

Reconstructing Cosmic Polarization Rotation with ResUNet-CMB

Eric Guzman¹ and Joel Meyers¹

¹*Department of Physics, Southern Methodist University, 3215 Daniel Ave, Dallas, Texas 75275, USA*

(Dated: December 13, 2022)

Cosmic polarization rotation, which may result from parity-violating new physics or the presence of primordial magnetic fields, converts E -mode polarization of the cosmic microwave background (CMB) into B -mode polarization. Anisotropic cosmic polarization rotation leads to statistical anisotropy in CMB polarization and can be reconstructed with quadratic estimator techniques similar to those designed for gravitational lensing of the CMB. At the sensitivity of upcoming CMB surveys, lensing-induced B -mode polarization will act as a limiting factor in the search for anisotropic cosmic polarization rotation, meaning that an analysis which incorporates some form of delensing will be required to improve constraints on the effect with future surveys. In this paper we extend the ResUNet-CMB convolutional neural network to reconstruct anisotropic cosmic polarization rotation in the presence of gravitational lensing and patchy reionization, and we show that the network simultaneously reconstructs all three effects with variance that is lower than that from the standard quadratic estimator nearly matching the performance of an iterative reconstruction method.

I. INTRODUCTION

Observations of the cosmic microwave background (CMB) have played a key role in establishing the current standard model of cosmology. Upcoming CMB surveys, including those with Simons Observatory [1], BICEP Array [2], CCAT-prime [3], CMB-S4 [4], Lite-BIRD [5], PICO [6], and CMB-HD [7] will provide exquisite measurements of the temperature and polarization anisotropies across a wide range of angular scales. These measurements will provide critical tests of the standard cosmological model and enable improved searches for novel cosmological signals.

One of the main targets for upcoming surveys is B -mode polarization sourced by primordial gravitational waves [8, 9]. Gravitational lensing of the CMB [10] acts as a source of confusion for primordial gravitational wave searches since lensing can convert E -mode polarization into B -mode polarization [11, 12]. Delensing the CMB, the process of using an estimate of the lensing field to reverse its effects, can improve constraints on primordial gravitational waves [13–15] and will be necessary to achieve the goals of upcoming surveys [4, 16, 17].

Cosmic polarization rotation, or cosmic birefringence, can also convert E -mode polarization to B -mode polarization. Cosmic birefringence can be caused by parity-violating new physics [18–21] or by the presence of primordial magnetic fields [22]. A well-motivated class of models containing parity violation is provided by axion-like fields that may be related to dark matter or dark energy, with Chern-Simons coupling to the electromagnetic field strength [23, 24].

Isotropic cosmic birefringence, characterized by a constant rotation angle of polarization vectors across the sky, leads to parity-violating EB and TB correlations in the CMB. Searches for these correlations are challenging due to a degeneracy of the effects of isotropic cosmic birefringence with calibration error of the polarization direction of the instrument [25–27]. This degeneracy can be miti-

gated by searching for the different effects of cosmic birefringence on photons propagating over different distances such as polarized foregrounds [28, 29] or CMB photons scattered during reionization [30]. A recent analysis of data from the Planck satellite [31] has shown a 2.4σ hint of non-zero isotropic cosmic birefringence [32].

Anisotropic cosmic birefringence, where the rotation angle takes different values along different directions, leads to non-stationary CMB statistics. The mode-coupling induced by anisotropic cosmic birefringence is analogous to the effect of deflection caused by gravitational lensing, and this change to CMB statistics can be exploited to reconstruct a map of the rotation angle [33–35] using techniques that are very similar to lensing reconstruction [36]. Searches for anisotropic cosmic birefringence with existing CMB data have placed upper limits on the angular power spectrum of the rotation angle [37–40], with the tightest constraints coming from ACTPol [41] and SPTpol [42] $\ell(\ell+1)C_\ell^{\alpha\alpha}/(2\pi) < 0.10 \times 10^{-4} \text{ rad}^2$ (0.033 deg^2 , 95% CL). Future CMB surveys are expected to dramatically improve constraints on cosmic birefringence [43].

Much like the search for primordial gravitational waves, B -mode polarization induced by CMB lensing acts as a source of confusion in searches for cosmic polarization rotation [34]. Therefore, delensing the CMB is necessary in order to achieve the tightest constraints on cosmic birefringence. Delensing also allows for tighter cosmological parameter constraints [44], reduces lensing-induced covariance, aids the search for primordial non-Gaussianity [45], and reduces the variance in the reconstruction of other CMB secondaries including patchy reionization [46].

CMB delensing using an estimate of the lensing field reconstructed from the observed CMB map will provide the best delensing performance at the sensitivity of upcoming surveys [47]. The process of internal delensing impacts the statistics of the delensed CMB in a way that can lead to biases if not treated carefully [48, 49]. Inter-

nal delensing has been demonstrated on real data [50–52], and improved delensing techniques are being developed and refined [53–59].

Machine learning has been shown to hold promise for CMB lensing reconstruction and delensing [60]. In a previous paper, we extended the network described in Ref. [60] and showed that the resulting network, ResUNet-CMB, performs very well at simultaneously delensing and reconstructing the modulation resulting from patchy reionization [61]. We showed that ResUNet-CMB achieved nearly optimal performance on patchy reionization reconstruction over a wide range of angular scales.

In this paper, we demonstrate that ResUNet-CMB is capable of achieving excellent performance when applied to the reconstruction of cosmic polarization rotation as well. In particular, we show that ResUNet-CMB reconstructs anisotropic cosmic birefringence with a variance that is lower than could be achieved by applying the standard quadratic estimator to the lensed and rotated CMB. Achieving this high fidelity reconstruction of cosmic polarization rotation required very little modification of the network architecture of ResUNet-CMB. It is very encouraging that ResUNet-CMB successfully and straightforwardly generalizes to incorporate cosmic birefringence reconstruction. The physical effects of gravitational lensing, patchy reionization, and cosmic polarization rotation are quite different, and still the machine learning architecture is able to handle them all.

This paper is organized as follows. In Sec. II, we review the quadratic estimator for fields that produce statistical anisotropy of the CMB and discuss iterative reconstruction techniques. Section III describes the ResUNet-CMB architecture, simulation data pipeline, and training methods. We present the predictions of ResUNet-CMB and draw comparisons with the quadratic estimator in Sec. IV. We conclude in Sec. V.

II. QUADRATIC ESTIMATOR

The standard technique for reconstruction of the CMB lensing potential makes use of a quadratic estimator [36]. A formalism similar to the lensing quadratic estimator can also be applied to reconstruct other distortion fields such as patchy reionization [62, 63] and cosmic birefringence [33–35]. In this section we will discuss the quadratic estimator and how reconstruction of the anisotropic rotation angle is impacted by the effects of lensing and patchy reionization.

Weak gravitational lensing of the CMB occurs when the photons traveling towards our telescopes are gravitationally deflected by cosmological structure along their path. The photons we receive appear to originate from a position that differs by a small angle from their true origin. In real space, this deflection alters the T , Q , and

U maps of the CMB as

$$T^{\text{len}}(\hat{\mathbf{n}}) = T^{\text{prim}}(\hat{\mathbf{n}} + \nabla\phi(\hat{\mathbf{n}})),$$

$$(Q^{\text{len}} \pm iU^{\text{len}})(\hat{\mathbf{n}}) = (Q^{\text{prim}} \pm iU^{\text{prim}})(\hat{\mathbf{n}} + \nabla\phi(\hat{\mathbf{n}})), \quad (1)$$

where $\hat{\mathbf{n}}$ is the line-of-sight direction, ϕ is the CMB lensing potential, the ‘prim’ superscript identifies the primordial CMB maps, and the ‘len’ superscript denotes the lensed maps. The deflection of fluctuations leads to non-stationary statistics of observed maps giving off-diagonal mode coupling in harmonic space. The induced mode coupling can be used to reconstruct a map of the lensing potential [36].

For a field θ that distorts CMB maps, the off-diagonal mode coupling is of the form

$$\langle X(\ell_1)Y(\ell_2) \rangle_{\text{CMB}} = f_{XY}^\theta(\ell_1, \ell_2)\theta(\ell), \quad (2)$$

for $\ell_1 \neq \ell_2$ and $\ell = \ell_1 + \ell_2$, where X and Y are CMB temperature or polarization fluctuations $X, Y \in \{T, E, B\}$. The general form of the quadratic estimator for any such distortion field can be written as

$$\hat{\theta}_{XY}(\ell) = N_{XY}^\theta(\ell) \int \frac{d^2\ell_1}{(2\pi)^2} X^{\text{obs}}(\ell_1) Y^{\text{obs}}(\ell_2) F_{XY}^\theta(\ell_1, \ell_2), \quad (3)$$

where the superscript ‘obs’ refers to the observed maps including instrumental noise, $N_{XY}^\theta(\ell)$ is a normalization constant chosen to make the estimator unbiased ($\langle \hat{\theta}_{XY}(\ell) \rangle_{\text{CMB}} = \theta(\ell)$, and F_{XY}^θ is a filter (often chosen to minimize the variance of the estimator). The CMB noise is characterized by power spectra given here by

$$C_\ell^{TT, \text{noise}} = \Delta_T^2 e^{\ell^2 \theta_{\text{FWHM}}^2 / (8 \ln 2)},$$

$$C_\ell^{EE, \text{noise}} = C_\ell^{BB, \text{noise}} = \Delta_P^2 e^{\ell^2 \theta_{\text{FWHM}}^2 / (8 \ln 2)}. \quad (4)$$

We will focus on the EB quadratic estimator for which the minimum variance filter is

$$F_{EB}^\theta(\ell_1, \ell_2) = \frac{f_{EB}^\theta(\ell_1, \ell_2)}{C_{\ell_1}^{EE, \text{obs}} C_{\ell_2}^{BB, \text{obs}}}, \quad (5)$$

The normalization for the minimum variance EB estimator is

$$N_{EB}^\theta(\ell) = \left[\int \frac{d^2\ell_1}{(2\pi)^2} f_{EB}^\theta(\ell_1, \ell_2) F_{EB}^\theta(\ell_1, \ell_2) \right]^{-1}. \quad (6)$$

For the minimum variance filter, the normalization also gives the noise power of the reconstructed map

$$\langle \hat{\theta}_{EB}(\ell) \hat{\theta}_{EB}(\ell') \rangle = (2\pi)^2 \delta(\ell + \ell') (C_\ell^{\theta\theta} + N_{EB}^\theta(\ell)). \quad (7)$$

The EB mode coupling induced by gravitational lensing deflection is

$$f_{EB}^\phi(\ell_1, \ell_2) = [C_{\ell_1}^{EE}(\ell_1 \cdot \ell) - C_{\ell_2}^{BB}(\ell_2 \cdot \ell)] \times \sin 2(\varphi(\ell_1) - \varphi(\ell_2)). \quad (8)$$

Cosmic polarization rotation leaves the temperature unchanged but mixes Q and U polarization

$$\begin{aligned} T^{\text{rot}} &= T^{\text{prim}} \\ Q^{\text{rot}}(\hat{\mathbf{n}}) &= Q^{\text{prim}}(\hat{\mathbf{n}}) \cos 2\alpha(\hat{\mathbf{n}}) - U^{\text{prim}}(\hat{\mathbf{n}}) \sin 2\alpha(\hat{\mathbf{n}}) \\ U^{\text{rot}}(\hat{\mathbf{n}}) &= Q^{\text{prim}}(\hat{\mathbf{n}}) \sin 2\alpha(\hat{\mathbf{n}}) + U^{\text{prim}}(\hat{\mathbf{n}}) \cos 2\alpha(\hat{\mathbf{n}}), \end{aligned} \quad (9)$$

which leads to an EB mode coupling of the form

$$f_{EB}^{\alpha}(\ell_1, \ell_2) = 2 [C_{\ell_1}^{EE} - C_{\ell_2}^{BB}] \cos 2(\varphi(\ell_1) - \varphi(\ell_2)). \quad (10)$$

Inhomogeneous reionization leads to a CMB optical depth that differs across the sky $\tau(\hat{\mathbf{n}})$, thereby causing a modulation of CMB fluctuations [62–73]

$$\begin{aligned} T^{\text{mod}}(\hat{\mathbf{n}}) &= T^{\text{prim}}(\hat{\mathbf{n}}) e^{-\tau(\hat{\mathbf{n}})} \\ (Q^{\text{mod}} \pm iU^{\text{mod}})(\hat{\mathbf{n}}) &= (Q^{\text{prim}} \pm iU^{\text{prim}})(\hat{\mathbf{n}}) e^{-\tau(\hat{\mathbf{n}})}. \end{aligned} \quad (11)$$

Patchy reionization also generates new polarization fluctuations through the scattering of remote temperature quadrupole anisotropies [62, 63], though we will neglect these scattering contributions here. The EB mode coupling from modulation is given by

$$f_{EB}^{\tau}(\ell_1, \ell_2) = [C_{\ell_1}^{EE} - C_{\ell_2}^{BB}] \sin 2(\varphi(\ell_1) - \varphi(\ell_2)). \quad (12)$$

Primordial B -mode polarization is not generated by density fluctuations at first order in perturbation theory, and the contribution from primordial gravitational waves is observationally constrained to be small [74]. Distortions of the CMB, such as deflection by gravitational lensing, cosmic polarization rotation, and modulation due to patchy reionization, convert E -mode polarization to B -mode polarization. As can be seen from Eqs. (5) and (6), the variance of the reconstruction of any distortion with the EB estimator depends upon the total observed B -mode power, including the secondary B -mode polarization produced by distortion fields. Lensing-induced B modes make the dominant contribution to the observed B -mode power in the absence of noise and foregrounds on the scales of interest.

The variance of reconstructions of distortion fields can thereby be improved if we can reverse the effects of all such distortions, and in particular if we can remove the additional B -mode power that these effects cause. For example, the B -mode power produced by gravitational lensing of E modes (to leading order in the gradient expansion) is

$$\begin{aligned} C_{\ell}^{BB, \text{len}} &= \int \frac{d^2 \ell_1}{(2\pi)^2} [(\ell_1 \cdot \ell_2) \sin 2(\varphi(\ell_1) - \varphi(\ell_2))]^2 \\ &\times C_{\ell_2}^{\phi\phi} C_{\ell_1}^{EE}. \end{aligned} \quad (13)$$

The B -mode power can be reduced by removing an estimate of the B modes constructed from the measurement of the E -mode polarization and the reconstructed lensing

potential

$$\begin{aligned} C_{\ell}^{BB, \text{len, res}} &= \int \frac{d^2 \ell_1}{(2\pi)^2} [(\ell_1 \cdot \ell_2) \sin 2(\varphi(\ell_1) - \varphi(\ell_2))]^2 \\ &\times C_{\ell_2}^{\phi\phi} C_{\ell_1}^{EE} \left[1 - \frac{C_{\ell_2}^{\phi\phi}}{C_{\ell_2}^{\phi\phi, \text{obs}}} \frac{C_{\ell_1}^{EE}}{C_{\ell_1}^{EE, \text{obs}}} \right], \end{aligned} \quad (14)$$

where $C_{\ell}^{\phi\phi, \text{obs}} = C_{\ell}^{\phi\phi} + N^{\phi}(\ell)$. Using the reduced B -mode power, one can obtain a lower variance estimate of the lensing field with the EB estimator, which can subsequently be used to remove a larger fraction of the lensed B modes [47]. Iterating this procedure to convergence gives a variance on the reconstructed lensing field which matches closely with the maximum likelihood reconstruction [47, 54].

A similar procedure can be followed to demodulate the CMB, giving

$$\begin{aligned} C_{\ell}^{BB, \text{mod, res}} &= \int \frac{d^2 \ell_1}{(2\pi)^2} [\sin 2(\varphi(\ell_1) - \varphi(\ell_2))]^2 \\ &\times C_{\ell_2}^{\tau\tau} C_{\ell_1}^{EE} \left[1 - \frac{C_{\ell_2}^{\tau\tau}}{C_{\ell_2}^{\tau\tau, \text{obs}}} \frac{C_{\ell_1}^{EE}}{C_{\ell_1}^{EE, \text{obs}}} \right], \end{aligned} \quad (15)$$

and also to derotate the CMB

$$\begin{aligned} C_{\ell}^{BB, \text{rot, res}} &= \int \frac{d^2 \ell_1}{(2\pi)^2} [\cos 2(\varphi(\ell_1) - \varphi(\ell_2))]^2 \\ &\times 4C_{\ell_2}^{\alpha\alpha} C_{\ell_1}^{EE} \left[1 - \frac{C_{\ell_2}^{\alpha\alpha}}{C_{\ell_2}^{\alpha\alpha, \text{obs}}} \frac{C_{\ell_1}^{EE}}{C_{\ell_1}^{EE, \text{obs}}} \right]. \end{aligned} \quad (16)$$

If we treat the effects of lensing, modulation, and rotation as entirely independent, the observed B -mode power spectrum takes the form $C_{\ell}^{BB, \text{obs}} \simeq C_{\ell}^{BB, \text{prim}} + C_{\ell}^{BB, \text{len}} + C_{\ell}^{BB, \text{mod}} + C_{\ell}^{BB, \text{rot}} + C_{\ell}^{BB, \text{noise}}$. The presence of lensing-induced B -mode power increases the variance of the modulation and rotation reconstruction, and similarly for each contribution to the total B -mode spectrum. We can improve the reconstruction of each field by delensing, demodulating, and derotating to remove B -mode power. This whole procedure of lensing reconstruction, delensing, modulation reconstruction, demodulation, polarization rotation reconstruction, and derotation can be iterated to convergence to estimate the maximum likelihood reconstruction noise in the presence of all three effects.

This iterative procedure is only an approximation, since it fails to account for the fact that the effects of lensing, modulation, and rotation are not independent. In the simulations we employ here, the polarization maps are first anisotropically rotated by α , then modulated by τ , and finally lensed by ϕ . The effect of this series of deformations is not fully captured by a simple sum of the B -mode power from each effect applied in isolation. Furthermore, the real physical effects are not so distinctly

separated in time, though we treat them as such for simplicity in the simulations we present.

Implementing an iterative reconstruction scheme at map level is a non-trivial task, though iterative lensing reconstruction schemes have been developed [15, 50, 53, 56, 58], and internal delensing of CMB maps has been demonstrated on data from Planck [50, 75], the Atacama Cosmology Telescope [76], and the South Pole Telescope [77]. Map-level iterative reconstruction of multiple sources of statistical anisotropy poses additional challenges, since one must deal at each step with biases that may arise. For example, the presence of lensing leads to a bias in the reconstruction of patchy reionization maps [46] and cosmic polarization rotation spectra [42, 78]. For the purpose of this work, we do not attempt to construct a map-level iterative reconstruction scheme based on the quadratic estimator. Instead, we use the scheme described above at the level of the power spectra to estimate the reconstruction noise that would be expected assuming the existence of a map-level iterative reconstruction procedure that was able to deal with these complications without a significant impact on the resulting reconstruction noise. In Sec. IV, we compare the noise of the cosmic polarization rotation reconstruction derived from this idealized iterative procedure to that obtained from ResUNet-CMB.

III. RESUNET-CMB ARCHITECTURE AND METHODS

Convolutional neural networks (CNNs) are a type of deep learning network often used for computer vision tasks. Since the data obtained from many cosmological observations can be processed as two- and three-dimensional images, CNNs are well-suited to several aspects of cosmological simulations and data analysis. For example, some current CNN applications in cosmology include producing full-sky CMB simulations [79], identification of HII regions in reionization [80], analysis of dark matter substructure [81, 82], and cosmic velocity field reconstruction [83].

In this paper we use a particular type of CNN, a ResUNet [84–86], to reconstruct anisotropic cosmic birefringence imprinted in the polarization of simulated CMB maps in the presence of both lensing and patchy reionization effects. As shown in Refs. [60, 61], ResUNets are capable of detecting and recovering the distortions in simulated CMB maps at high fidelity.

In this section we will review the design and training of the ResUNet-CMB architecture introduced in Ref. [61], noting the modifications we made to the network and data pipeline in order to apply the framework to the new task of reconstructing cosmic birefringence.

A. ResUNet-CMB

We start from the ResUNet-CMB network described in Ref. [61] which is a modified version of the ResUNet from Ref. [60]. We extend the architecture in order to simultaneously reconstruct anisotropic cosmic polarization rotation α , the lensing convergence κ , the patchy reionization modulation τ , and the primordial E -mode polarization E^{prim} . The ResUNet-CMB architecture¹ is shown in Fig. 1. The network is implemented using the Keras package [87] of Tensorflow 2.0 [88].

The inputs to the network are simulated maps of the rotated, modulated, and lensed polarization plus noise described with Stokes parameters (Q^{obs} , U^{obs}) concatenated along the color channel. The network has four output maps, (α^{biased} , κ^{biased} , τ^{biased} , E^{biased}), each with its own branch. The superscripts ‘biased’ indicate that the outputs of the network exhibit a multiplicative bias, the correction of which will be described in Sec. IV.

Only two outputs are shown in Fig. 1 for clarity, and the other output branches have an identical structure. For the remainder of this paper we will refer to the network used in Ref. [61] as the 3-output network and that in the current paper as the 4-output network. The 4-output network has a total of 5 494 028 trainable parameters, 206 915 more than the 3-output network.

The primary building block of the network is the convolution block shown in the legend of Fig. 1. Each convolution block, apart from the first and last blocks in the network, consists of a dropout layer with a drop value of 0.3, a convolutional layer using ‘same’ padding, an activation layer using the scaled exponential linear unit (SELU) function [89], and a batch normalization layer in that order. The first convolution block does not contain a dropout layer in order to prevent the irreversible loss of initial information. In the final convolution block of each output branch, the activation layer is set to a linear activation function instead of SELU.

Accuracy of the predictions from CNNs can be negatively affected with increasing network depth [90–92]. Residual connections, which can help mitigate this problem, take the input of a convolution block and add it element-wise to the output of a different convolution block [91, 93, 94]. In the ResUNet-CMB architecture, the input of a convolution block is added element-wise to the output of the subsequent block. The residual connection over two convolution blocks is called a residual block and is shown in the legend of Fig. 1. The first residual connection of the network is from the concatenated (Q^{obs} , U^{obs}) layer to the input of the third convolution block. The next residual connection is from the input of the third convolution block to that of the fifth convolution block. This pattern of a connection every two

¹ The code for the updated ResUNet-CMB architecture and the new data pipeline is located at <https://github.com/EEGuzman/resunet-cmb>.




FIG. 1. ResUNet-CMB architecture showing two of four output branches and with all residual connections excluded for visual clarity. The input layer consists of the $(Q^{\text{obs}}, U^{\text{obs}})$ CMB polarization maps concatenated along the color channel. The maps $(\alpha^{\text{biased}}, \kappa^{\text{biased}}, \tau^{\text{biased}}, E^{\text{biased}})$ are the output of the final batch normalization layers of their respective branches. The numbers appearing on the bottom of each convolutional layer represent the number of filters that layer contains and the image size is displayed on side of every third batch normalization layer. In the encoder phase down-sampling is done by changing the stride of convolution layers to a value of 2. (The graphic was made with publicly available code from <https://github.com/HarisIqbal88/PlotNeuralNet>.)

blocks continues with the exception of the final convolution block in each output branch where no residual connection is placed. Like the architecture from Ref. [61], we include a batch normalization layer in each residual connection, placed after a convolutional layer if one is present, since we found this additional layer reduced validation loss across all outputs.

The network includes skip connections in addition to residual connections. The skip connections are concatenations along the color channel of the output of convolution blocks in the encoder phase of the network with the input of convolution blocks of the same size in the decoder phase of the network [95]. The skip connections help mitigate some localization information lost during down-sampling by providing high resolution information to convolution blocks in the decoder phase.

B. Data pipeline and network training

We generate maps using the same set of cosmological parameters as in Ref. [61] in order to ease comparisons among the results. The fiducial cosmology is defined by: $H_0 = 67.9 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_b h^2 = 0.0222$, $\Omega_c h^2 = 0.118$,

$n_s = 0.962$, $\tau = 0.0943$, and $A_s = 2.21 \times 10^{-9}$. The primordial CMB and lensing power spectra resulting from the chosen cosmology are produced with the software CAMB² [96]. For the patchy reionization spectrum, $C_\ell^{\tau\tau}$, we use the $\bar{\tau} = 0.058$ model of Ref. [97]. We include anisotropic cosmic birefringence described by a scale-invariant power spectrum of the form $\ell^2 C_\ell^{\alpha\alpha}/(2\pi) = 0.014 \text{ deg}^2$, which is about a factor of 2 smaller than the current 95% upper limit from ACTPol [41] and SPT-pol [42].

We simulate two-dimensional maps that cover a $5^\circ \times 5^\circ$ patch of sky with 128×128 pixels by using a modified version of Orphics³. In total eight different types of maps are generated, $(Q^{\text{prim}}, U^{\text{prim}}, E^{\text{prim}}, \alpha, \tau, \kappa, Q^{\text{obs}}, U^{\text{obs}})$. To obtain $(Q^{\text{obs}}, U^{\text{obs}})$ we first take the maps $(Q^{\text{prim}}, U^{\text{prim}})$ and rotate them using α according to Eq. (9). Next, the $(Q^{\text{rot}}, U^{\text{rot}})$ maps are modulated by τ then lensed by κ . Lastly, we smooth the maps with a Gaussian beam, add a noise map, and apply a 1.5° cosine taper to the result.

² <https://camb.info>

³ <https://github.com/msyriac/orphics>

The procedure to create the training, validation, test, and prediction data sets as well as the details of the pre-processing of the training data and post-processing of the predictions follows closely the discussion of Ref. [61] modified only by the definition of the null maps and by including the additional output for α . We restate the key points here.

We initially generate 70 000 sets of the 8 maps described above for four different experimental configurations. We simulate the same set of experiments as in Ref. [61]: a noiseless experiment; one with $\Delta_T = 0.2 \mu\text{K-arcmin}$ and $\theta_{\text{FWHM}} = 1.0'$; and experiments with $\Delta_T = 1 \mu\text{K-arcmin}$ and $2 \mu\text{K-arcmin}$ each with $\theta_{\text{FWHM}} = 1.4'$. For each experiment, $\Delta_P = \sqrt{2}\Delta_T$. Note that due to the size and resolution of our maps, the beam scale is smaller than our pixel scale and plays little role. However, we make this choice to match the design of future experiments like CMB-S4 [16] and to ease comparisons with previous work [46, 60, 61]. Of the 70 000 sets of maps generated, for each noise level, 20% are randomly selected to be null maps. We define the null maps to be (Q, U) maps that are not rotated by α but are modulated by τ then lensed by κ . Note that this is different than the definition of null maps used in Ref. [61], where in that case both τ and κ were set to zero in the null maps.

The 70 000 simulated maps are then split into training, validation, and test sets with a ratio of 80:10:10, respectively. In addition to the above three data sets, we produce an additional set of 7000 simulations called the prediction set which does not include null maps. Most of the analyses reported in Sec. IV are derived from predictions made on the prediction set.

The training process also follows that of Ref. [61]. Training was done on a single Nvidia Tesla V100 32GB GPU with a batch size of 32. We use the Adam optimizer [98] set with default parameters and choose a mean squared error loss function. An initial learning rate of 0.25 is used with a 50% reduction in its value after three epochs of no improvement in total validation loss. Training is stopped if there is no decrease in total validation loss after 10 consecutive epochs, and the model with the lowest loss is saved. In total, four networks are trained – one for each experiment.

Network training time for each experiment varies due to the random initialization of the learning process. On average, the noiseless, $0.2 \mu\text{K-arcmin}$, and $1 \mu\text{K-arcmin}$ networks each take approximately 15 hours to train and converge after about 60 epochs while the $2 \mu\text{K-arcmin}$ network takes only 7 hours and about 30 epochs on average to finish training.

While the training in this paper was done on a V100 GPU, we found using a Nvidia Tesla P100 16GB GPU with a batch size of 16 was also viable. Training with a batch size of 16 increased training time by a few hours, but the VRAM usage peaked at approximately 13.2 GB for both 3- and 4-output ResUNet-CMB models. Once the model is trained, generating a single set of predictions

of the maps, $(\alpha, \kappa, \tau, E)$, on the V100 GPU with a batch size of 32 in the prediction function takes 0.005 seconds.

IV. RESULTS

In this section we present results for the reconstructed maps, spectra, and noise curves for predictions made from the 4-output architecture and assess the performance of ResUNet-CMB. We compare the results to the predictions of the quadratic estimator and also to the 3-output network described in Ref. [61].

As previously described in Refs. [61] and [60], all of the output maps generated by the network are biased. For the rest of this paper we will refer to the outputs of the network prior to bias correction as $(\alpha^{\text{biased}}, \kappa^{\text{biased}}, \tau^{\text{biased}}, E^{\text{biased}})$, and the predicted maps after applying the bias correction described below will be labeled with a hat, e.g. $\hat{\alpha}$.

A. Reconstructed signal

First we examine the α maps reconstructed by ResUNet-CMB. Figure 2 shows the map-level α^{biased} predictions compared to the noiseless truth map for all four experiments. The bottom row of residual maps in the figure are defined as the truth minus the predictions. The large scale features appear to be faithfully reconstructed at all four noise levels, while only the lowest noise levels capture structure on small scales. As expected, as the instrument noise increases, the reconstruction worsens leaving larger residuals. The residual power spectrum, which we define as the power spectrum of the residual maps in Fig. 2, is an average of 14% larger over the range of $56 < \ell < 560$ for the $1 \mu\text{K-arcmin}$ experiment than in the noiseless case.

Map-level results for the remaining three outputs, $(\kappa^{\text{biased}}, \tau^{\text{biased}}, E^{\text{biased}})$, are visually very similar to those found in Ref. [61] despite the additional source of statistical anisotropy. A more quantitative comparison of the predictions of the 4-output network and the 3-output network will be described in Sec. IV D.

Next, we present the reconstructed average pseudo-power spectra for α , $\langle C_\ell^{\alpha^{\text{biased}} \alpha^{\text{biased}}} \rangle$, in Fig. 3. To calculate $\langle C_\ell^{\alpha^{\text{biased}} \alpha^{\text{biased}}} \rangle$ for each noise level, we take the α predictions made on the prediction data set and find the power spectrum of each individual α^{biased} map. We then average over all 7000 spectra and divide the result by the mean of the squared taper applied to each map to get a rough window-corrected average spectrum. The true spectrum shown in Fig. 3 and below, is calculated by finding the power spectrum of the truth α maps in the prediction data set, averaging the result, and applying the same window correction. The peak of the reconstructed power is in the range $56 < \ell < 170$ for all four noise levels. At $\ell > 170$, the reconstructed power




FIG. 2. Example ResUNet-CMB predictions of anisotropic cosmic polarization rotation α^{biased} from fully trained networks for each of four experimental configurations along with the noiseless true α map. Residual maps are computed as the truth minus the prediction.




FIG. 3. Power spectra of α^{biased} prediction maps from the ResUNet-CMB network averaged over the 7000 simulated CMB maps from the prediction data set for each of the four experiments (solid lines) plotted against the average power spectrum of the true α maps (black dashed-dot).

decreases with increasing ℓ . As with the map-level results, the other outputs (κ^{biased} , τ^{biased} , E^{biased}) show reconstruction power similar to those from Ref. [61].

Finally, we compare the performance of the standard quadratic estimator in reconstructing the α power spectrum to the predictions of ResUNet-CMB. For the quadratic estimator we use the EB estimator for anisotropic cosmic birefringence, $\hat{\alpha}_{EB}(\ell)$ (see Eqs. (3), (5), (6), and (10)), which is implemented using the soft-

ware package **Symlens**⁴.

As previously mentioned, the direct outputs of ResUNet-CMB exhibit a multiplicative bias.

We must correct this bias in order to make comparisons with the predictions from the quadratic estimator. We follow the same steps that were used for the ResUNet lensing reconstruction by Ref. [60] and for patchy reionization in Ref. [61], modified appropriately for cosmic polarization rotation. The biased maps of each distortion field θ are rescaled by a factor

$$A_{\ell}^{\theta} = \left[\frac{\langle C_{\ell}^{\theta\theta\text{biased}} \rangle}{\langle C_{\ell}^{\theta\theta} \rangle} \right]^{-1}, \quad (17)$$

where the $\langle \dots \rangle$ is an average over the entire prediction data set.

The final unbiased estimate of the anisotropic polarization rotation is given by

$$\hat{\alpha}(\ell) = A_{\ell}^{\alpha} \alpha^{\text{biased}}(\ell). \quad (18)$$

The reconstructed α power spectrum predicted by ResUNet-CMB is then

$$C_{\ell}^{\hat{\alpha}\hat{\alpha}} = (A_{\ell}^{\alpha})^2 (C_{\ell}^{\alpha^{\text{biased}}\alpha^{\text{biased}}}). \quad (19)$$

In Fig. 4 we compare the reconstructed α power spectra produced with the quadratic estimator and with

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




FIG. 4. Average of reconstructed α power spectra produced with ResUNet-CMB (blue) and quadratic estimator (red) computed with 7000 simulated CMB maps, plotted for four experimental configurations. We present results for maps rotated with two realizations of the true α , one using the nominal amplitude $\ell^2 C_\ell^{\alpha\alpha}/(2\pi) = 0.014 \text{ deg}^2$ (solid lines) and the other using a reduced amplitude $\ell^2 C_\ell^{\alpha\alpha}/(2\pi) = 0.0016 \text{ deg}^2$ (dotted). We also show the power spectra of the true α maps used for the nominal (black dash-dot) and reduced (black dashed) cases.

ResUNet-CMB for all four experimental configurations and for two choices of fiducial $C_\ell^{\alpha\alpha}$. We start by producing a single map of α with a scale invariant power spectrum $\ell^2 C_\ell^{\alpha\alpha}/(2\pi) = 0.014 \text{ deg}^2$. A second map of α is produced by rescaling the first map such that it has a reduced power spectrum $\ell^2 C_\ell^{\alpha\alpha}/(2\pi) = 0.0016 \text{ deg}^2$ but the same random fluctuations. A single realization of each of κ and τ are also produced. Next we generate 7000 realizations of the CMB polarization. Each of these 7000 polarization maps is then rotated, modulated, and lensed by the fixed (α, τ, κ) fields and noise is added. This results in 14 000 maps of $(Q^{\text{obs}}, U^{\text{obs}})$, 7000 for each amplitude of the α map. We apply the quadratic estimator to each observed map, the power spectrum of each reconstructed α is computed, and the power spectra are averaged over the 7000 CMB realizations for each amplitude of α to produce the red curves shown in Fig. 4. The fully trained ResUNet-CMB is applied to the same set of maps, except that a cosine taper is applied before predictions are made. The average of the window-corrected power spectra of the bias-corrected predictions are shown

in blue in Fig. 4. The process is repeated for each of the four experimental configurations we consider.

It is clear from Fig. 4 that ResUNet-CMB outperforms the standard quadratic estimator in reconstructing anisotropic cosmic birefringence, especially at low noise. The power spectra produced from the predictions of the quadratic estimator are almost independent of the true α spectrum, except at low experimental noise and on large scales $\ell < 200$. This is suggestive that the reconstructed power spectra from the quadratic estimator are dominated by noise rather than by signal for the configurations we considered here. On the other hand, for all but the highest noise experiment we considered, the predictions of ResUNet-CMB are clearly sensitive to the amplitude of the input signal. The ResUNet-CMB spectra also exhibit lower total (signal plus noise) power across a wide range of scales. For the $2 \mu\text{K-arcmin}$ case, the quadratic estimator and ResUNet-CMB perform nearly identically, with a reconstructed power that does not depend significantly on the signal amplitude, suggesting that the reconstruction is dominated by noise.




FIG. 5. Cosmic birefringence reconstruction noise power spectra for ResUNet-CMB (blue); the standard quadratic estimator using lensed, rotated, and modulated CMB spectra (red dashed); and a quadratic estimator using the iterative reconstruction technique described in Sec. II (red dotted). Also plotted is the theory α spectrum (black dash-dot).

We performed a set of tests to check whether A_ℓ^α depends on the cosmology that is used to train the network. We trained models for different amplitudes of the birefringence spectrum and compared the value of A_ℓ^α from each and found them to generally agree, though there is a mild dependence of A_ℓ^α on the input amplitude of $C_\ell^{\alpha\alpha}$, particularly at the lowest experimental noise levels. We also checked the dependence of A_ℓ^α on the values of other cosmological parameters. For significant changes to the assumed cosmology, there were non-trivial changes to A_ℓ^α . This is to be expected, since the reconstruction noise inevitably depends on the properties of the primordial CMB maps. This dependence on cosmology appears also in the normalization of the standard quadratic estimator.

When used in a likelihood analysis on real data, one could correct for the cosmology dependence of A_ℓ^α by accounting for how it depends upon the inputs used to train the network. A strategy to achieve this is discussed in Ref. [60] and briefly summarized here. First, the network could be trained for a set of cosmologies in the neighborhood of the best fit. The outputs of this set of trained models could be used to calculate numerical derivatives

of the A_ℓ^α with respect to the fiducial cosmological parameters or input spectra. It would then be possible to apply a correction at every sampled point of the likelihood, as is done with the Planck lensing likelihood, for example (see Appendix C of Ref. [99]). This allows one to avoid biases that may arise from cosmology dependence of A_ℓ^α without having to re-train the network for each likelihood sample.

B. Noise

We now specify the procedure for computing the noise of the ResUNet-CMB reconstruction of cosmic polarization rotation and compare to the reconstruction noise obtained with the quadratic estimator.

The α reconstruction noise for ResUNet-CMB is defined in terms of the bias-corrected power spectrum as

$$N_\ell^{\alpha\alpha} = \langle C_\ell^{\hat{\alpha}\hat{\alpha}} \rangle - \langle C_\ell^{\alpha\alpha} \rangle. \quad (20)$$

where the angle brackets again represent an average over the prediction data set. Using Eq. (20), we calculate the reconstruction noise power spectrum obtained with

ResUNet-CMB for each experiment, and we show the results in Fig. 5.

We also show in Fig. 5 two estimates of the reconstruction noise using the quadratic estimator. One estimate, labeled “Lensed CMB” uses the observed CMB spectra without delensing, demodulation, or derotation. In this case, the lensing-induced B modes act a source of noise for the sake of the birefringence reconstruction, and $N_\ell^{\alpha\alpha}$ does not decrease much with decreasing noise. For our choice of parameters, patchy reionization has only a very small effect on the reconstruction noise computed here. The other case, labeled “Iter. Delensed, Demodulated, and Derotated CMB” uses the iterative reconstruction procedure sketched in Sec. II. Decreasing the experimental noise gives significant improvements in the α reconstruction noise obtained with the iterative estimator. Both versions of the quadratic estimator noise are computed using $\ell_{\max} = 6100$ in order to match the range of angular scales used in the simulations.

For the noiseless and $0.2 \mu\text{K-arcmin}$ experiments the ResUNet-CMB reconstruction noise follows closely the noise curve for the iterative reconstruction across a wide range of ℓ values. At higher noise levels, the ResUNet-CMB reconstruction noise approaches that obtained from the idealized iterative procedure over a smaller range of ℓ , though it still outperforms the standard quadratic estimator with no delensing or derotation over a wide range of scales. These results indicate that the deep learning network is able to successfully disentangle and reconstruct the effects of α , κ , and τ simultaneously, in a nearly optimal way.

C. Null test

A null test was performed to ensure that the ResUNet-CMB α estimates are due to the presence of a true signal and are not simply an artifact of the network or training process. We produce 7000 simulated CMB polarization maps, (Q, U) , that are modulated by τ and lensed by κ but are not rotated (the α spectrum is set to zero). We add noise to these maps and use ResUNet-CMB to make predictions as previously described. The power spectra of the bias-corrected α predictions are averaged and the results for each experiment are plotted in Fig. 6. For comparison, we also plot the bias-corrected reconstructed power spectra predicted from maps with non-zero α (i.e. the bias-corrected version of the spectra shown in Fig. 3).

The reconstructed α power spectra from the null maps are consistent with the reconstruction noise spectra shown in Fig. 5. The bias-corrected power spectra predicted from maps with non-zero α are consistent with signal plus noise. We therefore see that ResUNet-CMB is performing as desired, and the network is recovering the true cosmic polarization rotation signal, rather than reproducing artifacts from the training process.




FIG. 6. Average bias-corrected reconstructed α power spectra of ResUNet-CMB predictions made with two different sets of simulated CMB maps: one set that includes modulation and lensing but zero rotation (null maps, solid) and on another set that includes rotation, modulation, and lensing (dotted). We also plot the true α spectrum used for the rotated maps (black dash-dot).

D. Comparison with the 3-output network

In this section we compare the reconstruction quality of the 3- and 4-output ResUNet-CMB networks by examining the reconstruction noise for the predictions they have in common, τ and κ . In Fig. 7 we show the reconstruction noise as calculated with Eq. (20), modified appropriately for τ and κ . The top row of figures features the $N_\ell^{\tau\tau}$ spectrum while the bottom row contains $N_\ell^{\kappa\kappa}$.

Generally, the reconstruction noise for the τ and κ predictions made by the 3- and 4-output networks are in very close agreement. The predictions of the $0.2 \mu\text{K-arcmin}$ experiment were particularly stable, with nearly identical reconstruction noise for the two networks over the entire ℓ range we consider. The noiseless experiment exhibited a modest increase in κ reconstruction noise for $\ell > 500$ with the addition of the α output branch.

The small scale τ predictions for the $1 \mu\text{K-arcmin}$ and $2 \mu\text{K-arcmin}$ experiments seemed to be most prominently affected by the addition of cosmic birefringence reconstruction. The 3-output network had been observed to perform worse than the quadratic estimator for these experiments on small angular scales, and the 4-output network seems to exhibit this breakdown at lower ℓ values. However, the τ reconstruction noise of the 3- and 4-output networks is essentially the same for $\ell < 1000$.

Figure 7 demonstrates that the τ and κ reconstruction have similar variance in the 3- and 4-output networks over a wide range of angular scales. Despite adding another source of statistical anisotropy to the polarization maps and adding an output branch to the network to incorporate its reconstruction, the predictions of τ and κ are mostly unaffected.




FIG. 7. Reconstruction noise power spectra of patchy reionization τ (top row) and lensing convergence κ (bottom row) for ResUNet-CMB 3-output (blue, cyan) and 4-output (coral, red) network architectures for each of the four experimental configurations we considered along with the fiducial theory spectra (black dash-dot).

V. CONCLUSION

We extended the ResUNet-CMB architecture from Ref. [61] by adding an additional output branch to the network for the reconstruction of anisotropic cosmic polarization rotation, $\alpha(\hat{\mathbf{n}})$. We showed that the ResUNet-CMB network is capable of α reconstruction from modulated, lensed, and rotated CMB polarization maps with a lower variance than what is achievable with the standard quadratic estimator over a wide range of angular scales. We also found the ResUNet-CMB estimates of α have a variance that approaches that of an idealized iterative reconstruction method. The ResUNet-CMB network is able to simultaneously reconstruct three sources of statistical anisotropy, (α, κ, τ) , while out-performing the standard quadratic estimator for each output.

The success of ResUNet-CMB with reconstruction of cosmic polarization rotation demonstrates the flexibility and generalizability of the architecture. We showed how the network faithfully reconstructs the α signal even as we vary the amplitude of the true α applied to the maps. Unlike gravitational lensing, for which the power spectrum is mostly determined by cosmological parameters that are already well-measured, cosmic polarization rotation results from physics beyond the standard cosmological model, and so we only have observational upper bounds on its amplitude. It is therefore encouraging that ResUNet-CMB produces results consistent with noise when no cosmic birefringence is present in the maps, implying that it would be a viable tool to search for a first detection of anisotropic cosmic polarization rotation. Furthermore, the patchy reionization signal that we focused on in Ref. [61] can be treated as a small perturbation to the lensed CMB, since only a statistical detection would be possible even for a noiseless experiment. On

the other hand, the cosmic birefringence we study here produces much larger effects on the CMB, allowing for a high signal-to-noise reconstruction of the rotation field on large angular scales. The flexibility of ResUNet-CMB is reinforced by the fact that the network is successful at reconstructing both patchy reionization and cosmic birefringence despite the differing size of their effects for the models we used.

In order to convert the 3-output model of Ref. [61] to the 4-output model studied here, the only modification to the ResUNet-CMB architecture was the addition of three convolution blocks that form an extra branch for the α map output. The data pipeline was altered to include the effects of cosmic birefringence, but no other major changes were required. The results show the variance on the reconstruction of α is smaller than that of the quadratic estimator, and the variances of the τ and κ predictions in the 4-output network are very similar to those of the 3-output network where cosmic polarization rotation was absent.

The success of ResUNet-CMB applied to the new task of cosmic birefringence reconstruction suggests that additional inputs and outputs can straightforwardly be added to the network. Experience here shows that these additions need not drastically increase VRAM requirements. Additional inputs and outputs would enable ResUNet-CMB to be applied to searches for other sources of secondary CMB anisotropies, including those that are most easily observed by cross-correlating CMB maps with other cosmological surveys.

One limitation of ResUNet-CMB is that it produces point estimates of reconstructed maps with uncertainties that are somewhat opaque. As a result, constraining cosmological parameters using the outputs of ResUNet-CMB requires some additional care. One path forward

would be to incorporate ResUNet-CMB into a Bayesian framework capable of quantifying uncertainties like the one in Ref. [58]. Alternatively, ResUNet-CMB could be used as the core component to construct a Bayesian neural network [100] that is capable of producing a probabilistic distribution of outputs. Bayesian neural networks have a growing presence in cosmology with applications such as parameter and uncertainty estimation with the CMB [101, 102] and with strong gravitational lensing [103, 104]. Incorporating ResUNet-CMB into a Bayesian neural network would allow for predictions with well-characterized posterior probabilities.

Building on the success of Refs. [60] and [61], we showed that the ResUNet-CMB architecture can be straightforwardly modified to reconstruct anisotropic cosmic polarization rotation with better performance

than the quadratic estimator. The results presented here give strong motivation to pursue additional applications of ResUNet-CMB.

ACKNOWLEDGMENTS

This work is supported by the US Department of Energy Office of Science under grant no. DE-SC0010129. EG is supported by a Southern Methodist University Computational Science and Engineering fellowship. Computations were carried out on ManeFrame II, a shared high-performance computing cluster at Southern Methodist University.

-
- [1] **Simons Observatory** Collaboration, P. Ade *et al.*, “The Simons Observatory: Science goals and forecasts,” *JCAP* **02** (2019) 056, [arXiv:1808.07445 \[astro-ph.CO\]](#).
 - [2] H. Hui *et al.*, “BICEP Array: a multi-frequency degree-scale CMB polarimeter,” *Proc. SPIE Int. Soc. Opt. Eng.* **10708** (2018) 1070807, [arXiv:1808.00568 \[astro-ph.IM\]](#).
 - [3] M. Aravena *et al.*, “The CCAT-Prime Submillimeter Observatory,” [arXiv:1909.02587 \[astro-ph.IM\]](#).
 - [4] **CMB-S4** Collaboration, K. N. Abazajian *et al.*, “CMB-S4 Science Book, First Edition,” [arXiv:1610.02743 \[astro-ph.CO\]](#).
 - [5] M. Hazumi *et al.*, “LiteBIRD: A Satellite for the Studies of B-Mode Polarization and Inflation from Cosmic Background Radiation Detection,” *J. Low Temp. Phys.* **194** no. 5-6, (2019) 443–452.
 - [6] **NASA PICO** Collaboration, S. Hanany *et al.*, “PICO: Probe of Inflation and Cosmic Origins,” [arXiv:1902.10541 \[astro-ph.IM\]](#).
 - [7] N. Sehgal *et al.*, “CMB-HD: An Ultra-Deep, High-Resolution Millimeter-Wave Survey Over Half the Sky,” [arXiv:1906.10134 \[astro-ph.CO\]](#).
 - [8] M. Kamionkowski, A. Kosowsky, and A. Stebbins, “Statistics of cosmic microwave background polarization,” *Phys. Rev. D* **55** (1997) 7368–7388, [arXiv:astro-ph/9611125](#).
 - [9] U. Seljak and M. Zaldarriaga, “Signature of gravity waves in polarization of the microwave background,” *Phys. Rev. Lett.* **78** (1997) 2054–2057, [arXiv:astro-ph/9609169](#).
 - [10] A. Lewis and A. Challinor, “Weak gravitational lensing of the CMB,” *Phys. Rept.* **429** (2006) 1–65, [arXiv:astro-ph/0601594](#).
 - [11] M. Zaldarriaga and U. Seljak, “Gravitational lensing effect on cosmic microwave background polarization,” *Phys. Rev. D* **58** (1998) 023003, [arXiv:astro-ph/9803150](#).
 - [12] A. Lewis, A. Challinor, and N. Turok, “Analysis of CMB polarization on an incomplete sky,” *Phys. Rev. D* **65** (2002) 023505, [arXiv:astro-ph/0106536](#).
 - [13] L. Knox and Y.-S. Song, “A Limit on the detectability of the energy scale of inflation,” *Phys. Rev. Lett.* **89** (2002) 011303, [arXiv:astro-ph/0202286](#).
 - [14] M. Kesden, A. Cooray, and M. Kamionkowski, “Separation of gravitational wave and cosmic shear contributions to cosmic microwave background polarization,” *Phys. Rev. Lett.* **89** (2002) 011304, [arXiv:astro-ph/0202434](#).
 - [15] U. Seljak and C. M. Hirata, “Gravitational lensing as a contaminant of the gravity wave signal in CMB,” *Phys. Rev. D* **69** (2004) 043005, [arXiv:astro-ph/0310163](#).
 - [16] K. Abazajian *et al.*, “CMB-S4 Science Case, Reference Design, and Project Plan,” [arXiv:1907.04473 \[astro-ph.IM\]](#).
 - [17] **CMB-S4** Collaboration, K. Abazajian *et al.*, “CMB-S4: Forecasting Constraints on Primordial Gravitational Waves,” [arXiv:2008.12619 \[astro-ph.CO\]](#).
 - [18] S. M. Carroll, G. B. Field, and R. Jackiw, “Limits on a Lorentz and Parity Violating Modification of Electrodynamics,” *Phys. Rev. D* **41** (1990) 1231.
 - [19] D. Harari and P. Sikivie, “Effects of a Nambu-Goldstone boson on the polarization of radio galaxies and the cosmic microwave background,” *Phys. Lett. B* **289** (1992) 67–72.
 - [20] S. M. Carroll, “Quintessence and the rest of the world,” *Phys. Rev. Lett.* **81** (1998) 3067–3070, [arXiv:astro-ph/9806099](#).
 - [21] A. Lue, L.-M. Wang, and M. Kamionkowski, “Cosmological signature of new parity violating interactions,” *Phys. Rev. Lett.* **83** (1999) 1506–1509, [arXiv:astro-ph/9812088](#).
 - [22] A. Kosowsky and A. Loeb, “Faraday rotation of microwave background polarization by a primordial magnetic field,” *Astrophys. J.* **469** (1996) 1–6, [arXiv:astro-ph/9601055](#).
 - [23] D. J. E. Marsh, “Axion Cosmology,” *Phys. Rept.* **643** (2016) 1–79, [arXiv:1510.07633 \[astro-ph.CO\]](#).
 - [24] E. G. M. Ferreira, “Ultra-Light Dark Matter,” [arXiv:2005.03254 \[astro-ph.CO\]](#).
 - [25] **QUaD** Collaboration, E. Y. S. Wu *et al.*, “Parity Violation Constraints Using Cosmic Microwave

- Background Polarization Spectra from 2006 and 2007 Observations by the QUaD Polarimeter,” *Phys. Rev. Lett.* **102** (2009) 161302, [arXiv:0811.0618 \[astro-ph\]](#).
- [26] **WMAP** Collaboration, E. Komatsu *et al.*, “Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation,” *Astrophys. J. Suppl.* **192** (2011) 18, [arXiv:1001.4538 \[astro-ph.CO\]](#).
- [27] B. G. Keating, M. Shimon, and A. P. S. Yadav, “Self-Calibration of CMB Polarization Experiments,” *Astrophys. J. Lett.* **762** (2012) L23, [arXiv:1211.5734 \[astro-ph.CO\]](#).
- [28] Y. Minami, H. Ochi, K. Ichiki, N. Katayama, E. Komatsu, and T. Matsumura, “Simultaneous determination of the cosmic birefringence and miscalibrated polarization angles from CMB experiments,” *PTEP* **2019** no. 8, (2019) 083E02, [arXiv:1904.12440 \[astro-ph.CO\]](#).
- [29] Y. Minami and E. Komatsu, “Simultaneous determination of the cosmic birefringence and miscalibrated polarization angles II: Including cross frequency spectra,” *PTEP* **2020** no. 10, (2020) 103E02, [arXiv:2006.15982 \[astro-ph.CO\]](#).
- [30] B. D. Sherwin and T. Namikawa, “Cosmic birefringence tomography and calibration-independence with reionization signals in the CMB,” [arXiv:2108.09287 \[astro-ph.CO\]](#).
- [31] **Planck** Collaboration, N. Aghanim *et al.*, “Planck 2018 results. III. High Frequency Instrument data processing and frequency maps,” *Astron. Astrophys.* **641** (2020) A3, [arXiv:1807.06207 \[astro-ph.CO\]](#).
- [32] Y. Minami and E. Komatsu, “New Extraction of the Cosmic Birefringence from the Planck 2018 Polarization Data,” *Phys. Rev. Lett.* **125** no. 22, (2020) 221301, [arXiv:2011.11254 \[astro-ph.CO\]](#).
- [33] M. Kamionkowski, “How to De-Rotate the Cosmic Microwave Background Polarization,” *Phys. Rev. Lett.* **102** (2009) 111302, [arXiv:0810.1286 \[astro-ph\]](#).
- [34] A. P. S. Yadav, R. Biswas, M. Su, and M. Zaldarriaga, “Constraining a spatially dependent rotation of the Cosmic Microwave Background Polarization,” *Phys. Rev. D* **79** (2009) 123009, [arXiv:0902.4466 \[astro-ph.CO\]](#).
- [35] V. Gluscevic, M. Kamionkowski, and A. Cooray, “De-Rotation of the Cosmic Microwave Background Polarization: Full-Sky Formalism,” *Phys. Rev. D* **80** (2009) 023510, [arXiv:0905.1687 \[astro-ph.CO\]](#).
- [36] W. Hu and T. Okamoto, “Mass reconstruction with cmb polarization,” *Astrophys. J.* **574** (2002) 566–574, [arXiv:astro-ph/0111606](#).
- [37] V. Gluscevic, D. Hanson, M. Kamionkowski, and C. M. Hirata, “First CMB Constraints on Direction-Dependent Cosmological Birefringence from WMAP-7,” *Phys. Rev. D* **86** (2012) 103529, [arXiv:1206.5546 \[astro-ph.CO\]](#).
- [38] **POLARBEAR** Collaboration, P. A. R. Ade *et al.*, “POLARBEAR Constraints on Cosmic Birefringence and Primordial Magnetic Fields,” *Phys. Rev. D* **92** (2015) 123509, [arXiv:1509.02461 \[astro-ph.CO\]](#).
- [39] G.-C. Liu and K.-W. Ng, “Axion Dark Matter Induced Cosmic Microwave Background *B*-modes,” *Phys. Dark Univ.* **16** (2017) 22–25, [arXiv:1612.02104 \[astro-ph.CO\]](#).
- [40] **BICEP2, Keck Array** Collaboration, P. A. R. Ade *et al.*, “BICEP2 / Keck Array IX: New bounds on anisotropies of CMB polarization rotation and implications for axionlike particles and primordial magnetic fields,” *Phys. Rev. D* **96** no. 10, (2017) 102003, [arXiv:1705.02523 \[astro-ph.CO\]](#).
- [41] T. Namikawa *et al.*, “Atacama Cosmology Telescope: Constraints on cosmic birefringence,” *Phys. Rev. D* **101** no. 8, (2020) 083527, [arXiv:2001.10465 \[astro-ph.CO\]](#).
- [42] **SPT** Collaboration, F. Bianchini *et al.*, “Searching for Anisotropic Cosmic Birefringence with Polarization Data from SPTpol,” *Phys. Rev. D* **102** no. 8, (2020) 083504, [arXiv:2006.08061 \[astro-ph.CO\]](#).
- [43] L. Pogosian, M. Shimon, M. Mewes, and B. Keating, “Future CMB constraints on cosmic birefringence and implications for fundamental physics,” *Phys. Rev. D* **100** no. 2, (2019) 023507, [arXiv:1904.07855 \[astro-ph.CO\]](#).
- [44] D. Green, J. Meyers, and A. van Engelen, “CMB Delensing Beyond the B Modes,” *JCAP* **1712** no. 12, (2017) 005, [arXiv:1609.08143 \[astro-ph.CO\]](#).
- [45] W. R. Coulton, P. D. Meerburg, D. G. Baker, S. Hotinli, A. J. Duivenvoorden, and A. van Engelen, “Minimizing gravitational lensing contributions to the primordial bispectrum covariance,” *Phys. Rev. D* **101** no. 12, (2020) 123504, [arXiv:1912.07619 \[astro-ph.CO\]](#).
- [46] M. Su, A. P. Yadav, M. McQuinn, J. Yoo, and M. Zaldarriaga, “An Improved Forecast of Patchy Reionization Reconstruction with CMB,” [arXiv:1106.4313 \[astro-ph.CO\]](#).
- [47] K. M. Smith, D. Hanson, M. LoVerde, C. M. Hirata, and O. Zahn, “Delensing CMB Polarization with External Datasets,” *JCAP* **06** (2012) 014, [arXiv:1010.0048 \[astro-ph.CO\]](#).
- [48] N. Sehgal, M. S. Madhavacheril, B. Sherwin, and A. van Engelen, “Internal Delensing of Cosmic Microwave Background Acoustic Peaks,” *Phys. Rev. D* **95** no. 10, (2017) 103512, [arXiv:1612.03898 \[astro-ph.CO\]](#).
- [49] A. Baleato Lizancos, A. Challinor, and J. Carron, “Impact of internal-delensing biases on searches for primordial B-modes of CMB polarisation,” *JCAP* **03** (2021) 016, [arXiv:2007.01622 \[astro-ph.CO\]](#).
- [50] J. Carron, A. Lewis, and A. Challinor, “Internal delensing of Planck CMB temperature and polarization,” *JCAP* **05** (2017) 035, [arXiv:1701.01712 \[astro-ph.CO\]](#).
- [51] **POLARBEAR** Collaboration, S. Adachi *et al.*, “Internal delensing of Cosmic Microwave Background polarization *B*-modes with the POLARBEAR experiment,” *Phys. Rev. Lett.* **124** no. 13, (2020) 131301, [arXiv:1909.13832 \[astro-ph.CO\]](#).
- [52] **SPTpol, BICEP, Keck** Collaboration, P. A. R. Ade *et al.*, “A demonstration of improved constraints on primordial gravitational waves with delensing,” *Phys. Rev. D* **103** no. 2, (2021) 022004, [arXiv:2011.08163 \[astro-ph.CO\]](#).
- [53] C. M. Hirata and U. Seljak, “Analyzing weak lensing of the cosmic microwave background using the likelihood function,” *Phys. Rev. D* **67** (2003) 043001, [arXiv:astro-ph/0209489 \[astro-ph\]](#).

- [54] C. M. Hirata and U. Seljak, “Reconstruction of lensing from the cosmic microwave background polarization,” *Phys. Rev. D* **68** (2003) 083002, [arXiv:astro-ph/0306354](#).
- [55] J. Carron and A. Lewis, “Maximum a posteriori CMB lensing reconstruction,” *Phys. Rev. D* **96** no. 6, (2017) 063510, [arXiv:1704.08230 \[astro-ph.CO\]](#).
- [56] M. Millea, E. Anderes, and B. D. Wandelt, “Bayesian delensing of CMB temperature and polarization,” *Phys. Rev. D* **100** no. 2, (2019) 023509, [arXiv:1708.06753 \[astro-ph.CO\]](#).
- [57] J. Carron, “Optimal constraints on primordial gravitational waves from the lensed CMB,” *Phys. Rev. D* **99** no. 4, (2019) 043518, [arXiv:1808.10349 \[astro-ph.CO\]](#).
- [58] M. Millea, E. Anderes, and B. D. Wandelt, “Bayesian delensing delight: sampling-based inference of the primordial CMB and gravitational lensing,” [arXiv:2002.00965 \[astro-ph.CO\]](#).
- [59] P. Diego-Palazuelos, P. Vielva, E. Martínez-González, and R. Barreiro, “Comparison of delensing methodologies and assessment of the delensing capabilities of future experiments,” [arXiv:2006.12935 \[astro-ph.CO\]](#).
- [60] J. Caldeira, W. Wu, B. Nord, C. Avestruz, S. Trivedi, and K. Story, “DeepCMB: Lensing Reconstruction of the Cosmic Microwave Background with Deep Neural Networks,” *Astron. Comput.* **28** (2019) 100307, [arXiv:1810.01483 \[astro-ph.CO\]](#).
- [61] E. Guzman and J. Meyers, “Reconstructing Patchy Reionization with Deep Learning,” *Phys. Rev. D* **104** no. 4, (2021) 043529, [arXiv:2101.01214 \[astro-ph.CO\]](#).
- [62] C. Dvorkin and K. M. Smith, “Reconstructing Patchy Reionization from the Cosmic Microwave Background,” *Phys. Rev. D* **79** (2009) 043003, [arXiv:0812.1566 \[astro-ph\]](#).
- [63] C. Dvorkin, W. Hu, and K. M. Smith, “B-mode CMB Polarization from Patchy Screening during Reionization,” *Phys. Rev. D* **79** (2009) 107302, [arXiv:0902.4413 \[astro-ph.CO\]](#).
- [64] W. Hu, “Reionization revisited: secondary cmb anisotropies and polarization,” *Astrophys. J.* **529** (2000) 12, [arXiv:astro-ph/9907103](#).
- [65] M. G. Santos, A. Cooray, Z. Haiman, L. Knox, and C.-P. Ma, “Small - scale CMB temperature and polarization anisotropies due to patchy reionization,” *Astrophys. J.* **598** (2003) 756–766, [arXiv:astro-ph/0305471](#).
- [66] O. Zahn, M. Zaldarriaga, L. Hernquist, and M. McQuinn, “The Influence of non-uniform reionization on the CMB,” *Astrophys. J.* **630** (2005) 657–666, [arXiv:astro-ph/0503166](#).
- [67] M. McQuinn, S. R. Furlanetto, L. Hernquist, O. Zahn, and M. Zaldarriaga, “The Kinetic Sunyaev-Zel’dovich effect from reionization,” *Astrophys. J.* **630** (2005) 643–656, [arXiv:astro-ph/0504189](#).
- [68] O. Dore, G. Holder, M. Alvarez, I. T. Iliev, G. Mellema, U.-L. Pen, and P. R. Shapiro, “The Signature of Patchy Reionization in the Polarization Anisotropy of the CMB,” *Phys. Rev. D* **76** (2007) 043002, [arXiv:astro-ph/0701784](#).
- [69] N. Battaglia, A. Natarajan, H. Trac, R. Cen, and A. Loeb, “Reionization on Large Scales III: Predictions for Low- ℓ Cosmic Microwave Background Polarization and High- ℓ Kinetic Sunyaev-Zel’dovich Observables,” *Astrophys. J.* **776** (2013) 83, [arXiv:1211.2832 \[astro-ph.CO\]](#).
- [70] H. Park, P. R. Shapiro, E. Komatsu, I. T. Iliev, K. Ahn, and G. Mellema, “The Kinetic Sunyaev-Zel’dovich effect as a probe of the physics of cosmic reionization: the effect of self-regulated reionization,” *Astrophys. J.* **769** (2013) 93, [arXiv:1301.3607 \[astro-ph.CO\]](#).
- [71] M. A. Alvarez, “The Kinetic Sunyaev-Zel’dovich Effect From Reionization: Simulated Full-sky Maps at Arcminute Resolution,” *Astrophys. J.* **824** no. 2, (2016) 118, [arXiv:1511.02846 \[astro-ph.CO\]](#).
- [72] S. Paul, S. Mukherjee, and T. R. Choudhury, “Inevitable imprints of patchy reionization on the cosmic microwave background anisotropy,” *Mon. Not. Roy. Astron. Soc.* **500** no. 1, (2020) 232–246, [arXiv:2005.05327 \[astro-ph.CO\]](#).
- [73] T. R. Choudhury, S. Mukherjee, and S. Paul, “CMB constraints on a physical model of reionization,” [arXiv:2007.03705 \[astro-ph.CO\]](#).
- [74] **BICEP2, Keck Array** Collaboration, P. A. R. Ade *et al.*, “BICEP2 / Keck Array x: Constraints on Primordial Gravitational Waves using Planck, WMAP, and New BICEP2/Keck Observations through the 2015 Season,” *Phys. Rev. Lett.* **121** (2018) 221301, [arXiv:1810.05216 \[astro-ph.CO\]](#).
- [75] **Planck** Collaboration, N. Aghanim *et al.*, “Planck 2018 results. VIII. Gravitational lensing,” *Astron. Astrophys.* **641** (2020) A8, [arXiv:1807.06210 \[astro-ph.CO\]](#).
- [76] **ACT** Collaboration, D. Han *et al.*, “The Atacama Cosmology Telescope: delensed power spectra and parameters,” *JCAP* **01** (2021) 031, [arXiv:2007.14405 \[astro-ph.CO\]](#).
- [77] M. Millea *et al.*, “Optimal Cosmic Microwave Background Lensing Reconstruction and Parameter Estimation with SPTpol Data,” *Astrophys. J.* **922** no. 2, (2021) 259, [arXiv:2012.01709 \[astro-ph.CO\]](#).
- [78] T. Namikawa, “Testing parity-violating physics from cosmic rotation power reconstruction,” *Phys. Rev. D* **95** no. 4, (2017) 043523, [arXiv:1612.07855 \[astro-ph.CO\]](#).
- [79] D. Han, N. Sehgal, and F. Villaescusa-Navarro, “MillimeterDL: Deep Learning Simulations of the Microwave Sky,” [arXiv:2105.11444 \[astro-ph.CO\]](#).
- [80] M. Bianco, S. K. Giri, I. T. Iliev, and G. Mellema, “Deep learning approach for identification of HII regions during reionization in 21-cm observations,” [arXiv:2102.06713 \[astro-ph.IM\]](#).
- [81] S. Alexander, S. Gleyzer, H. Parul, P. Reddy, M. W. Toomey, E. Usai, and R. Von Klar, “Decoding Dark Matter Substructure without Supervision,” [arXiv:2008.12731 \[astro-ph.CO\]](#).
- [82] K. Vattis, M. W. Toomey, and S. M. Koushiappas, “Deep learning the astrometric signature of dark matter substructure,” [arXiv:2008.11577 \[astro-ph.CO\]](#).
- [83] Z. Wu, Z. Zhang, S. Pan, H. Miao, X. Wang, C. G. Sabiu, J. Forero-Romero, Y. Wang, and X.-D. Li, “Cosmic Velocity Field Reconstruction Using AI,” *Astrophys. J.* **913** no. 1, (2021) 2, [arXiv:2105.09450 \[astro-ph.CO\]](#).

- [84] B. Kayalibay, G. Jensen, and P. van der Smagt, “Cnn-based segmentation of medical imaging data,” [arXiv:1701.03056 \[cs.CV\]](#).
- [85] F. Milletari, N. Navab, and S.-A. Ahmadi, “V-net: Fully convolutional neural networks for volumetric medical image segmentation,” [arXiv:1606.04797 \[cs.CV\]](#).
- [86] Z. Zhang, Q. Liu, and Y. Wang, “Road extraction by deep residual u-net,” *IEEE Geoscience and Remote Sensing Letters* **15** no. 5, (May, 2018) 749–753. <http://dx.doi.org/10.1109/LGRS.2018.2802944>.
- [87] F. Chollet *et al.*, “Keras.” <https://github.com/fchollet/keras>, 2015.
- [88] M. Abadi *et al.*, “TensorFlow: Large-scale machine learning on heterogeneous systems,” 2015. <https://www.tensorflow.org/>.
- [89] G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter, “Self-normalizing neural networks,” [arXiv:1706.02515 \[cs.LG\]](#).
- [90] K. He and J. Sun, “Convolutional neural networks at constrained time cost,” [arXiv:1412.1710 \[cs.CV\]](#).
- [91] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” [arXiv:1512.03385 \[cs.CV\]](#).
- [92] R. K. Srivastava, K. Greff, and J. Schmidhuber, “Highway networks,” [arXiv:1505.00387 \[cs.LG\]](#).
- [93] K. He, X. Zhang, S. Ren, and J. Sun, “Identity mappings in deep residual networks,” [arXiv:1603.05027 \[cs.CV\]](#).
- [94] D. Balduzzi, M. Frean, L. Leary, J. Lewis, K. W.-D. Ma, and B. McWilliams, “The shattered gradients problem: If resnets are the answer, then what is the question?,” [arXiv:1702.08591 \[cs.NE\]](#).
- [95] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” [arXiv:1505.04597 \[cs.CV\]](#).
- [96] A. Lewis, A. Challinor, and A. Lasenby, “Efficient computation of CMB anisotropies in closed FRW models,” *Astrophys. J.* **538** (2000) 473–476, [arXiv:astro-ph/9911177](#).
- [97] A. Roy, A. Lapi, D. Spergel, and C. Baccigalupi, “Observing patchy reionization with future CMB polarization experiments,” *JCAP* **05** (2018) 014, [arXiv:1801.02393 \[astro-ph.CO\]](#).
- [98] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” [arXiv:1412.6980 \[cs.LG\]](#).
- [99] Planck Collaboration, P. A. R. Ade *et al.*, “Planck 2015 results. XV. Gravitational lensing,” *Astron. Astrophys.* **594** (2016) A15, [arXiv:1502.01591 \[astro-ph.CO\]](#).
- [100] Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” [arXiv:1506.02142 \[stat.ML\]](#).
- [101] H. J. Hortua, R. Volpi, D. Marinelli, and L. Malagò, “Parameter estimation for the cosmic microwave background with Bayesian neural networks,” *Phys. Rev. D* **102** no. 10, (2020) 103509, [arXiv:1911.08508 \[astro-ph.IM\]](#).
- [102] S. He, S. Ravanbakhsh, and S. Ho, “Analysis of cosmic microwave background with deep learning,” in *ICLR*. 2018.
- [103] L. Perreault Levasseur, Y. D. Hezaveh, and R. H. Wechsler, “Uncertainties in Parameters Estimated with Neural Networks: Application to Strong Gravitational Lensing,” *Astrophys. J. Lett.* **850** no. 1, (2017) L7, [arXiv:1708.08843 \[astro-ph.CO\]](#).
- [104] W. R. Morningstar, Y. D. Hezaveh, L. Perreault Levasseur, R. D. Blandford, P. J. Marshall, P. Putzky, and R. H. Wechsler, “Analyzing interferometric observations of strong gravitational lenses with recurrent and convolutional neural networks,” [arXiv:1808.00011 \[astro-ph.IM\]](#).