

Cosmic Microwave Background Recovery: A Graph-Based Bayesian Convolutional Network Approach

Jadie Adams^{1,2}, Steven Lu¹, Krzysztof M. Gorski¹, Graça Rocha¹, Kiri L. Wagstaff¹

¹ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109-8099, USA

² Scientific Computing and Imaging Institute, University of Utah, 201 Presidents' Cir, Salt Lake City, UT 84112, USA
jadie@sci.utah.edu, {you.lu,krzysztof.m.gorski,graca.m.rocha}@jpl.nasa.gov, wkiri@wkiri.com

Abstract

The cosmic microwave background (CMB) is a significant source of knowledge about the origin and evolution of our universe. However, observations of the CMB are contaminated by foreground emissions, obscuring the CMB signal and reducing its efficacy in constraining cosmological parameters. We employ deep learning as a data-driven approach to CMB cleaning from multi-frequency full-sky maps. In particular, we develop a graph-based Bayesian convolutional neural network based on the U-Net architecture that predicts cleaned CMB with pixel-wise uncertainty estimates. We demonstrate the potential of this technique on realistic simulated data based on the Planck mission. We show that our model accurately recovers the cleaned CMB sky map and resulting angular power spectrum while identifying regions of uncertainty. Finally, we discuss the current challenges and the path forward for deploying our model for CMB recovery on real observations.

Introduction

Cosmic Microwave Background (CMB) is remnant electromagnetic radiation from the early universe and a fundamental probe of cosmology. Small anisotropies, or irregularities, in the temperature fluctuations and polarization of CMB sky maps, can constrain models of structure formation and fundamental physics. However, such anisotropies are hidden by foreground emissions (i.e., contamination from Galactic emission, dust emission, synchrotron emission, etc.) as shown in Figure 1. Thus correctly discriminating between CMB and foreground emissions via observations made at multiple frequencies is crucial to cosmological parameter estimation.

Recent success in applying machine learning techniques to astrophysical tasks has motivated using artificial neural networks to remove foreground and reduce systematic errors in CMB cleaning. For such techniques to be maximally valuable for cosmological study, capturing uncertainty or an estimate of model confidence is crucial. In this work, we propose a graph-based Bayesian convolutional neural network (CNN) for recovering the underlying CMB signal from multi-frequency temperature maps of the full-sky in a probabilistic manner. To avoid edge effects imposed by project-

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

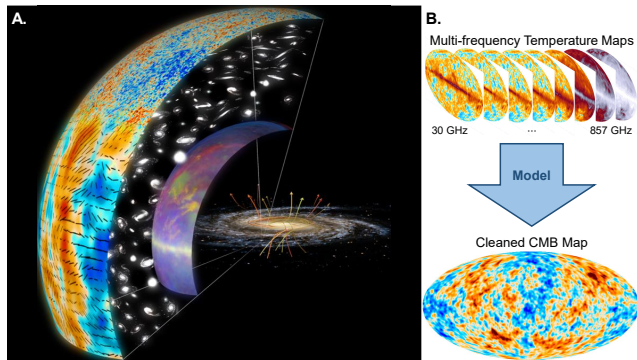


Figure 1: (A) The CMB. This graphic displays our galaxy at the center (including magnetic fields), the intergalactic medium, and the CMB in the outer shell. The color map shows temperature fluctuation and the vector lines represent the polarization amplitude and direction of the radiation. The inner shell displays foreground emissions which block telescopes from observing the CMB. Figure from a Keck Institute for Space Studies (KISS) workshop report (Rocha et al. 2019a). (B) CMB Cleaning Overview. CMB cleaning models take multi-frequency full-sky temperature maps as input and recover the cleaned CMB map.

ing spherical sky maps onto a two-dimensional plane, we employ HEALPix (Gorski et al. 2005) - a graph representation that enables local operations while minimizing discretization error. We reformulate the state-of-the-art image-to-image network, U-Net (Ronneberger, Fischer, and Brox 2015), as a Bayesian framework that operates on the graph-based representation and provides probabilistic predictions. The proposed model shows promise in accurately recovering CMB temperature maps and quantifying data and model-dependent uncertainty associated with the prediction.

The contributions of this paper are (1) a novel approach to CMB extraction that combines a graph-based representation with a Bayesian U-Net CNN, (2) initial results on low-resolution full-sky data with a comparison to existing methods for CMB extraction that do not employ machine learning, (3) an analysis of uncertainty quantification calibration and quality, and (4) a discussion of the strengths and weaknesses of the proposed approach.

Related Work

Sky component separation, i.e., extraction of primordial CMB anisotropy signals from multi-frequency measurements of the microwave sky, is the key data processing step to identify the cosmological model that best describes the observed properties of the universe. The most comprehensive comparison of component separation methods applied to data collected by the Planck satellite was provided by the Planck Collaboration (2014, 2016a, 2020). These component separation methods range from linear data combination techniques to multiparameter Bayesian model fitting (Rocha et al. 2019b). In this initial work, we compare our results to those derived using the simplest implementation of the Internal Linear Combination (ILC) technique (Bennett et al. 1992).

Recent work has explored employing neural networks as a data-driven approach to CMB extraction. Early approaches utilized fully-connected multilayer perceptrons, neglecting the inherent spatial correlation of the input maps (Nørgaard-Nielsen and Jørgensen 2008). CNNs retain spatial correlations but are formulated for Euclidean input. There has been progress in training CNNs using flattened 2D patches of full and partial-sky maps (Wang et al. 2022; Casas et al. 2022); however, these 2D projections are prone to edge artifacts. DeepSphere (Perraudin et al. 2019) presents an efficient spherical CNN that leverages the hierarchical HEALPix (Gorski et al. 2005) graph representations for pooling and performs convolution in the spectral domain. A recent model (Petroff et al. 2020) combines DeepSphere with a U-Net architecture (Ronneberger, Fischer, and Brox 2015) for full-sky map CMB recovery. Additionally, this model employs concrete dropout for UQ. We employ a similar technique, but in (Petroff et al. 2020) independent copies of the U-Net encoder are trained for each of the different frequency input maps, then features are combined and passed through the decoder. In our formulation, we treat the frequencies as channels of the same input and have a single encoder, making the model more scalable and preserving the cross-frequency correlation.

Deep neural network uncertainty quantification (UQ) is an active line of research. The most straightforward technique to acquire such measures of model confidence is to train an ensemble of networks and consider the variance in predictions (Lakshminarayanan, Pritzel, and Blundell 2017). However, ensembles are computationally expensive and are not guaranteed to produce diverse output. The M -heads framework was proposed to remedy this by using one common network that produces M different output (Lee et al. 2015; Rupprecht et al. 2017). Nevertheless, the M -heads method similarly scales poorly and lacks a principled theoretical foundation.

A more principled approach is Bayesian deep learning, which uses Bayes' rule to learn a distribution over network weights, providing both aleatoric and epistemic uncertainty quantification. Aleatoric (or data-dependent) uncertainty is the irreducible uncertainty inherent in the input data (for example, noise resulting from measurement imprecision). It can be learned as a function of the input data via MAP inference by adjusting the loss function (Kendall and Gal

2017). Epistemic (or model-dependent) uncertainty is the reducible uncertainty in model structure and parameters and is captured by placing a probability distribution over model weights. (Kendall and Gal 2017) showed that Monte Carlo (MC) dropout is a scalable, principled approach to approximate variational inference for acquiring epistemic uncertainty measures. Concrete dropout is a continuous relaxation of dropouts discrete masks that allows updating the optimization operation to automatically tune the dropout probabilities of a Bayesian network (Gal, Hron, and Kendall 2017). An alternative approach to dropout combines a conditional variational autoencoder (Kingma and Welling 2013) with a U-Net (Ronneberger, Fischer, and Brox 2015) to learn a conditional density model over output conditioned on the input (Kohl et al. 2018). However, this approach does not differentiate between aleatoric and epistemic uncertainty and requires training auxiliary networks making it challenging to scale. For these reasons, we elect to use concrete dropout with heteroscedastic uncertainty loss in adapting our graph-based CNN to be Bayesian for UQ.

Methods

Our model operates on full-sky maps, predicting the cleaned CMB from nine temperature observations (frequency bands in GHz: 30, 44, 70, 100, 143, 217, 353, 545, and 857). We use Bayesian graph-based CNN with U-Net architecture that preserves the spherical nature of the data and provides uncertainty reasoning.

Graphical Representation and Convolution

To address this, we perform convolutions in the spectral domain and represent sky maps as graphs using the Hierarchical Equal Area isoLatitude Pixelization (HEALPix (Gorski et al. 2005)) scheme as is done in DeepSphere (Perraudin et al. 2019). The Fourier transform of these HEALPix graph signals is the projection of the signal onto the set of eigenvectors of the graph Laplacian. In this work, we use the combinatorial Laplacian, which is defined as:

$$\mathbf{L} = \mathbf{D} - \mathbf{A} \quad (1)$$

where \mathbf{A} is the weighted adjacency matrix of the graph and \mathbf{D} is the diagonal degree matrix, which contains the sum of the weights of all the edges for each node on its diagonal.

Traditional graph convolutions are often performed in the spectral domain using spherical harmonics, and the computational complexity of such convolution methods is $O(N_{pix}^3)$. In convolutional neural networks (CNN), there are often tens to thousands of convolution operators, and each convolution operation is done in every forward pass during the training process. A computational complexity of $O(N_{pix}^3)$ convolution operation will almost make the training process intractable. In this work, we use the efficient convolution operation proposed in (Perraudin et al. 2019), which reduces the computational complexity of graph convolution to $O(N_{pix})$. This efficiency is obtained by representing each convolutional kernel, h , as a polynomial approximation using Chebyshev polynomials. Convolution for

the CMB signal, f , is then defined as:

$$h_{\phi}(\hat{\mathbf{L}})f = \sum_{k=0}^K \theta_k T_k(\hat{\mathbf{L}})f \quad (2)$$

where $T_k(\cdot)$ is the Chebyshev polynomial of degree k defined by the recursive relation, $\hat{\mathbf{L}}$ is the normalized form of the Laplacian matrix, and ϕ is the polynomial coefficients that are learned during the training process. In addition to the convolution operation, CNN architectures involve pooling and up-sampling layers. Pooling refers to an operation that summarizes (e.g., down-samples) a feature map by a pre-defined factor, and up-sampling is the inverse operation that expands a feature map by a given factor. This work uses maximum pooling and nearest neighbor up-sampling defined on HEALPix grids. The hierarchical tree structure of the HEALPix graph easily allows for such operations.

Network Architecture and Loss

We utilize a Bayesian formulation of the U-Net architecture (Ronneberger, Fischer, and Brox 2015) shown in Figure 2. It consists of an encoding and a decoding path. The encoding path consists of repeated operations of two Chebyshev convolutions with a Chebyshev polynomial order of 3, each followed by a 1-dimensional batch normalization and a rectified linear unit (ReLU) (Agarap 2018). After the two Chebyshev convolutions, the HEALPix-based max-pooling operation is applied to down-sample the feature maps by a factor of 4. Each depth of the decoding path consists of a HEALPix-based nearest neighbor up-sampling operation, concatenation with the corresponding feature maps from the encoding path, two Chebyshev convolutions, each followed by 1-dimensional batch normalization and ReLU.

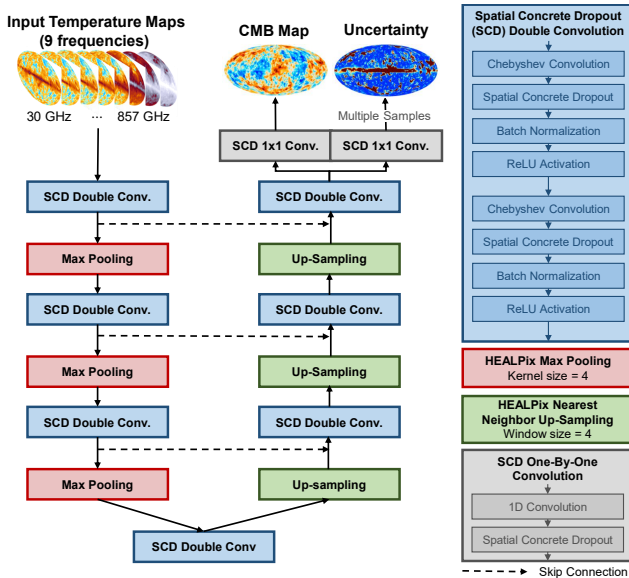


Figure 2: Model Architecture. The graph-based CNN has a U-Net based architecture. The down arrows represent the encoding path and the up arrows represent the decoding path.

The network is formulated as Bayesian using concrete dropout for approximate variational inference and heteroscedastic aleatoric loss. Monte Carlo dropout was initially formulated as a technique for stochastic variational inference that uses a Bernoulli distribution for the approximate posterior distribution, $q_{\theta}(\omega)$ (where θ are the variational parameters and ω the network weights), adapting the model output stochastically (Kendall and Gal 2017). However, obtaining well-calibrated uncertainty estimates requires a tedious grid search to tune dropout probabilities layer-wise. Spatial concrete dropout (Gal, Hron, and Kendall 2017) presents a continuous relaxation of Monte Carlo dropout’s discrete masks that allows updating the optimization operation to tune layer-wise dropout probabilities automatically. We apply spatial concrete dropout to every Chebyshev convolutional layer in the network in training and testing. We sample multiple predicted clean CMB maps with different dropout masks during testing and quantify the epistemic uncertainty as the variance across those maps.

The model captures aleatoric uncertainty by directly predicting a pixel-wise Gaussian distribution for the output CMB map. There are two one-by-one convolutional output layers; one predicts the mean, \hat{y} and the other the log variance, $\log(\hat{\sigma}^2)$ (note the log is predicted for numerical stability). The model is trained to predict both by maximizing the logarithm of the likelihood function under the observed data. The heteroscedastic aleatoric loss with concrete dropout is defined as:

$$\mathcal{L} = \frac{1}{D} \sum_{i=1}^D \left(\frac{1}{2} \hat{\sigma}_i^{-2} \|\mathbf{y}_i - \hat{\mathbf{y}}_i\|^2 + \frac{1}{2} \log(\hat{\sigma}_i^2) \right) + KL(q_{\theta}(\omega) \| p(\omega)) \quad (3)$$

where D is the number of pixels, \mathbf{y} is the target CMB map, $\hat{\mathbf{y}}$ is the predicted CMB map, $\hat{\sigma}^2$ is the predicted variance, and the last term is the Kullback–Leibler (KL) regularization term (Kendall and Gal 2017; Gal, Hron, and Kendall 2017). The KL regularization term ensures that the approximate posterior $q_{\theta}(\omega)$ does not deviate too far from the prior distribution $p(\omega)$, where $p(\omega)$ is selected to be standard Gaussian to allow for analytical evaluation.

Prediction uncertainty, denoted $\text{Var}(\mathbf{y})$, is quantified as the sum of the epistemic and aleatoric uncertainty. $\text{Var}(\mathbf{y})$ is acquired by sampling predictions with different concrete dropout masks T times, then computing:

$$\text{Var}(\mathbf{y}) \approx \frac{1}{T} \sum_{t=1}^T \hat{\mathbf{y}}_t^2 - \left(\frac{1}{T} \sum_{t=1}^T \hat{\mathbf{y}}_t \right)^2 + \frac{1}{2} \sum_{t=1}^T \hat{\sigma}_t^2 \quad (4)$$

where $\hat{\mathbf{y}}_t, \hat{\sigma}_t^2$ are the network predictions with the t -th weights sampled from the approximate posterior, $\omega_t \sim q_{\theta}(\omega)$. The first part of Equation 4, the variance in the predicted means, is the epistemic uncertainty, and the second part, the average predicted variance, is the aleatoric uncertainty.

Model Initialization

As noted in the Bayesian deep learning literature, using the negative log-likelihood based loss in equation 3 in conjunction with a gradient-based optimizer can lead to very poor

but stable parameter estimates (Seitzer et al. 2022). This can result in poor predictive accuracy. We found that this pitfall can be avoided using a sensible transfer learning approach. We initialize the weights of our graph-based Bayesian CNN using a trained deterministic version of the CNN. The deterministic version has an identical architecture to the proposed model (Figure 2), except the predictive variance output layer and all spatial concrete dropout layers are removed. We train the deterministic network using mean square error loss:

$$\mathcal{L}_{\text{deterministic}} = \frac{1}{D} \sum_{i=1}^D \|\mathbf{y}_i - \hat{\mathbf{y}}_i\|^2 \quad (5)$$

until it reaches convergence. We then initialize the Chebyshev convolutional weights in the Bayesian model using the weights learned by the deterministic model. This pre-training helps the Bayesian CNN achieve good predictive accuracy while learning to quantify uncertainty.

The output layer that predicts the variance, which is not pre-trained, is initialized to have small random values sampled from the Gaussian distribution, $\mathcal{N}(0, 1e-6)$. This initialization prevents the variance term from overshooting in early Bayesian training epochs. The layer-wise spatial concrete dropout probabilities are all initialized to be $1e-3$. This small initialization also makes model training more stable and helps account for any covariance shift caused by batch normalization in conjunction with dropout (Li et al. 2019).

Experimental Results

Simulated Data

We generate simulated sky observations following established practice that combines three elements:

1. C_i : A randomly generated realization of primordial CMB anisotropies (the target we seek to recover) from the distribution specified by Planck 2015 estimates (Planck Collaboration 2016b) of the cosmological model parameters.
2. F : Existing (fixed) maps, at each of the nine Planck frequency bands, of the composite foreground emission provided in the 2015 Planck data release (Planck Collaboration 2016a).
3. N_i : Randomly generated, spatially modulated noise realizations consistent with Planck instrument sensitivity per frequency band.

Each complete simulation is the sum of elements $C_i + F + N_i$. We reduced the angular resolution of the signal maps (CMB and foregrounds) to FWHM=150 arcmin (unlike Planck, the same resolution per frequency band, to simplify the analysis), and rendered the sky maps at HEALPix resolution of NSIDE = 64. The resulting simulated sky maps contain 49,152 pixels.

We generated 1,000 instances of noisy multi-frequency maps and clean CMB pairs, each with a unique foreground sky map and unique noise realization. We split these instances into a training set, validation set, and held-out testing sets using a random 80%, 10%, 10% split, resulting in

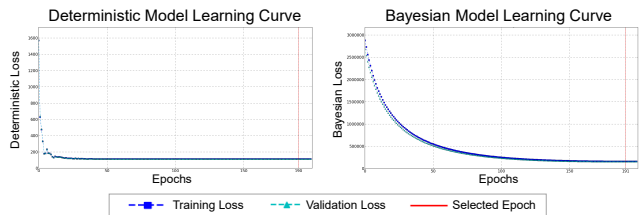


Figure 3: Model Training Learning Curves. The curves show training and validation loss over the epochs as well as the selected epoch to use for initialization or prediction. Note the validation loss is very close to the training loss, suggesting the models are not over-fitting to the training data.

a training set of 800, a validation set of 100, and a testing set of 100. The input temperature maps were normalized channel/frequency-wise to be zero-centered with a standard deviation of one. This standard input normalization improves model training. We found that applying normalization techniques to the target CMB maps did not improve model accuracy, so the model learns to predict un-scaled (not normalized) CMB maps.

Model Training

The model was implemented using PyTorch and trained on an NVIDIA Tesla M60 RAF GPU. The deterministic model used for weight initialization was trained with a learning rate of $1e-3$ with SGD optimization and MSE loss (Equation 5). A batch size of 10 was selected as this was the largest batch size the GPU memory would allow. In both the deterministic and Bayesian training, we saved the model after every training epoch and used the validation set to determine which model to use. We selected the epoch that minimized combined loss defined as: $0.8 \times (\text{validation loss}) + 0.2 \times (\text{train loss})$. This metric allowed us to select the ideal stopping epoch to prevent over-fitting. The selected epoch for the deterministic model was 190, and it took 21 hours to train. The Bayesian model was found to train better with a lower learning rate ($1e-5$) and Adam optimization (Kingma and Ba 2014). A length scale of $1e-4$ was used for concrete dropout. The batch size was reduced to 7 in training the Bayesian model to account for the memory required by the additional concrete dropout parameters. The selected epoch for the Bayesian model was 191, and it took 51 hours to train. The combined training required 382 epochs and took about three days. The learning curves and selected model epochs are shown in Figure 3.

Analysis of Results

We provide qualitative and quantitative metrics for analyzing the prediction accuracy and uncertainty calibration on the simulated data. Figure 4 displays three examples from the test set. On the left, we can see the true CMB temperature map versus the one predicted by the model. On the right, we can see the absolute error, or the absolute difference between the true and predicted maps, and the predicted uncertainty. The predicted uncertainty is the standard deviation of the pixel-wise predicted distribution (or the square

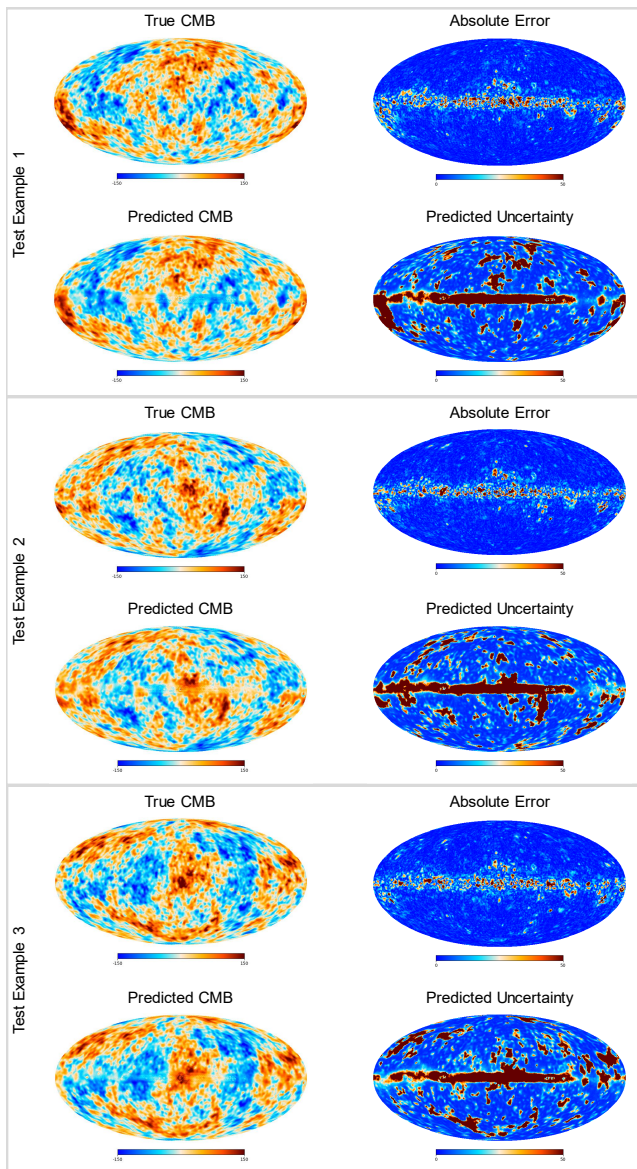


Figure 4: Model Prediction. An example from the test set is provided with the true and predicted CMB maps, absolute error or absolute difference between the true and predicted, and predicted uncertainty.

root of the variance given by Equation 4). As expected, we can see the error and uncertainty are highest along the galactic plane. The predicted uncertainty should correlate with the error, and we observe that it does so, but it is over-estimated in some regions.

To analyze prediction accuracy, we compare against the ILC method. We consider the overall root mean square error (RMSE) as well as the Pearson correlation coefficient between the true and predicted pixel values within a ± 30 degree sky cut. On the test set, the ILC method predictions yield $RMSE=7.597$ and $r=0.993$. The predictions from the proposed CNN yield $RMSE=4.449$ and $r=0.997$, suggest-

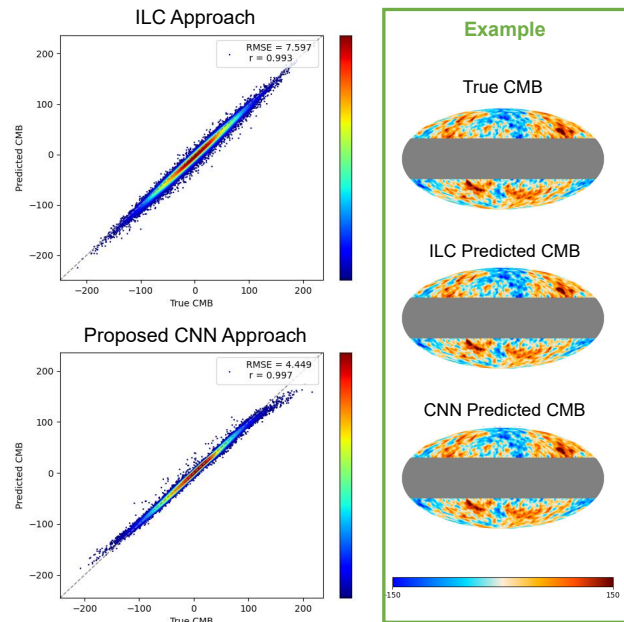


Figure 5: Pixel-wise Prediction Accuracy. Plots show correlation between true and predicted CMB pixel values (where values that lie on the diagonal are more accurately predicted) for the ILC and proposed method. An example from the test set is provided to illustrate the ± 30 degree sky cut used in evaluation.

ing improved accuracy over the ILC baseline, Figure 5 illustrates how closely predicted pixel values correlate with the true CMB pixel values over the entire test for both methods. We note that the CNN predicted values are closer to the diagonal (i.e., more accurate) in general than the ILC values, they are farther from the diagonal for pixels with an actual CMB value of large magnitude. This can be seen in the slight "S" shape of the distribution in Figure 5, where the tails are under-estimated by the model. This issue is also manifested in the sky maps in Figure 4. Regions of large temperature magnitude in the true CMB correlate to regions of slightly higher error in the absolute error maps and notably higher uncertainty in the predicted uncertainty maps. This indicates that the model has less confidence in predicting values of high magnitude, suggesting such regions are under-represented in the training data.

In Figure 6, we can see a comparison in the spectral domain between the test maps predicted by the proposed method and those generated from the ILC method. The spectra were computed on ± 30 degree sky cuts and averaged over the 100 test realizations. In comparing the spectra from the true (input CMB) and the predictions from the ILC and proposed CNN, we can see that both approaches deviate. Notably, the ILC approach picks up noise at high ℓ values, and we see a drop in the power from the CNN predictions after $\approx 100\ell$. This suggests that using the two approaches in tandem could provide better accuracy.

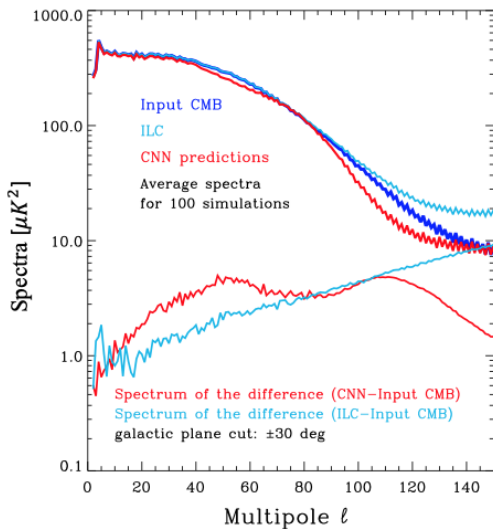


Figure 6: Spectral Comparison. The upper lines show the input CMB (true) spectra as well as those from the ILC and proposed CNN predictions. The lower lines show the spectra of difference maps between ILC and CNN predictions vs. input CMB (smaller is better).

Conclusion and Future Work

We presented a Bayesian deep learning approach to CMB recovery via a probabilistic graph-based CNN. We demonstrated the efficacy of this approach in accurately predicting cleaned CMB maps using simulated low-resolution temperature maps of the full-sky. These results demonstrate the potential for using the proposed model to recover the CMB signal from Planck observations. Three major challenges need to be addressed to make this possible:

1. **Improving accuracy in regions of high magnitude.** As previously noted, the model underestimates regions where the CMB temperature has a high magnitude and predicts high uncertainty in these areas. We find this issue in related works (Petroff et al. 2020; Wang et al. 2022; Casas et al. 2022) as well, where there are structured differences in the true and recovered spectra, which cannot be accounted for by randomness or noise. Resolving this issue will improve prediction accuracy and uncertainty quantification calibration. We believe the root of this issue is that regions of high signal-to-noise are under-represented in the training data, resulting in low model confidence in predicting them. To resolve this, we plan to explore adapting the loss to weigh these regions to a greater extent as well as data augmentation techniques to reduce pixel-value imbalance in simulated data.
2. **Scaling up to full-resolution sky maps.** Currently, we are using a HEALPix map resolution of NSIDE = 64 for faster experimentation and model training. However, the Planck observations provide a resolution of NSIDE 1024 to 4096. Acquiring the most precise estimation of the CMB requires using this full resolution. We plan to implement parallelization or distributed training across

multiple GPUs to address the computational expense associated with scaling up to this high-resolution data.

3. **Predicting both temperature and polarization.** CMB anisotropies lie both in the temperature and polarization of the radiation. Currently, our approach and other existing works consider only fluctuations in temperature. Adapting the model to predict polarization amplitude and direction is required to recover the full CMB signal. Accurately predicting polarization will be more challenging than temperature because the polarization of the foreground emissions is an order of magnitude larger than the CMB. Integrating traditional techniques for CMB cleaning, such as Commander (Eriksen et al. 2008), could improve our approach and help address this challenge. Such a technique could be applied first and provided as additional input to the graph-based CNN, which could then learn to refine the CMB cleaning further.

Future work may also involve exploring different techniques for uncertainty quantification. The concrete dropout approach to variational approximate inference relies on the assumption that the weight prior is standard Gaussian. Methods that learn a prior, such as (Kohl et al. 2018), rather than assuming the prior, may provide more flexibility and thus better calibrated uncertainty quantification.

Path to Deployment. The graph-based Bayesian CNN approach proposed in this work has the potential to be used by current and future cosmology missions from NASA and ESA (the European Space Agency). Separating CMB signals from foreground emission and systematic noise provides the key input needed to constrain models of how the universe began and has evolved. Currently the maps available on mission websites and archives lack pixel-wise uncertainties. Deployment would include integration into mission pipelines and publicly posting the CMB maps with uncertainties (e.g., <https://pla.esac.esa.int/pla/#maps>). Two authors of this paper (Krzysztof M. Gorski and Graça Rocha) are cosmologists who worked as project scientists on ESA’s Planck mission and can inform and guide integration. All cosmologists could use these maps for their own further modeling. The pixel-wise uncertainties will enable fine-grained identification of reliable and problematic sky regions as well as filtering for users who wish to exclude less confidently modeled regions.

Acknowledgments

We thank the Jet Propulsion Laboratory Research and Technology Development program for funding support and the Science Understanding from Data Science initiative leads, Susan Owen and Lukas Mandrake, for this project’s continuous support. The High Performance Computing resources used in this investigation were provided by funding from the JPL Information and Technology Solutions Directorate. This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Copyright 2022. All rights reserved.

References

- Agarap, A. F. 2018. Deep Learning using Rectified Linear Units (ReLU). *arXiv preprint arXiv:1803.08375*.
- Bennett, C. L.; Smoot, G. F.; Hinshaw, G.; Wright, E. L.; Kogut, A.; de Amici, G.; Meyer, S. S.; Weiss, R.; Wilkinson, D. T.; Gulkis, S.; Janssen, M.; Boggess, N. W.; Cheng, E. S.; Hauser, M. G.; Kelsall, T.; Mather, J. C.; Moseley, J., S. H.; Murdock, T. L.; and Silverberg, R. F. 1992. Preliminary Separation of Galactic and Cosmic Microwave Emission for the COBE Differential Microwave Radiometers. *The Astrophysics Journal*, 396(L7-L12).
- Casas, J.; Bonavera, L.; González-Nuevo, J.; Baccigalupi, C.; Cueli, M.; Crespo, D.; Goitia, E.; Santos, J.; Sánchez, M.; and de Cos, F. 2022. CENN: A Fully Convolutional Neural Network for CMB Recovery in Realistic Microwave Sky Simulations. *arXiv preprint arXiv:2205.05623*.
- Eriksen, H.; Jewell, J.; Dickinson, C.; Banday, A.; Górski, K.; and Lawrence, C. 2008. Joint Bayesian Component Separation and CMB Power Spectrum Estimation. *The Astrophysical Journal*, 676(1): 10.
- Gal, Y.; Hron, J.; and Kendall, A. 2017. Concrete Dropout. *Advances in Neural Information Processing Systems*, 30.
- Gorski, K. M.; Hivon, E.; Banday, A. J.; Wandelt, B. D.; Hansen, F. K.; Reinecke, M.; and Bartelmann, M. 2005. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *The Astrophysical Journal*, 622(2): 759.
- Kendall, A.; and Gal, Y. 2017. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? *Advances in Neural Information Processing Systems*, 30.
- Kingma, D.; and Ba, J. 2014. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations*.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding Variational Bayes. *arXiv preprint arXiv:1312.6114*.
- Kohl, S.; Romera-Paredes, B.; Meyer, C.; De Fauw, J.; Ledsam, J. R.; Maier-Hein, K.; Eslami, S.; Jimenez Rezende, D.; and Ronneberger, O. 2018. A Probabilistic U-Net for Segmentation of Ambiguous Images. *Advances in Neural Information Processing Systems*, 31.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. *Advances in Neural Information Processing Systems*, 30.
- Lee, S.; Purushwalkam, S.; Cogswell, M.; Crandall, D.; and Batra, D. 2015. Why M Heads are Better than One: Training a Diverse Ensemble of Deep Networks. *arXiv preprint arXiv:1511.06314*.
- Li, X.; Chen, S.; Hu, X.; and Yang, J. 2019. Understanding the Disharmony between Dropout and Batch Normalization by Variance Shift. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2682–2690.
- Nørgaard-Nielsen, H. U.; and Jørgensen, H. 2008. Foreground Removal from CMB Temperature Maps using an MLP Neural Network. *Astrophysics and Space Science*, 318(3): 195–206.
- Perraudin, N.; Defferrard, M.; Kacprzak, T.; and Sgier, R. 2019. DeepSphere: DeepSphere: Efficient spherical Convolutional Neural Network with HEALPix sampling for cosmological applications. *Astronomy and Computing*, 27: 130–146.
- Petroff, M. A.; Addison, G. E.; Bennett, C. L.; and Weiland, J. L. 2020. Full-sky Cosmic Microwave Background Foreground Cleaning Using Machine Learning. *The Astrophysical Journal*, 903(104).
- Planck Collaboration. 2014. Planck 2013 results. XII. Component separation. *Astronomy & Astrophysics*, 571(A12).
- Planck Collaboration. 2016a. Planck 2015 Results. IX. Diffuse component separation: CMB maps. *Astronomy & Astrophysics*, 594(A9).
- Planck Collaboration. 2016b. Planck 2015 Results. XIII. Cosmological parameters. *Astronomy & Astrophysics*, 594(A13).
- Planck Collaboration. 2020. Planck 2018 Results. IV. Diffuse component separation. *Astronomy & Astrophysics*, 641(A4).
- Rocha, G.; Banday, A.; Barreiro, R. B.; Challinor, A.; Gorski, K. M.; Hensley, B.; Jae, T.; Jewell, J.; Keating, B.; Kogut, A.; Lawrence, C.; Panopoulou, G.; Partridge, B.; Pearson, T.; Silk, J.; Steinhardt, P.; Wehus, I.; Bock, J.; Crill, B.; Delabrouille, J.; Dore, O.; Raul, F. R.; Ijjas, A.; Keskitalo, R.; Kritsuk, A.; Mangilli, A.; Moncelsi, L.; Myers, S.; Steinbach, B.; and Tristram, M. 2019a. Designing Future CMB Experiments: Think. Forthcoming.
- Rocha, G.; Banday, A.; Barreiro, R. B.; Challinor, A.; Górski, K. M.; Hensley, B.; Jaffe, T.; Jewell, J.; Keating, B.; and Kogut, A. 2019b. Astro2020 APC White Paper: The need for better tools to design future CMB experiments. *Bulletin of the American Astronomical Society*, 51, Issue 7(221).
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention*, volume 9351, 234–241.
- Rupprecht, C.; Laina, I.; DiPietro, R.; Baust, M.; Tombari, F.; Navab, N.; and Hager, G. D. 2017. Learning in an Uncertain World: Representing Ambiguity Through Multiple Hypotheses. In *Proceedings of the IEEE International Conference on Computer Vision*, 3591–3600.
- Seitzer, M.; Tavakoli, A.; Antic, D.; and Martius, G. 2022. On the Pitfalls of Heteroscedastic Uncertainty Estimation with Probabilistic Neural Networks. *arXiv preprint arXiv:2203.09168*.
- Wang, G.-J.; Shi, H.-L.; Yan, Y.-P.; Xia, J.-Q.; Zhao, Y.-Y.; Li, S.-Y.; and Li, J.-F. 2022. Recovering the CMB Signal with Machine Learning. *The Astrophysical Journal Supplement Series*, 260(1): 13.