



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Deep learning simulations of the microwave sky

Dongwon Han, Neelima Sehgal, and Francisco Villaescusa-Navarro

Phys. Rev. D **104**, 123521 — Published 7 December 2021

DOI: [10.1103/PhysRevD.104.123521](https://doi.org/10.1103/PhysRevD.104.123521)

MillimeterDL: Deep Learning Simulations of the Microwave Sky

Dongwon Han,¹ Neelima Sehgal,¹ and Francisco Villaescusa-Navarro^{2,3}

¹*Physics and Astronomy Department, Stony Brook University, Stony Brook, NY 11794*

²*Department of Astrophysical Sciences, Princeton University, 4 Ivy Lane, Princeton, NJ, USA 08544*

³*Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY, USA 10010*

We present 500 high-resolution, full-sky millimeter-wave Deep Learning (DL) simulations that include lensed CMB maps and correlated foreground components. We find that these MillimeterDL simulations can reproduce a wide range of non-Gaussian summary statistics matching the input training simulations, while only being optimized to match the power spectra. The procedure we develop in this work enables the capability to mass produce independent full-sky realizations from a single expensive full-sky simulation, when ordinarily the latter would not provide enough training data. We also circumvent a common limitation of high-resolution DL simulations that they be confined to small sky areas, often due to memory or GPU issues; we do this by developing a “stitching” procedure that can recover the large-scale, high-order statistics and avoid discontinuities or repeated features in the maps. In addition, since our network takes as input a full-sky lensing convergence map, it can in principle take a full-sky lensing convergence map from any large-scale structure (LSS) simulation and generate the corresponding lensed CMB and correlated foreground components at millimeter wavelengths; this is especially useful in the current era of combining results from both CMB and LSS surveys, which require a common set of simulations.

I. INTRODUCTION

The Cosmic Microwave Background (CMB) has been a cornerstone of modern precision cosmology. The study of the primordial anisotropy imprinted on the CMB at the surface of last scattering has shed light on the physics of the early Universe. A rich set of secondary anisotropies induced by intervening Large Scale Structures (LSS) and complex astrophysical phenomena has revealed a picture of the evolving Universe. As the sensitivity of CMB experiments has improved, more information has been obtained from these features. A few recent examples include the constraints on cosmological parameters from CMB power spectra [1–5], the measurement of CMB lensing auto spectra [6–8], and the study of kinematic and thermal Sunyaev-Zel’dovich effects [9–12].

Current CMB data analyses rely heavily on simulations to (1) model the underlying physical processes, (2) verify the data analysis pipeline, (3) investigate potential biases, and (4) generate covariance matrices used in likelihood analyses. As a result, there have been efforts to improve simulations of the microwave sky [13–22]. Current and future high-resolution CMB experiments, such as the Atacama Cosmology Telescope (ACT) [23, 24], the South Pole Telescope (SPT) [25], the Simons Observatory (SO) [26], CMB-S4 [27], and CMB-HD [28, 29] will observe large sky areas at arcminute resolution or better with unprecedented sensitivity. Analyzing the data from these observations will need many realizations of realistic simulations of comparable resolution and sky coverage. In particular, as the sensitivity of CMB experiments improves, simulations capturing the correlations between foreground components and non-Gaussian features will become ever more important.

Simulations of the millimeter-wave sky can be di-

vided into several components: (1) primordial CMB maps, (2) extragalactic foregrounds, and (3) Galactic foregrounds. The procedure to generate simulations of primordial CMB maps is well-established and computationally efficient at arcminute resolution [30–34]. The extragalactic foreground components include the lensing convergence field (κ), the kinetic and thermal Sunyaev-Zel’dovich effects (kSZ and tSZ), the Cosmic Infrared Background (CIB), and radio galaxies (Radio). There are a few extragalactic foreground simulations publicly available that include all the above non-Gaussian, correlated foregrounds at high-resolution over a large fraction of the sky [19, 22]. These simulations are generated by first making three-dimensional realizations of cosmological features, and then projecting them along the line-of-sight to make two-dimensional maps. To a large extent, these existing simulations can reproduce the statistics of the latest CMB observations. However, there are only a limited number of realizations available due to the high computational cost of generating initial three-dimensional realizations. As a result, many CMB analyses model extragalactic foregrounds as Gaussian random fields matched to the observed foreground power spectra (e.g. [4–6]).

There have also been a number of efforts to simulate Galactic foregrounds in both intensity and polarization [35–43]. Modeling the small-scale features of Galactic foregrounds is particularly challenging given the current lack of observations on comparable scales in millimeter wavebands, though there is progress in this direction [44]. In this work, we focus on generating simulations including the primordial CMB and extragalactic foreground components, and leave inclusion of Galactic foregrounds to future work.

Recently, Machine Learning has found a broad application in cosmology; parameter estimation [45–52], foreground cleaning [42, 43, 53, 54], feature extraction [55–

58], inpainting [59–61], filtering [62], and modeling of physical processes [43, 63–69]. Deep Learning (DL) is a subset of Machine Learning that fits *multiple* non-linear functions to the input data using neural networks. In particular, generative DL models aim at sampling the manifold (which can be high-dimensional) where the input data lives, so that new points in that manifold are associated with new complex data with the same statistical properties as the desired input data. Recently, there have been many approaches in this direction [42–44, 70–81]. However, these methods are not yet widely used in CMB data analyses due to their limited sky-coverage or summary statistics mismatch.

In this work, we generate for the first time, many independent full-sky millimeter-wave DL simulations that include all correlated extragalactic foreground components appropriate for CMB analyses. We find that our MillimeterDL (hereafter mmDL) simulations are able to reproduce well the non-Gaussian statistics of the original input maps. Several recent works have focused on using various DL techniques to generate two-dimensional projected microwave fields from three-dimensional N-body simulations [70, 71, 73–78, 80, 81]; in this work, we use two-dimensional projected simulations at 148 GHz described in [19] (hereafter S10) as the input data to train our DL algorithm, as opposed to three-dimensional N-body simulations. This is suitable for generating maps for a CMB imaging survey, since most of the relevant information for CMB data analyses is contained in the two-dimensional projected maps. The advantage of our approach is that DL training in two dimensions is computationally more efficient than that in three dimensions; therefore, for a given computing resource, one can train on a larger footprint at higher-resolution.

In particular, we use DL methods to generate 500 realizations of full-sky millimeter-wave simulations that match the statistics of the S10 simulations. Each of these 500 simulation realizations includes:

- a full-sky simulation at half-arcminute resolution at six different frequencies (30, 90, 148, 219, 277, and 350 GHz)
- the lensed CMB in both temperature and polarization
- extragalactic foreground components appropriately correlated with each other, including:
 - the lensing convergence map (κ)
 - the kinetic Sunyaev-Zel’dovich effect (kSZ)
 - the thermal Sunyaev-Zel’dovich effect (tSZ)
 - the Cosmic Infrared Background (CIB)
 - the radio galaxies (Radio)
- non-Gaussian information that reproduces statistics such as the one-point function, bispectra, and trispectra.

These mmDL simulations are trained using the cosmology and extragalactic model described in S10 [19]. However, the method presented here is general and can be applied to reproduce other two-dimensional simulations with minor modifications. In Section III, we present our procedure to generate the mmDL simulations. The comparison between the statistical properties of the mmDL and the S10 simulations is shown in Section IV. We discuss potential applications and the data products released in Section VI. These mmDL simulations are publicly available in both HEALPix¹ and CAR formats on the Legacy Archive for Microwave Background Data Analysis (LAMDA)² and the National Energy Research Scientific Computing Center (NERSC) cluster³. We also provide the code to produce additional realizations at <https://github.com/dwhan89/cosmikyu>.

II. DL PRIMARY INPUT DATA

We generate our *Primary Input Data* from the S10 simulations, which include κ , kSZ, tSZ, CIB and Radio [19] extragalactic foreground components. The S10 simulations were generated by post-processing the output of a three-dimensional simulation, which has *WMAP5* cosmology [82]. The S10 simulations have correlated foreground components, and have been shown to have non-Gaussian statistics that match observations [83–87]. The S10 simulations and source catalogs are publicly available on NASA’s LAMBDA website (see Footnote 2). The simulations are saved in HEALPix format [88] with $N_{\text{side}} = 8192$, and have units of Jy/sr. The S10 κ map has $N_{\text{side}} = 4096$, and is dimensionless.

We describe our procedure to prepare the mmDL Primary Input data below. Our procedure includes (1) pre-processing the original S10 148 GHz simulations, (2) dividing the pre-processed S10 simulations into small patches, and (3) applying a normalization to the data sets to stabilize and speed up the training of a neural network.

A. Pre-processing of the S10 Simulations

Below we detail the steps in our procedure to pre-process the original HEALPix S10 simulations for the network training.

- Following [26, 87], we scale the flux of CIB sources and the tSZ map by a factor of 0.75 to make the

¹ <http://healpix.sourceforge.net>

² https://lambda.gsfc.nasa.gov/simulation/tb_sim_ov.cfm

³ <https://crd.lbl.gov/departments/computational-science/c3/c3-research/cosmic-microwave-background/cmb-data-analysis-at-nersc/>

S10 simulations more closely match recent measurements from ACT, Planck and SPT [89–93].

- We convert the units of the maps from Jy/sr to μK , except for the κ map which is dimensionless.
- We reproject the κ , tSZ, and kSZ foreground components of the S10 simulations from HEALPix pixelization to plate carrée (CAR) pixelization at half arcminute resolution using the public *pixell*⁴ and *libsharp*⁵ [94] libraries. Since CAR maps intrinsically consist of a grid of rectangular pixels, there is a natural way to extract two-dimensional submaps; this makes it simpler to train DL networks and apply two-dimensional Fourier transform operations. On the other hand, CAR pixels are increasingly over-sampled farther away from the equator (the image is stretched), which we discuss below. The reprojection is done by converting the original HEALPix maps to spherical harmonic a_{lm} up to an ℓ_{max} of 10,000. These a_{lm} are then reprojected onto full-sky CAR pixelized maps.
- For the CIB and Radio foreground components, we place sources directly from the S10 catalogs into full-sky CAR pixelized maps to avoid spectral leakage (“ringing”) around point-like sources when manipulating them in harmonic space. We apply a flux-cut of 7 mJy at 148 GHz to both CIB and Radio CAR maps in order to match current and future CMB analyses that typically detect and remove bright sources.
- Since CAR maps require the same number of pixels at each declination, the image appears stretched at the poles. To prevent the network from learning non-physical features induced by this stretching, we use only the region of the simulations near the equator to train the network (from -10° to 10° in declination). However, in order to make use of all the sky area available in the S10 simulations, we rotate the full S10 map many times, redefining the equator of the map each time. In particular, we rotate each full-sky map by 15 degrees in ψ and 20 degrees in θ , where ψ and θ are the first two Euler angles. The successive rotations in ψ and θ described above generate 25 unique cylindrical strips around each new equator, which strips in CAR projection have minimal image distortion compared to the sphere. For κ , tSZ, and kSZ, the rotations are done in spherical harmonic space. For CIB and Radio, we simply remap source positions for given rotation angles.

B. Division of Full Sky Simulations into Smaller Patches

To train a network, we need many training samples. Since we start with a single realization of the full sky S10 simulation, we increase the number of samples by dividing this simulation into smaller overlapping patches. Below we detail our steps.

- We randomly pull out patches of 128×128 pixels (roughly $1^\circ \times 1^\circ$ patches) between the declination of -10° and 10° . We cut out patches in this way from our 25 cylindrical strips (described above) for κ , tSZ, kSZ, CIB and Radio maps. Thus each patch has five foreground components totaling an array of size $5 \times 128 \times 128$. We repeat this step until we collect 30,000 validation samples (to pick initial parameters for the training), 200,000 training samples, and 30,000 test samples (to check the performance of the training). We ensure that we have an equal number of samples from the 25 rotated maps. We call the resulting data set the Primary Input Data. Note that these samples will necessarily have overlapping regions among them. We discuss our method to overcome this issue in Section III.
- During the network training, we randomly flip a sample either vertically or horizontally, or both as a natural way to augment the training data set.

C. Normalization to Speed up and Stabilize the Training

We need to normalize the samples to stabilize and to speed up the neural network training. Since we have multiple foreground components in a single sample, the normalization needs to account for two issues. First, the pixel intensity Probability Density Function (pixel PDF) of different foreground components follows different distributions. The pixel PDFs of tSZ, CIB, and Radio roughly follow a Poisson distribution, whereas those of κ and kSZ are better described by a Gaussian distribution. Second, these pixel PDFs have varying dynamic ranges. For example, all the pixel values of the κ map fall between -1.5 and 1.5; on the other hand, the Radio map has pixel values ranging from 0 to 852 μK . Ideally, we would like the pixel PDFs to follow a standard Gaussian distribution (i.e. with unit variance) for faster and more stable network training. To address the shape of the pixel PDF, we first perform a pixel-wise natural log scaling adopted from [74], to the tSZ, CIB, and Radio maps. In particular, we replace the value of each pixel, given by v , with

$$f(v) = \text{sgn}(v) \ln \left[\frac{\text{abs}(v)}{\sigma_v} + 1 \right] \quad (1)$$

where $\text{sgn}(v)$ is sign function, and σ_v the standard deviation of the pixel PDF of the input map. After this normalization, the pixel PDFs of the tSZ, CIB, and Radio maps are closer to a Gaussian distribution.

⁴ <https://github.com/simonsobs/pixell>

⁵ <https://github.com/Libsharp/libsharp>

To mitigate the dynamic range issue, we then standardize the maps using the standard score method (i.e. subtract the mean of all the pixels from each pixel value and then divide that by the variance of the map).⁶

$$s(v') = (v' - \mu_{v'}) / \sigma_{v'} \quad (2)$$

Here, v' is the pixel value, and $\mu_{v'}$ and $\sigma_{v'}$ are the mean and standard deviation of the pixel values respectively.⁷ Note that each foreground component is normalized independently from the others. The parameter values used in the normalization procedure are shown in Appendix A. Note that each of these normalization operations has a well-defined inverse operation. When we generate the final products, we undo these normalizations prior to the interpolation procedures discussed in Section III D.

III. METHOD

Generative Adversarial Networks (GANs) are a class of networks that will generate new data matching the statistics of the training data set [95]. Fundamental to a GAN is a set of two neural networks, called the generator and discriminator. A generator makes new data, while a discriminator tries to distinguish between the real data and the generated data. By iteratively training a generator and a discriminator, a GAN reaches an equilibrium where the discriminator can no longer distinguish new data from training data. For more discussion about GANs, we refer the reader to [96, 97].

GAN models can only generate random samples limited to the same footprint as the training data. Since our Primary Input Data consists of $1^\circ \times 1^\circ$ patches, the random samples output by the GAN do not have the full information about modes on scales larger than 1 degree (roughly $\ell \leq 200$). This presents a challenge in stitching together the small patches to make a full-sky simulation with the correct large-scale fluctuations. Furthermore, there is no straightforward way to tile these random patches together to make a full-sky map without having discontinuities at the tile edges.

Instead, we train a conditional GAN network to predict the four extragalactic foreground components (tSZ, kSZ, CIB, and Radio) from a κ map. Once this network training step is completed, we generate a full-sky Gaussian κ map, and convert it to the other four full-sky extragalactic components tile by tile. Since the input Gaussian

κ maps is continuous, the resulting extragalactic foreground maps are also continuous. This image prediction approach also solves another major concern about using DL to generate new samples, which is mode collapse. Mode collapse is when the network predicts only a subset of all possible outcomes, instead of a continuous distribution. However, with our approach, since we provide a unique input full-sky κ map to the network for each full-sky realization the network outputs, we can ensure that the output maps are also unique.

A challenge of this predictive approach, which uses a conditional GAN model, is that it has a large number of parameters that need to be trained; thus this conditional GAN requires a much larger training data set than simple GANs. Since our Primary Input Data is relatively small and has over-lapping patches (see Section II for details), we first use the Primary Input Data to train a simple GAN network (which is easier to train) to make a large sample of non over-lapping patches of κ , tSZ, kSZ, CIB and Radio maps (Intermediate Product 1 in Figure 1). This larger data set is then used to train a conditional GAN (Training Step 2 in Figure 1). After Training Step 2 is done, the network can now predict extragalactic foregrounds given a κ map. We input a unique full-sky Gaussian κ map into the network to generate the other extragalactic foreground components (tSZ, kSZ, CIB and Radio) tile by tile (Intermediate Product 2 in Figure 1). The extragalactic components predicted from the Gaussian κ map will be missing some non-Gaussian information due to the missing non-Gaussian information in the κ map; thus, we introduce a simpler version of the conditional GAN that can restore this missing non-Gaussian information, which we train on the Primary Input Data plus Intermediate Product 2. Once this training is done, the network is equipped to go from unique input full-sky Gaussian κ maps and unlensed CMB T, Q , and U maps to a full-sky simulation including non-Gaussian extragalactic foregrounds, a non-Gaussian κ map, and the corresponding lensed CMB T, Q , and U maps (T, Q , and U are the temperature and two polarization fields of the CMB). The schematic of the overall procedure is shown in Figure 1.

Since the extragalactic maps are statistically invariant under translation, a natural choice is to use a deep convolutional generative adversarial network (DCGAN) since it has convolution kernels that are also invariant under translation. In this work, we use three variations of the DCGAN: 1) Deep Convolutional Wasserstein GAN with gradient penalty (DCWGAN-GP) to make Intermediate Product 1 (this is Training Step 1 discussed in Section III A), 2) Pix2Pix GAN (PIXGAN) to make to make Intermediate Product 2 (this is Training Step 2 discussed in Section III B), and 3) Variational Auto-Encoder GAN (VAEGAN) to augment the missing non-Gaussian information (this is Training Step 3 discussed in Section III C).

⁶ Another popular approach is to scale an image such that all the pixel values lie between -1 and 1. During our network training, we found that using the the standard score method leads to more stable training compared to this rescaling.

⁷ Note that we distinguish v and v' here because the pixel values of the tSZ, CIB, and Radio maps are first scaled by Equation 1 (i.e. $p' = f(p)$), whereas $p' = p$ for the κ and kSZ maps.

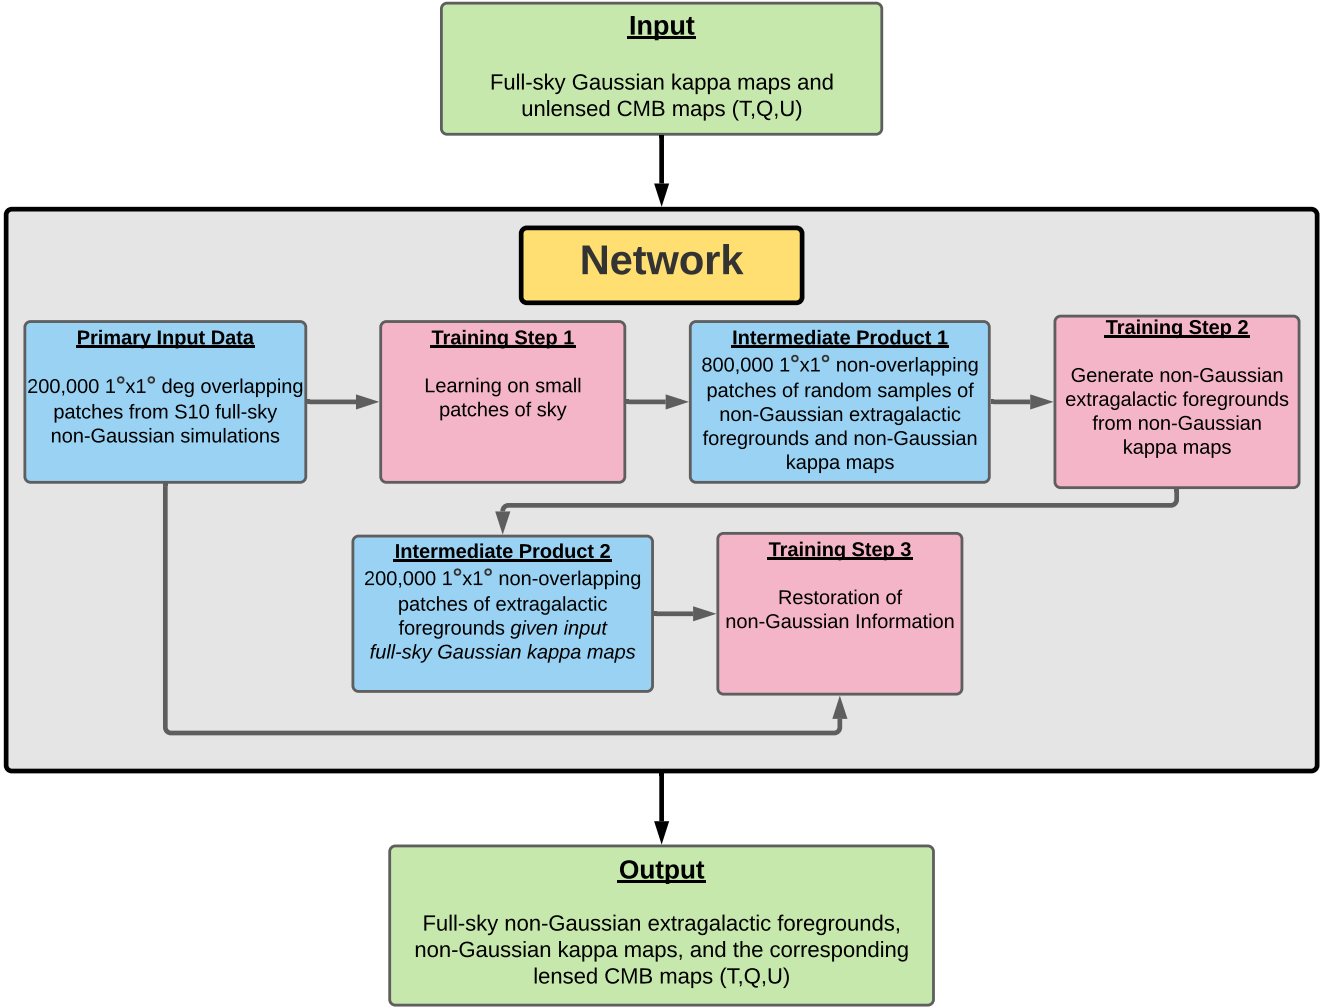


Figure 1: Shown is a schematic of the overall procedure to train the network and generate the output maps from the input data. The network is represented by the shaded gray box in the center. The mmDL procedure starts with the Primary Input Data consisting of 200,000 $1^\circ \times 1^\circ$ overlapping patches cut out from the original S10 simulations. The pink boxes represent the three major training steps described in Section III. The blue boxes represent inputs into each training step. The dark grey arrows indicate the input and output of each step. Once the network training is completed, we feed full-sky Gaussian kappa maps and unlensed CMB maps (T, Q, U) into the network (top green box); given these inputs, the network generates the final output products (bottom green box), which are full-sky millimeter-wave simulations including lensed $T, Q,$ and U maps, non-Gaussian kappa maps, and non-Gaussian extragalactic foregrounds correlated with the kappa map and each other.

A. Training Step 1: Learning on Small Patches with Wasserstein GAN with Gradient Penalty (DCWGAN-GP)

The DCWGAN-GP model we use is summarized by the generator depicted in Figure 2 and the discriminator described in Table I. We take the *CosmoGAN* network architecture used in [72] as a starting point of our DCWGAN-GP network because it has been demonstrated to reproduce non-Gaussian statistics for a one-component map. From this starting point, we make the following adjustments. First, we switch the simple GAN loss function used in *CosmoGAN* with the Wasserstein loss function with gradient penalty introduced in [98];

the loss function is the statistic that all GANs optimize. We find that this change of the loss-function substantially improves the stability of our network against mode-collapsing (i.e. against the network missing critical features). To be consistent with the standard Wasserstein GAN (WGAN) architecture, we remove the batch normalization layers (i.e. the intermediate normalization steps within the GAN) from the discriminator to stabilize the loss function gradient, and replace the sigmoid function activation layer with a linear function activation layer, following [98]. In addition, we increase the depth of the convolution layers in both generators and discriminators from four to five, and we increase

the latent-vector length (i.e. array size of input random numbers) from the original 64 to 256 in order to have a more expressive network. We also change the last activation layer in the generator model from $\tanh x$ to $a * \tanh b/a * x$, which we call a Scaled Tanh. Here a is chosen to be 15 to accommodate the dynamic range of the normalized pixel values we obtained using Equation 1, while b is chosen to be 2 so that the gradient around the boundary values (i.e. -15 and 15) do not vanish too rapidly; we find the form of the Scaled Tanh function and its parameters by checking the training statistics when using the validation set of patches (see Section II B). We also add a custom linear layer that linearly re-scales each extragalactic component independently so that the network has more freedom to re-scale the components relative to each other. We call this custom layer the “Linear Channel” (LC). We find that using this LC helps the network to converge faster. We reduce the two-dimensional convolution layer kernel size (i.e. the width and the height of a convolution filter) from 5×5 , used in *CosmoGAN*, to 4×4 to account for the difference in the input image size from 256×256 pixels to 128×128 pixels (the latter is for $1^\circ \times 1^\circ$ patches at 0.5 arcmin resolution). Lastly, we replace the Rectified Linear Unit (ReLU) layers in the generator with the Leaky Rectified Linear Unit (Leaky ReLU) with $\alpha = 0.2$ to achieve better performance, as suggested by [99]. (For more discussion about various types of activation functions, we refer to [100]).

We train the DCWGAN-GP network with the Primary Input Data using up to 100 epochs (i.e. the number of cycles through the entire training data set used by the network during training). Following [98], we adopt a gradient penalty coefficient (λ) of 10 and the number of discriminator iterations per generator iteration (n_{critic}) of 5. For gradient descent, we use the Adam optimizer [101] with a learning rate (lr) of 10^{-4} and $(\beta_1, \beta_2) = (0.5, 0.9)$. Throughout this training, we check not only the loss function values, but also the visual images, Minkowski functionals, power spectra, cross spectra, and pixel PDFs of the network outputs; we will discuss each metric in Section IV. If the network shows signs of over-fitting or mode-collapsing, we restart the training either from the beginning or after reducing the learning rate by half (i.e. effectively reducing the step size). We terminate the training cycle once all the statistics considered are satisfactorily reproduced. We use the trained DCWGAN-GP network to produce *Intermediate Product 1* ($800,000 1^\circ \times 1^\circ$ patches of non-Gaussian κ maps and correlated non-Gaussian foregrounds) which we use in Training Step 2. Note that the samples in *Intermediate Product 1* are no longer overlapping and are statistically independent of each other.

Layer	Activation Function	Output Shape of Tensor	# of Trainable Parameters
Input Map	N/A	$N \times 128 \times 128$	N/A
Conv 4x4	LReLU	$64 \times 64 \times 64$	5,184
Conv 4x4	LReLU	$128 \times 32 \times 32$	131,200
Conv 4x4	LReLU	$256 \times 16 \times 16$	524,544
Conv 4x4	LReLU	$512 \times 8 \times 8$	2,097,664
Conv 4x4	LReLU	$1024 \times 4 \times 4$	8,389,632
Linear	N/A	1	16,385
Total Trainable Parameters			11,164,609

Table I: Summary of the discriminator model described in Section III A. We use the same discriminator model throughout all three training steps described in Section III. The first and second columns list the layers and the corresponding activation functions. The third and the fourth columns give the shapes of the output arrays and the number of trainable parameters in each layer. Here, N is the number of input components, i.e. $N = 5$ for DCWGAN-GP and PIXGAN and $N = 10$ for VAEGAN. Conv 4×4 indicates a two-dimensional convolution layer with a 4×4 kernel, and LReLU stands for a leaky rectified linear unit activation with $\alpha = 0.2$ (see Section III A for details). Note that Conv 4×4 downsamples an image by a factor of two, and increases the number of features to consider.

B. Training Step 2: Generate Extragalactic Foregrounds from Kappa Map with Pix2Pix GAN (PIXGAN)

The PIXGAN model we use is summarized by the generator depicted in Figure 3 and the discriminator described in Table I. A PIXGAN can convert input images to other images by implementing a U-NET generator [102, 103]. We use the PIXGAN network to convert a given κ map to the other four non-Gaussian extragalactic foregrounds consisting of tSZ, kSZ, CIB, and Radio components. We start with the original PIXGAN architecture presented in [103], and make the following modifications. We first replace the \tanh activation layer in the generator with a sequence of a Linear Channel and a Scaled Tanh; the latter two are defined in III A. Since we terminate the training well before the network shows signs of over-fitting, we also remove the drop-out layers from the generator to speed up the training. Lastly, instead of using the PatchGAN discriminator as in [103], we use the same DCWGAN-GP discriminator model described in Section III A, and summarized in Table I.

We use *Intermediate Product 1* ($800,000 1^\circ \times 1^\circ$ non-overlapping patches of non-Gaussian κ , tSZ, kSZ, CIB and Radio maps) generated by Training Step 1 to train the PIXGAN network.⁸ We use the same optimizer setup as in Training Step 1. During the training process, we

⁸ We still use the test and the validation data sets from our Primary



Figure 2: Shown is a diagram of the DCWGAN-GP generator model described in Section III A and indicated in the schematic of Figure 1 as Training Step 1. The input to this generator is a latent-vector consisting of 256 random numbers drawn from a Gaussian distribution with mean equal to zero and variance equal to one ($\mu = 0, \sigma = 1$). The output is one set of $1^\circ \times 1^\circ$ maps consisting of the five foreground components. The direction of data flow is indicated by arrows. A solid arrow is a fully connected Linear layer. A short-dashed arrow, labeled TC-BN-LRU, represents a sequence of 1) a 4×4 two-dimensional transposed convolution, 2) a batch normalization, and 3) a leaky rectified linear unit activation function with $\alpha = 0.2$. Lastly, a long-dashed arrow, labeled TC-LC-STanh, is a sequence of 1) a 4×4 two-dimensional transposed convolution, 2) a Linear Channel, and 3) a Scaled Tanh activation function. The Linear Channel and Scaled Tanh are defined in Section III A. Note that a transposed convolution upsamples an image by a factor of two.

find that the PIXGAN network tends to overestimate the number of massive tSZ clusters, which results in an excess of tSZ power at large scales. To mitigate this effect, we add a linear loss term to the DCWGAN-GP generator loss function evaluation defined as the second term in the right side of the equation below.

$$\mathcal{L}_G^{\text{PIXGAN}} = \mathcal{L}_G^{\text{DCWGAN-GP}} + \gamma \sum_i |\sigma^2(S_{\text{network}}^i) - \sigma^2(S_{\text{input}}^i)| \quad (3)$$

where $\mathcal{L}_G^{\text{DCWGAN-GP}}$ is the standard DCWGAN-GP generator loss function defined in [98], and $\sigma^2(S^i)$ is the total pixel variance of the i -th foreground component in a sample. Since the sum of power spectra values across multipoles is the equal to the total pixel variance of the map (a corollary of Parseval's theorem), it follows that penalizing the difference in the total pixel variance is equivalent to penalizing the mismatch in the summed power spectra. We train the model with $\gamma = 100$, and find that it helps to regulate the excess power in the tSZ map.

We train the PIXGAN network using up to 10 epochs⁹. We validate the network every 200,000 training samples (four times per epoch) with the updated loss function in Equation 3, checking the same metrics as in Section III A. We intentionally terminate early the training of the PIXGAN network for two reasons. First, once the PIXGAN network starts to overfit, the quality of the output degrades quickly. We find it more conservative

to under-train the PIXGAN network rather than to risk over-fitting. Second, we find that the VAEGAN network described in the next training step (Training Step 3) can absorb some imperfection in the simulations and correct for them. Therefore, the output from the PIXGAN network can have minor defects without affecting the quality of the final products. We terminate the training once we see signs of over-fitting; these signs are 1) minimal improvement of the loss function over the training samples, 2) large fluctuations of the other metrics with respect to the S10 simulations, and 3) images missing significant features. At this point, we choose the snapshot of the network that reproduces the summary statistics the best. In particular, for the power spectra, we check that the PIXGAN network can reproduce the shape of each foreground component power spectra up to a constant multiplication factor.¹⁰ When this network snapshot reproduces the shape of the power spectra to within 10% for $\ell \in (2000, 8000)$, and the cross spectra and Minkowski functionals also show similar agreement, we consider the training complete. Otherwise, we repeat the training procedure until we reach this criteria. After this training step is done, we feed in several full-sky Gaussian κ maps through the trained PIXGAN network, producing 200,000 $1^\circ \times 1^\circ$ non-overlapping patches of extragalactic foregrounds consisting of tSZ, kSZ, CIB, and Radio components; we call this *Intermediate Product 2*. The full-sky Gaussian κ maps are generated using *pixell*; we use the power spectra of the S10 kappa map as the input spectra to these Gaussian kappa simulations.

Input Data to fine-tune the network parameters and to validate the network outputs.

⁹ Note that each epoch here has four times more samples than the epochs in Section III A

¹⁰ Note that the next network (Training Step 3) can re-scale each foreground component by a constant factor as described in Section III C.




Figure 3: Shown is a diagram of the UNET generator model described in Section III B and indicated in the schematic of Figure 1 as Training Step 2. The generator takes as input a $1^\circ \times 1^\circ \kappa$ map, and it outputs a set of $1^\circ \times 1^\circ$ extragalactic foreground maps consisting of tSZ, kSZ, CIB, and Radio components. The direction of the data flow is indicated by the arrows. A solid arrow, labeled C-BN-LRU, represents a sequence of 1) a 4×4 two-dimensional convolution, 2) a batch normalization, and 3) a leaky rectified linear unit activation function. A short-dashed arrow, labeled TC-BN-RU, represents a sequence of 1) a 4×4 two-dimensional transposed convolution, 2) a batch normalization, and 3) a rectified linear unit activation function. A two-dimensional convolution downgrades an image by a factor of two, while a transposed convolution upsamples an image by the same factor. A dotted arrow is a concatenation operation, which concatenates the output of an originating box to the output of a destination box. A long-dashed arrow, labeled LC + STanh, on the top right represents a sequence of a Linear Channel and a Scaled Tanh activation function. Linear Channel and Scaled Tanh are defined in Section III A. Deviating from our notation, we do not apply a batch normalization to the first and the last convolution layers, indicated by a star over a line.

C. Training Step 3: Restoration of Non-Gaussian Information with Variational Autoencoder GAN (VAEGAN)

The VAEGAN model we use is summarized by the generator depicted in Figure 4 and the discriminator described in Table I. Neural networks can be used to improve existing simulations, either by improving the resolution or by adding missing small-scale information [44, 104]. In this work, we adopt the Variational Autoencoder GAN (VAEGAN) in order to add back missing non-Gaussian information [105]. Our VAEGAN generator architecture is similar to the U-NET generator architecture discussed in Section III B except for the following. Unlike the PIXGAN, the VAEGAN generator does not have skip-connections for each layer (i.e. there are no array concatenations for each layer like as shown by the dotted lines in Figure 3), except that we add back input maps from Intermediate Product 2 at the very end (see dotted line in Figure 4) as done in [80]. This addition at the end allows the VAEGAN network to focus just on

the missing non-Gaussian information (i.e. to focus on the residual difference between the network product of Training Step 3 and Intermediate Product 2). Another difference between the PIXGAN and VAEGAN generators are that the latter has fewer trainable parameters, which makes it easier to train; thus we only need 200,000 samples, as opposed to 800,000, to train it. Regarding the discriminator, we use the same DCWGAN-GP discriminator (i.e. the same Wasserstein loss function with the gradient penalty) discussed in Section III A, and summarized in Table I.

We use as input both the Primary Input Data and Intermediate Product 2 to train the VAEGAN network. The output of the VAEGAN network is the restoration of the missing non-Gaussian information in Intermediate Product 2. The VAEGAN network does this by comparing the original Intermediate Product 2 to the S10 simulations and learning the difference. To do this we use the same loss function and training setup as in Training Step 2. Since the output from Training Step 3 is our final product, except for a few corrections described later




Figure 4: Shown is a diagram of the VAEAN generator model described in Section III C and indicated in the schematic of Figure 1 as Training Step 3. The input to this generator is a set of $1^\circ \times 1^\circ$ Gaussian κ maps and their corresponding extragalactic foreground maps from the PIXGAN network in Figure 3 (i.e. Intermediate Product 2). The output are a set of non-Gaussian κ maps and their corresponding non-Gaussian extragalactic foreground maps. The direction of the data flow is indicated by arrows. A solid arrow, labeled C-BN-LRU, represents a sequence of 1) a 4×4 two-dimensional convolution, 2) a batch normalization, and 3) a leaky rectified linear unit activation function. A short-dashed arrow, labeled TC-BN-RU, represents a sequence of 1) a 4×4 two-dimensional transposed convolution, 2) a batch normalization, and 3) a rectified linear unit activation function. The two-dimensional convolution downgrades the image by a factor of two, while the transposed convolution upsamples the image by the same factor. Shown as a dotted arrow is a concatenation operation between the input Intermediate Product 2 and the VAEAN network output, which focuses the network on the residual difference between the two; note that we apply the inverse of Scaled Tanh to this input prior to the concatenation. A long-dashed arrow, labeled LC + STanh, represents a sequence of a Linear Channel and a Scaled Tanh activation function; they are described in Section III A. Deviating from our notation, we do not apply a batch normalization to the first and the last convolutions layers, indicated by a star below the arrow.

in Section III D, we further fine-tune the Training Step 3 training procedure to achieve optimal results. In particular, we revisit the issue of overestimation of massive tSZ clusters that we described in Section III B. In order to further reduce the chance of overestimation, we make the following adjustments to our training procedure: 1) first, we sort the training samples by the total variance of the tSZ component, 2) then we start training with the bottom eighty percent of samples in terms of the total tSZ variance; 3) Once the training is stabilized, we restart the training with the full set of training samples. We closely monitor the tSZ statistics along with the loss function and other summary statistics. We terminate the training when all the statistics considered are within 5% for $\ell \in (2000, 8000)$.

D. Post-processing the Network Simulations

As discussed in Section III, we convert a full-sky Gaussian κ map to other foreground components tile by tile, where each tile is a $1^\circ \times 1^\circ$ flat patch of sky. Here, we describe our method to divide full-sky κ maps into tiles and to reproject the processed tiles back on to a full-sky map. When we sample a $1^\circ \times 1^\circ$ patch from a full-sky κ map, we rotate that patch to the celestial equator and then cut out a tile of size 128×128 pixels with a 20-pixel overlap region between each tile as shown in Figure 15 in Appendix B. This is similar to how we cut the original S10 simulations as discussed in Section II B. This results in a grid of tiles with 200 tiles along a longitudinal di-

rection and 400 tiles along a latitude direction.¹¹ When we reproject the tiles back on to a full-sky map, we interpolate them back to the correct location on the sphere according to their coordinates. Here, we take advantage of the fact that CAR pixels in a full-sky map get stretched only along the latitude direction. As a result, the pixels in the longitudinal direction can be remapped without any interpolation. This effectively reduces the dimensions of the interpolation from two to one, which lowers the overall computational cost.

We use two different interpolation methods depending on which foreground component we would like to reproject. We use a linear interpolation to reproject the κ , tSZ, and kSZ foreground components because they are continuous fields.¹² To reproject the CIB and Radio galaxies, we use the nearest neighbor method (i.e. each CAR pixel in the full-sky map obtains a value from the nearest pixel in the flat-sky tile). We use this method because linear interpolation can “blur” sharp variations in amplitude, which is not ideal for interpolating discrete point sources. We remove any superfluous neg-

¹¹ Around the poles of the sphere (i.e. the first and the last longitudinal rows of the grid), there is not enough physical area to sample a $1^\circ \times 1^\circ$ tile. Hence, we take the second and the second-to-the-last rows of the grid, mirror them longitudinally, and replace the first and the last rows of the grid with these mirrored rows. This ensures that we have continuous physical tiles around the poles at the cost of them not being unique. We find that this pole correction does not affect the overall statistics of the network simulations.

¹² We tried both linear interpolation and popular cubic spline methods, and found no noticeable difference in the resulting summary statistics. Thus, we chose to use the linear interpolation method which is more computationally efficient.

ative flux sources before reprojecting point source tiles back onto the full-sky. Since we initially sample the full-sky κ map with overlapping regions, one pixel in the full-sky map can be mapped onto more than one flat-sky tile. Therefore, we precompute the number of mapped tiles for each full-sky map pixel and construct an effective “hit-counts” map. If more than one tile corresponds to a single full-sky pixel, then we add the tile pixel values and divide by the hit-counts map to obtain an average. In addition, we apodize each tile with a cosine apodization mask prior to the full-sky reprojection as done in [44] to suppress any discontinuity around the edges; this is done by weighting the overlapping regions by the apodization mask to downweight pixels closer to the edges of a given tile.¹³ Lastly, following Section 4.3 in [106], we compute two-dimensional transfer functions to correct for the reprojection effects for the κ , tSZ, and kSZ maps. Note that we do not apply two-dimensional transfer functions to the CIB and Radio maps to avoid spectral leakage when shifting to Fourier space.

We find that the raw full-sky outputs from the network can reproduce the statistics of the S10 simulations overall quite well. However, we do find some small mismatches in the summary statistics between the network and the S10 simulations; for example, for the power spectra in the range $\ell \in (2000, 8000)$, the agreement is better than 3%, 10%, 0.5%, 17% and 0.3% for κ , kSZ, tSZ, CIB and Radio respectively, however, it is not perfect. Therefore, we make a few corrections to the network simulations to match the S10 simulations more closely. Here, we describe the corrections we apply. First, to correct for the small mismatch in the power spectra, we apply one-dimensional transfer functions to the κ , kSZ, and tSZ maps. These transfer functions are computed as the square root of the ratio of the average of 10 network power spectra to the S10 power spectra. We compute the transfer functions at each multipole, and apply to these transfer functions a one-dimensional Gaussian filter with $\sigma_\ell = 10$ to smooth them out. We then convert these to isotropic two-dimensional transfer functions, and apply them to the spherical harmonics of the full-sky map.

We find larger deviations in the power spectrum between the network and the S10 simulations than specified above at larger ($\ell \leq 2000$) and smaller scales ($\ell \geq 8000$). The large scale deviation is due to the limited size of the sample patches as discussed in Section III. For the small scales, we find a roll-off in the power spectra starting at $\ell = 8000$. We find that this roll-off starts at larger scales ($\ell \leq 8000$) if we train the network with lower resolution simulations (for example, 1 arcminute resolution), which indicates that the roll-off is related to the resolu-

tion of the initial training samples.

A second correction we apply is to adjust the CIB and Radio maps as follows. We re-scale the CIB map by a factor of 1.1 to match the S10 power spectra for $\ell \in (4000, 8000)$. We also re-scale the network CIB and Radio source fluxes to match the re-scaled S10 source counts.¹⁴ After this re-scaling, we still find that the network overestimates the number counts of CIB sources with flux values greater than 1 mJy; thus we reduce the flux of those sources by replacing the original flux value, S , with $S^{0.63}$. For Radio sources, we find that replacing the original flux by $S^{(\lambda)} = ((S + 1)^\lambda - 1)/\lambda$ with $\lambda = 1.25$ [107] improves the match at the high-flux end.

E. Generating Lensed CMB and Multi-frequency Maps

In order to generate lensed CMB (T, Q, U) maps, we first generate full sky unlensed CMB (T, Q, U) maps as Gaussian Random Fields on a CAR full-sky grid using the CAMB theory spectra used in S10.¹⁵ Then, these unlensed CMB maps are lensed with a network full-sky κ map using the lensing method described in [4]. As in [4], the lensing operation is performed at 1 arcminute resolution and agrees with the theory to better than few percent up to $\ell = 10,000$.

Up to this point, all the training and corrections are done with the 148 GHz network simulations. Here, we discuss how we simulate the other frequency maps. We assume to leading order that the lensed CMB and kSZ perfectly follow a black-body spectrum (i.e. they do not vary with frequency in units of differential CMB temperature). The κ map is frequency independent. In order to simulate the tSZ map at different frequencies, we use the analytic calculations in [108, 109] to compute the frequency dependent scaling factor, $a(\nu)$, and scale the 148 GHz tSZ map following $M^{\text{tSZ}}(\nu) = a(\nu)/a(148) * M^{\text{tSZ}}(148 \text{ GHz})$, where M^{tSZ} is the real-space tSZ map. For CIB and Radio sources, we find the mean spectral indices between each frequency and the reference frequency of 148 GHz, directly from the S10 catalogues; this results in the spectral indices shown in Table III of Appendix C. For each source in the 148 GHz map, we then randomly sample a spectral index from a Gaussian distribution with the mean and the standard deviation given in Table III, and scale its flux accordingly. In this way, we generate full-sky maps of the combined lensed CMB, kSZ, tSZ, CIB, and Radio signals at 30, 90, 148, 219, 277, and 350 GHz.

¹³ We set the pixel values at the outermost edges of the apodization mask to the values of pixels one pixel away inside the map to avoid losing information.

¹⁴ Note that the original S10 CIB source fluxes have been re-scaled by a factor of 0.75 as discussed in Section II A.

¹⁵ The CAMB theory spectra used in this work can be found under the resources directory at <https://github.com/dwhan89/cosmikyu>.

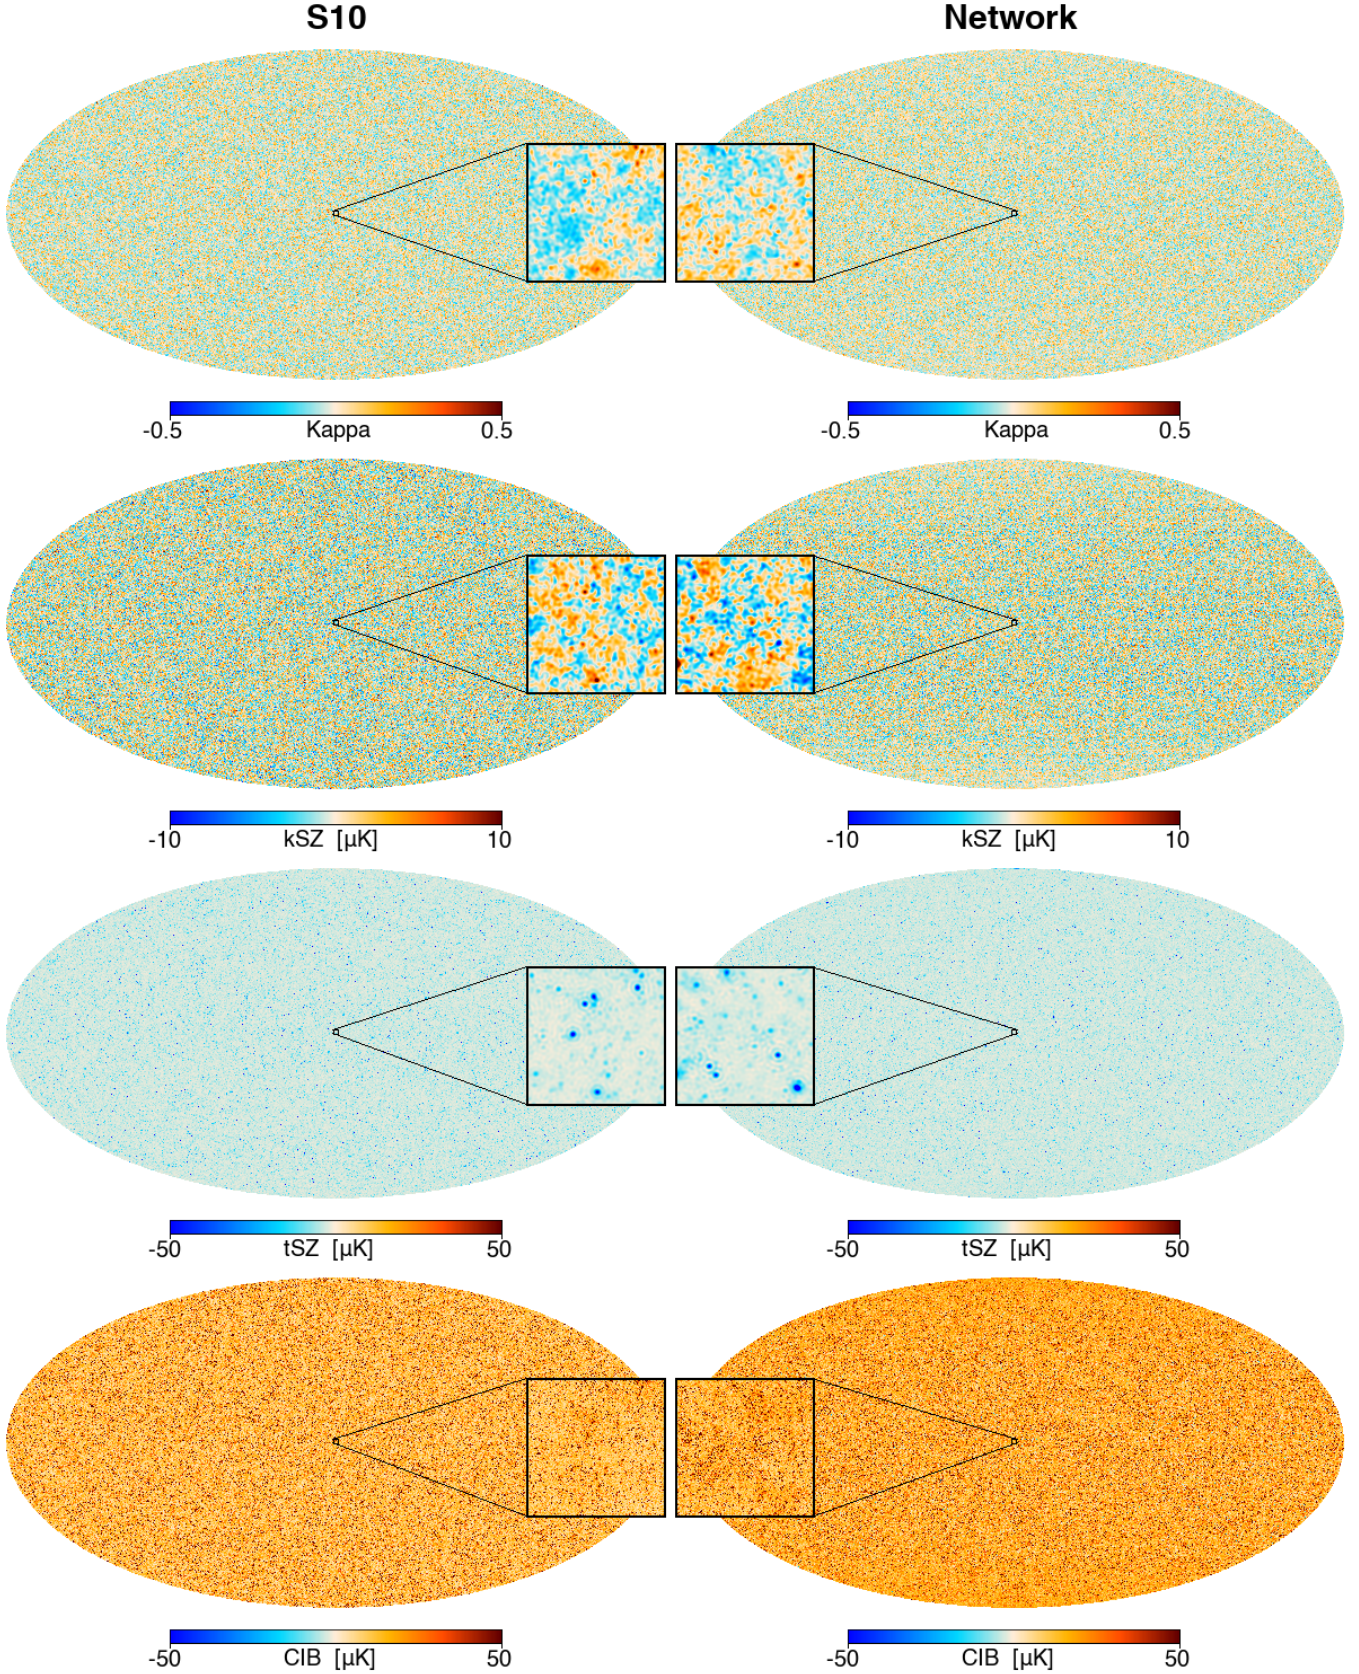


Figure 5: From top to bottom, shown are the lensing convergence (κ), the kinetic Sunyaev-Zel'dovich effect (kSZ), the thermal Sunyaev-Zel'dovich effect (tSZ), and the Cosmic Infrared Background (CIB) maps at 148 GHz from the S10 simulations (left column) and from the network (right column). A flux cut of 7 mJy at 148 GHz is applied to the CIB maps. Full-sky maps are shown in the Mollweide projection in the background, while center panels show zoom-ins of $1^\circ \times 1^\circ$ patches. All maps have the units of μK , except for the κ map which is dimensionless. Note that the S10 simulations are unique for only one octant of the sky; on the other hand, the network simulations do not have any repeated tiles.

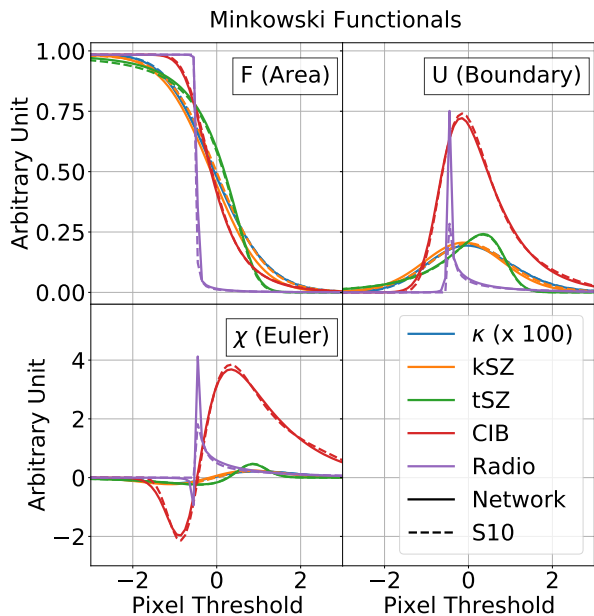


Figure 6: Shown are the Minkowski functionals of the extragalactic foreground components at 148 GHz. The solid curves are calculated using the average of $100 1^\circ \times 1^\circ$ tiles output by the VAEGAN described in Section III C. The dashed curves are calculated using the full-sky S10 simulation described in Section II. Both the S10 and the network simulations are normalized following the procedure detailed in Section II C. To compute the Minkowski functionals, a binary map is first constructed with a value of one above a certain pixel threshold, and zero otherwise. Then, this binary map is used to compute the area above the threshold (F), its boundary (U), and the dimensionless curvature known as the Euler characteristic (χ). These three functionals, give an estimate of the non-Gaussianity of each component [110]. We find that the network can reproduce well the Minkowski functionals calculated using the S10 simulations. This indicates that the network learned the non-Gaussian information in S10.

IV. RESULTS

In this section, we study the statistical properties of the mmDL simulations generated by our network. Unless otherwise noted, all statistics are computed using full-sky maps after applying a flux-cut of 7 mJy at 148 GHz to Radio and CIB maps, and applying the corrections discussed in Section III D. We first visually compare the full-sky mmDL and S10 simulations in Figure 5. Here we show the full-sky maps in the Mollweide projection, as well as zoom-ins of $1^\circ \times 1^\circ$ patches. Visually, the mmDL simulations and the S10 simulations are indistinguishable.

A. Pixel-level Summary Statistics

There are a number of pixel-level summary statistics used in the literature. Such examples are 1) Minkowski

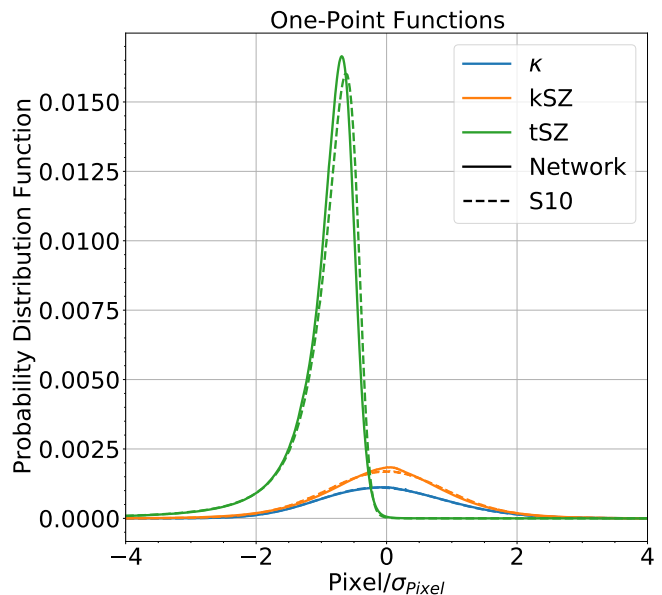


Figure 7: Shown are one-point probability distribution functions of the extragalactic foregrounds at 148 GHz. Solid curves are calculated from a random sample of a full-sky simulation from the network, and dashed curves are calculated using the full-sky S10 simulation. A flux-cut of 7 mJy is applied to both CIB maps. Each map is divided by the standard deviation of its pixel values in order to have unit variance. We find good agreement between the network curves and the S10 curves.

functionals [110], 2) one-point probability distributions, and 3) source counts. For two dimensional images, Minkowski functionals measure the area (F), the contour length (U), and the dimensionless curvature known as the Euler characteristic (χ) as a function of varying pixel threshold value. Combined together these three functionals give an estimate of the topology of an image, which is sensitive to the presence of non-Gaussianity. We compute Minkowski functionals using the *minkfncts2d* software, which is publicly available at <https://github.com/moutazhaq/minkfncts2d>. Figure 6 shows as solid curves the functionals computed using the average of $100 1^\circ \times 1^\circ$ tiles output by the VAEGAN described in Section III C.¹⁶ The dashed curves are calculated using the full-sky S10 simulation described in Section II. Here, we compute the functionals from the network without the corrections discussed in Section III D. Hence, the Minkowski functionals shown in Figure 6 give an estimate of how well the network performs before applying any corrections. Overall, we find very good agreement between the functionals computed from the network and the S10 simulations. The only no-

¹⁶ We recomputed the Minkowski functionals on a 20 degree wide band around the celestial equator after making the corrections described in Section III D; we did not find any significant change, except a slight shift of the CIB curves as compared to Figure 6.




Figure 8: Shown is a histogram of the source number counts per square degree per flux bin for CIB and Radio maps. The solid curves are generated using a random simulation from the network, while the dashed curves are generated using the S10 simulation. A flux cut of 7 mJy is applied to both CIB and Radio maps for both network and S10 simulations. We find good overall agreement.

ticeable mismatch is in the U and χ functionals of the Radio sources, which indicates a mismatch in the real-space shape of Radio point sources between the mmDL and S10 simulations. The S10 point sources are “true” point sources by construction (i.e. each source falls into a single pixel). On the other hand, the network point sources are constructed as the summation of multiple convolution kernels, hence they can potentially extend beyond a single pixel. In practice, the simulations will be convolved with a beam with a full width at half maximum (FWHM) larger than 0.5 arcminute. So this minor mismatch in the Radio source shape will have a negligible effect in actual analyses.

Another pixel-level summary statistic we consider is the one-point probability distribution function (one-point PDF) shown in Figure 7; this one-point PDF is the histogram of pixel values. Before computing this histogram, we divide the pixels in each component map by the standard deviation of pixel values in that map in order to have unit variance. One-point PDFs have been studied previously in the context of CMB weak lensing and tSZ analyses (e.g. [80, 111, 112]), and have been shown to be good probes of non-Gaussian information. We find good overall agreement between a random realization of a full-sky network simulation and the S10 simulations.

The last pixel-level statistics we consider is the source number counts. In Figure 8, we show the histogram of the source number counts per square degree per flux bin




Figure 9: Shown are the power spectra of simulations at 148 GHz. The solid curves are calculated from a random sample of a full-sky simulation from the network, and dashed curves are calculated using the full-sky S10 simulation. A flux-cut of 7 mJy is applied to both CIB and Radio maps. For the purpose of visualization, we multiply the κ map by a factor of 100 before taking the power spectrum. Since we apply one-dimensional transfer functions and scale factors, computed by comparing the network power spectra to the S10 power spectra (see Section III D for details), the power spectra of the network and S10 match by construction.

for the simulated CIB and Radio maps. As discussed in Section III D, the network overestimates the number of CIB sources with flux levels greater than 1 mJy, and we correct for this by scaling down the network flux levels at the high flux tail. This scaling substantially improves the overall match. This scaling has negligible impact on the other CIB statistics discussed below, since the other statistics are largely driven by lower flux sources.

B. Two-point Statistics

Next, we show the two-point statistics of the simulations in Figure 9 and Figure 10. Since we apply one-dimensional transfer functions and scale factors, computed by comparing the network power spectra to the S10 power spectra (see Section III D for details), the power spectra of the network and S10 match by construction, as shown in Figure 9. Figure 10 shows the cross spectra between the κ map and the extragalactic foreground components at 148 GHz. Overall, we find a good match between the two sets of cross spectra. We find only a small extra correlation between κ and tSZ at small scales ($l \geq 8000$). We also check the cross correlation among the other components (for instance the tSZ and CIB), and similarly find good matches. Note that we




Figure 10: Shown via cross correlation coefficients are the cross spectra between the κ map and the extragalactic foreground components at 148 GHz. The solid curves are generated with a full-sky map from our network, while the dashed curves are generated from the S10 simulations. A flux-cut of 7 mJy is applied to both the CIB and Radio maps. There is a good overall match between the curves. Note that we do not apply any correction to the network cross spectra; the matches between the network and the S10 cross spectra indicate that the network has correctly learned the correlations among the foreground components.

do not apply any correction to the network cross spectra. The matches between the network and the S10 cross spectra indicate that the network has correctly learned the correlations among the foreground components.

C. Three-point and Four-point Functions

We also compute higher order statistics, in particular the three-point and four-point functions. We compute the bispectra of foregrounds following [87, 113, 114] using the *PiInTheSky* software.¹⁷ The binned bispectra of the simulations are shown in Figure 11 and Figure 12. One advantage of the binned bispectrum method is that it requires no explicit modeling of the signal. For visualization purposes, we plot the absolute value of the kSZ bispectra, and we plot the tSZ bispectra multiplied by a factor of -1. Overall we find a good match in the shape and amplitude of the bispectra, even though our training procedure was not explicitly optimized to reproduce the bispectra. With minor modifications to the training pro-




Figure 11: Shown are the binned bispectra of simulations at 148 GHz. The solid curves are calculated from a random sample of a full-sky simulation from the network, and dashed curves are calculated using the full-sky S10 simulation. A flux-cut of 7 mJy is applied to both CIB and Radio maps. Here $\ell = \sqrt{\ell_1^2 + \ell_2^2 + \ell_3^2}$, where ℓ_1, ℓ_2 , and ℓ_3 form a triangle.

cedure, we expect this match can be further improved.

In Figure 12, we show a squeezed configuration bispectra, where one leg of the triangle is $\ell_1 = 35$, and the other two legs have the same length ($\ell_2 = \ell_3$). As discussed in Section III D, we generate the full-sky realizations by “stitching” up $1^\circ \times 1^\circ$ tiles. Naively, the network simulations are not expected to reproduce the statistics for modes larger than the tile size (i.e. roughly $\ell \leq 200$). However, we find that the network correctly reproduces the correlations between the super tile modes ($\ell_1 = 35$) and the smaller scale modes ($\ell \geq 200$). This indicates that our overall procedure can generate faithful full-sky realizations, which contain the correct information at all scales.

Figure 13 and Figure 14 show the comparison for the four-point functions used to reconstruct the lensing potential. To do this comparison, we first generate an unlensed Gaussian CMB realization using the *WMAP5* cosmology [82]; we then make two lensed versions of this CMB map, one lensed by the the network κ and another lensed by the S10 κ map. Since our lensing reconstruction filter requires some instrument noise and a finite resolution, we add to each simulation a white noise level of $10\mu\text{K-arcminute}$ and a 1.3 arcminute Gaussian beam, which matches current ground-based CMB experiment sensitivity. We reconstruct the lensing potential using the temperature-only quadratic estimator [115] in order to focus on the bias induced from extragalactic foregrounds. We include CMB multipoles in the range of $\ell \in [100, 3000]$ for the reconstruction. Figure 13 shows the cross spectra between the reconstructed lensing con-

¹⁷ <https://github.com/simonsobs/PiInTheSky>




Figure 12: Similar to Figure 11, but for a squeezed configuration, where one leg is $\ell_1 = 35$, and the other two legs are the same length ($\ell_2 = \ell_3$). We find that the network can correctly reproduce the correlations between modes larger than the tile size (i.e. $1^\circ \times 1^\circ$ or roughly $\ell \leq 200$ in the harmonic domain) and modes within the tiles (i.e. roughly $\ell \geq 200$).

vergence and the input lensing convergence from both the network and S10 simulations, in the absence of other foreground contamination. We find a good match for all the scales considered. Figure 14 shows the contamination of the reconstructed lensing convergence from simulated lensed CMB maps plus the extragalactic foregrounds as a percent bias (i.e. $C_{L,bias}^{\kappa\kappa} / C_{L,unbiased}^{\kappa\kappa} \times 100$, where $C_L^{\kappa\kappa}$ is computed from the cross spectra between the reconstructed and the input lensing convergences as described above). We find good agreement between the biases computed from the network and those from the S10 simulations.

V. FINAL NETWORK PRODUCT

We make 500 mmDL realizations in total using our network, and release them publicly on the Legacy Archive for Microwave Background Data Analysis (LAMDA)¹⁸ and the National Energy Research Scientific Computing Center (NERSC) cluster¹⁹. These are full-sky simulations at half-arcminute resolution at six different frequencies (30, 90, 148, 219, 277, and 350 GHz), which include:




Figure 13: Shown are the reconstructed lensing auto spectra. The red solid curve is generated from a CMB map lensed with a full-sky κ map from the network. The blue dashed curve is generated using the same CMB map, instead lensed with the full-sky S10 κ . No foregrounds are added. Lensing reconstruction is done using the temperature-only quadratic estimator, including multipoles of $\ell \in [100, 3000]$. We compute the lensing auto spectrum as the cross spectra of the reconstructed kappa map and the input kappa map in order to avoid reconstruction noise bias.

- the lensing convergence map (κ),
- the kinetic Sunyaev-Zel'dovich effect (kSZ),
- the thermal Sunyaev-Zel'dovich effect (tSZ),
- the Cosmic Infrared Background (CIB), and
- the radio galaxies (Radio).

For each realization, we generate six maps of the frequency-dependent components (tSZ, CIB, and Radio), and one map of the frequency-independent components (lensed CMB (T, Q, U), κ , and kSZ). We release all of these components separately, as well as a combined final microwave sky map in temperature at each of the six frequencies. In total, each realization contains 29 ($6 \times 3 + 5 + 6$) full sky maps. The simulations are natively made in the CAR pixelization and converted to HEALPix by reprojecting the CAR maps to spherical harmonic a_{lm} up to an ℓ_{max} of 10,000 for κ , kSZ, and tSZ maps, and an ℓ_{max} of 25,000 for CIB and Radio maps; note that we chose a higher ℓ_{max} of 25,000 for CIB and Radio maps to reduce the effect of spectral leakage around the sources and to reproduce the source number counts shown in Figure 8. These a_{lm} are then reprojected onto full-sky HEALPix maps of $N_{side} = 8192$. We release this set of simulations in both HEALPix and CAR pixel formats,

¹⁸ https://lambda.gsfc.nasa.gov/simulation/tb_sim_ov.cfm

¹⁹ <https://crd.lbl.gov/departments/computational-science/c3/c3-research/cosmic-microwave-background/cmb-data-analysis-at-nersc/>




Figure 14: Shown is the percent bias to the reconstructed lensing auto spectrum induced by extragalactic foregrounds at 148 GHz. The solid curves are calculated with a random full-sky simulation from the network, while the dashed curves are computed using the S10 simulation. Lensed CMB maps are made as discussed in Figure 13. A flux-cut of 7 mJy is applied to both CIB and Radio maps. The lensing auto spectra are computed as the cross spectra of the reconstructed and input kappa maps as in Figure 13. We find good agreement between the S10 simulation and the network.

however, for space reasons we provide 500 HEALPix realizations and only 30 CAR realizations. We also provide scripts to convert between HEALPix and CAR maps. Each HEALPix fits file is 3.2 GB for double-precision, and each CAR file is 3.6 GB for single-precision. Thus the full package is 50 TB, for $500 \times 29 \times 3.2 \text{ GB} + 30 \times 29 \times 3.6 \text{ GB}$. We also release the code used in this work at <https://github.com/dwhan89/cosmikyu>.

VI. DISCUSSION

We have presented for the first time high-resolution full-sky DL simulations at millimeter wavelengths that include correlated foreground components. We find that these mmDL simulations can reproduce not only the power spectra of the original simulations, but also can naturally reproduce the other non-Gaussian statistics considered in this work (i.e. one-point PDFs, cross-spectra, bispectra, lensing reconstructions, and lensing reconstruction biases) with good overall agreement. This is notable since we only amend the mmDL simulations to reproduce the power spectra and the source number counts above 1 mJy.

The procedure we developed can be used to generate an unlimited number of full-sky realizations starting

from just a single realization of a full-sky simulation. Ordinarily there would not be enough independent small tiles from a single full-sky simulation to train a network to predict new foreground maps given a kappa map. However, in this work we achieve this by first generating an augmented data set (Intermediate Product 2 in Figure 1) by training our network on the original simulations. This augmented data set is more than four times larger than the original one and each patch is statistically independent; thus it can be used to train the next network step (Training Step 2 in Figure 1) that requires a larger amount of training data. This enables the capability to mass produce independent full-sky realizations from a single expensive full-sky simulation.

Currently, DL training is mostly done with small tiles due to limited computing resources (both memory and GPU). Though these resource issues will be alleviated as hardware improves overtime, we expect that they will persist for the foreseeable future. Unfortunately, these small DL simulations have a limited use when analyzing current and future wide-field CMB survey data. In this work, we outline a procedure for generating many realizations of full-sky maps, stitching together many individual tiles, that can faithfully recover the high-order statistics of a full-sky map without discontinuities or repeated features.

An advantageous feature of our network is that the input is a full-sky kappa map (top green box in Figure 1). Thus one can now in principle take a full-sky kappa map from any large-scale structure (LSS) simulation and generate the corresponding lensed CMB and correlated foreground components at millimeter wavelengths. This is especially useful in the current era of combining results from both CMB and LSS surveys, which require a common set of simulations. One can, for example, input in a full-sky kappa map taken from a simulation developed for the Rubin or Roman surveys [116, 117], and generate the corresponding CMB maps, resulting in simulations with both galaxies and CMB secondary components.²⁰

We envision these mmDL simulations useful for all parts of CMB data analyses, such as verification of pipelines, generation of covariance matrices, and analyses of foreground biases. Furthermore, since our procedure is numerically efficient, especially for generating foreground realizations on a small patch of the sky, we can use our DL method as a part of a forward modeling pipeline (such as discussed in [118]) where one varies the initial conditions to match observations.

Before we conclude, we briefly comment on the potential origin of the corrections we make in Section III D. The properties of generative DL models, in particularly those of GAN, are active areas of research (see [119] for more discussion), which is out of the scope of this work.

²⁰ If the given LSS kappa map includes non-Gaussian structure, we would just remove Training Step 3 from our network.

However, we discuss a couple considerations relevant to the applications of generative DL models in physics. First, there is no simple way to impose physical constraints (for instance, symmetry requirements) to neural networks, though there are studies in this direction (for example [120, 121]). This prevents DL practitioners from leveraging some physical insights important to the traditional modeling of physical processes. Second, even though the Wasserstein distance used throughout this work is the most general way to compute the distance between two distributions (i.e. the most general discriminator), we can only compute an approximated form of the Wasserstein metric due to numerical reasons. It is not intuitively clear whether optimizing this approximated version should always lead to physical results. These limitations are unlikely fundamental shortcomings of generative DL models, and we expect that they will be eventually overcome.

We also acknowledge that additional external independent datasets, similar to S10, would be useful for the further validation of the method presented in this work, in particular to assess both the statistical and systematic error inherent to the method. This would also potentially allow us to tune network parameters to match high-order statistics with greater precision. Given the expense of generating such simulations via traditional methods, it is not clear when many such datasets will become available.

Finally, we note that our network is trained to the specific cosmology and baryonic model used to generate the S10 simulations. In principle, a GAN network should be able to generate simulations at different cosmologies and with different gas physics by extrapolating the latent space variables (see [79] for discussion). This is potentially a very powerful technique that can vastly expand the applicability of these mmDL simulations.

Acknowledgments

DH would like to thank the Center for Computational Astrophysics (CCA) at the Flatiron Institute for support as a CCA pre-doctoral fellow. NS would like to thank the CCA for support during her sabbatical, during which time the idea for this project originated. DH and NS thank Julian Borrill and Andrea Zonca for their help with the product release on NERSC, and Graeme Addison for the product release on NASA LAMBDA. DH and NS acknowledge support from NSF grant numbers AST-1513618 and AST-1907657, and DOE grant number DE-SC0020441. FVN acknowledges funding from the WFIRST program through NNG26PJ30C and NNN12AA01C. Some of the results in this paper have been derived using the HEALPix package [88].

component	σ_p	$\mu_{p'}$	$\sigma_{p'}$
κ	N/A	5.342e-19	0.081
kSZ	N/A	5.173e-17	2.347
tSZ	3.6809	-0.629	0.286
CIB	17.134	0.605	0.349
Radio	6.392	0.199	0.421

Table II: The normalization parameters described in Section II C. Here σ_p is the standard deviation of pixel values used in the log-normalization (Equation 1). $\mu_{p'}$ and $\sigma_{p'}$ are related to the mean and the standard deviation of the pixel values used in the standard score (Equation 2). Note that the log-normalization is not applied to κ and kSZ maps as discussed in Section II C.

Appendix A: Normalization Parameters

In the DL literature, normalization refers to a set of techniques to modify the training data in order to stabilize and to speed up the training. A few popular normalization methods are the Box-Cox transformation [122] (which includes a log scaling), clipping, min-max, and the standard score. We tried a number of different normalization methods, and found that the combination of log scaling (adopted from [74]), and the standard score yielded the best results in our case. In Table II, we list the parameter values used in our normalization procedure. These parameter values are derived from the validation data set described in Section II B. Each of these parameters are described in detail in Section II C.

Appendix B: Grid schematic

Since our network can only process relatively small tiles, we need a systemic way to divide a full-sky map into a grid of tiles. For numerical reasons, the natural choice of tile width and height is 2^n pixels. In this work, we use the tile size of 128×128 pixels, which is roughly $1^\circ \times 1^\circ$. The schematic of our tile grid is shown in Figure 15. Note that we include 20 pixels overlap between tiles, which ensures continuity when we reproject tiles on to a full-sky map. This grid procedure results in a grid of tiles with 200 tiles along a longitudinal direction and 400 tiles along a latitude direction.

Appendix C: Frequency Dependence of CIB and Radio Sources

As discussed in Section III D, we estimate the frequency dependence of CIB and Radio sources using the S10 catalogues. For each source in the catalogues, we compute its spectral index between a given frequency and the reference frequency of 148 GHz by comparing




Figure 15: We show a schematic of the tiling grid described in Section III D. Each tile is 128×128 pixels (roughly $1^\circ \times 1^\circ$) in the CAR pixellization and corresponds to a patch effectively centered at the celestial equator. We cutout each tile with a 20-pixel overlap region around the edges to ensure that the final maps are continuous when we reproject the tiles back on to a full-sky CAR map.

the flux of that source at the given frequency (30, 90, 219, 277, and 350 GHz) to the flux at the reference frequency. In Table III, we show the means and standard deviations of the spectral indices for CIB and Radio sources. Note that we only consider the sources that meet our flux-cut criteria (≤ 7 mJy at 148 GHz) in this estimation.

For each source in the network 148 GHz map, we randomly assign it a spectral index for each frequency, with each index drawn from a Gaussian distribution with a

Frequency (GHz)	β_{CIB}	β_{Radio}
30	3.304 ± 0.138	-0.799 ± 0.158
90	3.202 ± 0.442	-0.811 ± 0.144
219	2.973 ± 0.567	-0.820 ± 0.135
277	2.870 ± 0.372	-0.822 ± 0.134
350	2.745 ± 0.301	-0.823 ± 0.133

Table III: The estimated spectral indices of CIB and Radio sources described in Section III D and Appendix C. The spectral indices are estimated using the S10 source catalogues by comparing the flux of a source at a given frequency (30, 90, 219, 277, and 350 GHz) to the flux at the reference frequency of 148 GHz; the mean and standard deviation of the source spectral indices are listed in the Table.




Figure 16: We show the power spectra of CIB maps at different frequencies. The solid curves are calculated from a random realization of a full-sky simulation from the network, and dashed curves are calculated using the full-sky S10 simulation. A flux-cut of 7 mJy at 148 GHz is applied to all CIB maps. For each source in the 148 GHz map, we simulate its frequency dependence as described in Section III D and Appendix C. We find good agreement for the CIB power spectra between the S10 simulation and the network.

mean and the standard deviation as specified in Table III; we then scale the source’s 148 GHz flux accordingly and place it in the new frequency map. We find that we can recover well the frequency dependence of the original S10 CIB and Radio maps using this method as shown in Figure 16 and Figure 17.

Appendix D: Realization Dependent Scatter

One common failure mode of GANs is mode collapse, which is when a GAN repeatedly generates the same or statistically similar images. In the context of our work, mode collapse will lead to a lower variance of the simulations than expected, due to the presence of repeated tiles. In this section, we empirically show that the variances of our network simulations match the expected variances. Figure 18 shows the diagonal elements of the covariance matrix derived from the κ power spectrum of 20 network simulations. Here, the expected variance is given by Equation 4 in [123]; in the absence of experimental noise, this equation gives the expected variance of an approximately Gaussian random field. Similarly, we show the estimated variance of 20 network kSZ and CIB maps in Figure 19 and Figure 20. For each of these components, we find that the variance from the network simulations matches the expected variance. Since




Figure 17: The same as Figure 16 above, but for Radio maps instead of CIB maps.

each lensed CMB realization is generated from independent unlensed CMB realizations, the lensed CMB maps should have the correct variance by the construction.

Appendix E: Supplementary Network Training Details

The details of our network architectures are discussed in Section III. In this section, we provide supplementary details of our training procedure. Table IV summarizes a few additional parameters of our network training procedure. As shown in the table, all three GAN layers share most of the common parameters including the learning rate and the Adam optimizer parameters. We check our metrics (images, loss function, power spectra, and others) every 100,000 training data samples processed (i.e. two times per epoch for DCWGAN-GP and VAEGAN; eight times per epoch for PIXGAN). If we find

	DCWGAN-GP	PIXGAN	VAEGAN
Type	Non Conditional	Conditional	Conditional
Loss function	Wasserstein		
Optimize	Adam Optimizer ($\beta_1 = 0.5, \beta_2 = 0.9$)		
Learning Rate	Initially $lr = 10^{-4}$. Halved if needed		
Batch Size	32		

Table IV: Shown are additional parameters used to train the three GAN layers discussed in Section III.




Figure 18: We show the diagonal elements of the covariance matrix derived from the κ power spectrum. The blue curves are estimated from 20 realizations of a full-sky simulation from the network, and the dashed curve is calculated using Equation 4 in [123] given the S10 κ power spectrum. We find good agreement between the variance from the network simulations and the expected variance.




Figure 19: The same as Figure 18 above, but for kSZ maps instead of κ maps.

a degradation of the output quality, we restart the training from the last epoch with the learning rate reduced by a factor of two. Also, if the model either can not reproduce the power spectra after the loss function plateaus or images are not visually correct, we start the training from scratch with the learning rate reduced by a factor of




Figure 20: The same as Figure 18 above, but for CIB maps at 148 GHz instead of κ maps. Note that a flux-cut of 7 mJy at 148 GHz is applied to the CIB maps.

two. Lastly, we do not explicitly check for mode collapse, but rely on the summary statistics (shown in Section IV) and the estimated variances (shown in Appendix D) to infer the presence of mode collapse.

- [1] Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, et al., *Astronomy and Astrophysics* **641**, A6 (2020), 1807.06209.
- [2] J. W. Henning, J. T. Sayre, C. L. Reichardt, P. A. R. Ade, A. J. Anderson, J. E. Austermann, J. A. Beall, A. N. Bender, B. A. Benson, L. E. Bleem, et al., *The Astrophysical Journal* **852**, 97 (2018), URL <https://doi.org/10.3847/2F1538-4357%2Faa9ff4>.
- [3] S. Aiola, E. Calabrese, L. Maurin, S. Naess, B. L. Schmitt, M. H. Abitbol, G. E. Addison, P. A. R. Ade, D. Alonso, M. Amiri, et al., *Journal of Cosmology and Astroparticle Physics* **2020**, 047 (2020), 2007.07288.
- [4] S. K. Choi, M. Hasselfield, S.-P. P. Ho, B. Koopman, M. Lungu, M. H. Abitbol, G. E. Addison, P. A. R. Ade, S. Aiola, D. Alonso, et al., *Journal of Cosmology and Astroparticle Physics* **2020**, 045 (2020), 2007.07289.
- [5] D. Han, N. Sehgal, A. MacInnis, A. van Engelen, B. D. Sherwin, M. S. Madhavacheril, S. Aiola, N. Battaglia, J. A. Beall, D. T. Becker, et al., *Journal of Cosmology and Astroparticle Physics* **2021**, 031 (2021), 2007.14405.
- [6] B. D. Sherwin, A. van Engelen, N. Sehgal, M. Madhavacheril, G. E. Addison, S. Aiola, R. Allison, N. Battaglia, D. T. Becker, J. A. Beall, et al., *Phys. Rev. D* **95**, 123529 (2017), 1611.09753.
- [7] Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, et al., *Astronomy and Astrophysics* **641**, A8 (2020), 1807.06210.
- [8] W. L. K. Wu, L. M. Mocuano, P. A. R. Ade, A. J. Anderson, J. E. Austermann, J. S. Avva, J. A. Beall, A. N. Bender, B. A. Benson, F. Bianchini, et al., *Astrophys. J.* **884**, 70 (2019), 1905.05777.
- [9] Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, et al., *Astronomy and Astrophysics* **594**, A24 (2016), 1502.01597.
- [10] S. Amodeo, N. Battaglia, E. Schaan, S. Ferraro, E. Moser, S. Aiola, J. E. Austermann, J. A. Beall, R. Bean, D. T. Becker, et al., *Phys. Rev. D* **103**, 063514 (2021), 2009.05558.
- [11] E. Schaan, S. Ferraro, S. Amodeo, N. Battaglia, S. Aiola, J. E. Austermann, J. A. Beall, R. Bean, D. T. Becker, R. J. Bond, et al., *Phys. Rev. D* **103**, 063513 (2021), 2009.05557.
- [12] N. Huang, L. E. Bleem, B. Stalder, P. A. R. Ade, S. W. Allen, A. J. Anderson, J. E. Austermann, J. S. Avva, J. A. Beall, A. N. Bender, et al., *Astronomical Journal* **159**, 110 (2020), 1907.09621.
- [13] A. H. Jaffe, E. Gawiser, D. Finkbeiner, J. C. Baker, A. Balbi, M. Davis, S. Hanany, W. Holzzapfel, M. Krumholz, L. Moustakas, et al., in *Microwave Foregrounds*, edited by A. de Oliveira-Costa and M. Tegmark (1999), vol. 181 of *Astronomical Society of the Pacific Conference Series*, p. 367, astro-ph/9903248.
- [14] J. González-Nuevo, L. Toffolatti, and F. Argüeso, *Astrophys. J.* **621**, 1 (2005), astro-ph/0405553.
- [15] B. M. Schäfer, C. Pfrommer, M. Bartelmann, V. Springel, and L. Hernquist, *Monthly Notices of the RAS* **370**, 1309 (2006), astro-ph/0407089.
- [16] M. Roncarelli, L. Moscardini, S. Borgani, and K. Dolag, *Monthly Notices of the RAS* **378**, 1259 (2007), astro-ph/0701680.
- [17] C. Carbone, V. Springel, C. Baccigalupi, M. Bartelmann, and S. Matarrese, *Monthly Notices of the RAS* **388**, 1618 (2008), 0711.2655.
- [18] N. Battaglia, J. R. Bond, C. Pfrommer, J. L. Sievers, and D. Sijacki, *The Astrophysical Journal* **725**, 91 (2010), URL <https://doi.org/10.1088/0004-637x/725/1/91>.
- [19] N. Sehgal, P. Bode, S. Das, C. Hernandez-Monteagudo, K. Huffenberger, Y.-T. Lin, J. P. Ostriker, and H. Trac, *Astrophys. J.* **709**, 920 (2010), 0908.0540.
- [20] S. Flender, L. Bleem, H. Finkel, S. Habib, K. Heitmann, and G. Holder, *Astrophys. J.* **823**, 98 (2016), 1511.02843.
- [21] R. Takahashi, T. Hamana, M. Shirasaki, T. Namikawa, T. Nishimichi, K. Osato, and K. Shiroyama, *Astrophys. J.* **850**, 24 (2017), 1706.01472.
- [22] G. Stein, M. A. Alvarez, J. R. Bond, A. van Engelen, and N. Battaglia, *Journal of Cosmology and Astroparticle Physics* **2020**, 012 (2020), 2001.08787.
- [23] J. W. Fowler, M. D. Niemack, S. R. Dicker, A. M. Aboobaker, P. A. R. Ade, E. S. Battistelli, M. J. Devlin, R. P. Fisher, M. Halpern, P. C. Hargrave, et al., *Appl. Opt.* **46**, 3444 (2007), astro-ph/0701020.
- [24] R. J. Thornton, P. A. R. Ade, S. Aiola, F. E. Angile, M. Amiri, J. A. Beall, D. T. Becker, H. M. Cho, S. K. Choi, P. Corlies, et al., *Astrophysical Journal, Supplement* **227**, 21 (2016), 1605.06569.
- [25] J. E. Carlstrom, P. A. R. Ade, K. A. Aird, B. A. Benson, L. E. Bleem, S. Busetti, C. L. Chang, E. Chauvin, H. M. Cho, T. M. Crawford, et al., *PASP* **123**, 568 (2011), 0907.4445.
- [26] P. Ade, J. Aguirre, Z. Ahmed, S. Aiola, A. Ali, D. Alonso, M. A. Alvarez, K. Arnold, P. Ashton, J. Austermann, et al., *Journal of Cosmology and Astroparticle Physics* **2019**, 056 (2019), 1808.07445.
- [27] K. Abazajian, G. Addison, P. Adshead, Z. Ahmed, S. W. Allen, D. Alonso, M. Alvarez, A. Anderson, K. S. Arnold, C. Baccigalupi, et al., arXiv e-prints arXiv:1907.04473 (2019), 1907.04473.
- [28] N. Sehgal, S. Aiola, Y. Akrami, K. Basu, M. Boylan-Kolchin, S. Bryan, S. Clesse, F.-Y. Cyr-Racine, L. Di Mascolo, S. Dicker, et al., in *Bulletin of the American Astronomical Society* (2019), vol. 51, p. 6, 1906.10134.
- [29] N. Sehgal, S. Aiola, Y. Akrami, K. moni Basu, M. Boylan-Kolchin, S. Bryan, C. M. Casey, S. Clesse, F.-Y. Cyr-Racine, L. Di Mascolo, et al., arXiv e-prints arXiv:2002.12714 (2020), 2002.12714.
- [30] A. Lewis, *Phys. Rev. D* **71**, 083008 (2005), astro-ph/0502469.
- [31] A. Amblard, C. Vale, and M. White, *New Astronomy* **9**, 687 (2004), astro-ph/0403075.
- [32] S. Basak, S. Prunet, and K. Benabed, arXiv e-prints arXiv:0811.1677 (2008), 0811.1677.
- [33] G. Fabbian and R. Stomp, *Astronomy and Astrophysics* **556**, A109 (2013), 1303.6550.
- [34] S. K. Naess and T. Louis, *Journal of Cosmology and Astroparticle Physics* **2013**, 001 (2013), 1307.0719.
- [35] C. Baccigalupi, *New Astronomy Reviews* **47**, 1127 (2003), astro-ph/0306181.
- [36] P. Hennebelle, R. Banerjee, E. Vázquez-Semadeni, R. S. Klessen, and E. Audit, *Astronomy and Astrophysics* **486**, L43 (2008), 0805.1366.
- [37] A. Waelkens, T. Jaffe, M. Reinecke, F. S. Kitaura, and T. A. Enßlin, *Astronomy and Astrophysics* **495**, 697 (2009),

- 0807.2262.
- [38] S. K. Choi and L. A. Page, *Journal of Cosmology and Astroparticle Physics* **2015**, 020 (2015), 1509.05934.
- [39] B. Thorne, J. Dunkley, D. Alonso, and S. Naess, *Monthly Notices of the RAS* **469**, 2821 (2017), 1608.02841.
- [40] F. Vansyngel, F. Boulanger, T. Ghosh, B. Wandelt, J. Aumont, A. Bracco, F. Levrier, P. G. Martin, and L. Montier, *Astronomy and Astrophysics* **603**, A62 (2017), 1611.02577.
- [41] D. Beck, J. Errard, and R. Stomp, *Journal of Cosmology and Astroparticle Physics* **2020**, 030 (2020), 2001.02641.
- [42] K. Aylor, M. Haq, L. Knox, Y. Hezaveh, and L. Perreault-Levasseur, *Monthly Notices of the RAS* **500**, 3889 (2021), 1909.06467.
- [43] B. Thorne, L. Knox, and K. Prabhu, *Monthly Notices of the RAS* **504**, 2603 (2021), 2101.11181.
- [44] N. Krachmalnicoff and G. Puglisi, *Astrophys. J.* **911**, 42 (2021), 2011.02221.
- [45] L. Perreault Levasseur, Y. D. Hezaveh, and R. H. Wechsler, *Astrophysical Journal, Letters* **850**, L7 (2017), 1708.08843.
- [46] S. Ravanbakhsh, J. Oliva, S. Fromenteau, L. C. Price, S. Ho, J. Schneider, and B. Poczós, arXiv e-prints arXiv:1711.02033 (2017), 1711.02033.
- [47] A. Gupta, J. M. Zorrilla Matilla, D. Hsu, and Z. Haiman, *Phys. Rev. D* **97**, 103515 (2018), 1802.01212.
- [48] A. Mathuriya, D. Bard, P. Mendygral, L. Meadows, J. Arnemann, L. Shao, S. He, T. Karna, D. Moise, S. J. Pennycook, et al., arXiv e-prints arXiv:1808.04728 (2018), 1808.04728.
- [49] J. Alsing, T. Charnock, S. Feeney, and B. Wandelt, *Monthly Notices of the RAS* **488**, 4440 (2019), 1903.00007.
- [50] M. Shirasaki, K. Moriwaki, T. Oogi, N. Yoshida, S. Ikeda, and T. Nishimichi, *Monthly Notices of the RAS* **504**, 1825 (2021), 1911.12890.
- [51] L. Lucie-Smith, H. V. Peiris, A. Pontzen, B. Nord, and J. Thiyaalingam, arXiv e-prints arXiv:2011.10577 (2020), 2011.10577.
- [52] F. Villaescusa-Navarro, B. D. Wandelt, D. Anglés-Alcázar, S. Genel, J. M. Zorrilla Matilla, S. Ho, and D. N. Spergel, arXiv e-prints arXiv:2011.05992 (2020), 2011.05992.
- [53] M. A. Petroff, G. E. Addison, C. L. Bennett, and J. L. Weiland, *Astrophys. J.* **903**, 104 (2020), 2004.11507.
- [54] K. Moriwaki, M. Shirasaki, and N. Yoshida, *Astrophysical Journal, Letters* **906**, L1 (2021), 2010.00809.
- [55] J. Caldeira, W. L. K. Wu, B. Nord, C. Avestruz, S. Trivedi, and K. T. Story, *Astronomy and Computing* **28**, 100307 (2019), 1810.01483.
- [56] V. Bonjean, *Astronomy and Astrophysics* **634**, A81 (2020), 1911.10778.
- [57] E. Guzman and J. Meyers, *Phys. Rev. D* **104**, 043529 (2021), 2101.01214.
- [58] Z. Lin, N. Huang, C. Avestruz, W. L. K. Wu, S. Trivedi, J. Caldeira, and B. Nord, *Monthly Notices of the RAS* **507**, 4149 (2021), 2102.13123.
- [59] G. Puglisi and X. Bai, *Astrophys. J.* **905**, 143 (2020), 2003.13691.
- [60] G. Montefalcone, M. H. Abitbol, D. Kodwani, and R. D. P. Grumitt, *Journal of Cosmology and Astroparticle Physics* **2021**, 055 (2021), 2011.01433.
- [61] A. Vafaei Sadr and F. Farsian, *Journal of Cosmology and Astroparticle Physics* **2021**, 012 (2021), 2004.04177.
- [62] M. Münchmeyer and K. M. Smith, arXiv e-prints arXiv:1905.05846 (2019), 1905.05846.
- [63] L. Lucie-Smith, H. V. Peiris, A. Pontzen, and M. Lochner, *Monthly Notices of the RAS* **479**, 3405 (2018), 1802.04271.
- [64] C. C. Lovell, V. Acquaviva, P. A. Thomas, K. G. Iyer, E. Gawiser, and S. M. Wilkins, *Monthly Notices of the RAS* **490**, 5503 (2019), 1903.10457.
- [65] R. Alves de Oliveira, Y. Li, F. Villaescusa-Navarro, S. Ho, and D. N. Spergel, arXiv e-prints arXiv:2012.00240 (2020), 2012.00240.
- [66] B. H. Chen, T. Goto, S. J. Kim, T. W. Wang, D. J. D. Santos, S. C.-C. Ho, T. Hashimoto, A. Poliszczuk, A. Pollo, S. Trippe, et al., *Monthly Notices of the Royal Astronomical Society* **501**, 3951 (2020), ISSN 0035-8711, <https://academic.oup.com/mnras/article-pdf/501/3/3951/35747741/staa3865.pdf>, URL <https://doi.org/10.1093/mnras/staa3865>.
- [67] M. Cranmer, A. Sanchez-Gonzalez, P. Battaglia, R. Xu, K. Cranmer, D. Spergel, and S. Ho, arXiv e-prints arXiv:2006.11287 (2020), 2006.11287.
- [68] D. Wadekar, F. Villaescusa-Navarro, S. Ho, and L. Perreault-Levasseur, arXiv e-prints arXiv:2012.00111 (2020), 2012.00111.
- [69] G.-J. Wang, X.-J. Ma, and J.-Q. Xia, *Monthly Notices of the RAS* **501**, 5714 (2021), 2004.13913.
- [70] P. Berger and G. Stein, *Monthly Notices of the RAS* **482**, 2861 (2019), 1805.04537.
- [71] S. He, Y. Li, Y. Feng, S. Ho, S. Ravanbakhsh, W. Chen, and B. Poczós, *Proceedings of the National Academy of Science* **116**, 13825 (2019), 1811.06533.
- [72] M. Mustafa, D. Bard, W. Bhimji, Z. Lukić, R. Al-Rfou, and J. M. Kratochvil, *Computational Astrophysics and Cosmology* **6**, 1 (2019), 1706.02390.
- [73] N. Perraudin, A. Srivastava, A. Lucchi, T. Kacprzak, T. Hofmann, and A. Réfrégier, *Computational Astrophysics and Cosmology* **6**, 5 (2019), 1908.05519.
- [74] T. Tröster, C. Ferguson, J. Harnois-Déraps, and I. G. McCarthy, *Monthly Notices of the RAS* **487**, L24 (2019), 1903.12173.
- [75] C. Chen, Y. Li, F. Villaescusa-Navarro, S. Ho, and A. Pullen, arXiv e-prints arXiv:2012.05472 (2020), 2012.05472.
- [76] O. Curtis and T. G. Brainerd, *Research Notes of the American Astronomical Society* **4**, 90 (2020), 2006.11359.
- [77] B. Dai and U. Seljak, *Proceedings of the National Academy of Science* **118**, 2020324118 (2021), 2010.02926.
- [78] Y. Li, Y. Ni, R. A. C. Croft, T. Di Matteo, S. Bird, and Y. Feng, *Proceedings of the National Academy of Science* **118**, 2022038118 (2021), 2010.06608.
- [79] A. Tamosiunas, H. A. Winther, K. Koyama, D. J. Bacon, R. C. Nichol, and B. Mawdsley, *Monthly Notices of the RAS* **506**, 3049 (2021), 2004.10223.
- [80] L. Thiele, J. C. Hill, and K. M. Smith, *Phys. Rev. D* **102**, 123545 (2020), 2009.06547.
- [81] M. Ullmo, A. Decelle, and N. Aghanim, *Astronomy and Astrophysics* **651**, A46 (2021), 2011.05244.
- [82] J. Dunkley, E. Komatsu, M. R. Nolta, D. N. Spergel, D. Larson, G. Hinshaw, L. Page, C. L. Bennett, B. Gold, N. Jarosik, et al., *Astrophysical Journal, Supplement* **180**, 306 (2009), 0803.0586.
- [83] F. Lacasa, N. Aghanim, M. Kunz, and M. Frommert, *Monthly Notices of the Royal Astronomical Society* **421**, 1982–1995 (2012), ISSN 0035-8711, URL <http://dx.doi.org/10.1111/j.1365-2966.2012.20415.x>.
- [84] F. Lacasa, A. Pénin, and N. Aghanim, *Monthly No-*

- tices of the Royal Astronomical Society **439**, 123–142 (2014), ISSN 0035-8711, URL <http://dx.doi.org/10.1093/mnras/stt2373>.
- [85] N. Hand, G. E. Addison, E. Aubourg, N. Battaglia, E. S. Battistelli, D. Bizyaev, J. R. Bond, H. Brewington, J. Brinkmann, B. R. Brown, et al., *Physical Review Letters* **109** (2012), ISSN 1079-7114, URL <http://dx.doi.org/10.1103/PhysRevLett.109.041101>.
- [86] J. C. Hill and B. D. Sherwin, *Physical Review D* **87** (2013), ISSN 1550-2368, URL <http://dx.doi.org/10.1103/PhysRevD.87.023527>.
- [87] A. van Engelen, S. Bhattacharya, N. Sehgal, G. P. Holder, O. Zahn, and D. Nagai, *Astrophys. J.* **786**, 13 (2014), 1310.7023.
- [88] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, *Astrophys. J.* **622**, 759 (2005), astro-ph/0409513.
- [89] J. Dunkley, E. Calabrese, J. Sievers, G. E. Addison, N. Battaglia, E. S. Battistelli, J. R. Bond, S. Das, M. J. Devlin, R. Dünner, et al., *Journal of Cosmology and Astroparticle Physics* **2013**, 025 (2013), 1301.0776.
- [90] J. L. Sievers, R. A. Hlozek, M. R. Nolta, V. Acquaviva, G. E. Addison, P. A. R. Ade, P. Aguirre, M. Amiri, J. W. Appel, L. F. Barrientos, et al., *Journal of Cosmology and Astroparticle Physics* **2013**, 060 (2013), 1301.0824.
- [91] Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, et al., *Astronomy and Astrophysics* **571**, A21 (2014), 1303.5081.
- [92] E. M. George, C. L. Reichardt, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, T. M. Crawford, A. T. Crites, et al., *Astrophys. J.* **799**, 177 (2015), 1408.3161.
- [93] Planck Collaboration, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, N. Bartolo, et al., *Astronomy and Astrophysics* **594**, A22 (2016), 1502.01596.
- [94] M. Reinecke and D. S. Seljebotn, *Astronomy and Astrophysics* **554**, A112 (2013), 1303.4945.
- [95] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, arXiv e-prints arXiv:1406.2661 (2014), 1406.2661.
- [96] I. Goodfellow, arXiv e-prints arXiv:1701.00160 (2016), 1701.00160.
- [97] A. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, and A. A. Bharath, *IEEE Signal Processing Magazine* **35**, 53 (2018).
- [98] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, arXiv e-prints arXiv:1704.00028 (2017), 1704.00028.
- [99] B. Xu, N. Wang, T. Chen, and M. Li, arXiv e-prints arXiv:1505.00853 (2015), 1505.00853.
- [100] C. Nwankpa, W. Ijomah, A. Gachagan, and S. Marshall, *Activation functions: Comparison of trends in practice and research for deep learning* (2018), 1811.03378.
- [101] D. P. Kingma and J. Ba, arXiv e-prints arXiv:1412.6980 (2014), 1412.6980.
- [102] O. Ronneberger, P. Fischer, and T. Brox, arXiv e-prints arXiv:1505.04597 (2015), 1505.04597.
- [103] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, arXiv e-prints arXiv:1611.07004 (2016), 1611.07004.
- [104] C. Dong, C. C. Loy, K. He, and X. Tang, in *Computer Vision – ECCV 2014*, edited by D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars (Springer International Publishing, Cham, 2014), pp. 184–199, ISBN 978-3-319-10593-2.
- [105] A. Boesen Lindbo Larsen, S. Kaae Sønderby, H. Larochelle, and O. Winther, arXiv e-prints arXiv:1512.09300 (2015), 1512.09300.
- [106] S. Naess, S. Aiola, J. E. Austermann, N. Battaglia, J. A. Beall, D. T. Becker, R. J. Bond, E. Calabrese, S. K. Choi, N. F. Cothard, et al., *Journal of Cosmology and Astroparticle Physics* **2020**, 046 (2020), 2007.07290.
- [107] I. Yeo and R. A. Johnson, *Biometrika* **87**, 954 (2000), ISSN 0006-3444, <https://academic.oup.com/biomet/article-pdf/87/4/954/633221/870954.pdf>, URL <https://doi.org/10.1093/biomet/87.4.954>.
- [108] Y. B. Zeldovich and R. A. Sunyaev, “ApSS” **4**, 301 (1969).
- [109] R. A. Sunyaev and Y. B. Zeldovich, “ApSS” **7**, 3 (1970).
- [110] H. Mantz, K. Jacobs, and K. Mecke, *Journal of Statistical Mechanics: Theory and Experiment* **2008**, P12015 (2008).
- [111] J. C. Hill, B. D. Sherwin, K. M. Smith, G. E. Addison, N. Battaglia, E. S. Battistelli, J. R. Bond, E. Calabrese, M. J. Devlin, J. Dunkley, et al., arXiv e-prints arXiv:1411.8004 (2014), 1411.8004.
- [112] J. Liu, J. C. Hill, B. D. Sherwin, A. Petri, V. Böhm, and Z. Haiman, *Phys. Rev. D* **94**, 103501 (2016), 1608.03169.
- [113] T. M. Crawford, K. K. Schaffer, S. Bhattacharya, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, A. T. Crites, et al., *Astrophys. J.* **784**, 143 (2014), 1303.3535.
- [114] W. R. Coulton and D. N. Spergel, *Journal of Cosmology and Astroparticle Physics* **2019**, 056 (2019), 1901.04515.
- [115] W. Hu and T. Okamoto, *Astrophys. J.* **574**, 566 (2002), astro-ph/0111606.
- [116] D. Spergel, N. Gehrels, C. Baltay, D. Bennett, J. Breckinridge, M. Donahue, A. Dressler, B. S. Gaudi, T. Greene, O. Guyon, et al., arXiv e-prints arXiv:1503.03757 (2015), 1503.03757.
- [117] Ž. Ivezić, S. M. Kahn, J. A. Tyson, B. Abel, E. Acosta, R. Allsman, D. Alonso, Y. AlSayyad, S. F. Anderson, J. Andrew, et al., *Astrophys. J.* **873**, 111 (2019), 0805.2366.
- [118] M. Millea, C. M. Daley, T.-L. Chou, E. Anderes, P. A. R. Ade, A. J. Anderson, J. E. Austermann, J. S. Avva, J. A. Beall, A. N. Bender, et al., arXiv e-prints arXiv:2012.01709 (2020), 2012.01709.
- [119] D. Bau, J.-Y. Zhu, H. Strobelt, B. Zhou, J. B. Tenenbaum, W. T. Freeman, and A. Torralba, arXiv e-prints arXiv:1811.10597 (2018), 1811.10597.
- [120] Y. Zhu, N. Zabarar, P.-S. Koutsourelakis, and P. Perdikaris, *Journal of Computational Physics* **394**, 56 (2019), 1901.06314.
- [121] J.-L. Wu, K. Kashinath, A. Albert, D. Chirila, Prabhat, and H. Xiao, *Journal of Computational Physics* **406**, 109209 (2020), 1905.06841.
- [122] G. E. P. Box and D. R. Cox, *Journal of the Royal Statistical Society. Series B (Methodological)* **26**, pp. 211 (1964), ISSN 00359246, URL <http://www.jstor.org/stable/2984418>.
- [123] L. Knox, *Physical Review D* **52**, 4307 (1995), astro-ph/9504054.