

Component Separation for Cosmic Microwave Background Studies



Richard David Paradorn Grumitt

Wolfson College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

September 9, 2020

Acknowledgements

Completing this DPhil has been a challenging undertaking to say the least. This has certainly not happened without a lot of support, and there are a great many to whom I owe my thanks. First, I would like to thank my supervisors, Angela, Mike and David. Your advice, covering topics ranging from cosmology to cookery, has been totally invaluable. I will definitely be back for more extravagant cookery evenings. I would also like to thank all my collaborators for the help they've given in pushing my work over the line, particularly Clive for his help on the NUTS paper.

I was lucky to have some amazing people in my research group. I'll always remember the many evenings spent pontificating in the Lamb & Flag with Jamie, and very much look forward to the return of the non-virtual pub night. To Luke, thank you for everything. Your advice and friendship throughout this DPhil have kept me sane, and I couldn't have done it without you.

Darsh and Zahra, you've been amazing and made those long days in the office all the better. Max, thank you for opening my eyes to the true value of r . To those excellent people from Hilda's, James, Jasper, Sarah, Steph, Joe, Alice and Wilson, thank you for still talking to me. James, you've been absolutely stellar in entertaining my frequent trips down to London. Jasper, I have always seen this as but a temporary pause on the great cowboy collaboration. And to Serena, I am so fortunate to have met someone as brilliant as you. Your companionship and continued tolerance of me has been the greatest comfort.

Finally, I must thank my family. To Mum and Dad, thank you for raising me, guiding me and feeding me. You quite literally made me into the person I am, and I cannot thank you enough. Jessica, Edward and Martina, thank you for accepting who the real MVP is, and just remember that it do be like that sometimes. And to all my family out in Thailand, thank you for being so wonderful over all these years. I can assure you that I will be back for a Sikhoraphum graduation photo-shoot as soon as possible.

Statement of Originality

This is to certify that no part of this thesis has been submitted in support of a degree or other qualification from the University of Oxford, other higher education institutes or any professional body. The work in this thesis is based on the following publications:

1. R. D. P. Grumitt et al. The C-Band All-Sky Survey (C-BASS): Total intensity point-source detection over the northern sky. *MNRAS*, June 2020, doi: 10.1093/mnras/staa1572.
2. Luke Jew and R. D. P. Grumitt. The spectral index of polarized diffuse Galactic emission between 30 and 44 GHz. *MNRAS*, May 2020, doi: 10.1093/mnras/staa1233.
3. R. D. P. Grumitt, Luke Jew and C. Dickinson. Hierarchical Bayesian CMB Component Separation with the No-U-Turn Sampler. *MNRAS*, June 2020, doi: 10.1093/mnras/staa1857.

This is to clarify that the thesis was written by myself, and the work presented is entirely my own with the following exceptions and clarifications:

- **Chapter 2** summarizes the background of the C-BASS project, including the scientific motivation, data reduction and calibration checks. The instrument and data reduction pipeline were largely developed by the C-BASS collaboration prior to my joining. I implemented the polarization angle calibration, calibration checks, and co-wrote the map deconvolution routine with Luke Jew.
- **Chapter 3** is based on Paper 1. I developed the point-source detection code, produced the C-BASS point-source catalogue and wrote the original draft of the paper. The C-BASS collaboration provided comments on the paper and performed validation checks on the C-BASS catalogue.

- **Chapter 4** presents a new analysis of the synchrotron power spectrum using C-BASS data, and summarizes some of the key scientific results from Paper 2. I performed the power spectrum analysis. The C-BASS spectral index map used in the tensor-to-scalar forecasts was produced by M. E. Jones. The core analysis code for Paper 2 was developed by Luke Jew. I contributed to the scientific interpretation of the results, and co-wrote the manuscript.
- **Chapter 5** is based on Paper 3. I developed the component separation code and wrote the original draft of the manuscript. Luke Jew developed the implementation of the mean-shift clustering algorithm. Luke Jew and C. Dickinson provided comments on the manuscript.

This work is the intellectual property of the author.

Abstract

The detection of primordial B -mode polarization in the cosmic microwave background (CMB) remains the major outstanding goal of CMB cosmology. However, current upper limits on the strength of the primordial signal mean that it will be sub-dominant to astrophysical foreground emission at all frequencies, over the entire sky. Detecting primordial B -modes therefore becomes a problem in component separation.

In this thesis I address this problem from two angles. First, I present my work analysing data from the C-Band All-Sky Survey (C-BASS), a 4.76 GHz survey covering the whole sky in total intensity and polarization. I describe the point-source detection algorithm developed for C-BASS, and the northern sky point-source catalogue that has been produced using it. This catalogue has allowed us to confirm the accuracy of the C-BASS pointing and flux-density scale, and will form the basis of any C-BASS point source masks. I also present an analysis of the synchrotron power spectra, using C-BASS, WMAP and *Planck* data. From this it is found the minimum synchrotron contamination to CMB B -modes corresponds to a tensor-to-scalar ratio of $r \sim 10^{-3}$, with a typical contamination at the level of $r \sim 10^{-2}$.

Alongside the analysis of C-BASS data, I present a novel implementation of Bayesian parametric component separation. This uses the No-U-Turn Sampler (NUTS) to explore the posterior distribution. NUTS is a gradient-based sampling algorithm with excellent scaling to high dimensions, and also contains important self-diagnostics of potential failures in geometric ergodicity. This is particularly important for the hierarchical foreground model introduced here, which can exhibit highly complex posterior geometries. The hierarchical foreground model fits for hyper-distributions over large sky regions, from which foreground spectral parameters are drawn. This is compared to a complete pooling model, where foreground spectral parameters are assumed to be constant in each sky region. The hierarchical model is able to remove artefacts from the recovered CMB maps without inflating parameter uncertainties, translating through to reduced biases on cosmological parameters.

Contents

1	Introduction	1
1.1	The Early Universe	1
1.1.1	Cosmological Inflation	3
1.1.2	The Cosmic Microwave Background	6
1.1.3	CMB Anisotropies	8
1.1.4	CMB polarization	11
1.2	CMB Foregrounds	18
1.2.1	Synchrotron	18
1.2.2	Thermal Dust	20
1.2.3	Anomalous Microwave Emission	21
1.2.4	Free-Free Emission	22
1.2.5	Extragalactic Sources	23
1.2.6	Other Astrophysical Foregrounds	24
1.3	CMB Component Separation Algorithms	25
1.3.1	Template Fitting	25
1.3.2	Internal Linear Combination	26
1.3.3	Independent Component Analysis	27
1.3.4	Parametric Component Separation	29
1.4	CMB Experiments	31
1.4.1	CMB Temperature and E -mode Observations	31
1.4.2	CMB B -mode Observations	33
1.4.3	Ancillary Surveys	34
1.5	Probability Distributions	36
1.6	Thesis Outline	37

2	The C-Band All-Sky Survey	39
2.1	Science Goals	40
2.1.1	Polarized CMB Foregrounds	40
2.1.2	Secondary Goals	41
2.2	Survey Constraints and Strategy	42
2.3	The C-BASS Instruments	44
2.4	Data Processing	50
2.4.1	Pointing	52
2.4.2	Gain and Polarization Angle Calibration	52
2.4.3	Non-Astrophysical Contamination	52
2.4.4	Astronomical Calibration	53
2.4.5	Map-Making	54
2.5	Data Checks: Calibration	54
2.5.1	Calibration Stability	54
2.5.2	Polarization Angles	57
2.6	Current Status	60
3	Point-Source Detection over the Northern Sky	62
3.1	Pre-Existing Source Catalogues	63
3.2	C-BASS Source Detection Algorithm	65
3.2.1	Source Detection Algorithm	66
3.2.2	The Spherical Mexican Hat Wavelet	69
3.2.3	Algorithm Validation	73
3.2.4	Catalogue Matching: Likelihood Ratios	75
3.2.5	Validation: Completeness, Position and Flux-Density Accuracy	77
3.3	The C-BASS Northern Sky Catalogue	80
3.3.1	Northern Sky Intensity Catalogue: Summary	81
3.3.2	GB6 and PMN Matching	83
3.3.3	Outlier Sources	89
3.3.4	Differential Source Counts	91
3.3.5	C-BASS Catalogue Sample	94
3.3.6	Source-Subtracted Map	98
3.4	Conclusions	98

4	The Power Spectra of Polarized Galactic Synchrotron Emission	102
4.1	Power Spectrum Estimation	103
4.1.1	Pseudo- C_ℓ Estimators	104
4.1.2	Bandpower Estimation	106
4.1.3	Validation Simulations	109
4.2	The C-BASS Polarization Power Spectra	115
4.2.1	E and B -mode power spectra	116
4.2.2	Synchrotron EB Correlation	120
4.3	Joint C-BASS Analysis with the WMAP and <i>Planck</i> Surveys	122
4.4	Synchrotron Contamination to CMB B -modes	125
4.4.1	Spatial Variation of Synchrotron Spectral Indices	125
4.4.2	Contamination to Primordial B -modes	130
4.5	Conclusions	132
5	Hierarchical Bayesian CMB Component Separation with the No-U-Turn Sampler	135
5.1	Approximate Bayesian Inference	137
5.2	Simulations	139
5.3	Mean-Shift Clustering of Sky Regions	143
5.4	Parametric Bayesian CMB Component Separation	146
5.4.1	Sky Model and Likelihood	147
5.4.2	The No-U-Turn Sampler	148
5.4.3	Complete Pooling of Spectral Parameters	151
5.4.4	Hierarchical Modelling of Spectral Parameters	154
5.4.5	Convergence Checks	157
5.5	Algorithm Validation	159
5.5.1	CMB Amplitude	163
5.5.2	Synchrotron and Dust Amplitudes	171
5.5.3	Synchrotron Spectral Parameters	174
5.5.4	Dust Spectral Parameters	177
5.6	Conclusions	178
6	Conclusions	181
6.1	Thesis Summary	181
6.2	Future Work	183
	Bibliography	186

List of Figures

1.1	An example single-field inflationary potential is plotted	5
1.2	The free electron fraction shown as a function of redshift on the lower x -axis, and of temperature on the upper x -axis	7
1.3	The effect of a quadrupole anisotropy on the linear polarization of the CMB.	12
1.4	The E and B -mode components of the polarization field	14
1.5	The polarization pattern, and the multipole power produced by scalar, vector and tensor perturbations	15
1.6	Plotted are the TT and EE power spectra, along with the BB power spectrum for a range of r values, and foreground contributions in ℓ -space for a selection of sky fractions and observing frequencies	17
1.7	The frequency dependence of CMB foregrounds for total intensity and polarization	19
1.8	The <i>Planck</i> 2018 CMB TT , TE and EE power spectra.	33
2.1	Photographs of the C-BASS north (<i>left</i>), and C-BASS south (<i>right</i>) telescopes	44
2.2	Panel (a): The C-BASS beam profile. Panel (b): The associated transfer functions	46
2.3	Simplified block diagram of the C-BASS front end, which is common to C-BASS north and south	48
2.4	Block diagrams of the C-BASS radiometer/polarimeter systems	49
2.5	Flow digram of the C-BASS data reduction pipeline	51
2.6	Recovered flux-densities for Cas A, Tau A, Cygnus A and Virgo A in monthly C-BASS maps	56
2.7	Panel (a): Polarization angle differences between C-BASS and WMAP K-band. Panel (b): Polarization angle differences between C-BASS and <i>Planck</i> 30 GHz	58

2.8	Histograms and corresponding KDE plots for the polarization angle differences between C-BASS and the WMAP K/ <i>Planck</i> 30 GHz maps	59
2.9	The current C-BASS Stokes I , Q , U , P and χ maps	61
3.1	Flowchart for the C-BASS point-source detection algorithm	66
3.2	Panel (a): An input 4.76 GHz simulated sky	70
3.3	Images illustrating the pixel tagging process	71
3.4	Panel (a): The real space SMHW2, shown for σ corresponding to a Gaussian beam FWHM of 45 arcmin and $R \in \{0.85, 0.95, 1.05, 1.15\}$. Panel (b): The corresponding window functions	72
3.5	The combined Galactic plane and C-BASS north mask used in the analysis of the C-BASS catalogue	77
3.6	Panel (a): A map of estimates of the flux density at 90 per cent completeness obtained from the validation simulations. Panel (b): A map of the corresponding reliability estimates	78
3.7	The histogram of measured absolute positional offsets, Δ , applying the CG30 mask, obtained for the 100 Monte Carlo simulations	79
3.8	Panel (a): The detected source flux densities against the matched input flux densities for the 100 Monte Carlo simulations, applying the CG30 mask. Panel (b): The histogram of the detected flux densities, divided by the matched input flux-densities	80
3.9	Positions of the northern sky, total intensity C-BASS point sources	83
3.10	$10^\circ \times 10^\circ$ gnomonic projection centred on Virgo A, showing spurious detections in the first sidelobe as white circles.	84
3.11	A histogram of the absolute positional offsets of C-BASS sources from matched GB6 and PMN sources	85
3.12	Flux-flux plots for matched C-BASS and GB6/PMN sources.	86
3.13	$10^\circ \times 10^\circ$ gnomonic projections, centred on the locations of sources in the C-BASS catalogue, identified as having a $\geq 3\sigma$ discrepancy with the matched GB6/PMN flux-densities.	90
3.14	Differential source counts calculated using the C-BASS northern sky total intensity catalogue	92
3.15	DETFLUX estimates plotted against APERFLUX estimates for the C-BASS catalogue	95
3.16	Panel (a): C-BASS northern sky map without source-subtraction. Panel (b): Source-subtracted C-BASS northern sky map	99

3.17	An example point-source mask derived from the C-BASS total intensity northern sky catalogue	100
4.1	The WMAP p06, p04 and p02 polarized emission masks, combined with the C-BASS northern sky mask covering declinations $\delta < -13^\circ$.	104
4.2	Recovered E and B auto-spectra at 4.76 GHz for the validation simulations	110
4.3	Recovered E -mode auto- and cross-spectra for the validation simulations, plotted as a function of frequency	112
4.4	Recovered B -mode auto- and cross-spectra for the validation simulations, plotted as a function of frequency	113
4.5	Recovered EB correlation coefficients for the validation simulations, shown for each sky mask	116
4.6	The C-BASS northern-sky E and B maps.	117
4.7	Recovered EE and BB auto-spectra for the C-BASS northern sky map, shown for each sky mask	118
4.8	Recovered EB correlation coefficients for the C-BASS northern sky map	121
4.9	Recovered E -mode auto- and cross-spectra, for C-BASS, WMAP and <i>Planck</i> maps, plotted as a function of frequency	123
4.10	Recovered B -mode auto- and cross-spectra, for C-BASS, WMAP and <i>Planck</i> maps, plotted as a function of frequency	124
4.11	Panel (a): The synchrotron spectral index map, estimated using the C-BASS, WMAP K, WMAP Ka and <i>Planck</i> 30 GHz maps. Panel (b): The corresponding standard deviation map	127
4.12	Gnomonic projections centred on the Galactic Centre, showing the mean (<i>top</i>) and standard deviation (<i>bottom</i>) of the posterior estimates of the spectral index estimates found in Jew and Grumitt (2020) . . .	129
4.13	Estimates of the synchrotron contamination to CMB B -modes	131
5.1	Input component parameter maps used in the validation simulations .	142
5.2	Regions defined using the mean-shift clustering algorithm	145
5.3	Graphical representation of the complete pooling model	152
5.4	Graphical representation of the hierarchical model	155
5.5	Histograms of the normalized deviations for the recovered parameter maps, obtained for the three validation sets	160
5.6	Output CMB amplitude maps obtained with the NUTS component separation algorithm	164

5.7	CMB amplitude residual maps for the three validation sets	165
5.8	CMB E -mode and B -mode power spectra for the three validation sets	167
5.9	Tensor-to-scalar ratio constraints obtained for the three validation sets	170
5.10	Synchrotron and dust amplitude maps obtained for the three validation sets	172
5.11	In panels (a) to (f) I show the synchrotron amplitude residual maps, and in panels (g) to (l) I show the dust amplitude residual maps . . .	173
5.12	Spectral parameter maps obtained for the three validation sets	175
5.13	Spectral parameter residual maps obtained for each validation set . .	176

Chapter 1

Introduction

In this chapter I discuss the origins of the cosmic microwave background (CMB), how we can use the CMB to constrain our understanding of the early Universe, and the key challenges in extracting the CMB signal from our noisy observations of the radio and microwave sky. In Section 1.1 I discuss cosmological inflation and the physics of the CMB. In Section 1.2 I discuss CMB foregrounds, focusing on the astrophysical emission components that complicate our observations of the CMB. In Section 1.3 I summarise common CMB component separation algorithms, that is methods used to extract the CMB signal from our observations of the radio and microwave sky. In Section 1.4 I give a brief summary of observational CMB studies, in Section 1.5 I define the probability distributions used in the statistical analyses in this thesis, and in Section 1.6 I outline the structure of this thesis.

1.1 The Early Universe

Soon after Einstein published his work on general relativity (Einstein, 1915, 1916), physicists were able to begin developing mathematically grounded models for the Universe. Initial investigations focused on a homogeneous and isotropic universe, described by the Friedmann-Lemaître-Robertson-Walker (FLRW) metric (Friedmann, 1922). Mathematically this takes the form

$$ds^2 = -c^2 dt^2 + a(t)^2 \left[\frac{dr^2}{1 - kr^2} + r^2 (d\theta^2 + \sin^2 \theta d\phi^2) \right], \quad (1.1)$$

where ds^2 is the space-time line element, $a(t)$ is the scale factor, k is a parameter describing the spatial curvature of the Universe, and (t, r, θ, ϕ) are the space-time coordinates. The scale factor is a dimensionless quantity that describes the expansion

history of the Universe, typically chosen to equal unity today. This metric can be used to derive the Friedmann equation that governs the evolution of the scale factor in a homogeneous and isotropic Universe, given by

$$H^2 \equiv \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{kc^2}{a^2}, \quad (1.2)$$

where H is known as the Hubble constant, and ρ is the energy density of the Universe. We can instead recast this equation in terms of fractional energy densities so that,

$$H^2 = H_0^2 \left(\frac{\Omega_M}{a^3} + \frac{\Omega_R}{a^4} + \Omega_\Lambda + \frac{\Omega_K}{a^2} \right), \quad (1.3)$$

where H_0 is the value of the Hubble constant today, Ω_M is the fractional energy density of matter today, Ω_R is the fractional energy density of radiation today, Ω_Λ is the fractional energy density of the cosmological constant today, and Ω_K is the effective fractional energy density in curvature today. The density parameter is defined such that $\Omega = \Omega_M + \Omega_R + \Omega_\Lambda$, and $\Omega = 1$ in a flat Universe. Each of these components of the energy density scale differently with $a(t)$, meaning they will dominate at different stages in the evolution of the Universe.

In general the Friedmann equation yields solutions that, depending on the energy content of the model Universe, undergo expansion and contraction at different points. Indeed, the cosmological constant term was originally introduced by Einstein in an attempt to construct a static Universe model. However, Einstein soon abandoned this notion when observations of galaxy redshifts by Edwin Hubble demonstrated that the Universe was indeed expanding (Hubble, 1929). Running the evolution of the Universe backwards it was realised that the Universe must have begun from some initial singularity, expanding ever since. This idea would develop into the Hot Big Bang model (Lemaître, 1931).

Through studying the thermal evolution of the early Universe, a number of key predictions were made that have gone on to be confirmed by observations. These include predictions of primordial helium abundances in stars (Alpher et al., 1948), and of the CMB radiation (Alpher and Herman, 1948). Despite its great success however, the standard Hot Big Bang model faces a number of fundamental problems. Three of these, the horizon, flatness, and magnetic monopole problems would motivate the development of the inflationary model for the early Universe (Guth, 1981; Linde, 1982; Albrecht and Steinhardt, 1982).

1.1.1 Cosmological Inflation

Cosmological observations have constrained the Universe to be very close to spatially flat ($k = 0$) (Planck Collaboration et al., 2018c), and the CMB radiation to be homogenous to approximately 1 part in 10^5 (Planck Collaboration et al., 2019). Moreover, Grand Unified Theories (GUTs), which posit that the electromagnetic, weak nuclear force and strong nuclear force are unified at sufficiently high energies, predict that magnetic monopoles should be formed in the early Universe ('t Hooft, 1974; Polyakov, 1974). These considerations lead to the flatness, horizon and magnetic monopole problems respectively, that were the original motivation behind the development of inflationary theory (Guth, 1981).

The flatness problem may be understood as follows. If we study the dynamics of solutions to the Friedmann equation, we find that a Universe with $\Omega = 1$ sits at an unstable fixed point. If there are slight deviations from $\Omega = 1$ after the Big Bang, such a Universe would be rapidly driven away from being spatially flat. In order for us to observe a spatially flat Universe today, we would require the value of Ω to be seemingly incredibly finely tuned to a value close to 1 at the time where the initial conditions are set. This problem is solved if the Universe is exactly spatially flat, i.e. $\Omega = 1$. However, if there are small deviations from $\Omega = 1$ today, we would seem to be presented with a fine tuning problem.

The horizon problem arises when we consider the size of the causal horizon at last scattering, that is the size of a region that could have been in causal contact up to the point at which the CMB was formed. Tracing the evolution of the scale factor through an early radiation dominated phase, followed by a matter dominated phase up to today, we can estimate the size of this apparent causal horizon to be $\sim 2^\circ$ on the sky. That is, when we observe the CMB, regions separated by more than $\sim 2^\circ$ could not have been in causal contact under such an expansion history. This again seems to present a fine tuning problem. The CMB is observed to be incredibly uniform, with deviations from the mean temperature at the level of 1 part in 10^5 (Planck Collaboration et al., 2019). We therefore require some explanation as to how apparently causally separated regions were able to reach thermal equilibrium with one another.

It is worth noting here that the flatness and horizon problems are not strict inconsistencies between the standard Hot Big Bang model and our observable Universe. If the initial conditions were such that the Universe was incredibly close to spatially flat and largely homogenous on superhorizon scales, the Universe would evolve to the state we see today. However, this requires the assumption that the initial conditions

were just so, and we may instead hope to construct some dynamical model with attractor solutions in the spatially flat and almost homogenous state we observe.

The magnetic monopole problem arises from theoretical considerations around GUTs, which predict that magnetic monopoles should have been formed at the very high energies present just after the Big Bang. However, none of these monopoles have been observed in nature. As such, for these GUTs to be consistent with reality we would require some mechanism through which we can reduce the density of magnetic monopoles to levels where they are unlikely to be observed (Peacock, 1999; Baumann, 2009).

Driven by these three problems, the theory of cosmological inflation was proposed (Guth, 1981; Linde, 1982; Albrecht and Steinhardt, 1982). The key idea behind inflation is that, in the period immediately after the Big Bang, the Universe underwent a period of super rapid expansion for a very short time. Provided there is sufficient expansion, the three problems outlined above are naturally resolved in this framework. For a non-flat Universe, we can re-write the Friedmann equation as

$$|1 - \Omega(a)| = \frac{1}{(aH)^2}, \quad (1.4)$$

this time writing the equation in terms of the time dependent density parameter, $\Omega(a)$. By construction, during inflation we have that aH is very large, meaning $1/(aH)^2 \rightarrow 0$, driving the non-flat Universe towards spatial flatness. The horizon problem is resolved in that, a small region of the Universe that is in causal contact prior to inflation massively grows in size during inflation. Large-scale modes that enter our causal horizon today were therefore in causal contact prior to inflation. Finally, should magnetic monopoles be produced in the early Universe, their density will be diluted enormously by the exponential expansion, providing a natural explanation for why they would not be observed today.

One can estimate the amount of inflation required to solve these problems to be ~ 60 e-folds, that is the Universe must expand by a factor of at least $\sim e^{60}$ during inflation (Peacock, 1999; Dodelson, 2003). It should be stressed here that there is no verified inflationary mechanism in particle physics. Instead inflationary theories typically propose some generic field or fields that drive inflation. The most common implementations are through scalar fields, which are able to naturally produce a period of inflation under certain conditions. It remains one the major outstanding theoretical challenges for inflation to understand the actual physical mechanism that drives it.

To understand how inflation can be implemented let us consider the case of a single, spatially uniform scalar field, ϕ with some potential, $V(\phi)$ and Lagrangian density, $\mathcal{L} = \frac{1}{2}(\partial_\mu\phi)^2 - V(\phi)$. The second time derivative of the scale factor is related to the energy density and pressure through the Raychaudhuri equation,

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3} \left(\rho + \frac{3P}{c^2} \right), \quad (1.5)$$

where P is the pressure. We can see immediately that for the Universe to undergo inflation where $\ddot{a} > 0$, we require the equation of state for our inflationary field to satisfy $P < -\rho c^2/3$. For our scalar field the equation of state is given by

$$w \equiv \frac{P}{\rho c^2} = \frac{\frac{1}{2}\dot{\phi}^2 - V(\phi)}{\frac{1}{2}\dot{\phi}^2 + V(\phi)}. \quad (1.6)$$

We can see here that if $\dot{\phi}^2/2 \ll V(\phi)$, i.e. the kinetic energy of the scalar field is much less than the potential energy, then the equation of state $w \approx -1$. It is precisely this slow-roll model that is the most common implementation of inflation. Physically we can interpret this as the inflaton (the particle associated with the inflationary field) slowly rolling down the potential slope during inflation (Peacock, 1999; Dodelson, 2003). This slow-roll picture is illustrated in Fig. 1.1.




Figure 1.1: An example single field inflationary potential is plotted, with the inflaton rolling down the potential slope. Initially the inflaton has a low kinetic energy in comparison to its potential energy and rolls slowly along the slope. It is during this period that inflation occurs. Eventually the kinetic energy of the inflaton increases as it approaches the potential minimum. Inflation ends at this stage and we enter the reheating phase. As the inflaton oscillates about the minimum, the couplings of the inflationary field to matter fields results in the production of particles. Figure taken from (Guzzetti et al., 2016).

In addition to providing a solution to the flatness, horizon and magnetic monopole problems, it was soon realised after its initial proposal that inflation provides a natural mechanism for generating the perturbations that would go on to seed large-scale structure. Quantum fluctuations in the early Universe generate background perturbations that are magnified during inflation. Perturbation modes with wavelengths much smaller than the horizon essentially behave as vacuum fluctuations on flat space, whilst modes that are magnified to scales much larger than the horizon are frozen in by causality. These super-horizon modes can then go on to act as the seeds for large-scale structure formation when they later re-enter the causal horizon (Starobinsky, 1982; Guth and Pi, 1982; Hawking, 1982; Bardeen et al., 1983; Brandenberger et al., 1983; Guth and Pi, 1985).

Since its initial proposal inflation has gone on to become a pillar of modern cosmology, receiving strong observational support. However, one of the key markers of inflation, the primordial B -mode polarization of the CMB, is yet to be observed. These observational tests and their associated challenges are discussed in detail in the remainder of the introduction.

1.1.2 The Cosmic Microwave Background

The discovery of the CMB was one of the greatest achievements in cosmology, and provided some of the most convincing evidence for the Big Bang theory. The CMB was observed by Penzias and Wilson with the Holmdel Horn Antenna (Penzias and Wilson, 1965). In attempting to remove sources of interference and noise from their observations, Penzias and Wilson found an isotropic excess antenna temperature of 3.5 ± 1.0 K. A joint announcement of the discovery was made with Dicke et al. (1965), who interpreted the measured radiation to be due to the CMB.

To understand the origins of the CMB we must consider the thermal history of the Universe after the Big Bang. Immediately after the Big Bang the Universe was in a very hot state, cooling as it expanded. At these high temperatures protons and electrons were decoupled, with the total ionization fraction of the primordial hydrogen being ~ 1 . When the Universe was filled with this high energy plasma the photon mean free path was much smaller than the cosmological horizon. Photons were unable to propagate for long before interacting with electrons through Compton scattering, resulting in the Universe being opaque.

Eventually the Universe cooled sufficiently that protons and electrons underwent recombination, forming neutral hydrogen. We can estimate the temperature at which

recombination occurs using the Saha equation,

$$\frac{X_e^2}{1 - X_e} = \frac{1}{n_e + n_H} \left[\left(\frac{m_e T}{2\pi} \right)^{3/2} \exp \left\{ -\frac{m_e + m_p - m_H}{T} \right\} \right], \quad (1.7)$$

where X_e is the free electron fraction, T is the temperature of the Universe, n_e is the number density of free electrons, n_H is the number density of neutral hydrogen, m_e is the electron mass, m_p is the proton mass and m_H is the mass of hydrogen. The Saha equation is derived by assuming the Universe is in a state where the reaction, $e^- + p \rightleftharpoons H + \gamma$ is in equilibrium.




Figure 1.2: The free electron fraction shown as a function of redshift on the lower x -axis, and of temperature on the upper x -axis. The full Boltzmann solution, and the Saha solution are plotted. The Boltzmann solution exhibits a longer period of recombination, and a freeze-out at $X_e \sim 10^{-4}$, with both solutions broadly agreeing on the start of recombination. CMB photons decouple from electrons during recombination, and we may view them as being released from a thin surface of last scattering. Figure taken from [Baumann \(2015\)](#).

At early times the Saha equation predicts $X_e \sim 1$, with appreciable recombination occurring when the temperature of the Universe falls to $T \sim 1/4$ eV, or a redshift of $z \sim 1000$. In reality, as X_e falls the rate of recombination falls and equilibrium conditions are no longer satisfied. To accurately track the evolution of X_e we must

solve the Boltzmann equation, given by

$$\frac{dX_e}{dt} = \langle \sigma v \rangle \left[(1 - X_e) \left(\frac{m_e T}{2\pi} \right)^{3/2} \exp \left\{ -\frac{m_e + m_p - m_H}{T} \right\} - X_e^2 n_b \right], \quad (1.8)$$

where $\langle \sigma v \rangle$ is the recombination rate and n_b is the Baryon number density. The solutions of the Saha and Boltzmann equations are shown in Fig. 1.2. The Saha equation is able to make a good prediction of the redshift at recombination, although the full Boltzmann solution reveals that recombination occurs at a slightly slower rate than predicted by the Saha equation, with a freeze-out at $X_e \sim 10^{-4}$. CMB photons decoupled from electrons during recombination, and can be viewed as having been released from a thin surface of last scattering (Dodelson, 2003).

The CMB is the most perfect blackbody that has been observed in nature (White, 1999). Observations with the FIRAS experiment have constrained the CMB blackbody temperature to be $T = 2.2760 \pm 0.0013 \text{K}$ (Fixsen, 2009). However, as mentioned previously, the CMB exhibits small anisotropies that provide a key window into constraining the physics of the early Universe. These anisotropies are discussed in the next section.

1.1.3 CMB Anisotropies

We can divide CMB anisotropies into primary, and secondary anisotropies. Primary anisotropies are those that result from processes prior to or at last scattering, whilst secondary anisotropies are due to processes that take place after last scattering (broadly, due to gravitational and scattering effects). Here the discussion is focused on primary anisotropies, for which I adopt much of the notation in Hu and Dodelson (2002). Secondary CMB anisotropies are discussed in detail in Aghanim et al. (2008).

In this section I will focus primarily on temperature perturbations, deferring detailed discussion of CMB polarization to the next section. For some sky direction, $\hat{\mathbf{n}}$ we can define the CMB temperature perturbation, $\Theta(\hat{\mathbf{n}}) = \Delta T / T_{\text{cmb}}$, where T_{cmb} is the mean CMB temperature and ΔT is some deviation from the mean temperature. Given that CMB data is analysed on the sphere, it is convenient to consider the harmonic decomposition of the temperature field,

$$\Theta_{\ell m} = \int d\hat{\mathbf{n}} Y_{\ell m}^*(\hat{\mathbf{n}}) \Theta(\hat{\mathbf{n}}), \quad (1.9)$$

where $Y_{\ell m}(\hat{\mathbf{n}})$ is the spherical harmonic function for multipole ℓ and order m . Assuming the CMB is a Gaussian field, these fluctuations are fully described by their power

spectrum,

$$\langle \Theta_{\ell m}^* \Theta_{\ell' m'} \rangle = \delta_{\ell\ell'} \delta_{mm'} C_\ell, \quad (1.10)$$

where C_ℓ are the power spectrum multipoles. The power spectrum predicted from physical models of the early Universe is the average power spectrum one would observe in an ensemble of universes. However, we have only one Universe to observe and for each multipole, at most $2\ell + 1$ modes we can average over. This presents us with a fundamental limit on the uncertainty we can achieve for each power spectrum multipole given by,

$$\Delta C_\ell = \sqrt{\frac{2}{2\ell + 1}} C_\ell. \quad (1.11)$$

This is known as the cosmic variance limit (Knox, 1995).

Given this formalism, we can now ask how we go about generating these anisotropies. In typical cosmological models, inflation acts as the mechanism which sets the initial cosmological perturbations. For a scalar field inflation model, quantum fluctuations result in perturbations to the scalar field, $\delta\phi$. These in turn induce curvature fluctuations in the early Universe. For scalar metric perturbations in the Newtonian gauge, these correspond to perturbations in the space-space part of the metric $\delta g_{ij} = 2a^2\Phi\delta_{ij}$. The time-time perturbation is given by $\delta g_{00} = 2\Psi$, where Ψ is the Newtonian gravitational potential and $\Psi \approx -\Phi$. Inflation predicts a nearly scale invariant power spectrum for the initial curvature perturbations i.e.,

$$\Delta_\Phi^2 \equiv \frac{k^3 P_\Phi(k)}{2\pi^2} \propto k^{n_s-1}, \quad (1.12)$$

where $P_\Phi(k)$ is the power spectrum of the curvature perturbations and the scalar spectral index, $n_s \approx 1$ for a nearly scale invariant power spectrum (Baumann, 2009).

We observe CMB anisotropies today on the spherical surface of last scattering, at $\mathbf{x} = D_* \hat{\mathbf{n}}$. $D_* = c \int_{t_*}^{t_0} dt/a(t)$ is the comoving distance to the surface of last scattering, where t_* is the time at recombination and t_0 is the time today. We can rewrite this in terms of the conformal time $\eta = \int dt/a(t)$, so that $D_* = c(\eta_0 - \eta_*)$. The observed CMB anisotropy today is given by

$$\Theta(\hat{\mathbf{n}}, \eta_0) = \sum_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}) \left[(-1)^\ell \int \frac{d^3 k}{(2\pi)^3} a_\ell(k) Y_{\ell m}^*(\hat{\mathbf{k}}) \right], \quad (1.13)$$

where $a_\ell(k)$ is the projected source term, which describes how some initial perturbation mode, k is mapped onto some ℓ -mode at recombination. The source term can be expressed in terms of the metric perturbations discussed earlier,

$$a_\ell(k) = [\Theta + \Psi](\eta_*) j_\ell(kD_*) + v_b(k, \eta_*) j'_\ell(kD_*) + \int_{\eta_*}^{\eta_0} d\eta (\dot{\Psi} - \dot{\Phi}) j_\ell(kD), \quad (1.14)$$

where $j_\ell(x)$ is the spherical Bessel function which acts to project our three-dimensional \mathbf{k} modes onto the surface of last scattering, v_b is the fluid velocity of baryons and $j'_\ell(x) = dj_\ell(x)/dx$.

The first term in Equation 1.14 is due to acoustic oscillations in the photon-baryon fluid up to recombination (Peebles and Yu, 1970). Perturbation modes experience oscillations due to the compression and rarefaction of the photon-baryon fluid in gravitational potential wells. The amplitudes of the Fourier modes are frozen in at recombination, setting a fundamental scale for the acoustic peaks in the angular power spectrum at

$$k_A = \frac{\pi}{s_*}, \quad (1.15)$$

where s_* is the size of the sound horizon at recombination. This corresponds to a fundamental angular scale for acoustic oscillations of

$$\ell_A = k_A \tilde{D}_*, \quad (1.16)$$

where \tilde{D}_* is the apparent distance to the last scattering surface, which will depend on the curvature of the Universe. In a closed Universe an object of a given size will be further away than it appears, and vice versa for an open Universe. The spacing of the acoustic peaks is given by $\ell_n = nk_A \tilde{D}_*$, where n is an integer. This means the observed scale of the acoustic peaks is very sensitive to the spatial curvature (Hu, 2003). The acoustic signatures in the angular power spectrum are further complicated by the fact that there are multiple components to the energy content of the Universe, with particularly important effects coming about from baryon loading, and the radiation dominated era prior to recombination.

Baryon loading refers to the effect that additional baryons have on the behaviour of the photon-baryon fluid. Baryons add mass but not pressure, meaning that additional baryons will increase the depth of the potential well the photons are undergoing acoustic oscillations in. The photons will fall further into the potential well before the oscillation is reversed, increasing the amplitude of peaks in the power spectrum representing compressions i.e., the odd numbered peaks (Hu and Sugiyama, 1995).

The Universe is radiation dominated prior to recombination, with the matter-to-radiation ratio being of order unity at recombination. Radiation changes the expansion rate of the Universe, and therefore changes the size of the sound horizon at recombination, further complicating the location of the peaks. However, radiation also plays an important role in providing a driving contribution at higher multipoles. As the photon fluid compresses, pressure eventually halts further compression, stabilizing the density perturbation. The gravitational potential will then decay as the

Universe expands and the background density decreases, increasing the amplitude of acoustic oscillations for smaller scale modes that started oscillating when the Universe was radiation dominated (Seljak, 1994; Hu and Sugiyama, 1995).

In addition to these contributions from radiation and matter, we also have damping effects at small scales. When we reach small enough scales the distances involved become comparable to the photon mean free path at recombination. If this is the case then hot and cold photons will be able to mix, smoothing out any anisotropies on small scales in a process known as Silk damping (Silk, 1968).

The second term in Equation 1.14 is due to the Doppler effect. The motion of the photon-baryon fluid as it undergoes oscillations Doppler shifts the radiation, which produces an additional temperature anisotropy. Finally, the third term in Equation 1.14 describes the integrated Sachs-Wolfe effect, which quantifies the gravitational red/blue-shifting of CMB photons as they travel from the surface of last scattering to us now. This includes contributions from radiation after decoupling, which leads to further decay of the gravitational potential. However, the dominant contribution here arises from the late-time decay of the gravitational potential caused by dark energy. Putting all this together, we obtain the predicted CMB power spectrum,

$$C_\ell = \frac{2}{\pi} \int dk k^2 a_\ell^2(k). \quad (1.17)$$

The full numerical calculation of the CMB power spectra has been implemented in sophisticated Boltzmann solvers (Seljak and Zaldarriaga, 1996; Lewis and Challinor, 2011; Lesgourgues, 2011), providing a detailed forward model for statistical inference.

1.1.4 CMB polarization

CMB polarization arises from the Thomson scattering of CMB photons off of electrons, in particular Thomson scattering in the presence of a quadrupole anisotropy. To understand this we can consider the Thomson scattering differential cross section,

$$\frac{d\sigma_T}{d\Omega} \propto |\hat{\epsilon}' \cdot \hat{\epsilon}|^2, \quad (1.18)$$

where $\hat{\epsilon}'$, and $\hat{\epsilon}$ are the incident and scattered photon polarization directions respectively. An incoming photon will cause an electron to oscillate parallel to the polarization direction. From Equation 1.18 we can see that the scattered radiation will peak at 90° to the incoming radiation, with the scattered polarization direction parallel to the incident polarization direction.

If the incoming radiation was isotropic the electron would be scattered in all directions, resulting in overall un-polarized scattered radiation. A dipole anisotropy would also be insufficient. In such a case hot radiation will scatter the electron with colder radiation approaching from the opposite direction. This will produce an outgoing photon at the mean temperature, with a polarization parallel to the incoming polarization. This will combine with scattering from photons perpendicular to these, which are at the mean temperature, producing no net polarization.




Figure 1.3: The effect of a quadrupole anisotropy on the linear polarization of the CMB. The red lines represent colder radiation, and the blue lines represent hotter radiation. The radiation field exhibits a quadrupole anisotropy where the hot and cold radiation approach at 90° to one another. The incoming radiation is scattered at 90° to the incident direction, with the scattered polarization parallel to the incident polarization direction. As shown in the diagram, the quadrupole anisotropy produces differing amplitudes at 90° , resulting in a net linear polarization. Figure taken from [Hu and White \(1997\)](#).

Fig. 1.3 illustrates the scattering scenario in the presence of a quadrupole anisotropy. In this case the radiation field possesses intensity peaks at 90° separations. If we have cold radiation approaching the electron at 90° to hot radiation, the scattered radiation produced can be viewed as a two outgoing waves polarized at 90° to one another. However, the intensities of the two waves will be unequal, resulting in an overall linear polarization along the outgoing directions ([Hu and White, 1997](#)).

Before moving to discuss how different types of primordial perturbation can generate different polarization fields, let us first understand how we can describe some polarization field on the sky. Consider an electromagnetic wave of frequency ω_0 prop-

agating along the z -direction, with components E_x and E_y given by

$$E_x = E_x(t) \cos(\omega_0 t - \theta_x(t)), \quad (1.19)$$

$$E_y = E_y(t) \sin(\omega_0 t - \theta_y(t)), \quad (1.20)$$

where $\theta_x(t)$ and $\theta_y(t)$ are the phase angles. We can describe the polarization state of the electromagnetic wave using the Stokes parameters (I, Q, U, V) , where I is the total intensity, Q and U describe the linear polarization, and V describes the circular polarization. In the (x, y) basis these are given by

$$I \equiv \langle E_x^2 \rangle + \langle E_y^2 \rangle, \quad (1.21)$$

$$Q \equiv \langle E_x^2 \rangle - \langle E_y^2 \rangle, \quad (1.22)$$

$$U \equiv \langle 2E_x E_y \cos(\theta_x - \theta_y) \rangle, \quad (1.23)$$

$$V \equiv \langle 2E_x E_y \sin(\theta_x - \theta_y) \rangle. \quad (1.24)$$

The I , and V Stokes parameters are both independent of the coordinate system. The Q and U Stokes parameters on the other hand are not invariant under a rotation of the coordinate axes. If we describe the linear polarization field using $Q \pm iU$, under an anti-clockwise rotation of the coordinate axes by an angle ψ , the polarization field transforms as $(Q \pm iU) \rightarrow (Q \pm iU) e^{\mp 2i\psi}$. The linear polarization therefore transforms under rotation as a spin-2 field. It is straightforward to show that the linearly polarized intensity, $P = \sqrt{Q^2 + U^2}$, is invariant under coordinate transforms. We can also define a polarization angle $\chi = \arctan(U/Q)/2$, which describes a constant orientation of the polarization parallel to the electric field vector (Kosowsky, 1999).

A description of the CMB polarization field using Q and U will be coordinate dependent. It is therefore helpful to introduce a coordinate independent representation. We can decompose the polarization field into E and B -mode components, using the spin-2 object $Q \pm iU$, as follows,

$$E_{\ell m} \pm iB_{\ell m} = - \int d\hat{\mathbf{n}} \, {}_{\pm 2}Y_{\ell m}^*(\hat{\mathbf{n}}) (Q(\hat{\mathbf{n}}) \pm iU(\hat{\mathbf{n}})), \quad (1.25)$$

where $E_{\ell m}$, and $B_{\ell m}$ are the spherical harmonic amplitudes for the E and B -modes, and ${}_{\pm 2}Y_{\ell m}(\hat{\mathbf{n}})$ are the second-rank tensor spherical harmonics. In general ${}_s Y_{\ell m}(\hat{\mathbf{n}})$ are the eigenfunctions of the Laplace operator acting on a rank s tensor (Hu, 2003). We then have the spin-0 E and B -fields defined as,

$$E(\hat{\mathbf{n}}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} E_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}), \quad (1.26)$$

$$B(\hat{\mathbf{n}}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} B_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}). \quad (1.27)$$




Figure 1.4: The E and B -mode components of the polarization field. The E -mode displays either a radial or tangential pattern, with no handedness. The B -mode on the other hand has a curl-like pattern, with a left or right handedness. The E -mode pattern is radial around cold spots and tangential around hot spots, whilst the B -mode displays differing handedness around hot and cold spots. Figure taken from [Krauss et al. \(2010\)](#).

The E and B -mode components are so named due to the similarity of their polarization patterns to electric and magnetic fields respectively. Fig. 1.4 shows the form of the E and B -mode polarization components. The E -mode has a radial or tangential pattern, and displays no handedness i.e., both E -mode patterns remain the same under reflection. The E -mode pattern is radial around cold spots in the CMB, and tangential around hot spots. For the B -mode we have a curl-like pattern, with the positive and negative B -mode patterns displaying a handedness i.e., the positive and negative B -modes are interchanged under reflection. The B -mode pattern displays differing handedness around hot and cold spots.

In Section 1.1.3 I discussed the impact of primordial perturbations on the CMB TT power spectrum. In general, perturbations to the metric can be decomposed into scalar, vector and tensor components, so called because of their transformation properties. Scalar perturbations can be interpreted as perturbations to the energy density of the photon-baryon fluid, and are the primary source of CMB temperature anisotropies. Vector perturbations represent vorticity perturbations to the photon-baryon fluid. Vector perturbations are not produced by inflation. Moreover, if there were primordial vector perturbations they would be rapidly damped by the expansion of the Universe during inflation, and so would be negligible anyway. Finally, tensor perturbations can be interpreted as primordial gravitational waves, and can

be sourced by inflation (Hu and White, 1997; Kamionkowski et al., 1997; Hu and Dodelson, 2002; Hu, 2003; Baumann et al., 2009).

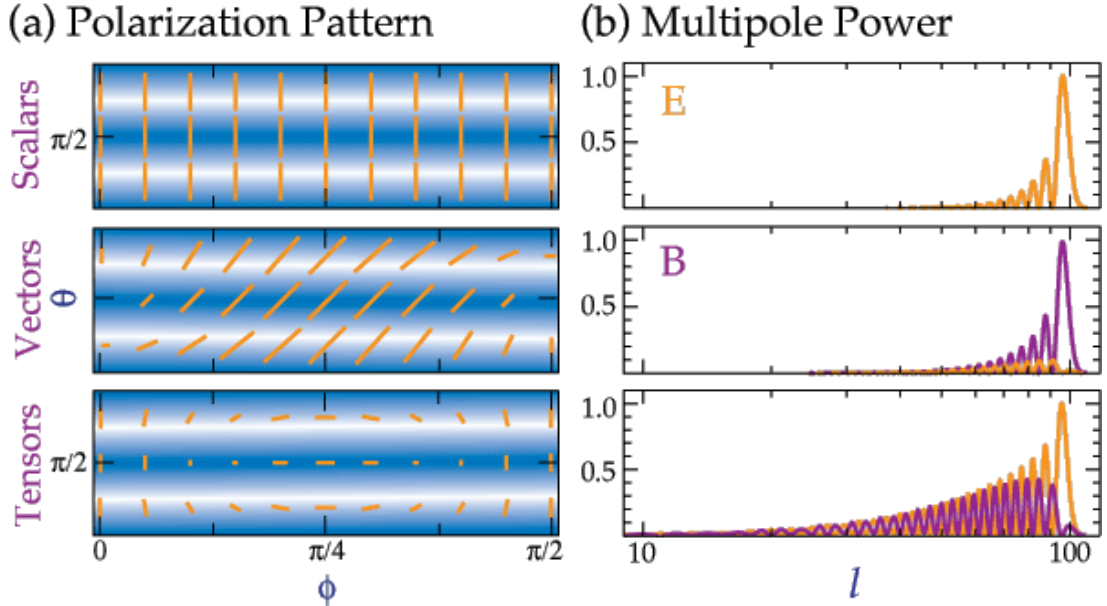


Figure 1.5: The polarization pattern, and the multipole power produced by scalar, vector and tensor perturbations. The local pattern is modulated by the plane wave spatial dependence of the temperature perturbation. For scalars we generate power only in the E -mode, for vectors mainly in the B -mode, and for tensors we generate similar power in E and B . Figure taken from Hu and White (1997).

The local polarization patterns produced by scalar, vector and tensor perturbations are shown in Fig. 1.5. Also shown is the overall power produced in E -modes and B -modes for the different perturbation modes. Thomson scattering can only produce a local E -mode polarization pattern. This is because the spherical harmonics describing the temperature perturbations have $(-1)^\ell$ parity, whilst the B -mode has $(-1)^{\ell+1}$ parity. The B -mode power from tensor perturbations comes about because the overall polarization pattern on the sky is the result of the local polarization pattern from Thomson scattering, modulated by the plane wave spatial dependence of the temperature perturbation. Modulation parallel or orthogonal to the polarization direction will preserve the E -mode nature of the local pattern. However, modulation that is crossed with (at 45° to) the polarization direction will generate a B -mode component.

For scalar perturbations the local pattern is curl free, so the modulation is always parallel or orthogonal to the polarization direction. Scalar perturbations therefore only generate power in the E -mode. For tensors on the other hand, the local pattern

has a curl component. The modulation is therefore crossed with the local pattern at certain points, generating power in both E and B . Vectors generate power mainly in the B -mode. However, as mentioned before, even if present vector perturbations will be damped away by the expansion of the Universe (Hu and White, 1997).

Tensor perturbations can be described by a metric with $g_{00} = -1$, $g_{0i} = 0$ and spatial elements

$$g_{ij} = a^2 \begin{pmatrix} 1 + h_+ & h_\times & 0 \\ h_\times & 1 - h_+ & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (1.28)$$

The perturbations here are described by h_+ and h_\times , with this particular choice of metric corresponding to perturbations in the x - y plane and the wavevector being along the z -axis. More generally, we can describe tensor perturbations with a divergenceless, traceless, symmetric tensor, \mathcal{H}_{ij} . For many inflationary models, the tensor power spectrum is expected to take the form of a power law

$$\Delta_h^2 \equiv \frac{k^3}{\pi^2} \langle |h_+(k)|^2 + |h_\times(k)|^2 \rangle \propto k^{n_t}, \quad (1.29)$$

where n_t is the tensor spectral index (Dodelson, 2003). The key observable of interest in the hunt for primordial gravitational waves is the tensor-to-scalar ratio, defined as

$$r \equiv \frac{\Delta_h^2(k_0)}{\Delta_\Phi^2(k_0)} \quad (1.30)$$

where k_0 is some pivot scale. For simple single-field, slow-roll inflationary models r can be directly related to the energy scale of inflation. For this class of model it can be shown that, to first order,

$$r = 16\epsilon_V = 8M_{\text{Pl}}^2 \left(\frac{\partial_\phi V}{V} \right)^2, \quad (1.31)$$

where $M_{\text{Pl}}^2 = 2.435 \times 10^{18} \text{ GeV}/c^2$ is the reduced Planck mass (Peiris et al., 2003). A measurement of r therefore not only provides powerful evidence for inflation, but can also provide a direct probe of the inflationary energy scale. Indeed, the detection of primordial B -mode polarization in the CMB is arguably the major outstanding goal of CMB cosmology. However, this goal faces significant challenges from polarized foreground contamination.

Fig. 1.6 shows plots of the theoretical TT , EE , and BB power spectra. The BB power spectrum has been plotted for a range r values. Also shown are a range of possible ℓ -space contributions from polarized dust and synchrotron emission, and the




Figure 1.6: Plotted are the TT and EE power spectra, along with the BB power spectrum for a range of r values, and foreground contributions in ℓ -space for a selection of sky fractions and observing frequencies. Current upper limits on the value of the tensor-to-scalar ratio are at $r < 0.056$ (Planck Collaboration et al., 2018d). We can see here that the expected contributions from polarized foregrounds will likely dominate over any primordial signal. Figure taken from Errard et al. (2016).

power from lensing B -modes. Lensing B -modes are produced by the gravitational lensing of CMB photons, which can mix E -modes into B -modes, and dominate at high ℓ .

The primordial B -mode signal exhibits two key features, the reionization bump at $\ell \sim 10$ and the recombination bump at $\ell \sim 100$. The reionization bump arises due to the re-scattering of CMB photons by free electrons during the reionization epoch. These free electrons came about from the emission of ultraviolet light by the first stars. The recombination bump is a result of gravitational waves entering the horizon at around the same time as recombination. At low ℓ the corresponding gravitational waves were still superhorizon, meaning their role in interactions within the horizon were suppressed. At high ℓ the corresponding gravitational waves had been redshifting prior to recombination, reducing their amplitude and hence the power at high ℓ (Kamionkowski and Kovetz, 2016).

1.2 CMB Foregrounds

In Section 1.1 I have surveyed the theoretical mechanisms that generate CMB anisotropies and CMB polarization. By constructing experiments to observe these CMB anisotropies we can place tight constraints on the physics of the early Universe. However, the practical exercise of taking these measurements is hampered by the presence of astrophysical foregrounds, particularly in contemporary searches for primordial B -mode polarization. Fig. 1.7 shows the frequency spectra of foreground components in total intensity and polarization. At first glance, it may seem that the foregrounds in polarization are simpler than in total intensity. However, given current upper limits on r , polarized foreground emission is expected to dominate over the CMB at all frequencies. In this section, I summarise the key astrophysical foregrounds that we must contend with.

1.2.1 Synchrotron

Synchrotron emission is caused by relativistic cosmic ray electrons spiralling in the Galactic magnetic field (GMF). For a single electron accelerating in a magnetic field, B , the frequency spectrum of the emitted radiation is sharply peaked around

$$\nu \approx \gamma^2 \nu_g, \quad (1.32)$$

where γ is the electron's Lorentz factor and $\nu_g = eB/2\pi m_e$ is the non-relativistic gyrofrequency (e is the electron charge and m_e is the electron rest mass). For an

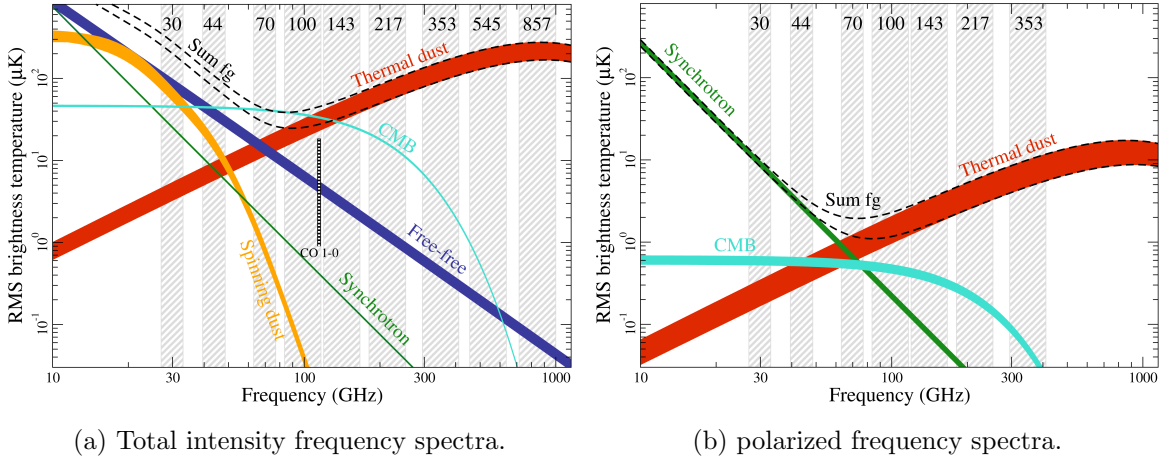


Figure 1.7: The frequency dependence of CMB foregrounds for total intensity and polarization. Also shown is the CMB spectrum, and the *Planck* observing bands in grey. Emission components are shown at a smoothing scale of 1° in total intensity and 40 arcmin in polarization. The total intensity foregrounds appear more complicated than in polarization, but are comparable to or lower in amplitude than the CMB. For *B*-mode polarization foreground emission dominates over the CMB at all frequencies. Figure taken from [Planck Collaboration et al. \(2016d\)](#).

ensemble of electrons with an isotropic distribution of pitch angles, the average power emitted per electron is given by

$$-\frac{dE}{dt} = \frac{4}{3}\sigma_T c \gamma^2 \frac{B^2}{2\mu_0}. \quad (1.33)$$

The distribution of electron energies in the interstellar medium (ISM) can be approximated by a power-law, $N(E)dE \propto E^{-p}dE$, where $N(E)dE$ is the number density of electrons with energies between E and $E + dE$. The energy radiated between frequencies ν and $\nu + d\nu$ is then given by

$$S_\nu d\nu = -\frac{dE}{dt} N(E)dE. \quad (1.34)$$

The frequency spectrum for this electron population therefore obeys the power law

$$S_\nu \propto \nu^{-(p-1)/2} = \nu^\alpha, \quad (1.35)$$

where α is the synchrotron flux-density spectral index, related to the temperature spectral index by $\beta_s = \alpha + 2$ (see e.g. [Longair, 2011](#)).

This power law approximation of the synchrotron spectrum holds well over a wide range of frequencies, from $\mathcal{O}(10 \text{ MHz})$ up to $\mathcal{O}(100 \text{ GHz})$ ([Lawson et al., 1987](#); [Reich and Reich, 1988](#); [Platania et al., 2003](#); [Davies et al., 2006](#); [Guzmán et al., 2011](#)).

Theoretical considerations and observations give a value of $\beta_s \approx -3$ for Galactic synchrotron emission, with observed spatial variations at the level of $\Delta\beta_s \approx 0.2$ (Dunkley et al., 2009; Planck Collaboration et al., 2016d; Dickinson, 2016).

This power-law behaviour is broken by two intrinsic effects, synchrotron self-absorption and spectral ageing (Mitton and Ryle, 1969; Harwood et al., 2016). Synchrotron self-absorption occurs when photons propagating out of some source are scattered by electrons in the plasma. This effect is much stronger at low frequencies (below ~ 1 GHz), resulting in a turn-off in the spectrum. Conversely, at high frequencies we see a cut-off in the spectrum caused by spectral ageing. High frequency synchrotron emission is caused by the highest energy electrons. From Equation 1.33 we can see higher energy electrons radiate their energy away faster. Therefore, as a population of synchrotron emitting electrons age the number of high energy electrons will fall faster than the number of low energy electrons, producing a high frequency cut-off (Longair, 2011).

In addition to these intrinsic effects, pixel and beam averaging will modify the observed synchrotron spectrum (Chluba et al., 2017; Remazeilles et al., 2018). These additional complications can be modelled through the introduction of curvature terms in the synchrotron power spectrum, or through the general moment expansion method in Tegmark (1998); Zaldarriaga et al. (2004); Chluba et al. (2017). However, in order to constrain these extra parameters we require additional low frequency data. These issues will be discussed in more detail in Chapter 5.

Synchrotron emission also displays a high level of polarization. The spiralling of cosmic ray electrons around magnetic field lines picks out a special direction. Broadly speaking, the electric field of the emitted radiation is aligned with the projected acceleration vector. In principle, the polarization fraction for synchrotron (the ratio of the polarized intensity to the total intensity) could be as high as 75%. However, depolarization effects along the line of sight and irregular magnetic fields reduce the observed polarization fraction, although polarization fractions in the range $\sim 10\text{--}40\%$ have been observed for synchrotron emission away from the Galactic plane (Dunkley et al., 2009; Dickinson, 2016).

1.2.2 Thermal Dust

Thermal dust emission is caused by the heating of dust grains in the ISM, causing them to vibrate and emit thermally. For frequencies above ~ 100 GHz thermal dust emission is the dominant foreground in both total intensity and polarization. The

thermal dust frequency spectrum can be approximated with a modified blackbody (MBB) model for frequencies below ~ 1 THz,

$$S_d = A_d \left(\frac{\nu}{\nu_0} \right)^{\beta_d+1} \frac{\exp(\gamma\nu_0) - 1}{\exp(\gamma\nu) - 1}, \quad (1.36)$$

where A_d is the dust amplitude, ν_0 is some reference frequency, $\gamma = h/k_B T_d$, T_d is the dust temperature and β_d is the dust spectral index. This model was adopted in the *Planck* foreground analysis, where typical values of the dust temperature were found to be $T_d \approx 19$ K, and the dust spectral index of $\beta_d \approx 1.6$ (Planck Collaboration et al., 2016d, 2018b). The MBB model is a simplification and in reality multiple dust populations will exist along the line of sight, contributing the same line of sight and beam averaging effects that modify the synchrotron spectrum. More complex emission models have been considered in the context of CMB component separation in Hensley and Bull (2018). However, these complex models raise the same issues as for synchrotron emission. The additional parameters introduced require additional data to constrain them, with current and planned CMB experiments already struggling to constrain dust spectral parameters for the simple MBB model (Jew et al., 2019).

Thermal dust emission also exhibits significant polarization. Dust grains are not spherically symmetric, with dust grains emitting preferentially along their longer axis, resulting in an intrinsic polarization. Large dust grains can also be aligned by the GMF, again causing the emission to be polarized. Observations from the *Planck* satellite have found that thermal dust polarization fractions can be as high as 20%, with a median value across the sky of approximately 8% (Planck Collaboration et al., 2016d, 2018b).

1.2.3 Anomalous Microwave Emission

The first signs of anomalous microwave emission (AME) were noted in Kogut et al. (1996); de Oliveira-Costa et al. (1997); Netterfield et al. (1997), with the first identification of an anomalous component to the Galactic emission being made in Leitch et al. (1997). The AME spectrum has been found to be peaked at a frequency of ~ 30 GHz, although observations of individual AME regions have found a range of peak frequencies, ranging from ~ 20 GHz, to as high as ~ 50 GHz.

The physical origin of AME is not completely understood. The most well developed model, originally proposed by Draine and Lazarian (1998), is the spinning dust model. The core idea here is that a dust grain with electric dipole moment μ will

radiate according to the Larmor formula

$$P = \frac{2}{3} \frac{\omega^2 \mu^2 \sin^2 \theta}{c^3}, \quad (1.37)$$

where P is the emitted power, ω is the angular velocity of the dust grain and θ is the angle between the angular velocity and dipole moment vectors. The challenge in developing theoretical spinning dust models is then in understanding the distribution of angular velocities and dipole moments of dust grains in the ISM (Dickinson et al., 2018). Alternative models have also been proposed, including the magnetic dust model from Draine and Lazarian (1999). Under this model thermal fluctuations can excite the magnetization of dust grains away from their minimum energy direction. The magnetization will then precess back to this minimum energy direction, with the dust grain emitting radiation in doing so.

AME measurements have generally found upper limits on the polarization fraction of a few percent. Significant polarization of spinning dust grains would require alignment of dust grains in the ISM. However, theoretical arguments in Draine and Hensley (2016) would suggest that such alignment is not achieved, and polarization fractions for spinning dust would be $\lesssim 10^{-6}$. Magnetic dust models typically predict much higher polarization fractions of up to $\sim 20\%$, which forms one of the main arguments against magnetic dust as one of the dominant components of AME. Additional studies with low frequency radio data ($\nu \lesssim 30$ GHz) will be needed to place tighter constraints on the level of AME polarization, such that we can properly assess the potential contamination to CMB polarization.

1.2.4 Free-Free Emission

Free-free emission is the result of electron-ion collisions, also known as bremsstrahlung radiation. Given the random nature of the Coulomb interactions between electrons and ions, free-free emission is expected to be intrinsically un-polarized, with upper limits on the polarization fraction of $\sim 1\%$ (Macellari et al., 2011). At the edges of bright H II regions, additional Thomson scattering of electrons can result in higher polarization fractions of $\sim 10\%$ (Rybicki and Lightman, 1986; Keating et al., 1998). However, for the purposes of CMB polarization studies, given the focus on higher Galactic latitudes and relatively low resolution, these effects are expected to be largely negligible (Dunkley et al., 2009).

Above a few GHz the frequency spectrum of free-free emission shows a power-law behaviour, with a spectral index of $\beta_{\text{ff}} \approx -2.1$. At lower frequencies, below ~ 1 GHz,

the medium for free-free emission becomes optically thick, leading to a break in the power-law frequency scaling. At these frequencies the mean free path of photons escaping from the free-free medium becomes much smaller than the source size. The full emission spectrum can be well described by the model of [Draine \(2011\)](#),

$$g_{\text{ff}} = \log \left\{ \exp \left[5.960 - \frac{\sqrt{3}}{\pi} \log \left(\frac{\nu}{10^9} \left(\frac{T_e}{10^4} \right)^{-3/2} \right) \right] + e \right\}, \quad (1.38)$$

$$\tau = 0.05468 T_e^{-3/2} \left(\frac{\nu}{10^9} \right)^{-2} \text{EM } g_{\text{ff}}, \quad (1.39)$$

$$S_{\text{ff}} = 10^6 T_e (1 - e^{-\tau}). \quad (1.40)$$

The two free parameters in this model are the electron temperature, T_e and the emission measure EM. The emission measure is defined as the integrated, squared electron density along the line of sight measured in cm^{-6}pc , and acts to set the amplitude of free-free emission. The spectral scaling of free-free is weakly set by the electron temperature. The term, g_{ff} is known as the Gaunt factor and accounts for quantum effects in the scattering. The form presented here assumes ions have charge $Z = +1$. A more general calculation of bremsstrahlung Gaunt factors can be found in [Chluba et al. \(2020\)](#).

1.2.5 Extragalactic Sources

Extragalactic sources are a key contaminant in CMB studies, with their detection and removal presenting an important challenge ([Taylor et al., 2001](#); [Waldram et al., 2003, 2007](#); [de Zotti et al., 2010](#)). For the brightest sources, they can be identified and either subtracted or masked from our sky maps. Assuming sources with flux densities $S_\nu > S_{\text{max}}$ have been successfully removed, and that sources are Poisson distributed on the sky, it can be shown that the point source contamination to the power spectrum, C_ℓ^{PS} is given by

$$C_\ell^{\text{PS}} = \int_0^{S_{\text{max}}} \frac{dN}{dS} S^2 dS, \quad (1.41)$$

where S is the flux density of sources and dN/dS is the differential source count, i.e., the number of sources per unit flux density, per unit steradian ([Tegmark and Efstathiou, 1996](#)).

For future CMB experiments, it has been shown that radio point-source contamination has the potential to significantly complicate measurements of the B -mode

polarization power spectrum (Mesa et al., 2002; Curto et al., 2013; Puglisi et al., 2018; Remazeilles et al., 2018). For values of the tensor-to-scalar ratio $r \sim 10^{-3}$, point sources are expected to be a dominant foreground after correcting for the gravitational lensing of E -modes on angular scales $\ell \gtrsim 50$ i.e., over angular scales corresponding to the recombination peak in the primordial B -mode power spectrum.

Galaxy clusters can also act to distort the observed CMB through the thermal Sunyaev-Zel'dovich (SZ) effect. As CMB photons pass through these galaxy clusters they are inverse Compton scattered off electrons in the hot intergalactic medium. This induces a distortion in the CMB blackbody spectrum given by

$$\frac{\delta I_\nu^{\text{cmb}}}{I_\nu^{\text{cmb}}} = -2 \int n_e \sigma_T \frac{k_B T_e}{m_e c^2} dl, \quad (1.42)$$

where n_e is the number density of electrons and the integral is taken along the line of sight (Sunyaev and Zeldovich, 1972). An additional distortion arises from the motion of galaxy clusters relative to the Hubble flow i.e., their peculiar velocities relative to the expansion of the Universe. This is known as the kinetic SZ effect and is given by

$$\frac{\delta I_\nu^{\text{cmb}}}{I_\nu^{\text{cmb}}} \approx -\tau_e \frac{v_z}{c}, \quad (1.43)$$

where $\tau_e = \int n_e \sigma_T dl$ is the optical depth of the cluster and v_z is the component of the cluster's peculiar velocity along the line of sight.

1.2.6 Other Astrophysical Foregrounds

Line emission can have a significant impact in narrow frequency bands, and was particularly apparent in the 100, 217 and 353 GHz *Planck* maps. In this case the line emission was the result of rotational transitions in carbon monoxide (CO) molecules. For the *Planck* mission, line emissions could be extracted during component separation by utilizing the differing bandpasses between different detectors. Whilst components with continuous spectra would vanish in difference maps between detectors, the line emission would not (Planck Collaboration et al., 2014a, 2016d). More generally, one can seek to mitigate this problem by designing the experimental bandpasses such that you avoid CO lines (Sekimoto et al., 2018).

In addition to the Galactic and extragalactic foregrounds discussed above, one must also account for zodiacal emission. This is caused by emission from interplanetary dust (IPD) in the solar system. The Sun heats IPD, causing it to emit thermally at infrared frequencies. The zodiacal emission can be distinguished from

other emission sources due to its strong seasonal dependence. In [Kelsall et al. \(1998\)](#) a full-sky parametric model of the zodiacal emission was constructed, which was used to correct the *Planck* sky maps. For the high frequency instrument (HFI) channels ($\nu \geq 100$ GHz) a model was subtracted from the time-ordered data (TOD) prior to map-making. For the low frequency instrument (LFI) maps a small correction was applied by extrapolating the effective 100 GHz zodiacal light map, using a power-law fit to the HFI maps between 100 GHz and 353 GHz. The primary effect of this correction was in removing a bias from the CO $J = 1 \rightarrow 0$ map at high Galactic latitudes, where J is the rotational quantum number ([Planck Collaboration et al., 2016d](#)).

1.3 CMB Component Separation Algorithms

Given a set of sky maps at a range of frequencies, one must then apply some algorithm to extract the CMB from these observations. A wide variety of techniques have been developed, including non-blind algorithms that make explicit assumptions regarding the spatial and/or frequency dependence of the foreground emission, and blind algorithms that attempt to extract the CMB with minimal assumptions about the foreground emission. In this section I summarise a set of common component separation methods i.e., template fitting, independent component analysis (ICA), internal linear combination (ILC) and parametric component separation.

1.3.1 Template Fitting

Template fitting is a simple component separation technique that proceeds by assuming the sky \mathbf{T} at some frequency ν can be described as the linear superposition of a set of spatial templates \mathbf{X}_i and noise \mathbf{n} . Here, we describe our sky and noise maps as vectors with a length given by the number of pixels. The noise term typically includes contributions from both instrumental noise and the CMB. The model sky is then given by

$$\mathbf{T} = \mathbf{n} + \sum_i \alpha_i \mathbf{X}_i, \quad (1.44)$$

where α_i are the template coefficients. The CMB map can be estimated as

$$\mathbf{T}_{\text{cmb}} = \mathbf{T} - \sum_i \alpha_i \mathbf{X}_i, \quad (1.45)$$

where the template coefficients are estimated by minimising,

$$\chi^2 = \left(\mathbf{T} - \sum_i \alpha_i \mathbf{X}_i \right)^\top \mathbf{C}^{-1} \left(\mathbf{T} - \sum_i \alpha_i \mathbf{X}_i \right). \quad (1.46)$$

The covariance matrix \mathbf{C} contains contributions from instrumental noise and the CMB. The CMB contribution is typically estimated from some fiducial CMB power spectrum.

Given that templates are fit across the whole sky, and by fully utilizing the spatial information in each template, the corresponding components can be constrained even in the regime that the individual pixel signal-to-noise ratio (SNR) is low. However, template fitting suffers from a key drawback in its assumption that the spatial and frequency scaling of each foreground component is completely separable. In reality, the spectral behaviour of foreground components show significant spatial dependence. Moreover, noise present in template maps will be aliased into the estimated CMB map, which can be particularly problematic in low signal regions of the sky (Dunkley et al., 2009).

Template fitting has seen numerous applications in the CMB literature, including in foreground removal for the COBE/DMR maps (Banday et al., 1996) and in the WMAP 5 year data (Komatsu et al., 2009). An implementation of template fitting, called SEVEM (Martínez-González et al., 2003), was also used in the component separation analysis for the *Planck* mission (Planck Collaboration et al., 2014a, 2016c,d, 2018b).

1.3.2 Internal Linear Combination

ILC is a blind component separation algorithm that was first employed in the harmonic domain in Tegmark and Efstathiou (1996). The method constructs a CMB estimate, $\hat{s}_{\ell m}$ as a linear combination of the observed $x_{\ell m}^i$ at each frequency,

$$\hat{s}_{\ell m} = \sum_i w_\ell^i x_{\ell m}^i, \quad (1.47)$$

where w_ℓ^i are a set of weights that are set such that we minimise the level of foreground contamination to the CMB. This can be achieved by choosing weights that minimise the power spectrum

$$\langle |x_{\ell m}|^2 \rangle = \mathbf{w}_\ell^\top \mathbf{C}_\ell \mathbf{w}_\ell, \quad (1.48)$$

where $C_\ell^{ij} = \langle (x_{\ell m}^i)^\dagger x_{\ell m}^j \rangle$. This is subject to the constraint that $\mathbf{w}_\ell \cdot \mathbf{e} = 1$, where \mathbf{e} is a vector containing the spectrum of the component we seek to isolate. In the

case of the CMB, working in CMB thermodynamic temperature units, this is simply a vector of ones. The weights obtained from this minimisation are given by

$$\mathbf{w}_\ell = \frac{\mathbf{C}_\ell^{-1} \mathbf{e}}{\mathbf{e}^\top \mathbf{C}_\ell^{-1} \mathbf{e}}. \quad (1.49)$$

The core advantages to this method are that it makes no assumptions about the frequency dependence of the foreground components and is computationally very efficient. However, it does make the unrealistic assumption that the spectral scaling of foregrounds is spatially uniform. This can be mitigated to some extent by performing the ILC on smaller sky regions. The algorithm also suffers drawbacks in that it minimises the total power spectrum, as opposed to the foreground power spectrum, the uncertainty in the recovered CMB map can be difficult to assess, and despite the apparent model-independence, the method leaves residuals in the recovered CMB map that must be modelled (Dunkley et al., 2009).

ILC can be improved by working in needlet space, a technique known as Needlet ILC (NILC) (Delabrouille et al., 2009; Basak and Delabrouille, 2013). Needlets are spherical wavelets that are characterised by the fact that they are perfectly localised in the harmonic domain and have potentially excellent localisation in the spatial domain. This means that weights estimated by NILC are able to smoothly vary as a function of position and angular scale. NILC was further extended in Remazeilles et al. (2011) to allow for the estimation of foreground components, and was used to construct maps of the thermal dust emission from *Planck* data (Planck Collaboration et al., 2018b).

1.3.3 Independent Component Analysis

Independent component analysis (ICA) proceeds in a similar fashion to ILC, constructing a CMB estimate as a linear combination of input sky maps, but utilizes different criteria for constructing the weights. One of the most developed implementations is Spectral Matching ICA (SMICA) (Delabrouille et al., 2003; Cardoso et al., 2008), which was used in the analysis of *Planck* data and produced the smallest residuals of any of the component separation methods (Planck Collaboration et al., 2014a).

SMICA performs component separation in the harmonic domain, constructing a CMB estimate as a linear combination of N_{chan} frequency maps,

$$\hat{s}_{\ell m} = \mathbf{w}_\ell^\dagger \mathbf{x}_{\ell m}. \quad (1.50)$$

$\mathbf{x}_{\ell m}$ is a $N_{\text{chan}} \times 1$ vector of spherical harmonic coefficients for the input maps and \mathbf{w}_ℓ is the $N_{\text{chan}} \times 1$ vector of weights given by

$$\mathbf{w}_\ell = \frac{\mathbf{C}_\ell^{-1} \mathbf{a}}{\mathbf{a}^\dagger \mathbf{C}_\ell^{-1} \mathbf{a}}, \quad (1.51)$$

where \mathbf{a} is the CMB spectrum evaluated at each frequency channel. If we were to set the power spectrum to

$$\hat{\mathbf{C}}_\ell = \frac{1}{2\ell + 1} \sum_m \mathbf{x}_{\ell m} \mathbf{x}_{\ell m}^\dagger, \quad (1.52)$$

we would simply perform a standard ILC. For SMICA the power spectrum is instead obtained by fitting a parametric model $\mathbf{C}_\ell(\theta)$, to $\hat{\mathbf{C}}_\ell$. The best fit parameters are obtained as,

$$\hat{\theta} = \arg \min_\theta \sum_\ell (2\ell + 1) \left[\text{Tr}(\hat{\mathbf{C}}_\ell \mathbf{C}_\ell^{-1}(\theta)) + \log \det \mathbf{C}_\ell(\theta) \right]. \quad (1.53)$$

This is the maximum likelihood solution under the assumption that the spherical harmonic coefficients are independent, normally distributed random variables with variance given by the corresponding power spectrum i.e., $\mathbf{x}_{\ell m} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_\ell)$.

For SMICA, the foreground emission is modelled by N_{fg} templates. The full parametric model for the power spectra is given by

$$\mathbf{C}_\ell(\theta) = \mathbf{a} \mathbf{a}^\dagger C_\ell^{\text{cmb}} + \mathbf{F} \mathbf{P}_\ell \mathbf{F}^\dagger + \mathbf{N}_\ell, \quad (1.54)$$

where C_ℓ^{cmb} is the CMB power spectrum, \mathbf{F} is a $N_{\text{chan}} \times N_{\text{fg}}$ matrix that encodes the frequency dependence of the foregrounds, \mathbf{P}_ℓ is a positive definite $N_{\text{fg}} \times N_{\text{fg}}$ foreground covariance matrix and \mathbf{N}_ℓ is the diagonal noise matrix.

To perform a fully flexible component separation one could fit the power spectrum model without imposing any constraints on the terms in Equation 1.54. However, in reality this typically proves too challenging due to the large dynamic range over the sky and multipoles considered for CMB experiments. For the *Planck* analysis a three step approach was adopted, first fitting for \mathbf{a} over quieter regions of the sky, and then fixing \mathbf{a} whilst fitting for \mathbf{F} over a large sky area. In the final step \mathbf{a} and \mathbf{F} were fixed, and the C_ℓ^{cmb} , \mathbf{P}_ℓ and \mathbf{N}_ℓ power spectra were fitted (Planck Collaboration et al., 2014a, 2016c).

1.3.4 Parametric Component Separation

At its core parametric component separation simply consists of fitting some model of the sky emission to our observed sky maps. For some pixel p at a frequency ν , we may write the observed pixel value as

$$d_{p,\lambda}(\nu) = s_{p,\lambda}(\nu) + n_{p,\lambda}(\nu), \quad (1.55)$$

where $\lambda = \{I, Q, U\}$ denotes the Stokes parameter under consideration, $s_{p,\lambda}(\nu)$ is the true sky signal and $n_{p,\lambda}(\nu)$ is the noise term. In the broadest sense, performing component separation can now be achieved by writing some parametrisation of the true sky emission, $s_{p,\lambda}(\nu, \Theta_{p,\lambda})$ where $\Theta_{p,\lambda}$ are the model parameters.

A range of methods and statistical frameworks can be applied to this problem of model fitting. One of the most notable implementations of parametric component separation is the COMMANDER algorithm (Eriksen et al., 2004, 2008). This adopts a Bayesian approach to the problem, constructing estimators for our model parameters by integrating over the posterior distribution

$$p(\Theta|\mathbf{d}) = \frac{p(\mathbf{d}|\Theta)p(\Theta)}{p(\mathbf{d})}. \quad (1.56)$$

Here I have used Θ to represent the full set of model parameters over the whole sky, and \mathbf{d} to represent the full dataset of sky maps. The term $p(\mathbf{d}|\Theta)$ is the likelihood, $p(\Theta)$ is the prior, which encodes our prior knowledge of the model parameters, and $p(\mathbf{d})$ is known as the evidence. For the purposes of parameter estimation with a single model, $p(\mathbf{d})$ can simply be viewed as a normalisation constant.

For simple, low dimensional models one can either write analytical posterior distributions or attempt brute force, grid-based calculations of the posterior. However, this quickly becomes infeasible as the dimension of the parameter space increases. COMMANDER employs Gibbs sampling, a particular flavour of Markov Chain Monte Carlo (MCMC) algorithm (Geman and Geman, 1984). Generally MCMC algorithms proceed by constructing a Markov chain with the desired stationary distribution. By simulating the Markov chain one obtains a set of correlated samples and, given sufficient computation and the satisfaction of certain ergodicity conditions, we are guaranteed to generate samples from our target distribution. The ergodicity conditions ensure that the MCMC algorithm is able to fully explore the target distribution. Poorly designed transitions between states in the Markov chain can lead to certain regions of the parameter space being inaccessible (Gelman et al., 2004). In Chapter 5

I discuss MCMC algorithms further, along with alternative methods for approximate Bayesian inference.

For Gibbs sampling, rather than drawing samples directly from the full CMB and foreground posterior, one instead draws successive parameter samples from conditional distributions,

$$C_\ell^{i+1} \leftarrow p(C_\ell | \mathbf{A}_{\text{cmb}}^i, \theta_{\text{fg}}^i, \mathbf{d}), \quad (1.57)$$

$$\mathbf{A}_{\text{cmb}}^{i+1} \leftarrow p(\mathbf{A}_{\text{cmb}} | C_\ell^{i+1}, \theta_{\text{fg}}^i, \mathbf{d}), \quad (1.58)$$

$$\theta_{\text{fg}}^{i+1} \leftarrow p(\theta_{\text{fg}} | C_\ell^{i+1}, \mathbf{A}_{\text{cmb}}^{i+1}, \mathbf{d}). \quad (1.59)$$

Here C_ℓ is the CMB power spectrum, \mathbf{A}_{cmb} denotes the CMB sky map and θ_{fg} is the set of foreground parameters (Dunkley et al., 2009). Given a sufficient number of samples this sampling scheme will converge on the full target distribution, $p(C_\ell, \mathbf{A}_{\text{cmb}}, \theta_{\text{fg}} | \mathbf{d})$. This target distribution is extremely high-dimensional, with large numbers of correlated parameters. By reducing the problem of sampling from the full posterior to one of sampling from conditional distributions, COMMANDER is able to take a computationally tractable approach. Further, COMMANDER obtains posterior samples of the CMB C_ℓ , allowing for the rigorous propagation of uncertainty from sky maps to the power spectrum. Recently, COMMANDER2 was used in the component separation analysis for the *Planck* 2018 release, which instead performs parametric modelling in the harmonic domain (Planck Collaboration et al., 2018b). In Chapter 5 I discuss a novel implementation of Bayesian parametric component separation presented in Grumitt et al. (2020a), utilizing hierarchical foreground modelling (Gelman and Hill, 2007) and the No-U-Turn Sampler (NUTS) (Hoffman and Gelman, 2014) for direct posterior sampling.

In addition to the Bayesian approach outlined above, one can also adopt a maximum likelihood approach (Stompor et al., 2009; Errard et al., 2011; Stompor et al., 2016). This is equivalent to assigning flat priors in the Bayesian approach, and performing a numerical optimization to find the parameter maximum likelihood estimates (MLEs). One could also approach this as an optimization problem with non-flat priors, and optimize for the maximum a-posteriori (MAP) estimate. The MAP approach can also be viewed as a form of regularized maximum-likelihood, with the prior in this case acting as an additional penalty term.

The key advantage of these approaches lies in their computational efficiency. In essence, the problem of performing numerical integration with MCMC is turned into a differentiation problem. However, this comes at the cost of not obtaining the same rigorous propagation of uncertainty from performing inference over the full posterior.

Further, whilst MLEs can sometimes be shown to be unbiased, minimum variance estimators, this is on the basis of asymptotic arguments i.e., in the regime of infinite data. Foreground analysis for CMB polarization experiments sits firmly outside of this asymptotic regime and without the regularizing effect of the prior, noisy observations can lead to poor MLEs. However, this is not necessarily alleviated by using the MAP estimate. In very high dimensional models the MAP value is often located far away from the typical set, where most of the posterior probability mass is located (Mackay, 2003; Gelman et al., 2008; Hamra et al., 2013; Cole et al., 2013). In this situation, one can obtain more reliable inferences by calculating expectations with respect to the posterior, which fully account for the distribution of probability mass.

1.4 CMB Experiments

Since the initial discovery of the CMB by Penzias and Wilson, there have been a large number of dedicated CMB experiments. These have consisted of ground, balloon, and space-based experiments. A comprehensive list of CMB experiments can be found at the Legacy Archive for Microwave Background Data Analysis (LAMBDA) (NASA, 2020). In this section I give an overview of past and present CMB experiments, along with related ancillary surveys.

1.4.1 CMB Temperature and E -mode Observations

The blackbody nature of the CMB spectrum was first confirmed by the COBE satellite, which also made the first statistical detection of temperature anisotropies in the CMB (Mather et al., 1990; Smoot et al., 1992; Mather et al., 1994; Wright et al., 1994b). Many of the key features of the CMB TT , TE and EE power spectra were first measured by a collection of ground and balloon experiments. The location of the first acoustic peak in the TT power spectrum was precisely determined by the TOCO (Devlin et al., 1999), BOOMERANG (de Bernardis et al., 2000) and MAXIMA (Hanany et al., 2000) experiments. Further measurements of the TT acoustic peaks and damping tail were made by the DASI (Pryke et al., 2002), VSA (Watson et al., 2003; Taylor et al., 2003; Scott et al., 2003), CBI (Padin et al., 2001, 2002) and Archeops (Benoît et al., 2002, 2003; Tristram et al., 2005a) experiments. The first measurements of CMB polarization anisotropy were made by DASI (Leitch et al., 2002; Kovac et al., 2002), and a significant detection of E -mode polarization by CBI

demonstrated that the peaks of the EE power spectrum are out of phase with the TT acoustic peaks (Readhead et al., 2004).

The WMAP satellite was launched in 2001 and made precise measurements of the CMB TT power spectrum over the multipole range $2 \leq \ell \leq 1000$, and detected the EE power spectrum over the range $24 \leq \ell \leq 800$ at more than 8σ . WMAP observations provided strong support for the standard Λ CDM model, making precise determinations of the six base parameters (Bennett et al., 2013; Hinshaw et al., 2013). The *Planck* satellite was launched in 2009 and made incredibly detailed measurements of the CMB TT power spectrum over multipoles $2 \leq \ell \leq 2508$, and the TE and EE power spectra over multipoles $2 \leq \ell \leq 1996$ (Planck Collaboration et al., 2019). Fig. 1.8 shows the measured CMB TT , TE and EE power spectra (Planck Collaboration et al., 2018c, 2019). Note that the TT and TE power spectra are plotted as


$$D_\ell^{XY} = \frac{\ell(\ell+1)}{2\pi} C_\ell^{XY}, \quad XY \in \{TT, TE\}. \quad (1.60)$$

Plotted alongside each measured power spectrum is the best fit Λ CDM model. This standard cosmological model can be described by six base parameters, which are summarised in Table 1.1. The *Planck* measurements again show excellent agreement with Λ CDM.


Parameter	Measured Value	Description
$\Omega_b h^2$	0.02237 ± 0.00015	Fractional density of baryons today
$\Omega_c h^2$	0.1200 ± 0.0012	Fractional density of cold dark matter today
$100\theta_{\text{MC}}$	1.04092 ± 0.00031	$100\times$ approximation to the angular size of the sound horizon at last scattering
τ	0.0544 ± 0.0073	Optical depth to reionization
$\ln(10^{10} A_s)$	3.044 ± 0.014	Log power of primordial curvature perturbations
n_s	0.9649 ± 0.0042	Scalar spectral index

Table 1.1: *Planck* 2018 constraints on the base Λ CDM parameters. Note that the power spectrum of primordial curvature perturbations was parametrised as $P_\Phi(k) = A_s(k/k_0)^{n_s-1}$ for this base Λ CDM analysis, where $k_0 = 0.05 \text{ Mpc}^{-1}$ was the chosen pivot scale. Table adapted from Planck Collaboration et al. (2018c).


In addition to these recent space missions, ground-based telescopes such as ACT (Swetz et al., 2011) and SPT Shirokoff et al. (2009) have made measurements of the TT power spectrum over small angular scales (up to $\ell = 10000$ for ACT). This has allowed for the characterisation of the Silk damping tail up to $\ell \approx 3000$, along with



(a) *Planck* 2018 TT spectrum.



(b) *Planck* 2018 TE spectrum.



(c) *Planck* 2018 EE spectrum.

Figure 1.8: The *Planck* 2018 CMB TT , TE and EE power spectra. A detailed description of the power spectrum estimation is given in [Planck Collaboration et al. \(2019\)](#). The lower panels show residuals from the fit. The *Planck* results are consistent with the base, six parameter Λ CDM model. Figures taken from [Planck Collaboration et al. \(2018c\)](#).

measurements of the contributions from the thermal and kinetic SZ effects to the TT power spectrum over multipoles $\ell \gtrsim 3000$ ([Keisler et al., 2011](#); [Das et al., 2014](#)).

1.4.2 CMB B -mode Observations

The core focus of modern CMB cosmology is on the detection of primordial B -mode polarization and hence a determination of r . Thus far, this effort has been led by ground and balloon experiments. The first direct detection of non-primordial CMB B -mode polarization was made by the POLARBEAR experiment, which measured the B -mode power spectrum over the multipole range $500 < \ell < 2100$ ([Polarbear Collaboration et al., 2014](#)). Further measurements of small scale polarization anisotropies have been made by ACTPol ([Louis et al., 2017](#)) and SPTPol ([Hanson et al., 2013](#);

Keisler et al., 2015), with the most precise measurements of lensing B -modes to date coming from the SPTPol experiment (Sayre et al., 2019).

The first experiment to specifically target measurements of the primordial B -mode power spectrum was the BICEP1 experiment (Barkats et al., 2014). This was followed by the BICEP2 and Keck Array experiments, which in 2014 reported a measurement of $r = 0.20_{-0.05}^{+0.07}$ (BICEP2 Collaboration et al., 2014). However, a subsequent joint analysis with *Planck* data found that the observed excess BB power could be entirely explained by polarized dust emission (Mortonson and Seljak, 2014; Flauger et al., 2014; BICEP2/Keck Collaboration et al., 2015). The most recent constraints on r were obtained through a joint analysis of the *Planck* 2018 data and the BICEP2/Keck 2015 data, giving an upper limit of $r < 0.056$ (Planck Collaboration et al., 2018d).

A range of next generation ground-based CMB experiments are under construction with a detection of r amongst their chief science goals. CMB-S4 is a set of telescopes based in the South Pole, the Atacama Desert and potentially additional northern hemisphere sites (Abazajian et al., 2016). This project aims to be able to achieve upper limits of $r < 2 \times 10^{-3}$. Alongside this, the Simons Observatory is being constructed in the Atacama desert and aims to measure r with a sensitivity of $\sigma(r) = 3 \times 10^{-3}$ (Simons Observatory Collaboration et al., 2019). In addition to these ground-based experiments, a range of space-based missions have been proposed, including *LiteBIRD* (Hazumi et al., 2019), *COrE* (The COrE Collaboration et al., 2011) and *PIXIE* (Kogut et al., 2011). Of these proposals, *LiteBIRD* has been selected as an L-class mission by the Japan Aerospace Exploration Agency (Hazumi, 2019), and aims to measure r with a sensitivity of $\sigma(r) = 10^{-3}$.

1.4.3 Ancillary Surveys

Conducting accurate CMB component separation, particularly for the challenge of detecting primordial B -modes, requires detailed information regarding the structure of foreground emission. This comes in the form of ancillary surveys, conducted at low and high frequencies either side of the CMB foreground minimum (~ 70 GHz). These surveys provide data on the spatial and spectral distributions of foreground components, allowing us to properly constrain foreground emission parameters.

Table 1.2 summarises a set of completed and ongoing surveys at frequencies relevant to CMB component separation analyses. Of key focus in this thesis is C-BASS (Jones et al., 2018), a 4.76 GHz experiment that will survey the entire sky in Stokes I , Q and U , with a 45 arcmin resolution. At this frequency, the sky emission is

Table 1.2: Existing and on-going large area surveys, at frequencies relevant to CMB studies, with angular resolutions $\lesssim 1^\circ$. Table taken from [Jones et al. \(2018\)](#).

Survey / Telescope	Frequency [GHz]	FWHM [arcmin]	Declination Coverage	Stokes ^a	Sensitivity ^b		Status ^c	Reference(s)
					noise	offsets		
Haslam (various)	0.408	51	All-sky	<i>I</i>	1 K	3 K	3	Haslam et al. (1982)
Dwingeloo	0.82	72	-7° to $+85^\circ$	<i>I</i>	0.2 K	0.6 K	3	Berkhuijsen (1972)
CHIPASS (Parkes)	1.394	14.4	$< +25^\circ$	<i>I</i>	0.6 mK	30 mK	3	Calabretta et al. (2014)
DRAO (26-m) ^d	1.4	36	$> -29^\circ$	<i>QU</i>	12 mK	30 mK	3	Wolleben et al. (2006)
Villa Elisa ^d	1.4	35.4	$< +10^\circ$	<i>IQU</i>	9 mK	50 mK	3	Testori et al. (2008)
Stockert ^d	1.42	35	$> -30^\circ$	<i>I</i>	9 mK	50 mK	3	Reich and Reich (1986)
GMIMS-HB N	1.28–1.75	30	$> -30^\circ$	<i>IQU</i>	12 mK	unknown	1	Wolleben et al. (2010)
STAPS (Parkes)	1.3–1.8	15	$< 0^\circ$	<i>IQU</i>	unknown	unknown	1	Haverkorn (priv. comm.)
HartRAO	2.326	20	-83° to $+13^\circ$	<i>I–Q</i>	25 mK	80 mK	3	Jonas et al. (1998)
S-PASS (Parkes)	2.3	9	$< 0^\circ$	<i>IQU</i>	0.1 mK	unknown	1	Carretti et al. (2013)
GEM	4.8–5.2	45	-52° to $+7^\circ$	<i>QU</i>	0.5 mK	unknown	0	Barbosa et al. (2006) ; Tello et al. (2013)
C-BASS	4.5–5.5	45	All-sky	<i>IQU</i>	0.1 mK	1 mK	0	Jones et al. (2018)
QUIJOTE	11–19,30,40	≈ 60	$\gtrsim 0^\circ$	<i>[I]QU</i>	25 μ K	unknown	1	Génova-Santos et al. (2015b)
WMAP	22.8–94	49–15	All-sky	<i>IQU</i>	4 μ K	1 μ K	3	Bennett et al. (2013)
Planck LFI	28.4–70	32–13	All-sky	<i>IQU</i>	3 μ K	1 μ K	3	Planck Collaboration et al. (2018a)
Planck HFI	100–353	10–5	All-sky	<i>IQU</i>	0.2–0.5 μ K	1–5 μ K	3	Planck Collaboration et al. (2018a)
Planck HFI	545, 857	5	All-sky	<i>I</i>	0.4, 0.8 μ K	1 μ K	3	Planck Collaboration et al. (2018a)
CLASS	38–217	90–18	-68° to $+22^\circ$	<i>QU</i>	0.4 μ K	unknown	0	Harrington et al. (2016)

^a [I]QU denotes surveys where total intensity (Stokes I) is measured but with much larger systematic errors than for the linear polarization (Stokes Q and U). I–Q denotes a single linear polarization.

^b Approximate average total intensity sensitivity in Rayleigh-Jeans temperature after convolution to 1° FWHM resolution: “noise” is local rms; “offsets” is global systematic uncertainty.

^c Status 0: observations ongoing; 1: observations complete, reduction in progress; 2: preliminary results released; 3: Final data released.

^d An all-sky 1.4 GHz map in IQU has been assembled from the Stockert, DRAO and Villa Elisa surveys ([Reich et al., 2004](#); [Testori et al., 2008](#)), but full details of its construction have not been published, and it is not clear if the currently-available version is the final one.

dominated by synchrotron and free-free emission in total intensity, and synchrotron emission in polarization. C-BASS will allow us to better constrain the spatial and spectral variations in synchrotron emission, which will be critical to avoid biases in any determinations of r at targeted sensitivities of $\sigma(r) = 10^{-3}$ ([Errard et al., 2016](#); [Remazeilles et al., 2016](#); [Jew et al., 2019](#)). To date, the northern sky survey has been completed, with the southern sky survey ongoing. Further details on the C-BASS survey and data are given in Chapter 2.

In addition to C-BASS, the S-PASS experiment has observed the southern sky at 2.3 GHz with the Parkes telescope, at a resolution of 9 arcmin ([Carretti, 2010](#)). Analysis of S-PASS data in [Krachmalnicoff et al. \(2018\)](#) has shown that the minimum level of synchrotron contamination to CMB B -modes over the southern sky is at the level of $r \sim 10^{-3}$. The QUIJOTE experiment is an ongoing survey, mapping several sky regions at frequencies between 10 and 40 GHz ([Génova-Santos et al., 2015b](#)). A

recent joint analysis of C-BASS and QUIJOTE data in [Cepeda-Arroita et al. \(2020\)](#) led to the first detection of spectral variations in AME across a single sky region.

1.5 Probability Distributions

In this section I define the probability distributions used in the statistical analyses presented throughout this thesis.

- **Gaussian:** The Gaussian or Normal distribution is parameterized by a mean, μ and standard deviation σ . I indicate Gaussian distributed random variables as $x \sim \mathcal{N}(\mu, \sigma)$, with the full probability density function (PDF) of the distribution being given by,

$$p(x|\mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right). \quad (1.61)$$

The support of the distribution is over $x \in \mathbb{R}$.

- **Half-Normal:** The Half-Normal distribution is parameterized by σ , which is the standard deviation of the unfolded Normal distribution. The standard deviation of the Half-Normal distribution is given by $\sigma\sqrt{1 - 2/\pi}$. I indicate random variables distributed according to the Half-Normal distribution as $x \sim \text{Half-Normal}(\sigma)$, with the full PDF of the distribution being given by,

$$p(x|\sigma) = \sqrt{\frac{2}{\pi\sigma^2}} \exp\left(-\frac{x^2}{2\sigma^2}\right). \quad (1.62)$$

The support of the Half-Normal distribution is over $x \in [0, \infty)$.

- **Exponential:** The Exponential distribution is parameterized by the rate parameter, λ . I indicate Exponentially distributed random variables as $x \sim \text{Exp}(\lambda)$, with the full PDF of the distribution being given by,

$$p(x|\lambda) = \lambda \exp(-\lambda x). \quad (1.63)$$

The support of the distribution is over $x \in [0, \infty)$.

- **von Mises:** The von Mises distribution is parameterized by a mean, μ and concentration parameter, κ . I indicate random variables distributed according to the von Mises distribution as $x \sim \text{VonMises}(\mu, \kappa)$, with the full PDF of the distribution being given by,

$$p(x|\mu, \kappa) = \frac{\exp(\kappa \cos(x - \mu))}{2\pi I_0(\kappa)}, \quad (1.64)$$

where $I_0(\kappa)$ is the modified Bessel function of order 0. The support of the von Mises distribution can be chosen to be over any interval of length 2π . For its use in this thesis the support is chosen to be over $x \in [-\pi, \pi]$.

- **Uniform:** The uniform distribution is parameterized by a lower limit, a and an upper limit b . I indicate uniformly distributed random variables as $x \sim \text{Unif}(a, b)$. The uniform distribution assigns equal probability to all values of the random variable between the upper and lower limits, and zero otherwise. The support of the uniform distribution is over $x \in [a, b]$. For $\text{Unif}(-\infty, 0)$, $\text{Unif}(0, \infty)$ and $\text{Unif}(-\infty, \infty)$ we obtain the improper half-flat or flat distributions.

1.6 Thesis Outline

This thesis addresses challenges in CMB component separation, particularly in the context of CMB polarization studies. The component separation challenge is approached from two core directions. In the first instance, I present work analysing a novel dataset from the C-BASS experiment. With additional, high-fidelity observations covering a wide range of frequencies, we are able to better characterize foreground emission and reliably extract CMB estimates from our observations. Alongside this, I present a new approach to parametric Bayesian CMB component separation, utilizing NUTS to perform scalable hierarchical foreground modelling.

In Chapter 2 I discuss the C-BASS experiment, summarising instrument and survey details. I describe key data processing steps involved in producing the final, calibrated and deconvolved maps of the northern sky. I also present data quality checks, performed to confirm the internal stability of the C-BASS calibration, and the current science-quality maps of the northern sky.

In Chapter 3 I discuss the work presented in [Grumitt et al. \(2020b\)](#), describing the production of the C-BASS northern sky point source catalogue. Producing a catalogue of point sources for the C-BASS map is important for any analyses of diffuse emission, such that we can construct properly masked and/or source subtracted C-BASS maps. Further, the source catalogue allows us to perform vital data checks on the C-BASS flux-density scale and pointing accuracy.

In Chapter 4 I present some of the first scientific analyses of C-BASS polarization data. Here, I calculate the auto and cross-spectra of C-BASS, WMAP and *Planck*

polarization maps, using these to constrain the spectral behaviour of polarized synchrotron emission in the harmonic domain. This analysis is combined with map domain estimates of the polarized synchrotron spectral index to estimate the level of synchrotron contamination to CMB B -modes over the northern sky.

In Chapter 5 I discuss the work presented in [Grmitt et al. \(2020a\)](#). This presents a novel approach to Bayesian CMB component separation, utilizing hierarchical modelling of foreground spectral parameters to eliminate biases on recovered CMB estimates without inflating the statistical uncertainty in these estimates. In order to perform scalable inference over such complex, hierarchical posterior distributions I utilize NUTS. This is a self tuning variant of Hamiltonian Monte Carlo (HMC) that offers excellent computational performance for such models.

In Chapter 6 I finish by summarizing the work presented in this thesis and future steps that can be taken to extend this work.

Chapter 2

The C-Band All-Sky Survey

Up to now, one of the most important synchrotron emission templates has been the 408 MHz Haslam map ([Haslam et al., 1981, 1982](#)). An improved version of the map, with stripe artefacts and bright point source emission removed was produced in [Remazeilles et al. \(2015\)](#). However, the Haslam map still suffers from scan-synchronous artefacts, and is only in total intensity. C-BASS is an experiment observing the whole sky at 4.76 GHz in total intensity and polarization ([Jones et al., 2018](#)). Northern survey observations have been completed using a telescope located at the Owens Valley Radio Observatory (OVRO) in California. Detailed scientific analyses of these data are underway, some of which is presented in Chapters [3](#) and [4](#). The southern survey is ongoing, using a telescope located at Klerefontein in South Africa, with observations ongoing. The primary purpose of the C-BASS experiment is to provide high signal-to-noise measurements of polarized Galactic synchrotron emission. These data will form a key element of component separation analyses for next-generation CMB *B*-mode studies, as well as providing a wealth of information for studies of Galactic astrophysics.

In this chapter I give a general overview of the C-BASS experiment. In Section [2.1](#) I summarize the core scientific motivation behind C-BASS. In Section [2.2](#) I discuss the survey constraints and strategy. In Section [2.3](#) I describe the C-BASS instruments. In Section [2.4](#) I describe the key data processing steps, and in Section [2.5](#) I present a number of data quality checks I have performed on the C-BASS data. I finish in Section [2.6](#) by summarizing the current survey status and presenting the latest science-quality maps of the northern sky. A detailed overview of the C-BASS experiment can be found in [Jones et al. \(2018\)](#).

2.1 Science Goals

2.1.1 Polarized CMB Foregrounds

Since observational studies of the CMB anisotropies began, foreground emission has been a major concern. In temperature, the CMB fluctuations are actually found to be dominant over foreground emission away from the Galactic plane, over a wide range of frequencies from ~ 15 GHz to ~ 150 GHz (Planck Collaboration et al., 2014a, 2016d, 2018b). However, for studies of the primordial B -mode signal, current constraints on the tensor-to-scalar ratio mean that the CMB is expected to be subdominant to foregrounds across all of the sky, at all frequencies (Remazeilles et al., 2018; BICEP2 Collaboration et al., 2018). Extracting the polarized CMB signal will therefore require careful algorithmic approaches that are able to produce high fidelity maps of the CMB, with an accurate quantification of their uncertainty. This must also be coupled with improved, high-sensitivity data covering a wide range of frequencies, such that we can accurately model foreground spectral energy distributions (SEDs). If it becomes clear that more complex foreground modelling is required to properly extract the CMB signal, this will inevitably require additional data in order to constrain the more complex emission models (see e.g. Remazeilles (2018); Jew et al. (2019)).

The primary goal of C-BASS is to provide high-sensitivity data at low frequencies, that will allow us to better constrain polarized synchrotron emission. Much of the focus regarding polarized CMB foregrounds has been on thermal dust emission, given that this is the dominant polarized component at frequencies close to the foreground minimum (~ 90 GHz) (BICEP2 Collaboration et al., 2014; BICEP2/Keck Collaboration et al., 2015; BICEP2 Collaboration et al., 2018). However, with targeted sensitivities of next-generation CMB experiments of $\sigma(r) \sim 10^{-3}$, it has become clear that synchrotron emission cannot be neglected without significantly biasing the extraction of any primordial signal (Remazeilles et al., 2016; Errard et al., 2016). In Jew et al. (2019) it was found that next-generation satellites such as *LiteBIRD* will not be able to constrain simple power-law models of synchrotron emission without additional low-frequency data from experiments such as C-BASS. This issue is explored in detail in Chapters 4 and 5. At its core, C-BASS will provide a vital new template of polarized synchrotron emission and allow the next-generation of CMB experiments to constrain the spectral behaviour of synchrotron emission.

2.1.2 Secondary Goals

In addition to the core motivation described above, C-BASS will also provide a wealth of additional astrophysical results e.g., [Dickinson et al. \(2019\)](#); [Cepeda-Arroita et al. \(2020\)](#). One particularly interesting feature is the WMAP Haze ([Finkbeiner, 2004a](#); [Dobler and Finkbeiner, 2008](#); [Dobler, 2012a](#); [Planck Collaboration et al., 2013](#)). This is a region of excess emission close to the Galactic Centre, which has been measured to have a total intensity spectral index of $\beta \approx -2.5$. The haze is thought to be caused by shallow spectrum synchrotron emission, and appears to be spatially correlated with the Fermi gamma-ray bubbles ([Dobler et al., 2010](#); [Su et al., 2010](#); [Dobler, 2012b](#)). In Chapter 4 I briefly summarize the results in [Jew and Grumitt \(2020\)](#) related to the WMAP Haze, where we studied its spectral properties using *Planck* polarization data. With the addition of C-BASS, we will be better able to constrain the spectral properties of the haze and the extent to which it is correlated with the Fermi Bubbles.

C-BASS total intensity data will also enable us to better constrain the properties of free-free emission and AME. In bright H II regions close to the Galactic plane, free-free emission dominates over synchrotron emission in total intensity at ~ 5 GHz ([Planck Collaboration et al., 2014c](#)). However, there is a degeneracy between the free-free spectrum parameters and the synchrotron spectrum parameters, which can make accurate component separation challenging. This is particularly problematic at high Galactic latitudes where the diffuse free-free background is much weaker ([Dickinson et al., 2003](#); [Finkbeiner, 2003](#); [Draine, 2011](#)). The addition of C-BASS will allow us to break this degeneracy and better constrain free-free emission across the whole sky. Whilst C-BASS observes at too low a frequency to make significant detections of AME, by providing a low-frequency template largely free of AME, C-BASS allows for a more accurate separation of AME from synchrotron and free-free emission ([Cepeda-Arroita et al., 2020](#)).

C-BASS polarization data will enable a detailed study of the Galactic Magnetic Field (GMF). This is manifested through Faraday rotation (rotation of electromagnetic wave polarization angles by magnetic fields) and the polarization of Galactic synchrotron emission, which is orthogonal to the projected magnetic field direction. The strength of Faraday rotation is given by the rotation measure, RM. In the Faraday-thin regime (where the polarized emission and Faraday rotation occur in two, physically separated locations) the polarization angle-wavelength relation is given by

$$\chi(\lambda) = \chi_0 + \text{RM} \lambda^2, \quad (2.1)$$

where $\chi(\lambda)$ is the polarization angle at a wavelength λ , and χ_0 is the intrinsic polarization angle. The Global Magneto-Ionic Medium Survey (GMIMS) is a project mapping the whole sky with continuous frequency coverage between 0.3 – 1.8 GHz, that will allow for detailed characterization of Galactic RMs (Wolleben et al., 2009, 2010). Combining GMIMS data with C-BASS and S-PASS data will allow us to produce maps of the Faraday rotation of polarized Galactic synchrotron emission and hence probe the GMF.

2.2 Survey Constraints and Strategy

Table 2.1: Key specifications of the C-BASS survey. Table taken from Jones et al. (2018).

	North	South
Location	OVRO California	Klerefontein South Africa
Latitude	37° 14'N,	30° 58'S,
Longitude	118° 17'W	21° 59'E
Telescope	6.1 m Gregorian	7.6 m Cassegrain
Sky Coverage	$\delta > -15^\circ.6$	$\delta < 28^\circ.6$
Frequency range	4.5 – 5.5 GHz	
Effective centre frequency	4.783 GHz	5 GHz
Effective bandwidth	0.499 GHz	1.0 GHz
Frequency channels	1	128
Angular resolution	45 arcmin FWHM	
Stokes coverage	$I, Q, U(V)$	
Sensitivity	$\lesssim 0.1$ mK r.m.s. (per beam)	

Table 2.1 gives the key specifications of the C-BASS survey, which are informed by the scientific motivation outlined above. The C-BASS observing frequency was chosen in order to provide polarized emission maps that are dominated by Galactic synchrotron emission, whilst having low levels of Faraday rotation. In Jones et al. (2018), the rms RM values for extragalactic sources in Taylor et al. (2009) were used to estimate the expected levels Faraday rotation at low frequencies. At 5 GHz we would expect rotations of $\sim 3^\circ$ at high Galactic latitudes ($|b| > 30^\circ$), and $\sim 9^\circ$ at low Galactic latitudes ($|b| < 30^\circ$). By comparison, at 2.3 GHz the expected rotations are $\sim 14^\circ$ at high Galactic latitudes and $\sim 42^\circ$ at low Galactic latitudes. The smaller

levels of Faraday rotation at 5 GHz motivate this choice of observing frequency, but also demonstrate the importance of correcting for this effect in any multi-frequency analyses.

The required resolution is determined by the fact that C-BASS needs to be able to constrain the synchrotron signals at angular scales corresponding to the peak of the primordial B -mode power spectrum, at $\ell \sim 80$. This sets the minimum angular resolution required at $\sim 1^\circ$. The C-BASS resolution is also constrained by the size of the available antenna and the need to under-illuminate it in order to minimize sidelobes. Given the northern survey was constrained to use an existing 6.1 m dish it was possible to achieve a 45 arcmin resolution for the northern survey.

The C-BASS sensitivity target was set by extrapolating the WMAP K-band polarization map to 5 GHz using a spatially uniform spectral index of $\beta = -3$. From this it was estimated that polarized emission is expected to be brighter than ~ 0.5 mK over 90 per cent of the sky. The C-BASS target sensitivity was therefore set to 0.1 mK per beam in polarization. This corresponds to a flux density sensitivity of 14 mJy. However, the confusion level in the C-BASS total intensity map is approximately 85 mJy (Jones et al., 2018). Below this level sources cannot be individually detected, acting as an additional noise contribution, meaning the C-BASS total intensity map is confusion limited.

The C-BASS survey requires two instruments, one in the northern and one in the southern hemisphere to cover the whole sky. There is significant overlap between the northern and southern surveys, which will allow for detailed calibration cross-checks. The C-BASS scan strategy consists of conducting a set of constant-elevation scans over the entire azimuth range, at the maximum possible slew rate (4 deg/sec). The telescope scans from 0° to 360° before decelerating, stopping, and turning around. This ensures that the entire sky is scanned at the maximum slew rate. For 60 per cent of the observing time, scans are conducted at an elevation equal to the latitude of the observing site. For these scans the telescope always passes through the celestial pole, allowing for cross-checks on any drifts in offsets during observations. However, these scans have very uneven sky coverage, being primarily focused at the poles and the lower declination limit. To ensure sufficient integration time in every pixel on the sky, additional constant-elevation scans are conducted at 10° , 30° and 40° above the elevation of the pole.

When scanning at a constant elevation, the ground spillover signal is a fixed function of azimuth. Further, a given declination on the sky is observed at two fixed azimuths, situated symmetrically about the meridian. This results in a degeneracy

between the ground spillover signal and the spherically symmetric $m = 0$ sky modes. This is partly mitigated by observing at different elevations, which have different ground spillover signals, and by using the overlap region between C-BASS north and C-BASS south, which will also have different ground spillover signals. However, the fitting of ground templates to the raw data, is a key part of the data reduction pipeline.

2.3 The C-BASS Instruments



(a) C-BASS north

(b) C-BASS south

Figure 2.1: Photographs of the C-BASS north (*left*), and C-BASS south (*right*) telescopes. The northern telescope consists of a 6.1m dish with absorbing baffles around the primary and secondary mirrors, in a Gregorian configuration. The southern telescope consists of an under-illuminated 7.6m dish, in a Cassegrain configuration.

In Fig. 2.1 I show the C-BASS north and C-BASS south telescopes. The northern telescope was donated to C-BASS by the Jet Propulsion Laboratory, having originally been designed as an array element for the Deep Space Network ([Imbriale and Abraham, 2004](#)). The telescope consists of a 6.1 m dish in a Gregorian configuration. The

southern telescope was donated by Telkom SA to SKA South Africa, and was originally designed for communication with low-earth orbit telecommunications satellites. It consists of an under-illuminated 7.6 m dish in a Cassegrain configuration. The telescope optics have been designed such that the northern and southern beams are well matched, and both surveys have the same intrinsic angular resolution (Holler et al., 2013).

Spillover at the edges of the mirrors means that the C-BASS beam has far-out sidelobes. As the telescope scans this results in the pickup of ground radiation, which can lead to significant scan-synchronous artefacts in the C-BASS maps. This was minimized for the northern telescope by adding absorbing baffles around the primary and secondary mirrors. This was not required for the southern telescope, given that it was under-illuminated in order to match the size of the northern telescope. This results in very low illumination at the edges of the telescope and negligible spillover. Far-out sidelobes were further reduced by supporting the secondary mirrors on radio-transparent dielectric foam, as opposed to using metal struts. This has the additional effect of maintaining circular symmetry, therefore minimizing cross-polarization.

In Fig. 2.2 I show the C-BASS beam profile, along with the corresponding spin-0 beam transfer function. This beam is derived from physical optics simulations of the telescope optics, and takes in to account the finite bandwidth of the receiver and the effective pointing jitter of the survey (Leech et al. in prep). Given some axisymmetric beam profile $\mathcal{B}(\theta)$, the spin-0 beam transfer function is given by the Legendre transform,

$$b_\ell = 2\pi \int_0^\pi d\theta \sin\theta P_\ell(\cos\theta) \mathcal{B}(\theta). \quad (2.2)$$


In polarization we must calculate the spin-2 Legendre transform. Assuming a pure co-polar beam, the E and B components of the spin-2 transform satisfy

$$b_{\ell,\pm 2}^E = \pm i b_{\ell,\pm 2}^B. \quad (2.3)$$

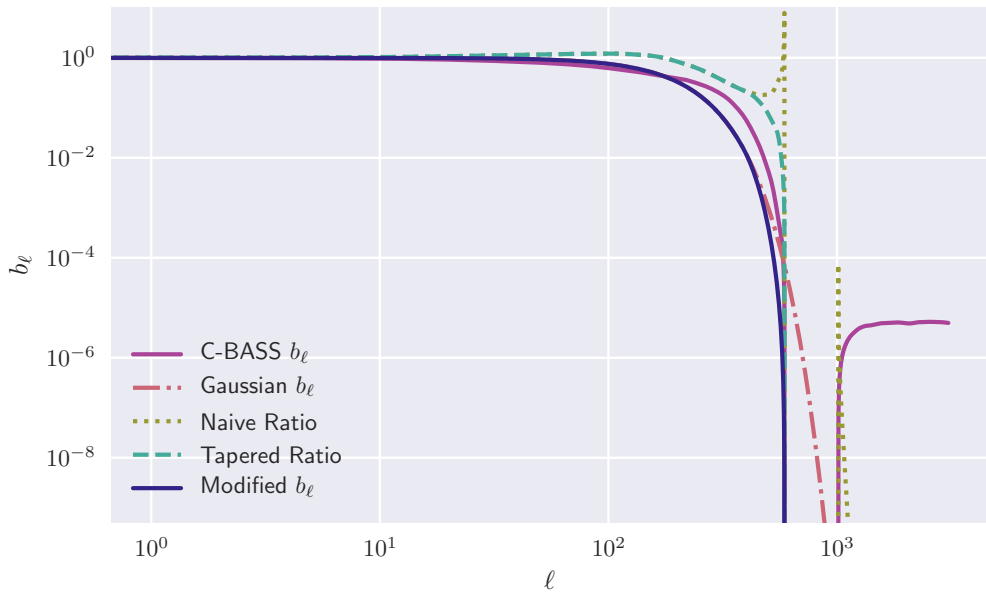
The beam response tensor for linear polarization is therefore entirely specified by $b_{\ell,2}^E$, which is given by

$$b_{\ell,2}^E = 2\pi N_\ell^2 \sqrt{\frac{4\pi}{2\ell+1}} \int_0^\pi d\theta \sin\theta \mathcal{B}(\theta) \left\{ (\ell+2)(\cos\theta-2) \frac{\partial^2 P_{\ell-1}}{\partial(\cos\theta)^2} + \left(2(\ell-1) - \frac{1}{2}l(\ell-1)\sin^2\theta - (\ell-4) \right) \frac{\partial^2 P_\ell}{\partial(\cos\theta)^2} \right\}, \quad (2.4)$$

where $N_\ell^2 = 2(\ell-2)!/(\ell+2)!$ (Challinor et al., 2000). Whilst I only show the spin-0 transfer function in Fig. 2.2b, the spin-2 transfer function only shows small



(a) The C-BASS beam profile.



(b) Beam transfer functions.

Figure 2.2: Panel (a): The C-BASS beam profile. Panel (b): The associated transfer functions; showing the C-BASS beam transfer function (*solid purple line*), the 1° FWHM Gaussian transfer function (*dot-dashed rose line*), the naive ratio between the C-BASS and Gaussian transfer functions (*dotted olive line*), the tapered ratio between the C-BASS and Gaussian transfer functions (*dashed teal line*), and the modified transfer function produced by multiplying the C-BASS transfer function and the tapered ratio (*solid indigo line*).

differences, except for $\ell \in [0, 1]$ where $b_{\ell,2}^E = 0$. At $\ell \sim 600$ the C-BASS transfer function drops off sharply. For scales beyond this the C-BASS map does not contain meaningful information. Note that the rise at $\ell \sim 1000$ is caused by the discrete sampling of the C-BASS beam in real-space, introducing small-scale artefacts into the transfer function.

Given the Legendre transform of the beam, we can attempt to naively deconvolve the C-BASS map by multiplying the C-BASS transfer function by the ratio,

$$r_\ell = \frac{g_\ell}{b_\ell}, \quad (2.5)$$

where g_ℓ is some target transfer function. In this case g_ℓ is the transfer function for a 1° FWHM Gaussian. However, performing this naive deconvolution can have the unintended effect of inflating noise at small-scales, introducing ringing artefacts into the map. This can be seen in Fig. 2.2b. Whilst the C-BASS transfer function drops sharply to zero at $\ell \sim 600$, the Gaussian beam corresponds to a continuous aperture (as opposed to the finite aperture of C-BASS) and therefore continues to drop off smoothly after this. This results in a crossing of the two transfer functions at these scales and a sharp spike in the naive ratio. To perform the deconvolution we therefore multiply r_ℓ by a Hanning window function over these scales, given by

$$H_n(\ell) = \begin{cases} 1, & \text{if } \ell < \ell_0, \\ \cos^2\left(\frac{\pi(\ell-\ell_0)}{2(\ell_c-\ell_0)}\right), & \text{if } \ell_0 \leq \ell \leq \ell_c, \\ 0, & \text{otherwise.} \end{cases} \quad (2.6)$$

The value of $\ell_0 = 387$ was chosen by defining an interpolating spline for r_ℓ over the range $300 \leq \ell \leq 700$ (which includes the spike), and selecting the ℓ value at which the derivative changes from negative to positive. The value of $\ell_c = 590$ was chosen by selecting the ℓ value at which the C-BASS transfer function drops to zero. Multiplying r_ℓ by Equation 2.6 results in the tapered ratio shown in Fig. 2.2b. It is this ratio function that is used in the C-BASS map deconvolution.

In Fig. 2.3 I show a simplified block diagram of the C-BASS front end and in Fig. 2.4 I show the block diagrams for the C-BASS north and south radiometer/polarimeter systems. C-BASS measures total intensity by comparing the power received by the antenna with a stabilized load. Polarization is measured by taking complex correlations of the left and right circularly polarized components of the

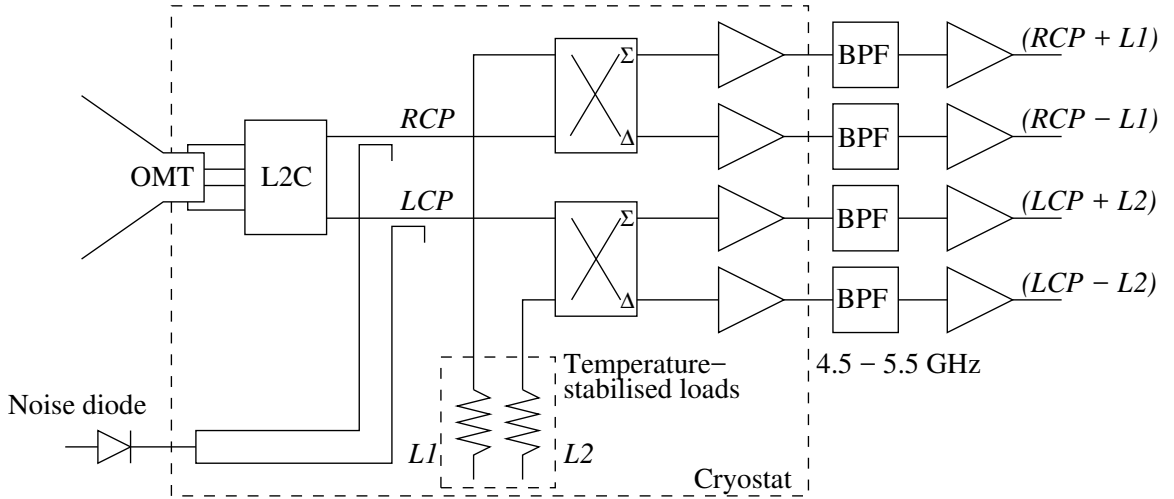


Figure 2.3: Simplified block diagram of the C-BASS front end, which is common to C-BASS north and south. Key: OMT = orthomode transducer, L2C = linear to circular converter, Σ, Δ = sum, differencing, BPF = bandpass filter, RCP = right circular polarization, LCP = left circular polarization. L1 and L2 are matched loads. Figure taken from [Jones et al. \(2018\)](#).

incoming wave. These correlations yield the Stokes parameters as follows,

$$\langle |E_R|^2 + |E_L|^2 \rangle = I, \quad (2.7)$$

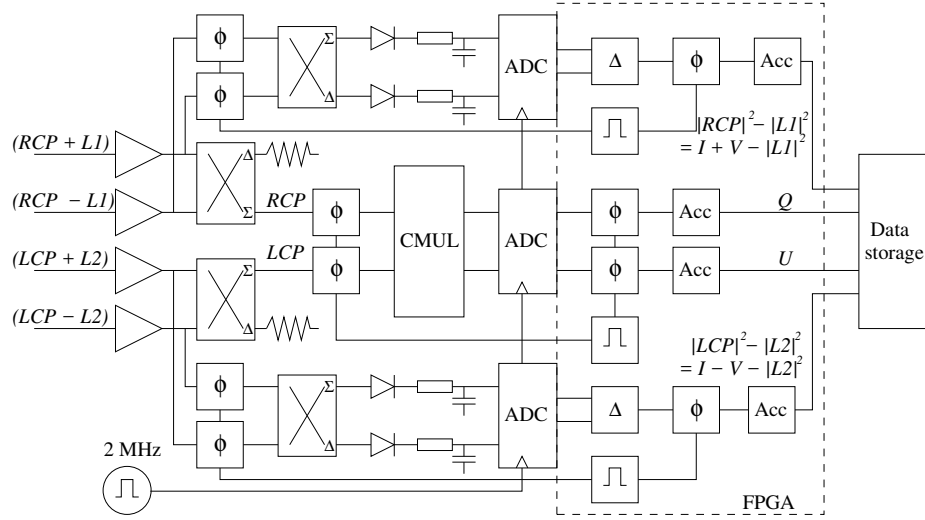
$$\langle E_R E_L^* \rangle = (Q + iU)/2, \quad (2.8)$$

$$\langle E_L E_R^* \rangle = (Q - iU)/2, \quad (2.9)$$

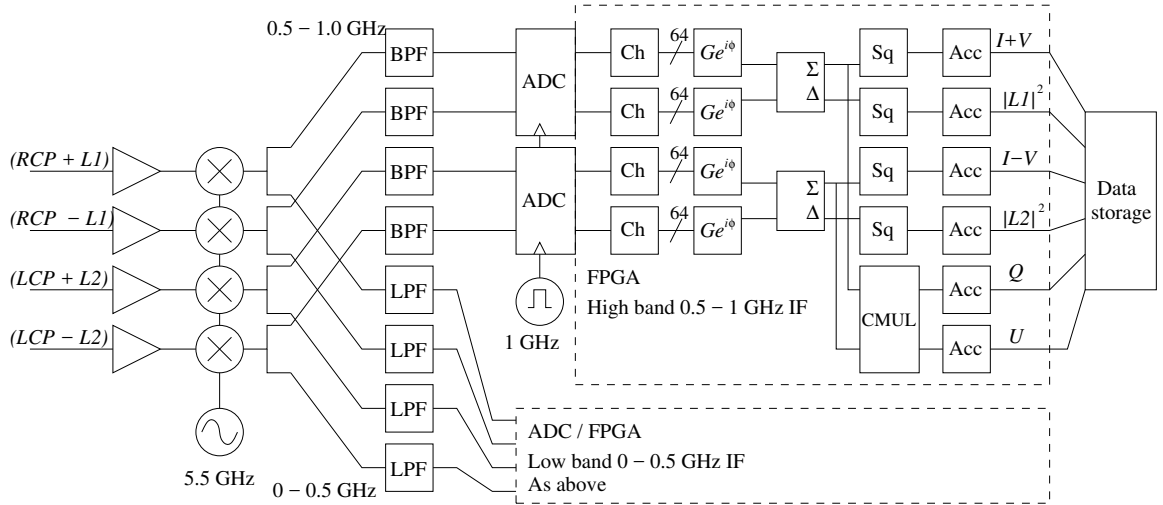
$$\langle |E_R|^2 - |E_L|^2 \rangle = V, \quad (2.10)$$

where E_L and E_R are the complex amplitudes of the left and right circular polarizations of the incident electric field respectively. To calculate Q using the linear polarizations (E_x and E_y) would require evaluating $Q = \langle |E_x|^2 - |E_y|^2 \rangle$. For most astrophysical sources $Q, U \ll I$, meaning this difference involves taking the difference of two large numbers to obtain the small polarized signal, which is an inherently inaccurate process compared to taking the correlations of the circular components. In principle, C-BASS can also measure Stokes V . However, astronomical circular polarization is expected to be extremely small, meaning this measurement would require a very careful calibration of any intensity measurements.

Radiation that is incident on the C-BASS horn is guided down to an orthomode transducer (OMT). The OMT couples to the E_x and E_y components of the incident field, with the signal then being passed through a linear to circular converter. This produces the right (RCP) and left (LCP) circular polarizations. The sky RCP and



(a) C-BASS north analogue backend.



(b) C-BASS south digital backend.

Figure 2.4: Block diagrams of the C-BASS radiometer/polarimeter systems. Panel (a): C-BASS north analogue backend, Panel (b): C-BASS south digital backend. Key: ϕ = phase switch modulation/demodulation, Σ, Δ = sum, differencing, ADC = analogue to digital converter, Acc = accumulator, CMUL = complex multiply, BPF = band-pass filter, LPF = low-pass filter, Ch = channeliser, $Ge^{i\phi}$ = complex gain correction, Sq = square detector (evaluates VV^* on complex voltages V). Figure taken from Jones et al. (2018).

LCP signals are combined with reference signals using 180° hybrids (i.e., the sum and difference operations shown in Fig. 2.3), which are amplified and passed through bandpass filters that define the 4.5 – 5.5 GHz C-BASS bandpass. This produces the four sky and reference signal combinations shown in Fig. 2.3.

Block diagrams of the radiometer and polarimeter systems for both C-BASS North and C-BASS South are shown in Fig. 2.4. For the radiometer part of the northern system, the sky and reference signals are separated using additional 180° hybrids which are then detected using diodes. This gives two measurements for the total intensity,

$$I_1 = |E_L|^2 = I - V, \quad (2.11)$$

$$I_2 = |E_R|^2 = I + V. \quad (2.12)$$

For the polarimeter, sky and reference signals are also separated using 180° hybrids, with LCP and RCP then being correlated to produce measurements of Q and U . Given measurements for I , Q and U the signals are then digitized using the analogue to digital converter (ADC). The ADC samples the data at 2 MHz, which are subsequently filtered and averaged down to 10 ms samples in a field-programmable gate array (FPGA). A detailed description of the northern system can be found in King et al. (2014); Stevenson et al. (2019).

The southern system performs the same core signal processing operations to produce intensity and polarization measurements. However, the signal digitization occurs much earlier, with the radiometry and polarimetry being performed digitally. The southern system also has 128 evenly spaced frequency channels, with the signals being combined on a channel-by-channel basis to produce the intensity and polarization measurements. A detailed description of the southern system can be found in Copley (2014).

2.4 Data Processing

In Fig. 2.5 I display a flowchart outlining the steps taken in the C-BASS data reduction pipeline. The data reduction pipeline was developed prior to my joining the C-BASS collaboration. In this section I outline the key steps necessary to take the raw time-ordered data (TOD) from the telescope and produce calibrated Stokes I , Q and U maps. Detailed discussions of the C-BASS data reduction and calibration can be found in Jew (2017), and will be given in the C-BASS survey and commissioning papers (Taylor et al., in prep. and Pearson et al., in prep).

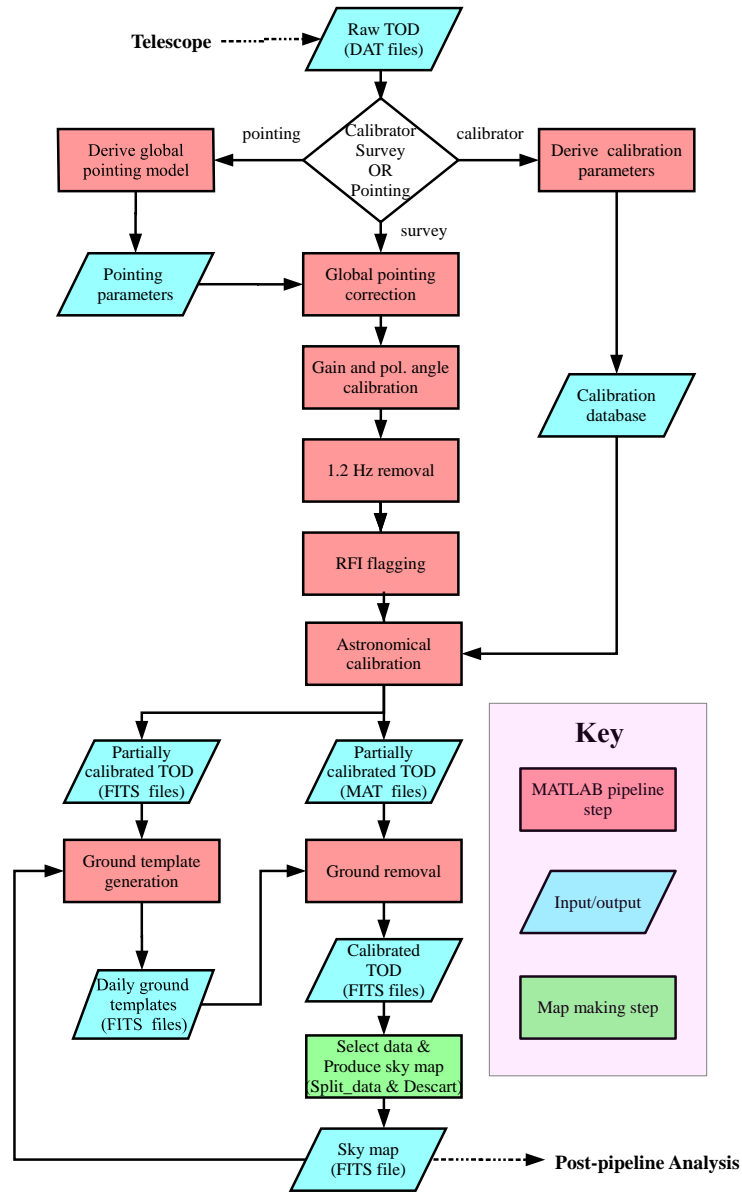


Figure 2.5: Flow diagram of the C-BASS data reduction pipeline. *Red squares* denote MATLAB subprocesses, *green squares* denote map-making steps and *cyan parallelograms* denote inputs and outputs. Figure adapted from Jew (2017).

2.4.1 Pointing

Small discrepancies exist between the directions the C-BASS telescope was pointed during observations and the recorded pointing directions. The C-BASS global pointing model accounts for this, and applies pointing corrections to the raw TOD. The model is periodically checked by conducting pointing observations of bright point sources. For this the telescope is first scanned through the source at constant elevation, and then at fixed azimuth. The C-BASS beam is fitted to the source data, with the distance between the fitted source location and the known source location being used to determine the pointing offset. If the pointing offset is found to be larger than 5 arcmin, the global pointing model is updated.

2.4.2 Gain and Polarization Angle Calibration

The raw C-BASS data contains gain and phase shifts that must be corrected for. These are monitored through the use of a noise diode. The noise diode is temperature stabilized in order to provide a fixed-amplitude reference signal in intensity and polarization. At the beginning of each telescope scan, the noise diode is turned on for a few seconds, allowing for gain measurements on the timescale of minutes and the removal of long timescale gain drifts. The phase of the noise diode signal is set such that it appears only in the instrumental Q channel. Phase variations in the RCP and LCP signals result in some of the noise diode signal appearing in the instrumental U channel. This is corrected for in the data reduction pipeline by applying a rotation such that the noise diode signal only appears in the instrumental Q channel.

2.4.3 Non-Astrophysical Contamination

The C-BASS north TOD contain a periodic contaminant at a frequency of ~ 1.2 Hz and its harmonics. This is thought to be caused by vibrations of the receiver components due to the cold head, which has a refrigeration cycle of 1.2 Hz. This contaminant is removed by fitting templates of the 1.2 Hz signal to the TOD and subtracting off its contribution. The 1.2 Hz removal is described in detail in [Stevenson \(2014\)](#).

Radio frequency interference (RFI) from human activities can lead to contaminating signals in the C-BASS band. RFI is flagged in the data by interpolating a sky model onto the TOD stream. This model is subtracted from the data, and any discrepant events are subsequently flagged and removed. Pointing offsets around bright sources can still lead to significant residuals. As such, RFI flagging is disabled for

bright source crossings. However, the impact of RFI around bright sources is proportionally lower than in fainter regions, and the fact each pixel is observed multiple times means the RFI in these regions is washed out in the final maps.

The last major non-astrophysical contaminant in the C-BASS TOD is caused by ground pick-up. As mentioned previously, the ground signal appears as a fixed function of azimuth, which varies on a timescale of days with variations in the ground temperature and emissivity. A sky model is again subtracted from the TOD, with regions of bright sky emission being excluded. The remaining data are averaged into azimuth bins to construct ground templates for each day, which are subsequently subtracted during the map-making stage. Note that this procedure also removes fixed RFI sources, such as geostationary satellites.

2.4.4 Astronomical Calibration

For the astronomical calibration, C-BASS data were calibrated against Tau A as a primary calibrator and Cas A as a secondary calibrator, with calibrator flux densities being calculated from the spectral forms in [Weiland et al. \(2011\)](#). Tau A is used as a calibrator except when it is unobservable due to being too close to the Sun. For Tau A the spectral model takes the form

$$\log S_\nu(\text{Jy}) = 2.506 - 0.302 \log \left(\frac{\nu}{40 \text{ GHz}} \right), \quad (2.13)$$

and for Cas A it takes the form

$$\log S_\nu(\text{Jy}) = 2.204 - 0.682 \log \left(\frac{\nu}{40 \text{ GHz}} \right) + 0.038 \log^2 \left(\frac{\nu}{40 \text{ GHz}} \right). \quad (2.14)$$

These spectral models give the flux density of Tau A at 4.76 GHz as $S_\nu = 609.8 \pm 4.2$ Jy in epoch 2005, and the flux-density of Cas A as $S_\nu = 736.1 \pm 3.4$ Jy in epoch 2000. The secular decreases adopted for Tau A and Cas A were the same as those used for the model fits in [Weiland et al. \(2011\)](#). For Tau A this was -0.167 per cent per year, which was taken from [Macías-Pérez et al. \(2010\)](#), and for Cas A the secular decrease was -0.53 per cent per year. The C-BASS calibration is accurate to ~ 5 per cent. The flux-temperature conversion is defined through

$$T_{\text{source}} = S_\nu(\text{Jy}) \frac{\pi}{2k_B \times 10^{26}} \left(\frac{6.1}{2} \right)^2 \alpha_{\text{eff}}, \quad (2.15)$$

where T_{source} is the source temperature, k_B is the Boltzmann constant and $\alpha_{\text{eff}} = 0.55$ is the theoretical aperture efficiency for C-BASS. The polarization angle calibration

was performed by matching the Tau A polarization angle to the fitted value in [Weiland et al. \(2011\)](#), accounting for Faraday rotation at 4.76 GHz. Given an intrinsic polarization angle of $\chi_0 = 148.5^\circ$, and a rotation measure of $\text{RM} = -24.1 \text{ rad m}^{-2}$, we obtain a calibrator polarization angle of $\chi(4.76 \text{ GHz}) = 143.0^\circ$.

2.4.5 Map-Making

C-BASS uses the destriping mapper, DESCART ([Sutton et al., 2010](#)) to produce its sky maps. This mapper is able to take account of $1/f$ noise in the C-BASS TOD. This is noise that has excess power at low frequencies, compared to white noise which has equal power at all frequencies. Whilst the C-BASS receiver was designed to minimize $1/f$ noise, long timescale variations in the background level, atmospheric emission and ground pick-up mean that $1/f$ noise is still present in the C-BASS TOD. Left unaccounted for, the $1/f$ noise leads to stripes in the C-BASS maps along the scan directions. DESCART solves for this by modelling the TOD signal, d_t , as

$$d_t = P_{tp}s_p + F_{ti}a_i + w_t. \quad (2.16)$$

The true sky emission is given by s_p , which is projected onto the TOD through the pointing matrix P_{tp} . The $1/f$ contribution is modelled through a series of offsets a_i , with the matrix F_{ti} defining the timescale over which the offsets vary. The white noise contribution is given by w_t . DESCART solves for the offsets, which are then subtracted from the TOD. This leaves clean TOD, consisting of the true sky signal plus white noise, which can then be binned into sky pixels.

2.5 Data Checks: Calibration

In order to check the stability of the C-BASS calibration, I have performed fits for the flux-densities of bright sources over monthly data-splits. I have also compared the polarization angles in the C-BASS map, with those in the WMAP K and *Planck* 30 GHz maps to assess the reliability of the polarization angle calibration. The results of these data-quality checks are presented here.

2.5.1 Calibration Stability

To check the stability of the C-BASS calibration as a function of time, I fit for the flux-densities of Tau A, Cas A, Cygnus A and Virgo A in monthly C-BASS maps i.e., maps produced using one month of data. I extract 90 arcmin Gnomonic

projections centred on the known source location, which is used at the origin of the local coordinate system for the fit. To these projections, I fit an axisymmetric model of the C-BASS beam, along with a local background offset and slope. The model value at a pixel position \mathbf{x}_p in the image plane is given by

$$m(\mathbf{x}_p) = T_s \mathcal{B}(|\mathbf{x}_p - \mathbf{x}_0|) + \mathbf{s} \cdot \mathbf{x}_p + C, \quad (2.17)$$

where T_s is the source temperature, $\mathbf{x}_0 = (x_0, y_0)^\top$ is the offset source position, $\mathbf{s} = (s_x, s_y)^\top$ is the background slope parameter and C is the background offset. Positions in the image plane are given in units of arcminutes. I assign the following weakly informative priors,

$$T_s \sim \text{Exp}(\lambda = 0.1 \text{ K}^{-1}), \quad (2.18)$$

$$x_0 \sim \mathcal{N}(\mu = 0 \text{ arcmin}, \sigma = 10 \text{ arcmin}), \quad (2.19)$$

$$y_0 \sim \mathcal{N}(\mu = 0 \text{ arcmin}, \sigma = 10 \text{ arcmin}), \quad (2.20)$$

$$s_x \sim \mathcal{N}(\mu = 0 \text{ arcmin}, \sigma = 10 \text{ arcmin}), \quad (2.21)$$

$$s_y \sim \mathcal{N}(\mu = 0 \text{ arcmin}, \sigma = 10 \text{ arcmin}), \quad (2.22)$$

$$C \sim \mathcal{N}(\mu = 0 \text{ K}, \sigma = 1 \text{ K}). \quad (2.23)$$

The priors are chosen to encapsulate the relevant physical scales for each parameter. I sample the parameter space using the NUTS algorithm (Hoffman and Gelman, 2014), implemented in PYMC3 (Salvatier et al., 2016). I assume pixel temperature values, $T(\mathbf{x}_p)$ follow a Gaussian distribution, $T(\mathbf{x}_p) \sim \mathcal{N}(\mu = m(\mathbf{x}_p), \sigma = \sigma(\mathbf{x}_p))$, where $\sigma(\mathbf{x}_p)$ is the standard deviation for the pixel p . Recovered peak temperatures are converted to flux-density units using Equation 2.15. Note that for some months there is very little data, which results in the sampler either failing to converge, or the occurrence of divergences, which are indicative of failures in geometric ergodicity. Divergences in NUTS are discussed further in Chapter 5. For these months I do not save the resulting posterior samples.

In addition to recovering individual flux-density measurements in each monthly map, I also fit for the additional scatter in the recovered flux-densities as a function of time. To do this, I model the flux-densities for each month as being drawn from some underlying Gaussian distribution,

$$S_\nu^{i,c} \sim \mathcal{N} \left(\mu = \mu_c, \sigma = \sqrt{\sigma(S_\nu^{i,c})^2 + \Sigma_c^2} \right). \quad (2.24)$$

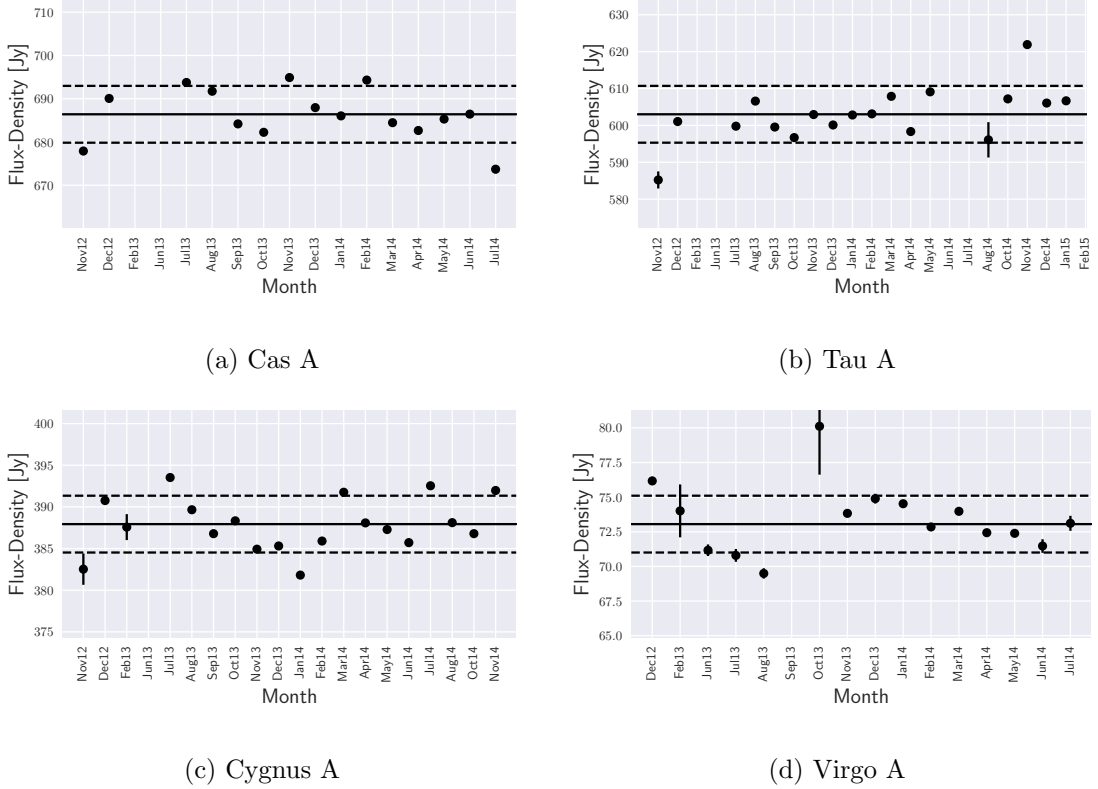


Figure 2.6: Recovered flux-densities for Cas A, Tau A, Cygnus A and Virgo A in monthly C-BASS maps. For months with no point marked, there was either no data available for that source, the sampler failed to converge, or divergences occurred during sampling. The *solid* line on each plot denotes the fitted mean flux-density, and the *dashed* lines are at $\mu_c \pm \Sigma_c$ for each source.

Here, $S_v^{i,c}$ is the fitted flux-density of source c for the i^{th} month, and $\sigma(S_v^{i,c})$ is the corresponding standard deviation. The parameter Σ_c is the additional scatter in flux-densities for the source c as a function of time, and μ_c is the corresponding mean flux-density. I assign weakly informative priors as follows,

$$\mu_c \sim \text{Exp}(\lambda = 10^{-3} \text{ Jy}^{-1}), \quad (2.25)$$

$$\Sigma_c \sim \text{Exp}(\lambda = 10^{-2} \text{ Jy}^{-1}). \quad (2.26)$$

These are again chosen to encapsulate the relevant physical scales for the bright sources under consideration.

In Fig. 2.6 I show the fitted flux-densities for each source in each month, along with the fitted mean and scatter as a function of time. In Table 2.2 I give the fitted mean and scatter as a function of time for each source. The mean flux-densities for Cas A and Tau A are consistent with the calibration models, after accounting for the secular

decrease. Note that different months are shown for each of the sources. The C-BASS observing strategy means that for some months there is very little data available for certain sources. This lack of data results in the poor flux-density constraint for Virgo A in October 2013. The scatter in source flux-densities over time is at the ~ 1 per cent level for Cas A, Tau A and Cygnus A, and ~ 3 per cent for Virgo A. This shows that the calibration of C-BASS is stable over the course of the survey.

Table 2.2: Fitted mean and scatter over time for Cas A, Tau A, Cygnus A and Virgo A. Mean flux-densities for Cas A and Tau A are consistent with the input calibration models. The additional scatter over time is at the ~ 1 per cent level for Cas A, Tau A and Cygnus A, and at the ~ 3 per cent level for Virgo A.

Source	μ_c	Σ_c
Cas A	686.4 ± 1.8 Jy	6.6 ± 1.4 Jy
Tau A	603.0 ± 1.9 Jy	7.7 ± 1.5 Jy
Cygnus A	387.9 ± 0.8 Jy	3.4 ± 0.6 Jy
Virgo A	73.0 ± 0.6 Jy	2.0 ± 0.5 Jy

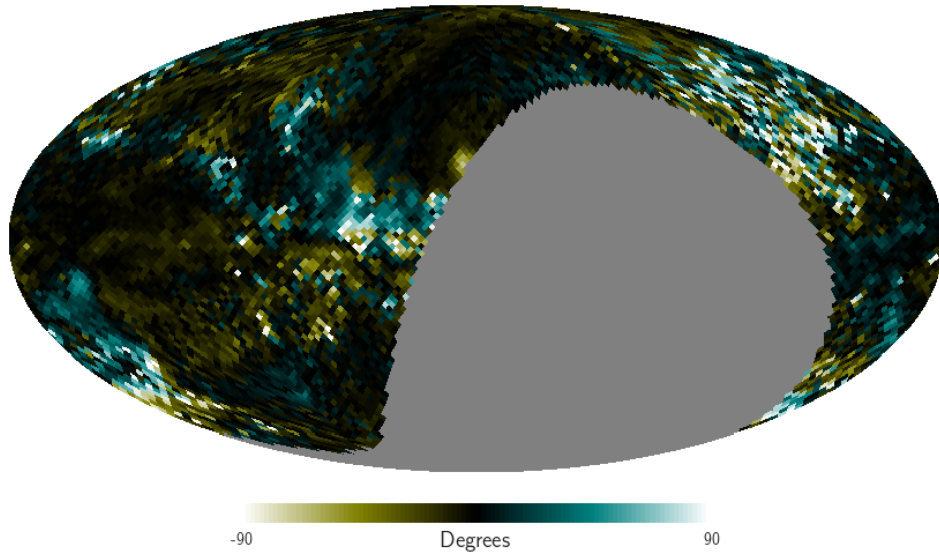
2.5.2 Polarization Angles

In order to check the astronomical polarization angle calibration, I compare the C-BASS polarization angles with the WMAP K and *Planck* 30 GHz polarization angles. Comparisons are all made using maps that are convolved with a 1° FWHM Gaussian. The polarization angles are calculated as

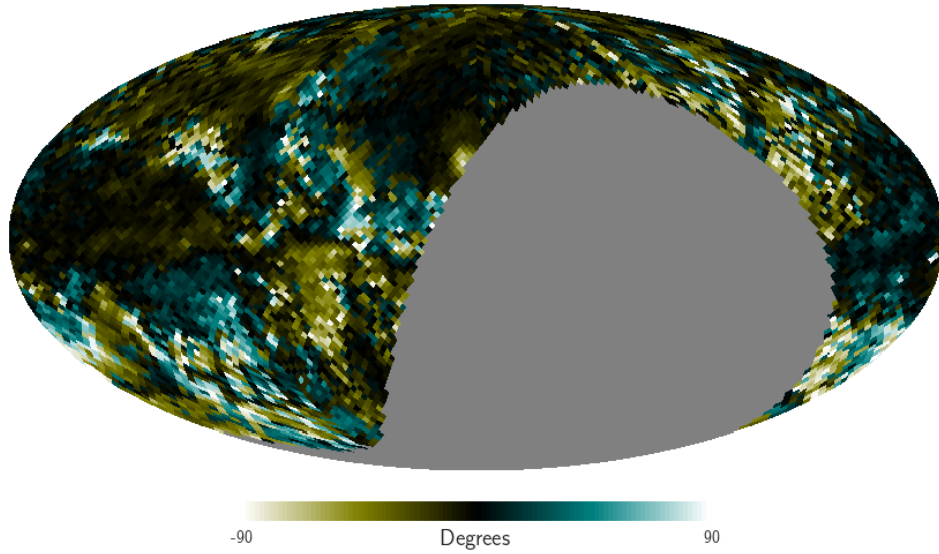
$$\chi = \frac{1}{2} \arctan_2(U, Q), \quad (2.27)$$

where the \arctan_2 function is used in order to deal with the quadrant ambiguity arising from the various possible signs on Q and U . In Fig. 2.7 I show maps of the polarization angle difference between C-BASS and the WMAP K and *Planck* 30 GHz maps, calculated at Healpix $N_{\text{SIDE}}=32$ (Górski et al., 2005). In regions close to the Galactic plane, strong Faraday rotation effects contaminate any direct comparisons between the C-BASS and WMAP/*Planck* polarization angles. Further, in fainter regions of the sky the WMAP and *Planck* maps have very low SNR, resulting in very noisy polarization angle estimates.

In Fig. 2.8 I show histograms and corresponding KDE plots of the polarization angle difference, but applying a mask over regions with strong Faraday rotation and



(a) C-BASS - WMAP K



(b) C-BASS - *Planck* 30 GHz

Figure 2.7: Panel (a): Polarization angle differences between C-BASS and WMAP K-band. Panel (b): Polarization angle differences between C-BASS and *Planck* 30 GHz. Close to the Galactic plane strong Faraday rotation effects contaminate any direct comparisons between C-BASS and WMAP/*Planck*. In fainter regions of the sky, the WMAP and *Planck* maps have very low SNR, meaning the corresponding polarization angle estimates are heavily contaminated by noise. In regions of high SNR, but with low levels of Faraday rotation, the C-BASS and WMAP/*Planck* polarization angles show broad agreement.

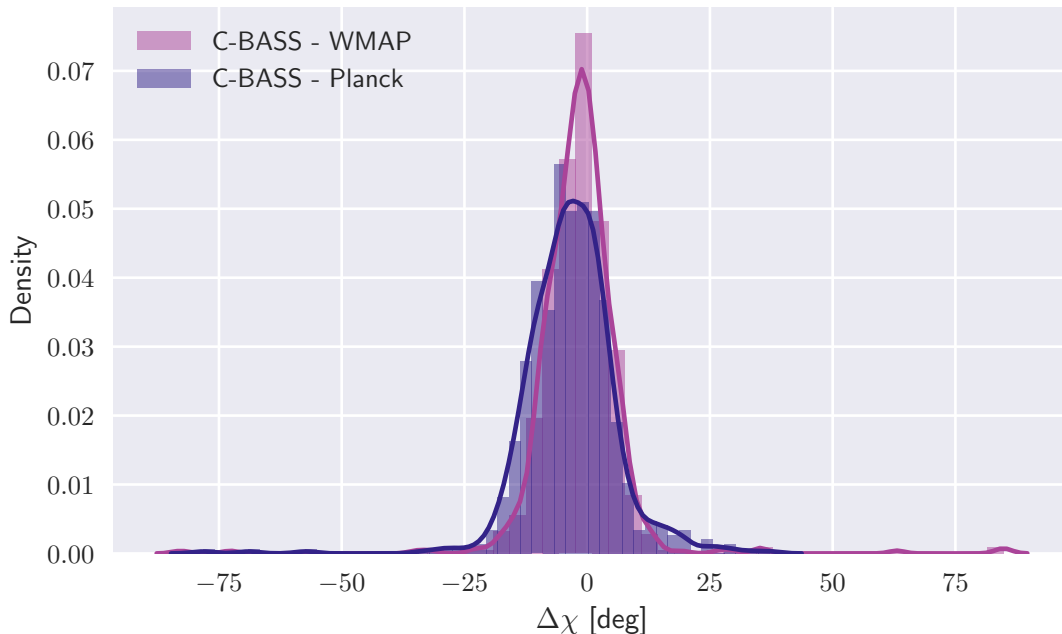


Figure 2.8: Histograms and corresponding KDE plots for the polarization angle differences between C-BASS and the WMAP K/*Planck* 30 GHz maps. Angle differences are given here in degrees. I only include polarization angle differences calculated in regions of the sky with high SNR and low levels of expected Faraday rotation. The C-BASS and WMAP/*Planck* polarization angles show broad agreement, with the mean polarization angle difference in both cases consistent with expectations from Faraday rotation.

low SNR. The mask was constructed by applying a conservative $\pm 30^\circ$ cut around the Galactic plane, and masking any pixels with a greater than 10° rotation in the Faraday depth map of [Oppermann et al. \(2012\)](#). To quantify the level of agreement between the C-BASS and WMAP/*Planck* polarization angles, I model the polarization angle differences as being drawn from a Gaussian distribution, and fit for the mean, μ_χ and standard deviation, σ_χ of these distributions. I assign weakly informative priors as follows,

$$\mu_\chi \sim \text{VonMises}(\mu = 0 \text{ rad}, \kappa = (18/\pi)^2 \text{ rad}^{-2}), \quad (2.28)$$

$$\sigma_\chi \sim \text{Exp}(\lambda = 18/\pi \text{ rad}^{-1}). \quad (2.29)$$

The von Mises distribution is assigned to μ_χ , which is the circular analogue of the Normal distribution. The parameter, κ is analogous to $1/\sigma^2$. The results from these fits are given in Table 2.3. As mentioned previously, using RM data from [Taylor et al. \(2009\)](#), one can estimate the expected level of Faraday rotation between 4.76 GHz and

~ 30 GHz, over Galactic latitudes $|b| > 30^\circ$, to be $\sim 3^\circ$. The mean angle differences found in the comparisons with WMAP and *Planck* maps are consistent with these expectations.

Table 2.3: Fitted means and standard deviations of the polarization angle difference distributions. Mean angle differences when comparing C-BASS with WMAP and *Planck* are consistent with expectations from Faraday rotation.

Survey Comparison	μ_χ	σ_χ
C-BASS - <i>Planck</i>	$-3.3^\circ \pm 0.4^\circ$	$9.5^\circ \pm 0.3^\circ$
C-BASS - WMAP	$-1.5^\circ \pm 0.4^\circ$	$9.6^\circ \pm 0.3^\circ$

2.6 Current Status

To date, the C-BASS northern survey has been completed. Detailed analyses of these data are underway, with some of the first scientific results being presented in [Dickinson et al. \(2019\)](#); [Grumitt et al. \(2020b\)](#); [Cepeda-Arroita et al. \(2020\)](#). The southern survey is ongoing and will require approximately two years of observations. In [Fig. 2.9](#) I show the current C-BASS northern sky maps, which have been deconvolved with the C-BASS beam and smoothed by a 1° FWHM Gaussian. These maps form the basis of the scientific analyses with C-BASS data presented in this thesis.

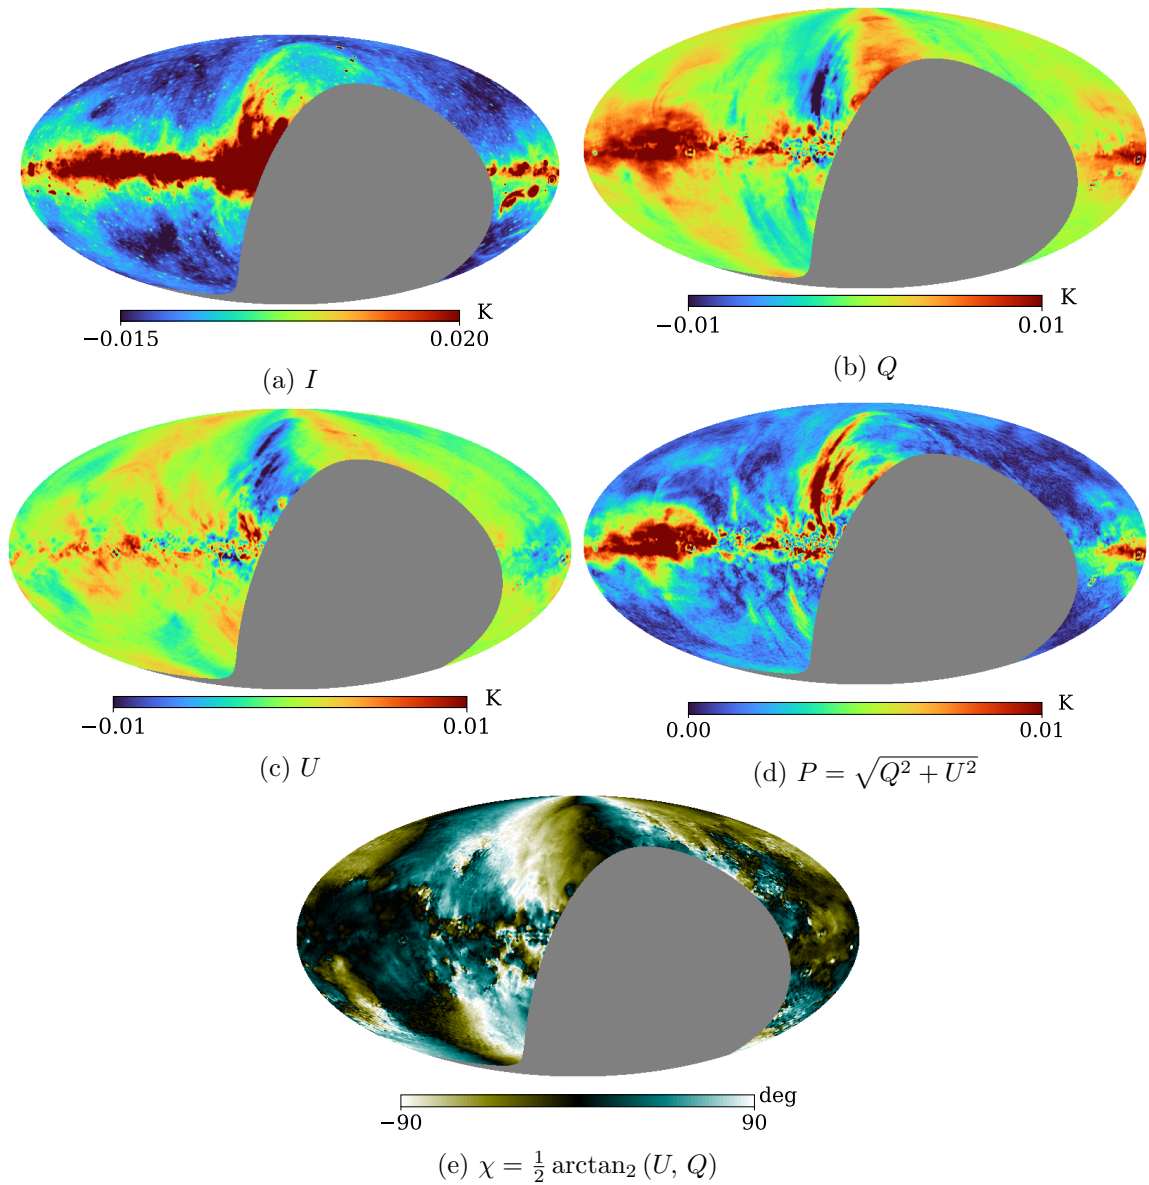


Figure 2.9: The current C-BASS Stokes I , Q , U , P and χ maps. These have been deconvolved with the C-BASS beam, and smoothed with a 1° FWHM Gaussian. These maps form the basis of scientific analyses with C-BASS northern sky data, and are used for the analysis in Chapters 3 and 4.

Chapter 3

Point-Source Detection over the Northern Sky

In this chapter I present a point-source catalogue obtained from the C-BASS northern sky maps. Given the large beam of C-BASS, the instrument is not well suited to point-source detection, with detections of all but the brightest sources ($\gtrsim 1$ Jy) being limited by source confusion and diffuse emission. Nevertheless, for the C-BASS analysis, performing a dedicated source detection enables important data quality checks on the C-BASS flux-density scale and pointing accuracy, by comparing the C-BASS catalogue with pre-existing catalogues around 4.76 GHz. The resulting catalogue will also allow us to construct an accurate mask of bright source emission in the C-BASS maps, which will be key to avoid biasing CMB component separation analyses with C-BASS. Fainter sources can either be subtracted using deeper, pre-existing source catalogues, or can be treated statistically at the power spectrum level. The work presented in this chapter had been published in [Grumitt et al. \(2020b\)](#).

Numerous source-detection algorithms have been developed for CMB experiments, see e.g., [Tegmark and de Oliveira-Costa \(1998\)](#); [Vielva et al. \(2003\)](#); [Argüeso et al. \(2009\)](#); [Bennett et al. \(2013\)](#); [Herranz et al. \(2009\)](#); [Carvalho et al. \(2012\)](#); [López-Caniego et al. \(2006\)](#); [González-Nuevo et al. \(2006\)](#); [Planck Collaboration et al. \(2014b, 2016f\)](#). The Planck Catalogue of compact sources (PCCS) used an algorithm based on filtering small patches of sky using the flat-sky second-order Mexican Hat wavelet as an approximation to a matched filter, followed by threshold detection based on the signal-to-noise ratio (SNR). Here I employ a similar scheme but using the equivalent spherical function on the whole sky at once, implemented in the Healpix pixelization scheme ([Górski et al., 2005](#)). Previous applications of the first-order SMHW (SMHW1) to searches for non-Gaussianity can be found in [Cayón](#)

et al. (2001); Martínez-González et al. (2002); Curto et al. (2011), and to point-source detection in Vielva et al. (2003). The details of this implementation are discussed in Section 3.2.

The outline of this chapter is as follows: in Section 3.1 I summarize pre-existing source catalogues directly relevant to the C-BASS analysis. In Section 3.2 I discuss the C-BASS point-source detection algorithm, giving an overview of the method and the Monte Carlo simulations used to validate the algorithm. In Section 3.3 I discuss the C-BASS total intensity, northern sky point-source catalogue obtained using the detection algorithm. I compare the C-BASS catalogue with the GB6 and PMN catalogues, calculate the differential source counts for the C-BASS sources as a cross-check on the statistical properties of the bright source population, and present a source-subtracted C-BASS northern sky map. I summarize the results in Section 3.4. Whilst analysis of polarized sources is important for future CMB polarization studies, this falls beyond the scope of this work where I focus on the C-BASS total intensity results. In addition, due to the low level of source polarization, only a small number of polarized sources ($\mathcal{O}(10)$) are expected to be detected.

3.1 Pre-Existing Source Catalogues

In order to construct an accurate template for masking out the point sources in the C-BASS maps it is necessary to construct a point-source catalogue based on the C-BASS observations themselves and any useful information that can be gleaned from other more sensitive, higher resolution surveys. Relevant to the work in this chapter are the GB6 (Gregory et al., 1996), PMN (Griffith and Wright, 1993; Wright et al., 1994a; Griffith et al., 1994, 1995; Wright et al., 1996a), Effelsberg S5 (Kuehr et al., 1981), RATAN-600 (Mingaliiev et al., 2007) and Combined Radio All-Sky Targeted Eight-GHz Survey (CRATES) (Healey et al., 2007) catalogues. These catalogues are primarily used in this chapter to make comparisons with the C-BASS catalogue, for the purposes of data validation and estimation of the statistical properties of the C-BASS catalogue. However, the catalogues are also useful for characterizing the faint point-source population in any C-BASS analysis (i.e., sources with flux-densities below ~ 1 Jy), where reliable extraction from the C-BASS map becomes more challenging. These radio surveys are summarized in Table 3.1.

The GB6 and PMN 4.85 GHz source catalogues cover declinations $-87.5^\circ \leq \delta \leq 75^\circ$. The GB6 catalogue was produced using the NRAO seven-beam receiver on the 91 m telescope, and the PMN catalogue was produced using the Parkes 64 m radio

Table 3.1: Summary of radio surveys relevant to the work in this chapter. In stating the sky coverage, δ denotes declination, and b denotes Galactic latitude.

Survey Name	Reference	Frequency [GHz]	Sky Coverage	FWHM [arcmin]	Flux Limit [mJy]	Number of Sources
C-BASS	Jones et al. (2018)	4.76	$-90^\circ \leq \delta \leq 90^\circ$ ^a	45	~ 500 ^b	1784
GB6	Gregory et al. (1996)	4.85	$0^\circ \leq \delta \leq 75^\circ$	3	~ 18	75,162
PMN	Wright et al. (1996b)	4.85	$-87.5^\circ \leq \delta \leq 10^\circ$	5	~ 35	50,814
Effelsberg S5	Kuehr et al. (1981)	4.9	$70^\circ \leq \delta \leq 90^\circ$	2.7	~ 250	476
RATAN-600	Mingaliev et al. (2007)	4.8 (1.1 – 21.7) ^c	$75^\circ \leq \delta \leq 88^\circ$	0.67×6.6 ^d	(200) ^e	504
CRATES	Healey et al. (2007)	8.4 (4.85) ^f	$ b > 10^\circ$	2.4	~ 65	11,131

^a Here I state the sky coverage of the whole C-BASS experiment. The results concerning point sources in this chapter were obtained for the C-BASS northern intensity data, covering declinations $-10^\circ \leq \delta \leq 90^\circ$.

^b For C-BASS I state the value $3.5\langle\sigma\rangle$, where $\langle\sigma\rangle$ is the mean background fluctuation level found across the map, applying the CG30 mask described in Section 3.2.5. Details on the estimation of background fluctuation levels are given in Section 3.2.1.

^c The RATAN-600 catalogue covers 6 frequencies from 1.1 – 21.7 GHz, including 4.8 GHz.

^d I state $\text{FWHM}_{\text{RA}} \times \text{FWHM}_\delta$ for RATAN-600, determined from the values given in [Kovalev et al. \(1999\)](#).

^e The RATAN-600 catalogue was produced by pre-selecting NVSS sources with $S_{1.4\text{GHz}} > 200$ mJy.

^f The CRATES catalogue is primarily at 8.4 GHz, with 4.85 GHz sources used as the basis for observations of 8.4 GHz counterparts. The properties of these 4.85 GHz sources are provided with the CRATES catalogue. Further observations were also made at 4.85 GHz to fill in gaps at $\delta > 88^\circ$.

telescope. The GB6 catalogue has a flux-density limit of approximately 18 mJy, whilst the PMN catalogue has an average flux-density limit of approximately 35 mJy over the sky. The GB6 and PMN catalogues provide far deeper flux-density coverage than can be achieved with C-BASS, given the differing resolutions and hence confusion levels. However, it is still necessary to obtain a dedicated source catalogue, such that an accurate mask can be constructed for the brightest sources in the C-BASS maps, and account for source variability between the C-BASS and GB6/PMN surveys. The GB6 and PMN catalogues remain useful in accounting for fainter sources (below ~ 1 Jy) in any C-BASS analysis.

Source catalogues covering the North Celestial Pole (NCP) region are more limited, with GB6 only covering declinations up to $\delta = 75^\circ$. The Effelsberg S5 catalogue covers declinations $\delta \geq 70^\circ$ and is complete down to 250 mJy ([Kuehr et al., 1981](#)). In comparing to the S5 catalogue there are significant issues from the variability of flat-spectrum sources, given the large separation in time between the S5 survey and the C-BASS survey. The RATAN-600 catalogue includes measurements at 4.8 GHz, and observed 504 sources in the NCP region with NRAO VLA Sky Survey (NVSS) flux-densities, $S_{1.4\text{GHz}} \geq 200$ mJy ([Mingaliev et al., 2001, 2007](#)). The RATAN-600 catalogue was used in a previous analysis of diffuse emission in the NCP region with C-BASS in [Dickinson et al. \(2019\)](#). Whilst this catalogue provides deeper flux-density

coverage than C-BASS, it was produced by pre-selecting sources for study from the NVSS catalogue at 1.4 GHz (Condon et al., 1998). This can potentially miss rising-spectrum sources that would otherwise be observable in the C-BASS catalogue. Given the C-BASS flux-density limit of approximately 500 mJy, we can use the NVSS source counts and the 1.4 – 4.85 GHz spectral index distributions in Tucci et al. (2011) to estimate that there may be $\mathcal{O}(1)$ sources that could be observed by C-BASS, whilst being missed by RATAN-600.

It is also worth noting the CRATES catalogue (Healey et al., 2007). CRATES is an 8.4 GHz catalogue of flat-spectrum sources with flux-densities $S_{\nu=4.85 \text{ GHz}} > 65 \text{ mJy}$, covering Galactic latitudes $|b| > 10^\circ$. The catalogue therefore serves as a useful proxy for flat-spectrum sources in the C-BASS catalogue, which are the primary contributor to source variability. Healey et al. (2009) made additional observations of the NCP region at declinations, $\delta > 88^\circ$ to supplement the original CRATES catalogue. The purpose of this was to bring the flux-density limit in this region down to the CRATES flux-density limit of $\sim 65 \text{ mJy}$. Three sources were observed in this region at 4.85 GHz, with flux-densities of 67 mJy, 58 mJy and 142 mJy.

3.2 C-BASS Source Detection Algorithm

To detect sources in a sky map we need to remove obscuring diffuse emission and noise. For a source with a known point-spread function (PSF) embedded in additive noise, the matched filter (MF) is the optimal filter that can be applied to maximize the source SNR. The matched filter is given by

$$\Psi_{\text{MF}}(k) = \left[2\pi \int dk k \frac{\tau^2(k)}{P(k)} \right]^{-1} \frac{\tau(k)}{P(k)}, \quad (3.1)$$

where k is the Fourier wavenumber, $\tau(k)$ is the source profile and $P(k)$ is the power spectrum of the unfiltered map (Tegmark and de Oliveira-Costa, 1998; López-Caniego et al., 2006). However, the calculation of the MF involves a number of complications. Chiefly, we are required to make a noisy estimate of the power spectrum from the unfiltered map, and integrate it. In constructing the PCCS, it was found that the Mexican Hat wavelet of the second kind (MHW2) achieved similar performance to the MF (López-Caniego et al., 2006; Planck Collaboration et al., 2014b, 2016f). For the present analysis I adapt the *Planck* algorithm, using instead the SMHW2 in place of the flat-space MHW2. The SMHW2 is straightforward to calculate, enables us

to filter the entire sky at once, and allows us to optimize a few free parameters as opposed to a noisy estimate of the full noise power spectrum.

3.2.1 Source Detection Algorithm

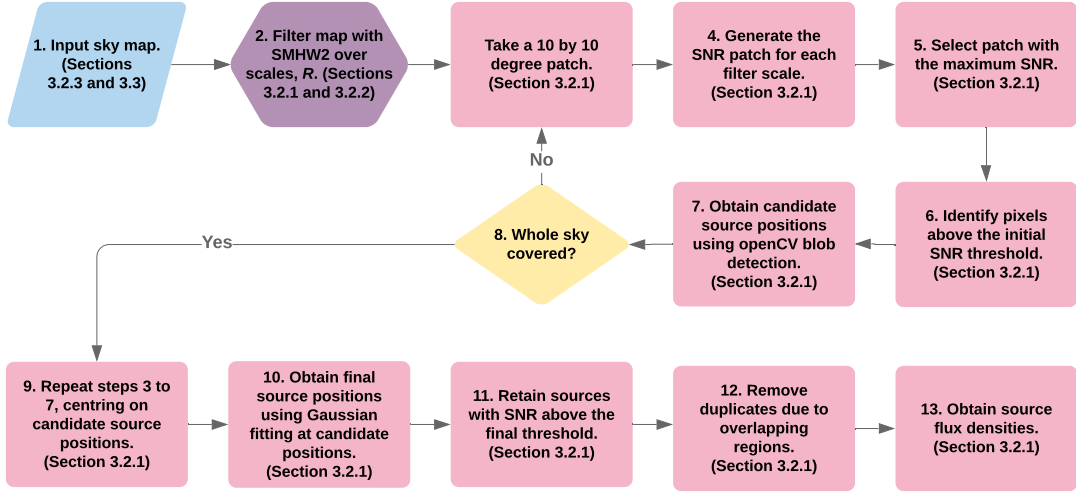


Figure 3.1: Flowchart for the C-BASS point-source detection algorithm. The SMHW2 filter is applied over the entire sky, whilst identification of point sources is performed on $10^\circ \times 10^\circ$ sky regions. Labelled sections refer to the relevant sections in this chapter.

Given a sky map consisting of point sources, diffuse emission and instrumental noise, we can enhance the SNR of point sources in the map by filtering with the SMHW2. I perform this filtering at a range of filter scales R , to maximize the number of sources extracted from the sky maps. In changing R , we effectively change the extent to which we down-weight large-scale ℓ -modes that are dominated by diffuse emission, and also small scales dominated by instrumental noise. This is particularly important in regions close to the Galactic plane where diffuse emission is very strong, meaning we must down-weight large-scale modes harshly in order to extract point sources. Conversely, in regions with little diffuse emission we may wish to be less extreme in the down-weighting, so that we do not excessively reduce point-source power and subsequently miss detection of fainter sources in these regions. I discuss the form of the SMHW2 filter and the effect of the filter scale, R in detail in Section 3.2.2.

The source-detection algorithm is described below, with the algorithm flowchart displayed in Fig. 3.1.

1. I begin by taking the spherical harmonic transform of the sky map, obtaining the $a_{\ell m}$ coefficients. I filter the map by weighting the $a_{\ell m}$ coefficients with the SMHW2 window function, $w_{\ell}^{S^2}(R)$ at some user-defined set of filter scales, i.e. I calculate

$$f_{\ell m}(R) = a_{\ell m} w_{\ell}^{S^2}(R). \quad (3.2)$$

The effect of this filtering is demonstrated in Fig. 3.2. Here I display a simulated sky consisting of diffuse emission, instrumental noise and point-sources, along with the corresponding filtered sky with scale $R = 1$. The details of the production of this simulated sky are discussed in Section 3.2.3. The filter has removed diffuse emission and enhanced the SNR of point-sources across the sky.

2. I break the maps, filtered at a range of R scales, into overlapping $10^\circ \times 10^\circ$ patches on the sky and calculate the corresponding SNR patches. I used R scales in the range $0.5 \leq R \leq 2$ in steps of 0.1. $10^\circ \times 10^\circ$ patches are used such that the flat-sky approximation holds, whilst capturing a sufficient sample of the sky background in the given sky region. The overlap is set to 5° in each direction to ensure that accurate positions are obtained for sources detected close to patch edges. The typical background fluctuation level in each patch is estimated by calculating the median absolute deviation (MAD) of the pixels, defined as

$$\text{MAD}(\mathbf{X}) = \text{median}(|X_i - \text{median}(\mathbf{X})|), \quad (3.3)$$

where \mathbf{X} is the set of pixel values. The MAD estimator avoids sensitivity to outlier pixel values, which can be caused by the presence of point sources in the sky patch. For each sky patch, I then select the filter scale that gives the maximal SNR for the patch i.e., the maximum of the peak pixel value divided by the MAD estimate.

3. Having selected the value of R that maximizes the patch SNR, an initial SNR threshold is applied to identify potential point sources. This produces a patch of tagged and un-tagged pixels. The tagging is performed on these small, overlapping regions so as to account for the varying background properties across the sky i.e., variations in residual diffuse emission and noise.
4. Given a patch of tagged pixels, I then use the OPENCV Simple Blob Detector¹ to obtain candidate source locations. The Simple Blob Detector algorithm

¹<https://opencv.org/>

works by applying thresholds to an input image and determining the centres of connected pixels, or sources. In addition to providing candidate source locations, this software allows for the filtering of detected sources according to their geometric properties i.e., their size, circularity, convexity and inertia ratio. For this analysis, I set a maximum size limit to help prevent spurious tagging of bright diffuse emission and disallow detections within one beam FWHM of one another. A detailed discussion of OPENCV and the Simple Blob Detector algorithm can be found in [Kaehler and Bradski \(2016\)](#). The processing steps involved in studying these small sky regions are illustrated in Fig. 3.3, where I display the typical output from the steps outlined above.

5. Given a set of tentative source locations obtained over the whole sky, I now repeat steps (ii) and (iii), this time centering on the tentative source locations. In order to refine the source position estimates, I fit an elliptical Gaussian to a 20 arcmin square region of the wavelet-filtered map, centred on the candidate location obtained using OPENCV. Whilst the C-BASS beam and SMHW2 filter are not Gaussian, by fitting to this small region around the candidate location the Gaussian fit is sufficient to make the necessary refinements (~ 1 arcmin) to the estimates of source peak locations.
6. A final SNR thresholding is applied by dividing the peak value of the source in the wavelet filtered map by the MAD value of pixels in an annulus centred on the source, with an inner radius of 3° and outer radius of 5° . The annulus was chosen to avoid the first sidelobe in the C-BASS beam, whilst capturing the local background fluctuations around the source. Sources are retained if this new SNR estimate exceeds some final threshold. The SMHW2 is constructed so as to preserve point-source amplitudes after being convolved with the C-BASS beam. As with the PCCS, I convert these wavelet amplitudes to Jy and report them as auxiliary flux-density (DETFLUX) estimates alongside the primary estimates from aperture photometry described in step (vii). The DETFLUX estimates are discussed in more detail in Section 3.3.5.
7. The detected source locations from all of the filtered maps are combined and duplicates are removed from the catalogue, defined here to be any reported source positions that are within the beam FWHM of one another. For matched sources in the internal catalogue, I retain the source with the greatest detected SNR.

8. Given the final set of source positions, source flux densities are obtained. As its primary method, the detection pipeline obtains flux densities using aperture photometry, adapting the method from [Planck Collaboration et al. \(2014b\)](#); [Génova-Santos et al. \(2015a\)](#). I begin by converting the map to Jy pix^{-1} , and define an aperture of radius 45 arcmin around the source position, and an annulus of inner radius 3° and outer radius 5° . The observed flux density is then given by

$$S_{\text{obs}} = \kappa (S_{\text{ap}} - \bar{S}_{\text{ann}} N_{\text{ap}}), \quad (3.4)$$

where $\kappa \approx 1.34$ is a correction applied to account for flux density missing from the aperture (calculated using Equation A.2 of [Planck Collaboration et al. \(2014b\)](#)), S_{ap} is the total flux density in the aperture, \bar{S}_{ann} is the median flux density in the annulus and N_{ap} is the number of pixels in the aperture. The uncertainty was estimated as

$$\sigma(S_{\text{obs}}) \approx \text{MAD}(\mathbf{X}_{\text{ann}}) \frac{N_{\text{ap}}}{\sqrt{N'_{\text{ap}}}}, \quad (3.5)$$

where $\text{MAD}(\mathbf{X}_{\text{ann}})$ is the MAD value of pixels in the annulus and N'_{ap} is the number of beams inside the aperture. The scaling applied to $\text{MAD}(\mathbf{X}_{\text{ann}})$ is to account for correlations in the background emission, approximately on the scale of the beam ([Génova-Santos et al., 2015a](#)).


3.2.2 The Spherical Mexican Hat Wavelet

The Mexican Hat Wavelet (MHW) of the n^{th} kind is defined in \mathbb{R}^2 (i.e., 2-dimensional Euclidean space) as


$$\Psi_n(\mathbf{x}, R) \propto \Delta^n \left\{ \exp\left(-\frac{|\mathbf{x}|^2}{2(R\sigma)^2}\right) \right\}, \quad (3.6)$$

where \mathbf{x} is the position in the image plane, Δ is the usual Euclidean Laplacian operator, σ is the Gaussian standard deviation and R is the filter scale factor. In changing R , I change the characteristic scale of the filter and hence the extent to which the filter down-weights large- and small-scale modes. Applying this expression we can arrive at the real-space expression for the MHW2, given by

$$\Psi_2(\mathbf{x}, R) \propto \left(8(R\sigma)^4 - 8(R\sigma)^2 |\mathbf{x}|^2 + |\mathbf{x}|^4 \right) \exp\left(-\frac{|\mathbf{x}|^2}{2(R\sigma)^2}\right). \quad (3.7)$$



(a) Input sky



(b) Filtered sky

Figure 3.2: Panel (a): An input 4.76 GHz simulated sky consisting of diffuse emission generated using PYSM, instrumental noise generated from a white-noise realization of the C-BASS sensitivity map, along with a point-source population generated by scattering the GB6, PMN and RATAN-600 sources at random positions over the sky. Alongside this I display a zoom-in centred on the location of the black cross on the map. Panel (b): The same map after being convolved with the SMHW2 filter, removing the large-scale diffuse emission and leaving behind a sky filled with point-sources. As before, I display a zoom-in of the filtered map alongside it, centred on the location of the black cross. It is important to note here that not all the diffuse emission has been removed at low Galactic latitudes. In regions of strong diffuse emission the source detection algorithm is significantly less reliable. In panel (b) it is apparent that the SMHW2 filtering can still leave a significant sidelobe around bright sources. This can lead to issues with the spurious tagging of sidelobe peaks as sources, particularly for the brightest sources on the sky (i.e., those with flux-densities $\gtrsim 20$ Jy). This problem can largely be side-stepped by allowing for larger exclusion zones around the brightest sources, as was the case with catalogues such as GB6, removing apparent detections in the first sidelobe of such sources.

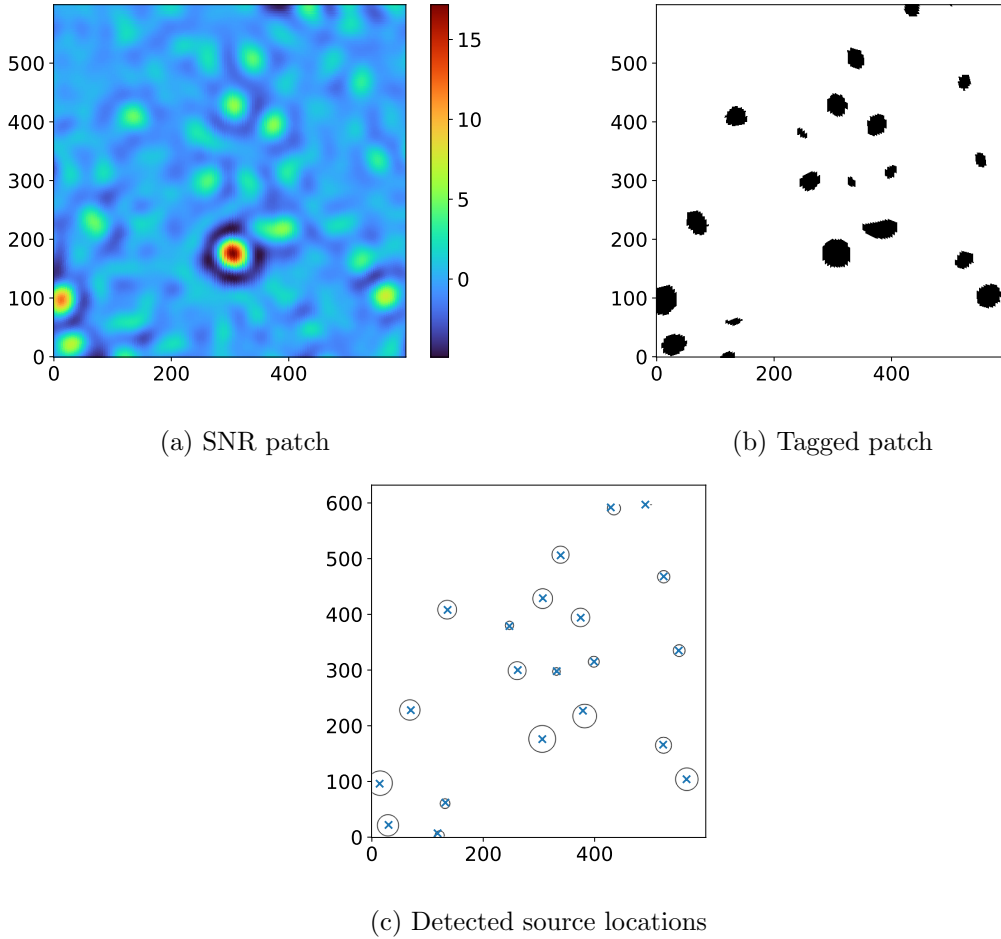


Figure 3.3: Images illustrating the pixel tagging process. Panel (a): An example $10^\circ \times 10^\circ$ SNR patch, produced by dividing the filtered sky patch by the MAD value of the pixels in the patch. Panel (b): The corresponding patch of tagged and un-tagged pixels. The tagged pixels are shown in black and un-tagged pixels in *white*, obtained by applying a threshold of $\text{SNR} \geq 3.5$ to the SNR patch. Panel (c): An image showing the detected source locations for this patch. I use the OPENCV blob detection algorithm to obtain preliminary source locations from this patch. These blob detections are shown as black circles centred on the preliminary source locations. These position estimates are then refined by fitting an elliptical Gaussian at the preliminary source locations, with the source positions obtained from the Gaussian fit shown as blue crosses.




Figure 3.4: Panel (a): The real space SMHW2, shown for σ corresponding to a Gaussian beam FWHM of 45 arcmin and $R \in \{0.85, 0.95, 1.05, 1.15\}$. Panel (b): The corresponding window functions. The SMHW2 down-weights large and small scales, with changes in R determining the severity with which large- and small-scale emission is down-weighted. Whilst the C-BASS beam is not exactly Gaussian, by constructing the SMHW2 around these ~ 45 arcmin scales I am able to filter the large- and small-scale modes around the characteristic beam scale, enhancing point-source SNR in the C-BASS sky map.

In applying the MHW2 to producing the PCCS, the sky was divided up into small patches, taking flat projections and applying the MHW2 filter to them. SNR thresholding was applied to the filtered patches to extract the source positions and flux densities were then obtained for the detected sources. Here I employ the SMHW2, the equivalent of the MHW2 on S^2 (i.e., the 2-sphere). For a given filter scale, this allows us to filter the entire sky at once, avoiding complications from edge effects when filtering small patches. In dealing with the C-BASS northern sky map, I apply a $5^\circ \cos^2$ apodization at the edge of the mask to mitigate edge effects.

In [Antoine and Vandergheynst \(1999\)](#) it was shown the continuous wavelet transform on S^2 can be constructed by taking the inverse stereographic projection of the \mathbb{R}^2 wavelet. This preserves the properties of the \mathbb{R}^2 wavelet and tends to the MHW in the small-angle limit. The spherical MHW of the n^{th} kind is given by

$$\Psi_{S^2}^{(n)} \propto \frac{1}{\cos^4 \frac{\theta}{2}} \Psi_n \left(|x| \equiv 2 \tan \frac{\theta}{2}, R \right), \quad (3.8)$$

where θ is the colatitude on the sphere ([Cayón et al., 2001](#); [Martínez-González et al., 2002](#); [Vielva et al., 2003](#); [Curto et al., 2011](#)).

Substituting the MHW2 into Equation 3.8, we obtain the SMHW2 as follows,

$$\Psi_{S^2}^{(2)}(\theta, R) \propto \frac{1}{\cos^4 \frac{\theta}{2}} \left[(R\sigma)^4 - 4(R\sigma)^2 \tan^2 \frac{\theta}{2} + 2 \tan^4 \frac{\theta}{2} \right] \exp\left(-\frac{2 \tan^2 \frac{\theta}{2}}{(R\sigma)^2}\right). \quad (3.9)$$

I normalize the SMHW2 such that point-source amplitudes are preserved after convolving with the filter, i.e.

$$\int_{S^2} \mathcal{B}(\theta) \Psi_{S^2}^{(2)}(\theta, R) \sin \theta \, d\theta d\phi = 1, \quad (3.10)$$

where $\mathcal{B}(\theta)$ is the C-BASS beam profile.

Given the analytic expression for the SMHW2 in real space on S^2 , we can obtain the SMHW2 window function by calculating the Legendre transform of Equation 3.9, i.e.

$$w_\ell^{S^2}(R) = 2\pi \int_0^\pi \Psi_{S^2}^{(2)}(\theta, R) P_\ell(\cos \theta) \sin \theta \, d\theta, \quad (3.11)$$

where $P_\ell(\cos \theta)$ is the Legendre polynomial of order ℓ .

In Fig. 3.4 I show the real-space SMHW2 and the corresponding window function, calculated for a σ value corresponding to a beam FWHM of 45 arcmin, with $R = \{0.85, 0.95, 1.05, 1.15\}$. The primary effect of the SMHW2 is to down-weight large scales, which are dominated by diffuse emission. In addition to this, the SMHW2 also acts to down-weight small-scale modes where instrumental noise becomes more significant. In the context of C-BASS, large-scale modes correspond to modes significantly larger than the C-BASS beam scale, $\ell \lesssim 100$, whilst small scales correspond to scales significantly smaller than the C-BASS beam scale, $\ell \gtrsim 1000$.

3.2.3 Algorithm Validation

I validate the algorithm by running it over 100 simulations of the sky at 4.76 GHz. The simulated skies were generated at NSIDE=1024, where the number of equal-area pixels on the Healpix sphere is given by $12 \times \text{NSIDE}^2$. This high resolution was chosen to allow for the precise determination of source positions, with the C-BASS PSF over-sampled by a factor of ~ 6.5 . The simulated sky includes diffuse emission generated by PYSM (Thorne et al., 2017) and an instrumental noise realization. The noise maps were created by generating white noise realizations of the C-BASS sensitivity map (including a simulated sensitivity map for the southern sky survey). This accounts for the declination dependent effects from the C-BASS scan strategy.

To the diffuse emission simulation I add a source population to the sky. These are created by taking source flux-densities from the PMN, GB6 and RATAN-600 surveys,

summarized in Section 3.1, and scattering the sources at random positions over the sky. Carrying this out 100 times I obtain 100 Monte Carlo simulations of the sky, consisting of the same diffuse background but with their source populations randomly distributed on the sky, and differing noise realizations. The simulated skies are all convolved with the C-BASS beam.

The diffuse components included in the simulation are as follows:

1. **Synchrotron Radiation** This was generated using the de-sourced, re-processed 408 MHz Haslam map (Haslam et al., 1981, 1982; Remazeilles et al., 2015) as a template and scaled using a spectral index and curvature model. This corresponds to the PYSM s3 model. Pixel values at 4.76 GHz range from $\sim 3000 \mu\text{K}$ to $\sim 8 \times 10^5 \mu\text{K}$, with a mean value of $\sim 2 \times 10^4 \mu\text{K}$.
2. **Free-Free Emission** This was generated with the analytic model used in the *Planck* COMMANDER analysis (Draine, 2011; Planck Collaboration et al., 2016d,c). The free-free template is scaled using a single power law. This corresponds to the PYSM f1 model. Pixel values at 4.76 GHz range from $\sim 90 \mu\text{K}$ to $\sim 5 \times 10^6 \mu\text{K}$, with a mean value of $\sim 1 \times 10^4 \mu\text{K}$.
3. **Anomalous Microwave Emission (AME)** This was generated by modelling emission caused by the sum of two spinning-dust populations. Scaling is calculated using the SPDUST2 code (Ali-Haïmoud et al., 2009). This corresponds to the PYSM a2 model. Pixel values at 4.76 GHz range from $\sim 5 \mu\text{K}$ to $\sim 2 \times 10^5 \mu\text{K}$, with a mean value of $\sim 300 \mu\text{K}$.
4. **Thermal Dust** This was generated by modelling the emission caused by a single modified black-body. This corresponds to the PYSM d1 model (Planck Collaboration et al., 2016d). Pixel values at 4.76 GHz range from $\sim 0.01 \mu\text{K}$ to $\sim 20 \mu\text{K}$, with a mean value of $\sim 0.4 \mu\text{K}$.
5. **CMB** A lensed CMB realization was generated using the Taylens code (Naess and Louis, 2013), as part of PYSM. This corresponds to the PYSM c1 model. Pixel values at 4.76 GHz range from $\sim -300 \mu\text{K}$ to $\sim 300 \mu\text{K}$, with a mean value of $\sim 0.1 \mu\text{K}$.

One of the simulated skies is shown in Fig. 3.2, along with the same sky filtered at the scale $R = 1$. After running the detection algorithm over the simulated skies, I compare the detected sources to the input catalogues. From this, I calculate the completeness level of the detected sources, averaged over all the simulations, and record

the absolute deviations of detected sources from their true positions. The validation results are presented in Section 3.2.5, produced using an initial SNR threshold of 2.5 to tag candidate sources on the first loop over the sky, and retaining sources with $\text{SNR} \geq 3.5$ after looping over the candidate sources. These thresholds were selected to give a catalogue reliability $\gtrsim 90$ per cent away from the Galactic plane. Catalogue completeness is defined as the fraction of sources above some given flux-density threshold that are recovered by the detection algorithm. The catalogue reliability is defined as the fraction of detected sources that are matched with at least one source in the input catalogue.

3.2.4 Catalogue Matching: Likelihood Ratios

The simulated validation catalogue is produced for a map with a resolution of 45 arcmin. Given the large number of sources used for the input catalogue, source confusion presents a major complication when trying to match the output catalogues with the corresponding input catalogues from the Monte Carlo simulations. One source in the output catalogue will likely result from the blending of multiple sources in the input catalogue, and a single source in the input catalogue can contribute to more than one source in the output catalogue. Identical issues are faced when matching the real C-BASS catalogue to higher resolution catalogues such as GB6 and PMN.

To find matches between the input and output catalogues I use a likelihood ratio test. This has previously been used in matching point-source catalogues, particularly when dealing with problems from source confusion, see e.g., Richter (1975); Sutherland and Saunders (1992); Mann et al. (1997); Rutledge et al. (2001); Chapin et al. (2011); Wang et al. (2014). The core of the likelihood ratio test consists in calculating the ratio of the probability of a match between an input and output source being a true match, to the probability of the match being a random association. The likelihood ratio is given by

$$\text{LR} = \frac{q(\Delta S_f) f(r)}{2\pi r \rho(\Delta S_f)}, \quad (3.12)$$

where $q(\Delta S_f)$ is the probability density function (PDF) for a given match being a true match as a function of the fractional flux-density difference (ΔS_f) between the two sources, $f(r)$ is the PDF for true matches as a function of the absolute positional offset between the two sources, r , $2\pi r$ is the PDF for random associations assuming a uniform spatial distribution of background sources, and $\rho(\Delta S_f)$ is the PDF for random associations as a function of ΔS_f .

For the positional PDF of true input-output matches, I assume a Gaussian distribution over the orthogonal coordinate axes. This gives a Rayleigh distribution for the PDF of true input-output matches as a function of r , i.e.

$$f(r) = \frac{r}{\sigma_r} \exp\left(-\frac{r^2}{2\sigma_r^2}\right), \quad (3.13)$$

where σ_r is a scaling parameter. The position and flux-density PDF parameters are estimated by comparing the output source catalogues with the corresponding input catalogues. I determine $f(r)$ by searching for all candidate matches between the input and output catalogues, using a search radius of 45 arcmin around each output source. Candidate matches will consist of true input-output matches, along with random background associations. Assuming a uniform distribution of background sources the random matches will follow a linear trend, scaling with the sky area enclosed by a given search radius. I therefore fit $f(r)$ plus a linear trend to the histogram of absolute positional offsets for candidate input-output matches to obtain an estimate for $f(r)$.

To determine $\rho(\Delta S_f)$, I generate randomized output catalogues by randomly distributing the output sources over the sky and randomizing source flux-densities from the parent distribution. I then find candidate matches between input and randomized output sources, again using a 45 arcmin search radius. I bin the candidate matches as a function of the fractional flux-density difference between the input and output source, given by

$$\Delta S_f = \frac{S_{\text{out}} - S_{\text{in}}}{S_{\text{in}}}, \quad (3.14)$$

where S_{out} is the output source flux-density and S_{in} is the input source flux-density. This gives the number of random matches as a function of ΔS_f and hence the estimate for $\rho(\Delta S_f)$. To obtain $q(\Delta S_f)$ I carry out the same procedure as above, but with the real output catalogues. I then estimate $q(\Delta S_f)$ from the difference between the number of matches as a function of ΔS_f for the real output catalogues, and the number of matches for the randomized catalogues, i.e. from the excess matches in the real output catalogues.

To determine the LR threshold for declaring a true match I calculate the likelihood ratios for all candidate matches between the input catalogues and randomized output catalogues. From this I determine the value of LR at which 10 per cent of the randomized matches would be declared to be true matches. This value is used as the LR threshold at which I declare a true input-output match. This is sufficient for limiting the fraction of false identifications with the input catalogue, whilst ensuring a low probability of spuriously identifying a source in the output catalogue as a true

positive, given that matched output sources are typically associated with multiple input sources.

3.2.5 Validation: Completeness, Position and Flux-Density Accuracy




Figure 3.5: The combined Galactic plane and C-BASS north mask used in the analysis of the C-BASS catalogue, shown here in Galactic coordinates. The mask was constructed to cover the brightest 30 per cent of the northern sky observed by C-BASS, along with declinations $\delta < -10^\circ$. I refer to this mask as the CG30 mask.

I obtain the catalogue completeness and reliability, averaged over the Monte Carlo simulations, to confirm the ability of the source-detection algorithm to accurately recover an input source population. Throughout this analysis I use the mask shown in Fig. 3.5 to characterize the catalogue properties away from the brightest diffuse emission along the Galactic plane. The mask was produced by smoothing the C-BASS northern sky map with a 10° FWHM Gaussian and masking the brightest 30 per cent of pixels. I also mask all pixels corresponding to declinations below $\delta < -10^\circ$ in order to match the sky area observed by the C-BASS northern-sky survey. I refer to this mask as the CG30 mask.

In determining the positional offset of the detected source from the input sources, I compare its position to the photocentre of matched input sources. That is, after matching to input sources using the LR test, I determine the weighted average position of the matched input sources, using the input source flux-densities as weights. In comparing flux densities, I sum flux densities for matched sources, weighted according to the C-BASS beam profile, i.e.

$$S_\nu^{\text{match}} = \sum_i S_\nu^{(i)} \mathcal{B}(|\mathbf{x}_i - \mathbf{x}_{\text{det}}|), \quad (3.15)$$

where the sum runs over matched sources in the input catalogues, $S_{\nu}^{(i)}$ is the flux-density of the i^{th} matched source, \mathbf{x}_i is the position of the i^{th} matched source, \mathbf{x}_{det} is the position of the detected source and \mathcal{B} is the C-BASS beam profile.

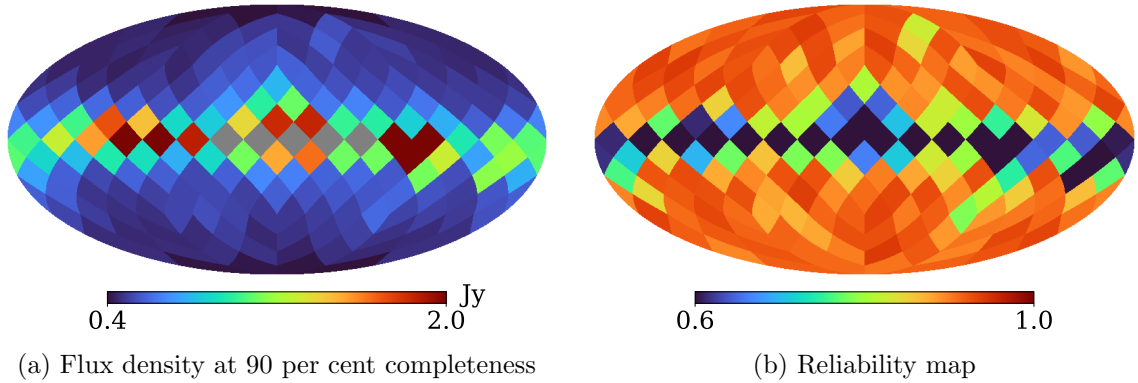


Figure 3.6: Panel (a): A map of estimates of the flux density at 90 per cent completeness obtained from the validation simulations. Panel (b): A map of the corresponding reliability estimates. These estimates were calculated on $N_{\text{SIDE}}=4$ sky pixels and are shown in Galactic coordinates. In pixels outside of the CG30 mask, the mean 90 per cent completeness level is approximately 500 mJy, with a corresponding mean reliability of 91 per cent. It can be seen that along the Galactic plane the completeness and reliability is significantly lower than at higher Galactic latitudes. This is to be expected given the stronger diffuse emission that is not removed by the SMHW2 filter. Grey pixels in the completeness map denote pixels where 90 per cent completeness was never achieved.

In Fig. 3.6 I show maps of the completeness and reliability in $N_{\text{SIDE}}=4$ pixels obtained for the validation simulations. We can see that close to the Galactic plane the 90 per cent completeness level is higher and the catalogue reliability is lower. Indeed, for some pixels along the Galactic plane 90 per cent completeness is never achieved. This is to be expected given the increased intensity of diffuse emission along the Galactic plane, which is not fully removed by the SMHW2 filter. Over regions of the sky left un-masked by the CG30 mask, I find a mean 90 per cent catalogue completeness of approximately 500 mJy, corresponding to a mean reliability of approximately 91 per cent.

From the validation simulations, I calculate the absolute positional offsets of detected sources from input sources. These cluster below 10 arcmin, peaking around zero, with a median offset of approximately 3.2 arcmin. Considering the 45 arcmin resolution of the maps used for source detection, and the significant confusion effects at this resolution, this is a small positional uncertainty for the source detection




Figure 3.7: The histogram of measured absolute positional offsets, Δ , applying the CG30 mask, obtained for the 100 Monte Carlo simulations. Over the un-masked sky area I find the detected positional offsets peak around zero, with a median positional offset of 3.2 arcmin, which is small compared to the C-BASS beam of 45 arcmin. Larger positional offsets are predominantly driven by fainter detected sources below ~ 1 Jy, and sources detected in regions of brighter diffuse emission. The positional offsets of all the detected sources will consist of samples drawn from Rayleigh distributions, which results in a tail of larger offsets beyond the peak of the positional offset distribution.

algorithm. Indeed, for the purposes of producing a point-source mask for the actual C-BASS maps this is more than sufficient. The histogram of absolute positional offsets of detected sources, applying the CG30 mask, is shown in Fig. 3.7.

I am able to recover input source flux-densities using aperture photometry, with a plot of detected flux densities against input flux densities being shown in Fig. 3.8. At flux densities greater than 1 Jy the majority of sources have flux-densities evenly distributed around the input flux-densities. A small fraction of sources (~ 1 per cent) are recovered with flux densities biased slightly high. This is to be expected given the 45 arcmin resolution of the C-BASS map, meaning that multiple input sources are associated with each detected source. The detected source flux density will be the summed contribution from associated input sources. Moreover, bright sources or extended emission in the vicinity of a source can spuriously contribute to the aperture photometry estimate of the source flux-density. For flux densities below ~ 1 Jy the scatter in the recovered flux densities reaches the same magnitude as the input flux densities. This is again to be expected; below this level sources become much more




Figure 3.8: Panel (a): The detected source flux densities against the matched input flux densities for the 100 Monte Carlo simulations, applying the CG30 mask. Panel (b): The histogram of the detected flux densities, divided by the matched input flux-densities, for sources with input flux densities greater than 10 Jy. We can see that for the majority of bright sources reliable flux-density estimates are obtained. However, there is a small population (~ 1 per cent) of detected sources with spuriously high flux-density estimates, due to confusion effects.

heavily obscured by surrounding diffuse emission, making photometric estimates of source flux densities significantly more unreliable. I note here that there is a small cluster of recovered sources at these low flux densities where the recovered flux-density estimate is spuriously high. This is largely caused by sources being detected in regions of bright diffuse and highly extended emission, leading to errors in the background estimation.

The validation simulations have demonstrated that the source-detection algorithm is able to accurately recover an underlying source population, from a sky consisting of additional contributions from diffuse emission and noise.

3.3 The C-BASS Northern Sky Catalogue

In this section I present the C-BASS northern sky, total-intensity point-source catalogue. The point-source catalogue was produced using data obtained from only night-time observations, combining data from the various elevation scans of the telescope. I perform the source detection on the C-BASS map at NSIDE=1024, using the high NSIDE map to allow for the precise determination of source positions.

The outline of this section is as follows: In Section 3.3.1 I summarize the catalogue properties. In Section 3.3.2 I present the results from matching the C-BASS and GB6/PMN catalogues, using these results to quantify the C-BASS pointing accuracy and flux-density scale. In Section 3.3.3 I comment briefly on sources detected with outlier flux densities. In Section 3.3.4 I show the differential source counts for the C-BASS catalogue, and compare these to the GB6 source counts. In Section 3.3.5 I provide a sample of the first ten sources in the C-BASS catalogue. In Section 3.3.6 I show a source-subtracted C-BASS northern sky map, produced by the C-BASS collaboration using the source catalogue discussed here. Additional comparisons were also made with sources in the RATAN-600 catalogue over the NCP region. In this case I recovered the expected bright source population, with the results obtained through these comparisons being consistent with those from comparisons with the GB6/PMN catalogues.

3.3.1 Northern Sky Intensity Catalogue: Summary

For the default catalogue, I use the same SNR thresholds as for the validation analysis i.e., $\text{SNR} \geq 2.5$ to tag candidate sources on the first loop over the sky, followed by a final detection threshold of $\text{SNR} \geq 3.5$ after looping over the candidate source locations. To estimate the completeness and reliability of the real C-BASS catalogue I compare it with the GB6 and PMN catalogues over their areas of common sky coverage. Applying the CG30 mask I obtain a catalogue reliability of approximately 98 per cent, with a corresponding 90 per cent completeness level of 610 mJy, recovering the expected bright source population at 4.76 GHz. This analysis has resulted in a catalogue of 1784 point sources covering declinations $\delta \geq -10^\circ$. The catalogue properties are summarized in Table 3.2, with the detected source positions for the catalogue shown in Fig. 3.9.

Beam sidelobes around bright sources can result in spurious source detections. The first sidelobe of the C-BASS beam has some azimuthal structure, and has a peak value of about 2.5 percent of the peak. Given the source detection level of around 500 mJy, we might therefore expect spurious detections from sources brighter than 20 Jy. Outside of the CG30 mask there are three sources brighter than this, 3C273, Virgo A and 3C279. Therefore I do not allow any additional detections within a 2° radius exclusion zone centred on these three sources. In constructing a point-source mask, any issues from spurious sidelobe detections can be further mitigated by masking larger areas around the brightest sources. An example of these spurious

Table 3.2: C-BASS northern sky, total intensity catalogue summary. Position uncertainties include intrinsic uncertainty from the source detection algorithm, along with additional errors from the C-BASS pointing.

C-BASS Catalogue	
Frequency	4.76 GHz
Map FWHM	45 arcmin
SNR threshold	3.5
Reliability:	
Full Sky	90%
CG30 mask	98%
Number of sources:	
Full sky	1784
CG30 mask	1136
90% completeness:	
Full sky	1000 mJy
CG30 mask	610 mJy
$N(S > S_{90})^a$:	
Full sky	793
CG30 mask	560
Position uncertainty	
Full sky	3.4 arcmin
CG30 mask	3.2 arcmin

^a This is the number of sources with measured flux densities greater than the 90 per cent completeness level.




Figure 3.9: Positions of the northern sky, total intensity C-BASS point sources. Positions shown here are in Galactic coordinates. *Filled purple circles* denote sources that fall within the CG30 mask. *Open cyan circles* denote sources outside the mask, with a fitted flux density below the 90 per cent completeness level (610 mJy), and *red crosses* denote the remaining sources with fitted flux densities above the 90 per cent completeness level.

detections can be seen in Fig. 3.10, where I show a gnomonic projection centred on Virgo A, highlighting the spurious detections in the first sidelobe.

It is important to note here that, whilst there are a large number of source detections in the Galactic plane, the reliability of these detections is significantly lower than at higher Galactic latitudes. This is a result of the more intense diffuse emission obscuring point sources, and highly extended emission in the Galaxy being spuriously identified as point sources. For the purposes of scientific analyses with other CMB datasets this is not a major issue, given that much of the Galactic plane will be masked out for such analyses.

3.3.2 GB6 and PMN Matching

To cross-check the C-BASS pointing and flux-density scale, I match the C-BASS point sources with the GB6 and PMN catalogues over their areas of common sky coverage. Similar to the situation encountered during the algorithm validation, the GB6 and PMN catalogues were produced using much higher resolution surveys than C-BASS. What may appear as a single source in the C-BASS maps could actually correspond to multiple sources in the GB6 and PMN catalogues. Matches between




Figure 3.10: $10^\circ \times 10^\circ$ gnomonic projection centred on Virgo A, showing spurious detections in the first sidelobe as white circles. Note that there is a real source located in the upper right of the Virgo A sidelobe. However, the tagging of the sidelobe means that the source position is not well determined. These detections are removed from the C-BASS catalogue, with a 2° radius exclusion zone being applied around 3C273, Virgo A and 3C279.

catalogues are found using the likelihood ratio technique described in Section 3.2.4. In characterizing the C-BASS pointing, I compare the C-BASS source position with the photocentre (defined in Section 3.2.4) of matched GB6 or PMN sources. In comparing flux densities I use the beam-weighted sum of matched source flux densities, defined in Equation 3.15.

In Fig. 3.11 I show a histogram of the absolute positional offsets of C-BASS sources from matched GB6 and PMN sources, using the photocentres of matched GB6 and PMN sources for comparison. The offsets are clustered below 10 arcmin, peaking at approximately 3.2 arcmin^2 . This is broadly consistent with the results obtained from the Monte Carlo simulations during the algorithm validation, with additional errors arising here from the C-BASS pointing. The distributions of the absolute positional offsets when matching to GB6 and PMN are also very similar. Larger positional offsets ($\gtrsim 10 \text{ arcmin}$) are predominantly driven by fainter detected sources below $\sim 1 \text{ Jy}$ and sources detected in regions of brighter diffuse emission. Further, given a set of positional offsets drawn from Rayleigh distributions, we expect

²Note that the position of the mode for a Rayleigh distribution is given by σ_r .




Figure 3.11: A histogram of the absolute positional offsets of C-BASS sources from matched GB6 and PMN sources, applying the CG30 mask, plotted alongside the corresponding histogram of positional offsets obtained from the validation simulations. I use the photocentres of matched GB6 and PMN sources for comparison. Positional offsets cluster below 10 arcmin, peaking at approximately 3.2 arcmin. We can see that the peak of the positional offsets obtained for the real C-BASS catalogue peak ~ 3 arcmin higher than for the validation simulations. This is to be expected given there is some level of pointing jitter in the real C-BASS data that is not present in the simulations. Similar distributions are obtained for the absolute positional offsets with both GB6 and PMN.

the overall distribution to have a tail of larger offsets. The typical positional offsets determined using the C-BASS source catalogue are consistent with the estimated pointing accuracy of C-BASS (Pearson et al. (in prep.)).

I also plot the C-BASS source flux densities (determined using aperture photometry, as described in Section 3.2.1) against matched GB6 and PMN flux-densities, using Equation 3.15 to determine the matched flux density, in Fig. 3.12. When no Galactic plane mask is applied, we can see that most sources cluster around the line $S_{\nu}^{\text{GB6}} = S_{\nu}^{\text{CBASS}}$ ($S_{\nu}^{\text{PMN}} = S_{\nu}^{\text{CBASS}}$). However, there is an apparent population of sources where the measured C-BASS flux density is significantly greater than the GB6 and PMN flux densities. This is largely due to the tagging of sources in the Galactic plane, where the additional presence of very bright, highly extended emission leads to spuriously high flux-density estimates. On applying a Galactic plane mask this population is largely removed. The small number of remaining sources where C-BASS significantly over-estimates the flux densities may either result from sources sitting in

regions of bright diffuse emission (e.g., within a Galactic spur) or from heavily confused sources in the C-BASS map. Most of the sources recovered with spuriously high densities are associated with low flux density GB6/PMN sources ($\lesssim 1$ Jy). At these flux densities the uncertainties in the flux-density estimates from aperture photometry become much larger, consistent with the larger spread in C-BASS to matched flux densities at these levels.




Figure 3.12: Flux-flux plots for matched C-BASS and GB6/PMN sources. Panel (a): A plot of C-BASS flux densities against matched GB6/PMN flux densities, with no mask applied. Panel (b): A plot of C-BASS flux densities against matched GB6/PMN flux densities, applying the CG30 mask. I also plot the line $S_{\nu}^{\text{GB6/PMN}} = S_{\nu}^{\text{CBASS}}$ in each panel. When I do not apply a Galactic mask, we can see a group of C-BASS sources with significantly higher flux densities than the matched GB6/PMN sources. This can be attributed to errors in fitting for source flux densities in the Galactic plane, particularly in the presence of bright highly extended emission. On applying a Galactic plane mask, this group of sources is largely removed and sources cluster more tightly around $S_{\nu}^{\text{GB6/PMN}} = S_{\nu}^{\text{CBASS}}$. Apparent outliers in the flux-density matches can largely be attributed to variable, flat-spectrum sources. Removing any C-BASS sources identified with flat-spectrum sources in the CRATES catalogue, I remove most of these outliers and improve the consistency between the C-BASS and GB6/PMN flux-density scales.

To quantify the agreement between the C-BASS and GB6/PMN flux-density scales I fit a straight line between the C-BASS and GB6/PMN flux-densities,

$$S_{\nu}^{\text{CBASS}} = b S_{\nu}^{\text{GB6/PMN}}, \quad (3.16)$$

where b is the gradient between the C-BASS and GB6/PMN flux-densities. I use a measurement-error model (Stan Development Team, 2012) to account for uncertain-

ties in the C-BASS and GB6/PMN flux densities by treating the true value of the GB6/PMN flux densities as missing data and marginalising over the distribution of the measured GB6/PMN flux densities. I also model the outlier distribution following [Hogg et al. \(2010\)](#). For this I assume outliers are drawn from some broad Gaussian distribution specified by a mean μ_{out} and standard deviation σ_{out} . The probability of a data point being an outlier is given by the parameter P_{out} . The full set of priors is given by,

$$b \sim \mathcal{N}(\mu = 1, \sigma = 0.25), \quad (3.17)$$

$$\mu_{\text{out}} \sim \mathcal{N}(\mu = 1 \text{ Jy}, \sigma = 100 \text{ Jy}), \quad (3.18)$$

$$\sigma_{\text{out}} \sim \text{Exp}(\lambda = 1/100 \text{ Jy}^{-1}), \quad (3.19)$$

$$P_{\text{out}} \sim \text{Unif}(0, 1), \quad (3.20)$$

$$S_{\nu}^{x,i} \sim \text{Unif}(0, \infty), \quad (3.21)$$

$$S_{\nu}^{\text{GB6/PMN},i} \sim \mathcal{N}(\mu = S_{\nu}^{x,i}, \sigma = \sigma_{\text{GB6/PMN}}^i), \quad (3.22)$$

where $S_{\nu}^{x,i}$ is the true GB6/PMN flux-density for the i^{th} matched source, $S_{\nu}^{\text{GB6/PMN},i}$ is the measured GB6/PMN flux-density for the i^{th} matched source and $\sigma_{\text{GB6/PMN}}^i$ is the corresponding uncertainty on the GB6/PMN flux density. The likelihood is given by

$$\mathcal{L} \propto \prod_{i=1}^{N_s} \left[\frac{1 - P_{\text{out}}}{\sqrt{2\pi} (\sigma_{\text{CBASS}}^i)^2} \exp \left\{ -\frac{(S_{\nu}^{\text{CBASS},i} - b S_{\nu}^{x,i})^2}{2 (\sigma_{\text{CBASS}}^i)^2} \right\} + \frac{P_{\text{out}}}{\sqrt{2\pi} (\sigma_{\text{out}}^2 + (\sigma_{\text{CBASS}}^i)^2)} \exp \left\{ -\frac{(S_{\nu}^{\text{CBASS},i} - \mu_{\text{out}})^2}{2 (\sigma_{\text{out}}^2 + (\sigma_{\text{CBASS}}^i)^2)} \right\} \right], \quad (3.23)$$

where N_s is the number of matched sources, $S_{\nu}^{\text{CBASS},i}$ is the measured C-BASS flux density for the i^{th} matched source and σ_{CBASS}^i is the corresponding uncertainty on the C-BASS flux density. I sample the model using NUTS algorithm, implemented in PYMC3.

[Pauliny-Toth et al. \(1978\)](#) showed that approximately 60 per cent of strong radio sources at 5 GHz have flat spectra and are therefore likely to be compact, variable sources. Indeed, previous studies of the GB6 source catalogue have found a significant fraction of variable sources ([Gregory et al., 1996, 1998, 2001](#)). As such, we would expect roughly half of the objects in 5 GHz surveys to be varying, meaning any comparison of individual flux densities observed at different epochs will be strongly

contaminated by source variability. A comparison of flux densities over the whole catalogue remains useful for checking the overall consistency of the different flux density scales. However, the CRATES catalogue can be used as a proxy for potentially variable sources in the C-BASS catalogue. By matching the C-BASS catalogue with the corresponding 4.85 GHz sources in the CRATES catalogue, I can remove any C-BASS sources that are positively identified with a CRATES source. In doing so, I obtain a set of CRATES cleaned C-BASS sources with the flat-spectrum sources largely removed, thereby removing the primary source of variability in the C-BASS catalogue. I perform the flux-density fits described above to both the full C-BASS catalogue matches, along with the CRATES cleaned matches. The latter set of fits will be better able to detect any instrumental or systematic errors in the C-BASS flux-density scale.

The results for the fitted values of b , comparing to GB6 and PMN, are given in Table 3.3. There is a slight discrepancy between the C-BASS flux-density scale and the PMN scale, when applying the CG30 mask and comparing to the full C-BASS catalogue. This discrepancy is removed when I remove any C-BASS sources associated with a CRATES source. A similar pattern is found for GB6, with the overall consistency being improved by cleaning the C-BASS catalogue of flat-spectrum sources. The C-BASS flux-density scale is consistent with the GB6 flux-density scale to within ~ 4 per cent. For the PMN catalogue there are fewer sources for comparison in the common overlap region, meaning consistency is only established in this case to ~ 10 per cent. Differences between the flux-density scales will result from a number of complicating factors e.g., the fact that C-BASS sources are really blends of multiple GB6/PMN sources, differing pass-bands, source variability etc. Given this, and the additional ~ 5 per cent uncertainties in the GB6 and PMN flux-density scales, I conclude that the C-BASS flux-density scale is consistent with these flux-density scales. I do not include any colour corrections in this analysis. The choice of 4.76 GHz as the effective frequency for C-BASS minimizes colour corrections to ~ 0.02 per cent for a source with spectral index $\alpha = 0$, and ~ 0.2 per cent for a source with spectral index $\alpha = -1$. Colour corrections for GB6/PMN are negligible due to their narrow bandwidths. Corrections due to flux-density extrapolation are also not considered here. For a typical source with spectral index, $\alpha \approx -0.83$ the correction is ~ 1.5 per cent. Given the uncertainties on the fitted values of b , and the additional uncertainties on the GB6/PMN flux-density scales, any such corrections to the flux-densities of C-BASS sources will not affect the overall conclusions here.

Table 3.3: Fitted values for b , the ratio between C-BASS and GB6 or PMN flux densities. Values were obtained by comparing flux densities between C-BASS and GB6/PMN, for both the full C-BASS catalogue and the CRATES-cleaned C-BASS catalogue. Alongside the mean fitted value for b , I also give the 95 per cent highest posterior density (HPD) interval. By removing sources from the C-BASS catalogue associated with sources in the CRATES catalogue, the overall consistency between the C-BASS and GB6/PMN catalogues is improved.

Comparison Catalogue	CRATES Cleaned?	$\langle b \rangle$	95% HPD[b]
GB6 no mask	No	0.95	{0.90, 1.00}
GB6 CG30 mask	No	0.95	{0.91, 1.00}
GB6 no mask	Yes	1.04	{0.97, 1.11}
GB6 CG30 mask	Yes	1.04	{0.97, 1.11}
PMN no mask	No	0.95	{0.89, 1.00}
PMN CG30 mask	No	0.82	{0.68, 0.97}
PMN no mask	Yes	1.05	{0.93, 1.17}
PMN CG30 mask	Yes	1.01	{0.83, 1.19}

3.3.3 Outlier Sources

In Fig. 3.12 we can see a number of C-BASS sources showing discrepancies with the GB6/PMN flux densities. As noted above, we expect a significant number of sources with flux densities $S_\nu \gtrsim 1$ Jy to be flat-spectrum, variable sources. If I extract sources from the CRATES cleaned C-BASS catalogue that show a $\geq 3\sigma$ discrepancy between the C-BASS and GB6/PMN flux densities, I find six sources. Gnomonic projections, centred on the source locations, are shown in Fig. 3.13. In four of the cases the matched GB6 flux densities are below 500 Jy. In this situation aperture photometry struggles to yield reliable point estimates of individual flux densities, contaminating any direct comparison between the C-BASS and GB6 flux densities. Considering instead the DETFLUX estimates, which are more reliable at these faint flux-densities, two of these sources remain as apparent outliers. In both these cases the detections are within a large extended structure, which can lead to poor peak identification and hence inaccurate DETFLUX estimates. The remaining two outlier sources were identified as QSO B0723-007 and Virgo A.

QSO B0723-007 was identified as having a discrepant flux density when matching to PMN (5.5 ± 1.2 Jy in C-BASS compared to 1.4 ± 0.07 Jy in PMN). This is a




Figure 3.13: $10^\circ \times 10^\circ$ gnomonic projections, centred on the locations of sources in the C-BASS catalogue, identified as having a $\geq 3\sigma$ discrepancy with the matched GB6/PMN flux-densities. Panel labels give the right ascension and declination values for the projection centres. The sources in panels in (a) to (d) were matched with GB6 sources with flux-densities below 500 mJy. The sources are also surrounded by bright diffuse and extended features, which act to make any aperture photometry estimates unreliable. Considering instead the DETFLUX estimates, only the sources in panels (b) and (d) are left as apparent outliers. However, both of these sources are located in large, extended structures. This can lead to poor peak identification and hence inaccurate DETFLUX estimates. The source in panel (e) is QSO B0723-007 and the source in panel (f) is Virgo A.

BL Lacertae object, which are known to exhibit a high degree of radio variability (Nieppola et al., 2009), but is not present in the CRATES catalogue. However, it is important to note that CRATES pre-selected flat-spectrum sources for observation on the basis of spectral indices calculated between the PMN and NVSS catalogues in this region. Given the large separation in time between these surveys and the CRATES survey, we may expect CRATES to miss a small number of variable sources (Healey et al., 2007). The final extracted source was Virgo A. Whilst Virgo A is not a variable source, it does possess a large-scale halo that is not picked up the small beam of GB6, leading to the seemingly anomalously high value of the C-BASS flux density (75.9 ± 1.08 Jy). The C-BASS flux density for Virgo A is consistent with measurements in Baars et al. (1977). Whilst a detailed study of source variability between the C-BASS and GB6/PMN catalogues is beyond the scope of this work, flat-spectrum sources do account for a significant degree of variability between these catalogues. This is evidenced by the improved consistency between the flux-density scales when matching to the CRATES-cleaned C-BASS catalogue, and the lack of a large additional population of variable sources beyond those identified with the CRATES catalogue.

3.3.4 Differential Source Counts

Differential source counts have been intensively studied (e.g., de Zotti et al. (2005, 2010); AMI Consortium et al. (2011)), having important implications for models of cosmic evolution (e.g., Condon (1984); Massardi et al. (2010)) and provided the first clear evidence against the Steady State Theory (Pooley and Ryle, 1968). However, for the purposes of the C-BASS analysis I calculate the differential source counts to provide an additional cross-check on the statistical properties of the bright-source population.

Given that I am concerned with the statistical properties of the catalogue at flux-densities $\gtrsim 1$ Jy, where the C-BASS catalogue is largely complete, I do not consider correction factors to the differential source counts necessary at lower flux densities to account for catalogue incompleteness and varying flux-density sensitivities across the sky. Sources are put into flux-density bins of width 0.2 dex. For each bin, source counts are determined against a central flux density, S_c , taken to be the geometric mean of flux densities in that bin. The differential source counts for a given flux-density bin are then given by

$$\left. \frac{dN}{dS} \right|_{S=S_c} = n(S_c) = \frac{N_c}{A\Delta S}, \quad (3.24)$$

where N_c is the number of sources in the flux-density bin, A is the sky area to which the catalogue corresponds and ΔS is the bin width. Errors are estimated for the source counts by assuming Poisson statistics. I only consider bins with $N_c \geq 9$ so as to maintain a $\text{SNR} \geq 3$ for each bin.




Figure 3.14: Differential source counts calculated using the C-BASS northern sky total intensity catalogue, plotted alongside the differential source counts for the GB6 catalogue. Source counts were calculated applying the CG30 mask. I also plot as grey-scale the fitted model source counts corresponding to 100 posterior samples for the C-BASS and GB6 models. At low flux densities we note the expected lower levels in the C-BASS counts. This arises due to the low 45 arcmin resolution of the C-BASS map, leading to confusion issues for faint sources, along with fainter sources being more heavily obscured by diffuse emission and noise. Good agreement between GB6 and C-BASS differential source counts is found over intermediate flux densities ($\sim 1 \text{ Jy} - 10 \text{ Jy}$).

The differential source counts for C-BASS north in total intensity, applying the CG30 mask, are plotted in Fig. 3.14, alongside the differential source counts for the GB6 catalogue. At low flux densities the C-BASS source counts are significantly lower. This is to be expected from the increasing catalogue incompleteness at these lower flux densities. From the differential source counts, we can again see the catalogue completeness level from the point at which the source count turn-off ends and there is good agreement between the C-BASS and GB6 counts, i.e., over flux densities $\sim 1 \text{ Jy} - 10 \text{ Jy}$. It is worth noting here that the errors obtained from assuming Poisson statistics are likely to be somewhat optimistic. These do not account for false positives in the source catalogues, or errors arising from incorrect binning due

to uncertain flux-density estimates. These effects become significant at lower flux densities, where the proportion of false positives in the catalogue is greater, and the photometry methods become less reliable.

I fit the usual power-law model to the differential source counts, given by

$$\frac{dN}{dS} = n(S) = \Lambda \left(\frac{S}{S_0} \right)^\gamma, \quad (3.25)$$

where Λ is the count amplitude, $S_0 = 3 \text{ Jy}$ is the reference flux density and γ is the power-law index for the source counts. For C-BASS sources I fit to data points above the turn-off in the source counts ($S_\nu > 750 \text{ Jy}$), where the C-BASS catalogue is largely complete. For GB6, I fit to data points over the range $S_\nu > 0.07 \text{ Jy}$, so as to avoid the need to apply correction factors to faint GB6 source counts. I assume a Gaussian likelihood, and assign the following weakly informative priors to Λ and γ ,

$$\Lambda \sim \mathcal{N}(\mu = 5 \text{ Jy}^{-1} \text{sr}^{-1}, \sigma = 5 \text{ Jy}^{-1} \text{sr}^{-1}), \quad (3.26)$$

$$\gamma \sim \mathcal{N}(\mu = -2.5, \sigma = 2.0). \quad (3.27)$$

To account for additional errors in the source counts arising from flux-density uncertainties and incorrect binning of sources, I marginalize over an additional percentage error, f_{extra} , so that the total error on the source count in the i^{th} bin is given by

$$\sigma_i = \sqrt{\sigma_{\text{Poisson}}^2 + (f_{\text{extra}} n(S_i))^2}, \quad (3.28)$$

where σ_{Poisson} is the Poisson error on the source count. I assign an exponential prior to the additional percentage error, $f_{\text{extra}} \sim \text{Exp}(\lambda = 20)$, corresponding to a characteristic value of 5 per cent. I again explore the parameter space using the NUTS algorithm.

The results from these power-law fits are summarized in Table 3.4. When I do not apply any Galactic plane mask to the C-BASS catalogue, I fit a noticeably shallower spectral index than for the GB6 catalogue. This is driven by the excess in bright source counts in the C-BASS catalogue. This is caused by the spurious tagging of bright, highly-extended emission in the Galactic plane, and the over-estimation of source flux-densities in the presence of such emission, as discussed in Section 3.3.2. On applying the CG30 mask, the 95 per cent HPD intervals for C-BASS and GB6 overlap as expected. This can be seen in Fig. 3.14, where I plot 100 posterior samples from the C-BASS and GB6 models. When I apply the CG30 mask the posterior samples for the C-BASS and GB6 models overlap significantly.

Table 3.4: Fitted power-law parameters for the differential source counts. I quote the mean parameter values and the 95 per cent HPD intervals. Results are given for cases where I apply no mask and the CG30 mask.

Catalogue	$\langle\Lambda\rangle$	95% HPD [Λ]	$\langle\gamma\rangle$	95% HPD [γ]
C-BASS:				
No mask	5.10	{4.43, 5.73}	-1.52	{-1.60, -1.43}
CG30 mask	4.73	{3.75, 5.70}	-2.37	{-2.59, -2.16}
GB6:				
No mask	3.15	{2.75, 3.53}	-2.41	{-2.46, -2.36}
CG30 mask	4.55	{3.80, 5.28}	-2.45	{-2.52, -2.39}

3.3.5 C-BASS Catalogue Sample

In Table 3.5 I list the first ten sources from the C-BASS catalogue. Columns are arranged as,

1. **SourceID**: The C-BASS source identifier.
2. **RA [deg]**: The source right ascension in degrees (J2000 coordinates).
3. **DEC [deg]**: The source declination in degrees (J2000 coordinates).
4. **l [deg]**: The source Galactic longitude in degrees.
5. **b [deg]**: The source Galactic latitude in degrees.
6. **APERFLUX [Jy]**: The source flux-density estimate obtained using aperture photometry.
7. **APERFLUX Error [Jy]**: The uncertainty on the APERFLUX estimate.
8. **DETFLUX [Jy]**: The source flux-density estimate obtained from the peak value in the SMHW2 filtered map.
9. **DETFLUX Error [Jy]**: The uncertainty on the DETFLUX estimate.
10. **SNR**: The SNR at which the source was detected in the SMHW2 filtered map.
11. **Matched**: Column denoting which catalogues contained at least one match to the C-BASS source. G denotes at least one GB6 match, P denotes at least one PMN match, R denotes at least one RATAN-600 match, and C denotes a source

that was matched with a CRATES source during the GB6/PMN comparisons. A 0 means the source was not matched with any source in these ancillary catalogues. For example, GP means the source was matched with at least one GB6 source, and at least one PMN source.

12. **CG30**: This column denotes whether the source was within the CG30 mask. A 1 means the source was inside the CG30 mask (i.e., it was masked for the analysis) and a 0 means the source was outside the CG30 mask.




Figure 3.15: DETFLUX estimates plotted against APERFLUX estimates for the C-BASS catalogue. I also plot the line DETFLUX=APERFLUX. For low flux-densities ($S_\nu \lesssim 500$ mJy) the DETFLUX estimates plateau at the C-BASS flux-density limit. At these low flux-densities the DETFLUX results typically provide more reliable flux-density estimates than the APERFLUX results. This is because sources are detected at greater SNR in the filtered maps, from which the DETFLUX peak estimates are extracted. However, for brighter sources the DETFLUX estimation procedure can miss extended emission, resulting in the flux-density estimates being systematically biased low compared to the APERFLUX estimates. This is more noticeable for sources close to the Galactic plane.

In Fig. 3.15 I show the DETFLUX estimates plotted against the APERFLUX estimates for the C-BASS catalogue. The DETFLUX estimates typically provide more reliable flux-density estimates for fainter sources ($S_\nu \lesssim 500$ mJy) and sources in regions of brighter diffuse emission. In this situation, the APERFLUX estimates become heavily contaminated by the diffuse background. This can result in the uncertainties on the APERFLUX estimates becoming larger than the APERFLUX

estimates, or yield negative APERFLUX estimates due to excessive background subtraction. However, for brighter sources ($S_\nu \gtrsim 1$ Jy) the DETFLUX estimates are systematically biased low compared to the APERFLUX estimates. This is due to the DETFLUX estimates missing any extended emission from brighter sources. It is for this reason that I use the APERFLUX estimates in the analysis of the catalogue flux-density scales, with the DETFLUX results provided as auxiliary flux-density estimates. These issues have previously been noted in [Planck Collaboration et al. \(2014b, 2016f\)](#), when constructing the PCCS.

I do not provide results from any Gaussian or PSF fitting at source locations in the C-BASS map. Due to the low resolution of the C-BASS survey, and the bright diffuse emission over much of the sky, fitting for detailed PSF parameters is a major challenge for all but the brightest sources ($S_\nu \gtrsim 10$ Jy).

Table 3.5: The first ten sources from the C-BASS northern sky catalogue. In column 1 I give the C-BASS source identifier. In columns 2 and 3 I give the right ascension and declination in degrees respectively (in J2000 coordinates). In columns 4 and 5 I give the Galactic longitude and latitude in degrees respectively. In column 6 I give the APERFLUX estimate for the source flux-density in Jy. In column 7 I give the uncertainty on the APERFLUX estimate in Jy. In column 8 I give the DETFLUX estimate in Jy, determined directly from the peak value of the source in the SMHW2 filtered map. In column 9 I give the uncertainty on the DETFLUX estimate in Jy. In column 10 I give the detected SNR of the source in the filtered map. This is equivalent to the DETFLUX estimate divided by its corresponding uncertainty. In column 11 I denote whether the C-BASS source was matched with sources in the GB6, PMN, RATAN-600 or CRATES catalogues. G means a source was matched with at least one GB6 source, P means a source was matched with at least one PMN source, R means a source was matched with at least one RATAN-600 source, and C means a source was matched with at least one CRATES source during the GB6/PMN comparisons. A 0 means the C-BASS source was not matched with any source in these ancillary catalogues. In column 12 I denote whether the C-BASS source was inside the CG30 mask. A 1 means the source was inside the CG30 mask and a 0 means the source was outside the CG30 mask.

SourceID	RA [deg] (J2000)	DEC [deg] (J2000)	l [deg]	b [deg]	APERFLUX [Jy]	APERFLUX Error [Jy]	DETFLUX [Jy]	DETFLUX Error [Jy]	SNR	Matches	CG30
CBASS_J000056+553448	0.23	+55.58	115.79	+6.59	0.43	5.23	0.72	0.19	3.8	G	1
CBASS_J000110+405055	0.29	+40.85	112.78	+21.03	0.11	0.93	0.66	0.13	5.2	G	0
CBASS_J000137+643717	0.40	+64.62	117.62	+2.26	6.30	3.84	2.41	0.44	5.4	0	1
CBASS_J000249+671313	0.71	+67.22	118.24	+4.79	127.65	4.29	40.92	0.31	133.2	0	1
CBASS_J000316+655840	0.82	+65.98	118.05	+3.56	32.00	4.03	1.75	0.35	5.1	0	1
CBASS_J000433+125050	1.14	+12.85	105.56	+48.44	1.69	1.20	1.04	0.14	7.3	G	0
CBASS_J000548+683650	1.45	+68.61	118.77	+6.11	42.21	4.16	6.21	0.26	24.1	0	1
CBASS_J000609+381223	1.54	+38.21	113.23	+23.81	1.11	0.93	0.86	0.14	6.0	GC	0
CBASS_J000618+722000	1.58	+72.33	119.48	+9.76	9.26	2.46	2.56	0.17	14.8	G	1
CBASS_J000622-062113	1.59	-6.35	93.63	-66.63	1.77	1.06	2.19	0.13	16.8	PC	0

3.3.6 Source-Subtracted Map

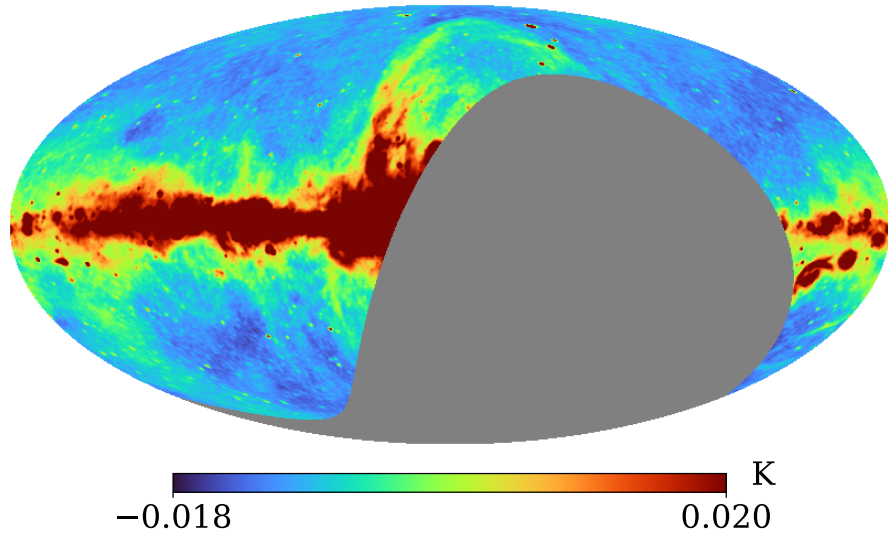
The C-BASS collaboration has used the northern sky point-source catalogue to produce a source-subtracted map of the northern sky. In Fig. 3.16 I show the C-BASS northern sky map, deconvolved with the C-BASS beam and smoothed by a 1° FWHM Gaussian, along with the corresponding source-subtracted map. The source-subtracted map was produced by taking sources outside the CG30 mask from the C-BASS catalogue with flux-densities greater than 1 Jy, and producing a map of the corresponding source emission. To avoid effects from the map pixelization, point sources were populated in the sky map in harmonic space, before transforming to pixel space. For sources fainter than 1 Jy, sources from the GB6, PMN and RATAN-600 catalogues were used. A 1° radius exclusion zone was applied around each C-BASS source, to avoid double counting sources in these external catalogues.

Comparing the two maps in Fig. 3.16, we can see that the bright point-source emission in the C-BASS map has been removed. In using this source-subtracted map for scientific analyses it is important to quantify the level of residual source contamination, and properly propagate uncertainties in the source flux-densities through to the source-subtracted map. This can be a significant challenge. To avoid such problems, a more conservative approach can be taken where the bright source population is masked out in any analyses of diffuse emission.


3.4 Conclusions

I have produced a total-intensity catalogue of point sources at 4.76 GHz, covering declinations $\delta \geq 10^\circ$. The catalogue was produced by filtering the whole sky using the SMHW2 filter, and contains 1784 sources. The C-BASS source catalogue has been characterized through comparisons with the pre-existing GB6 and PMN catalogues over their common declination ranges. Through these direct comparisons I estimate the C-BASS catalogue to have a 90 per cent completeness level of approximately 610 mJy, when applying the CG30 mask, with a corresponding reliability of 98 per cent.

Comparisons with the GB6 and PMN catalogues also enabled a number of checks to be performed on the C-BASS data. I find a modal absolute positional offset of approximately 3.2 arcmin when applying the CG30 mask. The low level of additional scatter in positional offsets, compared to the positional uncertainty found during the algorithm validation, acts to confirm the accuracy of the C-BASS pointing. I



(a) Non-source subtracted map



(b) Source-subtracted map

Figure 3.16: Panel (a): C-BASS northern sky map without source-subtraction. Panel (b): Source-subtracted C-BASS northern sky map. The source-subtracted map was produced by subtracting any sources from the C-BASS catalogue that were outside the CG30 mask, and had flux-densities greater than 1 Jy. For fainter sources, the external GB6, PMN and RATAN-600 catalogues were used, applying a 1° radius exclusion zone around any C-BASS sources to avoid double counting bright sources. We can see that the bright source population is largely removed through the process, although care must be taken in evaluating the level of source residuals and in propagating uncertainties to the source-subtracted map.

compare the flux densities of C-BASS sources with matched sources in the GB6 and PMN catalogues, using a measurement error model to fit for the gradient between the flux-density scales. Applying the CG30 mask, and cleaning flat-spectrum sources from C-BASS using the CRATES catalogue, I find a mean gradient of $\langle b \rangle = 1.04$ comparing to GB6, and $\langle b \rangle = 1.01$ comparing to PMN. In both cases the 95 per cent HPD intervals overlap with 1. Considering the ~ 5 per cent uncertainties in the GB6 and PMN flux-density scales, I conclude that the C-BASS and GB6/PMN flux-density scales are consistent.




Figure 3.17: An example point-source mask derived from the C-BASS total intensity northern sky catalogue, shown here in Galactic coordinates with masked pixels in black. In this example I mask all sources with flux densities greater than 610 mJy i.e., above the 90 per cent completeness level obtained when applying the CG30 mask. I mask all pixels within 45 arcmin of a source, and further mask all pixels within 2° of sources with flux-densities greater than 10 Jy. In addition to masking around sources, I have masked all pixels within the CG30 mask.

As an additional check on the statistical properties of the bright source population in the C-BASS catalogue, I calculate the differential source counts for C-BASS sources. Applying the CG30 mask, the C-BASS source counts show excellent agreement with the GB6 source counts over the flux-density range $\sim 1 - 10$ Jy. At lower

flux densities, there is the expected turn-off in the C-BASS source counts as the C-BASS catalogue becomes more incomplete.

For the purposes of C-BASS analysis, the primary products of this work are the checks on the data quality, and a mask of the bright point-source emission in the C-BASS map. In Fig. 3.17 I show an example point-source mask derived from the C-BASS catalogue, combined with the CG30 mask. Here I mask all sources with flux densities greater than 610 mJy. I mask all pixels within 45 arcmin of a source, and further mask all pixels within 2° of sources brighter than 10 Jy. Such masks will be essential for the scientific analysis of C-BASS intensity data.

Chapter 4

The Power Spectra of Polarized Galactic Synchrotron Emission

The detection of primordial B -mode polarization is one of the major outstanding goals of CMB cosmology (BICEP2/Keck Collaboration et al., 2015; Kamionkowski and Kovetz, 2016). At its core, this is a problem of accurate component separation. That is, the extraction of the primordial CMB signal from foreground contaminated observations of the radio and microwave sky (Betoule et al., 2009; Dunkley et al., 2009; Errard et al., 2016). Given current constraints on r , along with the targeted sensitivities of next-generation CMB experiments of $\sigma(r) \sim 10^{-3}$, future CMB experiments must be able to detect a CMB signal that is sub-dominant to foreground emission across all of the sky, at all frequencies (Kogut et al., 2011; Abazajian et al., 2016; Remazeilles et al., 2018; Sekimoto et al., 2018; BICEP2 Collaboration et al., 2018; Simons Observatory Collaboration et al., 2019; Shandera et al., 2019; Hanany et al., 2019). This presents two primary challenges. First we must be sure to have data of a sufficient sensitivity with enough frequency coverage to be able to model foreground spectral energy distributions (SEDs) with sufficient accuracy and precision. Secondly, our component separation algorithms must be able to extract the CMB signal from our noisy observations with high fidelity, and properly quantify the uncertainty in the extracted signal.

In this chapter I focus on the first challenge, presenting the first analysis of the power spectrum of polarized Galactic synchrotron emission in the C-BASS northern sky maps, along with cross-correlations with low-frequency WMAP and *Planck* data. A similar analysis has previously been conducted over the southern sky using S-PASS data in Krachmalnicoff et al. (2018), where it was found that the minimum level of

synchrotron contamination to CMB B -modes corresponds to an equivalent tensor-to-scalar ratio of $r \sim 10^{-3}$. Indeed, I obtain similar results here using C-BASS data over the northern sky. Due to the observing frequency of 2.3 GHz, S-PASS is more severely affected by de-polarization effects due to Faraday rotation. This limits the utility of cross-correlations between the S-PASS and WMAP/*Planck* maps. In the analysis of auto- and cross-spectra between S-PASS and WMAP/*Planck*, cross-spectra between S-PASS and WMAP/*Planck* were not included in any model fits for this reason. C-BASS observes the sky at an effective frequency of 4.76 GHz, where Faraday rotation is expected to be limited to $\sim 3^\circ$ away from the Galactic plane (compared to $\sim 14^\circ$ at 2.3 GHz) (Taylor et al., 2009; Jones et al., 2018), enabling a more direct joint analysis with WMAP and *Planck* data.

The outline of this chapter is as follows: In Section 4.1 I describe the power spectrum estimation and the corresponding validation simulations. In Section 4.2 I present the polarized C-BASS auto-spectra. In Section 4.3 I present a cross-spectrum based analysis with WMAP and *Planck* data, using it to place constraints on the polarized synchrotron SED. In Section 4.4 I use the C-BASS B -mode auto-spectrum, and constraints on the synchrotron spectral index, to estimate the level of synchrotron contamination to primordial B -modes. I summarize and discuss the results in Section 4.5.

4.1 Power Spectrum Estimation

For the analysis in this chapter I evaluate the auto- and cross-spectra of the C-BASS (Jones et al., 2018), WMAP K (23 GHz), WMAP Ka (33 GHz) (Bennett et al., 2013), and *Planck* 30 GHz maps (central frequency 28.4 GHz) (Planck Collaboration et al., 2016b). Note that for C-BASS I used the map that had been deconvolved with the instrumental beam and smoothed with a 1° FWHM Gaussian. I use the NAMASTER library (Alonso et al., 2019), which calculates pseudo- C_ℓ estimators for the power spectra. NAMASTER is able to account for the effects of applying a sky mask, and perform E and B -mode purification. B -mode purification is particularly important when the strength of the B -mode signal is much weaker than the E -mode signal. In this situation, E -to- B leakage (contamination from E -modes in the B -mode component) can lead to the variance of the B -mode estimator becoming dominated by the cosmic variance of the E -mode estimator. This is not a major concern here, given that we are not in the regime where $B \ll E$. However, this is an important issue when studying CMB B -modes, as is the case in Chapter 5. Throughout this chapter

I use the WMAP p06, p04 and p02 polarized emission masks¹. These mask 26.5, 33.1 and 44.6 per cent of the sky respectively. In Fig. 4.1 I show these sky masks, combined with the C-BASS northern sky mask covering declinations $\delta < -13^\circ$.

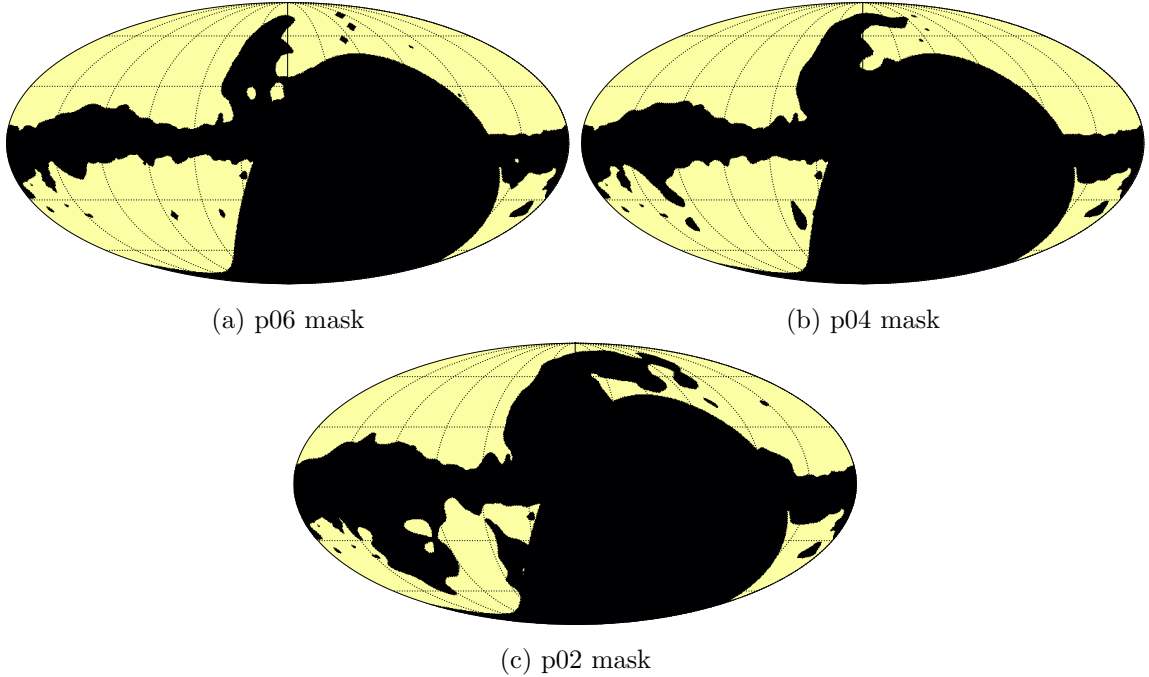


Figure 4.1: The WMAP p06, p04 and p02 polarized emission masks, combined with the C-BASS northern sky mask covering declinations $\delta < -13^\circ$. Without the addition of the C-BASS mask, the WMAP masks cover 26.5, 33.1 and 44.6 per cent of the sky respectively. Masked regions are shown in black.

4.1.1 Pseudo- C_ℓ Estimators

Estimating the power spectrum of a masked sky map is a non-trivial exercise, given the couplings between ℓ -modes caused by the mask. A variety of methods have been developed for the calculation of power spectra in this regime, including maximum-likelihood (Bond et al., 1998; Wandelt and Hansen, 2003) and quadratic estimators (Tegmark, 1997; Tegmark and de Oliveira-Costa, 2001; Vanneste et al., 2018), which are able to estimate the power spectra with almost no loss of information. Bayesian methods have also been employed in Eriksen et al. (2004, 2008); Taylor et al. (2008); Alsing et al. (2016). With these Bayesian methods the samples generated from the joint posterior of the map-based CMB amplitudes and the CMB power spectra, act to

¹The WMAP sky masks can be found at https://lambda.gsfc.nasa.gov/product/map/dr2/masks_info.cfm

constrain the behaviour of the CMB fields inside the mask. In doing so, the Bayesian methods do not have to treat the masked region as being totally removed. This can be particularly effective at low multipoles, where the cut-sky would otherwise yield poor power spectrum estimates. Moreover, careful use of priors and hierarchical model structures can yield improved power spectrum estimates when we are far from any asymptotic regime (as is the case at low multipoles). However, all of these methods become prohibitively expensive at high multipoles, scaling as $\mathcal{O}(\ell_{\max}^6)$ in the worst case (where ℓ_{\max} is the maximum multipole).

As an alternative, pseudo- C_ℓ estimators have been developed to preserve the $\mathcal{O}(\ell_{\max}^3)$ scaling of the spherical harmonic transform (Wandelt et al., 1998; Szapudi et al., 2001; Hivon et al., 2002; Chon et al., 2004; Tristram et al., 2005b; Alonso et al., 2019). This comes at the cost of a loss of optimality, which is particularly apparent at low multipoles. However, at large scales the methods mentioned above become tractable. In this section I summarize the core results relevant to pseudo- C_ℓ estimation and B -mode purification, using much of the same notation as Alonso et al. (2019).

Given the focus here on E and B -mode power spectra, I will restrict this discussion to spin-2 fields. For spin-2 fields $\mathbf{a}(\hat{\theta})$ and $\mathbf{b}(\hat{\theta})$, observed with sky masks $v(\hat{\theta})$ and $w(\hat{\theta})$, we can define $\mathbf{a}^v \equiv v(\hat{\theta})\mathbf{a}(\hat{\theta})$ and $\mathbf{b}^w \equiv w(\hat{\theta})\mathbf{b}(\hat{\theta})$. Note that the spin-2 field consists of a Q and U component, such that $\mathbf{a}(\hat{\theta}) = (Q_a, U_a)^\top$ and $\mathbf{b}(\hat{\theta}) = (Q_b, U_b)^\top$. A power spectrum estimate can then be constructed as

$$\hat{C}_\ell(\mathbf{a}^v, \mathbf{b}^w) = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} \mathbf{a}_{\ell m}^v (\mathbf{b}_{\ell m}^w)^\dagger, \quad (4.1)$$

where $\mathbf{a}_{\ell m}$ and $\mathbf{b}_{\ell m}$ are the spherical harmonic coefficients of the \mathbf{a} and \mathbf{b} fields. However, this does not account for the couplings of ℓ -modes induced by the mask. The necessary correction for this can be calculated analytically, giving the unbiased pseudo- C_ℓ estimator, $\langle \tilde{C}_\ell(\mathbf{a}_{\ell m}^v, \mathbf{b}_{\ell m}^w) \rangle$, as

$$\text{vec} \left[\langle \tilde{C}_\ell(\mathbf{a}_{\ell m}^v, \mathbf{b}_{\ell m}^w) \rangle \right] = \sum_{\ell'} \mathbf{M}_{\ell\ell'}^{22} \text{vec} \left[\hat{C}_{\ell'}(\mathbf{a}^v, \mathbf{b}^w) \right], \quad (4.2)$$

where the vectorization operator is defined as

$$\text{vec} \left[\begin{pmatrix} C_\ell^{EE} & C_\ell^{EB} \\ C_\ell^{BE} & C_\ell^{BB} \end{pmatrix} \right] \equiv (C_\ell^{EE}, C_\ell^{EB}, C_\ell^{BE}, C_\ell^{BB})^\top. \quad (4.3)$$

The mode-coupling matrix for spin-2 fields is given by

$$\mathbf{M}_{\ell\ell'}^{22} = \begin{pmatrix} M_{\ell\ell'}^+ & 0 & 0 & M_{\ell\ell'}^- \\ 0 & M_{\ell\ell'}^+ & -M_{\ell\ell'}^- & 0 \\ 0 & -M_{\ell\ell'}^- & M_{\ell\ell'}^+ & 0 \\ M_{\ell\ell'}^- & 0 & 0 & M_{\ell\ell'}^+ \end{pmatrix}, \quad (4.4)$$

where

$$M_{\ell\ell'}^\pm = \frac{2\ell' + 1}{4\pi} \sum_{\ell''} (2\ell'' + 1) \hat{\mathbf{C}}_{\ell''}(v, w) \begin{pmatrix} \ell & \ell' & \ell'' \\ 2 & -2 & 0 \end{pmatrix}^2 \frac{1 \pm (-1)^{\ell+\ell'+\ell''}}{2}. \quad (4.5)$$

When calculating the E and B -mode power spectra on a cut-sky we encounter an additional complication due to E -to- B leakage. To correct for this NAMASTER can perform additional B -mode purification. For some spin-2 field \mathbf{P} , observed on the sky with some mask or weight function $w(\hat{\theta})$, the standard pseudo B -mode is given by

$$\tilde{\mathbf{B}}_{\ell m} = \int d\hat{\theta} w(\hat{\theta}) (\mathbf{D}_s^B Y_{\ell m})^\dagger \mathbf{P}(\hat{\theta}), \quad (4.6)$$

where \mathbf{D}_s^B is a differential operator whose action on a scalar field generates a B -mode, and s denotes the spin (in this case, $s = 2$). An explicit definition can be found in Appendix A of [Alonso et al. \(2019\)](#), with further detail on pseudo- C_ℓ estimation in [Wandelt et al. \(1998\)](#); [Szapudi et al. \(2001\)](#); [Hivon et al. \(2002\)](#); [Chon et al. \(2004\)](#); [Tristram et al. \(2005b\)](#). Whilst the term $\mathbf{D}_s^B Y_{\ell m}$ is a B -mode, $w(\hat{\theta})(\mathbf{D}_s^B Y_{\ell m})^\dagger$ is not a B -mode. In the presence of a sky mask, the action of Equation 4.6 does not project a spin-2 field to a pure B -mode. The pure B component of the field can instead be defined as

$$\mathbf{B}_{\ell m}^p = \int d\hat{\theta} (\mathbf{D}_s^B (w(\hat{\theta}) Y_{\ell m}))^\dagger \mathbf{P}(\hat{\theta}). \quad (4.7)$$

B -mode purification then proceeds by evaluating the pure B component of the input field as defined in Equation 4.7, and then calculating the pseudo- C_ℓ estimator using the purified field in Equation 4.1. It is important to note that calculating the purified field requires the first and second derivatives of the mask, meaning care must be taken to ensure the mask is differentiable up to second order. This is typically achieved by applying a differentiable apodization at the edge of the mask.

4.1.2 Bandpower Estimation

Given the C-BASS, WMAP K, WMAP Ka and *Planck* 30 GHz maps, there are a total of ten separate auto- and cross-spectra to evaluate. For auto-spectra I calculate

the cross-spectra between data-splits from the survey of interest, in order to avoid complications from noise bias. The data-splits used for each survey map are as follows:

1. C-BASS: The C-BASS auto-spectra are evaluated using even/odd data-splits. These were produced by listing the TOD fits files alphanumerically by start date (using the format DD-*MMM*-YYYY:hh:mm:ss) and separating them into odd and even splits, according to whether the files were listed on odd or even lines. Systematic differences in these data-splits would have to manifest on approximately 100 minute timescales, and as a function of the alphanumerical ordering. No such effects have been identified for the C-BASS data (Jew, 2017).
2. WMAP K: The WMAP K data-split was produced by co-adding the single year 1 to 4 maps for one part of the split, and the year 5 to 9 maps to produce the other part.
3. WMAP Ka: The WMAP Ka data-split was produced in the same manner as for the WMAP K map.
4. Planck 30 GHz: The Planck 30 GHz data-split was produced by co-adding the year 1 and 3 maps for one part of the split, and the year 2 and 4 maps for the other part.

Cross-spectra between maps at different frequencies were evaluated using the full-mission maps at each frequency. Bandpowers were estimated using a bin width of $\Delta\ell = 20$ (excluding the first bin which contains multipoles $2 \leq \ell < 20$), with the central bandpower estimates for each auto- and cross-spectrum being evaluated using the maps outlined above. It is important to note that the Galactic synchrotron emission is not Gaussian. However, the study of its two-point correlations is motivated in part by the fact that CMB foregrounds are often treated using a Gaussian assumption at the power spectrum level, and is complementary to map-based analyses (Planck Collaboration et al., 2016a).

Bandpower covariances were estimated using a set of 100 signal plus noise simulations. To account for synchrotron cosmic variance I generated Gaussian realizations of fiducial synchrotron power spectra, given by

$$C_\ell^{XX} = A_s^{XX} \left(\frac{\ell}{80} \right)^{\alpha_s^{XX}}, \quad XX = \{EE, BB\}, \quad (4.8)$$

where A_s^{XX} is the reference synchrotron amplitude at $\ell = 80$ and α_s^{XX} is the synchrotron spectral tilt. The fiducial synchrotron parameter values at 4.76 GHz are

given in Table 4.1. Reference amplitudes were chosen to correspond with the approximate amplitude scales at 4.76 GHz, and α_s^{XX} values were selected based on the typical values found in Krachmalnicoff et al. (2018). After generating a synchrotron realization at 4.76 GHz, the simulated emission was scaled to the corresponding frequencies for the maps under consideration, using a spatially uniform spectral index of $\beta_s = -3$. A complete analysis would account for parameter dependencies in the covariance matrix. However, in Kodwani et al. (2019) it was found that ignoring this parameter dependence has negligible impact on subsequent inferences with two-point correlation functions (provided the fiducial model is close to the recovered model, as is the case for the model considered here). In addition to the synchrotron simulations, I also add CMB realizations from power spectra generated using CAMB (Lewis et al., 2000), with cosmological parameters set to the fiducial Λ CDM values in Planck Collaboration et al. (2016e).

Table 4.1: Fiducial synchrotron model parameters at 4.76 GHz, used for each sky mask. For each Monte Carlo simulation, a Gaussian realization of the power spectra was generated, which was then extrapolated to the map frequencies under consideration using a spatially uniform spectral index of $\beta_s = -3$.

Mask	A_s^{EE}	A_s^{BB}	α_s^{EE}	α_s^{BB}
p06	$8 \times 10^{-11} \text{ K}^2$	$2 \times 10^{-11} \text{ K}^2$	-3.15	-3.15
p04	$8 \times 10^{-11} \text{ K}^2$	$2 \times 10^{-11} \text{ K}^2$	-3.15	-3.15
p02	$6 \times 10^{-11} \text{ K}^2$	$2 \times 10^{-11} \text{ K}^2$	-3.15	-3.15

Noise realisations were generated for each Monte Carlo simulation by obtaining white noise realisations of the variance maps at the relevant frequencies (and for the relevant data-split in the case of auto-spectrum calculations). Whilst the maps do display large-scale features in the noise power spectrum, particularly for C-BASS where atmospheric and ground-based signals generate large-scale correlated noise, this is primarily a concern at lower multipoles. Indeed, in Jew (2017) jackknife tests found that the C-BASS noise power spectrum is white noise dominated over multipoles $\ell \gtrsim 40$. In Section 4.1.3, the validation simulations show that reliable results are only obtained for the power spectra in this analysis over $\ell \gtrsim 40$. Given this, the white noise simulations are sufficient for estimating the bandpower covariances at higher multipoles. Bandpower covariances were obtained by evaluating the power spectra over these 100 signal plus noise simulations, and calculating their correspond-

ing covariances. Error bars shown in power spectrum plots are given by the standard deviation of the bandpowers evaluated over the Monte Carlo simulations.

4.1.3 Validation Simulations

In order to validate the performance of the power spectrum estimator, and assess the multipoles over which reliable inferences can be conducted, I calculated the auto- and cross-spectra of a set of Gaussian simulations. These were generated to allow for a direct comparison between the recovered power spectra and the input power spectra. The E and B power spectra at 4.76 GHz used for the simulations are given by the fiducial p06 synchrotron model, specified in Table 4.1. Moreover, I included zero synchrotron EB correlation in the input power spectra.


Given these power spectra, Gaussian realisations were then generated at 4.76 GHz and extrapolated to the relevant WMAP and *Planck* frequencies using a global spectral index of $\beta_s = -3$. To each frequency map I added a white noise realization generated using the corresponding variance maps for each frequency channel and data-split as appropriate. Each map was smoothed using the instrumental beam for each frequency channel. Note that for the Monte Carlo simulations here, synchrotron realizations were generated as described above, but in this case the fiducial model corresponds exactly with the underlying input model.

In Fig. 4.2 I show the recovered E and B auto-spectra at 4.76 GHz, alongside the corresponding input power spectra. For $\ell \lesssim 40$ the sky mask results in the pseudo- C_ℓ estimator becoming sub-optimal, resulting in higher variance bandpower estimates. This is particularly noticeable for the B -mode power spectra, where poor point estimates are obtained at low multipoles. It is also worth noting that the real C-BASS data contains significant contributions from correlated noise at these angular scales. For $\ell \gtrsim 120$, the noise and 1° beam correction results in higher variance bandpower estimates, and poor individual point estimates for the bandpowers. To check the ability of the power spectrum estimation to recover the input model, I fit Equation 4.8 to the E and B auto-spectra, over the range $40 \leq \ell < 120$. I assign weakly informative priors as follows,


$$A_s^{XX} \sim \text{Exp}(\lambda = 10^{10} \text{ K}^{-2}), \quad (4.9)$$

$$\alpha_s^{XX} \sim \mathcal{N}(\mu = -3, \sigma = 1). \quad (4.10)$$


The scale factor for the amplitude prior was chosen to encapsulate the relevant synchrotron amplitude scales at $\ell = 80$, whilst helping to regularize inferences by down-



(a) p06



(b) p04



(c) p02

Figure 4.2: Recovered E and B auto-spectra at 4.76 GHz for the validation simulations, plotted alongside the corresponding input models. Below $\ell \lesssim 40$ the sky mask results in poor point estimates for the bandpowers, with the pseudo- C_ℓ estimator being sub-optimal at these scales. For $\ell \gtrsim 160$, the beam correction and noise results in higher variance bandpower estimates. Including these higher multipoles in model fits results in poor parameter recovery.

weighting extremely large amplitude values. The α_s^{XX} prior was chosen to capture the relevant scales at which the synchrotron spectral index has been measured.

Mask	A_s^{EE}	A_s^{BB}	α_s^{EE}	α_s^{BB}
p06	$(8.2 \pm 0.2) \times 10^{-11} \text{ K}^2$	$(2.1 \pm 0.1) \times 10^{-11} \text{ K}^2$	-3.10 ± 0.11	-3.17 ± 0.13
p04	$(8.1 \pm 0.2) \times 10^{-11} \text{ K}^2$	$(2.1 \pm 0.1) \times 10^{-11} \text{ K}^2$	-3.12 ± 0.10	-3.23 ± 0.12
p02	$(8.1 \pm 0.3) \times 10^{-11} \text{ K}^2$	$(2.1 \pm 0.1) \times 10^{-11} \text{ K}^2$	-3.05 ± 0.16	-3.24 ± 0.16

Table 4.2: Recovered auto-spectrum parameters at 4.76 GHz, for the validation simulations. The recovered model parameters are consistent with the input model i.e., $A_s^{EE} = 8 \times 10^{-11} \text{ K}^2$, $A_s^{BB} = 2 \times 10^{-11} \text{ K}^2$ and $\alpha_s^{XX} = -3.15$. Including multipoles $\ell \geq 120$ was found to result in poor model fits, showing tensions with the input model.

I assume that bandpowers are distributed according to a multivariate Gaussian likelihood,

$$\ln \mathcal{L}_{XX} = \text{const.} - \frac{1}{2} \sum_{\ell_b \ell'_b} \left[\tilde{C}_{\ell_b}^{XX} - \mu_{\ell_b}^{XX} \right] (\Sigma^{-1})_{\ell_b \ell'_b} \left[\tilde{C}_{\ell'_b}^{XX} - \mu_{\ell'_b}^{XX} \right] \quad (4.11)$$

where $\mu_{\ell_b}^{XX}$ denotes the model value in the bin ℓ_b , and Σ is the bandpower covariance matrix estimated from the Monte Carlo simulations. Whilst the assumption of Gaussianity for the bandpower distribution is not strictly correct for the scales of interest ($\ell > 40$), at these multipoles the averaging over larger numbers of m modes, and the bin-averaging mean that the multivariate Gaussian approximation provides a reliable likelihood through the central limit theorem. I sample the posterior using the NUTS algorithm, implemented in PYMC3. In Table 4.2 I state the results for the auto-spectrum fits, which are consistent with the input model.

In Fig. 4.3 and 4.4 I show the recovered E and B auto- and cross-spectra as a function of frequency respectively, for bandpowers between $40 \leq \ell < 100$. For multipoles $\ell \geq 100$, the bandpower estimates at WMAP and *Planck* frequencies became severely degraded by noise, meaning I failed to recover the input power-law scaling. The recovered point estimates shown below are in broad agreement with the input power-law scaling. Note that the simulations were generated by assuming a spatially uniform spectral index, meaning there are no decorrelation effects due to varying spectral indices. Further, the simulations considered here are purely Gaussian, which is not the case for the real Galactic emission. The WMAP and *Planck* auto- and cross-spectra become progressively more noisy as the severity of the sky mask is increased. This is particularly apparent in the B -mode estimates. It can also be seen

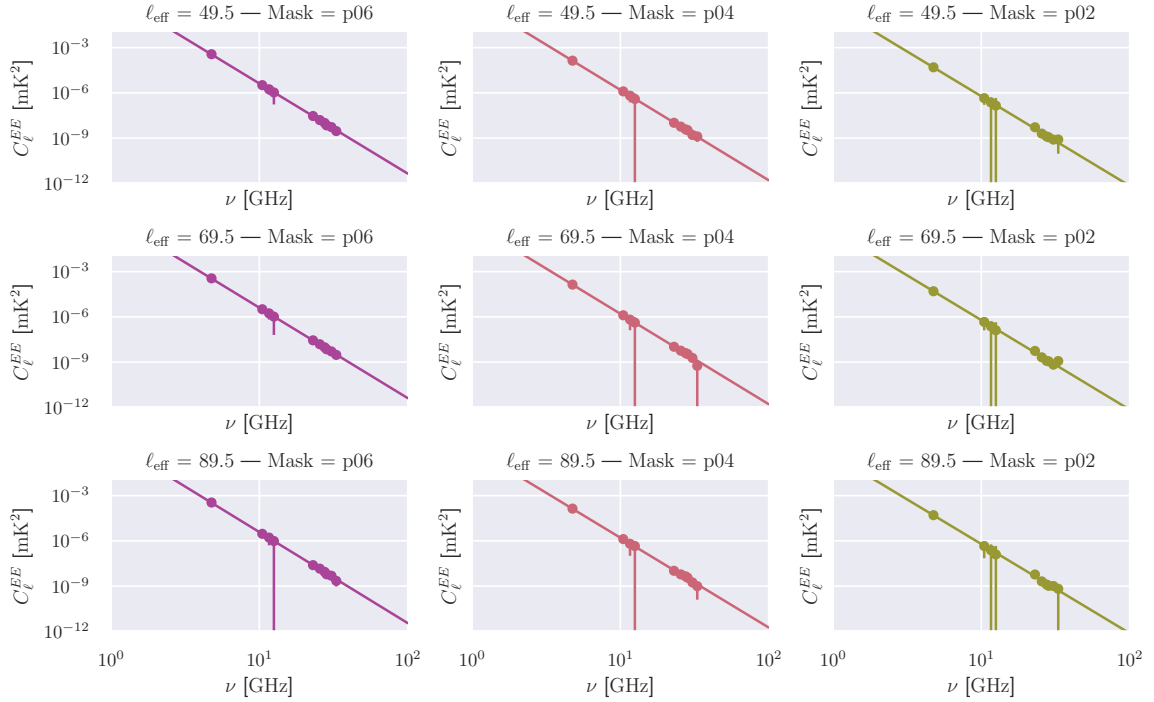


Figure 4.3: Recovered E -mode auto- and cross-spectra for the validation simulations, plotted as a function of frequency. Moving horizontally, I show the recovered bandpowers for each mask. Moving vertically, I show the recovered bandpowers for each multipole bin, labelled by the effective multipole. Points are plotted at the geometric mean of the two map frequencies ($\sqrt{\nu_1 \nu_2}$). Alongside the bandpower estimates, I show the input power-law spectrum. The point estimates for the auto- and cross-spectra broadly follow the spatially uniform $\beta_s = -3$ frequency scaling used in the validation simulations. However, for $\ell \gtrsim 100$, the higher noise levels in the WMAP and *Planck* maps result in poor auto- and cross-spectrum estimates. This in turn leads to unreliable parameter estimates during any spectral model fitting.

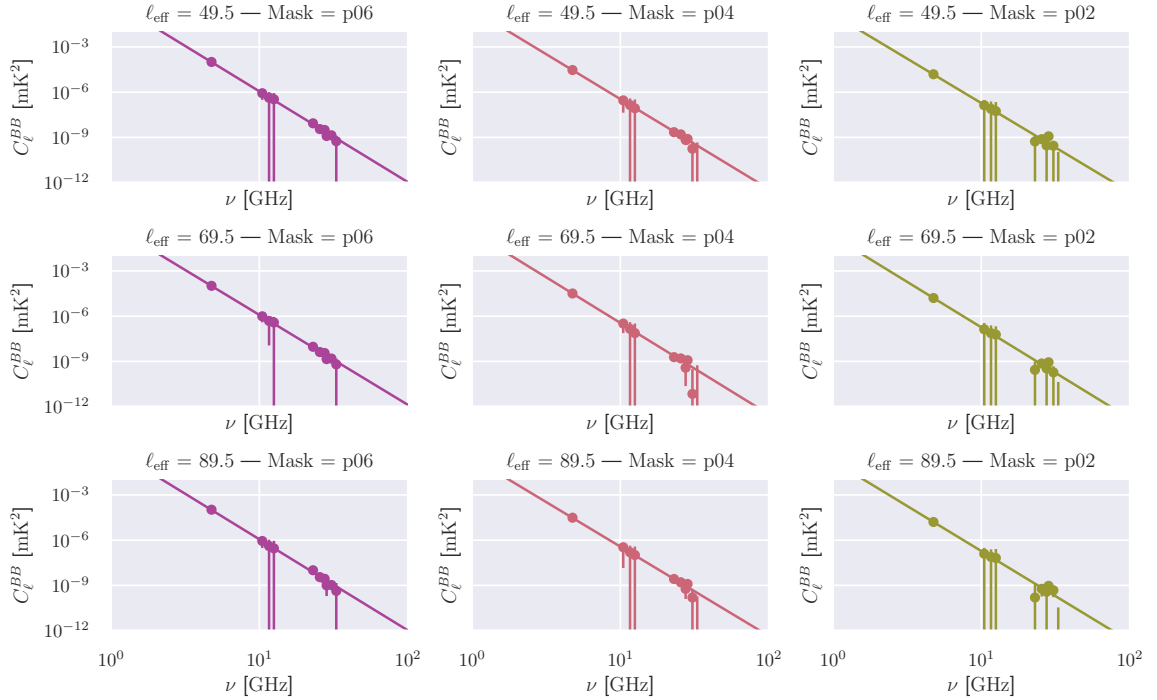


Figure 4.4: Recovered B -mode auto- and cross-spectra for the validation simulations, plotted as a function of frequency. Moving horizontally, I show the recovered bandpowers for each mask. Moving vertically, I show the recovered bandpowers for each multipole bin, labelled by the effective multipole. Points are plotted at the geometric mean of the two map frequencies ($\sqrt{\nu_1 \nu_2}$). Alongside the bandpower estimates, I show the input power-law spectrum. As for the E -mode power spectra, the recovered point estimates broadly follow the simulated frequency scaling, with bandpower estimates for WMAP and *Planck* data-points becoming heavily degraded beyond $\ell \gtrsim 100$. Poor bandpower estimates are also obtained for the WMAP and *Planck* data-points when applying the more aggressive p02 mask.

that cross-spectra between C-BASS and WMAP/*Planck* data points are typically poorly constrained, particularly for the B -mode power spectra.

In order to check the frequency scaling of the recovered bandpower estimates I fit the following power-law model,

$$\tilde{C}_{\ell_b}^{XX}(\nu_1 \times \nu_2) = A_{\nu_0}^{XX} \left(\frac{\nu_1 \nu_2}{\nu_0^2} \right)^{\beta_s^{XX}}, \quad (4.12)$$

where $\tilde{C}_{\ell_b}^{XX}(\nu_1 \times \nu_2)$ is the bandpower estimate for the cross-spectrum between maps at frequencies ν_1 and ν_2 , and $A_{\nu_0}^{XX}$ is the reference amplitude for a given bin at the frequency $\nu_0 = 4.76$ GHz. I assign weakly informative priors, similarly to the ℓ -space

model,

$$A_{\nu_0}^{XX} \sim \text{Exp}(\lambda = 10^{10} \text{ K}^{-2}), \quad (4.13)$$

$$\beta_s^{XX} \sim \mathcal{N}(\mu = -3, \sigma = 1). \quad (4.14)$$

I assume bandpowers follow a Gaussian likelihood, so that

$$\tilde{C}_{\ell_b}^{XX}(\nu_1 \times \nu_2) \sim \mathcal{N}(\mu = \mu_{\ell_b}^{XX}(\nu_1 \times \nu_2), \sigma = \sigma_{\ell_b}^{XX}(\nu_1 \times \nu_2)), \quad (4.15)$$

where $\mu_{\ell_b}^{XX}(\nu_1 \times \nu_2)$ denotes the model value for the bandpower estimate of the $\nu_1 \times \nu_2$ cross-spectrum, and $\sigma_{\ell_b}^{XX}(\nu_1 \times \nu_2)$ denotes the corresponding covariance between the cross-spectra. I sample the posterior using the NUTS algorithm, implemented in PYMC3.

In addition to the Gaussian prior on β_s^{XX} , I multiply this by the corresponding Jeffreys independence-rule prior (Jeffreys, 1946, 1961). In general, for some likelihood \mathcal{L} and parameter θ , the Jeffreys independence-rule prior is given by,

$$\pi_J(\theta) \propto \sqrt{\mathbb{E}_\theta \left[-\frac{\partial^2 \ln \mathcal{L}}{\partial \theta^2} \right]}, \quad (4.16)$$

where $\mathbb{E}_\theta[f(\theta)]$ denotes the expectation of $f(\theta)$. For the synchrotron spectral index this gives

$$\pi_J(\beta_s^{XX}) \propto \sqrt{\sum_i \left[\left(\frac{\nu_i}{\nu_0} \right)^{\beta_s^{XX}} \frac{1}{\sigma_i} \ln \left(\frac{\nu_i}{\nu_0} \right) \right]^2}. \quad (4.17)$$

In the low signal regime, where the power-spectrum constraints are consistent with there being no synchrotron emission at high frequencies, the spectral index estimates can be driven to unphysical, large-magnitude negative values using the Gaussian prior alone. It is important to note that the general application of the Jeffreys prior can be problematic e.g., with improper prior distributions which can potentially result in improper posteriors. However, in this case the Jeffreys prior alleviates the problem regarding synchrotron spectral index estimates in the low signal regime, by counteracting the volume and hence probability mass associated with these extreme parameter values. In Eriksen et al. (2004, 2008); Alonso et al. (2017) it was found that applying this prior distribution had the effect of removing biases on recovered synchrotron spectral index estimates.

The recovered parameter estimates are shown in Table 4.3, which are consistent with the input validation model.

Multipole Bin	Mask	$A_{\nu_0}^{EE}$	$A_{\nu_0}^{BB}$	β_s^{EE}	β_s^{BB}
$40 \leq \ell < 60$	p06	$(35.6 \pm 2.9) \times 10^{-11} \text{ K}^2$	$(9.9 \pm 0.8) \times 10^{-11} \text{ K}^2$	-2.99 ± 0.03	-3.01 ± 0.04
$60 \leq \ell < 80$	p06	$(13.6 \pm 0.8) \times 10^{-11} \text{ K}^2$	$(2.9 \pm 0.2) \times 10^{-11} \text{ K}^2$	-3.00 ± 0.02	-2.98 ± 0.06
$80 \leq \ell < 100$	p06	$(4.9 \pm 0.3) \times 10^{-11} \text{ K}^2$	$(1.5 \pm 0.1) \times 10^{-11} \text{ K}^2$	-2.97 ± 0.04	-2.95 ± 0.09
$40 \leq \ell < 60$	p04	$(34.9 \pm 2.7) \times 10^{-11} \text{ K}^2$	$(10.1 \pm 0.8) \times 10^{-11} \text{ K}^2$	-3.00 ± 0.03	-2.98 ± 0.03
$60 \leq \ell < 80$	p04	$(13.8 \pm 0.8) \times 10^{-11} \text{ K}^2$	$(3.2 \pm 0.2) \times 10^{-11} \text{ K}^2$	-3.00 ± 0.02	-3.03 ± 0.06
$80 \leq \ell < 100$	p04	$(5.0 \pm 0.3) \times 10^{-11} \text{ K}^2$	$(1.6 \pm 0.1) \times 10^{-11} \text{ K}^2$	-2.97 ± 0.04	-3.02 ± 0.10
$40 \leq \ell < 60$	p02	$(33.8 \pm 2.7) \times 10^{-11} \text{ K}^2$	$(10.1 \pm 1.3) \times 10^{-11} \text{ K}^2$	-3.02 ± 0.03	-3.02 ± 0.05
$60 \leq \ell < 80$	p02	$(13.5 \pm 1.0) \times 10^{-11} \text{ K}^2$	$(3.0 \pm 0.3) \times 10^{-11} \text{ K}^2$	-2.98 ± 0.03	-2.96 ± 0.06
$80 \leq \ell < 100$	p02	$(5.0 \pm 0.4) \times 10^{-11} \text{ K}^2$	$(1.6 \pm 0.1) \times 10^{-11} \text{ K}^2$	-2.96 ± 0.05	-2.99 ± 0.12

Table 4.3: Recovered spectral model parameters for the validation simulations. Recovered model parameters are consistent with the input model i.e., $A_s^{EE}(\ell = 80) = 8 \times 10^{-11} \text{ K}^2$ and $A_s^{BB}(\ell = 80) = 2 \times 10^{-11} \text{ K}^2$ and $\beta_s^{XX} = -3$. Note that the reference amplitudes in each multipole bin, $A_{\nu_0}^{XX}$ correspond to the $A_s^{XX}(\ell = 80)$ amplitudes, scaled accordingly using Equation 4.8.

In Fig. 4.5 I show the recovered EB correlation coefficients for bandpowers between $40 \leq \ell < 120$. The EB correlation coefficient is defined as

$$\rho_\ell^{EB} = \frac{C_\ell^{EB}}{\sqrt{C_\ell^{EE} C_\ell^{BB}}}. \quad (4.18)$$

For all three sky masks, the recovered EB correlations are consistent with zero as expected.

The results presented above confirm the ability of the power spectrum estimation procedure to recover the underlying power spectra. I obtain reliable E and B power spectra over the multipole range $40 \leq \ell < 120$, which is the multipole range studied when analysing the real C-BASS, WMAP and *Planck* maps in the remainder of this chapter.

4.2 The C-BASS Polarization Power Spectra

Given the C-BASS Q and U maps shown in Fig. 2.9, it is possible to produce the C-BASS E and B maps, shown in Fig. 4.6. To avoid ringing caused by the partial sky coverage I apply a $4^\circ \cos^2$ apodization at the edge of the observed sky area. It can be seen that the E -mode component traces Galactic loops, whereas the B -mode component is comparatively weak in these regions. Similar behaviour has previously been noted in the WMAP and *Planck* maps in Liu et al. (2018). It is also clear from the E and B





Figure 4.5: Recovered EB correlation coefficients for the validation simulations, shown for each sky mask. Points for each mask are slightly offset for clarity, with the central point in each triple being shown at the effective multipole for that bin. The recovered correlation coefficients are consistent with the input of $\rho_\ell^{EB} = 0$.

maps that the Galactic synchrotron emission exhibits highly non-Gaussian features. Given this, it might be expected that the properties of the synchrotron E and B -mode power spectra will differ, and that there may be a non-zero synchrotron EB correlation. In the remainder of this section I present constraints on the synchrotron angular power spectra using C-BASS northern-sky data.


4.2.1 E and B -mode power spectra

In Fig. 4.7 I show the C-BASS E and B auto-spectra, obtained by evaluating the cross-spectrum between odd/even data-splits as described in Section 4.1.2. Over the multipole range considered here, I fit a power-law model to the power spectra. Note that I do not consider point source contamination to the power spectrum, which is only expected to be significant at smaller scales ($\ell \gtrsim 200$). To the auto-spectra, I fit the same ℓ -space power-law model described in Section 4.1.3, with the corresponding parameter estimates given in Table 4.4.

Note that, whilst there are physical arguments justifying the approximation of




(a) C-BASS E map




(b) C-BASS B map


Figure 4.6: The C-BASS northern-sky E and B maps. The E -mode component of the synchrotron emission can be seen to trace Galactic loops, most prominently the North Polar Spur. The B -mode component, by contrast, is relatively weak in these regions. It is evident from these maps that the Galactic synchrotron emission exhibits highly non-Gaussian structures, and that there are significant differences in the distribution of E and B -mode intensities.




(a) C-BASS EE : p06




(b) C-BASS BB : p06




(c) C-BASS EE : p04



(d) C-BASS BB : p04



(e) C-BASS EE : p02



(f) C-BASS BB : p02

Figure 4.7: Recovered EE and BB auto-spectra for the C-BASS northern sky map, shown for each sky mask. Dashed lines indicate the 1σ boundaries of the posterior predictive distributions. This graphical posterior predictive check does not identify any systematic deviations from the predictive distribution in the observed data, indicating the power-law model provides an adequate description of the C-BASS auto-spectra over the multipoles considered here.

the synchrotron frequency spectrum as a power-law (see e.g. Longair (2011)), it is not necessarily obvious that the synchrotron power spectrum should follow a power-law. In order to assess the adequacy of the power-law model I therefore draw a set of samples from the posterior predictive distribution. In general, the posterior predictive distribution of some new data \mathbf{x}_{new} , given the observed data \mathbf{x} , is given by

$$p(\mathbf{x}_{\text{new}}|\mathbf{x}) = \int d\theta p(\theta|\mathbf{x})p(\mathbf{x}_{\text{new}}|\theta), \quad (4.19)$$

where θ are the model parameters. The posterior predictive distribution averages over all of the generative processes consistent with the model and therefore quantifies the set of possible observations consistent with the observed data and our prior knowledge (Gelman et al., 1996; Betancourt, 2020). The posterior predictive distribution can be used to perform graphical assessments of the model adequacy, by plotting the new, simulated data alongside the observed data. For high-dimensional datasets, this typically involves comparing the distribution of low-dimensional summary statistics (e.g., the mean of the data) under the posterior predictive distribution with the observed value of the summary statistic. For the power-law fits performed here, I plot the 1σ bounds of the posterior predictive distributions alongside the observed power spectra in Fig. 4.7. From this it can be seen that the observed data show no systematic deviations from the posterior predictive distribution, indicating that the power-law model provides an adequate description of the data.

Mask	A_s^{EE}	A_s^{BB}	α_s^{EE}	α_s^{BB}
p06	$(8.4 \pm 0.2) \times 10^{-11} \text{ K}^2$	$(2.3 \pm 0.1) \times 10^{-11} \text{ K}^2$	-3.31 ± 0.09	-3.50 ± 0.13
p04	$(7.7 \pm 0.2) \times 10^{-11} \text{ K}^2$	$(1.9 \pm 0.1) \times 10^{-11} \text{ K}^2$	-3.34 ± 0.10	-3.52 ± 0.14
p02	$(6.1 \pm 0.2) \times 10^{-11} \text{ K}^2$	$(1.5 \pm 0.1) \times 10^{-11} \text{ K}^2$	-3.55 ± 0.13	-3.43 ± 0.25

Table 4.4: Recovered E and B -mode auto-spectrum parameters for C-BASS. The E and B -mode spectral tilts are broadly consistent for each sky mask. The measured synchrotron E -to- B ratios range from approximately 3.6 to 4.1.

The weighted average of the E -mode spectral tilts across the sky masks is $\alpha_s^{EE} = -3.37 \pm 0.06$, and for the B -mode spectral tilts is $\alpha_s^{BB} = -3.50 \pm 0.09$. The weighted average of the E -to- B ratios across the sky masks is $A_s^{EE}/A_s^{BB} = 3.9 \pm 0.01$. These results differ slightly from those found using S-PASS data at 2.3 GHz. In Krachmalnicoff et al. (2018) the synchrotron auto-spectra were found to have a weighted mean spectral tilt of approximately -3.15 at Galactic latitudes $|b| > 30^\circ$, with slightly steeper spectral tilts for the E -mode power spectra. For the same Galactic latitudes, the

S-PASS derived results found E -to- B ratios of ~ 2 . Direct comparisons between the C-BASS derived synchrotron power spectra, and the S-PASS derived power spectra will be complicated by Faraday rotation effects between their observing frequencies. Further, the analysis in [Krachmalnicoff et al. \(2018\)](#) used symmetric Galactic latitude cuts to define sky masks, whereas the analysis in this chapter uses more realistic Galactic emission masks. A full joint analysis between the C-BASS and S-PASS data, accounting for all such complications, is beyond the scope the work presented here.

4.2.2 Synchrotron EB Correlation

For inflationary models where parity is conserved, the EB correlation of the CMB is expected to vanish due to the opposite parities of the E and B fields ([Zaldarriaga and Seljak, 1997](#)). This fact can be exploited to perform CMB self-calibration. In order to define the Stokes parameters on the sky, one must be able to transform from the instrument frame to the sky frame. This is determined through a relative orientation angle ψ . However, a host of systematic effects can lead to small deviations $\Delta\psi$, which in turn lead to a rotation of the observed Stokes parameters ([Keating et al., 2013](#); [Kaufman et al., 2014](#)). CMB self-calibration proceeds by assuming any measured EB correlation is due to systematic effects. By setting the EB correlation to zero one can therefore correct for any misalignment. However, this also assumes that there are no residual polarized foregrounds in the recovered CMB maps, which could in principle display some level of EB correlation. Following the standard self-calibration procedure described above can therefore lead to biases in subsequent cosmological analyses, which can be particularly important when considering sensitive next-generation searches for primordial B -modes ([Abitbol et al., 2016](#)).

In order to assess the potential impact of this effect it is therefore important to obtain constraints on the level of EB correlation in our polarized foreground components. In [Fig. 4.8](#) I show the measured EB correlation coefficients obtained for C-BASS. There is some evidence for significant EB correlation, particularly at $\ell > 100$. As noted above, the E and B maps do show highly non-Gaussian features e.g., the E -mode traces Galactic loops, where the B -mode is relatively weak. It is therefore not necessarily surprising that there should be some degree of synchrotron EB correlation. However, the measured correlation could potentially be caused by systematic effects related to the C-BASS polarization angle calibration. Work assessing the accuracy and stability of the C-BASS polarization angle calibration is ongoing.




Figure 4.8: Recovered EB correlation coefficients for the C-BASS northern sky map. Points are slightly offset for clarity, with the central point in each triple being shown at the effective multipole for that bin. There is some evidence for EB correlation, particularly at $\ell > 100$, which can be caused by the highly non-Gaussian features seen in the Galactic synchrotron emission. However, this can also be caused by systematic effects related to the C-BASS polarization angle calibration. Work assessing the accuracy and stability of this calibration is ongoing. Horizontal dashed lines indicate the $\mu_\rho \pm \Sigma_\rho$ values for each mask. These constraints are consistent with correlation coefficients of $|\rho_\ell^{EB}| \sim 0.2$ in the worst case.

These measurements can still be used to estimate limits on the level of synchrotron EB correlation. To do this I assume the correlation coefficients for a given mask are drawn from an underlying Gaussian distribution,

$$\rho_\ell^{EB} \sim \mathcal{N}\left(\mu = \mu_\rho, \sigma = \sqrt{\sigma(\rho_\ell^{EB})^2 + \Sigma_\rho^2}\right), \quad (4.20)$$

where μ_ρ is the mean of the distribution and Σ_ρ is the additional scatter in the correlation coefficients. I assign weakly informative priors that encapsulate the relevant parameter scales as follows,

$$\mu_\rho \sim \mathcal{N}(\mu = 0, \sigma = 1), \quad (4.21)$$

$$\Sigma_\rho \sim \text{Exp}(\lambda = 2). \quad (4.22)$$

The results of these fits are given in Table 4.5. In the worst case, the limits obtained here are consistent with synchrotron EB correlation at the level of $|\rho_\ell^{EB}| \sim 0.2$. It is clear that a careful analysis of residual synchrotron contamination in CMB maps for next-generation experiments will be important to avoid introducing biases through CMB self-calibration.

Mask	μ_ρ	Σ_ρ
p06	0.02 ± 0.08	0.12 ± 0.09
p04	0.03 ± 0.09	0.16 ± 0.11
p02	-0.08 ± 0.09	0.14 ± 0.11

Table 4.5: Mean EB correlation coefficients, and the estimated additional scatter. In the worst case, the measurements obtained here are consistent with correlation coefficients of $|\rho_\ell^{EB}| \sim 0.2$.

4.3 Joint C-BASS Analysis with the WMAP and *Planck* Surveys

To conduct a full assessment of synchrotron spectral properties requires a map-based analysis where spatial variations in synchrotron spectral parameters can be given a more complete treatment. This work is being conducted by the C-BASS collaboration, alongside the power-spectrum based analysis presented here, and is discussed briefly in Section 4.4.1. Nonetheless, a harmonic-space analysis of the synchrotron SED does bear relevance to next-generation CMB B -mode searches. Many foreground modelling approaches are conducted at the harmonic-level, particularly with ground-based experiments targeting small sky areas. In this section I present the results from a joint harmonic-space analysis of synchrotron emission with C-BASS and WMAP/*Planck* low-frequency data.

In Fig. 4.9 I show the recovered E -mode auto- and cross-spectra, and in Fig. 4.10 I show the corresponding B -mode auto- and cross-spectra. Frequency plots are shown for multipole bins with $\ell < 100$. For higher multipoles reliable model fits could no longer be obtained in the validation simulations and hence are not included here. Broadly, the recovered auto- and cross-spectra can be seen to follow a power-law scaling. Note that deviations from a pure power-law scaling are to be expected, given that spatial variations in synchrotron spectral parameters can lead to decorrelation

effects at higher frequencies, and the Galactic synchrotron emission is not a pure Gaussian random field.

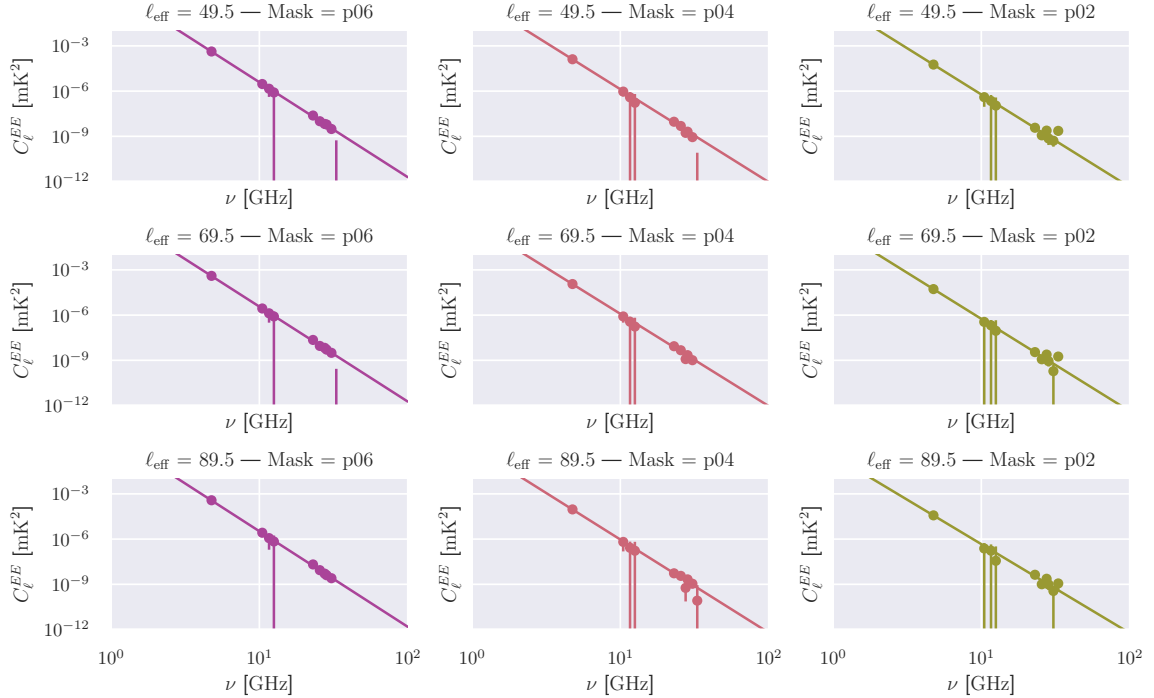


Figure 4.9: Recovered E -mode auto- and cross-spectra, for C-BASS, WMAP and *Planck* maps, plotted as a function of frequency. Moving horizontally, I show the recovered bandpowers for each mask. Moving vertically, I show the recovered bandpowers for each multipole bin, labelled by the effective multipole. Each point is plotted at the geometric mean of the two map frequencies used in the power-spectrum calculation ($\sqrt{\nu_1\nu_2}$). Alongside the bandpower estimates, I show the corresponding best-fit power law spectra, defined by the mean parameter values in Table 4.6.

In Table 4.6 I state the results obtained by fitting the model in Equation 4.12 to the auto- and cross-spectra as a function of frequency for each multipole bin and mask. The weighted average spectral index across the masks is $\beta_s^{EE} = -3.10 \pm 0.01$ for the E -mode power spectra, and $\beta_s = -3.20 \pm 0.03$ for the B -mode power spectra. The B -mode SED appears to be somewhat steeper than the E -mode SED. Looking again at the E and B maps, it can be seen that the E -mode tends to trace loops in the synchrotron emission. It has been proposed that these loops are relatively recent supernova remnants, which could possess systematically shallower synchrotron spectra (Mertsch and Sarkar, 2013; Vidal et al., 2015). The S-PASS analysis in Krachmalnicoff et al. (2018) found an average spectral index of $\beta_s = -3.22 \pm 0.08$

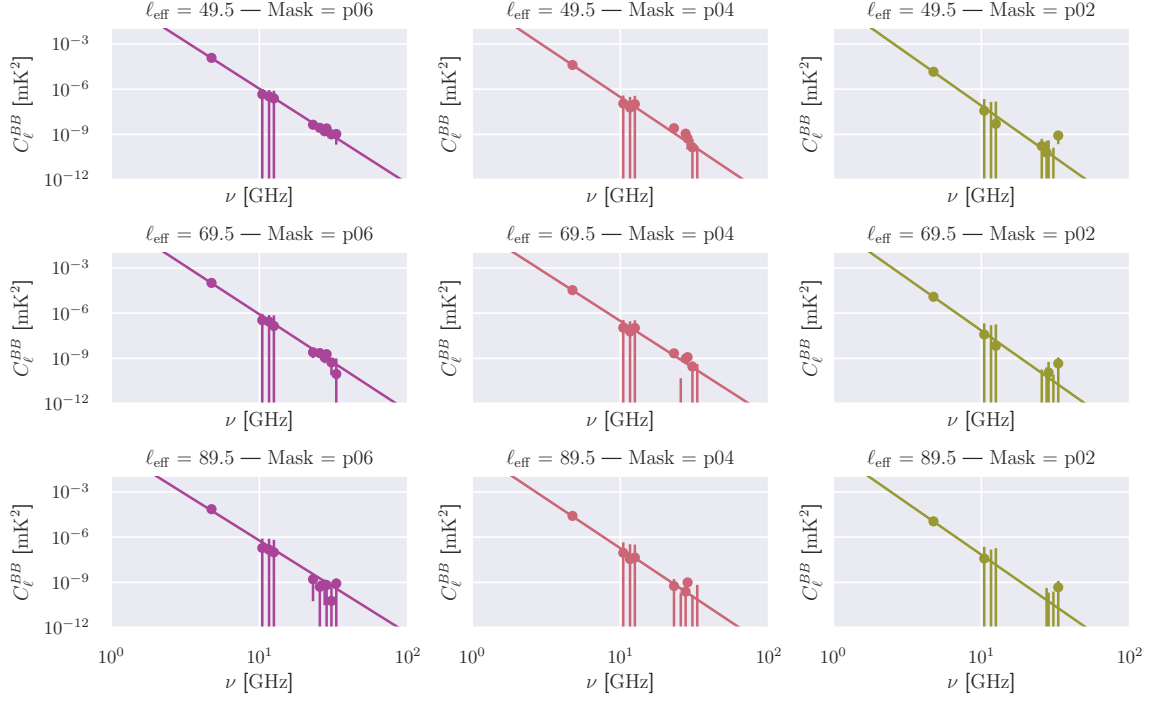


Figure 4.10: Recovered B -mode auto- and cross-spectra, for C-BASS, WMAP and *Planck* maps, plotted as a function of frequency. Moving horizontally, I show the recovered bandpowers for each mask. Moving vertically, I show the recovered bandpowers for each multipole bin, labelled by the effective multipole. Each point is plotted at the geometric mean of the two map frequencies used in the power-spectrum calculation ($\sqrt{\nu_1\nu_2}$). As for the E -mode power spectra, the recovered point estimates broadly follow the simulated frequency scaling, with bandpower estimates for WMAP and *Planck* data-points becoming heavily degraded beyond $\ell \gtrsim 100$. Poor bandpower estimates are also obtained for the WMAP and *Planck* data-points when applying the more aggressive p02 mask. Alongside the bandpower estimates, I show the corresponding best-fit power-law spectra, defined by the mean parameter values in Table 4.6.

over both E and B , and are broadly consistent with the spectral indices found in the C-BASS analysis presented here.

Multipole Bin	Mask	$A_{\nu_0}^{EE}$	$A_{\nu_0}^{BB}$	β_s^{EE}	β_s^{BB}
$40 \leq \ell < 60$	p06	$(41.9 \pm 2.4) \times 10^{-11} \text{ K}^2$	$(11.5 \pm 1.1) \times 10^{-11} \text{ K}^2$	-3.15 ± 0.03	-3.15 ± 0.05
$60 \leq \ell < 80$	p06	$(12.9 \pm 0.6) \times 10^{-11} \text{ K}^2$	$(4.0 \pm 0.3) \times 10^{-11} \text{ K}^2$	-3.08 ± 0.03	-3.30 ± 0.12
$80 \leq \ell < 100$	p06	$(5.7 \pm 0.3) \times 10^{-11} \text{ K}^2$	$(1.4 \pm 0.1) \times 10^{-11} \text{ K}^2$	-3.07 ± 0.04	-3.47 ± 0.24
$40 \leq \ell < 60$	p04	$(39.3 \pm 2.5) \times 10^{-11} \text{ K}^2$	$(9.9 \pm 1.0) \times 10^{-11} \text{ K}^2$	-3.15 ± 0.03	-3.21 ± 0.07
$60 \leq \ell < 80$	p04	$(11.4 \pm 0.8) \times 10^{-11} \text{ K}^2$	$(3.3 \pm 0.3) \times 10^{-11} \text{ K}^2$	-3.06 ± 0.03	-3.16 ± 0.10
$80 \leq \ell < 100$	p04	$(5.3 \pm 0.3) \times 10^{-11} \text{ K}^2$	$(1.2 \pm 0.1) \times 10^{-11} \text{ K}^2$	-3.06 ± 0.05	-3.46 ± 0.23
$40 \leq \ell < 60$	p02	$(36.7 \pm 2.8) \times 10^{-11} \text{ K}^2$	$(5.8 \pm 1.5) \times 10^{-11} \text{ K}^2$	-3.17 ± 0.03	-3.08 ± 0.92
$60 \leq \ell < 80$	p02	$(9.2 \pm 0.6) \times 10^{-11} \text{ K}^2$	$(2.4 \pm 0.4) \times 10^{-11} \text{ K}^2$	-3.09 ± 0.04	-3.28 ± 0.18
$80 \leq \ell < 100$	p02	$(3.7 \pm 0.2) \times 10^{-11} \text{ K}^2$	$(1.1 \pm 0.1) \times 10^{-11} \text{ K}^2$	-2.95 ± 0.05	-3.42 ± 0.24

Table 4.6: Recovered synchrotron spectral parameters, obtained using the C-BASS, WMAP K, WMAP Ka and *Planck* 30 GHz maps. Spectral index estimates in each multipole bin are largely consistent between masks.

4.4 Synchrotron Contamination to CMB B -modes

C-BASS data can be used to obtain estimates of the synchrotron contamination to primordial B -modes. To do this requires us to be able to extrapolate synchrotron amplitudes at 4.76 GHz to frequencies close to the foreground minimum (~ 90 GHz), where CMB experiments are typically targeted. In Section 4.4.1 I describe the synchrotron spectral index map that has been produced by the C-BASS collaboration. In Section 4.4.2 I then present estimates of the synchrotron contamination to CMB B -modes obtained by combining the C-BASS power spectrum estimates with the spectral index map.

4.4.1 Spatial Variation of Synchrotron Spectral Indices

The C-BASS collaboration has used the C-BASS, WMAP K, WMAP Ka and *Planck* 30 GHz maps to obtain estimates for the map-based synchrotron spectral indices. These were obtained by fitting a power-law model to the polarized intensity maps,

$$P_s(\nu) = A_s \left(\frac{\nu}{\nu_0} \right)^{\beta_s}. \quad (4.23)$$

Assuming the uncertainties on the Gaussian distributed Stokes parameters Q and U are equal, the observed polarized intensities are distributed as Rician random variables. This is a reasonable approximation for the C-BASS polarized emission maps and gives the likelihood for a given pixel as

$$\mathcal{L} = \prod_i \frac{P_i}{\sigma_i^2} \exp\left(-\frac{P_i^2 + P_s(\nu_i)}{2\sigma_i^2}\right) I_0\left(\frac{P_i P_s(\nu_i)}{\sigma_i^2}\right), \quad (4.24)$$

where P_i is the observed polarized intensity at a frequency ν_i , σ_i is the corresponding standard deviation of the noise and $I_0(z)$ is the zeroth order modified Bessel function of the first kind.

To estimate the synchrotron spectral indices, a weakly informative Gaussian prior on the spectral index was used, so that $\beta_s \sim \mathcal{N}(\mu = -3, \sigma = 0.5)$. This was selected to encompass previously measured values of the synchrotron spectral index, and to help regularize corresponding inferences. In regions of low SNR, the combination of C-BASS, WMAP and *Planck* maps struggle to constrain the synchrotron spectral index. In this situation, using an apparently uninformative prior actually places most of its probability mass at extreme parameter values, leading to poor point estimates on the synchrotron spectral indices. The Gaussian prior was multiplied by the Jeffreys independence-rule prior to further down-weight the posterior volume associated with large-magnitude negative spectral indices. For the Rician likelihood considered here, the Jeffreys prior on the spectral index is given by,

$$\pi_J(\beta_s) \propto \sqrt{\sum_i P_s(\nu_i) \ln\left(\frac{\nu_i}{\nu_0}\right) \left[\frac{2P_s(\nu_i)}{\sigma_i^2} - \frac{P_s^3(\nu_i)}{\sigma_i^4} \left(1 - \frac{I_1(P_s^2(\nu_i)/\sigma_i^2)}{I_0(P_s^2(\nu_i)/\sigma_i^2)}\right) \right]}, \quad (4.25)$$

where $I_1(z)$ is the first order modified Bessel function of the first kind. For the synchrotron amplitude the prior was chosen so that $p(A_s) \propto 1/A_s$. This corresponds to the Jeffreys prior for the polarized amplitudes, acting here to down-weight extreme amplitude values at $A_s \rightarrow \infty$.

For the C-BASS analysis, the posterior distribution was calculated through brute-force gridding, made possible by the fact the model only has two free parameters. The resulting synchrotron spectral index map is shown in Fig. 4.11 at $N_{\text{SIDE}}=8$, along with the corresponding standard deviation map. Significant small-scale variations are found in the synchrotron spectral index. In the harmonic domain, these can induce decorrelation effects that have the potential to corrupt power-law fits of the power spectra. In the map domain, neglecting these spatial variations has the potential to




Figure 4.11: Panel (a): The synchrotron spectral index map, estimated using the C-BASS, WMAP K, WMAP Ka and *Planck* 30 GHz maps. Panel (b): The corresponding standard deviation map. With the addition of C-BASS it is possible to place improved constraints on the synchrotron spectral index, particularly in regions of low SNR. Significant small-scale variations are found in the synchrotron spectral index, that must be accounted for in order to avoid biasing foreground analyses.

significantly bias CMB component separation analyses. These issues are discussed in more detail in Chapter 5.

A similar method was previously applied in [Jew and Grumitt \(2020\)](#) to estimate spectral indices between the *Planck* 30 GHz and 40 GHz maps. Here, we fitted directly for the polarized amplitudes in the two maps, using the resulting posterior samples to estimate the distribution over the spectral indices. Through this analysis, it was found that these low-frequency spectral indices show additional power on all angular scales, that is not present in the default simulations of PYSM. This result is consistent with analysis in [Krachmalnicoff et al. \(2018\)](#).

Of particular note in [Jew and Grumitt \(2020\)](#) was that the full-sky coverage enabled a detailed study of the spectral properties of the WMAP Haze ([Finkbeiner, 2004a](#); [Dobler and Finkbeiner, 2008](#); [Dobler, 2012a](#); [Planck Collaboration et al., 2013](#)). The WMAP Haze is a diffuse structure, consisting of a northern and southern bubble, located at the Galactic Centre with a total intensity spectral index of $\beta_s \approx -2.5$. The WMAP Haze has been found to be spatially correlated with the *Fermi* Bubbles, a symmetric double-lobed structure located at the Galactic Centre, observed at gamma-ray frequencies in Fermi-LAT data ([Dobler et al., 2010](#); [Su et al., 2010](#); [Dobler, 2012b](#)). A range of mechanisms have been proposed for the generation of the WMAP Haze and *Fermi* Bubbles, including dark matter annihilation ([Finkbeiner, 2004b](#)) and AGN-like activity in the Galactic Centre ([Su and Finkbeiner, 2012](#); [Yang et al., 2013](#)).

Previous analyses of the WMAP Haze in total intensity had fitted elliptical Gaussian templates to the haze that assumed its spectral properties to be uniform ([Planck Collaboration et al., 2013](#)). In Fig. 4.12 I show the spectral index estimates we obtained around the Galactic Centre. By performing a pixel-by-pixel analysis we found that the two bubbles have asymmetric spectral properties, with an average spectral index in the northern bubble of $\beta_s = -3.00 \pm 0.05$ and in the southern bubble of $\beta_s = -2.36 \pm 0.09$. It was further noted that the average spectral index in the two bubbles ($\beta_s = -2.57 \pm 0.08$) was remarkably close to the value obtained from the *Planck* template fitting analysis ($\beta_s = -2.56 \pm 0.05$). Whilst the analysis in [Jew and Grumitt \(2020\)](#) enabled us to study synchrotron spectral properties in regions of high SNR, away from the Galactic plane the *Planck* low-frequency maps alone are severely limited by noise. In this situation, the addition of C-BASS helps greatly in constraining the synchrotron spectral index.

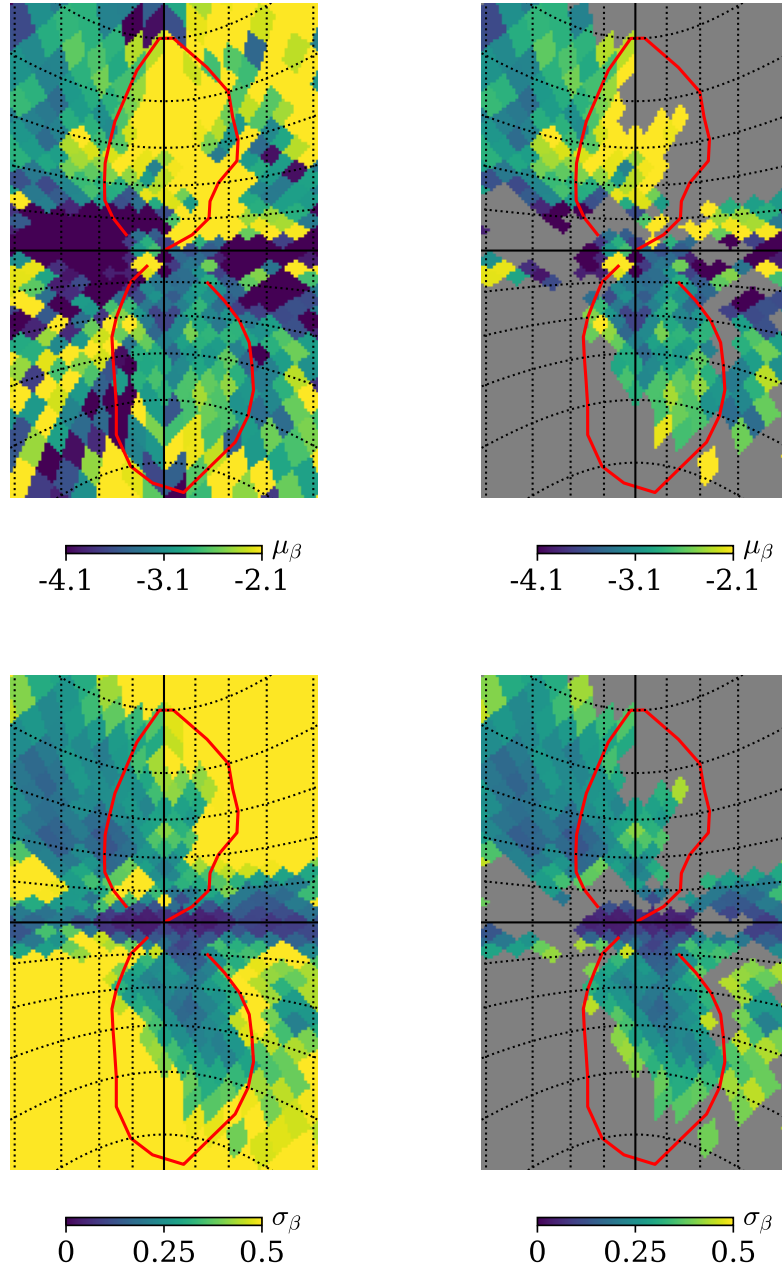


Figure 4.12: Gnomonic projections centred on the Galactic Centre, showing the mean (*top*) and standard deviation (*bottom*) of the posterior estimates of the spectral index estimates found in [Jew and Grumitt \(2020\)](#). In the *left column* no pixels have been masked, in the *right column* pixels identified as having poor detections of the spectral index are masked. The graticules show lines of constant latitude and longitude separated by 10° . The *red* lines mark the edges of the *Fermi* bubbles, as defined by [Su et al. \(2010\)](#). Figure taken from [Jew and Grumitt \(2020\)](#).

4.4.2 Contamination to Primordial B -modes

To estimate the synchrotron contamination to primordial B -modes I follow the method in [Krachmalnicoff et al. \(2018\)](#). Similar methods have also been used to assess the contamination of thermal dust emission to the CMB in [Planck Collaboration et al. \(2016a\)](#). I begin by dividing the $N_{\text{SIDE}}=128$ C-BASS map into a set of circular regions, with radius 15° , centred on $N_{\text{SIDE}}=8$ super-pixels. For each of these circular patches, I calculate the B -mode power spectrum, using a bin width of $\Delta\ell = 20$, and fit for the B -mode amplitude at $\ell = 80$. This is done using the power-law model in Equation 4.8 but fixing the spectral tilt to $\alpha_s^{BB} = -3.50$ (corresponding to the weighted average of the values found in Section 4.2.1). The increased uncertainties on the bandpower estimates, due to the smaller regions used, mean it is not possible to simultaneously constrain the synchrotron amplitude and spectral index.

Given the amplitude is a linear parameter, and assuming a Gaussian likelihood, I obtain MLEs given by


$$A_s^{\text{MLE}} = \frac{\sum_i (P_i/\sigma_i^2)(\nu_i/\nu_0)^{\beta_s}}{\sum_i (P_i/\sigma_i^2)(\nu_i/\nu_0)^{2\beta_s}}. \quad (4.26)$$

The uncertainty on this estimate can be calculated using the second derivative of the log-likelihood at the MLE, giving


$$\sigma(A_s^{\text{MLE}}) = \left\{ \sigma_i \left[\frac{1}{\sigma_i} \left(\frac{\nu_i}{\nu_0} \right)^{\beta_s} \right]^2 \right\}^{-1/2}. \quad (4.27)$$

Whilst improved constraints could potentially be obtained by fitting a full Bayesian model, as in the previous sections, this is not necessary here. The goal of this section is to estimate the approximate level of synchrotron contamination to B -modes, as opposed to obtaining detailed model constraints.


Given an estimate for the synchrotron amplitude at $\ell = 80$, I extrapolate this to 90 GHz using the spectral index map described in the previous section. After extrapolating to 90 GHz, I average the results in pixels common to multiple regions before downgrading the map from $N_{\text{SIDE}}=128$ to $N_{\text{SIDE}}=8$. In addition to a central estimate, I also obtain estimates for plausible upper and lower limits on the synchrotron contamination to B -modes. For the lower limit I extrapolate $A_s^{\text{MLE}} - \sigma(A_s^{\text{MLE}})$, using a lower limit on the spectral index in each pixel. For the upper limit I extrapolate $A_s^{\text{MLE}} + \sigma(A_s^{\text{MLE}})$, using an upper limit on the spectral index in each pixel. Finally, given the extrapolated synchrotron amplitudes, I divide these by the primordial CMB B -mode amplitude at $\ell = 80$ ($A_{\text{cmb}}^{BB}(\ell = 80) = 7.54 \times 10^{-2} \mu\text{K}^2$) to obtain estimates for the equivalent tensor-to-scalar ratio r_{eff} .




(a) $\langle r_{\text{eff}} \rangle$: Map




(b) $\langle r_{\text{eff}} \rangle$: Histogram/KDE plot




(c) r_{eff}^+ : Map



(d) r_{eff}^+ : Histogram/KDE plot



(e) r_{eff}^- : Map



(f) r_{eff}^- : Histogram/KDE plot

Figure 4.13: Estimates of the synchrotron contamination to CMB B -modes, obtained using the C-BASS auto-spectrum, expressed as equivalent tensor-to-scalar ratios. In the left column I show maps of the r_{eff} values in $N_{\text{SIDE}}=8$ pixels outside the p06 mask. In the right column I show the corresponding histogram/KDE plots. The minimum level of synchrotron contamination to primordial B -modes is found to be at the level of $r_{\text{eff}} \sim 10^{-3}$.

Maps of the estimated r_{eff} values are shown in Fig. 4.13, alongside histograms and KDE plots of the r_{eff} values. I only show values for pixels outside of the WMAP p06 mask. Regions within the mask are heavily contaminated by diffuse Galactic emission, and are typically masked out for any CMB polarization analyses. The minimum level of the synchrotron contamination to CMB B -modes corresponds to $r_{\text{eff}} \sim 10^3$, with typical synchrotron contamination at the level of $r_{\text{eff}} \sim 10^{-2}$. This is consistent with results obtained over the southern sky, using S-PASS data in [Planck Collaboration et al. \(2016a\)](#). Although these results are only approximate estimates for the synchrotron contamination, they are robust to uncertainties in the extrapolation. The minimum synchrotron contamination therefore corresponds to the targeted sensitivity of next-generation satellite missions ([Sekimoto et al., 2018](#)), with the typical synchrotron contamination already being at the levels targeted by next-generation ground-based CMB experiments ([Abazajian et al., 2016](#); [Simons Observatory Collaboration et al., 2019](#)). It is clear that to avoid biasing any potential detections of primordial B -modes, synchrotron emission cannot be neglected, with additional low-frequency channels being necessary to properly constrain synchrotron spectral properties.

4.5 Conclusions

In this chapter I have presented the results from an early polarization analysis, using the C-BASS northern sky, WMAP K, WMAP Ka and *Planck* 30 GHz maps. Auto- and cross-spectra of these maps were calculated using the NAMASTER library, which was validated against a set of Gaussian simulations. Using these simulations, I found I was able to obtain reliable results for the C-BASS auto-spectra over the multipole range $40 \leq \ell < 120$. When studying the synchrotron SED, using auto- and cross-spectra with WMAP/*Planck* maps, I was able to recover reliable results over the multipole range $40 \leq \ell < 100$, being limited at higher multipoles by noise in the WMAP and *Planck* maps.

To the C-BASS E and B -mode auto-spectra, I fitted a power-law model. The adequacy of the power-law model was assessed using posterior predictive checks, with the observed data being consistent with the posterior predictive distribution in all cases. From this analysis I found a weighted average spectral tilt of $\alpha_s^{EE} = -3.37 \pm 0.06$ for the E -mode auto-spectrum, and $\alpha_s^{BB} = -3.50 \pm 0.09$ for the B -mode auto-spectrum. The weighted average of the E -to- B ratio was $A_s^{EE}/A_s^{BB} = 3.9 \pm 0.01$. Given the differing spatial distributions of the C-BASS E and B -mode maps, and

their highly non-Gaussian features, unequal power in E and B is not surprising. These results differ somewhat from those in [Krachmalnicoff et al. \(2018\)](#), who found an average spectral tilt of approximately -3.15 at Galactic latitudes $|b| > 30^\circ$, and slightly steeper spectral tilts for the E -mode power spectra. Over the same Galactic latitudes, the S-PASS analysis found E -to- B ratios of ~ 2 . However, it is important to note that Faraday rotation between S-PASS and C-BASS frequencies complicates any direct comparisons of these results.

I also obtained constraints on the synchrotron EB correlation. There was some evidence for non-zero EB correlation, particularly at $\ell \geq 100$, which can be caused by non-Gaussian features in the Galactic synchrotron emission. In the worst case, the C-BASS data are consistent with correlation coefficients of up to $|\rho_\ell^{EB}| \sim 0.2$. Non-zero EB correlations have the potential to bias the CMB self-calibration process, meaning a careful accounting of residual synchrotron emission in any recovered CMB maps is required before conducting any self-calibration with next-generation CMB experiments. Note that spurious EB correlations can potentially be caused by systematic errors in the C-BASS polarization angle calibration. Work to assess the accuracy and stability of this is ongoing.

In studying the auto- and cross-spectra between the C-BASS, WMAP K, WMAP Ka and *Planck* 30 GHz maps I was able to obtain constraints on synchrotron spectral behaviour in the harmonic domain. Across the multipole bins considered, I obtained a weighted average spectral index of $\beta_s = -3.10 \pm 0.01$ for the E -mode power spectra, and $\beta_s = -3.20 \pm 0.03$ for the B -mode power spectra. The spectral index constraints were broadly consistent with those found in [Krachmalnicoff et al. \(2018\)](#). The shallower E -mode spectral indices could potentially result from the different Galactic structures traced by the E and B -mode components of the Galactic synchrotron emission e.g., the E -mode preferentially traces the Galactic loops.

Finally, I used the C-BASS B -mode auto-spectrum to obtain estimates for the level of synchrotron contamination to B -modes. Outside the p06 mask it was found that the minimum level of synchrotron contamination corresponded to a tensor-to-scalar ratio of $r_{\text{eff}} \sim 10^{-3}$, with typical values across the sky of $r_{\text{eff}} \sim 10^{-2}$. The level of the minimum contamination corresponds to the sensitivities of planned next-generation CMB satellites, and the typical level of contamination corresponds to the already targeted sensitivities of next-generation ground-based experiments. It is clear from this that synchrotron emission cannot be neglected in any B -mode foreground analyses, with additional low-frequency channels being required to properly constrain synchrotron spectral properties. In the next chapter, I consider how the addition of

low-frequency information from C-BASS to that expected from future higher frequency experiments such as *LiteBIRD* improves the extraction of the CMB B -mode signal, using a new implementation of Bayesian parametric component separation.

Chapter 5

Hierarchical Bayesian CMB Component Separation with the No-U-Turn Sampler

In this chapter I focus on algorithmic challenges in detecting primordial CMB B -modes, in particular studying Bayesian parametric component separation. The material in this chapter has been published in [Grmitt et al. \(2020a\)](#). One of the main motivations for using Bayesian parametric component separation is the ability to obtain properly motivated probability distributions for our model parameters, and hence a proper quantification of the uncertainty. This does come at the cost of Bayesian inference being computationally expensive, especially when using MCMC techniques. Further, uncertainties in our modelling of foreground emission can lead to biases in our inference. This is potentially highly problematic in the case of B -mode science, where the potential biases can be of the same order as the value of r we are attempting to measure. The impact of such modelling errors have received significant previous attention, see e.g., [Remazeilles et al. \(2016\)](#); [Chluba et al. \(2017\)](#); [Remazeilles et al. \(2018\)](#). These modelling problems are very closely linked with the available data. If it becomes apparent that more complex modelling is required, more comprehensive data covering a wide range of frequencies will be required to constrain the additional model parameters ([Jew et al., 2019](#)).

In this chapter I present a new implementation of Bayesian parametric CMB component separation, using the NUTS algorithm to explore the parameter space. NUTS is a self-tuning variant of HMC, originally presented in [Hoffman and Gelman \(2014\)](#). Since this original exposition, the algorithm has undergone a number of developments, with state-of-the-art implementations in STAN ([Stan Development Team, 2012](#); [Car-](#)

penter et al., 2017) and PYMC3 (Salvatier et al., 2016). HMC algorithms make use of first-order gradient information to generate efficient proposal steps. This allows HMC to avoid the random-walk behaviour of standard Metropolis-Hastings and Gibbs sampling methods, which becomes particularly problematic as the dimension of the parameter space increases.

To validate the performance of the component separation algorithm, I apply it to simulated observations of *LiteBIRD* and C-BASS. Previous CMB component separation analyses have considered the extent to which foreground spectral parameters should be allowed to vary, such that a balance can be struck between model realism and simplicity. These studies have considered various approaches to defining regions on the sky, over which foreground spectral parameters are typically assumed to be constant (Stompor et al., 2009; Errard et al., 2011; Stompor et al., 2016; Alonso et al., 2017; Irfan et al., 2019; Khatri, 2019; Thorne et al., 2019). In this chapter I use the mean-shift algorithm to define regions on the sky, clustering according to naively estimated synchrotron and dust spectral parameters. I initially fit a complete pooling model where I assume foreground spectral parameters to be constant in each region. I then fit a hierarchical foreground model. Hierarchical modelling has recently been employed in the context of blind CMB component separation in Wagner-Carena et al. (2020). For the hierarchical analysis in this chapter, I assume individual pixel spectral parameters are drawn from underlying Gaussian distributions, jointly fitting for the mean and variance of the Gaussian hyper-distributions, and the individual pixel-by-pixel spectral parameters in each region. In doing so, the hierarchical model is able to provide a faithful generative description of the underlying foreground emission, whilst reducing the propensity for fitting noisy outliers when assuming total independence between pixel spectral parameters (Gelman, 2006a; Gelman and Hill, 2007).

Bench-marking against the rate at which the algorithm generates effective CMB amplitude samples, I find that NUTS offers performance improvements of $\sim 10^3$ over sampling with the Metropolis-Hastings algorithm for the complete pooling model. Sampling from the posterior distribution of the hierarchical model is particularly challenging. Hierarchical models are known to exhibit geometrical pathologies that make it extremely difficult to achieve convergence using non-gradient based sampling algorithms. In these situations variants of HMC are often the only tractable approaches to sampling from the posterior (Betancourt and Girolami, 2015).

The outline of this chapter is as follows: In Section 5.1 I discuss methods for performing approximate Bayesian inference with high-dimensional models. In Section

5.2 I describe the polarized emission simulations for the C-BASS and *LiteBIRD* experiments used in the algorithm validation. In Section 5.3 I describe the mean-shift clustering algorithm used to define regions on the sky. In Section 5.4 I describe the component separation algorithm, give a general description of the NUTS algorithm, and discuss the complete pooling and hierarchical foreground models used during validation. In Section 5.5 I present the results from the algorithm validation. I conclude in Section 5.6. Throughout this chapter I work in units of Rayleigh-Jeans brightness temperature unless otherwise stated.

5.1 Approximate Bayesian Inference

Attempting to perform parametric Bayesian CMB component separation over the whole sky presents us with a formidable problem, in attempting to perform integration over a very high-dimensional posterior distribution. In general, the goal of an applied Bayesian analysis is to calculate expectations of functions of the model parameters $f(\Theta)$ with respect to the target distribution $\pi(\Theta)$, i.e.,

$$\mathbb{E}_\pi[f] = \int \pi(\Theta) f(\Theta) d\Theta. \quad (5.1)$$

For the purposes of our discussion here, Θ represents some set of model parameters. Beyond $\mathcal{O}(10)$ dimensions, evaluating these integrals becomes intractable using brute-force methods or deterministic integration algorithms. Fundamentally, these methods fail in high dimensions because of the way they attempt to explore the entirety of the parameter space. In high dimensions most of the space that these algorithms explore makes a negligible contribution to any expectations. Contributions to expectations are determined through the product of the probability density and volume, $\pi(\Theta) d\Theta$, i.e., the probability mass. As the number of dimensions increases, the probability mass increasingly concentrates in a small region known as the typical set (Mackay, 2003; Betancourt, 2017). For example, for a multivariate Gaussian, as the number of dimensions grows the probability mass concentrates in a thin shell away from the mode of the distribution.

In order to evaluate the integral in Equation 5.1 efficiently we require an algorithm that spends most of its time exploring the typical set. A range of methods have been developed for this purpose, but one of the most effective tools has been MCMC. MCMC algorithms work by generating correlated samples from the target distribution, through simulating a Markov chain whose stationary distribution is the target distribution. A Markov chain can be specified through an initial probability

distribution $p^0(\Theta)$ and a transition probability $T(\Theta'; \Theta) = p(\Theta'|\Theta)$. The state of the Markov chain at the $(t + 1)$ th step is given by

$$p^{t+1}(\Theta) = \int T(\Theta'; \Theta)p^t(\Theta) d\Theta. \quad (5.2)$$

We can generate a Markov chain by repeatedly applying the transition operator, drawing successive parameter samples as $\tilde{\Theta}^t \leftarrow T(\Theta^t|\Theta^{t-1})$. To construct an MCMC algorithm we need to select the transition probability such that the target distribution is an invariant distribution i.e.,

$$\pi(\Theta') = \int T(\Theta'; \Theta)\pi(\Theta) d\Theta. \quad (5.3)$$

The Markov chain must also be ergodic so that $p^t(\Theta) \rightarrow \pi(\Theta)$ as $t \rightarrow \infty$ for any $p^0(\Theta)$. The transition is typically chosen to satisfy detailed balance so that

$$T(\Theta'; \Theta)\pi(\Theta) = T(\Theta; \Theta')\pi(\Theta'). \quad (5.4)$$

If detailed balance is satisfied then $\pi(\Theta)$ is an invariant distribution of the Markov chain. A detailed discussion of MCMC algorithms and Markov chains can be found in [Mackay \(2003\)](#). In this chapter I explore the application of the NUTS algorithm to high-dimensional inference in CMB component separation, with the algorithm being discussed in more detail in [Section 5.4.2](#).

Whilst variants of HMC in particular have shown great success in enabling Bayesian inference over very high-dimensional spaces, for problems involving extremely large datasets or very complex models these MCMC methods can prove computationally intractable. Recently, extensive research has been conducted into variational inference (VI), with a detailed review in [Blei et al. \(2016\)](#). VI approximates the posterior through some family of approximating distributions \mathcal{Q} , and finds the member of this family that minimizes the Kullback-Leibler (KL) divergence with the posterior $p(\Theta|\mathbf{d})$,

$$q^*(\Theta) = \arg \min_{q(\Theta) \in \mathcal{Q}} \text{KL}(q(\Theta)||p(\Theta|\mathbf{d})) = \arg \min_{q(\Theta) \in \mathcal{Q}} \{ \mathbb{E}_q[\ln q(\Theta)] - \mathbb{E}_q[\ln p(\Theta|\mathbf{d})] \}. \quad (5.5)$$

However, the KL divergence depends on $p(\mathbf{d})$ through the posterior expectation, which is intractable to calculate for the high-dimensional problems of interest. Instead, VI maximises the evidence lower bound (ELBO),

$$\text{ELBO}(q) = \mathbb{E}_q[\ln p(\Theta, \mathbf{d})] - \mathbb{E}_q[\ln q(\Theta)]. \quad (5.6)$$

This is equivalent to the negative KL divergence plus $\ln p(\mathbf{d})$, which does not depend on $q(\Theta)$. By turning the problem of sampling from the posterior into an optimization problem, VI can potentially scale to much more high-dimensional problems. However, whilst MCMC algorithms are guaranteed to asymptotically converge on the target distribution, VI does not have the same guarantees. Poorly chosen approximating families can fail to properly capture the target geometry, leading to poor inference. I do not consider applications of VI to the problem of CMB component separation here. However, this would be an interesting avenue for scaling Bayesian inference to more complex sky models e.g., those including global instrumental corrections and more complex hierarchical structures than those presented in Section 5.4.4.

5.2 Simulations

To validate the performance of the component separation algorithm I generate a set of simulated Stokes Q and U maps, corresponding to the frequencies and sensitivities of the C-BASS and *LiteBIRD* experiments (Jones et al., 2018; Sekimoto et al., 2018). The frequencies and polarization sensitivities of the C-BASS and *LiteBIRD* simulations are given in Table 5.1.

LiteBIRD is a planned next-generation CMB satellite, aiming to measure the tensor-to-scalar ratio with a sensitivity of $\sigma(r) \sim 10^{-3}$. To accomplish this task *LiteBIRD* will target large angular scales up to $\ell \sim 200$, covering the reionization peak at $\ell \sim 10$ and the recombination peak at $\ell \sim 80$ in the primordial B -mode power spectrum. The experiment is proposed to cover frequencies from 40 GHz to 402 GHz, with the lowest resolution 40 GHz channel having a resolution of approximately 70 arcmin. I smooth all of the simulated maps to this 70 arcmin resolution, which is sufficient for the validation analysis targeting angular scales $\ell \lesssim 200$.

The focus here is on component separation for CMB polarization studies, and as such I consider only spectral models for synchrotron, thermal dust and CMB emission. I simulate maps of polarized emission using PYSM (Thorne et al., 2017). Polarized synchrotron maps were generated at each frequency using the PYSM s1 model. This used the 9-year *WMAP* 23 GHz maps (Bennett et al., 2013), smoothed to 3° , as synchrotron Q/U templates. Small scales are added to these templates by extrapolating the map power spectra to high ℓ and obtaining Gaussian realizations of the power spectra. Details of the implementation can be found in Thorne et al. (2017). These synchrotron templates were then extrapolated to higher frequencies

Table 5.1: Frequencies and polarization sensitivities for the C-BASS and *LiteBIRD* experiments, used in the simulations for algorithm validation. Sensitivities are given in CMB thermodynamic temperature units. Simulated maps were smoothed to a common resolution of 70 arcmin, corresponding to the lowest resolution, 40 GHz *LiteBIRD* channel. This resolution is sufficient for targeting multipoles $\ell \lesssim 200$, i.e., angular scales corresponding to the reionization and recombination peaks of the primordial *B*-mode power spectrum.

Experiment Name	Frequency [GHz]	Polarization Sensitivity [μK_{cmb} -arcmin]
C-BASS ^a	5	4320
	40	27.9
	50	19.6
	60	15.6
	68	12.3
	78	10.0
	89	9.4
<i>LiteBIRD</i> ^b	100	7.6
	119	6.4
	140	5.1
	166	7.0
	195	5.8
	235	8.0
	280	9.1
	337	11.4
	402	19.6

^aJones et al. (2018)

^bSekimoto et al. (2018)

using a power-law frequency spectrum,

$$S_s = A_s \left(\frac{\nu}{\nu_0} \right)^{-\beta_s}, \quad (5.7)$$

where A_s is the reference synchrotron amplitude, ν is the observing frequency, $\nu_0 = 23$ GHz is the reference frequency and β_s is the synchrotron spectral index. This was done across the whole sky using a spatially varying spectral index map taken from [Miville-Deschênes et al. \(2008\)](#). Across the whole sky, the synchrotron spectral index map has a mean of $\langle \beta_s \rangle \approx -3.0$ and a standard deviation of $\sigma(\beta_s) \approx 0.06$. It is worth noting that analysis of the synchrotron angular power spectrum in [Krachmalnicoff et al. \(2018\)](#) found that the spectral index map used in PYSM lacks power on all angular scales. However, for the purposes of the validation analysis it is sufficient.

For thermal dust I use the PYSM d1 model, which uses the *Planck* 353 GHz maps as Q/U thermal dust templates. The templates are scaled using the MBB model defined in Equation 1.36, with spatially varying dust temperature and spectral index maps obtained from the *Planck* COMMANDER analysis ([Planck Collaboration et al., 2016d](#)). Across the whole sky the mean of the dust temperature map is $\langle T_d \rangle \approx 20.9$ K and the standard deviation is $\sigma(T_d) \approx 2.2$ K. For the dust spectral index the mean value is $\langle \beta_d \rangle \approx 1.54$ and the standard deviation is $\sigma(\beta_d) \approx 0.04$. The dust templates are smoothed to 2.6° , with small scales being added using the same prescription as for synchrotron emission.

The CMB map was generated as a realization of lensed CMB power spectra, $\mathbf{C}_\ell = (C_\ell^{TT}, C_\ell^{EE}, C_\ell^{BB}, C_\ell^{TE})$, calculated using CAMB ([Lewis et al., 2000](#)). For the B -mode simulation I set $r = 5 \times 10^{-3}$, and the lensing amplitude $A_L = 1$. The CMB follows a blackbody spectrum given by,

$$S_{\text{cmb}} = A_{\text{cmb}} \frac{x^2 \exp(x)}{(\exp(x) - 1)^2}, \quad (5.8)$$

where A_{cmb} is the CMB amplitude, $x = h\nu/k_B T_{\text{cmb}}$ and $T_{\text{cmb}} = 2.7255$ K is the mean CMB temperature ([Fixsen, 2009](#)). In this validation study I do not consider the detailed impact of delensing on recovered B -mode estimates. In future work it would be valuable to combine the foreground analysis presented here, with a Bayesian delensing scheme such as in [Millea et al. \(2020\)](#), to form a complete forward model for our cosmological observations. The component amplitude templates and spectral parameter maps used in these simulations are shown in Fig. 5.1.

It is important to note here that the true underlying foreground SEDs are more complicated than those used in this analysis. Mis-modelling the underlying sky emission can result in biases in recovered cosmological parameter estimates. This has

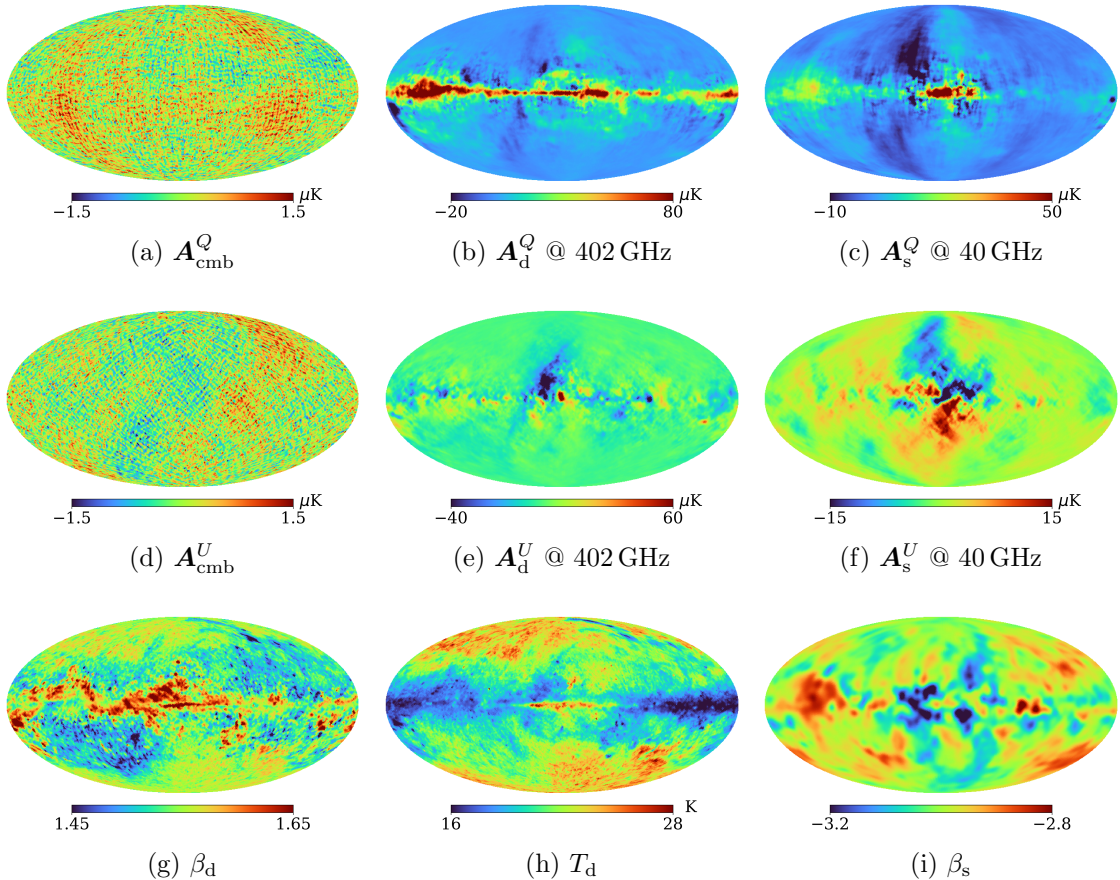


Figure 5.1: Input component parameter maps used in the validation simulations. The synchrotron and dust amplitude maps are shown here at reference frequencies of 40 GHz and 402 GHz respectively. When performing the component separation with C-BASS and *LiteBIRD* simulations, I set the synchrotron reference frequency to $\nu_0 = 5$ GHz.

received significant attention in B -mode forecasting analyses, see e.g., [Errard et al. \(2016\)](#); [Remazeilles et al. \(2016\)](#); [Remazeilles et al. \(2018\)](#). In this chapter I present a hierarchical foreground model (described in Section 5.4.4), that enables us to model spatial variations in foreground spectral parameters, without inflating parameter uncertainties. Given the simulated datasets considered here, fitting foreground spectral parameters with any parametric method will ultimately be limited by the number of frequency channels available with sufficient sensitivity. For example, in [Jew et al. \(2019\)](#), it was found that to constrain models including the synchrotron curvature would require additional low-frequency channels between 5 GHz and 30 GHz. Given a maximum observed frequency of ~ 400 GHz, dust spectral parameters in the modified blackbody (MBB) model are also very poorly constrained ([Remazeilles, 2018](#); [Jew et al., 2019](#)). This problem could be alleviated to some extent by considering

model re-parametrizations that yield more tractable posterior geometries. However, at its core this is a problem in the data.

5.3 Mean-Shift Clustering of Sky Regions

The modelling of spectral parameters in CMB component separation presents a number of challenges. In the face of limited data it can prove difficult to properly constrain spectral parameters. Attempting to allow full pixel-by-pixel variations in the spectral parameters in this situation can result in significant increases in post component separation noise and in the prior dominating the posterior. Given this, one may seek to reduce the number of degrees of freedom in the sky model by fitting for global spectral parameters. However, this approach will inevitably lead to modelling errors that have the potential to bias cosmological measurements made with the derived CMB map (Thorne et al., 2019). These challenges have motivated modelling approaches where spatially uniform spectral parameters are assumed over a set of defined sky regions. These have included regions defined as super-pixels on low N_{SIDE} HEALPIX maps, and regions defined according to similarities in their spectral properties (Górski et al., 2005; Stompor et al., 2009; Errard et al., 2011; Stompor et al., 2016; Alonso et al., 2017; Irfan et al., 2019; Khatri, 2019; Thorne et al., 2019). A detailed discussion of the modelling approaches that can be adopted for spectral parameters in these sky regions is given in Sections 5.4.3 and 5.4.4. In this section I describe the mean-shift clustering algorithm, implemented in SCIKIT-LEARN (Comaniciu and Meer, 2002; Pedregosa et al., 2011), that was used to define regions on the sky for the component separation analyses. The mean-shift clustering algorithm has previously been used in Jew and Grumitt (2020) to identify pixels with good detections of the synchrotron spectral index.

I use the mean-shift clustering algorithm to construct sky regions according to their location on the sky and their spectral properties. Specifically, I cluster according to the Cartesian coordinates of pixel centres on the unit sphere, (x^1, x^2, x^3) , and the naive spectral indices between the C-BASS 5 GHz and *LiteBIRD* 40 GHz channels, and the *LiteBIRD* 337 GHz and 402 GHz channels in polarized intensity. The two sets of frequency maps are used as synchrotron and thermal dust tracers respectively, with the naive spectral indices in a pixel, p , being given by,

$$\beta_p^{i,j} = \frac{\ln(m_i(p)/m_j(p))}{\ln(\nu_i/\nu_j)}, \quad (5.9)$$

where $m_i(p)$ and $m_j(p)$ are the map values in the pixel p , at the frequencies ν_i and ν_j respectively. Given these parameters we may form the feature vector,

$$\mathbf{z}_p = (x_p^1/\zeta, x_p^2/\zeta, x_p^3/\zeta, \beta_p^{5,40}, \beta_p^{337,402}), \quad (5.10)$$

where ζ is a spatial vector scaling factor. Setting the value of ζ to be less than 1 allows us to preferentially weight proximity on the sky as being favourable over proximity in spectral index space. The function of the mean-shift algorithm is then to cluster points in this five-dimensional feature space.

The mean-shift algorithm proceeds by assigning a walker to each pixel, giving us the starting vectors, $\mathbf{z}_p^0 = \mathbf{z}_p$. The walkers then step through the feature space towards regions of higher density, with the t^{th} update being calculated as,

$$\mathbf{z}_p^t = \mathbf{z}_p^{t-1} + \mathbf{s}(\mathbf{z}_p^{t-1}), \quad (5.11)$$

where \mathbf{s} is the mean-shift vector, given by,

$$\mathbf{s}(\mathbf{z}_p^t) = \frac{\sum_q K(\mathbf{z}_q^t - \mathbf{z}_p^t) \mathbf{z}_q^t}{\sum_q K(\mathbf{z}_q^t - \mathbf{z}_p^t)}. \quad (5.12)$$

For this analysis I choose the kernel, K , to be a top-hat defined by,

$$K(\Delta) = \begin{cases} 1, & |\Delta| \leq \omega, \\ 0, & |\Delta| > \omega, \end{cases} \quad (5.13)$$

where ω is the bandwidth parameter. Walkers take steps until they converge i.e., walker positions in feature space no longer change with new updates (up to some threshold). Regions are then defined as a set of pixels whose walkers have converged on the same position in the 5-dimensional feature space. Any regions containing fewer pixels than some arbitrary minimum are reassigned to the nearest region in feature space containing a sufficient number of pixels.

In Fig. 5.2 I show the regions obtained using the frequency channels outlined above, which are used as the region definitions for the component separation analyses in this chapter. I have used the simulated maps at a HEALPIX N_{SIDE} of 64, setting $\zeta = 0.5$, $\omega = 0.3$ and the minimum number of pixels in a region to 10 (Górski et al., 2005). Using the foreground tracers and parameters described here, I obtain 171 regions. The smallest region on the sky contains 138 pixels, and the largest region contains 645 pixels. The mean number of pixels in a region is $\langle N_{\text{pix}} \rangle \approx 290$, and the standard deviation in the number of pixels is $\sigma(N_{\text{pix}}) \approx 95$. In regions of low SNR the borders of regions become less smooth. This can be mitigated to some extent by defining




Figure 5.2: Regions defined using the mean-shift clustering algorithm, produced using $N_{\text{SIDE}}=64$ maps. I set $\zeta = 0.5$, $\omega = 0.3$ and the minimum number of pixels in a region to 10. Given the set-up here, I obtain 171 regions. The smallest region contains 138 pixels, and the largest region contains 645 pixels. The mean number of pixels is $\langle N_{\text{pix}} \rangle \approx 290$, and the standard deviation of the number of pixels is $\sigma(N_{\text{pix}}) \approx 95$. In regions of low SNR we can see that the region borders become less smooth, due to the naive spectral indices used to cluster pixels on the sky becoming contaminated by noise.

regions on lower N_{SIDE} maps, at the cost of the larger pixel size meaning regions become more coarse. Using lower N_{SIDE} maps can lead to a significant degradation in computational performance when assuming foreground spectral parameters in a given region are related in some manner.

A range of possible, non-trivial, extensions to the clustering algorithm exist that could help to alleviate some of the issues surrounding region definition in areas of the sky with low SNR. Improved parameter constraints could likely be obtained for the hierarchical model discussed in Section 5.4.4 by defining larger regions in areas of low SNR. Instead of using naive spectral indices as a tracer of the spectral properties of diffuse emission over the sky, a more sophisticated estimation of the spectral indices could be performed, accounting for noise properties across the sky. The spectral index estimation could also be improved with better tracers of synchrotron and thermal dust emission. It is worth noting here that the use of the C-BASS map as a synchrotron template is particularly important. If instead I had only used low-frequency *LiteBIRD* channels as the synchrotron tracers the naive synchrotron spectral index estimates would have been heavily noise dominated. Using the methods outlined here, I found

it would only be possible to define very coarse regions on $N_{\text{SIDE}}=8$ maps.

5.4 Parametric Bayesian CMB Component Separation

I have developed a new implementation of Bayesian pixel-by-pixel CMB component separation, using the NUTS algorithm to explore the parameter space (Hoffman and Gelman, 2014). The primary benefit in using NUTS to sample from the target distribution is in its avoidance of the random walk behaviour that slows more conventional sampling algorithms such as Metropolis-Hastings and Gibbs sampling. The component separation code is written in the PYTHON programming language, with the NUTS algorithm being implemented through the PYMC3 library (Salvatier et al., 2016). In its current form, the whole-sky component separation is parallelized over the sky regions defined using the mean-shift clustering algorithm. For the validation analysis in this chapter, I do not consider monopole and dipole corrections, or instrumental factors such as colour corrections in the modelling. The extension of the modelling to include such complications is left to future work.

Bench-marking the algorithm performance against the rate of effective sample generation, NUTS offers potential performance improvements of $\sim 10^3$ compared to sampling with Metropolis-Hastings¹. Close to the Galactic plane sampling becomes more difficult, with the CMB completely sub-dominant to foregrounds. In this situation computational performance can be degraded such that the sampler exhibits undesirable random walk behaviour. By parallelizing over sky regions, and masking the most contaminated sky regions close to the Galactic plane, it should be possible to achieve rapid convergence. These bench-marking tests have been performed without extensive optimization of the component separation code. By exploring re-parametrizations, model prior choice, optimizing sampling parameters etc. it is likely that further performance improvements could be attained. It would also be worthwhile considering the potential for GPU acceleration. For the hierarchical model, the posterior exhibits geometrical pathologies that make sampling with non-gradient based algorithms essentially intractable (Betancourt and Girolami, 2015). However, using NUTS I am able to achieve comparable computational performance to the complete pooling model.

¹These simple bench-marking tests were performed on a single Intel Xeon CPU, running at ~ 2.6 GHz.

The outline of this section is as follows: In Section 5.4.1 I describe the general sky model and likelihood used in the component separation analysis. In Section 5.4.2 I give an overview of the NUTS algorithm. In Section 5.4.3 I discuss the complete pooling model, where I fit for constant spectral parameters over sky regions. In Section 5.4.4 I discuss the hierarchical modelling approach, where I directly fit for the underlying hyper-distributions of the spectral parameters. In Section 5.4.5 I describe the convergence checks and diagnostics performed during sampling.

5.4.1 Sky Model and Likelihood

In attempting to observe the CMB we actually observe multiple sky components. It is the goal of component separation to extract the CMB from these additional confusing components. For a given sky pixel, p , we may write the observed value in that pixel as

$$d_{p,\lambda}(\nu) = s_{p,\lambda}(\nu) + n_{p,\lambda}(\nu), \quad (5.14)$$

where ν is the observing frequency, $\lambda = \{I, Q, U\}$ represents one of the Stokes parameters, $s_{p,\lambda}(\nu)$ is the true sky signal and $n_{p,\lambda}(\nu)$ is the noise term. The analysis in this chapter is restricted to $\lambda = \{Q, U\}$.

The true sky signal includes contributions from diffuse emission, compact sources, line emission etc. For the current modelling purposes I focus on diffuse emission. Additional contributions from compact sources, line emission etc. are important considerations when performing component separation on real experimental data. A variety of strategies exist for mitigating their contribution, e.g. through masking point sources, and direct modelling during the fitting process. However, this sits beyond the scope of the algorithm validation.

The noise term consists largely of contributions from instrumental white noise and $1/f$ noise, which acts to introduce large-scale correlated noise in the sky maps. For ground-based experiments mitigating atmospheric noise is a significant challenge, largely limiting accessible angular scales to $\ell \gtrsim 30$. The ability of such experiments to accurately recover low multipoles will be vital for future B -mode experiments (Alonso et al., 2017).

Given a set of sky maps at frequencies $\nu = [\nu_1, \dots, \nu_N]$, the Gaussian likelihood for a pixel, p , is given by,

$$-\ln \mathcal{L}_{p,\lambda}(\Theta_{p,\lambda}) = \sum_i \frac{1}{2} \ln \left(2\pi\sigma_{p,\lambda}^2(\nu_i) \right) + \frac{\left(d_{p,\lambda}(\nu_i) - s_{p,\lambda}(\nu_i, \Theta_{p,\lambda}) \right)^2}{2\sigma_{p,\lambda}^2(\nu_i)}, \quad (5.15)$$

where $\Theta_{p,\lambda}$ are the model parameters, and $\sigma_{p,\lambda}$ is the pixel noise. I assume for simplicity that the noise is independent between pixels and frequency channels. This is not necessarily optimal when we consider the complications described above. However, the approximation is sufficient for the purposes of algorithm validation.

5.4.2 The No-U-Turn Sampler

The key measure of the efficiency of a sampling algorithm is in its ability to produce effective/independent samples, or equivalently reduce the correlation between samples. Indeed, this is where NUTS significantly out-performs standard Metropolis-Hastings and Gibbs sampling algorithms. Even though an individual step in these simpler sampling algorithms is less computationally expensive, their random-walk behaviour results in highly correlated samples and hence very inefficient generation of effective samples. At its core, the NUTS algorithm is an extension of HMC, which was originally developed for performing calculations in lattice field theory (Duane et al., 1987). In this section I give an overview of the HMC algorithm, and the additional tuning procedures NUTS implements to avoid the need for hand-tuned HMC implementations. For a detailed discussion of NUTS and HMC see Hoffman and Gelman (2014); Monnahan et al. (2017); Betancourt (2017); Betancourt et al. (2017).

HMC essentially proceeds by generating physical trajectories through parameter space, akin to simulating particle trajectories through a potential. A simple implementation of HMC may proceed as follows:

1. Given a set of parameters, $\Theta = [\Theta_1, \dots, \Theta_d]$, with corresponding joint density, $\pi(\Theta)$, we introduce a set of auxiliary momentum variables, $\mathbf{r} = [r_1, \dots, r_d]$. We take the distribution over the momenta to be a Gaussian distribution centered on zero, i.e., $\mathbf{r} \sim \mathcal{N}(\mathbf{0}, \mathbf{M})$, where \mathbf{M} is the mass matrix. This defines a kinetic energy term,

$$K(\Theta, \mathbf{r}) = -\ln p(\mathbf{r}|\Theta) = \frac{1}{2}\mathbf{r}^\top \mathbf{M}^{-1}\mathbf{r} + \ln |\mathbf{M}| + \text{const.} \quad (5.16)$$

We may then define the Hamiltonian of the system as,

$$\mathcal{H}(\mathbf{r}, \Theta) = -\ln p(\mathbf{r}|\Theta) - \ln \pi(\Theta) = K(\Theta, \mathbf{r}) + V(\Theta), \quad (5.17)$$

where we define the potential energy term, $V(\Theta) = -\ln \pi(\Theta)$.

2. We then evolve the position in parameter space by integrating Hamilton's equations,

$$\frac{d\Theta}{dt} = \frac{\partial \mathcal{H}}{\partial \mathbf{r}}, \quad (5.18)$$

$$\frac{d\mathbf{r}}{dt} = -\frac{\partial \mathcal{H}}{\partial \Theta}. \quad (5.19)$$

Practically this is done through a leapfrog algorithm. To generate a new sample in the Markov chain we draw \mathbf{r} from $\mathcal{N}(\mathbf{0}, \mathbf{M})$. The leapfrog steps then proceed as,

$$\mathbf{r}^{t+\epsilon/2} = \mathbf{r}^t - \frac{\epsilon}{2} \nabla_{\Theta} V(\Theta^t), \quad (5.20)$$

$$\Theta^{t+\epsilon} = \Theta^t + \epsilon \mathbf{M} \mathbf{r}^{t+\epsilon/2}, \quad (5.21)$$

$$\mathbf{r}^{t+\epsilon} = \mathbf{r}^{t+\epsilon/2} - \frac{\epsilon}{2} \nabla_{\Theta} V(\Theta^{t+\epsilon}), \quad (5.22)$$

where ϵ is the leapfrog step-size. The leapfrog steps used to update the position have the convenient property of being a symplectic integrator. That is, the numerical trajectory generated by the leapfrog steps preserve the volume of phase space, as is the case for the Hamiltonian trajectories they approximate. A more detailed discussion of the numerical integration of Hamilton's equations can be found in [Leimkuhler and Reich \(2004\)](#). These leapfrog steps are performed L times to generate a new proposal position, (\mathbf{r}^*, Θ^*) .

3. The new proposal position is then accepted with a probability of acceptance given by

$$p_{\alpha} = \min \left\{ 1, \exp \left(\mathcal{H}(\mathbf{r}, \Theta) - \mathcal{H}(\mathbf{r}^*, \Theta^*) \right) \right\}. \quad (5.23)$$

What we have done here is essentially generate a Metropolis-Hastings proposal step with a very high chance of being accepted.

4. By repeating this sampling procedure N times we may generate the parameter samples for the Markov chain.

The mass matrix used to define the distribution over the momenta acts to rotate and re-scale parameter space. Choosing \mathbf{M}^{-1} to be the covariance of the target distribution will help to de-correlate the target distribution, which can lead to significant performance improvements when dealing with highly correlated parameters. For practical implementations the mass matrix can be estimated during a tuning phase. Starting with the identity matrix we can generate an initial sample set, from which

we update the estimate of the mass matrix using the sample covariance. We may then iterate over this tuning process to generate an accurate estimate of the covariance of the target distribution (Betancourt, 2017). Whilst estimating off-diagonal elements of the mass matrix does help in de-correlating the target distribution, using the off-diagonal elements does not necessarily scale well to high-dimensional problems given the need to invert the mass matrix at the end of tuning, and perform matrix multiplications during leapfrog steps. For the analysis in this chapter I only tune diagonal elements of the mass matrix. However, this can significantly improve sampling efficiency for single pixel analyses, or analyses assuming complete independence between pixels. In this case, one can employ the tuning steps in Foreman-Mackey et al. (2019), using the default tuning schedule described in Stan Development Team (2012).

In principle, the HMC algorithm described above could be dramatically improved by allowing the mass matrix to vary with location in the parameter space. This implementation is known as Riemannian HMC (RHMC) (Girolami and Calderhead, 2011). By allowing the mass matrix to vary with location, RHMC is able to make local corrections to the target distribution that can greatly improve computational performance. Moreover, the $\ln |\mathbf{M}|$ term now acts like an energy sink, absorbing and releasing energy, so that the sampler is able to more efficiently explore the typical set of distributions with complex geometries. However, RHMC has proven extremely challenging to implement in practice, especially as an automatic implementation (Betancourt, 2017).

The efficiency of HMC as described above critically depends on the choice of ϵ and L used in the leapfrog steps. If ϵ is chosen to be too small, the sampler will waste computation taking very small steps along the Hamiltonian trajectories. In contrast, if ϵ is chosen to be too large, the simulation of the Hamiltonian trajectory will become inaccurate and the sampler will produce proposal steps with low acceptance probabilities. If L is chosen to be too small, the sampler will generate samples close to one another, resulting in undesirable random walk behaviour. If L is chosen to be too large on the other hand, the sampler will generate paths through parameter space that loop back on themselves. This results in proposal steps close to the starting value, with the additional waste of generating the extended trajectory. In even more severe scenarios, a poor choice of L that results in the sampler jumping from one side of parameter space to another at each iteration can result in a non-ergodic chain i.e. a chain that is not guaranteed to converge on the target distribution (Neal, 2012).

The need to finely tune ϵ and L means that standard implementations of HMC typically require costly tuning runs. This can significantly reduce the utility and general applicability of HMC in realistic problems. The NUTS algorithm overcomes these problems by automatically tuning these sampling parameters. The value of ϵ is tuned during an initial tuning phase to meet some target acceptance probability. The target acceptance probability can be adjusted depending on the degree of curvature in the posterior, with a higher acceptance rate (or equivalently, smaller step-size) being needed for highly curved distributions. The value of L is modified during sampling to meet a No-U-Turn criterion. That is, the leapfrog integrator is iterated over until the simulated trajectory begins to turn back on itself, or some maximum number of leapfrog simulations are performed. In doing so, the sampler is able to maximize the distance between the proposal step and the initial position, before looping back on itself and wasting computation. It is worth noting that the actual implementation of NUTS involves simulating a Hamiltonian trajectory both forwards and backwards in time. This is because naively applying the No-U-Turn criterion after only evolving forwards in time fails to satisfy detailed balance. Details on these tuning procedures can be found in [Hoffman and Gelman \(2014\)](#); [Stan Development Team \(2012\)](#); [Salvatier et al. \(2016\)](#).

5.4.3 Complete Pooling of Spectral Parameters

As discussed in Section 5.3, in the face of limited data and low SNR, allowing spectral parameters to vary completely from pixel-to-pixel is sub-optimal, resulting in increased levels of post component separation noise and the posterior potentially becoming prior dominated. As a first alternative to allowing full pixel-by-pixel variations we may instead assume spectral parameters to be constant over the regions defined as in Section 5.3 i.e., we assume a complete pooling of the spectral parameters. In Fig. 5.3 I show the directed graphical representation of the complete pooling model, which displays the explicit conditional dependencies between random variables in the Bayesian model.

The priors used in the complete pooling model are given in Table 5.2. I assign informative Normal priors to the spectral parameters. The standard deviations on the priors for β_s and β_d are chosen to be 0.3, corresponding to the bandwidth used in clustering sky regions and encompassing most of the range over which these parameters have been measured ([Planck Collaboration et al., 2016d](#); [Krachmalnicoff et al., 2018](#)). For the dust temperature I set a prior based on constraints on the dust

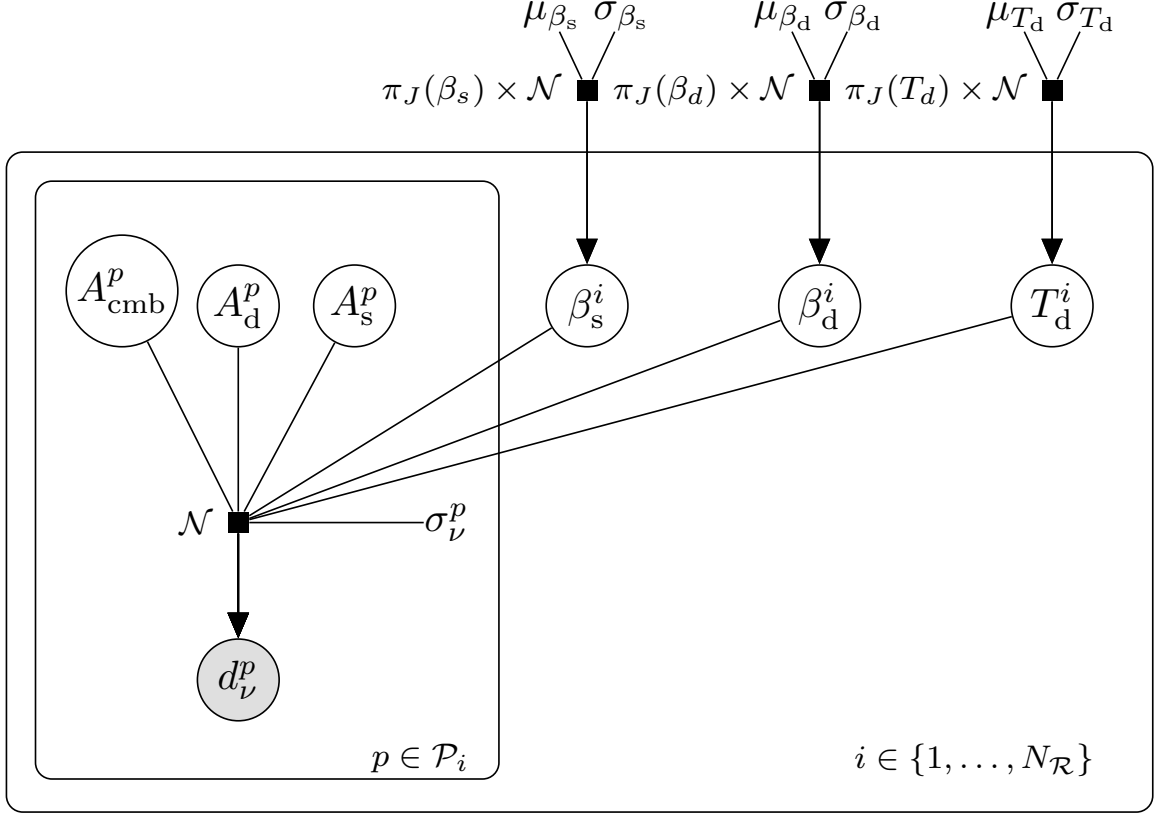


Figure 5.3: Graphical representation of the complete pooling model. Un-shaded circles represent random variables, shaded circles denote observed variables and un-circled parameters denote constants. Square nodes denote factors in the graph, in this case representing the distributions of model variables. Here I use plate notation to represent repeated sub-graphs, with a plate being repeated once for each index indicated for the plate. Directed lines (known as edges) joining variables/constants represent conditional dependencies in the model. Edges crossing plate boundaries are repeated once for each plate index that they cross into. Outside the plates, we have the global constants $\{\mu_X, \sigma_X\}$, $X = \{\beta_s, \beta_d, T_d\}$. Each pair denotes the mean and standard deviation of the Normal priors for the region-level spectral parameters, $\{\beta_s^i, \beta_d^i, T_d^i\}$. The larger plate contains these region-level spectral parameters, with $i \in \{1, \dots, N_{\mathcal{R}}\}$ denoting the region index. The smaller plate corresponds to the pixel-level, indexed by pixels $p \in \mathcal{P}_i$, where \mathcal{P}_i denotes the set of pixels contained within the region i . Here we have the pixel-level amplitude parameters, $\{A_{\text{cmb}}^p, A_s^p, A_d^p\}$, and the pixel standard deviation, σ_ν^p at the frequency ν . The observed pixel value at the frequency ν is given by d_ν^p .

temperature found in the *Planck* COMMANDER analysis (Planck Collaboration et al., 2016c,d). These help to down-weight the more extreme regions of parameter space, offering significant computational performance improvements, and helping to regularize the posterior by reducing the biasing effect of the probability mass associated with extreme parameter values. Detailed discussion of prior choice, in particular around the use of weakly informative priors, can be found in Gelman (2006b); Gelman et al. (2008); Evans and Jang (2011); Polson and Scott (2012); Gelman and Hennig (2017); Simpson et al. (2017); Gelman et al. (2017). In addition to the Normal priors on the spectral parameters we multiply these by the associated Jeffreys priors (Jeffreys, 1946, 1961) as in the *Planck* COMMANDER analysis (Eriksen et al., 2008; Planck Collaboration et al., 2014a, 2016c,d). For a single pixel the spectral Jeffreys priors are given by,

$$\pi_J(\beta_s) \propto \left(\sum_i \left[\frac{1}{\sigma_i} \left(\frac{\nu_i}{\nu_0} \right)^{\beta_s} \ln \left(\frac{\nu_i}{\nu_0} \right) \right]^2 \right)^{1/2}, \quad (5.24)$$

$$\pi_J(\beta_d) \propto \left(\sum_i \left[\frac{1}{\sigma_i} \left(\frac{\nu_i}{\nu_0} \right)^{\beta_d+1} \frac{\exp(\gamma\nu_0) - 1}{\exp(\gamma\nu_i) - 1} \ln \left(\frac{\nu_i}{\nu_0} \right) \right]^2 \right)^{1/2}, \quad (5.25)$$

$$\pi_J(T_d) \propto \left(\sum_i \left[\frac{1}{\sigma_i} \left(\frac{\nu_i}{\nu_0} \right)^{\beta_d+1} \frac{\exp(\gamma\nu_0) - 1}{\exp(\gamma\nu_i) - 1} \left\{ \frac{\nu_0}{1 - \exp(-\gamma\nu_0)} - \frac{\nu_i}{1 - \exp(-\gamma\nu_i)} \right\} \frac{1}{T_d^2} \right]^2 \right)^{1/2}. \quad (5.26)$$

The sums run over the observed frequencies, with σ_i being the pixel uncertainty at a frequency ν_i .

The emission model for a pixel, p in some sky region is given by,

$$s_{p,\lambda}(\nu) = A_s^{p,\lambda} f_s(\nu, \beta_s) + A_d^{p,\lambda} f_d(\nu, \beta_d, T_d) + A_{\text{cmb}}^{p,\lambda} f_{\text{cmb}}(\nu). \quad (5.27)$$

The functions, $f_s(\nu, \beta_s)$, $f_d(\nu, \beta_d, T_d)$ and $f_{\text{cmb}}(\nu)$ are the spectral forms of the synchrotron, dust and CMB components, as defined in Section 5.2. Note that I assume the spectral parameters to be identical for $\lambda = \{Q, U\}$. I model the data as being Normally distributed i.e., I assume the Gaussian likelihood in Equation 5.15.

Complete pooling offers a potentially effective approach to account for the spatial variation in spectral parameters whilst avoiding the generation of excessive post component separation noise. However, with additional data points at low and/or

Table 5.2: Priors for parameters in the complete pooling model. We assign informative Normal priors to the spectral parameters, and flat priors to the amplitude parameters. In addition to the Normal priors for the spectral parameters, we multiply these by their associated Jeffreys priors as in the *Planck* COMMANDER analysis. Spectral parameters are assumed to be constant over a given region, whilst amplitude parameters are allowed to vary within each pixel, p . The analysis is restricted to polarization, so that $\lambda = \{Q, U\}$.

Θ	$p(\Theta)$
β_s	$\mathcal{N}(\mu = -3, \sigma = 0.3) \times \pi_J(\beta_s)$
β_d	$\mathcal{N}(\mu = 1.6, \sigma = 0.3) \times \pi_J(\beta_d)$
T_d	$\mathcal{N}(\mu = 21.0 \text{ K}, \sigma = 2.0 \text{ K}) \times \pi_J(T_d)$
$A_s^{p,\lambda}$	$\text{Unif}(-\infty, \infty)$
$A_d^{p,\lambda}$	$\text{Unif}(-\infty, \infty)$
$A_{\text{cmb}}^{p,\lambda}$	$\text{Unif}(-\infty, \infty)$

high frequencies it is possible to adopt a more sophisticated, hierarchical model of the spectral parameters in these regions. It is worth noting that for the complete pooling model one can analytically marginalize over the amplitude parameters as in [Alonso et al. \(2017\)](#). This greatly reduces the dimension of parameter space and hence improves the sampling efficiency. I have not implemented sampling of this marginal distribution for the analysis here, where I study the computational performance of NUTS in sampling the full posterior. Indeed, the intrinsic efficiency of NUTS makes this unnecessary.

5.4.4 Hierarchical Modelling of Spectral Parameters

In statistical modelling, we often encounter scenarios where the model contains a set of latent variables that are related in some way. In such a scenario it is neither ideal to treat the latent variables as being entirely independent, or to simply fit for a single, global variable. Instead we can take a hierarchical approach. In a hierarchical Bayesian model we introduce a set of population-level hyper-parameters, which define the distribution from which the individual latent variables are drawn ([Gelman, 2006a](#); [Gelman and Hill, 2007](#)). The graphical representation of the hierarchical model is displayed in [Fig. 5.4](#).

In the context of CMB component separation, the pixel-level spectral parameters in each region can be modelled as being drawn from some underlying hyper-distributions. In this particular case I model the spectral parameters as being drawn

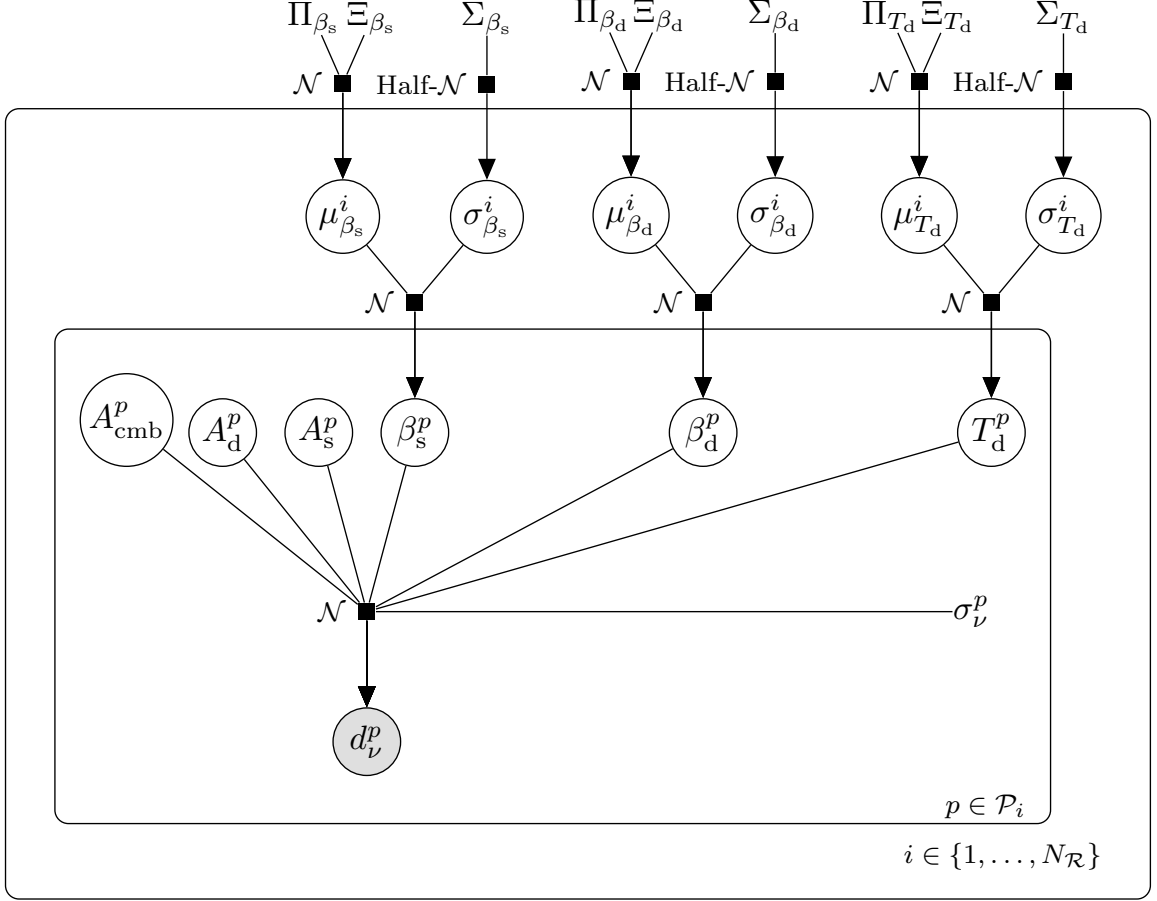


Figure 5.4: Graphical representation of the hierarchical model. Outside of the plates, we have the global constants $\{\Pi_X, \Xi_X, \Sigma_X\}$, $X = \{\beta_s, \beta_d, T_d\}$. Π_X denotes the mean of the Normal prior for the μ_X^i hyper-parameter, and Ξ_X denotes the corresponding standard deviation. Σ_X denotes the standard deviation of the Half-Normal prior for the σ_X^i hyper-parameter. At the region-level, denoted again by the larger plate, we have the set of hyper-parameters $\{\mu_X^i, \sigma_X^i\}$. The smaller plate again denotes the pixel-level, which now includes individual spectral parameters, $\{\beta_s^p, \beta_d^p, T_d^p\}$. These pixel-level spectral parameters are constrained through their conditional dependence on the connected region-level hyper-parameters.

from underlying Normal distributions, parametrized by the hyper-parameters, μ_X and σ_X , where $X \in \{\beta_s, \beta_d, T_d\}$. Each pair corresponds to the mean and standard deviation of the underlying Normal hyper-distribution for the synchrotron spectral index, the dust spectral index and the dust temperature respectively. During component separation I jointly fit for the population-level hyper-parameters and the associated pixel-level spectral parameters. The hierarchical approach allows me to model the pixel-level variations in the spectral parameters, with the hyper-distributions reducing the propensity of the model to overreact to noise, as would be the case if I had assumed total independence between pixel-level spectral parameters (Katahira, 2016).

The emission model takes the same form as in Equation 5.27, and I again assume the Gaussian likelihood in Equation 5.15. The priors for the hierarchical model are listed in Table 5.3. In analogy to the complete pooling model, I set informative priors on the spectral hyper-parameters. For the means of the hyper-distributions I set the same Normal priors as for the spectral parameters in the complete pooling model. For the standard deviations of the hyper-distributions I set Half-Normal priors, with scale parameters set to correspond to the standard deviations of the mean priors. The Half-Normal prior constrains the standard deviations to be positive, with the scale parameters chosen to encapsulate the likely degree of variation of spectral parameters in a given region. Setting informative priors on the hyper-parameters in a hierarchical model can be particularly important in ensuring the robust computational performance of the sampling algorithm. Hyper-parameters are highly correlated with the associated pixel-level parameters, and small changes in the values of the hyper-parameters can induce large changes in the target distribution. This can result in funnel-like geometries in the posterior when the data is limited i.e., a region of high density but low volume below a region of low density but high volume. The funnel regions are highly curved, which can lead to major computational difficulties during sampling, in the worst case leading to a failure in geometric ergodicity. This problem can be partly mitigated through setting informative priors that down-weight more extreme parameter values as I have done here (Gelman, 2006b; Betancourt and Girolami, 2015).

In addition to the choice of informative priors, I re-parametrize the spectral parameters by introducing the auxiliary variables,

$$\Gamma_c^p \sim \mathcal{N}(0, 1), c = \{s, d, \text{cmb}\}. \quad (5.28)$$

For the synchrotron spectral index, this can be re-expressed as,

$$\beta_s^p = \mu_{\beta_s} + \Gamma_s^p \sigma_{\beta_s}, \quad (5.29)$$

Table 5.3: Priors for parameters in the hierarchical model. For the means of the spectral hyper-distributions I assign Normal priors, and for the standard deviations of the hyper-distributions I assign Half-Normal priors. The priors for the pixel-level spectral parameters are then defined through their conditional dependence on the hyper-parameters. As with the complete pooling model, I assign flat priors to the amplitude parameters and allow them to vary from pixel to pixel. The analysis is restricted to polarization.

Θ	$p(\Theta)$
μ_{β_s}	$\mathcal{N}(\mu = -3, \sigma = 0.3)$
σ_{β_s}	Half-Normal($\sigma = 0.35$)
μ_{β_d}	$\mathcal{N}(\mu = 1.6, \sigma = 0.3)$
σ_{β_d}	Half-Normal($\sigma = 0.35$)
μ_{T_d}	$\mathcal{N}(\mu = 21 \text{ K}, \sigma = 2 \text{ K})$
σ_{T_d}	Half-Normal($\sigma = 2.5 \text{ K}$)
β_s^p	$\mathcal{N}(\mu = \mu_{\beta_s}, \sigma = \sigma_{\beta_s})$
β_d^p	$\mathcal{N}(\mu = \mu_{\beta_d}, \sigma = \sigma_{\beta_d})$
T_d^p	$\mathcal{N}(\mu = \mu_{T_d}, \sigma = \sigma_{T_d})$
$A_s^{p,\lambda}$	Unif($-\infty, \infty$)
$A_d^{p,\lambda}$	Unif($-\infty, \infty$)
$A_{\text{cmb}}^{p,\lambda}$	Unif($-\infty, \infty$)

with analogous expressions for β_d^p and T_d^p . Thus, instead of directly sampling the $\{\beta_s^p, \beta_d^p, T_d^p\}$, I instead sample a set of Gaussian latent variables and obtain the pixel-by-pixel spectral parameters through a translation and scaling with the hyper-parameters. This is known as the non-centred parametrization and has the convenient effect of reducing correlations between the hyper-parameters and the pixel-level spectral parameters. A detailed discussion of the geometrical pathologies of hierarchical models and practical approaches to their mitigation can be found in [Betancourt and Girolami \(2015\)](#).

5.4.5 Convergence Checks

Given an infinite number of samples it can be shown the the NUTS algorithm will converge on the target distribution. However, it remains important to perform a number of checks to reassure ourselves of convergence after a finite number of samples. To this end, I output a number of convergence diagnostics that are described below.

The first convergence diagnostic I output is the Gelman-Rubin statistic ([Gelman](#)

and Rubin, 1992; Brooks and Gelman, 1998). This compares the variance between multiple, independently initialized chains with the variance within each chain, and is defined as,

$$\hat{R} = \frac{\hat{V}}{\hat{W}}, \quad (5.30)$$

where \hat{V} is the between-chain variance and \hat{W} is the within-chain variance. If convergence has been achieved the between-chain and within-chain variance will be equal. In reality I apply the threshold, $\hat{R} \leq 1.1$ as a check that the parameter chains satisfy the necessary geometric ergodicity conditions.

I also output the number of effective samples in each chain, n_{eff} (Geyer, 1992; Brooks et al., 2011). When sampling from a target distribution using some MCMC algorithm, we may draw a total of N samples, but these samples are not totally independent. The effective sample size provides a measure of the number of independent samples in a chain, defined as,

$$n_{\text{eff}} \equiv \frac{N}{\sum_{t=-\infty}^{\infty} \rho_t} = \frac{N}{1 + 2 \sum_{t=1}^{\infty} \rho_t}, \quad (5.31)$$

where ρ_t is the auto-correlation within a chain at a lag t . For a chain with joint distribution $\pi(\Theta)$ with mean μ and standard deviation σ , this is given by

$$\rho_t = \frac{1}{\sigma^2} \int_{\Theta} (\Theta^n - \mu)(\Theta^{n+t} - \mu) \pi(\Theta) d\Theta. \quad (5.32)$$

Details on the estimation of the auto-correlation can be found in Stan Development Team (2012)². The appropriate number of effective samples to be able to properly capture the target distribution is to some extent a question of judgement. However, in Kruschke (2011) a threshold of ~ 1000 effective samples is proposed to be confident in expectations calculated with parameter chains. As such, I adopt this as a confidence threshold for the sampling output.

Finally, I also output warnings when divergences occur during sampling. A divergence takes place when the sampler encounters a region of the target distribution where the curvature is too high to be resolved given the tuned step-size. In practice, divergences are detected when the value of the Hamiltonian diverges from its initial value when simulating trajectories through parameter space. This is significant in that divergences can mean that the conditions for geometric ergodicity are not met, and therefore using the resultant chains to construct statistical estimators can lead

²It is worth noting that, if a chain is estimated to have a negative auto-correlation between samples, one can obtain $n_{\text{eff}} > N$.

to biased inferences (Betancourt and Girolami, 2015; Betancourt, 2016, 2017). The sensitivity to divergences is one of the core advantages of HMC algorithms when sampling from complex high-dimensional geometries. Using non-gradient based samplers it can be challenging to diagnose these situations where geometric ergodicity is not satisfied, whereas HMC comes with a sensitive, in-built diagnostic.

5.5 Algorithm Validation


I validate the component separation algorithm’s performance against the simulated experimental observations described in Section 5.2. For ease of discussion, I assign the following labels to the three validation sets (i.e., simulation and modelling runs):

- CP(L): The *LiteBIRD* only analysis, fitting the complete pooling model described in Section 5.4.3.
- CP(LC): The C-BASS and *LiteBIRD* analysis, fitting the complete pooling model.
- H(LC): The C-BASS and *LiteBIRD* analysis, fitting the hierarchical model described in Section 5.4.4.


In all three validation sets I used regions defined on $N_{\text{SIDE}}=64$ maps, as described in Section 5.3, for the spectral modelling. I do not fit the hierarchical model to the simulation set consisting of just *LiteBIRD* observations. It was found to be very challenging to control for the occurrence of divergences when fitting the hierarchical model to *LiteBIRD*-only observations, leaving the convergence properties of the resulting MCMC chains suspect. These problems around controlling divergences can be understood when we consider the lack of low-frequency channels in *LiteBIRD*. Given limited available information to constrain synchrotron spectral parameters, the posterior geometry for the hierarchical model becomes extremely difficult to sample. Applying the hierarchical model to *LiteBIRD*-only observations likely requires a careful study of prior choice for model hyper-parameters and extended tuning phases to help mitigate the occurrence of divergences.

For each validation set, I evaluate the normalized deviations of the recovered parameters. For some parameter Θ , the normalized deviations are defined as,


$$\eta_\lambda = \frac{\Theta^{\text{in},\lambda} - \Theta^{\text{out},\lambda}}{\sigma_\Theta^\lambda}, \quad \lambda = \{Q, U\}, \quad (5.33)$$




(a) η : CMB amplitude




(b) η : Synchrotron amplitude




(c) η : Dust amplitude



(d) η : Synchrotron spectral index



(e) η : Dust spectral index



(f) η : Dust temperature

Figure 5.5: Histograms of the normalized deviations for the recovered parameter maps. Alongside these histograms I plot the standard Gaussian, $\mathcal{N}(0, 1)$. Mismatch between the normalized deviation histograms and $\mathcal{N}(0, 1)$ indicate departures from Gaussianity in the marginal parameter posteriors. In itself, this is not surprising, given the posterior distributions include contributions from non-Gaussian priors, and in the case of the hierarchical model contains complex correlations between hyper-parameters and pixel-level parameters.

Parameter	Validation Set	Med(η_Q)	Med(η_U)	MAD(η_Q)	MAD(η_U)
A_{cmb}	CP(L)	0.21	-0.10	2.60	2.18
A_{cmb}	CP(LC)	0.12	-0.02	2.66	2.25
A_{cmb}	H(LC)	-0.03	2×10^{-3}	1.30	1.29
A_s	CP(L)	2×10^{-3}	-0.02	2.17	1.98
A_s	CP(LC)	0.08	-2×10^{-3}	2.14	2.00
A_s	H(LC)	0.01	1×10^{-3}	1.89	1.88
A_d	CP(L)	0.01	0.01	0.74	0.62
A_d	CP(LC)	0.02	0.01	0.81	0.68
A_d	H(LC)	0.03	2×10^{-3}	0.45	0.43
β_s	CP(L)	-0.44	...	7.05	...
β_s	CP(LC)	-1.03	...	16.8	...
β_s	H(LC)	0.09	...	1.46	...
β_d	CP(L)	0.09	...	4.18	...
β_d	CP(LC)	-0.35	...	4.96	...
β_d	H(LC)	-0.35	...	1.42	...
T_d	CP(L)	0.61	...	4.33	...
T_d	CP(LC)	1.23	...	5.11	...
T_d	H(LC)	0.63	...	1.96	...

Table 5.4: In columns 2 and 3 I state the medians of the normalized deviations for the recovered model parameters in each validation set, and in columns 4 and 5 I state the corresponding MAD values. For spectral parameters the normalized deviations in Q and U are identical. As such, I only state results for spectral parameters under their corresponding Q columns.

Parameter	Validation Set	Med(\mathcal{R}_Q)	Med(\mathcal{R}_U)	MAD(\mathcal{R}_Q)	MAD(\mathcal{R}_U)
A_{cmb}	CP(L)	$-6 \times 10^{-3} \mu\text{K}$	$3 \times 10^{-3} \mu\text{K}$	$0.07 \mu\text{K}$	$0.06 \mu\text{K}$
A_{cmb}	CP(LC)	$-3 \times 10^{-3} \mu\text{K}$	$4 \times 10^{-4} \mu\text{K}$	$0.07 \mu\text{K}$	$0.05 \mu\text{K}$
A_{cmb}	H(LC)	$1 \times 10^{-3} \mu\text{K}$	$-6 \times 10^{-5} \mu\text{K}$	$0.05 \mu\text{K}$	$0.04 \mu\text{K}$
A_s	CP(L)	$-4 \times 10^{-4} \mu\text{K}$	$3 \times 10^{-3} \mu\text{K}$	$0.4 \mu\text{K}$	$0.4 \mu\text{K}$
A_s	CP(LC)	$-2.8 \mu\text{K}$	$0.09 \mu\text{K}$	$80.5 \mu\text{K}$	$75.2 \mu\text{K}$
A_s	H(LC)	$-0.5 \mu\text{K}$	$-0.03 \mu\text{K}$	$74.1 \mu\text{K}$	$72.8 \mu\text{K}$
A_d	CP(L)	$-8 \times 10^{-5} \mu\text{K}$	$-1 \times 10^{-4} \mu\text{K}$	$5 \times 10^{-3} \mu\text{K}$	$4 \times 10^{-3} \mu\text{K}$
A_d	CP(LC)	$-1 \times 10^{-4} \mu\text{K}$	$-7 \times 10^{-5} \mu\text{K}$	$6 \times 10^{-3} \mu\text{K}$	$5 \times 10^{-3} \mu\text{K}$
A_d	H(LC)	$2 \times 10^{-4} \mu\text{K}$	$-1 \times 10^{-5} \mu\text{K}$	$3 \times 10^{-3} \mu\text{K}$	$3 \times 10^{-3} \mu\text{K}$
β_s	CP(L)	6×10^{-3}	...	0.1	...
β_s	CP(LC)	3×10^{-3}	...	0.03	...
β_s	H(LC)	-2×10^{-3}	...	0.03	...
β_d	CP(L)	1×10^{-3}	...	0.04	...
β_d	CP(LC)	3×10^{-3}	...	0.04	...
β_d	H(LC)	6×10^{-3}	...	0.02	...
T_d	CP(L)	-0.2 K	...	1.5 K	...
T_d	CP(LC)	-0.4 K	...	1.3 K	...
T_d	H(LC)	-0.3 K	...	1.0 K	...

Table 5.5: In columns 3 and 4 I state the medians of the residuals for the given parameter and validation set, and in columns 5 and 6 I state the corresponding MAD values. Spectral parameters are identical in Q and U . As such, I state the median and MAD values under the Q columns only. Note that, in comparing synchrotron amplitude residuals, the residuals for the CP(L) set are calculated at a reference frequency of 40 GHz, compared to 5 GHz for the CP(LC) and H(LC) sets. This results in the seemingly low level of synchrotron amplitude residuals, given the lower level of synchrotron emission at 40 GHz.

where $\Theta^{\text{out},Q/U}$ is the output parameter Q/U map, $\Theta^{\text{in},Q/U}$ is the input parameter Q/U map, and $\sigma_{\Theta}^{Q/U}$ is the corresponding Q/U standard deviation map. For spectral parameters the Q and U normalized deviations will be identical. If the observed $\Theta^{\text{out},Q/U}$ are drawn from a Gaussian distribution with mean given by $\Theta^{\text{in},Q/U}$ and standard deviation given by $\sigma_{\Theta}^{Q/U}$, the normalized deviations should be distributed as a standard Gaussian, $\mathcal{N}(0, 1)$.

Histograms of the normalized deviations for the model parameters recovered from each validation set are shown in Fig. 5.5. In Table 5.4 I state the median of the normalized deviations for each parameter and validation set, along with the corresponding MAD values. The distributions of the normalized deviations for each parameter are discussed in the relevant subsections outlined below.

The outline of the remainder of this section is as follows: in Section 5.5.1 I discuss the CMB amplitude output, in Section 5.5.2 I present the synchrotron and dust amplitude constraints, in Section 5.5.3 I discuss constraints on the synchrotron spectral index, and in Section 5.5.4 I show results for the dust spectral parameters. In Fig. 5.6 I show maps of the recovered CMB amplitudes and effective sample size for each validation set. In Fig. 5.7 I show the corresponding CMB amplitude residuals. In Fig. 5.10 I show the recovered synchrotron and dust amplitude maps, and in Fig. 5.11 I show the corresponding residual maps. In Fig. 5.12 I show the recovered synchrotron and dust spectral parameter maps, and in Fig. 5.13 I show the corresponding residual maps. In Table 5.5 I state the median residual level for each parameter and validation set, along with the corresponding MAD values.

5.5.1 CMB Amplitude

The primary output from the NUTS component separation are CMB amplitude maps in Q and U . This consists of a set of maps corresponding to individual posterior samples, along with the summary maps of the mean and standard deviation of the amplitude maps. In Fig. 5.6 I show the mean CMB Q and U amplitude maps obtained for the three validation sets, along with the associated maps of the effective sample size. When using a complete pooling model obvious artefacts can be seen in the recovered CMB amplitude maps near the Galactic plane. This is to be expected, given the bright diffuse emission in these regions makes the extraction of weak CMB signals extremely challenging. When using a hierarchical model these artefacts are no longer present. By accounting for the real variation in spectral parameters in each region, whilst constraining this variation through the fitted hyper-distributions, the

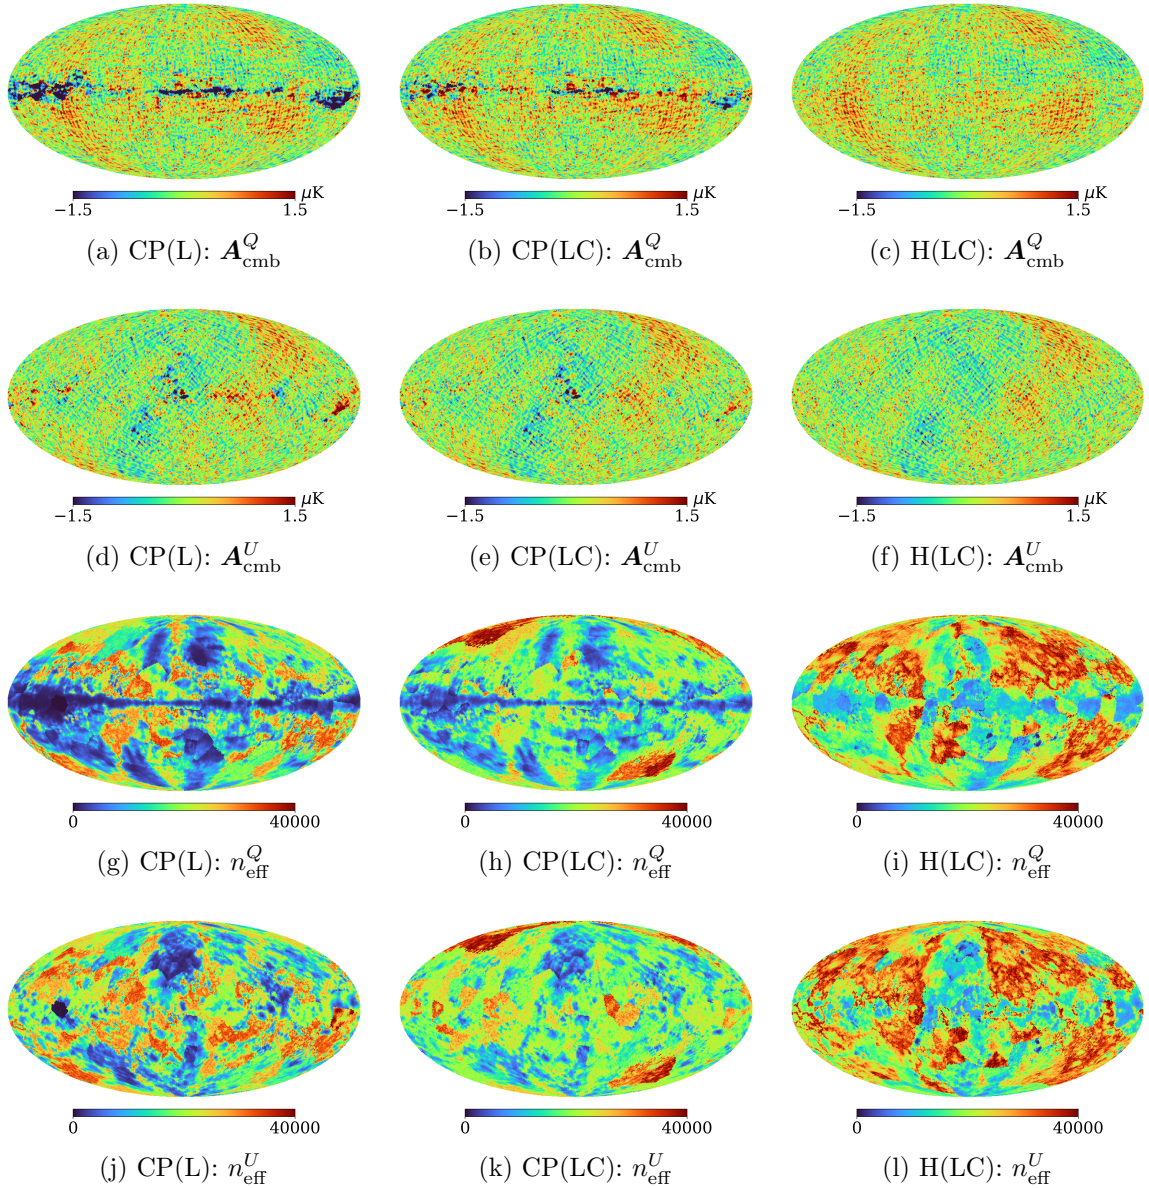


Figure 5.6: Output CMB amplitude maps obtained with the NUTS component separation algorithm. In panels (a) to (f) I show the output $\mathbf{A}_{\text{cmb}}^{Q/U}$ maps for the various validation sets, and in panels (g) to (l) I show the corresponding maps of the effective sample size. Obvious artefacts can be seen in the output $\mathbf{A}_{\text{cmb}}^{Q/U}$ maps when using a complete pooling model. This is the result of extremely bright foreground emission dominating over the CMB signal, making the extraction of the CMB signal very challenging. These artefacts are well traced by the effective sample size, being most apparent in bright regions close to the Galactic plane and North Polar Spur. By using a hierarchical model I am able to remove these obvious artefacts from the recovered CMB maps.




Figure 5.7: CMB amplitude residual maps for the three validation sets. When using the complete pooling model, strong residuals are apparent close to the Galactic plane and along bright diffuse features such as the North Polar Spur. These residuals are greatly reduced when using the hierarchical model. Large-scale residuals present in the CMB amplitude maps when using the complete pooling model can ultimately translate into significant biases on recovered cosmological parameters.

hierarchical model is able to achieve a more accurate foreground removal and thereby remove the biases apparent from assuming constant spectral parameters. Regions of the CMB amplitude maps containing component separation artefacts are well traced by the effective sample size. In regions of low effective sample size ($n_{\text{eff}} \lesssim 1000$), the parameter chains exhibit a high degree of auto-correlation, indicative of the sampler struggling to draw independent posterior samples. Thresholding maps of the effective sample size can be used to construct confidence masks for the CMB amplitude maps, which can be used in combination with standard Galactic emission masks.

In Fig. 5.5a I show histograms of the normalized deviations of the $\mathbf{A}_{\text{cmb}}^{Q/U}$ maps. For the validation sets studied here the distributions of the normalized deviations are slightly wider than the standard Gaussian. This means the uncertainties reported by the CMB amplitude standard deviation maps underestimate the errors on the recovered CMB amplitudes, and the CMB amplitude posterior shows slight departures from Gaussianity. The complete pooling model results in small biases away from zero in the median of the normalized deviations, which are not present with the hierarchical model.

I perform power spectrum estimation using the NAMASTER library (Alonso et al., 2019). Note that for the regime considered here $B \ll E$, meaning B -mode purification is very important in order to avoid the variance of the B -mode estimator being dominated by the E -mode cosmic variance. I generate the Galactic emission mask following the procedure in Remazeilles et al. (2018), with a common mask being used to enable direct comparison between the recovered power spectra. I extrapolated 10° smoothed 5 GHz and 402 GHz polarized intensity maps to 70 GHz, and applied a 5σ threshold against the standard deviation of the 10° smoothed CMB polarized intensity map at 70 GHz. The 5 GHz map was extrapolated using a constant spectral index of $\beta_s = -3$ and the 402 GHz map was extrapolated using an MBB SED, setting $\beta_d = 1.6$ and $T_d = 19.4$ K. I further mask all pixels with $n_{\text{eff}} \lesssim 1000$ in either the Q or U maps, although most of these pixels are already contained within the Galactic emission masks. The Galactic emission mask produced by masking all pixels below the 5σ threshold excludes approximately 60% of the sky.

In Fig. 5.8 I show the E and B -mode power spectra derived for the three validation sets. For the purposes of power spectrum estimation, I perform component separation on two splits of the input data with differing noise realisations. I then evaluate the cross-spectra of the recovered CMB maps from the two data-splits, avoiding complications from noise bias. I show power spectra, covering multipoles $2 \leq \ell < 180$, using a bin width of $\Delta\ell = 10$ (excluding the first bin, which covers multipoles $2 \leq \ell < 10$).

For multipoles $\ell \lesssim 30$ it was found that the purified pseudo- C_ℓ estimator no longer recovered accurate point estimates for the bandpowers. At these large angular scales, the sky mask results in the pseudo- C_ℓ estimator becoming sub-optimal. The pseudo- C_ℓ estimator being unbiased provides a guarantee that the expectation of the estimator will converge on the true value in the asymptotic regime. However, for low multipoles on a cut-sky we are far from this asymptotic regime, which can lead to poor point estimates. For these multipoles, one can obtain better power spectrum estimates by directly sampling the C_ℓ from the joint distribution, $P(C_\ell, \mathbf{A}_{\text{cmb}}|\mathbf{d})$, as is performed in the COMMANDER component separation code (Wandelt et al., 2004; Eriksen et al., 2004, 2008; Taylor et al., 2008). As noted in Alsing et al. (2016), such Bayesian sampling techniques have the advantage of inducing probabilistic information on the CMB fields inside the mask, and allow for the imposition of priors on the power spectra that can help in regularizing corresponding parameter estimates. I leave the implementation of this additional power spectrum estimation technique to future work.




Figure 5.8: CMB E -mode and B -mode power spectra for the three validation sets, covering the multipole range, $2 \leq \ell < 180$. Power spectra are binned, using a bin width of $\Delta\ell = 10$ (excluding the first bin, which contains multipoles $2 \leq \ell < 10$). I plot the power spectra for the three validation sets together, slightly offset from one another for display purposes. The central point within each triple is at the effective multipole value for that bin. We can see that the E -mode power spectrum is recovered with high accuracy for $\ell > 30$, demonstrating the overall fidelity of the algorithm. For the B -mode power spectrum we can see the results obtained using a complete pooling model contain significant large-scale foreground residuals, that are not present when using the hierarchical model. For $\ell \lesssim 30$ the sky mask results in the pseudo- C_ℓ estimator becoming sub-optimal, which is particularly apparent in the $10 \leq \ell < 20$ bin, where the point estimates for the E -mode bandpowers are negative. I plot a dotted vertical line at $\ell = 30$ to separate these angular scales. For $\ell \gtrsim 150$ the beam correction results in higher variance pseudo- C_ℓ estimates. This beam correction, along with the limiting pixel resolution at $N_{\text{SIDE}}=64$, results in bandpower estimates becoming unreliable for $\ell \gtrsim 180$.

For multipoles $\ell \gtrsim 150$, the 70 arcmin beam correction resulted in higher variance pseudo- C_ℓ estimates. At these small scales the beam correction begins to inflate the noise present in the CMB amplitude maps. These effects, in combination with the pixel resolution at $N_{\text{SIDE}}=64$, result in bandpower estimates becoming unreliable for $\ell \gtrsim 180$.

To estimate the bandpower covariances, noise realizations were obtained by taking the difference between individual CMB amplitude posterior samples and the mean CMB amplitude map for the respective data-splits, i.e.,

$$\delta \mathbf{A}_{\text{cmb}}^{k,j,\lambda} = \mathbf{A}_{\text{cmb}}^{k,j,\lambda} - \langle \mathbf{A}_{\text{cmb}}^{j,\lambda} \rangle, \quad \lambda = \{Q, U\}, \quad (5.34)$$

where k denotes the k^{th} posterior sample and $j \in \{1, 2\}$ denotes the relevant data-split. The covariance matrix was then estimated by calculating the power spectra of 1000 signal plus noise simulations. For the signal simulations I generated 1000 realizations of the theoretical input CMB power spectrum to account for cosmic variance.

The E -mode power spectrum is obtained with high accuracy for $\ell \geq 30$, confirming the overall fidelity of the component separation algorithm. The recovery of the B -mode power spectrum is more challenging, given this signal is significantly weaker than the E -mode signal and potentially sub-dominant to foregrounds over most of the sky, at all frequencies. It can be seen that when using a complete pooling model the recovered B -mode power spectrum contains large-scale foreground residuals, biasing the power spectrum high. In real experimental applications one could attempt to mitigate this by applying even more aggressive Galactic emission masks, although this comes at the cost of increasing the uncertainty in the recovered power spectra. For multipoles $\ell < 30$ we can see that the apparent residuals in the B -mode power spectrum are reduced when using the hierarchical model. However, as discussed above, the pseudo- C_ℓ estimates at these multipoles become sub-optimal, demonstrated clearly here by the negative point estimates obtained for the E -mode bandpowers in the $10 \leq \ell < 20$ bin.

To quantify the impact of foreground residuals in the B -mode power spectrum, we can study the tensor-to-scalar ratio constraints that would be obtained from these power spectra for $\ell > 30$. To do this I approximate the likelihood for the CMB power spectra as a multivariate Gaussian,

$$-2 \ln \mathcal{L} = \text{const.} + \sum_{\ell_b, \ell'_b} \left(\tilde{C}_{\ell_b} - \tilde{C}_{\ell_b}^{\text{th}} \right) (\Sigma^{-1})_{\ell_b, \ell'_b} \left(\tilde{C}_{\ell'_b} - \tilde{C}_{\ell'_b}^{\text{th}} \right). \quad (5.35)$$

The \tilde{C}_{ℓ_b} are the binned power spectra or bandpowers, with $\tilde{C}_{\ell_b}^{\text{th}}$ being the corresponding theoretical bandpowers, and Σ is the bandpower covariance matrix. The sum

here runs over the effective ℓ values for each bin. It is worth noting that, in general, the CMB likelihood is non-Gaussian. However, for the higher multipoles I consider here the power spectrum estimates are formed by averaging over the individual $a_{\ell m}$'s corresponding to a given multipole, justifying the use of the Gaussian approximation through the central limit theorem.

I parametrize the theoretical power spectrum as,

$$C_{\ell}^{\text{th}} = \frac{r}{0.01} C_{\ell}^{BB}(r = 0.01) + A_L C_{\ell}^{\text{lens}}, \quad (5.36)$$

where A_L is the lensing amplitude, $C_{\ell}^{BB}(r = 0.01)$ is a fiducial primordial B -mode power spectrum corresponding to $r = 0.01$, and C_{ℓ}^{lens} is the fiducial lensing B -mode power spectrum. In a general B -mode analysis I would fit jointly for r and A_L . However, it is challenging to constrain A_L solely through the B -mode power spectrum here. In a realistic experiment, tight constraints can be put on the lensing B -mode through analysis of the E -mode, TE correlations and the lensing potential power spectrum. Combined with constraints from external data sets and delensing, one can expect to be able to place a tight prior on A_L . For the sake of simplicity here, I fix A_L to the input value of 1.

Sampling from this likelihood using PYMC3, I obtain the constraints on r shown in Fig. 5.9. Using the complete pooling model I recover biased estimates of the tensor-to-scalar ratio, obtaining $r = (12.9 \pm 1.4) \times 10^{-3}$ for the CP(L) set and $r = (9.0 \pm 1.1) \times 10^{-3}$ for the CP(LC) set. This can be understood when we consider the large residuals and artefacts present in the recovered CMB when using a complete pooling model. By comparison, the bias is effectively removed for the H(LC) set, obtaining $r = (5.2 \pm 1.0) \times 10^{-3}$. It is worth noting that, despite the increased degrees of freedom in using a hierarchical model, the uncertainties obtained for the H(LC) set are smaller than those for the CP(LC) set. By adopting a multi-level structure for the spectral parameters in each region we prevent the model from overreacting to noise, whilst still capturing the spatial variation in spectral parameters. This in turn removes many of the foreground residuals present in the CMB amplitude maps obtained with the complete pooling model. These residuals result in increased uncertainties and biases on r , caused by the misspecified power spectrum model in the presence of foreground residuals.

As an aside, I do caution against interpreting these results as direct forecasts for the proposed *LiteBIRD* experiment (and/or some combination with additional experiments). *LiteBIRD* will also obtain constraints on lower multipoles around the reionization peak in the B -mode power spectrum, which I have not considered




Figure 5.9: Tensor-to-scalar ratio constraints obtained for the three validation sets, using derived power spectra over the range $30 \leq \ell < 180$ and assuming a multivariate Gaussian likelihood. For the CP(L) set I find $r = (12.9 \pm 1.4) \times 10^{-3}$, for the CP(LC) set I find $r = (9.0 \pm 1.1) \times 10^{-3}$, and for the H(LC) set I find $r = (5.2 \pm 1.0) \times 10^{-3}$. With the addition of C-BASS we see a reduction in the bias on r with the complete pooling model, although a $\sim 3.6\sigma$ bias still remains for the CP(LC) set. This bias is removed when using a hierarchical model. The uncertainty on the measured tensor-to-scalar ratio is smaller for the hierarchical model. The hyper-distributions over foreground spectral parameters allow the model to fit spatial variations without increasing parameter uncertainties. Further, foreground residuals that result from the complete pooling model mean the power spectrum model is misspecified, resulting in increased biases and uncertainties on cosmological parameters.

here due to the sub-optimal pseudo- C_ℓ estimator used, increasing the sensitivity of any tensor-to-scalar ratio measurement. Further, for simplicity in this validation analysis I smoothed all channels to the 70 arcmin resolution of the lowest frequency *LiteBIRD* channel. This is likely a somewhat pessimistic approach. However, it is beyond the scope of this work to study the impact of the low resolution of the low-frequency *LiteBIRD* channels on the ability to recover the CMB power spectra at higher multipoles. I have also not considered the effect of mis-modelling foreground SEDs or the impact of experimental systematics, both of which would significantly complicate any B -mode measurements.

As discussed previously, given the frequency coverage considered here I already struggle to constrain spatial variations in foreground spectral parameters for these simple models. Fitting more complex models will require additional low and high-frequency data. I leave the impact of mis-modelling complex foreground SEDs to

more detailed forecasting analyses, focusing here on validating the ability of Bayesian hierarchical modelling to reduce biases in recovered CMB estimates without inflating parameter uncertainties. It is worth noting that I consider the performance of the hierarchical model here for an experimental frequency coverage much broader than that of next-generation ground-based experiments such as the Simons Observatory (Simons Observatory Collaboration et al., 2019) and CMB-S4 (Abazajian et al., 2016). Given the reduced frequency coverage of these experiments, fitting a full hierarchical model will be significantly more challenging, and will likely require a careful study of the prior choice for model hyper-parameters, and an exploration of model re-parametrizations.

5.5.2 Synchrotron and Dust Amplitudes

In Fig. 5.10 I show the dust and synchrotron amplitude maps. Dust amplitude maps are shown at the reference frequency of $\nu_0 = 402$ GHz for all three validation sets. Synchrotron amplitude maps are shown at a reference frequency of $\nu_0 = 40$ GHz for the CP(L) validation set, and at $\nu_0 = 5$ GHz for the CP(LC) and H(LC) validation sets. In all cases we can see the recovered component amplitude maps trace the input component amplitude maps well. However, this is to some extent a result of the choice of reference frequency, with the amplitudes being constrained by the pixel values at those frequencies. The overall level of residuals in the synchrotron amplitude maps is reduced by ~ 5 per cent for the H(LC) set compared to the CP(LC) set, and the typical residuals in the dust amplitude maps are reduced by ~ 40 per cent. These reductions were estimated using the MAD values of the residuals. Whilst the absolute value of these reductions is small at synchrotron and dust frequencies, they ultimately propagate through to significant biases in the recovered CMB maps, as seen in the tensor-to-scalar ratio estimates.

In Fig. 5.5b I show histograms of the normalized deviations for the synchrotron amplitude, and in Fig. 5.5c I show histograms of the normalized deviations for the dust amplitude. For the synchrotron amplitude I obtain similar distributions across all three validation sets. The normalized deviations for the synchrotron amplitudes are wider than the standard Gaussian, indicating the standard deviations of the posterior samples underestimate the uncertainties on the synchrotron amplitudes. Conversely, for the dust amplitudes the normalized deviations are narrower than the standard Gaussian.

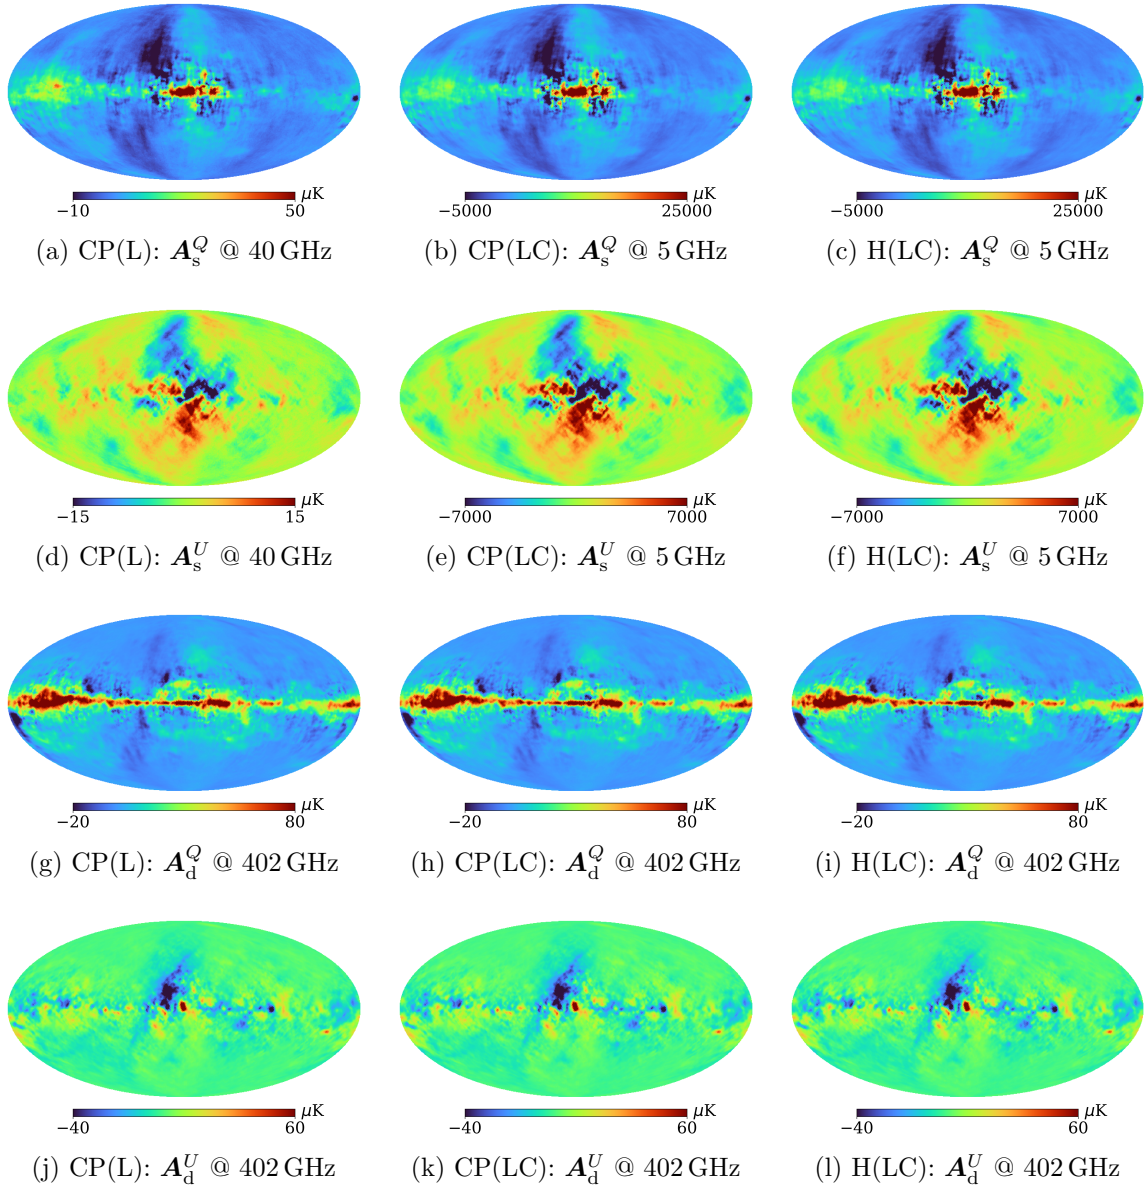


Figure 5.10: Synchrotron and dust amplitude maps obtained for the three validation sets. In panels (a) to (f) I show the synchrotron amplitude maps, shown at the reference frequency of $\nu_0 = 40$ GHz for the CP(L) validation set, and $\nu_0 = 5$ GHz for the CP(LC) and H(LC) validation sets. In panels (g) to (l) I show the dust amplitude maps, displayed at the reference frequency of $\nu_0 = 402$ GHz. In all three cases the component separation algorithm recovers accurate estimates of the foreground amplitude maps at their reference frequencies. This is to be expected, and is to some extent a result of the choice of synchrotron and dust reference frequencies. The typical residuals in synchrotron amplitude maps are $\sim 5\%$ lower for the H(LC) set compared to the CP(LC) set. Similarly, the dust amplitude residuals are $\sim 40\%$ lower for the H(LC) set compared to the CP(LC) set.

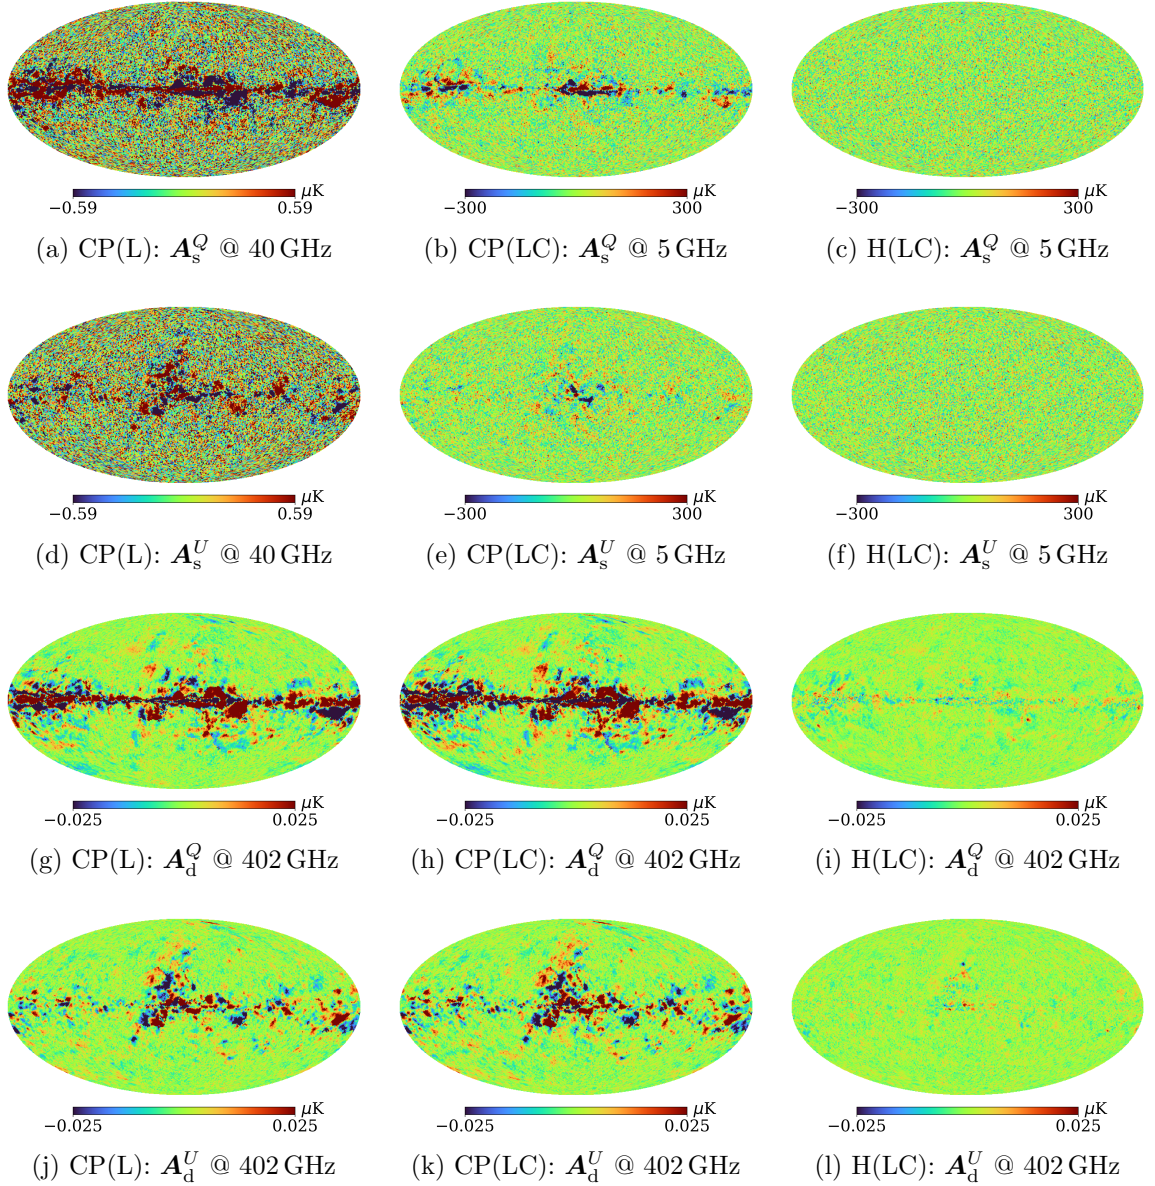


Figure 5.11: In panels (a) to (f) I show the synchrotron amplitude residual maps, and in panels (g) to (l) I show the dust amplitude residual maps. In displaying the synchrotron amplitude residuals for the CP(L) set, I have scaled the $\pm 300 \mu\text{K}$ axis limits used for the CP(LC) and H(LC) sets from 5 GHz to 40 GHz, using a spectral index of -3 . For both the synchrotron and dust amplitudes we can see that the residuals, particularly in regions of bright diffuse emission, are greatly reduced when using a hierarchical model. Whilst the level of these residuals is small relative to the total synchrotron and dust emission at their reference frequencies, these residuals ultimately propagate through to significant foreground residuals in the recovered CMB amplitude maps.

5.5.3 Synchrotron Spectral Parameters

In Fig. 5.12a, 5.12b and 5.12c, I show the synchrotron spectral indices obtained for the three validation sets. In the case of the CP(L) set, I struggle to place accurate constraints on the synchrotron spectral index. This is to be expected, given the lack of low-frequency channels below 40 GHz, and is consistent with expectations from the single-pixel component separation analysis presented in Jew et al. (2019). For the CP(LC) validation set, I am able to place improved constraints on the synchrotron spectral index, with the variations in the synchrotron spectral index from region-to-region tracing the variations in the input spectral index map shown in Fig. 5.1.

When using a hierarchical model I am able to more finely model variations in the synchrotron spectral index across the sky. As can be seen in Fig. 5.12c, the spectral index in regions of high SNR traces the same variations seen in the input synchrotron spectral index map. In regions of low SNR, away from the Galactic plane, we do see additional features not present in the input β_s map. This is a result of the noisier estimates of the synchrotron spectral index obtained in these regions. These noisy variations are constrained by the hyper-distribution, which penalises individual estimates of β_s being too far from the population mean, μ_{β_s} . Further, even in these noisier regions of the sky we can see that many of the large-scale variations in the spectral index are still traced by the individual spectral indices. This is a well known property of hierarchical models, known as posterior shrinkage, and is one of the main advantages of adopting the hierarchical approach i.e., we obtain improved point estimates of the latent variables (Katahira, 2016).

For the simulations I have considered here, the typical residuals in the synchrotron spectral index maps are similar between the CP(LC) and H(LC) validation sets. This is partly a result of the input synchrotron spectral index map being highly idealized, lacking in small-scale features (Krachmalnicoff et al., 2018). Improvements can also likely be obtained for the hierarchical model by refining the region definition, such that we define larger regions in areas of low SNR, increasing the smoothing effect of the hyper-distributions.

In Fig. 5.5d I show histograms of the normalized deviations for the synchrotron spectral index. When using a complete pooling model, the distributions of the normalized deviations are significantly wider than the standard Gaussian. This is particularly apparent for the CP(LC) set. With the addition of C-BASS, there is greater constraining power on the average spectral index across each sky region. However, normalized deviations are evaluated at the pixel-level. The uncertainties obtained for

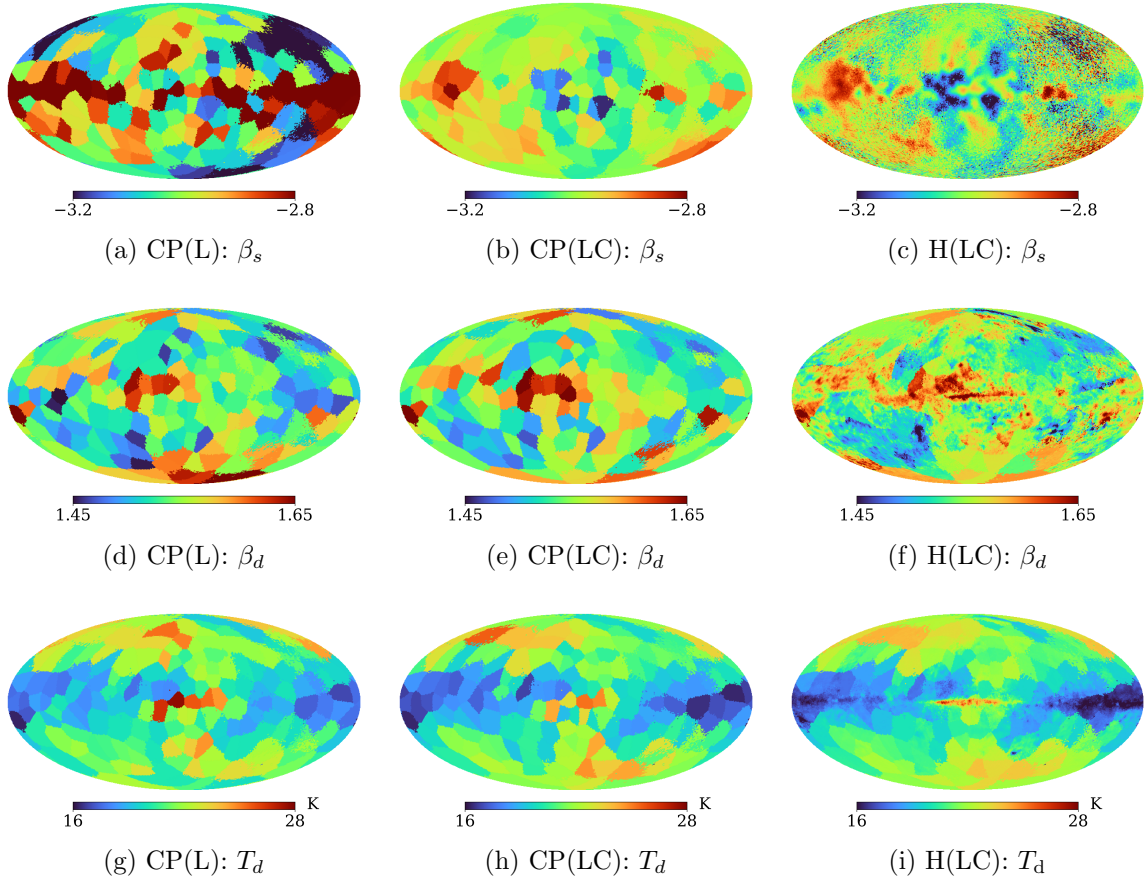


Figure 5.12: Spectral parameter maps obtained for the three validation sets. Panels (a) to (c): Recovered synchrotron spectral index maps. Panels (d) to (f): Recovered dust spectral index maps. Panels (g) to (i): Recovered dust temperature maps. For the CP(L) set, I struggle to accurately constrain the synchrotron spectral index, due to the absence of low-frequency channels below 40 GHz. With the addition of a 5 GHz C-BASS channel for the CP(LC) and H(LC) sets, I am able to place improved constraints on the synchrotron spectral index. Given the frequency coverage of *LiteBIRD* I struggle to constrain dust spectral parameters for all three validation sets, reflected in low effective sample sizes for the dust spectral parameters. The typical residuals for the synchrotron spectral index maps are comparable between the CP(LC) and H(LC) sets. However, the performance of the CP(LC) set is likely exaggerated by the lack of small-scale features in the input synchrotron spectral index map. The typical residuals for the dust spectral index maps are reduced by ~ 50 per cent for the H(LC) set compared to the CP(LC) set. The typical residuals in the recovered dust temperature maps are ~ 25 per cent lower for the H(LC) set compared to the CP(LC) set. The simulated dust spectral parameter maps contain additional small-scale features compared to the synchrotron spectral index map.




Figure 5.13: Spectral parameter residual maps obtained for each validation set. Panels (a) to (c): Synchrotron spectral index residual maps. Panels (d) to (f): Dust spectral index residual maps. Panels (g) to (i): Dust temperature residual maps. For the CP(L) validation set I find significant residuals for the synchrotron spectral index over much of the sky. Due to the lack of low-frequency channels below 40 GHz, *LiteBIRD* data alone is unable to properly constrain the synchrotron spectral index. In comparing the residuals for the CP(LC) and H(LC) sets, we can see that residuals are reduced for the H(LC) set in regions of bright diffuse emission. However, there are regions of low SNR away from the Galactic plane where the residuals for the H(LC) set are noticeably larger. The fitted hyper-distributions in this case do help to reduce the propensity for the model to overreact to noise. Improved point estimates could likely be obtained here by defining larger regions in areas of low SNR. It is also worth noting that the performance of the complete pooling model is likely exaggerated by the lack of small-scale features in the simulated synchrotron spectral index map. For dust spectral parameters, I obtain lower typical residuals for the H(LC) set due to modelling real spatial variations in the parameters. However, given a maximum observed frequency of 402 GHz, I still struggle to place constraints on dust spectral parameters.

the average spectral index in a given region will not be representative of the pixel-level uncertainties. For the range of values shown in Fig. 5.5d, the CP(L) set appears to show a positive bias. However, the median is driven to -0.44 by a set of large magnitude, negative normalized deviations. Given the lack of low-frequency channels below $\nu = 40$ GHz, *LiteBIRD* data alone is unable to properly constrain synchrotron spectral parameters.

5.5.4 Dust Spectral Parameters

Most of the constraining power for dust spectral parameters comes from high-frequency channels i.e., $\nu \gtrsim 100$ GHz. In the three validation sets these remained identical, being the high-frequency *LiteBIRD* channels. In Jew et al. (2019) it was found that, given a *LiteBIRD*-like frequency coverage, it is difficult to constrain dust spectral parameters. Indeed, this was the case for the analysis here, where informative priors were needed on the dust spectral parameters. The difficulty in constraining dust spectral parameters is reflected in their low effective sample size, which was typically $\lesssim 1000$ over much of the sky in all three validation sets.

In Fig. 5.12d, 5.12e and 5.12f I show the recovered dust spectral index maps, and in Fig. 5.12g, 5.12h and 5.12i I show the recovered dust temperature maps. In using a hierarchical model for the dust spectral parameters, I am able to fit for variations in the spectral parameters in regions of high SNR close to the Galactic plane. Away from the Galactic plane, the individual variations become much smaller in each region, with the resulting maps of β_d and T_d very obviously tracing out the crude structure of the regions used in the component separation. Given the limited frequency coverage, there is simply not enough information to constrain the low-level variations of the dust spectral parameters in each region. In this case the marginal distributions for σ_{β_d} and σ_{T_d} have a large fraction of their probability mass close to zero, constraining individual spectral parameters to be very close to their population means, μ_{β_d} and μ_{T_d} . The typical residuals for the dust spectral index maps are reduced by ~ 50 per cent for the H(LC) set compared to the CP(LC) set. The typical residuals in the recovered dust temperature maps are ~ 25 per cent lower for the H(LC) set compared to the CP(LC) set. By allowing for some of the low-level variations in the dust spectral parameters I am able to obtain smaller residuals in the recovered parameter maps, propagating through to reduced biases in the CMB amplitude estimates.

In Fig. 5.5e I show histograms of the normalized deviations for the dust spectral index, and in Fig. 5.5f I show histograms of the normalized deviations for the dust

temperature. For both the dust spectral index and the dust temperature I find biases in the medians of the normalized deviations. The MAD values obtained with a complete pooling model are significantly larger than one. This is again a result of the fact that the uncertainty obtained on the average spectral parameters in a region are not representative of the pixel-level uncertainty in those parameters. The distributions of normalized deviations for the H(LC) set are closer to the standard Gaussian than for the CP(LC) and CP(L) sets. Improved point estimates could potentially be obtained by conducting a more detailed exploration of informative priors on the hyper-parameters of dust spectral parameters, or considering model re-parametrizations. However, the hierarchical model is still limited by the lack of additional frequency channels at $\nu > 402$ GHz.

5.6 Conclusions

I have developed a new implementation of Bayesian CMB component separation, using the NUTS algorithm to draw samples from the full posterior distribution. The NUTS algorithm is a self-tuning variant of HMC, that avoids the random walk behaviour that leads to slow convergence when using standard Metropolis-Hastings and Gibbs sampling algorithms. Measured against the rate of effective sample generation, NUTS offers performance improvements of $\sim 10^3$ compared to Metropolis-Hastings when fitting the complete pooling model. Geometrical pathologies typical of hierarchical models often make variants of HMC the only reliable option for the diagnosis of divergences and biased inferences (Betancourt and Girolami, 2015).

I apply this component separation algorithm to simulations of the *LiteBIRD* and C-BASS experiments to validate the algorithm performance and fidelity. These simulations use a tensor-to-scalar ratio of $r = 5 \times 10^{-3}$ and a lensing amplitude of $A_L = 1$. Component separation is performed over a set of separate sky regions, defined using the mean-shift algorithm. This clusters sky regions according to the similarity in their synchrotron and dust spectral properties, and their location on the sky. In each region I adopt two different modelling approaches, namely complete pooling and a hierarchical model. In the complete pooling model I assume the spectral parameters in each region are constant. In the hierarchical model I assume spectral parameters are drawn from underlying Gaussian distributions, fitting for the hyper-parameters defining the mean and variance of the Gaussian hyper-distributions, along with the individual pixel-by-pixel spectral parameters constrained by these hyper-distributions.

When using the complete pooling model I am able to recover accurate estimates of the CMB over much of the sky. However, component separation artefacts are present close to the Galactic plane where the CMB is highly sub-dominant to foregrounds. Using the hierarchical model, these artefacts are removed from the recovered CMB. Estimating the CMB power spectra with these maps, I find the complete pooling model induces large-scale foreground residuals in the recovered power spectra. Using multipoles between $30 \leq \ell \leq 180$ and fixing the lensing amplitude, I was able to translate these power spectrum estimates into tensor-to-scalar ratio constraints. With only *LiteBIRD* frequency channels, and using the complete pooling model, I find $r = (12.9 \pm 1.4) \times 10^{-3}$. Applying the complete pooling model with an additional C-BASS channel at 5 GHz I find $r = (9.0 \pm 1.1) \times 10^{-3}$, and using the hierarchical model with C-BASS and *LiteBIRD* I find $r = (5.2 \pm 1.0) \times 10^{-3}$. The addition of C-BASS reduces the bias in the recovered tensor-to-scalar ratio for the complete pooling model. However, the crude assumptions made regarding the behaviour of spectral parameters still leaves a $\sim 3.6\sigma$ bias in the estimate of r . I find that the hierarchical model offers an effective generative approach to the modelling of spectral parameters, that helps to mitigate the propensity for fitting outliers when assuming total independence between spectral parameters. The reduced foreground residuals in the recovered CMB maps also means that the hierarchical model does not suffer from the same issues around model misspecification at the power spectrum level, which results in increased uncertainties and biases on recovered cosmological parameters for the complete pooling model.

For the analysis in this chapter I have developed a simple proof-of-concept implementation of the component separation algorithm. Potential future work includes the extension of the algorithm to allow for the joint fitting of dipole and monopole corrections, along with instrumental effects such as bandpass corrections. It would also be prudent to extend the power spectrum estimation, to include the direct joint sampling of CMB amplitudes and power spectra at low multipoles, and expand the code to allow for multi-resolution analyses. From the perspective of modelling, one may also consider more sophisticated approaches to clustering. For the purposes of validation in this study I used naive spectral indices as tracers of the foreground spectral properties, which will be contaminated by noise. In future work, it would be worthwhile examining improvements in region definition by using more sophisticated estimators of foregrounds spectral properties, along with studying the optimal datasets to be used as foreground templates. Finally, significant gain can potentially

be made by exploring optimizations to the component separation code e.g., through re-parametrizations and GPU acceleration.

Chapter 6

Conclusions

Accurate component separation will be vital in order to detect any primordial B -mode signal. If present, this would be buried deep within Galactic foreground emission. The core challenge is two-fold. First, we have a data challenge i.e., we require high sensitivity data, with a frequency coverage that is wide enough for us to constrain foreground emission across the sky. Second, we have an algorithmic challenge. Our component separation algorithms must be able to accurately recover a primordial signal at targeted sensitivities of $\sigma(r) \sim 10^{-3}$, with an accurate quantification of the corresponding uncertainty. In this chapter I summarize the work presented in this thesis, and discuss relevant future work.

6.1 Thesis Summary

In Chapter 2 I gave an overview of the C-BASS experiment, a survey covering the whole sky in total intensity and polarization at 4.76 GHz. C-BASS will provide vital constraints on polarized synchrotron emission, that will allow next-generation CMB experiments to constrain synchrotron spectral behaviour. I described the key survey specifications and data reduction steps, and presented cross-checks of the calibration stability and polarization angle calibration, confirming that the C-BASS calibration is both accurate and stable.

In Chapter 3 I presented the work described in [Grumitt et al. \(2020b\)](#), regarding the detection of point sources in the C-BASS northern sky map. I described the SMHW2 source detection algorithm, which adapted the source detection algorithm of [Planck Collaboration et al. \(2014b, 2016f\)](#) to perform harmonic filtering on the sphere. This detection algorithm was used to produce a catalogue of 1784 point-sources covering declinations $\delta > -10^\circ$. Important data-quality checks were

performed, by matching the C-BASS northern sky catalogue with the GB6 and PMN catalogues. The typical positional offsets of C-BASS sources were ~ 3 arcmin from GB6/PMN sources, which is consistent with the C-BASS pointing residuals. Further, the C-BASS flux-density scale was found to be consistent with the GB6/PMN flux-density scales to within ~ 4 per cent, and the C-BASS and GB6 differential source counts are consistent. The C-BASS source catalogue can now be used for the construction of point-source masks in any analysis of diffuse emission, and has been used by the C-BASS collaboration to produce a source-subtracted map of the northern sky.

In Chapter 4 I presented an early analysis of the polarized power spectra of Galactic synchrotron emission in the C-BASS northern sky maps. By conducting a joint, cross-spectrum based analysis with low-frequency WMAP and *Planck* data I was able to place constraints on the synchrotron SED in harmonic space, which corresponds to the domain of many CMB component separation analyses. Over the multipoles and masks considered, the weighted average of the E -mode spectral index was $\beta_s^{EE} = -3.10 \pm 0.01$, and the weighted average of the B -mode spectral index was $\beta_s^{BB} = -3.20 \pm 0.04$. The spectral index constraints were broadly consistent with those found in [Krachmalnicoff et al. \(2018\)](#). There was also some evidence for synchrotron EB correlation, which has the potential to complicate the self-calibration of next-generation CMB experiments. Using constraints on the C-BASS B -mode auto-spectrum, I was able to estimate the approximate level of synchrotron contamination to primordial B -modes. The minimum level of contamination corresponded to an effective tensor-to-scalar ratio of $r_{\text{eff}} \sim 10^{-3}$, with a typical synchrotron contamination of $r_{\text{eff}} \sim 10^{-2}$. This is consistent with analysis in [Krachmalnicoff et al. \(2018\)](#) over the southern sky, and clearly demonstrates that synchrotron emission cannot be neglected in next-generation B -mode searches.

In Chapter 5 I moved to consider algorithmic challenges in CMB component separation, presenting the work described in [Grumitt et al. \(2020a\)](#). Here, I described a new implementation of Bayesian parametric component separation, utilizing the highly efficient NUTS algorithm for posterior sampling. NUTS is a gradient-based sampling algorithm, first presented in [Hoffman and Gelman \(2014\)](#), that avoids the random walk behaviour that affects non-gradient based sampling algorithms such as Metropolis-Hastings and Gibbs sampling. In addition to its favourable scaling to high-dimensional problems, NUTS also provides powerful self-diagnostics of potential failures in geometrical ergodicity when sampling from complex target distributions. This is particularly important for the hierarchical foreground model, with this class

of model being known to exhibit complex, highly-curved posterior geometries. For foreground modelling I compared two approaches. The first was a complete pooling model, which assumed foreground spectral parameters to be constant over large sky regions. This method has been applied in several CMB component separation analyses as a means through which one can reduce the statistical uncertainties on recovered CMB amplitudes (Górski et al., 2005; Stompor et al., 2009; Errard et al., 2011; Stompor et al., 2016; Alonso et al., 2017; Irfan et al., 2019; Khatri, 2019; Thorne et al., 2019). The second was a hierarchical foreground model, which modelled hyperdistributions over these large sky regions, which subsequently induce a prior on pixel-level spectral parameters. The complete pooling model was found to result in significant artefacts in the recovered CMB maps, that were removed by the hierarchical model without inflating parameter uncertainties. This could be translated to power spectrum constraints. For an input tensor-to-scalar ratio of $r = 5 \times 10^{-3}$, the hierarchical model recovered $r = (5.2 \pm 1.0) \times 10^{-3}$ using multipoles $30 \leq \ell < 180$. For the same dataset, the complete pooling model recovered $r = (9.0 \pm 1.1) \times 10^{-3}$. Hierarchical modelling was found to provide a powerful generative framework through which foreground spatial variations can be modelled without inflating the final parameter uncertainties.

6.2 Future Work

Significant work remains in the scientific analysis of C-BASS northern sky data. In this thesis I have presented an early analysis of polarized synchrotron emission in the harmonic domain, in combination with map-based spectral index estimates. In areas of the sky with low SNR, these spectral index estimates become heavily contaminated by noise, which can introduce unphysical small-scale artefacts into the spectral index maps. Improved synchrotron spectral index maps can likely be obtained by using the hierarchical modelling techniques discussed in Chapter 5. In the future, it will also be important to conduct joint analyses with other low-frequency surveys such as S-PASS (Carretti, 2010) and QUIJOTE (Génova-Santos et al., 2015b), and integrate C-BASS into the component separation pipelines for next-generation experiments such as the Simons Observatory (Simons Observatory Collaboration et al., 2019), CMB-S4 (Abazajian et al., 2016) and *LiteBIRD* (Matsumura et al., 2014; Hazumi et al., 2019). Analysis with S-PASS and QUIJOTE can be particularly powerful in assessing the strength of Faraday rotation effects. Constraints on RM values across

the sky will be very important when fully integrating C-BASS into any component separation pipelines.

The C-BASS northern-sky survey overlaps with S-PASS along a $\sim 10^\circ$ strip at its lower declination bounds. Whilst this can be utilized for initial scientific analyses, southern survey data will be vital to fully exploit such a joint analysis. Full-sky coverage will also be key for both space-based missions, as well as ground-based experiments that will observe primarily in the southern hemisphere. The C-BASS southern survey is ongoing, and will require at least two years of observations. Significant future work on C-BASS will involve the development of the C-BASS south data reduction pipeline, particularly to fully utilize the spectral resolution of the southern system. Moreover, the overlap between the northern and southern surveys can be used to conduct cross-checks and improve the C-BASS calibration. Looking beyond C-BASS, additional low-frequency experiments between 5 – 30 GHz can be invaluable for providing detailed constraints on the synchrotron SED, and reducing biases on recovered cosmological parameters. As discussed in Chapter 5, the addition of C-BASS data enables us to constrain the synchrotron spectral index. However, in order to constrain more realistic and complex models, we will require additional data. ELFS has been proposed to provide the necessary frequency coverage to constrain more sophisticated synchrotron models (Hill-Valler, 2019; Ganga et al., 2019).

The NUTS component separation algorithm presented in Chapter 5 has been developed as a proof-of-concept implementation. The code should be modified to include realistic complications such as bandpass corrections, monopole and dipole fitting etc. It would also be natural to extend the algorithm to allow for multi-resolution analyses by performing fitting in the harmonic domain, as is the case for COMMANDER2 (Planck Collaboration et al., 2018b). The success of hierarchical foreground modelling is intricately tied to the choice of regions over which hyper-distributions are modelled. This could likely be improved by adopting clustering schemes that adaptively define larger regions where there is low SNR, thereby increasing the smoothing effects of the hyper-distributions in these regions. More complex hierarchical structures could also be adopted e.g., by defining an additional global hyper-distribution that could help to smooth the transitions between regions. The foreground modelling approach could also be combined with a Bayesian delensing and power spectrum estimation scheme, such as that in Millea et al. (2020), to form a complete forward model from our sky maps to key cosmological statistics.

It is important to note that the hierarchical foreground model pushes the limits of standard implementations of NUTS. Including the complications described above

would further strain the algorithm’s computational performance. It would therefore be prudent to consider extensions to approximate Bayesian inference algorithms. This has numerous applications beyond cosmology, and significant work is being conducted into ways through which Bayesian inference can be scaled to very high-dimensional problems. In this thesis I have primarily considered applications of HMC (through NUTS), and also briefly summarized variational inference (VI). Both of these methods sit at the forefront of modern Bayesian computation, but there is major potential for improvements.

HMC algorithms sample very effectively from strongly log-concave distributions such as Gaussians (Mangoubi and Smith, 2017). However, when faced with more complex target distributions, such as the funnel geometries typical of hierarchical models, the performance of HMC can suffer dramatically. Recent work has considered the application of normalizing flows to finding mappings between general target geometries and easy-to-sample geometries (Hoffman et al., 2019). Normalizing flows are able to construct incredibly complex and rich mappings by stacking a set of invertible maps, typically parametrized through neural network architectures (Papamakarios et al., 2019). Other approaches have considered computational speed-ups for particular model classes e.g., through the automatic marginalization over certain model parameters (Margossian et al., 2020). For very high-dimensional models, such as Bayesian neural networks, VI has often been favoured over MCMC. In these models, MCMC algorithms are often intractable. By casting posterior inference as an optimization problem, VI is able to scale much better to such problems. However, VI is often prone to overly compact representations of the target geometry, and a poor choice of approximating family can drastically fail to capture the target geometry. Much work remains to be done in understanding the regimes in which one can obtain reliable posterior approximations through VI. Improvements can also potentially be obtained by considering alternative objective functions to the Kullback-Leibler divergence, and by exploring the interface between VI and MCMC methods (Blei et al., 2016).

Bibliography

- K. N. Abazajian et al. CMB-S4 Science Book, First Edition. *arXiv e-prints*, art. arXiv:1610.02743, Oct. 2016.
- M. H. Abitbol, J. C. Hill, and B. R. Johnson. Foreground-induced biases in CMB polarimeter self-calibration. *MNRAS*, 457(2):1796–1803, Apr. 2016. doi: 10.1093/mnras/stw030.
- N. Aghanim, S. Majumdar, and J. Silk. Secondary anisotropies of the CMB. *Reports on Progress in Physics*, 71(6):066902, June 2008. doi: 10.1088/0034-4885/71/6/066902.
- A. Albrecht and P. J. Steinhardt. Cosmology for Grand Unified Theories with Radiatively Induced Symmetry Breaking. *Phys. Rev. Lett.*, 48(17):1220–1223, Apr. 1982. doi: 10.1103/PhysRevLett.48.1220.
- Y. Ali-Haïmoud, C. M. Hirata, and C. Dickinson. A refined model for spinning dust radiation. *MNRAS*, 395:1055–1078, May 2009. doi: 10.1111/j.1365-2966.2009.14599.x.
- D. Alonso et al. Simulated forecasts for primordial B -mode searches in ground-based experiments. *Phys. Rev. D*, 95:043504, Feb 2017. doi: 10.1103/PhysRevD.95.043504.
- D. Alonso et al. A unified pseudo- C_ℓ framework. *MNRAS*, 484:4127–4151, Apr. 2019. doi: 10.1093/mnras/stz093.
- R. A. Alpher and R. Herman. Evolution of the Universe. *Nature*, 162(4124):774–775, Nov. 1948. doi: 10.1038/162774b0.
- R. A. Alpher, H. Bethe, and G. Gamow. The Origin of Chemical Elements. *Physical Review*, 73(7):803–804, Apr. 1948. doi: 10.1103/PhysRev.73.803.

- J. Alsing et al. Hierarchical cosmic shear power spectrum inference. *MNRAS*, 455 (4):4452–4466, Feb. 2016. doi: 10.1093/mnras/stv2501.
- AMI Consortium et al. 10C survey of radio sources at 15.7 GHz - II. First results. *MNRAS*, 415:2708–2722, Aug. 2011. doi: 10.1111/j.1365-2966.2011.18925.x.
- J.-P. Antoine and P. Vandergheynst. Wavelets on the 2-Sphere: A Group-Theoretical Approach. *Applied and Computational Harmonic Analysis*, 7(3):262–291, nov 1999. ISSN 1063-5203. doi: 10.1006/ACHA.1999.0272.
- F. Argüeso et al. Detection/estimation of the modulus of a vector. Application to point- source detection in polarization data. *MNRAS*, 395:649–656, May 2009. doi: 10.1111/j.1365-2966.2009.14549.x.
- J. W. M. Baars et al. The absolute spectrum of CAS A - an accurate flux density scale and a set of secondary calibrators. *A&A*, 61:99–106, Oct. 1977.
- A. J. Banday et al. Noncosmological Signal Contributions to the COBE DMR 4 Year Sky Maps. *ApJ*, 468:L85, Sept. 1996. doi: 10.1086/310235.
- D. Barbosa et al. A Polarized Galactic Emission Mapping Experiment at 5-10 GHz. In A. M. Mourão et al., editors, *New Worlds in Astroparticle Physics: Proceedings of the Fifth International Workshop*, pages 233–237, Feb. 2006. doi: 10.1142/9789812774439_0024.
- J. M. Bardeen, P. J. Steinhardt, and M. S. Turner. Spontaneous creation of almost scale-free density perturbations in an inflationary universe. *Phys. Rev. D*, 28:679–693, Aug. 1983. doi: 10.1103/PhysRevD.28.679.
- D. Barkats et al. Degree-scale Cosmic Microwave Background Polarization Measurements from Three Years of BICEP1 Data. *ApJ*, 783(2):67, Mar. 2014. doi: 10.1088/0004-637X/783/2/67.
- S. Basak and J. Delabrouille. A needlet ILC analysis of WMAP 9-year polarization data: CMB polarization power spectra. *MNRAS*, 435(1):18–29, Oct. 2013. doi: 10.1093/mnras/stt1158.
- D. Baumann. TASI Lectures on Inflation. *arXiv e-prints*, art. arXiv:0907.5424, July 2009.
- D. Baumann. Cosmology - Part III Mathematical Tripos, September 2015.

- D. Baumann et al. Probing Inflation with CMB Polarization. In S. Dodelson et al., editors, *American Institute of Physics Conference Series*, volume 1141 of *American Institute of Physics Conference Series*, pages 10–120, June 2009. doi: 10.1063/1.3160885.
- C. L. Bennett et al. Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results. *ApJS*, 208:20, Oct. 2013. doi: 10.1088/0067-0049/208/2/20.
- A. Benoît et al. Archeops: a high resolution, large sky coverage balloon experiment for mapping cosmic microwave background anisotropies. *Astroparticle Physics*, 17(2):101–124, May 2002. doi: 10.1016/S0927-6505(01)00141-4.
- A. Benoît et al. Cosmological constraints from Archeops. *A&A*, 399:L25–L30, Mar. 2003. doi: 10.1051/0004-6361:20021722.
- E. M. Berkhuijsen. A survey of the continuum radiation at 820 MHz between declinations -7° and $+85^\circ$. I. Observations and reductions. *A&AS*, 5:263, Mar. 1972.
- M. Betancourt. Diagnosing Suboptimal Cotangent Disintegrations in Hamiltonian Monte Carlo. *arXiv e-prints*, art. arXiv:1604.00695, Apr 2016.
- M. Betancourt. A Conceptual Introduction to Hamiltonian Monte Carlo. *arXiv e-prints*, art. arXiv:1701.02434, Jan 2017.
- M. Betancourt. Towards A Principled Bayesian Workflow (RStan). https://github.com/betanalpha/knitr_case_studies/tree/master/principled_bayesian_workflow,commit23eb263be4cfb44278d0dfb8ddb593a4b142506, 2020. (Accessed on 21/07/2020).
- M. Betancourt et al. The geometric foundations of hamiltonian monte carlo. *Bernoulli*, 23(4A):2257–2298, 11 2017. doi: 10.3150/16-BEJ810.
- M. J. Betancourt and M. Girolami. Hamiltonian Monte Carlo for Hierarchical Models. In S. K. Upadhyay, U. Singh, D. K. Dey, and A. Loganathan, editors, *Current Trends in Bayesian Methodology with Applications*, chapter 4. Chapman and Hall/CRC, New York, 2015.
- M. Betoule et al. Measuring the tensor to scalar ratio from CMB B-modes in the presence of foregrounds. *A&A*, 503(3):691–706, Sep 2009. doi: 10.1051/0004-6361/200911624.

- BICEP2 Collaboration, Keck Array Collaboration, et al. Constraints on Primordial Gravitational Waves Using Planck, WMAP, and New BICEP2/Keck Observations through the 2015 Season. *Phys. Rev. Lett.*, 121(22):221301, Nov. 2018. doi: 10.1103/PhysRevLett.121.221301.
- BICEP2 Collaboration et al. Detection of B-Mode Polarization at Degree Angular Scales by BICEP2. *Phys. Rev. Lett.*, 112(24):241101, June 2014. doi: 10.1103/PhysRevLett.112.241101.
- BICEP2/Keck Collaboration, Planck Collaboration, et al. Joint Analysis of BICEP2/Keck Array and Planck Data. *Phys. Rev. Lett.*, 114(10):101301, Mar 2015. doi: 10.1103/PhysRevLett.114.101301.
- D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational Inference: A Review for Statisticians. *arXiv e-prints*, art. arXiv:1601.00670, Jan. 2016.
- J. R. Bond, A. H. Jaffe, and L. Knox. Estimating the power spectrum of the cosmic microwave background. *Phys. Rev. D*, 57(4):2117–2137, Feb. 1998. doi: 10.1103/PhysRevD.57.2117.
- R. Brandenberger, R. Kahn, and W. H. Press. Cosmological perturbations in the early universe. *Phys. Rev. D*, 28(8):1809–1821, Oct. 1983. doi: 10.1103/PhysRevD.28.1809.
- S. Brooks et al. *Handbook of Markov Chain Monte Carlo*. CRC press, 2011.
- S. P. Brooks and A. Gelman. General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7(4):434–455, 1998. doi: 10.1080/10618600.1998.10474787.
- M. R. Calabretta, L. Staveley-Smith, and D. G. Barnes. A New 1.4 GHz Radio Continuum Map of the Sky South of Declination +25°. *Publ. Astron. Soc. Australia*, 31:e007, Jan. 2014. doi: 10.1017/pasa.2013.36.
- J. Cardoso et al. Component separation with flexible models. application to the separation of astrophysical emissions. 2008.
- B. Carpenter et al. Stan : A probabilistic programming language. *Journal of Statistical Software*, 76(1), 1 2017. ISSN 1548-7660. doi: 10.18637/jss.v076.i01.

- E. Carretti. In *The Dynamic Interstellar Medium: A Celebration of the Canadian Galactic Plane Survey*, 2010.
- E. Carretti et al. Giant magnetized outflows from the centre of the Milky Way. *Nature*, 493:66–69, Jan. 2013. doi: 10.1038/nature11734.
- P. Carvalho et al. PowellSnakes II: a fast Bayesian approach to discrete object detection in multi-frequency astronomical data sets. *MNRAS*, 427:1384–1400, Dec. 2012. doi: 10.1111/j.1365-2966.2012.22033.x.
- L. Cayón et al. Spherical Mexican hat wavelet: an application to detect non-Gaussianity in the COBE-DMR maps. *MNRAS*, 326:1243–1248, Oct. 2001. doi: 10.1111/j.1365-2966.2001.04641.x.
- R. Cepeda-Arroita et al. First Detection of Spectral Variations of Anomalous Microwave Emission with QUIJOTE and C-BASS. *arXiv e-prints*, art. arXiv:2001.07159, Jan. 2020.
- A. Challinor et al. All-sky convolution for polarimetry experiments. *Phys. Rev. D*, 62(12):123002, Dec. 2000. doi: 10.1103/PhysRevD.62.123002.
- E. L. Chapin et al. A joint analysis of BLAST 250-500 μm and LABOCA 870 μm observations in the Extended Chandra Deep Field-South. *MNRAS*, 411:505–549, Feb 2011. doi: 10.1111/j.1365-2966.2010.17697.x.
- J. Chluba, J. C. Hill, and M. H. Abitbol. Rethinking CMB foregrounds: systematic extension of foreground parametrizations. *MNRAS*, 472:1195–1213, Nov. 2017. doi: 10.1093/mnras/stx1982.
- J. Chluba, A. Ravenni, and B. Bolliet. Improved calculations of electron-ion bremsstrahlung Gaunt factors for astrophysical applications. *MNRAS*, 492(1):177–194, Feb. 2020. doi: 10.1093/mnras/stz3389.
- G. Chon et al. Fast estimation of polarization power spectra using correlation functions. *MNRAS*, 350(3):914–926, May 2004. doi: 10.1111/j.1365-2966.2004.07737.x.
- S. Cole, H. Chu, and S. Greenland. Maximum likelihood, profile likelihood, and penalized likelihood: A primer. *American journal of epidemiology*, 179, 10 2013. doi: 10.1093/aje/kwt245.

- D. Comaniciu and P. Meer. Mean shift: a robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(5): 603–619, may 2002. ISSN 01628828. doi: 10.1109/34.1000236.
- J. J. Condon. Cosmological evolution of radio sources. *ApJ*, 287:461–474, Dec. 1984. doi: 10.1086/162705.
- J. J. Condon et al. The NRAO VLA Sky Survey. *AJ*, 115:1693–1716, May 1998. doi: 10.1086/300337.
- C. J. Copley. The C-Band All Sky Survey. DPhil thesis, University of Oxford, 2014.
- A. Curto, E. Martínez-González, and R. B. Barreiro. On the optimality of the spherical Mexican hat wavelet estimator for the primordial non-Gaussianity. *MNRAS*, 412:1038–1046, Apr. 2011. doi: 10.1111/j.1365-2966.2010.17961.x.
- A. Curto et al. Forecasts on the contamination induced by unresolved point sources in primordial non-Gaussianity beyond Planck. *MNRAS*, 432:728–742, June 2013. doi: 10.1093/mnras/stt511.
- S. Das et al. The Atacama Cosmology Telescope: temperature and gravitational lensing power spectrum measurements from three seasons of data. *JCAP*, 2014(4): 014, Apr. 2014. doi: 10.1088/1475-7516/2014/04/014.
- R. D. Davies et al. A determination of the spectra of Galactic components observed by the Wilkinson Microwave Anisotropy Probe. *MNRAS*, 370(3):1125–1139, Aug. 2006. doi: 10.1111/j.1365-2966.2006.10572.x.
- P. de Bernardis et al. A flat Universe from high-resolution maps of the cosmic microwave background radiation. *Nature*, 404(6781):955–959, Apr. 2000. doi: 10.1038/35010035.
- A. de Oliveira-Costa et al. Galactic Microwave Emission at Degree Angular Scales. *ApJ*, 482:L17–L20, June 1997. doi: 10.1086/310684.
- G. de Zotti et al. Predictions for high-frequency radio surveys of extragalactic sources. *A&A*, 431:893–903, Mar 2005. doi: 10.1051/0004-6361:20042108.
- G. de Zotti et al. Radio and millimeter continuum surveys and their astrophysical implications. *A&A Reviews*, 18:1–65, Feb. 2010. doi: 10.1007/s00159-009-0026-0.

- J. Delabrouille, J. F. Cardoso, and G. Patanchon. Multidetector multicomponent spectral matching and applications for cosmic microwave background data analysis. *MNRAS*, 346(4):1089–1102, Dec. 2003. doi: 10.1111/j.1365-2966.2003.07069.x.
- J. Delabrouille et al. A full sky, low foreground, high resolution CMB map from WMAP. *A&A*, 493(3):835–857, Jan. 2009. doi: 10.1051/0004-6361:200810514.
- M. J. Devlin et al. A Measurement of the Angular Power Spectrum of the CMB from $l = 100$ to 400. In *American Astronomical Society Meeting Abstracts*, volume 195 of *American Astronomical Society Meeting Abstracts*, page 55.05, Dec. 1999.
- R. H. Dicke et al. Cosmic Black-Body Radiation. *ApJ*, 142:414–419, July 1965. doi: 10.1086/148306.
- C. Dickinson. CMB foregrounds - A brief review. *ArXiv e-prints*, June 2016.
- C. Dickinson, R. D. Davies, and R. J. Davis. Towards a free-free template for CMB foregrounds. *MNRAS*, 341:369–384, May 2003. doi: 10.1046/j.1365-8711.2003.06439.x.
- C. Dickinson et al. The State-of-Play of Anomalous Microwave Emission (AME) research. *New Astronomy Review*, 80:1–28, Feb. 2018. doi: 10.1016/j.newar.2018.02.001.
- C. Dickinson et al. The C-Band All-Sky Survey (C-BASS): constraining diffuse Galactic radio emission in the North Celestial Pole region. *MNRAS*, 485:2844–2860, May 2019. doi: 10.1093/mnras/stz522.
- G. Dobler. A Last Look at the Microwave Haze/Bubbles with WMAP. *ApJ*, 750(1):17, May 2012a. doi: 10.1088/0004-637X/750/1/17.
- G. Dobler. Identifying the Radio Bubble Nature of the Microwave Haze. *ApJL*, 760(1):L8, Nov. 2012b. doi: 10.1088/2041-8205/760/1/L8.
- G. Dobler and D. P. Finkbeiner. Extended Anomalous Foreground Emission in the WMAP Three-Year Data. *ApJ*, 680(2):1222–1234, June 2008. doi: 10.1086/587862.
- G. Dobler et al. The Fermi Haze: A Gamma-ray Counterpart to the Microwave Haze. *ApJ*, 717(2):825–842, July 2010. doi: 10.1088/0004-637X/717/2/825.
- S. Dodelson. *Modern cosmology*. 2003.

- B. T. Draine. *Physics of the Interstellar and Intergalactic Medium*. 2011.
- B. T. Draine and B. S. Hensley. Quantum Suppression of Alignment in Ultrasmall Grains: Microwave Emission from Spinning Dust will be Negligibly Polarized. *ApJ*, 831(1):59, Nov. 2016. doi: 10.3847/0004-637X/831/1/59.
- B. T. Draine and A. Lazarian. Electric Dipole Radiation from Spinning Dust Grains. *ApJ*, 508(1):157–179, Nov. 1998. doi: 10.1086/306387.
- B. T. Draine and A. Lazarian. Magnetic Dipole Microwave Emission from Dust Grains. *ApJ*, 512(2):740–754, Feb. 1999. doi: 10.1086/306809.
- S. Duane et al. Hybrid monte carlo. *Physics Letters B*, 195(2):216 – 222, 1987. ISSN 0370-2693. doi: [https://doi.org/10.1016/0370-2693\(87\)91197-X](https://doi.org/10.1016/0370-2693(87)91197-X).
- J. Dunkley et al. Prospects for polarized foreground removal. In S. Dodelson et al., editors, *American Institute of Physics Conference Series*, volume 1141 of *American Institute of Physics Conference Series*, pages 222–264, June 2009. doi: 10.1063/1.3160888.
- A. Einstein. Die Feldgleichungen der Gravitation. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin)*, pages 844–847, Jan. 1915.
- A. Einstein. Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik*, 354(7):769–822, Jan. 1916. doi: 10.1002/andp.19163540702.
- H. K. Eriksen et al. Power Spectrum Estimation from High-Resolution Maps by Gibbs Sampling. *ApJS*, 155(2):227–241, Dec. 2004. doi: 10.1086/425219.
- H. K. Eriksen et al. Joint Bayesian Component Separation and CMB Power Spectrum Estimation. *ApJ*, 676(1):10–32, Mar. 2008. doi: 10.1086/525277.
- J. Errard, F. Stivoli, and R. Stompor. Framework for performance forecasting and optimization of CMB B-mode observations in the presence of astrophysical foregrounds. *Phys. Rev. D*, 84(6):063005, Sep 2011. doi: 10.1103/PhysRevD.84.063005.
- J. Errard et al. Robust forecasts on fundamental physics from the foreground-obscured, gravitationally-lensed CMB polarization. *JCAP*, 2016(3):052, Mar 2016. doi: 10.1088/1475-7516/2016/03/052.
- M. Evans and G. H. Jang. Weak informativity and the information in one prior relative to another. *Statist. Sci.*, 26(3):423–439, 08 2011. doi: 10.1214/11-STS357.

- D. P. Finkbeiner. A Full-Sky $H\alpha$ Template for Microwave Foreground Prediction. *ApJS*, 146:407–415, June 2003. doi: 10.1086/374411.
- D. P. Finkbeiner. Microwave Interstellar Medium Emission Observed by the Wilkinson Microwave Anisotropy Probe. *ApJ*, 614(1):186–193, oct 2004a. ISSN 0004-637X. doi: 10.1086/423482.
- D. P. Finkbeiner. WMAP Microwave Emission Interpreted as Dark Matter Annihilation in the Inner Galaxy. *ArXiv e-prints*, sep 2004b.
- D. J. Fixsen. The Temperature of the Cosmic Microwave Background. *ApJ*, 707(2): 916–920, Dec. 2009. doi: 10.1088/0004-637X/707/2/916.
- R. Flauger, J. C. Hill, and D. N. Spergel. Toward an understanding of foreground emission in the BICEP2 region. *JCAP*, 2014(8):039, Aug. 2014. doi: 10.1088/1475-7516/2014/08/039.
- D. Foreman-Mackey, G. Barentsen, and T. Barclay. dfm/exoplanet: exoplanet v0.1.4, Feb. 2019. URL <https://doi.org/10.5281/zenodo.2561395>.
- A. Friedmann. Über die Krümmung des Raumes. *Zeitschrift für Physik*, 10:377–386, Jan. 1922. doi: 10.1007/BF01332580.
- K. Ganga et al. European Work on Future Ground-Based CMB Experiments. In *Bulletin of the AAS*, volume 51, page 111, Sep 2019.
- A. Gelman. Multilevel (hierarchical) modeling: What it can and cannot do. *Technometrics*, 48(3):432–435, 2006a. doi: 10.1198/004017005000000661.
- A. Gelman. Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper). *Bayesian Anal.*, 1(3):515–534, 09 2006b. doi: 10.1214/06-BA117A.
- A. Gelman and C. Hennig. Beyond subjective and objective in statistics. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 180(4):967–1033, 2017. doi: 10.1111/rssa.12276.
- A. Gelman and J. Hill. *Data analysis using regression and multilevel/hierarchical models*, volume Analytical methods for social research. Cambridge University Press, New York, 2007.

- A. Gelman and D. B. Rubin. Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*, 7:457–472, Jan 1992. doi: 10.1214/ss/1177011136.
- A. Gelman, X. li Meng, and H. Stern. Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica*, pages 733–807, 1996.
- A. Gelman, D. Simpson, and M. Betancourt. The Prior Can Often Only Be Understood in the Context of the Likelihood. *Entropy*, 19:555, Oct 2017. doi: 10.3390/e19100555.
- A. Gelman et al. *Bayesian Data Analysis*. Chapman and Hall/CRC, 2nd ed. edition, 2004.
- A. Gelman et al. A weakly informative default prior distribution for logistic and other regression models. *Ann. Appl. Stat.*, 2(4):1360–1383, 12 2008. doi: 10.1214/08-AOAS191.
- S. Geman and D. Geman. Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6):721–741, Nov 1984. ISSN 1939-3539. doi: 10.1109/TPAMI.1984.4767596.
- R. Génova-Santos et al. QUIJOTE scientific results - I. Measurements of the intensity and polarisation of the anomalous microwave emission in the Perseus molecular complex. *MNRAS*, 452:4169–4182, Oct. 2015a. doi: 10.1093/mnras/stv1405.
- R. Génova-Santos et al. The QUIJOTE experiment: project overview and first results. In *Highlights of Spanish Astrophysics VIII*, pages 207–212, May 2015b.
- C. J. Geyer. Practical markov chain monte carlo. *Statist. Sci.*, 7(4):473–483, 11 1992. doi: 10.1214/ss/1177011137.
- M. Girolami and B. Calderhead. Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011. doi: 10.1111/j.1467-9868.2010.00765.x.
- J. González-Nuevo et al. The Mexican hat wavelet family: application to point-source detection in cosmic microwave background maps. *MNRAS*, 369:1603–1610, July 2006. doi: 10.1111/j.1365-2966.2006.10442.x.

- K. M. Górski et al. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *ApJ*, 622:759–771, Apr. 2005. doi: 10.1086/427976.
- P. C. Gregory, W. K. Scott, and B. Poller. Variable Radio Sources in the GB6 Catalog. In J. A. Zensus, G. B. Taylor, and J. M. Wrobel, editors, *IAU Colloq. 164: Radio Emission from Galactic and Extragalactic Compact Sources*, volume 144 of *Astronomical Society of the Pacific Conference Series*, page 283, Jan 1998.
- P. C. Gregory et al. The GB6 Catalog of Radio Sources. *ApJS*, 103:427, Apr. 1996. doi: 10.1086/192282.
- P. C. Gregory et al. The GB6 4.85 GHz radio variability catalog. In R. T. Schilizzi, editor, *Galaxies and their Constituents at the Highest Angular Resolutions*, volume 205 of *IAU Symposium*, page 98, Jan 2001.
- M. R. Griffith and A. E. Wright. The Parkes-MIT-NRAO (PMN) surveys. I - The 4850 MHz surveys and data reduction. *AJ*, 105:1666–1679, May 1993. doi: 10.1086/116545.
- M. R. Griffith et al. The Parkes-MIT-NRAO (PMN) surveys. 3: Source catalog for the tropical survey (-29 deg less than δ less than -9 deg .5). *ApJS*, 90:179–295, Jan. 1994. doi: 10.1086/191863.
- M. R. Griffith et al. The Parkes-MIT-NRAO (PMN) surveys. 6: Source catalog for the equatorial survey (-9.5 deg less than δ less than $+10.0$ deg). *ApJS*, 97: 347–453, Apr. 1995. doi: 10.1086/192146.
- R. D. P. Grumitt, L. R. P. Jew, and C. Dickinson. Hierarchical Bayesian CMB Component Separation with the No-U-Turn Sampler. *MNRAS*, June 2020a. doi: 10.1093/mnras/staa1857.
- R. D. P. Grumitt et al. The C-Band All-Sky Survey (C-BASS): Total intensity point-source detection over the northern sky. *MNRAS*, June 2020b. doi: 10.1093/mnras/staa1572.
- A. H. Guth. Inflationary universe: A possible solution to the horizon and flatness problems. *Phys. Rev. D*, 23(2):347–356, Jan. 1981. doi: 10.1103/PhysRevD.23.347.
- A. H. Guth and S. Y. Pi. Fluctuations in the New Inflationary Universe. *Phys. Rev. Lett.*, 49(15):1110–1113, Oct. 1982. doi: 10.1103/PhysRevLett.49.1110.

- A. H. Guth and S.-Y. Pi. Quantum mechanics of the scalar field in the new inflationary universe. *Phys. Rev. D*, 32(8):1899–1920, Oct. 1985. doi: 10.1103/PhysRevD.32.1899.
- A. E. Guzmán et al. All-sky Galactic radiation at 45 MHz and spectral index between 45 and 408 MHz. *A&A*, 525:A138, Jan 2011. doi: 10.1051/0004-6361/200913628.
- M. C. Guzzetti et al. Gravitational waves from inflation. *Nuovo Cimento Rivista Serie*, 39(9):399–495, Sept. 2016. doi: 10.1393/ncr/i2016-10127-1.
- G. Hamra, R. Maclehorse, and S. Cole. Sensitivity analyses for sparse-data problems - using weakly informative bayesian priors. Mar. 2013.
- S. Hanany et al. MAXIMA-1: A Measurement of the Cosmic Microwave Background Anisotropy on Angular Scales of $10'$ - 5° . *ApJ*, 545(1):L5–L9, Dec. 2000. doi: 10.1086/317322.
- S. Hanany et al. PICO: Probe of Inflation and Cosmic Origins. In *Bulletin of the AAS*, volume 51, page 194, Sep 2019.
- D. Hanson et al. Detection of B-Mode Polarization in the Cosmic Microwave Background with Data from the South Pole Telescope. *Phys. Rev. Lett.*, 111(14):141301, Oct. 2013. doi: 10.1103/PhysRevLett.111.141301.
- K. Harrington et al. *The Cosmology Large Angular Scale Surveyor*, volume 9914 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 99141K. 2016. doi: 10.1117/12.2233125.
- J. J. Harwood et al. FR II radio galaxies at low frequencies - I. Morphology, magnetic field strength and energetics. *MNRAS*, 458:4443–4455, June 2016. doi: 10.1093/mnras/stw638.
- C. G. T. Haslam et al. A 408 MHz all-sky continuum survey. I - Observations at southern declinations and for the North Polar region. *A&A*, 100:209–219, July 1981.
- C. G. T. Haslam et al. A 408 MHz all-sky continuum survey. II - The atlas of contour maps. *A&AS*, 47:1, Jan. 1982.
- S. Hawking. The development of irregularities in a single bubble inflationary universe. *Physics Letters B*, 115(4):295 – 297, 1982. ISSN 0370-2693. doi: [https://doi.org/10.1016/0370-2693\(82\)90373-2](https://doi.org/10.1016/0370-2693(82)90373-2).

- M. Hazumi. B-mode from Space. <https://wwwmpa.mpa-garching.mpg.de/~komatsu/meetings/bmode2019/schedule/talks/hazumi.pdf>, December 2019.
- M. Hazumi et al. LiteBIRD: A Satellite for the Studies of B-Mode Polarization and Inflation from Cosmic Background Radiation Detection. *Journal of Low Temperature Physics*, 194(5-6):443–452, Mar. 2019. doi: 10.1007/s10909-019-02150-5.
- S. E. Healey et al. CRATES: An All-Sky Survey of Flat-Spectrum Radio Sources. *ApJS*, 171(1):61–71, Jul 2007. doi: 10.1086/513742.
- S. E. Healey et al. Filling in the Gaps in the 4.85 GHz Sky. *AJ*, 138(4):1032–1036, Oct 2009. doi: 10.1088/0004-6256/138/4/1032.
- B. S. Hensley and P. Bull. Mitigating Complex Dust Foregrounds in Future Cosmic Microwave Background Polarization Experiments. *ApJ*, 853:127, Feb 2018. doi: 10.3847/1538-4357/aaa489.
- D. Herranz et al. A novel multifrequency technique for the detection of point sources in cosmic microwave background maps. *MNRAS*, 394:510–520, Mar. 2009. doi: 10.1111/j.1365-2966.2008.14336.x.
- J. Hill-Valler. NextBASS: a new low-frequency CMB foregrounds instrument. DPhil thesis, University of Oxford, 2019.
- G. Hinshaw et al. Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results. *ApJS*, 208(2):19, Oct. 2013. doi: 10.1088/0067-0049/208/2/19.
- E. Hivon et al. MASTER of the Cosmic Microwave Background Anisotropy Power Spectrum: A Fast Method for Statistical Analysis of Large and Complex Cosmic Microwave Background Data Sets. *ApJ*, 567(1):2–17, Mar 2002. doi: 10.1086/338126.
- M. Hoffman et al. NeuTra-lizing Bad Geometry in Hamiltonian Monte Carlo Using Neural Transport. *arXiv e-prints*, art. arXiv:1903.03704, Mar. 2019.
- M. D. Hoffman and A. Gelman. The No-U-turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. *JMLR*, 15(1):1593–1623, Jan. 2014. ISSN 1532-4435.

- D. W. Hogg, J. Bovy, and D. Lang. Data analysis recipes: Fitting a model to data. *arXiv e-prints*, art. arXiv:1008.4686, Aug. 2010.
- C. M. Holler et al. A Circularly Symmetric Antenna Design With High Polarization Purity and Low Spillover. *IEEE Transactions on Antennas and Propagation*, 61: 117–124, Jan. 2013. doi: 10.1109/TAP.2012.2219843.
- W. Hu. CMB temperature and polarization anisotropy fundamentals. *Annals of Physics*, 303(1):203–225, Jan. 2003. doi: 10.1016/S0003-4916(02)00022-2.
- W. Hu and S. Dodelson. Cosmic Microwave Background Anisotropies. *ARA&A*, 40: 171–216, Jan. 2002. doi: 10.1146/annurev.astro.40.060401.093926.
- W. Hu and N. Sugiyama. Anisotropies in the Cosmic Microwave Background: an Analytic Approach. *ApJ*, 444:489, May 1995. doi: 10.1086/175624.
- W. Hu and M. White. A CMB polarization primer. *New Astronomy*, 2(4):323–344, Oct. 1997. doi: 10.1016/S1384-1076(97)00022-5.
- E. Hubble. A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae. *Proceedings of the National Academy of Science*, 15(3):168–173, Mar. 1929. doi: 10.1073/pnas.15.3.168.
- W. A. Imbriale and R. Abraham. Radio Frequency Optics Design of the Deep Space Network Large Array 6-Meter Breadboard Antenna. *Interplanetary Network Progress Report*, 157:1–8, May 2004.
- M. O. Irfan et al. Determining thermal dust emission from Planck HFI data using a sparse, parametric technique. *A&A*, 623:A21, Mar 2019. doi: 10.1051/0004-6361/201834394.
- H. Jeffreys. An Invariant Form for the Prior Probability in Estimation Problems. *Proceedings of the Royal Society of London Series A*, 186:453–461, Sep 1946. doi: 10.1098/rspa.1946.0056.
- H. Jeffreys. *Theory of Probability*. Oxford, Oxford, England, third edition, 1961.
- L. Jew. Measurements of diffuse galactic emission at 5 GHz with C-BASS. DPhil thesis, University of Oxford, 2017.

- L. Jew and R. D. P. Grumitt. The spectral index of polarized diffuse Galactic emission between 30 and 44 GHz. *MNRAS*, 495(1):578–593, May 2020. doi: 10.1093/mnras/staa1233.
- L. Jew et al. The C-Band All-Sky Survey (C-BASS): Simulated parametric fitting in single pixels in total intensity and polarization. *MNRAS*, Oct. 2019. doi: 10.1093/mnras/stz2697.
- J. L. Jonas, E. E. Baart, and G. D. Nicolson. The Rhodes/HartRAO 2326-MHz radio continuum survey. *MNRAS*, 297(4):977–989, July 1998. doi: 10.1046/j.1365-8711.1998.01367.x.
- M. E. Jones et al. The C-Band All-Sky Survey (C-BASS): design and capabilities. *MNRAS*, 480(3):3224–3242, Nov. 2018. doi: 10.1093/mnras/sty1956.
- A. Kaehler and G. Bradski. *Learning OpenCV 3: Computer Vision in C++ with the OpenCV Library*. O’Reilly Media, 2016. ISBN 9781491938003.
- M. Kamionkowski and E. D. Kovetz. The Quest for B Modes from Inflationary Gravitational Waves. *Annual Review of Astron and Astrophys*, 54:227–269, Sep 2016. doi: 10.1146/annurev-astro-081915-023433.
- M. Kamionkowski, A. Kosowsky, and A. Stebbins. Statistics of cosmic microwave background polarization. *Phys. Rev. D*, 55(12):7368–7388, June 1997. doi: 10.1103/PhysRevD.55.7368.
- K. Katahira. How hierarchical models improve point estimates of model parameters at the individual level. *Journal of Mathematical Psychology*, 73:37 – 58, 2016. ISSN 0022-2496. doi: <https://doi.org/10.1016/j.jmp.2016.03.007>.
- J. P. Kaufman et al. Self-calibration of BICEP1 three-year data and constraints on astrophysical polarization rotation. *Phys. Rev. D*, 89(6):062006, Mar. 2014. doi: 10.1103/PhysRevD.89.062006.
- B. Keating et al. Large Angular Scale Polarization of the Cosmic Microwave Background Radiation and the Feasibility of Its Detection. *ApJ*, 495:580–596, Mar 1998. doi: 10.1086/305312.
- B. G. Keating, M. Shimon, and A. P. S. Yadav. Self-calibration of Cosmic Microwave Background Polarization Experiments. *ApJL*, 762(2):L23, Jan. 2013. doi: 10.1088/2041-8205/762/2/L23.

- R. Keisler et al. A Measurement of the Damping Tail of the Cosmic Microwave Background Power Spectrum with the South Pole Telescope. *ApJ*, 743(1):28, Dec. 2011. doi: 10.1088/0004-637X/743/1/28.
- R. Keisler et al. Measurements of Sub-degree B-mode Polarization in the Cosmic Microwave Background from 100 Square Degrees of SPTpol Data. *ApJ*, 807(2):151, July 2015. doi: 10.1088/0004-637X/807/2/151.
- T. Kelsall et al. The COBE Diffuse Infrared Background Experiment Search for the Cosmic Infrared Background. II. Model of the Interplanetary Dust Cloud. *ApJ*, 508(1):44–73, Nov. 1998. doi: 10.1086/306380.
- R. Khatri. Data driven foreground clustering approach to component separation in multifrequency CMB experiments: a new Planck CMB map. *JCAP*, 2019(2):039, Feb 2019. doi: 10.1088/1475-7516/2019/02/039.
- O. G. King et al. The C-Band All-Sky Survey (C-BASS): design and implementation of the northern receiver. *MNRAS*, 438(3):2426–2439, Mar. 2014. doi: 10.1093/mnras/stt2359.
- L. Knox. Determination of inflationary observables by cosmic microwave background anisotropy experiments. *Phys. Rev. D*, 52(8):4307–4318, Oct. 1995. doi: 10.1103/PhysRevD.52.4307.
- D. Kodwani, D. Alonso, and P. Ferreira. The effect on cosmological parameter estimation of a parameter dependent covariance matrix. *The Open Journal of Astrophysics*, 2(1):3, Mar. 2019. doi: 10.21105/astro.1811.11584.
- A. Kogut et al. Microwave Emission at High Galactic Latitudes in the Four-Year DMR Sky Maps. *ApJ*, 464:L5, June 1996. doi: 10.1086/310072.
- A. Kogut et al. The Primordial Inflation Explorer (PIXIE): a nulling polarimeter for cosmic microwave background observations. *JCAP*, 2011(7):025, July 2011. doi: 10.1088/1475-7516/2011/07/025.
- E. Komatsu et al. Five-Year Wilkinson Microwave Anisotropy Probe Observations: Cosmological Interpretation. *ApJS*, 180(2):330–376, Feb. 2009. doi: 10.1088/0067-0049/180/2/330.
- A. Kosowsky. Introduction to microwave background polarization. *New Astronomy Reviews*, 43(2-4):157–168, July 1999. doi: 10.1016/S1387-6473(99)00009-3.

- J. M. Kovac et al. Detection of polarization in the cosmic microwave background using DASI. *Nature*, 420(6917):772–787, Dec. 2002. doi: 10.1038/nature01269.
- Y. Y. Kovalev et al. Survey of instantaneous 1-22 GHz spectra of 550 compact extragalactic objects with declinations from -30^{deg} to $+43^{deg}$. *A&AS*, 139:545–554, Nov 1999. doi: 10.1051/aas:1999406.
- N. Krachmalnicoff et al. S-PASS view of polarized Galactic synchrotron at 2.3 GHz as a contaminant to CMB observations. *A&A*, 618:A166, Oct. 2018. doi: 10.1051/0004-6361/201832768.
- L. M. Krauss, S. Dodelson, and S. Meyer. Primordial Gravitational Waves and Cosmology. *Science*, 328(5981):989, May 2010. doi: 10.1126/science.1179541.
- J. K. Kruschke. *Doing Bayesian data analysis : a tutorial with R and BUGS*. Academic Press, Burlington, MA, 2011. ISBN 9780123814852 0123814855.
- H. Kuehr et al. The 5-GHz strong source surveys. V - Survey of the area between declinations 70 and 90 deg. *AJ*, 86:854–863, June 1981. doi: 10.1086/112957.
- K. D. Lawson et al. Variations in the spectral index of the galactic radio continuum emission in the northern hemisphere. *MNRAS*, 225:307–327, Mar 1987. doi: 10.1093/mnras/225.2.307.
- B. J. Leimkuhler and S. Reich. *Simulating Hamiltonian dynamics*. Cambridge monographs on applied and computational mathematics. Cambridge Univ., Cambridge, 2004.
- E. M. Leitch et al. An Anomalous Component of Galactic Emission. *ApJ*, 486:L23–L26, Sept. 1997. doi:10.1086/310823.
- E. M. Leitch et al. Measurement of polarization with the Degree Angular Scale Interferometer. *Nature*, 420(6917):763–771, Dec. 2002. doi: 10.1038/nature01271.
- G. Lemaître. The Beginning of the World from the Point of View of Quantum Theory. *Nature*, 127(3210):706, May 1931. doi: 10.1038/127706b0.
- J. Lesgourgues. The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview. *arXiv e-prints*, art. arXiv:1104.2932, Apr. 2011.
- A. Lewis and A. Challinor. CAMB: Code for Anisotropies in the Microwave Background, Feb. 2011.

- A. Lewis, A. Challinor, and A. Lasenby. Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models. *ApJ*, 538(2):473–476, Aug 2000. doi: 10.1086/309179.
- A. D. Linde. A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B*, 108:389–393, Feb. 1982. doi: 10.1016/0370-2693(82)91219-9.
- H. Liu, J. Creswell, and P. Naselsky. E and B families of the Stokes parameters in the polarized synchrotron and thermal dust foregrounds. *JCAP*, 2018(5):059, May 2018. doi: 10.1088/1475-7516/2018/05/059.
- M. S. Longair. *High Energy Astrophysics*. 2011.
- M. López-Caniego et al. Comparison of filters for the detection of point sources in Planck simulations. *MNRAS*, 370:2047–2063, Aug. 2006. doi: 10.1111/j.1365-2966.2006.10639.x.
- T. Louis et al. The Atacama Cosmology Telescope: two-season ACTPol spectra and parameters. *JCAP*, 2017(6):031, June 2017. doi: 10.1088/1475-7516/2017/06/031.
- N. Macellari et al. Galactic foreground contributions to the 5-year Wilkinson Microwave Anisotropy Probe maps. *MNRAS*, 418:888–905, Dec 2011. doi: 10.1111/j.1365-2966.2011.19542.x.
- J. F. Macías-Pérez et al. Global Spectral Energy Distribution of the Crab Nebula in the Prospect of the Planck Satellite Polarization Calibration. *ApJ*, 711(1):417–423, Mar. 2010. doi: 10.1088/0004-637X/711/1/417.
- D. J. C. Mackay. *Information Theory, Inference and Learning Algorithms*. 2003.
- O. Mangoubi and A. Smith. Rapid Mixing of Hamiltonian Monte Carlo on Strongly Log-Concave Distributions. *arXiv e-prints*, art. arXiv:1708.07114, Aug. 2017.
- R. G. Mann et al. Observations of the Hubble Deep Field with the Infrared Space Observatory - IV. Association of sources with Hubble Deep Field galaxies. *MNRAS*, 289:482–489, Aug 1997. doi: 10.1093/mnras/289.2.482.
- C. C. Margossian et al. Hamiltonian Monte Carlo using an adjoint-differentiated Laplace approximation. *arXiv e-prints*, art. arXiv:2004.12550, Apr. 2020.

- E. Martínez-González et al. The performance of spherical wavelets to detect non-Gaussianity in the cosmic microwave background sky. *MNRAS*, 336(1):22–32, Oct 2002. doi: 10.1046/j.1365-8711.2002.05648.x.
- E. Martínez-González et al. Cosmic microwave background power spectrum estimation and map reconstruction with the expectation-maximization algorithm. *MNRAS*, 345(4):1101–1109, Nov. 2003. doi: 10.1046/j.1365-2966.2003.06885.x.
- M. Massardi et al. A model for the cosmological evolution of low-frequency radio sources. *MNRAS*, 404:532–544, May 2010. doi: 10.1111/j.1365-2966.2010.16305.x.
- J. C. Mather et al. A Preliminary Measurement of the Cosmic Microwave Background Spectrum by the Cosmic Background Explorer (COBE) Satellite. *ApJ*, 354:L37, May 1990. doi: 10.1086/185717.
- J. C. Mather et al. Measurement of the cosmic microwave background spectrum by the COBE FIRAS instrument. *ApJ*, 420:439–444, Jan. 1994. doi: 10.1086/173574.
- T. Matsumura et al. Mission Design of LiteBIRD. *Journal of Low Temperature Physics*, 176(5-6):733–740, Sept. 2014. doi: 10.1007/s10909-013-0996-1.
- P. Mertsch and S. Sarkar. Loops and spurs: the angular power spectrum of the Galactic synchrotron background. *JCAP*, 2013(6):041, June 2013. doi: 10.1088/1475-7516/2013/06/041.
- D. Mesa et al. Polarization properties of extragalactic radio sources and their contribution to microwave polarization fluctuations. *A&A*, 396:463–471, Dec 2002. doi: 10.1051/0004-6361:20021392.
- M. Millea, E. Anderes, and B. D. Wandelt. Bayesian delensing delight: sampling-based inference of the primordial CMB and gravitational lensing. *arXiv e-prints*, art. arXiv:2002.00965, Feb. 2020.
- M. G. Mingaliev et al. Observations of the bright radio sources in the North Celestial Pole region at the RATAN-600 radio telescope. *A&A*, 370:78–86, Apr. 2001. doi: 10.1051/0004-6361:20010215.
- M. G. Mingaliev et al. Spectral characteristics of radio sources near the North Celestial Pole. *Astronomy Reports*, 51:343–363, May 2007. doi: 10.1134/S1063772907050010.

- S. Mitton and M. Ryle. High resolution observations of Cygnus A at 2.7 GHz and 5 GHz. *MNRAS*, 146:221, Jan 1969. doi: 10.1093/mnras/146.3.221.
- M. A. Miville-Deschênes et al. Separation of anomalous and synchrotron emissions using WMAP polarization data. *A&A*, 490(3):1093–1102, Nov 2008. doi: 10.1051/0004-6361:200809484.
- C. C. Monnahan, J. T. Thorson, and T. A. Branch. Faster estimation of bayesian models in ecology using hamiltonian monte carlo. *Methods in Ecology and Evolution*, 8(3):339–348, 2017. doi: 10.1111/2041-210X.12681.
- M. J. Mortonson and U. Seljak. A joint analysis of Planck and BICEP2 B modes including dust polarization uncertainty. *JCAP*, 2014(10):035, Oct. 2014. doi: 10.1088/1475-7516/2014/10/035.
- S. K. Næss and T. Louis. Lensing simulations by Taylor expansion — not so inefficient after all. *Journal of Cosmology and Astro-Particle Physics*, 2013:001, Sept. 2013. doi: 10.1088/1475-7516/2013/09/001.
- NASA. LAMBDA - CMB Experiments. <https://lambda.gsfc.nasa.gov/product/>, January 2020.
- R. Neal. MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 06 2012. doi: 10.1201/b10905-6.
- C. B. Netterfield et al. A Measurement of the Angular Power Spectrum of the Anisotropy in the Cosmic Microwave Background. *ApJ*, 474(1):47–66, Jan. 1997. doi: 10.1086/303438.
- E. Nieppola et al. Long-Term Variability of Radio-Bright BL Lacertae Objects. *AJ*, 137(6):5022–5036, June 2009. doi: 10.1088/0004-6256/137/6/5022.
- N. Oppermann et al. An improved map of the Galactic Faraday sky. *A&A*, 542:A93, June 2012. doi: 10.1051/0004-6361/201118526.
- S. Padin et al. First Intrinsic Anisotropy Observations with the Cosmic Background Imager. *ApJ*, 549(1):L1–L5, Mar. 2001. doi: 10.1086/319142.
- S. Padin et al. The Cosmic Background Imager. *PASP*, 114(791):83–97, Jan. 2002. doi: 10.1086/324786.

- G. Papamakarios et al. Normalizing Flows for Probabilistic Modeling and Inference. *arXiv e-prints*, art. arXiv:1912.02762, Dec. 2019.
- I. I. K. Pauliny-Toth et al. The 5 GHz strong source surveys. IV. Survey of the area between declination 35 and 70 degrees and summary of source counts, spectra and optical identifications. *AJ*, 83:451–474, May 1978. doi: 10.1086/112223.
- J. A. Peacock. *Cosmological Physics*. 1999.
- F. Pedregosa et al. Scikit-learn: Machine learning in python. *J. Mach. Learn. Res.*, 12:2825–2830, Nov. 2011. ISSN 1532-4435.
- P. J. E. Peebles and J. T. Yu. Primeval Adiabatic Perturbation in an Expanding Universe. *ApJ*, 162:815, Dec. 1970. doi: 10.1086/150713.
- H. V. Peiris et al. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Implications For Inflation. *ApJS*, 148(1):213–231, Sept. 2003. doi: 10.1086/377228.
- A. A. Penzias and R. W. Wilson. A Measurement of Excess Antenna Temperature at 4080 Mc/s. *ApJ*, 142:419–421, July 1965. doi: 10.1086/148307.
- Planck Collaboration et al. Planck intermediate results. IX. Detection of the Galactic haze with Planck. *A&A*, 554:A139, June 2013. doi: 10.1051/0004-6361/201220271.
- Planck Collaboration et al. Planck 2013 results. XII. Diffuse component separation. *A&A*, 571:A12, Nov. 2014a. doi: 10.1051/0004-6361/201321580.
- Planck Collaboration et al. Planck 2013 results. XXVIII. The Planck Catalogue of Compact Sources. *A&A*, 571:A28, Nov. 2014b. doi: 10.1051/0004-6361/201321524.
- Planck Collaboration et al. Planck intermediate results. XIV. Dust emission at millimetre wavelengths in the Galactic plane. *A&A*, 564:A45, Apr. 2014c. doi: 10.1051/0004-6361/201322367.
- Planck Collaboration et al. Planck intermediate results. XXX. The angular power spectrum of polarized dust emission at intermediate and high Galactic latitudes. *A&A*, 586:A133, Feb. 2016a. doi: 10.1051/0004-6361/201425034.
- Planck Collaboration et al. Planck 2015 results. VI. LFI mapmaking. *A&A*, 594:A6, Sep 2016b. doi: 10.1051/0004-6361/201525813.

- Planck Collaboration et al. Planck 2015 results. IX. Diffuse component separation: CMB maps. *A&A*, 594:A9, Sept. 2016c. doi: 10.1051/0004-6361/201525936.
- Planck Collaboration et al. Planck 2015 results. X. Diffuse component separation: Foreground maps. *A&A*, 594:A10, Sept. 2016d. doi: 10.1051/0004-6361/201525967.
- Planck Collaboration et al. Planck 2015 results. XIII. Cosmological parameters. *A&A*, 594:A13, Sept. 2016e. doi: 10.1051/0004-6361/201525830.
- Planck Collaboration et al. Planck 2015 results. XXVI. The Second Planck Catalogue of Compact Sources. *A&A*, 594:A26, Sept. 2016f. doi: 10.1051/0004-6361/201526914.
- Planck Collaboration et al. Planck 2018 results. I. Overview and the cosmological legacy of Planck. *arXiv e-prints*, art. arXiv:1807.06205, July 2018a.
- Planck Collaboration et al. Planck 2018 results. IV. Diffuse component separation. *arXiv e-prints*, art. arXiv:1807.06208, July 2018b.
- Planck Collaboration et al. Planck 2018 results. VI. Cosmological parameters. *arXiv e-prints*, art. arXiv:1807.06209, July 2018c.
- Planck Collaboration et al. Planck 2018 results. X. Constraints on inflation. *arXiv e-prints*, art. arXiv:1807.06211, July 2018d.
- Planck Collaboration et al. Planck 2018 results. V. CMB power spectra and likelihoods. *arXiv e-prints*, art. arXiv:1907.12875, July 2019.
- P. Platania et al. Full sky study of diffuse Galactic emission at decimeter wavelengths. *A&A*, 410:847–863, Nov 2003. doi: 10.1051/0004-6361:20031125.
- Polarbear Collaboration et al. A Measurement of the Cosmic Microwave Background B-mode Polarization Power Spectrum at Sub-degree Scales with POLARBEAR. *ApJ*, 794(2):171, Oct. 2014. doi: 10.1088/0004-637X/794/2/171.
- N. G. Polson and J. G. Scott. On the half-cauchy prior for a global scale parameter. *Bayesian Anal.*, 7(4):887–902, 12 2012. doi: 10.1214/12-BA730.
- A. M. Polyakov. Particle spectrum in quantum field theory. *Soviet Journal of Experimental and Theoretical Physics Letters*, 20:194, Sept. 1974.

- G. G. Pooley and M. Ryle. The extension of the number-flux density relation for radio sources to very small flux densities. *MNRAS*, 139:515, 1968. doi: 10.1093/mnras/139.4.515.
- C. Pryke et al. Cosmological Parameter Extraction from the First Season of Observations with the Degree Angular Scale Interferometer. *ApJ*, 568(1):46–51, Mar. 2002. doi: 10.1086/338880.
- G. Puglisi et al. Forecasting the Contribution of Polarized Extragalactic Radio Sources in CMB Observations. *ApJ*, 858:85, May 2018. doi: 10.3847/1538-4357/aab3c7.
- A. C. S. Readhead et al. Polarization Observations with the Cosmic Background Imager. *Science*, 306(5697):836–844, Oct. 2004. doi: 10.1126/science.1105598.
- P. Reich and W. Reich. A radio continuum survey of the northern sky at 1420 MHz. II. *A&AS*, 63:205, Feb. 1986.
- P. Reich and W. Reich. A map of spectral indices of the galactic radio continuum emission between 408 MHz and 1420 MHz for the entire northern sky. *A&AS*, 74: 7–23, Jul 1988.
- P. Reich, W. Reich, and J. C. Testori. Spectral Index Variations of Galactic Emission. In B. Uyaniker, W. Reich, and R. Wielebinski, editors, *The Magnetized Interstellar Medium*, pages 63–68, Feb. 2004.
- M. Remazeilles. Component separation for future CMB B-mode satellites. *arXiv e-prints*, art. arXiv:1806.01026, June 2018.
- M. Remazeilles, J. Delabrouille, and J.-F. Cardoso. Foreground component separation with generalized Internal Linear Combination. *MNRAS*, 418(1):467–476, Nov. 2011. doi: 10.1111/j.1365-2966.2011.19497.x.
- M. Remazeilles et al. An improved source-subtracted and destriped 408-MHz all-sky map. *MNRAS*, 451:4311–4327, Aug. 2015. doi: 10.1093/mnras/stv1274.
- M. Remazeilles et al. Sensitivity and foreground modelling for large-scale cosmic microwave background B-mode polarization satellite missions. *MNRAS*, 458(2): 2032–2050, May 2016. doi: 10.1093/mnras/stw441.

- M. Remazeilles et al. Exploring cosmic origins with CORE: B-mode component separation. *JCAP*, 2018:023, Apr. 2018. doi: 10.1088/1475-7516/2018/04/023.
- G. A. Richter. Search for Optical Identifications in the 5C3 Radio Survey. II. Statistical Treatment and Results. *Astronomische Nachrichten*, 296(2):65, Jan 1975. doi: 10.1002/asna.19752960203.
- R. E. Rutledge et al. The Quiescent X-Ray Spectrum of the Neutron Star in Centaurus X-4 Observed with Chandra/ACIS-S. *ApJ*, 551:921–928, Apr 2001. doi: 10.1086/320247.
- G. B. Rybicki and A. P. Lightman. *Radiative Processes in Astrophysics*. 1986.
- J. Salvatier, T. V. Wiecki, and C. Fonnesbeck. Probabilistic programming in python using pymc3. *PeerJ Computer Science*, 2:e55, 2016.
- J. T. Sayre et al. Measurements of B-mode Polarization of the Cosmic Microwave Background from 500 Square Degrees of SPTpol Data. *arXiv e-prints*, art. arXiv:1910.05748, Oct. 2019.
- P. F. Scott et al. First results from the Very Small Array - III. The cosmic microwave background power spectrum. *MNRAS*, 341(4):1076–1083, June 2003. doi: 10.1046/j.1365-8711.2003.06354.x.
- Y. Sekimoto et al. Concept design of the LiteBIRD satellite for CMB B-mode polarization. In *Proc. SPIE*, volume 10698 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 106981Y, Aug 2018. doi: 10.1117/12.2313432.
- U. Seljak. A Two-Fluid Approximation for Calculating the Cosmic Microwave Background Anisotropies. *ApJ*, 435:L87, Nov. 1994. doi: 10.1086/187601.
- U. Seljak and M. Zaldarriaga. A Line-of-Sight Integration Approach to Cosmic Microwave Background Anisotropies. *ApJ*, 469:437, Oct. 1996. doi: 10.1086/177793.
- S. Shandera et al. Probing the origin of our Universe through cosmic microwave background constraints on gravitational waves. *Bulletin of the AAS*, 51(3):338, May 2019.
- E. Shirokoff et al. The south pole telescope sz-receiver detectors. *IEEE Transactions on Applied Superconductivity*, 19(3):517–519, 2009.

- J. Silk. Cosmic Black-Body Radiation and Galaxy Formation. *ApJ*, 151:459, Feb. 1968. doi: 10.1086/149449.
- Simons Observatory Collaboration et al. The Simons Observatory: science goals and forecasts. *JCAP*, 2019(2):056, Feb. 2019. doi: 10.1088/1475-7516/2019/02/056.
- D. Simpson et al. Penalising model component complexity: A principled, practical approach to constructing priors. *Statist. Sci.*, 32(1):1–28, 02 2017. doi: 10.1214/16-STS576.
- G. F. Smoot et al. Structure in the COBE Differential Microwave Radiometer First-Year Maps. *ApJ*, 396:L1, Sept. 1992. doi: 10.1086/186504.
- Stan Development Team. *Stan Modeling Language User’s Guide and Reference Manual, Version 1.0*. 2012.
- A. A. Starobinsky. Dynamics of phase transition in the new inflationary universe scenario and generation of perturbations. *Physics Letters B*, 117(3-4):175–178, Nov. 1982. doi: 10.1016/0370-2693(82)90541-X.
- M. A. Stevenson. *Observational and theoretical advances in cosmological foreground emission*. PhD thesis, California Institute of Technology, Jan 2014.
- M. A. Stevenson et al. The C-Band All-Sky Survey (C-BASS): digital backend for the northern survey. *MNRAS*, 484(4):5377–5388, Apr. 2019. doi: 10.1093/mnras/stz313.
- R. Stompor, J. Errard, and D. Poletti. Forecasting performance of CMB experiments in the presence of complex foreground contaminations. *Phys. Rev. D*, 94(8):083526, Oct 2016. doi: 10.1103/PhysRevD.94.083526.
- R. Stompor et al. Maximum likelihood algorithm for parametric component separation in cosmic microwave background experiments. *MNRAS*, 392(1):216–232, Jan. 2009. doi: 10.1111/j.1365-2966.2008.14023.x.
- M. Su and D. P. Finkbeiner. Evidence for Gamma-Ray Jets in the Milky Way. *ApJ*, 753(1):61, July 2012. doi: 10.1088/0004-637X/753/1/61.
- M. Su, T. R. Slatyer, and D. P. Finkbeiner. Giant Gamma-ray Bubbles from Fermi-LAT: Active Galactic Nucleus Activity or Bipolar Galactic Wind? *ApJ*, 724(2):1044–1082, Dec. 2010. doi: 10.1088/0004-637X/724/2/1044.

- R. A. Sunyaev and Y. B. Zeldovich. The Observations of Relic Radiation as a Test of the Nature of X-Ray Radiation from the Clusters of Galaxies. *Comments on Astrophysics and Space Physics*, 4:173, Nov. 1972.
- W. Sutherland and W. Saunders. On the likelihood ratio for source identification. *MNRAS*, 259:413–420, Dec 1992. doi: 10.1093/mnras/259.3.413.
- D. Sutton et al. Fast and precise map-making for massively multi-detector CMB experiments. *MNRAS*, 407:1387–1402, Sept. 2010. doi: 10.1111/j.1365-2966.2010.16954.x.
- D. S. Swetz et al. Overview of the Atacama Cosmology Telescope: Receiver, Instrumentation, and Telescope Systems. *ApJS*, 194(2):41, June 2011. doi: 10.1088/0067-0049/194/2/41.
- I. Szapudi et al. Fast Cosmic Microwave Background Analyses via Correlation Functions. *ApJL*, 548(2):L115–L118, Feb. 2001. doi: 10.1086/319105.
- G. 't Hooft. Magnetic monopoles in unified gauge theories. *Nuclear Physics B*, 79(2): 276–284, Sept. 1974. doi: 10.1016/0550-3213(74)90486-6.
- A. C. Taylor et al. The radio source counts at 15 GHz and their implications for cm-wave CMB imaging. *MNRAS*, 327:L1–L4, Oct. 2001. doi: 10.1046/j.1365-8711.2001.04877.x.
- A. C. Taylor et al. First results from the Very Small Array - II. Observations of the cosmic microwave background. *MNRAS*, 341(4):1066–1075, June 2003. doi: 10.1046/j.1365-8711.2003.06493.x.
- A. R. Taylor, J. M. Stil, and C. Sunstrum. A Rotation Measure Image of the Sky. *ApJ*, 702(2):1230–1236, Sept. 2009. doi: 10.1088/0004-637X/702/2/1230.
- J. F. Taylor, M. A. J. Ashdown, and M. P. Hobson. Fast optimal CMB power spectrum estimation with Hamiltonian sampling. *MNRAS*, 389(3):1284–1292, Sep 2008. doi: 10.1111/j.1365-2966.2008.13630.x.
- M. Tegmark. How to measure CMB power spectra without losing information. *Phys. Rev. D*, 55(10):5895–5907, May 1997. doi: 10.1103/PhysRevD.55.5895.
- M. Tegmark. Removing Real-World Foregrounds from Cosmic Microwave Background Maps. *ApJ*, 502(1):1–6, July 1998. doi: 10.1086/305905.

- M. Tegmark and A. de Oliveira-Costa. Removing Point Sources from Cosmic Microwave Background Maps. *ApJ*, 500:L83–L86, June 1998. doi: 10.1086/311410.
- M. Tegmark and A. de Oliveira-Costa. How to measure CMB polarization power spectra without losing information. *Phys. Rev. D*, 64(6):063001, Sept. 2001. doi: 10.1103/PhysRevD.64.063001.
- M. Tegmark and G. Efstathiou. A method for subtracting foregrounds from multifrequency CMB sky maps**. *MNRAS*, 281:1297–1314, Aug. 1996. doi: 10.1093/mnras/281.4.1297.
- C. Tello et al. The 2.3 GHz continuum survey of the GEM project. *A&A*, 556:A1, Aug. 2013. doi: 10.1051/0004-6361/20079306.
- J. C. Testori, P. Reich, and W. Reich. A fully sampled $\lambda 21$ cm linear polarization survey of the southern sky. *A&A*, 484(3):733–742, June 2008. doi: 10.1051/0004-6361:20078842.
- The CORE Collaboration et al. CORE (Cosmic Origins Explorer) A White Paper. *arXiv e-prints*, art. arXiv:1102.2181, Feb. 2011.
- B. Thorne et al. The Python Sky Model: software for simulating the Galactic microwave sky. *MNRAS*, 469:2821–2833, Aug. 2017. doi: 10.1093/mnras/stx949.
- B. Thorne et al. Removal of Galactic foregrounds for the Simons Observatory primordial gravitational wave search. *arXiv e-prints*, art. arXiv:1905.08888, May 2019.
- M. Tristram et al. The CMB temperature power spectrum from an improved analysis of the Archeops data. *A&A*, 436(3):785–797, June 2005a. doi: 10.1051/0004-6361:20042416.
- M. Tristram et al. XSPECT, estimation of the angular power spectrum by computing cross-power spectra with analytical error bars. *MNRAS*, 358(3):833–842, Apr. 2005b. doi: 10.1111/j.1365-2966.2005.08760.x.
- M. Tucci et al. High-frequency predictions for number counts and spectral properties of extragalactic radio sources. New evidence of a break at mm wavelengths in spectra of bright blazar sources. *A&A*, 533:A57, Sept. 2011. doi: 10.1051/0004-6361/201116972.

- S. Vanneste et al. Quadratic estimator for CMB cross-correlation. *Phys. Rev. D*, 98 (10):103526, Nov. 2018. doi: 10.1103/PhysRevD.98.103526.
- M. Vidal et al. Polarized radio filaments outside the Galactic plane. *MNRAS*, 452: 656–675, Sep 2015. doi: 10.1093/mnras/stv1328.
- P. Vielva et al. Point source detection using the Spherical Mexican Hat Wavelet on simulated all-sky Planck maps. *MNRAS*, 344:89–104, Sept. 2003. doi: 10.1046/j.1365-8711.2003.06792.x.
- S. Wagner-Carena et al. A novel CMB component separation method: hierarchical generalized morphological component analysis. *MNRAS*, 494(1):1507–1529, Mar. 2020. doi: 10.1093/mnras/staa744.
- E. M. Waldram et al. 9C: a survey of radio sources at 15 GHz with the Ryle Telescope. *MNRAS*, 342(3):915–925, Jul 2003. doi: 10.1046/j.1365-8711.2003.06628.x.
- E. M. Waldram et al. Some estimates of the source counts at Planck Surveyor frequencies, using the 9C survey at 15 GHz. *MNRAS*, 379(4):1442–1452, Aug 2007. doi: 10.1111/j.1365-2966.2007.12025.x.
- B. D. Wandelt and F. K. Hansen. Fast, exact CMB power spectrum estimation for a certain class of observational strategies. *Phys. Rev. D*, 67(2):023001, Jan. 2003. doi: 10.1103/PhysRevD.67.023001.
- B. D. Wandelt, E. Hivon, and K. M. Gorski. Cosmic microwave background anisotropy power spectrum statistics for high precision cosmology. *arXiv e-prints*, art. astro-ph/9808292, Aug. 1998.
- B. D. Wandelt, D. L. Larson, and A. Lakshminarayanan. Global, exact cosmic microwave background data analysis using Gibbs sampling. *Phys. Rev. D*, 70(8): 083511, Oct 2004. doi: 10.1103/PhysRevD.70.083511.
- L. Wang et al. HerMES: point source catalogues from Herschel-SPIRE observations II. *MNRAS*, 444:2870–2883, Nov 2014. doi: 10.1093/mnras/stu1569.
- R. A. Watson et al. First results from the Very Small Array - I. Observational methods. *MNRAS*, 341(4):1057–1065, June 2003. doi: 10.1046/j.1365-8711.2003.06338.x.

- J. L. Weiland et al. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Planets and Celestial Calibration Sources. *ApJS*, 192(2):19, Feb 2011. doi: 10.1088/0067-0049/192/2/19.
- M. J. White. Anisotropies in the CMB. In *Proceedings, American Physical Society (APS) Meeting of the Division of Particles and Fields (DPF 99): Los Angeles, CA, January 5-9, 1999*, 1999.
- M. Wolleben et al. An absolutely calibrated survey of polarized emission from the northern sky at 1.4 GHz. Observations and data reduction. *A&A*, 448(1):411–424, Mar. 2006. doi: 10.1051/0004-6361:20053851.
- M. Wolleben et al. GMIMS: the Global Magneto-Ionic Medium Survey. In K. G. Strassmeier, A. G. Kosovichev, and J. E. Beckman, editors, *Cosmic Magnetic Fields: From Planets, to Stars and Galaxies*, volume 259 of *IAU Symposium*, pages 89–90, Apr. 2009. doi: 10.1017/S1743921309030117.
- M. Wolleben et al. Rotation Measure Synthesis of Galactic Polarized Emission with the DRAO 26-m Telescope. *AJ*, 139(4):1681–1690, Apr. 2010. doi: 10.1088/0004-6256/139/4/1681.
- A. E. Wright et al. The Parkes-MIT-NRAO (PMN) surveys. 2: Source catalog for the southern survey (delta greater than -87.5 deg and less than -37 deg). *ApJS*, 91: 111–308, Mar. 1994a. doi: 10.1086/191939.
- A. E. Wright et al. The Parkes-MIT-NRAO (PMN) Surveys. VIII. Source Catalog for the Zenith Survey (-37.0 degrees less than delta less than -29.0 degrees). *ApJS*, 103:145, Mar. 1996a. doi: 10.1086/192272.
- A. E. Wright et al. VizieR Online Data Catalog: The Parkes-MIT-NRAO 4.85GHz (PMN) Surveys (Griffith+ 1993-1996). *VizieR Online Data Catalog*, 8038, Oct. 1996b.
- E. L. Wright et al. Interpretation of the COBE FIRAS CMBR spectrum. *ApJ*, 420: 450–456, Jan. 1994b. doi: 10.1086/173576.
- H.-Y. K. Yang, M. Ruszkowski, and E. Zweibel. The Fermi bubbles: gamma-ray, microwave and polarization signatures of leptonic AGN jets. *MNRAS*, 436(3): 2734–2746, dec 2013. ISSN 1365-2966. doi: 10.1093/mnras/stt1772.

M. Zaldarriaga and U. Seljak. All-sky analysis of polarization in the microwave background. *Phys. Rev. D*, 55(4):1830–1840, Feb. 1997. doi: 10.1103/PhysRevD.55.1830.

M. Zaldarriaga, S. R. Furlanetto, and L. Hernquist. 21 Centimeter Fluctuations from Cosmic Gas at High Redshifts. *ApJ*, 608(2):622–635, June 2004. doi: 10.1086/386327.