

Semi-blind Bayesian inference of CMB map and power spectrum

Flavien Vansyngel^{1,2}, Benjamin D. Wandelt^{1,2,3,4}, Jean-François Cardoso^{1,5,6}, and Karim Benabed^{1,2}

¹ Sorbonne Universités, UPMC Univ Paris 06, UMR 7095, Institut d'Astrophysique de Paris, 75014 Paris, France
e-mail: vansyng@iap.fr

² CNRS, UMR 7095, Institut d'Astrophysique de Paris, 75014 Paris, France

³ Department of Physics, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

⁴ Department of Astronomy, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

⁵ Laboratoire Traitement et Communication de l'Information, CNRS, UMR 5141 and Télécom ParisTech, 46 rue Barrault, 75634 Paris Cedex 13, France

⁶ APC, AstroParticule et Cosmologie, Université Paris Diderot, 75013 Diderot, France

Received 31 August 2014 / Accepted 24 January 2016

ABSTRACT

We present a new blind formulation of the cosmic microwave background (CMB) inference problem. The approach relies on a phenomenological model of the multifrequency microwave sky without the need for physical models of the individual components. For all-sky and high resolution data, it unifies parts of the analysis that had previously been treated separately such as component separation and power spectrum inference. We describe an efficient sampling scheme that fully explores the component separation uncertainties on the inferred CMB products such as maps and/or power spectra. External information about individual components can be incorporated as a prior giving a flexible way to progressively and continuously introduce physical component separation from a maximally blind approach. We connect our Bayesian formalism to existing approaches such as Commander, spectral mismatch independent component analysis (SMICA), and internal linear combination (ILC), and discuss possible future extensions.

Key words. cosmic background radiation – methods: data analysis – methods: statistical

1. Introduction

Observations of the cosmic microwave background (CMB) constrain cosmological models. In particular, the CMB fluctuations are very sensitive to the parameters of the current standard model of cosmology (Jungman et al. 1996). Current and future experiments designed for CMB analysis are signal dominated (Planck Collaboration III 2014; Schaffer et al. 2011; The CORe Collaboration et al. 2011; Baumann et al. 2009; André et al. 2014). Therefore, the remaining issue in deriving cosmological information from CMB is the separation between the CMB signal and the foregrounds signals. The ability to propagate component separation uncertainties to final constraints on fundamental physics is a leading issue in CMB analysis.

Most CMB experiments, such as the ongoing *Planck* mission (Planck Collaboration I 2014), observe in the microwave domain. The CMB is not the only emission that is received when observing from the solar system at these frequencies. Free-free, synchrotron, and thermal dust emissions emanating from our galaxy are among the most intense signals in the microwave domain (Sehgal et al. 2010). An observation of the sky at these frequencies is therefore a mixture of the photons from the different sources and the CMB must be extracted through component separation techniques.

Apart from the assumption that the CMB emission law follows a black body, the method presented in the present paper makes no assumption about foreground emission. We make use of independent component analysis (ICA) after assuming the mutual independence of the different signals constituting the data. Blind separation of independent sources (e.g., Cardoso 1998) is a very general process that finds applications in various

fields, from telecommunication to biomedical signals. Blind ICA has previously been applied in cosmology, particularly in CMB analysis. Fast ICA (FastICA; Baccigalupi et al. 2000; Maino et al. 2002) and spectral mismatch ICA (SMICA; Cardoso et al. 2002, 2008; Delabrouille et al. 2003) are two examples of that class of methods. Other methods, like generalized morphological component analysis GMCA (Starck et al. 2004, 2013), exploit sparsity rather than independence to discriminate between different signals. In this paper we adopt the first approach and we propose a Bayesian instance of semi-blind ICA.

The different component separation methods are mainly characterized by two important aspects, the basis on which the data are expressed and the parametrization of the data. Current methods perform the separation on different bases such as pixel space (Eriksen et al. 2006), spherical harmonic space (Tegmark 1998; Delabrouille et al. 2003), or needlet space (Delabrouille et al. 2009; Moudden et al. 2005; Fernandez-Cobos et al. 2012). The description of the data involves either a non-parametric model and exploits the independence between the CMB and the non-CMB component only – e.g., needlet ILC (NILC; Delabrouille et al. 2009), spectral estimation via expectation maximization (SEVEM; Fernandez-Cobos et al. 2012) – or a parametric model that is fitted to the data – e.g., Commander (Eriksen et al. 2006). SMICA belongs to an intermediate category of methods that assume coherence through frequency and complete independence of the components; it fits a non-parametric foreground model to the data via likelihood maximization.

One step forward is to choose a generic statistical model of the components based on generic assumptions (e.g., statistical independence of the components, spatial coherence between

frequencies, spatial or angular scale statistical independence), which then allows a full Bayesian exploration of the posterior density. The introduction of a simple but full generative model that approximates the stochastic model of the component permits propagating the uncertainties within that model. The simplifying assumptions allow either a numerical marginalization over all nuisance parameters (Gratton 2008) or, as in this paper, a full exploration of the model and a joint sampling of both the component maps and power spectra. The goal is to infer a CMB map and power spectrum, not to produce physical maps of the non-CMB components.

This paper is organized as follows. The method and the simulations to which it is applied are described in Sects. 2 and 3. Section 4 presents the results. The robustness of the method is analyzed in Sect. 5. The method is compared to previous component separation methods in Sect. 6. We conclude and comment on future directions for development of these ideas in Sects. 7 and 8.

2. Method

2.1. Data model

We model the data as signal plus noise and the observed signal is assumed to be a linear mixture of several diffuse emissions. Hence the following decomposition for the piece of data d_{ilm} contained in the spherical harmonic coefficient (l, m) of the observation map at frequency i (over N_f frequencies)

$$d_{ilm} = \sum_{k=1}^{N_c} A_{ik} s_{k\ell m} + n_{ilm}, \quad (1)$$

where the sum runs over the assumed number of components N_c , $s_k = \{s_{k\ell m}; \ell = \ell_{\min} \dots \ell_{\max}, m = -\ell \dots \ell\}$ is the spherical harmonic transform of the k th component map, A_{ik} is the amount of component k at frequency i , and n_{ilm} is the amount of instrumental noise present in d_{ilm} . We assume the data to be beam-corrected since deconvolution in harmonic space is performed by just dividing the data by the beam transfer functions. Equation (1) reads in matrix form:

$$d_{\ell m} = A s_{\ell m} + n_{\ell m}. \quad (2)$$

The $N_c \times N_f$ matrix A , which gathers all the coefficients A_{ik} , is called the *mixing matrix*. The noise is assumed to be Gaussian with zero mean and $\langle n_{ilm} n_{i' \ell' m'} \rangle = v_{i\ell}^2 \delta_{i i'} \delta_{\ell \ell'} \delta_{m m'}$. We use C to denote the power spectra of the components and $\langle s_{k\ell m} s_{k' \ell' m'} \rangle = C_{k\ell} \delta_{k k'} \delta_{\ell \ell'} \delta_{m m'}$. The covariance of the data predicted by the model of Eq. (2) is then

$$R_{\ell} = A C_{\ell} A^T + N_{\ell}, \quad (3)$$

where $(C_{\ell})_{k k'} = \delta_{k k'} C_{k\ell}$ and $(N_{\ell})_{i i'} = \delta_{i i'} v_{i\ell}^2$ is the noise covariance at multipole ℓ .

The contribution from point sources is neglected in this work, but is discussed in Sect. 7.

2.2. SMICA likelihood

SMICA (Cardoso et al. 2008) is a blind method working in harmonic space that provides an estimate of auto and cross power spectra and frequency spectra of the various components in the data. The parameters are determined by fitting the empirical spectral covariance of the data to the covariance of the model of Eq. (3). This procedure is equivalent to the maximization of

the SMICA likelihood $\mathcal{L}_{\text{SMICA}}$ that is obtained after assuming the independence and the Gaussianity of the components, and

$$-2 \log \mathcal{L}_{\text{SMICA}}(d|A, C) = \sum_{\ell, m} \log |2\pi R_{\ell}| + d_{\ell m}^T R_{\ell}^{-1} d_{\ell m}. \quad (4)$$

Links between SMICA and the present method are detailed in Sect. 6.

2.3. Bayesian formulation

The aim of the method is to provide a joint probability distribution over the mixing matrix A , the component maps s , and component power spectra $\{A, s, C\}$, knowing the data. If desired, prior knowledge on the parameters can be added. This formulation can be quantitatively written using Bayes' theorem:

$$P(A, s, C|d) \propto \mathcal{L}(d|A, s) P(A, s, C). \quad (5)$$

The likelihood $\mathcal{L}(d|A, s)$ encodes the data model chosen together with instrumental properties and the prior distribution $P(A, s, C)$ encodes the information already available about A , s , and C . The power spectra are hyperparameters of the model since they parametrize the prior distribution on the component s .

There are two motivations for considering the power spectra as additional stochastic parameters. First, it introduces more flexibility into the modeling of the components. Second and more importantly, the posterior provides an inference of the power spectra jointly with the mixing matrix and the component maps. Thus, errors introduced by the component separation step are propagated to the power spectra.

2.3.1. Likelihood and marginalization

If A and $s_{\ell m}$ in Eq. (2) are kept fixed then the stochasticity of $d_{\ell m}$ relies on the instrumental noise $n_{\ell m}$ alone. Noise properties of an instrument are usually well determined but complex. Considering the noise to be Gaussian, independent from frequency to frequency is a good approximation for the WMAP and *Planck* missions (Planck Collaboration III 2014; Jarosik et al. 2011). Thus the probability distribution function (PDF) for the noise is taken to be a multivariate normal distribution with zero mean and covariance N . The noise is supposed to be stationary, which implies that N is diagonal in harmonic space. Therefore, the likelihood is a product of Gaussians with mean $A s_{\ell m}$ and covariance N_{ℓ} , i.e., $\mathcal{L}(d|A, s) = \prod_{\ell m} \mathcal{N}(d_{\ell m}; A s_{\ell m}, N_{\ell})$.

If high resolution maps are required, then the posterior PDF is defined on a $\sim 10^7$ -dimension parameter space. The usual sampling schemes used to draw from intractable PDFs are then inefficient because of high correlation length in the chains. One solution is to split the sampling problem into several sampling problems:

$$P(A, s, C|d) = P(s|A, C, d) P(A, C|d). \quad (6)$$

Sampling $\{A, C\}$ from the marginal $P(A, C|d)$ and then post-processing these samples to sample the component maps s is equivalent to sampling $\{A, s, C\}$ from $P(A, s, C|d)$. The marginal is defined on a parameter space with a considerably reduced number of dimensions, the number of parameters being on the order of 10^3 , and is therefore simpler to sample from. In addition, sampling the maps is much more time-consuming than sampling the mixing matrix and the covariance.

We use the definition of a marginal distribution and the Bayes theorem of Eq. (5) and obtain

$$P(A, C | d) \propto \int ds \mathcal{L}(d | A, s) P(A, s, C). \quad (7)$$

If A and $\{s, C\}$ are independent in the prior distribution, and if the prior on s is taken to be a Gaussian with zero mean and covariance C , then

$$P(A, C | d) \propto \mathcal{L}_{\text{SMICA}}(d | A, C) P(A) P(C). \quad (8)$$

This marginalization implicitly assumes that C and N are diagonal on the same basis. The components are described by their power spectrum only. In the case of a very large data set like the *Planck* data, the Gaussian prior does not enforce isotropy of the components in the posterior; however, in the prior the components are isotropic in the absence of other information. Exploring the marginal posterior amounts to exploring the SMICA likelihood weighted by prior distributions on the mixing matrix and the power spectra.

2.3.2. Choice of priors

As stated above, the mixing matrix A and the components s are independent in the prior PDF, i.e., $P(A, s, C) = P(A)P(s|C)P(C)$. In this section we describe and justify the choice of the prior distributions of our analysis. Prior information must be chosen with care as it can introduce biases. We opted for non-informative or mildly informative priors in order to keep the analysis as blind as possible, although the Bayesian formalism does allow inclusion of further information.

Component maps As a prior on component maps, we put a zero mean Gaussian PDF uncorrelated from multipole to multipole and from component to component that is parametrized by its diagonal covariance C .

The simplest inflationary theories predict that CMB fluctuations are very nearly a Gaussian random field (e.g., Mukhanov 2013). Therefore, choosing a Gaussian random field as a prior for the CMB is physically well motivated. The choice of a Gaussian prior is for computational reasons. We will assess the impact of this approximation in Sect. 5.

Mixing matrix The spectral behavior of the CMB is supposed to be perfectly known and to follow a black-body law, which means that the CMB signal will have constant response through frequencies in the data. Thus, when expressed in thermodynamical units, the elements of the column related to the CMB are fixed to the same unit constant. In the other columns, the elements are normalized with regard to a reference frequency. The normalization is necessary to break a continuous degeneracy between each column of the mixing matrix and the corresponding component spectrum. The mixing coefficients, i.e., how much a component is present in the data, and the standard deviation of this component can be chosen up to an arbitrary factor. In order to fix this degree of freedom, one element in each column of the mixing matrix is fixed to unity. A flat prior is applied on the remaining elements, such that $P(A) \propto 1$. Negative values for the entries of the mixing matrix are allowed.

Component power spectra All components are uncorrelated a priori. For all multipole ℓ , C_ℓ is diagonal and thus contains

the power spectra of the components on the diagonal. The non-informative Jeffreys prior and a positivity constraint are applied on the power spectra, i.e., $P(C) \propto \prod_{k,\ell} \mathbf{H}(C_{k\ell}) C_{k\ell}^{-1}$, where \mathbf{H} is the Heaviside step function.

2.4. Sampling techniques

The marginal posterior in Eq. (8) is not a common PDF in $\{A, C\}$. It cannot be sampled from directly and we need to explore the parameter space with a sampling algorithm. We adopt the Metropolis-Hastings formalism to draw samples of $\{A, C\}$ and then estimate the marginal posterior. Then the full posterior over $\{A, s, C\}$ is recovered sampling conditionally on $\{A, C\}$ as in Eq. (6).

In practice, we let the Metropolis-Hastings sampler evolve until it converges and draws a sufficient number (e.g., N_{sam}) of uncorrelated samples of $\{A, C\}$. The chain $\{\{A_n, C_n\}; n = 1 \dots N_{\text{sam}}\}$ provides an estimation of the marginal posterior. Then, for each sample $\{A_n, C_n\}$, we draw a sample of the component maps s_n from the conditional $P(s | A=A_n, C=C_n, d)$. The chain $\{\{A_n, s_n, C_n\}; n = 1 \dots N_{\text{sam}}\}$ provides an estimate of the full posterior $P(A, s, C | d)$.

The sampling of the s_n is straightforward since the conditional $P(s | A, C, d)$ is a Gaussian distribution, independent from one harmonic coefficient to another. For each piece of the map $s_{\ell m}$ the mean and covariance are

$$\mu_{\ell m} = \Sigma_\ell A^T N_\ell^{-1} d_{\ell m} \quad (9)$$

$$\Sigma_\ell = \left(A^T N_\ell^{-1} A + C_\ell^{-1} \right)^{-1}. \quad (10)$$

Thus, the values of $\mu_{\ell m}$ are obtained by Wiener-filtering the data. There is a loss of power at high multipole in the mean owing to the filter. Sampling the maps corrects this and the samples of the maps have the correct covariance. Following Wandelt et al. (2004), the sampling of the maps is done by solving the system

$$\Sigma_\ell^{-1} s_{\ell m} = A^T N_\ell^{-1} d_{\ell m} + \xi_{\ell m},$$

where $\xi_{\ell m}$ are independent and identically distributed for each $s_{\ell m}$ from a Gaussian distribution with zero mean and covariance Σ_ℓ^{-1} .

2.5. Summary of the method

The context of the method is set by several hypotheses. The cosmic microwave background and the foregrounds are assumed to be described as a fixed number of uncorrelated components, all coherent through frequencies. The mixing matrix is assumed to be spatially constant. The noise is supposed to be Gaussian, uncorrelated from pixel to pixel and from frequency to frequency, and with isotropic statistics. In the prior model, the components are a set of uncorrelated random Gaussian fields and the prior for the mixing matrix is a flat distribution. Also, masking and beaming are not addressed in this work. Several extensions of the method will be studied in a future work. The rest of this subsection gives a summary of our Bayesian method from these assumptions.

Given the model of Eq. (2) and since we have chosen a flat prior for the mixing matrix A and a Gaussian prior for the component maps s and the Jeffreys prior for the power spectra C , the

full posterior we want to sample from is

$$P(A, s, C | d) \propto \prod_{\ell, m} |N_\ell|^{-\frac{1}{2}} \exp -\frac{1}{2} (d_{\ell m} - A s_{\ell m})^T N_\ell^{-1} (d_{\ell m} - A s_{\ell m}) \prod_{\ell, m} |C_\ell|^{-\frac{1}{2}} \exp -\frac{1}{2} s_{\ell m}^T C_\ell^{-1} s_{\ell m} \prod_{k, \ell} C_{k\ell}^{-1}. \quad (11)$$

In order to draw a sample from the full posterior, we split the problem (see Eq. (6)), where in our case

$$P(A, C | d) \propto \prod_{\ell} |R_\ell|^{-\frac{2\ell+1}{2}} \exp -\frac{1}{2} \text{tr} \left((2\ell+1) R_\ell^{-1} V_\ell \right) \prod_{k, \ell} C_{k\ell}^{-1} \quad (12)$$

$$P(s | A, C, d) \propto \prod_{\ell, m} |\Sigma_\ell|^{-\frac{1}{2}} \exp -\frac{1}{2} (s_{\ell m} - \mu_{\ell m})^T \Sigma_\ell^{-1} (s_{\ell m} - \mu_{\ell m}) \quad (13)$$

with R_ℓ , $\mu_{\ell m}$, and Σ_ℓ respectively defined in Eqs. (3), (9), and (10), and $V_\ell = \frac{1}{2\ell+1} \sum_m d_{\ell m} d_{\ell m}^T$.

The method consists in sampling $\{A, C\}$ from Eq. (12) using Metropolis-Hastings sampling and then choosing a subset of uncorrelated samples of $\{A, C\}$ to conditionally draw s from the Gaussian in Eq. (13). The factorization in Eq. (6) assures that this process amounts to sampling from the full posterior of Eq. (11).

3. Simulations

For the purpose of this paper, we consider $N_c = 3$ components to be separated, observed at $N_f = 4$ *Planck*'s high frequency instrument (HFI) frequency channels: 100 GHz, 143 GHz, 217 GHz, and 353 GHz. The data maps are simulated according to the linear model of Eq. (2). The set of simulations is a noisy composite of CMB, thermal dust emission, and free-free emission.

Figure 1 shows the four simulated observation maps. The plot shows the CMB power spectrum in black and the power spectra of all the foregrounds at each frequency channel in color, from red to purple being from 100 GHz to 353 GHz.

3.1. The component maps and their power spectra

The data model of Eq. (2) implies that the components are coherent through frequencies. Therefore, one map only of each component is simulated.

The CMB map is simulated using the HEALPix software¹ (Górski et al. 2005) from a power spectrum computed by the CAMB software (Lewis et al. 2000) in a standard Λ CDM model. The spatial distributions of the foregrounds are simulated using the publicly available version of the *Planck* Sky Model (PSM; Delabrouille et al. 2013). The free-free map from the PSM has an electron temperature of 7000 K, a power law $I(\nu) \propto \nu^\beta$ with spectral index β close to -0.15 , and is a composite of maps from Dickinson et al. (2003) and from the WMAP MEM map (Bennett et al. 2013). The thermal dust map from the PSM is simulated from the Schlegel-Finkbeiner-Davies map from which the ultra-compact HII regions are subtracted. More details can be found in Delabrouille et al. (2013).

¹ <http://healpix.sourceforge.net>

3.2. Mixing matrix

The CMB column of the mixing matrix is fixed to 1, and the normalizing elements are also fixed to 1. The other elements agree with the PSM.

3.3. Noise maps

The noise is simulated at the map level, and is uncorrelated from pixel to pixel. The noise standard deviation maps are designed to be consistent with the scanning strategy chosen for the *Planck* spacecraft and is therefore anisotropic. In the harmonic domain the noise is characterized by one white power spectrum per frequency derived from the noise standard deviation maps. Our inference approximates the noise as isotropic. The impact of this approximation will be assessed in Sect. 5.

4. Joint inference of CMB map and power spectrum

Power spectrum is the main interest in CMB analysis. Considering the power spectra as hyperparameters of the model has two effects: (a) removing the bias that would have been present if the power spectra were fixed and (b) going beyond component separation only in order to jointly separate the CMB and to infer its power spectrum. Instead of being static, we let the power spectra be constrained by the data, free of any informative prior. De facto, the errors from component separation are encoded in the posterior distribution because the maps and their power spectra are jointly inferred with the mixing matrix. Thus, the PDF over the maps and the power spectra takes into account the systematic errors due to the presence of multiple components in the data and their inference.

4.1. Power spectrum inference

All elements in one column of the mixing matrix are fixed to the same arbitrary constant (we chose 1). This prior information leads the sampler to distribute the information contained in the data. Any emission that is constant through frequency is transferred into the power spectrum corresponding to the constant mixing matrix column, and any other emission is transferred into the other power spectra. Since the CMB is the only coherent signal with constant response across all frequencies, our analysis amounts to inferring a CMB power spectrum in the presence of foregrounds systematics.

Regarding the other components, the unmixing is not unique. Since there is no physical information on either the power spectra or the frequency spectra, the outcome of these parameters is a mixture of the input parameters that obtains the most probable configuration. A priori, we don't expect the individual input power spectra of the non-CMB components to be identifiable as dust and free-free because we force no correlation between the two spatial distributions². Likewise, the recovered mixing matrix is not physical and negative entries might occur in the reconstruction.

Figure 3 shows the inference of the power spectra, the CMB on top and the sum of the non-CMB components below. In order to visualize the inference, we plot the peak of the marginal PDF for each multipole. The peaks are represented by black dots in the upper panel of each plot of the figure. As expected for the CMB, the shape around the peaks at low multipole is

² Physically, dust and free-free are spatially correlated since they are both more prevalent in the galactic disk than at high latitudes.




Fig. 1. Simulated data maps at four of the *Planck* HFI frequencies from 100 GHz to 353 GHz using realistic spatial distributions of free-free and thermal dust emissions from the PSM. In these simulated data maps the templates of the component maps are scaled through frequency according to the mixing matrix. This set of channels was chosen because the CMB is the least contaminated by foregrounds and noise in this frequency range. The plot shows the power spectrum of the CMB (black line) and the level of foreground at each frequency channel (from red to purple is 100 GHz to 353 GHz).

not symmetric. An inverse-gamma distribution fits the individual marginal distributions well (see Fig. 2) as expected (see, e.g., Wandelt et al. 2004). The gray region represents the shape as if the distribution were a two-sided Gaussian distribution: the upper error is one upper standard deviation and the lower error is one lower standard deviation. The solid red line shows the input power spectrum. The lower panel of each plot in the figure shows the relative error to the input power spectrum. The input power spectrum of the CMB lies within the error bars and the recovery is accurate at better than the percent level. For the sum of the non-CMB components, small biases appear from $\ell = 1000$. The biases occur because the correlation between the components is not taken into account, as explained in Sect. 5. These biases are

small compared to the CMB power and therefore have no significant effect on the CMB inference.

4.2. Map inference

Marginalizing the posterior over all parameters but one pixel of one component map leads to a distribution that is consistent with a Gaussian distribution. We therefore consider the sample mean of the map samples, which is an estimate of the mean posterior CMB map, as a reference for a recovered CMB map. Figure 4 shows the input CMB map and the absolute value of the residual map. There is more residual error in the galactic plane because of foreground contamination. Pixels are correlated in the posterior,




Fig. 2. Posterior PDF marginalized over all parameters except the CMB power spectrum at multipole $\ell = 5$. The histogram is an estimation of the PDF and the solid red curve is the best fit of an inverse-gamma function to the histogram.

but qualitative errors on the recovered map are given by the sample variance of each pixel.

Figure 5 shows a map containing the sample standard deviation of each pixel of the CMB map. As stated above, the errors on the CMB map include the uncertainty due to the presence of galactic emission. We also plot in Fig. 5 the standardized error map, i.e., the ratio between the residual map and the standard deviation map. The isotropic noise approximation leads to an overestimation of the error bars in the regions of low noise. The per-pixel error is underestimated in highly contaminated regions. The residuals have strong spatial correlations, see Fig. 6. Another explanation could be that the Gaussian model is too coarse an approximation in regions where the foregrounds are the most intense and highly non-Gaussian. If this is the case, we could use these results to construct masks from the sample variance map to mask the observation maps where necessary, since these regions are correlated to the regions of high variance in the sample variance map (top of Fig. 5).

Pixels in the posterior are correlated and the sample variance map alone is not sufficient to fully describe the error on the reconstructed CMB map. To show the correlation, we compute the correlation matrix of the CMB map samples on a lower resolution map. The correlation matrix has a row for every pixel showing the correlation of this pixel to all other pixels. Figure 6 shows the correlation maps of two pixels in the galaxy plane. The pixels in the galactic plane are highly correlated, which explains at least part of the excess of standardized error in the galactic plane of Fig. 5. An eigenvalue analysis of the correlation matrix shows that, in addition to noise uncertainties on small scales, the foreground subtraction uncertainties are dominated by a few, highly correlated modes (see Fig. 7).

5. Model checking

In Fig. 3, the uncertainties on the reconstruction directly rely on the shape of the posterior. Therefore the errors are correctly estimated if the a priori model correctly describes the data. It is therefore important to assess the quality of the fit achieved by the model through model checking (Gelman & Meng 1996). In order to check for biases in the reconstruction due to assumptions on the statistical properties of the components, we measure the mismatch between the empirical covariance of the data and the




Fig. 3. Input and inferred power spectra of the CMB (*top*) and the sum of the non-CMB components (*bottom*). In the *upper panel* of each figure, the black dot at each multipole represents the peaks of the marginal posterior, the gray region shows the asymmetric $\pm 1\sigma$ -error bar derived from the marginal posterior, the red line is the input power spectrum. The *lower panel* represents the relative error to the input power spectrum. The sampler accurately recovers the power spectra of the CMB.

covariance of the data model of Eq. (3). Following the analysis of Delabrouille et al. (2003), for each ℓ we examine the quantity

$$D_\ell = (2\ell + 1) \left[\text{tr}(\hat{R}_\ell R_\ell^{-1}) - \log |\hat{R}_\ell R_\ell^{-1}| - N_f \right]$$

with \hat{R}_ℓ the estimation of the data covariance at multipole ℓ . The parameter D_ℓ is simply the log-likelihood, rewritten as twice the Kullback-Leibler divergence between the two PDF $p_\ell(d | \hat{R}_\ell)$ and $p_\ell(d | R_\ell)$ with

$$p_\ell(d | \Sigma) = \prod_{m=-\ell}^{\ell} |2\pi\Sigma|^{-1/2} \exp -\frac{1}{2} d_{\ell m}^T \Sigma^{-1} d_{\ell m}.$$

The posterior PDF on the CMB power spectrum approaches a Gaussian as ℓ increases, such that D_ℓ has the properties of the chi-squared distribution for sufficiently large ℓ .

The number of degrees of freedom of the distribution is obtained by subtracting the number of stochastic parameters per multipole from the number of degrees of freedom of a symmetric matrix $N_f \times N_f$. A number of spectra (N_c) are sampled per multipole and the mixing matrix is sampled once for all multipoles. Thus, if the number of degrees of freedom within a mixing




Fig. 4. Input and residual CMB map. This residual map represents the absolute value of the difference between the sample mean and the input map. The errors are wider in the galactic plane but the uncertainties in this region of the sky are also larger, as shown in Fig. 5. To show the noise in the residual map, it is shown on a decimal logarithm scale.

matrix is distributed over all multipoles, the number of degrees of freedom of the chi-squared distribution followed by each D_ℓ is

$$N_{\text{d.o.f.}} = N_f (N_f + 1) / 2 - \left[N_c + \frac{(N_f - 1)(N_c - 1)}{(\ell_{\text{max}} - \ell_{\text{min}} + 1)} \right] \\ \simeq N_f (N_f + 1) / 2 - N_c.$$

In particular $\mathbb{E}[D_\ell] = N_{\text{d.o.f.}}$, which does not depend on ℓ . In our test case $N_{\text{d.o.f.}} = 7$.

In Fig. 8 we plot the values of D_ℓ for R_ℓ containing the inferred value of $\{A, C\}$. For comparison, we also plot the values of D_ℓ in the case where A and C are set to their input values. Although the input parameters are the true parameters to be recovered, the inferred values have a lower mismatch because the components are correlated in the data but not in the model and the sampler finds uncorrelated components that fit the data better. If the foreground modeling matches the statistical properties of the input foregrounds, the values of D_ℓ should follow a chi-squared distribution with a number of degrees of freedom $N_{\text{d.o.f.}}$,

whose mean $N_{\text{d.o.f.}}$ is represented by the horizontal red line in Fig. 8.

We performed a separation where the cross power spectra of the component are taken into account during the sampling. We do not sample the cross power spectra, but we use the covariance of the input components instead, i.e., each C_ℓ is a non-diagonal matrix, but only the diagonal is stochastic. The prior on the power spectra is modified in order to enforce positive definiteness of the covariance, and

$$P(C) \propto \prod_{\ell} |C_{\ell}|^{-(N_c+1)/2} \prod_{\ell,k} \mathbf{H}(\lambda_{\ell k}),$$

where the $\lambda_{\ell k}$ are the eigenvalues of C .

In Fig. 9 we plot D_ℓ with the output values of $\{A, C\}$ of such a run. Taking the correlation of the input component maps into account erases the discrepancy at low ℓ . A chi-squared distribution fits the histogram of D_ℓ when ℓ is large enough ($\ell > 700$ for this plot). In the figure, the fit is shown by three blue solid lines, the solution of the fit and the $\pm 1\sigma$ error on the fit. The




Fig. 5. *Top:* standard deviation map of the CMB map samples. The posterior distribution is wider in the region of the galactic plane. *Bottom:* standardized error map, all red pixels have a value of 4 or more. This map is the ratio between the residual map of the CMB and its standard deviation map (*top*). The posterior standard deviation map only represents the part of the uncertainty that is uncorrelated from pixel to pixel, while the Bayesian analysis returns a fully correlated error model for the recovered map shown in Fig. 4. Standardizing with the uncorrelated errors reveals two things: the isotropic noise approximation leads to overestimated uncertainties in low noise regions, and an uncorrelated error model does not capture the uncertainties in regions where foregrounds dominate. See Fig. 6 for a visualization of correlated uncertainties.

red dashed line represents the chi-squared distribution with the expected number of degrees of freedom $N_{\text{d.o.f.}}$ and it lies within the error bars. The remaining deviations from an ℓ -independent distribution are due to the differences between the data model and the data actually used. The mismatch is completely flat and has the appropriate degree of freedom if a set of simulations completely consistent with the data model is used. Also, taking the correlations between components into account erases the biases in the inference of the power spectrum of the sum of the components.

Considering a Gaussian model for the components and neglecting the correlations between the components do not affect the reconstruction of the CMB since the CMB is not correlated to the foregrounds. The power of the non-CMB components is distributed into the effective components and the mismatch D_ℓ .

The power is distributed differently when the cross-spectra are included or not. On the other hand, neglecting the correlation between the non-CMB components is too coarse an approximation if the galactic components also need to be recovered.

6. Comparison to previous component separation methods

6.1. Comparison to Commander

As is our method, Commander (Eriksen et al. 2006, 2008) is a Bayesian formulation of the joint component separation and CMB power spectrum inference problem. The main difference between the two approaches is the parametrization of the problem.




Fig. 6. Two rows of the posterior correlation matrix for the two pixels marked by a black cross in each map at HEALPix resolution parameter $N_{\text{side}} = 16$. The inferred uncertainties due to foreground removal are highly correlated in the galactic plane and must be taken into account in a meaningful statistical interpretation of the recovered CMB map.

Commander makes intensive use of parametric models to describe the physical emissions while we adopt a quasi-blind and a phenomenological description of the different components. Thus Commander infers maps and CMB power spectrum within a constraining physical model, and therefore the most probable values of the parameters and their errors depend on this model. We do not rely on any physical assumption, except for the constant response of CMB signal across frequencies. The drawback is that only the recovered map and power spectrum of the CMB have a physical meaning, but the non-trivial difference is that our results do not depend on physical modeling assumptions.

In addition, Commander works at the map level, whereas our method works at the multipole level. Thus, our sampler is fast and allows treatment of high resolution data maps.

6.2. Comparison to SMICA

Spectral matching independent component analysis (SMICA; Cardoso et al. 2008) is a method that provides component power

spectra and mixing coefficients. The parameters are estimated by finding the best match between an A - and C -dependent covariance and an empirical data covariance. Depending on the binning of the power spectrum, and following the original formulation of SMICA, the method is equivalent to a maximization with respect to $\{A, C\}$ of the SMICA likelihood of Eq. (4). One of several applications of this method is to apply a Wiener filter to the data with the solution of the maximization in order to recover the component maps.

We can understand SMICA analysis from a Bayesian perspective as follows:

1. begin with the Bayesian formulation, Eq. (5);
2. choose Gaussian priors for the components, flat priors for the mixing matrix and the power spectra;
3. marginalize the posterior over all component maps;
4. maximize the obtained marginal distribution with respect to A and C .




Fig. 7. One hundred of the highest eigenvalues of the posterior correlation matrix of the low resolution CMB map. Two modes dominate.

In the Bayesian formulation, maximizing the marginal posterior $P(A, C | d)$ in Eq. (8) is equivalent to maximizing the likelihood in Eq. (4). Thus, instead of just finding the peak of the distribution as SMICA does, the Bayesian sampler explores the whole distribution over A and C , and therefore returns an error model for the cleaned CMB map that includes foreground cleaning uncertainties. In addition, the power spectrum inference is marginalized over foreground cleaning uncertainties.

6.3. Comparison to standard ILC

The internal linear combination (ILC) method (Bennett et al. 1992) provides a map of a component, given its frequency dependence. The method involves a weighted average of the observation maps in order to cancel all components except the one of interest. Usually, the ILC is derived by minimizing the variance of a linear combination of the observation maps. The remaining fluctuations are the CMB anisotropies.

The ILC can also be derived by adopting a Bayesian point of view:

1. begin with the Bayesian formulation, Eq. (5);
2. choose Gaussian priors for the component, flat priors for the mixing matrix and the power spectra;
3. marginalize the posterior over all component maps except the CMB map;
4. maximize the obtained marginal distribution with regard to the CMB map.

The solution is

$$\hat{s}_{\text{cmb},\ell m} = \frac{e^T (N_\ell + AC_\ell A^T)^{-1}}{C_{\text{cmb},\ell}^{-1} + e^T (N_\ell + AC_\ell A^T)^{-1} e} d_{\ell m}. \quad (14)$$

Here the $N_\ell \times (N_c - 1)$ mixing matrix A has no column dedicated to the CMB, C_ℓ is the covariance of all components except the CMB at multipole ℓ and $e = (1 \dots 1)^T$, i.e., the frequency response of the CMB. The standard ILC formula is

$$\hat{s}_{\text{cmb},\ell m} = \frac{e^T C_d^{-1}}{e^T C_d^{-1} e} d_{\ell m}, \quad (15)$$

where C_d is an estimate of the data covariance

$$C_{d_{\ell m}} = C_{\text{cmb},\ell} e e^T + A_{\text{true}} C_{\text{true},\ell m} A_{\text{true}}^T + N_{\ell m}. \quad (16)$$





Fig. 8. Mismatch between the data and the data modeling. *Top*: divergence between the data covariance and input parameters. The large mismatch at low ℓ is due to correlations between the input component maps. *Bottom*: divergence between the data covariance and the recovered parameters. During sampling, we impose no cross-correlation. Thus, the sampler converges towards components that are uncorrelated but whose power spectra are almost capable of capturing the covariance of the input component maps. The red line is the mean of the expected chi-squared distribution that D_ℓ should approximate.

If the CMB fluctuations are neglected in the ILC formula, and if the prior variance of the CMB is infinite (i.e., flat prior) in the Bayesian formula 14 then the two approaches are equivalent. The ILC method naturally chooses an estimation of the true mixing matrix and the true power spectra of the components, which is why, despite its very simple formulation and its approximations, the ILC method is physically consistent.

7. Discussion and future works

The fact that the sampler cannot distinguish between non-CMB emissions is due to additional degrees of freedom. Without any information about the physical emissions, all the recovered components other than CMB are a mixture of the true signals. Putting a prior either on the mixing matrix, i.e., importing knowledge about the frequency spectra, or on the shapes of the component power spectra would break the degeneracies. Although the constraints of the priors can be controlled, the blindness of the method would be lost with this introduction of a priori information. Furthermore, in this paper we impose decorrelation between the component maps. The foregrounds could be modeled




Fig. 9. Mismatch between the data and the result of a sampler that includes the input correlations between the components during the sampling process (*top*) and its PDF (*bottom*). In the upper panel the red line represents the mean of the expected chi-squared distribution that D_ℓ should follow. In the bottom panel, the blue lines are the best χ^2 -fit to a chi-squared curve and the ± 1 standard deviation error from the fit. The red dashed line represents the chi-squared distribution with the expected number of degrees of freedom $N_{\text{d.o.f.}} = 7$. The introduction of the correlation between the input component erases the large discrepancy at low multipole.

in more detail to get a full component separation method, but the focus here is on CMB reconstruction.

The next step is to apply the method to real data. The main problem is the instrumental noise. In this paper we assume a simple noise model. To deal with real data noise requires developing a more realistic noise model. A full model of correlated noise involves very large covariance matrices. Therefore, alternative ways to deal with noise like half-ring half difference maps or noise simulations should be considered.

Unresolved point sources appear as extra power at small angular scales of the inferred CMB power spectrum. Masking the listed point sources and inpainting in the mask would be a way to address the point source issue. In addition, allowing the components to mix differently in different regions of the sky or in different angular scale ranges by allocating different mixing matrices in each range would reduce mismatch due to lack of coherence. Also, foregrounds that are not fully coherent from frequency to frequency can be modeled by increasing the number of components in our model.

In this work we reconstructed the CMB maps on the full sky. It remains to be seen whether this is achievable on realistic data. Since our approach is similar to the one by SMICA, which

provides a clean map on a large fraction of the sky in *Planck* analysis (Planck Collaboration XII 2014), treatment of almost full sky data should be feasible. Since we work under the assumption of diagonal covariances in ℓ -space the effect of a small mask needs to be tested. It may be possible to avoid masking using inpainting. If necessary a Wiener filter method such as Elsner & Wandelt (2013) could be used to implement a full Gibbs sampling approach. If joint modeling of foregrounds allows working with a large part of the sky, we may be able to ignore mode coupling effects due to the mask with high accuracy.

Although in this paper we chose to address component separation with a blind analysis, we can use the flexibility of the Bayesian formalism in order to introduce physical parametrization of the problem. The current understanding of physical galactic and extra-galactic phenomena can be progressively introduced by a more detailed data model and by assigning a parametrized prior PDF on the foregrounds. The CMB power spectrum is also highly parametrizable since its shape depends on a small number of cosmological parameters (Planck Collaboration XVI 2014). Thus a joint inference of cleaned CMB map, CMB power spectrum, and cosmological parameters is conceivable thanks to fast and accurate Boltzmann code emulators like parameters for the impatient cosmologist (PICO; Fendt & Wandelt 2007). A less parametric approach would be to exploit the smoothness of the power spectra through binning or representation in terms of a smooth basis function, such as splines.

In addition, as in other component separation methods, the Bayesian formulation presented in this paper can be extended to infer the CMB polarization power spectrum.

We leave these further explorations to future work.

8. Conclusion

We have presented a new formulation for the CMB foreground cleaning. In our analysis, we avoid physical parametrizations, except that the CMB behaves like a black body and we model the components as Gaussian random fields. The CMB is then cleaned by jointly inferring CMB, galactic residuals, and frequency spectra. This Bayesian method provides an evaluation of a posterior PDF for the CMB power spectrum which thus takes into account uncertainties due to the removal of foreground contamination. Samples of maps of CMB anisotropies are drawn and reveal that the dominant foreground residuals are captured in terms of a small number of error modes. We also showed that previous component separation methods can be derived as special cases of our Bayesian formulation, which thus provides a unified approach for semi-blind foreground cleaning from multifrequency CMB data.

Acknowledgements. F.V. thanks Franz Elsner for useful discussions in the early stages of the work. This work made in the ILP LABEX (under reference ANR-10-LABX-63) was supported by French state funds managed by the ANR within the Investissements d’Avenir program under reference ANR-11-IDEX-0004-02. This work was also partially supported by NSF AST 07-08849. B.D.W. is supported by a senior Excellence Chair by the Agence Nationale de Recherche (ANR-10-CEXC-004-01) and a Chaire Internationale at the Université Pierre et Marie Curie

References

- André, P., Baccigalupi, C., Banday, A., et al. 2014, *J. Cosmol. Astropart. Phys.*, 1402, 006
- Baccigalupi, C., Bedini, L., Burigana, C., et al. 2000, *MNRAS*, 318, 769
- Baumann, D. et al. 2009, *AIP Conf. Proc.*, 1141, 10

- Bennett, C., Smoot, G., Hinshaw, G., et al. 1992, *ApJ*, 396, L7
- Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, *ApJS*, 208, 20
- Cardoso, J.-F. 1998, *Proceedings of the IEEE*, 86, 2009
- Cardoso, J.-F., Snoussi, H., Delabrouille, J., & Patanchon, G. 2002, ArXiv eprints [[arXiv:astro-ph/0209466](https://arxiv.org/abs/astro-ph/0209466)]
- Cardoso, J.-F., Martin, M., Delabrouille, J., Betoule, M., & Patanchon, G. 2008, ArXiv eprints [[arXiv:0803.1814](https://arxiv.org/abs/0803.1814)]
- Delabrouille, J., Cardoso, J.-F., & Patanchon, G. 2003, *MNRAS*, 346, 1089
- Delabrouille, J., Cardoso, J.-F., Jeune, M. L., et al. 2009, *A&A*, 493, 835
- Delabrouille, J., Betoule, M., Melin, J.-B., et al. 2013, *A&A*, 553, A96
- Dickinson, C., Davies, R. D., & Davis, R. J. 2003, *MNRAS*, 341, 369
- Elsner, F., & Wandelt, B. D. 2013, *A&A*, 549, A111
- Eriksen, H. K., Dickinson, C., Lawrence, C., et al. 2006, *ApJ*, 641, 665
- Eriksen, H., Jewell, J., Dickinson, C., et al. 2008, *ApJ*, 676, 10
- Fendt, W. A., & Wandelt, B. D. 2007, *ApJ*, 654, 2
- Fernandez-Cobos, R., Vielva, P., Barreiro, R., & Martinez-Gonzalez, E. 2012, *MNRAS*, 3, 2162
- Gelman, A., & Meng, X.-L. 1996, in *Model checking and model improvement* (Springer), 189
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759
- Gratton, S. 2008, ArXiv e-prints [[arXiv:0805.0093](https://arxiv.org/abs/0805.0093)]
- Jarosik, N., Bennett, C. L., Dunkley, J., et al. 2011, *ApJS*, 192, 14
- Jungman, G., Kamionkowski, M., Kosowsky, A., & Spergel, D. N. 1996, *Phys. Rev.*, D54, 1332
- Lewis, A., Challinor, A., & Lasenby, A. 2000, *ApJ*, 538, 473
- Maino, D., Farusi, A., Baccigalupi, C., et al. 2002, *MNRAS*, 334, 53
- Moudden, Y., Cardoso, J.-F., Starck, J.-L., & Delabrouille, J. 2005, *EURASIP J. Appl. Signal Process.*, 15, 2437
- Mukhanov, V. 2013, *EPJC*, 73, 2486
- Planck Collaboration I. 2014, *A&A*, 571, A1
- Planck Collaboration III. 2014, *A&A*, 571, A3
- Planck Collaboration XII. 2014, *A&A*, 571, A12
- Planck Collaboration XVI. 2014, *A&A*, 571, A16
- Schaffer, K., Crawford, T., Aird, K., et al. 2011, *ApJ*, 743, 90
- Sehgal, N., Bode, P., Das, S., et al. 2010, *ApJ*, 709, 920
- Starck, J.-L., Elad, M., & Donoho, D. L. 2004, *Adv. Imag. Elect. Phys.*, 132, 287
- Starck, J.-L., Donoho, D. L., Fadili, M. J., & Rassat, A. 2013, *A&A*, 552, A133
- Tegmark, M. 1998, *ApJ*, 502, 1
- The COre Collaboration, Armitage-Caplan, C., Avillez, M., et al. 2011, ArXiv eprints [[arXiv:1102.2181](https://arxiv.org/abs/1102.2181)]
- Wandelt, B. D., Larson, D. L., & Lakshminarayanan, A. 2004, *Phys. Rev. D*, 70, 083511