

PAPER • OPEN ACCESS

A hybrid map- C_ℓ component separation method for primordial CMB B -mode searches

To cite this article: S. Azzoni *et al* JCAP03(2023)035

View the [article online](#) for updates and enhancements.

You may also like

- [Dipole modulation of cosmic microwave background temperature and polarization](#)
Shamik Ghosh, Rahul Kothari, Pankaj Jain et al.
- [LARGE PECULIAR MOTION OF THE SOLAR SYSTEM FROM THE DIPOLE ANISOTROPY IN SKY BRIGHTNESS DUE TO DISTANT RADIO SOURCES](#)
Ashok K. Singal
- [A minimal power-spectrum-based moment expansion for CMB \$B\$ -mode searches](#)
S. Azzoni, M.H. Abitbol, D. Alonso et al.

A hybrid map- C_ℓ component separation method for primordial CMB B -mode searches

S. Azzoni,^{a,b} D. Alonso,^a M.H. Abitbol,^a J. Errard^c
and N. Krachmalnicoff^{d,e,f}

^aDepartment of Physics, University of Oxford, Denys Wilkinson Building,
Keble Road, Oxford OX1 3RH, United Kingdom

^bKavli Institute for the Physics and Mathematics of the Universe (Kavli IPMU, WPI),
UTIAS, The University of Tokyo, Kashiwa, Chiba 277-8583, Japan

^cUniversité Paris Cité, CNRS, Astroparticule et Cosmologie, F-75013 Paris, France

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

^eInstitute for Fundamental Physics of the Universe (IFPU),
Via Beirut 2, 34151 Grignano (TS), Italy

^fNational Institute for Nuclear Physics (INFN), Sezione di Trieste,
Via Valerio 2, I-34127 Trieste, Italy

E-mail: susanna.azzoni@physics.ox.ac.uk, david.alonso@physics.ox.ac.uk,
maximilian.abitbol@physics.ox.ac.uk, josquin@apc.in2p3.fr, nkrach@sisssa.it

Received October 28, 2022

Revised February 2, 2023

Accepted February 6, 2023

Published March 14, 2023

Abstract. The observation of the polarised emission from the Cosmic Microwave Background (CMB) from future ground-based and satellite-borne experiments holds the promise of indirectly detecting the elusive signal from primordial tensor fluctuations in the form of large-scale B -mode polarisation. Doing so, however, requires an accurate and robust separation of the signal from polarised Galactic foregrounds. We present a component separation method for multi-frequency CMB observations that combines some of the advantages of map-based and power-spectrum-based techniques, and which is directly applicable to data in the presence of realistic foregrounds and instrumental noise. We demonstrate that the method is able to reduce the contamination from Galactic foregrounds below an equivalent tensor-to-scalar ratio $r_{\text{FG}} \lesssim 5 \times 10^{-4}$, as required for next-generation observatories, for a wide range of foreground models with varying degrees of complexity. This bias reduction is associated with a mild ~ 20 – 30% increase in the final statistical uncertainties, and holds for large sky areas, and for experiments targeting both the reionisation and recombination bumps in the B -mode power spectrum.

Keywords: CMBR experiments, CMBR polarisation, cosmological parameters from CMBR, gravitational waves and CMBR polarization

ArXiv ePrint: [2210.14838](https://arxiv.org/abs/2210.14838)



Contents

1	Introduction	1
2	Methodology	4
2.1	Map-domain cleaning	4
2.1.1	Map-level parametrization	4
2.1.2	Map-level spectral likelihood	6
2.1.3	Map-level residuals	6
2.2	C_ℓ likelihood on deprojected spectral modes	7
2.2.1	Power spectrum model	8
2.2.2	Power spectrum likelihood	9
2.2.3	Residual C_ℓ s and covariance	10
2.3	Summary of the method	11
3	Simulations	12
3.1	Instrumental effects	12
3.2	Gaussian foreground simulations	14
3.3	Realistic foreground simulations	15
3.3.1	“d1sm” simulations	15
3.3.2	Additional simulations	15
4	Validation	16
4.1	Gaussian foreground simulations	16
4.2	Realistic foreground simulations	18
4.3	Larger sky fractions	19
5	Conclusion	23

1 Introduction

The search for B -mode polarization of a primordial origin in the Cosmic Microwave Background (CMB) is currently one of the most compelling fundamental science cases for cosmology. Within the inflationary paradigm, B -mode polarisation can be generated via Thomson scattering in the early-Universe in the presence of tensor perturbations of the spacetime metric. In this context, a detection of the B -mode signal would therefore imply measuring the amplitude of primordial gravitational waves (through the so-called tensor-to-scalar ratio parameter r), which itself contains information about the details of the physical processes giving rise to inflation [1]. Thus, if detected with sufficient sensitivity, this signal would allow us to rule out several families of inflationary and non-inflationary models predicting small (or no) tensor fluctuations. A sufficiently tight upper bound, at the level of $r \lesssim 10^{-3}$, would also allow us to rule out important regions of model space, such as the so-called R^2 models, often based on quantum corrections to gravity [2, 3], which predict $r \sim 1/N^2 \sim 10^{-3}$, where N is the number of e -folds inflation lasts for.

This has motivated the design of several next-generation experiments, both ground-based and space-borne, which will make precise measurements of the polarised CMB sky over a large range of frequencies in the next decade. Some of these facilities, such as the BICEP Array [4], or the Simons Observatory [5], have just started taking data (or will do so within the next couple of years), and current bounds on r will see a sharp improvement before the end of the decade. Due to atmospheric noise and other systematics, which contaminate the largest angular scales, these ground-based experiments will be able to constrain B -modes on intermediate, degree-sized scales ($\ell \sim 100$) through the so-called “recombination bump” of the CMB power spectrum. The LiteBIRD space mission [6–8], on the other hand, will start taking data in 2029, and will be able to reach tighter constraints on r by targetting both the recombination and reionization bumps, the latter of which occurs at larger scales ($\ell \lesssim 10$).

Arguably the largest source of systematic uncertainty for both space and ground experiments will be contamination by the astrophysical B -mode signal from the Milky Way. This polarized Galactic emission is dominated at low frequencies ($\nu \lesssim 60$ GHz) by synchrotron, while polarized thermal dust dominates at higher frequencies [9, 10]. The combination of both dominates the CMB polarization signal on the scales of interest by at least a factor ~ 5 [10, 11] over the whole frequency range. Devising methods to separate this contamination from the CMB signal in a robust and efficient manner, in the presence of practical imperfections in the data taken by real instruments, is therefore of central importance in this quest.

A successful strategy, followed by various ground-based CMB experiments, is the so-called “multi-frequency power spectrum likelihood” approach [12–16]. In these C_ℓ -based methods, the data vector is composed of all the auto- and cross-power spectra (in this case B -power spectra) between sky maps at different frequencies, which is then fitted to a model that incorporates the correlations between all sky components, Galactic and cosmological. Current state-of-the-art constraints from the BICEP-Keck collaboration have been obtained using this methodology. In its simplest incarnation, the method assumes foregrounds are perfectly correlated in frequency (i.e. they have a homogeneous spectrum across the observed sky), and a functional form is assumed for the power spectrum of the foreground amplitudes (generally in the form of a power law, [e.g. 4]). The model is then characterised by a set of constant spectral parameters (e.g. spectral indices, dust temperature), and by the parameters controlling the amplitude power spectra. Spatial variation of spectral parameters, which is physically expected, should however lead to an imperfect correlation between foregrounds at different frequencies. The simplest way to account for this is by allowing for a simple scale-independent frequency decorrelation parameter, as would be the case for uncorrelated spatial variations [15, 17]. Recently a generalisation of this approach was proposed by [18], using a minimal moment expansion of the foreground SEDs to account for both the amplitude and scale dependence of this effect within a consistent physical model. The method was shown to significantly improve over the vanilla approach, being able to mitigate the foreground bias for realistic sky models over areas much larger than that covered by BICEP-Keck. This minimal expansion can be fully generalised, including all potential correlations between different spectral parameter variations and amplitudes, and avoiding any assumptions about their non-Gaussianity. This has been shown [19, 20] to yield unbiased results given a sufficiently large number of frequency channels. This methodology suffers from two main shortcomings: first, it implicitly treats foregrounds as statistically isotropic fields by compressing their information into the power spectrum, ignoring couplings between different angular modes, and approximating the power spectrum covariance as that of a Gaussian field. Secondly, by adopting a simple functional form for the power spectrum of the foreground amplitudes, which are the dominant sky

components, the model may lead to biases if the real sky deviates from this functional form on the range of scales explored. This is particularly relevant for experiments targeting the low- ℓ regime, but it may be important at any scale given enough sensitivity.

At the other end of the spectrum, parametric *map-based methods* aim to describe the sky at the map level, separating the CMB component and using it to constrain cosmological parameters [10, 21–27]. Note that, while these methods mostly exploit the frequency information to separate components at the map level, recently studies have been carried out on techniques to separate out Galactic foreground purely based on their spatial statistics [28, 29]. Map-based methods have several advantages: first, they are manifestly optimal for any specified sky model.¹ Secondly, since all components are modelled at the map level, it is natural and simple to account for spatial variations in the foreground spectra. Finally, foreground amplitudes are reconstructed in each pixel without assuming a specific model for their spatial correlations, thus avoiding potential biases inherent in assuming a functional form for the foreground power spectra. This approach is not without shortcomings, however: first, while it is easy to incorporate spatially varying spectra in the model, controlling the scales over which spectral parameters vary is not straightforward. The most conservative approach, allowing for uncorrelated variations over all pixels in the map, can lead to significant degradation in the final cosmological constraints, especially if only a small number of frequency channels are available [23, 30, 31]. Several methods have been proposed to select larger regions over which spectral parameters can be held constant or develop strong correlations, leading to rich and complex Bayesian forward modelling schemes [23, 32, 33]. The computational complexity of the resulting frameworks, particularly if simultaneously sampling the CMB power spectra, is another disadvantage of these methods. While computing time is often a fair price to pay for improved or more robust constraints, it is often necessary to carry out multiple analyses of the same data before a clear understanding of all relevant sources of systematic uncertainty is achieved. Thus, fast, lightweight methods are extremely useful when confronted with new data. Finally, incorporating the impact of real-world imperfections in the data (e.g. filtering, gain or polarization angle variations) is often not straightforward in map-based methods, or may increase their computational complexity significantly.

In this paper, we explore an alternative method combining features from map-based and C_ℓ -based analyses. In this hybrid approach, foreground amplitudes are determined at the map level, free of assumptions regarding their statistical correlations, assuming homogeneous spectral properties. We then subtract the contribution from these foregrounds from the data, and model the residual foreground contamination (due to imperfect subtraction or spatial SED variations), at the power spectrum level. The method is thus robust to most assumptions (Gaussianity, power spectrum form) about the dominant foreground amplitudes, for which we drop any assumptions about the spatial dependence. We then impose those assumptions only on the smaller foreground residuals, while retaining information about their frequency dependence. The residuals, by their nature, have a much lower amplitude and SNR, and therefore an approximate model for their spatial dependence may be sufficient to describe them. In principle, this will limit the degradation in the final constraints caused by spectral index variations.

We will study the applicability of this method for next-generation ground-based B -mode experiments targetting large sky areas ($f_{\text{sky}} \gtrsim 0.1$) in the presence of inhomogeneous, correlated instrumental noise, for foreground models with varying degrees of complexity.

¹Although C_ℓ methods will always achieve equivalent constraints if all components are Gaussian.

We will also study the potential limitations of the method when employed on larger sky fractions, with consequently larger foreground contamination, quantifying its applicability for future space-based missions. The paper is structured as follows. Section 2 presents the method in detail. The simulations used to validate it are described in section 3, including the instrumental setup assumed. The method is then validated in section 4, where we explore its performance as a function of foreground complexity and sky fraction. The main conclusions of this study are summarised in section 5.

2 Methodology

The component separation method proposed here consists of two stages:

1. We find a set of well-educated constant spectral indices β_S and β_D and clean out the spatially-constant part of the foregrounds at the map-level using the methods described in [23, 30, 32, 34, 35]. In practice, this step returns a filter \mathbf{Q} that de-projects the linear combinations of the data that follow the best-fit SEDs of the foreground sources under consideration (dust and synchrotron in this case). This step is described in section 2.1.
2. We use the results of the previous step to measure the auto- and cross-power spectra of the map residuals at all frequencies after removing the constant-SED foreground components. Then, under the assumption that the residual foreground contamination is small, we can model its contribution to the multi-frequency power spectrum assuming a relatively simple model, with the frequency dependence given by the derivatives of the foreground SEDs with respect to the spectral indices. As done in [18], we assume a power law model for the scale dependence of the foreground residual power spectra. This step is described in section 2.2.

2.1 Map-domain cleaning

2.1.1 Map-level parametrization

At the map level, we model the sky signal using matrix notation as:

$$\mathbf{m} = \mathbf{S} \mathbf{T} + \mathbf{n}, \quad (2.1)$$

where

- \mathbf{m} is the map vector containing the measured polarized sky signal $(Q_\nu(\hat{\boldsymbol{\theta}}), U_\nu(\hat{\boldsymbol{\theta}}))$ at angular coordinates $\hat{\boldsymbol{\theta}}$ in each observing frequency ν .
- \mathbf{S} is the so-called mixing matrix, with dimension $N_\nu \times N_c$, where N_c is the number of components, and N_ν is the number of frequencies. This matrix encodes the emission law $S_\nu^c(\beta_c(\hat{\boldsymbol{\theta}}))$ of each component. We assume the same SEDs for both Q and U .
- \mathbf{T} is a vector of size N_c containing the amplitude of each component for a given pixel.
- \mathbf{n} is a vector containing the instrumental noise at different frequencies for a given pixel. Its covariance matrix will be labelled $\mathbf{N} \equiv \langle \mathbf{n} \mathbf{n}^T \rangle$.

Note that, implicitly, we have assumed that all quantities in eq. (2.1), including the foreground spectra, depend on sky coordinates $\hat{\boldsymbol{\theta}}$.

We consider three components, with the following spectral dependence:

1. *Thermal dust*: dust grains in the interstellar medium are heated by stellar UV radiation, producing emission on microwave frequencies. The alignment of elongated dust grains with the Galactic magnetic field produces a linear polarization perpendicular to the to both the magnetic field and the direction of propagation, making dust the most relevant foreground for B -mode searches on frequencies $\nu \gtrsim 100$ GHz. Thermal dust emission is thought to be well-characterized by a modified black-body spectrum of the form

$$S_\nu^{\text{D}} = \left(\frac{\nu}{\nu_0^{\text{D}}} \right)^{\beta_{\text{D}}} \frac{B_\nu(\Theta_{\text{D}})}{B_{\nu_0^{\text{D}}}(\Theta_{\text{D}})}, \quad (2.2)$$

where β_{D} and Θ_{D} are the dust spectral index and temperature.

2. *Synchrotron*: galactic synchrotron emission is caused by the interaction of high-energy cosmic ray electrons with the Galactic magnetic field. Synchrotron is strongly polarised, and is characterised by a smooth power-law spectrum tracing the energy distribution of cosmic ray electrons. The synchrotron spectrum used here is therefore

$$S_\nu^{\text{S}} = \left(\frac{\nu}{\nu_0^{\text{S}}} \right)^{\beta_{\text{S}}}, \quad (2.3)$$

where β_{S} is the synchrotron spectral index.

3. *CMB*: the spectrum of CMB temperature anisotropies in antenna temperature units is

$$S_\nu^{\text{CMB}} = e^x \left(\frac{x}{e^x - 1} \right)^2, \quad x = \frac{h\nu}{k_B \Theta_{\text{CMB}}}, \quad (2.4)$$

where h is the Planck constant, k_B is the Boltzmann constant, and $\Theta_{\text{CMB}} = 2.7255$ K is the CMB monopole temperature [36]. The CMB spectrum is isotropic and is not normalized at any pivot frequency.

We will consider spatial variations in β_{D} , which takes values $\beta_{\text{D}} \sim 1.6$. The restricted frequency range available to most ground-based experiments, including the Simons Observatory ($\nu \lesssim 280$ GHz) makes B -mode studies almost insensitive to the value of Θ_{D} , and therefore we fix it to $\Theta_{\text{D}} = 19.6$ K here. β_{S} is the synchrotron spectral index, which takes values $\beta_{\text{S}} \sim -3$.

We can separate the spectra \mathbf{S} into two parts: a spatially-independent ‘‘mean’’ $\bar{\mathbf{S}}$ and a perturbation $\Delta\mathbf{S}$, such that eq. (2.1) becomes

$$\mathbf{m} = \bar{\mathbf{S}}\mathbf{T} + \Delta\mathbf{S}\mathbf{T} + \mathbf{n}. \quad (2.5)$$

In the first term, $\bar{\mathbf{S}}$ represents our best guess of the mean spectra of all components. Here, we will determine best-fit constant spectral indices ($\beta_{\text{S}}, \beta_{\text{D}}$) using the map-based spectral likelihood described in section 2.1.2, and $\bar{\mathbf{S}}$ is given by evaluating the spectra above at these best-fit values. The residual $\Delta\mathbf{S}$ includes the spatially-dependent part of the foreground spectra, as well as any difference between the estimated best-fit spectral indices and their true spatial average. The subsequent treatment of $\Delta\mathbf{S}$ does not change as long as these differences are small.

2.1.2 Map-level spectral likelihood

The procedure to obtain the best-fit spectral indices at the map level follows the methods described in [23, 30, 34]. Given the model in eq. (2.1), the posterior marginalised distribution for $\beta \equiv (\beta_D, \beta_S)$ is given by

$$-2 \ln \mathcal{L}_\beta = \left(\bar{\mathbf{S}}^T \mathbf{N}^{-1} \mathbf{m} \right)^T \left(\bar{\mathbf{S}}^T \mathbf{N}^{-1} \bar{\mathbf{S}} \right)^{-1} \left(\bar{\mathbf{S}}^T \mathbf{N}^{-1} \mathbf{m} \right), \quad (2.6)$$

where \mathbf{N} is the noise covariance matrix. As described in [30], this results from analytically marginalising over the linear amplitude parameters \mathbf{T} , and imposing a Jeffreys-like prior that corrects for volume effects in the spectral index parameters due to their non-linear nature in the model.

We obtain the best-fit spectral indices β_{BF} , defining $\bar{\mathbf{S}}$, as the values that maximize this likelihood, $\beta_{\text{BF}} \equiv \text{argmax}(\mathcal{L}_\beta)$, assuming that they are homogeneous across the sky. Furthermore, we use a Fisher-matrix approach to estimate the uncertainty on the spectral parameters. The inverse covariance of β_{BF} is therefore approximated as the Hessian of the log-posterior at the best fit:

$$(\text{Cov}_\beta)_{ij}^{-1} \simeq - \left. \frac{\partial^2 \ln \mathcal{L}_\beta}{\partial \beta_i \partial \beta_j} \right|_{\beta=\beta_{\text{BF}}} \quad (2.7)$$

This approximation is accurate for sufficiently constraining data [37, 38] (as is the case of SO), and allows us to avoid sampling the full two-dimensional parameter space.

A number of technical aspects regarding the implementation of this method in practice must be noted at this stage. First, we estimate that the data is provided in the form of splits, i.e. a set of maps containing the same sky but different noise realisations (e.g. corresponding to observations at different times). We use differences between these splits, containing only noise, to estimate the noise covariance matrix \mathbf{N} . Moreover, we assume, for simplicity, that the covariance matrix is diagonal in pixel space (i.e. noise different pixels are uncorrelated). In this paper we also assume that the noise is diagonal in frequency space, although accounting for off-diagonal correlations would pose no additional difficulty to the method. It is worth emphasizing that, even though our implementation of the method assumes uncorrelated noise for computational efficiency, this is never the case in real observations, nor in the simulations used here. As we will see, this does not spoil the validity of the method, since this first step only requires us to find spectral index values that are sufficiently close to the spatial mean over the observed footprint. Note that we do account for inhomogeneity in the noise component, using the hits count map. After calculating \mathbf{N} , we coadd all data splits into a single set of maps, which we then process as described above to find $\bar{\mathbf{S}}$.

2.1.3 Map-level residuals

After determining the best-fit spectral indices and the corresponding mixing matrix $\bar{\mathbf{S}}$, the corresponding best-fit amplitudes can be found analytically by solving eq. (2.5) as a least-squared problem:

$$\mathbf{T}_{\text{BF}} = \left(\bar{\mathbf{S}}^T \mathbf{N}^{-1} \bar{\mathbf{S}} \right)^{-1} \bar{\mathbf{S}}^T \mathbf{N}^{-1} \mathbf{m}. \quad (2.8)$$

It is worth emphasizing that, at this stage, we have not imposed any priors on the statistical correlations of foreground amplitudes, and instead treated them as an independent parameter in each pixel. This is one of the key differences with the usual power-spectrum-based methods, where a functional form for the foreground power spectra is often assumed.

From \mathbf{T}_{BF} in eq. (2.8) we can get a best-fit estimate of the foreground contribution to our observations as $\mathbf{m}_{\text{BF}}^{\text{FG}} = \bar{\mathbf{S}}\mathbf{P}\mathbf{T}_{\text{BF}}$, where $\mathbf{P} \equiv \text{diag}(1, 1, \dots, 1, 0)$ is an $N_c \times N_c$ diagonal matrix that selects only the non-CMB components of \mathbf{T}_{BF} (we have assumed the CMB to be the last component). We can then subtract $\mathbf{m}_{\text{BF}}^{\text{FG}}$ from the data to obtain the residual maps:

$$\mathbf{r} \equiv \mathbf{m} - \mathbf{m}_{\text{BF}}^{\text{FG}} = \mathbf{Q}\mathbf{m}, \quad (2.9)$$

where the matrix \mathbf{Q} is

$$\mathbf{Q} \equiv \mathbf{1} - \bar{\mathbf{S}}\mathbf{P} \left(\bar{\mathbf{S}}^T \mathbf{N}^{-1} \bar{\mathbf{S}} \right)^{-1} \bar{\mathbf{S}}^T \mathbf{N}^{-1}. \quad (2.10)$$

The matrix \mathbf{Q} is a projector ($\mathbf{Q}^2 = \mathbf{Q}$), and satisfies the nice property that $\mathbf{Q}\bar{\mathbf{S}} = \bar{\mathbf{S}}(\mathbf{1} - \mathbf{P})$, where $\mathbf{1} - \mathbf{P}$ selects only the CMB component. Note also that, since the spatial and frequency dependence of \mathbf{N} is factorisable, \mathbf{Q} is constant across the sky. Finally, it is straightforward to show that $\text{Tr}(\mathbf{Q}) = N_\nu - N_c + 1$. Thus, in our case, with $N_c = 3$, \mathbf{Q} projects the N_ν -dimensional dataset onto a subspace of dimension $N_\nu - 2$, eliminating two independent modes corresponding to synchrotron and dust assuming constant spectral indices (i.e. $\mathbf{Q}\bar{\mathbf{S}}^{\text{D}} = \mathbf{Q}\bar{\mathbf{S}}^{\text{S}} = 0$).

Substituting our model for \mathbf{m} (eq. (2.1)) into eq. (2.9), and using the properties of \mathbf{Q} , we obtain

$$\mathbf{r} = \mathbf{Q} \left(\bar{\mathbf{S}}\mathbf{T} + \Delta\mathbf{S}\mathbf{T} + \mathbf{n} \right), \quad (2.11)$$

$$= \mathbf{m}_{\text{CMB}} + \mathbf{Q}\Delta\mathbf{S}\mathbf{T} + \mathbf{Q}\mathbf{n}, \quad (2.12)$$

where $(\mathbf{m}_{\text{CMB}})_\nu(\hat{\boldsymbol{\theta}}) \equiv S_\nu^{\text{CMB}} T_{\text{CMB}}(\hat{\boldsymbol{\theta}})$ is the CMB component of the maps. Since the CMB row of $\Delta\mathbf{S}$ is zero (given that the CMB spectrum is constant), the second term in the equation above only contains foreground sources.

Assuming that the fluctuations in the spectral indices with respect to their constant best-fit values are small, we can use a linear expansion of $\Delta\mathbf{S}$ around β_{BF} (as in [18]):

$$\Delta S_\nu^c(\hat{\boldsymbol{\theta}}) = \partial_{\beta_c} \bar{S}_\nu^c \delta\beta_c(\hat{\boldsymbol{\theta}}), \quad (2.13)$$

where $\partial_{\beta_c} \equiv \partial/\partial\beta_c$, and we do not implicitly sum over c . Then, we can rewrite eq. (2.12) as:

$$r_\nu(\hat{\boldsymbol{\theta}}) = S_\nu^{\text{CMB}} T_{\text{CMB}}(\hat{\boldsymbol{\theta}}) + \tilde{S}_\nu^c \tilde{T}_c(\hat{\boldsymbol{\theta}}) + \tilde{n}_\nu(\hat{\boldsymbol{\theta}}), \quad (2.14)$$

where the modified spectra and amplitudes are:

$$\tilde{S}_\nu^c \equiv Q_\nu^{\nu'} \partial_{\beta_c} \bar{S}_{\nu'}^c, \quad \tilde{T}_c(\hat{\boldsymbol{\theta}}) \equiv \delta\beta_c(\hat{\boldsymbol{\theta}}) T_c(\hat{\boldsymbol{\theta}}), \quad (2.15)$$

and we implicitly sum over ν' .

In other words, if the spectral indices are truly constant across the sky, and if the inferred value of β is equal to the true one, then \mathbf{r} in eq. (2.9) is a set of multi-frequency maps, each containing the same copy of the CMB. Otherwise each map in eq. (2.9) contains a spectral residual, with a frequency dependence given by the first-order expansion of each component's SED. We will then model and marginalise over these residuals at the power spectrum level, as described in the next section.

2.2 C_ℓ likelihood on deprojected spectral modes

To separate the CMB contribution at the level of the power spectrum, we use a multi-frequency power spectrum likelihood, similar to that used in the latest analysis carried out by the BICEP-*Keck* collaboration [4, 14]. In this case, the data vector is the full matrix of cross-power spectra $C_\ell^{\nu\nu'}$ between all pairs from a set of frequency maps.² The model used to describe this data vector, and the procedure used to estimate it from the data are described here.

²The code is available at <https://github.com/simonsobs/BBPower>.

2.2.1 Power spectrum model

Here, the map model for the residual described above gets propagated through to power spectrum. We start from a linear model of the form:

$$d_\nu(\hat{\theta}) = \mathcal{S}_\nu^c \mathcal{T}_c(\hat{\theta}) + n_\nu, \quad (2.16)$$

where $d_\nu(\hat{\theta})$ is a set of frequency maps, $\mathcal{T}_c(\hat{\theta})$ is a set of component amplitude maps, and \mathcal{S}_ν^c is a *constant* matrix characterising the spectra of the different components. Here we will explore two possible approaches:

1. In the *baseline* case, d_ν is the raw set of frequency maps, and \mathcal{S}_ν^c and \mathcal{T}_c are the SEDs and amplitude maps of dust, synchrotron, and the CMB.
2. In the *hybrid* case, d_ν is the set of residual maps described in the previous section, and $(\mathcal{S}_\nu^c, \mathcal{T}_c)$ are given by the modified spectra and amplitudes $(\tilde{\mathcal{S}}_\nu^c, \tilde{\mathcal{T}}_c)$ in eq. (2.15), complemented with the CMB SED and amplitude map.

Since we assume the mixing matrix \mathcal{S}_ν^c to be spatially constant, the theoretical prediction for the multi-frequency power spectrum of the map-level model in eq. (2.16) is simply given by

$$C_\ell^{\nu\nu'} = \sum_{c,c'} \mathcal{S}_\nu^c \mathcal{S}_{\nu'}^{c'} C_\ell^{cc'}, \quad (2.17)$$

where $C_\ell^{cc'}$ is the cross-power spectrum between the amplitude maps of components c and c' .

Following the model used by [14], we parametrize $C_\ell^{cc'}$ for the different components described above as follows.

- Foreground auto-correlations are modelled as power-laws of the form:

$$\frac{\ell(\ell+1)}{2\pi} C_\ell^{cc} = A_c \left(\frac{\ell}{\ell_0} \right)^{\alpha_c}, \quad (2.18)$$

with $c \in \{\text{D}, \text{S}\}$, $\ell_0 \equiv 80$.

- The dust-synchrotron cross-correlation is parametrized through a scale-independent correlation coefficient ϵ_{DS} :

$$C_\ell^{\text{DS}} = \epsilon_{\text{DS}} \sqrt{C_\ell^{\text{DD}} C_\ell^{\text{SS}}}, \quad (2.19)$$

- The CMB B -mode power spectrum is parametrized as

$$C_\ell^{\text{CMB}} = A_{\text{lens}} C_\ell^{\text{lens}} + r C_\ell^{\text{tens}} \Big|_{r=1}, \quad (2.20)$$

where C_ℓ^{lens} and $C_\ell^{\text{tens}}|_{r=1}$ are templates for the B -mode power spectrum caused by gravitational lensing and by primordial tensor fluctuations with tensor-to-scalar ratio $r = 1$ respectively. We assume no correlation between the CMB and foreground components.

In both the “baseline” and “hybrid” approaches, the model has a total of 9 free parameters: $\theta = \{A_{\text{lens}}, r, A_{\text{D}}, \alpha_{\text{D}}, \beta_{\text{D}}, A_{\text{S}}, \alpha_{\text{S}}, \beta_{\text{S}}, \epsilon_{\text{DS}}\}$. However, since the bulk of the foreground contamination is removed at the map level in the hybrid approach, we will fix $\epsilon_{\text{DS}} = 0$ in that case, and explore the impact of freeing up this parameter. Note that, although the best-fit

β_D and β_S have been determined in the first step of the “hybrid” method, we allow for them to vary in the power spectrum likelihood, with a prior given by the uncertainties calculated in that step.

It is important to note that the power-law form assumed for the foregrounds power spectra, although routinely adopted by C_ℓ -based methods, is not grounded on any physical principles. The method proposed here addresses this concern, by dropping this assumption for the dominant constant-spectrum component, and only adopting it for the residuals. This is justified if the foreground residual angular power spectra are proportional to those of the total foregrounds, as [31, 39] derive to be accurate to first order. Two qualitative results follow. First, since the foreground residuals are intrinsically smaller in amplitude, we assume that imposing a power spectrum form to the residuals will be less constraining than doing so for the bulk of the foreground contamination (as done by the “baseline” method). We test the validity of this assumption for various sky models in section 4. Secondly, summarising the residuals at the power spectrum level allows us to compress more information into only a few parameters, compared to standard map-based approaches. We then use a physically motivated moment expansion (see eq. (2.14)) to predict the power spectrum shape of the residual, and allow for the parameters of the higher order terms to be free.

2.2.2 Power spectrum likelihood

Given a set of measured power spectra $\hat{C}_\ell^{\nu\nu'}$, and their theoretical prediction outlined in the previous section, construct a Gaussian likelihood of the form

$$\chi^2 \equiv -2 \ln p(\hat{C}_\ell | \theta) = \sum_{a,b} [C_\ell^{ij} - \hat{C}_\ell^{ij}] (\Sigma^{-1})_{ij, mn\ell'} [C_{\ell'}^{mn}(\theta) - \hat{C}_{\ell'}^{mn}], \quad (2.21)$$

where Σ is the covariance matrix of the power spectrum measurements. Note that, strictly speaking, power spectra of Gaussian fields follow a Wishart distribution, which can be approximated as proposed by [40], but we verified that the Gaussian approximation is sufficiently accurate on the scales used here ($30 \leq \ell \leq 300$). This approximation becomes inaccurate on larger scales, where the small number of available modes invalidates the central limit theorem and, as quadratic functions of Gaussian fields, power spectra follow a markedly non-Gaussian distribution. To account for this effect, we employ the method of [40] (“HL” hereon) when including scales $\ell < 30$ (see section 4.3).

As described in [18], we estimate the cross-correlation of the data splits introduced in section 2.1.1 in order to avoid modelling and subtracting the noise bias. We also make use of 100 Gaussian simulations, described in section 3.2, to estimate the covariance matrix Σ . In order to avoid numerical noise in the covariance, we set all off-diagonal elements coupling two ℓ s separated by more than 2 bandpowers to zero. When applying our method to simulations with larger sky areas, in section 4.3, we will approximate the covariance using the “Knox formula” [41]:

$$\Sigma_\ell^{a,b} \equiv \langle \Delta \hat{C}_\ell^{a,b} \Delta \hat{C}_\ell^{c,d} \rangle = \delta_{\ell\ell'} \frac{C_\ell^{a,c} C_\ell^{b,d} + C_\ell^{a,d} C_\ell^{b,c}}{(2\ell + 1) \Delta\ell f_{\text{sky}}}, \quad (2.22)$$

where $\Delta\ell$ is the number of multipoles in the bandpower labelled with multipole ℓ , and f_{sky} is the usable sky fraction.

We measure power spectra using bandpowers of width $\Delta\ell = 10$. By default, we will only include scales in the range $30 \leq \ell \leq 300$, appropriate for a ground-based experiment, affected by atmospheric noise, targeting a B -mode detection from the recombination bump.

2.2.3 Residual C_ℓ s and covariance

In both the “baseline” and “hybrid” approaches, the parameter inference starts from the full set of cross-correlations between all raw frequency maps, \hat{C}_ℓ . In the “baseline” approach, these are compared directly with the model in section 2.2.1. In the “hybrid” approach, however, we must first construct the power spectra of the residual maps.

The residual maps \mathbf{r} are related to the raw maps \mathbf{m} via eq. (2.12), where \mathbf{Q} is the projector matrix defined in eq. (2.10), and estimated in the first step of the method from the map-level likelihood. The power spectrum of the residual maps, C_ℓ^r is thus related to the power spectrum of the raw maps simply via

$$C_\ell^r = \mathbf{Q} \cdot C_\ell \cdot \mathbf{Q}^T. \quad (2.23)$$

Although it would be tempting to simply apply the modelling and likelihood described in the preceding sections to \hat{C}_ℓ^r calculated as above from the raw power spectrum measurements, this is, unfortunately, not possible. Since \mathbf{Q} projects out the two independent modes corresponding to synchrotron and dust with constant spectral indices, the rank of C_ℓ^r is $N_\nu - 2$, and not N_ν , and therefore has null eigenvalues. This will also be the case for the power spectrum covariance, which is therefore not invertible. This makes the Gaussian likelihood in eq. (2.21) numerically unstable.

To solve this problem, rather than making use of pseudo-inverses, we explicitly construct a new basis for our data in the subspace defined by the image of \mathbf{Q} , and express the data and covariance in this lower-dimensional space. Since \mathbf{Q} is a projector, and $\text{Tr}(\mathbf{Q}) = N_\nu - 2$, we can diagonalise it as

$$\mathbf{Q} = \mathbf{B} \mathbf{D} \mathbf{B}^{-1}, \quad (2.24)$$

where \mathbf{B} is a matrix containing the eigenvectors of \mathbf{Q} as columns, and $\mathbf{D} = \text{diag}(1, \dots, 1, 0, 0)$ is the eigenvalue matrix. We have arbitrarily selected the eigenvector order such that the two null eigenvectors are at the end. Let $\tilde{\mathbf{D}}$ be the matrix that results from removing the last two null rows of \mathbf{D} , and define the matrix $\mathbf{R} \equiv \tilde{\mathbf{D}} \mathbf{B}^{-1}$. \mathbf{R} has dimensions $(N_\nu - 2) \times N_\nu$, and preserves the important property of \mathbf{Q} that $\mathbf{R} \tilde{\mathbf{S}}^{\mathbf{D}} = \tilde{\mathbf{R}} \tilde{\mathbf{S}}^{\mathbf{D}} = 0$. Thus \mathbf{R} projects general N_ν -dimensional vectors onto the $(N_\nu - 2)$ -dimensional image of \mathbf{Q} , explicitly reducing the dimensionality of the resulting vector. In order to obtain non-singular residual power spectra and covariances, we therefore use \mathbf{R} instead of \mathbf{Q} .

In summary, given the $N_\nu \times N_\nu$ matrix of multi-frequency power spectra $C_\ell^{\nu\nu'}$, and its covariance matrix, we construct the $(N_\nu - 2) \times (N_\nu - 2)$ residual power spectrum matrix

$$\tilde{C}_\ell^{\alpha\beta} \equiv R_{\alpha\nu} R_{\beta\nu'} C_\ell^{\nu\nu'}, \quad (2.25)$$

and its covariance matrix

$$\text{Cov}(\tilde{C}_\ell^{\alpha\beta}, \tilde{C}_\ell^{\gamma\delta}) = R_{\alpha\nu_1} R_{\beta\nu_2} R_{\gamma\nu_3} R_{\delta\nu_4} \text{Cov}(C_\ell^{\nu_1\nu_2} C_\ell^{\nu_3\nu_4}). \quad (2.26)$$

Note that the transformation in eq. (2.25) is applied to both the input power spectrum data, and to the theoretical prediction described in section 2.2.1.

2.3 Summary of the method

The hybrid method proposed here can be summarised as follows.

1. Starting from a set of N_ν frequency maps \mathbf{m} , we obtain the best-fit spectral index parameters β_{BF} , the associated spectra, and the projector \mathbf{Q} that removes the best-fit dust and synchrotron components from the data, using the map-level spectral likelihood of section 2.1.2. We also construct the $(N_\nu - 2) \times N_\nu$ matrix

$$\mathbf{R} \equiv \tilde{\mathbf{D}} \mathbf{B}^{-1}, \quad (2.27)$$

where \mathbf{B} is the matrix of eigenvectors of \mathbf{Q} , and $\tilde{\mathbf{D}}$ selects only the rows of \mathbf{B}^{-1} corresponding to the non-zero eigenvalues of \mathbf{Q} .

2. We estimate the full set of power spectra between all pairs of frequencies, $\hat{C}_\ell^{\nu\nu'}$, from the same maps, and their covariance matrix.
3. We use \mathbf{R} to deproject the constant dust and synchrotron modes from $\hat{C}_\ell^{\nu\nu'}$ and its covariance, according to eqs. (2.25) and (2.26).
4. We calculate the theoretical prediction for these power spectra according to eq. (2.17), using simple power-law models to describe the scale dependence of the foreground residuals, and the derivative of the foreground SEDs to describe their frequency dependence (see section 2.1.3). We apply the \mathbf{R} matrix to this prediction using again eq. (2.25).
5. We use a Gaussian likelihood to obtain parameter constraints by comparing the data against this theoretical prediction. When inferring parameter constraints, by default we vary the CMB parameters $\{A_{\text{lens}}, r\}$, the foreground power spectrum parameters $\{A_c, \alpha_c\}$ ($c \in \{\text{D}, \text{S}\}$), and the spectral indices $\{\beta_{\text{D}}, \beta_{\text{S}}\}$, imposing a prior on the latter based on the map-level constraints.

The main differences with respect to the “baseline” power-spectrum-based approach are that the latter skips steps 1 and 3 above, applies the power-law modelling for the full foreground power spectra (instead of the residuals), assuming constant spectral indices, and therefore does not make use of the \mathbf{R} matrix to project onto the space of residuals. The assumption of constant spectral indices can be relaxed at the power spectrum level by implementing a moment expansion method, as done in [18], or by allowing for frequency decorrelation [17]. The main shortcomings of these purely C_ℓ -based methods, which the hybrid approach is able to address, are the need to assume a model for the scale dependence of the foreground amplitudes, and to either ignore spectral index variations, or to assume a model for their statistics (e.g. Gaussian, uncorrelated with amplitude variations, or with a flat scale dependence). In turn, the hybrid approach only assumes a model for the scale dependence for the much smaller foreground residuals, limiting the number of parameters needed to model them.

It must be noted that, as described above, the method proposed is not rigorous from a statistical point of view, as it does not propagate the uncertainties of all parameters self-consistently. Instead of splitting the parameter inference into two stages (map-based and C_ℓ -based), a more principled approach, similar to [21], would simultaneously sample the linear foreground and CMB amplitudes, the constant spectral indices, the power spectrum of the foreground residuals, and the CMB power spectrum parameters, in a joint map-level likelihood,

Frequency (GHz)	FWHM (arcmin)	Noise (baseline) (μK -arcmin)	ℓ_{knee} —	α_{knee} —
27	91	35	15	-2.4
39	63	21	15	-2.4
93	30	2.6	25	-2.5
145	17	3.3	25	-3.0
225	11	6.3	35	-3.0
280	9	16	40	-3.0

Table 1. Summary of the beam Full Width at Half Maximum (FWHM) apertures, baseline and goal sensitivity levels for each band of the SO Small Aperture Telescope (SAT) from [5]. The correlated noise power spectrum is parametrized as $N_\ell = N_{\text{white}}[(\ell/\ell_{\text{knee}})^{\alpha_{\text{knee}}} + 1]$.

effectively treating the foreground residuals as a two-dimensional Gaussian process. While fully self-consistent, this kind of approach is significantly more computationally demanding, although not intractable. Instead, we first use the maps to obtain tight constraints on the average spectral indices, and then reuse the data at the power spectrum level to marginalise over the foreground residuals. As we will show, however, in practice we find that the method proposed here is able to produce reliable constraints when applied to simulated data. The main reason for this is the fact that the information contained in the maps, for the instrumental sensitivities explored here, allows for a very precise measurement of the mean spectral indices in the first step. This first step is therefore akin to obtaining an external, high-precision measurement of β , which can be used to deproject the constant foreground amplitudes from the data, rendering the procedure fully self-consistent.

3 Simulations

In order to test the validity of the hybrid component separation method described in the previous section, we test it on a suite of sky simulations. These simulations include the most relevant sky components, as discussed in section 2. In all simulations, the CMB contribution was generated as a Gaussian random field drawn from the power spectrum in eq. (2.20). We use fiducial values ($A_{\text{lens}} = 1$, $r = 0$), unless otherwise stated. The simulations also incorporate the contribution from instrumental noise and the effects of a limited sky coverage as described in section 3.1. We generate simulations of the polarized Stokes Q and U sky maps in six frequency bands as described in table 1. These maps are generated using the HEALPix pixelization scheme [42] with resolution parameter $N_{\text{side}} = 256$.

3.1 Instrumental effects

All the simulations produced in this work include a model of the instrumental noise which replicates the characteristics of the Simons Observatory (SO), as described in [5] (SO19 hereon). The most relevant effects are the scale dependence and inhomogeneity of the noise, and the limited sky coverage.

- *Scale-dependent noise.* We model the noise using the noise calculator released with the data supplement of SO19. This uses a noise power spectrum of the form:

$$N_\ell = N_{\text{white}} \left[\left(\frac{\ell}{\ell_{\text{knee}}} \right)^\alpha + 1 \right], \quad (3.1)$$




Figure 1. Sky mask used in the analysis, proportional to the map of hit counts in Equatorial coordinates used for the SO SATs.

where N_{white} is the white noise variance (in units of $\mu K_{\text{CMB}}^2 \text{srad}$), while ℓ_{knee} and α_{knee} parametrize the scale at which $1/f$ noise starts dominating, and the steepness with which it rises on large scales, respectively. The values for these three parameters used here are listed in table 1.

- *Sky footprint and inhomogeneous noise.* Except in section 4.3, we will assume a sky footprint that follows the hits count map provided with the SO19 data supplement, and shown in figure 1. Inhomogeneous noise realisations are created by generating a homogeneous Gaussian random field with the power spectrum of eq. (3.1), and then scaling the resulting noise maps with the inverse square-root of this hits map. We also use the hits map to construct a differentiable sky mask by first smoothing the hits map with a 1° Gaussian beam, and then applying a 5° ‘C1’ apodization (see [5]). We use this mask to compute all power spectra (except in section 4.3) using a pseudo- C_ℓ algorithm with B -mode purification as implemented in NaMaster³ [43].

Note that we have not included the effects of an instrumental beam in these simulations in order to simplify the analysis. The impact of beams is straightforward to incorporate in the power spectrum analysis, and thus does not impact this method. At the map-level stage, maps should be pre-smoothed to a common beam before obtaining the best-fit spectral parameters. Again, this should not invalidate any of the results presented here.

As mentioned in section 2.1, for each sky simulation, we generate 4 independent noise maps, each with a noise amplitude twice as large as the target sensitivity. This allows us to use these 4 realizations as independent data splits. Therefore, each simulation consists of 24 pairs of (Q, U) maps, corresponding to the 6 frequency channels and 4 data splits. In the map-domain cleaning described in section 2.1, the data vector which enters the spectral likelihood in eq. (2.6) is given by

$$\hat{\mathbf{m}} = \frac{1}{N_{\text{splits}}} \sum_i^{N_{\text{splits}}} (\mathbf{s} + \mathbf{n}_i), \quad (3.2)$$

i.e. the map of the coadded splits. When computing the multi-frequency power spectrum matrix that forms the basis of the C_ℓ -cleaning stage, we start by computing all BB auto-

³<https://github.com/LSSTDESC/NaMaster>.

and cross-correlations between different pairs of maps, for a total of 300 different BB power spectra between all $N_{\text{splits}} = 4$ splits and $N_\nu = 6$ frequency bands. In order to avoid modelling the noise bias as part of the likelihood described in section 2.2), we then generate a coadded power spectrum matrix by averaging over all power spectra combining different data splits for each frequency pair. This results in a collection of 21 distinct coadded power spectra, which form the independent elements of the 6×6 multi-frequency power spectrum matrix.

3.2 Gaussian foreground simulations

We generate a large suite of ‘‘Gaussian’’ foreground realizations. For these, we simulate sky maps for the amplitude ($\mathbf{T}_c(\hat{\boldsymbol{\theta}})$) of each component as Gaussian random fields governed by power spectra following the power-law model in eq. (2.18).

The aim of these Gaussian simulations is twofold. First, we employ them to compute the power-spectrum covariance as described in section 2.2.2. To estimate covariance matrices, we generate a suite of 100 Gaussian simulations with constant spectral indices.

Secondly, we use these Gaussian simulations to study the performance of the method proposed here with different levels of spectral index variation. We parametrize this variation in terms of the standard deviation of $\delta\beta_c(\hat{\boldsymbol{\theta}})$ at the pixel level. This is given in terms of the power spectrum of $\delta\beta_c$ as:

$$\sigma^2(\beta_c) = \sum_{\ell=\ell_{\min}}^{\ell_{\max}} \frac{2\ell+1}{4\pi} C_\ell^{\beta_c}. \quad (3.3)$$

As in [18], we model the power-spectrum of the spectral index variations $C_\ell^{\beta_c}$ as a power-law of the form

$$C_\ell^{\beta_c} = B_c \left(\frac{\ell}{\ell_0} \right)^{\gamma_c}, \quad (3.4)$$

with $\ell_0 = 80$ in all cases. This model is in reasonable agreement with constraints from current measurements of the synchrotron and dust spectral index [44]. The spectral index maps are generated as Gaussian realizations of $C_\ell^{\beta_c}$ in eq. (3.4) with varying values for the amplitude B_c . Instead of varying B_c directly, we generate maps of $\delta\beta_c(\hat{\boldsymbol{\theta}})$ with an arbitrary amplitude and then renormalize them to enforce a given per-pixel standard deviation $\sigma(\beta_c)$. We generate simulations for $\sigma(\beta_c) = \{0, 0.1, 0.2, 0.3\}$ for both synchrotron and dust, which covers. We fix the spectral tilt γ_c to $(\gamma_D, \gamma_S) = (-3.5, -2.5)$. We then add the mean spectral indices $\bar{\beta}_c$ to $\delta\beta_c(\hat{\boldsymbol{\theta}})$ and use the Python Sky Model software (PySM,⁴ [45]) to generate observed sky maps in the six frequency channels of table 1.

All the Gaussian simulations presented here use the same values for the constant spectral indices ($\bar{\beta}_D = 1.6$, $\bar{\beta}_S = -3$) and the same pivot frequencies ($\nu_0^D = 353$ GHz, $\nu_0^S = 23$ GHz). The dust-synchrotron correlation coefficient was set to $\epsilon_{SD} = 0$. We use the foreground parameters that best fit the dust and synchrotron template maps used in the realistic set of simulations, described in section 3.3, within the footprint used here:

$$A_D = 28 \mu\text{K}^2, \quad \alpha_D = -0.16, \quad A_S = 1.6 \mu\text{K}^2, \quad \alpha_S = -0.93. \quad (3.5)$$

⁴https://github.com/bthorne93/PySM_public.

3.3 Realistic foreground simulations

3.3.1 “d1sm” simulations

We produce an additional set of 20 simulations with a higher level of complexity, including the foreground templates provided by PySM. These models include “realistic” non-Gaussian amplitude and spectral index maps that are based on existing observations.

Specifically, for the dust amplitude and spectral index map we adopt the templates included in the **d1** model. The spatially-varying dust spectral index map in **d1** is consistent with the estimate from *Planck* data using the **Commander** component separation code [46]. Similarly, we use the synchrotron amplitude map assumed by the **s1** model. However, we use a modified version of the **s1** model for the synchrotron spectral index β_S . The original **s1** spectral-index map was determined by combining the Haslam 408 GHz map [47, 48] with the WMAP 23 GHz map [49, 50]. It was shown by [11], that the value of β_S and its level of variation is too low compared to the data acquired by the SPASS experiment. We therefore generate a synchrotron spectral index map by re-scaling the **s1** map and extending it to smaller scales using a power-law spectral index power spectrum $C_\ell^{\beta_S}$ consistent with the measurements of [11]. The level of realism included in these simulations is compatible with the one used to assess the performance of forthcoming *B*-mode experiments in e.g. [5, 51, 52]. We label the resulting set of sims “**d1sm**”. Assuming free amplitudes at the first step of our pipeline, as described in section 2.1, should allow us to account for the higher complexity of these models.

Although the level of complexity increases compared to the Gaussian simulations, these simulations contain a comparatively low level of spectral index variation for both dust and synchrotron, compared to the worst-case scenario explored in Gaussian case. For context, the rms fluctuation around the mean of the spectral index maps are

$$\sigma_{\beta_D} = 0.04, \quad \sigma_{\beta_S} = 0.22. \quad (3.6)$$

The foreground amplitude maps, shown in figure 2, used in these simulations reproduce the properties of the power spectrum used to generate the Gaussian realisations described in the previous section (eq. (3.5)) within the footprint use here, and within the scales ($\ell \in [30, 300]$) we use in our main analysis. It should be pointed out that, on scales larger than those analyzed here, the foreground power spectra present obvious departures from a perfect power law [46], but this does not affect our analysis. Therefore, our model of the scale dependence of foreground amplitudes described in section 2.2 will not induce any bias on r for these simulations; the bias will arise because of the spatially-varying spectral indices.

3.3.2 Additional simulations

To further test the validity of the method, we have generated two additional sets of 20 simulations using other models that introduce different aspects of foreground complexity. In these simulations, the synchrotron SED still corresponds to the **sm** model described in section 3.3.

- *Hensley & Draine (HD)*. The thermal dust contribution is modelled using the specifications provided by [53], which we label “HD” hereafter. This is based on a micro-physical model of dust grains, taking into account the strength of the local radiation field, the grain compositions (carbonaceous, and silicate with varying degrees of iron abundance) and the temperature distribution depending on grain size. In these simulations, the synchrotron SED still corresponds to the **sm** model described in section 3.3.




Figure 2. “d1sm” dust and synchrotron polarized amplitude maps (*top row*) and spectral index maps (*bottom row*), in equatorial coordinates.

- *Vansyngel (VS)*. An additional suite of 20 simulations was run using the statistical model described in [17] (labelled “VS”). This model takes into account the spin orientation of the dust grains, in order to simulate the diffuse polarized emission from galactic dust. The three-dimensional structure of the GMF is described by a finite number of layers (we use $N_{\text{layer}} = 7$ layers). The coherent component of the magnetic field is the same in all layers, while its turbulent part is generated as a Gaussian random field. Once the direction of the GMF is determined, maps of the Q and U Stokes parameters are generated by scaling the dust intensity map found by *Planck* [46]. As an additional level of complexity, and in an attempt to describe the three-dimensional distribution of the dust spectral index, we associate each layer with a different Gaussian realization of $\delta\beta_D$ with standard deviation $\sigma_{\beta_D} = 0.13$ (the combined rms variation for 7 layers is $\sigma_{\beta_D} \simeq 0.35$).

4 Validation

We now report on the performance of the hybrid component separation methodology proposed here, as a function of different aspects of foreground complexity. We will present the best-fit and standard deviation on r , obtained with the *hybrid* method, and compare it with the results found using the *baseline* C_ℓ -based method, and with its *moments*-based extension, described in [18].

4.1 Gaussian foreground simulations

We use 20 Gaussian simulations with varying values of σ_β , representative of the range allowed by current data [46], to explore the impact of increased spatial variability of the spectral indices.




Figure 3. *Upper panel:* best-fit values of r for 20 realizations of the sky calculated at different values of σ_{β_c} (the same for synchrotron and dust). Results are shown for the baseline cross- C_ℓ pipeline assuming constant spectral indices (*turquoise*), adding moments (*purple*) and using the hybrid method (*orange*). The position of each simulation in the x axis is shifted slightly from its true σ_{β_c} for clarity. The larger, solid dots, at the centre of each σ_{β_c} value show the mean and standard deviation of each suite of simulations. *Lower panel:* statistical uncertainty σ_r averaged over the twenty realizations in the case of constant spectral indices (*turquoise*), adding moments (*purple*) and using the hybrid method (*orange*). Both the hybrid and the moments method are able to correct the bias on r for all values of σ_{β_c} considered, at the cost of increased final uncertainties with respect to a model with constant spectral indices (which themselves increase monotonically with σ_{β_c}).

The results are shown in figure 3 as a function of σ_β (equal for synchrotron and dust). The upper panel shows the best-fit and standard deviation of β for each simulation, for the baseline, moments, and hybrid methods in turquoise, purple, and orange, respectively. The darker, larger points, show the overall mean and standard deviation of all simulations for each value of σ_β . The lower panel then shows the standard deviation obtained by each method. The mean measured value of r effectively correspond to the bias we get at different levels of spectral variation, since all the simulations were run with an input $r = 0$.

We find that, as σ_β increases, the bias on r obtained by the baseline method grows up to $\delta r = 0.017$ for the highest value of $\sigma_\beta = 0.3$ explored. Since the baseline method achieves uncertainties $\sigma_r \simeq 0.002$, this corresponds to a large, $\sim 8\sigma$ parameter bias. The hybrid method is then able to correct for this bias in all cases, at the cost of a slight increase in the final parameter uncertainty, with σ_r growing to ~ 0.0024 – 0.0033 (a ~ 20 – 50% increase). The moments expansion method is also able to correct this bias, as was shown in [18], although at a slightly higher cost in terms of σ_r , which rises by ~ 50 – 70% .

At SO-like sensitivities, we are able to detect the impact of spatially-varying foreground SEDs, particularly in the case of dust. Figure 4 shows the Signal-to-Noise Ratio (SNR) of the foreground amplitudes using the hybrid method for the Gaussian foreground simulations with increasing spectral index variation ($\sigma_\beta = 0$ – 0.3 , the same for dust and synchrotron). The SNR is calculated for each foreground component (c) as the ratio of the amplitude (A_d^{BB} for dust, A_s^{BB} for synchrotron) to the respective σ obtained from the posterior distribution. We




Figure 4. Signal-to-Noise Ratio (SNR) of the foreground amplitudes, dust A_d^{BB} (in blue) and synchrotron A_s^{BB} (in red), using the hybrid method for the Gaussian foreground simulations with increasing spectral index variation ($\sigma_\beta = 0-0.3$, the same for dust and synchrotron). For both components, the dots correspond to the SNR of each realization, the line is the mean and the shaded area covers the 1σ standard deviation. We find that, at values of $\sigma_\beta > 0.1$, it is possible to significantly detect the effect of the spatial variation of the spectral parameters, particularly the dust spectral index, on the foregrounds multi-frequency power spectra for SO-like sensitivities.

find that, at values of $\sigma_\beta > 0.1$, it is possible to significantly detect the effect of the spatial variation of the dust spectral index for SO-like sensitivities, while a scatter at the level of $\sigma_\beta \gtrsim 0.3$ is required to detect variations in the synchrotron spectral index.

4.2 Realistic foreground simulations

Having used the Gaussian simulations to explore the performance of the method as a function of spatial spectral index variability, we now make use of the realistic foreground simulations described in sections 3.3.1 and 3.3.2 to validate the method in the presence of other aspects of foreground complexity. The `d1sm` simulation includes non-Gaussian and statistically inhomogeneous foreground amplitudes, and realistic spectral index maps. The HD simulations make use of a dust spectral law that is not a modified black-body in detail. Finally, the VS simulations account for the three-dimensional variability in the dust spectral index, and include an enhanced level of non-Gaussianity in the dust amplitude. The results are shown in figure 5 for the three different models (shown in the x axis).

For `d1sm`, the mean best-fit r from 20 realisations using the baseline method is centered at 0.0059 ± 0.0023 , corresponding to a $\sim 2.5\sigma$ bias. With the hybrid method, this reduces to $r = -0.0004 \pm 0.003$. The bias is eliminated at the cost of increasing uncertainties by $\sim 25\%$, as expected from the results found with the Gaussian simulations. Similarly, the moments method recovers unbiased results with a $\sim 40\%$ increase in σ_r .

The conclusions are similar when analysing the HD simulations. The different dust SED used in this model introduces a bias at the $1-2\sigma$ level when analysed with the baseline method, obtaining $r = 0.0033 \pm 0.002$. The hybrid method corrects this bias (we obtain a mean $r = 4.3 \cdot 10^{-5}$), with a widening of the final constraints by $\sim 40\%$. The analysis with the moments method introduces instead a negative 3σ bias. This must be due to the foreground bias departing from the theoretical model used here. The two main assumptions are that the

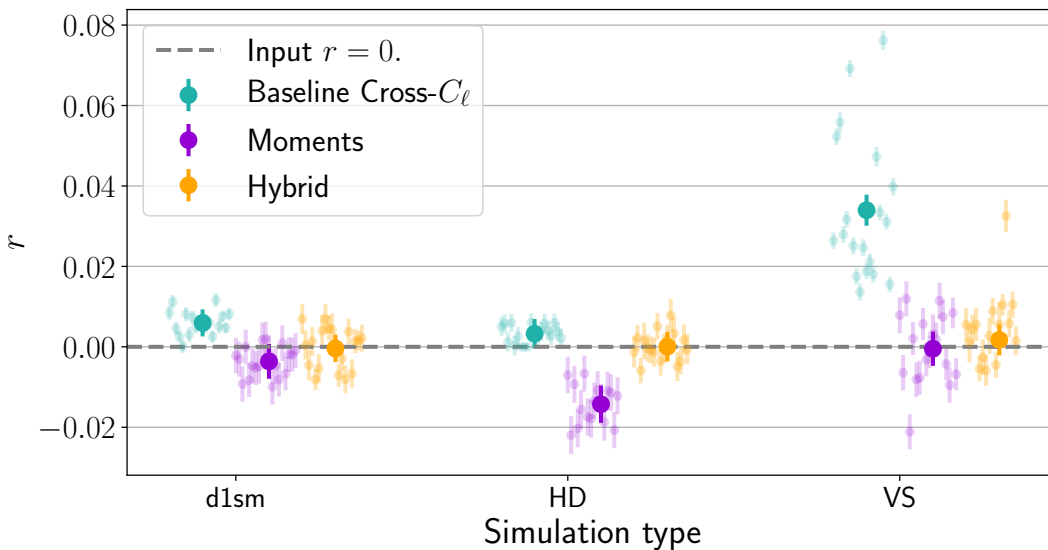


Figure 5. Best-fit values of r for 20 sky simulations for different types of realistic simulations, displayed on the x -axis (d1sm, HD, VS). The orange dots show the baseline cross- C_ℓ results assuming constant spectral indices, while the purple and orange dots show the results of the hybrid method and the moments method, respectively. The larger, solid dots at the centre of each simulation type show the mean of each suite of simulation. The hybrid method is able to correct the bias on r for all the different types of realistic simulations that appears when assuming constant spectral indices. The moments method recovers unbiased results for the d1sm and VS simulations, but introduces a negative 3σ bias in the HD case.

variations in β_D are small enough that the residual spectrum can be described by the Taylor expansion of the modified black-body SED, and that the residuals follow a power-law scale dependence. The additional complexity in the way the non-Gaussian foreground amplitude maps are generated might contribute to either of these causes. It could also be that the simulated data in this case differs significantly from the Gaussian assumption used to construct the covariance matrix and the resulting likelihood is ill behaved.

The significantly more complex VS simulations introduce a much larger bias on r , at the level of 16σ , when using the baseline method. As in the d1sm case, both the hybrid and the moments method eliminate this bias, with an increase in uncertainties by $\sim 30\%$ and $\sim 40\%$ respectively. This is in agreement with the findings of [18]. Note that one of the key advantages of the hybrid method is that it does not rely on any assumption regarding the Gaussianity of the spectral index variations, unlike the moments method, and therefore it provides more flexibility to handle the additional complexity of the VS model. The impact of this additional complexity, however, does not seem to hamper the moments expansion method significantly by comparison with the hybrid approach explored here.

4.3 Larger sky fractions

The ability of the hybrid method to remove the bulk of the foreground contamination at the map level, without assuming a model for their spatial correlations, or for the correlation between them and the foreground residuals resulting from spatially-varying spectra, could potentially allow us to obtain constraints on sky areas significantly larger than the footprint explored so far. This is relevant in the context of future space missions, such as LiteBIRD [6, 8, 55],

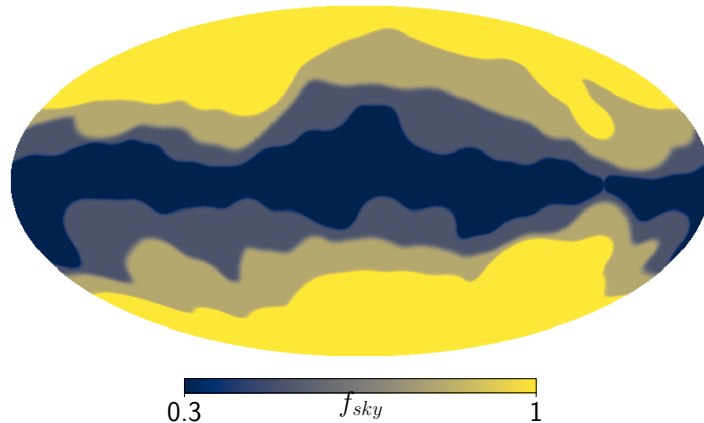


Figure 6. Masks applied to the simulations in section 4.3, which exclude the Galactic plane from the estimation of the data. Three masks are created by setting different threshold on the polarization intensity ($P = \sqrt{Q^2 + U^2}$) of the Planck 353 GHz map [54], to obtain sky cuts $f_{\text{sky}} = 0.3$ (dark blue), 0.6 (lighter blue) and 0.8 (grey). The yellow area corresponds to the observed footprint in galactic coordinates.

or when B -mode constraints become cosmic-variance dominated due to imperfect delensing. We have therefore generated and analysed a new set of simulations covering larger sky areas.

In this case, we generate three new sky masks three masks by setting different thresholds on the polarized intensity ($P = \sqrt{Q^2 + U^2}$) of the Planck 353 GHz map [54], smoothed on a scale of 1° . The thresholds are defined to yield observable sky fractions $f_{\text{sky}} = \{0.3, 0.6, 0.8\}$. The resulting sky masks, avoiding the Galactic plane, are displayed in figure 6. These are smoothed with a 10° FWHM beam and made differentiable with the application of a “C1” apodization with a scale of 5° to the resulting map [43]. The result is similar to the sky mask used in [56]. The larger sky fractions cover regions where foreground contamination is higher, and where spectral index variation is potentially more damaging. For example, in figure 7 we find that the mean SNR of the foreground amplitudes recovered with the hybrid method in the `d1sm` model is $\text{SNR} = 0.2, 2.1, 8.9$ (for dust) and $\text{SNR} = 1.2, 2.7, 7.9$ (for synchrotron) in the $f_{\text{sky}} = 0.3, 0.6$, and 0.8 footprints, respectively.

In order to avoid the computational cost of re-calculating the power spectrum covariance matrix from hundreds of simulations for each sky fraction, here we use the approximate formula in eq. (2.22). Note that this approximation is not accurate enough for actual data analysis, which should involve extensive simulations, but it serves the needs of this forecasting exercise, as illustrated in [5]. The power spectra introduced in the formula contain both signal and noise contributions. Although the approximation accounts for the impact of incomplete sky observations in reducing the number of available independent modes, it does not account for the effects of mode-mode coupling, or the mixing of E and B modes. To get around this limitation, we have also modified the simulations used for this part of the analysis. We remove any E -to- B leakage from the simulations by hand, by transforming the original Q and U maps into E and B , setting the E -mode component to zero, and transforming back to Q and U . Although this degrades the level of realism used in the previous sections, it allows us to account for the main impact of foreground complexity: the large dynamic range of foregrounds on large scale areas, and the departure from perfect power-law behaviour on large scales. For simplicity, in these simulations we assume the same noise power spectrum used before (eq. (3.1)), with a flat sky coverage (i.e. the noise properties are homogeneous throughout the observed footprint).




Figure 7. Signal-to-Noise Ratio (SNR) of the foreground amplitudes, dust A_d^{BB} (in blue) and synchrotron A_s^{BB} (in red), using the hybrid method for the `d1sm` foreground simulations with increasing observed sky fractions ($f_{\text{sky}} 0.3\text{--}0.8$). For both components, the dots correspond to the SNR of each realization, the line is the mean and the shaded area covers the 1σ standard deviation. We find that at increasingly larger sky fractions, the foreground contamination is higher, and we can significantly detect the effect of the spatial variation of the spectral parameters on the foregrounds multi-frequency power spectra at $f_{\text{sky}} > 30\%$.

Figure 8 shows the constraints on r found with the “baseline”, “moments”, and “hybrid” methods, as a function of sky fraction, for the `d1sm` foreground model. Over the 30% footprint, all methods obtain unbiased constraints on r . This, however, changes rapidly as we increase the sky area. At $f_{\text{sky}} = 0.6$ and 0.8 , the baseline method yields $r = 0.004 \pm 0.0008$ and $r = 0.013 \pm 0.0009$, corresponding to a $\sim 5\sigma$ and a $\sim 16\sigma$ bias, respectively. Remarkably, the hybrid method is able to absorb this bias in all cases ($r = 0.00001 \pm 0.0011$ and -0.0007 ± 0.0012 for 60% and 80% f_{sky} respectively), even when using 80% of the sky. The validity of this result, obviously depends on the level of realism of the foreground simulations used here, which may not be sufficiently accurate close to the Galactic plane. However, the result highlights the robustness of the method to high levels of foreground contamination, in comparison with standard approaches. Similarly, the moments method is able to recover unbiased results, recovering $r = 0.0 \pm 0.0013$ and -0.0001 ± 0.0015 for 60% and 80% f_{sky} , but with a 20–30% increase in σ_r compared to the hybrid method constraints.

The availability of a larger sky area allows us to constrain primordial B -modes using larger scales. Figure 9 shows the constraints obtained when using the $f_{\text{sky}} = 0.8$ footprint as a function of the minimum multipole, ℓ_{min} , included in the analysis. The two panels in this figure show the results for simulations with an input $r = 0$ (top), and for $r = 0.01$ (bottom), which thus allows us to study the performance of the method in the presence of a detectable B -mode signal. For our default $\ell_{\text{min}} = 30$, we obtain unbiased constraints on r with the hybrid approach in both cases ($r = -0.0007 \pm 0.0012$ and $r = 0.009 \pm 0.0013$), whereas the baseline approach recovers a biased r at the same level as that displayed in figure 8. When including larger angular scales ($\ell_{\text{min}} < 30$), we use the HL non-Gaussian likelihood [40], as described in section 2.2.2. When extending the scale range to $\ell_{\text{min}} = 10$, the hybrid method obtains $r = 0.0005 \pm 0.0013$ (input $r = 0$) and $r = 0.0099 \pm 0.0013$ (input $r = 0.01$). When including even larger scales, down to $\ell_{\text{min}} = 2$, the hybrid method recovers again unbiased results, i.e.

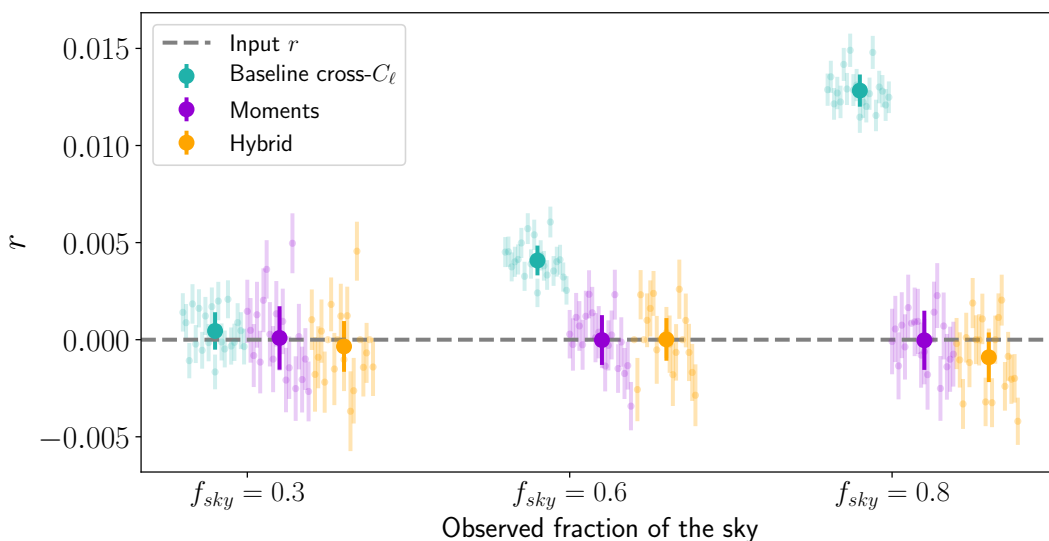


Figure 8. Constraints on r obtained by the “baseline”, “moments”, and “hybrid” methods (in turquoise, purple and orange), as a function of sky fraction, for the `d1sm` foreground model. With sky areas wider than 30% ($f_{\text{sky}} > 0.3$), the baseline pipeline yields biased results. Both the hybrid method and the moments method are able to eliminate this bias, with an increase in σ_r by 30–40% (hybrid) and by 50–80% (moments) with respect to the baseline constraints.

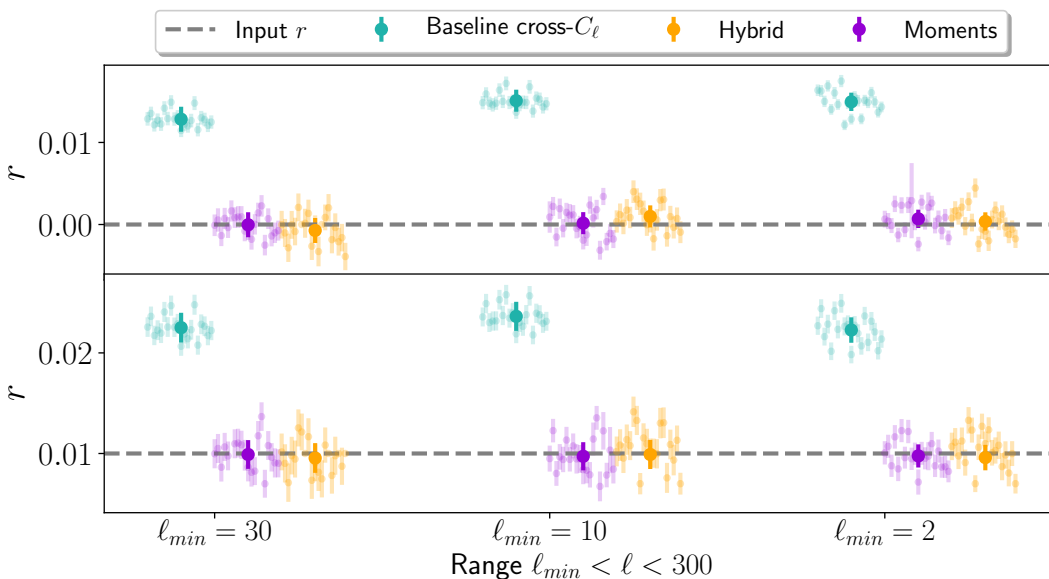


Figure 9. Best-fit values of r for 20 sky simulations observing 80% of the sky ($f_{\text{sky}} = 0.8$) across different ranges of angular scales (from ℓ_{min} up to $\ell = 300$) for input $r = 0$ (*upper panel*) and $r = 0.01$ (*lower panel*). The x -axis displays increasing values of angular scales (lower ℓ_{min}). The red dots show the baseline cross- C_ℓ results, while the green dots show the results of the hybrid method. The larger, solid dots at the centre of each suite of simulation show the mean across simulations. The hybrid method is able to correct the bias on r even at higher angular scales down to $\ell_{\text{min}} = 2$.

$r = 0.0002 \pm 0.0011$ (input $r = 0$) and $r = 0.0096 \pm 0.0012$ (input $r = 0.01$), compared to the 19σ bias introduced by the baseline cross- C_ℓ pipeline. Similarly, the moments method recovers the input r value on all scales, with a $\sim 10\%$ increase in σ_r compared to the hybrid method, and introducing a small negative $< 0.2\sigma$ bias.

5 Conclusion

Detecting the signature of inflation in the CMB polarized B -mode signal is one of the most compelling goals of cosmology, where a strict upper constraint on the large scales amplitudes has the potential to further support or disprove various inflationary models [57].

Many sources of systematic uncertainty contaminate the faint primordial B -modes with the creation of “spurious” B -mode signal. Among these, the most significant astrophysical source is coming from Galactic polarized foregrounds [9]. Therefore, one of the most important steps in the analysis of CMB B -mode data is the separation of the multi-frequency data into its different components. In principle, the separation of CMB radiation and foregrounds can be performed either in real space, dealing with maps of the Q/U Stokes parameters, or in Fourier space, processing the E/B modes power spectra. Different CMB experiments have followed either approach, revealing their advantages and caveats. Whilst achieving satisfactory results on simulated data, current methods are facing additional challenges put forth by the increasing sensitivity of the next-generation CMB experiments, where the precision of the different approaches will become crucial.

Here, we investigate a hybrid approach that combines characteristics from both map-based and C_ℓ -based studies. In the first step of the method, we determine foreground amplitudes at the map level without making any assumptions about their statistical correlations and assuming spatially uniform spectral characteristics. We next take this spatially-homogeneous foreground contribution away from the original data, and the residual foreground contamination (due to inaccurate subtraction or spatial SED fluctuations) is modelled at the power spectrum level. Therefore, only at this stage we This allows us to retain the same C_ℓ -level assumptions, imposing a model for the scale dependence of the foreground amplitudes, while also limiting the degradation in the final constraints brought on by spectral index variations, by applying those assumptions to the smaller foreground residuals (instead of the full foreground spectra) and retaining information about their frequency dependence. Therefore this method allows us to overcome the limitations imposed by C_ℓ approaches (Gaussianity, power spectrum form), limiting the number of parameters needed to model the smaller foreground residuals.

We have applied this method to simulated data including realistic foreground models and instrumental noise. We have shown that we are able to obtain unbiased constraints on r in the presence of realistic (and often pessimistic) foreground complexity. We compare this with the results of a “baseline” fiducial cross- C_ℓ analysis (analogue to the method used to derive the current state-of-the-art constraints on the tensor-to-scalar-ratio r from B -modes [4]), and with the extended “moments” approach described in [18], which is able to further mitigate contamination from foreground residuals. Generally, we find that the hybrid method outperforms the baseline one in terms of bias, and the moments method in terms of σ_r (and bias for some of the most complicated foreground models).

In the simpler foreground scenario, which includes no spatial variation of the spectral parameters (corresponding to the Gaussian simulations with $\sigma_\beta = 0$), we obtain unbiased tensor-to-scalar ratio r with all the three pipelines, with a respective increase in σ_r by $\sim 30\%$ (moments) and $< 10\%$ (hybrid) compared to the baseline approach. In section 4.1, we find

that increasing spectral index variations ($\sigma_\beta = 0.1\text{--}0.3$) can introduce a bias on r of up to 8σ with the baseline method, which both the hybrid and the moments methods are able to correct for at the cost of an increase in statistical uncertainty (the former by $\sim 20\text{--}50\%$, the latter by $\sim 50\text{--}70\%$).

We have also explored more realistic foreground simulations based on various models proposed in the literature (section 4.2) with different levels of complexity. Overall, we find that the hybrid method is able to correct the bias on r for all the different types of realistic simulations that appears when assuming constant spectral indices, at the cost of a limited σ_r increase by $< 30\%$. The moments analysis recovers unbiased results for most of these simulations. However, the reliance of the moments method on simplifying assumptions (Gaussianity of the spectral index variations, uncorrelation among spectral indices and foreground amplitudes) limits its ability to recover unbiased results for one type of complex foreground model explored here.

Unlike the moments method, the hybrid approach has the advantage of not relying on any assumptions about the shape of the spectral index fluctuations, and proceeds by removing the bulk of the foreground contamination at the map level, without assuming a model for their spatial correlations, or for the correlation between them and the foreground residuals resulting from spatially-varying spectra. This allows us to obtain constraints on increasingly larger sky areas and at larger angular scales in section 4.3. Here we find that the results remain largely unbiased, even for experiments covering $\sim 80\%$ of the sky, suggesting the potential application of this method to future satellite missions such as LiteBIRD.

Acknowledgments

We thank Carlo Baccigalupi, Kevin Wolz, Davide Poletti for many helpful discussions. SA is funded by a Kavli/IPMU PhD Studentship. MHA acknowledges support from the Beecroft Trust and Dennis Sciama Junior Research Fellowship at Wolfson College. DA acknowledges support from the Beecroft Trust, and from the Science and Technology Facilities Council through an Ernest Rutherford Fellowship, grant reference ST/P004474. This publication arises from research funded by the John Fell Oxford University Press Research Fund. Some of the results in this paper have been derived using the HEALPix package [42]. The contour plots were generated using the GetDist package [58]. We made extensive use of the numpy [59, 60], scipy [61], healpy [62], and matplotlib [63] python packages.

References

- [1] D.H. Lyth, *What would we learn by detecting a gravitational wave signal in the cosmic microwave background anisotropy?*, *Phys. Rev. Lett.* **78** (1997) 1861 [[hep-ph/9606387](#)] [[INSPIRE](#)].
- [2] A.A. Starobinsky, *Spectrum of relict gravitational radiation and the early state of the universe*, *JETP Lett.* **30** (1979) 682 [[INSPIRE](#)].
- [3] F.L. Bezrukov and M. Shaposhnikov, *The Standard Model Higgs boson as the inflaton*, *Phys. Lett. B* **659** (2008) 703 [[arXiv:0710.3755](#)] [[INSPIRE](#)].
- [4] BICEP2 and KECK ARRAY collaborations, *Improved Constraints on Cosmology and Foregrounds from BICEP2 and Keck Array Cosmic Microwave Background Data with Inclusion of 95 GHz Band*, *Phys. Rev. Lett.* **116** (2016) 031302 [[arXiv:1510.09217](#)] [[INSPIRE](#)].
- [5] SIMONS OBSERVATORY collaboration, *The Simons Observatory: Science goals and forecasts*, *JCAP* **02** (2019) 056 [[arXiv:1808.07445](#)] [[INSPIRE](#)].

- [6] M. Hazumi et al., *LiteBIRD: A Satellite for the Studies of B-Mode Polarization and Inflation from Cosmic Background Radiation Detection*, *J. Low Temp. Phys.* **194** (2019) 443 [INSPIRE].
- [7] H. Sugai et al., *Updated Design of the CMB Polarization Experiment Satellite LiteBIRD*, *J. Low Temp. Phys.* **199** (2020) 1107 [arXiv:2001.01724] [INSPIRE].
- [8] LITEBIRD collaboration, *Probing Cosmic Inflation with the LiteBIRD Cosmic Microwave Background Polarization Survey*, arXiv:2202.02773 [DOI:10.1093/ptep/ptac150] [INSPIRE].
- [9] BICEP2 collaboration, *Detection of B-Mode Polarization at Degree Angular Scales by BICEP2*, *Phys. Rev. Lett.* **112** (2014) 241101 [arXiv:1403.3985] [INSPIRE].
- [10] PLANCK collaboration, *Planck 2018 results. IV. Diffuse component separation*, *Astron. Astrophys.* **641** (2020) A4 [arXiv:1807.06208] [INSPIRE].
- [11] N. Krachmalnicoff et al., *S-PASS view of polarized Galactic synchrotron at 2.3 GHz as a contaminant to CMB observations*, *Astron. Astrophys.* **618** (2018) A166 [arXiv:1802.01145] [INSPIRE].
- [12] J. Dunkley et al., *The Atacama Cosmology Telescope: likelihood for small-scale CMB data*, *JCAP* **07** (2013) 025 [arXiv:1301.0776] [INSPIRE].
- [13] BICEP2 and PLANCK collaborations, *Joint Analysis of BICEP2/KeckArray and Planck Data*, *Phys. Rev. Lett.* **114** (2015) 101301 [arXiv:1502.00612] [INSPIRE].
- [14] BICEP2 and KECK ARRAY collaborations, *BICEP2/Keck Array x: Constraints on Primordial Gravitational Waves using Planck, WMAP, and New BICEP2/Keck Observations through the 2015 Season*, *Phys. Rev. Lett.* **121** (2018) 221301 [arXiv:1810.05216] [INSPIRE].
- [15] BICEP and KECK collaborations, *Improved Constraints on Primordial Gravitational Waves using Planck, WMAP, and BICEP/Keck Observations through the 2018 Observing Season*, *Phys. Rev. Lett.* **127** (2021) 151301 [arXiv:2110.00483] [INSPIRE].
- [16] ACT collaboration, *The Atacama Cosmology Telescope: a measurement of the Cosmic Microwave Background power spectra at 98 and 150 GHz*, *JCAP* **12** (2020) 045 [arXiv:2007.07289] [INSPIRE].
- [17] F. Vansyngel et al., *Statistical simulations of the dust foreground to cosmic microwave background polarization*, *Astron. Astrophys.* **603** (2017) A62 [Erratum *ibid.* **618** (2018) C4] [arXiv:1611.02577] [INSPIRE].
- [18] S. Azzoni et al., *A minimal power-spectrum-based moment expansion for CMB B-mode searches*, *JCAP* **05** (2021) 047 [arXiv:2011.11575] [INSPIRE].
- [19] M. Remazeilles, A. Rotti and J. Chluba, *Peeling off foregrounds with the constrained moment ILC method to unveil primordial CMB B-modes*, *Mon. Not. Roy. Astron. Soc.* **503** (2021) 2478 [arXiv:2006.08628] [INSPIRE].
- [20] LITEBIRD collaboration, *Moment expansion of polarized dust SED: A new path towards capturing the CMB B-modes with LiteBIRD*, *Astron. Astrophys.* **660** (2022) A111 [arXiv:2111.07742] [INSPIRE].
- [21] H.K. Eriksen et al., *Joint Bayesian component separation and CMB power spectrum estimation*, *Astrophys. J.* **676** (2008) 10 [arXiv:0709.1058] [INSPIRE].
- [22] J. Dunkley et al., *Five-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Bayesian Estimation of CMB Polarization Maps*, *Astrophys. J.* **701** (2009) 1804 [arXiv:0811.4280] [INSPIRE].
- [23] R. Stompor, J. Errard and D. Poletti, *Forecasting performance of CMB experiments in the presence of complex foreground contaminations*, *Phys. Rev. D* **94** (2016) 083526 [arXiv:1609.03807] [INSPIRE].

- [24] B.S. Hensley and P. Bull, *Mitigating Complex Dust Foregrounds in Future Cosmic Microwave Background Polarization Experiments*, *Astrophys. J.* **853** (2018) 127 [[arXiv:1709.07897](#)] [[INSPIRE](#)].
- [25] A. Mangilli et al., *Dust moments: towards a new modelling of the galactic dust emission for CMB B-modes analysis*, *Astron. Astrophys.* **647** (2021) A52 [[arXiv:1912.09567](#)] [[INSPIRE](#)].
- [26] L. McBride, P. Bull and B.S. Hensley, *Characterizing line-of-sight variability of polarized dust emission with future CMB experiments*, *Mon. Not. Roy. Astron. Soc.* **519** (2023) 4370 [[arXiv:2207.14213](#)] [[INSPIRE](#)].
- [27] L. Vacher et al., *High precision modeling of polarized signals: Moment expansion method generalized to spin-2 fields*, *Astron. Astrophys.* **669** (2023) A5 [[arXiv:2205.01049](#)] [[INSPIRE](#)].
- [28] O.H.E. Philcox, B.D. Sherwin and A. van Engelen, *Detection and Removal of B-mode Dust Foregrounds with Signatures of Statistical Anisotropy*, *Mon. Not. Roy. Astron. Soc.* **479** (2018) 5577 [[arXiv:1805.09177](#)] [[INSPIRE](#)].
- [29] N. Jeffrey et al., *Single frequency CMB B-mode inference with realistic foregrounds from a single training image*, *Mon. Not. Roy. Astron. Soc.* **510** (2022) L1 [[arXiv:2111.01138](#)] [[INSPIRE](#)].
- [30] D. Alonso, J. Dunkley, B. Thorne and S. Naess, *Simulated forecasts for primordial B-mode searches in ground-based experiments*, *Phys. Rev. D* **95** (2017) 043504 [[arXiv:1608.00551](#)] [[INSPIRE](#)].
- [31] J. Errard and R. Stompor, *Characterizing bias on large scale CMB B-modes after galactic foregrounds cleaning*, *Phys. Rev. D* **99** (2019) 043529 [[arXiv:1811.00479](#)] [[INSPIRE](#)].
- [32] R.D.P. Grumitt, L.R.P. Jew and C. Dickinson, *Hierarchical Bayesian CMB Component Separation with the No-U-Turn Sampler*, *Mon. Not. Roy. Astron. Soc.* **496** (2020) 4383 [[arXiv:1910.14170](#)] [[INSPIRE](#)].
- [33] G. Puglisi et al., *Improved galactic foreground removal for B-mode detection with clustering methods*, *Mon. Not. Roy. Astron. Soc.* **511** (2022) 2052 [[arXiv:2109.11562](#)] [[INSPIRE](#)].
- [34] R. Stompor, S.M. Leach, F. Stivoli and C. Baccigalupi, *Maximum Likelihood algorithm for parametric component separation in CMB experiments*, *Mon. Not. Roy. Astron. Soc.* **392** (2009) 216 [[arXiv:0804.2645](#)] [[INSPIRE](#)].
- [35] F. Stivoli et al., *Maximum likelihood, parametric component separation and CMB B-mode detection in suborbital experiments*, *Mon. Not. Roy. Astron. Soc.* **408** (2010) 2319 [[arXiv:1004.4756](#)] [[INSPIRE](#)].
- [36] D.J. Fixsen, *The Temperature of the Cosmic Microwave Background*, *Astrophys. J.* **707** (2009) 916 [[arXiv:0911.1955](#)] [[INSPIRE](#)].
- [37] J. Errard, F. Stivoli and R. Stompor, *Publisher's Note: Framework for performance forecasting and optimization of CMB B-mode observations in the presence of astrophysical foregrounds* [*Phys. Rev. D* **84**, 063005 (2011)], *Phys. Rev. D* **84** (2011) 069907.
- [38] J. Errard, S.M. Feeney, H.V. Peiris and A.H. Jaffe, *Robust forecasts on fundamental physics from the foreground-obscured, gravitationally-lensed CMB polarization*, *JCAP* **03** (2016) 052 [[arXiv:1509.06770](#)] [[INSPIRE](#)].
- [39] R. Stompor et al., *Making maps of the cosmic microwave background: The MAXIMA example*, *Phys. Rev. D* **65** (2002) 022003 [[astro-ph/0106451](#)] [[INSPIRE](#)].
- [40] S. Hamimeche and A. Lewis, *Likelihood Analysis of CMB Temperature and Polarization Power Spectra*, *Phys. Rev. D* **77** (2008) 103013 [[arXiv:0801.0554](#)] [[INSPIRE](#)].
- [41] L. Knox, *Cosmic microwave background anisotropy observing strategy assessment*, *Astrophys. J.* **480** (1997) 72 [[astro-ph/9606066](#)] [[INSPIRE](#)].
- [42] K.M. Górski et al., *HEALPix — A Framework for high resolution discretization, and fast analysis of data distributed on the sphere*, *Astrophys. J.* **622** (2005) 759 [[astro-ph/0409513](#)] [[INSPIRE](#)].

- [43] LSST DARK ENERGY SCIENCE collaboration, *A unified pseudo- C_ℓ framework*, *Mon. Not. Roy. Astron. Soc.* **484** (2019) 4127 [[arXiv:1809.09603](#)] [[INSPIRE](#)].
- [44] PLANCK collaboration, *Planck 2015 results. IX. Diffuse component separation: CMB maps*, *Astron. Astrophys.* **594** (2016) A9 [[arXiv:1502.05956](#)] [[INSPIRE](#)].
- [45] B. Thorne, J. Dunkley, D. Alonso and S. Naess, *The Python Sky Model: software for simulating the Galactic microwave sky*, *Mon. Not. Roy. Astron. Soc.* **469** (2017) 2821 [[arXiv:1608.02841](#)] [[INSPIRE](#)].
- [46] PLANCK collaboration, *Planck 2015 results. X. Diffuse component separation: Foreground maps*, *Astron. Astrophys.* **594** (2016) A10 [[arXiv:1502.01588](#)] [[INSPIRE](#)].
- [47] C. Haslam et al., *A 408 mhz all-sky continuum survey. I. Observations at southern declinations and for the north polar region*, *Astron. Astrophys.* **100** (1981) 209.
- [48] C.G.T. Haslam, C.J. Salter, H. Stoffel and W.E. Wilson, *A 408 MHz all-sky continuum survey. II. The atlas of contour maps*, *Astron. Astrophys. Suppl. Ser.* **47** (1982) 1 [[INSPIRE](#)].
- [49] J. Dunkley et al., *CMBPol Mission Concept Study: Prospects for polarized foreground removal*, *AIP Conf. Proc.* **1141** (2009) 222 [[arXiv:0811.3915](#)] [[INSPIRE](#)].
- [50] WMAP collaboration, *Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results*, *Astrophys. J. Suppl.* **208** (2013) 20 [[arXiv:1212.5225](#)] [[INSPIRE](#)].
- [51] K. Abazajian et al., *CMB-S4 Science Case, Reference Design, and Project Plan*, [arXiv:1907.04473](#) [FERMILAB-PUB-19-431-AE-SCD] [[INSPIRE](#)].
- [52] CMB-S4 collaboration, *CMB-S4: Forecasting Constraints on Primordial Gravitational Waves*, *Astrophys. J.* **926** (2022) 54 [[arXiv:2008.12619](#)] [[INSPIRE](#)].
- [53] B.S. Hensley and B.T. Draine, *Thermodynamics and Charging of Interstellar Iron Nanoparticles*, *Astrophys. J.* **834** (2017) 134.
- [54] PLANCK collaboration, *Planck intermediate results. LVII. Joint Planck LFI and HFI data processing*, *Astron. Astrophys.* **643** (2020) A42 [[arXiv:2007.04997](#)] [[INSPIRE](#)].
- [55] T. Matsumura et al., *Mission design of LiteBIRD*, *J. Low Temp. Phys.* **176** (2014) 733 [[arXiv:1311.2847](#)] [[INSPIRE](#)].
- [56] PLANCK collaboration, *Planck 2018 results. XI. Polarized dust foregrounds*, *Astron. Astrophys.* **641** (2020) A11 [[arXiv:1801.04945](#)] [[INSPIRE](#)].
- [57] J. Martin, C. Ringeval, R. Trotta and V. Vennin, *The Best Inflationary Models After Planck*, *JCAP* **03** (2014) 039 [[arXiv:1312.3529](#)] [[INSPIRE](#)].
- [58] A. Lewis, *GetDist: a Python package for analysing Monte Carlo samples*, [arXiv:1910.13970](#) [[INSPIRE](#)].
- [59] T. Oliphant, *Guide to NumPy*, <https://web.mit.edu/dvp/Public/numpybook.pdf> (2006).
- [60] S. van der Walt, S.C. Colbert and G. Varoquaux, *The NumPy Array: A Structure for Efficient Numerical Computation*, *Comput. Sci. Eng.* **13** (2011) 22 [[arXiv:1102.1523](#)] [[INSPIRE](#)].
- [61] P. Virtanen et al., *SciPy 1.0 — Fundamental Algorithms for Scientific Computing in Python*, *Nature Meth.* **17** (2020) 261 [[arXiv:1907.10121](#)] [[INSPIRE](#)].
- [62] A. Zonca et al., *healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python*, *J. Open Source Softw.* **4** (2019) 1298 [[INSPIRE](#)].
- [63] J.D. Hunter, *Matplotlib: A 2D Graphics Environment*, *Comput. Sci. Eng.* **9** (2007) 90 [[INSPIRE](#)].