

Debiasing cosmological parameters from large-scale foreground contamination in Cosmic Microwave Background data

Alessandro Carones

International School for Advanced Studies (SISSA), Via Bonomea 265, Trieste 34136, Italy

INFN Sezione di Trieste, Via Valerio 2, Trieste 34127, Italy

IFPU, Via Beirut 2, Trieste 34151, Italy

E-mail: acarones@sisssa.it

Abstract. Current and future Cosmic Microwave Background (CMB) experiments aim to achieve high-precision reconstruction of the CMB polarization signal, with the most ambitious objective being the detection of primordial B modes sourced by cosmic inflation. Given the expected low amplitude of the signal, its estimate—parametrized by the tensor-to-scalar ratio r —is highly susceptible to contamination from Galactic foreground residuals that remain after component separation.

In this work, we introduce an agnostic, model-independent procedure to construct a spectral template of residual foreground contamination in the observed angular power spectrum. Specifically, a cleaned multifrequency set of foreground-emission maps is blindly reconstructed from the observed data using the Generalized Needlet Internal Linear Combination (GNILC) technique. These maps are then combined with the weights adopted for CMB reconstruction, yielding an estimate of the spatial distribution of foreground residuals after component separation. The power spectrum of this estimated residual map, after proper noise debiasing, is incorporated into the spectral model of the cosmological likelihood, thereby enabling unbiased inference of cosmological parameters.

We validate the proposed method using realistic simulations of a *LiteBIRD*-like experiment processed with two alternative Internal Linear Combination (ILC) component-separation techniques, focusing on constraints on the tensor-to-scalar ratio. When the foreground contribution is not included in the spectral model, the resulting posteriors of the cosmological parameter are biased, irrespective of its input value, the assumed foreground model, or the adopted masking strategy. Conversely, when the residual template is included in the likelihood, the analysis yields unbiased estimates of r for all considered cases, thereby demonstrating the robustness of the proposed procedure.

The pipeline has been made publicly available as part of the BROOM Python package.

Keywords: CMBR polarization – gravitational waves and CMBR polarization – cosmological parameters from CMBR – CMBR experiments

Contents

1	Introduction	1
2	Background methodology	2
2.1	Microwave sky simulations	2
2.2	Component separation	4
2.3	Tensor-to-scalar ratio inference	5
3	Construction of a spectral template of foreground residuals	6
3.1	Derivation of the template of the foreground residuals	7
3.2	Spectral model	9
4	Results	10
5	Conclusions	14

1 Introduction

Upcoming measurements of the polarization field of the Cosmic Microwave Background (CMB), with unprecedented sensitivity, promise to deepen our understanding of the Universe [1, 2]. For instance, they will provide new insights into cosmic inflation [3, 4], the reionization history [5], neutrino properties [6, 7], additional relativistic relics [8, 9], and cosmic birefringence [10, 11]. These advances rely on precise reconstruction of both the CMB E -mode and B -mode polarization patterns across the sky [12, 13]. Achieving this, however, is hindered by two major challenges. First, residual contamination from polarized foreground emission inevitably remains in the data due to imperfect component separation. Second, instrumental systematics can leak into the reconstructed CMB signal. Both effects can mimic genuine cosmological signatures at the power spectrum level, thereby introducing systematic biases in the inference of cosmological parameters from CMB polarization measurements. To mitigate the impact of foreground residuals, three complementary (though not mutually exclusive) strategies can be pursued: (i) improving the effectiveness and robustness of component-separation methods in the data-analysis pipeline; (ii) refining the masking of the most contaminated sky regions after component separation; and (iii) explicitly modelling the residual contamination at the spectral level and incorporating it into the likelihood.

In this work, we focus on the last approach. We present a new methodology for constructing a spectral template of foreground residuals after component separation, which is then consistently incorporated into the spectral likelihood to mitigate biases in cosmological parameter estimation. The approach relies on obtaining cleaned templates of foreground emission at all observed frequency channels using the Generalized Needlet Internal Linear Combination (GNILC) technique [14, 15]. These templates are subsequently combined with the component separation weights employed to recover the CMB signal, yielding an estimate of the overall residual contamination in the reconstructed CMB maps. The angular power spectrum of this residual template is finally incorporated into the cosmological harmonic likelihood as a spectral model of the contamination in the observed power spectrum. The procedure is fully model-independent, ensuring robustness against the specific properties of polarized foregrounds that will be revealed by future surveys, and is broadly applicable across different map-based component separation pipelines.

The approach is validated using realistic simulations of the microwave sky based on a *LiteBIRD*-like instrumental configuration as proposed in [2], designed to achieve high-precision reconstruction of the CMB polarization field over a large fraction of the sky. As a proof of concept, we apply the method to the recovered B -mode power spectrum, demonstrating its ability to debias the estimation of the amplitude of primordial tensor perturbations, quantified by the tensor-to-scalar ratio r [4]. The component-separation techniques adopted for CMB reconstruction in this work are state-of-the-art Internal Linear Combination (ILC) methods [16, 17].

Details of the simulated data set, the employed component separation pipelines, and the inference of the tensor-to-scalar ratio, which form the basis of the analysis, are presented in Section 2. A full description of the methodology used to derive a template of foreground residuals and its incorporation into the likelihood is provided in Section 3. The results are reported in Section 4, followed by concluding remarks in Section 5.

2 Background methodology

In this section, we describe the considered sets of microwave sky simulations (Section 2.1), which are subsequently processed with a realistic component separation pipeline (outlined in Section 2.2). To assess the quality of the cosmological information that can be extracted from the reconstructed CMB maps, we rely on the angular power spectrum, the standard statistical measure in this context. The procedure for computing the angular power spectra, including the adopted masking strategies, and the subsequent cosmological inference are detailed in Section 2.3.

2.1 Microwave sky simulations

In this work, we simulate realistic observations expected from a *LiteBIRD*-like satellite. This instrumental configuration is designed to observe the sky in 15 frequency channels, achieving an overall polarization sensitivity of $2.2\,\mu\text{K-arcmin}$. For our analysis, we adopt the angular resolution and sensitivity specifications for each frequency channel as reported in Table 13 of [2].

The simulated observations include contributions from the CMB signal, Galactic foreground emission, and instrumental noise. We generate 100 independent simulations of the microwave sky, each featuring different realizations of the CMB signal and noise. All components are simulated as HEALPix-pixelized maps [18, 19] at a spatial resolution of $N_{\text{side}} = 64$, in terms of the linear polarization Stokes parameters Q and U , and expressed in the standard CMB temperature units of μK_{CMB} . The choice of this HEALPix resolution is motivated by the focus on large angular scales. Indeed, in this proof-of-concept study, we apply the proposed debiasing technique to the estimation of the tensor-to-scalar ratio r , whose constraining power lies primarily at multipoles $\ell < 180$, which can be fully sampled with the chosen N_{side} . In addition, the most significant contamination from Galactic diffuse emission arises precisely on these angular scales, making them the regime where cosmological parameter inference is most susceptible to astrophysical systematic biases.

The instrumental beams are assumed to be Gaussian and the bandpasses are modeled as δ functions (i.e., with no bandwidth). These simplifying assumptions do not affect either the core methodology presented in this work or its main results.

The CMB polarization anisotropies are modeled as Gaussian random fields drawn from theoretical angular power spectra computed with the CAMB Boltzmann code¹ [20], assuming the Λ CDM best-fit cosmological parameters from the Planck Data Release 3 angular power spectra [9]. For this analysis, we consider three sets of CMB maps that differ only in their B -mode component: one

¹<https://github.com/cmbant/camb>

with a B -mode angular power spectrum corresponding to lensed E -modes only ($r = 0$) and two with additional tensor perturbations at the levels of $r = 0.004$ and $r = 0.01$, respectively. These input tensor-to-scalar ratios are chosen to represent scientifically and observationally relevant benchmarks for future CMB data analyses [1, 2, 21].

Simulated Galactic foreground emission includes synchrotron [22, 23] and thermal dust radiation [24, 25], generated using the PySM Python package [26–28]. Synchrotron radiation, $X_s(\hat{n}, \nu)$, along the line-of-sight direction $\hat{n}$ and at frequency ν , is modeled as the dominant low-frequency emission and assumed to follow a power-law spectral energy distribution (SED) in brightness temperature units:

$$X_s(\hat{n}, \nu) = X_s(\hat{n}, \nu_s) \cdot \left(\frac{\nu}{\nu_s} \right)^{\beta_s(\hat{n})}, \quad (2.1)$$

where $X = \{Q, U\}$ are Stokes parameters, $\beta_s(\hat{n})$ is the synchrotron spectral index and $X_s(\hat{n}, \nu_s)$ is the synchrotron template at the reference frequency ν_s . We consider two PySM synchrotron models: s1 and s5. Both adopt as polarization template the *WMAP* 9-year 23-GHz Q and U maps [29], smoothed with a Gaussian kernel with Full Width at Half Maximum (FWHM) equal to 5° . In s1, small-scale fluctuations are added as Gaussian realizations drawn from the extrapolated large-scale angular power spectrum [26]. In contrast, s5 introduces non-Gaussian small-scale modes using the `logpol tens` formalism [28]. In both cases, the synchrotron spectral index map is obtained from a spectral fit to the Haslam 408-MHz and *WMAP* 23-GHz 7-year data [30], but in s5 it is further rescaled to account for the increased variability observed in S-PASS data [22]. Overall, the s5 model represents a more complex and realistic description of Galactic synchrotron emission.

Thermal dust emission, $X_d(\hat{n}, \nu)$, is modeled as a modified blackbody (MBB) spectrum in intensity units:

$$X_d(\hat{n}, \nu) = X_d(\hat{n}, \nu_d) \cdot \left(\frac{\nu}{\nu_d} \right)^{\beta_d(\hat{n})} \cdot \frac{B_\nu(T_d(\hat{n}))}{B_{\nu_d}(T_d(\hat{n}))}, \quad (2.2)$$

where $B_\nu(T)$ is the Planck blackbody function, $\beta_d(\hat{n})$ is the dust spectral index, $T_d(\hat{n})$ the dust temperature, and $X_d(\hat{n}, \nu_d)$ the dust template at a reference frequency ν_d . As for synchrotron, we consider two PySM dust models featuring increasing complexity: d1 and d10. In both cases, the β_d , T_d , and $X_d(\hat{n}, \nu_d)$ maps are obtained from MBB fits to *Planck* data. These fits are performed with the `Commander` algorithm [26, 31] for d1 and using the GNILC-denoised data [15, 28] for d10. Large-scale modes are directly extracted from the data, while small-scale features are added following the same prescriptions used for the s1 and s5 synchrotron models, respectively.

From these ingredients, we construct two distinct sets of foreground maps: a simpler one, d1s1 (combining d1 dust and s1 synchrotron), and a more complex one, d10s5 (including d10 dust and s5 synchrotron).

Finally, instrumental noise in the Q and U maps is simulated as white, isotropic, and independent Gaussian realizations, with the pixel standard deviation determined from the sensitivity specifications reported in [2]. The scanning strategy is not included, as it is not expected to introduce any significant impact on the results. Moreover, no correlated $1/f$ noise component is added, under the assumption of proper functioning of the Half-Wave Plate polarization modulator.

All these microwave components are generated separately through the `BROOM` Python package² (corresponding paper in preparation) and then co-added together.

²<https://github.com/alecarones/broom>

2.2 Component separation

To forecast a realistic reconstruction of the CMB B -mode polarization pattern from *LiteBIRD*-like simulated maps, we apply a state-of-the-art component separation procedure. Specifically, we implement an advanced ILC technique [32], applied independently at different angular scales taking advantage of needlet filtering of the input data [33], leading to the so-called Needlet ILC (NILC) [16]. The procedure consists of the following steps:

1. The multifrequency data-set is brought to a common angular resolution, chosen to be the one of the lowest frequency channel in this case: FWHM = 70.5 arcmin. Q and U maps are then transformed into B -mode maps via a full-sky harmonic decomposition.
2. The B -mode maps are convolved with a set of needlet filters $b^{(j)}$ [33], where the index j labels the different kernels, resulting in multiple multifrequency sets of needlet coefficients $\beta^{(j)}(\hat{n}, \nu)$, each sensitive to a specific range of angular scales. The adopted set of needlet bands in harmonic space is shown in the left panel of Figure 1.
3. For each needlet scale j , the multifrequency needlet maps are linearly combined with frequency- and pixel-dependent weights to yield a cleaned estimate of the CMB signal $\tilde{\beta}_{\text{CMB}}^{(j)}(\hat{n})$:

$$\tilde{\beta}_{\text{CMB}}^{(j)}(\hat{n}) = \mathbf{w}_{(j)}(\hat{n}) \cdot \beta^{(j)}(\hat{n}, \nu) = \mathbf{w}_{(j)}(\hat{n}) \cdot (\beta_{\text{CMB}}^{(j)} + \beta_{\text{fgds}}^{(j)} + \beta_{\text{noi}}^{(j)})(\hat{n}, \nu), \quad (2.3)$$

where the contributions from CMB, foregrounds, and noise are explicitly separated. The weights $\mathbf{w}_{(j)}$ are derived so as to minimize the local variance of the solution while preserving the CMB signal, i.e. $\mathbf{w}_{(j)}^T(\hat{n}) \cdot \mathbf{A}_{\text{CMB}} = 1$, with $\mathbf{A}_{\text{CMB}}$ denoting the CMB SED in the adopted units [16]. The combination coefficients satisfying these conditions are obtained as:

$$\mathbf{w}_{(j)}(\hat{n}) = \left[\mathbf{A}_{\text{CMB}}^T \left(C_{(j)}^{-1}(\hat{n}) \right) \mathbf{A}_{\text{CMB}} \right]^{-1} \mathbf{A}_{\text{CMB}}^T C_{(j)}^{-1}(\hat{n}), \quad (2.4)$$

where $C_{(j)}$ is the empirical covariance matrix:

$$(C_{(j)})_{\nu\nu'}(\hat{n}) = \left\langle \beta^{(j)}(\hat{n}', \nu) \cdot \beta^{(j)}(\hat{n}', \nu') \right\rangle_{\hat{n}' \in \mathcal{D}_{\hat{n}}}, \quad (2.5)$$

estimated as a local average of cross-frequency input needlet maps over a sky domain $\mathcal{D}_{\hat{n}}$. In practice, $\mathcal{D}_{\hat{n}}$ is a circular region centered on the pixel $\hat{n}$, with pixel weights that decrease smoothly and symmetrically toward the borders. The size of each domain depends on the needlet scale j : low- j values (corresponding to larger angular scales) require broader domains to include a sufficient number of modes for an accurate covariance estimation.

4. The recovered CMB coefficients from the different needlet scales are finally recombined through an inverse needlet transform [33] to produce a real-space CMB map.

In addition to the baseline NILC implementation, we also consider an alternative methodology. The analysis in [17] demonstrated that residual foreground and noise contamination in the CMB B -mode solution can be further reduced by refining the concept of locality in the variance minimization, moving from a purely spatial basis to one based on foreground properties. Specifically, the sky is partitioned into distinct regions each characterized by similar spectral properties of the B -mode foregrounds. The ILC minimization and weight estimation of Equation 2.4 are then implemented independently within each region. This defines the so-called Multi-Clustering NILC (MC-NILC) technique.

To trace the variability of the foreground spectral properties in a model-independent fashion, an optimal blind tracer has been identified: the ratio of two frequency channels, one dust-dominated (in the considered instrumental configuration, the 337 GHz channel) and one with strong constraining power for the CMB reconstruction (here the 119 GHz channel) [17]. This tracer must be reconstructed directly from the observed data. For this purpose, two cleaned templates of the foreground emission—with substantially reduced contributions from CMB and noise—are derived at the relevant frequencies through the application of the GNILC technique (further details of this methodology are provided in Section 3.1).

At present, a successful partitioning can only be achieved on large angular scales. Consequently, MC-NILC is applied exclusively to the first needlet band ($j = 0$), while standard NILC is used for all other scales. We refer to this hybrid approach as (MC-)NILC.

Both NILC and (MC-)NILC are applied to the input simulations described in Section 2.1, making use of dedicated component-separation modules implemented in the BROOM package. The output of these pipelines consists of full-sky B -mode maps of the reconstructed CMB signal.

2.3 Tensor-to-scalar ratio inference

To infer the posterior distribution of the tensor-to-scalar ratio from the reconstructed CMB signal, we compute the B -mode angular power spectrum of the maps, which is then used as input to the cosmological likelihood. Prior to the power spectrum estimation, the reconstructed maps are masked to suppress leakage from the brightest foreground residuals along the Galactic plane. Two masking strategies are considered in this work, corresponding to different retained sky fractions, in order to test the robustness of the proposed “marginalization” procedure. Specifically, we adopt the GAL60 and GAL40 Galactic masks released by the Planck Collaboration³, which retain 60% and 40% of the sky, respectively. These masks are constructed from intensity data and may therefore be suboptimal for polarization analyses. However, this actually represents an advantage for our purposes, as it allows us to validate the marginalization approach under extreme conditions, where foreground residuals are expected to be significant and could therefore severely bias the inferred tensor-to-scalar ratio.

For simplicity, the angular power spectra are computed using the *anafast* routine of the *healpy* Python package [19]. No E – B leakage correction is required, since the spectra are computed directly from the masked B -mode maps. The loss of power induced by sky masking is corrected by a factor $1/f_{\text{sky}}$, while that due to finite resolution is accounted for by dividing by the corresponding harmonic transfer function. The only effect of masking not explicitly accounted for in this framework is the induced correlation between multipoles. As a consistency check, we repeat the analysis for selected cases using a pseudo- C_ℓ correction following the MASTER approach [34, 35], as implemented in the NaMaster package [36], and obtain consistent results.

We consider different cases corresponding to all combinations of sky models, input values of r , adopted component separation pipelines, and masking strategies. For each case, as described in Section 2.1, we have 100 cleaned CMB maps and their associated angular power spectra (C_ℓ^{out}). To derive the expected posterior distribution of the tensor-to-scalar ratio, we compute the average spectrum across simulations ($\langle C_\ell^{\text{out}} \rangle$) and plug it into the following cosmological inverse-Wishart log-likelihood [38, 39]:

$$\log \mathcal{L}(r) = -f_{\text{sky}} \frac{2\ell + 1}{2} \left[\frac{\langle C_\ell^{\text{out}} \rangle}{C_\ell(r)} + \log C_\ell(r) - \frac{2\ell - 1}{2\ell + 1} \log \langle C_\ell^{\text{out}} \rangle \right], \quad (2.6)$$

³<https://pla.esac.esa.int/>




Figure 1. Harmonic needlet filters used in the component separation step (left) and in the GNILC derivation of the multifrequency cleaned foreground maps (right). Both configurations employ mexican-hat needlet bands [37].

where

$$C_\ell(r) = C_\ell^{\text{lens}} + r \cdot C_\ell^{\text{tens}}(r=1) + \langle C_\ell^{\text{noi}} \rangle \quad (2.7)$$

is the theoretical B -mode angular power spectrum including lensing, tensor perturbations, and residual instrumental noise. The latter is estimated as the average of the power spectra of noise-only residuals after component separation, obtained by combining noise-only input maps with the component separation weights. The posterior of cosmological parameters (here, only the tensor-to-scalar ratio r) is sampled with a Markov Chain Monte Carlo (MCMC) algorithm using the `emcee` Python package [40], adopting a uniform prior $0 \leq r \leq 0.05$. Although only one parameter is varied, we employ MCMC sampling to ensure consistency with the case where two parameters are sampled (presented in Section 3.2). We note that the lensing amplitude is kept fixed, so any residual foreground contamination in the cleaned CMB maps is absorbed as a bias in r .

Strictly speaking, the likelihood in Equation 2.6 is exact only for full-sky, signal-dominated power spectra [38, 39]. This condition is not met in our case, where sky masking introduces non-trivial correlations among multipoles and residuals can contribute significantly to the total power budget. Nonetheless, we adopt this likelihood for three main reasons: (i) it is straightforward to implement, and the main conclusions are not expected to be sensitive to the exact likelihood formalism; (ii) it allows for direct comparison with similar analyses in the literature [2, 17, 41]; and (iii) our primary objective is to compare the posteriors of r obtained when foreground residuals are explicitly included in the spectral model with those that neglect them, making the internal consistency of the chosen likelihood the most relevant criterion. Overall, the adoption of this likelihood in this context is expected to yield slightly inaccurate estimates of the error bounds—whose precise determination is not the main objective of this work—and, in some cases, to introduce a bias in the inferred cosmological parameters. However, we anticipate that no statistically significant bias is observed in any of the considered marginalization cases.

3 Construction of a spectral template of foreground residuals

In this section, we present our novel procedure for deriving a cleaned template of the foreground residuals (Section 3.1) and describe how its angular power spectrum is incorporated into the cosmological likelihood to debias the tensor-to-scalar ratio (Section 3.2).

3.1 Derivation of the template of the foreground residuals

The key step in deriving a B -mode template of foreground residuals is to obtain cleaned estimates of the Galactic foreground emission across the observed frequency range. These foreground B -mode maps, with reduced contamination from CMB and instrumental noise, are then combined using the same component separation weights employed for CMB recovery in Equation 2.3. This procedure naturally produces a cleaned estimate of the distribution of foreground residuals that remain in the reconstructed CMB solution after component separation. In the proposed approach, the multifrequency Galactic emission data set is obtained by applying the GNILC technique [14, 15] to the simulated microwave maps. Although GNILC has been extensively described in the literature, we briefly summarize here the key steps required to successfully implement the pipeline, given that it constitutes the core methodology of our procedure.

Step 1 and Step 2 of the component separation procedure (see Section 2.2) are first applied to the input Q/U maps. The adopted set of needlet bands in this case is shown in the right panel of Fig. 1.

The GNILC technique is then applied to the pre-processed B -mode needlet maps. At each needlet scale, the following steps are carried out:

1. Two covariance matrices are computed: one from the data as in Equation 2.5,

$$(C_{(j)})_{\nu\nu'}(\hat{n}) = \langle \beta^{(j)}(\hat{n}', \nu) \cdot \beta^{(j)}(\hat{n}', \nu') \rangle_{\hat{n}' \in \mathcal{D}_{\hat{n}}} = (C_{(j)}^{\text{CMB}})_{\nu\nu'}(\hat{n}) + (F_{(j)})_{\nu\nu'}(\hat{n}) + (N_{(j)})_{\nu\nu'}(\hat{n}), \quad (3.1)$$

and one from realistic simulations of the nuisance components to be deprojected (CMB and noise in this case):

$$(\tilde{N}_{(j)})_{\nu\nu'}(\hat{n}) = (C_{(j)}^{\text{CMB}})_{\nu\nu'}(\hat{n}) + (N_{(j)})_{\nu\nu'}(\hat{n}), \quad (3.2)$$

where $C_{(j)}^{\text{CMB}}$, $F_{(j)}$, and $N_{(j)}$ denote the CMB, foreground, and noise covariance contributions, respectively. For simplicity, in the following we omit the (j) and $\nu\nu'$ subscripts.

2. The *whitened* data covariance is then constructed:

$$\tilde{C}(\hat{n}) = (\tilde{N}^{-1/2} C \tilde{N}^{-1/2})(\hat{n}) \simeq (\tilde{N}^{-1/2} F \tilde{N}^{-1/2})(\hat{n}) + \mathbb{1}, \quad (3.3)$$

whose diagonalization yields

$$\tilde{C}(\hat{n}) \simeq (U_F(\hat{n}) \mid U_N(\hat{n})) \begin{pmatrix} \lambda_1(\hat{n}) & & & \\ & \ddots & & \\ & & \lambda_{m(\hat{n})}(\hat{n}) & \\ & & & 1 \\ & & & & \ddots \\ & & & & & 1 \end{pmatrix} \begin{pmatrix} U_F(\hat{n})^T \\ U_N(\hat{n})^T \end{pmatrix}, \quad (3.4)$$

where eigenvalues significantly larger than unity ($\lambda_1, \dots, \lambda_m$) correspond to statistically independent Galactic emission modes that dominate over the CMB and instrumental noise variance at a given sky position $\hat{n}$. These foreground modes are represented by the set of eigenvectors $U_F(\hat{n})$. The rank of the foreground sub-space $m(\hat{n})$, i.e. the number of eigenvalues significantly departing from 1, is derived by minimizing the Akaike Information Criterion (AIC) [15]:

$$\text{AIC}(m, \hat{n}) = 2m + \sum_{k=m+1}^{N_{\text{freq}}} [\lambda_k(\hat{n}) - \log \lambda_k(\hat{n}) - 1], \quad (3.5)$$

with N_{freq} the number of observed frequency channels.

3. Finally, the weights that reconstruct the Galactic emission across all frequency channels are obtained by deprojecting eigenmodes with eigenvalues close to unity in Equation 3.4 (i.e., dominated by CMB and noise) [14, 15]:

$$W(\hat{n}) = [\tilde{F} (\tilde{F}^T C^{-1} \tilde{F})^{-1} \tilde{F}^T C^{-1}](\hat{n}), \quad (3.6)$$

with $\tilde{F} = \tilde{N}^{1/2} U_F$. The weight matrix W has dimension $(N_{\text{freq}} \times N_{\text{freq}} \times N_{\text{pix}})$, where N_{pix} denotes the number of pixels. Applying W to the input needlet maps with dimension $(N_{\text{freq}} \times N_{\text{pix}})$ yields a multifrequency set of cleaned foreground B -mode maps.

For each frequency channel, the cleaned foreground needlet maps at the different scales are then recombined through an inverse needlet transformation [33], following the same procedure adopted in Step 4 of the component-separation pipeline (see Section 2.2). This final step enables the full-scale reconstruction of the Galactic emission.

The GNILC procedure is robust against foreground uncertainties, as no explicit modeling of Galactic emission is implemented or required. On the other hand, the GNILC technique implicitly assumes a perfect characterization of the two-point statistics of the contaminants to be deprojected, since the nuisance covariance $\tilde{N}$ is a required input. In practice, this assumption is quite robust: the expected CMB power spectrum is very well known and the instrumental noise properties can be characterized with high accuracy. The only subtle uncertainty in the B -mode implementation arises from the contribution of primordial tensor modes. In this analysis, we do not explore the impact of potential mismatches between the true tensor contribution and the one assumed in the nuisance covariance. However, we do not expect this to have any significant impact because: (i) tensor modes, given the current upper limits on their amplitude [42, 43], contribute appreciably to the CMB B -mode covariance only on the largest angular scales, which are never considered in isolation within the first needlet scale; and (ii) even if tensor modes are not explicitly included, they are effectively deprojected together with the lensing B -modes, as both share the same frequency scaling.

The cleaned maps of B -mode foreground emission, already homogenized to a common angular resolution of FWHM= 70.5 arcmin, are processed in exactly the same way as the input total maps within the CMB reconstruction pipeline. Specifically, they are convolved with the set of needlet filters used in Step 2 of the component separation procedure (shown in the left panel of Fig. 1). The resulting needlet maps represent our best estimate of the foreground contribution to the input needlet maps used in component separation. At each needlet scale, these maps are combined with the corresponding component separation weights $w_{(j)}(\hat{n})$ as done for input data in Equation 2.3, yielding a scale-specific estimate of the residual foreground contamination in the reconstructed CMB solution. Finally, the residual foreground template in pixel space, $\tilde{f}_{\text{res}}(\hat{n})$, is obtained through an inverse needlet transformation. This procedure, like the component-separation routines, is included in the BROOM package.

For each simulation, an independent set of 100 multifrequency realizations containing only instrumental noise is first propagated through the GNILC pipeline and then combined with the component-separation weights, following the same procedure used for the GNILC templates. The power spectrum of the resulting *noise residuals*, $C_{\ell}^{\tilde{n}}$, provides an estimate of the noise bias in the angular power spectrum of the residuals template.

To decouple the outcomes of the needlet component separation used for CMB reconstruction (NILC) and for foreground estimation (GNILC), we adopt different hyperparameters for the two pipelines. Specifically, we use distinct sets of needlet bands (as shown in the two panels of Fig. 1) and different values of the residual ILC bias (see Appendix A of [16] for further details), which effectively control the size of the domains $\mathcal{D}_{\hat{n}}$ over which covariances are independently computed

(see Equations 2.5 and 3.1). In practice, lower values of the ILC bias correspond to larger circular domains.

Although not fully explored in this work, the GNILC implementation can be further modified to improve the reconstruction of the foreground-residual template. One relevant hyperparameter is the rank of the Galactic subspace, $m(\hat{n})$ (see Equation 3.4), which determines the number of modes retained in the signal subspace. In practice, one can choose to preserve more (or fewer) modes than those indicated by the AIC criterion in Equation 3.5, thereby achieving a better (or poorer) reconstruction of the foreground emission at the cost of enhanced (or reduced) contamination from noise and CMB. A second hyperparameter available to the user is the possibility to enforce a specific level of CMB deprojection. This corresponds to constraining the GNILC weights in Equation 3.6 to remove a fixed fraction of the CMB signal at all frequency channels. However, imposing a stronger deprojection of the CMB signal can, in turn, increase the reconstruction noise due to the additional constraint applied in the weight estimation. Both the rank of the foreground subspace and the level of CMB deprojection can be manually adjusted when applying this methodology within BROOM and may assume different values for distinct needlet bands. In this work, we configure GNILC to preserve $m(\hat{n}) + 1$ independent modes in its reconstruction and to fully deproject the CMB signal at all needlet scales except for $j = 0$ (corresponding to the largest angular scales). This configuration was found to be robust across all considered scenarios.

Finally, we note that, although this work focuses on propagation through the NILC and MC-NILC pipelines, the multifrequency set of foreground templates obtained from the GNILC pipeline can be combined with the weights from any map-based component-separation method—either parametric or blind—implemented in any domain, including pixel, harmonic, or needlet space.

3.2 Spectral model

The spectral template of the foreground residuals, $C_\ell^{f_{res}}$, contaminating C_ℓ^{out} , is derived for each simulation as

$$C_\ell^{f_{res}} = C_\ell^{\tilde{f}_{res}} - \langle C_\ell^{\tilde{n}} \rangle, \quad (3.7)$$

where $C_\ell^{\tilde{f}_{res}}$ is the angular power spectrum of the foreground residuals template obtained with the procedure described in Section 3.1, and $\langle C_\ell^{\tilde{n}} \rangle$ is the average over 100 simulations of the corresponding noise contribution. This second term is required to debias the spectral template from reconstruction noise. This correction is robust, as the debiasing can be performed either by estimating the noise bias from realistic instrumental simulations (as done in this work), or by computing cross-spectra of the map-based templates obtained from different data splits with uncorrelated noise. All spectra are computed using the same routine and masking strategies as described in Section 2.3. The derivation of the tensor-to-scalar ratio posterior is then modified. While the likelihood function and sampling procedure remain unchanged, the power spectrum model of Equation 2.6 is replaced by:

$$C_\ell(r, A_f) = C_\ell^{\text{lens}} + r \cdot C_\ell^{\text{tens}}(r = 1) + \langle C_\ell^{\text{noi}} \rangle + A_f \cdot \langle C_\ell^{f_{res}} \rangle, \quad (3.8)$$

so that the MCMC now samples the two-dimensional posterior $P(r, A_f)$. The additional parameter A_f quantifies any mismatch between the amplitude of the actual foreground residuals and that of the derived template in Equation 3.7. By construction, if the spectral shape of $C_\ell^{f_{res}}$ matches that of the actual foreground residuals, A_f absorbs their contribution to C_ℓ^{out} , thereby mitigating the bias that would otherwise affect the inferred tensor-to-scalar ratio. In this case, we impose the prior $0 \leq A_f < \infty$.

A schematic flowchart summarizing the full procedure proposed and implemented in this work, as applied to each simulation, is shown in Figure 2.




Figure 2. Flowchart of the pipeline used to process each simulated data set, leading to the derived average observed and model angular power spectra, which are then used as inputs to the cosmological likelihood in Equation 2.6.

4 Results

As outlined in Section 2.1, we consider six sets of input microwave simulations, corresponding to three values of the tensor-to-scalar ratio ($r = 0, 0.004, 0.01$) combined with two foreground sky models (d1s1 and d10s5). For each set, 100 independent realizations are generated assuming a *Lite-BIRD*-like instrumental configuration. These simulations are then processed with two component separation pipelines, NILC and (MC-)NILC (see Section 2.2), and the resulting power spectra are computed using two masking strategies (GAL60 and GAL40). Altogether, this amounts to 24 analysis cases. For each case, we report constraints on the tensor-to-scalar ratio, both with and without inclusion of the spectral template of foreground residuals in the model of Equation 2.6.

We begin by comparing the power spectrum of the actual foreground residuals with that of the corresponding template, derived according to Equation 3.7 using both NILC and (MC-)NILC component separation routines. This comparison is, of course, only possible when dealing with simulations rather than real data, but it provides a robust validation test for the proposed methodology. For each simulation and case, a map of the foreground residuals is obtained by processing the input foreground-only maps in the same way as the total data (where all components have been co-added) and by combining them with the component-separation weights of Equation 2.4.

Figure 3 presents the average spectra over all simulations, computed with the GAL60 and GAL40 masks. As an illustrative example, we focus here on the case based on the d1s1 model with $r = 0$ as input. A very good agreement is found between the spectral shape of the power spectrum of the actual foreground residuals and that of the template across the full multipole range, for both component-separation methods and both masking strategies. The same conclusions hold for the additional cases considered but not shown here. As discussed in Section 3.2, the agreement in spectral shape is the most relevant aspect, since any mismatch in amplitude can be accounted for by the additional parameter A_f in Equation 3.8.

We first propagate the observed power spectrum, obtained after component separation, into the cosmological likelihood without including any foreground contribution in the spectral model as in Equation 2.7. The resulting posteriors of r for NILC and (MC-)NILC, obtained for both masking strategies and all considered input values of r , are shown in Figure 4. As in Figure 3, we report results for the d1s1 scenario as a representative example. In all cases, we find that foreground residuals significantly bias the best-fit value of r relative to the input value in the simulations. As expected, the




Figure 3. Angular power spectra, averaged over 100 realizations, of the noise residuals (dotted lines), foreground residuals (solid lines) and denoised template of foreground residuals (dashed lines, obtained with Equation 3.7). Left and right panels report results for NILC and (MC-)NILC, respectively. Orange and red lines correspond to results obtained when GAL60 and GAL40 masks are used to compute angular power spectra. The reported results refer to the case where input $r = 0$. The grey shaded area denotes the range of primordial tensor CMB B -mode power spectra corresponding to tensor-to-scalar ratio values $r \in [0.004, 0.01]$.




Figure 4. Posterior distributions of the tensor-to-scalar ratio derived from the average observed angular power spectrum after applying NILC (top) and (MC-)NILC (bottom) to simulated data sets with the d1s1 foreground model and input CMB signals with $r = 0$ (left), $r = 0.004$ (center), and $r = 0.01$ (right). In these cases, no contribution from foreground residuals is included in the spectral model, as defined in Equation 2.7. The vertical black dashed lines indicate the input nonzero values of the tensor-to-scalar ratio.




Figure 5. Two-dimensional and one-dimensional posteriors (r and A_f , shown in the top and right sub-panels of each panel, respectively) obtained from sampling the likelihood of Equation 2.6 when the spectral template of foreground residuals is included in the model of Equation 3.8. The top and bottom rows correspond to the application of NILC and (MC-)NILC, respectively, to the simulated d1s1 data set. Different input values of the tensor-to-scalar ratio are shown: $r = 0$ (left), $r = 0.004$ (center), and $r = 0.01$ (right). The vertical black dashed lines indicate the input nonzero values of the tensor-to-scalar ratio.

use of the more aggressive GAL40 mask and/or of the (MC-)NILC methodology partially mitigates this bias. This improvement arises, respectively, from the exclusion of broader regions close to the Galactic plane and from the more effective subtraction of foreground contaminants on large angular scales achieved by the MC-NILC framework (see Section 2.2). In all cases shown in Figure 4, the input fiducial value r_{in} lies outside the 95% confidence region, indicating a statistically significant level of contamination.

We then re-sample the cosmological likelihood of Equation 2.6, this time including the spectral template derived in Equation 3.7 within the model of Equation 3.8. The resulting posterior is two-dimensional, with parameters r and A_f , where A_f parametrizes the relative amplitude of the template with respect to the actual foreground residuals in the observed power spectrum. The choice of introducing a simple amplitude rescaling factor A_f in Equation 3.8 to account for potential mismatches between the actual foreground residual contamination and the derived template represents the most agnostic and minimal assumption in the spectral model. This choice is strongly supported by the excellent agreement in the spectral shapes of the two angular power spectra shown in Figure 3.

For the same illustrative cases as those shown without marginalization in Figure 4, the 2D and 1D posterior distributions, obtained when the spectral model of the foreground residuals is included, are presented in Figure 5. As expected, we find a clear detection of foreground residuals, since A_f is not statistically compatible with zero in any case. This, in turn, yields an unbiased estimate of

	$r_{\text{in}} = 0$	$r_{\text{in}} = 0.004$	$r_{\text{in}} = 0.01$
NILC - d1s1 - GAL60	0.0067 ± 0.0009	$0.01^{+0.0011}_{-0.001}$	0.015 ± 0.0012
NILC - d1s1 - GAL60 (marg.)	< 0.00106	0.0039 ± 0.0013	$0.01^{+0.0015}_{-0.0014}$
NILC - d1s1 - GAL40	$0.0042^{+0.0010}_{-0.0009}$	$0.0072^{+0.0013}_{-0.0012}$	$0.0125^{+0.0015}_{-0.0014}$
NILC - d1s1 - GAL40 (marg.)	< 0.00128	0.0039 ± 0.0015	0.01 ± 0.0017
(MC-)NILC - d1s1 - GAL60	$0.0035^{+0.0008}_{-0.0007}$	$0.0077^{+0.0009}_{-0.0008}$	0.0135 ± 0.001
(MC-)NILC - d1s1 - GAL60 (marg.)	< 0.00083	$0.0044^{+0.0012}_{-0.0011}$	0.0104 ± 0.0014
(MC-)NILC - d1s1 - GAL40	$0.0021^{+0.0008}_{-0.0007}$	0.0063 ± 0.001	0.012 ± 0.0012
(MC-)NILC - d1s1 - GAL40 (marg.)	< 0.00084	$0.0041^{+0.0013}_{-0.0012}$	0.0101 ± 0.00015
NILC - d10s5 - GAL60	$0.0084^{+0.0013}_{-0.0012}$	0.012 ± 0.0014	0.0173 ± 0.0016
NILC - d10s5 - GAL60 (marg.)	< 0.002	$0.004^{+0.0021}_{-0.002}$	0.01 ± 0.0023
NILC - d10s5 - GAL40	$0.0065^{+0.0015}_{-0.0014}$	$0.01^{+0.0017}_{-0.0016}$	0.0153 ± 0.0019
NILC - d10s5 - GAL40 (marg.)	< 0.0023	$0.004^{+0.0023}_{-0.0022}$	0.0098 ± 0.0028
(MC-)NILC - d10s5 - GAL60	$0.0056^{+0.001}_{-0.0009}$	0.01 ± 0.001	$0.016^{+0.0012}_{-0.0011}$
(MC-)NILC - d10s5 - GAL60 (marg.)	< 0.0011	$0.0044^{+0.0017}_{-0.0016}$	$0.01^{+0.0025}_{-0.0024}$
(MC-)NILC - d10s5 - GAL40	0.0045 ± 0.0011	0.0088 ± 0.0012	0.015 ± 0.0014
(MC-)NILC - d10s5 - GAL40 (marg.)	< 0.0012	$0.0043^{+0.0019}_{-0.0018}$	0.01 ± 0.0026

Table 1. Constraints on the tensor-to-scalar ratio derived from the posteriors obtained by sampling the likelihood in Equation 2.6. Cases where a spectral template of foreground residuals is included in the model (Equation 3.8) are marked with “(marg.)”. When the null value does not lie within the 68% confidence interval, a detection is reported as the median value with its corresponding 68% confidence interval; otherwise, an upper bound at 68% confidence level is provided.

the tensor-to-scalar ratio, thus effectively fully mitigating the impact of foreground residuals into cosmological inference of the amplitude of primordial B modes. The result proves robust, holding across different masking strategies, component separation approaches, and input tensor amplitudes. The best-fit value of A_f remains stable across the two masking strategies for all the cases considered, demonstrating the correct rescaling of the templates with the adopted mask, in agreement with that observed in the actual foreground residuals. At the same time, the marginalization procedure retains its effectiveness in fully mitigating the foreground-induced bias across all cases. Interestingly, we note that the correlation pattern between the two fitted parameters varies with the adopted amplitude of the input tensor perturbations (r_{in}), and, to a lesser extent, with the employed masking strategy.

The summary statistics of the r posteriors, obtained with and without including the template of foreground residuals in the spectral model, are reported in Table 1 for all the analysed 24 cases, including those based on the d10s5 scenario. We find that in every case foreground residuals affect the r estimation, leading to significant biases in the tensor-to-scalar ratio. Once these residuals are properly accounted for, we recover an unbiased upper bound for $r = 0$ or a consistent detection when the input r is nonzero. We note that the debiasing of the tensor-to-scalar ratio is accompanied, as expected, by an increase in uncertainty, stemming from the additional parameter introduced in the model and sampled in the MCMC analysis. From this perspective, achieving a good match between the spectral features of the actual foreground residuals and those of the template is crucial, as it minimizes the number of nuisance parameters that must be included in the spectral model of Equation 3.8, thereby reducing the impact on the uncertainty of the inferred cosmological parameters.

An overall comparison between the r constraints obtained for the GAL60 and GAL40 cases, when the foreground spectral model is included, suggests that in terms of overall sensitivity, it is preferable to retain a larger sky fraction and thus marginalize over a correspondingly larger level of contam-

ination rather than exclude it. This conclusion, however, is expected to hold as long as there is a good agreement between the template and the actual foreground residuals. The gain in sensitivity achieved by retaining a larger sky fraction—at the cost of a higher residual amplitude to be marginalized over—is actually underestimated by the likelihood framework adopted in this work. Indeed, every term included in the model of the inverse-Wishart likelihood in Equation 2.6 contributes as additional variance in the parameter inference. This applies, for instance, to the enhanced power from the foreground residuals absorbed by the term $A_f \cdot \langle C_\ell^{\text{res}} \rangle$ in Equation 3.8. Yet this effect arises from the use of an idealized likelihood: in practice, such a contribution is not statistical but almost entirely systematic, and thus should not translate into a corresponding increase in the parameter uncertainties. Consequently, in a more realistic likelihood framework, the impact of the larger foreground residual contribution on the spectrum variance would be smaller, further enhancing the difference between optimal marginalization with a larger retained sky fraction and the opposite case.

The size of the bias on the cosmological parameter naturally depends on the chosen component separation method and masking strategy. However, since the procedure is fully model-independent, it provides a general and robust basis for treating potential foreground residual contamination in future reconstructions of the CMB polarization field.

To further evaluate the robustness of the proposed procedure, we repeat the same analysis for all the considered cases, this time separately examining the power at the reionization and recombination peaks of the tensor spectrum. In practice, this is achieved by sampling the likelihood in Equation 2.6, restricting the multipole ranges to $\ell \leq 10$ and $\ell > 10$, respectively. The corresponding results, obtained with and without marginalization, are presented in Figure 6, separately for the two peaks and for the different values of r_{in} considered in this work. Overall, we find that applying marginalization yields unbiased estimates of the tensor-to-scalar ratio for both peaks and across all examined foreground models, masking strategies, and input amplitudes of the tensor perturbations. This demonstrates the effectiveness of the proposed marginalization approach in mitigating residual foreground contamination in both the reionization and recombination regimes. We observe that, as expected, the overall uncertainty on the tensor-to-scalar ratio is significantly reduced at reionization in the (MC-)NILC framework, owing to the substantial decrease in reconstruction noise achieved through the introduction of tailored sky patches, as also illustrated in Figure 3. The difference in error bars between NILC and (MC-)NILC becomes much smaller at recombination, where the same component-separation approach is employed in both pipelines. We also note the presence of a small residual bias, particularly in the reionization regime at large angular scales, likely arising from inaccuracies in the adopted likelihood formalism, which is not fully valid at very low multipoles in the presence of masking and the resulting multipole correlations.

5 Conclusions

Several previous forecast analyses for future CMB experiments have shown evidence of residual foreground contamination in the reconstructed CMB polarization field, which can bias the inferred cosmological parameters [2, 17, 44–48]. As discussed in Section 1, this issue can be addressed in three distinct (though not mutually exclusive) ways: (i) by improving component-separation techniques; (ii) by refining the masking strategy to better identify and exclude sky regions strongly affected by residuals; (iii) by explicitly accounting for such contamination at the power-spectrum level.

In this work, we mainly focus on the third approach. We note that, in other similar forecast analyses employing alternative map-based component-separation techniques [47, 49], a spectral template of foreground residuals corresponding to the power spectrum of an observed dust-dominated frequency channel has already been successfully incorporated into the cosmological likelihood [2, 45].




Figure 6. Constraints on the tensor-to-scalar ratio, shown as best-fit values with their corresponding 68% confidence intervals. The left and right panels respectively display results obtained when considering only the reionization ($\ell \leq 10$) and recombination ($\ell > 10$) bumps of the tensor spectrum. Results are reported for all combinations of component-separation pipelines (NILC and (MC-)NILC), foreground models (d1s1 and d10s5), and masking strategies (GAL60 and GAL40). Different rows correspond to the various input tensor-to-scalar ratios r_{in} used in the simulations: 0 (top), 0.004 (middle), and 0.01 (bottom). Constraints shown with circles and dashed lines correspond to cases without marginalization, whereas those with squares and solid lines correspond to cases in which a template of the foreground residuals is included in the spectral model. Some of the constraints for the NILC cases in the reionization regime are missing, as the corresponding biases are significant and shift the high-probability interval beyond the considered range for the vertical axis.

In this context, our work represents an advancement in the treatment of foreground residuals by introducing a more agnostic and coherent method for deriving the residual foreground power-spectrum template directly from the data. This template is then incorporated into the cosmological likelihood to mitigate potential biases on cosmological parameters induced by residual foregrounds.

We focus on the specific case of constraining the tensor-to-scalar ratio from cleaned CMB B -mode maps, as this represents the most contamination-prone observable in future analyses of microwave data. The proposed procedure is simple and intuitive, yet proves both effective and robust when validated on realistic simulations of the microwave sky expected for a *LiteBIRD*-like satellite [2].

Cleaned templates of foreground emission at the different observed frequency channels are derived from the multifrequency simulated data set using the well-established blind technique GNILC [14, 15]. These templates are then combined with the component separation weights, yielding a cleaned estimate of the spatial distribution of foreground residuals. The corresponding power spectrum, properly debiased from residual noise, is subsequently included in the spectral model of the cosmological likelihood to account for contamination from component-separation residuals in the observed power spectrum (see Equation 3.7). We consider two slightly different pipelines for the component-separation step, NILC and (MC-)NILC, which differ in the implementation of the minimum-variance combination at large angular scales (see Section 2.2 for further details).

We demonstrate in this analysis that the proposed approach, when applied to *LiteBIRD*-like B -mode simulations, allows us to recover an unbiased posterior of the tensor-to-scalar ratio, regardless of the assumed input value of r (see Figure 5 and Table 1). In the same cases, posteriors obtained without including the foreground residual component in the model show significant biases, highlighting both the effectiveness and the necessity of the proposed approach. This conclusion holds independently of the simulated Galactic foreground model or the adopted masking strategy, underscoring the robustness of the results under varying conditions.

The proposed procedure is general and can be readily extended to other map-based component-separation pipelines, whether parametric [47, 50–52] or blind [44, 53, 54], or instrumental configurations, such as the Simons Observatory [1], whose latest forecasts reported residual biases in the tensor-to-scalar ratio arising from foreground contamination when using the NILC pipeline [45]. Such bias has been shown to be mitigated through tailored (idealized) masking [55]; therefore, the proposed procedure could serve as an alternative means to retain a larger sky fraction, potentially reducing the statistical uncertainty.

The marginalization methodology is based on the reconstruction of cleaned templates of the foreground emission at the various frequency channels used for CMB reconstruction. In practice, it relies on the GNILC approach (see Section 2.2). However, other techniques available in the literature may offer further improvements or enhanced robustness, such as the Wavelet Scattering Transform framework [56, 57] and non-local means methods [58], owing to their ability to exploit higher-order statistics.

The introduced technique promises to be also valuable for other CMB science cases where large-scale foreground contamination, although less significant compared to the CMB signal, may still affect cosmological analyses—for instance, the inference of the optical depth to the last-scattering surface, τ , from the low- ℓ E -mode power spectrum.

The inclusion of marginalization can therefore be applied in any situation where contamination from foreground residuals becomes evident, for instance as indicated by higher-order statistical assessments of the reconstructed CMB maps, such as those explored in [59].

Acknowledgments

AC acknowledges Nicoletta Krachmalnicoff and Carlo Baccigalupi for their valuable comments throughout the development of this project. In this work we made use of the following additional software/packages: [numpy](#)⁴, [astropy](#)⁵, [matplotlib](#)⁶. AC acknowledges partial support by the Italian Space Agency (ASI Grants No. 2020-9-HH.0 and 2016-24-H.1-2018), as well as the Radio-ForegroundsPlus Project HORIZON-CL4-2023-SPACE-01, GA 101135036 and through the Project SPACE-IT-UP by the Italian Space Agency and Ministry of University and Research, Contract Number 2024-5-E.0. This work has also received support by the European Union’s Horizon 2020 research and innovation programme under grant agreement no. 101007633 CMB-Inflate.

References

- [1] P. Ade, J. Aguirre, Z. Ahmed, S. Aiola, A. Ali, D. Alonso et al., *The Simons Observatory: science goals and forecasts*, *J. Cosmology Astropart. Phys.* **2019** (2019) 056 [[1808.07445](#)].
- [2] LiteBIRD Collaboration, E. Allys, K. Arnold, J. Aumont, R. Aurlien, S. Azzoni et al., *Probing cosmic inflation with the LiteBIRD cosmic microwave background polarization survey*, *Progress of Theoretical and Experimental Physics* **2023** (2023) 042F01 [[2202.02773](#)].
- [3] K.N. Abazajian, K. Arnold, J. Austermann, B.A. Benson, C. Bischoff, J. Bock et al., *Inflation physics from the cosmic microwave background and large scale structure*, *Astroparticle Physics* **63** (2015) 55 [[1309.5381](#)].
- [4] M. Kamionkowski and E.D. Kovetz, *The Quest for B Modes from Inflationary Gravitational Waves*, *Annual Review of Astronomy and Astrophysics* **54** (2016) 227 [[1510.06042](#)].
- [5] M. Zaldarriaga, *Polarization of the microwave background in reionized models*, *Phys. Rev. D.* **55** (1997) 1822 [[astro-ph/9608050](#)].
- [6] S. Bashinsky and U. Seljak, *Signatures of relativistic neutrinos in CMB anisotropy and matter clustering*, *Phys. Rev. D.* **69** (2004) 083002 [[astro-ph/0310198](#)].
- [7] J. Lesgourgues and S. Pastor, *Massive neutrinos and cosmology*, *Phys. Rep.* **429** (2006) 307 [[astro-ph/0603494](#)].
- [8] D. Baumann, D. Green, J. Meyers and B. Wallisch, *Phases of new physics in the CMB*, *J. Cosmology Astropart. Phys.* **2016** (2016) 007 [[1508.06342](#)].
- [9] Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi et al., *Planck 2018 results. VI. Cosmological parameters*, *Astronomy and Astrophysics* **641** (2020) A6 [[1807.06209](#)].
- [10] M. Kamionkowski, *How to Derotate the Cosmic Microwave Background Polarization*, *Physical Review Letters* **102** (2009) 111302 [[0810.1286](#)].
- [11] Y. Minami and E. Komatsu, *New Extraction of the Cosmic Birefringence from the Planck 2018 Polarization Data*, *Physical Review Letters* **125** (2020) 221301 [[2011.11254](#)].
- [12] M. Kamionkowski, A. Kosowsky and A. Stebbins, *A Probe of Primordial Gravity Waves and Vorticity*, *Physical Review Letters* **78** (1997) 2058 [[astro-ph/9609132](#)].
- [13] M. Zaldarriaga and U. Seljak, *All-sky analysis of polarization in the microwave background*, *Phys. Rev. D.* **55** (1997) 1830 [[astro-ph/9609170](#)].
- [14] M. Remazeilles, J. Delabrouille and J.-F. Cardoso, *Foreground component separation with generalized Internal Linear Combination*, *Monthly Notices of the Royal Astronomical Society* **418** (2011) 467 [[1103.1166](#)].

⁴[numpy.org/](#)

⁵[astropy.org/](#)

⁶[matplotlib.org/](#)

- [15] Planck Collaboration, N. Aghanim, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini et al., *Planck intermediate results. XLVIII. Disentangling Galactic dust emission and cosmic infrared background anisotropies*, *Astronomy and Astrophysics* **596** (2016) A109 [[1605.09387](#)].
- [16] J. Delabrouille, J.-F. Cardoso, M. Le Jeune, M. Betoule, G. Fay and F. Guilloux, *A full sky, low foreground, high resolution CMB map from WMAP*, *Astronomy and Astrophysics* **493** (2009) 835 [[0807.0773](#)].
- [17] A. Carones, M. Migliaccio, G. Puglisi, C. Baccigalupi, D. Marinucci, N. Vittorio et al., *Multiclustering needlet ILC for CMB B-mode component separation*, *Monthly Notices of the Royal Astronomical Society* **525** (2023) 3117 [[2212.04456](#)].
- [18] K.M. Górski, E. Hivon, A.J. Banday, B.D. Wandelt, F.K. Hansen, M. Reinecke et al., *HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere*, *The Astrophysical Journal* **622** (2005) 759 [[astro-ph/0409513](#)].
- [19] A. Zonca, L. Singer, D. Lenz, M. Reinecke, C. Rosset, E. Hivon et al., *healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python*, *The Journal of Open Source Software* **4** (2019) 1298.
- [20] A. Lewis, A. Challinor and A. Lasenby, *Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models*, *The Astrophysical Journal* **538** (2000) 473 [[astro-ph/9911177](#)].
- [21] A. Starobinsky, *A new type of isotropic cosmological models without singularity*, *Physics Letters B* **91** (1980) 99.
- [22] N. Krachmalnicoff, E. Carretti, C. Baccigalupi, G. Bernardi, S. Brown, B.M. Gaensler et al., *S-PASS view of polarized Galactic synchrotron at 2.3 GHz as a contaminant to CMB observations*, *Astronomy and Astrophysics* **618** (2018) A166 [[1802.01145](#)].
- [23] J.L. Weiland, G.E. Addison, C.L. Bennett, M. Halpern and G. Hinshaw, *Polarized Synchrotron Foreground Assessment for CMB Experiments*, *The Astrophysical Journal* **936** (2022) 24 [[2203.11445](#)].
- [24] Planck Collaboration, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini et al., *Planck 2018 results. IV. Diffuse component separation*, *Astronomy and Astrophysics* **641** (2020) A4 [[1807.06208](#)].
- [25] S.K. Choi and L.A. Page, *Polarized galactic synchrotron and dust emission and their correlation*, *J. Cosmology Astropart. Phys.* **2015** (2015) 020 [[1509.05934](#)].
- [26] B. Thorne, J. Dunkley, D. Alonso and S. Naess, *The Python Sky Model: software for simulating the Galactic microwave sky*, *Monthly Notices of the Royal Astronomical Society* **469** (2017) 2821 [[1608.02841](#)].
- [27] A. Zonca, B. Thorne, N. Krachmalnicoff and J. Borrill, *The Python Sky Model 3 software*, *The Journal of Open Source Software* **6** (2021) 3783 [[2108.01444](#)].
- [28] The Pan-Experiment Galactic Science Group, :, J. Borrill, S.E. Clark, J. Delabrouille, A.V. Frolov et al., *Full-sky Models of Galactic Microwave Emission and Polarization at Sub-arcminute Scales for the Python Sky Model*, *arXiv e-prints* (2025) [arXiv:2502.20452](#) [[2502.20452](#)].
- [29] C.L. Bennett, D. Larson, J.L. Weiland, N. Jarosik, G. Hinshaw, N. Odegard et al., *Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results*, *ApJ Supplements* **208** (2013) 20 [[1212.5225](#)].
- [30] M.A. Miville-Deschênes, N. Ysard, A. Lavabre, N. Ponthieu, J.F. Macías-Pérez, J. Aumont et al., *Separation of anomalous and synchrotron emissions using WMAP polarization data*, *Astronomy and Astrophysics* **490** (2008) 1093 [[0802.3345](#)].
- [31] Planck Collaboration, R. Adam, P.A.R. Ade, N. Aghanim, M.I.R. Alves, M. Arnaud et al., *Planck 2015 results. X. Diffuse component separation: Foreground maps*, *Astronomy and Astrophysics* **594** (2016) A10 [[1502.01588](#)].

- [32] C.L. Bennett, R.S. Hill, G. Hinshaw, M.R. Nolta, N. Odegard, L. Page et al., *First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Foreground Emission*, *ApJ Supplements* **148** (2003) 97 [[astro-ph/0302208](#)].
- [33] D. Marinucci, D. Pietrobon, A. Balbi, P. Baldi, P. Cabella, G. Kerkycharian et al., *Spherical needlets for cosmic microwave background data analysis*, *Monthly Notices of the Royal Astronomical Society* **383** (2008) 539 [[0707.0844](#)].
- [34] E. Hivon, K.M. Górski, C.B. Netterfield, B.P. Crill, S. Prunet and F. Hansen, *MASTER of the Cosmic Microwave Background Anisotropy Power Spectrum: A Fast Method for Statistical Analysis of Large and Complex Cosmic Microwave Background Data Sets*, *The Astrophysical Journal* **567** (2002) 2 [[astro-ph/0105302](#)].
- [35] G. Polenta, D. Marinucci, A. Balbi, P. de Bernardis, E. Hivon, S. Masi et al., *Unbiased estimation of an angular power spectrum*, *J. Cosmology Astropart. Phys.* **2005** (2005) 001 [[astro-ph/0402428](#)].
- [36] D. Alonso, J. Sanchez, A. Slosar and LSST Dark Energy Science Collaboration, *A unified pseudo- C_ℓ framework*, *Monthly Notices of the Royal Astronomical Society* **484** (2019) 4127 [[1809.09603](#)].
- [37] D. Geller and A. Mayeli, *Continuous wavelets on compact manifolds*, *Mathematische Zeitschrift* **262** (2008) 895.
- [38] S. Hamimeche and A. Lewis, *Likelihood analysis of CMB temperature and polarization power spectra*, *Phys. Rev. D* **77** (2008) 103013 [[0801.0554](#)].
- [39] M. Gerbino, M. Lattanzi, M. Migliaccio, L. Pagano, L. Salvati, L. Colombo et al., *Likelihood methods for CMB experiments*, *Frontiers in Physics* **8** (2020) 15 [[1909.09375](#)].
- [40] D. Foreman-Mackey, D.W. Hogg, D. Lang and J. Goodman, *emcee: The MCMC Hammer*, *Publications of the Astronomical Society of the Pacific* **125** (2013) 306 [[1202.3665](#)].
- [41] G. Puglisi, G. Mihaylov, G.V. Panopoulou, D. Poletti, J. Errard, P.A. Puglisi et al., *Improved galactic foreground removal for B-mode detection with clustering methods*, *Monthly Notices of the Royal Astronomical Society* **511** (2022) 2052 [[2109.11562](#)].
- [42] M. Tristram, A.J. Banday, K.M. Górski, R. Keskitalo, C.R. Lawrence, K.J. Andersen et al., *Improved limits on the tensor-to-scalar ratio using BICEP and Planck data*, *Phys. Rev. D* **105** (2022) 083524 [[2112.07961](#)].
- [43] G. Galloni, N. Bartolo, S. Matarrese, M. Migliaccio, A. Ricciardone and N. Vittorio, *Updated constraints on amplitude and tilt of the tensor primordial spectrum*, *J. Cosmology Astropart. Phys.* **2023** (2023) 062 [[2208.00188](#)].
- [44] M. Remazeilles, A. Rotti and J. Chluba, *Peeling off foregrounds with the constrained moment ILC method to unveil primordial CMB B modes*, *Monthly Notices of the Royal Astronomical Society* **503** (2021) 2478 [[2006.08628](#)].
- [45] K. Wolz, S. Azzoni, C. Hervías-Caimapo, J. Errard, N. Krachmalnicoff, D. Alonso et al., *The Simons Observatory: Pipeline comparison and validation for large-scale B-modes*, *Astronomy and Astrophysics* **686** (2024) A16 [[2302.04276](#)].
- [46] R. Aurlien, M. Remazeilles, S. Belkner, J. Carron, J. Delabrouille, H.K. Eriksen et al., *Foreground separation and constraints on primordial gravitational waves with the PICO space mission*, *J. Cosmology Astropart. Phys.* **2023** (2023) 034 [[2211.14342](#)].
- [47] A. Rizzieri, C. Leloup, J. Errard and D. Poletti, *Cleaning Galactic foregrounds with spatially varying spectral dependence from CMB observations with fgbuster*, *arXiv e-prints* (2025) arXiv:2510.08534 [[2510.08534](#)].
- [48] F. Bianchini, D. Beck, W.L. Kimmy Wu, Z. Ahmed, S. Belkner, J. Carron et al., *CMB-S4: Foreground-Cleaning Pipeline Comparison for Measuring Primordial Gravitational Waves*, *arXiv e-prints* (2025) arXiv:2502.04300 [[2502.04300](#)].

- [49] R. Stompor, S. Leach, F. Stivoli and C. Baccigalupi, *Maximum likelihood algorithm for parametric component separation in cosmic microwave background experiments*, *Monthly Notices of the Royal Astronomical Society* **392** (2009) 216 [[0804.2645](#)].
- [50] R. de Belsunce, S. Gratton and G. Efstathiou, *Testing for spectral index variations in polarized CMB foregrounds*, *Monthly Notices of the Royal Astronomical Society* **517** (2022) 2855 [[2205.13968](#)].
- [51] R. de Belsunce, S. Gratton and G. Efstathiou, *B-mode constraints from Planck low-multipole polarization data*, *Monthly Notices of the Royal Astronomical Society* **518** (2023) 3675 [[2207.04903](#)].
- [52] H.K. Eriksen, J.B. Jewell, C. Dickinson, A.J. Banday, K.M. Górski and C.R. Lawrence, *Joint Bayesian Component Separation and CMB Power Spectrum Estimation*, *The Astrophysical Journal* **676** (2008) 10 [[0709.1058](#)].
- [53] R. Fernández-Cobos, P. Vielva, R.B. Barreiro and E. Martínez-González, *Multiresolution internal template cleaning: an application to the Wilkinson Microwave Anisotropy Probe 7-yr polarization data*, *Monthly Notices of the Royal Astronomical Society* **420** (2012) 2162 [[1106.2016](#)].
- [54] A. Carones and M. Remazeilles, *Optimization of foreground moment deprojection for semi-blind CMB polarization reconstruction*, *J. Cosmology Astropart. Phys.* **2024** (2024) 018 [[2402.17579](#)].
- [55] A. Carones, M. Migliaccio, D. Marinucci and N. Vittorio, *Analysis of Needlet Internal Linear Combination performance on B-mode data from sub-orbital experiments*, *Astronomy and Astrophysics* **677** (2023) A147 [[2208.12059](#)].
- [56] J.M. Delouis, E. Allys, E. Gauvrit and F. Boulanger, *Non-Gaussian modelling and statistical denoising of Planck dust polarisation full-sky maps using scattering transforms*, *Astronomy and Astrophysics* **668** (2022) A122 [[2207.12527](#)].
- [57] L. Mousset, E. Allys, M.A. Price, J. Aumont, J.M. Delouis, L. Montier et al., *Generative models of astrophysical fields with scattering transforms on the sphere*, *Astronomy and Astrophysics* **691** (2024) A269 [[2407.07007](#)].
- [58] G.F. Quispe Peña and A.V. Frolov, *Astrophysical foreground cleanup using non-local means*, *arXiv e-prints* (2023) arXiv:2306.00211 [[2306.00211](#)].
- [59] C. Ranucci, A. Carones, L. Vacher, N. Krachmalnicoff and C. Baccigalupi, *B-sure I: Minkowski functionals as robustness test for tensor-to-scalar ratio detection from CMB observations*, *arXiv e-prints* (2025) arXiv:2510.18868 [[2510.18868](#)].