

Fast End-to-End Framework for Cosmological Parameter Inference from CMB Data Using Machine Learning

Larissa Santos

*Center for Gravitation and Cosmology, Yangzhou University,
No. 180, Siwangting Road, Yangzhou, 225002, P.R.China**

Camila P. Novaes

*Instituto Nacional de Pesquisas Espaciais,
Av. dos Astronautas 1758, Jardim da Granja,
São José dos Campos, SP, Brazil*

Elisa G. M. Ferreira

*Kavli IPMU (WPI), UTIAS,
The University of Tokyo, 5-1-5 Kashiwanoha,
Kashiwa, Chiba 277-8583, Japan
and*

*Center for Data-Driven Discovery, Kavli IPMU (WPI), UTIAS,
The University of Tokyo, Kashiwa, Chiba 277-8583, Japan*

Carlo Baccigalupi

*The International School for Advanced Studies (SISSA),
via Bonomea 265, I-34136 Trieste,
Italy; The National Institute for Nuclear Physics (INFN),
via Valerio 2, I-34127, Trieste,
Italy; The Institute for Fundamental Physics of the Universe (IFPU),
Via Beirut 2, I-34151, Trieste, Italy*

(Dated: November 4, 2025)

Precise estimation of cosmological parameters from the cosmic microwave background (CMB) remains a central goal of modern cosmology and a key test of inflationary physics. However, this task is fundamentally limited by strong foreground contamination, primarily from Galactic emissions, which obscure the faint CMB B-mode polarization signal. In this Letter, we introduce a fast, simulation-based, end-to-end pipeline that integrates a robust component separation technique with machine-learning, leading to cosmological parameter estimation. Our approach combines the Analytical Blind Separation (ABS) method for foreground removal with a neural network (NN) framework optimized to extract cosmological parameters directly from full-sky simulations. We assess the performance of this methodology for the forthcoming LiteBIRD and PICO satellite missions, designed to detect CMB B-modes with unprecedented sensitivity. Applying the pipeline to realistic sky simulations, we obtain 1σ errors of 0.0035 (LiteBIRD) and 0.0030 (PICO) for the optical depth τ , and 0.005 (LiteBIRD) and 0.0014 (PICO) for the tensor-to-scalar ratio, r . In all cases, the recovered parameters are consistent with input values within 1σ across most of the parameter space. Results for LiteBIRD are in excellent agreement with the latest forecasts from the collaboration. Our findings establish this combined ABS–NN pipeline as a competitive, accurate, and computationally efficient alternative for cosmological parameter inference, offering a powerful framework for forthcoming CMB experiments.

I. INTRODUCTION

The cosmic microwave background (CMB), relic radiation from when the Universe was about 380,000 years old, encodes fundamental information about the early Universe. High-precision measurements of its temperature and polarization anisotropies provide stringent tests of cosmological models and the physics driving cosmic

inflation. Upcoming CMB space missions, such as LiteBIRD [1], PICO [2], and CORE [3], aim to measure B-mode polarization with unprecedented accuracy, probing the energy scale of inflation. A secondary goal is to achieve cosmic-variance-limited E-mode measurements at low multipoles, improving constraints on the optical depth to reionization, τ , one of the least precisely determined Λ CDM parameters.

At microwave frequencies, however, the CMB signal is strongly contaminated by foreground emissions from the Galactic interstellar medium and extragalactic sources. Accurate cosmological inference thus depends on effec-

* larissa@yzu.edu.cn

tive foreground modeling and removal. Component separation techniques based on linear combinations of multifrequency data exploit two key properties: (i) distinct frequency dependences of astrophysical emissions and (ii) their statistical independence. Modeling the sky as a mixture of emission components allows component separation through inversion of a linear system. However, residual contamination in maps and power spectra remains a persistent challenge for future experiments.

Beyond foreground cleaning, precise cosmological parameter estimation and model testing remain central goals of CMB research. Given current tensions within the Λ CDM framework [4], alternative inference approaches are both timely and essential. Machine learning (ML), especially neural networks (NNs), has recently emerged as a powerful alternative to traditional likelihood-based methods, enabling simulation-based inference that avoids assumptions about the functional form of the likelihood of the data [5–8]. However, NN-based inference requires extensive training datasets representative of the data distribution, making the repeated application of component separation methods across large ensembles of simulations computationally prohibitive.

In this Letter, we propose a fast end-to-end methodology combining cosmological parameter inference via ML with component separation method applied consistently to all simulations. Specifically, we employ the Analytical Blind Separation (ABS) technique [9–11] to estimate cleaned CMB E - and B -mode power spectra, coupled with a NN architecture optimized to infer cosmological parameters. We demonstrate this framework on simulated LiteBIRD and PICO observations, focusing on two key parameters for next-generation CMB polarization data: the optical depth τ and the tensor-to-scalar ratio r .

II. ANALYSIS

A. Simulations

We generate synthetic polarization maps using the Planck Sky Model (PSM) [12], which simulates CMB, polarized foregrounds, and instrumental noise (Gaussian noise, uncorrelated from pixel to pixel and from channel to channel) according to LiteBIRD and PICO specifications (Tables I and II). All maps are produced at a HEALPix resolution of $N_{\text{side}} = 256$ [13].

1. CMB maps

We compute theoretical EE and BB angular power spectra using the Boltzmann solver CLASS [14, 15] under the Λ CDM framework. To train and test the NN for cosmological inference, we calculate C_ℓ^{EE}, C_ℓ^{BB} for a set of 700 different $\{r, \tau\}$ pairs of values (each pair refereed as a cosmology). The training set spans $0 < r < 0.05$

and $0.01 < \tau < 0.13$ (500 samples), while the test set covers $0 < r < 0.03$ and $0.038 < \tau < 0.070$ (200 samples). This strategy ensures that the NNs are trained over sufficiently broad parameter ranges, mitigating the prior edge effects that could otherwise introduce biases in the inferred parameters [16, 17]. The narrower parameter range adopted for the test set, while avoiding such boundary-induced artifacts, confines the inference to physically plausible regions of the parameter space consistent with current cosmological constraints [18].

We sample the parameter space using the Latin Hypercube approach [19] and generate 10 realizations of Q and U polarization maps for each cosmology using the power spectra. All other Λ CDM parameters are fixed to the Planck 2018 best-fit values: $H_0 = 67.27 \text{ km s}^{-1}$, $\Omega_c = 0.363$, $\Omega_b = 0.049$, $n_s = 0.9649$. Given the constraint $10^9 A_s e^{-2\tau} = 1.884$ [20], we fix this combination while varying A_s and τ . The resulting E - and B -mode spectra (Figure 1, left) illustrate the impact of varying r and τ .

2. Foreground contaminants

In addition to the CMB, several astrophysical processes contribute to sky emission across 21–462 GHz. We model the two dominant sources of polarized Galactic foregrounds: synchrotron radiation from relativistic charged particles in Galactic magnetic fields, dominant at low frequencies, and thermal dust emission from heated dust grains, prevailing at high frequencies. Extragalactic contributions from radio and far-infrared galaxies, along with cluster emission, are also included to ensure a comprehensive description of polarized contaminants.

TABLE I. Litebird specifications [1, 21].

Band center (GHz)	Beam FWHM (arcmin)	noise level ($\mu\text{K}_{\text{CMB-arcmin}}$)
040	70.5	37.42
050	58.5	33.46
060	51.1	21.31
068	47.1	16.87
078	43.8	12.07
089	41.5	11.3
100	37.8	6.56
119	33.6	4.58
140	30.8	4.79
166	28.9	5.57
195	28.6	5.85
235	24.7	10.79
280	22.5	13.80
337	20.9	21.95
402	17.9	47.45




FIG. 1. Angular power spectra of the input clean CMB simulations (left) and the spectra recovered by ABS from LiteBIRD- (middle) and PICO-like (right) simulations. ABS spectra include the effect of the corresponding largest beam (first row of Tables I and II). The green and blue regions represent the noise levels (frequency dependent) for LiteBIRD and PICO cases, respectively.

TABLE II. PICO specifications [2].

Band center (GHz)	Beam FWHM (arcmin)	noise level ($\mu K_{\text{CMB}}\text{-arcmin}$)
021	38.4	16.9
025	32.0	13.0
030	28.3	8.7
036	23.6	5.6
043	22.2	5.6
052	18.4	4.0
062	12.8	3.8
075	10.7	3.0
090	9.5	2.0
108	7.9	1.6
129	7.4	1.5
155	6.2	1.3
186	4.3	2.8
223	3.6	3.2
268	3.2	2.2
321	2.6	3.0
385	2.5	3.2
462	2.1	6.4

B. Component separation

A novel approach [10] enables direct estimation of the CMB angular power spectrum from multivariate spectra of multifrequency data. Its computational simplicity and efficiency make it ideal for integration with neural network-based parameter inference in an end-to-end framework. The ABS method has been validated on simu-

lated datasets [9, 11] and implemented in the AliCPT-1 pipeline [22], a high-altitude CMB polarization experiment, demonstrating its robustness and readiness for future satellite missions.

The ABS formalism assumes a data model in which the observations in N_f different frequency channels contain a superposition of CMB, foreground emission, and noise as

$$\mathcal{D}_{ij}^{\text{obs}}(\ell) = f_i f_j \mathcal{D}^{\text{cmb}}(\ell) + \mathcal{D}_{ij}^{\text{fore}}(\ell) + \delta \mathcal{D}_{ij}^{\text{noise}}(\ell), \quad (1)$$

where $i, j = 1, 2, \dots, N_f$. We use thermodynamic units for the observations, so that the CMB emission pattern is constant across frequencies, and $f_i = 1, \forall i$. $\mathcal{D}_{ij}^{\text{obs}}$ represents the cross-band power spectrum of the observations in the i - and j -th frequency channels. The three main contributions to the data are the CMB signal $\mathcal{D}^{\text{cmb}}$, $\mathcal{D}_{ij}^{\text{fore}}$ and $\delta \mathcal{D}_{ij}^{\text{noise}}$, which are the cross band power matrices of the foreground and residual instrumental noise, respectively. The ensemble-averaged noise power spectrum is assumed to be known and has been subtracted beforehand from the measured cross-power spectrum. The ABS solution can be written as:

$$\hat{\mathcal{D}}^{\text{cmb}} = \left(\sum_{\tilde{\lambda}_\mu \geq \lambda_{\text{cut}}} \tilde{G}_\mu^2 \tilde{\lambda}_\mu^{-1} \right)^{-1} - \mathcal{S}, \quad (2)$$

where

$$\tilde{\mathcal{D}}_{ij}^{\text{obs}} \equiv \frac{\mathcal{D}_{ij}^{\text{obs}}}{\sqrt{\sigma_{\mathcal{D},i}^{\text{noise}} \sigma_{\mathcal{D},j}^{\text{noise}}}} + \tilde{f}_i \tilde{f}_j \mathcal{S} \quad (3)$$

$$\tilde{f}_i \equiv \frac{f_i}{\sqrt{\sigma_{\mathcal{D},i}^{\text{noise}}}}, \quad \tilde{G}_\mu \equiv \tilde{\mathbf{f}} \cdot \tilde{\mathbf{E}}^\mu. \quad (4)$$

Here $\tilde{\mathbf{E}}^\mu$ and $\tilde{\lambda}_\mu$ are the μ -th eigenvector and corresponding eigenvalue of $\tilde{\mathcal{D}}_{ij}^{\text{obs}}$. The ABS method thresholds the eigenvalues $\tilde{\lambda}_\mu$, only keeping those signal-dominated modes. We choose $\tilde{\lambda}_{\text{cut}} = 1$ for EE and BB spectra according to [9]. In Eq. 2, the free parameter $\mathcal{S}$ corresponds to an amplitude shift of the input CMB power spectrum from $\mathcal{D}^{\text{cmb}}$ to $\mathcal{D}^{\text{cmb}} + \mathcal{S}$, particularly important for low signal-to-noise regime.

C. Cosmological constraints

We use NNs to map the EE and BB power spectra recovered by ABS in a full-sky analysis for LiteBIRD and PICO (second and third columns of Figure 1) into the parameters of interest, r and τ . The cosmological inference follows a procedure fully based on simulations, as described below (for further details, see [23, 24]).

1. Training and test data sets

The robustness of simulation-based cosmological inference critically depends on both the amount of synthetic data available and the degree to which these simulations accurately reproduce observational reality. We address these challenges through the ABS method, which operates directly at the level of the angular power spectrum rather than on sky maps, greatly improving computational efficiency over conventional component-separation techniques. This enables the rapid generation of the extensive and statistically representative training and test data sets required for robust ML-based cosmological inference.

Here, we use a set of 10 recovered BB, EE spectra for each of the 500+200 cosmologies, or pairs $\{r, \tau\}$, considered for simulations. We split the first 500 into 400/100 cosmologies training/validation of the NN model, testing it on the other 200 cosmologies (narrower intervals of parameters). This ensures that cosmologies used for testing are entirely independent from those employed in training and validation. A cross-validation procedure is implemented following [23].

The training+validation set, $\mathcal{T}(X^i, y^i)$, is defined so that the NN can learn the relation between the BB (EE) power spectrum, $X^i = C_\ell^{BB,i}$ ($X^i = C_\ell^{EE,i}$) and r (τ) parameter, $y^i = r^i$ ($y^i = \tau^i$), for the i th simulation. Notice that the training is performed over (i) $\mathcal{T}(C_\ell^{BB,i}, r^i)$ and (ii) $\mathcal{T}(C_\ell^{EE,i}, \tau^i)$ separately, resulting in an NN model for each case. The efficiency of the trained NN models is evaluated over the test set, $t(X^j)$, outputting the

TABLE III. Error estimates evaluating the performance of our predictions of the cosmological parameters. In parenthesis we show RMSE estimated from the training set.

	LiteBIRD		PICO	
	$\langle\sigma\rangle \times 10^2$	RMSE $\times 10^2$	$\langle\sigma\rangle \times 10^2$	RMSE $\times 10^2$
r	0.50	0.56 (0.52)	0.10	0.15 (0.15)
τ	0.30	0.35 (0.33)	0.30	0.30 (0.27)

predicted values of the corresponding cosmological parameter x^j for the j th simulation of the test set.

2. Neural network implementation

We employ a fully connected NN, using the OPTUNA package to sample the hyperparameters space and automatically define the optimal architecture [25]. The maximal number of trials (tested architectures) is fixed to 500, regardless of the case, (i) or (ii), which usually takes no more than 24 hours on 56 cores of a processor Intel Xeon Gold 5120 2.20 GHz and 512 GB of RAM. After defining the architecture, the training+validation process provides an NN model in a few minutes. The loss function is chosen to be the mean square error, $\mathcal{L} = \langle (y^{\text{Pred}} - y^{\text{True}})^2 \rangle$, averaged over the simulations, where y^{Pred} and y^{True} are the predicted and true values of the cosmological parameter to be constrained.

3. Parameter inference

Figure 2 presents the results of applying the trained NN models to the test sets, comparing the predicted and true values of r and τ . Model performance is quantified using two error estimates: (i) the average standard deviation over the 200 cosmologies of the test set, $\langle\sigma\rangle = \sigma(y^{\text{Pred}})$, where σ is the standard deviation over the 10 predictions for one cosmology (ii) the root-mean-square error comparing predicted and true values of the 10×200 simulations, $\text{RMSE} = \sqrt{\mathcal{L}}$, corresponding to the 1σ error from all predictions. Table III summarizes these results. The close agreement between the RMSE calculated from the training and test sets confirms that our results are not affected by overfitting.

III. DISCUSSION AND CONCLUSION

In this work, we demonstrate the feasibility and scientific value of integrating the ABS method with a NN-based cosmological inference pipeline trained entirely on simulations. The exceptional efficiency of ABS, capable of cleaning foregrounds from hundreds of simulated skies within a few days, makes it ideally suited for simulation-based inference in CMB analyses.




FIG. 2. Predicted versus true cosmological parameters (r and τ) for LiteBIRD-like (blue) and PICO-like (orange) instruments. Dots and error bars show the mean and standard deviation from 10 simulations for each of the 200 cosmologies of the test set. The black dotted and dashed lines denote the linear fits to LiteBIRD and PICO results, while the thin diagonal indicates the identity $y^{Pred} = y^{True}$. The corresponding root-mean-square errors are also shown. Bottom panels display the statistical significance of the predictions as a function of the true values, with the dashed line marking the 1σ level.

As shown in Figure 1, ABS robustly recovers both EE and BB power spectra, with PICO-like instruments slightly outperforming LiteBIRD-like configurations in

B-mode reconstruction, consistent with their noise levels. The resulting parameter constraints (Figure 2) confirm this performance: r is recovered with 1σ uncertainties of 0.0056 for LiteBIRD and 0.0015 for PICO, fully consistent with mission specifications. Both instruments show strong potential for B-mode detection. More specifically, our method predicts r values within 1σ for PICO-like sensitivities, and within 1σ (3σ) for LiteBIRD in the regime $r > 0.01$ ($r < 0.01$).

From the EE spectra, both missions yield consistent τ constraints with 1σ errors of 0.0035 (LiteBIRD) and 0.0030 (PICO), with recovered values matching true parameters within 1σ across the entire parameter space of the test set. The agreement between $\langle\sigma\rangle$ and the 1σ (RMSE) confirms the robustness and internal consistency of the method. Moreover, our LiteBIRD results align with the latest mission forecasts [21].

In summary, this end-to-end, simulation-based framework provides a fast, accurate, and scalable alternative to traditional likelihood analyses. The synergy between ABS's speed and NN inference offers a powerful tool for forecasting and optimizing next-generation CMB experiments, well suited for large simulation ensembles and realistic data pipelines. This approach represents a promising direction for future missions where precision, computational efficiency, and adaptability are paramount.

ACKNOWLEDGMENTS

L. S. is supported by the National Key R&D Program of China (2020YFC2201600). C. P. N. and E. F. thank the Serrapilheira Institute for financial support. This research used computing resources at Kavli IPMU. The Kavli IPMU is supported by the WPI (World Premier International Research Center) Initiative of the MEXT (Japanese Ministry of Education, Culture, Sports, Science and Technology). C. B. acknowledges partial support by the Italian Space Agency LiteBIRD Project (ASI Grants No. 2020-9-HH.0 and 2016-24-H.1-2018), and the Italian Space Agency Euclid Project, as well as the InDark and LiteBIRD Initiative of the National Institute for Nuclear Physics, and the Radio-ForegroundsPlus Project HORIZON-CL4-2023-SPACE-01, GA 101135036, and Project SPACE-IT-UP by the Italian Space Agency and Ministry of University and Research, Contract Number 2024-5-E.0 and The CMB-Inflate project funded by the European Union's Horizon 2020 Research and Innovation Staff Exchange under the Marie Skłodowska-Curie grant agreement No 101007633.

[1] M. Hazumi *et al.* (LiteBIRD), Proc. SPIE Int. Soc. Opt. Eng. **11443**, 114432F (2020), arXiv:2101.12449 [astro-ph.IM].

[2] S. Hanany *et al.* (NASA PICO), (2019), arXiv:1902.10541 [astro-ph.IM].

[3] J. e. a. Delabrouille, J. Cosmology Astropart. Phys.

- 2018**, 014 (2018), arXiv:1706.04516 [astro-ph.IM].
- [4] E. Di Valentino, J. Levi Said, and E. N. Saridakis, arXiv e-prints , arXiv:2509.25288 (2025), arXiv:2509.25288 [astro-ph.CO].
- [5] K. Cranmer, J. Brehmer, and G. Louppe, Proceedings of the National Academy of Sciences **117**, 30055 (2020).
- [6] J. Alsing, B. Wandelt, and S. Feeney, Monthly Notices of the Royal Astronomical Society **477**, 2874 (2018).
- [7] J. Alsing and B. Wandelt, Monthly Notices of the Royal Astronomical Society **488**, 5093 (2019).
- [8] N. Jeffrey, J. Alsing, and F. Lanusse, Monthly Notices of the Royal Astronomical Society **501**, 954 (2021).
- [9] J. Yao, L. Zhang, Y. Zhao, P. Zhang, L. Santos, and J. Zhang, ApJS **239**, 36 (2018), arXiv:1807.07016 [astro-ph.CO].
- [10] P. Zhang, J. Zhang, and L. Zhang, MNRAS **484**, 1616 (2019).
- [11] L. Santos, J. Yao, L. Zhang, S. Ghosh, P. Zhang, W. Zhao, T. Villela, J. Chen, and J. Delabrouille, arXiv e-prints , arXiv:1908.07862 (2019), arXiv:1908.07862 [astro-ph.CO].
- [12] J. Delabrouille *et al.*, Astron. Astrophys. **553**, A96 (2013), arXiv:1207.3675 [astro-ph.CO].
- [13] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, Astrophys. J. **622**, 759 (2005), arXiv:astro-ph/0409513 [astro-ph].
- [14] J. Lesgourgues, ArXiv e-prints (2011), arXiv:1104.2932 [astro-ph.IM].
- [15] D. Blas, J. Lesgourgues, and T. Tram, JCAP **2011** (07), 034.
- [16] F. Villaescusa-Navarro, B. D. Wandelt, D. Anglés-Alcázar, S. Genel, J. M. Z. Matilla, S. Ho, and D. N. Spergel, The Astrophysical Journal **928**, 44 (2022).
- [17] L. A. Perez, S. Genel, F. Villaescusa-Navarro, R. S. Somerville, A. Gabrielpillai, D. Anglés-Alcázar, B. D. Wandelt, and L. Yung, arXiv preprint arXiv:2204.02408 (2022).
- [18] N. Aghanim *et al.*, Astron. Astrophys **641**, A6 (2020).
- [19] M. D. McKay, R. J. Beckman, and W. J. Conover, Technometrics **21**, 239 (1979), <https://doi.org/10.1080/00401706.1979.10489755>.
- [20] K. Wolz, N. Krachmalnicoff, and L. Pagano, Astronomy & Astrophysics **676**, A30 (2023).
- [21] E. Allys *et al.* (LiteBIRD), PTEP **2023**, 042F01 (2023), arXiv:2202.02773 [astro-ph.IM].
- [22] J. Zhang, S. Ghosh, J. Dou, Y. Liu, S. Li, J. Chen, J. Wang, Z. Zhang, J. Delabrouille, M. Remazeilles, C. Feng, B. Hu, H. Liu, L. Santos, P. Zhang, W. Zhao, L. Zhang, Z.-Q. Huang, H. Li, and X. Zhang, ApJS **274**, 26 (2024), arXiv:2402.01233 [astro-ph.CO].
- [23] C. P. Novaes, E. J. de Mericia, F. B. Abdalla, C. A. Wuensche, L. Santos, J. Delabrouille, M. Remazeilles, V. Luccardo, E. Abdalla, L. Barosi, *et al.*, Monthly Notices of the Royal Astronomical Society **528**, 2078 (2024).
- [24] C. P. Novaes, L. Thiele, J. Armijo, S. Cheng, J. A. Cowell, G. A. Marques, E. G. M. Ferreira, M. Shirasaki, K. Osato, and J. Liu, arXiv e-prints , arXiv:2409.01301 (2024), arXiv:2409.01301 [astro-ph.CO].
- [25] We fix the optimizer to **Adam**, with $\beta_1, \beta_2 = 0.9, 0.999$. The hyperparameters and the respective searching spaces considered by OPTUNA are: number of neurons in the hidden layers: [1, 500]; number of layers: [1, 3]; activation function: [ReLU, tanh]; learning rate: [10^{-4} , 10^{-2}]; batch size: [50, 500]; and number of epochs: [50, 500].