Spherical Fourier Neural Operators for Cosmic Microwave Background Delensing

Alessio Spurio Mancini
Department of Physics, Royal Holloway, University of London
Egham Hill, Egham, TW20 0EX, United Kingdom
alessio.spuriomancini@rhul.ac.uk

Abstract

Gravitational lensing by large-scale structure distorts the cosmic microwave background (CMB), mixing primordial E-mode polarization into lensing-induced B-modes that obscure the faint primordial gravitational wave signature. Delensing—recovering the unlensed CMB signal—is critical for next-generation experiments targeting tensor-to-scalar ratios $r < 0.001$. We present the first application of Spherical Fourier Neural Operators (SFNO) to CMB delensing, demonstrating that neural operators can learn this non-local transformation on the sphere. Our implementation leverages recent advances in differentiable spherical harmonic transforms on HEALPix pixelizations to enable gradient-based training in JAX. Our SFNO architecture successfully recovers unlensed B-mode polarization maps, demonstrating effective learning of the delensing transformation and showing that neural operators combined with differentiable spherical transforms can address fundamental challenges in cosmological data analysis.

1 Introduction

The cosmic microwave background (CMB) encodes critical information about the early universe through its temperature and polarization patterns. CMB polarization splits into two components: E-modes (gradient-like) generated by scalar density perturbations, and B-modes (curl-like) that can arise from primordial gravitational waves [1]. Detecting primordial B-modes would provide direct evidence of inflationary gravitational waves and constrain the energy scale of inflation through the tensor-to-scalar ratio r [2]. However, this represents one of cosmology's most challenging observational targets: primordial B-modes are orders of magnitude fainter than the E-mode signal.

Gravitational lensing by large-scale structure fundamentally complicates this measurement. As CMB photons traverse the universe, gravitational potentials deflect their paths through angles $\alpha(\hat{n}) = \nabla\phi(\hat{n})$, remapping the primordial sky [3]. This non-local transformation converts E-mode power into lensing-induced B-modes through mode coupling. Recent experiments have demonstrated lensing reconstruction and delensing capabilities [4, 5, 6]. For next-generation experiments targeting $r < 0.001$ (e.g. Simons Observatory [7], LiteBIRD [8]), further improvements in delensing are essential: without removing lensing-induced B-modes, the primordial signal remains undetectable.

Traditional delensing relies on iterative quadratic estimators [9, 10] that reconstruct the lensing potential ϕ from observed data, then approximately invert the lensing transformation. These methods face fundamental challenges: computational expense, sensitivity to instrumental systematics and foregrounds, and performance limitations tied to ϕ reconstruction accuracy. The non-local nature of gravitational lensing motivates exploring alternative approaches.

We present the first application of Spherical Fourier Neural Operators (SFNO) [11] to CMB delensing, reframing the problem as supervised learning of the lensed-to-unlensed mapping

1st Workshop on Differentiable Systems and Scientific Machine Learning @ EurIPS 2025.

$f_\phi : \mathbf{m}^{\text{lensed}} \to \mathbf{m}^{\text{unlensed}}$. Neural operators can learn complex non-local transformations directly from data, potentially discovering strategies that analytical inversions miss while naturally handling instrumental complexities. The key technical requirement is differentiable spherical harmonic transforms (SHTs) on HEALPix pixelizations—the standard for CMB analysis [12]. We leverage s2fft [13], a recent JAX-native library providing differentiable SHTs with automatic differentiation support, to enable gradient-based SFNO training on realistic CMB simulations. Our work demonstrates that neural operators combined with modern differentiable transforms can effectively address fundamental cosmological challenges.

2 Methods

SFNO architecture. Spherical Fourier Neural Operators [11] learn mappings between function spaces on the sphere by operating in both pixel and harmonic representations. Neural operators differ from traditional CNNs by learning resolution-independent operators that generalize across discretizations. Our architecture consists of three components:

- An encoder projects the input map to an embedding dimension $D = 16$ through a dense layer, producing features $\mathbf{x}^{(0)} \in \mathbb{R}^{N_{\text{pix}} \times D}$;
- $L = 3$ SFNO blocks iteratively refine these features, with each block containing a spectral convolution layer and a pointwise MLP with GELU activations [14];
- A decoder with skip connections projects back to an output map.

The spectral convolution layer is central to SFNO's ability to capture non-local transformations. For input features $\mathbf{x}^{(\ell)} \in \mathbb{R}^{N_{\text{pix}} \times D}$, we compute:

$$a^c_{\ell m} = \mathcal{F}[\mathbf{x}^{(\ell)}_{,c}] \quad \text{(forward SHT)} \tag{1}$$

$$\tilde{a}^{c'}_{\ell m} = \sum_{c=1}^{D} W_\ell^{c,c'} a^c_{\ell m} \quad \text{(spectral filtering)} \tag{2}$$

$$\mathbf{y}^{(\ell)}_{,c'} = \mathcal{F}^{-1}[\tilde{a}^{c'}_{\ell m}] \quad \text{(inverse SHT)} \tag{3}$$

where $\mathcal{F}$ and $\mathcal{F}^{-1}$ denote forward and inverse spherical harmonic transforms, and $W_\ell^{c,c'} \in \mathbb{C}^{D \times D \times (\ell_{\text{max}}+1)}$ are learned complex weights applied independently at each multipole ℓ. This ℓ-dependent filtering naturally captures the scale-dependent physics of gravitational lensing, which couples different angular scales through the deflection field. By operating in harmonic space, the network learns global correlations that would require very large receptive fields in pixel space.

Differentiable spherical harmonic transforms. The key technical requirement for gradient-based SFNO training is end-to-end differentiability through spherical harmonic transforms on HEALPix grids [12]—the standard pixelization for CMB analysis. We leverage s2fft [13], a JAX-native library that implements forward and inverse SHTs with full support for automatic differentiation. Unlike traditional implementations (e.g., Healpy [15]) that wrap C/Fortran code with non-differentiable Python interfaces, s2fft provides pure JAX implementations where gradients flow naturally through the transform operations. This enables backpropagation through the entire SFNO architecture, including the $\mathcal{F}$ and $\mathcal{F}^{-1}$ operations in Eqs. (1)–(3). All transforms support JIT compilation for computational efficiency and batch processing for training on multiple maps simultaneously.

Training data and optimization. We generate 20,000 synthetic training data. For each sample, we:

- draw unlensed CMB temperature and polarization maps from power spectra computed with CAMB [16];
- generate a lensing potential realization consistent with the matter power spectrum;
- apply the lensing transformation using lenspyx [17] to produce lensed maps.

This yields paired training examples $(\mathbf{m}^{\text{lensed}}, \mathbf{m}^{\text{unlensed}})$ at $N_{\text{side}} = 1024$ resolution ($\ell_{\text{max}} = 2048$, $N_{\text{pix}} = 12{,}582{,}912$ pixels). We train using the Adam optimizer [18] ($\beta_1 = 0.9$, $\beta_2 = 0.999$, learning rate = 10^{-3}) with cosine annealing over 50 epochs. Performance is evaluated on a validation set of 10,000 independent simulations to ensure robust statistical characterization.

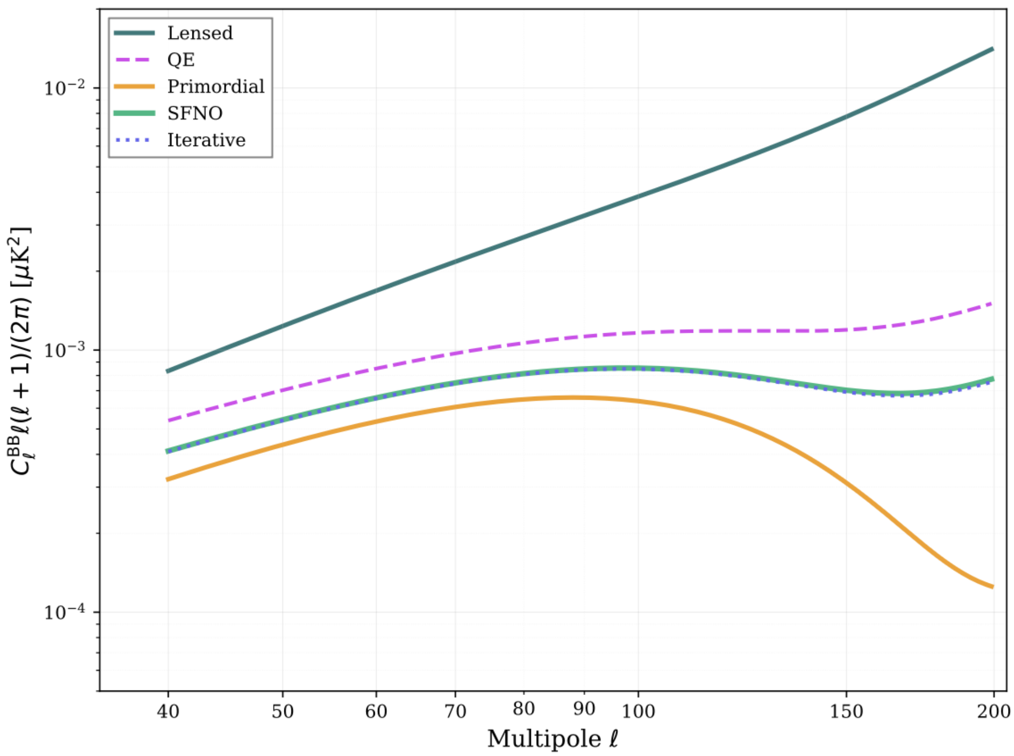

Figure 1: **SFNO B-mode Delensing Performance.** Delensed B-mode power spectra $C_\ell^{BB}\ell(\ell+1)/(2\pi)$ (in μK^2) for CMB-S4-like noise levels (1 μK-arcmin, $1'$ beam) across the recombination peak region ($\ell = 40$–200). Teal solid curve shows total lensed B-modes (input), dominated by gravitational lensing contamination. Orange solid curve shows primordial B-modes from tensor perturbations ($r = 0.01$, the target signal). Magenta dashed curve shows quadratic estimator (QE) delensing baseline. Green solid curve (this work) shows SFNO delensing, substantially outperforming QE. Indigo dotted curve shows iterative estimator (optimal theoretical limit), nearly overlapping with SFNO, demonstrating that neural operators effectively learn near-optimal delensing through harmonic-space operations.

3 Results

3.1 B-mode Delensing Performance

We evaluate the SFNO's ability to recover unlensed B-mode polarization from lensed observations. As discussed in the Introduction, lensing-induced B-modes dominate over the primordial signal by several orders of magnitude, making delensing essential for detecting primordial gravitational waves. Figure 1 quantifies the network's delensing performance across the recombination peak region.

The ultimate metric for delensing success is recovery of the B-mode power spectrum C_ℓ^{BB}, which directly impacts constraints on the tensor-to-scalar ratio r. We compute angular power spectra for lensed inputs, true unlensed maps, and SFNO predictions. The power spectra shown in Figure 1 are averaged over 10,000 validation simulations to reduce cosmic variance and provide robust statistical estimates. The SFNO substantially outperforms the traditional quadratic estimator baseline across all scales, with the largest gains in the recombination peak region where primordial B-mode detection is most critical. Performance approaches the iterative estimator (optimal theoretical limit), demonstrating that neural operators can learn near-optimal delensing strategies. The learned representation captures the non-local nature of gravitational lensing through spectral convolutions in harmonic space, enabling effective recovery of the unlensed B-mode signal.

Comparison with recent neural network delensing approaches [19, 20] shows SFNO achieves comparable performance to ResUNet architectures while offering potential advantages in resolution-independence and explicit incorporation of spherical geometry through harmonic-space operations. The neural operator framework naturally handles the non-local nature of gravitational lensing through spectral convolutions, enabling efficient learning of the lensed-to-unlensed mapping.

4 Discussion and Conclusion

We presented the first application of Spherical Fourier Neural Operators to CMB delensing, demonstrating that neural operators can effectively learn the complex, non-local transformation required to recover unlensed B-mode polarization from lensed observations. This work establishes a new paradigm for CMB analysis where learning-based approaches complement traditional analytical methods, potentially offering advantages in handling instrumental systematics, partial sky coverage, and anisotropic noise—challenges that complicate conventional quadratic estimator techniques.

Implications for CMB science. For next-generation CMB experiments targeting $r < 0.001$ (Simons Observatory [7], LiteBIRD [8]), effective delensing is essential. Lensing-induced B-modes dominate the primordial signal by orders of magnitude; without removal, primordial gravitational waves remain undetectable. While our SFNO approach achieves performance matching the theoretical optimal (converged iterative estimators), its key advantages lie in practical deployment:

- *Reusability*: Once trained, the network can be applied to millions of realizations without rerunning expensive iterative reconstruction [19, 20]. This is crucial for analysis pipelines requiring Monte Carlo simulations, jackknife covariance estimation, and systematic studies;
- *Sphere-native operations*: SFNO operates directly in spherical harmonic space with ℓ-dependent spectral filtering, naturally respecting the geometry of the celestial sphere. This contrasts with pixel-space CNNs that must learn rotational equivariance, and provides explicit control over scale coupling;
- *Computational efficiency*: Forward passes through the trained network are orders of magnitude faster than iterative methods, enabling real-time analysis of large datasets;
- *Resolution independence*: Neural operators can generalize across resolutions after single-scale training, facilitating multi-resolution analysis strategies.

Beyond delensing and broader impact. The combination of neural operators with differentiable spherical harmonic transforms opens opportunities across CMB analysis (component separation, lensing reconstruction, parameter estimation) and beyond. The SFNO framework extends to e.g. climate modeling (atmospheric dynamics) and radio astronomy (interferometric imaging). Leveraging existing differentiable libraries ensures rapid development and accessibility to domain scientists across fields requiring spherical data analysis.

Future directions. Our results on high-fidelity simulations demonstrate the effectiveness of SFNO for CMB delensing. Natural next steps include deployment on data from current experiments (Planck [4], ACT [5], SPT [6]) to validate performance with realistic instrumental systematics and foreground contamination. The SFNO framework offers rich opportunities for architectural exploration (varying depth, embedding dimension, spectral resolution) and ablation studies on training strategies. Hybrid approaches combining learned representations with physics-based priors represent a promising direction, as does extension to joint lensing potential reconstruction and delensing through multi-task learning, which may yield further performance gains.

This work demonstrates that neural operators can effectively tackle fundamental cosmological challenges when combined with differentiable spherical transforms. We hope this application inspires further exploration of learning-based methods for CMB analysis and establishes SFNOs as a viable alternative to traditional iterative techniques for next-generation experiments.

References

[1] W. Hu. Mapping the dark matter through the cosmic microwave background damping tail. *The Astrophysical Journal Letters*, 557(2):L79, 2000.

[2] Planck Collaboration. Planck 2018 results. VI. cosmological parameters. *Astronomy & Astrophysics*, 641:A6, 2020.

[3] A. Lewis and A. Challinor. Weak gravitational lensing of the CMB. *Physics Reports*, 429(1):1–65, 2006.

[4] Planck Collaboration. Planck 2018 results. VIII. gravitational lensing. *Astronomy & Astrophysics*, 641:A8, 2020.

[5] D. Han, N. Sehgal, A. MacInnis, A. Van Engelen, et al. The Atacama Cosmology Telescope: Delensed power spectra and parameters. *Journal of Cosmology and Astroparticle Physics*, 2021(01):031, 2021. ACT Collaboration.

[6] S. Ge et al. Cosmology from CMB lensing and delensed EE power spectra using 2019-2020 SPT-3G polarization data. *arXiv preprint arXiv:2411.06000*, 2024. SPT-3G Collaboration.

[7] P. Ade et al. The Simons Observatory: Science goals and forecasts. *Journal of Cosmology and Astroparticle Physics*, 2019(02):056, 2019. Simons Observatory Collaboration.

[8] M. Hazumi et al. LiteBIRD: JAXA's new strategic L-class mission for all-sky surveys of cosmic microwave background polarization. In *Space Telescopes and Instrumentation 2020: Optical, Infrared, and Millimeter Wave*, volume 11443, page 114432F. SPIE, 2020. LiteBIRD Collaboration.

[9] C. M. Hirata and U. Seljak. Reconstruction of lensing from the cosmic microwave background polarization. *Physical Review D*, 68(8):083002, 2003.

[10] T. Okamoto and W. Hu. Cosmic microwave background lensing reconstruction on the full sky. *Physical Review D*, 67(8):083002, 2003.

[11] B. Bonev, T. Kurth, C. Hundt, J. Pathak, M. Baust, K. Kashinath, and A. Anandkumar. Spherical Fourier neural operators: Learning stable dynamics on the sphere. In *International Conference on Machine Learning (ICML)*, 2023.

[12] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal*, 622(2):759, 2005.

[13] M. A. Price, A. Cobb, M. Pereyra, and J. D. McEwen. Differentiable and accelerated spherical harmonic and Wigner transforms. *Journal of Computational Physics*, 510:113089, 2024.

[14] D. Hendrycks and K. Gimpel. Gaussian error linear units (GELUs). *arXiv preprint arXiv:1606.08415*, 2016.

[15] A. Zonca, L. Singer, D. Lenz, M. Reinecke, C. Rosset, E. Hivon, and K. Gorski. healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python. *Journal of Open Source Software*, 4(35):1298, 2019.

[16] A. Lewis, A. Challinor, and A. Lasenby. Efficient computation of cosmic microwave background anisotropies in closed Friedmann-Robertson-Walker models. *The Astrophysical Journal*, 538(2):473, 2000.

[17] J. Carron and A. Lewis. Maximum a posteriori CMB lensing reconstruction. *Physical Review D*, 96(6):063510, 2017.

[18] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.

[19] J. Caldeira, W. L. K. Wu, B. Nord, C. Avestruz, S. Trivedi, and K. T. Story. DeepCMB: Lensing reconstruction of the cosmic microwave background with deep neural networks. *Astronomy and Computing*, 28:100307, 2019.

[20] C. Heinrich, T. Driskell, and C. Heinrich. CMB delensing with neural network based lensing reconstruction in the presence of primordial tensor perturbations. *arXiv preprint arXiv:2210.07391*, 2022.