, B. P., Cuttaia, F., de Bernardis, P., de Zotti, G., Delabrouille, J., Di Valentino, E., Diego, J. M., Doré, O., Douspis, M., Ducout, A., Dupac, X., Efstathiou, G., Elsner, F., Enßlin, T. A., Eriksen, H. K., Fantaye, Y., Fernandez-Cobos, R., Finelli, F., Forastieri, F., Frailis, M., Fraisse, A. A., Franceschi, E., Frolov, A., Galeotta, S., Galli, S., Ganga, K., Génova-Santos, R. T., Gerbino, M., Ghosh, T., González-Nuevo, J., Górski, K. M., Gratton, S., Gruppuso, A., Gudmundsson, J. E., Hamann, J., Handley, W., Hansen, F. K., Herranz, D., Hivon, E., Huang, Z., Jaffe, A. H., Jones, W. C., Karakci, A., Keihänen, E., Keskitalo, R., Kiiveri, K., Kim, J., Knox, L., Krachmalnicoff, N., Kunz, M., Kurki-Suonio, H., Lagache, G., Lamarre, J.-M., Lasenby, A., Lattanzi, M., Lawrence, C. R., Le Jeune, M., Levrier, F., Lewis, A., Liguori, M., Lilje, P. B., Lindholm, V., López-Caniego, M., Lubin, P. M., Ma, Y.-Z., Macías-Pérez, J. F., Maggio, G., Maino, D., Mandolesi, N., Mangilli, A., Marcos-Caballero, A., Maris, M., Martin, P. G., Martínez-González, E., Matarrese, S., Mauri, N., McEwen, J. D., Melchiorri, A., Mennella, A., Migliaccio, M., Miville-Deschênes, M.-A., Molinari, D., Moneti, A., Montier, L., Morgante, G., Moss, A., Natoli, P., Pagano, L., Paoletti, D., Partridge, B., Patanchon, G., Perrotta, F., Pettorino, V., Piacentini, F., Polastri, L., Polenta, G., Puget, J.-L., Rachen, J. P., Reinecke, M., Remazeilles, M., Renzi, A., Rocha, G., Rosset, C., Roudier, G., Rubiño-Martín, J. A., Ruiz-Granados, B., Salvati, L., Sandri, M., Savelainen, M., Scott, D., Sirignano, C., Sunyaev, R., Suur-Uski, A.-S., Tauber, J. A., Tavagnacco, D., Tenti, M., Toffolatti, L., Tomasi, M., Trombetti, T., Valiviita, J., Van Tent, B., Vielva, P., Villa, F., Vittorio, N., Wandelt, B. D., Wehus, I. K., White, M., White, S. D. M., Zacchei, A., and Zonca, A. Planck2018 results: Viii. gravitational lensing. Astronomy & Astrophysics, 641:A8, September 2020. ISSN 1432-0746. doi: 10.1051/0004-6361/201833886. URL http://dx. doi.org/10.1051/0004-6361/201833886.

- Bellm, E. C., Ford, E. B., Tohuvavohu, A., Coughlin, M. W., Morris, B., Miller, B., Sobeck, J., Riddle, R., Dong, C., and Yoachim, P. Scheduling discovery in the 2020s, 2019. URL https://arxiv.org/abs/1907.07817.
- Biewald, L. Experiment tracking with weights and biases, 2020. URL https://www.wandb.com/. Software available from wandb.com.
- Bond-Taylor, S., Leach, A., Long, Y., and Willcocks, C. G.
  Deep generative modelling: A comparative review of vaes, gans, normalizing flows, energy-based and autore-gressive models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):7327–7347, November 2022. ISSN 1939-3539. doi: 10.1109/tpami.2021. 3116668. URL http://dx.doi.org/10.1109/TPAMI.2021.3116668.
- Bourdin, A., Legin, R., Ho, M., Adam, A., Hezaveh, Y., and Perreault-Levasseur, L. Inpainting galaxy counts onto n-body simulations over multiple cosmologies and astrophysics, 2024. URL https://arxiv.org/abs/ 2408.00839.
- Bowman, S. R., Vilnis, L., Vinyals, O., Dai, A. M., Jozefowicz, R., and Bengio, S. Generating sentences from a continuous space, 2016. URL https://arxiv.org/ abs/1511.06349.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. JAX: composable transformations of Python+NumPy programs, 2018. URL http://github.com/google/jax.
- Caldeira, J., Wu, W., Nord, B., Avestruz, C., Trivedi, S., and Story, K. DeepCMB: Lensing reconstruction of the cosmic microwave background with deep neural networks. *Astronomy and Computing*, 28:100307, jul 2019. doi: 10.1016/j.ascom.2019.100307. URL https://doi. org/10.1016%2Fj.ascom.2019.100307.
- Cranmer, K., Brehmer, J., and Louppe, G. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, May 2020. ISSN 1091-6490. doi: 10.1073/ pnas.1912789117. URL http://dx.doi.org/10. 1073/pnas.1912789117.
- Dewdney, P. E., Hall, P. J., Schilizzi, R. T., and Lazio, T. J. L. W. The Square Kilometre Array. *IEEE Proceedings*, 97(8):1482–1496, August 2009. doi: 10.1109/ JPROC.2009.2021005.
- Durkan, C., Murray, I., and Papamakarios, G. On contrastive learning for likelihood-free inference, 2020. URL https://arxiv.org/abs/2002.03712.

- Flöss, T., Coulton, W. R., Duivenvoorden, A. J., Villaescusa-Navarro, F., and Wandelt, B. D. Denoising diffusion delensing delight: Reconstructing the non-gaussian cmb lensing potential with diffusion models, 2024. URL https://arxiv.org/abs/2405.05598.
- Fu, H., Li, C., Liu, X., Gao, J., Celikyilmaz, A., and Carin, L. Cyclical annealing schedule: A simple approach to mitigating kl vanishing, 2019. URL https://arxiv. org/abs/1903.10145.
- Galitzki, N., Tsan, T., Spisak, J., Randall, M., Silva-Feaver, M., Seibert, J., Lashner, J., Adachi, S., Adkins, S. M., Alford, T., Arnold, K., Ashton, P. C., Austermann, J. E., Baccigalupi, C., Bazarko, A., Beall, J. A., Bhimani, S., Bixler, B., Coppi, G., Corbett, L., Crowley, K. D., Crowley, K. T., Day-Weiss, S., Dicker, S., Dow, P. N., Duell, C. J., Duff, S. M., Gerras, R. G., Groh, J. C., Gudmundsson, J. E., Harrington, K., Hasegawa, M., Healy, E., Henderson, S. W., Hubmayr, J., Iuliano, J., Johnson, B. R., Keating, B., Keller, B., Kiuchi, K., Kofman, A. M., Koopman, B. J., Kusaka, A., Lee, A. T., Lew, R. A., Lin, L. T., Link, M. J., Lucas, T. J., Lungu, M., Mangu, A., McMahon, J. J., Miller, A. D., Moore, J. E., Morshed, M., Nakata, H., Nati, F., Newburgh, L. B., Nguyen, D. V., Niemack, M. D., Page, L. A., Sakaguri, K., Sakurai, Y., Rao, M. S., Saunders, L. J., Shroyer, J. E., Sugiyama, J., Tajima, O., Takeuchi, A., Bua, R. T., Teply, G., Terasaki, T., Ullom, J. N., Lanen, J. L. V., Vavagiakis, E. M., Vissers, M. R., Walters, L., Wang, Y., Xu, Z., Yamada, K., and Zheng, K. The simons observatory: Design, integration, and testing of the small aperture telescopes, 2024. URL https://arxiv.org/abs/2405.05550.
- Gardner, J. P., Mather, J. C., Clampin, M., Doyon, R., Greenhouse, M. A., Hammel, H. B., Hutchings, J. B., Jakobsen, P., Lilly, S. J., Long, K. S., Lunine, J. I., Mccaughrean, M. J., Mountain, M., Nella, J., Rieke, G. H., Rieke, M. J., Rix, H.-W., Smith, E. P., Sonneborn, G., Stiavelli, M., Stockman, H. S., Windhorst, R. A., and Wright, G. S. The james webb space telescope. *Space Science Reviews*, 123(4):485–606, apr 2006. doi: 10.1007/s11214-006-8315-7. URL https:// doi.org/10.1007%2Fs11214-006-8315-7.
- Hanson, D., Challinor, A., and Lewis, A. Weak lensing of the CMB. *General Relativity and Gravitation*, 42(9):2197–2218, jun 2010. doi: 10.1007/ s10714-010-1036-y. URL https://doi.org/10. 1007%2Fs10714-010-1036-y.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard,

K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E. Array programming with NumPy. *Nature*, 585(7825):357–362, sep 2020. doi: 10.1038/ s41586-020-2649-2. URL https://doi.org/10. 1038%2Fs41586-020-2649-2.

- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models, 2020.
- Hunter, J. D. Matplotlib: A 2d graphics environment. Computing in Science & Engineering, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55.
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., Abel, B., Acosta, E., Allsman, R., Alonso, D., AlSayyad, Y., Anderson, S. F., Andrew, J., Angel, J. R. P., Angeli, G. Z., Ansari, R., Antilogus, P., Araujo, C., Armstrong, R., Arndt, K. T., Astier, P., Aubourg, É., Auza, N., Axelrod, T. S., Bard, D. J., Barr, J. D., Barrau, A., Bartlett, J. G., Bauer, A. E., Bauman, B. J., Baumont, S., Bechtol, E., Bechtol, K., Becker, A. C., Becla, J., Beldica, C., Bellavia, S., Bianco, F. B., Biswas, R., Blanc, G., Blazek, J., Blandford, R. D., Bloom, J. S., Bogart, J., Bond, T. W., Booth, M. T., Borgland, A. W., Borne, K., Bosch, J. F., Boutigny, D., Brackett, C. A., Bradshaw, A., Brandt, W. N., Brown, M. E., Bullock, J. S., Burchat, P., Burke, D. L., Cagnoli, G., Calabrese, D., Callahan, S., Callen, A. L., Carlin, J. L., Carlson, E. L., Chandrasekharan, S., Charles-Emerson, G., Chesley, S., Cheu, E. C., Chiang, H.-F., Chiang, J., Chirino, C., Chow, D., Ciardi, D. R., Claver, C. F., Cohen-Tanugi, J., Cockrum, J. J., Coles, R., Connolly, A. J., Cook, K. H., Cooray, A., Covey, K. R., Cribbs, C., Cui, W., Cutri, R., Daly, P. N., Daniel, S. F., Daruich, F., Daubard, G., Daues, G., Dawson, W., Delgado, F., Dellapenna, A., de Peyster, R., de Val-Borro, M., Digel, S. W., Doherty, P., Dubois, R., Dubois-Felsmann, G. P., Durech, J., Economou, F., Eifler, T., Eracleous, M., Emmons, B. L., Fausti Neto, A., Ferguson, H., Figueroa, E., Fisher-Levine, M., Focke, W., Foss, M. D., Frank, J., Freemon, M. D., Gangler, E., Gawiser, E., Geary, J. C., Gee, P., Geha, M., Gessner, C. J. B., Gibson, R. R., Gilmore, D. K., Glanzman, T., Glick, W., Goldina, T., Goldstein, D. A., Goodenow, I., Graham, M. L., Gressler, W. J., Gris, P., Guy, L. P., Guyonnet, A., Haller, G., Harris, R., Hascall, P. A., Haupt, J., Hernandez, F., Herrmann, S., Hileman, E., Hoblitt, J., Hodgson, J. A., Hogan, C., Howard, J. D., Huang, D., Huffer, M. E., Ingraham, P., Innes, W. R., Jacoby, S. H., Jain, B., Jammes, F., Jee, M. J., Jenness, T., Jernigan, G., Jevremović, D., Johns, K., Johnson, A. S., Johnson, M. W. G., Jones, R. L., Juramy-Gilles, C., Jurić, M., Kalirai, J. S., Kallivayalil, N. J., Kalmbach, B., Kantor, J. P., Karst, P., Kasliwal, M. M., Kelly, H., Kessler, R., Kinnison, V., Kirkby, D., Knox, L., Kotov, I. V., Krabbendam, V. L., Krughoff, K. S., Kubánek, P., Kuczewski, J., Kulkarni, S., Ku, J.,

Kurita, N. R., Lage, C. S., Lambert, R., Lange, T., Langton, J. B., Le Guillou, L., Levine, D., Liang, M., Lim, K.-T., Lintott, C. J., Long, K. E., Lopez, M., Lotz, P. J., Lupton, R. H., Lust, N. B., MacArthur, L. A., Mahabal, A., Mandelbaum, R., Markiewicz, T. W., Marsh, D. S., Marshall, P. J., Marshall, S., May, M., McKercher, R., McQueen, M., Meyers, J., Migliore, M., Miller, M., Mills, D. J., Miraval, C., Moeyens, J., Moolekamp, F. E., Monet, D. G., Moniez, M., Monkewitz, S., Montgomery, C., Morrison, C. B., Mueller, F., Muller, G. P., Muñoz Arancibia, F., Neill, D. R., Newbry, S. P., Nief, J.-Y., Nomerotski, A., Nordby, M., O'Connor, P., Oliver, J., Olivier, S. S., Olsen, K., O'Mullane, W., Ortiz, S., Osier, S., Owen, R. E., Pain, R., Palecek, P. E., Parejko, J. K., Parsons, J. B., Pease, N. M., Peterson, J. M., Peterson, J. R., Petravick, D. L., Libby Petrick, M. E., Petry, C. E., Pierfederici, F., Pietrowicz, S., Pike, R., Pinto, P. A., Plante, R., Plate, S., Plutchak, J. P., Price, P. A., Prouza, M., Radeka, V., Rajagopal, J., Rasmussen, A. P., Regnault, N., Reil, K. A., Reiss, D. J., Reuter, M. A., Ridgway, S. T., Riot, V. J., Ritz, S., Robinson, S., Roby, W., Roodman, A., Rosing, W., Roucelle, C., Rumore, M. R., Russo, S., Saha, A., Sassolas, B., Schalk, T. L., Schellart, P., Schindler, R. H., Schmidt, S., Schneider, D. P., Schneider, M. D., Schoening, W., Schumacher, G., Schwamb, M. E., Sebag, J., Selvy, B., Sembroski, G. H., Seppala, L. G., Serio, A., Serrano, E., Shaw, R. A., Shipsey, I., Sick, J., Silvestri, N., Slater, C. T., Smith, J. A., Smith, R. C., Sobhani, S., Soldahl, C., Storrie-Lombardi, L., Stover, E., Strauss, M. A., Street, R. A., Stubbs, C. W., Sullivan, I. S., Sweeney, D., Swinbank, J. D., Szalay, A., Takacs, P., Tether, S. A., Thaler, J. J., Thayer, J. G., Thomas, S., Thornton, A. J., Thukral, V., Tice, J., Trilling, D. E., Turri, M., Van Berg, R., Vanden Berk, D., Vetter, K., Virieux, F., Vucina, T., Wahl, W., Walkowicz, L., Walsh, B., Walter, C. W., Wang, D. L., Wang, S.-Y., Warner, M., Wiecha, O., Willman, B., Winters, S. E., Wittman, D., Wolff, S. C., Wood-Vasey, W. M., Wu, X., Xin, B., Yoachim, P., and Zhan, H. LSST: From Science Drivers to Reference Design and Anticipated Data Products. Astrophysical Journal, 873(2):111, March 2019. doi: 10.3847/1538-4357/ab042c.

- Kaipio, J. and Somersalo, E. Statistical and Computational Inverse Problems. Springer, Dordrecht, 2005. doi: 10.1007/b138659. URL https://cds.cern.ch/ record/1338003.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes, 2022.
- Kohl, S. A. A., Romera-Paredes, B., Maier-Hein, K. H., Rezende, D. J., Eslami, S. M. A., Kohli, P., Zisserman, A., and Ronneberger, O. A hierarchical probabilistic unet for modeling multi-scale ambiguities, 2019. URL https://arxiv.org/abs/1905.13077.

- Lemos, P., Coogan, A., Hezaveh, Y., and Perreault-Levasseur, L. Sampling-based accuracy testing of posterior estimators for general inference, 2023.
- Lewis, A. and Bridle, S. Cosmological parameters from CMB and other data: A Monte Carlo approach. *Physical Review D*, 66:103511, 2002. doi: 10.1103/PhysRevD.66. 103511.
- Lewis, A. and Challinor, A. Weak gravitational lensing of the CMB. *Physics Reports*, 429(1):1–65, jun 2006. doi: 10.1016/j.physrep.2006.03.002. URL https://doi. org/10.1016%2Fj.physrep.2006.03.002.
- Li, P., Onur, I. I., Dodelson, S., and Chaudhari, S. Highresolution cmb lensing reconstruction with deep learning, 2022.
- Ntampaka, M., Avestruz, C., Boada, S., Caldeira, J., Cisewski-Kehe, J., Stefano, R. D., Dvorkin, C., Evrard, A. E., Farahi, A., Finkbeiner, D., Genel, S., Goodman, A., Goulding, A., Ho, S., Kosowsky, A., Plante, P. L., Lanusse, F., Lochner, M., Mandelbaum, R., Nagai, D., Newman, J. A., Nord, B., Peek, J. E. G., Peel, A., Poczos, B., Rau, M. M., Siemiginowska, A., Sutherland, D. J., Trac, H., and Wandelt, B. The role of machine learning in the next decade of cosmology, 2021. URL https://arxiv.org/abs/1902.10159.
- Pan, Z., Bianchini, F., Wu, W. L. K., Ade, P. A. R., Ahmed, Z., Anderes, E., Anderson, A. J., Ansarinejad, B., Archipley, M., Aylor, K., Balkenhol, L., Barry, P. S., Basu Thakur, R., Benabed, K., Bender, A. N., Benson, B. A., Bleem, L. E., Bouchet, F. R., Bryant, L., Byrum, K., Camphuis, E., Carlstrom, J. E., Carter, F. W., Cecil, T. W., Chang, C. L., Chaubal, P., Chen, G., Chichura, P. M., Cho, H.-M., Chou, T.-L., Cliche, J.-F., Coerver, A., Crawford, T. M., Cukierman, A., Daley, C., de Haan, T., Denison, E. V., Dibert, K. R., Ding, J., Dobbs, M. A., Doussot, A., Dutcher, D., Everett, W., Feng, C., Ferguson, K. R., Fichman, K., Foster, A., Fu, J., Galli, S., Gambrel, A. E., Gardner, R. W., Ge, F., Goeckner-Wald, N., Gualtieri, R., Guidi, F., Guns, S., Gupta, N., Halverson, N. W., Harke-Hosemann, A. H., Harrington, N. L., Henning, J. W., Hilton, G. C., Hivon, E., Holder, G. P., Holzapfel, W. L., Hood, J. C., Howe, D., Huang, N., Irwin, K. D., Jeong, O., Jonas, M., Jones, A., Kéruzoré, F., Khaire, T. S., Knox, L., Kofman, A. M., Korman, M., Kubik, D. L., Kuhlmann, S., Kuo, C.-L., Lee, A. T., Leitch, E. M., Levy, K., Lowitz, A. E., Lu, C., Maniyar, A., Menanteau, F., Meyer, S. S., Michalik, D., Millea, M., Montgomery, J., Nadolski, A., Nakato, Y., Natoli, T., Nguyen, H., Noble, G. I., Novosad, V., Omori, Y., Padin, S., Paschos, P., Pearson, J., Posada, C. M., Prabhu, K., Quan, W., Raghunathan, S., Rahimi, M., Rahlin, A., Reichardt, C. L., Riebel, D., Riedel, B., Ruhl, J. E., Sayre,

J. T., Schiappucci, E., Shirokoff, E., Smecher, G., Sobrin, J. A., Stark, A. A., Stephen, J., Story, K. T., Suzuki, A., Takakura, S., Tandoi, C., Thompson, K. L., Thorne, B., Trendafilova, C., Tucker, C., Umilta, C., Vale, L. R., Vanderlinde, K., Vieira, J. D., Wang, G., Whitehorn, N., Yefremenko, V., Yoon, K. W., Young, M. R., and Zebrowski, J. A. Measurement of gravitational lensing of the cosmic microwave background using spt-3g 2018 data. *Phys. Rev. D*, 108:122005, Dec 2023. doi: 10.1103/PhysRevD.108.122005. URL https://link.aps.org/doi/10.1103/PhysRevD.108.122005.

- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, 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. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips. cc/paper\_files/paper/2019/file/ bdbca288fee7f92f2bfa9f7012727740-Paper. pdf.
- Qu, F. J., Sherwin, B. D., Madhavacheril, M. S., Han, D., Crowley, K. T., Abril-Cabezas, I., Ade, P. A. R., Aiola, S., Alford, T., Amiri, M., Amodeo, S., An, R., Atkins, Z., Austermann, J. E., Battaglia, N., Battistelli, E. S., Beall, J. A., Bean, R., Beringue, B., Bhandarkar, T., Biermann, E., Bolliet, B., Bond, J. R., Cai, H., Calabrese, E., Calafut, V., Capalbo, V., Carrero, F., Carron, J., Challinor, A., Chesmore, G. E., mei Cho, H., Choi, S. K., Clark, S. E., Rosado, R. C., Cothard, N. F., Coughlin, K., Coulton, W., Dalal, R., Darwish, O., Devlin, M. J., Dicker, S., Doze, P., Duell, C. J., Duff, S. M., Duivenvoorden, A. J., Dunkley, J., Dünner, R., Fanfani, V., Fankhanel, M., Farren, G., Ferraro, S., Freundt, R., Fuzia, B., Gallardo, P. A., Garrido, X., Gluscevic, V., Golec, J. E., Guan, Y., Halpern, M., Harrison, I., Hasselfield, M., Healy, E., Henderson, S., Hensley, B., Hervías-Caimapo, C., Hill, J. C., Hilton, G. C., Hilton, M., Hincks, A. D., Hložek, R., Ho, S.-P. P., Huber, Z. B., Hubmayr, J., Huffenberger, K. M., Hughes, J. P., Irwin, K., Isopi, G., Jense, H. T., Keller, B., Kim, J., Knowles, K., Koopman, B. J., Kosowsky, A., Kramer, D., Kusiak, A., Posta, A. L., Lague, A., Lakey, V., Lee, E., Li, Z., Li, Y., Limon, M., Lokken, M., Louis, T., Lungu, M., MacCrann, N., MacInnis, A., Maldonado, D., Maldonado, F., Mallaby-Kay, M., Marques, G. A., McMahon, J., Mehta, Y., Menanteau, F., Moodley, K., Morris, T. W., Mroczkowski, T., Naess, S., Namikawa, T., Nati, F., Newburgh, L., Nicola, A., Niemack, M. D., Nolta, M. R., Orlowski-Scherer, J., Page, L. A., Pandey, S., Partridge, B., Prince, H., Puddu, R., Radiconi, F., Robertson, N.,

Rojas, F., Sakuma, T., Salatino, M., Schaan, E., Schmitt, B. L., Sehgal, N., Shaikh, S., Sierra, C., Sievers, J., Sifón, C., Simon, S., Sonka, R., Spergel, D. N., Staggs, S. T., Storer, E., Switzer, E. R., Tampier, N., Thornton, R., Trac, H., Treu, J., Tucker, C., Ullom, J., Vale, L. R., Engelen, A. V., Lanen, J. V., van Marrewijk, J., Vargas, C., Vavagiakis, E. M., Wagoner, K., Wang, Y., Wenzl, L., Wollack, E. J., Xu, Z., Zago, F., and Zheng, K. The atacama cosmology telescope: A measurement of the dr6 cmb lensing power spectrum and its implications for structure growth. *The Astrophysical Journal*, 962(2):112, feb 2024. doi: 10.3847/1538-4357/acfe06. URL https: //dx.doi.org/10.3847/1538-4357/acfe06.

- Racca, G. D., Laureijs, R., Stagnaro, L., Salvignol, J.-C., Alvarez, J. L., Criado, G. S., Venancio, L. G., Short, A., Strada, P., Bönke, T., Colombo, C., Calvi, A., Maiorano, E., Piersanti, O., Prezelus, S., Rosato, P., Pinel, J., Rozemeijer, H., Lesna, V., Musi, P., Sias, M., Anselmi, A., Cazaubiel, V., Vaillon, L., Mellier, Y., Amiaux, J., Berthé, M., Sauvage, M., Azzollini, R., Cropper, M., Pottinger, S., Jahnke, K., Ealet, A., Maciaszek, T., Pasian, F., Zacchei, A., Scaramella, R., Hoar, J., Kohley, R., Vavrek, R., Rudolph, A., and Schmidt, M. The euclid mission design. In MacEwen, H. A., Fazio, G. G., Lystrup, M., Batalha, N., Siegler, N., and Tong, E. C. (eds.), *SPIE Proceedings*. SPIE, jul 2016. doi: 10.1117/12.2230762. URL https://doi.org/10.1117%2F12.2230762.
- Rezende, D. J. and Viola, F. Taming vaes, 2018. URL https://arxiv.org/abs/1810.00597.
- Ronneberger, O., Fischer, P., and Brox, T. U-net: Convolutional networks for biomedical image segmentation, 2015. URL https://arxiv.org/abs/1505.04597.
- Tachella, J. and Pereyra, M. Equivariant bootstrapping for uncertainty quantification in imaging inverse problems, 2023. URL https://arxiv.org/abs/ 2310.11838.
- von Platen, P., Patil, S., Lozhkov, A., Cuenca, P., Lambert, N., Rasul, K., Davaadorj, M., Nair, D., Paul, S., Berman, W., Xu, Y., Liu, S., and Wolf, T. Diffusers: State-of-the-art diffusion models. https://github. com/huggingface/diffusers, 2022.
- Yan, Y.-P., Wang, G.-J., Li, S.-Y., and Xia, J.-Q. Delensing of cosmic microwave background polarization with machine learning. *The Astrophysical Journal Supplement Series*, 267(1):2, jun 2023a. doi: 10.3847/1538-4365/acd2ce. URL https://doi.org/10.3847%2F1538-4365%2Facd2ce.
- Yan, Y.-P., Wang, G.-J., Li, S.-Y., Yan, Y.-J., and Xia, J.-Q. Lensing reconstruction from the cosmic microwave background polarization with machine learning. *The*

Astrophysical Journal, 952(1):15, jul 2023b. doi: 10. 3847/1538-4357/acdb72. URL https://doi.org/10.3847%2F1538-4357%2Facdb72.

## A. Data Generation

#### A.1. GRF Rotation

Figure 3 illustrates the data generation process for the GRF rotation experiment. We train (test) the model on  $2^{16}$  ( $2^{11}$ ) examples with  $32 \times 32$  resolution. The data was generated by sampling from the prior distribution ( $\mu_{pri} = 0, \sigma_{pri} = 1$ ), rotating the samples by  $90^{\circ}$ , and adding noise ( $\sigma_n = 0.5$ ) to the rotated samples. In order to increase the effective size of the training set and enhance the model's generalization, each training epoch uses different noise realizations.



Figure 3. Data generation process for the GRF rotation experiment.

## A.2. CMB Delensing

Figure 4 illustrates the data generation process for the CMB delensing experiment. We begin by calculating the lensed and unlensed TT power spectra using CAMB, based on the cosmological parameters listed in Table 2. We then apply the same noise realization to these power spectra to generate lensed and unlensed maps in Fourier space, and then transform the resulting maps to real space. The training (test) set contains  $2^{13}$  ( $2^{11}$ ) maps, each covering a  $160' \times 160'$  region with  $32 \times 32$  resolution. To simulate observational noise, we add random noise ( $\sigma_n = 20$ ) to the lensed maps. Similar to the GRF rotation experiment, we apply noise during training and use different realizations for each epoch. The input to the framework is a lensed map, and the target is the **difference** between the lensed and unlensed maps.

Parameter				
$H_0$	$H_0$ Hubble parameter (km s <sup>-1</sup> Mpc <sup>-1</sup> )			
$\Omega_m$	Matter density parameter	0.27		
$\Omega_k$	Curvature density parameter	0		
au	Thomson scattering optical depth	0.06		
$\ln\left(10^{10}A_s\right)$	Primordial curvature perturbations	2.996		
$n_s$	Scalar spectrum power-law index	0.965		
$\Sigma m_{\nu}$	Sum of active neutrino masses (eV)	0.06		
r	Tensor-to-scalar ratio	0		

Table 2. Cosmological parameters for the simulated CMB maps.

A Fast Generative Framework for High-dimensional Posterior Sampling



Figure 4. Data generation process for the CMB delensing experiment.

# **B.** Training

### **B.1.** Hyperparameters

The following hyperparameters were used for the experiments presented in the results:

Data Normalization. We use Gaussian standardization for the GRF rotation experiment and maximum absolute scaling for the CMB delensing experiment. The normalization coefficients are calculated globally, i.e., based on the entire dataset rather than individually for each example. We use the training set's normalization coefficients during inference.

Architecture. All U-Nets have a depth of 4, with a downsampling and upsampling scale factor of 2. For  $32 \times 32$  inputs, this results in a feature map dimensionality of  $2 \times 2$  at the U-Net bottleneck. We employ average pooling for downsampling and nearest-neighbor interpolation for upsampling, except for the GRF rotation Mean Network, which uses strided convolution and transpose convolution. The number of channels for different models is summarized in Table 3. All models use the following kernel sizes for their respective convolutional layers: [7, 7, 7, 5, 3].

Model		Input Channels	Intermediate Channels <sup>+</sup>
GRF Rotation	Mean Network	1	[32, 64, 128, 128, 128]
	Dispersion Network	2*	[8, 16, 32, 32, 32]
CMB Delensing	Mean Network	1	[64, 128, 256, 256, 256]
	Dispersion Network	2*	[16, 32, 64, 64, 128]

Table	3.	Model	architecture.

<sup>+</sup> Ordered from input to bottleneck. The expanding path mirrors this structure.

\* The Dispersion Network receives both the observation and the posterior mean.

Loss Function. When training the Dispersion Networks, we draw 10 samples from the model to estimate the posterior mean and covariance.

**Optimizer.** We employ the Adamax optimizer with a weight decay of  $10^{-5}$  and an initial learning rate of  $10^{-4}$ . All other optimizer hyperparameters are set to their PyTorch default values. The learning rate is then reduced by a factor of 5 at specific milestones: epochs 200 and 300 for the Mean Networks, every 20 epochs for the GRF rotation Dispersion Network, and every 250 epochs for the CMB delensing Dispersion Network. We use a batch size of 128.

## **B.2.** Compute Resources

Table 4 summarizes the model size and training time on a single NVIDIA V100 GPU. The training set size and number of epochs are also included to ensure meaningful comparisons.

Model		# Parameters	Dataset Size	Epochs	Training Time
GRF Rotation	Mean Network	7.5 M	$2^{16}$	400	25 hours
	Dispersion Network	780 k	$2^{16}$	50	13 hours
CMB Delensing	Mean Network	27.8 M	$2^{13}$	400	87 hours
	Dispersion Network	3.9 M	$2^{13}$	600	42 hours

Table 4.	Model	size	and	training	time.
inche it	11100001	OILC	unu	uanning	unic.

## **B.3.** Baseline

We establish our diffusion baselines using the HuggingFace Diffusers library (von Platen et al., 2022), specifically instantiating the UNet2DConditionModel. This architecture follows a symmetrical encoder-decoder U-Net design. The encoder consists of three standard convolutional blocks, followed by a cross-attention block that integrates conditioning information. The decoder mirrors this structure in reverse. The number of channels increases progressively through the convolutional blocks: 64, 128, and 256. The conditioning inputs, comprising the observation y and the Mean Network's prediction  $\hat{x}$ , are processed by a separate convolutional encoder. This encoder is composed of three sequential convolutional blocks. The resulting feature maps are restructured and used as keys and values in the U-Net's cross-attention mechanism.

We train this model using the denoising diffusion probabilistic model (DDPM) framework of Ho et al. (2020), with the Adamax optimizer and a cosine learning rate schedule with a peak learning rate of  $10^{-4}$  and and 500 warmup steps.

## **B.4.** Software

This research made use of CAMB 1.3.6 (Lewis & Bridle, 2002), PyTorch (Paszke et al., 2019), JAX (Bradbury et al., 2018), NumPy (Harris et al., 2020), Matplotlib (Hunter, 2007), TensorBoard (Abadi et al., 2016), and Weights&Biases (Biewald, 2020). We acknowledge the use of ChatGPT for guidance and clarification during the course of this research.