

A Novel CMB Component Separation Method: Hierarchical Generalized Morphological Component Analysis

Sebastian Wagner-Carena^{1,2★}, Max Hopkins^{3,4†}, Ana Diaz Rivero¹, Cora Dvorkin^{1‡}

¹ *Department of Physics, Harvard University, Cambridge, MA 02138, USA*

² *Department of Physics, Stanford University, Stanford, CA 94305, USA*

³ *Department of Mathematics, Harvard University, Cambridge, MA 02138, USA*

⁴ *Department of Computer Science and Engineering, University of California San Diego, San Diego, CA 92092, USA*

28 April 2020

ABSTRACT

We present a novel technique for Cosmic Microwave Background (CMB) foreground subtraction based on the framework of blind source separation. Inspired by previous work incorporating local variation to Generalized Morphological Component Analysis (GMCA), we introduce Hierarchical GMCA (HGMCA), a Bayesian hierarchical graphical model for source separation. We test our method on $N_{\text{side}} = 256$ simulated sky maps that include dust, synchrotron, free-free and anomalous microwave emission, and show that HGMCA reduces foreground contamination by 25% over GMCA in both the regions included and excluded by the *Planck* UT78 mask, decreases the error in the measurement of the CMB temperature power spectrum to the 0.02–0.03% level at $\ell > 200$ (and $< 0.26\%$ for all ℓ), and reduces correlation to all the foregrounds. We find equivalent or improved performance when compared to state-of-the-art Internal Linear Combination (ILC)-type algorithms on these simulations, suggesting that HGMCA may be a competitive alternative to foreground separation techniques previously applied to observed CMB data. Additionally, we show that our performance does not suffer when we perturb model parameters or alter the CMB realization, which suggests that our algorithm generalizes well beyond our simplified simulations. Our results open a new avenue for constructing CMB maps through Bayesian hierarchical analysis.

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

1 INTRODUCTION

The Cosmic Microwave Background (CMB) is the earliest observable light in the universe and one of the most important observables for constraining models of early-universe physics. The Lambda Cold Dark Matter (Λ CDM) model, inflation, and large-scale structure can all be better understood by more precise measurements of the CMB, as can be seen from the range of scientific results obtained with data from the recent *Planck* satellite (Planck Collaboration I 2018). A wide range of applications require an accurate reconstruction of the CMB sky maps: extraction of the gravitational lensing and integrated Sachs-Wolfe (ISW)

signals, constraints on isotropy, global geometry, topological defects, and searches for primordial non-Gaussianity (Planck Collaboration XIX 2014; Planck Collaboration VIII 2018; Planck Collaboration VII 2019; Planck Collaboration IX 2019; McEwen et al. 2017; Saadeh et al. 2016). Unfortunately, several sources, some within our own galaxy, others distant celestial objects, emit radiation at the same frequencies as the CMB. Separating these foreground sources from the CMB has become a critical problem for modern cosmology.

Traditionally, cosmologists have used template fitting techniques and second order statistical methods to isolate the CMB (Górski et al. 1996; Hobson et al. 1998; Bedini et al. 2005; Martínez-González et al. 2003; Vio & Andreani 2008; Delabrouille & Guillaux 2009; Rogers et al. 2016a; Leach et al. 2008; Bennett et al. 2013; Tegmark et al. 2003).

★ E-mail: swagnerc@stanford.edu

† E-mail: nmhopkin@eng.ucsd.edu

‡ E-mail: cdvorkin@g.harvard.edu

However, recent work has suggested that approaching the problem from a Bayesian perspective could yield CMB maps that are more accurate (Pham & Garat 1997; Delabrouille et al. 2003; Moudden et al. 2005; Bobin et al. 2008; Rapin et al. 2013; Bobin et al. 2013). Bayesian techniques relax assumptions about second or higher-order statistics, and can be adapted to leverage alternative properties of the CMB and foreground components. In this paper we focus on methods that exploit sparsity - the assumption that the CMB and foreground components have small support in a specific basis. Our particular basis of interest is constructed from scale-discretized wavelets - oscillating functions defined both in real and harmonic space that can be used to decompose signals on the sphere. Extensive work has been done on constructing and computing accurate wavelets on the sphere (Antoine & Vanderghyest 1999; Narcowich et al. 2006; Wiaux et al. 2008), and there are numerous examples of the use of wavelets in the context of CMB data (Moudden et al. 2005; Delabrouille & Guilloix (2009); Planck Collaboration XII (2014); Planck Collaboration X (2016); McEwen et al. (2017); Rogers et al. (2016a); Planck Collaboration IV (2018)).

Adding a Bayesian prior on the sources that enforces sparsity is the cornerstone of the Generalized Morphological Component Analysis (GMCA) algorithm (Bobin et al. 2008), which effectively separates the CMB from foregrounds. An extension to GMCA, coined Local GMCA (LGMCA) (Bobin et al. 2013), attempts to better capture the local variance of the foreground sources, circumventing one of the main limitations of vanilla GMCA. By running multiple independent GMCA models on different partitions of a sky map, the LGMCA method allows the foreground sources to display differing frequency dependence by position. In Bobin et al. (2013), the authors showed that this method excels at producing full sky maps with minimal foreground, especially thermal Sunyaev-Zeldovich (SZ) contaminants, but struggles with cosmic infrared background contamination at high multipole moments. While LGMCA improves on GMCA by incorporating local variation, it no longer falls into the Bayesian framework: there is no single probabilistic model that describes the data. Instead, LGMCA is the combination of many individual models that generate the data in different patches/at different scales and are therefore not necessarily consistent with one another¹. To reconstruct a final map, LGMCA is forced to use heuristic measurements to select from these reconstructions (for examples of possible heuristics see Bobin et al. 2013).

In this work we propose a new algorithm, Hierarchical Generalized Morphological Component Analysis (HGMCA), which captures the key improvement of LGMCA (incorporating local variation) in a single, formal generative probabilistic model. HGMCA runs GMCA across a number of partitions of the map while enforcing global consistency through the use of hierarchical priors. This enhances the reliability of the local fits and the ability to consider higher degrees of local variation compared to LGMCA.

To highlight HGMCA's performance we compare it with GMCA and a wavelet-based internal linear combina-

tion (ILC) implementation on simulated CMB and foregrounds. We leave LGMCA out of the comparison because the code is not publicly available, and our implementation of LGMCA was not able to reproduce the improvements shown in Bobin et al. (2013). The simulations are generated using the Python Sky Model (PySM) code (Thorne et al. 2017) at a resolution of $N_{\text{side}} = 256$ and include synchrotron, thermal dust, anomalous microwave emission (AME), and free-free emission. These foreground emissions are modeled from *Planck* templates and make conservative assumptions about their spectral dependence. While these simulations are of moderate resolution and realism, lacking differing beam resolutions and contaminants such as SZ clusters, they are sufficient to demonstrate the improvements resulting from the new algorithm: by using hierarchical priors, our model fits to local variations in the foreground frequency scaling without sacrificing the constraining power offered by the full map. This reduces the presence of foreground contamination in the reconstructed CMB signal as compared to GMCA and ILC-based methods (Section 5).

In addition to showing HGMCA's reduction in foreground emissions in the reconstructed CMB, we also provide rigorous testing of its generalizability and robustness (Section 5.3). We provide two types of testing for generalization: differing realizations of the CMB, and differing foreground templates. Since the hyperparameters of HGMCA are trained using PySM simulations, to ensure that we accurately represent the performance of HGMCA, we present all results across simulations with a different CMB realization than used in training. This includes our presented comparisons to GMCA and wavelet ILC (WavILC), an ILC implementation based on the work of Planck Collaboration IV (2018); Basak & Delabrouille (2013); Rogers et al. (2016a). To ensure that we have not overfit hyperparameters to specific foreground templates, we also show HGMCA's performance over simulations which use different foreground templates with smaller-scale local variation. Our results suggest that when the map is subdivided to a sufficient extent, HGMCA is robust both to changes in CMB and foreground realizations. Finally, we show HGMCA's robustness to changes in the hierarchical prior and sparsity hyperparameters by demonstrating that altering these parameters by up to 20% in either direction has a negligible effect on performance.

In the following sections, we briefly review the existing foreground separation techniques (Section 2), describe our model from a probabilistic perspective (Section 3), explain the implementation of our optimization algorithm (Section 4), and compare the results of our method to those of GMCA and a state of the art ILC CMB reconstruction on a set of simulations (Section 5). Finally, we show that these improvements do not appear to be sensitive to perturbations in our model parameters, the random seed, or the specifics of our simulation (Section 5.3). We conclude in Section 6, and mention possible extensions and applications of our HGMCA method to future CMB experiments.

2 BACKGROUND

Our work follows a long line of research in source separation for the CMB. All of the techniques that we discuss (and

¹ A detailed discussion of this property (and LGMCA) can be found in Section 2.3.

which are used by the *Planck* collaboration, [Planck Collaboration XVI 2014](#); [Planck Collaboration XIII 2016](#); [Planck Collaboration VI 2018](#)) view the sky as a linear combination of the CMB, foreground components, and instrumental noise. In matrix form, this can be written as

$$X = AS + N, \quad (1)$$

where X is the observed data vector, A is referred to as the mixing matrix, S is the source vector, and N is the noise. Each row of X is a sky map at a different frequency, each row of S corresponds to a different source – for instance, independent components of signal in the sky, including both the CMB and sources of foreground – and the coefficients a_{ij} of the mixing matrix A dictate how much source j contributes at frequency i . The dimensions of A and S depend on the number of sources (N_S). Note that assuming a single mixing matrix A implies that the amount a source contributes at a given frequency does not change spatially, which is known to not hold true for the foreground sources ([Dunkley et al. 2009](#); [Fuskeland et al. 2014](#); [Planck Collaboration X 2016](#)). Incorporating local frequency dependence requires a modification to the framework and presents a unique challenge: finding a source separation model which allows local variation while maintaining global consistency.

Many different techniques have been developed to separate the CMB from observed data based on Equation (1). In Appendix A, we review algorithms used by *Planck* from the three main paradigms: template fitting, internal linear combination (ILC), and Bayesian inference. In the following sections we provide some background on the lattermost technique, and discuss GMCA and LGMCA – the main predecessors to HGMCA.

2.1 Bayesian Source Separation

Bayesian methods aim to isolate not just the CMB, but examine the entire posterior distribution $P(A, S|X)$. When done in a blind manner (*i.e.*, with few assumptions about both A and S), this reduces to a Bayesian version of a problem which has received wide attention in the machine learning and statistics communities known as Blind Source Separation (BSS). Popular methods for examining the posterior include solving for the maximum a posteriori estimate, including variants of independent component analysis (ICA) and GMCA ([Hyvärinen & Oja 2000](#); [Pham & Garat 1997](#); [Bobin et al. 2008, 2013](#)), or estimation via sampling ([Geman & Geman 1984](#)). Bayesian techniques such as these also allow for sophisticated priors on A and S such as non-negativity (a problem referred to as non-negative matrix factorization) ([Rapin et al. 2013](#)), or compact and simple representation in particular bases, like wavelets ([Bobin et al. 2008](#)). The widespread ILC-type algorithms discussed in Appendix A2 do not have analytic solutions that satisfy the constraints provided by such priors; therefore, Bayesian techniques are better suited to take advantage of known properties of the CMB and foregrounds.

2.2 GMCA

While previous source separation techniques, such as wavelet SMICA ([Moudden et al. 2005](#)), had already exploited the

distinction of the CMB and foreground sources in wavelet bases, none had used the fact that these sources are also sparse. [Bobin et al. \(2008\)](#) introduced Generalized Morphological Component Analysis (GMCA) to leverage the sparsity of the distinct signals as an additional constraint. GMCA begins by assuming that the noise N in the BSS model is Gaussian, such that we can rewrite the model as a conditional probability distribution²:

$$X \sim \mathcal{N}(AS, \sigma_X I), \quad (2)$$

where I is the identity matrix and $\mathcal{N}$ is a multivariate Normal distribution. Here the covariance matrix $\sigma_X I$ is determined by the level of instrumental noise. Enforcing the sparsity of S amounts to adding a leptokurtic prior to S of the form ([Bobin et al. 2008](#)):

$$P(S) \propto \prod_{i,j} e^{-\lambda |s_{i,j}|^\gamma}, \gamma \leq 1. \quad (3)$$

As γ goes to 0, this returns a negative-log-likelihood (up to an additive constant) of $\lambda \|S\|_0$, where $\|S\|_0$ counts the number of non-zero values of S and is known as the 0-norm.

To solve for A and S , the authors compute its Maximum a Posteriori (MAP) estimates, which in the limit of small γ amounts to solving the following optimization problem:

$$\operatorname{argmin}_{A,S} \left(\|X - AS\|_F^2 + \lambda_S \|S\|_0 \right), \quad (4)$$

where S and X have been transformed into the wavelet domain³, a convention that will be used throughout the paper unless otherwise noted. Appendix B gives a summary of the wavelet transform on the sphere, although the needlet wavelets used by the authors (see [Bobin et al. 2014](#)) differ from the scale-discretized wavelets used in this work. Note that $\|M\|_F \equiv \sqrt{\operatorname{Tr}(M^T M)}$ is the Frobenius norm of a matrix M , and λ_S is determined by the relative strength of the constant λ in $P(S)$ along with the covariance $\sigma_X I$ of $P(X)$.

Optimization for Equation (4) comes with a few inherent difficulties. First, the linear mixture model $X = AS$ has a scaling degeneracy—if A and S are inversely scaled by a constant, their product remains invariant. To circumvent this issue, [Bobin et al. \(2008\)](#) assume that the columns of A have a 2-norm of 1, and enforce this in their algorithm. We will make the same assumption in our optimization for HGMCA (see Appendix D), along with introducing a Bayesian interpretation of the assumption in Section 3.

Because $\|S\|_0$ is a combinatorial object, this optimization problem is, in general, NP-Hard, and thus intractable ([Bobin et al. 2008](#)). However, in many cases it is possible to either approximate or exactly solve the minimization problem through a convex relaxation—replacing the 0-norm with a 1-norm ($\|S\|_1 = \sum_{i,j} |s_{i,j}|$). The problem is then tackled via optimization techniques from compressed sensing such as LASSO shooting ([Santosa & Symes 1986](#)), or, more generally, optimization techniques involving soft or hard thresholding. Another option is to directly approximate the 0-

² Throughout this work, when we write $X \sim \mathcal{N}(AS, \sigma_X I)$ we are implicitly invoking a vector notation of the form $\vec{X} \sim \mathcal{N}(\vec{A}\vec{S}, \sigma_X I)$ where $\vec{X}$ is the vectorized form of X and $\sigma_X I$ encodes the covariance of $\vec{A}\vec{S}$.

³ Note that this conserves the linearity of the product AS .




Figure 1. Graphical model for the base GMCA algorithm. Variables are denoted by un-shaded circles, constants by dots, and observed data by shaded circles. Because our graphical model contains repeated sub-graphs, we use plate notation. A plate denotes a sub-graph repeated several times in the diagram, once per index provided in the plate label. Here, the plate stands for as many sub-graphs as there are wavelet scales μ . Arrows, which indicate probabilistic dependence between variables, cross plate lines and are assumed to be repeated once per index. For example, here each observed X^μ is generated by a corresponding S^μ and the global mixing matrix A (which is the same for all wavelet scales).

norm solution through k-thresholding⁴. We will outline the thresholding and LASSO strategies in Section 4.1. For a more detailed discussion of these techniques in the context of GMCA, see (Bobin et al. 2008).

As a visual aid, Figure 1 depicts the graphical model associated to the GMCA algorithm.

2.3 LGMCA

GMCA assumes that the mixing matrix is the same at every pixel, an assumption that does not hold for the heterogeneous emission in the sky, but is made for good reason: letting the mixing matrix vary across pixels is computationally intractable for problems such as the CMB. LGMCA (Bobin et al. 2013), the “local” modification of GMCA, incorporates locally-variable mixing matrices in a tractable manner, by exploiting a quad-tree decomposition to compromise between local adaptability and global consistency. This decomposition conveniently uses the hierarchical nature of the HEALPix pixelization scheme (Gorski et al. 2005). Figure 2 describes how this decomposition works on a full HEALPix map. At a given level l , the full map is decomposed into 4^l patches, meaning that, as l increases, the patches probe smaller areas of the map.⁵

In LGMCA, each level l is treated as an independent probabilistic model, where GMCA is applied to each patch. Because each level is treated independently, the algorithm produces multiple reconstructions for the same patch of data. To merge these reconstructions into a final map, LGMCA dynamically chooses the level at which to reconstruct each patch of the sky by selecting the reconstruction with the lowest variance. Bobin et al. (2013) argue that this heuristic will select the level with the least foreground contamination due to the statistical independence of the CMB from the instrument noise and the foreground components.

The local variation permitted by LGMCA has been shown to reduce galactic foreground contamination overall




Figure 2. HEALPix decomposition of sky map used by LGMCA at levels $l = 1, 2, 3$, and 4 , from left to right. Orange regions highlight the size of the patches of sky on which mixing matrices are calculated at a given level. Original HEALPix image credit: Gorski et al. (2005).

compared to GMCA and ILC-type methods, with a particularly large improvement for Sunyaev-Zel’dovich contamination (Bobin et al. 2013). However, a drawback of LGMCA is that it does not admit a unified probabilistic structure. To see why, consider the simple case of a set of data divided into four equal patches. LGMCA could propose the global mixing matrix solution (i.e. the mixing matrix acquired by fitting to all four patches at once) for three of the patches and the local mixing matrix solution (i.e. the mixing matrix acquired from fitting to each patch separately) for the remaining patch. The issue with a probabilistic interpretation for LGMCA lies with this final patch, which must now show up twice in a probabilistic model - once associated to the global mixing matrix and once to its own local mixing matrix. Because LGMCA treats the generation of the data by the global and local mixing matrices separately, the same data are generated by two different, independent processes, which is inconsistent with a probabilistic model. Additionally, because deeper (higher l) levels have no dependence on their predecessors, small patches are ignorant of the constraining power offered by the rest of the sky and are thus prone to overfitting.

3 HIERARCHICAL GMCA

Rather than heuristically choosing between multiple independent models run at each level, we propose expanding LGMCA’s notion of a quadtree decomposition to a single Bayesian hierarchical graphical model. Similar to GMCA, HGMCA aims to provide the MAP estimate of a posterior distribution. Unlike GMCA, this posterior distribution is spread across a hierarchy of mixing and source matrices, representing the entire probabilistic structure of HGMCA at subdivision level $l_{\max}$ (denoted HGMCA- $l_{\max}$):

$$P(\{A_p^l\}, \{S_p^\mu\} | X) = \underbrace{\left(\prod_{\mu} P(S_p^\mu) \prod_l \prod_{p \in \{p_l\}} \mathcal{N}(X_p^{l,\mu} | A_p^l S_p^\mu, \sigma_X l) P(A_p^l) \right)}_{\text{likelihood and priors}} \times \underbrace{\left(\prod_{l>0} \prod_{p \in \{p_{l-1}\}} \prod_{p' \in [1,4]} \mathcal{N}_{S,+}(A_{p,p'}^l | A_p^{l-1}, \sigma_A l) \right)}_{\text{Mixing matrix dependence across levels}}. \quad (5)$$

Figure 3 shows the corresponding graphical model to this posterior distribution for $l_{\max} = 3$ (HGMCA-3). Note that both Figure 3 and Equation 5 ignore some complexity related to the support of the wavelet scales that we will in-

⁴ Here we are referring to hard thresholding optimization schemes where all but k of the entries of a source are zeroed out.

⁵ This decomposition is done in the wavelet space (see Appendix B), meaning that the patches contain wavelet coefficients.

introduce in Figure 4 and Equation (12). Here the patch index p at level l corresponds to a tuple $(p_1, \dots, p_l)$ and $\{p_l\}$ is the set of all such tuples at level l . The tuple notation corresponds to the structure of the hierarchy: for an $l-1$ tuple p , the level l mixing matrix $A_{p,p'}^l$ is drawn from $A_{p'}^{l-1}$ with $p' \in [1,4]$. $X_p^{l,\mu}$ is a sub-matrix of X corresponding to the patch p at level l and wavelet scale μ , and A_p^l and $S_p^{l,\mu}$ are likewise the patch-specific mixing and source matrices. Note that the mixing matrix should not vary across different wavelet scales, and therefore there is only one mixing matrix per patch for all wavelet scales μ . $X_{p_1,p_2,p_3}^{3,\mu}$, which is our observed variable, is generated from A_{p_1,p_2,p_3}^3 and $S_{p_1,p_2,p_3}^{3,\mu}$. The value of $P(S_p^{l,\mu})$ enforces the sparsity of the coefficients, with the exact distribution discussed in Section 2.2. $P(A_p^l)$ is a Heaviside function prior that enforces the non-negativity constraints discussed in Appendix C (in short, A cannot be negative because it quantifies the contribution of a given source at a given frequency, which is positive for the CMB and its foregrounds⁶).

We have made the assumption that a mixing matrix is drawn from a Gaussian defined by its prior and a covariance matrix $\sigma_A I$. More flexible choices can be made, but they would introduce additional parameters to the algorithm. The remainder of this section will be spent detailing the hierarchical dependence represented by Equation (5) and the optimization strategy for HGMCA.

Figure 3 shows the basic plate diagram describing the generative probabilistic structure of our model. In this example, we assume that the map is divided up to level $l=3$, which means that it is being divided into 4^3 patches⁷. At the level $l=0$, there is a mixing matrix A^0 that encodes the most general mixing matrix from which all the local mixing matrices are thenceforth drawn. While this mixing matrix does not directly generate the data, it is useful to think of it as the global mixing matrix for the full map. From this mixing matrix, a set of four additional mixing matrices $A_{p_1}^1$ are drawn from a Gaussian restricted to the positive orthant of the unit sphere, denoted $\mathcal{N}_{S,+}$ ⁸. Each $A_{p_1}^1$ corresponds to one of the four subdivisions of the full map at level 1, $p_1 \in [1,4]$. If we were only running our analysis up to level 1, these mixing matrices would then be responsible for generating the data. However, in this example we go up to level $l=3$, therefore each patch p_1 is further subdivided into 4 patches indexed by p_2 , giving the corresponding mixing matrices $A_{p_1,p_2}^2 \sim \mathcal{N}_{S,+}(A_{p_1}^1, \sigma_A I)$. The process is repeated one final time to get the set of 64 mixing matrices $A_{p_1,p_2,p_3}^3 \sim \mathcal{N}_{S,+}(A_{p_1,p_2}^2, \sigma_A I)$ that each correspond to one patch of the map at subdivision level $l=3$.

By constructing this hierarchical structure we can better encode and quantify the trade-off between local adapt-




Figure 3. Graphical model for Hierarchical GMCA at subdivision level $l=3$. $X_{p_1,p_2,p_3}^{3,\mu}$ represents the input maps at wavelet scale μ and patch p_1, p_2, p_3 of subdivision level 3. $S_{p_1,p_2,p_3}^{3,\mu}$ represents the source at the same patch and wavelet scale, and A_{p_1,p_2,p_3}^3 represents the mixing matrix for that patch. The hierarchical nature of the mixing matrices is akin to a Hierarchical Bayesian Model (Gelman et al. 2013), but the addition of the source matrix S distances our model from this framework. For details on the patch notation see Section 3. See Figure 1 for a description of the probabilistic graphical model notation.

ability and global consistency. Each patch is constrained by the data on the rest of the map through the tree of priors, and the strength of this constraint is dictated by the ratio of σ_X and σ_A .

So far, we have assumed that we can subdivide the map at all scales and frequencies down to subdivision level 3. However, the support of the wavelets being used at scale μ effectively limits the resolution of the map X^μ : the level l of analysis used for X^μ cannot subdivide the map into patches that are smaller than that resolution limit.

To account for this, we modify the model such that the data at wavelet scale X^μ is analyzed at the highest l permitted by the scale's resolution limit (recall that patch size decreases with increasing l). We call the set of scales whose resolution limits them to level l $\{\mu_l\}$. Note that there is no overlap between $\{\mu_0\}$, $\{\mu_1\}$, $\{\mu_2\}$, and $\{\mu_3\}$, as each map at each scale can only be allowed to contribute once in the graphical model. Further, it is permissible that a set $\{\mu_l\}$ be empty. A visual representation of HGMCA incorporating this additional source of complexity is shown in Figure 4.

Note that the mixing matrix does not vary within $\{\mu_l\}$, and therefore lacks the superscript μ . On the other hand, we do not constrain the mixing matrix A_{p_1,p_2}^2 of patch (p_1, p_2) to be the same as its parent $A_{p_1}^1$. This accounts for the higher order corrections to the local variation of the mixing matrix that scales at level $l=2$ are sensitive to but scales at $l=1$ are not.

⁶ This does not hold true for SZ clusters, but one could easily remove the non negativity prior for a single source and add an SZ frequency scaling prior.

⁷ In practice, HEALPix requires that the first level be divided into twelve rather than four patches. For notational simplicity, we will use four throughout the paper.

⁸ In more detail, each column is being drawn from a Von Mises-Fisher distribution restricted to the positive orthant.: $A_{p_1}^1 \sim \mathcal{N}_{S,+}(A^0, \sigma_A I)$




Figure 4. Graphical model for Hierarchical GMCA at subdivision level $l = 3$. Note that, unlike Figure 3, this version of the model accounts for the resolution limit of each wavelet scale. This requires analyzing different wavelet scales at different levels of subdivision. $X_{p_1, \dots, p_l}^{l, \mu}$ represents the input maps at subdivision level l , wavelet scale μ , and patch $p_1, \dots, p_l$ of subdivision level l . $S_{p_1, \dots, p_l}^{l, \mu}$ represents the source at the same patch and wavelet scale, and $A_{p_1, \dots, p_l}^l$ represents the mixing matrix for patch $p_1, \dots, p_l$. $\{\mu_l\}$ is the set of scales whose resolution limits them to being analyzed at level l . There is no overlap between $\{\mu_i\}$ and $\{\mu_j\}$ for $i \neq j$. See Figure 1 for a description of the probabilistic graphical model notation.

4 IMPLEMENTATION

4.1 GMCA Optimization

Computing the maximum a posteriori (MAP) of HGMCA reduces to running GMCA in each patch with additional priors from the mixing matrix hierarchy. Bobin et al. (2008) computes the MLE of GMCA via coordinate descent. By fixing all but a single column of the mixing matrix A or row of the source matrix S , one can derive a closed form update that minimizes the negative log-posterior of the model. The results we present in this work are based on our own GMCA optimization strategy. While the mixing matrix optimization remains the same, we relax the 0-norm sparsity prior used in GMCA to a 1-norm, as suggested in Bobin et al. (2008) and discussed in Section 2.2. This allows us to write the closed form solution for a single row of our source matrix:

$$s_i = (a^i)^T R_i - \frac{\partial}{\partial s_i} \frac{\lambda_s}{2} \|s_i\|_1, \quad (6)$$

and a single column of our mixing matrix:

$$a^i = \left[\frac{R_i s_i^T}{\|R_i s_i^T\|_F} \right]_+. \quad (7)$$

Here s_i is a row of the source matrix, a^i is the corresponding column of the mixing matrix, $R_i = X_p^l - \sum_{j \neq i} a^j s_j$, and $[\dots]_+$ indicates a non-negativity constraint enforced by setting negative values to 0. See Appendix D and Bobin et al. (2008) for a more detailed derivation of these equations.

Note that Equation (6) contains a derivative with respect to the 1-norm $\|s_i\|_1$. There are a number of approximate methods to solve this type of optimization problem in the literature. Due to its relative simplicity of implementation, we use the LASSO shooting technique (Santosa & Symes 1986), which gives the update equation:

$$s_i = \begin{cases} (a^i)^T R_i - \frac{\lambda_s}{2} \text{sign}((a^i)^T R_i) & |(a^i)^T R_i| > \frac{\lambda_s}{2} \\ 0 & |(a^i)^T R_i| \leq \frac{\lambda_s}{2} \end{cases}. \quad (8)$$

While standard LASSO shooting would suggest repeatedly conducting this update step until convergence, this is wasteful in the case of (H)GMCA as updating the remaining sources moves the point of convergence. Instead this update is only applied once before moving to the next source, and convergence is reached over many iterations over the rows of the source matrix and columns of the mixing matrix. We also empirically find that updating S with the least squares solution (equivalent to setting $\lambda_s = 0$ in the loss function) every ~ 100 steps improves the rate of convergence and mitigates issues with local minima (tests of convergence are discussed in more detail in Section 5.2). We will call this version of the algorithm GMCA^{LASSO}.

A public version of the GMCA implementation can be found online⁹. This version of GMCA implements an optimization procedure in line with what is described in Bobin et al. (2008): it utilizes soft or hard median absolute deviation (MAD) thresholding¹⁰ to optimize the sparsity term and implements an alternating least squares version of GMCA known as “Fast GMCA.” By assuming that the mixing matrix has nearly orthogonal columns, the update rule for this optimization only needs to fix one matrix at a time. Fixing S , the solution for the mixing matrix becomes

$$A = X S^T (S S^T)^{-1}. \quad (9)$$

When used for CMB estimation, Bobin et al. (2008) suggest fixing one column of the mixing matrix to the theoretical frequency dependence for the CMB, which we denote as $\hat{a}^{CMB}$. This involves resetting that column back to $\hat{a}^{CMB}$ after applying Equation (9). Fixing A , the solution for the source matrix in the case of soft thresholding is:

$$s_i = \begin{cases} [A_{\text{pinv}} X]_i - \lambda_{\text{MAD}}(t) \text{sign}([A_{\text{pinv}} X]_i) & |[A_{\text{pinv}} X]_i| > \lambda_{\text{MAD}}(t) \\ 0 & |[A_{\text{pinv}} X]_i| \leq \lambda_{\text{MAD}}(t) \end{cases}. \quad (10)$$

where $\lambda_{\text{MAD}}(t)$ is a thresholding parameter scaled by the MAD of $[A_{\text{pinv}} X]_i$ at iteration t . $[M]_i$ refers to the i^{th} row of matrix M . Note that, unlike in GMCA^{LASSO}, $\lambda_{\text{MAD}}(t)$ is allowed to decrease with increasing iterations. The solution for hard thresholding does not subtract from the surviving source components:

$$s_i = \begin{cases} [A_{\text{pinv}} X]_i & |[A_{\text{pinv}} X]_i| > \lambda_{\text{MAD}}(t) \\ 0 & |[A_{\text{pinv}} X]_i| \leq \lambda_{\text{MAD}}(t) \end{cases}. \quad (11)$$

In Appendix F we compare Fast GMCA to our GMCA^{LASSO} implementation. While the optimum performance of both approaches is comparable, we find that

⁹ <https://github.com/jbobin/pyGMCA/tree/master/pyGMCA>

¹⁰ Bobin et al. (2008) suggest that while soft thresholding is more rigorous, the code returns better results with hard thresholding




Figure 5. From top left to bottom right: the CMB, dust, synchrotron, anomalous microwave emission (AME), and free-free input maps that were combined to create the mock data used in this work. The full map with all signals combined is in the bottom right. Note that each input is presented at 90 GHz and has a different scaling for the color bar so that the structure is clearly visible.

GMCA^{LASSO} is significantly more stable to perturbations in the random seed and CMB realization. Since the main purpose of this work is to present our hierarchical approach, and we intend to compare the generalizability and performance of HGMCA to a wavelet-based ILC-type algorithm, we will default to the more stable GMCA^{LASSO}.

4.2 HGMCA Optimization

To compute the MAP estimate for HGMCA, we adopt a similar strategy. We first convert the posterior in Equation (5) into a loss function defined by the negative log-likelihood

(with additive constants removed). We can write this as

$$\begin{aligned} \mathcal{L} = & \underbrace{\sum_l \sum_{\mu_l \in \{\mu_l\}} \sum_{p \in \{p_l\}} i_S^+(A_p^l) + \lambda_S \|S_p^{l, \mu_l}\|_1 + \|X_p^{l, \mu_l} - A_p^l S_p^{l, \mu_l}\|_F^2}_{\text{non-negativity prior, source prior, and reconstruction error}} \\ & + \underbrace{\sum_{l>0} \sum_{p \in \{p_{l-1}\}} \sum_{i \in \{1, \dots, 4\}} \lambda_A \|A_p^l - A_{p,i}^l\|_F^2}_{\text{A matrix dependence across levels}}. \end{aligned} \quad (12)$$

where $i_S^+(A)$ is the characteristic function of positive matrices with normalized columns:

$$i_S^+(A) = \begin{cases} 0 & A \geq 0 \text{ and } \|A^i\|_2 = 1 \\ \infty & \text{else.} \end{cases} \quad (13)$$

The factors λ_A and λ_S come from multiplying the log-likelihood by σ_X^2 and are ratios of the uncertainty in the draws of X_p^{l, μ_l} to the uncertainty in our priors. For example, $\lambda_A = \sigma_X^2 / \sigma_A^2$.




Figure 6. *Top:* the ratio of the 90 GHz map to the 23 GHz map for a realization of synchrotron radiation. *Bottom:* the same ratio for a realization of dust radiation. Both show local variation: without it, we would expect the ratios to return a constant value across the sky, but in these maps it varies spatially.

As in the case of GMCA, while no tractable closed form solution for the mixing and source matrices in the loss function exists, it is possible to optimize by iterating over an approximate closed form¹¹. We use a coordinate descent strategy that is executed as follows. First, we initialize the sources and mixing matrices at all levels of subdivision. We find that using a short run of GMCA^{LASSO} as an initialization point rather than a random initialization improves convergence. We then extract an approximate closed form solution by fixing all but one row of the source matrix or column of the mixing matrix and minimizing the loss function. Here we state these solutions; their full derivation can be found in Appendix D.

For a given row of the source matrix, the closed form solution is still given by Equation (10). Note, however, that in this case the source matrix is level-dependent. For a given subdivision level l and patch $p \in \{p_l\}$, S_p^l is the concatenation of source matrices for all scales in $\{\mu_l\}$, as each shares the same source matrix. S_p^l then has dimensions of (N_S) by (number of pixels times number of scales in μ_l), and the concatenated data matrix X_p^l has dimensions (number of frequencies) by (number of pixels times number of scales in μ_l).

For a column of the mixing matrix, the closed form now

includes priors from the mixing matrix hierarchy:

$$a^i = \left[\frac{R_i s_i^T + \lambda_A \tilde{a}^i + \lambda_A \sum_j b_j^i}{\|R_i s_i^T + \lambda_A \tilde{a}^i + \lambda_A \sum_j b_j^i\|_F} \right]_+. \quad (14)$$

Here $\tilde{A}$ denotes the mixing matrix that defines the distribution from which A is drawn. As an example, for $A = A_{p_1}^1$ this would be $\tilde{A} = A^0$. B_j denotes the j^{th} matrix that is drawn from A . As an example, for $A = A_{p_1}^1$ this would be $B_j = A_{p_1,j}^2$ for $j \in [1, 4]$. The terms b_j^i and $\tilde{a}^i$ are the i^{th} column of B_j and $\tilde{A}$ respectively. Finally, recall that $[\dots]_+$ indicates a non-negativity constraint enforced by setting negative values to 0. Note that setting λ_A to 0 returns the GMCA update equation (7).

At each level, we optimize the sources and mixing matrices iteratively before proceeding to the next level, iterating over all the levels until convergence is reached.

4.3 CMB Prior

While (H)GMCA is blind in theory, in practice it improves performance to include theoretical knowledge such as the frequency dependence of the CMB. [Bobin et al. \(2008\)](#) suggest enforcing this on the column of the mixing matrix corresponding to the CMB (a^{CMB}) throughout optimization, which is equivalent to setting an infinitely strong (delta function) prior on the column. Because the CMB spectrum is not measured to infinite precision, in our GMCA^{LASSO} implementation we choose to use a Gaussian prior. For HGMCA, this means that the a^{CMB} column of the mixing matrix is drawn from this Gaussian prior rather than the hierarchy, reflecting the fact that the CMB has no local variation to the frequency dependence. The update equation for the CMB column then becomes

$$a^{\text{CMB}} = \left[\frac{R_i s_i^T + \lambda_{\text{CMB}} \tilde{a}^{\text{CMB}}}{\|R_i s_i^T + \lambda_{\text{CMB}} \tilde{a}^{\text{CMB}}\|_F} \right]_+, \quad (15)$$

where λ_{CMB} is the strength of the CMB prior, and $\tilde{a}^{\text{CMB}}$ is the theoretical frequency dependence of the CMB.

4.4 Simulated Data and Wavelet Transformation

In order to understand the benefits of a hierarchical modeling strategy, we compare HGMCA with a number of baselines on simulated CMB and foregrounds. The simulations are generated using the Python Sky Model (PySM) code ([Thorne et al. 2017](#)) at a resolution of $N_{\text{side}} = 256$ and bandlimited to $\ell_{\text{CMB}} = 620$. PySM offers different models to simulate the foregrounds. These are built from *Planck* templates and make conservative assumptions about their spectral dependence. We use Model 1 for synchrotron, thermal dust, anomalous microwave emission (AME), and free-free emission. The CMB random seed parameter was set to 2 for the final results and 1111 for hyperparameter tuning. The exact details of this model, its parameters, and the underlying assumptions can be found in [Thorne et al. \(2017\)](#). Note that while the PySM code is capable of modeling both instrument noise and distortions from the beam, for simplicity we do not include these effects in our simulated data. The maps

¹¹ Stochastic gradient descent would be another possibility, but it performed poorly in GMCA^{LASSO}

are generated at frequencies corresponding to those measured by *Planck* (Planck Collaboration I 2014). In this work we focus on intensity maps, and leave the polarization maps for future work.

The individual CMB and foreground maps that are used to create the mock full sky map for the results presented in this work are shown in Figure 5. The bottom right panel shows the full sky image at 90 GHz. The top and bottom panels of Figure 6 show the ratio between the 90 GHz and 23 GHz maps for synchrotron radiation (top) and dust emission (bottom), which showcase the local variation in the foreground signal. If there were no spatial dependence to the frequency scaling, the ratio would be a constant value everywhere; instead, the ratio changes considerably as we move across the sky. We draw special attention to the spatial dependence of the frequency scaling because HGMCA has been built to better model local variations.

All of the inputs to the algorithms we present were converted into the scale-discretized wavelet basis. Appendix B describes this transform and the accompanying code in detail. For the results in this work we used $\lambda = 3$, a bandlimit of $\ell_\mu = 621$ (having a bandlimit one larger than the bandlimit of the data gives effectively zero reconstruction error), and an analysis depth for the scaling function of $J_0 = 1$.

4.5 Wavelet ILC

To compare HGMCA to the foreground correction algorithms used by *Planck*, we have implemented a wavelet ILC (WavILC) code similar to what is done in Planck Collaboration IV (2018); Basak & Delabrouille (2013); Rogers et al. (2016a). We use the same scale-discretized wavelets as for our HGMCA and GMCA implementations, but the minimum J_0 is set to 6 (we find a larger analysis depth for the scaling function is important for ILC-type algorithms). To calculate the covariance matrix we follow the same approach as used in Basak & Delabrouille (2013); Rogers et al. (2016a) and convolve our map with a Gaussian whose FWHM is set to be $9000/N_{\text{side}}$ radians. More precise details of how the covariance matrix and weights are calculated can be found in Appendix A4. Note that both WavILC and the Needlet Internal Linear Combination (NILC) method on which it is based calculate different ILC weights for each pixel, meaning that they are capable of capturing local variation in the mixing matrix.

While our WavILC code uses the same scale-discretized wavelets as Rogers et al. (2016a), our implementation has notable differences. Rogers et al. (2016a)'s implementation of Scale-discretized directional wavelet ILC (SILC) is unique among ILC algorithms in that it does not use the HEALPix pixelization scheme. Instead it uses the sampling scheme developed in McEwen & Wiaux (2011), which requires fewer samples for a bandlimited signal. SILC also varies the wavelet dilation parameter λ by multipole, whereas here we keep the same value of $\lambda = 3$ for all multipoles. This is true both for the wavelet coefficients we use for WavILC as well as the wavelet coefficients we use for GMCA and HGMCA. Similarly, SILC takes advantage of the directional extension to scale-discretized wavelets, while we do not. We also do not mask point sources or utilize inpainting because the simulations we present here are not as complex as the Full Focal

Plane 8 simulations used by SILC (Planck Collaboration XII 2016).

4.6 Pipeline

The HGMCA results presented in this work were constructed through the following steps:

- (i) The simulation data is generated using PySM. The parameters for the PySM code are discussed in Section 4.4. The data is bandlimited to $\ell_{\text{CMB}} = 620$.
- (ii) The region of the map corresponding to the *Planck* UT78 mask is scaled by 0.5 for all frequencies. We found that this preprocessing step improves convergence.
- (iii) The input data is transformed into the wavelet basis using the S2LET¹² (McEwen et al. 2015b; McEwen 2015; McEwen et al. 2015a; Chan et al. 2016) code. A discussion of the basis can be found in Appendix B. For the results in this work we used $\lambda = 3$, a bandlimit of $\ell_\mu = 621$, and a minimum $J_0 = 1$.
- (iv) HGMCA is run on the transformed input data until convergence. This consists of:

- (a) Grouping the scales by their resolution limit. Since we do not have varying beam resolutions, this is dictated by the wavelet scales themselves (Section 3), i.e. group data into the sets $\{\mu_l\}$.
- (b) Initialize the mixing matrix hierarchy and the source matrix with one epoch of the base GMCA^{LASSO} algorithm. By providing a smart initialization, this step cuts back on the total runtime of HGMCA.
- (c) For each epoch and each level, update the source and mixing matrices for each patch using Equations (10) and (14). For results presented in Section 5, we use the hyperparameter values $\lambda_S = 50$, $N_S = 5$, $\lambda_{\text{CMB}} = 10^{11}$, and $\lambda_A = 5 \times 10^{10}$. A discussion of how these values are chosen and the sensitivity of our results to the hyperparameters can be found in Section 5.3.
- (d) Every 100 epochs, rather than using Equation (10) to update the sources, use the least squares solution $S = A_{\text{pinv}}X$ where A_{pinv} is the pseudo-inverse of A . As we discuss in Section 4.1, we found empirically that this step improves convergence and helps avoid local minima.
- (v) Use the mixing matrix hierarchy and the source matrix to reconstruct the CMB signal in the wavelet basis. If the known frequency dependence of the CMB was used as a prior, the column on which this prior was applied is selected as the CMB. Otherwise, whichever column most closely resembles the CMB frequency dependence is selected.
- (vi) Transform the wavelet representation of the CMB back into the pixel basis using S2LET.

5 RESULTS

The principal advantage of HGMCA is that it has the ability to fit to local variations without sacrificing the robustness offered by having access to the full map. To showcase the strength of the algorithm, we empirically compare its performance to both GMCA^{LASSO} (which produces one global

¹² <http://www.s2let.org>

Algorithm	RMSE [μK]	RMSE _{UT78} [μK]	% Error CMB C_ℓ 's, all ℓ 's	% Error CMB C'_ℓ 's, $\ell > 200$
GMCA ^{LASSO}	10.4430	1.2004	0.3815	0.1439
HGMCA - 1	8.2860	1.1235	0.2582	0.0314
HGMCA - 2	7.7935	0.9044	0.1663	0.0322
HGMCA - 3	6.9835	0.9467	0.2311	0.0289
WavILC	4.3412	1.8538	0.9416	0.0487

Table 1. Comparison of RMSE (μK) and average percent error in the CMB angular power spectrum including and excluding the *Planck* UT78 mask region for WavILC, GMCA^{LASSO}, and HGMCA at several levels. GMCA^{LASSO} and HGMCA maps are selected via the input CMB blind criterion described in Section 5.3.2. HGMCA outperforms GMCA^{LASSO} across all metrics and WavILC across all metrics except unmasked RMSE. The comparison between HGMCA-2 and HGMCA-3 is less clear, with HGMCA-2 winning on masked RMSE and overall C_ℓ error, and HGMCA-3 performing better on unmasked RMSE and larger C_ℓ 's.

fit) and WavILC (which can capture local variation). Using separate CMB sky realizations to tune parameters and evaluate results, we find that HGMCA is able to consistently return the highest fidelity CMB reconstruction (Section 5.1). In addition, we show that GMCA^{LASSO} and HGMCA converge to a stable solution (Section 5.2), and confirm that our results hold under a number of generalization tests (Section 5.3). Finally, we show that GMCA^{LASSO}, HGMCA, and WavILC do not introduce substantial non-Gaussianity into the reconstructed map, suggesting that such algorithms may be of interest for non-Gaussianity studies (Section 5.4).

5.1 GMCA^{LASSO}, HGMCA, and WavILC Comparison

We run a comparison between GMCA^{LASSO}, WavILC, and HGMCA at level 1 (HGMCA-1), level 2 (HGMCA-2), and level 3 (HGMCA-3)¹³. All of the algorithms are run on the input maps described in Section 4.4. We leave LGMCA out of the comparison because the code is not publicly available, and our implementation was not able to reproduce the improvements shown in Bobin et al. (2013)¹⁴. To determine the quality of the reconstructions, we use three metrics: (1) the root mean square error (RMSE) of the residual between the reconstructed output CMB and the input CMB maps, (2) the cross-power spectrum between the residual map and the input foregrounds, and (3) the difference between the power spectra of the input CMB and the reconstructed CMB.

The results for each algorithm are presented in Table 1, where we show the RMSE between the input and output CMB maps (both with and without applying the UT78 mask used by *Planck* (see Planck Collaboration XI 2016)¹⁵, as well

as the difference in their power spectra. The HGMCA algorithm generates a better reconstruction of the CMB than GMCA^{LASSO} across all metrics and WavILC across all metrics except unmasked RMSE. In Figure 7 we present the masked residual maps for the five algorithms. The residuals show that all three levels of HGMCA offer a marked improvement over GMCA^{LASSO} and WavILC across the masked map. Notably, HGMCA shows little residual signal at the poles of the map, with almost none at all for HGMCA-2. For the HGMCA maps, the masked RMSE is roughly equivalent for levels $l = 3$ and $l = 2$, but the residual increases as we go to level $l = 1$. Figures 7 and 8 show that this increased RMSE comes from additional dust contamination (see Figure 5 for the shape of the dust foreground).

To further understand the remaining foreground contamination, we show the correlation coefficients between the residuals of the five reconstructed maps (GMCA^{LASSO}, WavILC, HGMCA-1, HGMCA-2, and HGMCA-3) and the input foregrounds in Figure 8. By using the residual maps, we can probe what contaminants still remain in the reconstructed maps and at which multipoles they contribute. For anomalous microwave emission (AME), dust, and synchrotron contamination, HGMCA shows a reduction by an order of magnitude over GMCA^{LASSO} and WavILC. These differences in source and level of contamination are further evidenced by the residual maps in Figure 7. Especially noticeable is the difference in the residual maps of WavILC and HGMCA-2. WavILC's residual map shows clear resemblance to the input synchrotron foreground from Figure 5, whereas HGMCA-2's residual contamination is focused near the core, likely coming from dust or free-free emissions. Similar structure in the residuals was also observed in the application of other ILC implementations to the *Planck* simulations (Rogers et al. 2016a). Large-scale residuals in the CMB reconstructions are typical of ILC-type methods since they result from chance correlations between the CMB and the foregrounds and the variance minimization built into the algorithm. This unintentionally cancels a few modes, proportionally to the power spectrum of the CMB, and the effect is therefore more pronounced on large scales. The improvement

¹³ The parameter values used for GMCA^{LASSO} and HGMCA are given in Appendix E and the values for WavILC are discussed in Section 4.5. The selection strategy for the parameters is discussed in Section 5.3 and Appendix E.

¹⁴ The results presented in Bobin et al. (2013) are run on higher resolution simulations with additional contaminants and a varying beam.

¹⁵ Foregrounds typically dominate regions of the Galactic plane to such an extent that they cannot be removed when reconstructing the CMB from a few frequency channels (many more channels, as well as external data, would be needed to successfully disentangle the CMB from foreground in those regions). There-

fore, most applications (*e.g.*, measurements of non-Gaussianity, isotropy, CMB lensing, Planck Collaboration XIX 2014; Planck Collaboration VIII 2018; Planck Collaboration VII 2019; Planck Collaboration IX 2019) conservatively mask those regions.




Figure 7. From top left to bottom: the residual maps for HGMCA-3, HGMCA-2, HGMCA-1, WavILC, and $\text{GMCA}^{\text{LASSO}}$. The dark gray region corresponds to the UT78 mask used by *Planck*, and the scale of the color bar has been significantly reduced to accentuate the structure of the residual. All HGMCA maps have less residual foreground contamination than $\text{GMCA}^{\text{LASSO}}$ and WavILC on the masked map, with HGMCA-2 as the lowest (see Table 1).

of HGMCA over $\text{GMCA}^{\text{LASSO}}$ seems to come mostly from decreasing the contamination from AME and synchrotron at the second level, and dust at the third level.

Within the different levels of HGMCA itself, as the error in the C_ℓ 's in Table 1 suggest, HGMCA-2 has the lowest correlation to AME and synchrotron contaminants, and the lowest overall contamination. Our analysis also suggests that the improvement of HGMCA-3 over HGMCA-1 comes from an order of magnitude improvement on dust contamination (reflected in the residual map itself in Figure 7) and the improvement at low multipoles. Despite this large improvement in dust contamination, HGMCA-3 has increased AME and synchrotron residuals when compared to HGMCA-2, causing HGMCA-2 to win across most metrics.

We also look at the power spectrum difference between the input CMB and reconstructed CMB for all five algorithms. Figure 9 shows the percent level error as a function of ℓ . As with the previous two tests, HGMCA at all levels offers a significant improvement over the base $\text{GMCA}^{\text{LASSO}}$

algorithm. This difference is most pronounced at $\ell > 300$, but persists all the way to the smallest ℓ . WavILC is generally competitive with HGMCA, except in the lowest ℓ bin, where it is an order of magnitude larger (this causes the average C_ℓ error to be higher). Of the three HGMCA variants, HGMCA-2 performs the best overall, with a $\sim 0.065\%$ improvement over the second best performer, HGMCA-3. At mid to high ℓ 's, HGMCA-3 slightly improves over HGMCA-2 by about 0.003% . $\text{GMCA}^{\text{LASSO}}$ performs poorly on the C_ℓ 's because it struggles to converge to the correct mixing matrix column for the CMB. Post-processing the $\text{GMCA}^{\text{LASSO}}$ solution by enforcing that this column equal the CMB prior makes GMCA more competitive on the C_ℓ 's but has a negligible impact on other metrics.

All three metrics show that the added versatility permitted by the HGMCA model leads to lower foreground contamination in CMB reconstructions. Interestingly, while the transition from GMCA to HGMCA-1 to HGMCA-2 leads to a consistent improvement in results, HGMCA-3 does not




Figure 8. Magnitude of the cross-spectra between the residual CMB for each algorithm and the input AME, dust, synchrotron, and free-free foreground contaminants. A lower value of the cross-spectrum indicates a smaller level of contamination from that foreground component. HGMCA outperforms WavILC across all contaminants, especially at high ℓ . HGMCA-1 shows similar performance to GMCA^{LASSO} throughout, and levels 2 and 3 outperform GMCA^{LASSO} across AME, dust, and synchrotron contaminants. In general, increasing the level of subdivision from $l = 1 \rightarrow l = 2$ minimizes foreground, but the step to $l = 3$ does not. The only exceptions to this is the improved performance of HGMCA-3 on dust, and of GMCA^{LASSO} and HGMCA-1 on free-free contamination.




Figure 9. Percent error in C_ℓ as a function of multipole moment ℓ for each of the algorithms. WavILC shows its strength for reconstruction at high ℓ over GMCA^{LASSO}, but HGMCA is able to produce a better reconstruction across almost all values of ℓ . The differences between the levels of HGMCA are more subtle, and are better understood quantitatively (see Table 1).

improve over HGMCA-2. This is likely because the additional local variation captured in transition between level 2 and level 3 of subdivision is small, and the low resolution of our maps limits how deep we can go before the optimization at the deepest levels is prior dominated. The generalization tests in Section 5.3.4 show how increasing the local variation at smaller scales causes HGMCA-3 to outperform HGMCA-2.

5.2 Convergence

In Figure 10 we show that the reconstruction error of GMCA^{LASSO} and HGMCA converge in the limit of many iterations. Here, reconstruction error is measured as the root mean squared error (RMSE) of the residual map. Note that we are measuring the convergence of the error on the residual map (which we cannot directly measure in our loss function), so we do not expect it to monotonically decrease. Indeed, at some points our algorithm achieves lower RMSE than its converged value. However, in order to make our algorithm blind to the input CMB, we always report on the converged map. This ensures that our results may still be applied in practice, where the input CMB is unknown.

Since our implementations of both algorithms reach a stable state after sufficient iterations, we do not need to tune any specific convergence parameter to achieve good performance. In this paper, all of the HGMCA results have been run for 800,000 iterations and all of the GMCA^{LASSO} results have been run for 400,000 iterations. These values were chosen to be well beyond the point of convergence for both algorithms and therefore ensure that all of our results are stable.

5.3 Generalization Errors

There are four parameters intrinsic to the HGMCA algorithm that must be set either theoretically or empirically: λ_S , λ_A , λ_{CMB} , and N_S . We ran a grid search over a wide range of possible values for the first three of these parameters and selected the configuration that yielded the lowest RMSE error for GMCA^{LASSO} and HGMCA (the exact val-




Figure 10. The reconstruction error as a function of iteration for the HGMCA and GMCA^{LASSO} maps presented in Table 1. Because we are measuring the RMSE of the residual map, we do not expect the final converged RMSE to be the lowest. However, all four algorithms converge in the limit of large numbers of iterations.

ues can be found in Appendix E). One potential concern is that the performance achieved by these parameter configurations is not generalizable beyond the simulations used to select them; this effect is often called overfitting. In this section, we discuss the four major steps we have taken to address this issue: 1) the independence of training and testing simulations, and the robustness of HGMCA 2) to random seeds, 3) to perturbations in parameters, and 4) to changes in foreground simulation.

5.3.1 CMB Realization

In order to show that our results are not an artifact of the specific simulation we used to tune our parameters, we exclusively report on a different simulation set. PySM generates the CMB from a Gaussian random field, and thus returns different maps depending on the seed for this process. For our hyperparameter tuning, this seed was set to 1111 instead of 2, as reported in Section 4.4. By changing the CMB instance, we train our parameters and report results on different maps, thus showing the robustness of our algorithm to changes in CMB simulation.

5.3.2 Random Seeds

One concern with the coordinate descent optimization strategy we employ for GMCA^{LASSO} and HGMCA is that the order in which the sources and levels are updated is arbitrary. We do not want the results of our algorithm to be strongly dependent on the specific order, since that too would be a sign of poor generalization. The standard approach is to choose the order of the sources or levels randomly at each step of optimization. This allows us to probe a wide range of orders simply by altering the seed used by the pseudo-random number generator of our algorithm. In Figure 11 we show the results for GMCA^{LASSO}, HGMCA-1, HGMCA-2, and HGMCA-3 on the percent error in the CMB C_ℓ metric for a number of random seeds. For each of the four plots, the transparent blue lines correspond to the results of different seeds. We are only showing the results of runs whose standard deviation of the reconstructed map falls in the bottom

75% of the random seeds to remove outliers. The red line corresponds to the performance of WavILC and the dark blue line corresponds to the maps we presented as our results in Section 5.

The maps we used for our results were selected by making two cuts. First, as we mentioned in Section 4.1, every 100 iterations of our algorithm we conduct a single step of least squares optimization on the source. This is equivalent to setting

$$S = A_{\text{pinv}}X, \quad (16)$$

where A_{pinv} is the Moore-Penrose pseudo-inverse of the matrix A . Since the matrix A_{pinv} is not a true inverse, $AS = AA_{\text{pinv}}X$ may not equal X . This means that some A matrices will return a large error for $\|X - AA_{\text{pinv}}X\|_F^2$ even without sparsity constraints, suggesting they will return a lower quality reconstruction. We therefore conservatively cut any run whose pseudo-inverse introduces error in the top 50% of runs. From the remaining runs, we select the one with the minimum standard deviation of the reconstructed map. Note that both of these cuts are completely blind to the input CMB map and are therefore reproducible outside of simulations.

While employing these two cuts does help select one of the better performing seeds, Figure 11 shows that selecting any seed with small standard deviation (in the bottom 75% of seeds) has better or equivalent performance to WavILC. This demonstrates that the results presented in this work are robust to the exact pattern of optimization that is employed. While the selection function we use for our results does extract one of the seeds with better performance, it is neither an outlier nor the best performer.

5.3.3 Parameter Perturbations

While we have made efforts to ensure that the parameter values used for our results are not overfit to the specific simulation, it is still important to understand the sensitivity of our simulations to parameter choices. In Figure 12 we show how our results are affected by a 20% decrease and a 20% increase in parameter values for all three levels of HGMCA. The lines plotted here were extracted using the same selection function described in Section 5.3.2 from a number of random seeds. As with the random seed test, there does not appear to be a significant change in performance even when our parameters are all simultaneously changed by a significant percentage. The results we present are also not an outlier; in fact, increasing all the parameter values by 20% seems to slightly improve our performance on this particular simulation.

5.3.4 Foreground Templates

We have shown that our results are robust to different instantiations of the CMB as well as perturbations in parameters and optimization, but we have not yet shown how HGMCA performs under different types of local variation and levels of sparsity in the foreground contaminants. To show this, we employ PySM’s Model 2 templates for dust and synchrotron emissions, which give more sophisticated models of local variation for these two sources. Especially noticeable is the fact that the dust variation is of a smaller scale,




Figure 11. Percent error in the CMB C_ℓ 's as a function of ℓ for varying seeds of each of the algorithms. The transparent blue lines are different random seeds with the same parameter configuration, the red lines are the results for WavILC, and the dark blue lines are the results we present in Section 5. The selection function used to pick the maps that return the dark blue lines is blind to the input CMB; it is explained in detail in section 5.3.2. Note that the dark blue lines (which correspond to our fiducial results) for all four algorithms is not an outlier compared to other random seeds.




Figure 12. Percent error in the CMB C_ℓ 's as a function of ℓ for varying parameter configurations. For all three HGMCA variants we increase and decrease all parameters by 20% to see how it affects the performance of the algorithm. As with the random seeds, the results we present are in line with what other parameter configurations return.




Figure 13. The difference between the kurtosis of the input CMB and the kurtosis of the outputs. None of the five algorithms has a clear advantage over the others, but all five do seem to introduce minimal kurtosis in their residuals. Note that the UT78 *Planck* mask has been applied to the maps (meaning that the regions within the mask were not used for the kurtosis calculation).

and cannot be captured until deeper levels of subdivision. We therefore expect not to see a noticeable improvement between $\text{GMCA}^{\text{LASSO}}$ and lower-level HGMCA. In addition, Model 2 changes the L1-norm (relaxed sparsity) of these sources in the wavelet space by 15 – 20%. To show the robustness of HGMCA to these changes, we continue to use the parameters trained on Model 1 blind to the new templates. However, because of the smaller scale variation in dust, we also present results on HGMCA at level 4 (HGMCA-4).

Table 2 shows a comparison between the performance of HGMCA, $\text{GMCA}^{\text{LASSO}}$, and WavILC on these simulations with respect to RMSE and difference in power spectra between the input and output maps. Table 3 compares the correlation of each algorithm’s residual error (the difference of input and output CMB) to each foreground source. Detailed figures of these comparisons can be found in Appendix H. All error metrics show that HGMCA-1 and HGMCA-2 struggle with the smaller scale local variation of model-2, and underperform when compared even to $\text{GMCA}^{\text{LASSO}}$. While HGMCA-1 and HGMCA-2 must have optima that match $\text{GMCA}^{\text{LASSO}}$, with little local variation to take advantage of at levels 1 and 2, it is expected that the simpler optimization space of $\text{GMCA}^{\text{LASSO}}$ will provide a better solution in practice. However, HGMCA-3 and HGMCA-4 are at a depth at which it is possible to take advantage of local variation and show a marked improvement across almost all error metrics over HGMCA-1, HGMCA-2, and $\text{GMCA}^{\text{LASSO}}$. Overall, HGMCA-4 and WavILC perform comparably, with WavILC providing better RMSE and C_ℓ ’s for $\ell > 200$, and HGMCA-4 providing lower residual contamination and C_ℓ error for low ℓ .

5.4 Non-Gaussianity

The input simulation we described in Section 4.4 treats the CMB as a Gaussian random field, so a perfect reconstruction would return a map with no non-Gaussianities. In this framework, we cannot fully test HGMCA’s ability to construct maps useful for non-Gaussianity studies, but we can probe the amount of non-Gaussianity introduced by our method. In Figure 13 we compare the change in kurtosis (the fourth standardized moment) between the input CMB maps and the reconstructed map for each algorithm along differ-

ent latitudes. This comparison was inspired by the analysis done in Bobin et al. (2013). Within the context of an input map without non-Gaussianity, none of the four algorithms appear to introduce substantial kurtosis from their residuals. Conducting an identical test with skewness rather than kurtosis gives similar results.

6 CONCLUSIONS

We have presented HGMCA, a hierarchical, Bayesian approach to Blind Source Separation for the CMB. The algorithm has been built to tackle one of the principal challenges in precision measurements of ISW signals, gravitational lensing, primordial non-Gaussianity, constraints on isotropy, topological defects, global geometry, etc: separating the CMB from its foregrounds at the map level. Work by Bobin et al. (2008) has shown that a sparsity-based BSS approach like GMCA can be effective in removing these contaminants. However, vanilla GMCA requires a fixed mixing matrix across the sky and therefore does not have the flexibility to capture the known local variation in the frequency dependence of foregrounds. To address this shortcoming, Bobin et al. (2013) introduced LGMCA. LGMCA subdivides the sky into a series of smaller patches, and independently runs GMCA across each of them. As a result, each patch’s frequency dependence is allowed to vary, but it is also blind to the constraining power of the rest of the sky. Further, since the same data must be re-used at different scales of subdivision, the strategy does not admit a generative probabilistic model and requires a heuristic method of choosing between multiple reconstructions.

HGMCA allows for local variation by decomposing the sky into small patches but, unlike LGMCA, it maintains global consistency by enforcing a hierarchy of mixing matrix priors. Further, by dividing the data by scale rather than repeating it at each level of subdivision, HGMCA admits a generative probabilistic model.

We have provided rigorous testing of HGMCA levels 1, 2, and 3 across simulations generated by PySM at $N_{\text{side}} = 256$. Our maps do not vary beam size as a function of frequency and include four different foreground components: dust, synchrotron radiation, free-free emission and anomalous microwave emission. We have shown that, given the Model 1 PySM templates, HGMCA provides cleaner CMB reconstructions than either GMCA or state-of-the-art Internal Linear Combination-type algorithms, with an order of magnitude less correlation to foreground sources, a $\sim 35\%$ (GMCA) / $\sim 75\%$ (WavILC) lower cross correlation error for all ℓ and $\sim 75\% / \sim 35\%$ for $\ell > 200$, and a $\sim 20\% / \sim 45\%$ reduction in RMSE to the input CMB. When extended to templates with much smaller scale local variation, HGMCA at deeper levels continues to improve over GMCA and performs competitively with WavILC.

Beyond demonstrating that HGMCA incorporates local variation effectively, we have also explored some of the generalization concerns related to sparsity based methods. GMCA and HGMCA require setting hyperparameters which, while theoretically motivated, are set by tuning to simulations. Since even the most advanced simulations will not perfectly reproduce the conditions of the data, it is a natural concern that GMCA and HGMCA may not generalize well to the

Algorithm	RMSE [μK]	RMSE _{UT78} [μK]	% Error CMB C_ℓ 's, all ℓ 's	% Error CMB C'_ℓ 's, $\ell > 200$
GMCA ^{LASSO}	23.4081	4.5756	81.7473	0.5335
WavILC	6.1665	1.4835	8.7473	0.04399
HGMCA - 1	35.0097	8.9987	382.5847	0.2703
HGMCA - 2	33.5835	8.8376	367.3803	0.2543
HGMCA - 3	22.5901	2.4231	10.9210	0.0919
HGMCA - 4	23.7845	1.6598	0.2706	0.0844

Table 2. Comparison of RMSE (μK) and average percent error in the CMB angular power spectrum including and excluding the UT78 mask for WavILC, GMCA^{LASSO}, HGMCA at several levels for sky maps with model 2 dust and synchrotron emissions. GMCA^{LASSO} and HGMCA maps are selected via the input CMB blind criterion described in Section 5.3.2. WavILC provides the construction with the lowest RMSE and C_ℓ for high ℓ , with HGMCA-4 offering a marked improvement at low ℓ . HGMCA levels 1 and 2 underperform across all metrics due to the smaller scale local variation of model-2. See Figure H1 for a plot across all ℓ .

Algorithm	$ C_\ell^{\text{Residual,AME}} $	$ C_\ell^{\text{Residual,Free-Free}} $	$ C_\ell^{\text{Residual,Synchrotron}} $	$ C_\ell^{\text{Residual,Dust}} $
GMCA ^{LASSO}	0.0024	0.1519	0.0702	0.1481
WavILC	0.0006	0.0437	0.0184	0.5459
HGMCA - 1	0.0051	0.3234	0.1474	1.1613
HGMCA - 2	0.0050	0.3165	0.1446	1.1380
HGMCA - 3	0.0009	0.0589	0.0283	0.2141
HGMCA - 4	0.0002	0.0107	0.0056	0.0369

Table 3. Comparison of correlation of residual CMB map of WavILC, GMCA^{LASSO}, and HGMCA to foreground contaminants. HGMCA levels 1 and 2 continue to underperform due to the small scale local variation of model-2. HGMCA-3 and WavILC have comparable residual contamination while HGMCA-4 improves in every category by at least a factor of three (and a full order of magnitude for dust contamination). See Figure H2 for plots across all ℓ .

observed sky. We address this concern by presenting results strictly on maps with different CMB realizations than in our training, and by presenting results on maps with different foreground templates for dust and synchrotron emissions. Further, we show that HGMCA is not particularly sensitive to the exact values of these hyperparameters.

In addition to demonstrating HGMCA's robustness to changes in the CMB, foreground, and hyperparameters, we have also addressed concerns about convergence to a stable CMB reconstruction. The publicly available implementation of GMCA includes a threshold parameter for convergence. In our experience, we found that both this parameter and the order of optimization had a strong effect on the CMB reconstruction even after many iterations. To address this we have introduced GMCA^{LASSO}, and shown that both it and HGMCA converge to a stable solution which depends only minimally on the order of optimization.

Despite these extensive tests, it is still an open question how the algorithm will behave when faced with the additional complexities of higher fidelity simulations. The algorithms on which WavILC is based have been shown to work well even when the data involves effects like varying beam size and CIB contamination, while HGMCA is still untested in these conditions. However, higher resolution will also allow for deeper levels of subdivision, potentially enabling HGMCA to further capitalize on its ability to capture local variations in the foregrounds.

HGMCA's rigorous probabilistic structure, favorable comparison to GMCA and WavILC, and robustness to changes in data suggest it is a well-suited method for use in current and upcoming CMB experiments. As a hierarchical framework, HGMCA also opens the door to incorporating local variation in other BSS-based algorithms with rigorous probabilistic structure. We plan to make the HGMCA code and the CMB reconstruction of the *Planck* dataset available in the near future. Additionally, by using directional spin wavelets we should be able to extend HGMCA to foreground separation on polarization maps (as an example of this, see Rogers et al. (2016b)). We leave this for future work.

REFERENCES

- Antoine J.-P., Vanderghenst P., 1999, *Applied and Computational Harmonic Analysis*, 7, 262
- Baldi P., Kerkycharian G., Marinucci D., Picard D., et al., 2009, *The Annals of Statistics*, 37, 1150
- Basak S., Delabrouille J., 2011, *Monthly Notices of the Royal Astronomical Society*, 419, 1163
- Basak S., Delabrouille J., 2013, *Monthly Notices of the Royal Astronomical Society*, 435, 18
- Bedini L., Herranz D., Salerno E., Baccigalupi C., Kuruoğlu E., Tonazzini A., 2005, *EURASIP Journal on Advances in Signal Processing*, 2005, 190845
- Bennett C. L., et al., 2013, *The Astrophysical Journal Supplement Series*, 208, 20
- Bobin J., Starck J.-L., Moudden Y., Fadili M. J., 2008, *Advances in Imaging and Electron Physics*, 152, 221
- Bobin J., Starck J.-L., Sureau F., Basak S., 2013, *Astronomy & Astrophysics*, 550, A73
- Bobin J., Sureau F., Starck J.-L., Rassat A., Paykari P., 2014, *Astronomy & Astrophysics*, 563, A105
- Chan J. Y., Leistedt B., Kitching T. D., McEwen J. D., 2016, *IEEE Transactions on Signal Processing*, 65, 5
- Delabrouille J., Cardoso J.-F., 2008, in *Data Analysis in Cosmology*. Springer, pp 159–205, doi:10.1007/978-3-540-44767-2_6
- Delabrouille J. C. L. J. B. F., Guilloix 2009, *Astronomy & Astrophysics*, 493, 835
- Delabrouille J., Cardoso J.-F., Patanchon G., 2003, *Monthly Notices of the Royal Astronomical Society*, 346, 1089
- Dunkley J., et al., 2009, *The Astrophysical Journal*, 701, 1804
- Eriksen H. K., et al., 2004, *The Astrophysical Journal Supplement Series*, 155, 227
- Fuskeland U., Wehus I. K., Eriksen H., Naess S. K., 2014, *The Astrophysical Journal*, 790, 104
- Gelman A., Carlin J. B., Stern H. S., Dunson D. B., Vehtari A., Rubin D. B., 2013, *Bayesian data analysis*. Chapman and Hall/CRC
- Geman S., Geman D., 1984, *IEEE Transactions on pattern analysis and machine intelligence*, pp 721–741
- Górski K., Banday A., Bennett C., Hinshaw G., Kogut A., Smoot G., Wright E., 1996, *The Astrophysical Journal Letters*, 464, L11
- Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelman M., 2005, *Astrophys. J.*, 622, 759
- Hobson M., Jones A., Lasenby A., Bouchet F., 1998, *Monthly Notices of the Royal Astronomical Society*, 300, 1
- Hyvärinen A., Oja E., 2000, *Neural networks*, 13, 411
- Jaffe T., Banday A., Eriksen H., Górski K., Hansen F., 2006, *The Astrophysical Journal*, 643, 616
- Land K., Magueijo J., 2006, *Monthly Notices of the Royal Astronomical Society*, 367, 1714
- Leach S. M., et al., 2008, *Astronomy & Astrophysics*, 491, 597
- Leistedt B., McEwen J. D., Vanderghenst P., Wiaux Y., 2013, *A&A*, 558, A128
- Martínez-González E., Diego J., Vielva P., Silk J., 2003, *Monthly Notices of the Royal Astronomical Society*, 345, 1101
- McEwen J. D., 2015, arXiv preprint, arXiv:1510.01595
- McEwen J. D., Wiaux Y., 2011, *IEEE Transactions on Signal Processing*, 59, 5876
- McEwen J. D., Büttner M., Leistedt B., Peiris H. V., Wiaux Y., 2015a, *IEEE Signal Processing Letters*, 22, 2425
- McEwen J. D., Leistedt B., Büttner M., Peiris H. V., Wiaux Y., 2015b, arXiv preprint, arXiv:1509.06749
- McEwen J. D., Feeney S. M., Peiris H. V., Wiaux Y., Ringeval C., Bouchet F. R., 2017, *Monthly Notices of the Royal Astronomical Society*, 472, 4081
- McEwen J. D., Durastanti C., Wiaux Y., 2018, *Applied and Computational Harmonic Analysis*, 44, 59
- Moudden Y., Cardoso J.-F., Starck J.-L., Delabrouille J., 2005, *EURASIP Journal on Advances in Signal Processing*, 2005, 484
- Narcowich F. J., Petrushev P., Ward J. D., 2006, *SIAM Journal on Mathematical Analysis*, 38, 574
- Pham D. T., Garat P., 1997, *IEEE transactions on Signal Processing*, 45, 1712
- Planck Collaboration I 2014, *Astronomy & Astrophysics*, 571, A1
- Planck Collaboration I 2018, arXiv preprint, arXiv:1807.06205
- Planck Collaboration IV 2018, arXiv preprint, arXiv:1807.06208
- Planck Collaboration IX 2019, arXiv preprint, arXiv:1905.05697
- Planck Collaboration VI 2018, arXiv preprint, arXiv:1807.06209
- Planck Collaboration VII 2019, arXiv preprint, arXiv:1906.02552
- Planck Collaboration VIII 2018, arXiv preprint, arXiv:1807.06210
- Planck Collaboration X 2016, *Astronomy & Astrophysics*, 594, A10
- Planck Collaboration XI 2016, *Astronomy & Astrophysics*, 594, A11
- Planck Collaboration XII 2014, *Astronomy & Astrophysics*, 571, A12
- Planck Collaboration XII 2016, *Astronomy & Astrophysics*, 594, A12
- Planck Collaboration XIII 2016, *Astronomy & Astrophysics*, 594, A13
- Planck Collaboration XIX 2014, *Astronomy & Astrophysics*, 571, A19
- Planck Collaboration XVI 2014, *Astronomy & Astrophysics*, 571, A16
- Rapin J., Bobin J., Larue A., Starck J.-L., 2013, *IEEE Transactions on Signal Processing*, 61, 5620
- Rogers K. K., Peiris H. V., Leistedt B., McEwen J. D., Pontzen A., 2016a, *Monthly Notices of the Royal Astronomical Society*, 460, 3014
- Rogers K. K., Peiris H. V., Leistedt B., McEwen J. D., Pontzen A., 2016b, *Monthly Notices of the Royal Astronomical Society*, 463, 2310
- Saaddeh D., Feeney S. M., Pontzen A., Peiris H. V., McEwen J. D., 2016, *Physical review letters*, 117, 131302
- Santosa F., Symes W. W., 1986, *SIAM Journal on Scientific and Statistical Computing*, 7, 1307
- Seljebotn D. S., Mardal K.-A., Jewell J., Eriksen H., Bull P., 2014, *The Astrophysical Journal Supplement Series*, 210, 24
- Tegmark M., de Oliveira-Costa A., Hamilton A. J., 2003, *Physical Review D*, 68, 123523
- Thorne B., Dunkley J., Alonso D., Naess S., 2017, *Monthly Notices of the Royal Astronomical Society*, 469, 2821
- Vio R., Andreani P., 2008, arXiv preprint, arXiv:0811.4277
- Wiaux Y., McEwen J. D., Vanderghenst P., Blanc O., 2008, *Monthly Notices of the Royal Astronomical Society*, 388, 770

ACKNOWLEDGMENTS

We thank David Alonso, Jason McEwen, Hiranya Peiris, Andrew Pontzen, and Keir Rogers for useful discussions. We would like to thank Benjamin Racine for his help with the simulations. We also thank Jerome Bobin and Jean-Luc Starck for helpful correspondence. We especially thank Boris Leistedt for various contributions to this project. CD was supported by NSF grant AST-1813694 and Department of Energy (DOE) grant DE-SC0020223. CD was also supported by the Star Family Challenge for Promising Scientific Research and the Dean’s Competitive Fund for Promising Scholarship at Harvard University. MH was supported

through NSF Award DGE-1650112. SWC was supported through the KIPAC-Chabolla fellowship.

APPENDIX A: CMB SOURCE SEPARATION TECHNIQUES

Here we highlight the four models used most recently by the *Planck* collaboration, SEVEM, SMICA, Commander, and NILC (Leach et al. 2008; Pham & Garat 1997; Seljebotn et al. 2014; Delabrouille & Guilloux 2009), as well as providing some general background on template fitting and internal linear combination.

A1 Template Fitting

In template fitting, Equation (1) is simplified to

$$d = \tilde{d} + \alpha t, \quad (\text{A1})$$

where the sky d is a linear combination of the CMB $\tilde{d}$, and a template t mapping anisotropies (Land & Magueijo 2006). To isolate the CMB, one must construct the template t and solve for α . A standard method is to solve for the maximum likelihood value of α by minimizing χ^2 . This technique can be extended to add multiple templates for different foreground components (Leach et al. 2008; Jaffe et al. 2006).

Template fitting techniques can be either parametric (using external information to parametrize the components of the sky) or blind. Parametric techniques develop templates iteratively, comparing results with external data and theory to develop new templates.

A2 Internal Linear Combination

Internal Linear Combination (ILC) is another technique that focuses solely on recovering the CMB, but relies on the assumption that the CMB signal is independent of frequency rather than using foreground templates. ILC observes that for any affine combination¹⁶ of input maps X_i :

$$\sum w_i X_i = C + \sum_i w_i (f_i + n_i), \quad (\text{A2})$$

where C is the CMB, f_i is foreground signal, and n_i is noise. Under the assumption that the CMB is independent from the foreground and noise, it is then possible to minimize the foreground contamination by finding the set of weights w_i that minimizes the variance of $\sum w_i X_i$, effectively setting the rightmost term of Equation (A2) to 0. A derivation of the formula used to calculate the weights can be found in Vio & Andreani (2008).

ILC's only major assumption is the frequency dependence of the CMB. It can operate in pixel, harmonic, or wavelet space, and can be solved for analytically. The reconstruction quality and noise properties are also well-understood theoretically (Vio & Andreani 2008). However, it does not tap into the extra information that can be gained by making more advanced assumptions about the foregrounds, such as their sparsity in the wavelet basis, which can enable more accurate reconstructions.

¹⁶ An affine combination is a linear combination with coefficients that sum to 1

A3 SEVEM

SEVEM is an internal template fitting technique currently used by the *Planck* collaboration. SEVEM creates templates without external data by subtracting *Planck* maps at close frequencies. Since the CMB is the same across frequencies, this isolates foreground components. SEVEM's template fitting is done in real space at every position in the sky. Given n such templates t_i , a position x and frequency ν the CMB is then given by the difference from the original map $d(x, \nu)$ and a linear combination of templates:

$$T_{CMB}(x, \nu) = d(x, \nu) - \sum_{i=1}^n \alpha_i(\nu) t_i(\nu).$$

No assumptions on foreground or noise are used in SEVEM, making it a robust technique (Leach et al. 2008). *Planck Collaboration X* (2016) suggests SEVEM is best used to deal with bright point sources in polarization, and notes that for a number of frequencies, it is the only technique that can provide independent CMB reconstructions.

A4 NILC

ILC can be performed either in harmonic or pixel space, but fails to take into account that noise and foreground dominate at different angular frequencies in harmonic space, and different Galactic latitudes in pixel space (Basak & Delabrouille 2011). In order to deal with this, Delabrouille & Guilloix (2009) introduced Needlet Internal Linear Combination (NILC), an extension of ILC to wavelet space. NILC uses a specific type of wavelet known as the needlet (Baldi et al. 2009) which avoids ILC's issue by having compact support in the harmonic domain along with localization in the pixel domain (Basak & Delabrouille 2011). More recently, Rogers et al. (2016a) proposed Scale-discretized, directional wavelet ILC (SILC). By using non-axisymmetric wavelets, SILC takes advantage of orientation in signal structure to further fine tune the weights of ILC. *Planck Collaboration X* (2016) note that NILC works well at high ℓ , with performance comparable to SMICA.

The covariance matrix used for NILC and SILC (as well as WavILC in this work) is empirically calculated by convolution with a Gaussian beam. First an approximate covariance matrix C_{ij}^{approx} is calculated by:

$$C_{ijk}^{\mu, \text{approx}} = x_{ik}^{\mu} x_{jk}^{\mu} \quad (\text{A3})$$

where i and j are two frequency bands, k is a pixel coordinate, and μ is the wavelet scale. The convolution is then calculated in the harmonic space as:

$$\tilde{C}_{ij\ell m}^{\mu} = \sum_{\ell=0}^{2\ell_{\text{max}}} \sum_{m=-\ell}^{\ell} g_{\ell} \tilde{C}_{ij\ell m}^{\mu, \text{approx}} \quad (\text{A4})$$

The final weights are then calculated as:

$$w_{ik}^{\mu} = \frac{\sum_{j=1}^{N_{\text{freq}}} C_{ijk}^{\mu} e_j}{\sum_{j'=1}^{N_{\text{freq}}} \sum_{j=1}^{N_{\text{freq}}} e_{j'} C_{j'jk}^{\mu} e_j} \quad (\text{A5})$$

where e is a vector of ones.

A5 SMICA

Independent Component Analysis (ICA) is a popular BSS method in the machine learning community (Hyvärinen & Oja 2000). ICA relies on two major assumptions: that the sources we wish to separate are both independent and non-Gaussian. Let W be the psuedo-inverse of A . Then a single source s_i may be expressed as

$$s_i = w_i X. \quad (\text{A6})$$

The basic concept of ICA is to compute w_i , and hence s_i , by maximizing the non-Gaussianity of $w_i X$ via measures such as kurtosis or negentropy. Further, since sources are assumed to be uncorrelated, optimizing over all sources is not difficult (Hyvärinen & Oja 2000).

While the foreground elements and CMB may be independent, the CMB is Gaussian. To avoid this issue, Pham & Garat (1997) introduced a modification of the method called spectral matching ICA (SMICA). SMICA fixes the issue of the CMB's gaussianity by trading non-Gaussianity for spectral distinction between sources. SMICA follows essentially the approach of ICA described above, but instead computes S and A via maximizing a spectral matching criterion exhibited by the CMB and foreground sources (Delabrouille & Cardoso 2008).

As initially proposed, SMICA analyzed spectral matching in the Fourier basis. This method is sub-optimal for capturing non-stationary foreground contaminants, even when using localized spectral statistics (Delabrouille et al. 2003). Moudén et al. (2005) propose solving this issue by using a wavelet rather than Fourier basis. Wavelets are better suited to deal with non-stationary components due to their locality in both pixel and harmonic space. *Planck Collaboration X* (2016) suggest using SMICA for mid to high ℓ .

A6 Commander

Commander is a Bayesian, forward parametric technique with the goal of computing the joint posterior distribution (Eriksen et al. 2004), $P(A, S|X)$. Commander is realized via Gibbs sampling, a statistical method for sampling joint distributions from iteratively sampling conditional distributions of A and S . In particular, A and S are updated by repeatedly sampling from:

$$A \sim P(A|S, X), \quad (\text{A7})$$

$$S \sim P(S|A, X). \quad (\text{A8})$$

Once the Gibbs sampling has converged and the posterior distribution appropriately sampled, one can compute desired empirical statistics for A and S , such as the mean and standard deviation maps for each component.

Originally, this technique only succeeded in producing low resolution maps, but recent incarnations such as Commander-Ruler (*Planck Collaboration XVI* 2014) and Commander 2 (*Planck Collaboration VI* 2018) fix this issue. *Planck Collaboration X* (2016) and *Planck Collaboration IV* (2018) recommend using the Commander framework over other techniques at low ℓ due to its lower "internal tension between low, intermediate, and high multipoles".

APPENDIX B: SCALE-DISCRETIZED WAVELET TRANSFORM AND PATCH DECOMPOSITION

The wavelet transform of a square integrable signal f on the two-dimensional sphere is a decomposition into wavelet coefficient maps $W^{\Psi^j}(\omega)$, themselves signals on the sphere, indexed with the wavelet scale j and calculated as

$$W^{\Psi^j}(\omega) = \int_{S^2} d\Omega(\omega') f(\omega') (\mathcal{R}_\omega \Psi^j)^*(\omega'). \quad (\text{B1})$$

Reconstruction of the input signal f is achieved through

$$f(\omega) = \sum_j \int_{S^2} d\Omega(\omega') W^{\Psi^j}(\omega') (\mathcal{R}_{\omega'} \Psi^j)(\omega). \quad (\text{B2})$$

Those operations are interpreted as convolutions on the sphere, relying on a rotation operator on the sphere $\mathcal{R}_\omega$. The transform is fully characterized by the choice of set of wavelets $\Psi^j(\omega)$, which are also signals on the sphere, typically compact and localized on the north pole. Some desirable properties of the transform include 1) invertibility, which amounts to the decomposition and reconstruction operations being lossless, 2) good localization properties of the wavelets in both pixel and harmonic space, 3) the computational cost of performing the convolutions above.

We adopt the framework of scale-discretized wavelets (McEwen et al. 2015b; Leistedt et al. 2013; Wiaux et al. 2008), as implemented in the S2LET code¹⁷. Central to this approach is the use of the SSHT sampling theorem (McEwen & Wiaux 2011), to convert signals on the sphere (and stored on the SSHT pixelization) into spherical harmonic coefficients, and conversely, in a theoretically exact fashion (*i.e.*, lossless for bandlimited signals). This can be exploited to perform the convolutions in harmonic space, which is fast and exact since it reduces to a discrete sum over harmonic coefficients. As a result, the wavelet transform itself is lossless and invertible, which in practice leads to errors of the order of floating point accuracy when transforming a signal back and forth. Furthermore, scale-discretized wavelets are based on wavelets $\Psi^j(\omega)$ that are simply defined in harmonic space and exactly compact, which leads to further computational improvements over existing approaches making use of wavelets with broad harmonic support (McEwen et al. 2015b; Leistedt et al. 2013; McEwen et al. 2018).

Note that in the concise presentation above we have omitted the scaling function, which is akin to an extra wavelet but captures the low-frequency, low-resolution information of the signal, typically not included in $\Psi^j(\omega)$. A full presentation can be found in (Leistedt et al. 2013). We also focused on axisymmetric wavelets, but scale-discretized wavelet can be made directional (McEwen et al. 2015b).

The main parameters of the wavelet transform are the band limit of the signals (for all spherical harmonic transform operations) and the details of the wavelets. Axisymmetric scale-discretized wavelets are characterized by a parameter λ . Each wavelet Ψ^j has a compact harmonic-space support in $[\lambda^{j-1}, \lambda^{j+1}]$, with the peak at λ^j . For the results in this work we used $\lambda = 3$, a bandlimit of $\ell = 621$, and a minimum $J_0 = 1$.

In our BSS formalism, the linearity of the wavelet transform allows us to write the signal X as a matrix product AS in the wavelet basis, just like we would do in pixel or harmonic space. The elements of X and the rows of the source matrix S would be wavelet coefficients, pixels, or harmonic coefficients, respectively. LGMCA and HGMCA use a decomposition of the map X into patches. This decomposition is done in the wavelet space (*i.e.* where X and S are wavelet coefficients) to avoid the computational cost of repeatedly conducting a forward and inverse wavelet transform. Because the patches are never allowed to probe scales near the wavelet support, these patches should have a close correspondence to the same patch in the original space. This decomposition also has the added benefit of smoothing out the harsh boundaries that would occur in a pixel-based decomposition.

APPENDIX C: NON-NEGATIVITY CONSTRAINTS FOR GMCA

For applications of GMCA, LGMCA, and HGMCA to CMB foreground separation our mixing matrix A cannot be negative. It is a representation of the contribution of a given source at a given frequency, and cosmological sources cannot contribute negative flux. With this in mind, Rapin et al. (2013) introduce a non-negative version of GMCA. Applying a prior to each column of A which is uniform over the positive orthant of the unit sphere¹⁸, the optimization problem becomes

$$\text{argmin}_{A,S} \lambda_S \|S\|_0 + \|X - AS\|_F^2 + i^+(A), \quad (\text{C1})$$

where $i^+(A)$ is the characteristic function of the positive orthant,

$$i^+(A) = \begin{cases} 0 & A \geq 0 \\ \infty & A < 0. \end{cases} \quad (\text{C2})$$

Rapin et al. (2013) present several optimization techniques for this problem, and empirically show that for the case of the sky where there are few sources to separate, it is the naive algorithm of setting any non-negative values to 0 which performs best. This same naive modification can be made to both LGMCA and HGMCA to enforce non-negativity.

APPENDIX D: HGMCA UPDATE RULES

We start by writing out the full loss function for HGMCA,

$$\begin{aligned} \mathcal{L} = & \underbrace{\sum_l \sum_{\mu_l} \sum_{p \in \{p_l\}} i^+(A_p^l) + \lambda_S \|S_p^{l,\mu_l}\|_1 + \|X_p^{l,\mu_l} - A_p^l S_p^{l,\mu_l}\|_F^2}_{\text{source prior and reconstruction error}} \\ & + \underbrace{\sum_l \sum_{p \in \{p_{l-1}\}} \sum_{i \in \{1, \dots, 4\}} \lambda_A \|A_p^l - A_{p,i}^l\|_F^2}_{\text{A matrix dependence across levels}}. \end{aligned} \quad (\text{D1})$$

¹⁸ This also accounts for the L2-normalization of A , and thus provides a Bayesian interpretation of (Rapin et al. 2013)

¹⁷ <https://github.com/astro-informatics/s2let>

The notation l denotes the level, μ_l denotes the wavelet scales with maximum subdivision level l , and $p \in \{p_l\}$ indicates the patches at subdivision level l (for example $p \in \{p_2\}$ are $p = p_1, p_2$ with $p_1, p_2 \in \{1, 2, 3, 4\}$). The norm $\|M\|_F^2$ is the Frobenius or 2-norm and $\|M\|_1$ is the 1-norm of a matrix M .

Since all the wavelet scales at a single level share the same mixing matrix, we concatenate the different μ_l scales together so that the source matrix S_p^l at subdivision level l on a patch $p \in \{p_l\}$. Therefore, S_p^l now has dimensions of N_S by (number of pixels times number of scales in μ_l), allowing us to drop the wavelet superscript. Similarly, for each level the dimensions of the data matrix X_p^l will be number of frequencies by (number of pixels times number of scales in μ_l).

We are interested in knowing the update rule for a specific row of a source matrix S_p^l or column of a mixing matrix A_p^l . We can therefore also drop the subscripts and superscripts related to the level and particular patches for simplicity. In the following, we denote s_i as the i^{th} row of S_p^l corresponding to source i and a^i as the i^{th} column of the mixing matrix A_p^l on a patch p corresponding to source i . Dropping all the terms that have no dependence on s_i we obtain the following loss function,

$$\mathcal{L}_{s_i} = \lambda_S \|s_i\|_1 + \|X - \sum_j a^j s_j\|_F^2, \quad (\text{D2})$$

which has contributions from our sparsity prior and the probability of generating X_p^l from A_p^l and S_p^l . For all $j \neq i$ the value of s_j is fixed, so we can further simplify the above loss function by introducing a new term $R_i = X - \sum_{j \neq i} a^j s_j$,

$$\begin{aligned} \mathcal{L}_{s_i} &= \lambda_S \|s_i\|_1 + \|R_i - a^i s_i\|_F^2 \\ &= \lambda_S \|s_i\|_1 + \text{Tr} \left[(R_i - a^i s_i)^T (R_i - a^i s_i) \right] \\ &= \lambda_S \|s_i\|_1 + \text{Tr} \left[R_i^T R_i - R_i^T a^i s_i - s_i^T (a^i)^T R_i + s_i^T (a^i)^T a^i s_i \right] \\ &= \lambda_S \|s_i\|_1 + \text{Tr} \left[s_i^T (a^i)^T a^i s_i - 2s_i^T (a^i)^T R_i \right], \end{aligned} \quad (\text{D3})$$

where we have dropped the $R_i^T R_i$ since it has no s_i dependence and taken advantage of the invariance of trace under transpositions.

We can now proceed to find the parameter s_i that minimizes the loss function,

$$0 = \frac{\partial \mathcal{L}_{s_i}}{\partial s_i} = \frac{\partial}{\partial s_i} \lambda_S \|s_i\|_1 + 2(a^i)^T a^i s_i - 2(a^i)^T R_i. \quad (\text{D4})$$

Note that $(a^i)^T a^i$ is simply the square of the 2-norm of the i^{th} column of A . Due to the degeneracy between scaling a column of A and a row of S by inverse factors it is common practice in the BSS literature to set the 2-norm of A to 1. With this additional simplification we are left with

$$s_i = (a^i)^T R_i - \frac{\partial}{\partial s_i} \frac{\lambda_S}{2} \|s_i\|_1. \quad (\text{D5})$$

We solve this equation through the LASSO shooting algorithm, which gives the following solution in matrix form (Santosa & Symes 1986),

$$s_i = \begin{cases} (a^i)^T R_i - \frac{\lambda_S}{2} \text{sign}((a^i)^T R_i) & |(a^i)^T R_i| > \frac{\lambda_S}{2} \\ 0 & |(a^i)^T R_i| \leq \frac{\lambda_S}{2} \end{cases}, \quad (\text{D6})$$

with sign being an element-wise sign function that returns the sign of the number or 0 if the number is 0.

We then have to obtain an analogous optimization equation for a^i . We use $\tilde{A}$ to denote the mixing matrix that defines the distribution from which A is drawn. As an example, for $A = A_{p_1}^1$ this would be $\tilde{A} = A^0$. We use B_j to denote the j^{th} matrix that is drawn from A . As an example, for $A = A_{p_1}^1$ this would be $B_j = A_{p_1, j}^2$ for $j \in \{1, 2, 3, 4\}$. We can now write the full loss function dropping all terms that have no dependence on a^i . Additionally, we add a Lagrange multiplier term λ_m that enforces that the two-norm of a column of A is 1,

$$\begin{aligned} \mathcal{L}_{a^i} &= \|R_i - a^i s_i\|_F^2 + \lambda_A \|A - \tilde{A}\|_F^2 + \sum_j \lambda_A \|A - B_j\|_F^2 \\ &\quad + \lambda_m (\|A\|_F^2 - 1) \\ &= \text{Tr} \left[(a^i)^T a^i s_i s_i^T - 2(a^i)^T R_i s_i^T \right] \end{aligned} \quad (\text{D7})$$

$$\begin{aligned} &+ \lambda_A \text{Tr} \left[(a^i)^T a^i - 2(a^i)^T \tilde{a}^i \right] \\ &+ \sum_j \lambda_A \text{Tr} \left[(a^i)^T a^i - 2(a^i)^T b_j^i \right] + \lambda_m \text{Tr} \left[(a^i)^T a^i \right]. \end{aligned} \quad (\text{D8})$$

This includes terms from the mixing matrix hierarchy prior as well the generative distribution of X_p^l . We can then find the parameter a^i that minimizes the loss function,

$$\begin{aligned} \frac{\partial \mathcal{L}_{a^i}}{\partial a^i} = 0 &= 2a^i \|s_i\|^2 - 2R_i s_i^T + 2\lambda_A \left((1 + N_j)a^i - \tilde{a}^i - \sum_j b_j^i \right) \\ &\quad + 2\lambda_m a^i \end{aligned} \quad (\text{D9})$$

$$a^i = \frac{1}{\|s_i\|^2 + \lambda_A(1 + N_j) - \lambda_m} \left[R_i s_i^T + \lambda_A \tilde{a}^i + \lambda_A \sum_j b_j^i \right], \quad (\text{D10})$$

where N_j is the number of matrices drawn from the distribution defined by A . Note that this is the most general solution for a column of the mixing matrix. We can then solve for λ_m by plugging in $\frac{\partial \mathcal{L}}{\partial \lambda_m} = 0 = \text{Tr}[(a^i)^T a^i] - 1$, which yields

$$\begin{aligned} (a^i)^T \frac{\partial \mathcal{L}_{a^i}}{\partial a^i} = 0 &= 2\|s_i\|^2 - 2(a^i)^T R_i s_i^T + 2\lambda_m \\ &\quad + 2\lambda_A \left((1 + N_j) - (a^i)^T \tilde{a}^i - (a^i)^T \sum_j b_j^i \right) \end{aligned} \quad (\text{D11})$$

$$\begin{aligned} \lambda_m &= -\|s_i\|^2 + (a^i)^T R_i s_i^T \\ &\quad - \lambda_A \left((1 + N_j) - (a^i)^T \tilde{a}^i - (a^i)^T \sum_j b_j^i \right). \end{aligned} \quad (\text{D12})$$

Let $[x]_+$ denote $\max(0, x)$, then plugging this back into our original equation and enforcing non-negativity we obtain

$$a^i = \left[\frac{R_i s_i^T + \lambda_A \tilde{a}^i + \lambda_A \sum_j b_j^i}{\|R_i s_i^T + \lambda_A \tilde{a}^i + \lambda_A \sum_j b_j^i\|_F} \right]_+, \quad (\text{D13})$$

which is the vector divided by its 2-norm.

It is possible that there are no wavelet scales at the level

Algorithm	GMCA	HGMCA-1/2/3
λ_S	50	50
N_S	5	5
λ_{CMB}	10^{11}	10^{11}
Map Normalization	None	Center
λ_A	N/A	5×10^{10}

Table E1. The parameters for best performing GMCA and HGMCA algorithms. As discussed in Section 5.3, these parameters are chosen with an initial set of simulations by testing which values return the lowest RMSE error. The final results presented throughout the rest of the paper are from a different set of simulations to show that we have not overfit to the initial simulation set.

corresponding to A in which case the solution would look like

$$a^i = \left[\frac{1}{\|\tilde{a}^i + \sum_j b_j^i\|_F} \left[\tilde{a}^i + \sum_j b_j^i \right] \right]_+, \quad (\text{D14})$$

which is the intuitive form we would expect (in the graphical model, an average of all the matrices it shares edges with). We can similarly deal with the case where there are no $\tilde{A}$ or B_j^i by dropping the corresponding terms.

APPENDIX E: GMCA AND HGMCA PARAMETERS

Table E1 describes the parameters used for the GMCA and HGMCA configurations presented in Section 5. Map Normalization refers to an extra pre-processing step where either nothing is done (None) or the region contained within the UT78 mask is scaled by 0.5 (Center). The overall effect on the results is small, but HGMCA appears to perform slightly better with this pre-processing. The parameters of all three algorithms are optimized to minimize RMSE on a set of training simulations. All the results presented in this work are from separate simulations where either the CMB realization is modified (Section 5.1) or all the source realizations are modified (Section 5.3.4)

APPENDIX F: GMCA OPTIMIZATION STRATEGIES

To test the performance and stability of Fast GMCA versus GMCA^{LASSO}, we follow the parameter selection criteria of section 5.3 and test each algorithm across a new instance of the CMB and across multiple random seeds. Our Fast GMCA implementation follows Bobin et al. (2008). This includes using hard-thresholding and manually enforcing that the CMB column match theoretical expectations. Figure F1 shows the best and worst performances of Fast GMCA and GMCA^{LASSO} across these seeds. The results show that while Fast GMCA's performance across C_ℓ 's is competitive with GMCA^{LASSO}, its optimization is significantly more sensitive to random initialization. In particular, the error in C_ℓ 's from Fast GMCA varies by nearly two orders




Figure F1. The variation in CMB C_ℓ error due to changes in the random seed for Fast GMCA and GMCA^{LASSO}. For both algorithms, the hyperparameters are fixed at an optimum value determined by the methodology described in Section 5.3. While both produce competitive best-case results, Fast GMCA has significantly larger variance.

of magnitude, whereas the error in GMCA^{LASSO} is within about a factor of 2 across all multipoles. The same results emerge if the CMB is perturbed and the random seed is held fixed.

APPENDIX G: RECONSTRUCTED MAPS

In this section, we present the reconstructed CMB maps from WavILC, GMCA, HGMCA-1, HGMCA-2, and HGMCA-3 corresponding to the results presented in Section 5.

APPENDIX H: FIGURES FOR TEMPLATE MODEL 2

In this section, we present the plots for the generalization test in Section 5.3.4. Specifically, we show the percent error in the CMB C_ℓ 's in Figure H1 and the magnitude of the cross spectra between the residual CMB and the foregrounds in Figure H2.

APPENDIX I: PROPAGATED FOREGROUND RESIDUALS AND SOURCE DEGENERACY

As an alternative to the cross-correlation coefficients we presented in Section 5, we can also directly propagate through the contributions from the different foreground components to our final CMB fit. If we have a solution for the mixing matrix A_p on patch p , then the final source matrix will be set by:

$$S_p = [A_p]_{\text{pinv}} X, \quad (\text{I1})$$

where $[A_p]_{\text{pinv}}$ is the pseudo-inverse of A_p . Since the CMB and foreground components add linearly to X , we can rewrite this as:

$$S_p = [A_p]_{\text{pinv}} \left(X_{\text{CMB}} + \sum_{\text{foreground}} X_{\text{foreground}} \right). \quad (\text{I2})$$




Figure G1. From top left to bottom: the reconstructed CMB sky for HGMCA-3, HGMCA-2, HGMCA-1, WavILC, and GMCA^{LASSO}.




Figure H1. For the generalization test described in Section 5.3.4, percent error in C_ℓ as a function of multipole moment ℓ for each of the algorithms.




Figure H2. For the generalization test described in Section 5.3.4, magnitude of the cross-spectra between the residual CMB for each algorithm and the input AME, dust, synchrotron, and free-free foreground contaminants.

In a simulation where we have access to each $X_{\text{foreground}}$, we can isolate the contribution to each source from a specific foreground simply by taking:

$$S_{p,\text{foreground}} = [A_p]_{\text{pinv}} X_{\text{foreground}}. \quad (13)$$

Examining the Dust and AME propagated residuals reveals an interesting degeneracy in the source fitting. In Figure 11, we show the propagated AME and Dust residuals for HGMCA-3 and GMCA. Individually, both algorithms show significant residual contributions from dust and AME. However, the HGMCA-3 residuals are anti-correlated - that is they have similar structure but opposite sign. When we add the residuals for both maps, we see that the HGMCA-3 sum cancels most of the large magnitude residuals whereas GMCA does not. Since both the dust and AME maps have similar footprints (they both roughly trace our galaxy) adding a bit more dust and subtracting a bit more AME leads to a similar CMB reconstruction. Thus any algorithm which is blind to these sources, i.e. includes no prior knowledge of frequency dependence, is susceptible to this degeneracy. While this does not create an issue for reconstructing the CMB, such a degeneracy does harm the reconstruction of foreground sources. It would be of interest in future work to explore how incorporating weak foreground priors could help break this degeneracy in the fitting to return more accurate reconstructions of sources besides the CMB.




Figure 11. Residual and sum of residual errors from dust and AME in the CMB reconstructions of GMCA and HGMCA-3. Most contamination for GMCA comes from AME (as seen in the middle right-hand image). While HGMCA-3 initially seems to have high residual error in Dust and AME, these contributions are due to a degeneracy in the blind sources separation of correlated sources. This is evidenced by the cancellation of dust and AME errors (as seen in the lower left-hand image), leading to a more accurate reconstruction than GMCA provides.