

POLARIZED DUST EMISSION AND THE MORPHOLOGY OF THE
INTERSTELLAR MEDIUM

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

George Halal

August 2024

© 2024 by George Halal. All Rights Reserved.
Re-distributed by Stanford University under license with the author.



This work is licensed under a Creative Commons Attribution-Noncommercial 3.0 United States License.
<http://creativecommons.org/licenses/by-nc/3.0/us/>

This dissertation is online at: <https://purl.stanford.edu/rn650nf7483>

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Susan Clark, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Patricia Burchat

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Chao-Lin Kuo

Approved for the Stanford University Committee on Graduate Studies.

Stacey F. Bent, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format.

Abstract

The interstellar medium (ISM) is a complex, multiphase system where various components, such as gas, dust, cosmic rays, and magnetic fields, are intricately interconnected. The same astrophysical phenomena that influence one component of the ISM often have far-reaching effects on others, creating a rich tapestry of interactions. This interconnectedness means that observations of one ISM component can provide valuable insights into the properties and behavior of others, offering multiple avenues to probe and understand the ISM's structure and evolution. This Thesis explores this concept by focusing on the relationship between the three-dimensional morphology of the ISM and polarized dust emission, which holds important implications for cosmic microwave background (CMB) studies and our understanding of astrophysical processes. By leveraging ancillary datasets that trace the 3D structure of the ISM, this Thesis provides novel insights into the characterization and modeling of polarized dust emission.

We investigate how the geometry of the Local Bubble and the complexity of the dust distribution along the line of sight affect the observed dust polarization statistics. Our

findings indicate that the extended 3D dust distribution, beyond just the Local Bubble, plays an important role in determining the observed polarization patterns. We utilize 3D neutral hydrogen ($\text{H}\alpha$) data, another tracer of the ISM, to characterize and model the polarized dust emission in the rest of the thesis. $\text{H}\alpha$ filaments trace the local 3D magnetic field structure, which polarizes the dust emission perpendicular to its orientation. We introduce a new approach for characterizing Galactic dust filaments by correlating BICEP/Keck and Planck data with 3D polarization templates based on $\text{H}\alpha$ observations. This method proves effective in detecting polarized dust emission at frequencies as low as 95 GHz and isolating contributions to the polarized dust emission from the Milky Way and the Magellanic Stream I and characterizing them. We further improve $\text{H}\alpha$ -based polarization templates through the development of the Spherical Rolling Hough Transform (Spherical RHT) algorithm, which efficiently quantifies filamentary structures on the sphere. We investigate how the morphology of magnetically aligned dusty filaments affects the polarized dust emission. This work demonstrates that the thinnest resolved filaments are most aligned with the magnetic field and how different filament geometries affect different polarization patterns. Finally, polarized dust emission is the dominant foreground for CMB polarization studies at high frequencies and its structure at small scales is unknown. Therefore, we introduce a novel approach to generating a high-resolution, non-Gaussian foreground model for CMB polarization studies using transformer-based deep learning techniques. This model fuses information from various sources to predict the small-scale dust structures.

Acknowledgments

First and foremost, thank you to my advisor, Susan Clark. You have absolutely spoiled me by being the most ideal PhD, life, and career mentor I could have ever asked for. You have provided the perfect balance of pushing me to be more productive and grow while being very supportive and understanding along the way. I deeply appreciate your going out of your way to respond to my Slack messages within an hour, no matter what day or time I send them, while not expecting me to do the same. I have much more to say, but I don't want the rest of the people I am acknowledging to feel jealous.

Next, I would like to thank my co-advisor, Chao-Lin Kuo, and my third reading committee member, Patricia Burchat. I have learned a great deal from Chao-Lin's seemingly boundless creative ideas and intuition for strategic project direction. Although Pat is not my advisor, her advice was invaluable when I felt lost during my first year at Stanford. I also had the pleasure of being her teaching assistant, which was my first teaching experience and a tremendous learning opportunity. I am also grateful to the other members of my thesis committee: Steven Allen and Gordon Wetzstein.

My sincere thanks go to all current and former members of the Clark and Kuo research groups and the BICEP/Keck collaboration. I am particularly indebted to my collaborators Dominic Beck, Ari Cukierman, Mehrnoosh Tahani, and Alejandro Dobles, without whom this thesis would not have been possible.

I would not have made it here without my undergraduate advisor, Rosi Reed. It's hard to overstate how deeply grateful I am for the impact you have had on my career as a scientist. You took me into your lab when I had no experience, were patient with me, taught me how to code and conduct research, and believed in me enough to give me challenging graduate-level projects. You were also the reason I was able to gain invaluable research experience in different research groups that significantly benefited my growth and graduate school applications. Also, a huge thank you to Professors Helen Caines, Sera Cremonini, and Michael Lisa for everything you have taught me.

I am very grateful to my family. Although stereotypical Syrian parents might love nothing more for their son than to be a doctor, lawyer, or worst case, an engineer, my parents love nothing more for me than to be happy. They knew they would not be able to see me as often if I moved to the U.S. for my education, and they still supported my decision. Little did they know that, due to various reasons, I would not be able to see them for 7 years. They had also never heard of physics as a career option, yet they supported me in choosing it once they saw my passion towards it. Thank you for being so supportive and caring.

To my friends, my chosen family: you are the source of my joy and resilience. I want

nothing more than to name you individually because your support and friendship have been integral to my journey, but I do not want to publicly pick favorites or risk forgetting anyone. You all know who you are and I hope to have lots of opportunities to show you my appreciation.

Finally, thank you, Vivi. Thank you for taking care of me whenever I needed it—and even when I didn't realize I did. Thank you for your patience during the stressful times when I was working hard throughout all my waking hours for months on end. Thank you for creating a fun and loving environment. And most of all, thank you for always ensuring that I never feel far from home.

Contents

Abstract	iv
Acknowledgments	vi
1 Introduction	1
1.1 3D Structure of the Magnetic ISM	3
1.1.1 Dust Mapping	3
1.1.2 Structure of the Local ISM	7
1.1.3 Neutral Hydrogen Gas	10
1.2 Dust Polarization and Magnetic Fields	14
1.3 The ISM as a Cosmological Foreground	22
1.4 Outline	25
2 Imprints of the Local Bubble and Dust Complexity on Polarized Dust Emission	27
2.1 Paper Status and External Contributions	28

2.2	Introduction	29
2.3	Data	32
2.3.1	Planck Data Products	32
2.3.2	3D Dust Maps and Local Bubble Geometries	35
2.3.3	Galactic Faraday Rotation Measure	37
2.4	No Imprint of the Local Bubble on the Dust Polarization Fraction	38
2.4.1	Motivation	38
2.4.2	Testing the Dependence of the Dust Polarization Fraction on the Magnetic Inclination Angle	42
2.4.3	Magnetic Fields in the Local Bubble Wall	46
2.4.4	No Detected Imprint of the Local Bubble Wall on the Dust Polarization Fraction	49
2.5	Imprint of Dust Complexity on Dust Polarization Fraction	50
2.5.1	Sightline Selection	51
2.5.2	Dust Complexity	54
2.5.3	Nearest-Neighbor Matching	58
2.5.4	Statistical Tests	61
2.5.5	Results	63
2.5.6	Validation	68
2.6	Discussion and Conclusions	70

2.7	Analysis Variations	76
3	Characterizing Dust Polarization Through Correlations with H I	80
3.1	Paper Status and External Contributions	81
3.2	Introduction	82
3.3	Data	87
3.3.1	Millimeter-wave Polarization	87
3.3.2	Neutral Hydrogen Emission	87
3.4	Methodology	88
3.4.1	Convolutional Rolling Hough Transform	88
3.4.2	RHT Transfer Function	90
3.4.3	BICEP/Keck and Planck Simulations Including Filamentary Dust	94
3.4.4	Cross Spectra	98
3.4.5	Covariance Matrices	99
3.4.6	Statistical Tests	101
3.4.7	Parameter Estimation	104
3.5	Velocity Decomposition	108
3.6	Results and Discussion	112
3.6.1	Tuning and Improving the RHT Model	112
3.6.2	Filamentary Polarization in the Local ISM	116
3.6.3	Frequency Decorrelation and the Polarized Dust SED	117

3.6.4	Individual Frequency Band Contribution	125
3.6.5	Polarized Dust in Magellanic Stream I	129
3.7	Summary and Outlook	134
3.8	Acknowledgments	137
3.9	Uncertainty Calculation	138
3.10	Analysis Variations	140
4	Filamentary Dust Polarization and the Morphology of Neutral Hydrogen Structures	145
4.1	Paper Status and External Contributions	146
4.2	Motivation	147
4.3	Data	150
4.3.1	Dust Emission	150
4.3.2	Neutral Hydrogen	151
4.4	H I-based Dust Polarization Prediction	152
4.4.1	RHT-based Angle Determination	152
4.4.2	Hessian-based Angle Determination	153
4.4.3	H I-based Polarization Template Construction	155
4.5	Improvements in H I-based Dust Polarization Prediction	158
4.5.1	Velocity Selection	158
4.5.2	Dust Map and Mask Comparisons	162

4.5.3	Effects of Resolution and Data Artifacts	164
4.6	Spherical Rolling Hough Transform	170
4.6.1	Spherical Convolutions	173
4.6.2	The Algorithm	175
4.6.3	Comparison with the RHT	177
4.7	Filament Morphologies	179
4.7.1	Morphological Parameter Space Exploration	179
4.7.2	Spherical RHT- and Hessian-based Template Comparison	184
4.8	Conclusions	192
4.9	Acknowledgments	196
5	Transformer-Based Polarized Dust Emission Super-Resolution	198
5.1	Paper Status and External Contributions	199
5.2	Introduction	199
5.3	Data	203
5.3.1	Planck Data Products	203
5.3.2	H I-based Dust Polarization Templates	204
5.4	Pre-processing	206
5.4.1	Masks	207
5.4.2	Patch Projections	208
5.4.3	Training, Validation, and Testing Data	212

5.4.4	Normalization	214
5.5	Model	215
5.5.1	Super-Resolution Techniques	215
5.5.2	Architecture	217
5.6	Results	223
5.7	Conclusions, Limitations, and Future Work	226
5.8	Acknowledgements	228
6	Conclusion	231
6.1	Summary of Results	231
6.2	Outlook	234

List of Tables

2.1	Results obtained from varying both the dendrogram parameters (Δ_n and $\#_{\text{voxels}}$) used to identify the density peaks in the line-of-sight dust distribution and from varying the upper and lower thresholds in \mathcal{N}_c used to divide the sightlines into low- and high-complexity groups. The results of the main analysis are bolded in the first row. The mean and standard deviation in the $\mu_{\Delta p_{353}}$ and $\sigma_{\Delta p_{353}}$ columns are over the 12 posterior samples of the Edenhofer et al. (2023) maps.	77
2.2	Results obtained from the same variations described in Table 2.1, where the sightline matching over the two complexity groups in this case includes the absolute value of the Galactic latitude as an additional confounding variable. The mean and standard deviation in the $\mu_{\Delta p_{353}}$ and $\sigma_{\Delta p_{353}}$ columns are over the 12 posterior samples of the Edenhofer et al. (2023) maps. . . .	78

2.3 Results obtained from sampling 1,000 sightlines from each of the $\mathcal{N}_c \leq 1.4$ and $\mathcal{N}_c \geq 2.3$ complexity groups before matching. The different rows are for different variations in the dendrogram parameters (Δ_n and $\#\text{voxels}$) used in identifying the density peaks in the line-of-sight dust distribution. The mean and standard deviation in the $\mu_{\Delta p_{353}}$ and $\sigma_{\Delta p_{353}}$ columns are over the 12 posterior samples of the Edenhofer et al. (2023) maps.	79
3.1 Statistical significance of the detection of V1 in units of equivalent Gaussian standard deviations as defined in Section 3.4.6 using the 95, 150, and 220 GHz bands of BICEP/Keck and the 353 GHz band of Planck. The column labeled "best" uses the parameters $D_W = 135'$, $\theta_{\text{FWHM}} = 4'$, and $Z = 0.75$, and the row labeled "default" uses the parameters $D_W = 75'$, $\theta_{\text{FWHM}} = 30'$, and $Z = 0.7$, which are used in Clark & Hensley (2019).	112
3.2 Comparison of the statistical significance of a detection of the cross correlation with the dust polarization in units of equivalent Gaussian standard deviations when including the channels in the IVC velocity range in the line-of-sight sum. The RHT parameters from Equation 3.13 are used here for the H \scriptsize I morphology template with the 95, 150, and 220 GHz bands of BICEP/Keck and the 353 GHz band of Planck.	119

3.3 Comparison of the statistical significance of a detection of the cross correlation between H ₁ morphology template and the dust polarization at different frequencies in units of equivalent Gaussian standard deviations as defined in Section 3.4.6.	125
3.4 Comparison of the statistical significance of a detection of the cross correlation between H ₁ morphology templates and the dust polarization in units of equivalent Gaussian standard deviations for V1, V2, and V3. We also add a column for V2 + V3, both of which are associated with Magellanic Stream 1. The 95, 150, and 220 GHz bands of BICEP/Keck and the 353 GHz band of Planck are used here.	132
3.5 Statistical significance of the detection of V1 in units of equivalent Gaussian standard deviations as defined in Section 3.4.6 using the 95, 150, and 220 GHz bands of BICEP/Keck and the 353 GHz band of Planck. The rows labeled "best" use the parameters $D_W = 135'$, $\theta_{\text{FWHM}} = 4'$, and $Z = 0.75$, and the rows labeled "default" use the parameters $D_W = 75'$, $\theta_{\text{FWHM}} = 30'$, and $Z = 0.7$, which are used in Clark & Hensley (2019). The bolded column (9) shows the main results.	141

3.6 Comparison of the statistical significance of a detection of the cross correlation with the 95, 150, and 220 GHz bands of BICEP/Keck and the 353 GHz band of Planck in units of equivalent Gaussian standard deviations when including the channels in the IVC velocity range in the line-of-sight sum. The bolded column (9) shows the main results.	142
3.7 The detection statistical significance of the cross correlation between H _I morphology templates and the dust polarization in units of equivalent Gaussian standard deviations as defined in Section 3.4.6. The bolded column (9) shows the main results.	143
3.8 Comparison of the statistical significance of a detection of the cross correlation between H _I morphology templates and the dust polarization in units of equivalent Gaussian standard deviations for V1, V2, and V3. We also add a column for V2 + V3, both of which are associated with Magellanic Stream 1. The 95, 150, and 220 GHz bands of BICEP/Keck and the 353 GHz band of Planck are used here. The bolded column (9) shows the main results. The other columns show the results for different variations of our model.	144
5.1 List of hyper-parameters used to train the model.	223

List of Figures

2.1 A map of the debiased signal-to-noise ratio of the Planck GNILC polarization fraction at $80'$. This is plotted with a diverging linear colorbar centered on 3, the cutoff we use as part of our sightline selections in Sections 2.4 and 2.5, with the allowed regions shown in blue.	33
---	----

2.2 Slices through the XY (left), XZ (middle), and YZ (right) planes of the 3D reconstructed differential extinction maps of Lallement et al. (2019) (top) and Edenhofer et al. (2023) (bottom). The Sun is at the origin. The positive X axis points towards the Galactic center at $l = 0^\circ$, the positive Y axis points towards $l = 90^\circ$ in the Galactic plane, and the positive Z axis points out of the plane in the direction of the Galactic North pole. The slices show the log of the differential extinction, which is in units of magnitudes per parsec. The subpanels only extend up to 400 pc in each direction for a direct comparison. The model for the Local Bubble surface geometry of Pelgrims et al. (2020) is overplotted in white in each subpanel. The white disk at the center of the bottom panel is due to missing data within 70 pc of the Sun in the Edenhofer et al. (2023) maps.

2.3	Maps used for selecting sightlines for the analysis described in Section 2.4.	
	<i>Top panel:</i> A map of the ratio of the Lallement et al. (2019) 3D dust differential extinction map integrated within 50 pc of the Local Bubble surface defined by Pelgrims et al. (2020), A_V^{LB} , over the Planck dust extinction, A_V^{Planck} . This is plotted with a diverging colorbar centered on the 75th percentile (0.2), the cutoff we use in our sightline selection, with the allowed regions shown in blue. <i>Bottom panel:</i> A map of the mask of the selected region, combining $A_V^{\text{LB}}/A_V^{\text{Planck}} > 0.2$ (top panel) with $\text{SNR}_p > 3$ (Figure 2.1).	43
2.4	Diagram of the angles described and used in Section 2.4.3 for projecting the plane-of-sky magnetic field orientation (orange line) onto the Local Bubble wall (blue). The coordinates in this diagram follow the COSMO (HEALPix) convention, which is used in the Planck GNILC maps. These are not the Galactic coordinates used in Figure 2.2. For each position in the sky looking outwards, the local horizontal axis points South, the local vertical axis points East, and the local z-axis points outwards.	46

2.7 The dust distribution in units of differential extinction (mag/pc, left vertical axis) and equivalent Hydrogen number density (cm^{-3} , right vertical axis) along a representative sightline through the Edensofer et al. (2023) maps. The Galactic coordinates of this sightline are $l = 163.12^\circ$ and $b = -11.42^\circ$. The region before the extinction reaches 50 mmag (hatched) is discarded from our analysis. The components identified by the dendrogram algorithm with $\Delta_n = 4.52 \times 10^{-3} \text{ cm}^{-3}$ ($A'_V = 7 \times 10^{-6}$), and $\#\text{voxels} = 3$ (21 pc) are shaded in different colors. 57

2.8 *Top panel:* A map of \mathcal{N}_c calculated using the dendrogram algorithm with $\Delta_n = 4.52 \times 10^{-3} \text{ cm}^{-3}$ ($A'_V = 7 \times 10^{-6}$), and $\#\text{voxels} = 3$ (21 pc). *Bottom panel:* A log-scale map of $N_{\text{H}}^{\text{Edensofer}}$ formed by integrating over the entire dataset (up to 1.25 kpc) with units of cm^{-2} . These maps are shown only for the sightlines selected in Section 2.5.1 and for one of the posterior samples of Edensofer et al. (2023). 59

2.9 The dust distribution in units of differential extinction (mag/pc, left vertical axis) and equivalent Hydrogen number density (cm ⁻³ , right vertical axis) along a pair of matched sightlines through the first posterior sample of the Edenhofer et al. (2023) map. These sightlines have the same $N_{\text{H}}^{\text{Edenhofer}}$, but the top one has a higher complexity than the bottom one. The Galactic coordinates, dust complexity, and dust polarization fraction of each of the sightlines are denoted on their subpanels.	63
2.10 Kernel density estimate plots of the Δp_{353} distributions over all matched pairs of sightlines of one posterior sample of the Edenhofer et al. (2023) maps. The p_{353} of the lower complexity sightline is always subtracted from that of the higher complexity one in the orange distribution. This has a mean of $\Delta p_{353} = -1.81 \times 10^{-2}$, which is plotted as an orange dot in Figure 2.11. The blue distribution contains the same pairs as the orange one with the sign randomly flipped for each pair, i.e., the distribution of the permutation-based null test described in Section 2.5.4.	64

3.2 Distributions of the best-fit values using E and B modes for 499 realizations of lensed- Λ CDM, noise, and Gaussian dust, added to the $\text{H}\alpha$ morphology template with fixed input values $a = 0.9$, $k = 0.7$, and $\beta_{\text{HI}} = 1.52$ that match the fit from the real data. The parameters a and β_{HI} are unitless, and k has units $\mu\text{K}_{\text{CMB}} / \text{K km s}^{-1}$. These known input values are plotted as dashed black vertical lines. The means of the distributions of the best-fit values are plotted as solid red vertical lines. The mean and standard deviation of each of the distributions are quoted above.	105
3.3 EE (top) and BB (bottom) correlation ratio of the integrated $\text{H}\alpha$ morphology template with individual $\text{H}\alpha$ morphology templates for the Hi4PI velocity channels across multipoles $37 < \ell < 579$. The 1D plots on top show the broadband correlation ratio calculated over one multipole bin spanning the entire multipole range. It is separated into 3 velocity regions, V1, V2, and V3. The LVC boundaries as defined in Panopoulou & Lenz (2020a) are indicated with dashed vertical lines. The broadband correlation ratio between the different pair combinations of the 3 velocity components is printed on the left of each histogram.	106

3.4 Integrated H ₁ intensity maps over the 3 different velocity components defined in Figure 3.3 in the BICEP/Keck region. The velocity boundaries for each component are printed on the bottom right of each map. The emission in V1 is dominated by the Milky Way, whereas the emission in V2 and V3 is dominated by Magellanic Stream 1 (Westmeier, 2018). The outlines of the BICEP3 and the BICEP2 and Keck Array observing fields are also plotted. The Small Magellanic Cloud (SMC) is indicated.	107
3.5 Polarized intensity maps of V1 in the BICEP/Keck region using RHT parameters that correlate $> 5\sigma$ (left) and $< 5\sigma$ (right) in <i>B</i> modes with BICEP/Keck and Planck data. Only the statistical significance in <i>B</i> modes is quoted in the title of each of the maps, because all of the RHT parameters we tried correlate well ($> 5\sigma$) in <i>E</i> modes. From top to bottom, the maps on the left have a $15.2\sigma_{EE}$, $12.6\sigma_{EE}$, and $14.9\sigma_{EE}$ detection significances, and the maps on the right have a $6.3\sigma_{EE}$, $8.6\sigma_{EE}$, and $8.2\sigma_{EE}$ detection significances.	113
3.6 Map of the first moment of the velocity distribution of the H ₁ structure in the BICEP/Keck region for $-12 \text{ km s}^{-1} < v_{\text{lsr}} < 10 \text{ km s}^{-1}$, the velocity range most correlated with the polarized dust emission. The texture is a line integral convolution of the magnetic field orientation as inferred by the H ₁ filaments.	114

3.7 The best-fit BB observables used in the $\Delta\chi^2$ statistic defined in Section 3.4.6 for the 95, 150, and 220 GHz bands of BICEP/Keck and the 353 GHz band of Planck. A modified blackbody frequency scaling, covariance matrix conditioning, and a transfer function for the $H\alpha$ morphology template with the RHT parameters from Equation 3.13 are used for the fit here. The cross spectrum between the real data and the $H\alpha$ morphology template (light blue), the best-fit cross spectrum between the $H\alpha$ morphology template and the modified $H\alpha$ -correlated component of the simulation (dark blue), and the mean of the cross spectra between the $H\alpha$ morphology template and the lensed- Λ CDM, noise, and Gaussian-dust components of the simulation (light green) are plotted.	120
3.8 Posteriors of k and $\beta_{H\alpha}$ fit using the Metropolis-Hastings algorithm on uniform priors and the χ^2 likelihood of the cross spectra of the real data with the $H\alpha$ morphology template. The parameter a is marginalized over. The E modes only (purple), B modes only (pink), and simultaneous E and B modes (navy) posteriors are shown. The units for k are $\mu K_{CMB} / K \text{ km s}^{-1}$ and $\beta_{H\alpha}$ is unitless.	121

3.9 Comparison of the posteriors for β_{HI} through a χ^2 likelihood using cross correlations with the H I morphology template (solid) to the ones of β_d using the Hamimeche and Lewis (HL) likelihood with a multicomponent model and no H I morphology template (dashed). We show the posteriors using B modes only (pink), and B and E modes (blue). The solid posteriors are the same as in Figure 3.8 plotted with the same colors. The B -mode-only total dust component posterior is identical to the posterior shown in black in Figure 4 of BK18.	122
3.10 Comparison of the posteriors for β_{HI} we get through a χ^2 likelihood using E - and B -mode cross correlations with the H I morphology template for different selections of frequency bands and for BICEP/Keck only and Planck only variations. The thick navy posterior labeled "All" is the same as the navy posterior in Figures 3.8 and 3.9.	127

3.12 *EE* (cross) and *BB* (circle) cross spectra as a function of multipole moment. The cross spectra between V1 and Planck data with 1σ variations are shown in red and brown to compare them to the cross spectra between V1 and BICEP/Keck data, which are shown in teal and turquoise. The errors are derived from spurious correlations between V1 and lensed- Λ CDM, Gaussian dust, and noise. Data points for similar frequencies between BICEP/Keck and Planck are plotted on the same panels for comparison. . . 131

4.1 The EE correlation ratio on the Planck 70% sky fraction Galactic plane mask of the Planck Commander dust maps with H I-based polarization templates produced by the Hessian algorithm applied to Hi4PI data. The leftmost point is for the H I-based polarization template at 2.03 km s^{-1} , and each successive point on each curve corresponds to the addition of information from the two adjacent velocity channels in the positive and negative directions. The labels on the horizontal axis correspond to the velocity centers of the positive velocity channels being added. The different curves correspond to different multipole bins shown in the color bar. The vertical dashed line corresponds to the integrated H I-based polarization template over the velocity range $-13 \text{ km s}^{-1} < v_{\text{lsr}} < 16 \text{ km s}^{-1}$, after which the correlation saturates and starts decreasing over most of the multipole bins considered as information from more velocity channels is added. 156

4.2 Comparison of the EE (solid) and BB (dashed) correlation ratios on the Planck 70% sky fraction Galactic plane mask of the Planck Commander dust maps (teal) and the Planck frequency maps (sandy brown) at 353 GHz with the H I-based polarization template constructed from applying the Hessian algorithm to Hi4PI data at each velocity channel and integrating the resulting maps over the velocity range $-13 \text{ km s}^{-1} < v_{\text{lsr}} < 16 \text{ km s}^{-1}$ 160

4.3 The EE (left) and BB (middle) correlation ratios of the Planck Commander dust maps with the H I -based polarization template constructed using the Hessian algorithm on the $\text{H}\alpha\text{4PI}$ intensity maps over the different non-overlapping masks shown on the right. The masks are the Planck 20% sky fraction Galactic plane mask (darkest), the inverted Planck 80% sky fraction Galactic plane mask (lightest), and the differences between the Planck 20%, 40%, 60%, and 80% sky fraction Galactic plane masks (other shades of blue) shown in a Mollweide projection in Galactic coordinates centered on the Galactic center. 161

4.4 Maps of a $15^\circ \times 15^\circ$ patch of sky centered on R.A. = $7^\circ 5$, decl. = $28^\circ 1$ and $v_{\text{lsr}} = 400.1 \text{ km s}^{-1}$ of the input intensity channel from GALFA- H I with width 0.74 km s^{-1} (left) and polarized intensity of H I -based polarization templates constructed using the Hessian (middle left) and RHT (middle right and right) algorithms applied to this channel. The three parameters listed in the titles of the polarized intensity maps produced with the RHT algorithm are D_W , θ_{FWHM} , and Z , respectively, explained in Section 4.4.1. 163

4.5	Polarized intensity map projections of a $400' \times 400'$ patch of sky, centered at $(l, b) = (15^\circ, 50^\circ)$, of H I-based polarization templates constructed using the Hessian algorithm applied to GALFA-H I intensity maps smoothed to different resolutions. The intensity map corresponding to the leftmost polarized intensity map is not smoothed, i.e., it has the native resolution of GALFA-H I. The intensity maps corresponding to the polarized intensity maps to the right of the first map are smoothed with Gaussian kernels to the resolution stated in their titles. The last map on the right corresponds to the Hessian algorithm applied to Hi4PI data at its native resolution, though it is integrated over a slightly different velocity range as mentioned in Section 4.5.3.	165
4.6	The EE (left), BB (middle), and TE (right) correlation ratios between the maps in Figure 4.5 and the Planck Commander dust maps at 353 GHz. The TE correlation ratio measures the correlation between the Planck total intensity and the templates' E modes. Correlations are computed on a combination of the GALFA-H I and Planck 70% sky fraction masks.	167

4.7 Comparison of the EE (solid) and BB (dashed) correlation ratios with the Planck Commander dust maps at 353 GHz of the H I-based polarization templates using GALFA-H I data smoothed to a FWHM of $7'$ (light green) and HI4PI data (teal). The H I-based polarization templates are integrated over a similar velocity range and constructed using the Hessian algorithm. The correlations are calculated on a combination of the GALFA-H I mask with the Planck 70% sky fraction mask. The teal lines are the same as those in Figure 4.2. 171

4.9 Stokes Q map projections of a $15^\circ \times 15^\circ$ patch of sky, centered at $(l, b) = (15^\circ, 50^\circ)$, of H I-based polarization templates for one velocity slice of the Hi4PI Survey centered at $v_{\text{lsr}} = 2.03 \text{ km s}^{-1}$ with width 1.3 km s^{-1} constructed using the RHT (left) and the Spherical RHT (middle) algorithms with parameters $D_W = 75'$, $\theta_{\text{FWHM}} = 30'$, and $Z = 0.7$. The map on the left is used in Clark & Hensley (2019). The map on the right is the percentage difference between the map on the left and the map in the middle. 173

4.10 Left: the EE (green) and BB (blue) correlation ratios between the H I-based polarization template used in Clark & Hensley (2019) and that reproduced using the Spherical RHT algorithm with the same parameters. Right: the EB correlation ratios, showing that the likely spurious positive correlation when the B modes of the Clark & Hensley (2019) H I-based polarization template are used (dashed) vanish when the curved sky is taken into account and the B modes produced with the Spherical RHT algorithm are used (solid). The orange (purple) curves represent correlation ratios where the E and B modes of the same algorithm (different algorithms) are used. Hi4PI data and the Planck 70% sky fraction mask are used in these plots. Note the difference in the y-scales between the left and right panels. 176

4.11 Polarized intensity map projections of a $16.7^\circ \times 16.7^\circ$ patch of sky, centered at $(l, b) = (15^\circ, 50^\circ)$, of H I-based polarization templates constructed using the Spherical RHT algorithm with different parameters applied to Hi4PI intensity maps. The parameter Z is fixed to 0.7 and the Spherical RHT is run on a grid of exponentially increasing parameters between $5'$ and $320'$ for θ_{FWHM} (to the right) and $20'$ and $640'$ for D_W (to the bottom).	178
4.12 The EE (left), BB (middle), and TE (right) correlation ratios between the maps in Figure 4.11 and the Planck Commander dust maps over a broadband multipole bin between $\ell = 20$ and $\ell = 600$. The Planck total intensity is correlated with the E modes of the Spherical RHT-based templates for the TE case.	181
4.13 Right panel: the TE (left), TB (middle), and EB (right) cross spectra (D_ℓ) of two synthetic filaments close to each other (blue, left panel) and far from each other (orange, left panel), where the arrangement of the filaments in the top and bottom rows have opposite parities. The headless vectors in the left panels show the polarization angle orientations perpendicular to the lengths of these filaments. The color scale in the left panels represents the polarized intensity.	183

4.14 The EE (solid) and BB (dashed) correlation ratios between the H I-based polarization template constructed using the Hessian algorithm and four H I-based polarization templates constructed using the Spherical RHT algorithm on Hi4PI data. The Planck 70% sky fraction mask was used for the spectra used for calculating these correlation ratios. The parameters listed in the legend are the D_W , θ_{FWHM} , and Z , respectively, defined in Section 4.4.1. The first set of parameters (sandy brown) is the one that correlates the best, and the other three are randomly selected.	186
4.15 Stokes Q (top) and U (bottom) map projections of a $16.^{\circ}7 \times 16.^{\circ}7$ patch of sky, centered at $(l, b) = (15^{\circ}, 50^{\circ})$, of the H I-based polarization templates constructed using the Hessian method (left) and Spherical RHT algorithm with parameters $D_W = 160'$, $\theta_{\text{FWHM}} = 10'$, and $Z = 0.7$ (middle) applied to Hi4PI intensity maps, and the full-mission Planck Commander dust map (right).	187

5.5	The full model architecture. The input images are each processed separately by an encoder, whose architecture is shown in Figure 5.6. The processed images along with an angular resolution embedding vector are then processed by a transformer fusion module, which consists of transformer layers (Figure 5.7). The transformer module's outputs corresponding to the Planck GNILC Stokes Q and U images are then processed through decoders, whose architecture is shown in Figure 5.6.	218
5.6	The architectures of the encoder (middle), of the decoder (bottom), and of their main component, the Skip Residual Block (top). The Skip Residual Block is made up of two convolutional layers with a non-linearity and a skip connection between them and a residual connection over the entire block. The encoder and decoder are made up of a convolutional layer, followed by a non-linearity, a number of Skip Residual Blocks, and a final convolutional layer. The encoder additionally has a residual connection from the output of the first convolutional layer to the end.	219
5.7	The transformer architecture used in this paper. It is made up of multi-head self-attention, where Q , K , and V are the query, key, and value tensors and a feed-forward layer followed by a non-linearity. Both components are preceded by a layer normalization and enveloped by a residual connection.	220

5.8 Example low-resolution input and high-resolution predictions and target 353 GHz Stokes Q (top 3 rows) and U (bottom 3 rows) patches of sky. The same patch of sky is shown across each column with its corresponding high angular resolution denoted at the top. The colorbars are centered at zero (darkest) and brighter red (blue) corresponds to higher positive (negative) values. . . 229

5.9 Attention map averaged over 80×80 pixels and over 100 test sample patches, 25 for each angular resolution. The values in each row sum up to 1 because of the softmax operation from Equation 5.2. 230

Chapter 1

Introduction

The space between stars is filled with magnetic fields, dust, gas, and cosmic rays. These components collectively form the interstellar medium (ISM). The ISM plays a crucial role in various astrophysical processes, including star formation. Stars, in turn, undergo supernova explosions, recycling material back into the ISM. This cyclical process is vital for the continuous evolution of galaxies. The ISM also acts as a foreground for cosmological observations, impacting the study of the cosmic microwave background (CMB). Understanding the ISM's properties and behavior is essential for interpreting these observations accurately.

Matter in the ISM is distributed into several phases, ranging from highly concentrated molecular clouds to diffuse ionized regions (McKee & Ostriker, 1977; Ferrière, 2001; Cox, 2005). Molecular clouds, which serve as stellar nurseries, comprise approximately 10%

of the ISM's mass (Miville-Deschénes et al., 2017). These structures maintain low temperatures (10-20 K) due to their high density (10^2 - 10^6 molecules cm^{-3}), which effectively attenuates ambient stellar radiation. In these regions, cosmic rays are the primary heating mechanism, interacting with gas molecules and leading to ionization and subsequent heating processes (Krumholz, 2011). Cooling in molecular clouds is primarily facilitated by CO line emission (Krumholz, 2011). A substantial portion of the remaining ISM mass exists as cold atomic hydrogen, often referred to as the cold neutral medium (CNM), with temperatures ranging from 50 to 100 K. This component typically manifests as diffuse clouds (Draine, 2011). The primary heating mechanism for the CNM is the photoelectric effect on small dust grains and polycyclic aromatic hydrocarbons (PAHs), while cooling is dominated by the [C II] fine-structure line emission (Dalgarno & Rudge, 1964; Wolfire et al., 1995). The CNM coexists with and forms out of the warm neutral medium (WNM), which has temperatures around 10^4 K. Heating in the WNM is primarily driven by photoionization from the diffuse ultraviolet background, and cooling occurs mainly via Lyman-alpha emission (Gould & Thakur, 1970; Dalgarno & McCray, 1972; Wolfire et al., 2003). There is also a substantial fraction of gas in an unstable phase between the CNM and WNM, maintained by a balance of these heating and cooling processes, with turbulence and shocks further contributing to the dynamics. The ISM also encompasses an ionized component that can reach temperatures up to 10^6 K (Ferrière, 2001). This diverse thermal and ionization structure highlights the complex interplay between various physical processes operating

within the ISM, including radiative cooling, heating by stellar radiation and cosmic rays, and dynamical phenomena such as shocks and turbulence.

Interstellar dust is an important component of the ISM despite comprising only 1% of its mass (Popescu & Tuffs, 2002). Dust grains are primarily composed of silicates, carbonaceous materials, and ices (Draine, 2003; Hensley & Draine, 2021). They catalyze the formation of molecular hydrogen, shield complex molecules from ultraviolet radiation, and regulate the heating and cooling of the ISM (Draine, 2011).

In this thesis, we use different measurements of the ISM to study the polarization of the thermal dust emission. We focus on the neutral phases of the ISM, and we show that neutral gas emission is a useful tool for this purpose. In this chapter, we describe the different ISM ingredients we use in the thesis and explain why they are useful for studying the polarized dust emission. We then focus on magnetic fields and their role in polarizing the dust emission. Finally, we describe the importance of characterizing the polarized dust emission for its role in obscuring cosmological measurements.

1.1 3D Structure of the Magnetic ISM

1.1.1 Dust Mapping

Dust in the interstellar medium interacts with light in various ways across the electromagnetic spectrum. At ultraviolet and visible wavelengths, dust primarily scatters and

absorbs light, while at infrared and longer wavelengths, it predominantly emits thermal radiation. This wavelength-dependent behavior, particularly the preferential absorption of shorter wavelength light, leads to the reddening of starlight and allows us to probe the three-dimensional distribution of dust in space.

Extinction, which refers to the absorption and scattering of starlight by dust along the line of sight, provides a powerful means to map the spatial distribution of dust throughout the Galaxy. It enables us to reconstruct not just the total integrated column of dust across the sky, but the density of dust throughout space. As stars are distributed in three dimensions throughout the volume of the Milky Way, they serve as tracers not just of the total integrated column of dust across the sky, but of the density of dust throughout space (Green et al., 2015). By measuring the reddening and extinction of starlight for numerous stars at various distances, it becomes possible to reconstruct the three-dimensional structure of dust in the interstellar medium.

The process of mapping dust using stellar observations relies on the principle that stars behind denser dust clouds appear more reddened relative to their intrinsic colors. The amount of reddening, combined with accurate distance measurements, allows for the de-projection of integrated extinction measurements into a 3D map of differential dust extinction (Zucker et al., 2019). This technique effectively uses stars as discrete sampling points of the continuous dust distribution, with the density of sampling increasing with the number of observed stars.

A significant challenge in dust mapping is the degeneracy between stellar type and dust reddening. The observed color of a star is influenced both by its intrinsic properties and the intervening dust, making it difficult to disentangle these effects. Early efforts were limited to relatively small samples of stars with known spectral types (Neckel & Klare, 1980). That is still true today with only millions of stars having spectroscopic measurements. However, the advent of large-scale photometric surveys and advanced statistical methods has revolutionized this field, enabling the inference of distances and reddening for hundreds of millions of stars (Gaia Collaboration et al., 2023; Zhang et al., 2023).

The Gaia mission, in particular, has significantly improved our ability to break the degeneracy between stellar properties and dust effects. It has provided accurate distances through precise parallax measurements for more than a billion stars (Gaia Collaboration et al., 2023). Also, the spectra from Gaia's "Blue Photometer" and "Red Photometer" instruments (BP/RP spectra), while low-resolution, offer crucial information for determining both stellar parameters and dust properties (Gaia Collaboration et al., 2023). The combination of Gaia BP/RP spectra with near-infrared photometry from surveys like 2MASS (Skrutskie et al., 2006) and WISE (Schlafly et al., 2019) further enhances the ability to break the degeneracy between stellar temperature and extinction because extinction has a far smaller effect at longer wavelengths (Zhang et al., 2023). This wealth of data has enabled the creation of 3D dust maps with unprecedented resolution and coverage.

Several approaches have been developed to leverage this data effectively. One method

involves using ab initio physical models which predict stellar spectra based on fundamental properties (Gaia Collaboration et al., 2023). Another approach is to develop empirical forward models of BP/RP spectra using a subset of stars with high-quality spectroscopic data from independent surveys (Zhang et al., 2023). A third method employs machine learning techniques to directly predict stellar parameters from BP/RP spectra (Rix et al., 2022; Andrae et al., 2023). These approaches have their respective strengths and limitations. Forward modeling techniques make full use of measurement uncertainties but require explicit modeling of all relevant parameters. In contrast, machine learning approaches may be more adept at handling systematic errors and unmodeled variables but may degrade more rapidly in low signal-to-noise regimes.

Approaches to constructing 3D dust maps can be broadly categorized into two main types: reconstructions on Cartesian grids and reconstructions on spherical grids. Cartesian grid reconstructions often exhibit less pronounced "fingers-of-god" artifacts but face challenges in scaling to large volumes. They typically either cover a limited volume at high resolution (Leike & Enßlin, 2019; Leike et al., 2020) or a larger volume at lower resolution (Capitanio et al., 2017; Lallement et al., 2018, 2019, 2022; Vergely et al., 2022a). Spherical grid reconstructions, on the other hand, can probe larger volumes at higher resolutions but often suffer from more pronounced "fingers-of-god" artifacts (Chen et al., 2019a; Green et al., 2018, 2019).

The "fingers-of-god" effect, characterized by the smearing of dust structures along the

line of sight, arises from the superior constraints on stars' plane-of-sky positions relative to their line-of-sight distance uncertainties. To mitigate this effect, various techniques have been employed, including the use of physical smoothness priors. Gaussian process priors, which are probabilistic models that assume smooth variations in the dust distribution based on spatial correlations, have been particularly useful in recent reconstructions of the 3D dust distribution (Leike et al., 2020; Edenhofer et al., 2023).

Alternative approaches have also been explored, such as using multiple small reconstructions (Leike et al., 2022), analytical methods (Rezaei Kh. et al., 2017, 2018, 2020; Rezaei Kh. & Kainulainen, 2022), and inducing point methods (Dharmawardena et al., 2022). However, these approaches have faced challenges in reconstructing dust at high resolution over large volumes without introducing artifacts.

As new data become available and computational techniques advance, the field of 3D dust mapping is likely to see further improvements in resolution, accuracy, and coverage, providing increasingly valuable tools for studying the structure and the complex interplay between dust, gas, and stellar populations in our Galaxy.

1.1.2 Structure of the Local ISM

The Local Bubble is a prominent feature of the ISM in the solar neighborhood, characterized by its low-density interior and surrounding shell of cold, dusty gas. Initially proposed based on extinction mapping and absorption line measurements (Welsh & Lallement, 2010; Frisch

et al., 2011), the Local Bubble's existence has been further supported by observations of diffuse soft X-ray emission (Snowden et al., 1997). Current evidence suggests that the Local Bubble was formed by stellar winds and supernovae from nearby massive stars (Smith & Cox, 2001; Breitschwerdt, 2001).

The advent of high-resolution 3D dust mapping techniques has considerably enhanced our understanding of the local ISM, particularly in relation to the Local Bubble. These advanced mapping methods have provided unprecedented insights into the Local Bubble's structure, evolution, and its role in shaping the solar neighborhood. Building upon these dust-based observations, we can now construct a more comprehensive picture of the Local Bubble as a prominent feature of the local ISM.

Recent studies have provided new insights into the Local Bubble's origin and structure. Kinematic analyses indicate that young star clusters entered the present Local Bubble region approximately 10-15 Myr ago, with about 14-20 high-mass members having exploded since then (Maíz-Apellániz, 2001; Fuchs et al., 2006). The detection of live ^{60}Fe on Earth has been successfully explained within this framework (Schulreich et al., 2017; Breitschwerdt et al., 2016), supporting the multiple supernovae origin of the Local Bubble.

The role of the Local Bubble in local star formation has been highlighted by Zucker et al. (2022), who found that nearly all star-forming regions in the solar neighborhood lie on the Local Bubble's surface, with young stars exhibiting outward expansion mainly perpendicular to this surface. This suggests that the Local Bubble's expansion may be

responsible for a large fraction of nearby star formation.

Three-dimensional mapping of the Local Bubble has employed various tracers, including NaI absorption measurements, stellar color excess, X-ray emission, and diffuse interstellar bands (Sfeir et al., 1999; Lallement et al., 2014; Snowden et al., 1998; Farhang et al., 2019). Recent high-resolution 3D dust maps have enabled more detailed reconstructions of the Local Bubble's shell geometry (Pelgrims et al., 2020; O'Neill et al., 2024).

The morphology of the Local Bubble, particularly at high latitudes, has been a subject of debate. Early models suggested an open "chimney" structure extending into the Galactic halo (Sfeir et al., 1999; Welsh et al., 1999), while more recent dust-based models presented a closed surface (Pelgrims et al., 2020). O'Neill et al. (2024) propose an asymmetric Local Chimney model, with an open, low-density Northern cap and a closed Southern cap.

The structure of the Local Bubble has implications for its evolution and interaction with the Galactic environment. Theoretical models and simulations suggest that superbubble breakout into the halo depends on factors such as gas density distribution, magnetic field orientation, and the positions of progenitor supernovae relative to the Galactic midplane (Mac Low et al., 1989; Orr et al., 2022).

The Local Bubble is part of a larger network of bubbles and shells in the local ISM. Nearby structures include the Gum Nebula, Antlia SNR, Monogem Ring, and larger-scale superbubbles such as the Orion-Eridanus superbubble (Reynolds & Ogden, 1979; Soler et al., 2018). The interaction between these structures, such as the adjacency of the Per-Tau

Shell to the Local Bubble, may play a role in triggering star formation (Zucker et al., 2022; Soler et al., 2023).

Understanding the Local Bubble’s structure and evolution contributes to our broader comprehension of the multiphase ISM shaped by supernova feedback (Cox & Smith, 1974; McKee & Ostriker, 1977). Further research, including high-resolution 3D mapping and detailed simulations, will be important for refining our model of the Local Bubble and its role in the local Galactic environment.

1.1.3 Neutral Hydrogen Gas

Neutral hydrogen gas (H I), observable through its 21 cm emission line, plays an important role in mapping the three-dimensional structure of the ISM. As the fundamental building block of galaxies, H I constitutes approximately 75% of the neutral gas in galaxies in the local Universe by mass (Duffy et al., 2012).

The 21 cm line of H I results from the hyperfine splitting of its ground state, caused by the interaction between the magnetic moments of the proton and the electron (Ewen & Purcell, 1951). This transition, although rare, is detectable due to the vast quantities of H I in the ISM. The emission at 21 cm is directly proportional to the column density of H I and can be used to map the spatial distribution of neutral hydrogen across the Galaxy (Kalberla & Kerp, 2009).

H I absorption occurs when background continuum sources, such as quasars or thermal

emission from ionized regions, are observed through H₁ clouds. The absorption profile provides information on the optical depth and temperature of the H₁ gas (Dickey & Lockman, 1990). Combining emission and absorption data allows for the determination of the spin temperature and column density, offering a more complete understanding of the physical state of H₁ (Heiles & Troland, 2003).

The Doppler effect on the 21 cm line allows for the measurement of the velocity of H₁ gas along the line of sight. When H₁ gas moves relative to the observer, the frequency of the emitted radiation is shifted. This Doppler shift provides a third dimension to the H₁ data: velocity with respect to the local standard of rest (Burton, 1988). By analyzing the Doppler-shifted frequencies, astronomers can construct velocity maps that reveal the kinematic structure of the ISM.

These velocity maps are essential for understanding the dynamic processes within the ISM, such as rotation, turbulence, and interactions between different gas phases. For instance, the study of rotational velocities helps in mapping the Galactic rotation curve, which in turn provides insights into the distribution of mass within the Milky Way (Sofue & Rubin, 2001). Additionally, velocity information can identify expanding shells, shock fronts, and other dynamic phenomena in the ISM (Heiles, 1979).

The ISM's structure is shaped by a balance of heating and cooling processes, resulting in a medium with varying temperatures and densities. In the standard theory for the thermal structures of interstellar H₁, it can be approximated as a two-phase medium,

with warm and cold phases coexisting in pressure equilibrium under certain conditions (Field et al., 1969; Wolfire et al., 1995; Cox, 2005). The ISM is not entirely atomic; a three-phase model has been proposed where cold neutral medium (CNM) clouds are surrounded by the warm neutral medium (WNM) and embedded within a supernova-heated hot, ionized medium (McKee & Ostriker, 1977; Cox & Smith, 1974). Recent theoretical models have further refined this understanding, incorporating factors such as turbulent pressure, leading to a more complex multi-phase structure of the ISM (Wolfire et al., 2003; Audit & Hennebelle, 2005; Bialy & Sternberg, 2019). Observationally, distinguishing between the CNM and WNM has been challenging due to the diffuse nature and large linewidth of the WNM, which can obscure the structure of the CNM in position-position-velocity cubes of $\text{H}\alpha$ emission (Kalberla & Haud, 2018). Techniques such as Gaussian decomposition of $\text{H}\alpha$ spectra and analysis of absorption-emission pairs have been developed to separate these components and study their individual properties (Kalberla & Haud, 2018).

Recent research has revealed a strong link between $\text{H}\alpha$ structure and magnetic field structure, as traced by starlight and dust polarization (Heiles, 1997a; Clark et al., 2014; Clark & Hensley, 2019). The $\text{H}\alpha$ gas is organized into complex structures, including filaments that trace the underlying magnetic field. These filaments are particularly evident in the CNM (Clark et al., 2014). The alignment of $\text{H}\alpha$ filaments with the magnetic field lines suggests that magnetic forces play an important role in shaping the structure of the ISM. The magnetic field can guide the flow of the gas, influence the formation of cold clouds,

and affect the propagation of shock waves through the ISM (Beck, 2015).

$\text{H}\alpha$ is often found in association with interstellar dust, a relationship important for understanding the ISM’s structure and dynamics (Draine, 2011). This association manifests in various ways with considerable implications for both components. For instance, regions with both high dust emission (Low et al., 1984; Boulanger et al., 1996a) and extinction (Sturch, 1969; Bohlin et al., 1978) tend to have higher $\text{H}\alpha$ column densities, allowing the dust distribution to be traced through $\text{H}\alpha$ observations. This linear relationship between the $\text{H}\alpha$ column density and dust reddening, with a mean ratio of $8.8 \times 10^{20} \text{ cm}^{-2} \text{ mag}^{-1}$, holds until $\text{H}\alpha$ column densities of about $4 \times 10^{20} \text{ cm}^{-2}$ (Lenz et al., 2017). Above that threshold, $\text{H}\alpha$ stops dominating the gas column along the line of sight and the presence of molecular hydrogen (H_2) breaks the linear relationship (Lenz et al., 2017). Another complication for using $\text{H}\alpha$ data as a tracer of reddening is the variations in the dust-to-gas ratio across the sky. This is addressed by using $\text{H}\alpha$ velocity information (Planck Collaboration et al., 2011a; Lenz et al., 2017). Lenz et al. (2017) find no significant reddening associated with gas at velocities higher than 90 km s^{-1} .

The interaction between $\text{H}\alpha$ and dust plays a vital role in the ISM’s physical and chemical processes. Dust grains shield $\text{H}\alpha$ from ultraviolet radiation, allowing it to remain neutral in regions where it might otherwise be ionized (Draine, 2011). Additionally, dust grains serve as catalysts for important chemical reactions in the ISM, most notably providing sites for the formation of H_2 (Hollenbach & Salpeter, 1971). This $\text{H}\alpha$ -to- H_2 conversion is a critical

step in the evolution of the ISM, as it leads to the formation of dense molecular clouds. These clouds, in turn, become the birthplaces of stars, linking the cycle of H I and dust to the broader process of galactic evolution (Klessen & Glover, 2016).

1.2 Dust Polarization and Magnetic Fields

Interstellar dust grains, ranging in size from nanometers to micrometers, play a crucial role in tracing the structure of interstellar magnetic fields. The physics of dust grain alignment with magnetic fields has seen significant advancements in the past two decades (Hoang et al., 2022; Hoang, 2022; Tram & Hoang, 2022).

The stability of grain alignment against perturbations from gas collisions is a critical factor in maintaining the observed polarization. Purcell (1979) proposed that grains rotating with kinetic energy significantly higher than any temperature in the system are needed to resist misalignment caused by gas bombardment. This suprothermal rotation can be achieved through torques that are fixed in the grain's coordinates, consistently inducing rotation in the same direction and minimizing the impact of gas collisions.

The Radiative Alignment Torque (RAT) mechanism (Dolginov & Mitrofanov, 1976; Draine & Weingartner, 1997; Lazarian & Hoang, 2007; Hoang & Lazarian, 2016) has emerged as the leading explanation for grain alignment. This process involves the interaction between anisotropic radiation and irregularly shaped dust grains, resulting in the grains spinning up and aligning their long axes perpendicular to the local magnetic field direction

(Andersson et al., 2015). The RAT mechanism operates through differential scattering of left- and right-hand circularly polarized components of incident radiation by helical grain surfaces. This interaction imparts angular momentum to the grains, inducing rotation. For paramagnetic grains, this rotation generates a magnetic moment via the Barnett effect, which subsequently interacts with the ambient magnetic field, leading to precession and eventual alignment (Dolginov & Mitrofanov, 1976; Draine & Weingartner, 1997). The RAT theory has been qualitatively validated through various observational studies (Andersson et al., 2015; Lazarian et al., 2015).

The alignment process involves both internal and external components. Internal alignment refers to the orientation of the grain’s angular momentum vector with its principal axis of maximum moment of inertia and takes on the order of a year, while external alignment describes the orientation of the grain’s angular momentum with respect to the magnetic field and takes on the order of 10^5 years. The timescales for these processes can vary, influencing the overall alignment efficiency in different interstellar environments (Purcell, 1979).

The efficiency of grain alignment is determined by the fraction of grains aligned exhibiting suprathermal rotation. This fraction is influenced by factors such as grain shape, radiation field intensity, and the angle between the radiation direction and the magnetic field (Lazarian & Hoang, 2007; Hoang & Lazarian, 2009). The presence of iron inclusions in dust grains can enhance the alignment efficiency through a process known as magnetically enhanced RAT (MRAT) alignment (Davis & Greenstein, 1951; Hoang & Lazarian, 2008;

Lazarian & Hoang, 2008; Hoang & Lazarian, 2016). This enhancement is due to the increased magnetic susceptibility of grains containing iron clusters, which is plausible given that iron is an abundant element in the universe, with up to 95% of iron missing from the gas phase (Jenkins, 2009; Dwek, 2016).

The MRAT alignment mechanism provides a framework for predicting the degree of grain alignment as a function of local gas density, radiation field intensity, grain sizes, and dust magnetic properties (Hoang & Lazarian, 2008, 2016; Lazarian & Hoang, 2019). This model helps explain the high observed polarization fractions (Panopoulou et al., 2019; Planck Collaboration et al., 2020a), which require alignment efficiencies of up to 60% (Draine & Hensley, 2021), beyond what can be achieved by RAT alone, which has a maximum efficiency of approximately 50% (Herranen et al., 2021). It is important to note that the composition of dust grains affects their alignment properties. Silicate grains exhibit higher paramagnetism compared to hydrocarbon grains, leading to more efficient alignment in the presence of magnetic fields (Lazarian & Hoang, 2019; Lazarian, 2020). The interplay between radiative torques, magnetic fields, and grain composition creates a complex system that governs the alignment of interstellar dust grains. This alignment process is fundamental to our understanding of interstellar magnetic field structures and the polarization of starlight observed in various astrophysical environments.

Interstellar dust grains function as both polarizing filters and sources of polarized light due to their elongated shape and alignment with magnetic fields. When starlight passes

through a cloud of aligned dust grains, it becomes polarized through dichroic extinction, with light polarized parallel to the grains' long axes being preferentially absorbed (Hiltner, 1949; Davis & Greenstein, 1951). Starlight polarization, resulting from the preferential absorption by magnetically aligned dust grains, has historically been instrumental in revealing the large-scale alignment of the Galactic magnetic field with the Galactic plane (Fosalba et al., 2002). This technique is particularly effective in probing the diffuse interstellar medium. The dispersion of polarization angles has been used in combination with probes of the density and velocity dispersion to estimate the strength of the plane-of-sky magnetic field component, as first proposed by (Davis, 1951) and (Chandrasekhar & Fermi, 1953), although it only gives a rough estimate of the magnitude of the magnetic field vector with large error bars (e.g., Ostriker et al., 2001; Heitsch, 2005; Falceta-Gonçalves et al., 2008; Hough et al., 2008; Cho & Yoo, 2016; Liu et al., 2021; Lazarian et al., 2022; Chen et al., 2022; Myers et al., 2024).

Conversely, these same aligned dust grains emit thermal radiation preferentially along their long axes, resulting in polarized emission perpendicular to the magnetic field orientation (Benoît et al., 2004). The grains absorb some of the interstellar radiation field, ambient ultraviolet, visible, and near-infrared radiation from stars in the Galaxy, which heats them up. To maintain thermal equilibrium, they emit primarily at far-infrared and submillimeter wavelengths to cool down (Planck Collaboration et al., 2020a). This thermal dust emission is typically optically thin and directly traces the dust temperature and column density

(Planck Collaboration et al., 2020a). Larger dust grains, which constitute the majority of the mass distribution, exhibit thermal radiation characterized by a modified blackbody function with emissivity that depends on the dust properties. Observational data indicate that the equilibrium temperature of interstellar dust grains is typically around 19.6 K in most of the ISM (Planck Collaboration et al., 2014a). The mean dust spectral index for the polarized dust emission is measured to be 1.53 ± 0.02 (Planck Collaboration et al., 2020b).

Measurements of the polarized dust emission have revealed an asymmetry in the *E*-mode and *B*-mode polarization components, quantified by the ratio of their power spectra (Planck Collaboration et al., 2020b). The angular power spectra of dust polarization at 353 GHz exhibit power-law behavior with statistically significant variations in exponents across sky regions, and a notable difference between *EE* and *BB* spectra exponents ($\alpha_{EE} = -2.42 \pm 0.02$ and $\alpha_{BB} = -2.54 \pm 0.02$ over 70% of the sky) (Planck Collaboration et al., 2020b). The *BB/EE* ratio is 0.53 ± 0.01 over the same sky area (Planck Collaboration et al., 2020b). These measurements, alongside a significant *TE* power, are explained by a statistical alignment between the local orientation of the Galactic magnetic field and filamentary density structures (Planck Collaboration et al., 2016a).

The dual nature of dust grains as both absorbers and emitters of polarized light provides complementary information about the interstellar magnetic field structure. Starlight polarization traces the magnetic field orientation between the observer and the background star, while polarized dust emission integrates contributions along the entire line of sight.

The orthogonality between absorption and emission polarization angles offers a valuable cross-check for magnetic field mapping techniques in different interstellar environments and at various wavelengths (Clemens et al., 2012). Comparing the thermal dust polarization with starlight polarization data also constrains Galactic dust models (Planck Collaboration et al., 2020a).

In addition to polarized dust emission and starlight polarization, several other observational techniques provide valuable insights into the structure of interstellar magnetic fields. Spectral line polarization, particularly through the Zeeman effect, allows for direct measurement of magnetic field strengths in denser regions of the interstellar medium. First detected in the ISM by Verschuur (1968), the Zeeman effect primarily measures the line-of-sight component of the magnetic field by measuring the difference between components of the circular polarization (Heiles et al., 1993; Crutcher, 2012). The Goldreich-Kylafis effect, another form of spectral line polarization, produces linearly polarized molecular line emission that can be either parallel or perpendicular to the projected magnetic field direction, providing additional information about field geometry in regions with anisotropic radiation (Goldreich & Kylafis, 1981; Greaves et al., 1999). Synchrotron radiation, produced by relativistic electrons spiraling in magnetic fields, traces the total and ordered components of the magnetic field perpendicular to the line of sight (Beck, 2001). This technique is particularly useful for studying magnetic fields in diffuse regions. Faraday rotation, which measures

the rotation of the plane of polarization of linearly polarized light passing through a magnetized plasma, provides information about the line-of-sight component of the magnetic field weighted by the electron density (Gardner & Whiteoak, 1966). When combined with dispersion measures from pulsars, Faraday rotation measurements can be used to derive average magnetic field strengths along different paths through the Galaxy, contributing to our understanding of the three-dimensional magnetic field structure (Han et al., 2006). Each of these techniques offers unique and complementary information, collectively enhancing our ability to map and understand the complex structure of interstellar magnetic fields across various scales and environments.

The three-dimensional nature of astrophysical magnetic fields necessitates precise 3D measurements to elucidate their dynamic influence on ISM evolution and the development of interstellar filaments and prestellar core collapse (Hull & Zhang, 2019; Hennebelle & Inutsuka, 2019; Pattle et al., 2023). Despite this requirement, accurate 3D magnetic field measurements in interstellar environments remain elusive. Recent years have witnessed the proposal of various methodologies aimed at constraining 3D magnetic fields through the integration of multiple observational tracers. These approaches include the synergistic use of polarized dust emission with Faraday rotation measurements (Tahani et al., 2018; Tahani, 2022), with velocity gradient information (Lazarian et al., 2022), and with polarized synchrotron emission (Jaffe et al., 2013).

Models of the Galactic magnetic field often decompose it into large-scale (regular) and

small-scale (random) components, with the latter sometimes further divided into "ordered random" and "isotropic random" subcomponents (Planck Collaboration et al., 2016b; Unger & Farrar, 2023). This decomposition reflects different physical processes: the large-scale component is associated with galaxy-scale dynamics, while the random components are linked to ISM turbulence and potential shearing or compression effects. Observations generally indicate that the large-scale field aligns with the spiral arms in the Galactic disk, a feature common to many spiral galaxies (Beck, 2015). One notable feature is the presence of at least one large-scale field reversal between the Sun and the Galactic Center, although the total number and locations of such reversals remain topics of debate (e.g., Simard-Normandin & Kronberg, 1980). The magnetic field's pitch angle, an important parameter given its spiral structure, is estimated to range from -5° to -30° locally, potentially varying with location due to interactions between spiral density waves and the magnetic field (Gómez & Cox, 2004).

The magnetic field strength based on Zeeman measurements shows interesting behavior across different density regimes, remaining relatively constant for densities of hydrogen nuclei lower than approximately 300 cm^{-3} , but scaling approximately as the two-thirds power of the density at higher values (Crutcher et al., 2010; Crutcher, 2012). In the Solar neighborhood, the total field strength is estimated at about $6 \mu\text{G}$, with contributions from both regular ($\sim 2 \mu\text{G}$) and random ($\sim 3 - 4 \mu\text{G}$) components (Haverkorn, 2015).

Despite these advances in characterizing the Galactic magnetic field, fundamental questions persist regarding its origin. Competing theories propose either a primordial field predating Galaxy formation or a field generated and sustained by a Galactic dynamo, with current observations insufficient to definitively support or rule out either scenario (Zweibel, 2005). This uncertainty extends to the broader question of magnetogenesis in the Universe, highlighting the need for continued research in this field (Durrer & Neronov, 2013).

1.3 The ISM as a Cosmological Foreground

The cosmic microwave background (CMB) radiation, a relic of the early Universe, plays a pivotal role in modern cosmology. Discovered serendipitously by Penzias & Wilson (1965) and Dicke et al. (1965), the CMB is thermal radiation peaking at microwave frequencies, with a characteristic blackbody spectrum corresponding to a temperature of approximately 2.7 K (Fixsen, 2009). This ubiquitous background radiation provides compelling evidence for the Big Bang theory, as it represents the cooled remnant of the hot, dense state of the early Universe (Dicke et al., 1965).

The CMB's uniform temperature across the sky, with only minute fluctuations on the order of 10^{-5} K, strongly supports the concept of a rapid expansion from a homogeneous initial state (Smoot et al., 1992). These temperature anisotropies, primarily observed at millimeter wavelengths, encode crucial information about the Universe's composition, geometry, and evolution (Planck Collaboration et al., 2020c).

While the standard cosmological model successfully explains many observed phenomena, it faced several challenges, including the flatness, horizon, and monopole problems (Guth, 1981). The inflationary paradigm was proposed to address these issues, positing a period of exponential expansion in the very early Universe. This theory not only resolves the aforementioned problems but also provides a framework for understanding the Universe's large-scale structure and isotropy (Baumann & Peiris, 2008).

A key prediction of inflationary models is the presence of primordial gravitational waves, which should leave an imprint on the CMB in the form of *B*-mode polarization (Seljak, 1997). While *E*-mode polarization, primarily produced by scalar density perturbations, has been observed (Kovac et al., 2002), the detection of primordial *B*-mode polarization remains elusive. This signal is distinct from the *B*-mode polarization caused by gravitational lensing, which has been detected at smaller angular scales (Hanson et al., 2013).

The search for primordial *B*-mode polarization has been complicated by foreground contamination, particularly from polarized dust emission in the ISM. This dust emission is most prominent at frequencies > 70 GHz and becomes increasingly dominant at higher frequencies (Planck Collaboration et al., 2020b). The BICEP2 collaboration's initial claim of detecting primordial *B*-mode polarization (BICEP2 Collaboration et al., 2014) was later shown to be consistent with the polarized dust signal (Flaumer et al., 2014; BICEP2/Keck Collaboration et al., 2015). This highlights the critical importance of accurately characterizing and modeling the ISM's contribution to the polarized emission.

The challenge posed by the ISM to cosmological studies extends beyond the search for inflationary signatures. It affects our ability to probe the history of the Universe across various scales and epochs. The ISM impacts measurements of the kinetic Sunyaev-Zeldovich effect, which probes the epoch of reionization and the distribution of ionized gas in the Universe (Reichardt, 2016). Galactic synchrotron emission, prominent at frequencies below 70 GHz, interferes with efforts to detect signals from the cosmic dawn and the epoch of reionization (Bernardi et al., 2009). The ISM’s magnetic fields and turbulence influence the propagation of ultra-high-energy cosmic rays, affecting our understanding of their origins and the most energetic processes in the Universe (Aharonian et al., 2019). Searches for cosmic birefringence, a phenomenon that could provide evidence for parity-violating physics or new scalar fields, require extremely precise foreground characterization. Mismodeling of polarized foregrounds could be mistaken for a birefringence signal, limiting our ability to constrain this subtle effect (Minami & Komatsu, 2020; Cukierman et al., 2023). As such, advancements in our understanding of the magnetized ISM are not merely tangential to cosmology but are integral to our ability to unravel the Universe’s evolutionary history.

While the ISM presents obstacles to cosmological investigations across a wide range of wavelengths, it also offers an opportunity for interdisciplinary research. Improving our knowledge of the ISM’s structure, composition, and magnetic properties will not only benefit astrophysics but will also be important for future cosmological discoveries, potentially including the definitive detection of primordial gravitational waves.

1.4 Outline

In this chapter, we have introduced the key concepts and context necessary for understanding the relationship between the three-dimensional structure of the ISM and polarized dust emission. The thesis aims to explore this relationship and what we can learn from it using various datasets and techniques, with the ultimate goal of improving our understanding of astrophysical processes and enhancing our ability to study the CMB. The research presented in this Thesis is structured into four main chapters, each addressing specific aspects of this relationship and its implications.

Chapter 2 investigates the influence of the 3D geometry of the nearby interstellar medium on the statistics of the polarized dust emission on large scales. It tests models of the Local Bubble geometry and explores how the complexity of the 3D dust distribution affects the dust polarization fraction. This chapter establishes the foundation for the subsequent investigations by highlighting the importance of considering the extended 3D distribution of the dust.

Chapter 3 focuses on developing a new method for characterizing Galactic dust filaments. This approach involves correlating observational data from BICEP/Keck and Planck with 3D polarization templates derived from H α observations. The chapter aims to push the detection limits of polarized dust emission to lower frequencies and identify the specific ISM components that contribute to this emission and characterize them.

Building on the methodologies established in Chapter 2, Chapter 4 further refines the

H_I-based polarization templates. This chapter introduces a new algorithm, the Spherical Rolling Hough Transform, designed to efficiently quantify filamentary structures on the sphere. The goal is to investigate how different aspects of filament morphology impact the observed polarization patterns and improve our ability to model these effects.

Chapter 5 represents the culmination of our research, focusing on the development of a transformer-based model for generating high-resolution, non-Gaussian foreground models of polarized dust emission. This chapter aims to fuse information from various sources to predict small-scale dust structures, advancing the field of foreground modeling for CMB studies.

Throughout these chapters, we explore how the 3D structure of the ISM fundamentally shapes the observed polarized dust emission. Each chapter builds upon the findings of the previous ones, creating a cohesive narrative that demonstrates the importance of considering 3D geometry in interpreting and modeling polarized dust emission data.

We conclude with Chapter 6, which provides a summary of our results and discusses the future outlook for this field of research.

Chapter 2

Imprints of the Local Bubble and Dust Complexity on Polarized Dust Emission

Abstract

Using 3D dust maps and Planck polarized dust emission data, we investigate the influence of the 3D geometry of the nearby interstellar medium (ISM) on the statistics of the dust polarization on large (80') scales. We test recent models that assume that the magnetic field probed by the polarized dust emission is preferentially tangential to the Local Bubble wall, but we do not find an imprint of the Local Bubble geometry on the dust polarization fraction. We also test the hypothesis that the complexity of the 3D dust distribution drives some of the measured variation of the dust polarization fraction. We compare sightlines with similar total column densities and find that, on average, the dust polarization fraction decreases when the dust column is substantially

distributed among multiple components at different distances. Conversely, the dust polarization fraction is higher for sightlines where the dust is more concentrated in 3D space. This finding is statistically significant for the dust within 1.25 kpc, but the effect disappears if we only consider dust within 270 pc. In conclusion, we find that the extended 3D dust distribution, rather than solely the dust associated with the Local Bubble, plays a role in determining the observed dust polarization fraction at 80'. This conclusion is consistent with a simple analytical prediction and remains robust under various modifications to the analysis. These results illuminate the relationship between the 3D geometry of the ISM and tracers of the interstellar magnetic field. We discuss implications for our understanding of the polarized dust foreground to the cosmic microwave background.

2.1 Paper Status and External Contributions

This chapter is based on the article published by The Astrophysical Journal under the title, "Imprints of the Local Bubble and Dust Complexity on Polarized Dust Emission" (Halal et al., 2024b). I am the first author of this paper. I performed all of the analysis, wrote all of the text, and produced all of the figures for this paper. However, this work is the result of weekly discussions, advising, and extensive editorial input from my advisor Susan Clark and postdoctoral scholar Mehrnoosh Tahani.

2.2 Introduction

Interstellar magnetic fields play an important role in various astrophysical processes (see, e.g., Ferrière, 2001; Heiles & Haverkorn, 2012; Pattle et al., 2023). However, little is known about the magnetic field structure in the nearby interstellar medium (ISM). Some works have suggested connections between the local magnetic field structure and other tracers of ISM morphology, perhaps due to dynamical influences, e.g. the formation of structures like superbubbles (e.g., Santos et al., 2011; Frisch et al., 2012; Berdyugin et al., 2014; Tahani et al., 2022a,b).

Aspherical dust grains in the ISM emit photons with an electric field oriented preferentially along their long axes (Purcell, 1975). The short axes of typical dust grains are preferentially aligned with the local magnetic field orientation (Andersson et al., 2015). As a result, their thermal emission is partially polarized perpendicular to the orientation of the magnetic field. Therefore, measurements of the polarized dust emission are used as a probe for the plane-of-sky magnetic field orientation in dusty regions of the ISM.

Recent evidence suggests that variations in the fractional polarization of the dust emission, i.e., the ratio of the polarized to the total intensity of the dust emission, over large angular scales in the diffuse sky are mainly driven by the structure of the magnetic field. Henceforth, we will refer to the fractional polarization of the dust emission as the dust polarization fraction. Planck Collaboration et al. (2020a) probed the influence of the magnetic field geometry on the dust polarization fraction by comparing the local polarization

angle dispersion with the 353 GHz polarization fraction. They calculated the polarization angle dispersion for an annulus with inner and outer radii of 40' and 120', and found that the 353 GHz polarization fraction is anti-correlated with the local polarization angle dispersion at 160' resolution. They showed that this relationship is consistent with models that only include topological effects of the turbulent magnetic field, but otherwise have uniform dust properties and alignment. Planck Collaboration et al. (2020a) conclude that the dust polarization fraction and the dispersion of polarization angles are similarly sensitive to the structure of the magnetic field. Hensley et al. (2019) further showed that some of the variability in the dispersion of polarization angles, and thus the dust polarization fraction, can be explained by the magnetic inclination angle, i.e., the angle between the magnetic field and the plane of the sky. The dust polarization fraction is maximized when the magnetic field is tangential to the plane of the sky and zero when it is parallel to the line of sight. Chen et al. (2019b) and Sullivan et al. (2021) used statistical properties of the observed dust polarization fraction to estimate the average inclination angle of molecular cloud-scale magnetic fields.

The observed dust polarization fraction also depends on other factors, such as the dust grain alignment efficiency (King et al., 2019; Medan & Andersson, 2019), the phase distribution of the neutral interstellar medium (Lei & Clark, 2023a), and measurement noise. However, the 3D structure of the magnetic field is one of the major factors (Clark, 2018; Hensley et al., 2019; Planck Collaboration et al., 2020a).

Our Sun's current location is near the center of a superbubble, which is thought to have been created by supernova explosions within the past $10\text{-}15 \times 10^6$ years (Cox & Reynolds, 1987; Maíz-Apellániz, 2001; Breitschwerdt et al., 2016; Schulreich et al., 2023). It is commonly known as the Local Bubble, Local Cavity, or Local Chimney (Welsh et al., 2004; Puspitarini & Lallement, 2012). Since star formation tends to be concentrated, sequential supernovae are common (Zucker et al., 2022; Watkins et al., 2023; Barnes et al., 2023; Sandstrom et al., 2023). Supernova explosions sweep up matter and magnetic field lines, leaving behind low-density superbubbles on the order of hundreds of parsecs in diameter (Kim & Ostriker, 2015). The swept-up matter is compressed into a shell surrounding the expanding superbubble, which is thought to trigger the formation of dense gas and stars (Elmegreen, 2011; Dawson, 2013; Inutsuka et al., 2015).

Thus, it is reasonable to expect that the formation of the Local Bubble dramatically influenced the magnetic field geometry in the nearby ISM. Some studies have aimed at modeling the geometry of the wall of cold neutral gas and dust surrounding the Local Bubble (e.g., Alves et al., 2018; Pelgrims et al., 2020). Since the geometry of the magnetic field affects the measured dust polarization fraction, we search for an imprint of the Local Bubble geometry on the dust polarization fraction in this paper.

Additionally, one of the probes of the 3D spatial distribution of the neutral ISM is dust extinction toward stars. This is due to the scattering and absorption of starlight by dust. The extinction of a star's apparent magnitude is correlated with the dust column density along

the line of sight from the observer to that star. The Gaia survey provided accurate distances to more than a billion stars within a few kiloparsecs from the Sun. Combining this distance information with the level of extinction towards each star has been transformative for the construction of 3D maps of the differential dust extinction (Lallement et al., 2019; Leike et al., 2020; Vergely et al., 2022a; Edenhofer et al., 2023). We use several 3D dust maps to quantify the complexity of the spatial distribution of the dust along the line of sight. We use that to explore the relationship between the 3D dust distribution and the measured dust polarization fraction.

In this work, we investigate the relationship between the 3D geometry of the nearby ISM and the dust polarization fraction. We start by introducing the data we use in Section 2.3. In Section 2.4, we search for an imprint of the Local Bubble geometry on the dust polarization fraction. In Section 2.5, we test how the dust polarization fraction is affected by the line-of-sight complexity of the dust. We discuss the implications of our results and conclude in Section 2.6.

2.3 Data

2.3.1 Planck Data Products

We use the 80' R3.00 Planck data processed with the Generalized Needlet Internal Linear Combination (GNILC; Remazeilles et al., 2011) method at 353 GHz to remove the Cosmic

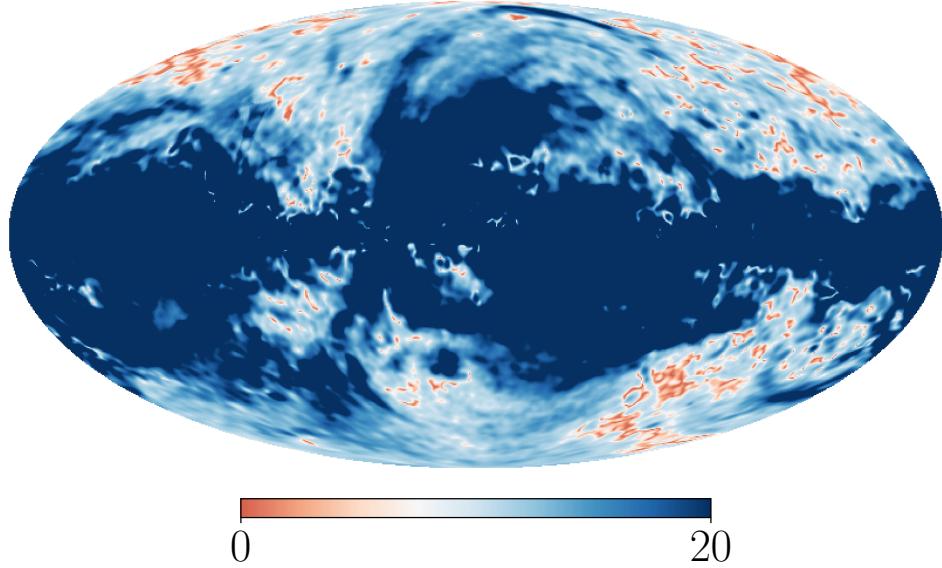


Figure 2.1: A map of the debiased signal-to-noise ratio of the Planck GNILC polarization fraction at $80'$. This is plotted with a diverging linear colorbar centered on 3, the cutoff we use as part of our sightline selections in Sections 2.4 and 2.5, with the allowed regions shown in blue.

Infrared Background (CIB) radiation from the Galactic dust emission (Planck Collaboration et al., 2016c). Following the fiducial offset corrections adopted by the Planck collaboration, we subtract $452 \mu\text{K}_{\text{CMB}}$ from the GNILC total intensity map to correct for the CIB monopole then add a Galactic offset correction of $63 \mu\text{K}_{\text{CMB}}$ (Planck Collaboration et al., 2020a). Because the GNILC data are at $\text{FWHM} = 80'$, we downgrade the maps from their native HEALPix pixelization (Górski et al., 2005) at $N_{\text{side}} = 2048$ to $N_{\text{side}} = 64$. We also use the R3.01 Planck data at 353 GHz, smoothed to $80'$ as a cross-check (Planck Collaboration et al., 2020d). All of these maps use the COSMO polarization convention. We do not convert to the IAU polarization convention.

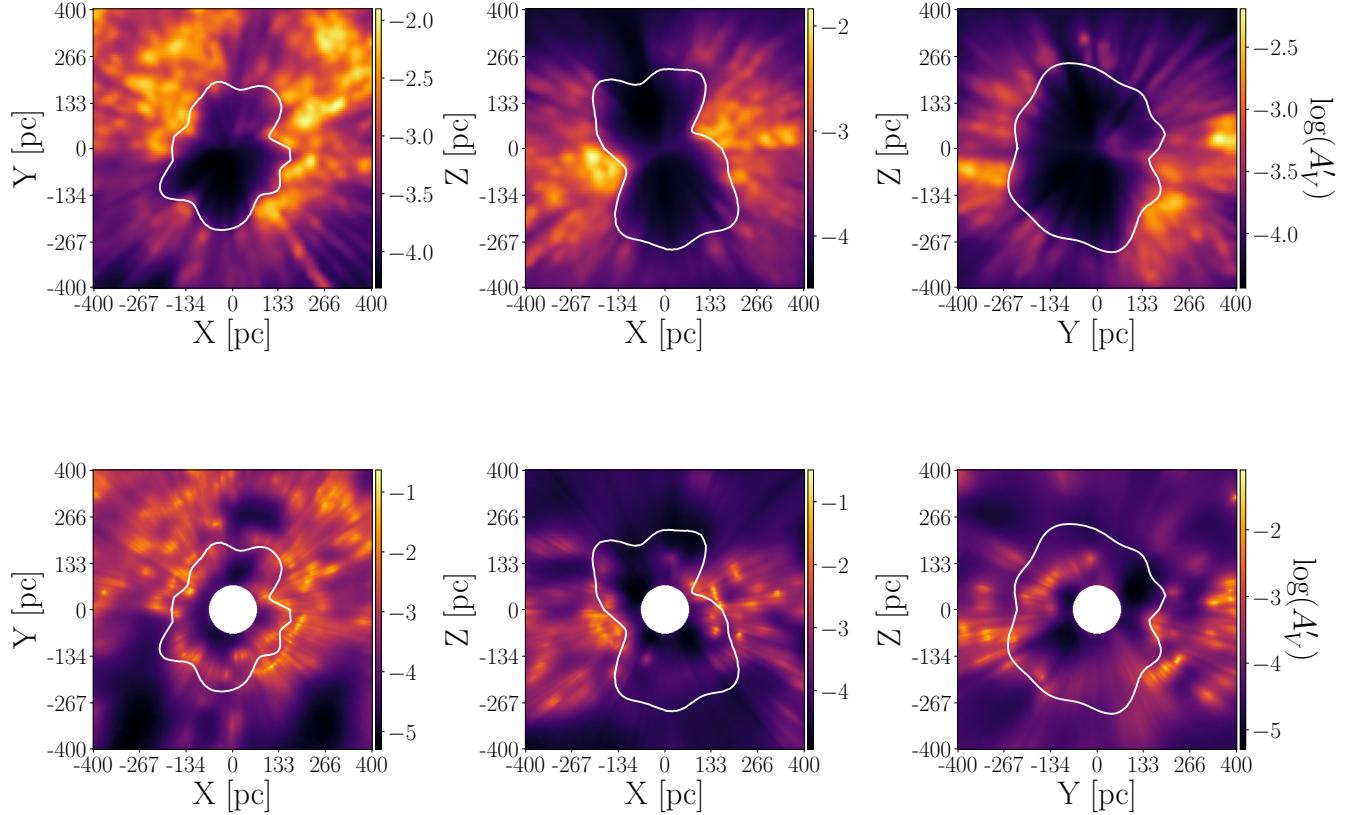


Figure 2.2: Slices through the XY (left), XZ (middle), and YZ (right) planes of the 3D reconstructed differential extinction maps of Lallement et al. (2019) (top) and Edenhofer et al. (2023) (bottom). The Sun is at the origin. The positive X axis points towards the Galactic center at $l = 0^\circ$, the positive Y axis points towards $l = 90^\circ$ in the Galactic plane, and the positive Z axis points out of the plane in the direction of the Galactic North pole. The slices show the log of the differential extinction, which is in units of magnitudes per parsec. The subpanels only extend up to 400 pc in each direction for a direct comparison. The model for the Local Bubble surface geometry of Pelgrims et al. (2020) is overplotted in white in each subpanel. The white disk at the center of the bottom panel is due to missing data within 70 pc of the Sun in the Edenhofer et al. (2023) maps.

We use the modified asymptotic estimator of Plaszczyński et al. (2014) to debias the polarized intensity of the GNILC dust emission map and the associated uncertainty. We obtain a signal-to-noise ratio map of the dust polarization fraction, SNR_p , shown in Figure 2.1.

We estimate the total dust extinction over the full sky using Planck data products. Planck Collaboration et al. (2014b) and Planck Collaboration et al. (2016c) fit a modified blackbody spectrum to the GNILC dust maps at different frequencies to estimate the dust temperature, spectral index, and optical depth over the sky. The dust optical depth is correlated with the reddening of quasars (Planck Collaboration et al., 2014b). Using this observation, Planck Collaboration et al. (2016c) multiply the GNILC dust optical depth map by a factor of 1.49×10^{-4} mag to construct a GNILC $E(B - V)$ map. We query the publicly available `dustmaps` Python package (Green, 2018) for the Planck Collaboration et al. (2016c) GNILC $E(B - V)$ map. Assuming a standard extinction law, we multiply the GNILC $E(B - V)$ map by 3.1 to obtain A_V^{Planck} .

2.3.2 3D Dust Maps and Local Bubble Geometries

In Section 2.4, we use the 3D model of the Local Bubble surface geometry constructed by Pelgrims et al. (2020). To create this model, Pelgrims et al. (2020) extract distances to the first high dust density regions around the Sun from the 3D Cartesian map of dust differential extinction constructed by Lallement et al. (2019). They smooth the map of the distances to the Local Bubble surface by filtering out spherical harmonic modes above some threshold to remove small-scale fluctuations that might appear due to an inhomogeneous distribution of the dust density on the small scales. In this work, we use the map filtered to $\ell_{\text{max}} = 10$, which is the map used in their analysis. The Lallement et al. (2019) dust map is based on

data from Gaia DR2 (Gaia Collaboration et al., 2018) and 2MASS (Skrutskie et al., 2006) and spans $6 \times 6 \times 0.8$ kpc³ in the Heliocentric right-handed Galactic-XYZ coordinates. It has a voxel volume of 125 pc³ and a spatial resolution of 25 pc. Example slices of this map with the Pelgrims et al. (2020) model overplotted are shown in Figure 2.2.

We also use the 3D model of the Local Bubble surface constructed by O’Neill et al. (2024) for a brief investigation in Section 2.4. They use the 3D dust differential extinction provided by Edenhofer et al. (2023) to extract the distance to the Local Bubble in all directions as a region of higher dust density around the Sun.

In Section 2.5, we query the 12 posterior samples for the 3D dust maps provided by Edenhofer et al. (2023) via the publicly available `dustmaps` Python package (Green, 2018) at their plane-of-sky native angular resolution of 14', which corresponds to a HEALPix pixelization scheme at $N_{\text{side}} = 256$. These 3D dust maps leverage distance and extinction estimates to stars from Zhang et al. (2023), which are derived from Gaia DR3 data (Gaia Collaboration et al., 2023). The distance resolution of these maps varies from 0.4 pc at 69 pc to 7 pc at 1.25 kpc. We query the map using uniform distance bins of 7 pc. The map is provided in unitless extinction values defined in Zhang et al. (2023). We multiply the map by a factor of 2.8 to convert it to Johnson’s V-band A'_V (Zhang et al., 2023). We then convert A'_V to volume density of hydrogen nuclei (n_H) using the extinction curve from Fitzpatrick et al. (2019) to convert $A_G = 0.796 A_V$ and the relationship $A_G / N_H = 4 \times 10^{-22}$ cm² mag from Zucker et al. (2021) and Bialy et al. (2021). To match the resolution of the dust

polarization fraction and smooth out small-scale fluctuations in the map, we smooth the HEALPix sphere at each distance bin to a FWHM= 80' then repixelate it to $N_{\text{side}} = 64$. Example slices of the raw map of the mean of the posterior samples are also shown in Figure 2.2 with the Pelgrims et al. (2020) model overplotted for comparison.

In Section 2.5, we also make comparisons with the 3D dust maps provided by Leike et al. (2020). These maps leverage distance and extinction estimates from the StarHorse catalog (Anders et al., 2019), which combines data from Gaia DR2 (Gaia Collaboration et al., 2018), ALLWISE (Cutri et al., 2013), PANSTARRS (Flewelling et al., 2020), and 2MASS (Skrutskie et al., 2006). The Leike et al. (2020) maps span $740 \times 740 \times 540 \text{ pc}^3$ in the Heliocentric Galactic-XYZ coordinates, respectively, with a voxel size of 1 pc^3 and spatial resolution of 1 pc. We also query the 12 posterior samples of the Leike et al. (2020) maps via the `dustmaps` package. The Leike et al. (2020) maps are given in optical depth in the Gaia G band per 1 pc. We convert to n_{H} following Zucker et al. (2021) and Bialy et al. (2021).

2.3.3 Galactic Faraday Rotation Measure

We use the all-sky Galactic Faraday rotation measure (RM) map produced by Hutschenreuter et al. (2023) for a brief investigation in Section 2.4. Using information field theory, a Bayesian inference framework for fields (Enßlin, 2019), Hutschenreuter et al. (2023) disentangle the Galactic contribution to the RM from the compiled RM catalogs of polarized

radio sources such as radio galaxies (Van Eck et al., 2023), supplemented by Galactic pulsar dispersion measures (Manchester et al., 2005), as well as data on Galactic bremsstrahlung emission (Planck Collaboration et al., 2016d) and the hydrogen α spectral line (Finkbeiner, 2003).

2.4 No Imprint of the Local Bubble on the Dust Polarization Fraction

2.4.1 Motivation

We begin the exploration of the effect of different geometrical factors on the dust polarization fraction by searching for an imprint of the geometry of the dust wall surrounding the Local Bubble on the Planck 353 GHz dust polarization fraction. In this subsection, we discuss the assumptions made in previous studies regarding the Local Bubble surface. These assumptions help us design a test for studying this effect in the next subsection. A significant detection of an imprint of the Local Bubble geometry on the dust polarization fraction would validate these assumptions.

The first assumption is that the observed polarized dust emission is dominated by dust in the Local Bubble wall at the relevant angular scales and Galactic latitudes. This assumption, made in several analyses (e.g., Alves et al., 2018; Pelgrims et al., 2020; O’Neill et al., 2023), is supported by several studies using optical starlight polarization data (Leroy, 1999;

Andersson & Potter, 2006; Santos et al., 2011; Frisch et al., 2015; Medan & Andersson, 2019; Cotton et al., 2019; Skalidis & Pelgrims, 2019). The alignment of neutral hydrogen structures at local velocities with starlight polarization toward stars at distances within a few hundred parsecs is consistent with a picture where most of these structures are positioned at comparable distances within the Local Bubble at high Galactic latitudes ($|b| > 30^\circ$) (Clark et al., 2014). Gontcharov & Mosenkov (2019) observed that starlight polarization fraction plateaus after 150-250 pc across the sky. By comparing the 353 GHz polarized emission with the polarized optical starlight, Skalidis & Pelgrims (2019) find that most of the 353 GHz polarized emission signal is captured within the first 250 pc at $|b| > 60^\circ$, suggesting the presence of a dust wall around that distance. They, however, find that this conclusion does not hold at $30^\circ < |b| < 60^\circ$.

The second assumption is that the magnetic field's inclination is tangential to the surface of the Local Bubble, which stems from the model assumed for the formation of the Local Bubble. To fit a model of the Local Bubble magnetic field, Alves et al. (2018) and Pelgrims et al. (2020) assume that all the swept-up matter and field lines due to the supernova explosions that formed the Local Bubble are squeezed into a thin layer that follows its surface, leading the magnetic field lines to be tangent to the surface. O'Neill et al. (2023) also make this assumption to project the observed polarization angles of the dust emission onto the Bubble's surface and build a 3D model of the Bubble wall magnetic field.

Other works have found magnetic field structure tangential to bubbles on supernova

scales, observationally (e.g., Kooches & Brown, 2009; West et al., 2016; Tahani et al., 2022a,b) and in simulations (Kim & Ostriker, 2015; Maconi et al., 2023). It has also been shown on the scales of HII regions, observationally (Tahani et al., 2023) and in simulations (Krumholz et al., 2007). On the scale of superbubbles, such as our Local Bubble, a comparison of the plane-of-the-sky and line-of-sight magnetic field strengths as well as measurements of the dust polarization fraction towards regions associated with the Orion-Eridanus superbubble suggest that the large-scale magnetic field in the region was primarily shaped by the expanding superbubble and is tangential to its surface (Heiles, 1997b; Soler et al., 2018). However, there is no direct evidence that the nearby magnetic field is preferentially tangential to the Local Bubble surface.

The shape of the Local Bubble wall has been modeled differently in different works. Some studies fit a generalized parametric geometry, such as an ellipsoid (Alves et al., 2018), in an attempt to fit general properties of the Local Bubble magnetic field on large scales. Others model the detailed boundary of the Local Bubble using 3D maps of the dust extinction (Pelgrims et al., 2020; O'Neill et al., 2024). These models therefore vary based on both the variations in the different 3D dust maps used as well as the methodology applied to these maps to define a Local Bubble surface. Pelgrims et al. (2020) model the radial distance of the Local Bubble wall from the Sun in each direction as the first distance where the second derivative with respect to the distance of the differential extinction constructed by Lallement et al. (2019) reaches zero, i.e., the first inflection point, $d^2A'_V(r)/dr^2 = 0$.

O’Neill et al. (2023) and Zucker et al. (2022) use the geometry defined by Pelgrims et al. (2020) in their work. O’Neill et al. (2024) employ a similar methodology to Pelgrims et al. (2020), using the differential extinction maps constructed by Edenhofer et al. (2023) instead to construct their model. Other studies use different tracers to model the geometry of the Bubble. Liu et al. (2017) assumes that the measured X-ray intensity is proportional to the distance to the Bubble in the considered direction. Several other tracers such as NaI absorption measurements (Sfeir et al., 1999; Lallement et al., 2003), stellar color excess measurements (Lallement et al., 2014), and diffuse interstellar bands (Farhang et al., 2019), have also been used for constructing models of the Bubble wall geometry. These geometries vary significantly from one to another. While some model the Local Bubble as a closed surface (Pelgrims et al., 2020), others describe the same structure as a Local Chimney, i.e., open in one or both directions away from the disk and funneling material into the Milky Way’s halo (e.g., Heiles, 1984; Sfeir et al., 1999; Lallement et al., 2003; Marchal & Martin, 2023; O’Neill et al., 2024). Also, the Local Bubble surface may have tunnels to surrounding cavities like the Gum Nebula and/or GSH238+00+09 rather than having a closed geometry (e.g., Welsh, 1991; Lallement et al., 2003; Marchal & Martin, 2023; O’Neill et al., 2024).

We aim to test whether the degree-scale structure of the Local Bubble is measurably imprinted in the statistics of polarized dust emission. We use the most recent models of the Local Bubble geometry constructed based on 3D dust mapping. For our main analysis, we use the Pelgrims et al. (2020) model, which has a closed geometry. However, we also

perform a brief test using the O’Neill et al. (2024) model, whose geometry is not fully connected. In each case, we test whether a magnetic field that is tangential to the Local Bubble wall is a statistical driver of the observed variation in the dust polarization fraction.

2.4.2 Testing the Dependence of the Dust Polarization Fraction on the Magnetic Inclination Angle

Using 3D dust extinction maps and a model of the Local Bubble surface geometry, we can measure the angle between the line of sight and the local magnetic field orientation projected onto the Bubble’s surface in each direction. If the magnetic field lines were tangential to the Local Bubble surface as discussed in Section 2.4.1, the angle between the line of sight and the Bubble’s surface should on average be correlated with the measured dust polarization fraction along sightlines where the dust is concentrated in the Local Bubble wall. Therefore, by quantifying the correlation between this angle and the dust polarization fraction for different sightlines, we can test whether there is measurable evidence that the magnetic field traced by the dust polarization is preferentially tangential to the Local Bubble surface.

The magnetic inclination angle γ , i.e., the angle between the magnetic field orientation and the plane of the sky, affects the measured dust polarization fraction. The polarization fraction, p , is defined as

$$p = P/I = \sqrt{Q^2 + U^2}/I, \quad (2.1)$$

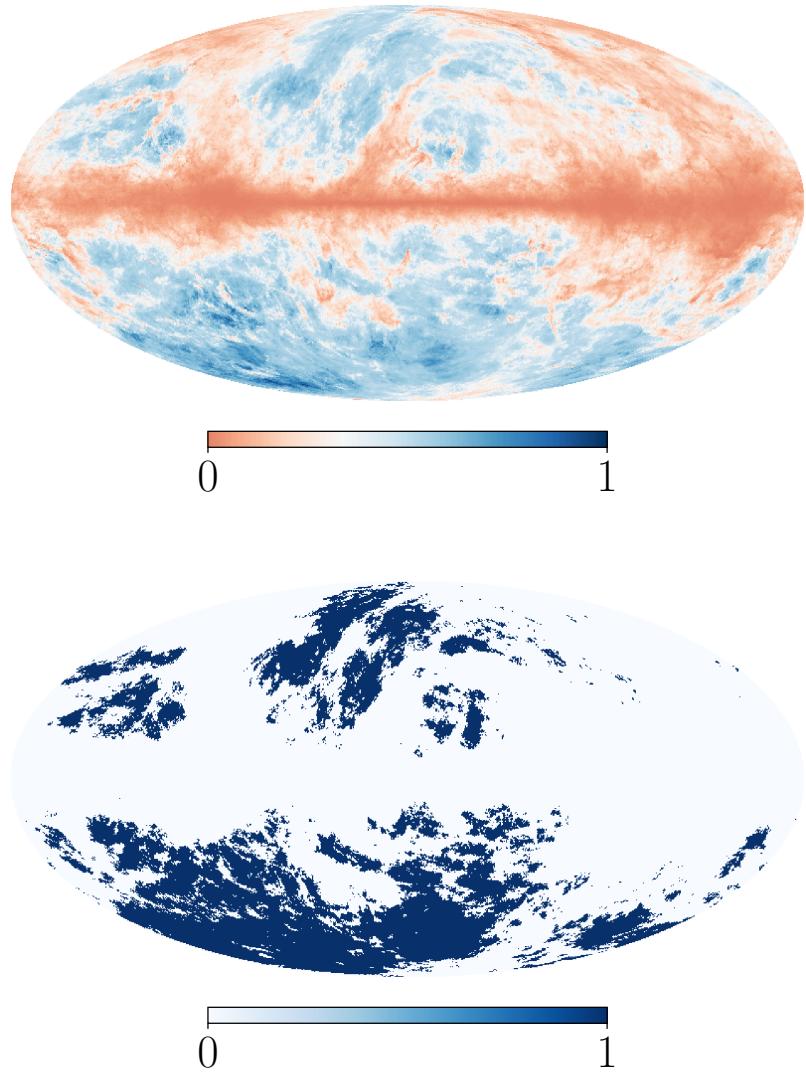


Figure 2.3: Maps used for selecting sightlines for the analysis described in Section 2.4. *Top panel*: A map of the ratio of the Lallement et al. (2019) 3D dust differential extinction map integrated within 50 pc of the Local Bubble surface defined by Pelgrims et al. (2020), A_V^{LB} , over the Planck dust extinction, A_V^{Planck} . This is plotted with a diverging colorbar centered on the 75th percentile (0.2), the cutoff we use in our sightline selection, with the allowed regions shown in blue. *Bottom panel*: A map of the mask of the selected region, combining $A_V^{\text{LB}}/A_V^{\text{Planck}} > 0.2$ (top panel) with $\text{SNR}_p > 3$ (Figure 2.1).

where P is the debiased polarized intensity as described in Section 2.3, Q and U are the Stokes parameters, and I is the total unpolarized intensity. Assuming uniform grain properties, the Stokes I , Q , and U parameters of the dust emission can be written as (Fiege & Pudritz, 2000; Padoan et al., 2001; Pelkonen et al., 2007)

$$I = \int \epsilon \kappa ds - \frac{1}{2} \int \alpha \epsilon \kappa \left(\cos^2 \gamma - \frac{2}{3} \right) ds, \quad (2.2)$$

$$Q = - \int \alpha \epsilon \kappa \cos 2\psi \cos^2 \gamma ds, \quad (2.3)$$

$$U = - \int \alpha \epsilon \kappa \sin 2\psi \cos^2 \gamma ds, \quad (2.4)$$

where κ is the volume density, ds is a distance segment along the line of sight, ϵ is the dust emissivity, α is a coefficient defined in Equation 15 of Padoan et al. (2001) that is a product of polarization efficiency factors, such as the degree of dust alignment and the dust grain polarization cross-section, and ψ is the angle between the projection of the magnetic field on the plane of the sky and South, and the Stokes parameters are given in the COSMO polarization convention. These equations show the dependence of the dust polarization fraction on γ , ψ , α , κ , and ϵ and their variations along the line of sight. At large angular scales and away from the Galactic plane, the variation in the observed dust polarization fraction is dominated by γ and ψ (Hensley et al., 2019; Planck Collaboration et al., 2020a).

To maximize the chances of detecting this correlation, we limit our analysis to sightlines where the dust in the Local Bubble wall contributes the most to the total extinction and where we have high signal-to-noise ratio measurements of the dust polarization fraction ($\text{SNR}_p > 3$; Figure 2.1). We then select sightlines in the highest quartile of $A_V^{\text{LB}} / A_V^{\text{Planck}}$, i.e., sightlines where the extinction in the Local Bubble wall has the highest contribution to the total observed extinction (Section 2.3.1). For A_V^{LB} , we integrate the 3D dust differential extinction maps of Lallement et al. (2019) in the Local Bubble wall. We integrate over 50 pc in each direction, starting with the distance to the inner Local Bubble surface as defined by the 3D model of Pelgrims et al. (2020). Note that the threshold that corresponds to the highest quartile is $A_V^{\text{LB}} / A_V^{\text{Planck}} \sim 0.2$. This does not necessarily mean that $\gtrsim 20\%$ of the total extinction is attributable to the Local Bubble wall, as the integration over 50 pc may not represent the true Bubble thickness for all sightlines, and A_V^{Planck} is estimated through a scale factor multiplied by the GNILC dust optical depth map (Section 2.3.1). Nevertheless, this represents a best estimate of the sky regions for which the Local Bubble wall accounts for the largest fraction of the total extinction. This map and the final mask are shown in Figure 2.3. While we define this mask to maximize the chance of detecting an imprint of the Local Bubble wall on the dust polarization, we also perform this analysis with different masks. If we redefine A_V^{Planck} using the integral from the Local Bubble surface to 100 pc, rather than 50 pc, 97% of the selected sightlines remain identical. The threshold that corresponds to the highest quartile increases to $A_V^{\text{LB}} / A_V^{\text{Planck}} \sim 0.4$ in

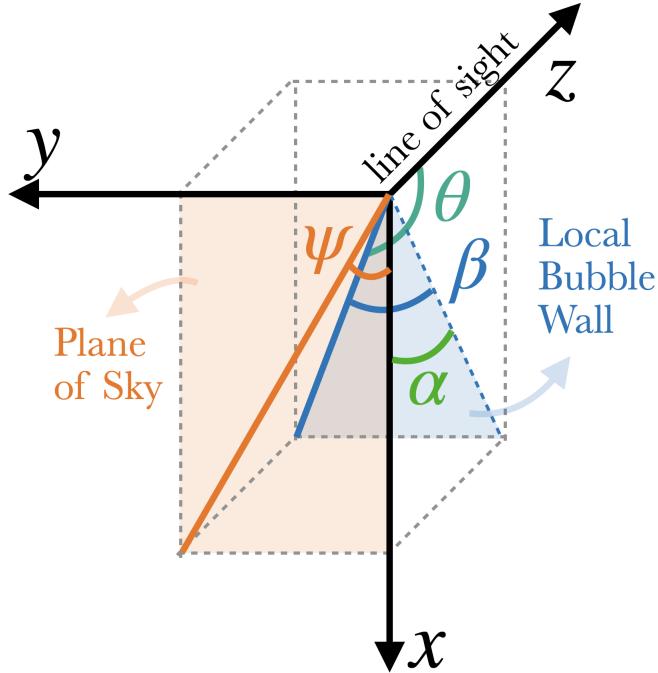


Figure 2.4: Diagram of the angles described and used in Section 2.4.3 for projecting the plane-of-sky magnetic field orientation (orange line) onto the Local Bubble wall (blue). The coordinates in this diagram follow the COSMO (HEALPix) convention, which is used in the Planck GNILC maps. These are not the Galactic coordinates used in Figure 2.2. For each position in the sky looking outwards, the local horizontal axis points South, the local vertical axis points East, and the local z-axis points outwards.

this case. We additionally perform this analysis with masks based on simple latitude cuts, including one focusing only on sightlines with $|b| > 60^\circ$. We find that the conclusions in Section 2.4.4 do not change for these different masks.

2.4.3 Magnetic Fields in the Local Bubble Wall

We perform the following steps to estimate the angle between the line of sight and the Local Bubble wall.

1. For the Local Bubble geometry, we use the 3D model developed by Pelgrims et al. (2020). Since this model is constructed based on the 3D dust map of Lallement et al. (2019), we use this dust map in this subsection. Example slices of the map with the model overlaid are shown at the top of Figure 2.2. The resolution of this dust map is 25 pc. We use a 3D Gaussian kernel with a standard deviation of 25 pc to smooth the data as done in Pelgrims et al. (2020). This smooths out spurious small-scale fluctuations in the data product that may affect the results.
2. We calculate the gradient of the differential extinction data cube, $\nabla A'_{V,i}$ for each voxel i . This is a vector in the direction of the steepest change in the 3D volume at a given voxel. For the voxels at the surface of the Local Bubble wall, $\nabla A'_{V,i}$ would therefore be orthogonal to that surface.
3. For each voxel i , we calculate the angle between the Local Bubble surface when projected onto the $x - z$ plane and the plane of the sky as

$$\alpha_i = \arccos \left(\frac{\nabla A'_{V,i} \cdot \mathbf{r}_i}{|\nabla A'_{V,i}| |\mathbf{r}_i|} \right), \alpha_i \in [0, \pi/2], \quad (2.5)$$

where \mathbf{r}_i is the line of sight vector, using the Sun, which is at the center of the data cube, as the origin. This angle is shown in Figure 2.4 in green.

4. Pelgrims et al. (2020) models the distance to the Local Bubble surface for each direction on a HEALPix sphere. We sample our 3D Cartesian cube of α_i at the radial

distance defined by the Pelgrims et al. (2020) model for each line-of-sight direction with a HEALPix pixelization scheme.

5. Using the Planck GNILC Stokes Q and U maps at 353 GHz, we calculate the plane-of-sky magnetic field orientation as

$$\psi = \frac{1}{2} \arctan \frac{-U}{-Q}. \quad (2.6)$$

This angle is shown in Figure 2.4 in orange.

6. We project the plane-of-sky magnetic field orientation onto the Local Bubble surface as

$$\beta = \arctan (\tan \psi \cos \alpha). \quad (2.7)$$

This angle is shown in Figure 2.4 in blue.

7. We calculate the angle between the line of sight and the magnetic field lines tangential to the surface of the Local Bubble as

$$\theta = \arccos (\cos \beta \sin \alpha). \quad (2.8)$$

This angle is shown in Figure 2.4 in teal.

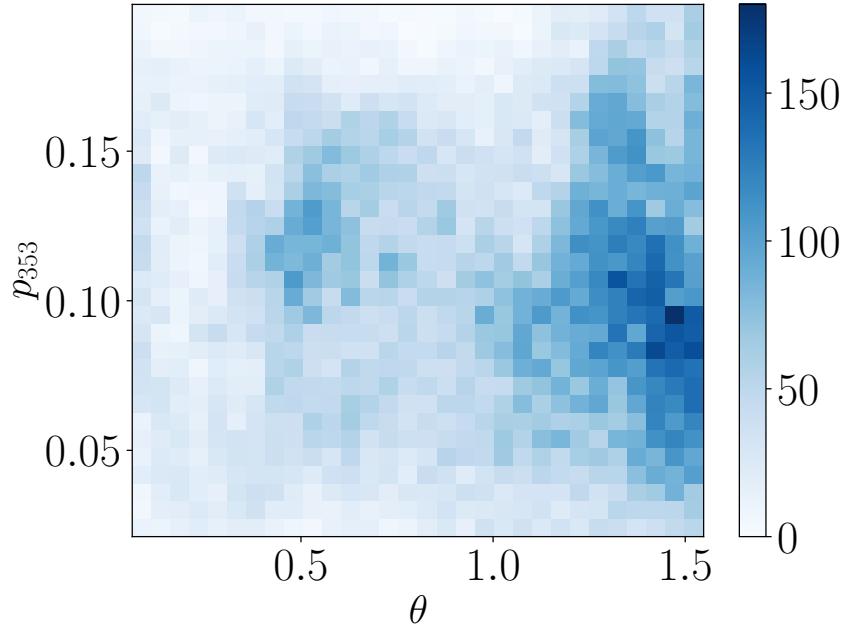


Figure 2.5: A 2D histogram showing the joint distribution of the dust polarization fraction and θ , the angle between the line of sight and the plane tangent to the surface of the Local Bubble (Equation 2.8) for the sightlines within the mask in Figure 2.3. The colorbar represents the number of sightlines in each bin. There is no significant correlation between these two quantities.

2.4.4 No Detected Imprint of the Local Bubble Wall on the Dust Polarization Fraction

We do not find any correlation between p_{353} and θ from Equation 2.8 over the mask defined in Figure 2.3. The joint distribution is shown as a 2D histogram in Figure 2.5. The Spearman rank coefficient, which is agnostic to the functional dependence between p_{353} and θ , is 3×10^{-3} . This lack in correlation persists for each of the masks described in Section 2.4.2. This indicates that at least one of the assumptions described in Section 2.4.1 is not valid. In other words, either the magnetic field is not generally tangent to the Local

Bubble surface, the dust polarization is not dominated by the dust within 50 pc of the surface as defined by the Pelgrims et al. (2020), or a combination of these possibilities.

We also test whether we find an anti-correlation between θ and the absolute value of the Faraday rotation measure (Section 2.3.3) but do not find any evidence for it. Therefore, we do not find an imprint of the detailed 3D geometry of the Local Bubble wall on the polarization statistics.

O'Neill et al. (2024) provide a map of the inclination angle between the Local Bubble wall as defined by their model and the plane of the sky. We also test for a correlation between this angle and the dust polarization fraction. We use a similar sightline selection criterion, replacing the Lallement et al. (2019) map with the Edenhofer et al. (2023) map and the Pelgrims et al. (2020) model with the O'Neill et al. (2024) model. We calculate a Spearman rank coefficient of 0.04, i.e., we do not find any correlation.

2.5 Imprint of Dust Complexity on Dust Polarization Fraction

We continue the investigation of how the 3D distribution of dust impacts the dust polarization fraction beyond the Local Bubble, taking into account the distribution of the dust in an extended volume around the Sun. We use the 3D dust maps constructed by Edenhofer et al. (2023) in this section. The benefits of these maps are that they have high resolution and

extend radially up to 1.25 kpc away from the Sun.

If contributions to the polarized dust emission originate from regions along the line of sight with differently oriented magnetic fields, the integrated signal will be depolarized relative to emission from a region with uniform magnetic fields. We postulate that sightlines with multiple dust components that are separated in distance and contribute similarly to the total column density are more likely to have substantial dust emission originating from regions with differently oriented magnetic fields. We test the hypothesis that on average, for sightlines with the same column density, the ones for which the dust is distributed into multiple components at different distances with similar contributions to the total column density are associated with higher levels of depolarization than those for which the dust contribution to the total column density is concentrated. Another way to state our hypothesis is that given two sightlines at the same column density, the one with a more complex 3D dust distribution will have a lower dust polarization fraction on average.

2.5.1 Sightline Selection

We select sightlines with high-fidelity dust polarization measurements ($\text{SNR}_p > 3$), that pass through regions with trustworthy 3D dust reconstruction, and that have a dominant contribution from the dust extinction within the 3D dust maps $A_V^{\text{Edenhofer}}$ to the total estimated extinction A_V^{Planck} (Section 2.3.1). The last constraint is to avoid sightlines at the lowest Galactic latitudes, where the dust extends in distance well beyond the regions where the dust

is mapped in 3D out to 1.25 kpc. We mask sightlines where the ratio $A_V^{\text{Edenhofer}}/A_V^{\text{Planck}} < 0.5$.

Edenhofer et al. (2023) use estimates of stellar distances and extinctions from Zhang et al. (2023) to construct their 3D dust maps. However, Edenhofer et al. (2023) found that their estimated extinction disagrees with Zhang et al. (2023) where the radially integrated differential extinction is below 50 mmag or above 4 mag. We therefore use those thresholds to mask sightlines where the dust differential extinction is likely to be significantly over- or under-estimated. The distance to those thresholds for a given sightline varies slightly for the 12 different posterior samples, so we apply a slightly different mask to each posterior sample. We only consider sightlines where the differential extinction integrated radially outwards reaches 50 mmag within 200 pc of the Sun. A threshold higher than 200 pc would include more sightlines in the selection, but it would shorten the minimum path length considered. Because in Section 2.5.5, we compare our results with the Leike et al. (2020) maps, which only extend out to 270 pc, we find 200 pc to be a good balance between having a large enough sample size and minimum path length. The minimum path length considered through the Edenhofer et al. (2023) maps for estimating the dust complexity is therefore 1.05 kpc.

A map of the distance at which the extinction reaches 50 mmag for different sightlines is shown at the top of Figure 2.6. This is for the first posterior sample, but the equivalent maps for the remaining 11 posterior samples look visually indistinguishable. We center the diverging colorbar in this subplot to 200 pc to show which sightlines pass the threshold.

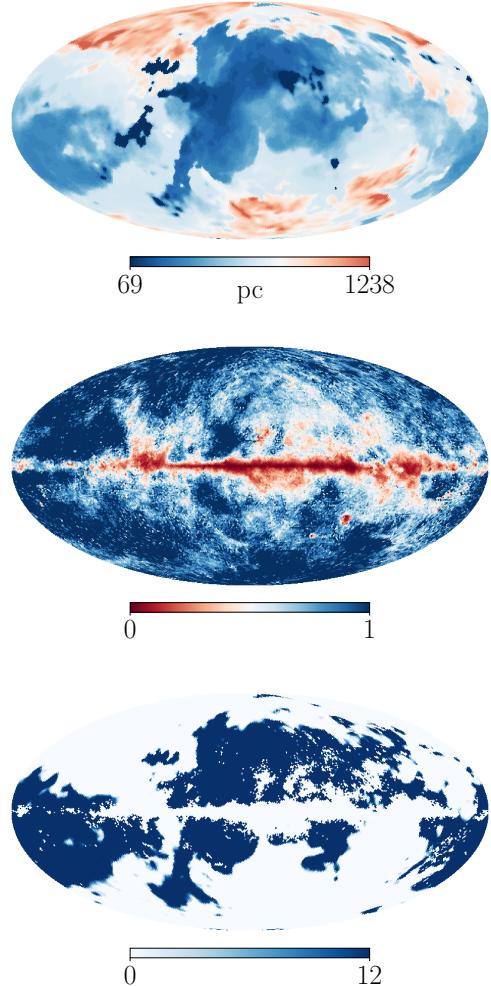


Figure 2.6: Maps used for selecting sightlines for the analysis described in Section 2.5. *Top panel*: A map of the distance at which the extinction in the first posterior sample of the Edenhofer et al. (2023) 3D dust maps reaches 50 mmag. This is plotted with a diverging colorbar centered on 200 pc, the cutoff we use in our sightline selection, with the allowed regions shown in blue. *Middle panel*: A map of the ratio of the Edenhofer et al. (2023) 3D dust differential extinction map integrated out to 1.25 kpc over the Planck dust extinction. This is plotted with a diverging colorbar centered on 0.5, the cutoff we use in our sightline selection, with the allowed regions shown in blue. *Bottom panel*: A sum of the masks of the selected regions over each of the 12 posterior samples, combining the selected regions from the quantities in the top panel, middle panel, and Figure 2.1.

After masking sightlines with $\text{SNR}_p < 3$ or with $A_V^{\text{Edenhofer}}/A_V^{\text{Planck}} < 0.5$ for the 12 posterior samples, the integrated differential extinction within 1.25 kpc is higher than 4 mag for only about 10 sightlines per posterior sample. We exclude these sightlines from our analysis.

After performing all the cuts, we are left with about 19,100 sightlines per posterior sample ($\sim 40\%$ of the sky) for this analysis. Those sightlines are shown at the bottom of Figure 2.6. The map shown is the sum of the binary masks over the 12 posterior samples. Note that the only discrepancies between the 12 posterior samples are at the edges of the selected regions. The rest of the map is either 12 or 0.

2.5.2 Dust Complexity

We aim to quantify the complexity of the 3D dust distribution along each line of sight. For this, we take inspiration from Panopoulou & Lenz (2020b), who perform a Gaussian decomposition of the neutral hydrogen (H I) emission spectra of each sightline to quantify its complexity. They use these components i , weighted by their column density N_{HI}^i , to define a metric

$$N_c^{\text{HI}} = \sum_{i=1}^{\text{nclouds}} \frac{N_{\text{HI}}^i}{N_{\text{HI}}^{\text{max}}}, \quad (2.9)$$

where $N_{\text{HI}}^{\text{max}}$ is the column density of the component with highest N_{HI} . This is a more relevant measurement of complexity than simply counting the number of components because it takes into account the relative contribution of each detected component, and the

dust polarization signal arises from a density-weighted integral along the line of sight. If there are two components along the line of sight and they have equal column densities, then $\mathcal{N}_c^{\text{HI}} = 2$, whereas if one has half the column density of the other, then $\mathcal{N}_c^{\text{HI}} = 1.5$, and so forth. Therefore, $\mathcal{N}_c^{\text{HI}}$ could be, for instance, 2 for any number of components larger than 1. This metric was used to detect an imprint of line-of-sight magnetic field tangling in dust polarization in Pelgrims et al. (2021a).

We use a metric inspired by Equation 2.9 to quantify the complexity of the 3D dust distribution, i.e., we decompose the 3D dust sightlines into "clouds," or components along the line of sight, and then use a version of Equation 2.9, replacing N_{HI} with N_{H} inferred from the dust extinction (Section 2.3), i.e.,

$$\mathcal{N}_c = \sum_{i=1}^{\text{n}_\text{clouds}} \frac{N_{\text{H}}^i}{N_{\text{H}}^{\text{max}}}. \quad (2.10)$$

To decompose each of the 19,100 sightlines into different components, we use the dendrogram technique (Rosolowsky et al., 2008). We use the Python package `astrodendro`¹. Cahlon et al. (2023) applied the 3D version of this technique on the Leike et al. (2020) 3D dust maps to produce a uniform catalog of molecular clouds in the Solar neighborhood. In this analysis, we apply the 1D version on individual sightlines.

Dendrogramming identifies density peaks in the data and connects them along hierarchical isosurfaces of constant n_{H} , forming a tree-like structure. We refer the reader to

¹<http://www.dendrograms.org/>

Rosolowsky et al. (2008) for a description of the core algorithm. We focus the explanation here on the algorithm's three parameters. The first parameter defines the minimum absolute Hydrogen number density n_{H} threshold for a structure to be included as part of the tree. We set this parameter to 0. This is because the tree-like structure constructed by this algorithm is not useful for our purposes. We only consider the density peaks identified.

The remaining two parameters, Δ_n and $\#\text{voxels}$, define the minimum prominence for a peak to be considered an independent component. Its largest n_{H} has to be Δ_n above the n_{H} of the adjacent isosurface for it to be considered an independent component from that isosurface. Similarly, $\#\text{voxels}$ defines the threshold number of voxels it has to span along a sightline to be considered an independent component, where each voxel spans 7 pc of the sightline. If a peak passes both of those thresholds, it is identified as a component by the algorithm.

For this analysis, we experiment with various values for Δ_n and $\#\text{voxels}$ and plot the resulting peaks identified for each variation. We find a range between $\Delta_n = 1.94 \times 10^{-3} \text{ cm}^{-3}$ ($A'_V = 3 \times 10^{-6} \text{ mag/pc}$) and $\Delta_n = 3.87 \times 10^{-3} \text{ cm}^{-3}$ ($A'_V = 9 \times 10^{-6} \text{ mag/pc}$) and between $\#\text{voxels} = 3$ (21 pc) and $\#\text{voxels} = 5$ (35 pc) recovers the visually identified peaks. We report our results for the fiducial values of $\Delta_n = 4.52 \times 10^{-3} \text{ cm}^{-3}$ ($A'_V = 7 \times 10^{-6}$), and $\#\text{voxels} = 3$ (21 pc). However, we find our results to be robust to all the variations we test for within the ranges mentioned.

We run the dendrogram algorithm on each sightline separately, considering the entire

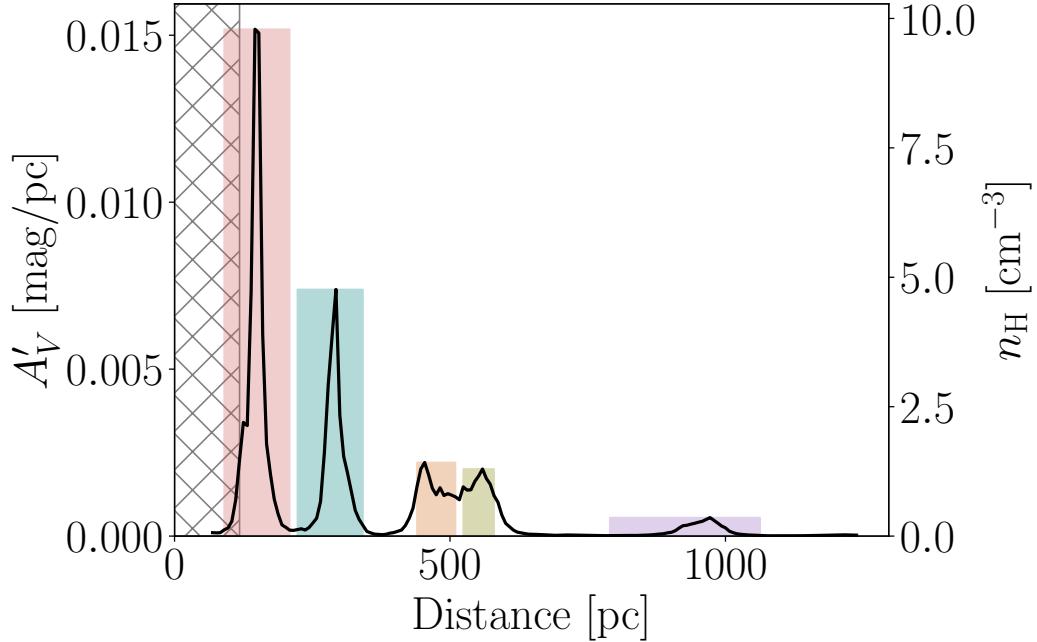


Figure 2.7: The dust distribution in units of differential extinction (mag/pc, left vertical axis) and equivalent Hydrogen number density (cm^{-3} , right vertical axis) along a representative sightline through the Edenhofer et al. (2023) maps. The Galactic coordinates of this sightline are $l = 163.12^\circ$ and $b = -11.42^\circ$. The region before the extinction reaches 50 mmag (hatched) is discarded from our analysis. The components identified by the dendrogram algorithm with $\Delta_n = 4.52 \times 10^{-3} \text{ cm}^{-3}$ ($A'_V = 7 \times 10^{-6}$), and $\#\text{voxels} = 3$ (21 pc) are shaded in different colors.

sightline from 69 pc to 1.25 kpc. However, we only keep identified components with peaks that are radially farther than the distance at which the extinction in that sightline reaches 50 mmag (Section 2.5.1). If the distance of the peak of a component is farther than this threshold but part of the component is below that threshold (Figure 2.7), we still consider the part of the component that is below that threshold when integrating over the n_{H} of that component. We find that the peak with the highest dust column density lies within 270 pc from the Sun for most sightlines within our mask.

We calculate Equation 2.10 for each sightline to quantify its dust complexity. For the sightline in Figure 2.7, for example, $\mathcal{N}_c=2.16$. We show a map of \mathcal{N}_c for the sightlines we select in Section 2.5.1 for one of the 12 posterior samples of the Edenhofer et al. (2023) maps at the top of Figure 2.8.²

2.5.3 Nearest-Neighbor Matching

To examine the effect of line-of-sight dust complexity on the dust polarization fraction, we compare sightlines that have the same total column densities but very different 3D dust distributions. We define low- and high-dust complexity sightlines as sightlines with $\mathcal{N}_c \leq 1.1$ and $\mathcal{N}_c \geq 1.5$, respectively. The goal is to compare sightlines with different dust distributions, so the particular \mathcal{N}_c threshold values are less important. We start with these values for our fiducial analysis because they are similar to the thresholds used in Pelgrims et al. (2021a). The Pelgrims et al. (2021a) analysis used a Gaussian decomposition of 3D maps of the neutral hydrogen emission line, where the third dimension is radial velocity. They used $\mathcal{N}_c=1$ and $\mathcal{N}_c \geq 1.5$ in their analysis. However, for the data and mask we use in our analysis, we find that across the 12 posterior sample maps from Edenhofer et al. (2023), only 8-31 sightlines have $\mathcal{N}_c=1$, i.e., a single dust component. Since this does not represent a large enough sample size, we define the low-complexity bin as sightlines with $\mathcal{N}_c \leq 1.1$. This increases the number of sightlines in that bin to around 1,000 per posterior sample.

²We make the \mathcal{N}_c maps for the 12 posterior samples of Edenhofer et al. (2023) publicly available at <https://doi.org/10.7910/DVN/IW09AE> (Halal et al., 2024)

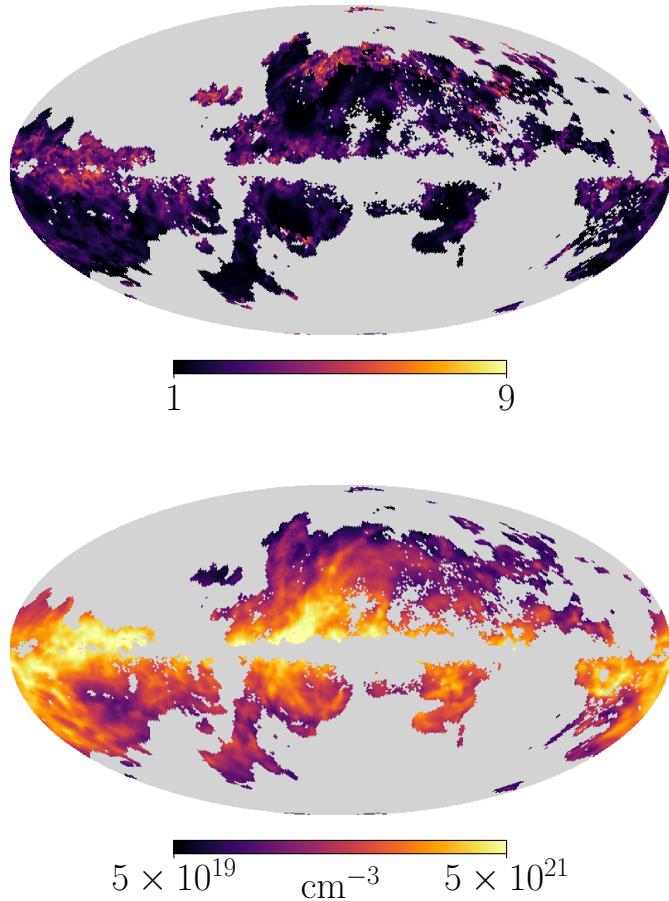


Figure 2.8: *Top panel:* A map of \mathcal{N}_c calculated using the dendrogram algorithm with $\Delta_n = 4.52 \times 10^{-3} \text{ cm}^{-3}$ ($A'_V = 7 \times 10^{-6}$), and $\#\text{voxels} = 3$ (21 pc). *Bottom panel:* A log-scale map of $N_H^{\text{Edenhofer}}$ formed by integrating over the entire dataset (up to 1.25 kpc) with units of cm^{-2} . These maps are shown only for the sightlines selected in Section 2.5.1 and for one of the posterior samples of Edenhofer et al. (2023).

The number of sightlines in the $\mathcal{N}_c \geq 1.5$ bin is around 13,000. However, we also confirm that our results are robust to variations in these thresholds (Section 2.5.6).

Because the \mathcal{N}_c bins are widely separated, small fluctuations in \mathcal{N}_c due to the choices of Δ_n and $\#\text{voxels}$ do not cause sightlines from one bin to shift to the other bin. However, we

also verify that our results are robust to different choices of Δ_n and $\#\text{voxels}$.

We treat the column density as a confounding variable when comparing the distribution of p_{353} for the low- and high-complexity bins. Any difference in the p_{353} distributions for the low- and high-complexity bins could potentially be explained by a difference in the column density integrated over the distance used to calculate \mathcal{N}_c or a difference in the total column density over a sightline. Therefore, both need to be taken into account for a fair comparison of p_{353} between the two complexity bins.

We integrate n_{H} for each sightline up to 1.25 kpc, starting from either the distance at which the extinction reaches 50 mmag or the minimum distance of the first detected component whose peak lies farther than that distance as explained in Section 2.5.2, whichever is closer. We call this $N_{\text{H}}^{\text{Edenhofer}}$ and show a map of it for the sightlines considered in our analysis in the bottom panel of Figure 2.8.

Since $N_{\text{H}}^{\text{Edenhofer}}$ only takes into account dust up to 1.25 kpc, we also consider the total column density over a sightline. We convert A_V^{Planck} (Section 2.3.1) to $N_{\text{H}}^{\text{Planck}}$ following the formalism in Section 2.3.2. In Section 2.5.6, we also experiment with adding the absolute value of the Galactic latitude as an additional confounding variable.

We perform nearest-neighbor matching with no replacement between the sightlines in the low- and high-complexity bins. We pair up each low-complexity sightline to the high-complexity sightline with the closest $N_{\text{H}}^{\text{Edenhofer}}$ and $N_{\text{H}}^{\text{Planck}}$ values based on the Manhattan distance using the ball-tree algorithm. The number of pairs is, therefore, equal to the

number of sightlines in the low-complexity group, which is the smaller group. For each matched pair, we subtract the p_{353} of the sightline with the lower complexity from the p_{353} of the sightline with the higher complexity. We take the average of the differences over the matched pairs, Δp_{353} , to test whether sightlines with higher dust complexity have higher depolarization levels on average than sightlines with lower dust complexity, i.e., $\Delta p_{353} < 0$.

2.5.4 Statistical Tests

We determine the statistical significance of our results through permutation tests. We perform the analysis described in Section 2.5.3 and obtain Δp_{353}^s as the mean over the pairs for each of the 12 posterior samples s . We perform a permutation-based null test in which we randomly choose one sightline to subtract from the other in each pair, rather than always subtracting p_{353} of the low-complexity sightline. We repeat this 10,000 times and obtain a distribution of $\Delta p_{353}^{s, \text{null}}$ for each posterior sample s . We calculate a p-value as the proportion of $\Delta p_{353}^{s, \text{null}}$ that are equal to or more extreme than Δp_{353}^s .

We additionally use an alternative null test. Instead of separating sightlines into high- and low-complexity groups, we randomly select 25% of sightlines (5,050 sightlines) to be in one group and 25% to be in the other group. We then run the same analysis on those 2 groups, pairing them up based on $N_{\text{H}}^{\text{Edenhofer}}$ and $N_{\text{H}}^{\text{Planck}}$ and subtracting p_{353} of one group from that of the other for each pair. To ensure we are not biasing the null test by matching neighboring sightlines which may have similar p_{353} , for each run, we randomly alternate

between selecting the sightlines in one group to be in the Northern Galactic hemisphere and the sightlines in the other group to be in the Southern Galactic hemisphere and the other way around. We find that sightlines in the Northern Galactic hemisphere tend to have higher p_{353} on average than sightlines in the Southern Galactic hemisphere, so $\Delta p_{353}^{s, \text{null}}$ will be biased towards positive or negative values based on which of the two groups is selected from which hemisphere. Using this version of the null test, therefore, takes this bias into account and yields a more conservative estimate of the significance of our hypothesis test. The permutation-based null test, however, guarantees the same number of pairs as the hypothesis test, so we use that version as the main null test when reporting the results. We find that our results are consistent regardless of which null test we use.

We consider results with a two-tailed p-value < 0.001 to be statistically significant. We repeat this analysis for several reasonable values for Δ_n and $\#\text{voxels}$ and for different thresholds of the \mathcal{N}_c bins to ensure that our results are independent of those choices.

In addition to checking the statistical significance of Δp_{353}^s for each of the 12 posterior samples, we also report the mean and standard deviation of Δp_{353}^s across the posterior samples. The error propagated due to the uncertainty on p_{353} is 2 orders of magnitude less than the result. The uncertainty is dominated by the scatter from the 12 posterior samples of the 3D dust map.

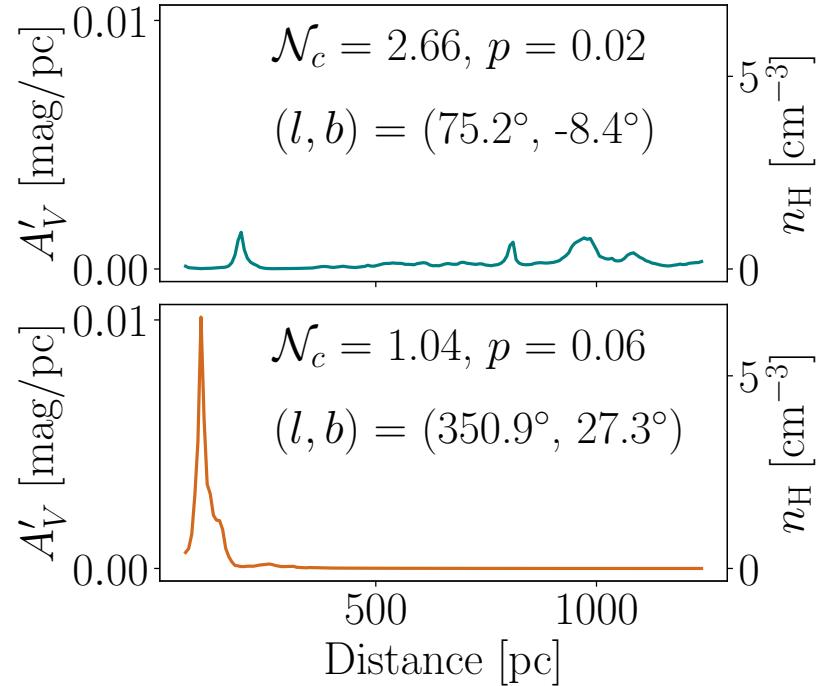


Figure 2.9: The dust distribution in units of differential extinction (mag/pc, left vertical axis) and equivalent Hydrogen number density (cm^{-3} , right vertical axis) along a pair of matched sightlines through the first posterior sample of the Edenhofer et al. (2023) map. These sightlines have the same $N_{\text{H}}^{\text{Edenhofer}}$, but the top one has a higher complexity than the bottom one. The Galactic coordinates, dust complexity, and dust polarization fraction of each of the sightlines are denoted on their subpanels.

2.5.5 Results

We perform the analysis described in the previous subsections on the Edenhofer et al. (2023) maps which extend radially to 1.25 kpc and compare the dust polarization fractions of sightlines with $\mathcal{N}_c \leq 1.1$ and those with $\mathcal{N}_c \geq 1.5$. An example of a pair of matched sightlines with the same $N_{\text{H}}^{\text{Edenhofer}}$ and $N_{\text{H}}^{\text{Planck}}$ is shown in Figure 2.9. Even though the dust polarization fraction of the higher-complexity sightline is lower than that of the lower-complexity sightline in this example, not all sightlines follow this trend. We are only looking

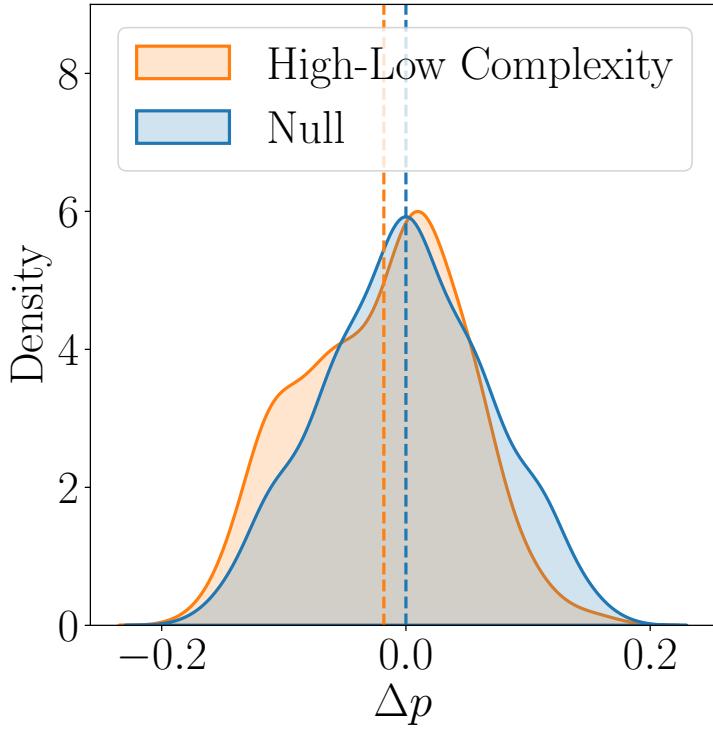


Figure 2.10: Kernel density estimate plots of the Δp_{353} distributions over all matched pairs of sightlines of one posterior sample of the Edenhofer et al. (2023) maps. The p_{353} of the lower complexity sightline is always subtracted from that of the higher complexity one in the orange distribution. This has a mean of $\Delta p_{353} = -1.81 \times 10^{-2}$, which is plotted as an orange dot in Figure 2.11. The blue distribution contains the same pairs as the orange one with the sign randomly flipped for each pair, i.e., the distribution of the permutation-based null test described in Section 2.5.4.

for a statistically significant average effect. An example of the distribution of the differences in p_{353} over all pairs in one of the map posterior samples is shown in Figure 2.10 along with the same distribution for the permutation-based null test.

For this test, Δp_{353}^s , the mean over the paired sightlines for each posterior sample s , is plotted in orange at the top of Figure 2.11. The mean and standard deviation of Δp_{353}^s over the 12 samples are -1.47×10^{-2} and 0.22×10^{-2} , respectively. All 12 posterior samples

pass both null tests described in Section 2.5.4 with a p-value < 0.001 . A random $\Delta p_{353}^{s,\text{null}}$ for each posterior sample is shown in blue at the top of Figure 2.11 as well. Therefore, we find that higher dust complexity at equivalent column densities is associated with depolarization at the 1.5% level. This is at the level of 6.8% of the maximum dust polarization fraction measured by Planck Collaboration et al. (2020a) at 353 GHz and $80'$.

To determine whether this result is uniquely enabled by the Edenhofer et al. (2023) dust maps since they extend radially to 1.25 kpc, we repeat the analysis using the 3D Leike et al. (2020) dust maps, which extend to 370 pc in the positive and negative Galactic-X and Y coordinates and 270 pc in the positive and negative Galactic-Z coordinate. For consistency between different sightlines in those maps, we truncate all sightlines at 270 pc. Also, we start each sightline at 70 pc since Leike et al. (2020) find that the reconstructed dust density closer than 70 pc resembles a smeared-out version of the farther dust, an artifact related to systematic data biases. Leike et al. (2020) also provide 12 posterior samples for their maps which we use for this analysis. Because we sample these data in increments of 2 pc as opposed to 7 pc as in the case of the Edenhofer et al. (2023) data, we set the dendrogram parameter $\#\text{voxels} = 10$, which corresponds to 20 pc, compared to $\#\text{voxels} = 3$, which corresponds to 21 pc in the case of the Edenhofer et al. (2023) data. We keep the dendrogram parameter Δ_n the same for both maps. This results in about 3,500 sightlines with $\mathcal{N}_c \leq 1.1$ and about 8,500 with $\mathcal{N}_c \geq 1.5$.

The results for repeating the analysis using the Leike et al. (2020) maps instead are

shown at the bottom of Figure 2.11. The mean and standard deviation of Δp_{353}^s over the 12 posterior samples are 2.75×10^{-3} and 3.06×10^{-3} , respectively. None of the 12 samples pass either of the null tests described in Section 2.5.4, i.e., they are indistinguishable from the distributions of $\Delta p_{353}^{s, \text{null}}$. We also show a random $\Delta p_{353}^{s, \text{null}}$ for each posterior sample in blue in the same subplot of Figure 2.11.

To determine whether the null result is attributed to using a different dataset or to the lower extent in radial distance, we run the analysis on the Edenhofer et al. (2023) maps up to 270 pc, the same distance used for the Leike et al. (2020) maps. The number of sightlines with $\mathcal{N}_c \leq 1.1$ and $\mathcal{N}_c \geq 1.5$ after this distance cut are about 10,000 and 3,500, respectively. The difference in the \mathcal{N}_c distributions between this data and the Leike et al. (2020) maps is due to several differences in the maps and the post-processing we perform on them, including having to use slightly different dendrogram parameters and not counting components whose peak is closer than the distance where the differential extinction integrated radially outwards reaches 50 mmag in the Edenhofer et al. (2023) maps among other differences. We plot the results in the middle of Figure 2.11. The mean and standard deviation of Δp_{353}^s over the 12 posterior samples are 1.05×10^{-3} and 1.36×10^{-3} , respectively. Again, none of the 12 samples pass either of the null tests described in Section 2.5.4, i.e., they are indistinguishable from the distributions of $\Delta p_{353}^{s, \text{null}}$. The consistency of this result with null as well indicates that the null result we found using the Leike et al. (2020) maps is due to only considering distances up to 270 pc, not the choice of 3D dust dataset. This illustrates the role that dust

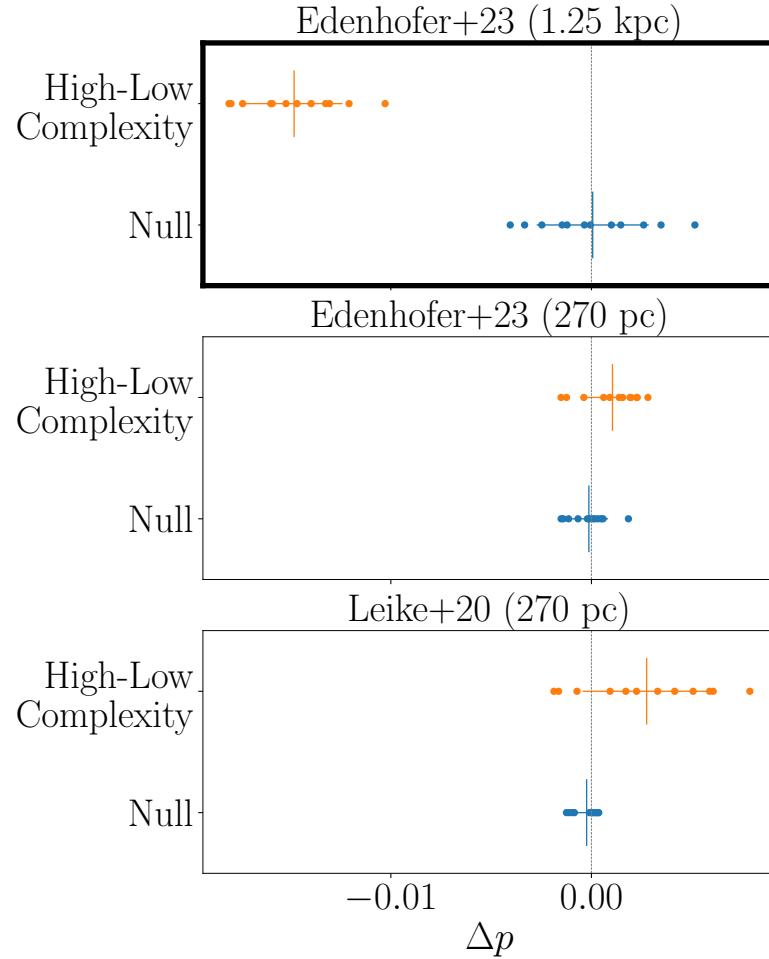


Figure 2.11: The mean difference in p_{353} over the paired sightlines for each of the 12 posterior samples, where for each pair, p_{353} for the lower-complexity sightline is subtracted from p_{353} for the higher-complexity sightline. \mathcal{N}_c is calculated on the Edenhofer et al. (2023) maps up to 1.25 kpc (top), the Edenhofer et al. (2023) maps up to 270 pc (middle), and the Leike et al. (2020) maps up to 270 pc (bottom). For each panel, the mean differences for the actual test are plotted in orange, and samples from the permutation-based null tests are plotted in blue. The mean and standard deviation of the 12 means for each test are also plotted. The top panel agrees with our hypothesis: that sightlines with similar column densities will, on average, exhibit lower dust polarization fractions when their 3D dust distribution is more complex.

components farther than 270 pc play in affecting polarization measurements.

2.5.6 Validation

In this subsection, we summarize some of the tests we performed to verify our results. As described in Section 2.5.2, we find that our results are robust to reasonable variations in Δ_n and $\#\text{voxels}$. These are summarized in Table 2.1 in Appendix 2.7. The results are also consistent when using the Planck R3.01 353 GHz maps instead of the GNILC maps.

To ensure that the pairs were matched correctly, we examine the distributions of the differences in $N_{\text{H}}^{\text{Edenhofer}}$ and $N_{\text{H}}^{\text{Planck}}$ between the matched pairs. We verify that those differences peak near zero and are not skewed toward the positive or negative values. We find that to be the case for both variables and over all 12 posterior samples of Edenhofer et al. (2023). We also find our results to be robust when including $|b|$ as an additional confounding variable to $N_{\text{H}}^{\text{Edenhofer}}$ and $N_{\text{H}}^{\text{Planck}}$ in the pair matching – in other words, the result is not attributable to a dependence of the dust polarization fraction on Galactic latitude. We show these results in Table 2.2 in Appendix 2.7.

We also examine the angular separations between the sightline pairs and the discrepancies in their overall path lengths. The distributions of angular distances are consistent with a random distribution of angular distances for all 12 posterior samples of Edenhofer et al. (2023). We also find no significant differences in the total path lengths between the paired sightlines across all 12 samples.

We experiment with varying the \mathcal{N}_c thresholds that define the bin edges of the low- and high-complexity sightlines. The 25th and 75th percentiles of the \mathcal{N}_c distribution over

our mask vary slightly over the 12 posterior samples of Edenhofer et al. (2023) and over variations in Δ_n and $\#\text{voxels}$. However, they are roughly 1.4 and 2.3, respectively. Therefore, we split our sightlines based on these values into low- ($\mathcal{N}_c \leq 1.4$) and high-complexity ($\mathcal{N}_c \geq 2.3$) bins to achieve a roughly equal number of sightlines in each group. This should improve the pair-matching outcomes since there are more sightlines to match from in the smaller group. We find that our results are robust to this change as shown in Table 2.1 in Appendix 2.7.

Since neighboring sightlines are likely to have similar \mathcal{N}_c and similar p_{353} values, we test whether large regions of neighboring sightlines belonging to either the low- or high-complexity groups bias our results. We randomly sample 1,000 sightlines out of about 5,000 sightlines from each of the low- and high-complexity groups before pair matching, where we use $\mathcal{N}_c \leq 1.4$ and $\mathcal{N}_c \geq 2.3$ for these groups, respectively, in this case. We find that our results are robust to this test as shown in Table 2.3 in Appendix 2.7.

Finally, we modify the definition of \mathcal{N}_c from Equation 2.10 to

$$\mathcal{N}_c = \frac{N_{\text{H}}^{\text{Edenhofer}}}{N_{\text{H}}^{\text{max}}}, \quad (2.11)$$

i.e., we use the total column density out to 1.25 kpc rather than a sum over the dust components in the numerator. With the new definition, \mathcal{N}_c is only sensitive to the dendrogram-identified component with the highest column density $N_{\text{H}}^{\text{max}}$ rather than also being sensitive

to the other dendrogram-identified components. Since the dendrogram parameters define the minimum prominence for a peak to be considered an independent component, we expect $N_{\text{H}}^{\text{max}}$ to be the least sensitive component to those parameters. Therefore, the new definition is much less sensitive to the dendrogram parameters. The 25th and 75th percentiles of this modified version of \mathcal{N}_c are about 0.05 and 0.09 for the 12 posterior samples, and we use these as the upper and lower thresholds for the low- and high-complexity bins, respectively. Over the 12 posterior samples, we find a mean and standard deviation of $\Delta p_{353} = 0.96 \times 10^{-2} \pm 0.34 \times 10^{-2}$. This passes the null test with a p-value < 0.001 .

2.6 Discussion and Conclusions

In this paper, we explore how different geometrical factors affect the fractional polarization of the dust emission. In Section 2.4, we test whether we detect an imprint of the Local Bubble geometry on the dust polarization fraction. Following the well-motivated assumptions that the magnetic field lines are tangential to the Local Bubble surface and that this surface is defined by the model of Pelgrims et al. (2020), we test for a correlation between the measured dust polarization fraction and the angle the line of sight makes with the tangential magnetic field lines. However, we do not find evidence for this in sightlines where the dust extinction is dominated by the Local Bubble. We also do not find a correlation between the dust polarization fraction and the inclination angle between the Local Bubble wall defined by the model of O'Neill et al. (2024) and the plane of the sky. Therefore, we conclude

that at least one of the commonly made assumptions must not hold. We hypothesize that dust structure beyond the Local Bubble wall plays a substantial role in determining the polarization structure of the dust emission. Our results show that simply projecting the Planck polarization data onto the Local Bubble geometry is not a well-motivated model for the magnetic field structure of the Local Bubble.

In Section 2.5, we test how dust complexity, i.e., how the 3D dust is distributed along the line of sight, affects the dust polarization fraction. We quantify the dust complexity for each sightline and group the sightlines into low- and high-complexity groups. We pair-match the sightlines across the two groups based on their column densities. For each pair, we subtract the dust polarization fraction of the sightline with low complexity from that of the sightline with high complexity. We find that on average, the dust polarization fraction of the sightlines with higher complexity is 2% lower than those with lower complexity. This is only true when considering dust out to 1.25 kpc. The result is not statistically different from null when considering dust out to 270 pc only. Note that our definition of complexity does not take into account the distance to different dust components. Future work could incorporate the effect of this distance into the analysis.

We test whether the order of magnitude of this result agrees with our expectation based on geometric depolarization. Padoan et al. (2001) model the polarized thermal dust emission from protostellar cores formed through supersonic turbulent flows within molecular clouds following the formalism in Fiege & Pudritz (2000). Fiege & Pudritz (2000)

develops this formalism to model the submillimeter polarization patterns for filamentary molecular clouds. Fiege & Pudritz (2000) and Padoan et al. (2001) ignore the effects of self-absorption and scattering since this model is for submillimeter wavelengths at which the diffuse interstellar medium is optically thin. Padoan et al. (2001) further assume that the dust grain properties are constant and the temperature is uniform. Since these assumptions are valid for our order of magnitude estimation, we follow the same formalism here.

For this test, we consider sightlines with 2 clouds but different \mathcal{N}_c values. For simplicity, we assume each cloud to have a constant volume density, plane-of-sky magnetic field angle, and magnetic inclination angle along a certain sightline. Therefore, for a given sightline, we write Equations 2.2, 2.3, and 2.4 as

$$I = \epsilon N_H - \frac{\alpha\epsilon}{2} \left(N_{H,a} \cos^2 \gamma_a + N_{H,b} \cos^2 \gamma_b - \frac{2}{3} N_H \right), \quad (2.12)$$

$$Q = -\alpha\epsilon(N_{H,a} \cos 2\psi_a \cos^2 \gamma_a + N_{H,b} \cos 2\psi_b \cos^2 \gamma_b), \quad (2.13)$$

$$U = -\alpha\epsilon(N_{H,a} \sin 2\psi_a \cos^2 \gamma_a + N_{H,b} \sin 2\psi_b \cos^2 \gamma_b), \quad (2.14)$$

where $N_{H,a}$ and $N_{H,b}$ are the column densities for each cloud, $N_H = N_{H,a} + N_{H,b}$, γ_a and γ_b are the magnetic inclination angles of the two clouds, and ϵ cancels out when calculating p .

We take $N_{H,a} > N_{H,b}$, i.e.,

$$\mathcal{N}_c = \frac{N_{H,a} + N_{H,b}}{N_{H,a}}. \quad (2.15)$$

We divide Equations 2.12, 2.13, and 2.14 by $N_{\text{H,a}}$ to be able to write them in terms of \mathcal{N}_c .

Therefore,

$$I/N_{\text{H,a}} = \epsilon \mathcal{N}_c - \frac{\alpha \epsilon}{2} \left[\cos^2 \gamma_a + (\mathcal{N}_c - 1) \cos^2 \gamma_b - \frac{2}{3} \mathcal{N}_c \right], \quad (2.16)$$

$$Q/N_{\text{H,a}} = -\alpha \epsilon \left[\cos 2\psi_a \cos^2 \gamma_a + (\mathcal{N}_c - 1) \cos 2\psi_b \cos^2 \gamma_b \right], \quad (2.17)$$

$$U/N_{\text{H,a}} = -\alpha \epsilon \left[\sin 2\psi_a \cos^2 \gamma_a + (\mathcal{N}_c - 1) \sin 2\psi_b \cos^2 \gamma_b \right], \quad (2.18)$$

where $N_{\text{H,a}}$ cancels out when calculating p .

We uniformly sample orientations in the range $[0, \pi]$ for ψ_a and ψ_b , values in the range $[0, 1]$ for $\cos \gamma_a$ and $\cos \gamma_b$, and a value in the range $[1, 1.1]$ for \mathcal{N}_c to calculate an instance of $p(\mathcal{N}_c \leq 1.1)$. We also sample different orientations and a value in the range $[1.5, 2]$ for \mathcal{N}_c to calculate an instance of $p(\mathcal{N}_c \geq 1.5)$. We then subtract $p(\mathcal{N}_c \leq 1.1)$ from $p(\mathcal{N}_c \geq 1.5)$ as in Section 2.5. We repeat this 10,000 times and average the results. We set $\alpha = 0.22$, which corresponds to a maximum dust polarization fraction across the sky $p_{\text{max}} = 0.22$. This is the value Planck Collaboration et al. (2020a) observe for p_{max} at 353 GHz and 80' resolution. We get

$$\langle p(\mathcal{N}_c \geq 1.5) - p(\mathcal{N}_c \leq 1.1) \rangle = -0.013. \quad (2.19)$$

This mean difference depends on the value for α . For instance, if we set $\alpha = 0.15$ as in Padoan et al. (2001) instead, we get a mean difference of -0.009. However, our estimate

for the mean difference agrees with the result we measure in Section 2.5 for all reasonable values of α . Thus, our empirical result is consistent with our theoretical estimate for the dust depolarization attributable to the line-of-sight dust complexity.

Variations in the orientation of magnetic fields along the line of sight induce differences in the polarization angles of different dust components along the same sightline (Lee & Draine, 1985; Tassis & Pavlidou, 2015; King et al., 2018). When the emission of those components has different spectral energy distributions (SEDs), a frequency-dependent variation of the observed dust polarization angle along that sightline emerges, a phenomenon known as line-of-sight frequency decorrelation. This decorrelation complicates the translation of polarized dust emission maps from one frequency to another. Current analysis within the BICEP/Keck field does not demonstrate evidence of this phenomenon (BICEP/Keck Collaboration et al., 2021, 2023a). However, a statistically significant detection of line-of-sight frequency decorrelation has been identified in larger sky areas across sightlines intersecting multiple dust clouds with varying magnetic field orientations (Pelgrims et al., 2021a). Since polarized dust emission is the major foreground for CMB polarization measurements at high frequencies, it is important to characterize how the spatial complexity of the magnetic field in the dust might affect the frequency dependence of the foreground signal. In this paper, we have presented evidence that the 3D spatial complexity of the dust affects the dust polarization signal, even at a fixed frequency.

The analysis in this paper highlights the importance of 3D dust mapping out to large

distances. The dust distribution affects the dust polarization fraction, which has implications for the 3D magnetic field distribution. Since we expect higher complexity sightlines to have a lower dust polarization fraction on average, we would infer that the magnetic field is more uniform along a sightline if it has both a highly complex dust distribution and a large dust polarization fraction. This analysis was performed on the sightlines shown in Figure 2.6. Improvements to 3D dust modeling will allow us to look for this effect at the very high Galactic latitudes excluded here.

Given the significance of the 3D dust distribution on measurements of the polarized dust emission, these data can be combined with position-position-velocity maps of the neutral hydrogen-based dust polarization templates. These templates, constructed based on the orientation of neutral hydrogen filaments, have been shown to correlate very well with the measured dust polarization (Clark & Hensley, 2019; Cukierman et al., 2023; BICEP/Keck Collaboration et al., 2023a; Halal et al., 2024a). Since the neutral hydrogen and dust trace similar volumes of the diffuse interstellar medium (Boulanger et al., 1996a; Lenz et al., 2017), future work could morphologically match the position-position-position dust maps with the position-position-velocity neutral hydrogen-based maps to form 4D position-position-position-velocity maps of the magnetic field and polarized dust emission. Starlight polarization can also be used to provide a tomographic view of the plane-of-the-sky magnetic field and polarized dust emission for sightlines with these measurements (Panopoulou et al., 2019; Tassis et al., 2018; Pelgrims et al., 2023, 2024). These data

combined with Faraday tomography (Van Eck et al., 2017) or rotation measures (Tahani et al., 2018) can be used to constrain the 3D magnetic field structure.

2.7 Analysis Variations

In this appendix, we list some of the results obtained from varying the main analysis choices in Section 2.5. These analysis variations are described and their results summarized in Section 2.5.6. In Table 2.1, we present some of the results obtained from varying the dendrogram parameters used in identifying the density peaks in the line-of-sight dust distribution and from varying the upper and lower thresholds in N_c used for splitting the sightlines into low- and high-complexity groups. Table 2.2 lists the results of the same variations performed in Table 2.1 but when including the absolute value of the Galactic latitude as an additional confounding variable in the sightline matching across the two complexity groups. Finally, we present the results of sampling 1,000 sightlines from each of the two complexity groups with $N_c \leq 1.4$ and $N_c \leq 2.3$ before matching in Table 2.3. We find that the result in our main analysis is robust to all of these analysis variations.

Δ_n	#voxels	low complexity	high complexity	#matchedpairs	$\mu_{\Delta p_{353}}$	$\sigma_{\Delta p_{353}}$	p-value
7×10^{-6}	3	$\mathcal{N}_c \leq 1.1$	$\mathcal{N}_c \geq 1.5$	753	-1.47×10^{-2}	0.22×10^{-2}	< 0.001
7×10^{-6}	3	$\mathcal{N}_c \leq 1.4$	$\mathcal{N}_c \geq 2.3$	4,674	-1.10×10^{-2}	0.08×10^{-2}	< 0.001
9×10^{-6}	3	$\mathcal{N}_c \leq 1.1$	$\mathcal{N}_c \geq 1.5$	776	-1.33×10^{-2}	0.26×10^{-2}	< 0.001
9×10^{-6}	3	$\mathcal{N}_c \leq 1.4$	$\mathcal{N}_c \geq 2.3$	4,719	-1.01×10^{-2}	0.12×10^{-2}	< 0.001
5×10^{-6}	5	$\mathcal{N}_c \leq 1.1$	$\mathcal{N}_c \geq 1.5$	1,005	-0.98×10^{-2}	0.20×10^{-2}	< 0.001
5×10^{-6}	5	$\mathcal{N}_c \leq 1.4$	$\mathcal{N}_c \geq 2.3$	5,311	-0.82×10^{-2}	0.14×10^{-2}	< 0.001
3×10^{-6}	5	$\mathcal{N}_c \leq 1.1$	$\mathcal{N}_c \geq 1.5$	984	-1.00×10^{-2}	0.21×10^{-2}	< 0.001
3×10^{-6}	5	$\mathcal{N}_c \leq 1.4$	$\mathcal{N}_c \geq 2.3$	5,284	-0.80×10^{-2}	0.13×10^{-2}	< 0.001

Table 2.1: Results obtained from varying both the dendrogram parameters (Δ_n and #voxels) used to identify the density peaks in the line-of-sight dust distribution and from varying the upper and lower thresholds in \mathcal{N}_c used to divide the sightlines into low- and high-complexity groups. The results of the main analysis are bolded in the first row. The mean and standard deviation in the $\mu_{\Delta p_{353}}$ and $\sigma_{\Delta p_{353}}$ columns are over the 12 posterior samples of the Edenhofer et al. (2023) maps.

Δ_n	#voxels	low complexity	high complexity	#matchedpairs	$\mu_{\Delta p_{353}}$	$\sigma_{\Delta p_{353}}$	p-value
7×10^{-6}	3	$\mathcal{N}_c \leq 1.1$	$\mathcal{N}_c \geq 1.5$	753	-1.45×10^{-2}	0.23×10^{-2}	< 0.001
7×10^{-6}	3	$\mathcal{N}_c \leq 1.4$	$\mathcal{N}_c \geq 2.3$	4,674	-1.02×10^{-2}	0.08×10^{-2}	< 0.001
9×10^{-6}	3	$\mathcal{N}_c \leq 1.1$	$\mathcal{N}_c \geq 1.5$	776	-1.31×10^{-2}	0.34×10^{-2}	< 0.001
9×10^{-6}	3	$\mathcal{N}_c \leq 1.4$	$\mathcal{N}_c \geq 2.3$	4,719	-1.00×10^{-2}	0.15×10^{-2}	< 0.001
5×10^{-6}	5	$\mathcal{N}_c \leq 1.1$	$\mathcal{N}_c \geq 1.5$	1,005	-0.60×10^{-2}	0.17×10^{-2}	< 0.001
5×10^{-6}	5	$\mathcal{N}_c \leq 1.4$	$\mathcal{N}_c \geq 2.3$	5,311	-0.82×10^{-2}	0.13×10^{-2}	< 0.001
3×10^{-6}	5	$\mathcal{N}_c \leq 1.1$	$\mathcal{N}_c \geq 1.5$	984	-0.67×10^{-2}	0.22×10^{-2}	< 0.001
3×10^{-6}	5	$\mathcal{N}_c \leq 1.4$	$\mathcal{N}_c \geq 2.3$	5,284	-0.80×10^{-2}	0.13×10^{-2}	< 0.001

Table 2.2: Results obtained from the same variations described in Table 2.1, where the sightline matching over the two complexity groups in this case includes the absolute value of the Galactic latitude as an additional confounding variable. The mean and standard deviation in the $\mu_{\Delta p_{353}}$ and $\sigma_{\Delta p_{353}}$ columns are over the 12 posterior samples of the Edenhofer et al. (2023) maps.

Δ_n	#voxels	$\mu_{\Delta p_{353}}$	$\sigma_{\Delta p_{353}}$	p-value
7×10^{-6}	3	-1.08×10^{-2}	0.17×10^{-2}	< 0.001
9×10^{-6}	3	-1.14×10^{-2}	0.24×10^{-2}	< 0.001
5×10^{-6}	5	-0.84×10^{-2}	0.17×10^{-2}	< 0.001
3×10^{-6}	5	-0.89×10^{-2}	0.18×10^{-2}	< 0.001

Table 2.3: Results obtained from sampling 1,000 sightlines from each of the $\mathcal{N}_c \leq 1.4$ and $\mathcal{N}_c \geq 2.3$ complexity groups before matching. The different rows are for different variations in the dendrogram parameters (Δ_n and #voxels) used in identifying the density peaks in the line-of-sight dust distribution. The mean and standard deviation in the $\mu_{\Delta p_{353}}$ and $\sigma_{\Delta p_{353}}$ columns are over the 12 posterior samples of the Edenhofer et al. (2023) maps.

Chapter 3

Characterizing Dust Polarization

Through Correlations with H I

Abstract

We characterize Galactic dust filaments by correlating BICEP/Keck and Planck data with polarization templates based on neutral hydrogen (H I) observations. Dust polarization is important for both our understanding of astrophysical processes in the interstellar medium (ISM) and the search for primordial gravitational waves in the cosmic microwave background (CMB). In the diffuse ISM, H I is strongly correlated with the dust and partly organized into filaments that are aligned with the local magnetic field. We analyze the deep BICEP/Keck data at 95, 150, and 220 GHz, over the low-column-density region of sky where BICEP/Keck has set the best limits on primordial gravitational waves. We separate the H I emission into distinct velocity components

and detect dust polarization correlated with the local Galactic H I but not with the H I associated with Magellanic Stream 1. We present a robust, multifrequency detection of polarized dust emission correlated with the filamentary H I morphology template down to 95 GHz. For assessing its utility for foreground cleaning, we report that the H I morphology template correlates in *B* modes at a \sim 10-65% level over the multipole range $20 < \ell < 200$ with the BICEP/Keck maps, which contain contributions from dust, CMB, and noise components. We measure the spectral index of the filamentary dust component spectral energy distribution to be $\beta = 1.54 \pm 0.13$. We find no evidence for decorrelation in this region between the filaments and the rest of the dust field or from the inclusion of dust associated with the intermediate velocity H I. Finally, we explore the morphological parameter space in the H I-based filamentary model.

3.1 Paper Status and External Contributions

This chapter is based on the article published by The Astrophysical Journal, Volume 945, Issue 1 under the title, "BICEP/Keck. XVI. Characterizing Dust Polarization through Correlations with Neutral Hydrogen." I am the corresponding author of this paper. I performed all of the analysis, wrote all of the text, and produced all of the figures for this paper. However, this work is the result of weekly discussions and advising from postdoctoral scholars Ari Cukierman and Dominic Beck and my co-advisors Susan Clark and Chao-Lin Kuo. I have also received extensive editorial input from Susan Clark, Ari Cukierman,

Dominic Beck, and other members of the BICEP/Keck Collaboration.

As part of the BICEP/Keck Collaboration, I additionally helped in building a forebaffle and in data reduction and operations of the BICEP3 and BICEP Array instruments. I also provide feedback on collaborators' projects and papers and am a co-author on collaboration papers (BICEP/Keck Collaboration et al., 2020, 2021, 2022a,b,c,d,e,f, 2023b,c,d, 2024a,b).

3.2 Introduction

An accurate characterization of polarized dust emission is important for understanding different astrophysical phenomena in the interstellar medium (ISM) and studying the polarization of the cosmic microwave background (CMB). The short axes of aspherical rotating dust grains are preferentially aligned with the local magnetic field. This causes their thermal emission to be linearly polarized (Purcell, 1975). Polarized dust emission is the dominant polarized CMB foreground at frequencies greater than approximately 70 GHz and at large scales (Planck Collaboration et al., 2016d). Characterizing and removing the dust contribution to CMB polarization measurements allows us to look for an excess signal generated by primordial gravitational waves, parameterized by the tensor-to-scalar ratio r , in order to constrain inflationary theories (Kamionkowski et al., 1997; Seljak & Zaldarriaga, 1997; Seljak, 1997).

Galactic neutral hydrogen (H I) gas has several advantages for tracing properties of the dust polarization. H I is strongly correlated with dust throughout the diffuse ISM (Boulanger

et al., 1996a; Lenz et al., 2017). The dust and H I are organized into filamentary structures (Clark et al., 2015; Planck Collaboration et al., 2014b). H I filaments are well aligned with the plane-of-sky magnetic field orientation (Clark et al., 2014, 2015). Moreover, since the H I measurements are spectroscopic, they provide 3D (position, position, and velocity) information about the H I emission, where velocity is inferred from the Doppler-shifted frequency of the 21-cm line. They are also independent from the broadband thermal dust millimeter-wave and far infrared emission observations, and therefore, do not contain correlated systematics. Finally, H I measurements are not contaminated by the cosmic infrared background (CIB; Chiang & Ménard, 2019). These advantages allow us to exploit cross correlations between the data collected by CMB experiments and H I surveys to better understand and characterize diffuse dust polarization. Clark & Hensley (2019) developed a formalism for modeling the linear polarization structure of Galactic dust emission solely from H I intensity measurements. They have shown that these H I morphology templates correlate at the $\sim 60\%$ ($\sim 50\%$) level in *E* modes (*B* modes) with Planck data at 353 GHz at multipole $\ell = 50$ over the high-Galactic latitude sky, and the correlation decays roughly monotonically to zero at around multipole moment $\ell \approx 1000$.

The BICEP2 and Keck Array CMB experiments target a $\sim 400 \text{ deg}^2$ patch of high-Galactic latitude sky (BICEP/Keck Collaboration et al., 2021, hereafter BK18). The instantaneous field of view of BICEP3 is larger and targets a $\sim 600 \text{ deg}^2$ patch, which encompasses that of BICEP2 and Keck Array (BICEP/Keck Collaboration et al., 2022f). These patches

were chosen to have relatively little dust emission in intensity (Finkbeiner et al., 1999). In this paper, we use BICEP/Keck maps using all data taken up to and including the 2018 observing season, the data set known as "BK18." These instruments have $\sim 30\%$ fractional bandwidths and have achieved great depths at different frequencies. The polarization maps at 95, 150, and 220 GHz reach depths of 2.8, 2.8, and $8.8 \mu\text{K}_{\text{CMB}}$ arcmin respectively (BICEP/Keck Collaboration et al., 2021). The signal-to-noise on polarized dust emission of the 220 GHz maps exceeds that of Planck at 353 GHz in the BICEP/Keck region (BICEP/Keck Collaboration et al., 2021). These data thus present an excellent opportunity to study the structure of the diffuse, magnetic ISM. Furthermore, this well-characterized region of sky will also be observed by future CMB experiments like CMB-S4 (CMB-S4 collaboration et al., 2022). In this paper, we make use of cross correlations of BK18 data with H I morphology maps. Because the H I morphology templates are defined solely from the morphology of linear H I structures, we refer to the component of the real dust field that is correlated with these templates as filamentary.

A motivation for using H I to study dust in the BICEP/Keck region is its promise as a tracer of the 3D structure of the magnetic ISM (Clark, 2018; Clark & Hensley, 2019). A differently oriented magnetic field along the line of sight will give rise to different dust polarization angles along that line of sight (Tassis & Pavlidou, 2015). If this dust is described by different spectral energy distributions (SEDs) in different locations along that sightline, the measured dust polarization angle will be frequency-dependent. This is referred to

as line-of-sight frequency decorrelation. Frequency decorrelation can also arise due to spatial variations of the dust SED in the plane of the sky, producing frequency-dependent variations in the dust polarization pattern. Decorrelation causes maps of dust emission at different frequencies to differ by more than just a multiplicative factor, complicating the ability to use dust maps at one frequency to constrain the dust emission at another frequency. The decorrelation parameter, Δ_d , defined as the ratio of the cross-spectrum between maps at 217 and 353 GHz to the geometric mean of the corresponding autospectra, is currently constrained to $\Delta_d > 0.98$ (68% C.L.) in the BICEP/Keck region (BICEP/Keck Collaboration et al., 2021). Therefore, we currently have no indication of dust decorrelation in this region. However, there is evidence for frequency decorrelation in data, either associated with superpositions of independent line-of-sight emission (Pelgrims et al., 2021a) or, at large scales, with spatial variations in the dust-polarization SED (Ritacco et al., 2022). Pelgrims et al. (2021a) measure evidence for line-of-sight frequency decorrelation. They make a statistically significant detection of a stronger frequency-dependent change of the polarization angle along lines of sight which intercept multiple dust clouds with different magnetic field orientations. Therefore, it is interesting to isolate and separately characterize the distinct H_I velocity components along the line of sight in the region observed by BICEP2, BICEP3, and the Keck Array instruments to look for evidence for this effect. Additionally, we look for evidence of decorrelation due to any variation in the polarized dust SED between dust filaments, identified by the H_I morphology model and generally

associated with the cold neutral medium (Clark et al., 2019; Kalberla et al., 2020), and the rest of the dust column.

In this paper, we perform cross correlations between the Stokes parameter maps of the H I morphology template and BICEP/Keck and Planck data and measure the statistical significance of the correlation as a function of frequency, instrument, and H I velocity component in the BICEP/Keck region. To clarify, the H I-based Stokes parameter maps are based on H I morphology and not on H I polarization. The cross correlations allow us to pick out the filamentary dust signal from the overall dust signal measured by BICEP/Keck and Planck in that region. We use our formalism to compare the sensitivities of Planck and BICEP/Keck in that region, to tune the H I morphology template, and to search for frequency decorrelation. We also measure the SED of the dust correlated with H I filaments. Knowledge of the dust SED is essential for CMB studies (Chluba et al., 2017; Hensley & Bull, 2018) and for providing constraints for physical models of dust composition (e.g. Hensley et al., 2022).

This paper is organized as follows. We introduce the data used in this work in Section 3.3. In Section 3.4, we introduce the methodology to estimate the statistical significance of the detection and to measure the filamentary dust SED. In Section 3.5, we present a method for separating the different velocity components in the BICEP/Keck regions using H I velocity information. Our results are presented and discussed in Section 3.6. We then conclude with a summary and outlook in Section 3.7.

3.3 Data

3.3.1 Millimeter-wave Polarization

In this paper, we use BICEP3 data at 95 GHz from 2016 to 2018, BICEP2 data at 150 GHz from 2010 to 2012, and Keck Array data at 150 and 220 GHz from 2012 to 2018 (BICEP/Keck Collaboration et al., 2021). We also use the Planck NPIPE processed maps at 143, 217, and 353 GHz (Planck Collaboration et al., 2020e). These are a subset of the maps we used in BK18 to set the most stringent upper limits on the tensor-to-scalar ratio, r . We do not consider the lower-frequency maps from CMB experiments, i.e. the 23 and 33 GHz bands of Wilkinson Microwave Anisotropy Probe (WMAP) and the 30 and 44 GHz bands of Planck, since we expect a negligible emission contribution from dust in those channels.

In Section 3.4.2, we use the Planck 70% sky fraction Galactic plane mask¹ (Planck Collaboration et al., 2015) for calculating a transfer function for the H I morphology template.

3.3.2 Neutral Hydrogen Emission

The HI4PI spectroscopic survey is the highest-resolution full-sky H I survey to date (HI4PI Collaboration et al., 2016). It has an angular resolution of $16'.2$, a spectral resolution of 1.49 km s^{-1} , and a velocity-bin separation of 1.29 km s^{-1} , achieved by merging data from the Effelsberg-Bonn H I Survey (EBHIS; Winkel et al., 2016) and the Parkes Galactic All-Sky Survey (GASS; McClure-Griffiths et al., 2009). We start out with the velocity channels in

¹Available for download at <http://pla.esac.esa.int> (HFI_Mask_GalPlane-apo0_2048_R2.00.fits)

the range $-120 \text{ km s}^{-1} < v_{\text{lsr}} < 230 \text{ km s}^{-1}$, because the HI4PI maps are noise dominated in the BICEP/Keck region outside that range. We use these data to form HI morphology templates as described in Section 3.4.1.

3.4 Methodology

3.4.1 Convolutional Rolling Hough Transform

Clark & Hensley (2019) used the Rolling Hough Transform (RHT; Clark et al., 2014, 2020) on the HI4PI data to construct 3D (position, position, and velocity) Stokes parameter maps. The mapping defined from HI emission to properties of the dust polarization is based on several observational facts, including that the HI column density correlates well with dust in the diffuse ISM (Boulanger et al., 1996a; Lenz et al., 2017). Also, HI gas contains substantial linear structures that are preferentially aligned with the plane-of-sky component of the local magnetic field (Clark et al., 2015). Therefore, the dust polarization angle is taken to be orthogonal to these filaments. Clark & Hensley (2019) have shown that these maps, integrated over the velocity dimension (Clark, 2018), are highly correlated with the Planck maps of the polarized dust emission at 353 GHz.

While recent work over large regions of high-Galactic latitude sky (not focused on the BICEP/Keck region) has shown that there may be a small aggregate misalignment between the filaments and the Planck-measured magnetic field orientation (Huffenberger et al., 2020;

Clark et al., 2021), the misalignment angle is only $\sim 2^\circ - 5^\circ$ and incorporating it increases the correlation by only an additive $\sim 0.1\%-0.5\%$ (Cukierman et al., 2023).

The first step of the RHT algorithm involves subtracting a smoothed version of the map from the original unsmoothed map. This is known as an unsharp mask and is used to remove the diffuse, large-scale H I emission. This introduces a free parameter that sets the scale of the Gaussian smoothing filter. We refer to this parameter as the smoothing radius (θ_{FWHM}). The second step is to quantize the pixels into a bit mask, where the pixels are turned into zeros and ones based on their sign in the unsharp-masked data. The third step is to apply the Hough transform (Hough, 1962) on a circular window of a given diameter centered on each pixel. The window diameter (D_W) is the second parameter of this algorithm. The fourth step is to retain only values above a certain threshold fraction of the window diameter, where the threshold fraction (Z) is the third and last parameter. Refer to Clark et al. (2014) for further details.

The RHT quantifies the intensity of linear structures as a function of orientation (Clark et al., 2014). Following Clark & Hensley (2019), we use the RHT output to construct Stokes Q and U polarization maps, weighted by the H I intensity. Together, the RHT parameters (θ_{FWHM} , D_W , Z) determine what H I filament morphologies most influence the H I morphology template. It is thus of interest to explore the RHT parameter space and cross correlate different H I morphology templates with the real dust polarization measurements, in order to determine what H I morphologies are most predictive of the true polarized dust

emission. Exploring the parameter space of the original RHT implementation was found to be computationally expensive, limited by the application of the Hough transform to each circular window of data. Other applications have used a convolutional implementation of the Hough transform (e.g., Kerbyson & Atherton, 1995). By rewriting the Hough transform step of the RHT as a series of convolutions, one for each orientation bin, we achieved a $\sim 35\times$ speedup in the RHT algorithm runtime. This convolutional implementation is made public via the RHT GitHub repository (Clark et al., 2020). In this work, we apply the convolutional RHT to the HI4PI data in the BICEP/Keck region to construct a 3D HI morphology template.

3.4.2 RHT Transfer Function

The HI morphology templates have different mode structures than the dust maps. As described in Section 3.4.1, one of the first steps of the RHT algorithm is an unsharp mask. This filter emphasizes small-scale features. For instance, the *E*- and *B*-mode autospectra of the templates constructed with the same RHT parameters as those used in Clark & Hensley (2019) peak in the multipole range $300 < \ell < 500$ and $150 < \ell < 350$, respectively. We denote these spectra by $D_\ell^{\text{HI}\times\text{HI}} = \ell(\ell+1)C_\ell^{\text{HI}\times\text{HI}}/(2\pi)$, where $C_\ell^{m_1\times m_2}$ is the cross spectrum bandpower between two maps, m_1 and m_2 , in the multipole bin ℓ . Correlation ratios are insensitive to this mode structure because the relative weightings of different multipole bins are normalized out of the calculation. Although the HI morphology template itself

shows a suppression of large-scale modes, the correlation with millimeter-wave polarization is strongest at large scales. The statistical tests defined in this paper, however, are based on cross spectra rather than correlation ratios. We form cross spectra between the data collected by CMB experiments and the $\text{H}\alpha$ morphology template defined in Section 3.4.1, and we denote these spectra by $D_{\ell}^{\text{data}\times\text{HI}}$.

We cannot make a direct comparison between $D_{\ell}^{\text{data}\times\text{HI}}$ and $D_{\ell}^{\text{HI}\times\text{HI}}$, because they are not, in general, proportional to each other. As in Cukierman et al. (2023), we model this effect as a multipole-dependent transfer function that describes the representation of the $\text{H}\alpha$ morphology template in the measured dust polarization. We denote the transfer function by t_{ℓ} . The goal in constructing t_{ℓ} is for $D_{\ell}^{\text{data}\times\text{HI}}$ to be approximately proportional to $t_{\ell}D_{\ell}^{\text{HI}\times\text{HI}}$. In our statistical tests, we will compare the former cross spectra to the latter multipole-filtered autospectra.

The aim in introducing the transfer function t_{ℓ} is to boost large-scale modes relative to small-scale modes in order to enhance the sensitivity of our statistical tests. The best estimate of t_{ℓ} would come from $D_{\ell}^{\text{data}\times\text{HI}}/D_{\ell}^{\text{HI}\times\text{HI}}$ (as in Cukierman et al., 2023), but this would lead to a fitting function ($t_{\ell}D_{\ell}^{\text{HI}\times\text{HI}}$) which is partly defined by the data itself. To avoid those complications, we use an ansatz based on the unsharp-mask filter, which produces most of the multipole distortion we wish to correct. This multipole correction is an ansatz and not a model of the true underlying reality. We use it in the same manner as a matched filter, i.e., to increase the sensitivity of our signal search by looking for a particular pattern

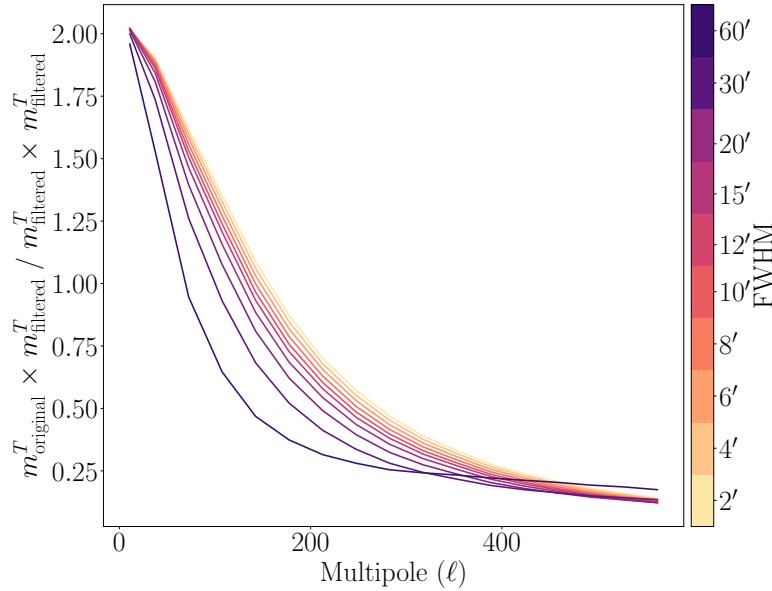


Figure 3.1: The RHT algorithm multipole-dependent unitless transfer function defined in Equation 3.1 for different Gaussian smoothing FWHM values, computed on the Planck 70% sky fraction Galactic plane mask.

rather than simply looking for deviations from zero. A discrepancy between the ansatz and the true reality would simply degrade our sensitivity.

To calculate this transfer function based on the unsharp-mask filter, we apply the following steps to the H I emission maps at each velocity channel:

1. Smooth the original H I intensity map with a Gaussian filter of a specific FWHM.
2. Subtract the smoothed map from the original map.
3. Quantize into a bit mask, i.e. set pixels with values > 0 to 1 and pixels with values < 0 to 0.
4. Multiply the bit mask by the original map.

These are the subset of the steps in the RHT algorithm that most substantially restrict the range of spatial scales of the H I emission that contributes to the measured H I orientation. The subsequent steps, the Hough transform and thresholding, introduce further scale-dependent effects that effectively set the minimum length of a detected linear feature.

We sum the filtered velocity channel maps and call this the *filtered* map. We refer to the velocity-integrated H I intensity as the *original* map. Because we do not expect this transfer function to vary dramatically over the sky, we use the Planck 70% sky fraction Galactic plane mask (Planck Collaboration et al., 2015) as opposed to the BICEP/Keck mask for calculating the transfer function in order to obtain higher signal-to-noise and to capture the filtering effect better over the lower multipole bins. We define the transfer function as

$$t_\ell = \frac{C_\ell^{\text{original} \times \text{filtered}}}{C_\ell^{\text{filtered} \times \text{filtered}}} = \frac{m_\ell^{\text{original}} \times m_\ell^{\text{filtered}}}{m_\ell^{\text{filtered}} \times m_\ell^{\text{filtered}}}. \quad (3.1)$$

We consider the standard 9 bins in the angular multipole range $20 < \ell < 335$ that we use in BICEP/Keck analyses. Note that the only free parameter of the RHT algorithm that is used in this filtering is the Gaussian smoothing radius θ_{FWHM} . In Figure 3.1, we plot this transfer function for the list of θ_{FWHM} values we analyze. This is applied to the H I-correlated component of the simulation in harmonic space. For the rest of this analysis, we present our results with the use of this transfer function. Repeating the analysis without the transfer function produces qualitatively similar results (see Appendix 3.10).

In the next subsection, we will describe a simulation construction that contains a component based on the HI morphology template. We incorporate the multipole correction in the simulation construction such that $D_\ell^{\text{data} \times \text{HI}}$ is approximately proportional to $t_\ell D_\ell^{\text{HI} \times \text{HI}}$. An explicit prescription is provided in the next section.

3.4.3 BICEP/Keck and Planck Simulations Including Filamentary Dust

We construct a set of mock realizations of the sky as observed by the BICEP/Keck and Planck instruments in order to check for biases and estimate uncertainties in the statistical tests introduced in subsequent sections. The baseline dust model in BICEP/Keck analyses is a statistically isotropic Gaussian-dust (GD) field and is our null-hypothesis dust model in this analysis. We call this model GD. It is uncorrelated with the HI morphology template. Simulations of this model are created as random Gaussian realizations with a power spectrum defined by its amplitude $A_{\text{d},353} = 3.75 \mu\text{K}_{\text{CMB}}^2$ at multipole moment $\ell = 80$ and frequency $\nu = 353$ GHz. The power spectrum scales spatially as a power law with index $\alpha_{\text{d}} = -0.4$ in multipole (BICEP/Keck Collaboration et al., 2021). In addition to the baseline dust model, we introduce a second component of filamentary dust that is perfectly correlated with the HI morphology template (HI). This is one realization based on real HI morphology that is added to 499 realizations of GD.

We modify the HI-correlated component in harmonic space according to the transfer function defined in Section 3.4.2 and inverse transform back to map space. We denote the

multipole-filtered version of the $\text{H}\alpha$ morphology template with a tilde ($\tilde{\text{HI}}$). It is important to note that the transfer function introduced in Section 3.4.2 is a phenomenological ansatz rather than a model for the true multipole dependence of the $\text{H}\alpha$ -correlated component of dust polarization. We use this ansatz as a fitting function in Section 3.4.6 in order to improve the sensitivity of our search for $\text{H}\alpha$ -correlated dust polarization, but the ansatz is likely only a rough approximation to the underlying reality. Indeed, we find moderate discrepancies between the measured $\text{H}\alpha$ -dust cross-spectra and the fitting-function ansatz (see Figure 3.7). Furthermore, there is no guarantee that the $\text{H}\alpha$ morphology template should appear in the dust field with a correction that depends only on multipole. If this assumption is made, however, a better estimate of the transfer function can be achieved by appealing to the $\text{H}\alpha$ -dust cross-spectra themselves, which is how a similar transfer function is constructed in Cukierman et al. (2023). As mentioned in Section 3.4.2, however, we wish for our fitting function to be independent of the data to which we are fitting, so we prefer, for the purposes of statistical tests, the ansatz based on the unsharp-mask filtering. For the purposes of constructing mock-sky realizations, it may be superior to use the data-based transfer function in order to keep the mean cross-spectrum bandpowers identical to those of the real data. For computational simplicity, however, we use only the transfer function of Section 3.4.2 for all of the results in this paper. When our mock-sky realizations are used with a nonzero $\text{H}\alpha$ -correlated component, we will only be interested in the variance of our fitting parameters. In the limit of relatively small perturbations, the variance in the fitting

parameters is independent of the mean, so we expect our variance estimates to be reliable in spite of the discrepancy between the measured bandpowers and the mean of the simulated bandpowers.

The full dust field at frequency ν is modeled as

$$m_{\nu}^d(\hat{\mathbf{n}}, a, k, \beta_{\text{HI}}) \equiv a \cdot f_{\nu}(\beta_{\text{GD}}) \cdot m^{\text{GD}}(\hat{\mathbf{n}}) + k \cdot f_{\nu}(\beta_{\text{HI}}) \cdot m^{\text{HI}}(\hat{\mathbf{n}}), \quad (3.2)$$

where $m(\hat{\mathbf{n}})$ represents a Stokes Q or U map, and a , k , and β_{HI} are free parameters. The amplitude a is unitless, and k acts as both an amplitude and a unit conversion factor with units $\mu\text{K}_{\text{CMB}} / \text{K km s}^{-1}$ because $m^{\text{GD}}(\hat{\mathbf{n}})$ has units μK_{CMB} and $m^{\text{HI}}(\hat{\mathbf{n}})$ has units K km s^{-1} . We use a modified blackbody scaling law f_{ν} with a fixed temperature, $T = 19.6$ K, and variable frequency spectral index β (e.g., Planck Collaboration et al., 2014b). The exact choice of dust temperature is of little consequence for our measurements, because we are measuring at frequencies far below the thermal peak. We fix $\beta_{\text{GD}} = 1.6$ in our fiducial model, which is close to the value inferred from data. The exact value does not affect the results because the observables we use in the statistical tests in Section 3.4.6 are cross correlations with the H I morphology template, and the GD and HI components are uncorrelated. In the baseline tensor-to-scalar ratio analysis of BICEP/Keck, we model the dust on the level of cross-frequency B -mode power spectra. In this context, the full dust model of this paper

would manifest itself as

$$D_{\ell}^{\nu_1 \times \nu_2} = a^2 A_d f_{\nu_1}(\beta_{\text{GD}}) f_{\nu_2}(\beta_{\text{GD}}) \left(\frac{\ell}{80} \right)^{\alpha_d} + k^2 f_{\nu_1}(\beta_{\text{HI}}) f_{\nu_2}(\beta_{\text{HI}}) D_{\ell}^{\text{HI} \times \text{HI}}. \quad (3.3)$$

We recover the standard dust model used in BICEP/Keck analyses (the null hypothesis) by setting $a = 1$, and $k = 0$. This hybrid model of GD and HI is continuously related to the GD null hypothesis because the null hypothesis is nested within the hybrid model. We also consider a variation of this model in Appendix 3.10, replacing f_{ν} with a power-law frequency scaling, and find that it does not affect the results, as expected in the Rayleigh-Jeans limit.

In this paper, we limit our analysis to the $\sim 400 \text{ deg}^2$ region mapped by BICEP2 and Keck Array, centered at R.A. 0^{h} , decl. $-57^{\circ}.5$ (hereafter the BICEP/Keck region). On this small region, we use a flat-sky approximation.

We convolve the HI morphology template with instrument-specific beams of different sizes. We also apply the instrument-specific observation matrices used in the BICEP/Keck cosmological analyses, \mathbf{R}_{ν} , capturing the linear filtering of Q and U maps, which includes data selection, polynomial filtering, scan-synchronous signal subtraction, weighting, binning into map pixels, and deprojection of leaked temperature signal (BICEP2 Collaboration et al., 2016). We define

$$\tilde{m}_{\nu}^{\text{HI}}(\hat{\mathbf{n}}) = \mathbf{R}_{\nu}(m^{\text{HI}}(\hat{\mathbf{n}})), \quad (3.4)$$

where $\tilde{m}_\nu^{\text{HI}}$ is the reobserved H I-correlated component of the simulation.

Following standard procedure in BICEP/Keck analyses, we add lensed- Λ CDM (Λ CDM) and noise (n) components to the dust realizations. Refer to BK18 for more details of these simulations. For Planck, we use the official noise simulations provided in the NPIPE data release (Planck Collaboration et al., 2020e).

The model for our total, observed map at frequency ν then becomes

$$\begin{aligned} \tilde{m}_\nu(\hat{\mathbf{n}}, a, k, \beta_{\text{HI}}) = & \tilde{m}_\nu^{\Lambda\text{CDM}}(\hat{\mathbf{n}}) + \tilde{m}_\nu^{\text{n}}(\hat{\mathbf{n}}) \\ & + a \cdot f_\nu(\beta_{\text{GD}}) \cdot \tilde{m}_\nu^{\text{GD}}(\hat{\mathbf{n}}) \\ & + k \cdot f_\nu(\beta_{\text{HI}}) \cdot \tilde{m}_\nu^{\text{HI}}(\hat{\mathbf{n}}). \end{aligned} \quad (3.5)$$

We also purify the maps at each observing frequency with a matrix operation such that the resulting B modes are cleaned of leakage from the much brighter E modes (BICEP2 Collaboration et al., 2016). We then apodize the maps with an inverse noise variance weighting, Fourier transform them, and rotate them from a Q/U to an E/B basis.

We refer to the real BICEP/Keck and Planck maps described in Section 3.3.1 as $\tilde{m}_\nu^{\text{real}}(\hat{\mathbf{n}})$.

3.4.4 Cross Spectra

The statistical tests defined in this paper are based on power spectra calculated using the standard power spectrum estimator of BICEP/Keck analyses as we described in BK18. We

consider 9 bins in the angular multipole range $20 < \ell < 335$ and compute both EE and BB autospectra. We then exploit the linearity of Equation 3.5 to decompose the full cross spectrum with the H I morphology template and calculate the binned bandpower expectation values as

$$\begin{aligned} D_\ell^{\text{data}\times\text{HI}}(a, k, \beta_{\text{HI}}) &= D_\ell^{\Lambda\text{CDM}\times\text{HI}} + D_\ell^{\text{n}\times\text{HI}} \\ &+ a \cdot f_\nu(\beta_{\text{GD}}) \cdot D_\ell^{\text{GD}\times\text{HI}} \\ &+ k \cdot f_\nu(\beta_{\text{HI}}) \cdot D_\ell^{\text{HI}\times\text{HI}}. \end{aligned} \quad (3.6)$$

We concatenate the 9 bandpowers of Equation 3.6 for a selection of frequencies over EE only, BB only, or EE and BB into $\mathbf{D}(a, k, \beta_{\text{HI}})$. The vector $\mathbf{D}(a, k, \beta_{\text{HI}})$ contains the observables from which we construct the covariance matrix in Section 3.4.5 and our statistical tests in Section 3.4.6. We similarly define the vector of cross spectra of the real data with the H I morphology template for a selection of frequencies over EE only, BB only, or EE and BB as \mathbf{D}^{real} .

3.4.5 Covariance Matrices

To construct covariance matrices, we start with 499 realizations of Equation 3.6 of the fiducial model, which coincides with the null-hypothesis model used in the standard BICEP/Keck analyses, i.e. $a = 1$ and $k = 0$. In the covariance matrix construction, we

neglect variances of the $\text{H}\,\text{i}$ -correlated dust component because we expect any uncertainty from the $\text{H}\,\text{i}$ data itself to be subdominant.

There are nonnegligible covariances between neighboring multipole bins and, because the lensed- ΛCDM and dust fields are broadband, between frequency channels. Therefore, we construct a covariance matrix of the form,

$$\mathbf{M} \equiv \frac{N}{N-1} \langle (\mathbf{D}(1, 0, 0) - \bar{\mathbf{D}}(1, 0, 0)) \otimes (\mathbf{D}(1, 0, 0) - \bar{\mathbf{D}}(1, 0, 0)) \rangle_{\text{rlz}}, \quad (3.7)$$

where $\bar{\mathbf{D}}$ is the mean of the vector of spectra over realizations, N is the number of realizations, \otimes is an outer product, and $\langle \rangle_{\text{rlz}}$ is a mean over realizations.

For the statistical test discussed in the next subsection, we use different combinations of the 95, 150, and 220 GHz channels of BICEP/Keck and the 143, 217, and 353 GHz channels of Planck. We use 9 bandpowers per spectrum and separately consider only B modes, only E modes, and E and B modes simultaneously. We condition the covariance matrix by forcing some entries to zero (e.g., Beck et al., 2022). We allow covariances between neighboring multipole bins and between any two frequencies (not just neighboring frequencies), and neglect the correlations between E and B modes in our covariance matrix construction.

3.4.6 Statistical Tests

In this subsection, we define the statistical tests that are used in Section 3.6 of this paper.

χ^2 Likelihood

We approximate the cross spectra defined in Section 3.4.4 between the simulations for our total, observed, maps and the H_I morphology templates as Gaussian distributed, so the natural choice for a test statistic to fit our model is

$$\begin{aligned}\chi^2(a, k, \beta_{\text{HI}}) &\equiv \left(\mathbf{D}^{\text{real}} - \bar{\mathbf{D}}(a, k, \beta_{\text{HI}}) \right)^T \\ &\quad \mathbf{M}^{-1} \left(\mathbf{D}^{\text{real}} - \bar{\mathbf{D}}(a, k, \beta_{\text{HI}}) \right),\end{aligned}\quad (3.8)$$

where, again, $\bar{\mathbf{D}}$ is the mean of the vector of spectra over 499 realizations.

To calibrate this test statistic through simulations, we input an ensemble of realizations from Equation 3.6 with $a = 1$ and $k = 0$ in place of \mathbf{D}^{real} . We fit the model by minimizing Equation 3.8 with respect to the three model parameters a , k , and β_{HI} . We form the test statistic

$$\hat{\chi}^2 \equiv \chi^2(\hat{a}, \hat{k}, \hat{\beta}_{\text{HI}}), \quad (3.9)$$

where \hat{a} , \hat{k} , and $\hat{\beta}_{\text{HI}}$ are the model parameters that minimize Equation 3.8 (e.g., Section 3.6.3).

Because our observables are cross-spectra between the H_I morphology template and the

dust polarization, we expect little sensitivity to the GD amplitude a . We retain a as a fitting parameter, however, so that our null hypothesis ($a = 1, k = 0$) is nested within the full fitting function. This will allow us to form the more sensitive $\Delta\chi^2$ test statistic in Section 3.4.6. Another approach to this analysis could have been to fit for k and β only and to report the statistical significance in terms of the number of standard deviations of \hat{k} from 0. However, we rely on the χ^2 distribution to estimate statistical significance.

When the data are drawn from the null-hypothesis model, the minimized test statistic $\hat{\chi}^2$ is expected to be χ^2 distributed with $n - 3$ degrees of freedom, where n is the number of observables used. For the cases where we only use one frequency band to estimate each band's contribution to the statistical significance of the detection, k and β_{HI} are degenerate. We therefore fit $\hat{k}f_\nu(\hat{\beta}_{\text{HI}})$ as one value. In those cases, there are only 2 effective parameters, a and $kf_\nu(\beta_{\text{HI}})$, and $\hat{\chi}^2$ is χ^2 distributed with $n - 2$ degrees of freedom.

We also use a Markov Chain Monte Carlo (MCMC) method to fully explore this parameter space and provide insight into the correlations and degeneracies between these parameters. We use noninformative uniform distributions for the priors, [-50, 50], [0, 5], and [0.8, 2.4], on a , k , β_{HI} , respectively. The range is large for a because the GD cross spectra with the H I morphology template have no constraining power for a . Using the χ^2 likelihood defined in Equation 3.8, we sample the posterior distributions using the Metropolis-Hastings algorithm implemented in the COBAYA MCMC Python package (Torrado & Lewis, 2019, 2021).

$\Delta\chi^2$ Detection Significance Metric

We form a $\Delta\chi^2$ statistic for measuring the statistical significance of detecting the H I morphology template. We compare $\hat{\chi}^2$ from Equation 3.9 to a model in which a , the amplitude of GD, is allowed to vary but for which $k = 0$. This comparison isolates the influence of the H I-related degrees of freedom.

We form the test statistic

$$\chi_{\text{GD}}^2(a) \equiv \chi^2(a, 0, 0) \quad (3.10)$$

and we minimize with respect to a to obtain

$$\hat{\chi}_{\text{GD}}^2 \equiv \chi_{\text{GD}}^2(\hat{a}^{\text{GD}}), \quad (3.11)$$

where \hat{a}^{GD} is the best-fit value for the model with GD only. The test statistic $\hat{\chi}_{\text{GD}}^2$ is expected to be χ^2 distributed with $n - 1$ degrees of freedom when the data are drawn from the null-hypothesis distribution.

We test for the added benefit of the H I-correlated component with the test statistic

$$\Delta\chi^2 = \hat{\chi}_{\text{GD}}^2 - \hat{\chi}^2, \quad (3.12)$$

which is expected to be χ^2 distributed with 2 degrees of freedom when the data are drawn from the null-hypothesis distribution. If only a single frequency band is used, then $\Delta\chi^2$ is

expected to be χ^2 distributed with only 1 degree of freedom.

The statistical significance of the correlation between the data and the H I morphology template can be estimated from $\Delta\chi^2$. The ensemble of $\Delta\chi^2$ measurements from the null-hypothesis simulations matches a χ^2 distribution with the given number of degrees of freedom. This allows us to calculate a p-value or a probability to exceed (PTE) as $\text{PTE} = 1 - \text{CDF}$, where CDF is the cumulative distribution function of the ensemble up to the $\Delta\chi^2$ value we get from the data. We convert the PTE to an equivalent Gaussian deviate to present the significance as a number of standard deviations from the mean. The reported significances, however, are less reliable $\gtrsim 3\sigma$, where there are no $\Delta\chi^2$ measurements from the null-hypothesis simulations.

3.4.7 Parameter Estimation

We perform a coverage test of our Bayesian model by computing the maximum-likelihood values of a simulation set of 499 realizations with fixed a , k , and β_{HI} values and compare their distributions to the posteriors obtained from real data. We use the best-fit results for k and β_{HI} from the real data. We fit a , such that the autospectrum of the total dust field is equivalent to the GD autospectrum used in BICEP/Keck analyses, and the cross spectrum of the total dust field with the H I morphology template is equivalent to the best-fit autospectrum of the H I morphology template. We call this best-fit \hat{a} to distinguish it from the best-fit \hat{a} we get from Section 3.4.6. We refer the reader to Appendix 3.9 for a detailed

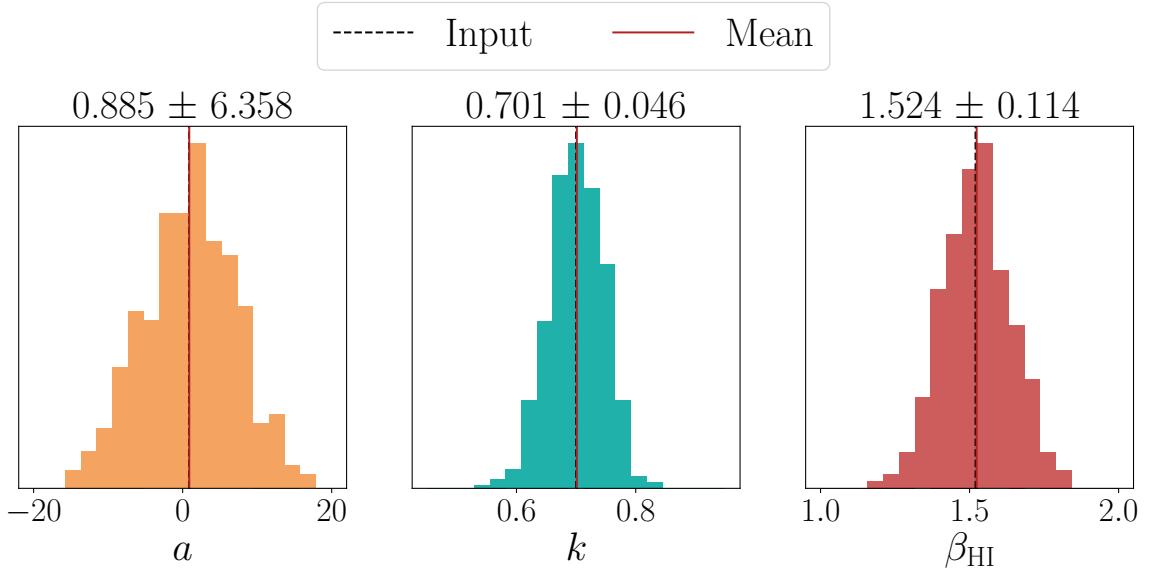


Figure 3.2: Distributions of the best-fit values using E and B modes for 499 realizations of lensed- Λ CDM, noise, and Gaussian dust, added to the $\text{H}\alpha$ morphology template with fixed input values $a = 0.9$, $k = 0.7$, and $\beta_{\text{HI}} = 1.52$ that match the fit from the real data. The parameters a and β_{HI} are unitless, and k has units $\mu\text{K}_{\text{CMB}} / \text{K km s}^{-1}$. These known input values are plotted as dashed black vertical lines. The means of the distributions of the best-fit values are plotted as solid red vertical lines. The mean and standard deviation of each of the distributions are quoted above.

description of this fit.

We then repeat the statistical test defined in Section 3.4.6, replacing \mathbf{D}^{real} with each of the cross spectra of these 499 realizations with the $\text{H}\alpha$ morphology template, and get a distribution of 499 best-fit values for each parameter. Example distributions of the best-fit values from these realizations are shown in Figure 3.2. The distributions shown here are from fitting E and B modes simultaneously using the 95, 150, and 220 GHz bands of BICEP/Keck and the 143, 217, and 353 GHz bands of Planck, conditioning the covariance

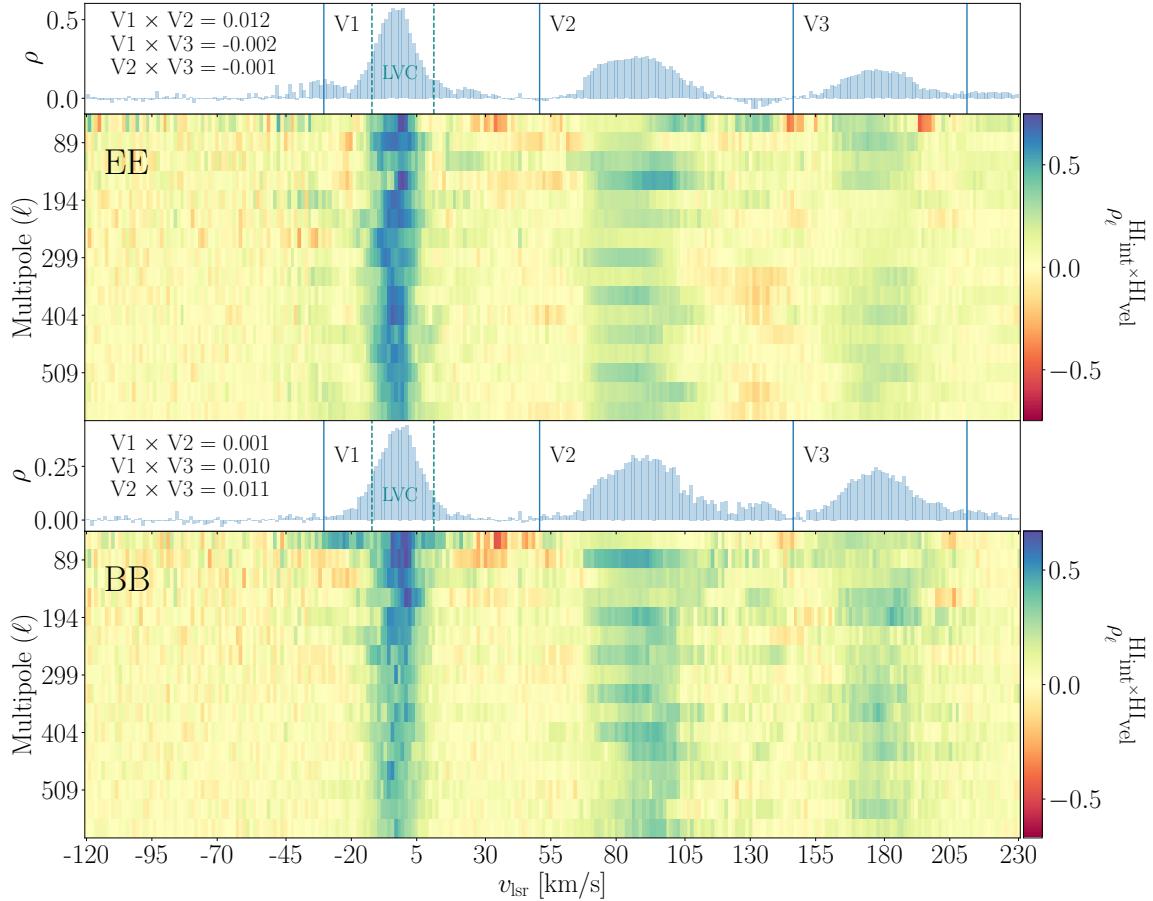


Figure 3.3: *EE* (top) and *BB* (bottom) correlation ratio of the integrated H I morphology template with individual H I morphology templates for the HI4PI velocity channels across multipoles $37 < \ell < 579$. The 1D plots on top show the broadband correlation ratio calculated over one multipole bin spanning the entire multipole range. It is separated into 3 velocity regions, V1, V2, and V3. The LVC boundaries as defined in Panopoulou & Lenz (2020a) are indicated with dashed vertical lines. The broadband correlation ratio between the different pair combinations of the 3 velocity components is printed on the left of each histogram.

matrix, and using a transfer function for the H I morphology template with RHT parameters

$$D_W = 135', \theta_{\text{FWHM}} = 4', \text{ and } Z = 0.75. \quad (3.13)$$

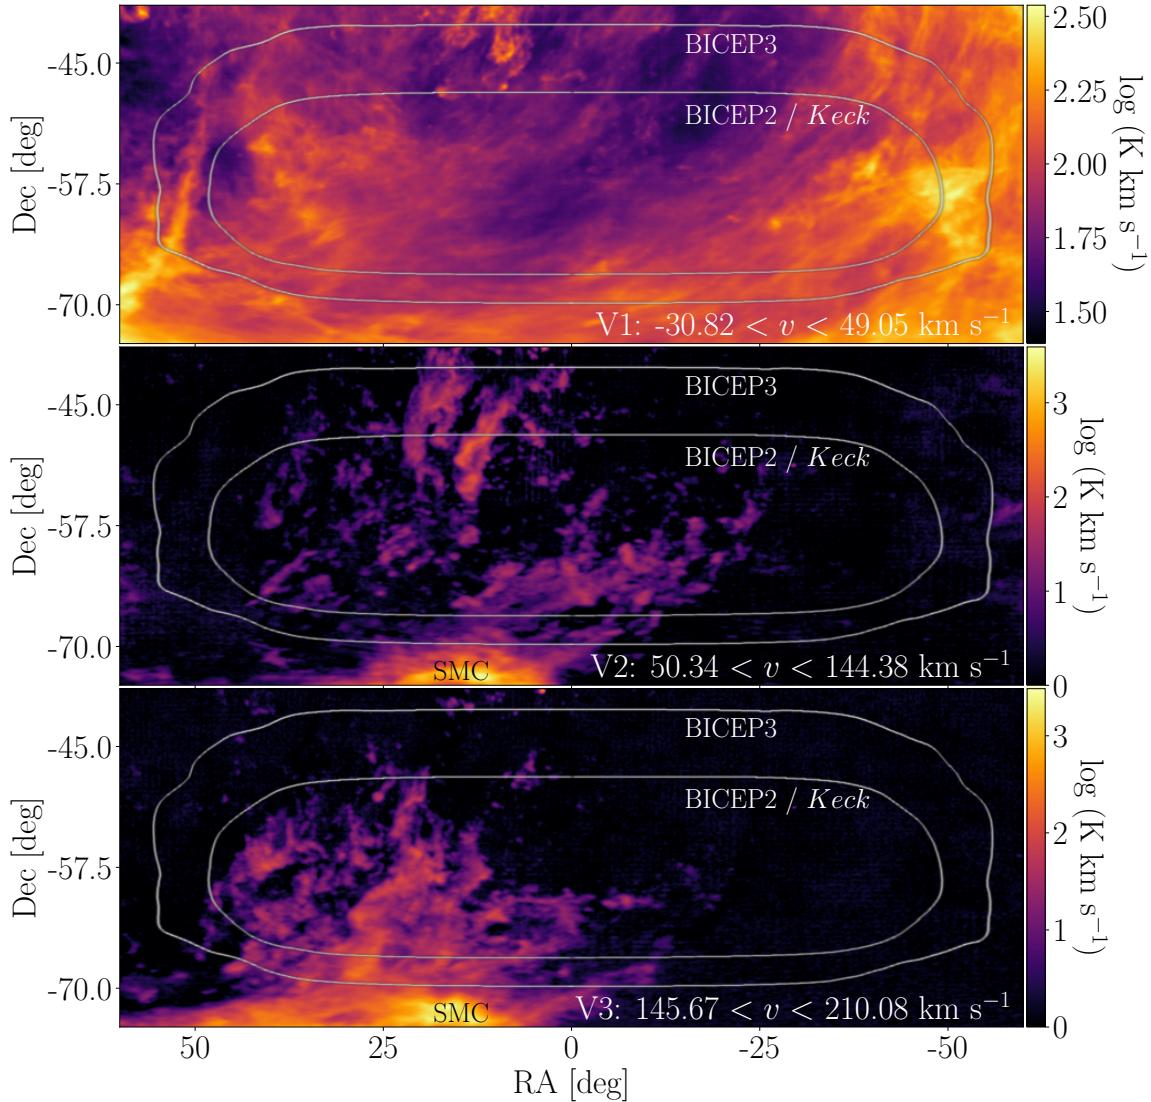


Figure 3.4: Integrated H I intensity maps over the 3 different velocity components defined in Figure 3.3 in the BICEP/Keck region. The velocity boundaries for each component are printed on the bottom right of each map. The emission in V1 is dominated by the Milky Way, whereas the emission in V2 and V3 is dominated by Magellanic Stream 1 (Westmeier, 2018). The outlines of the BICEP3 and the BICEP2 and Keck Array observing fields are also plotted. The Small Magellanic Cloud (SMC) is indicated.

This RHT parameter selection is motivated in Section 3.6.1 and is the fiducial set we use in the results of this paper unless otherwise mentioned. For these choices, the fixed input values

used for constructing the simulation set are 0.9, 0.7, and 1.52 for a , k , and β_{HI} , respectively. We find that our parameter estimation method is unbiased. The sample mean of a is notably close to the input value relative to the standard error, but we checked the p-value and found it to be 4.1%, which we deem to be small but acceptable. We conclude that our fits are unbiased, and we use the spread of the distributions for the 499 realizations to obtain an estimate of the parameter uncertainties. These are consistent with the uncertainties inferred from the marginalized posterior distributions in Section 3.6.3, which are 6.7, 0.050, and 0.13 for a , k , and β_{HI} , respectively. The standard deviation for a is relatively large because the GD cross spectra with the H_I morphology template have no constraining power for a , and this parameter is marginalized over in our analysis.

3.5 Velocity Decomposition

At the high-Galactic latitudes considered here, there is no simple one-to-one mapping between the Galactic H_I emission's velocity along the line of sight and the distance to the H_I gas. However, the bulk velocity of clouds at various distances will often differ, resulting in distinct kinematic components in the H_I spectra. Utilizing the velocity dimension of the 3D H_I morphology Stokes parameter maps in the BICEP/Keck region, we can separate the different velocity components contributing the most to the polarization of the H_I morphology template along the line of sight.

We integrate the H_I morphology Stokes parameter maps in the BICEP/Keck region

across the velocity dimension over the range $-120 \text{ km s}^{-1} < v_{\text{lsr}} < 230 \text{ km s}^{-1}$ (see Section 3.3.2) to form the maps Q_{int} and U_{int} . This is analogous to the line-of-sight integration inherent in thermal dust emission measurements. We then correlate this integrated map with the maps for each velocity channel, $H_{\text{I}_{\text{vel}}}$, using the correlation ratio defined as

$$\rho_{\ell}^{X_{\text{int}} \times X_{\text{vel}}} = \frac{D_{\ell}^{X_{\text{int}} \times X_{\text{vel}}}}{\sqrt{D_{\ell}^{X_{\text{int}} \times X_{\text{int}}} \times D_{\ell}^{X_{\text{vel}} \times X_{\text{vel}}}}}, \quad (3.14)$$

where X denotes either the E or B modes of the $\text{H}\alpha$ morphology templates, and D_{ℓ} is the cross spectra over multipole moment ℓ . This metric quantifies the contribution of each velocity channel map to the polarization signal of the line-of-sight integrated template. We use the RHT parameters in Equation 3.13 for this plot; though the results are qualitatively similar when varying those parameters.

We plot $\rho_{\ell}^{X_{\text{int}} \times X_{\text{vel}}}$ in Figure 3.3, where each column represents the correlation of each velocity channel map with the integrated map, and each row represents a multipole moment bin. We expect neighboring velocity channels to be correlated on physical grounds. Therefore, the consistent horizontal bands at each multipole bin in the 2D plots are due to the similarity between adjacent velocity channels.

We also calculate a broadband correlation coefficient that is binned into one multipole bin that spans the entire range ($37 < \ell < 579$) and plot it above the 2D plots in Figure 3.3. We clearly see distinct peaks in three different velocity ranges, which we refer to as V1, V2, and V3. These peaks are in roughly the same locations as the peaks we see when

plotting the H_I intensity as a function of velocity but have different relative amplitudes, with the second peak having a much lower amplitude in intensity than the third peak. We plot vertical lines to define roughly where the boundaries between those components are. As we will show in Section 3.6 and Table 3.2, the exact boundaries do not affect the results, which are dominated by the velocity channels at the peaks.

The H_I line emission at high-Galactic latitudes is conventionally divided into low-velocity clouds (LVCs), intermediate-velocity clouds (IVCs), and high-velocity clouds (HVCs) based on its radial velocity with respect to the local standard of rest (v_{lsr}) or the Galactic standard of rest (v_{gsr}), or on its deviation from a simple model of Galactic rotation (see, e.g., Putman et al. (2012) for more details). The boundaries between these classes vary by tens of kilometers per second in the literature. For instance, Magnani & Smith (2010), Wakker (1991), and Wakker (2001) define the boundary between LVCs and IVCs at $|v_{\text{lsr}}| = 20, 30$, and 40 km s^{-1} , respectively. Panopoulou & Lenz (2020a) propose $-12 \text{ km s}^{-1} < v_{\text{lsr}} < 10 \text{ km s}^{-1}$ as the range for LVCs based on the first and 99th percentiles of the distribution of cloud velocities that pass a certain threshold in the H_I column density in the Northern and Southern Galactic Polar regions.

The boundary between IVCs and HVCs is usually taken to be at $|v_{\text{lsr}}| = 70 \text{ km s}^{-1}$ (Wakker & Boulanger, 1986) or 90 km s^{-1} (Richter & De Boer, 2005). The boundaries for V1 defined here encompass the range of LVCs adopted by Panopoulou & Lenz (2020a) as shown in Figure 3.3. We limit the higher end of the IVC range to $|v_{\text{lsr}}| = 50 \text{ km s}^{-1}$ in

the BICEP/Keck region such that V2, which is primarily associated with the Magellanic System (Westmeier, 2018), is excluded. As already mentioned, the results are dominated by the velocity channels at the peaks, and the exact boundaries do not affect the results.

Our interpretation of the peaks in Figure 3.3 is that each corresponds to a substantial contribution of that velocity component to the integrated map. As a sanity check, however, we test whether the V2 and V3 peaks in the correlation with the integrated map are due to spurious correlations with each other or with V1 by calculating $\rho^{X_{Vi} \times X_{Vj}}$, where i and $j \in \{1, 2, 3 \mid i \neq j\}$. We report those values in Figure 3.3 and find that the correlation is less than approximately 1%.

We integrate the velocity channel maps in each range and plot the resulting H I intensity maps in Figure 3.4 on a log color scale. V1 is dominated by H I emission from the Galaxy, whereas V2 and V3 are dominated by H I emission from Magellanic Stream 1, a stream of high-velocity gas associated with the Magellanic System (Westmeier, 2018). The outlines of the BICEP3 and the BICEP2 and Keck Array observing fields are included in the figure to distinguish the H I structure that lies inside and outside each of the observing fields. For consistency in our statistical tests defined in Section 3.4.6, we analyze the smaller field as mentioned in Section 3.4.3. The bright emission in V2 and V3 directly below the BICEP3 observing field in decl. is from the Small Magellanic Cloud (SMC).

	Default	Best
<i>BB</i>	4.7	6.7
<i>EE</i>	12.3	14.6
<i>BB + EE</i>	12.9	16.1

Table 3.1: Statistical significance of the detection of V1 in units of equivalent Gaussian standard deviations as defined in Section 3.4.6 using the 95, 150, and 220 GHz bands of BICEP/Keck and the 353 GHz band of Planck. The column labeled "best" uses the parameters $D_W = 135'$, $\theta_{\text{FWHM}} = 4'$, and $Z = 0.75$, and the row labeled "default" uses the parameters $D_W = 75'$, $\theta_{\text{FWHM}} = 30'$, and $Z = 0.7$, which are used in Clark & Hensley (2019).

3.6 Results and Discussion

In this section, we tune the RHT parameters to increase the correlation between BICEP/Keck and Planck data with the H_I morphology template (Section 3.6.1). Using the tuned parameters, we quantify the detection of filamentary dust polarization in the Galactic component of H_I (Section 3.6.2). We look for evidence of frequency decorrelation in the BICEP/Keck region from the inclusion of the IVC component in the line-of-sight sum and between the filamentary dust component and the total dust component (Section 3.6.3). We also quantify the contribution of each of the datasets used in this measurement (Section 3.6.4). Finally, we look for a detection of filamentary dust polarization in the higher-velocity H_I components associated with Magellanic Stream 1 (Section 3.6.5).

3.6.1 Tuning and Improving the RHT Model

Due to computational expense, the RHT parameter space has not been explored before in the context of building dust polarization templates. However, limiting the sky area to the

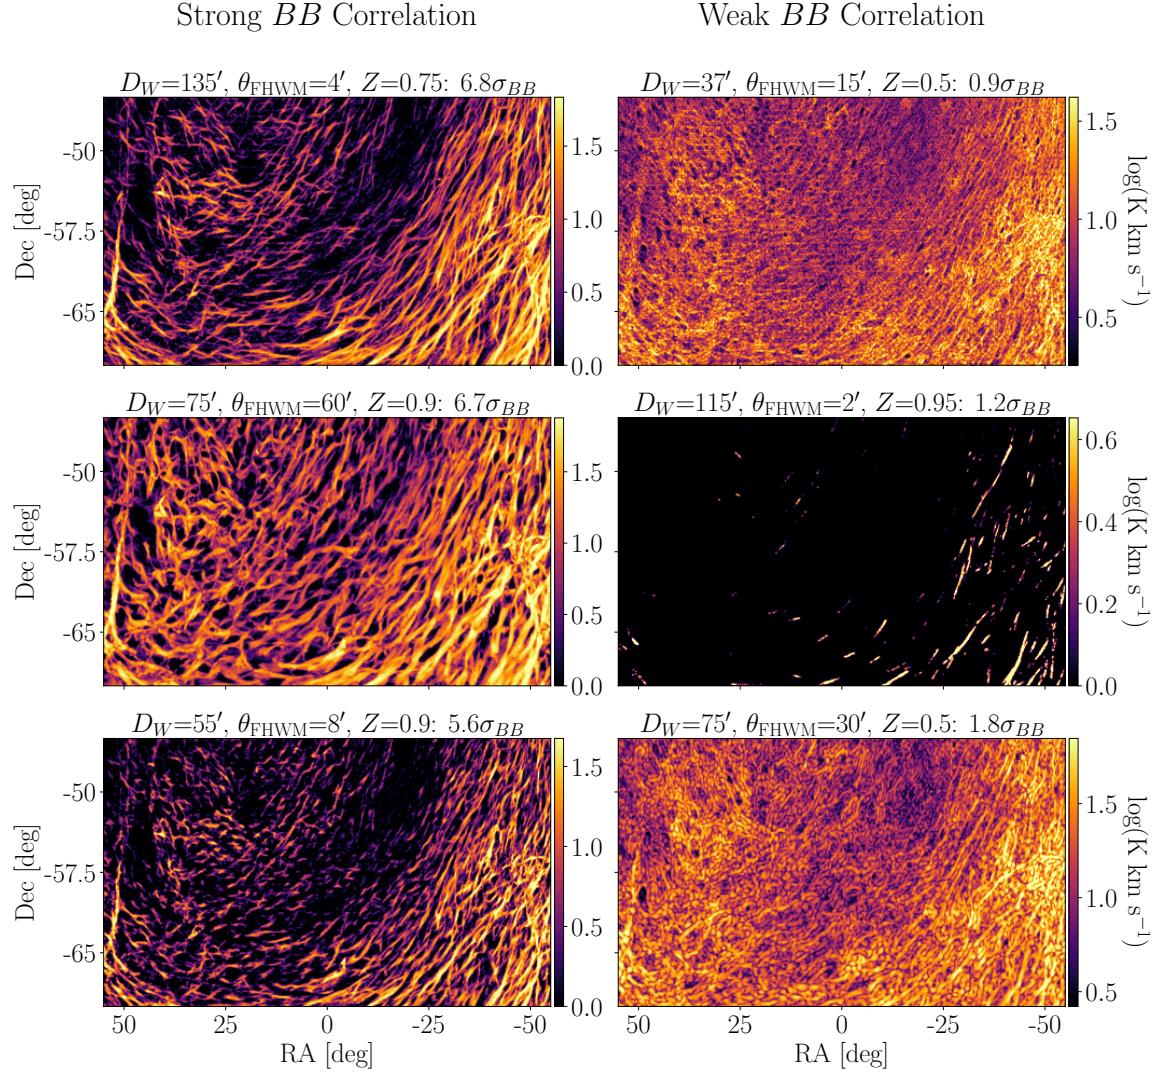


Figure 3.5: Polarized intensity maps of V1 in the BICEP/Keck region using RHT parameters that correlate $> 5\sigma$ (left) and $< 5\sigma$ (right) in B modes with BICEP/Keck and Planck data. Only the statistical significance in B modes is quoted in the title of each of the maps, because all of the RHT parameters we tried correlate well ($> 5\sigma$) in E modes. From top to bottom, the maps on the left have a $15.2\sigma_{EE}$, $12.6\sigma_{EE}$, and $14.9\sigma_{EE}$ detection significances, and the maps on the right have a $6.3\sigma_{EE}$, $8.6\sigma_{EE}$, and $8.2\sigma_{EE}$ detection significances.

BICEP/Keck region and speeding up the algorithm by $\sim 35\times$, as described in Section 3.4.1,

have allowed us to search the parameter space more efficiently. We evaluate the $\Delta\chi^2$ metric

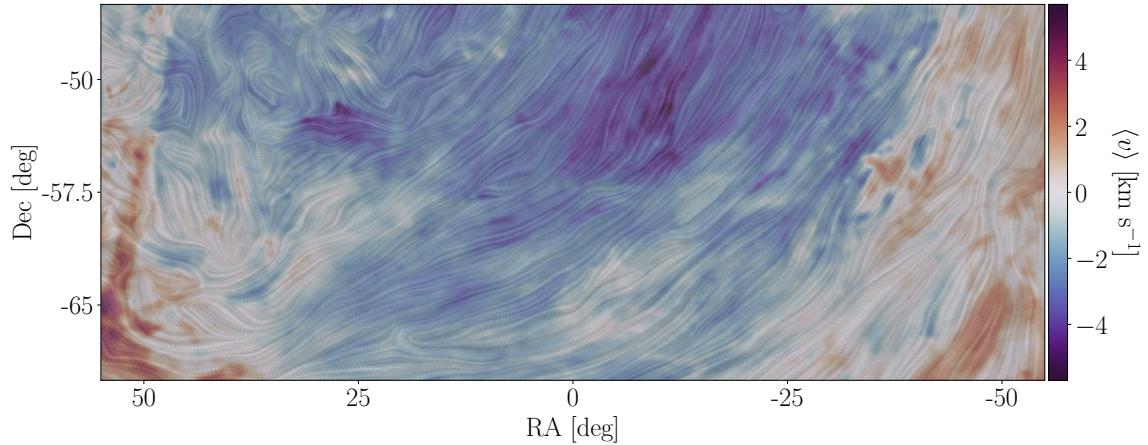


Figure 3.6: Map of the first moment of the velocity distribution of the H₁ structure in the BICEP/Keck region for $-12 \text{ km s}^{-1} < v_{\text{lsr}} < 10 \text{ km s}^{-1}$, the velocity range most correlated with the polarized dust emission. The texture is a line integral convolution of the magnetic field orientation as inferred by the H₁ filaments.

from Section 3.4.6 in parallel on a grid of values spanning a reasonable range of interest in each of the RHT parameters. We consider $D_W = 37', 55', 75', 95', 115', 135',$ and $149'$; $\theta_{\text{FWHM}} = 2', 4', 6', 8', 10', 12', 15', 30',$ and $60'$; and $Z = 0.5, 0.7, 0.75, 0.8, 0.85, 0.9,$ and 0.95 .

We find that the RHT parameters that maximize the statistical significance of the detection among the ones we tried are $D_W = 135'$, $\theta_{\text{FWHM}} = 4'$, and $Z = 0.75$. These parameters maximize the statistical significance when fitting the metric using *B* modes only, *E* modes only, and *E* and *B* modes simultaneously. We compare the results we get using these parameters to the results we get using the RHT parameters used in Clark & Hensley (2019) ($D_W = 75'$, $\theta_{\text{FWHM}} = 30'$, and $Z = 0.7$) in Table 3.1. Our results improve by $\sim 2\sigma$ in *BB* and in *EE* and by $\sim 3\sigma$ when *EE* and *BB* are combined.

We look for trends in the detection significance with BICEP/Keck and Planck data when varying each of the RHT parameters monotonically, but do not find any. Instead, we find that the correlation is robust for a wide range of parameter choices. The exceptions are at the extremes of the parameter space. We show examples of polarized intensity maps of the H I morphology templates made with parameters that correlate well or poorly with the dust in *BB* in Figure 3.5. The polarized intensity is defined as

$$P = \sqrt{Q^2 + U^2}, \quad (3.15)$$

where Q and U are the Stokes parameters of the H I morphology template. We quote the statistical significance of the detection in *BB* in the title of each panel. All of the variations we tried correlate well ($> 5\sigma$) in *EE*, including the ones shown in Figure 3.5 with their detection significances stated in the caption. Note that the examples that are weakly correlated with the dust in *BB* either have a high Z ($Z \gtrsim 0.95$) and $D_W \gg \theta_{\text{FWHM}}$ or have a low Z ($Z \lesssim 0.5$). While the significance is larger than 5σ in *EE* for the examples on the right, it is still fairly low by *E*-mode standards compared to the examples on the left with a lower correlation ratio.

The cases with a high Z limit the RHT-detected linear structure to longer, more connected filaments, while lower Z decomposes the H I intensity into numerous shorter filaments. The choppiness of the filaments affects the predicted *B*-mode power more than it does the *E*-mode power because the *B*-mode structure of this template is affected by the finite extent

of the filaments. Real-space maps of the E - and B -mode amplitudes support this intuition (Huffenberger et al., 2020). The net signal arising from choppy, colinear filaments produces a constructive interference for E modes but a destructive interference for B modes. Also, because θ_{FWHM} affects the largest spatial scales of the H_I emission and the product of the D_W and the Z parameters defines an effective lower limit on the length of the filaments, the combination of high Z with $D_W \gg \theta_{\text{FWHM}}$, such as the middle right panel of Figure 3.5, discards most of the structure in the map and is only sensitive to the most prominent filaments. The B -mode-correlated H_I structure is related to the overall distribution of filaments, such that annihilating all but a few substantially weakens the correlation with the dust B modes.

We defer a more comprehensive interpretation of the RHT parameters and their implications to a future study. For now, we propose the parameters in Equation 3.13 as the recommended ones when using the RHT in future analyses on Hi4PI data for making dust polarization or magnetic field templates in the diffuse, high-Galactic latitude ISM. However, these parameters might be sensitive to the BICEP/Keck filtering or to the specific sky region. These effects will be explored in future work.

3.6.2 Filamentary Polarization in the Local ISM

Using the $\Delta\chi^2$ statistical test defined in Section 3.4.6, we find a significant correlation between the H_I morphology templates and the first velocity component, V1, as shown in

Table 3.1. These results are insensitive to covariance matrix conditioning, frequency scaling law, or use of a transfer function for the H I morphology template as shown in Table 3.5.

Above a certain threshold in the column density of H I, Panopoulou & Lenz (2020a) find an agreement between the Northern and Southern Galactic Polar regions in the first and 99th percentiles of the H I cloud velocity distributions. They therefore use those percentiles to adopt the boundaries $-12 \text{ km s}^{-1} < v_{\text{lsr}} < 10 \text{ km s}^{-1}$ between LVCs and IVCs. We use this range to visualize the first moment map of the velocity distribution of the H I structure in the BICEP/Keck region in Figure 3.6. That is, we plot the intensity-weighted mean velocity,

$$\langle v \rangle = \frac{\sum_v v \cdot I(v)}{\sum_v I(v)}, \quad (3.16)$$

to highlight the regions in the map where the emission is dominated by different velocities. This is the velocity range that exhibits the most substantial contribution to the dust-correlated template as we show in Section 3.6.3. We perform a line integral convolution (Cabral & Leedom, 1993) on the H I morphology Q and U maps in that velocity range, smoothed to the RHT window diameter scale, to visualize the magnetic field orientation inferred by the H I filaments and overplot it as the texture in Figure 3.6.

3.6.3 Frequency Decorrelation and the Polarized Dust SED

Dust components along the same line of sight with different polarization angles and SEDs give rise to a phenomenon called line-of-sight frequency decorrelation. We test for evidence

of this phenomenon in the BICEP/Keck region between the LVC and IVC components and between the filamentary and total dust components.

LVCs and IVCs are known to contain dust (Boulanger et al., 1996a; Reach et al., 1998; Planck Collaboration et al., 2011b). The velocity range of V1 spans both LVCs and IVCs using the velocity boundaries defined in Panopoulou & Lenz (2020a). These are the same boundaries that Pelgrims et al. (2021a) use in their analysis of line-of-sight frequency decorrelation in Planck data. Panopoulou & Lenz (2020a) use a Gaussian decomposition of the H I emission profiles to estimate the number of distinct clouds along each sightline. While they show that most sightlines in the BICEP/Keck region are dominated by one LVC cloud on average, they do detect more than one cloud along some sightlines. Pelgrims et al. (2021a) detect line-of-sight frequency decorrelation in the sightlines that contain LVCs and IVCs with different polarization angles predicted by H I morphology. While we know from Panopoulou & Lenz (2020a) that IVCs are not an important fraction of the H I column in the BICEP/Keck region, we check whether that is also true in polarization, i.e., whether the polarization inferred from the H I morphology templates in the IVC velocity range contributes significantly to the correlation with dust polarization. We find that the IVC emission integrated over the BICEP/Keck region is $\sim 25\%$ of the V1 column in intensity and $\sim 10\%$ of the V1 column in polarized intensity. Table 3.2 shows that the detection significance is not strongly changed by the inclusion of IVC-associated H I morphology template in the line-of-sight sum, as expected on account of the amplitude ratios. The shifts

	Range for LVCs + IVCs	Range for LVCs
<i>BB</i>	6.7	6.8
<i>EE</i>	14.6	14.3
<i>BB + EE</i>	16.1	16.1

Table 3.2: Comparison of the statistical significance of a detection of the cross correlation with the dust polarization in units of equivalent Gaussian standard deviations when including the channels in the IVC velocity range in the line-of-sight sum. The RHT parameters from Equation 3.13 are used here for the H I morphology template with the 95, 150, and 220 GHz bands of BICEP/Keck and the 353 GHz band of Planck.

in detection significance are $\lesssim 0.3\sigma$ in all cases.

Therefore, we do not have good reason to expect strong decorrelation from the IVC population in the BICEP/Keck region. However, there could be frequency decorrelation arising from different dusty regions along the line of sight that are all associated with gas within the LVC range. The kinematic substructure of the LVC H I could in principle be used to further explore the 3D distribution and phase structure of the gas in this region, and its possible association with different contributions to the total dust SED.

Since the H I morphology template is filamentary, the *E* and *B* modes of this template are sourced by the same filaments (Huffenberger et al., 2020), although variations in the 3D dust properties could still give rise to SED differences between *E* and *B* modes (Vacher et al., 2022). Minimizing the χ^2 test statistic defined in Equation 3.8, we fit β using both *E* and *B* modes simultaneously.

For the most sensitive measurement of β_{HI} in V1, we use both *E* and *B* modes, the best-fit RHT parameters from Equation 3.13, the 95, 150, and 220 GHz bands of BICEP/Keck, and the 143, 217, and 353 GHz bands of Planck. We condition the covariance matrix and use a

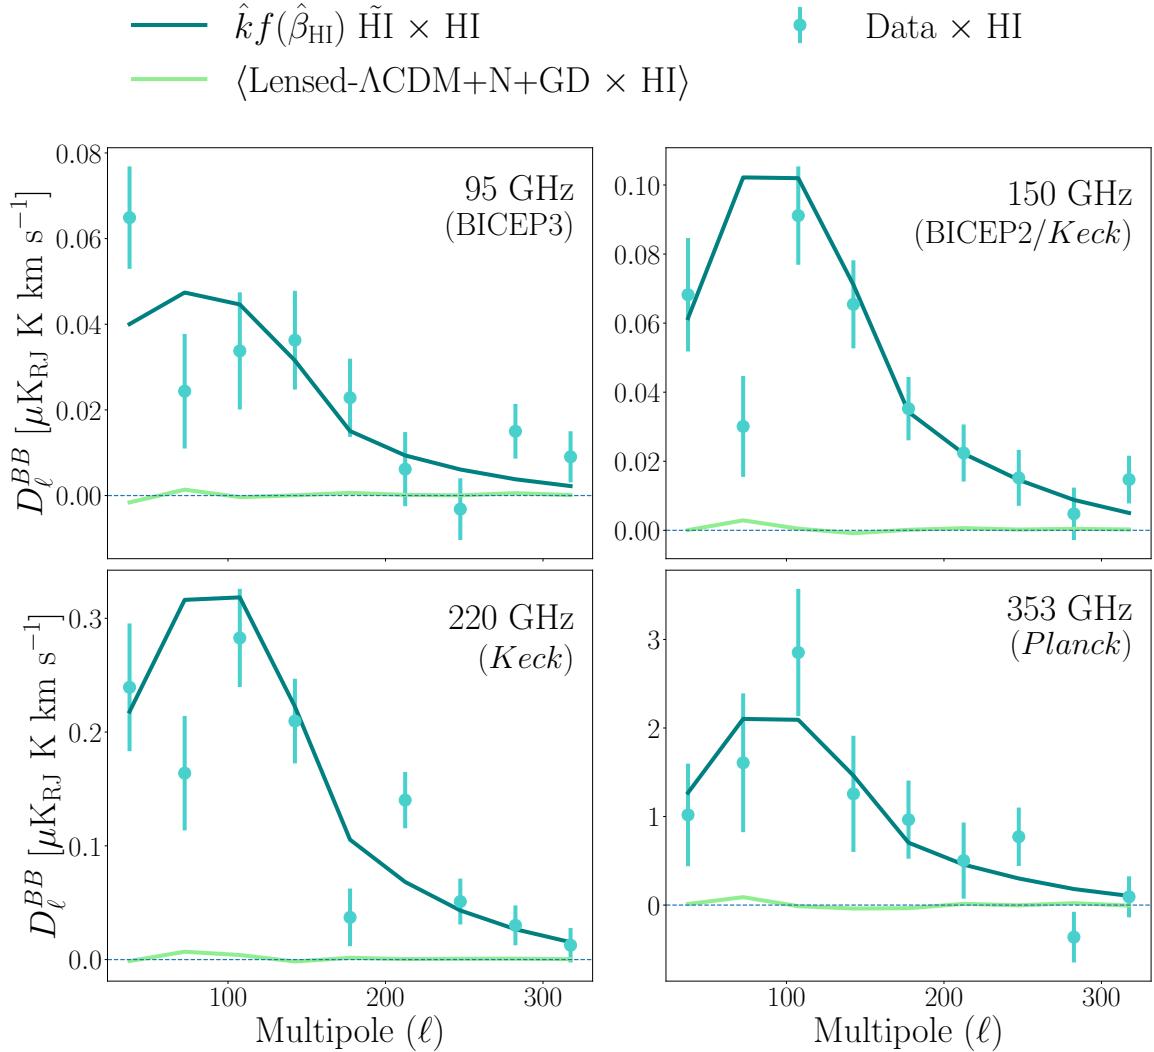


Figure 3.7: The best-fit BB observables used in the $\Delta\chi^2$ statistic defined in Section 3.4.6 for the 95, 150, and 220 GHz bands of BICEP/Keck and the 353 GHz band of Planck. A modified blackbody frequency scaling, covariance matrix conditioning, and a transfer function for the $\text{H}\alpha$ morphology template with the RHT parameters from Equation 3.13 are used for the fit here. The cross spectrum between the real data and the $\text{H}\alpha$ morphology template (light blue), the best-fit cross spectrum between the $\text{H}\alpha$ morphology template and the modified $\text{H}\alpha$ -correlated component of the simulation (dark blue), and the mean of the cross spectra between the $\text{H}\alpha$ morphology template and the lensed- ΛCDM , noise, and Gaussian-dust components of the simulation (light green) are plotted.

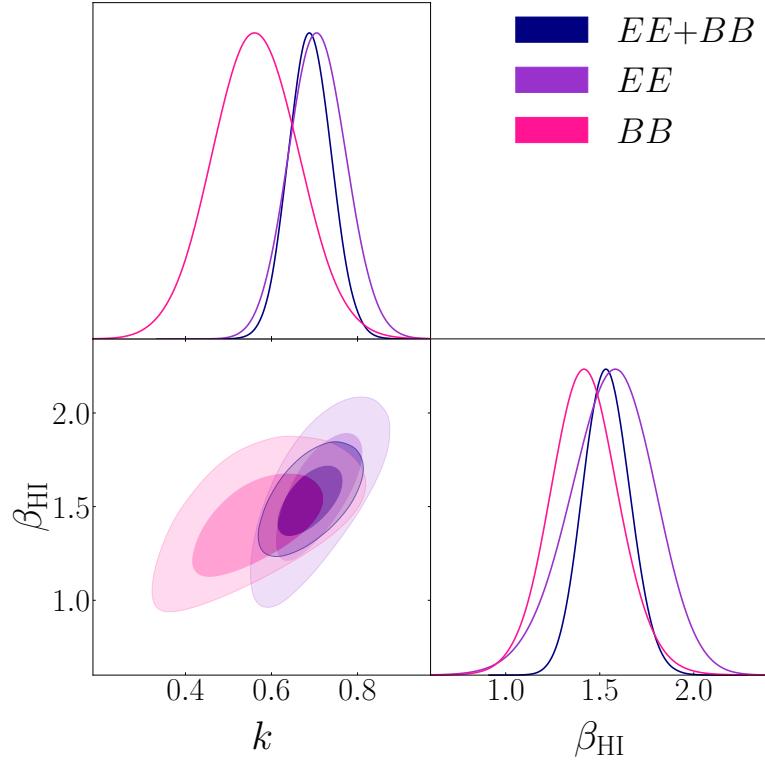


Figure 3.8: Posteriors of k and β_{HI} fit using the Metropolis-Hastings algorithm on uniform priors and the χ^2 likelihood of the cross spectra of the real data with the $\text{H}\,\text{i}$ morphology template. The parameter a is marginalized over. The E modes only (purple), B modes only (pink), and simultaneous E and B modes (navy) posteriors are shown. The units for k are $\mu\text{K}_{\text{CMB}} / \text{K km s}^{-1}$ and β_{HI} is unitless.

transfer function for the $\text{H}\,\text{i}$ morphology template; though those choices do not substantially

affect the result as shown in Appendix 3.10.

From the χ^2 minimization described in Section 3.4.6, we get $\chi^2/\text{d.o.f.} = 1.4$, where d.o.f. is the number of degrees of freedom. We find $\hat{\beta}_{\text{HI}} = 1.52 \pm 0.11$ and plot the best-fit BB observables for the 4 most sensitive bands used in this measurement in Figure 3.7. The error bars are the square root of the diagonal elements of the covariance matrix used in the χ^2 fit. Since the $\text{H}\,\text{i}$ morphology template does not correlate with the lensed- ΛCDM , noise,

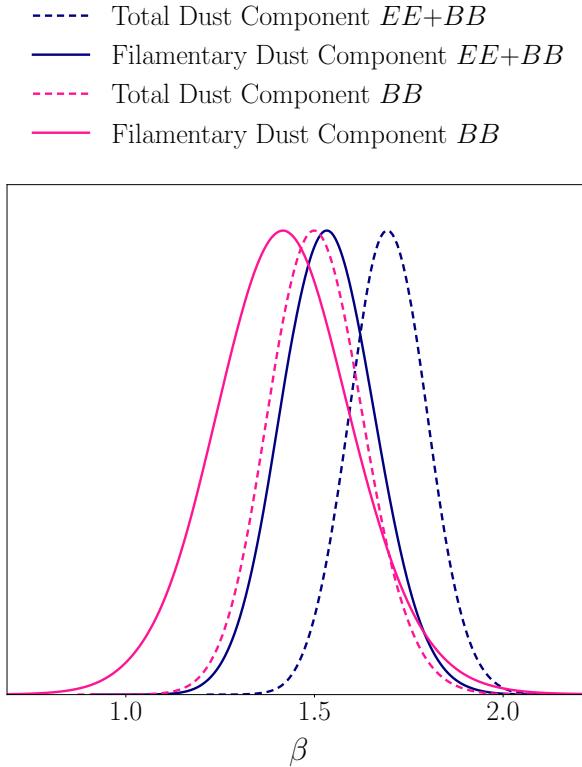


Figure 3.9: Comparison of the posteriors for β_{HI} through a χ^2 likelihood using cross correlations with the $\text{H}\alpha$ morphology template (solid) to the ones of β_d using the Hamimeche and Lewis (HL) likelihood with a multicomponent model and no $\text{H}\alpha$ morphology template (dashed). We show the posteriors using B modes only (pink), and B and E modes (blue). The solid posteriors are the same as in Figure 3.8 plotted with the same colors. The B -mode-only total dust component posterior is identical to the posterior shown in black in Figure 4 of BK18.

and GD components, the mean of these cross spectra plotted in light green is statistically consistent with zero. Any visible deviations are due to the sample variance in the finite simulation ensemble. The $55 < \ell < 90$ bandpower of the cross spectrum between the real data and the $\text{H}\alpha$ morphology template fluctuates low relative to the cross spectrum between the $\text{H}\alpha$ morphology template and the $\text{H}\alpha$ -correlated component of the simulation,

which is modified by the transfer function defined in Section 3.4.2. This is consistent across frequencies because each multipole bin bandpower is well correlated with the bandpower of the same multipole bin at the different frequencies. The modified H I-correlated component of the simulation is not guaranteed to match the real data, because we do not have a data-driven model for the multipole-dependent representation of the H I morphology template in the real dust field (Section 3.4.2). Note that the cross correlations with the real data highly exceed the spurious correlations across all frequencies.

Taking a Bayesian approach, we use `COBAYA` (Torrado & Lewis, 2019, 2021) to run MCMC and compute the posteriors on a , k , and β_{HI} as described in Section 3.4.6. We marginalize over a because the GD cross spectra with the H I morphology template have no constraining power for a , and show the contour plots for the more interesting k and β_{HI} in Figure 3.8 for E modes only, B modes only, and E and B modes simultaneously. The value for k folds in the normalization of the H I morphology template. However, the more standard deviations away from zero it is, the stronger the detection of an H I-correlated component there is in the cross spectra of the real data with the H I morphology template. The posterior of $\beta_{\text{HI}} = 1.54 \pm 0.13$ when using E and B modes simultaneously is consistent with the best-fit value and standard deviation we get using the frequentist maximum-likelihood approach.

We find consistency between the spectral index of the filamentary dust SED, β_{HI} , and the total dust SED, β_{d} , as obtained in BK18 by fitting BICEP/Keck, WMAP and Planck B -mode auto and cross spectra to a GD model. That work used a multicomponent parametric

model with the Hamimeche and Lewis (HL) likelihood that includes auto and cross spectra across frequencies. The posteriors are shown in Figure 3.9 with repeated posteriors from Figure 3.8 for comparison. The posteriors plotted are measuring a related but different quantity, because we are correlating with a filament-based template in this paper. The results obtained in BK18 are based on a dust model that assumes a constant ratio between the dust *EE* and *BB* power spectra. In this paper, we modify this model to allow the dust *EE* and *BB* power spectra to have independent power-law spectral behavior. We find a slight shift to higher values when *E* modes are included in the fit. The best-fit values and 1σ deviations for the filamentary and total dust components, respectively, are 1.42 ± 0.19 and 1.49 ± 0.13 for *BB* and 1.54 ± 0.13 and 1.70 ± 0.10 for *EE + BB*. We do not find significant tension between the filamentary and total dust SEDs. However, it would be interesting to check whether the differences become statistically significant with tighter uncertainties, which would have important implications for *B*-mode cosmology.

Since the H_I morphology model is identifying only filamentary contributions to the dust polarization, the similarity in the best-fit values and posteriors for β between the two methods indicates that there is no evidence of decorrelation between the filamentary structures that are preferentially associated with the cold neutral medium (Clark et al., 2019; Kalberla et al., 2020) and the rest of the dust column in the BICEP/Keck region. If the H_I morphology method yielded a different SED, the combination of H_I and GD would produce different polarization angles at different frequencies due to the changing relative

	<i>BB</i>	<i>EE</i>	<i>BB + EE</i>
BICEP3 95 GHz	4.53	1.22	4.72
Planck 143 GHz	0.05	0.72	0.12
BICEP2/Keck 150 GHz	5.31	2.43	5.98
Planck 217 GHz	3.50	2.37	4.02
Keck 220 GHz	5.82	7.13	9.26
Planck 353 GHz	3.18	7.99	8.59

Table 3.3: Comparison of the statistical significance of a detection of the cross correlation between H I morphology template and the dust polarization at different frequencies in units of equivalent Gaussian standard deviations as defined in Section 3.4.6.

weighting between the two components. We also find that the results for β_{HI} are consistent for different RHT parameters.

The fact that we find a similar SED fit for the filamentary component and for the total dust in the BICEP/Keck region does not have to be the case in other regions of the sky. The dust associated with the warmer, more diffuse H I component may scale differently in frequency in other regions. Because the H I morphology templates use the orientation of filamentary structures, a data-driven model for the dust polarization associated with the diffuse, nonfilamentary dust is currently lacking.

3.6.4 Individual Frequency Band Contribution

We study the contribution of each band and instrument used in the results of Sections 3.6.2 and 3.6.3 and measure the statistical significance of the detection of filamentary dust polarization as a function of frequency.

We measure a significant detection of dust down to 95 GHz as shown in Table 3.3.

These results are insensitive to the covariance matrix conditioning, frequency scaling law, or use of a transfer function for the $\text{H}\,\text{i}$ morphology template as shown in Table 3.7. We find that, in the BICEP/Keck region, the BICEP3 95 GHz band is more sensitive to dust polarization than any of the Planck bands below 353 GHz when using both E and B modes and is more sensitive than any Planck band when using B modes only. When using both E and B modes, the Planck 353 GHz band is the only Planck band that exceeds 5σ , while the 150 and 220 GHz bands of BICEP/Keck both exceed 5σ , and the 95 GHz band is correlated with the $\text{H}\,\text{i}$ morphology template at $\sim 5\sigma$. This shows the power of the BICEP/Keck bands for characterizing the dust in this field, and especially, for measuring its SED. The detection at 95 GHz is also interesting because it provides a low-frequency lever arm for the dust SED, and it is the band where the ΛCDM component starts to dominate over the dust component in polarized emission at smaller scales (BICEP/Keck Collaboration et al., 2021). These results are consistent with our expectations from the map depths we have shown in BK18 and with the number of standard deviations away from zero the peak of the posterior for k is for each case.

At frequencies lower than 220 GHz, almost all of the detection significance is coming from B modes. That is, the statistical significance of the detection is equivalent at lower frequencies when including E modes. This is because at lower frequencies, in E modes, we are limited by the sample variance of the CMB, i.e., the statistical significance of the detection will not improve unless we remove the CMB component or increase the observed

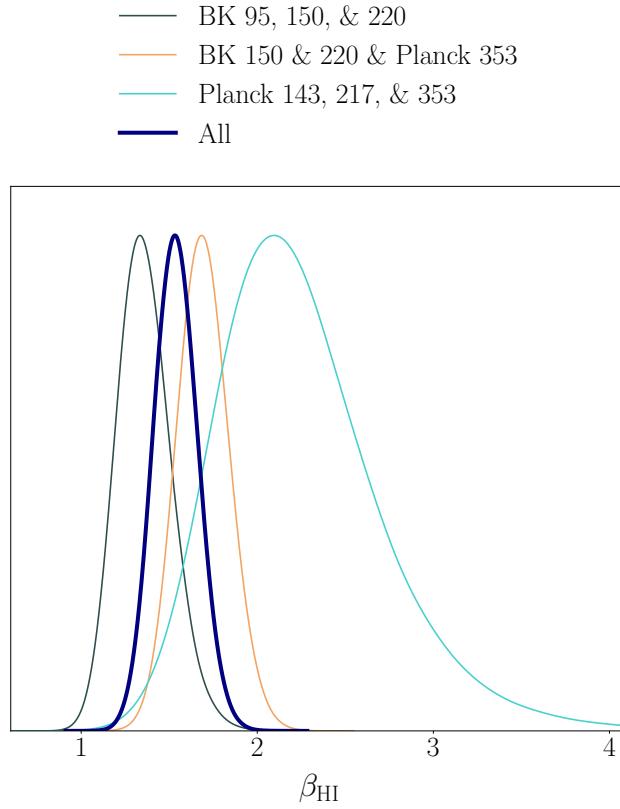


Figure 3.10: Comparison of the posteriors for β_{HI} we get through a χ^2 likelihood using E - and B -mode cross correlations with the $\text{H}\alpha$ morphology template for different selections of frequency bands and for BICEP/Keck only and Planck only variations. The thick navy posterior labeled "All" is the same as the navy posterior in Figures 3.8 and 3.9.

sky area. The E modes at those frequencies produce a negligible change in the overall significance estimates because they are downweighted by our statistical metrics.

Moreover, we can measure whether the SED changes when we omit the low- or high-frequency channels from our analysis. We show the β_{HI} posteriors, using both E and B modes in the fits, in Figure 3.10. For the BICEP/Keck-only case, we find $\beta_{\text{HI}} = 1.36^{+0.14}_{-0.17}$. For the Planck-only case, we find $\beta_{\text{HI}} = 2.26^{+0.32}_{-0.54}$. Using similar frequencies to the Planck-only case but replacing Planck's 143 and 217 GHz bands with the 150 and 220 GHz bands

of BICEP/Keck, we find $\beta_{\text{HI}} = 1.69 \pm 0.15$. Finally, we also plot the posterior using all the frequency bands, for which $\beta_{\text{HI}} = 1.54 \pm 0.13$, with the same color as in Figures 3.8 and 3.9 for comparison.

Note that, although two of the cases cover approximately the same frequency range, the Planck-only case has a wider posterior that is shifted slightly toward higher values of β_{HI} . This is because Planck's 143 and 217 GHz bands are not very sensitive to filamentary dust polarization when restricted to the BICEP/Keck region as compared to BICEP/Keck's 150 and 220 GHz bands. That said, the four posteriors are statistically consistent with each other to within 2σ . The results are qualitatively similar when fitting E modes and B modes separately.

Finally, we also calculate the correlation ratio as a function of multipole ℓ between BICEP/Keck or Planck data and V1 with RHT parameters from Equation 3.13. The correlation ratio is defined as

$$\rho_{\ell}^{\text{data} \times \text{HI}} = \frac{D_{\ell}^{\text{data} \times \text{HI}}}{\sqrt{D_{\ell}^{\text{data} \times \text{data}} \times D_{\ell}^{\text{HI} \times \text{HI}}}}. \quad (3.17)$$

The autospectra in the denominator contain noise biases. It would be possible to debias, but this would change the interpretation of the resulting correlation ratio. With noise debiasing, the correlation ratio would reflect the fraction of the sky signal that is accounted for by the HI morphology template. Without noise debiasing, as in Equation 3.17, the correlation ratio reflects the fraction of the data (including noise) that is accounted for by the HI morphology

template. For the purposes of forecasting sensitivity to r , we wish to retain the diluting effects of noise.

We plot the results in Figure 3.11. The error bars show the 1σ deviation of the correlation of 499 realizations of lensed- Λ CDM, GD, and noise with V1. Comparing BICEP/Keck data points with Planck bands of similar frequencies, we note that the BICEP/Keck bands correlate better in B modes with the H_I morphology template in this region. Also, the BICEP/Keck 220 GHz data is only slightly less correlated with V1 in EE but much more correlated in BB than the Planck 353 GHz data. This is consistent with the dust sensitivity estimates from BK18 that show that the BICEP/Keck 220 GHz data is more sensitive to dust than the Planck 353 GHz data (Figure 6 of BK18). The correlation ratio is larger in B modes than that in E modes for BICEP/Keck bands due to the CMB sample variance at lower frequencies. For a direct comparison of the error bars between BICEP/Keck and Planck bands of similar frequencies, we plot the numerator of the correlation ratio ρ_ℓ , i.e., the cross spectra D_ℓ in Figure 3.12. The error bars are clearly smaller for the BICEP/Keck bands, especially in B modes.

3.6.5 Polarized Dust in Magellanic Stream I

Passing through the BICEP/Keck region is a stream of high-velocity gas, known as Magellanic Stream 1 (Westmeier, 2018). The metallicity and abundance measurements of the Magellanic Stream are consistent with an origin in the SMC, created by a gravitational tug

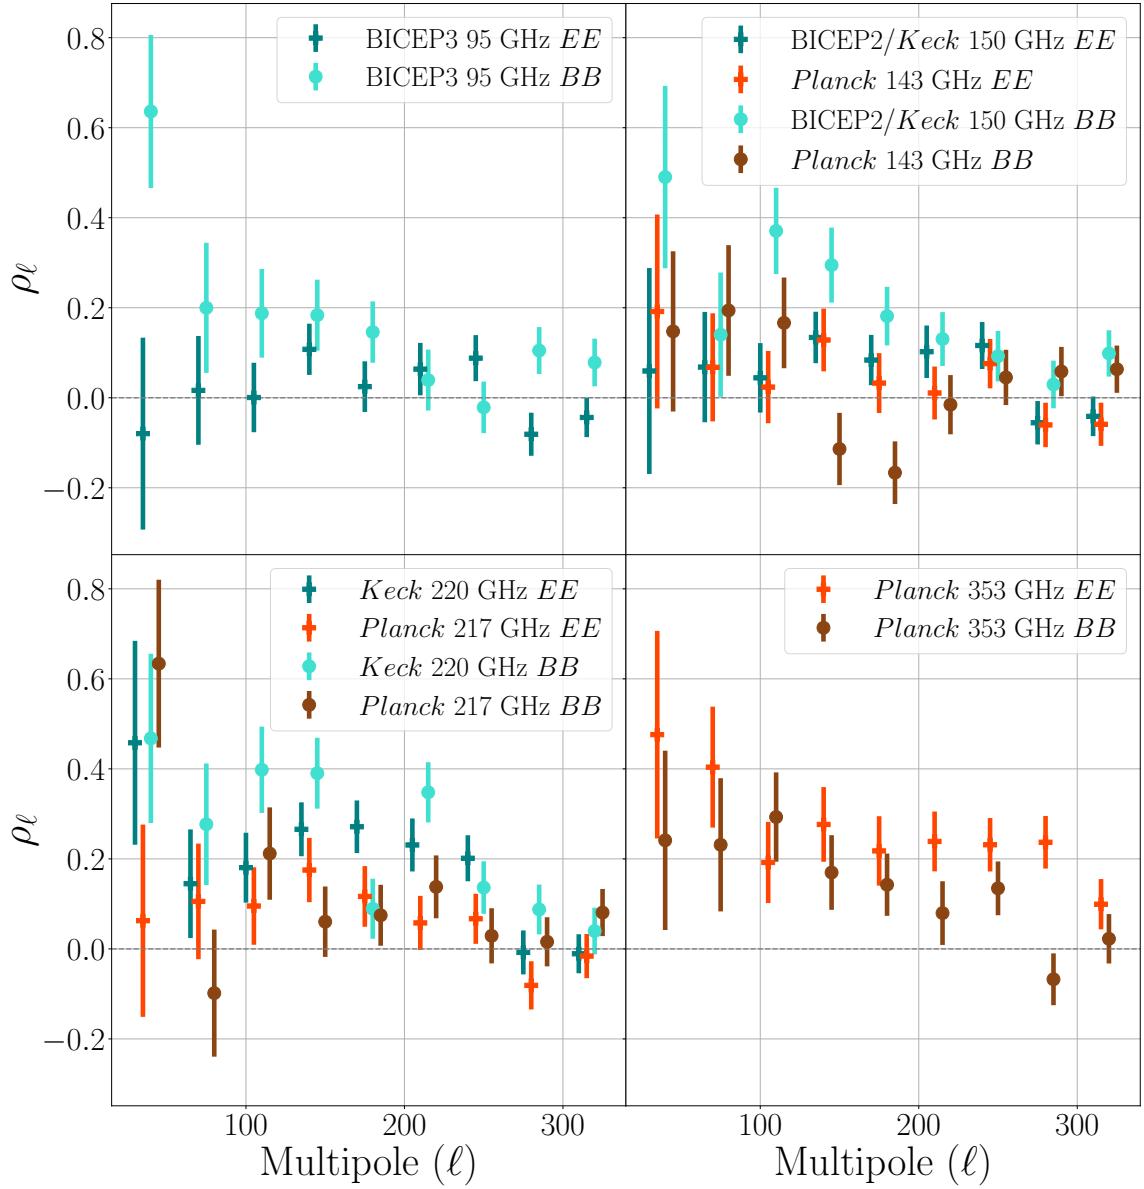


Figure 3.11: EE (cross) and BB (circle) unitless correlation ratios as a function of multipole moment. The correlation ratios between V1 and Planck data with 1σ variations are shown in red and brown to compare them to the correlation ratios between V1 and BICEP/Keck data, which are shown in teal and turquoise. The errors are derived from spurious correlations between V1 and lensed- Λ CDM, Gaussian dust, and noise. Data points for similar frequencies between BICEP/Keck and Planck are plotted on the same panels for comparison.

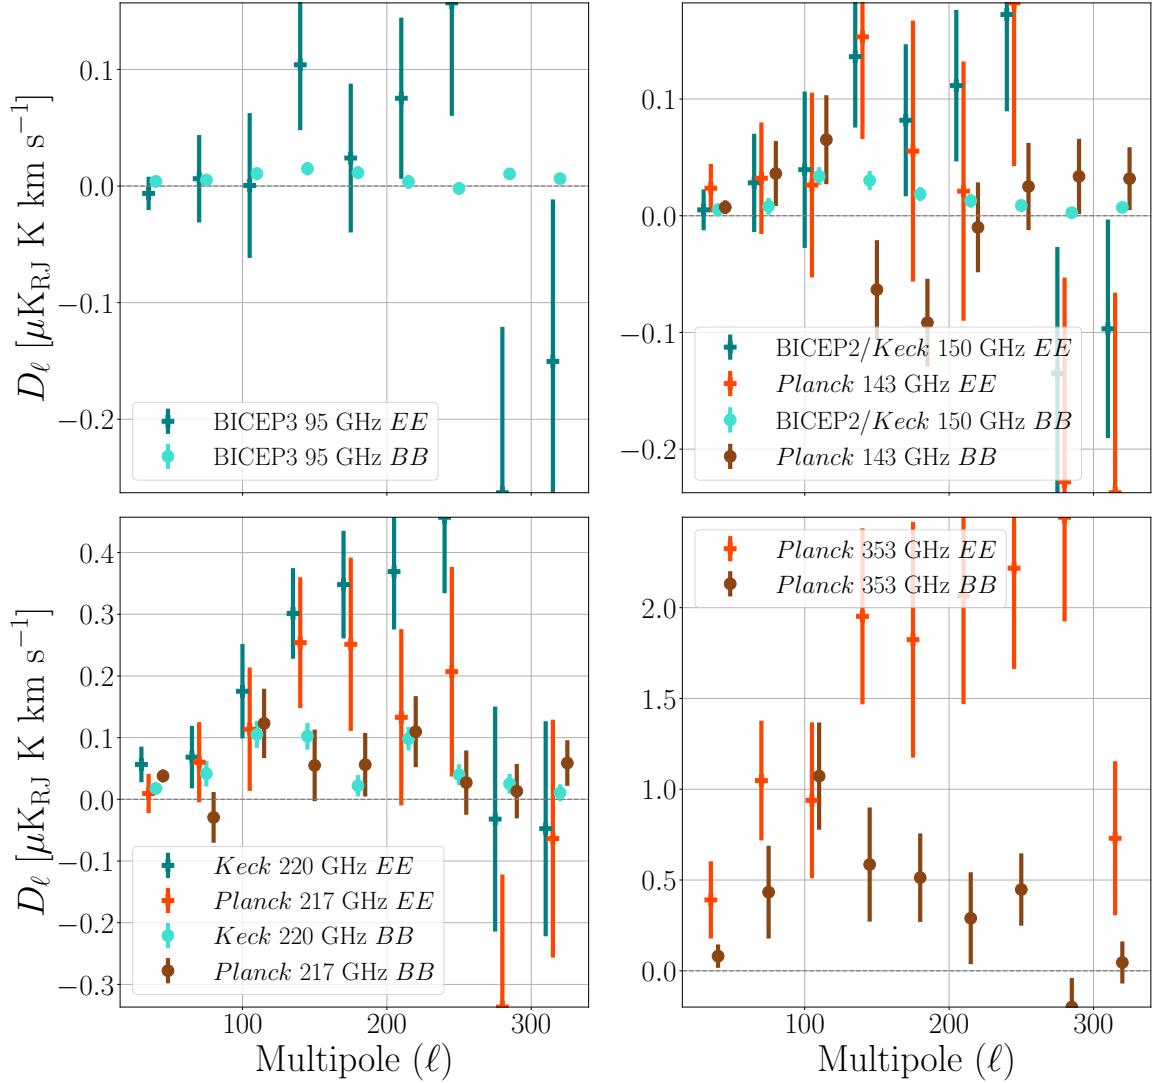


Figure 3.12: *EE* (cross) and *BB* (circle) cross spectra as a function of multipole moment. The cross spectra between V1 and Planck data with 1σ variations are shown in red and brown to compare them to the cross spectra between V1 and BICEP/Keck data, which are shown in teal and turquoise. The errors are derived from spurious correlations between V1 and lensed- Λ CDM, Gaussian dust, and noise. Data points for similar frequencies between BICEP/Keck and Planck are plotted on the same panels for comparison.

from the Large Magellanic Cloud (Fox et al., 2018). The Magellanic Stream and Clouds are part of the Magellanic System, along with the Magellanic Bridge and the Leading Arm

	V1	V2	V3	V2 + V3
<i>BB</i>	6.7	1.3	0.6	0.9
<i>EE</i>	14.6	2.4	1.4	2.5
<i>BB + EE</i>	16.1	1.6	1.2	0.3

Table 3.4: Comparison of the statistical significance of a detection of the cross correlation between H I morphology templates and the dust polarization in units of equivalent Gaussian standard deviations for V1, V2, and V3. We also add a column for V2 + V3, both of which are associated with Magellanic Stream 1. The 95, 150, and 220 GHz bands of BICEP/Keck and the 353 GHz band of Planck are used here.

(see, e.g., D’Onghia & Fox, 2016, for a review). The nature of dust in the Magellanic Stream is not well constrained by observations. Measurements of the gas-to-dust ratio in the Magellanic Clouds indicate a much lower dust content than that in the Milky Way (Fong et al., 1987; Richter, 2000; Tumlinson et al., 2002).

However, there is good reason to believe that the Stream may contain some dust since the same processes that inject metals, such as Mg II and Fe II, into clouds should also inject dust (Benjamin, 2005; Wakker, 2001). Constraints on the dust content of the Magellanic Stream can thus have important implications for dust survival and destruction in the Stream environment. Although efforts to detect dust emission from the Magellanic Stream have not yielded positive results yet in intensity or reddening (Wakker & Boulanger, 1986; Lenz et al., 2017), we test whether we can detect it in polarization, assuming the dust there is polarized due to a coherent magnetic field. While not yet directly detected in the Magellanic Stream, a coherent magnetic field is plausible given the detections in other tidal features and in the Magellanic Bridge using Faraday rotation measurements (Kaczmarek et al., 2017).

Using the $\Delta\chi^2$ statistical test defined in Section 3.4.6, we find no statistically significant

correlation with the second and third velocity components, V2 and V3. The correlation metric does not exceed $\sim 2.5\sigma$ for any of the choices in Table 3.4, including the addition of V2 and V3. This is also true for all the different variations of RHT parameters we tried.

We also try looking for a correlation in total intensity (TT) between BICEP/Keck or Planck T and V2 or V3 T (the H I intensity integrated over the V2 and V3 velocity ranges) and find no correlation. Since the standard BICEP/Keck simulations constrain T to the well-measured Planck T map, we only use that one realization for computational simplicity and only look for a visual correlation rather than making statistical inferences. Furthermore, adding V2, V3, or both to V1 decreases the TT correlation with BICEP/Keck and Planck.

We therefore do not detect evidence for dust in Magellanic Stream 1. The Magellanic Stream's distance may limit our sensitivity to resolving the local magnetic field orientations there because structures on the plane of the sky of the same angular scale as the Galactic gas correspond to much larger structures at the distance of the Magellanic Stream. The Stream's distance is fairly uncertain. Lucchini et al. (2021) recently estimated it to be ~ 20 kpc away from the Sun at its closest point through the use of simulations. For comparison, at these high-Galactic latitudes, the dust associated with V1 is likely at a distance of order 100 pc (e.g., Vergely et al., 2022b; Guo et al., 2021). Furthermore, our analysis is restricted to only the section of the Stream that intersects the BICEP/Keck region. Extending the sky area to include the entire Stream, running the RHT on forthcoming H I emission data from the Galactic Australian Square Kilometre Array Pathfinder (GASKAP) Survey (Dickey et al.,

2013) with $30''$ angular resolution, and using higher angular resolution dust polarization data (CCAT-Prime collaboration et al., 2021; Hensley et al., 2022; CMB-S4 collaboration et al., 2022) are all possible extensions of this work that can improve the sensitivity of this method for detecting or setting limits on dust polarization from the Stream.

3.7 Summary and Outlook

We characterize the filamentary dust polarization in the BICEP/Keck observing region through correlations with template maps based on measurements of H I. A detection of primordial gravitational wave B modes depends on reliable component separation because the polarized dust emission is the dominant foreground at frequencies $\gtrapprox 70$ GHz (Dunkley et al., 2009; Planck Collaboration et al., 2016d) and has a higher amplitude than that of the polarization associated with primordial gravitational waves (Flauger et al., 2014; BICEP2/Keck Collaboration et al., 2015; Errard et al., 2016). Therefore, polarized dust emission must be characterized to great accuracy and precision. We concentrate on the BICEP/Keck region as a test case for the diffuse high-Galactic latitude sky with deep data across several frequencies.

We summarize the conclusions of this work below.

- We separate the H I emission in the BICEP/Keck region into three distinct velocity components that together account for the bulk of the polarized intensity in the H I morphology template. One is associated with the Milky Way, while the other two are

associated with Magellanic Stream I.

- We explore the RHT parameter space to increase the correlation with BICEP/Keck and Planck by $\sim 2\sigma$ in *BB* and $\sim 3\sigma$ in *EE* and *EE+BB* with respect to the parameters used in Clark et al. (2019). The parameters we recommend using on Hi4PI data in the BICEP/Keck region for producing H I morphology templates are $D_W = 135'$, $\theta_{\text{FWHM}} = 4'$, and $Z = 0.75$.
- Using polarization data from BICEP/Keck and Planck, we find a statistically significant detection of filamentary dust polarization in the Galactic component of H I at $\sim 7\sigma$ in *BB*, $\sim 15\sigma$ in *EE*, and $\sim 16\sigma$ in *EE+BB*.
- We show that the overwhelming majority of the contribution comes from the LVC velocity range, $-12 \text{ km s}^{-1} < v_{\text{lsr}} < 10 \text{ km s}^{-1}$, and find no evidence of frequency decorrelation in the BICEP/Keck region as defined in Pelgrims et al. (2021a). The inclusion of the IVC component to the line-of-sight sum affects the correlation by $\lesssim 0.1\sigma$ in *BB*, $\lesssim 0.3\sigma$ in *EE*, and $\lesssim 0.2\sigma$ in *EE + BB*. We note that the dust structure associated with H I kinematic substructure within the LVC range could still produce frequency decorrelation.
- We fit an SED with $\hat{\beta}_{\text{HI}} = 1.54 \pm 0.13$ in the BICEP/Keck region for the filamentary dust polarization component associated with the Galactic component. This is consistent with the SED fit in BK18 for the total dust component in the BICEP/Keck

region. The similarity between the SED of the filamentary contributions to the dust polarization and the SED of the rest of the dust field indicates that there is no evidence for decorrelation between the filamentary dust and the rest of the dust column in the BICEP/Keck region.

- We present the first multifrequency detection of filamentary dust polarization in cross-correlation with H_I filaments down to 95 GHz. We show that the 95 GHz band of BICEP3 is more sensitive than any Planck band to the *B*-mode correlation in the BICEP/Keck region, providing a low-frequency lever arm for the dust SED. We also find that, at low frequencies, the brightness of the CMB in *E* modes limits our sensitivity but that the correlation could improve in *B* modes with more data. As a consistency check, we also omit certain frequency bands in the multifrequency correlations to compare the contribution of the different bands to our measurements.
- We do not find evidence for dust polarization in the higher-velocity H_I components associated with Magellanic Stream 1. This confirmation is important for future CMB observations whose field-of-view intercepts the Magellanic Stream.

In addition to facilitating foreground removal for *B*-mode cosmology, this type of H_I-based characterization of the dust polarization can also be a method for removing the Milky Way foreground contribution for studies of the Magellanic Clouds in dust polarization. Such a study is planned with CCAT-prime (CCAT-Prime collaboration et al., 2021).

3.8 Acknowledgments

We thank the anonymous referee for a thoughtful review.

This work was supported by the National Science Foundation under grant No. AST-2106607.

The BICEP/Keck projects have been made possible through a series of grants from the National Science Foundation including 0742818, 0742592, 1044978, 1110087, 1145172, 1145143, 1145248, 1639040, 1638957, 1638978, and 1638970, and by the Keck Foundation. We thank the staff of the U.S. Antarctic Program and in particular the South Pole Station and the heroic winter-overs without whose help this research would not have been possible. We also thank all those who have contributed past efforts to the BICEP/Keck series of experiments.

This publication utilizes data from Planck, an ESA science mission funded by ESA Member States, NASA, and Canada.

This work makes use of data from the HI4PI Survey, which is constructed from the Effelsberg-Bonn HI Survey (EBHIS), made with the 100 m radio telescope of the MPIfR at Effelsberg/Germany, and the Galactic All-Sky Survey (GASS), observed with the Parkes Radio Telescope, part of the Australia Telescope National Facility, which is funded by the Australian Government for operation as a National Facility managed by CSIRO. EBHIS

was funded by the Deutsche Forschungsgemeinschaft (DFG) under the grants KE757/7-1 to 7-3.

The computations in this paper were run on the Sherlock cluster, supported by the Stanford Research Computing Center at Stanford University, and on the Odyssey/Cannon cluster, supported by the FAS Science Division Research Computing Group at Harvard University.

3.9 Uncertainty Calculation

To measure the uncertainty on the best-fit a , k , and β_{HI} values, we construct a simulation set of 499 filtered dust realization Stokes Q/U maps as

$$\tilde{m}_\nu^{\text{dust}}(\hat{\mathbf{n}}, \alpha) \equiv \alpha \cdot f_\nu(\beta_{\text{GD}}) \cdot \tilde{m}_\nu^{\text{GD}}(\hat{\mathbf{n}}) + \hat{k} \cdot f_\nu(\hat{\beta}_{\text{HI}}) \cdot \tilde{m}_\nu^{\text{HI}}(\hat{\mathbf{n}}), \quad (3.18)$$

where f_ν is a modified blackbody scaling law with a fixed temperature, $T = 19.6$ K, as in Section 3.4.3, \hat{k} and $\hat{\beta}_{\text{HI}}$ are the best-fit results from the real data, and $\tilde{m}_\nu^{\text{HI}}$ is the result of applying the transfer function defined in Section 3.4.2 in harmonic space to $\tilde{m}_\nu^{\text{HI}}$ and then inverse transforming back to map space. The free parameter α is chosen such that

$$\overline{\mathbf{D}}^{\text{dust} \times \text{dust}}(\alpha) = f_\nu^2(\beta_{\text{GD}}) \overline{\mathbf{D}}^{\text{GD} \times \text{GD}}, \quad (3.19)$$

where $\bar{\mathbf{D}}$ is the mean over realizations of the vector of autospectra over EE , BB , and multipole bins. One frequency, 353 GHz, is sufficient for the fit here.

Therefore, we fit for α using a Gaussian likelihood approximation, i.e. a χ^2 -minimization

$$-2 \log \mathcal{L} = \left(\hat{\mathbf{S}}(\alpha) - \mathbf{S} \right)^T \mathbf{Z}^{-1} \left(\hat{\mathbf{S}}(\alpha) - \mathbf{S} \right), \quad (3.20)$$

where, from Equation 3.19,

$$\begin{aligned} \hat{\mathbf{S}}(\alpha) - \mathbf{S} = & \quad (\alpha^2 f_{353 \text{ GHz}}^2(\beta_{\text{GD}}) - 1) \cdot \bar{\mathbf{D}}^{\text{GD} \times \text{GD}} + \\ & + \hat{k}^2 \cdot f_{353 \text{ GHz}}^2(\hat{\beta}_{\text{HI}}) \cdot \mathbf{D}^{\tilde{\text{HI}} \times \tilde{\text{HI}}} + \\ & + 2 \cdot \alpha \cdot \hat{k} \cdot f_{353 \text{ GHz}}(\hat{\beta}_{\text{HI}}) \cdot f_{353 \text{ GHz}}(\beta_{\text{GD}}) \cdot \bar{\mathbf{D}}^{\text{GD} \times \tilde{\text{HI}}} \end{aligned} \quad (3.21)$$

and \mathbf{Z} is the covariance matrix due to variations in the GD.

After fitting α , we define

$$d_\nu = \tilde{m}_\nu^{\Lambda\text{CDM}}(\hat{\mathbf{n}}) + \tilde{m}_\nu^n(\hat{\mathbf{n}}) + \tilde{m}_\nu^{\text{dust}}(\hat{\mathbf{n}}, \hat{\alpha}), \quad (3.22)$$

where $\hat{\alpha}$ is the best-fit value, and repeat the process in Section 3.4.6, replacing \mathbf{D}^{real} with the cross spectra of d_ν with the H I morphology template.

Expecting the fits for α , k , and β_{HI} to yield the inputs $\hat{\alpha}$, \hat{k} , and $\hat{\beta}_{\text{HI}}$, we use the spread of the best-fit distributions for the 499 realizations to calculate the uncertainty on our fitting

method for a , k , and β_{HI} , respectively. An example of this is shown in Figure 3.2 and described in Section 3.4.7.

3.10 Analysis Variations

For the main results presented in Tables 3.1, 3.2, 3.3, and 3.4, we condition the covariance matrix and use a transfer function and a modified blackbody scaling for the H I morphology template. In this appendix, we present those same results for different variations of those choices in Tables 3.5, 3.6, 3.7, 3.8, respectively. The main results are shown in the bolded columns of these tables. Note that the results are not qualitatively affected by these variations.

Covariance Matrix:		Not Conditioned				Conditioned			
Frequency Scaling:		Power Law		Modified Blackbody		Power Law		Modified Blackbody	
Transfer Function:		Used	Not Used	Used	Not Used	Used	Not Used	Used	Not Used
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<i>BB</i>	best	7.0	6.3	6.8	6.2	6.9	6.2	6.7	6.2
	default	4.8	3.9	4.7	3.9	4.7	4.0	4.7	4.0
<i>EE</i>	best	15.2	14.9	15.2	14.9	14.6	14.2	14.6	14.2
	default	12.2	10.7	12.2	10.7	12.3	10.8	12.3	10.8
<i>BB + EE</i>	best	17.2	16.6	17.1	16.5	16.2	15.7	16.1	15.6
	default	13.8	12.3	13.8	12.3	12.9	11.6	12.9	11.6

Table 3.5: Statistical significance of the detection of V1 in units of equivalent Gaussian standard deviations as defined in Section 3.4.6 using the 95, 150, and 220 GHz bands of BICEP/Keck and the 353 GHz band of Planck. The rows labeled "best" use the parameters $D_W = 135'$, $\theta_{\text{FWHM}} = 4'$, and $Z = 0.75$, and the rows labeled "default" use the parameters $D_W = 75'$, $\theta_{\text{FWHM}} = 30'$, and $Z = 0.7$, which are used in Clark & Hensley (2019). The bolded column (9) shows the main results.

Covariance Matrix:		Not Conditioned						Conditioned					
		Power Law		Modified Blackbody		Power Law		Modified Blackbody		Used		Not Used	
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(7)	(8)	(9)	(10)
		BB	LVCs + IVCs	7.0	6.3	6.8	6.2	6.9	6.2	6.7	6.2	6.8	6.2
EE	LVCs	LVCs	LVCs	7.1	6.4	7.0	6.3	7.0	6.2	6.7	6.2	6.8	6.2
		EE	LVCs + IVCs	15.2	14.9	15.2	14.9	14.6	14.2	14.6	14.2	14.3	14.2
BB + EE	LVCs + IVCs	LVCs	LVCs	14.9	14.8	14.9	14.7	14.4	14.2	14.6	14.2	14.3	14.2
		BB + EE	LVCs + IVCs	17.2	16.6	17.1	16.5	16.2	15.7	16.1	15.6	16.1	15.6
		LVCs	LVCs	16.9	16.4	16.9	16.3	16.2	15.7	16.1	15.6	16.1	15.6

Table 3.6: Comparison of the statistical significance of a detection of the cross correlation with the 95, 150, and 220 GHz bands of BICEP/Keck and the 353 GHz band of Planck in units of equivalent Gaussian standard deviations when including the channels in the IVC velocity range in the line-of-sight sum. The bolded column (9) shows the main results.

Covariance Matrix:			Not Conditioned						Conditioned					
Frequency Scaling:			Power Law			Modified Blackbody			Power Law			Modified Blackbody		
Transfer Function:			Used	Not Used	Used	Not Used	Used	Not Used	Used	Not Used	Used	Not Used	Used	Not Used
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)					
<i>BB</i>	BICEP3 95 GHz	4.44	3.84	4.44	3.84	4.53	3.98	4.53	3.98					
	Planck 143 GHz	0.16	0.36	0.16	0.36	0.05	0.40	0.05	0.40					
	BICEP2/Keck 150 GHz	5.13	4.76	5.13	4.76	5.31	4.97	5.31	4.97					
	Planck 217 GHz	3.50	3.31	3.50	3.31	3.50	3.14	3.50	3.14					
	Keck 220 GHz	5.90	5.65	5.90	5.65	5.82	5.60	5.82	5.60					
	Planck 353 GHz	3.18	2.53	3.18	2.53	3.18	2.60	3.18	2.60					
<i>EE</i>	BICEP3 95 GHz	1.21	1.18	1.21	1.18	1.22	1.25	1.22	1.25					
	Planck 143 GHz	0.76	0.20	0.76	0.20	0.72	0.20	0.72	0.20					
	BICEP2/Keck 150 GHz	2.42	2.00	2.42	2.00	2.43	2.01	2.43	2.01					
	Planck 217 GHz	2.28	1.52	2.28	1.52	2.37	1.62	2.37	1.62					
	Keck 220 GHz	7.35	6.79	7.35	6.79	7.13	6.61	7.13	6.61					
	Planck 353 GHz	7.92	8.13	7.92	8.13	7.99	8.12	7.99	8.12					
<i>BB + EE</i>	BICEP3 95 GHz	4.61	3.93	4.61	3.93	4.72	4.05	4.72	4.05					
	Planck 143 GHz	0.13	1.32	0.13	1.32	0.12	1.55	0.12	1.55					
	BICEP2/Keck 150 GHz	5.83	5.32	5.83	5.32	5.98	5.49	5.98	5.49					
	Planck 217 GHz	3.72	2.80	3.72	2.80	4.02	3.05	4.02	3.05					
	Keck 220 GHz	9.02	8.46	9.02	8.46	9.26	8.76	9.26	8.76					
	Planck 353 GHz	8.65	8.59	8.65	8.59	8.59	8.57	8.59	8.57					

Table 3.7: The detection statistical significance of the cross correlation between H I morphology templates and the dust polarization in units of equivalent Gaussian standard deviations as defined in Section 3.4.6. The bolded column (9) shows the main results.

Covariance Matrix:			Not Conditioned						Conditioned		
Frequency Scaling:		Power Law		Modified Blackbody		Power Law		Modified Blackbody		Not Used	
Transfer Function:	Used	Not Used	Used	Not Used	Used	Not Used	Used	Not Used	Used	Not Used	(10)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)		
<i>BB</i>	V1	7.0	6.3	6.8	6.2	6.9	6.2	6.7	6.2		
	V2	1.0	1.0	1.1	1.1	1.1	1.3	1.3	1.4		
	V3	0.8	0.2	0.8	0.2	0.7	0.1	0.6	0.1		
	V2 + V3	0.7	0.5	0.8	0.5	0.9	0.7	0.9	0.7		
<i>EE</i>	V1	15.2	14.9	15.2	14.9	14.6	14.2	14.6	14.2		
	V2	2.4	2.3	2.4	2.4	2.3	2.3	2.4	2.3		
	V3	1.4	1.1	1.4	0.8	1.5	0.8	1.4	1.1		
	V2 + V3	1.6	0.5	1.6	0.2	2.5	2.3	2.5	2.3		
<i>BB + EE</i>	V1	17.2	16.6	17.1	16.5	16.2	15.7	16.1	15.6		
	V2	2.0	2.2	2.1	2.3	1.5	2.1	1.6	2.1		
	V3	1.1	1.1	1.2	1.1	1.2	1.1	1.2	1.1		
	V2 + V3	0.7	0.2	0.7	0.2	0.4	0.2	0.3	0.1		

Table 3.8: Comparison of the statistical significance of a detection of the cross correlation between H I morphology templates and the dust polarization in units of equivalent Gaussian standard deviations for V1, V2, and V3. We also add a column for V2 + V3, both of which are associated with Magellanic Stream 1. The 95, 150, and 220 GHz bands of BICEP/Keck and the 353 GHz band of Planck are used here. The bolded column (9) shows the main results. The other columns show the results for different variations of our model.

Chapter 4

Filamentary Dust Polarization and the Morphology of Neutral Hydrogen Structures

Abstract

Filamentary structures in neutral hydrogen (HI) emission are well aligned with the interstellar magnetic field, so HI emission morphology can be used to construct templates that strongly correlate with measurements of polarized thermal dust emission. We explore how the quantification of filament morphology affects this correlation. We introduce a new implementation of the Rolling Hough Transform (RHT) using spherical harmonic convolutions, which enables efficient quantification of filamentary structure on the sphere. We use this Spherical RHT algorithm along with a Hessian-based

method to construct HI-based polarization templates. We discuss improvements to each algorithm relative to similar implementations in the literature and compare their outputs. By exploring the parameter space of filament morphologies with the Spherical RHT, we find that the most informative HI structures for modeling the magnetic field structure are the thinnest resolved filaments. For this reason, we find a $\sim 10\%$ enhancement in the *B*-mode correlation with polarized dust emission with higher-resolution HI observations. We demonstrate that certain interstellar morphologies can produce parity-violating signatures, i.e., nonzero *TB* and *EB*, even under the assumption that filaments are locally aligned with the magnetic field. Finally, we demonstrate that *B* modes from interstellar dust filaments are mostly affected by the topology of the filaments with respect to one another and their relative polarized intensities, whereas *E* modes are mostly sensitive to the shapes of individual filaments.

4.1 Paper Status and External Contributions

This chapter is based on the article published by The Astrophysical Journal, Volume 961, Issue 1 under the title, "Filamentary Dust Polarization and the Morphology of Neutral Hydrogen Structures" (Halal et al., 2024a). I am the lead author of this paper. I performed all of the analysis, wrote all of the text, and produced all of the figures for this paper. However, this work is the result of weekly discussions and advising from postdoctoral scholars Ari Cukierman and Dominic Beck and my co-advisors Susan Clark and Chao-Lin

Kuo. I have also received extensive editorial input from Susan Clark, Ari Cukierman, and Dominic Beck.

Part of the work presented in this chapter utilizes methodology described in Cukierman et al. (2023), which I am a co-author on. Cukierman et al. (2023) present evidence for a scale-independent misalignment of interstellar dust filaments and magnetic fields. We use HI-based polarization templates I developed, which are described in this chapter. I also ran some tests to confirm some of the conclusions in Cukierman et al. (2023).

4.2 Motivation

Modeling polarized dust emission is crucial for studying various astrophysical phenomena in the interstellar medium (ISM) and for analyzing the polarization of the cosmic microwave background (CMB). Aspherical rotating dust grains preferentially align their short axes with the local magnetic field, resulting in their thermal emission being linearly polarized (Purcell, 1975). Dust polarization thus traces the plane-of-sky magnetic field orientation and is widely used to trace magnetic field structure in the Galaxy (e.g., Han, 2017). At large scales and frequencies greater than approximately 70 GHz, polarized dust emission is the predominant polarized CMB foreground (Planck Collaboration et al., 2016e). The accurate modeling and elimination of the dust contribution to CMB polarization measurements are essential to search for an excess polarization signal induced by primordial gravitational waves (Kamionkowski et al., 1997; Seljak & Zaldarriaga, 1997; Seljak, 1997).

Galactic neutral hydrogen (H_I) emission is a tracer of the neutral medium that can be fruitfully compared to the dust distribution. H_I and dust trace similar volumes of the diffuse ISM (Boulanger et al., 1996b; Lenz et al., 2017). Much of the diffuse H_I emission is organized into filamentary structures that show significant alignment with the plane-of-sky magnetic field orientation (Clark et al., 2014, 2015). The spectroscopic nature of H_I measurements means that these structures can be studied in 3D, namely as a function of longitude, latitude, and radial velocity with respect to the local standard of rest v_{lsr} , i.e., the Doppler-shifted frequency of the 21-cm line (Clark, 2018). Furthermore, because H_I and broadband thermal dust emission are independently observed, cross correlations between the two are free from correlated telescope systematics. H_I data are also not contaminated by the cosmic infrared background (Chiang & Ménard, 2019).

Using these insights, Clark & Hensley (2019) developed a model of polarized dust emission based solely on H_I measurements. Cross correlations between this dust polarization model and millimeter-wave polarization data have proven useful for characterizing dust properties such as the spectral index (BICEP/Keck Collaboration et al., 2023e). Clark & Hensley (2019) used the Rolling Hough Transform (RHT; Clark et al., 2014, 2020) algorithm as a first step for quantifying the orientations of linear dust filaments. The RHT has free parameters that set the scale and shape of the identified filaments. This is ideal for the exploration of different filament morphologies and their polarization effects.

The RHT algorithm runs on images or flat-sky projections of small patches of the sky.

It is possible to construct an H_I-based polarization template on the full sky by projecting a small patch of the spherical map around each pixel to an image, running the algorithm, and projecting the result back to the sphere, as done in Clark & Hensley (2019) for a single set of parameters. However, this is computationally expensive to perform for multiple sets of parameters. In this paper, we develop an algorithm for running the RHT directly on the sphere using spherical harmonic convolutions.

Another filament-finding algorithm that can be used to construct polarization templates from H_I emission is based on the Hessian matrix (e.g., Cukierman et al., 2023). We explore the advantages and disadvantages of the Hessian-based algorithm relative to the RHT-based algorithm.

The paper is organized as follows. We introduce the dust emission and H_I data used in this work in Section 4.3. We explore how different modifications to the Hessian-based polarization template affect the correlation with polarized dust emission in Section 4.5. We introduce a spherical convolution version of the RHT in Section 4.6. We use it to explore the polarization effects of different filament morphologies and how filament-finding algorithms can be used for determining morphologies that produce parity-violating polarization signatures in Section 4.7. We summarize and conclude in Section 4.8.

4.3 Data

4.3.1 Dust Emission

We make use of two sets of Stokes I , Q , and U Planck data products at 353 GHz provided by the Planck Legacy Archive¹. The first is the set of Planck Commander dust maps with an angular resolution of 5', constructed by component separation applied to the Planck frequency maps (Planck Collaboration et al., 2020f). The second is the set of 353 GHz maps from Planck data release R3.01 with an angular resolution of 5' (Planck Collaboration et al., 2020d). While the former is processed to remove emission other than dust, the latter contains contributions from multiple components. We compare these two data products in Section 4.5.2 and use the Commander dust maps for all subsequent analyses. The reported results are insensitive to smoothing the Planck data, so we use them at their native resolution.

For cross-spectrum calculations between these maps and other polarization data products in this paper, we use the full-mission maps. When calculating autospectra of these maps, we compute cross spectra of the half-mission splits to avoid noise bias.

For most of the analysis in this paper, we use the Planck 70% sky fraction Galactic plane mask (Planck Collaboration et al., 2015). However, in Section 4.5.2, we also employ the 20%, 40%, 60%, and 80% sky fraction Galactic plane masks. The higher the sky fraction, the greater the contribution from lower Galactic latitudes.

¹pla.esac.esa.int

4.3.2 Neutral Hydrogen

For the Galactic neutral hydrogen (H I) emission, we use the H I 4π Survey (Hi4PI; HI4PI Collaboration et al., 2016), which has the highest-resolution full-sky measurements of the 21 cm hyperfine transition to date. Hi4PI merges data from the Effelsberg-Bonn H I Survey (EBHIS; Winkel et al., 2016) and the Parkes Galactic All-Sky Survey (GASS; McClure-Griffiths et al., 2009) to achieve an angular resolution of $16.^{\prime}2$, a spectral resolution of 1.49 km s^{-1} , and a normalized brightness temperature noise of $\sim 53 \text{ mK}$ for a 1 km s^{-1} velocity channel. We use the publicly available H I intensity data described in Clark & Hensley (2019), binned into velocity channels of equal integrated intensity in each pair of channels moving symmetrically outward from the local standard of rest.

In Section 4.5.3, we also use H I emission data from the Galactic Arecibo L-Band Feed Array H I Survey (GALFA-H I; Peek et al., 2018), which is higher resolution but only covers $\sim 32\%$ of the sky. GALFA-H I has an angular resolution of $4.^{\prime}1$, a spectral resolution of 0.184 km s^{-1} , and a normalized brightness temperature noise of 150 mK for a 1 km s^{-1} channel. We also use the publicly available H I intensity data described in Clark & Hensley (2019), binned into velocity channels of equal width of 3.7 km s^{-1} . These maps span the range $1.^{\circ}5 < \text{decl.} < 35.^{\circ}5$ to avoid telescope scan artifacts at the edges of the Arecibo declination range.

4.4 H I-based Dust Polarization Prediction

The H I-based polarization templates are constructed by measuring the orientation of linear structures to determine the polarization angle and combining this information with some weighting representing the polarized intensity at different locations in the map. The orientation can be determined by different algorithms applied to the H I intensity maps. In this section, we summarize the two algorithms we use in this paper for polarization angle determination and describe how their outputs are used along with different polarized intensity weighting schemes to construct H I-based polarization templates.

4.4.1 RHT-based Angle Determination

The RHT is a computer vision algorithm that identifies linear structures and their orientations in images (Clark et al., 2014, 2020). The steps involved in this process are:

1. Unsharp masking, which involves subtracting a version of the map smoothed to a given scale, θ_{FWHM} , from the original map. This step acts as a high-pass filter of the map to remove larger-scale emission.
2. Bit masking, which converts all pixels with negative values to zero and all pixels with positive values to one.
3. Applying the Hough transform (Hough, 1962) on a circular window of a given diameter, D_W , centered on each pixel to quantify the relative intensities of differently

oriented linear structures passing through that pixel.

4. Storing only the linear intensities over a certain threshold fraction, Z , of the window diameter. The output is stored as linear intensity as a function of orientation $R(\hat{\mathbf{n}}, \theta, v)$.

This algorithm, therefore, has three free parameters, θ_{FWHM} , D_{W} , and Z , which can be tuned to different values for different applications.

4.4.2 Hessian-based Angle Determination

Hessian-based filament identification has been applied to different maps, e.g., Planck 353 GHz total intensity maps (Planck Collaboration et al., 2016f,a), Hi4PI H_I intensity and Planck 857 GHz total intensity maps (Kalberla et al., 2021), Herschel images of molecular clouds (Polychroni et al., 2013), and simulations of the cosmic web (Colombi et al., 2000; Forero-Romero et al., 2009). In this work, we use the version of the Hessian-based filament-finding algorithm described in Cukierman et al. (2023).

The Hessian matrix serves as a tool to determine the orientation of filaments. It contains information about the local second derivatives. A negative curvature indicates the presence of at least one negative Hessian eigenvalue. By examining a map for areas exhibiting negative curvature, we can identify possible filaments.

We apply the Hessian to the H_I intensity maps in individual velocity bins I . We work in spherical coordinates with polar angle θ and azimuthal angle ϕ . The local Hessian matrix

is given by

$$H \equiv \begin{pmatrix} H_{xx} & H_{xy} \\ H_{yx} & H_{yy} \end{pmatrix}, \quad (4.1)$$

where

$$H_{xx} = \frac{\partial^2 I}{\partial \theta^2}, \quad (4.2)$$

$$H_{yy} = \frac{1}{\sin^2 \theta} \frac{\partial^2 I}{\partial \phi^2}, \quad (4.3)$$

$$H_{xy} = H_{yx} = -\frac{1}{\sin \theta} \frac{\partial^2 I}{\partial \phi \partial \theta}. \quad (4.4)$$

The eigenvalues are

$$\lambda_{\pm} = \frac{1}{2} (H_{xx} + H_{yy} \pm \alpha), \quad (4.5)$$

where

$$\alpha \equiv \sqrt{(H_{xx} - H_{yy})^2 + 4H_{xy}^2}. \quad (4.6)$$

For the local curvature to be negative along at least one axis, we require $\lambda_- < 0$.

We also require λ_- to be the larger of the two eigenvalues in magnitude such that this negative curvature is the dominant local morphology. In constructing HI-based polarization templates, we define a weighting w_H for each pixel at each velocity that is equal to λ_- when the eigenvalues satisfy our two requirements and equal to zero when they do not.

The orientation of the filaments is determined by the local eigenbasis. The polarization

angle is determined as

$$\theta_H = \arctan \left(\frac{H_{xx} - H_{yy} + \alpha}{2H_{xy}} \right). \quad (4.7)$$

4.4.3 H I-based Polarization Template Construction

The polarization angle determined using the RHT and Hessian algorithms can be combined with some weighting representing the local contribution to the H I-based polarized intensity to construct H I-based polarization templates. For instance, Clark & Hensley (2019) normalize the RHT-measured linear intensity $R(\hat{\mathbf{n}}, \theta, v)$ over different orientation bins such that

$$\sum_{\theta} R(\hat{\mathbf{n}}, \theta, v) = 1. \quad (4.8)$$

They use the normalized $R(\hat{\mathbf{n}}, \theta, v)$ and the H I intensity maps $I_{HI}(\hat{\mathbf{n}}, v)$ as the weighting to produce Stokes Q_{RHT} and U_{RHT} maps as

$$Q_{RHT}(\hat{\mathbf{n}}, v) = I_{HI}(\hat{\mathbf{n}}, v) \sum_{\theta} R(\hat{\mathbf{n}}, \theta, v) \cos 2\theta, \quad (4.9)$$

$$U_{RHT}(\hat{\mathbf{n}}, v) = I_{HI}(\hat{\mathbf{n}}, v) \sum_{\theta} R(\hat{\mathbf{n}}, \theta, v) \sin 2\theta. \quad (4.10)$$

These maps have the same units (K km s^{-1}) as the intensity maps, $I_{HI}(\hat{\mathbf{n}}, v)$. To construct the H I-based polarization template, the Stokes parameter maps are integrated over velocity

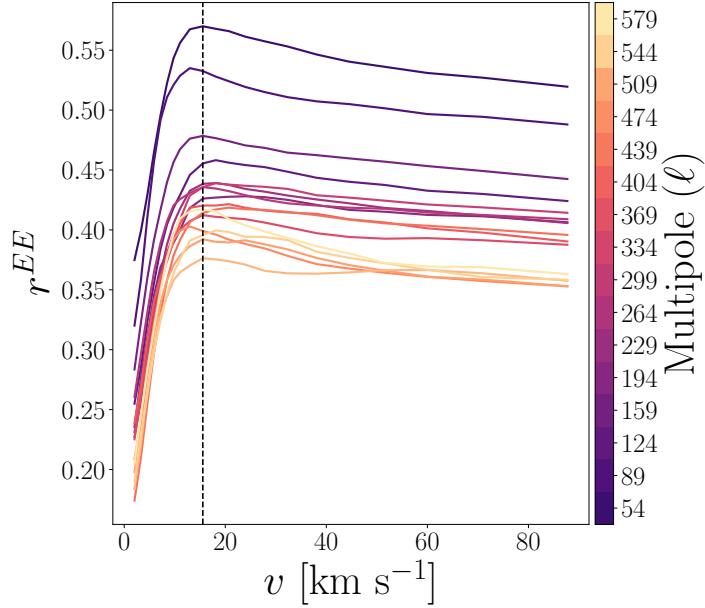


Figure 4.1: The EE correlation ratio on the Planck 70% sky fraction Galactic plane mask of the Planck Commander dust maps with $\text{H}\alpha$ -based polarization templates produced by the Hessian algorithm applied to $\text{H}\alpha$ 4PI data. The leftmost point is for the $\text{H}\alpha$ -based polarization template at 2.03 km s^{-1} , and each successive point on each curve corresponds to the addition of information from the two adjacent velocity channels in the positive and negative directions. The labels on the horizontal axis correspond to the velocity centers of the positive velocity channels being added. The different curves correspond to different multipole bins shown in the color bar. The vertical dashed line corresponds to the integrated $\text{H}\alpha$ -based polarization template over the velocity range $-13 \text{ km s}^{-1} < v_{\text{lsr}} < 16 \text{ km s}^{-1}$, after which the correlation saturates and starts decreasing over most of the multipole bins considered as information from more velocity channels is added.

channels as

$$Q_{\text{RHT}}(\hat{\mathbf{n}}) = \sum_v Q_{\text{RHT}}(\hat{\mathbf{n}}, v), \quad (4.11)$$

$$U_{\text{RHT}}(\hat{\mathbf{n}}) = \sum_v U_{\text{RHT}}(\hat{\mathbf{n}}, v). \quad (4.12)$$

The Hessian-based filament-finding algorithm (hereafter Hessian algorithm), by contrast, uses the local eigenbasis of the Hessian matrix to determine the orientation of linear

structures, and the negative eigenvalues to determine the Stokes weighting (Cukierman et al., 2023). See Section 4.4.2 for details. Equations 4.9 and 4.10, therefore, become

$$Q_H(\hat{\mathbf{n}}, v) = w_H(\hat{\mathbf{n}}, v) \cos 2\theta_H(\hat{\mathbf{n}}, v), \quad (4.13)$$

$$U_H(\hat{\mathbf{n}}, v) = w_H(\hat{\mathbf{n}}, v) \sin 2\theta_H(\hat{\mathbf{n}}, v), \quad (4.14)$$

where θ_H is the polarization angle perpendicular to the orientation of the local linear structure determined by the Hessian, and w_H is formed from the negative eigenvalues as described in Section 4.4.2. These maps are then summed over velocity as in Equations 4.11 and 4.12 to produce $Q_H(\hat{\mathbf{n}})$ and $U_H(\hat{\mathbf{n}})$.

Although work conducted across extensive portions of the high-Galactic latitude sky indicates that there could be a minor uniform misalignment between the filaments and the orientation of the magnetic field as measured by Planck, the angle of misalignment amounts to approximately $\sim 2^\circ$ - 5° only (Huffenberger et al., 2020; Clark et al., 2021; Cukierman et al., 2023). Rotating the H I-based polarization template angles by this amount to emulate this misalignment effect leads to only a slight enhancement in the correlation at the level of $\sim 0.1\%-0.5\%$ (Cukierman et al., 2023), and we do not apply this rotation here.

4.5 Improvements in HI-based Dust Polarization Prediction

In this section, we employ the Hessian algorithm on HI4PI and GALFA-HI data to construct HI-based polarization templates. Working within the Hessian-based framework, we examine how to construct templates that correlate most strongly with Planck polarized dust emission maps. We later contrast these templates with alternative maps based on the Spherical RHT.

4.5.1 Velocity Selection

Since we expect the correlation between HI and dust to vanish for high-velocity clouds (Wakker & Boulanger, 1986; Planck Collaboration et al., 2011b; Lenz et al., 2017), we restrict the velocity range over which we integrate the HI-based polarization template in Figure 4.1. More generally, any contribution to the HI-based polarization template from noise, data artifacts, or HI emission that is not correlated with dust structure will tend to decrease the measured correlation between the template and the polarized dust emission. To restrict the velocity range, we calculate the correlation ratio defined as

$$r_\ell^{\text{data}\times\text{HI}} = \frac{D_\ell^{\text{fm}\times\text{HI}}}{\sqrt{D_\ell^{\text{hm1}\times\text{hm2}} \times D_\ell^{\text{HI}\times\text{HI}}}}, \quad (4.15)$$

where $D_\ell^{m_1 \times m_2}$ is the cross-spectrum bandpower between two maps, m_1 and m_2 , in the multipole bin ℓ . All power spectra in this paper are computed with the pspy² code (Louis et al., 2020). We use the Planck Commander dust maps as the *data*, where the full-mission maps (*fm*) are used for the cross spectra with the H I-based polarization template and the half-mission splits (*hm1* and *hm2*) are used in the denominator. The *HI* in this equation refers to the H I-based polarization template, which in this case is constructed using the Hessian algorithm on the Hi4PI dataset.

The linear polarization field described by the Stokes *Q* and *U* maps can be decomposed into *E*-mode and *B*-mode components (Seljak & Zaldarriaga, 1997; Zaldarriaga, 2001). In Equation 4.15, *data*×*HI* can be the correlation of any combination of the *E*-mode, *B*-mode, or intensity components of the Planck Commander dust maps and the H I-based polarization templates. For example, in Figure 4.1, we calculate Equation 4.15 for their *E*-mode components. We do not show the *BB* correlation ratio because it exhibits similar behavior. We use the Planck 70% sky fraction Galactic plane mask for the correlation ratio calculations in this figure. We use the same multipole binning used in BICEP/Keck Collaboration et al. (2023e) for a direct comparison of results and because we did not find that multipoles higher than 600 provide additional insights. We start with the H I-based polarization template of the individual velocity channel at $v_{\text{lsr}} = 2.03 \text{ km s}^{-1}$. We pick this velocity channel because it has the highest H I intensity integrated over the unmasked sky. We plot the correlation ratio as a function of velocity integration range and spatial

²<https://github.com/simonsobs/pspy>

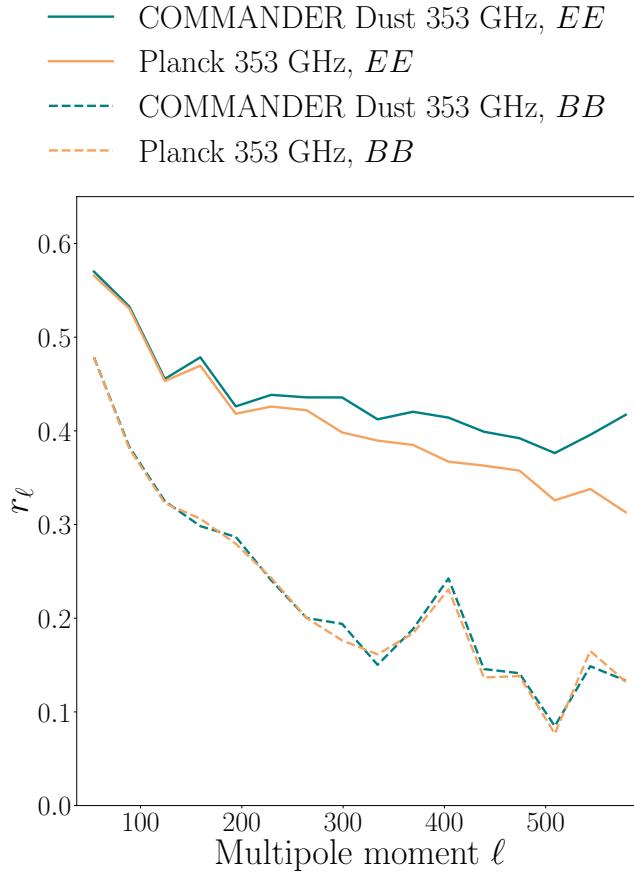


Figure 4.2: Comparison of the *EE* (solid) and *BB* (dashed) correlation ratios on the Planck 70% sky fraction Galactic plane mask of the Planck Commander dust maps (teal) and the Planck frequency maps (sandy brown) at 353 GHz with the H I-based polarization template constructed from applying the Hessian algorithm to H I4PI data at each velocity channel and integrating the resulting maps over the velocity range $-13 \text{ km s}^{-1} < v_{\text{lsr}} < 16 \text{ km s}^{-1}$.

scale in Figure 4.1. The horizontal axis on this figure is cumulative, i.e., moving toward higher velocities on these plots corresponds to symmetrically integrating outwards in the positive and negative directions from the starting velocity channel, adding one H I-based polarization template from each direction to the previous H I-based polarization template.

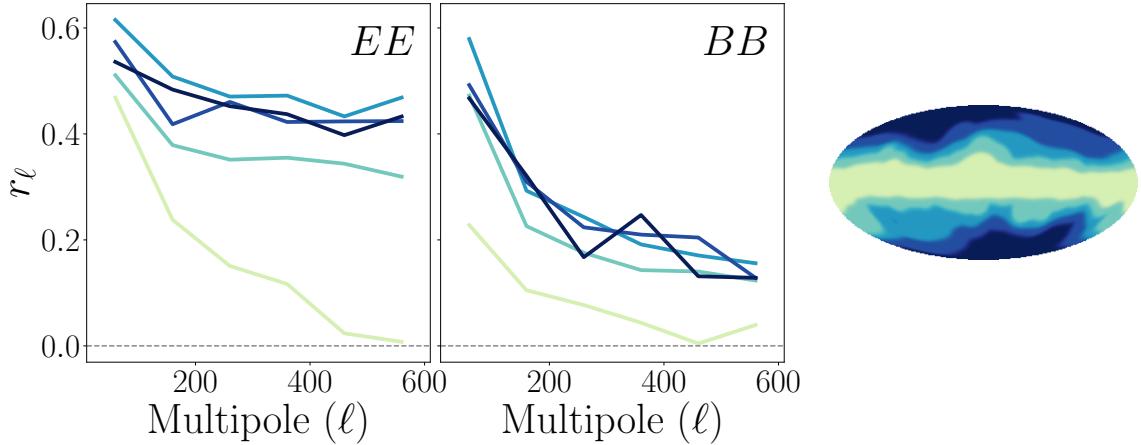


Figure 4.3: The *EE* (left) and *BB* (middle) correlation ratios of the Planck Commander dust maps with the H α -based polarization template constructed using the Hessian algorithm on the H α 4PI intensity maps over the different non-overlapping masks shown on the right. The masks are the Planck 20% sky fraction Galactic plane mask (darkest), the inverted Planck 80% sky fraction Galactic plane mask (lightest), and the differences between the Planck 20%, 40%, 60%, and 80% sky fraction Galactic plane masks (other shades of blue) shown in a Mollweide projection in Galactic coordinates centered on the Galactic center.

The correlation ratio in all multipole bins saturates and even starts decreasing as information from more channel maps is added after a certain velocity. We conclude from this analysis that the H α -based polarization template is most strongly correlated with the polarized dust emission in the range $-13 \text{ km s}^{-1} < v_{\text{lsr}} < 16 \text{ km s}^{-1}$ over most of the multipole bins considered. We find the same range for *B* modes as well. We use this cut for the rest of the analysis in this paper. Note that the correlation is already at the $\sim 20\%-40\%$ level for the H α -based polarization template of the individual velocity channel at $v_{\text{lsr}} = 2.03 \text{ km s}^{-1}$.

Cukierman et al. (2023) restrict the velocity range to $-15 \text{ km s}^{-1} < v_{\text{lsr}} < 4 \text{ km s}^{-1}$. With our velocity selection, we achieve an additional $\sim 5\%$ increase in the *EE* and *BB* correlation ratios with the Planck Commander dust maps relative to Cukierman et al. (2023).

Panopoulou & Lenz (2020a) propose the range $-12 \text{ km s}^{-1} < v_{\text{lsr}} < 10 \text{ km s}^{-1}$ for low-velocity clouds (LVCs). These are the 1st and 99th percentiles of the velocity distribution of clouds with H I column density $N_{\text{HI}} > 2.5 \times 10^{20} \text{ cm}^{-2}$ located in the Northern and Southern Galactic Polar regions. The velocity selection we make is close to this range. A benefit of restricting our analysis to this velocity range is that our template is less likely to include contributions from gas at very different distances, which decreases the likelihood of mixing different physical scales.

Using the RHT algorithm for determining the polarization angle, the H I intensity maps as the weighting (see Section 4.5), and the Spearman rank correlation coefficient and mean angle alignment as the correlation metrics, Clark & Hensley (2019) did not see the decrease in the correlation after a certain velocity that we see in Figure 4.1. The correlation asymptotes instead. There is also evidence that the intermediate velocity cloud (IVC) gas is organized into filaments that are aligned with their local magnetic fields (Panopoulou et al., 2019; Pelgrims et al., 2021b). The difference could be caused by the Hessian algorithm being more sensitive than the RHT to artifacts in low-signal velocity channels. We explore this in Section 4.5.3.

4.5.2 Dust Map and Mask Comparisons

We examine how the choice of dust emission maps and sky masks affects the correlation with the H I-based polarization template constructed using the Hessian algorithm on the

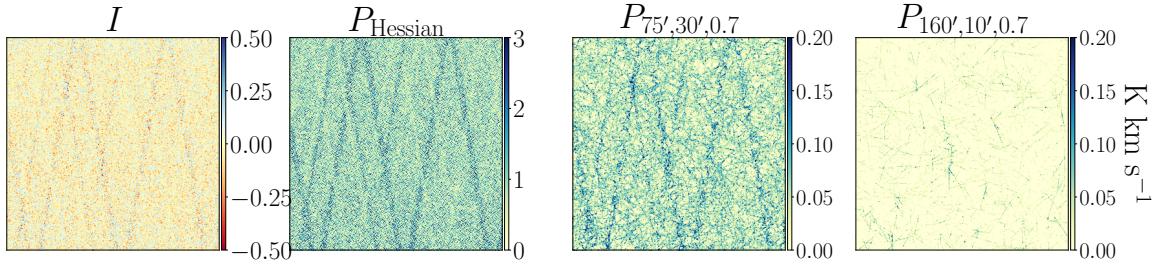


Figure 4.4: Maps of a $15^\circ \times 15^\circ$ patch of sky centered on R.A. = $7^\circ 5$, decl. = $28^\circ 1$ and $v_{\text{lsr}} = 400.1 \text{ km s}^{-1}$ of the input intensity channel from GALFA-HI with width 0.74 km s^{-1} (left) and polarized intensity of HI-based polarization templates constructed using the Hessian (middle left) and RHT (middle right and right) algorithms applied to this channel. The three parameters listed in the titles of the polarized intensity maps produced with the RHT algorithm are D_W , θ_{FWHM} , and Z , respectively, explained in Section 4.4.1.

aforementioned velocity selection in HI4PI data. We compare the E - and B -mode correlation ratios with the HI-based polarization template on the 70% sky fraction mask between the Planck Commander dust maps at 353 GHz and the Planck frequency maps at 353 GHz in Figure 4.2. We note that the difference is negligible in B modes and at low multipoles in E modes. The Commander dust maps correlate more strongly than the Planck frequency maps at higher multipoles in E modes. This is due to the CMB E modes, which contribute $\sim 10\%$ of the E -mode power to the 353 GHz frequency maps at $\ell > 300$. We therefore use the Commander dust maps for the rest of the analysis in this paper.

We test how the correlation ratio between the Planck Commander dust maps and the HI-based polarization template changes at different Galactic latitudes by utilizing the Planck sky fraction masks mentioned in Section 4.3. We invert the 80% sky fraction Galactic plane mask by switching ones to zeros, and vice versa, and call this the low-Galactic-latitude 20%

mask. We analyze the results for the high-Galactic latitude 20% sky fraction mask, the low-Galactic-latitude 20% mask, and the differences between the 20% and 40%, 40% and 60%, and 60% and 80% sky fraction Galactic plane masks. These five masks are shown in Figure 4.3, along with the EE and BB correlation ratios calculated over these masks. Although a significant portion of the dust column stops being traced by H I at lower Galactic latitudes because it is associated with molecular gas there (Lenz et al., 2017), we find a $\sim 20\%$ correlation with the low-Galactic-latitude 20% mask up to multipoles of $\ell \sim 400$. Since the velocity selection was optimized for the 70% sky fraction Galactic plane mask, we test whether the reported correlations calculated with the low-Galactic-latitude 20% mask increase when the H I-based polarization templates are integrated over a wider velocity range. We find that the correlation steadily increases over the entire multipole range considered with each template added out to $\pm 90 \text{ km s}^{-1}$, reaching $\sim 68\%$ at $\ell \sim 50$ in EE . This increase is expected because our initial velocity selection includes less than 47% of the total Galactic H I column density in this mask.

4.5.3 Effects of Resolution and Data Artifacts

In the previous subsections, we have only applied the Hessian to Hi4PI data, which has an angular resolution of $16.^{\circ}2$, using $N_{\text{side}} = 1024$. In this subsection, we apply the Hessian algorithm to GALFA-H I data, which has a much finer angular resolution of $4'$, using $N_{\text{side}} = 2048$. The velocity binning between the Hi4PI and GALFA-H I datasets is different, so we

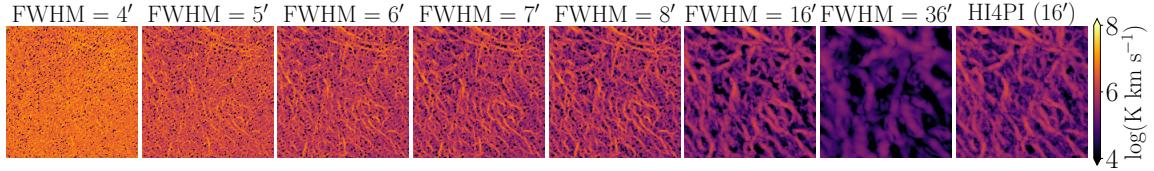


Figure 4.5: Polarized intensity map projections of a $400' \times 400'$ patch of sky, centered at $(l, b) = (15^\circ, 50^\circ)$, of H_I-based polarization templates constructed using the Hessian algorithm applied to GALFA-H_I intensity maps smoothed to different resolutions. The intensity map corresponding to the leftmost polarized intensity map is not smoothed, i.e., it has the native resolution of GALFA-H_I. The intensity maps corresponding to the polarized intensity maps to the right of the first map are smoothed with Gaussian kernels to the resolution stated in their titles. The last map on the right corresponds to the H_I4PI data at its native resolution, though it is integrated over a slightly different velocity range as mentioned in Section 4.5.3.

use the closest velocity range we can define for the GALFA-H_I dataset to the velocity selection we found using the H_I4PI dataset, which is $-15 \text{ km s}^{-1} < v_{\text{lsr}} < 18 \text{ km s}^{-1}$. However, as shown in Figure 4.1, most of the correlation comes from the H_I emission near the local standard of rest, and we find the difference in dust correlation between the velocity ranges $-15 \text{ km s}^{-1} < v_{\text{lsr}} < 18 \text{ km s}^{-1}$ and $-11 \text{ km s}^{-1} < v_{\text{lsr}} < 15 \text{ km s}^{-1}$ to be negligible.

We find that the Hessian algorithm highlights any structure with significant local curvature, which includes scan-pattern artifacts and other emission that is irrelevant to the physical gas filament distribution. To demonstrate this, we apply the Hessian method to a high-velocity channel from GALFA-H_I centered on $v_{\text{lsr}} = 400.1 \text{ km s}^{-1}$, which has relatively little emission above the noise level in the region analyzed. We compare the polarized intensity maps of an example patch of sky of H_I-based polarization templates

constructed using the Hessian and RHT algorithms applied to this channel in Figure 4.4.

Polarized intensity is defined as

$$P = \sqrt{Q^2 + U^2}, \quad (4.16)$$

where Q and U are the Stokes parameter maps of the HI-based polarization template at $v_{\text{lsr}} = 400.1 \text{ km s}^{-1}$ in this case. The RHT algorithm is discussed in 4.4.1. The input intensity map in this figure clearly shows the scan-pattern artifacts. The same artifacts are also obvious in the polarized intensity maps of the Hessian-based template and the first RHT-based template. However, they become much less obvious in the second RHT-based template with different parameters. This shows a limitation of the Hessian method, which the RHT algorithm can be tuned to avoid. This supports the hypothesis made in Section 4.5.1 concerning the decrease in the correlation with polarized dust emission when incorporating higher velocity channels beyond a specific threshold, which is observed only when using the Hessian method (Figure 4.1) but not when using the RHT algorithm (Clark & Hensley, 2019).

To mitigate this limitation, we experimented with various sensitivity-based weighting schemes for downweighting the low-intensity map pixels in each HI channel. We apply the Hessian algorithm to these new maps and find this weighting to have a negligible effect on the correlation ratio with polarized dust emission. This is true when using both GALFA-HI and HI4PI data for the HI intensity. This implies that the low-intensity HI pixels do not strongly affect the template, i.e., it is not necessarily the low-intensity pixels that contain

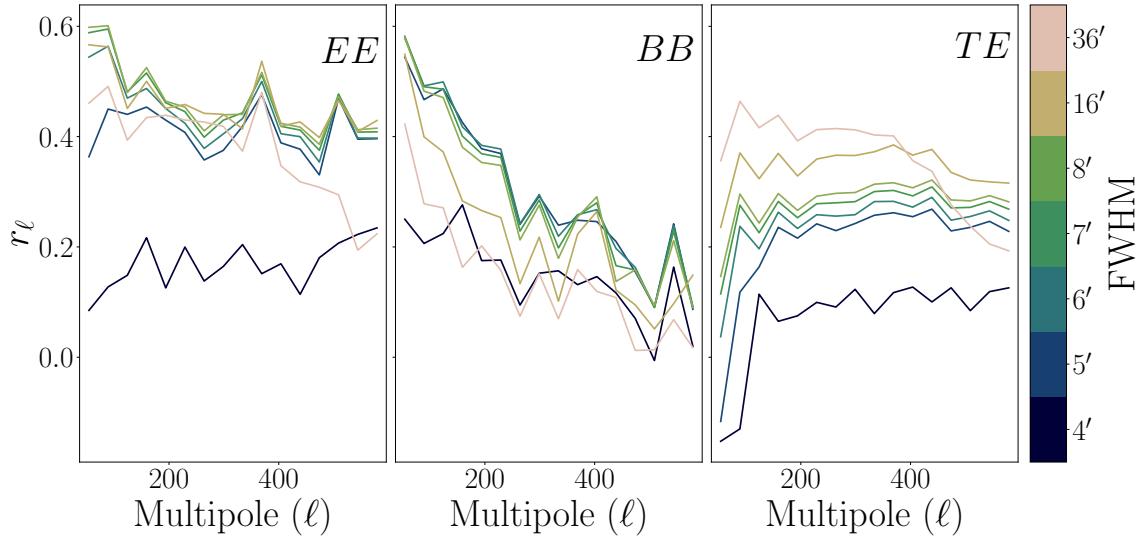


Figure 4.6: The EE (left), BB (middle), and TE (right) correlation ratios between the maps in Figure 4.5 and the Planck Commander dust maps at 353 GHz. The TE correlation ratio measures the correlation between the Planck total intensity and the templates' E modes. Correlations are computed on a combination of the GALFA-H I and Planck 70% sky fraction masks.

emission that is uninformative about filament orientations. Rather, this uninformative emission is likely scan-pattern artifacts and other data systematics.

To modify the scale of structure that the Hessian algorithm is most sensitive to, we apply a Gaussian smoothing kernel to each of the GALFA-H I intensity maps before applying the Hessian algorithm. Figure 4.5 shows the effect of smoothing before applying the Hessian. The maps shown are projections of the polarized intensity over the velocity selection described in Section 4.5.1 of an example patch of sky. The first map on the left corresponds to the Hessian algorithm run on the GALFA-H I intensity maps at their native resolution of $4'$. Note that it is difficult to see the filamentary structure because the Hessian algorithm is sensitive to other local variations, such as noise and scan-pattern artifacts. Each of the maps

to the right of the first one corresponds to the Hessian algorithm applied to the GALFA-H_I intensity maps smoothed with a Gaussian kernel to the labeled FWHM resolution. Increasing the FWHM emphasizes real filamentary structure but also makes the filaments wider. We include maps for FWHM=16' and 36', which are the native resolutions of Hi4PI and the Leiden/Argentine/Bonn (LAB) surveys (Kalberla et al., 2005), respectively. We include the projection of the polarized intensity of an example patch of the sky when the Hessian algorithm is applied to Hi4PI data at their native resolution of 16' on the right of the figure for comparison with the map titled FWHM=16'. The velocity binning and the velocity range over which the templates of the two datasets are integrated are not identical, so we do not expect their resulting polarized intensity maps to be identical. When smoothed to the Hi4PI beam, the GALFA-H_I maps are modestly more sensitive.

Smoothing the H_I data before constructing the H_I-based polarization templates deemphasizes small-scale noise at the cost of sensitivity to real small-scale H_I structure that may correlate well with the measured polarized dust emission. We explore this trade-off in GALFA-H_I data by calculating the *EE*, *BB*, and *TE* correlation ratios of the different maps in Figure 4.5 with the Planck Commander dust maps and plot the results in Figure 4.6. For the *TE* case, we correlate the Planck dust total intensity with the templates' *E* modes.

The trend is not consistent between the three panels. The *TE* correlation simply increases as the GALFA-H_I data are smoothed, with an expected dip at the smoothing scale, which is only within the multipole range considered for the FWHM=36' case. However,

the EE and BB correlations are maximized when the H_I data are smoothed to intermediate resolutions. While an increase in the EE and BB correlation ratios is what we should aim for, an increase in the TE correlation ratio is not necessarily better since we expect the real TE correlation ratio to be ~ 0.36 (Planck Collaboration et al., 2020g). Therefore, a near-ideal E -mode template should correlate with the Planck dust total intensity at about that level. However, in all cases, we achieve a significant improvement in the correlation with polarized dust emission by smoothing the map before applying the Hessian algorithm. The H_I-based polarization template based on Hi4PI data at their native 16' resolution is similarly correlated with the dust polarization to the template based on the GALFA-H_I data smoothed to 16'. The map smoothed to 36' decreases both the EE and BB correlations with the Planck Commander dust maps. This shows the utility of the higher-resolution H_I intensity data from the GALFA-H_I and Hi4PI surveys in modeling polarized dust emission over the lower-resolution LAB survey.

We compare the H_I-based polarization templates based on GALFA-H_I data smoothed to $\text{FWHM} = 7'$ and on Hi4PI data by plotting their EE and BB correlation ratios with the Planck Commander dust maps at 353 GHz in Figure 4.7. We use a $\sim 23\%$ sky fraction mask by combining the GALFA-H_I mask with the Planck 70% sky fraction mask. The H_I-based polarization template constructed from the GALFA-H_I dataset is more strongly correlated in B modes with the polarized dust emission than that constructed using the Hi4PI dataset. The improvement is at the $\sim 10\%$ level at multipoles $\ell < 350$. This implies that

higher-resolution H I data are useful for better modeling the polarized dust foreground in B modes. If we smooth the GALFA-H I dataset to a resolution of $\text{FWHM} = 16'$ (the resolution of Hi4PI) instead of $7'$, the BB improvement becomes negligible.

The EE and BB correlation ratios between the H I-based polarization template constructed from Hi4PI data and the Planck Commander dust maps at 353 GHz do not increase after smoothing the Hi4PI dataset. For the rest of the analysis in this paper, we use the Hi4PI dataset at its native resolution.

4.6 Spherical Rolling Hough Transform

As discussed in the previous section, the Hessian algorithm is sensitive to the local curvature in images and thus is most sensitive to structure at the image resolution. To explore different filament morphologies, however, we need an algorithm with free parameters that help set the scale and shape of the identified filaments. One such algorithm is the RHT (Clark et al., 2014, 2020) described in Section 4.4.1.

The code for the RHT algorithm is publicly available and has been applied to a variety of astronomical images, including molecular clouds (Malinen et al., 2016; Panopoulou et al., 2016), magnetohydrodynamic simulations (Inoue & Inutsuka, 2016), depolarization canals (Jelić et al., 2018), the solar corona (Boe et al., 2020), and supernova remnants (Raymond et al., 2020). The algorithm has been adapted to work on resolved stars for stellar stream detection (Pearson et al., 2022) and extended to add the ability to identify filaments with

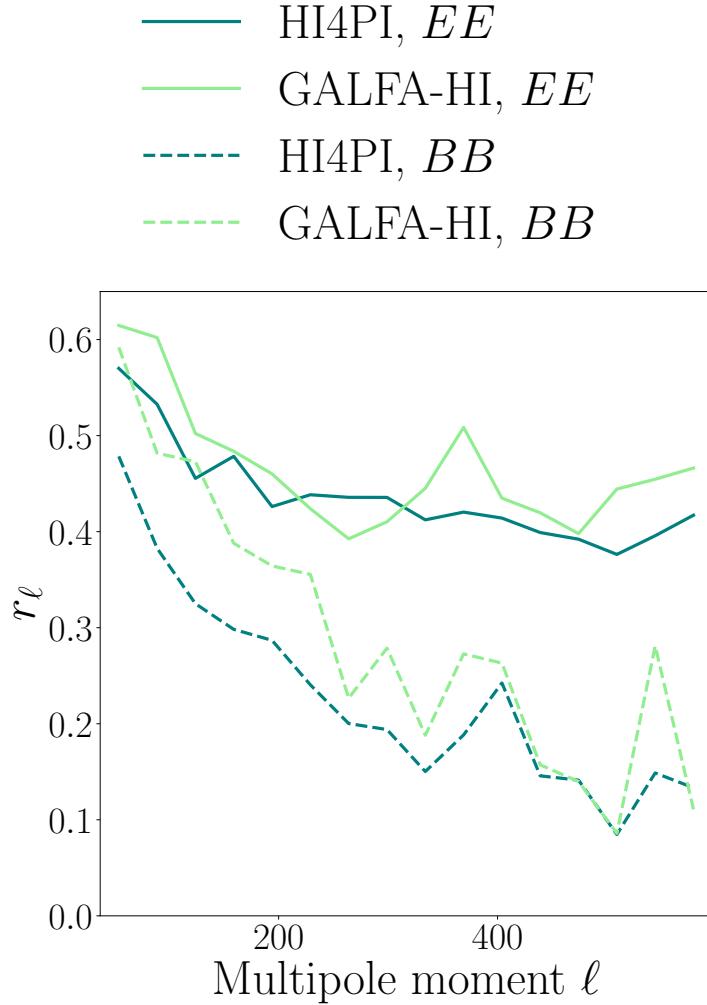


Figure 4.7: Comparison of the EE (solid) and BB (dashed) correlation ratios with the Planck Commander dust maps at 353 GHz of the $\text{H}\alpha$ -based polarization templates using GALFA- $\text{H}\alpha$ data smoothed to a FWHM of $7'$ (light green) and HI4PI data (teal). The $\text{H}\alpha$ -based polarization templates are integrated over a similar velocity range and constructed using the Hessian algorithm. The correlations are calculated on a combination of the GALFA- $\text{H}\alpha$ mask with the Planck 70% sky fraction mask. The teal lines are the same as those in Figure 4.2.

user-specified widths (Carrière et al., 2022).

The RHT algorithm currently runs on flat-sky projections of small patches of the sky.

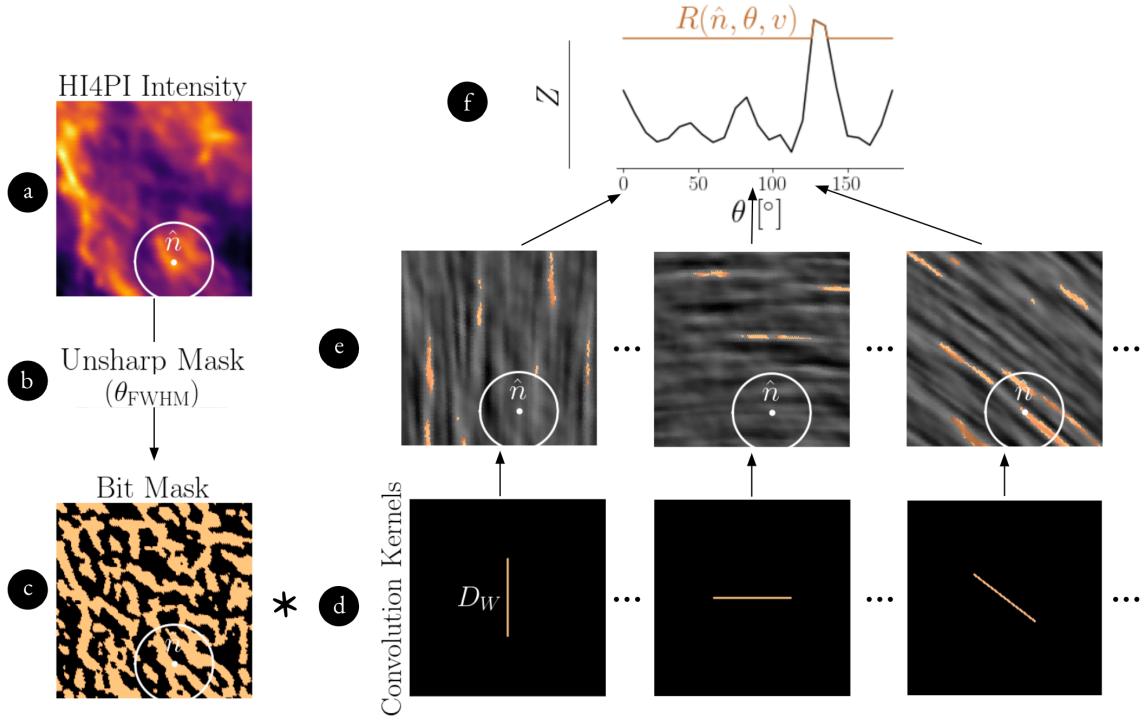


Figure 4.8: Diagram of the Spherical RHT procedure. (a) A flat-sky projection of a $400' \times 400'$ patch of sky, centered at $(l, b) = (15^\circ, 50^\circ)$ and $v_{\text{lsr}} = 2.03 \text{ km s}^{-1}$ for the initial H I intensity channel with width 1.3 km s^{-1} from the HI4PI Survey. The diameter of the white circle drawn around pixel \hat{n} is equal to D_W , the length of the convolution kernels, three of which are shown to scale (d). We chose $D_W = 160'$ in this case. (c) The resulting binary map of the preprocessing steps (b; Steps 1 and 2 in Section 4.4.1) with $\theta_{\text{FWHM}} = 10'$ applied to (a). The convolution kernels (d) are both rotated and convolved with (c) in spherical harmonic space. (e) The results of the convolutions between (c) and (d). (f) The result of the convolutions for pixel \hat{n} over orientations θ . A threshold, $Z = 0.7$ in this case, is applied to the result of the convolutions (Step 4 in Section 4.4.1), leaving $R(\hat{n}, \theta, v)$ (copper). The colors in (e) and (f) are set to match, i.e., the pixels (e) have a copper-like color scale where the resulting intensities (f) pass Z and a gray color scale otherwise.

Therefore, to achieve results over the full sky, small patches of the spherical map need to be projected into separate flat-sky images. This is time-consuming and the reprojection step may produce distortion effects. Hence, we implement an alternative algorithm that enables RHT computation directly on the sphere by utilizing spherical harmonic convolutions.

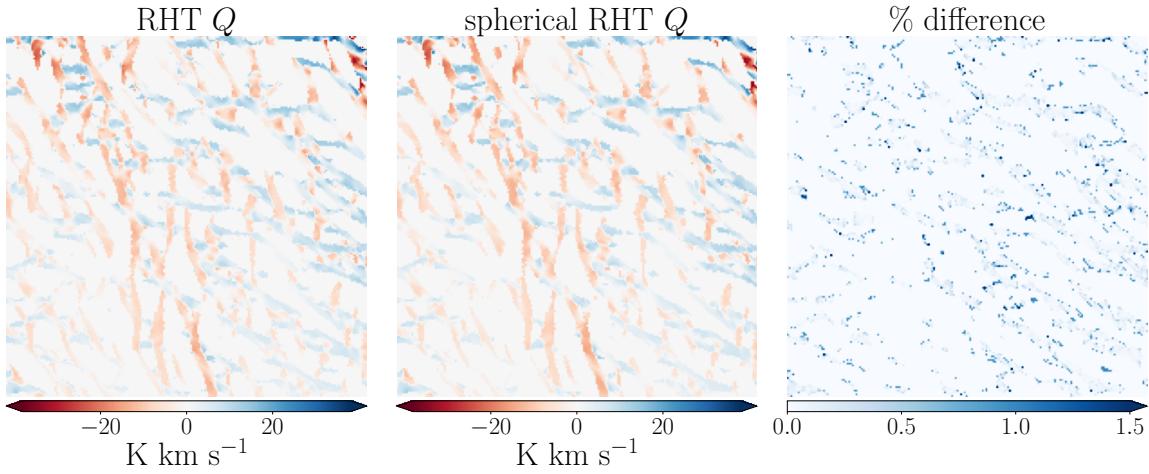


Figure 4.9: Stokes Q map projections of a $15^\circ \times 15^\circ$ patch of sky, centered at $(l, b) = (15^\circ, 50^\circ)$, of H I-based polarization templates for one velocity slice of the Hi4PI Survey centered at $v_{\text{lsr}} = 2.03 \text{ km s}^{-1}$ with width 1.3 km s^{-1} constructed using the RHT (left) and the Spherical RHT (middle) algorithms with parameters $D_W = 75'$, $\theta_{\text{FWHM}} = 30'$, and $Z = 0.7$. The map on the left is used in Clark & Hensley (2019). The map on the right is the percentage difference between the map on the left and the map in the middle.

4.6.1 Spherical Convolutions

In a flat geometry, a convolution between two maps can be computed from the product of their Fourier representations. On the sphere, the convolution can be expressed similarly as a product of their spherical harmonic representations. A major difference, however, is that the spherical harmonic representation is weighted by Wigner matrices (Wandelt & Górski, 2001; Prézeau & Reinecke, 2010). The convolution of a map with spherical harmonics $a_{\ell m}$ and a convolution kernel defined on the sphere with spherical harmonics $b_{\ell m'}$ for Euler

angles (α, β, γ) can be written as

$$c(\alpha, \beta, \gamma) = \sum_{m'=-m'_{\max}}^{m'_{\max}} \sum_{m=-\ell_{\max}}^{\ell_{\max}} e^{im'\alpha} e^{im\gamma} C_{m'm}(\beta), \quad (4.17)$$

where

$$C_{m'm}(\beta) \equiv \sum_{\ell=0}^{\ell_{\max}} b_{\ell m'}^* D_{m'm}^{\ell}(\beta) a_{\ell m}, \quad (4.18)$$

and $D_{m'm}^{\ell}$ are the so-called Wigner matrices. In our case, α would represent the orientation of the convolution kernel, and β and γ represent the latitude and longitude of the sky, respectively. Refer to Prézeau & Reinecke (2010) for more details.

The computation of the Wigner matrices is usually the bottleneck of convolution algorithms. This is a major problem in CMB analyses in the context of beam convolutions (e.g., Challinor et al., 2000). If the convolution kernel is restricted to a small set of m' values, then the algorithm can run faster. An example is a symmetric beam, which is restricted to $m = 0$. Our filamentary kernels are not symmetric, but we can limit the maximum m' value and retain the intended shape as described in Section 4.6.2. We use `ducc`³, a computationally efficient code for performing convolutions with axially asymmetric convolution kernels.

³<https://gitlab.mpcdf.mpg.de/mtr/ducc>

4.6.2 The Algorithm

We implement the steps described in Section 4.4.1 directly on the sphere rather than on flat-sky image projections. We call this new implementation the spherical Rolling Hough Transform (Spherical RHT). This implementation replaces Step 3 of the RHT algorithm in Section 4.4.1 with the spherical harmonic convolutions described in Section 4.6.1. We show a diagram of the full procedure in Figure 4.8. This diagram shows how the parameters θ_{FWHM} , D_W , and Z are used to transform the HI intensity at each velocity channel v and pixel \hat{n} into $R(\hat{n}, \theta, v)$, which is used in Equations 4.8, 4.9, and 4.10 to construct the HI-based Stokes Q and U maps for that velocity channel.

We define a convolution kernel as a line of neighboring nonzero pixels of length D_W on a HEALPix grid of a higher resolution than the maps we convolve it with. We then smooth this line of pixels so that the pixelization of the lower-resolution maps captures all of the information in the kernel. This prevents aliasing from small scales that the pixelization of the lower-resolution maps is insensitive to. For instance, when run on Hi4PI data at $N_{\text{side}} = 1024$, we define the kernel at $N_{\text{side}} = 4096$ and smooth it to a FWHM of $3.4'$, the width of a pixel at $N_{\text{side}} = 1024$.

We find that limiting the m'_{max} of the kernel to 50 retains the intended shape of the kernel visually. Also, the results in this paper are identical when m'_{max} is increased to 100. Limiting the m'_{max} of the convolution kernel increases the computational efficiency of the algorithm.

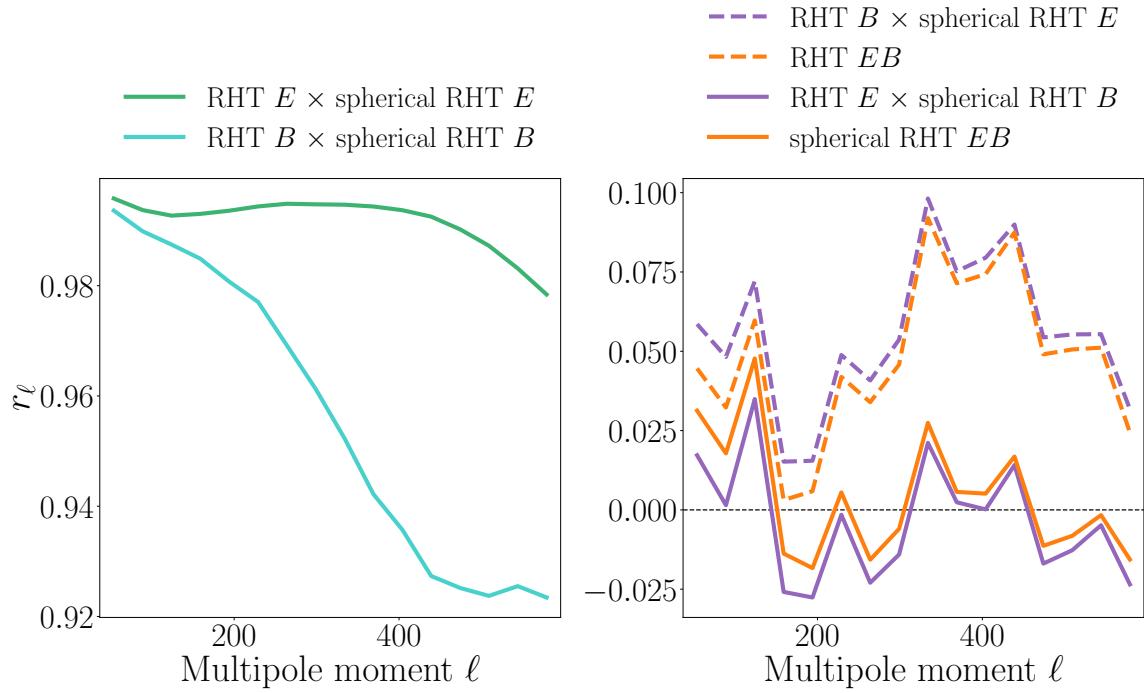


Figure 4.10: Left: the EE (green) and BB (blue) correlation ratios between the HI-based polarization template used in Clark & Hensley (2019) and that reproduced using the Spherical RHT algorithm with the same parameters. Right: the EB correlation ratios, showing that the likely spurious positive correlation when the B modes of the Clark & Hensley (2019) HI-based polarization template are used (dashed) vanish when the curved sky is taken into account and the B modes produced with the Spherical RHT algorithm are used (solid). The orange (purple) curves represent correlation ratios where the E and B modes of the same algorithm (different algorithms) are used. HI4PI data and the Planck 70% sky fraction mask are used in these plots. Note the difference in the y-scales between the left and right panels.

We convolve this kernel at different orientations with the map as described in Section 4.6.1. In the standard RHT algorithm, the number of orientations depends on D_W . For all D_W used in this paper, dividing the kernel orientations into 25 bins yields consistent results to dividing them into 300 bins. Therefore, we use 25 orientations for computational efficiency. However, the number of orientations is left as a free parameter in the code in

case more orientations are necessary for different applications. The code is made publicly available on GitHub⁴ (Halal et al., 2023).

4.6.3 Comparison with the RHT

We follow the prescription described in Clark & Hensley (2019), replacing the RHT on small flat-sky projections with the Spherical RHT, to construct full-sky HI-based polarization templates and compare the results of the two algorithms.

Once the distribution $R(\hat{\mathbf{n}}, \theta, v)$ over orientations θ is obtained for each pixel $\hat{\mathbf{n}}$ at each velocity channel v after Step 4 in Section 4.4.1, Clark & Hensley (2019) construct Stokes Q_{HI} and U_{HI} maps as in Equations (4.8), (4.9), (4.10), (4.11), and (4.12).

We run the Spherical RHT on HI4PI data with the same velocity binning as in Clark & Hensley (2019) and with the same free parameters, $\theta_{\text{FWHM}} = 30'$, $D_W = 75'$, and $Z = 0.7$, and compare the resulting Stokes Q polarization maps to those of Clark & Hensley (2019) in Figure 4.9. We show flat-sky projections of an example patch of sky of these maps. We also show the percentage difference between the maps and note that the results are qualitatively the same but not numerically identical, and we do not expect them to be. We note that the difference is mostly concentrated at the edges of the filaments. We do not show the Stokes U results because the conclusions are the same.

⁴<https://github.com/georgehalal/sphericalrht>

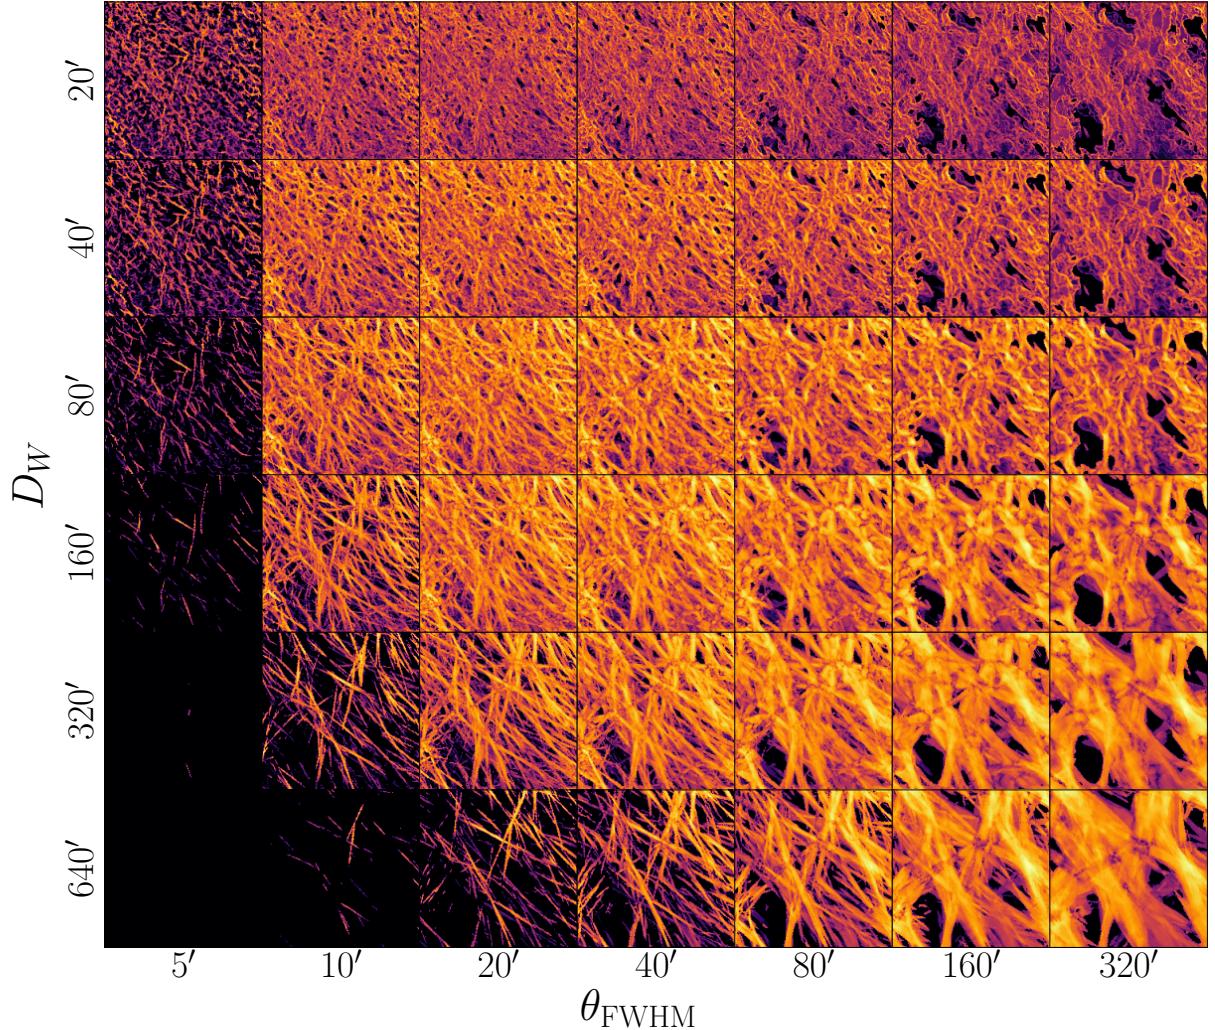


Figure 4.11: Polarized intensity map projections of a $16.7^\circ \times 16.7^\circ$ patch of sky, centered at $(l, b) = (15^\circ, 50^\circ)$, of H I-based polarization templates constructed using the Spherical RHT algorithm with different parameters applied to H I4PI intensity maps. The parameter Z is fixed to 0.7 and the Spherical RHT is run on a grid of exponentially increasing parameters between 5' and 320' for θ_{FWHM} (to the right) and 20' and 640' for D_W (to the bottom).

To quantitatively test the differences between the two algorithms for constructing H I-based polarization templates, we plot the correlation ratio of all different combinations of the E and B modes of the two algorithms in Figure 4.10 using the Planck 70% sky

fraction mask. The correlations between the Clark & Hensley (2019) and the Spherical RHT templates are higher than 98% in E modes and higher than 92% in B modes across all multipoles considered. Since both algorithms assume no misalignment between the filament orientations and the local magnetic fields as described in Section 4.5, we expect the EB correlation to be zero unless the morphology of the filaments across the sky has a preferred chirality as described in Section 4.7.1. The EB correlations are negligible in this figure when the B modes are predicted by the Spherical RHT algorithm, i.e., when the curved sky is taken into account. However, the EB correlations are at the $\sim 5\%$ level across the multipole range when the B modes are predicted by the Clark & Hensley (2019) maps. These results are not latitude dependent. Note that this is not a bug in the RHT code itself. Rather, it is an artifact of projecting each window of the map onto a flat-sky image. While it does not affect any of the results presented in Clark & Hensley (2019), it is preferable for our morphology investigation that the Spherical RHT does not have this property.

4.7 Filament Morphologies

4.7.1 Morphological Parameter Space Exploration

The Spherical RHT enables efficient exploration of the D_W , θ_{FWHM} , and Z parameter space that governs how the geometry of the filamentary H I structure is mapped into an H I-based polarization template. We use this to investigate what H I filament morphologies are most

predictive of the measured polarized dust emission. This is a continuation of the parameter space exploration performed in BICEP/Keck Collaboration et al. (2023e). That work used the RHT on a small patch that covered $\sim 1\%$ of the sky and assessed the cross-correlation between the HI-based polarization template and multifrequency polarized dust emission data from BICEP/Keck and Planck. Utilizing the Spherical RHT, we extend this analysis to the full sky in this paper.

Using results from the exploration in BICEP/Keck Collaboration et al. (2023e), we fix the Z parameter to 0.7, such that the algorithm is only sensitive to structures larger than 70% of D_W . We run the Spherical RHT on a grid of exponentially increasing parameters between 5' and 320' for θ_{FWHM} and 20' and 640' for D_W . We show the polarized intensity maps corresponding to these parameters in Figure 4.11. We calculate the EE , BB , and TE correlation ratios between these maps and the Planck Commander dust maps over a broadband multipole bin between $\ell = 20$ and $\ell = 600$ in Figure 4.12. For the TE case, we compute the correlation between the Planck 353 GHz total intensity and the template E modes.

For the lowest θ_{FWHM} and highest D_W case, the RHT intensity is zero because no linear structures cover at least 70% of D_W after the unsharp mask step removes structure on scales greater than 5'. By contrast, for the highest θ_{FWHM} and lowest D_W case, the template morphology is more sensitive to lower-intensity diffuse HI structure, which tends to be less filamentary. This is reflected in the morphology of the HI-based polarization template in the upper right-hand corner of Figure 4.11. At fixed θ_{FWHM} , the filaments tend

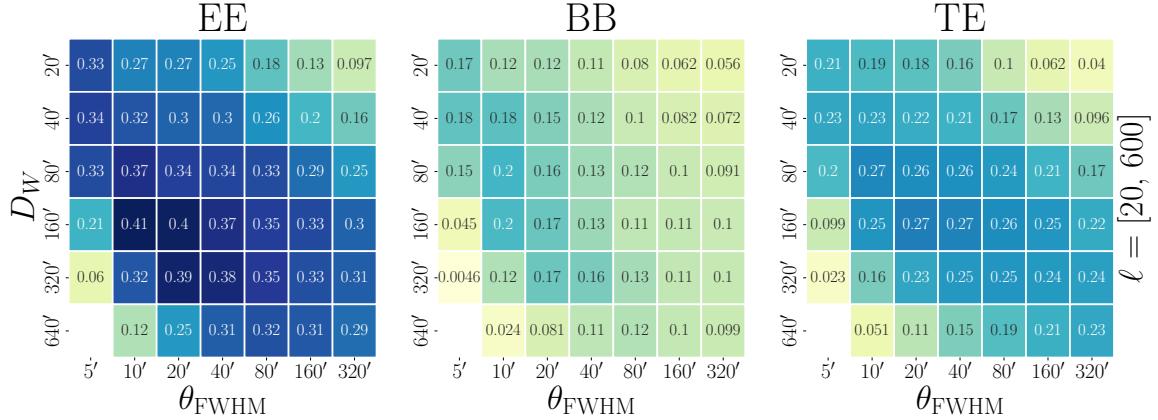


Figure 4.12: The *EE* (left), *BB* (middle), and *TE* (right) correlation ratios between the maps in Figure 4.11 and the Planck Commander dust maps over a broadband multipole bin between $\ell = 20$ and $\ell = 600$. The Planck total intensity is correlated with the *E* modes of the Spherical RHT-based templates for the *TE* case.

to be longer with increasing D_W , and at fixed D_W , the filaments tend to be wider with increasing θ_{FWHM} . We find a clear gradient in the correlation coefficient over the parameter space explored with a preference toward $\theta_{\text{FWHM}} \sim 10' - 20'$, i.e., near the $16'$ HI4PI beam scale, and $D_W \sim 80' - 160'$. We repeat this exercise for smaller broadband multipole bins and find that the $\theta_{\text{FWHM}} \sim 10' - 20'$ preferred scale does not change, while the preferred scale for D_W increases slightly when considering larger scales and decreases slightly when considering smaller scales.

The fact that the preferred θ_{FWHM} range is approximately at the beam scale indicates that the thinnest resolved filaments are the most informative about the magnetic field orientation. This means that the HI filaments that are best-correlated with the polarized dust emission are somewhat thinner and longer than the structures that the Clark & Hensley (2019) analysis was most sensitive to at $D_W = 75'$ and $\theta_{\text{FWHM}} = 30'$. We find that the structures that are

more qualitatively filamentary in Figure 4.11 correlate best in all the metrics in Figure 4.12. The less-linear morphologies that populate the upper right-hand portion of Figure 4.11 correlate poorly with the Planck data. This indicates that elongated linear structures are genuinely the geometry that best describes the polarized dust emission field within the morphological parameter space we can explore.

In the isolated filament case, *E* modes are primarily sourced along the length of the filament, while *B* modes are primarily sourced at the edges of the filament (Huffenberger et al., 2020). Therefore, the H I-based polarization templates that maximize the *EE* correlation with polarized dust emission are likely the ones that most closely recover the geometry of the dust filaments that dominate the polarized intensity, i.e., their lengths, widths, and orientations. The similarity in the correlation dependence on D_W and θ_{FWHM} between the *EE* and *TE* panels of Figure 4.12 also supports this conclusion. The H I-based polarization templates that maximize the *BB* correlation with polarized dust emission are likely the ones whose filaments are at the correct distances from each other, which affects the constructive and destructive interference of the *B*-mode patterns.

To illustrate the *E*- and *B*-mode patterns produced in the multifilament case, we run the Spherical RHT algorithm with the parameters used in Clark & Hensley (2019, $D_W = 75'$, $\theta_{FWHM} = 30'$, and $Z = 0.7'$) on maps of two synthetic filaments, slightly offset from one another in longitude. Figure 4.13 shows the results when these filaments are positioned at different distances relative to each other in latitude and when they are reflected

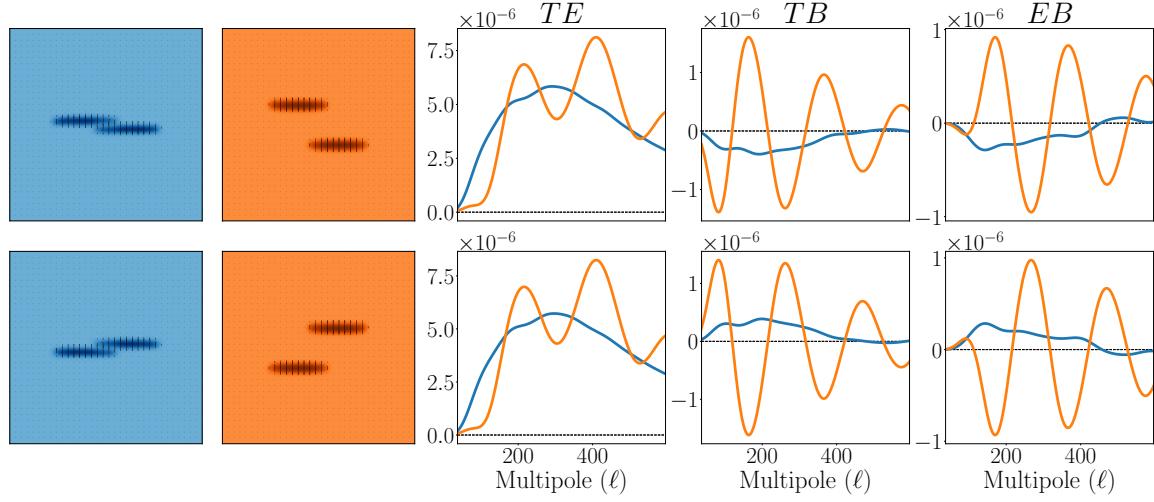


Figure 4.13: Right panel: the *TE* (left), *TB* (middle), and *EB* (right) cross spectra (D_ℓ) of two synthetic filaments close to each other (blue, left panel) and far from each other (orange, left panel), where the arrangement of the filaments in the top and bottom rows have opposite parities. The headless vectors in the left panels show the polarization angle orientations perpendicular to the lengths of these filaments. The color scale in the left panels represents the polarized intensity.

relative to each other in longitude. The *TE* spectra are positive whether the filaments are close to or far from one another and are unaffected by the parity of the filaments' relative positions. By contrast, the *TB* spectra fluctuate around zero when the filaments are far from one another but are only positive or only negative, depending on parity, when the filaments are close to one another. This shows how *E* modes mainly depend on the individual filament geometries, whereas *B* modes mainly depend on the geometry of the filaments relative to one another.

Figure 4.13 is also a demonstration of how the Spherical RHT algorithm can be used as a tool to extend our intuition about morphological features that produce parity-violating signatures, i.e., nonzero *TB* and *EB*. We show positive-only *TB* and *EB* signals when

the two filaments are positioned close to one another in one handedness and negative-only signals when those filaments are close to one another in the opposite handedness. Characterizing parity-violating signatures in polarized dust emission is interesting for several applications, including confounding cosmic birefringence searches (Minami & Komatsu, 2020) and biasing CMB polarization "self-calibration," which assumes that TB and EB signals are due to systematic errors because they must vanish in the standard cosmological model (Abitbol et al., 2016). Huffenberger et al. (2020) and Clark et al. (2021) have shown that a misalignment between dust filaments and the local magnetic field orientations can produce parity-violating signatures. The HI-based polarization templates assume perfect alignment between the filament morphologies and the magnetic field orientations. However, even without local misalignment, chirality in the distribution of filaments or in the non-filamentary polarizing structures could produce these signatures. These may be especially significant on small patches of the sky, whereas the signal may average down if large sky areas show no preferred "handedness" of the dust intensity distribution.

4.7.2 Spherical RHT- and Hessian-based Template Comparison

We compare the HI-based polarization templates produced using the Spherical RHT algorithm to those produced using the Hessian algorithm to understand the factors that affect the correlation of each with the measured polarized dust emission. We start by cross correlating

the Hessian-based template with Spherical RHT-based templates constructed using different parameters. We show the EE and BB correlation ratios of four of these in Figure 4.14. All of the spectra use the Planck 70% sky fraction mask. The BB correlation ratios are much weaker than the EE correlation ratios of the same parameters at small scales – in other words, the Spherical RHT and Hessian templates are more similar to one another in E modes than in B modes. While we find a set of Spherical RHT parameters that produces a template that correlates at the $\sim 90\%$ level in E modes with the Hessian-based template, none of the parameter sets we test produces a template that correlates higher than $\sim 60\%$ at $\ell > 500$ in B modes.

We plot the Stokes Q and U projections of an example patch of sky for the HI-based polarization templates based on the Hessian method and the Spherical RHT algorithm using the best-correlating parameters we found in Section 4.7.1 and the full-mission Planck Commander dust map in Figure 4.15. This gives a sense of the different polarization structures predicted by these filamentary templates as compared to the observed total dust field. The Planck maps are noise dominated at small scales. However, the large-scale correspondence is visible in this figure.

We compare the correlation of the HI-based polarization templates from Figure 4.15 with polarized dust emission. For a broadband multipole bin between $\ell = 20$ and $\ell = 140$ and centered at $\ell = 80$, a range relevant to primordial B -mode detection, we find that the EE (BB) correlation ratios are 0.5 and 0.6 (0.42 and 0.48) for the templates constructed

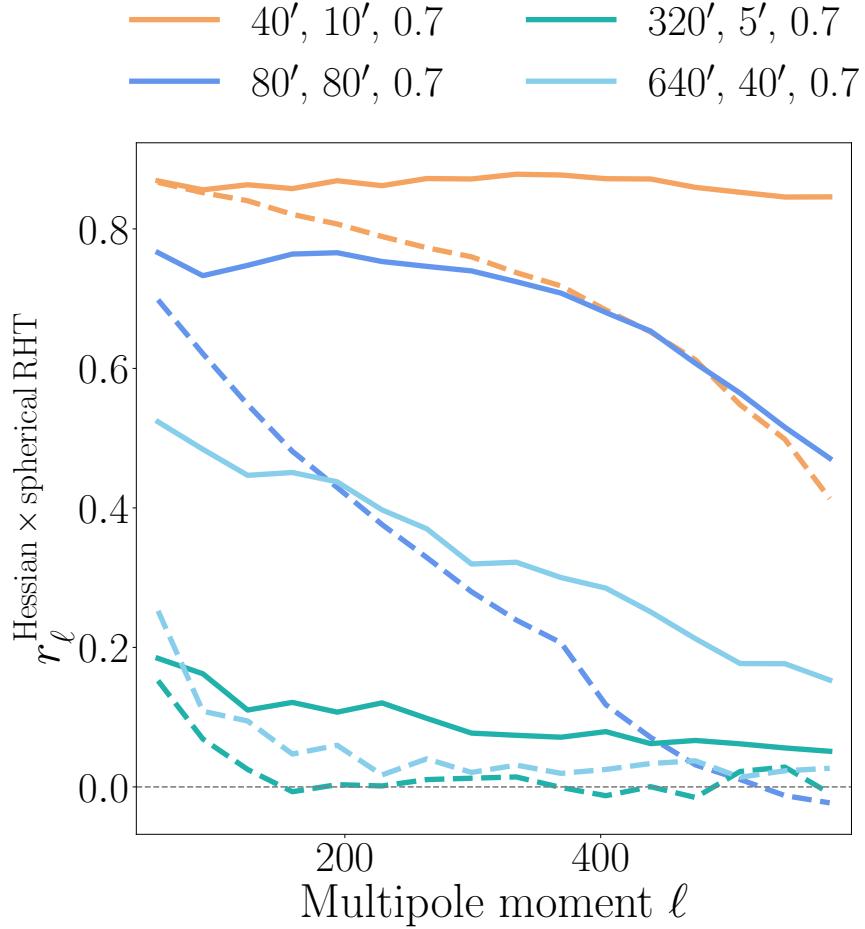


Figure 4.14: The *EE* (solid) and *BB* (dashed) correlation ratios between the H I-based polarization template constructed using the Hessian algorithm and four H I-based polarization templates constructed using the Spherical RHT algorithm on HI4PI data. The Planck 70% sky fraction mask was used for the spectra used for calculating these correlation ratios. The parameters listed in the legend are the D_w , θ_{FWHM} , and Z , respectively, defined in Section 4.4.1. The first set of parameters (sandy brown) is the one that correlates the best, and the other three are randomly selected.

using the Hessian and Spherical RHT algorithms, respectively. We plot the *EE*, *BB*, and *TE* correlation ratios with the Planck Commander dust maps as well as the *EE*-to-*BB* autospectra ratio in Figure 4.16. All of the plots in Figure 4.16 use the Planck 70% sky

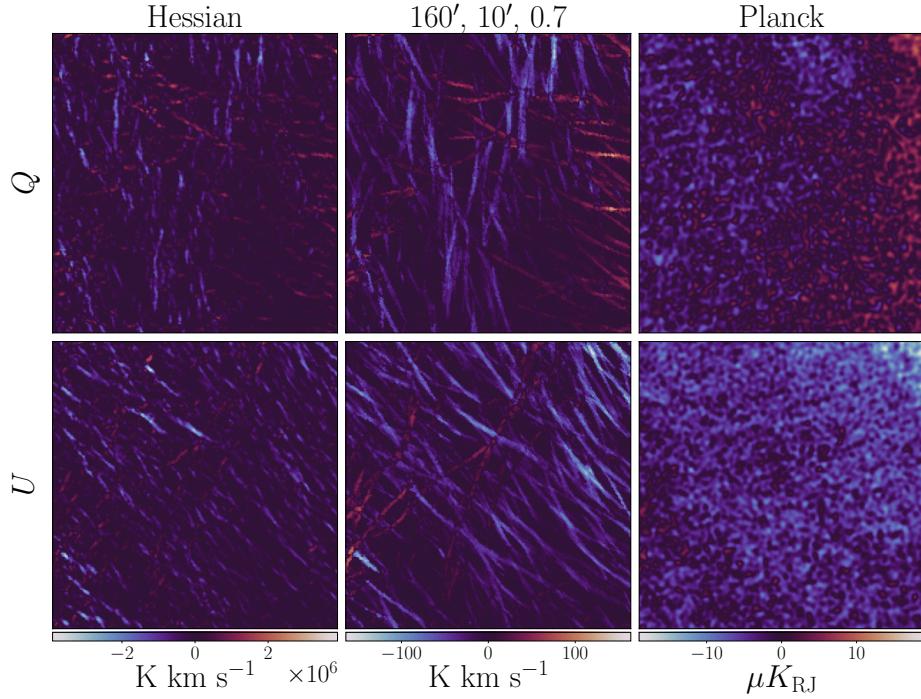


Figure 4.15: Stokes Q (top) and U (bottom) map projections of a $16.7^\circ \times 16.7^\circ$ patch of sky, centered at $(l, b) = (15^\circ, 50^\circ)$, of the HI-based polarization templates constructed using the Hessian method (left) and Spherical RHT algorithm with parameters $D_W = 160'$, $\theta_{\text{FWHM}} = 10'$, and $Z = 0.7$ (middle) applied to HI4PI intensity maps, and the full-mission Planck Commander dust map (right).

fraction mask and HI4PI data. However, the results are qualitatively similar for different sky fraction masks.

For the templates constructed with the Spherical RHT algorithm, the EE -to- BB autospectrum ratio peaks at different multipoles for different parameter sets, which determine the typical size of the measured filaments. The peaks are driven by the EE autospectra in the numerators of these ratios. This is because E modes predominantly originate along the filaments, so the most sensitive scale of the filament quantification method sets the dominant scale of the E -mode power. Note that these templates correctly predict an excess

of E modes over B modes, supporting the idea that a preference for filaments in the real sky to be magnetically aligned results in an observed EE -to- BB ratio higher than unity (Clark et al., 2015; Planck Collaboration et al., 2016a). These templates, however, may overpredict the fraction of power in E modes on certain scales. Any method of quantifying filament orientations will be more or less sensitive to structure on particular scales (see the discussion in Hacar et al., 2022). This means that an H I-based polarization template built from filament orientations will have H I-based polarized intensity concentrated at particular angular scales, as is the case in this work. Some of the overprediction of E modes is likely due to this scale dependence not being representative of the hierarchical filamentary morphology of the real sky. Additionally, the H I-based polarization template may have excess E -mode power in part because they only model the filamentary component of the polarized dust emission, whereas the real sky contains additional polarized emission in extended structures that may not resemble filaments. The BB -to- EE ratio of the real polarized dust emission observed by Planck over large sky areas is $\sim 0.53 \pm 0.01$ over the multipole range $40 \leq \ell \leq 600$ (Planck Collaboration et al., 2020g).

Although we found that the Spherical RHT template correlates better with polarized dust emission than the Hessian template between $\ell = 20$ and $\ell = 140$, the Hessian algorithm correlates better at higher multipoles, especially in B modes, as shown in Figure 4.16. Because the RHT and Hessian templates predict different distributions of both polarization angles and polarized intensity, we investigate whether one of these factors is driving the

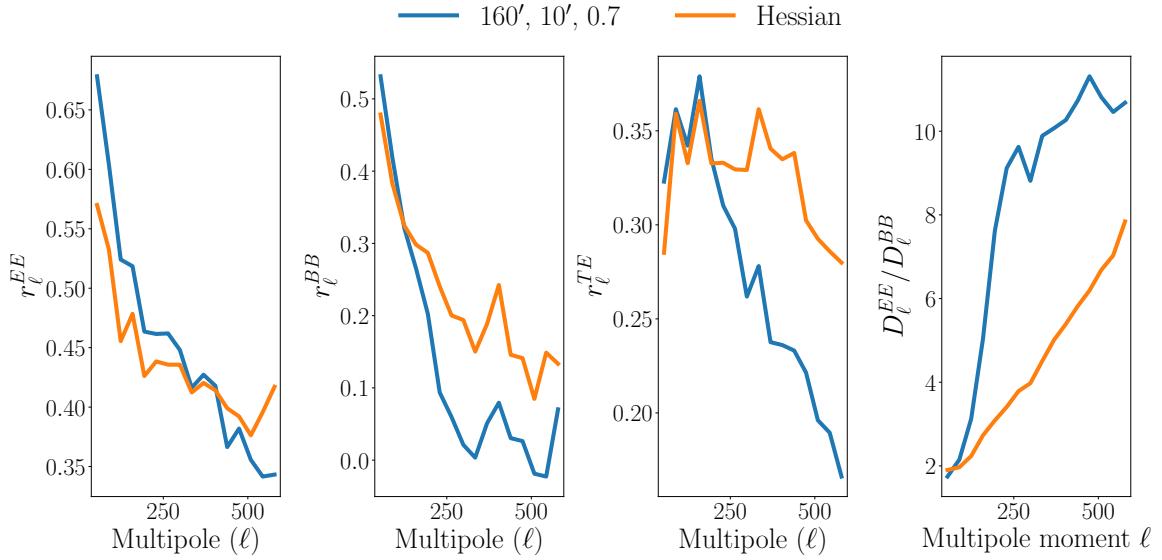


Figure 4.16: The EE (left), BB (middle left), and TE (middle right) correlation ratios of the Planck Commander dust maps with the $\text{H}\alpha$ -based polarization template constructed using the Hessian algorithm (orange) and those constructed using the Spherical RHT with $D_W = 160'$, $\theta_{\text{FWHM}} = 10'$, and $Z = 0.7$ (blue). The right panel shows the EE -to- BB ratios of the autospectra of the aforementioned $\text{H}\alpha$ -based polarization templates. All spectra in this figure are calculated using the Planck 70% sky fraction mask.

stronger correlation between the Hessian-based template and the polarized dust emission data.

We implement a modification to the Spherical RHT-based template construction to make its orientation angle selection more similar to that of the Hessian algorithm. While the Hessian algorithm determines the orientation of the filaments based on the local eigenbasis, the $\text{H}\alpha$ -based polarization template constructed with the RHT algorithm computes a mean over all orientations weighted by the result of the convolutions at those orientations $R(\hat{\mathbf{n}}, \theta, v)$; (see Section 4.5). Therefore, we instead take the angle at the peak of $R(\hat{\mathbf{n}}, \theta, v)$ as the orientation of the filament. We find that for small window diameters D_W , this increases

the E -mode correlation with the polarized dust emission by $\sim 5\%$. However, the effect is negligible in B modes and for large window diameters. This behavior is expected because larger window diameters are more likely to have singly peaked $R(\hat{\mathbf{n}}, \theta, v)$, such that the weighted mean of the angles and the angle at peak $R(\hat{\mathbf{n}}, \theta, v)$ are similar.

Each Stokes Q/U template is constructed from $\text{H}\alpha$ -based orientations and $\text{H}\alpha$ -based polarized intensity. We construct hybrid templates where the orientations are derived from the Spherical RHT and the polarized intensities are derived from the Hessian, and vice versa. Using the Spherical RHT parameters that produce the best correlation with the Hessian algorithm, we construct additional templates using Spherical RHT orientations, but weighting the Stokes Q and U maps by the Hessian eigenvalue-based weighting $w_H(\hat{\mathbf{n}}, v)$ described in Section 4.4.3 instead of the $\text{H}\alpha$ -intensity-based weighting $I_{\text{HI}}(\hat{\mathbf{n}}, v)$. Similarly, we construct another hybrid template using Hessian-derived orientations and Stokes Q and U maps weighted by the $\text{H}\alpha$ -intensity-based weighting instead of the Hessian eigenvalue-based weighting. The main difference between the two polarized intensity weighting maps is that the Hessian eigenvalue-based one has a more uniform weighting across different filaments than the $\text{H}\alpha$ -intensity-based one. The standard deviation divided by the mean of the logarithm of the Hessian eigenvalue-based weighting maps is ~ 0.1 and ~ 0.25 for the $\text{H}\alpha$ -intensity-based one. Also, for a wide filament with relatively abrupt edges, the Hessian eigenvalue-based weighting upweights those edges relative to the rest of the filament. This is not the case for the $\text{H}\alpha$ -intensity-based weighting.

We compare the EE and BB correlation ratios with the Planck Commander dust maps of these hybrid maps compared to those of the maps constructed with the original polarized intensity weighting in Figure 4.17 to isolate the effect of the weighting scheme. We find that the Hessian eigenvalue-based weighting increases the correlation at higher multipoles. This is especially the case for B modes, where the improvement is at the level of $\sim 10\%$. The improvement is less obvious in E modes, especially when the Hessian algorithm is used for the orientation angle calculation. In Figure 4.17, we repeat the exercise with the set of parameters that we found in Figure 4.14 to maximize the correlation with the Hessian-based template and confirm the same qualitative conclusion. These results are not latitude dependent. They are also consistent with our results in Figure 4.16, where we see that the improvement in the B -mode correlation at small scales for the Hessian-based template over the Spherical RHT-based template is more obvious than the E -mode correlation. Therefore, we attribute most of the enhancement in the B -mode correlation of the template constructed using the Hessian algorithm to the polarized intensity weighting applied to the different filaments in the Stokes Q and U maps. This indicates that the relative weighting of the filaments relative to one another affects B modes more than it affects E modes. We correlate the weighting maps directly with the Planck Commander dust polarized emission. The correlation with the HI-intensity-based weighting map is stronger at large scales and weaker at small scales compared to the correlation with the Hessian eigenvalue-based weighting. This is consistent with the results for the HI-based

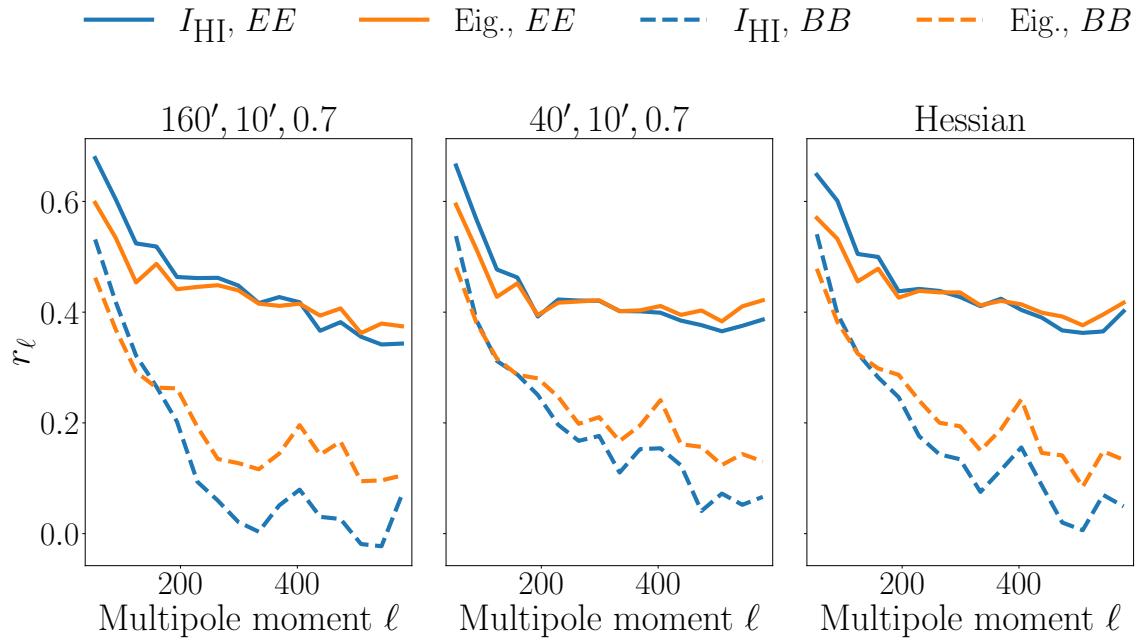


Figure 4.17: The EE (solid) and BB (dashed) correlation ratios of the Planck Commander dust maps with the $\text{H}\alpha$ -based polarization template constructed using different algorithms for the orientation angle calculations (different panels) and different weighting schemes (different colors). The left and middle panels use the Spherical RHT algorithm with different parameters for the orientation angle calculations, and the right panel uses the Hessian algorithm for those calculations. The parameters listed in the titles of the left and middle panels are the D_W , θ_{FWHM} , and Z , respectively, defined in Section 4.4.1. In each panel, the $\text{H}\alpha$ intensity-based polarized intensity weighting (blue) is compared with the Hessian eigenvalue-based polarized intensity weighting (orange).

polarization templates themselves.

4.8 Conclusions

We examine the impact of various alterations to $\text{H}\alpha$ -based polarization templates on the correlation with polarized millimeter-wave observations. This correlation probes the relationship between filamentary ISM structures and the magnetic field, and we investigate

what H_I structures are most predictive of the magnetic field orientation. We also use this framework to quantify the contribution of ISM filaments to the polarized dust emission power spectra. This is useful for CMB foreground separation. We make the H_I-based polarization templates discussed in this work publicly available at doi:10.7910/DVN/74MEMX.

We summarize the conclusions of this work below.

1. We improve the *B*-mode correlation between the Hessian-based template and the polarized dust emission by $\sim 5\%$ over that in Cukierman et al. (2023) by limiting the H_I velocity range used to $-13 \text{ km s}^{-1} < v_{\text{lsr}} < 16 \text{ km s}^{-1}$. This is similar to the LVC range proposed in Panopoulou & Lenz (2020a). The correlation with dust polarization is worse for wider velocity ranges when using the Hessian method because the Hessian is sensitive to artifacts in low-signal, high-absolute velocity channels.
2. We quantify the correlation between the Hessian-based template and the polarized dust emission in different masks of the sky and find it to be highest (at the $\sim 30\% - 60\%$ level) in the region between the Planck 20% and 40% sky fraction masks and lowest (at the $\sim 10\% - 20\%$ level) in the region that covers the 20% of the sky with the highest integrated dust intensity, i.e., the lowest Galactic latitudes.
3. We introduce the Spherical RHT algorithm, an efficient version of the RHT algorithm that uses spherical harmonic convolutions to run directly on the sphere. We find that the Spherical RHT fixes a spurious *EB* signal present at the $\sim 5\%$ level in the Clark

& Hensley (2019) H I-based polarization template due to projection effects.

4. We use the Spherical RHT to explore the parameter space of filament morphologies and their resulting polarization patterns. We find that the thinnest resolved H I filaments are the most informative for determining the magnetic field orientation. We also find that when using the Hessian method, the H I-based polarization template constructed from the GALFA-H I data smoothed to $7'$ correlates $\sim 10\%$ better with the *B*-mode polarized dust emission field than a template constructed from the $16'$ Hi4PI data. This motivates the use of even higher resolution H I data, such as the forthcoming Galactic Australian Square Kilometre Array Pathfinder (GASKAP; Dickey et al., 2013) and the Deep Synoptic Array (DSA-2000; Hallinan et al., 2019).
5. We use the Spherical RHT to demonstrate that parity-violating morphologies in the ISM can give rise to nonzero *TB* and *EB* even when local structures are perfectly aligned with the magnetic field. Since cosmic birefringence could lead to parity-odd polarization signals in the CMB, it is important to quantify parity-odd polarized dust emission as a foreground to those signals.
6. We apply the Spherical RHT to maps of synthetic filaments. We show how individual filament geometries mainly affect the *E*-mode pattern, whereas the positions and orientations of the filaments relative to one another mainly affect the *B*-mode pattern. We encourage the reader to use the Spherical RHT for exploring the polarization

signatures of other synthetic filament morphologies.

7. We compare the Spherical RHT- and Hessian-based polarization templates. We find that the most significant difference in the correlation with polarized dust emission is in B modes at small scales, where the Hessian-based template produces a higher correlation. We find that this is due to the difference in the polarized intensity weightings of the templates. The Spherical RHT-based template uses the H I intensity distribution, while the Hessian-based template uses the Hessian eigenvalue map, which tends to be more uniform. This indicates that B modes are more sensitive to the polarized intensity weighting of different filaments relative to one another than E modes are.

The correlation ratio between the integrated H I4PI intensity map and the Planck total intensity map at 353 GHz over the Planck 70% sky fraction Galactic plane mask decreases from $\sim 80\%$ at $\ell \sim 40$ to $\sim 35\%$ at $\ell \sim 600$ (Cukierman et al., 2023). This imperfect correlation is partly due to the cosmic infrared background (CIB) in the Planck total intensity map, as well as uncorrelated data noise and systematics. The correlation between the H I-based polarization templates and the dust polarization in EE decreases from $\sim 60\%$ at $\ell \sim 40$ to $\sim 45\%$ at $\ell \sim 600$. The strength of this correlation is partly affected by the fact that the H I-based polarization templates do not quantify the diffuse, non-filamentary component of the dust, nor the small local misalignments between the orientations of filaments and magnetic fields.

In addition to providing intuition on the filamentary polarized dust emission patterns, these conclusions provide a step forward in modeling dust polarization using H I data. This work has focused on comparisons to the observed polarized dust emission using polarization power spectra. Future work could consider additional metrics, including those sensitive to non-Gaussian structures in the dust, e.g., Minkowski functionals (Mantz et al., 2008) or the scattering transform (Mallat, 2011). These techniques have recently been used to quantify structures in dust (Delouis et al., 2022) and in H I emission (Lei & Clark, 2023b) individually.

4.9 Acknowledgments

This work was supported by the National Science Foundation under grant No. AST-2106607 (PI S.E.C.).

This publication utilizes data from Planck, an ESA science mission funded by ESA Member States, NASA, and Canada.

This work makes use of data from the Hi4PI Survey, which is constructed from the Effelsberg-Bonn H I Survey (EBHIS), made with the 100 m radio telescope of the MPIfR at Effelsberg/Germany, and the Galactic All-Sky Survey (GASS), observed with the Parkes Radio Telescope, part of the Australia Telescope National Facility, which is funded by the Australian Government for operation as a National Facility managed by CSIRO. EBHIS was funded by the Deutsche Forschungsgemein-schaft (DFG) under the grants KE757/7-1

to 7-3.

This work makes use of data from the Galactic Arecibo L -band Feed Array H I (GALFA-H I) Survey. It is with the Arecibo 305 m telescope, which is operated by SRI International under a cooperative agreement with the National Science Foundation (AST-1100968), and in alliance with Ana G. Méndez-Universidad Metropolitana and the Universities Space Research Association. The GALFA-H I surveys have been funded by the NSF through grants to Columbia University, the University of Wisconsin, and the University of California.

The computations in this paper were run on the Sherlock cluster, supported by the Stanford Research Computing Center at Stanford University.

Chapter 5

Transformer-Based Polarized Dust Emission Super-Resolution

Abstract

This study introduces a novel approach to generating high-resolution, non-Gaussian foreground models for cosmic microwave background (CMB) polarization studies. We develop a transformer-based model to increase the resolution of dust polarized emission images by a factor of 4, utilizing limited data and applying the same model across different resolutions. Our method fuses information from various sources, including Planck dust optical depth at 353 GHz (τ_{353}), and H I-based Stokes Q and U templates (Q^{HI} and U^{HI}), to predict small-scale dust structure. We quantify the relative importance of each input dataset, finding that τ_{353} , Q^{HI} , and U^{HI} contribute almost equally to the prediction of small-scale features. The model's attention map analysis

supports the assumption of scale-independence in dust polarization, consistent with the power-law approximation of dust polarization power spectra across angular scales. While our predictions serve as realistic non-Gaussian extrapolations useful as simulations, we emphasize that they may not represent actual small-scale dust polarized emission structure. This work contributes to the development of more accurate foreground models, potentially improving component separation, lensing reconstruction, and the detection of primordial B modes in future CMB polarization studies.

5.1 Paper Status and External Contributions

I wrote all the code, performed all the analysis, wrote all the text, and produced all the figures for this chapter. However, this work was performed under supervision from and the text has received extensive editorial input from my advisor Susan Clark.

5.2 Introduction

The cosmic microwave background (CMB) serves as a crucial probe into the early universe, offering insights into its initial conditions and subsequent evolution (Hu & Dodelson, 2002). While observations of CMB temperature anisotropies have yielded significant cosmological constraints, current research efforts are increasingly focused on measuring CMB polarization, particularly the B -mode component at large angular scales (Kamionkowski

& Kovetz, 2016). *B*-mode polarization is of particular interest due to its potential to detect primordial gravitational waves, a key prediction of inflationary theories (Guth, 1981; Linde, 1982). Numerous experiments, including ground-based (e.g., Liu et al., 2022; Ade et al., 2019; Abazajian et al., 2016; ACT Collaboration et al., 2024; SPT Collaboration et al., 2023; POLARBEAR Collaboration et al., 2022; BICEP/Keck Collaboration et al., 2021), balloon-borne (e.g., SPIDER Collaboration et al., 2022), and satellite-based (e.g., Collaboration et al., 2023) instruments, are actively pursuing inflationary *B*-mode detection.

However, the pursuit of primordial *B*-modes faces significant challenges. Gravitational lensing of CMB photons by large-scale structures converts some *E* modes to *B* modes that dominate over the primordial signal at small angular scales (Planck Collaboration et al., 2020h). While this lensing signal carries valuable cosmological information, it necessitates precise reconstruction and removal techniques to access the underlying primordial *B*-modes (Hu & Okamoto, 2002).

A more formidable obstacle is the presence of Galactic foregrounds, primarily thermal dust and synchrotron emission, which dominate the polarized signal at high ($\gtrsim 70$ GHz) and low ($\lesssim 70$ GHz) frequencies, respectively (Dunkley et al., 2009; Planck Collaboration et al., 2016e). These foregrounds exceed the expected CMB *B*-mode signal across all frequencies and sky positions, necessitating sophisticated component separation techniques for their removal (Delabrouille et al., 2009; Stompor et al., 2009).

Critically, Galactic foregrounds exhibit significant non-Gaussianity, particularly at large

scales (Ade et al., 2019). This non-Gaussianity is expected to persist at smaller scales due to the complex distribution of the interstellar medium (ISM) structure and turbulent interstellar magnetic fields. The presence of non-Gaussian foregrounds introduces mode coupling in angular power spectra and potentially biases both primordial *B*-mode detection and lensing reconstruction (Beck et al., 2020).

Current foreground models are limited by the lack of high-resolution, large-area observations. Existing templates, primarily derived from extrapolating Planck and Wilkinson Microwave Anisotropy Probe (WMAP) data to other frequency bands through assumptions about their spectral energy distributions (SEDs), characterize foregrounds down to approximately 1° resolution (Planck Collaboration et al., 2020f; Bennett et al., 2013). Simulation packages like the Python Sky Model (PySM) extrapolate these templates to smaller scales using power-law fits and Gaussian realizations (Thorne et al., 2017; Zonca et al., 2021). However, this approach fails to capture the expected non-Gaussianity at small angular scales.

Alternative approaches to foreground modeling include data-driven methods (Clark & Hensley, 2019; BICEP/Keck Collaboration et al., 2023e; Halal et al., 2024a), phenomenological models (Hervías-Caimapo & Huffenberger, 2022), and magnetohydrodynamic (MHD) simulations (Kim et al., 2019). While each offers unique insights, they face limitations in either reproducing the observed morphology or achieving high resolution efficiently.

Recent advancements in machine learning techniques offer promising avenues for generating high-resolution, non-Gaussian foreground simulations. Generative Adversarial Networks (GANs) have been employed to inject small-scale features into low-resolution dust observations while preserving statistical properties (Krachmalnicoff & Puglisi, 2021; Yao et al., 2024). However, these works use high-resolution dust total intensity rather than polarized emission maps as the ground truth of the output. They, therefore, assume that the statistical properties of the small-scale thermal dust emission in polarization are the same as in total intensity.

In this work, we pursue a different approach to produce realistic non-Gaussian high-resolution maps of the polarized dust emission at small scales. To accomplish this, we introduce several techniques. We utilize ancillary high-resolution datasets described in Section 5.3 as additional inputs to the model. We perform different smoothings and projections on multi-resolution maps as described in Section 5.4. In Section 5.5, we describe how our model fuses high-angular-resolution information from ancillary datasets with the low-angular-resolution information of the polarized dust emission maps to generate high-angular-resolution maps of the polarized dust emission. We describe our results in Section 5.6 and conclude in Section 5.7.

5.3 Data

In this section, we describe the different datasets used in the inputs and outputs of our model and explain the motivation behind using each of them.

5.3.1 Planck Data Products

For the polarized dust emission, we use the multi-resolution R3.00 Planck data at 353 GHz, processed with the Generalized Needlet Internal Linear Combination (GNILC; Remazeilles et al., 2011) method to eliminate the Cosmic Infrared Background (CIB) radiation from the Galactic dust emission (Planck Collaboration et al., 2016c). In line with the fiducial offset corrections used by the Planck collaboration, we add a Galactic offset correction of $63 \mu\text{K}_{\text{CMB}}$ then correct for the CIB monopole by subtracting $452 \mu\text{K}_{\text{CMB}}$ from the GNILC total intensity map (Planck Collaboration et al., 2020a). These maps are provided in the COSMO convention, and we use them in that convention. These data are used to produce the low-angular-resolution input and high-angular-resolution output Stokes Q and U images as explained in Section 5.4.

We also use the accompanying masks indicating the angular resolution of each sky region to determine the ideal methodology for dividing the sky into patches as explained in Section 5.4. The mask is shown in Figure 5.1 with the Galactic plane masked out.

We use the R1.20 Planck dust optical depth at 353 GHz (τ_{353} ; Planck Collaboration et al., 2014b) as an additional input to our model. Planck Collaboration et al. (2014b) fit

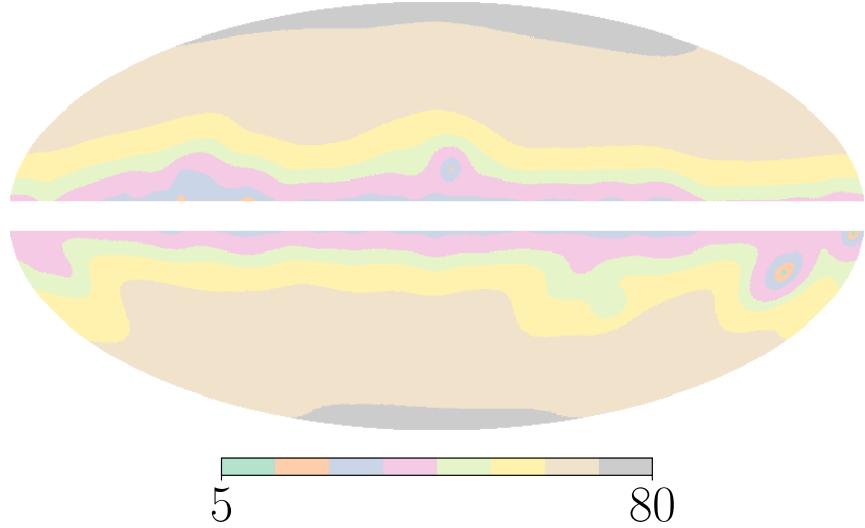


Figure 5.1: Mollweide map projection of the variable angular resolutions of the Planck GNILC polarization data. The colors on the colorbar correspond to 5', 7', 10', 15', 20', 30', 60', and 80'. The region with Galactic latitudes $|b| < 5^\circ$ is masked out.

a modified blackbody spectrum to Planck data at different frequencies to obtain a map of the optical depth. This is useful as a high-angular-resolution (FWHM=5') tracer of the dust spatial distribution at 353 GHz.

5.3.2 H_I-based Dust Polarization Templates

Galactic neutral hydrogen (H_I) emission serves as an invaluable proxy for studying dust polarization in the interstellar medium (ISM). The strong correlation between H_I and dust distributions in the diffuse ISM has been well-established (Boulanger et al., 1996b; Lenz et al., 2017). Both components exhibit filamentary structures that demonstrate significant alignment with the plane-of-sky magnetic field orientation (Clark et al., 2014, 2015).

The spectroscopic nature of H_I observations, utilizing the 21 cm line, provides three-dimensional information (longitude, latitude, and radial velocity) about the ISM structure (Clark, 2018). This characteristic offers a distinct advantage over traditional dust emission measurements. Additionally, H_I data are free from contamination by the cosmic infrared background (Chiang & Ménard, 2019) and are independent of broadband thermal dust emission observations, eliminating concerns about correlated instrumental systematics. Exploiting these properties, Clark & Hensley (2019) developed a model for dust polarization based solely on H_I intensity measurements. Their approach, utilizing the Rolling Hough Transform algorithm (Clark et al., 2014, 2020), demonstrated significant correlations with Planck 353 GHz data, particularly at large angular scales. Specifically, they observed correlations of approximately 60% and 50% for *E* modes and *B* modes, respectively, at multipole $\ell = 50$ over high Galactic latitudes, with the correlation diminishing towards zero at $\ell = 1000$. This methodology has proven valuable for characterizing dust properties, including spectral indices, through cross-correlations with millimeter-wave polarization data (BICEP/Keck Collaboration et al., 2023e).

Given H_I's advantages as a high-resolution tracer of the polarized dust emission, we utilize H_I-based Stokes *Q* and *U* templates (Q^{HI} and U^{HI}) as additional inputs to the model. Although there exist higher angular resolution H_I surveys than the H_I 4 π Survey (HI4PI; HI4PI Collaboration et al., 2016), such as the Galactic Arecibo L-Band Feed Array H_I Survey (GALFA-H_I; Peek et al., 2018) with an angular resolution of 4. $'$ 1, we

utilize templates constructed with Hi4PI data with an angular resolution of $16.^{\prime}2$ in this analysis. This is because we are limited by the highest angular resolution polarized dust emission maps we use, which are at $15'$, and the Hi4PI survey covers the full sky.

Hi4PI combines data from the Effelsberg-Bonn H α Survey (EBHIS; Winkel et al., 2016) and the Parkes Galactic All-Sky Survey (GASS; McClure-Griffiths et al., 2009). It achieves a spectral resolution of 1.49 km s^{-1} and a normalized brightness temperature noise of $\sim 53 \text{ mK}$ for a 1 km s^{-1} velocity channel.

For the H α -based polarization template we use the publicly available maps described in Halal et al. (2024a) constructed utilizing the version of the Hessian-based filament-finding algorithm described in Cukierman et al. (2023). These templates are divided into velocity bins with respect to the local standard of rest, inferred from the Doppler-shifted frequency along the line of sight. We use the version that integrates the templates along the range $-13 \text{ km s}^{-1} < v < 16 \text{ km s}^{-1}$, which is found by Halal et al. (2024a) to result in the highest correlation with the Planck polarization data at 353 GHz. We convert these maps to the COSMO convention to match the convention of the Planck GNILC maps by flipping the sign of the Stokes U template.

5.4 Pre-processing

We divide the sky into overlapping patches and project them onto square images to perform the training. For the ground truth, we do not use simulations or a factor derived from the

dust total intensity as in Krachmalnicoff & Puglisi (2021); Yao et al. (2024) in this analysis. Instead, we use the polarized dust emission data itself. We use the highest-resolution Planck GNILC data available as the high-resolution ground truth and an artificially smoothed version of this data as the input.

We train a model to increase the angular resolution of the input images by a factor of 4, no matter what the input angular resolution is. The motivation behind this approach is that the highest angular resolution ground truth data available has variable resolution across the sky as shown in Figure 5.1. Since deep learning models require a large sample size to train well, we train the same model on the different variable-resolution regions of the sky. To accomplish this, we vary the size of the projected patch based on its angular resolution, with larger patches for lower-angular-resolution regions and vice versa. Note that this approach assumes scale invariance. To weaken this assumption, we add an input to the network that is dependent on the resolution used for a given patch. This allows the network to learn a scale-dependent term and is explained further in Section 5.5.

5.4.1 Masks

We use the angular resolution map shown in Figure 5.1 to inform our projection and training strategies. Note from Figure 5.1 that the 5', 7', and 10' regions of the sky are not large enough to create enough training patches, so we limit the training to output 353 GHz Stokes Q and U images at 15', 20', 25', and 30', given the same images at 60', 80', 100', and 120',

respectively, i.e., a $4\times$ increase in resolution. We mask out Galactic latitudes $|b| < 5^\circ$ as shown in Figure 5.1. This is to remove regions associated with the inner Galactic plane. We exclude the $60'$ and $80'$ regions from the training samples. Therefore, the resulting sky area from which we project onto square patches covers about 40% of the sky.

For each angular resolution $\theta \in \{15', 20', 25', 30'\}$, we project patches of size $20\theta \times 20\theta$ to 80×80 pixel images. The longitudinal and latitudinal distances between neighboring patch centers are 8θ , i.e., twice the low angular resolution associated with the patches. We consider projections within 3 overlapping sky masks. The first is a combination of all the masks with an angular resolution equal to or higher than $15'$ and is used for predicting images with angular resolution $15'$, given images with angular resolution $60'$. The second is a combination of all the masks with an angular resolution higher than or equal to $20'$, which includes the first mask, and is used for predicting images with angular resolutions $20'$ and $25'$, given images with angular resolutions $80'$ and $100'$, respectively. The third and last is a combination of all the masks with an angular resolution higher or equal to $30'$, which includes the first and second masks. This is used for predicting images with angular resolution $30'$, given images with angular resolution $120'$.

5.4.2 Patch Projections

Within each of the 3 masks described in the previous subsection, we smooth the data to a constant resolution equal to the lowest angular resolution within that mask. For instance,

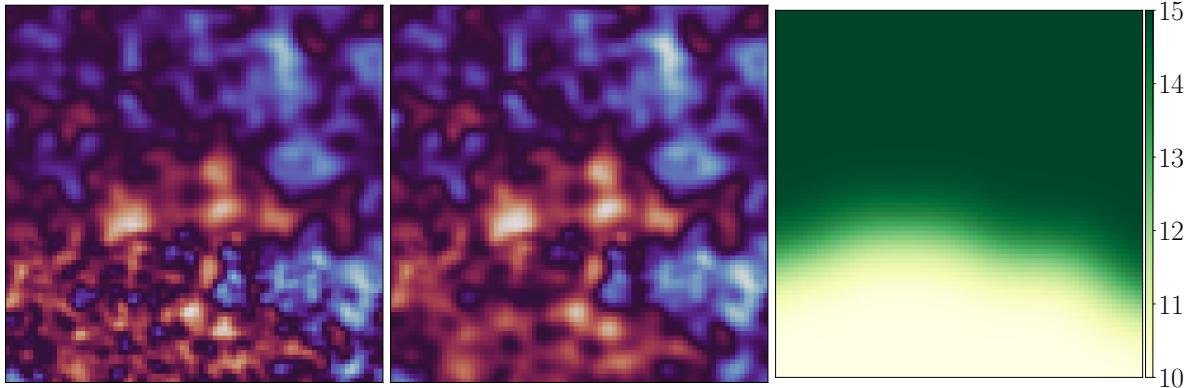


Figure 5.2: Example 80×80 pixel projected patch centered at $(l, b) = (210^\circ, -26^\circ)$ with a pixel width of $3.75'$. The projected patch spans both the $10'$ and $15'$ angular resolution regions of the variable resolution Planck GNILC map. Left: a projection of the original variable resolution GNILC Stokes Q map. Right: the Planck GNILC variable resolution mask, smoothed by a 1° Gaussian smoothing kernel. Middle: a weighted average of the map in the left panel and the same map smoothed with a Gaussian smoothing kernel from $10'$ to $15'$ according to the weights in the right panel, where $10'$ corresponds to the smoothed map and $15'$ corresponds to the original map.

we need to use the data in the regions corresponding to the $5'$, $7'$, and $10'$ resolutions as part of the mask used for predicting $15'$ resolution images. We start by smoothing the angular resolution map shown in Figure 5.1 using a Gaussian kernel with $\text{FWHM} = 1^\circ$ as shown in an example patch on the right in Figure 5.2. We do this to avoid sharp transitions when combining maps as follows.

We smooth the entire multi-resolution map with a Gaussian kernel that corresponds to the smoothing scale that would degrade a map resolution from $\text{FWHM}_{\text{current}}$ to $\text{FWHM}_{\text{desired}}$, i.e., a kernel with

$$\text{FWHM} = \sqrt{\text{FWHM}_{\text{desired}}^2 - \text{FWHM}_{\text{current}}^2}, \quad (5.1)$$

where $\text{FWHM}_{\text{desired}}$ is $7'$ and $\text{FWHM}_{\text{current}}$ is $5'$ in this case. This is useful in the $5'$ region, where the data are correctly smoothed to $7'$. We blend the original and smoothed maps in the regions corresponding to the $5'$ and $7'$ resolutions, using the smoothed resolution map to determine the weights for the weighted average. In the regions of the smoothed resolution map where the pixel values are $7'$ or higher, we use the original map. In the regions where the pixel values are $5'$, we use the map smoothed from $5'$ to $7'$. For the other pixel values between $5'$ and $7'$ in the smoothed resolution map, i.e., at the transition between the two angular resolution regions, we replace the pixels with a weighted average between the original and the smoothed data, where the weights are determined by the relative distances of the pixel values in the smoothed resolution map from $5'$ and $7'$. We iteratively repeat this process, combining the $7'$ region, which now includes the original $5'$ region, with the $10'$ region, and so on. An example patch where the $10'$ and $15'$ angular resolution regions are combined is shown in Figure 5.2. Through this process, we can project patches that overlap with regions of different angular resolution after having combined them into a common resolution.

When performing the patch projections for different projection sizes and centers, we ignore patches that contain regions outside the 3 aforementioned masks, i.e., if parts of the patch are at Galactic latitudes $|b| < 5^\circ$ or in angular resolution regions outside the defined mask. For the projection schema, we perform a bilinear interpolation with a zenithal equal area (ZEA) projection. This projection preserves the area of the sky region being projected

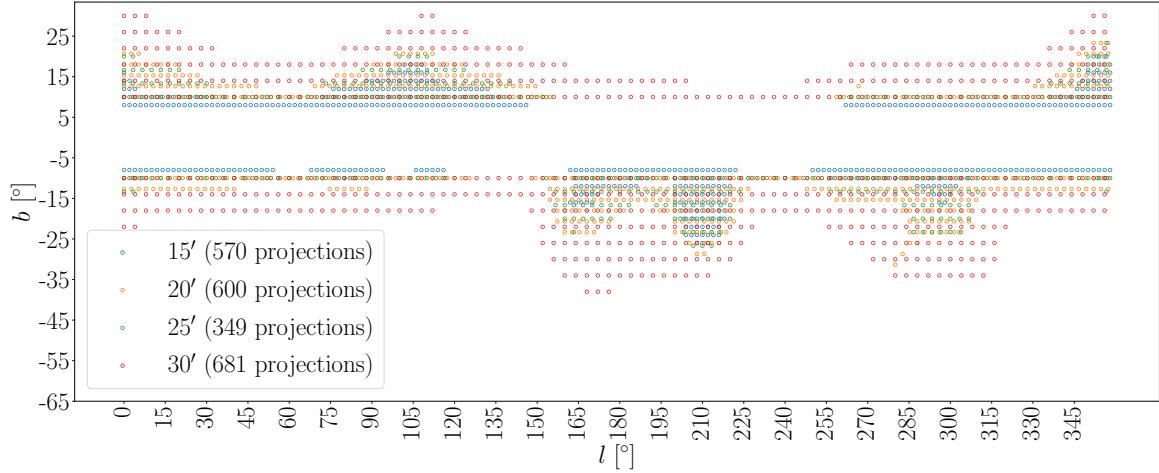


Figure 5.3: Scatter plot of the projection centers of patches used for different resolutions. The legend has the number of patches for each resolution. These are not the total number of patches used for the training since augmentations including flipping and rotations are also used.

while minimizing distortions near the center of the projection. Figure 5.3 shows the resulting projection centers for the patches corresponding to the different angular resolutions used in the training. We repeat this projection procedure for patches rotated by 45° . This is a form of augmentation used to increase the size of the dataset. Unlike 90° rotations, which can be performed after the projections, the 45° rotations include different parts of the sky. Therefore, not every 45° rotated counterpart of every patch is included in the training set if the rotated patch contains a part of the sky that is outside the mask. Figure 5.4 shows the resulting projection centers for the 45° rotated patches corresponding to the different angular resolutions used in the training.

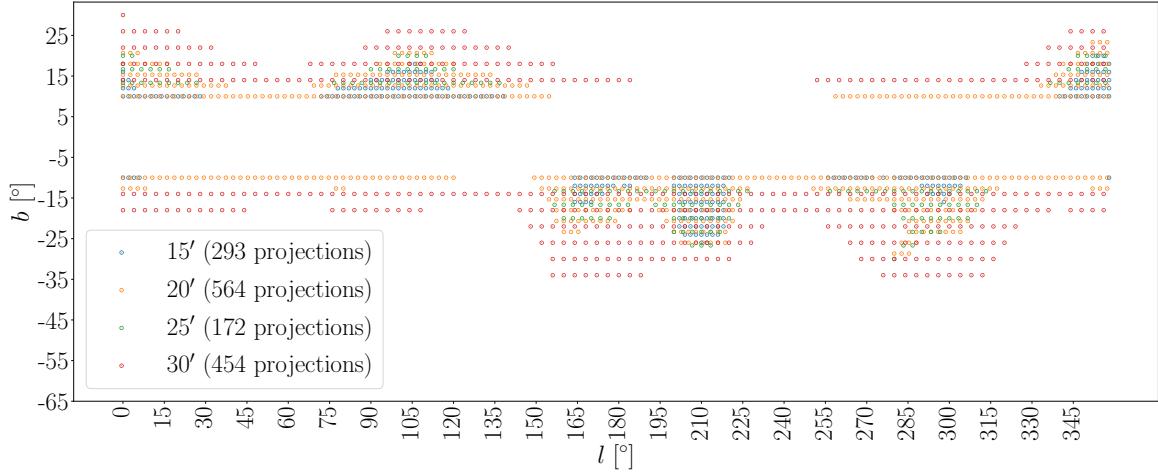


Figure 5.4: Scatter plot of the projection centers of patches rotated by 45° used for different resolutions. The legend has the number of patches for each resolution. These are not the total number of patches used for the training since augmentations including flipping and rotations are also used.

5.4.3 Training, Validation, and Testing Data

For each of the patches in Figures 5.3 and 5.4, we perform 7 augmentations, resulting in 8 images. These include 90° , 180° , and 270° rotations and a mirror flip for each rotation including the original image. Therefore, we have 3,683 patches with 8 augmentations for each, resulting in 29,464 training samples. Normally, not all augmentations are applied to all images in computer vision tasks. Instead, images are randomly augmented by the CPU before being passed to the GPU for training (Buslaev et al., 2020). This avoids overfitting the model on the training data. However, in this case, we do use all augmentations for all images because the number of images available is small otherwise. We test for overfitting by splitting the dataset into 80% (23,576 samples) for training, 10% (2,944 samples) for validation, and 10% (2,944 samples) for testing. We ensure all augmentations of the same

patch are in the same data split to avoid data leakage between the different splits.

Each training sample consists of 7 images and a scale-dependent term that indicates to the model the desired output angular resolution. The 5 images used as inputs in addition to the angular resolution are the smoothed low-resolution Planck GNILC Stokes Q and U 353 GHz images (Q_{353}^{LR} and U_{353}^{LR}), the H α -based Stokes Q and U templates (Q^{HI} and U^{HI}), and the optical depth at 353 GHz (τ_{353}). The output consists of the high-resolution Planck GNILC Stokes Q and U 353 GHz images (Q_{353}^{HR} and U_{353}^{HR}). Therefore, we have 206,248 images in total. We use float-32 as the data type. This corresponds to 80 pixels \times 80 pixels \times 2 bytes = 2.6 GB for all the images. Note that the input and output images have the same pixelization, so no downsampling and upsampling is required.

The reasoning behind including the τ_{353} map as an input is to provide the network with a high-resolution tracer of the spatial distribution of the dust at 353 GHz. Similarly, Q^{HI} and U^{HI} provide the network with a high-resolution polarization template that is highly correlated with the polarized dust emission.

For different input and output angular resolutions for the Planck GNILC 353 GHz Stokes images, we also use differently smoothed τ_{353} , Q^{HI} , and U^{HI} images. For the desired output resolution θ , we smooth the τ_{353} , Q^{HI} , and U^{HI} images according to Equation 5.1, where $\text{FWHM}_{\text{current}}$ is 5', 16.2', and 16.2', respectively, and $\text{FWHM}_{\text{desired}}$ is $\sqrt{\theta^2 - 15^2 + 5^2}$, $\sqrt{\theta^2 - 15^2 + 16.2^2}$ and $\sqrt{\theta^2 - 15^2 + 16.2^2}$, respectively. That is, when the desired output resolution is 15', we use the τ_{353} , Q^{HI} , and U^{HI} images at their native

resolutions. When the desired output resolution θ is a value larger than $15'$, the τ_{353} , Q^{HI} , and U^{HI} images are smoothed proportional to the distance between θ^2 and $15'^2$. The smoothing is done on the full-sky maps before projecting the maps into patches to avoid any edge effects.

5.4.4 Normalization

We normalize the images to a common scale to accelerate the optimization and convergence of the model. It is common in computer vision tasks to normalize images to the range $[0, 1]$ (Krizhevsky et al., 2012). However, since Stokes Q and U can be positive or negative, we scale, Q^{HI} , and U^{HI} to the range $[-1, 1]$, where 0 in the original maps remains 0 after the normalization. Since τ_{353} is positive, we scale those data to the range $[0, 1]$. This scaling is done on the sky map level rather than on the patch level, i.e., the ratio of the pixel values across different patches remains constant after scaling. We only consider pixels within the sky mask from which we perform the patch projections. Also, we scale the Stokes Q and U maps of the same dataset together, i.e., we consider the pixel values of both the Stokes Q and U maps together when scaling. This preserves the ratio of Q to U and thus the polarization angle. For the Planck GNILC data, we determine the normalization based on the unsmoothed maps, Q_{353}^{HR} and U_{353}^{HR} , and use the same normalization for the smoothed maps, Q_{353}^{LR} and U_{353}^{LR} . Therefore, we only need 3 values to perform the normalization for our datasets. These are the values that will map to 1 in the normalizations, while 0 remains 0

for all the maps. One is obtained from the distribution of the τ_{353} pixel values, one from the distribution of the Q_{353}^{HR} and U_{353}^{HR} pixel values, and one from the distribution of the Q^{HI} and U^{HI} pixel values. These 3 distributions contain outlier pixel values. We empirically find that the 99.9995th percentiles of the absolute values of these distributions separates the bulk of the distributions from these outlier pixel values in the tails of the distributions. This percentile corresponds to 10,020 outliers in each distribution. We set all pixel values with magnitudes higher than this percentile to the value at that percentile in the positive and negative directions before scaling. These thresholds are 1.0×10^{-3} K, 3.2×10^{-3} K, and 4.8×10^7 K km s⁻¹ for the τ_{353} data, the Q_{353}^{HR} and U_{353}^{HR} data, and the Q^{HI} and U^{HI} data, respectively.

5.5 Model

5.5.1 Super-Resolution Techniques

Image super-resolution, a fundamental task in computer vision, encompasses two primary approaches: single image super-resolution (SISR) and multi-image super-resolution (MISR). SISR techniques aim to enhance the resolution of a single low-quality image, while MISR methods leverage multiple low-resolution images of the same scene to reconstruct a high-resolution output. MISR is particularly prevalent in remote sensing applications, where satellites capture multiple temporal views of the same geographical area (e.g., An

et al., 2022).

Deep-learning-based MISR techniques have significantly advanced the field by introducing sophisticated mechanisms for fusing information from multiple low-resolution inputs. These fusion approaches vary in their architectural designs and information integration strategies (e.g., Deudon et al., 2020; Bordone Molini et al., 2020; Dorr, 2020). For instance, Arefin et al. (2020) employed a Gated Recurrent Unit (GRU), a type of recurrent neural network, as the fusion module.

More recent work has explored the use of transformers and the attention mechanism as fusion modules. The attention mechanism, originally introduced in natural language processing, allows a model to focus on different parts of the input when producing each part of the output. In the context of image processing, attention enables the model to selectively emphasize or suppress different spatial regions or feature channels. Mathematically, the attention mechanism can be expressed as

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{Q K^T}{\sqrt{d_k}} \right) V \quad (5.2)$$

where Q , K , and V are the query, key, and value matrices, respectively, and d_k is the dimension of the key vectors (Vaswani et al., 2017). The query, key, and value matrices are learnable linear weights applied to the input sequence.

Several studies have applied attention mechanisms to MISR tasks. Salvetti et al. (2020)

introduced a 3D convolutional feature attention mechanism with embedded residual connections, enabling the network to focus on extracting high-frequency information in both the temporal and spatial dimensions. Valsesia & Magli (2021) utilized self-attention mechanisms and permutation invariance for temporal images, allowing their model to process input frames in any order. An et al. (2022) were the first to apply a transformer, which is based on self-attention, to the remote sensing image MISR task. Li et al. (2023) extracted image features from different scales, and used attention in channel and spatial dimensions and across images, allowing their model to capture multi-scale information and inter-image relationships.

5.5.2 Architecture

Rather than training two separate networks, one for the Stokes Q and one for the Stokes U image, we train the same model on both. This is because information in the Stokes Q and U images are not physically independent. We utilize convolution-based encoders and decoders, that sandwich a transformer-based fusion module, which also takes in an embedding vector corresponding to the desired output resolution as an input. The overall architecture of the model we use in this work is shown in Figure 5.5. All of the convolutions used in this model have a kernel size of 3×3 pixels, a stride of 1 pixel, and a zero-padding of 1 pixel. The result of this convolution has the same dimensions as the original image. Therefore, the output images of this model have 80×80 pixels, which is the same as the

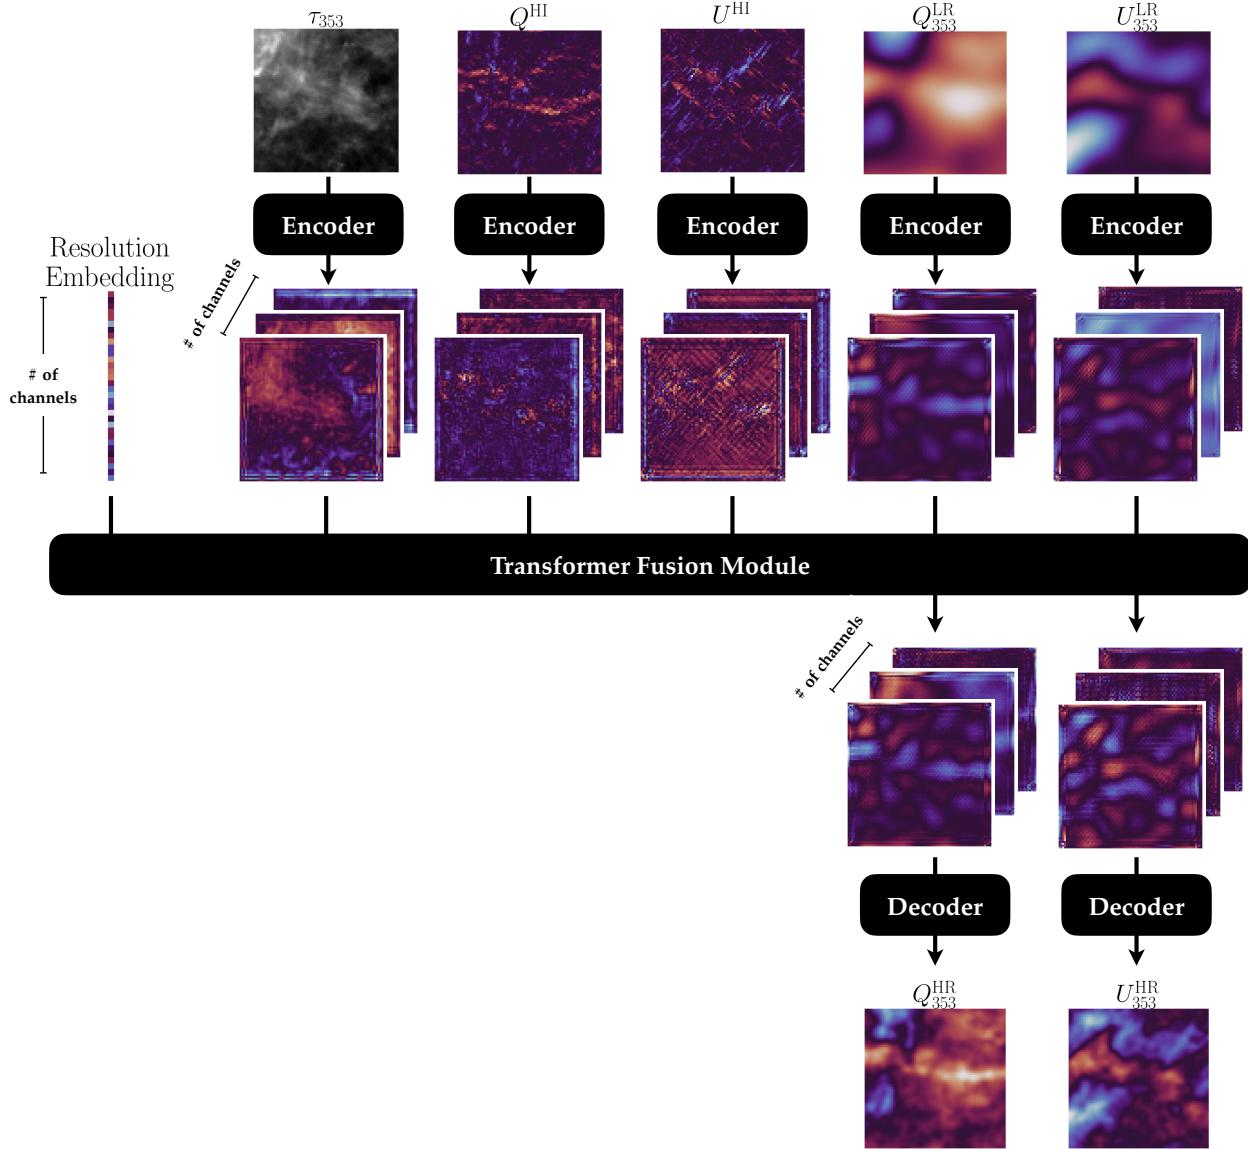


Figure 5.5: The full model architecture. The input images are each processed separately by an encoder, whose architecture is shown in Figure 5.6. The processed images along with an angular resolution embedding vector are then processed by a transformer fusion module, which consists of transformer layers (Figure 5.7). The transformer module's outputs corresponding to the Planck GNILC Stokes Q and U images are then processed through decoders, whose architecture is shown in Figure 5.6.

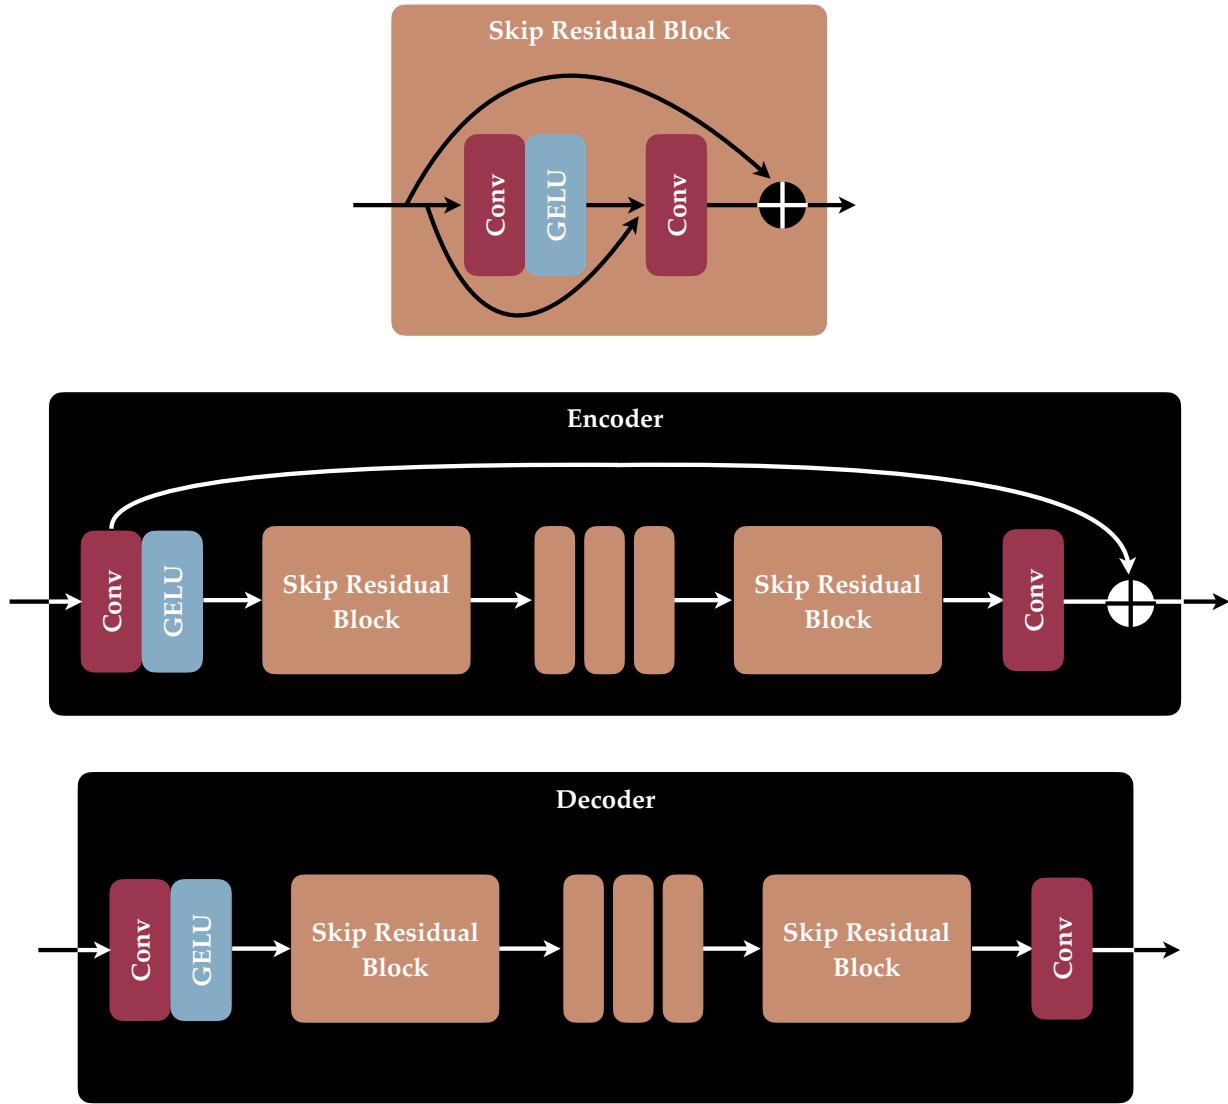


Figure 5.6: The architectures of the encoder (middle), of the decoder (bottom), and of their main component, the Skip Residual Block (top). The Skip Residual Block is made up of two convolutional layers with a non-linearity and a skip connection between them and a residual connection over the entire block. The encoder and decoder are made up of a convolutional layer, followed by a non-linearity, a number of Skip Residual Blocks, and a final convolutional layer. The encoder additionally has a residual connection from the output of the first convolutional layer to the end.

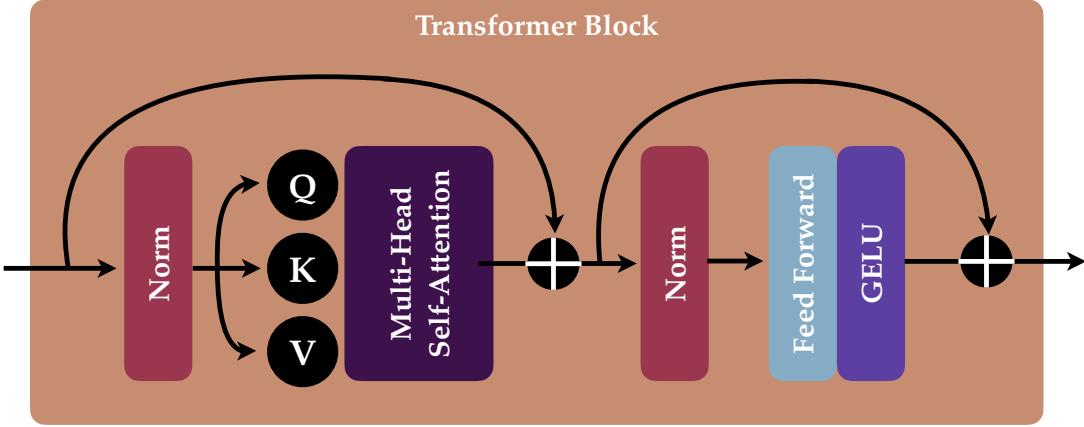


Figure 5.7: The transformer architecture used in this paper. It is made up of multi-head self-attention, where Q , K , and V are the query, key, and value tensors and a feed-forward layer followed by a non-linearity. Both components are preceded by a layer normalization and enveloped by a residual connection.

input images. No downsampling or upsampling is performed in this model. While each convolutional layer only has a receptive field of 3×3 pixels, stacking multiple convolutional layers in a sequence allows later convolutional layers to have a wider receptive field of the original images. We use the Gaussian Error Linear Unit (GELU; Hendrycks & Gimpel, 2016) as the non-linearity throughout this model. It applies a smooth approximation to the identity function, which enhances model performance by enabling better gradient flow during training.

The main building block in the encoder and decoder architectures is the Skip Residual Block. The architecture of this block is shown in Figure 5.6. It consists of convolutional layers, a non-linearity, and skip and residual connections. Skip and residual connections have several useful properties, including accelerating the convergence of the training process

and avoiding the common vanishing and exploding gradient problems (He et al., 2015). Skip connections concatenate the output of an earlier layer with the output of the current layer before feeding them into the following layer. Residual connections add the output of an earlier layer with the output of the current layer before feeding them into the following layer. The first convolutional layer has the same number of input and output channels, N_{channels} . Due to the concatenation performed before the second convolutional layer, the input of the second convolutional layer is double the size of the input of the first convolutional layer, i.e., $2 N_{\text{channels}}$, but its output has N_{channels} . Therefore, the residual connection is performed element-wise between two sets of N_{channels} . The non-linearity only follows the first convolutional layer.

The same patch of the sky for each dataset is processed by the same encoder architecture, resulting in a set of encoded images equal in number to the number of convolutional channels in the last convolutional layer of the encoder, N_{channels} . The encoder architecture is shown in Figure 5.6. The images are first passed through a convolutional layer with N_{channels} channels. The output of this layer is both passed through subsequent layers of the encoder and added to the output of the encoder through a residual connection. The results of the first convolutional layer are passed through a non-linearity, followed by a number of Skip Residual Blocks, $N_{\text{encoder blocks}}$. The output is then passed through a convolution layer with an equal number of input and output channels, N_{channels} .

The desired output resolution is embedded by an embedding layer into a vector with

size N_{channels} , equal to the number of channels in the last convolutional layer of the encoder. The embedding layer learns a separate embedding vector for each of the 4 angular resolutions used in the training: 15', 20', 25', and 30'.

A transformer module is used to fuse information from the different datasets. It accepts a sequence of 6 input vectors of size N_{channels} . The first vector in this sequence is always the embedding of the angular resolution. The five vectors after that correspond to the five input datasets. Each training sample is split into 80×80 training samples, one for each pixel. For each pixel, a vector of size N_{channels} is constructed from the pixels at the same location across the channel dimension in the output of the encoder. Therefore, the five vectors in the transformer sequence after the resolution embedding vector correspond to a pixel in the same location across the five input datasets. In practice, we treat the 80×80 pixels as different samples in a batch, i.e., the batch size is scaled by a factor of 80×80 and passed through the transformer in a single pass rather than stacking the results of 80×80 passes. The output of the transformer also has a sequence length of 6 vectors. We stack a number of transformer layers $N_{\text{transformer layers}}$ back-to-back. For the last transformer layer, we discard the first 4 outputs and only pass the output vectors corresponding to the Planck GNILC 353 GHz Stokes Q and U images through decoders.

The architecture of each transformer layer we use is shown in Figure 5.7. Its main components are the multi-head self-attention, the feed-forward layer, and residual connections. The feed-forward layer is followed by a non-linearity. The number of self-attention

Hyper-parameter	Value
N_{channels}	32
$N_{\text{encoder blocks}}$	4
$N_{\text{decoder blocks}}$	5
$N_{\text{transformer layers}}$	1
N_{heads}	1
Batch size	1
Learning rate	0.0001
Loss function	L1
Optimizer	AdamW ($\beta_1=0.9, \beta_2=0.999$)
N_{epochs}	209

Table 5.1: List of hyper-parameters used to train the model.

heads N_{heads} is a hyper-parameter. We apply layer normalization (Lei Ba et al., 2016) before each layer to control the scale of the gradients.

The decoder architecture is shown in Figure 5.6. It is very similar to the encoder architecture but with three differences. First, the number of Skip Residual Blocks $N_{\text{decoder blocks}}$ is another hyper-parameter that does not need to be equal to $N_{\text{encoder blocks}}$. Second the output number of channels of the last convolutional layer is 1, the target image. Because this is not equal to N_{channels} , the third difference is that we do not use a residual connection here.

5.6 Results

We experiment with different hyper-parameters and loss functions while attempting to overfit the model on 100 training samples. Once we find a set of hyper-parameters for which the model is able to overfit, we use those hyper-parameters for training the model on

the full training dataset. The set of hyper-parameters we use, including the loss function, are summarized in Table 5.1. With these hyper-parameters, the model has 928,194 tunable parameters. We find that increasing the batch size degrades the model’s performance so we use a batch size of 1. We also see no improvement in the model’s loss when increasing $N_{\text{transformer layers}}$ and N_{heads} from 1 to 2 each for the same number of epochs. Therefore, we use 1 for both since that decreases the number of tunable parameters and enables interpretability of the attention weights.

While training this model, we save its weights whenever the validation loss reaches a new minimum. We also log the output of the same 4 validation patches (1 for each angular resolution) once every 50 training steps. We stopped the training when those model outputs visually look similar enough to the ground truth with no visible artifacts. The last saved model checkpoint is at epoch 209 after the training is stopped. We note that the validation loss was still decreasing when we stopped the training. The model can, therefore, be trained for longer for more optimal results. Using the last saved model weights, we show the model’s outputs for 4 randomly selected test patches, 1 for each angular resolution, in Figure 5.8. Note that there is no visible distinction in the model’s performance on the different angular resolutions.

The attention weights, i.e., the softmax $\left(\frac{Q K^T}{\sqrt{d_k}}\right)$ part of Equation 5.2, is useful for interpreting how much attention each input element receives from each output element. The softmax ensures that these weights sum to 1 across all positions, effectively creating

a probability distribution. Because the query and key matrices are dependent on the input encoded images, the attention weights vary pixel-by-pixel and patch-by-patch. Therefore, to visualize the attention map, we average the attention weights over all the pixels and over 100 different patches. We find no measurable difference in the attention weights between the different angular resolution patches. Therefore, for the 100 different patches we average over for visualizing the attention map in Figure 5.9, we select 25 randomly from each angular resolution.

Figure 5.9 shows how important each of the 6 inputs to the transformer block are to each of the 2 outputs that are passed through decoders. The importance is quantified as weights that sum to 1. Note that the encoder outputs corresponding to Q_{353}^{LR} and U_{353}^{LR} are the least important for predicting U_{353}^{HR} and Q_{353}^{HR} , respectively. We also note that the angular resolution embedding is also unimportant. If the resolution embedding adequately captures the scale dependence in the data, then this result supports the assumption of scale independence. Conversely, if the resolution embedding is not sufficient, the model would be imposing scale independence and the result is evidence that the imposed scale independence is working as intended. The most important encoder outputs for predicting Q_{353}^{HR} and U_{353}^{HR} are those corresponding to Q_{353}^{LR} and U_{353}^{LR} , respectively. This is expected since the output images are just higher-resolution versions of the input images. It is interesting that both Q^{HI} and U^{HI} are important for predicting both Q_{353}^{HR} and U_{353}^{HR} , and τ_{353} is also as important. This result highlights the importance of these datasets in predicting dust structure at small

scales. It is still the case, however, that Q^{HI} is slightly more informative than U^{HI} for predicting Q_{353}^{HR} and vice versa, as expected.

5.7 Conclusions, Limitations, and Future Work

The development of accurate high-resolution foreground models is crucial for the future of CMB polarization studies. Such models will enable more robust component separation, improved lensing reconstruction, and ultimately, a clearer path to detecting primordial B modes. In this work, we contribute to this effort by exploring novel techniques for generating realistic, non-Gaussian foreground simulations at small angular scales. We demonstrated how the attention mechanism can be used to fuse information from images of different sources to increase the resolution of polarized dust emission images by a factor of 4. We also demonstrated how to achieve this with limited data, using the same model for different resolutions. We quantify the importance of each input dataset in predicting the small-scale dust structure and find the τ_{353} , Q^{HI} , and U^{HI} to all be almost as important as the low-resolution version of the map.

Given that no ground truth is available for the actual small-scale structure of the dust polarized emission, our predictions are only to be considered as realistic non-Gaussian extrapolations from large-scale structure that follow the correct statistics. We emphasize that these predictions are to be used as simulations rather than actual predictions of the small-scale structure of the dust polarized emission. One of the weak assumptions made in our

modeling approach is that the mapping between two scales which are a multiplicative factor apart is similar to the mapping between two different scales with the same multiplicative factor between them. We weaken this assumption by adding a scale-dependent term as an input to the model. However, by studying the attention map of our model, we find that this term is given little attention. This either implies that the assumption of scale independence is justified or that encoding the scale dependence in this term is insufficient for capturing the scale dependence in the data. This is in line with the fact that the power spectra of the observed dust polarization can, at first order, be approximated as a power law as a function of angular scales (Córdova Rosado et al., 2024).

The next step in this analysis is to run the model on patches of the entire sky to create a higher-resolution version of the Planck GNILC map. Given that the full sky Planck GNILC map includes regions with angular resolutions other than the 4 resolutions we train our model on and the fact that the attention weights corresponding to the resolution embedding are very low, we could retrain the model without this embedding.

Several options are possible for extending the analysis described in this chapter. The main one is to turn this model into a generator by adding randomness to the input images, adding a critic model, and performing adversarial training using techniques like the Wasserstein generative adversarial network with gradient penalty (WGAN-GP). This would allow the model to generate different possible small-scale predictions for the same patch of sky, which would be useful when different realizations are needed. Moreover, we have

not performed extensive hyper-parameter tuning of the model architecture. We only experimented with different configurations until we found one that produces satisfactory results. Therefore, extensive hyper-parameter tuning, including experimenting with different loss functions and image quality metrics, can be helpful for producing the optimal predictions. Another possible analysis extension would be to experiment with changes to the input datasets. For instance, the Stokes Q and U images can be combined into $Q + iU$ images with complex-valued neural networks. Another example is that the integrated H α -based polarization templates can be replaced with different H α -based polarization templates divided across the velocity dimension. Additionally, the τ_{353} data could be replaced with data from the WISE experiment, which measures emission from polycyclic aromatic hydrocarbon (PAH) molecules at 12 microns (Meisner & Finkbeiner, 2014). These molecules are intermixed with the dust in the interstellar medium (Córdova Rosado et al., 2024). The benefit of this data is its angular resolution at 5''. The possibilities listed here may not improve the model's predictions but are worth experimenting with in case they do.

5.8 Acknowledgements

This work was supported by the National Science Foundation under grant No. 2106607. This work utilizes data from Planck, an ESA science mission funded by ESA Member States, NASA, and Canada.

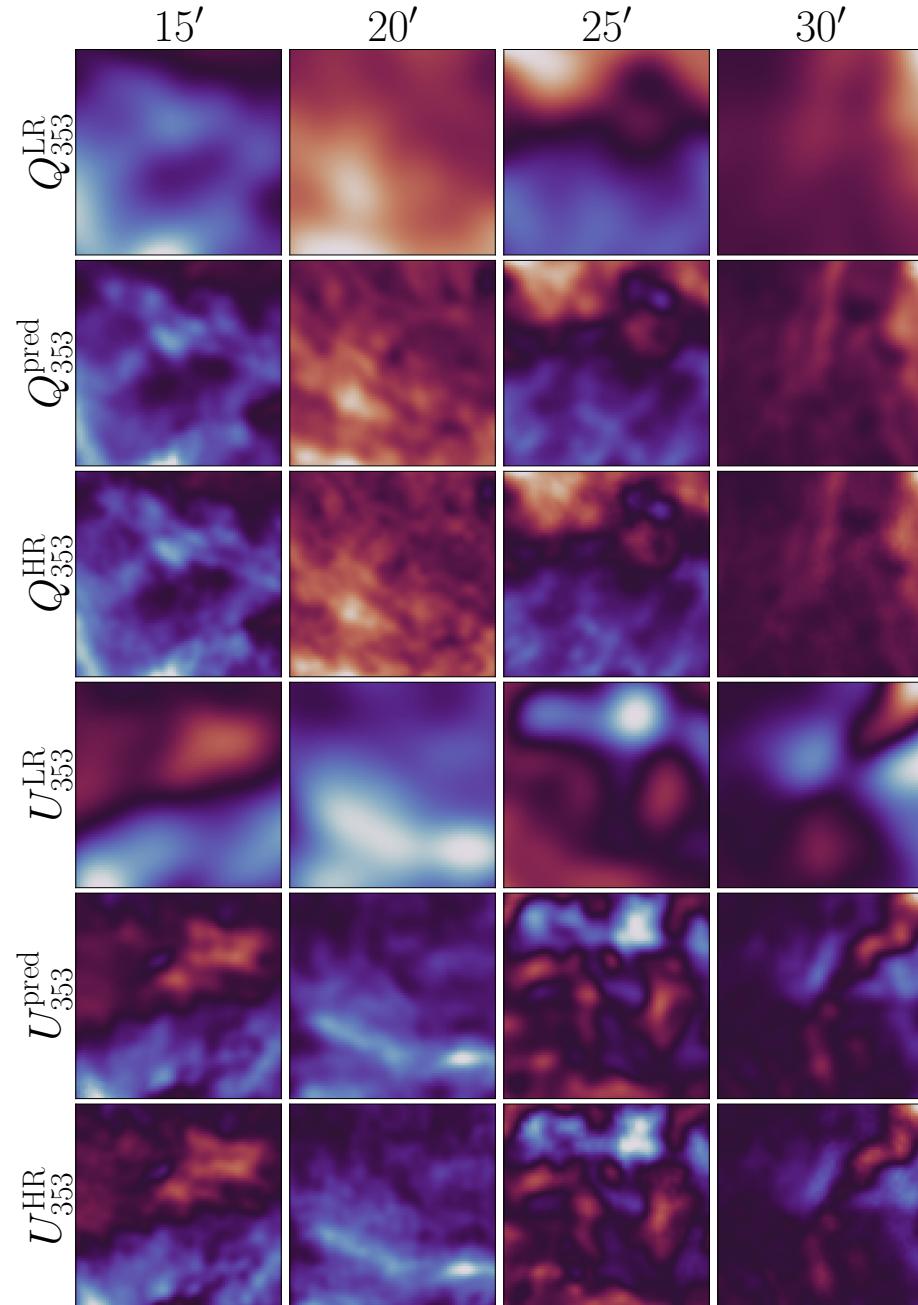


Figure 5.8: Example low-resolution input and high-resolution predictions and target 353 GHz Stokes Q (top 3 rows) and U (bottom 3 rows) patches of sky. The same patch of sky is shown across each column with its corresponding high angular resolution denoted at the top. The colorbars are centered at zero (darkest) and brighter red (blue) corresponds to higher positive (negative) values.

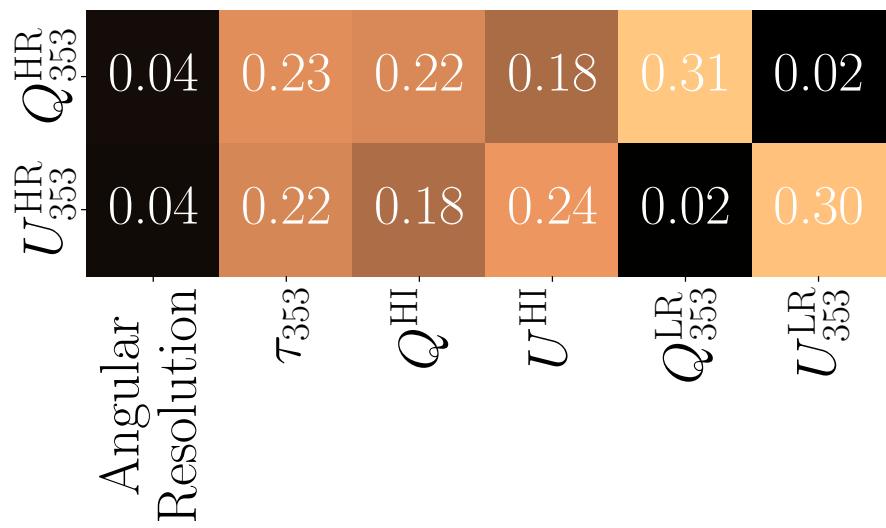


Figure 5.9: Attention map averaged over 80×80 pixels and over 100 test sample patches, 25 for each angular resolution. The values in each row sum up to 1 because of the softmax operation from Equation 5.2.

Chapter 6

Conclusion

This Thesis establishes an important connection between the 3D morphology of the interstellar medium (ISM) and the polarized dust emission, leveraging ancillary datasets that trace the 3D structure of the ISM to characterize, model, and enhance our understanding of polarized dust emission. This work is significant not only for advancing our knowledge of astrophysical processes but also for improving our ability to study the early universe through observations of the cosmic microwave background (CMB).

6.1 Summary of Results

The research presented here demonstrates the importance of considering the 3D geometry of the ISM when interpreting polarized dust emission data. Chapter 2 reveals that the complexity of the 3D dust distribution along the line of sight influences the observed dust

polarization fraction. Sightlines with dust distributed across multiple distance components exhibit, on average, lower polarization fractions compared to those where the dust is more concentrated. This finding highlights the limitations of models that assume a simple relationship between the Local Bubble geometry and the observed polarization patterns.

Chapter 3 introduces a novel approach to characterizing Galactic dust filaments by correlating BICEP/Keck and Planck data with 3D polarization templates based on neutral hydrogen ($\text{H}\,\text{i}$) observations. This method proves effective in detecting polarized dust emission down to 95 GHz and in identifying which ISM components contribute measurably to the polarized dust emission. We demonstrate this in the BICEP/Keck patch, where we find no detectable contribution from Magellanic Stream 1 and isolate the detected polarized emission to dust associated with the Milky Way. The analysis also reveals that the dust associated with filamentary structures shares a similar spectral energy distribution with the overall dust field in the region, suggesting no significant decorrelation between filamentary and non-filamentary dust components.

Chapter 4 further refines the $\text{H}\,\text{i}$ -based polarization templates by exploring the impact of filament morphology quantification on their correlation with polarized dust emission. The introduction of the Spherical Rolling Hough Transform (Spherical RHT) algorithm enables efficient characterization of filamentary structures on the sphere. This work demonstrates that the thinnest resolved $\text{H}\,\text{i}$ filaments are the most informative for determining magnetic field orientation, and higher-resolution $\text{H}\,\text{i}$ data can lead to improved correlations with

polarized dust emission. The chapter also demonstrates that the relative topology and polarized intensities of these filaments predominantly influence the *B*-mode patterns in the dust emission, while the shapes of individual filaments affect the *E*-mode patterns. We also show how, even when filaments are perfectly aligned with the local magnetic field, certain shapes can produce parity-violating signatures. This has implications for searches for cosmic birefringence.

The Thesis culminates in Chapter 5 with the development of a transformer-based model for generating high-resolution, non-Gaussian foreground models of the polarized dust emission. This innovative approach fuses information from various sources, including Planck dust optical depth and H I-based Stokes *Q* and *U* templates, to predict small-scale dust structures. The model's success in extrapolating realistic non-Gaussian features from large-scale structures, while maintaining consistent statistics, represents a considerable advancement in foreground modeling for CMB studies.

Throughout this work, the consistent theme emerges that the 3D structure of the ISM plays a fundamental role in shaping the observed polarized dust emission. The results from Chapter 2, which demonstrate the influence of dust distribution complexity on polarization fraction, motivate the subsequent chapters' focus on incorporating 3D information through H I data to model and study polarized dust emission. This approach proves fruitful, as evidenced by the improved correlations and insights gained in Chapters 3 and 4, and the successful generation of high-resolution models in Chapter 5.

6.2 Outlook

The methodologies and findings presented in this Thesis have significant implications for future CMB polarization studies, particularly in the quest to detect primordial gravitational waves through *B*-mode polarization. By providing more accurate and detailed models of the polarized dust foreground, this work contributes to improving component separation techniques, enhancing lensing reconstruction, and ultimately increasing the sensitivity of searches for primordial *B*-modes.

Moreover, the techniques developed here, such as the Spherical RHT algorithm and the transformer-based super-resolution model, offer versatile tools that can be applied to a wide range of astrophysical problems beyond dust polarization studies. These methods open new avenues for exploring the intricate relationships between ISM structure, magnetic fields, and observable phenomena across various scales and environments.

This Thesis advances our understanding of the complex interplay between the 3D structure of the ISM and polarized dust emission, providing both theoretical insights and practical tools for future research. By bridging the gap between high-resolution tracers of ISM morphology and the observed polarized dust emission, this work lays a foundation for more precise foreground modeling in CMB studies and deeper exploration of astrophysical processes in the Galactic ISM.

Bibliography

Abazajian, K. N., Adshead, P., Ahmed, Z., et al. 2016, arXiv e-prints, arXiv:1610.02743,
doi: 10.48550/arXiv.1610.02743

Aabitbol, M. H., Hill, J. C., & Johnson, B. R. 2016, Mon. Not. R. Astron. Soc., 457, 1796,
doi: 10.1093/mnras/stw030

ACT Collaboration, Madhavacheril, M. S., Qu, F. J., et al. 2024, *Astrophys. J.*, 962, 113,
doi: 10.3847/1538-4357/acff5f

Ade, P., Aguirre, J., Ahmed, Z., et al. 2019, *Journal of Cosmology and Astroparticle Physics*,
2019, 056

Aharonian, F., Yang, R., & de Oña Wilhelmi, E. 2019, *Nature Astronomy*, 3, 561, doi: 10.
1038/s41550-019-0724-0

Alves, M. I. R., Boulanger, F., Ferrière, K., & Montier, L. 2018, *Astron. Astrophys.*, 611,
L5, doi: 10.1051/0004-6361/201832637

An, T., Zhang, X., Huo, C., et al. 2022, IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 15, 1373, doi: 10.1109/JSTARS.2022.3143532

Anders, F., Khalatyan, A., Chiappini, C., et al. 2019, Astron. Astrophys., 628, A94, doi: 10.1051/0004-6361/201935765

Andersson, B. G., Lazarian, A., & Vaillancourt, J. E. 2015, Annu. Rev. Astron. Astrophys., 53, 501, doi: 10.1146/annurev-astro-082214-122414

Andersson, B. G., & Potter, S. B. 2006, Astrophys. J. Lett., 640, L51, doi: 10.1086/503199

Andrae, R., Rix, H.-W., & Chandra, V. 2023, Astrophys. J. Suppl. Ser., 267, 8, doi: 10.3847/1538-4365/acd53e

Arefin, M. R., Michalski, V., St-Charles, P.-L., et al. 2020, in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, 206–207

Audit, E., & Hennebelle, P. 2005, Astron. Astrophys., 433, 1, doi: 10.1051/0004-6361:20041474

Barnes, A. T., Watkins, E. J., Meidt, S. E., et al. 2023, Astrophys. J. Lett., 944, L22, doi: 10.3847/2041-8213/aca7b9

Baumann, D., & Peiris, H. V. 2008, arXiv e-prints, arXiv:0810.3022, doi: 10.48550/arXiv.0810.3022

Beck, D., Cukierman, A., & Wu, W. L. K. 2022, *Mon. Not. R. Astron. Soc.*, 515, 229, doi: 10.1093/mnras/stac1775

Beck, D., Errard, J., & Stompor, R. 2020, *Journal of Cosmology and Astroparticle Physics*, 2020, 030

Beck, R. 2001, *Space Sci. Rev.*, 99, 243, doi: 10.1023/A:1013805401252

—. 2015, *A&A Rev.*, 24, 4, doi: 10.1007/s00159-015-0084-4

Benjamin, R. A. 2005, *Unsolved Mysteries of High-velocity Clouds* (Dordrecht: Springer Netherlands), 371–389, doi: 10.1007/1-4020-2579-3_17

Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, *The Astrophysical Journal Supplement Series*, 208, 20

Benoît, A., Ade, P., Amblard, A., et al. 2004, *Astron. Astrophys.*, 424, 571, doi: 10.1051/0004-6361:20040042

Berdyugin, A., Piirola, V., & Teerikorpi, P. 2014, *Astron. Astrophys.*, 561, A24, doi: 10.1051/0004-6361/201322604

Bernardi, G., de Bruyn, A. G., Brentjens, M. A., et al. 2009, *Astron. Astrophys.*, 500, 965, doi: 10.1051/0004-6361/200911627

Bialy, S., & Sternberg, A. 2019, *Astrophys. J.*, 881, 160, doi: 10.3847/1538-4357/ab2fd1

Bialy, S., Zucker, C., Goodman, A., et al. 2021, *Astrophys. J. Lett.*, 919, L5, doi: 10.3847/2041-8213/ac1f95

BICEP2 Collaboration, Ade, P. A. R., Aikin, R. W., et al. 2014, *Phys. Rev. Lett.*, 112, 241101, doi: 10.1103/PhysRevLett.112.241101

BICEP2 Collaboration, Keck Array Collaboration, Ade, P. A. R., et al. 2016, *Astrophys. J.*, 825, 66, doi: 10.3847/0004-637X/825/1/66

BICEP2/Keck Collaboration, Planck Collaboration, Ade, P. A. R., et al. 2015, *Phys. Rev. Lett.*, 114, 101301, doi: 10.1103/PhysRevLett.114.101301

BICEP/Keck Collaboration, Cornelison, J., Ade, P. A. R., et al. 2020, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 11453, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy X, ed. J. Zmuidzinas & J.-R. Gao, 1145327, doi: 10.1117/12.2562905

BICEP/Keck Collaboration, Ade, P. A. R., Ahmed, Z., et al. 2021, *Phys. Rev. Lett.*, 127, 151301, doi: 10.1103/PhysRevLett.127.151301

BICEP/Keck Collaboration, Cornelison, J., Vergès, C., et al. 2022a, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12190, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI, ed. J. Zmuidzinas & J.-R. Gao, 121901X, doi: 10.1117/12.2620212

BICEP/Keck Collaboration, Goldfinger, D. C., Ade, P. A. R., et al. 2022b, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12190, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI, ed. J. Zmuidzinas & J.-R. Gao, 121901V, doi: 10.1117/12.2629490

BICEP/Keck Collaboration, Soliman, A., Ade, P. A. R., et al. 2022c, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12190, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI, ed. J. Zmuidzinas & J.-R. Gao, 1219014, doi: 10.1117/12.2628058

—. 2022d, arXiv e-prints, arXiv:2203.16556, doi: 10.48550/arXiv.2203.16556

—. 2022e, Phys. Rev. D, 105, 022006, doi: 10.1103/PhysRevD.105.022006

—. 2022f, Astrophys. J., 927, 77, doi: 10.3847/1538-4357/ac4886

—. 2023a, Astrophys. J., 945, 72, doi: 10.3847/1538-4357/acb64c

BICEP/Keck Collaboration, Dierickx, M., Ade, P. A. R., et al. 2023b, Journal of Low Temperature Physics, 211, 366, doi: 10.1007/s10909-023-02967-1

BICEP/Keck Collaboration, Schillaci, A., Ade, P. A. R., et al. 2023c, Journal of Low Temperature Physics, 213, 317, doi: 10.1007/s10909-023-03005-w

BICEP/Keck Collaboration, Ade, P. A. R., Ahmed, Z., et al. 2023d, Astrophys. J., 949, 43, doi: 10.3847/1538-4357/acc85c

—. 2023e, *Astrophys. J.*, 945, 72, doi: 10.3847/1538-4357/acb64c

—. 2024a, arXiv e-prints, arXiv:2405.19469, doi: 10.48550/arXiv.2405.19469

BICEP/Keck Collaboration, Fatigoni, S., Ade, P. A. R., et al. 2024b, *Journal of Low Temperature Physics*, doi: 10.1007/s10909-024-03100-6

Boe, B., Habbal, S., & Druckmüller, M. 2020, *Astrophys. J.*, 895, 123, doi: 10.3847/1538-4357/ab8ae6

Bohlin, R. C., Savage, B. D., & Drake, J. F. 1978, *Astrophys. J.*, 224, 132, doi: 10.1086/156357

Bordone Molini, A., Valsesia, D., Fracastoro, G., & Magli, E. 2020, *IEEE Transactions on Geoscience and Remote Sensing*, 58, 3644, doi: 10.1109/TGRS.2019.2959248

Boulanger, F., Abergel, A., Bernard, J. P., et al. 1996a, *Astron. Astrophys.*, 312, 256

—. 1996b, *Astron. Astrophys.*, 312, 256

Breitschwerdt, D. 2001, *Ap&SS*, 276, 163

Breitschwerdt, D., Feige, J., Schulreich, M. M., et al. 2016, *Nature*, 532, 73, doi: 10.1038/nature17424

Burton, W. B. 1988, in *Galactic and Extragalactic Radio Astronomy*, ed. K. I. Kellermann & G. L. Verschuur, 295–358

Buslaev, A., Iglovikov, V. I., Khvedchenya, E., et al. 2020, *Information*, 11, doi: 10.3390/info11020125

Cabral, B., & Leedom, L. C. 1993, in *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '93* (New York, NY, USA: Association for Computing Machinery), 263–270, doi: 10.1145/166117.166151

Cahlon, S., Zucker, C., Goodman, A., Lada, C., & Alves, J. 2023, arXiv e-prints, arXiv:2308.14794, doi: 10.48550/arXiv.2308.14794

Capitanio, L., Lallement, R., Vergely, J. L., Elyajouri, M., & Monreal-Ibero, A. 2017, *Astron. Astrophys.*, 606, A65, doi: 10.1051/0004-6361/201730831

Carrière, J. S., Ferrière, K., Ristorcelli, I., & Montier, L. 2022, *Astron. Astrophys.*, 668, A42, doi: 10.1051/0004-6361/202244550

CCAT-Prime collaboration, Aravena, M., Austermann, J. E., et al. 2021, arXiv e-prints, arXiv:2107.10364. <https://arxiv.org/abs/2107.10364>

Challinor, A., Fosalba, P., Mortlock, D., et al. 2000, *Phys. Rev. D*, 62, 123002, doi: 10.1103/PhysRevD.62.123002

Chandrasekhar, S., & Fermi, E. 1953, *Astrophys. J.*, 118, 113, doi: 10.1086/145731

Chen, B. Q., Huang, Y., Yuan, H. B., et al. 2019a, *Mon. Not. R. Astron. Soc.*, 483, 4277, doi: 10.1093/mnras/sty3341

Chen, C.-Y., King, P. K., Li, Z.-Y., Fissel, L. M., & Mazzei, R. R. 2019b, *Mon. Not. R. Astron. Soc.*, 485, 3499, doi: 10.1093/mnras/stz618

Chen, C.-Y., Li, Z.-Y., Mazzei, R. R., et al. 2022, *Mon. Not. R. Astron. Soc.*, 514, 1575, doi: 10.1093/mnras/stac1417

Chiang, Y.-K., & Ménard, B. 2019, *Astrophys. J.*, 870, 120, doi: 10.3847/1538-4357/aaf4f6

Chluba, J., Hill, J. C., & Abitbol, M. H. 2017, *Mon. Not. R. Astron. Soc.*, 472, 1195, doi: 10.1093/mnras/stx1982

Cho, J., & Yoo, H. 2016, *Astrophys. J.*, 821, 21, doi: 10.3847/0004-637X/821/1/21

Clark, S. E. 2018, *Astrophys. J. Lett.*, 857, L10, doi: 10.3847/2041-8213/aabb54

Clark, S. E., & Hensley, B. S. 2019, *Astrophys. J.*, 887, 136, doi: 10.3847/1538-4357/ab5803

Clark, S. E., Hill, J. C., Peek, J. E. G., Putman, M. E., & Babler, B. L. 2015, *Phys. Rev. Lett.*, 115, 241302, doi: 10.1103/PhysRevLett.115.241302

Clark, S. E., Kim, C.-G., Hill, J. C., & Hensley, B. S. 2021, *Astrophys. J.*, 919, 53, doi: 10.3847/1538-4357/ac0e35

Clark, S. E., Peek, J., Putman, M., Schudel, L., & Jaspers, R. 2020, RHT: Rolling Hough Transform. <http://ascl.net/2003.005>

Clark, S. E., Peek, J. E. G., & Miville-Deschénes, M. A. 2019, *Astrophys. J.*, 874, 171, doi: 10.3847/1538-4357/ab0b3b

Clark, S. E., Peek, J. E. G., & Putman, M. E. 2014, *Astrophys. J.*, 789, 82, doi: 10.1088/0004-637X/789/1/82

Clemens, D. P., Pinnick, A. F., Pavel, M. D., & Taylor, B. W. 2012, *Astrophys. J. Suppl. Ser.*, 200, 19, doi: 10.1088/0067-0049/200/2/19

CMB-S4 collaboration, Abazajian, K., Addison, G. E., et al. 2022, *The Astrophysical Journal*, 926, 54, doi: 10.3847/1538-4357/ac1596

Collaboration, L., Ally, E., Arnold, K., et al. 2023, *Progress of Theoretical and Experimental Physics*, 2023, 042F01

Colombi, S., Pogosyan, D., & Souradeep, T. 2000, *Phys. Rev. Lett.*, 85, 5515, doi: 10.1103/PhysRevLett.85.5515

Córdova Rosado, R., Hensley, B. S., Clark, S. E., et al. 2024, *Astrophys. J.*, 960, 96, doi: 10.3847/1538-4357/ad05cd

Cotton, D. V., Marshall, J. P., Frisch, P. C., et al. 2019, *Mon. Not. R. Astron. Soc.*, 483, 3636, doi: 10.1093/mnras/sty3318

Cox, D. P. 2005, *Annu. Rev. Astron. Astrophys.*, 43, 337, doi: 10.1146/annurev.astro.43.072103.150615

Cox, D. P., & Reynolds, R. J. 1987, *Annu. Rev. Astron. Astrophys.*, 25, 303, doi: 10.1146/annurev.aa.25.090187.001511

Cox, D. P., & Smith, B. W. 1974, *Astrophys. J. Lett.*, 189, L105, doi: 10.1086/181476

Crutcher, R. M. 2012, *Annu. Rev. Astron. Astrophys.*, 50, 29, doi: 10.1146/annurev-astro-081811-125514

Crutcher, R. M., Wandelt, B., Heiles, C., Falgarone, E., & Troland, T. H. 2010, *Astrophys. J.*, 725, 466, doi: 10.1088/0004-637X/725/1/466

Cukierman, A. J., Clark, S. E., & Halal, G. 2023, *Astrophys. J.*, 946, 106, doi: 10.3847/1538-4357/acb0c4

Cutri, R. M., Wright, E. L., Conrow, T., et al. 2013, Explanatory Supplement to the AllWISE Data Release Products, Explanatory Supplement to the AllWISE Data Release Products, by R. M. Cutri et al.

Dalgarno, A., & McCray, R. A. 1972, *Annu. Rev. Astron. Astrophys.*, 10, 375, doi: 10.1146/annurev.aa.10.090172.002111

Dalgarno, A., & Rudge, M. R. H. 1964, *Astrophys. J.*, 140, 800, doi: 10.1086/147974

Davis, Leverett, J., & Greenstein, J. L. 1951, *Astrophys. J.*, 114, 206, doi: 10.1086/145464

Davis, L. 1951, *Physical Review*, 81, 890, doi: 10.1103/PhysRev.81.890.2

Dawson, J. R. 2013, PASA, 30, e025, doi: 10.1017/pas.2013.002

Delabrouille, J., Cardoso, J. F., Le Jeune, M., et al. 2009, Astron. Astrophys., 493, 835, doi: 10.1051/0004-6361:200810514

Delouis, J. M., Ally, E., Gauvrit, E., & Boulanger, F. 2022, Astron. Astrophys., 668, A122, doi: 10.1051/0004-6361/202244566

Deudon, M., Kalaitzis, A., Goytom, I., et al. 2020, arXiv preprint arXiv:2002.06460

Dharmawardena, T. E., Bailer-Jones, C. A. L., Fouesneau, M., & Foreman-Mackey, D. 2022, Astron. Astrophys., 658, A166, doi: 10.1051/0004-6361/202141298

Dicke, R. H., Peebles, P. J. E., Roll, P. G., & Wilkinson, D. T. 1965, Astrophys. J., 142, 414, doi: 10.1086/148306

Dickey, J. M., & Lockman, F. J. 1990, Annu. Rev. Astron. Astrophys., 28, 215, doi: 10.1146/annurev.aa.28.090190.001243

Dickey, J. M., McClure-Griffiths, N., Gibson, S. J., et al. 2013, Publications of the Astronomical Society of Australia, 30, e003, doi: 10.1017/pasa.2012.003

Dolginov, A. Z., & Mitrofanov, I. G. 1976, Ap&SS, 43, 291, doi: 10.1007/BF00640010

D’Onghia, E., & Fox, A. J. 2016, Annu. Rev. Astron. Astrophys., 54, 363, doi: 10.1146/annurev-astro-081915-023251

Dorr, F. 2020, *Remote Sensing*, 12, 3812

Draine, B. T. 2003, *Annu. Rev. Astron. Astrophys.*, 41, 241, doi: 10.1146/annurev.astro.41.011802.094840

—. 2011, *Physics of the Interstellar and Intergalactic Medium*

Draine, B. T., & Hensley, B. S. 2021, *Astrophys. J.*, 919, 65, doi: 10.3847/1538-4357/ac0050

Draine, B. T., & Weingartner, J. C. 1997, *Astrophys. J.*, 480, 633, doi: 10.1086/304008

Duffy, A. R., Kay, S. T., Battye, R. A., et al. 2012, *Mon. Not. R. Astron. Soc.*, 420, 2799, doi: 10.1111/j.1365-2966.2011.19894.x

Dunkley, J., Amblard, A., Baccigalupi, C., et al. 2009, in *American Institute of Physics Conference Series*, Vol. 1141, *CMB Polarization Workshop: Theory and Foregrounds: CMBPol Mission Concept Study*, ed. S. Dodelson, D. Baumann, A. Cooray, J. Dunkley, A. Fraisse, M. G. Jackson, A. Kogut, L. Krauss, M. Zaldarriaga, & K. Smith, 222–264, doi: 10.1063/1.3160888

Durrer, R., & Neronov, A. 2013, *A&A Rev.*, 21, 62, doi: 10.1007/s00159-013-0062-7

Dwek, E. 2016, *Astrophys. J.*, 825, 136, doi: 10.3847/0004-637X/825/2/136

Edenhofer, G., Zucker, C., Frank, P., et al. 2023, arXiv e-prints, arXiv:2308.01295, doi: 10.48550/arXiv.2308.01295

Elmegreen, B. G. 2011, in EAS Publications Series, Vol. 51, EAS Publications Series, ed. C. Charbonnel & T. Montmerle, 45–58, doi: 10.1051/eas/1151004

Enßlin, T. A. 2019, *Annalen der Physik*, 531, 1800127, doi: 10.1002/andp.201800127

Errard, J., Feeney, S. M., Peiris, H. V., & Jaffe, A. H. 2016, *JCAP*, 2016, 052, doi: 10.1088/1475-7516/2016/03/052

Ewen, H. I., & Purcell, E. M. 1951, *Nature*, 168, 356, doi: 10.1038/168356a0

Falceta-Gonçalves, D., Lazarian, A., & Kowal, G. 2008, *Astrophys. J.*, 679, 537, doi: 10.1086/587479

Farhang, A., van Loon, J. T., Khosroshahi, H. G., Javadi, A., & Bailey, M. 2019, *Nature Astronomy*, 3, 922, doi: 10.1038/s41550-019-0814-z

Ferrière, K. M. 2001, *Reviews of Modern Physics*, 73, 1031, doi: 10.1103/RevModPhys.73.1031

Fiege, J. D., & Pudritz, R. E. 2000, *Astrophys. J.*, 544, 830, doi: 10.1086/317228

Field, G. B., Goldsmith, D. W., & Habing, H. J. 1969, *Astrophys. J. Lett.*, 155, L149, doi: 10.1086/180324

Finkbeiner, D. P. 2003, *Astrophys. J. Suppl. Ser.*, 146, 407, doi: 10.1086/374411

Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999, *Astrophys. J.*, 524, 867, doi: 10.1086/307852

Fitzpatrick, E. L., Massa, D., Gordon, K. D., Bohlin, R., & Clayton, G. C. 2019, *Astrophys. J.*, 886, 108, doi: 10.3847/1538-4357/ab4c3a

Fixsen, D. J. 2009, *Astrophys. J.*, 707, 916, doi: 10.1088/0004-637X/707/2/916

Flaeger, R., Hill, J. C., & Spergel, D. N. 2014, *JCAP*, 2014, 039, doi: 10.1088/1475-7516/2014/039

Flewelling, H. A., Magnier, E. A., Chambers, K. C., et al. 2020, *Astrophys. J. Suppl. Ser.*, 251, 7, doi: 10.3847/1538-4365/abb82d

Fong, R., Jones, L., Shanks, T., et al. 1987, *Monthly Notices of the Royal Astronomical Society*, 224, 1059

Forero-Romero, J. E., Hoffman, Y., Gottlöber, S., Klypin, A., & Yepes, G. 2009, *Mon. Not. R. Astron. Soc.*, 396, 1815, doi: 10.1111/j.1365-2966.2009.14885.x

Fosalba, P., Lazarian, A., Prunet, S., & Tauber, J. A. 2002, in *American Institute of Physics Conference Series*, Vol. 609, *Astrophysical Polarized Backgrounds*, ed. S. Cecchini, S. Cortiglioni, R. Sault, & C. Sbarra (AIP), 44–50, doi: 10.1063/1.1471821

Fox, A. J., Barger, K. A., Wakker, B. P., et al. 2018, *The Astrophysical Journal*, 854, 142, doi: 10.3847/1538-4357/aaa9bb

Frisch, P. C., Redfield, S., & Slavin, J. D. 2011, *Annu. Rev. Astron. Astrophys.*, 49, 237, doi: 10.1146/annurev-astro-081710-102613

Frisch, P. C., Andersson, B. G., Berdyugin, A., et al. 2012, *Astrophys. J.*, 760, 106, doi: 10.1088/0004-637X/760/2/106

Frisch, P. C., Berdyugin, A., Pirola, V., et al. 2015, *Astrophys. J.*, 814, 112, doi: 10.1088/0004-637X/814/2/112

Fuchs, B., Breitschwerdt, D., de Avillez, M. A., Dettbarn, C., & Flynn, C. 2006, *Mon. Not. R. Astron. Soc.*, 373, 993, doi: 10.1111/j.1365-2966.2006.11044.x

Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, *Astron. Astrophys.*, 616, A1, doi: 10.1051/0004-6361/201833051

Gaia Collaboration, Vallenari, A., Brown, A. G. A., et al. 2023, *Astron. Astrophys.*, 674, A1, doi: 10.1051/0004-6361/202243940

Gardner, F. F., & Whiteoak, J. B. 1966, *Annu. Rev. Astron. Astrophys.*, 4, 245, doi: 10.1146/annurev.aa.04.090166.001333

Goldreich, P., & Kylafis, N. D. 1981, *Astrophys. J. Lett.*, 243, L75, doi: 10.1086/183446

Gómez, G. C., & Cox, D. P. 2004, *Astrophys. J.*, 615, 744, doi: 10.1086/424499

Gontcharov, G. A., & Mosenkov, A. V. 2019, *Mon. Not. R. Astron. Soc.*, 483, 299, doi: 10.1093/mnras/sty2978

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *Astrophys. J.*, 622, 759, doi: 10.1086/427976

Gould, R. J., & Thakur, R. K. 1970, Annals of Physics, 61, 351, doi: 10.1016/0003-4916(70)90289-7

Greaves, J. S., Holland, W. S., Friberg, P., & Dent, W. R. F. 1999, *Astrophys. J. Lett.*, 512, L139, doi: 10.1086/311888

Green, G. M. 2018, *Journal of Open Source Software*, 3, 695, doi: 10.21105/joss.00695

Green, G. M., Schlafly, E., Zucker, C., Speagle, J. S., & Finkbeiner, D. 2019, *Astrophys. J.*, 887, 93, doi: 10.3847/1538-4357/ab5362

Green, G. M., Schlafly, E. F., Finkbeiner, D. P., et al. 2015, *Astrophys. J.*, 810, 25, doi: 10.1088/0004-637X/810/1/25

Green, G. M., Schlafly, E. F., Finkbeiner, D., et al. 2018, *Mon. Not. R. Astron. Soc.*, 478, 651, doi: 10.1093/mnras/sty1008

Guo, H. L., Chen, B. Q., Yuan, H. B., et al. 2021, *Astrophys. J.*, 906, 47, doi: 10.3847/1538-4357/abc68a

Guth, A. H. 1981, *Phys. Rev. D*, 23, 347, doi: 10.1103/PhysRevD.23.347

Hacar, A., Clark, S., Heitsch, F., et al. 2022, Initial Conditions for Star Formation: A Physical Description of the Filamentary ISM, *arXiv*, doi: 10.48550/ARXIV.2203.09562

Halal, G., Clark, S. E., Cukierman, A., Beck, D., & Kuo, C.-L. 2023, Spherical Rolling Hough Transform, v2.0.1, *Zenodo*, doi: 10.5281/zenodo.8025777

Halal, G., Clark, S. E., Cukierman, A., Beck, D., & Kuo, C.-L. 2024a, *Astrophys. J.*, 961, 29, doi: 10.3847/1538-4357/ad06aa

Halal, G., Clark, S. E., & Tahani, M. 2024b, arXiv e-prints, arXiv:2404.11009, doi: 10.48550/arXiv.2404.11009

Halal, G., Clark, S. E., & Tahani, M. 2024, Halal et al. 2024 3D Dust Complexity Maps, V1, Harvard Dataverse, doi: 10.7910/DVN/IW09AE

Hallinan, G., Ravi, V., Weinreb, S., et al. 2019, The DSA-2000 – A Radio Survey Camera. <https://arxiv.org/abs/1907.07648>

Han, J. L. 2017, *Annu. Rev. Astron. Astrophys.*, 55, 111, doi: 10.1146/annurev-astro-091916-055221

Han, J. L., Manchester, R. N., Lyne, A. G., Qiao, G. J., & van Straten, W. 2006, *Astrophys. J.*, 642, 868, doi: 10.1086/501444

Hanson, D., Hoover, S., Crites, A., et al. 2013, *Phys. Rev. Lett.*, 111, 141301, doi: 10.1103/PhysRevLett.111.141301

Haverkorn, M. 2015, in *Astrophysics and Space Science Library*, Vol. 407, *Magnetic Fields in Diffuse Media*, ed. A. Lazarian, E. M. de Gouveia Dal Pino, & C. Melioli, 483, doi: 10.1007/978-3-662-44625-6_17

He, K., Zhang, X., Ren, S., & Sun, J. 2015, arXiv e-prints, arXiv:1512.03385, doi: 10.48550/arXiv.1512.03385

Heiles, C. 1979, *Astrophys. J.*, 229, 533, doi: 10.1086/156986

—. 1984, *Astrophys. J. Suppl. Ser.*, 55, 585, doi: 10.1086/190970

—. 1997a, *Astrophys. J.*, 481, 193, doi: 10.1086/304033

—. 1997b, *Astrophys. J. Suppl. Ser.*, 111, 245, doi: 10.1086/313010

Heiles, C., Goodman, A. A., McKee, C. F., & Zweibel, E. G. 1993, in *Protostars and Planets III*, ed. E. H. Levy & J. I. Lunine, 279

Heiles, C., & Haverkorn, M. 2012, *Space Sci. Rev.*, 166, 293, doi: 10.1007/s11214-012-9866-4

Heiles, C., & Troland, T. H. 2003, *Astrophys. J.*, 586, 1067, doi: 10.1086/367828

Heitsch, F. 2005, in *Astronomical Society of the Pacific Conference Series*, Vol. 343, *Astrophysical Polarimetry: Current Status and Future Directions*, ed. A. Adamson, C. Aspin, C. Davis, & T. Fujiiyoshi, 166

Hendrycks, D., & Gimpel, K. 2016, arXiv e-prints, arXiv:1606.08415, doi: 10.48550/arXiv.1606.08415

Hennebelle, P., & Inutsuka, S.-i. 2019, *Frontiers in Astronomy and Space Sciences*, 6, 5, doi: 10.3389/fspas.2019.00005

Hensley, B. S., & Bull, P. 2018, *Astrophys. J.*, 853, 127, doi: 10.3847/1538-4357/aaa489

Hensley, B. S., & Draine, B. T. 2021, *Astrophys. J.*, 906, 73, doi: 10.3847/1538-4357/abc8f1

Hensley, B. S., Zhang, C., & Bock, J. J. 2019, *Astrophys. J.*, 887, 159, doi: 10.3847/1538-4357/ab5183

Hensley, B. S., Clark, S. E., Fanfani, V., et al. 2022, *Astrophys. J.*, 929, 166, doi: 10.3847/1538-4357/ac5e36

Herranen, J., Lazarian, A., & Hoang, T. 2021, *Astrophys. J.*, 913, 63, doi: 10.3847/1538-4357/abf096

Hervías-Caimapo, C., & Huffenberger, K. M. 2022, *The Astrophysical Journal*, 928, 65

HI4PI Collaboration, Ben Bekhti, N., Flöer, L., et al. 2016, *Astron. Astrophys.*, 594, A116, doi: 10.1051/0004-6361/201629178

Hiltner, W. A. 1949, *Science*, 109, 165, doi: 10.1126/science.109.2825.165

Hoang, T. 2022, *Astrophys. J.*, 928, 102, doi: 10.3847/1538-4357/ac5408

Hoang, T., & Lazarian, A. 2008, *Mon. Not. R. Astron. Soc.*, 388, 117, doi: 10.1111/j.1365-2966.2008.13249.x

—. 2009, *Astrophys. J.*, 697, 1316, doi: 10.1088/0004-637X/697/2/1316

—. 2016, *Astrophys. J.*, 831, 159, doi: 10.3847/0004-637X/831/2/159

Hoang, T., Tram, L. N., Minh Phan, V. H., et al. 2022, *Astron. J.*, 164, 248, doi: 10.3847/1538-3881/ac9af5

Hollenbach, D., & Salpeter, E. E. 1971, *Astrophys. J.*, 163, 155, doi: 10.1086/150754

Hough, J. H., Aitken, D. K., Whittet, D. C. B., Adamson, A. J., & Chrysostomou, A. 2008, *Mon. Not. R. Astron. Soc.*, 387, 797, doi: 10.1111/j.1365-2966.2008.13274.x

Hough, P. V. 1962. <https://www.osti.gov/biblio/4746348>

Hu, W., & Dodelson, S. 2002, *Annu. Rev. Astron. Astrophys.*, 40, 171, doi: 10.1146/annurev.astro.40.060401.093926

Hu, W., & Okamoto, T. 2002, *The Astrophysical Journal*, 574, 566

Huffenberger, K. M., Rotti, A., & Collins, D. C. 2020, *The Astrophysical Journal*, 899, 31, doi: 10.3847/1538-4357/ab9df9

Hull, C. L. H., & Zhang, Q. 2019, *Frontiers in Astronomy and Space Sciences*, 6, 3, doi: 10.3389/fspas.2019.00003

Hutschenreuter, S., Haverkorn, M., Frank, P., Raycheva, N. C., & Enßlin, T. A. 2023, arXiv e-prints, arXiv:2304.12350, doi: 10.48550/arXiv.2304.12350

Inoue, T., & Inutsuka, S.-i. 2016, *Astrophys. J.*, 833, 10, doi: 10.3847/0004-637X/833/1/10

Inutsuka, S.-i., Inoue, T., Iwasaki, K., & Hosokawa, T. 2015, *Astron. Astrophys.*, 580, A49, doi: 10.1051/0004-6361/201425584

Jaffe, T. R., Ferrière, K. M., Banday, A. J., et al. 2013, *Mon. Not. R. Astron. Soc.*, 431, 683, doi: 10.1093/mnras/stt200

Jelić, V., Prelogović, D., Havercorn, M., Remeijn, J., & Klindžić, D. 2018, *Astron. Astrophys.*, 615, L3, doi: 10.1051/0004-6361/201833291

Jenkins, E. B. 2009, *Astrophys. J.*, 700, 1299, doi: 10.1088/0004-637X/700/2/1299

Kaczmarek, J. F., Purcell, C. R., Gaensler, B. M., McClure-Griffiths, N. M., & Stevens, J. 2017, *Mon. Not. R. Astron. Soc.*, 467, 1776, doi: 10.1093/mnras/stx206

Kalberla, P. M. W., Burton, W. B., Hartmann, D., et al. 2005, *A&A*, 440, 775, doi: 10.1051/0004-6361:20041864

Kalberla, P. M. W., & Haud, U. 2018, *Astron. Astrophys.*, 619, A58, doi: 10.1051/0004-6361/201833146

Kalberla, P. M. W., & Kerp, J. 2009, *Annu. Rev. Astron. Astrophys.*, 47, 27, doi: 10.1146/annurev-astro-082708-101823

Kalberla, P. M. W., Kerp, J., & Haud, U. 2020, *Astron. Astrophys.*, 639, A26, doi: 10.1051/0004-6361/202037602

—. 2021, *Astron. Astrophys.*, 654, A91, doi: 10.1051/0004-6361/202140274

Kamionkowski, M., Kosowsky, A., & Stebbins, A. 1997, Physical Review Letters, 78, 2058–2061, doi: 10.1103/physrevlett.78.2058

Kamionkowski, M., & Kovetz, E. D. 2016, Annu. Rev. Astron. Astrophys., 54, 227, doi: 10.1146/annurev-astro-081915-023433

Kerbyson, D., & Atherton, T. 1995, in Fifth International Conference on Image Processing and its Applications, 1995., 370–374, doi: 10.1049/cp:19950683

Kim, C.-G., Choi, S. K., & Flauger, R. 2019, Astrophys. J., 880, 106, doi: 10.3847/1538-4357/ab29f2

Kim, C.-G., & Ostriker, E. C. 2015, Astrophys. J., 802, 99, doi: 10.1088/0004-637X/802/2/99

King, P. K., Chen, C.-Y., Fissel, L. M., & Li, Z.-Y. 2019, Mon. Not. R. Astron. Soc., 490, 2760, doi: 10.1093/mnras/stz2628

King, P. K., Fissel, L. M., Chen, C.-Y., & Li, Z.-Y. 2018, Mon. Not. R. Astron. Soc., 474, 5122, doi: 10.1093/mnras/stx3096

Klessen, R. S., & Glover, S. C. O. 2016, Saas-Fee Advanced Course, 43, 85, doi: 10.1007/978-3-662-47890-5_2

Kothes, R., & Brown, J.-A. 2009, in Cosmic Magnetic Fields: From Planets, to Stars and

Galaxies, ed. K. G. Strassmeier, A. G. Kosovichev, & J. E. Beckman, Vol. 259, 75–80, doi: 10.1017/S1743921309030087

Kovac, J. M., Leitch, E. M., Pryke, C., et al. 2002, Nature, 420, 772, doi: 10.1038/nature01269

Krachmalnicoff, N., & Puglisi, G. 2021, The Astrophysical Journal, 911, 42

Krizhevsky, A., Sutskever, I., & Hinton, G. E. 2012, in Advances in Neural Information Processing Systems, ed. F. Pereira, C. Burges, L. Bottou, & K. Weinberger, Vol. 25 (Curran Associates, Inc.)

Krumholz, M. R. 2011, in American Institute of Physics Conference Series, Vol. 1386, XV Special Courses at the National Observatory of Rio de Janeiro, ed. E. Telles, R. Dupke, & D. Lazzaro (AIP), 9–57, doi: 10.1063/1.3636038

Krumholz, M. R., Stone, J. M., & Gardiner, T. A. 2007, *Astrophys. J.*, 671, 518, doi: 10.1086/522665

Lallement, R., Babusiaux, C., Vergely, J. L., et al. 2019, *Astron. Astrophys.*, 625, A135, doi: 10.1051/0004-6361/201834695

Lallement, R., Vergely, J. L., Babusiaux, C., & Cox, N. L. J. 2022, *Astron. Astrophys.*, 661, A147, doi: 10.1051/0004-6361/202142846

Lallement, R., Vergely, J. L., Valette, B., et al. 2014, *Astron. Astrophys.*, 561, A91, doi: 10.1051/0004-6361/201322032

Lallement, R., Welsh, B. Y., Vergely, J. L., Crifo, F., & Sfeir, D. 2003, *Astron. Astrophys.*, 411, 447, doi: 10.1051/0004-6361:20031214

Lallement, R., Capitanio, L., Ruiz-Dern, L., et al. 2018, *Astron. Astrophys.*, 616, A132, doi: 10.1051/0004-6361/201832832

Lazarian, A. 2020, *Astrophys. J.*, 902, 97, doi: 10.3847/1538-4357/abb1b4

Lazarian, A., Andersson, B. G., & Hoang, T. 2015, in *Polarimetry of Stars and Planetary Systems*, 81, doi: 10.48550/arXiv.1511.03696

Lazarian, A., & Hoang, T. 2007, *Mon. Not. R. Astron. Soc.*, 378, 910, doi: 10.1111/j.1365-2966.2007.11817.x

—. 2008, *Astrophys. J. Lett.*, 676, L25, doi: 10.1086/586706

—. 2019, *Astrophys. J.*, 883, 122, doi: 10.3847/1538-4357/ab3d39

Lazarian, A., Yuen, K. H., & Pogosyan, D. 2022, *Astrophys. J.*, 935, 77, doi: 10.3847/1538-4357/ac6877

Lee, H. M., & Draine, B. T. 1985, *Astrophys. J.*, 290, 211, doi: 10.1086/162974

Lei, M., & Clark, S. E. 2023a, arXiv e-prints, arXiv:2312.03846, doi: 10.48550/arXiv.2312.03846

—. 2023b, *Astrophys. J.*, 947, 74, doi: 10.3847/1538-4357/acc02a

Lei Ba, J., Kiros, J. R., & Hinton, G. E. 2016, arXiv e-prints, arXiv:1607.06450, doi: 10.48550/arXiv.1607.06450

Leike, R. H., Edenhofer, G., Knollmüller, J., et al. 2022, arXiv e-prints, arXiv:2204.11715, doi: 10.48550/arXiv.2204.11715

Leike, R. H., & Enßlin, T. A. 2019, *Astron. Astrophys.*, 631, A32, doi: 10.1051/0004-6361/201935093

Leike, R. H., Glatzle, M., & Enßlin, T. A. 2020, *Astron. Astrophys.*, 639, A138, doi: 10.1051/0004-6361/202038169

Lenz, D., Hensley, B. S., & Doré, O. 2017, *Astrophys. J.*, 846, 38, doi: 10.3847/1538-4357/aa84af

Leroy, J. L. 1999, *Astron. Astrophys.*, 346, 955

Li, J., Lv, Q., Zhang, W., et al. 2023, *Remote Sensing*, 15, 4183, doi: 10.3390/rs15174183

Linde, A. D. 1982, *Physics Letters B*, 108, 389, doi: 10.1016/0370-2693(82)91219-9

Liu, J., Zhang, Q., Commerçon, B., et al. 2021, *Astrophys. J.*, 919, 79, doi: 10.3847/1538-4357/ac0cec

Liu, J., Sun, Z., Han, J., et al. 2022, *Science China Physics, Mechanics, and Astronomy*, 65, 109511, doi: 10.1007/s11433-022-1966-4

Liu, W., Chiao, M., Collier, M. R., et al. 2017, *Astrophys. J.*, 834, 33, doi: 10.3847/1538-4357/834/1/33

Louis, T., Naess, S., Garrido, X., & Challinor, A. 2020, *Phys. Rev. D*, 102, 123538, doi: 10.1103/PhysRevD.102.123538

Low, F. J., Beintema, D. A., Gautier, T. N., et al. 1984, *Astrophys. J. Lett.*, 278, L19, doi: 10.1086/184213

Lucchini, S., D’Onghia, E., & Fox, A. J. 2021, *Astrophys. J. Lett.*, 921, L36, doi: 10.3847/2041-8213/ac3338

Mac Low, M.-M., McCray, R., & Norman, M. L. 1989, *Astrophys. J.*, 337, 141, doi: 10.1086/167094

Maconi, E., Soler, J. D., Reissl, S., et al. 2023, *Mon. Not. R. Astron. Soc.*, 523, 5995, doi: 10.1093/mnras/stad1854

Magnani, L., & Smith, A. J. 2010, *Astrophys. J.*, 722, 1685, doi: 10.1088/0004-637X/722/2/1685

Maíz-Apellániz, J. 2001, *Astrophys. J. Lett.*, 560, L83, doi: 10.1086/324016

Malinen, J., Montier, L., Montillaud, J., et al. 2016, *Mon. Not. R. Astron. Soc.*, 460, 1934, doi: 10.1093/mnras/stw1061

Mallat, S. 2011, arXiv e-prints, arXiv:1101.2286, doi: 10.48550/arXiv.1101.2286

Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, *Astron. J.*, 129, 1993, doi: 10.1086/428488

Mantz, H., Jacobs, K., & Mecke, K. 2008, *Journal of Statistical Mechanics: Theory and Experiment*, 2008, 12015, doi: 10.1088/1742-5468/2008/12/P12015

Marchal, A., & Martin, P. G. 2023, *Astrophys. J.*, 942, 70, doi: 10.3847/1538-4357/aca4d2

McClure-Griffiths, N. M., Pisano, D. J., Calabretta, M. R., et al. 2009, *Astrophys. J. Suppl. Ser.*, 181, 398, doi: 10.1088/0067-0049/181/2/398

McKee, C. F., & Ostriker, J. P. 1977, *Astrophys. J.*, 218, 148, doi: 10.1086/155667

Medan, I., & Andersson, B. G. 2019, *Astrophys. J.*, 873, 87, doi: 10.3847/1538-4357/ab063c

Meisner, A. M., & Finkbeiner, D. P. 2014, *Astrophys. J.*, 781, 5, doi: 10.1088/0004-637X/781/1/5

Minami, Y., & Komatsu, E. 2020, *Progress of Theoretical and Experimental Physics*, 2020, 103E02, doi: 10.1093/ptep/ptaa130

Minami, Y., & Komatsu, E. 2020, *PhRvL*, 125, 221301, doi: 10.1103/PhysRevLett.125.221301

Miville-Deschénes, M.-A., Murray, N., & Lee, E. J. 2017, *Astrophys. J.*, 834, 57, doi: 10.3847/1538-4357/834/1/57

Myers, P. C., Stephens, I. W., & Coudé, S. 2024, *Astrophys. J.*, 962, 64, doi: 10.3847/1538-4357/ad1596

Neckel, T., & Klare, G. 1980, *A&AS*, 42, 251

O'Neill, T. J., Zucker, C., Goodman, A. A., & Edenhofer, G. 2024, arXiv e-prints, arXiv:2403.04961, doi: 10.48550/arXiv.2403.04961

Orr, M. E., Fielding, D. B., Hayward, C. C., & Burkhardt, B. 2022, *Astrophys. J. Lett.*, 924, L28, doi: 10.3847/2041-8213/ac479f

Ostriker, E. C., Stone, J. M., & Gammie, C. F. 2001, *Astrophys. J.*, 546, 980, doi: 10.1086/318290

O'Neill, T., Goodman, A., Soler, J., Han, J. J., & Zucker, C. 2023, doi: 10.22541/au.167303779.92162611/v3

Padoan, P., Goodman, A., Draine, B. T., et al. 2001, *Astrophys. J.*, 559, 1005, doi: 10.1086/322504

Panopoulou, G. V., & Lenz, D. 2020a, *Astrophys. J.*, 902, 120, doi: 10.3847/1538-4357/abb6f5

—. 2020b, *Astrophys. J.*, 902, 120, doi: 10.3847/1538-4357/abb6f5

Panopoulou, G. V., Psaradaki, I., & Tassis, K. 2016, *Mon. Not. R. Astron. Soc.*, 462, 1517, doi: 10.1093/mnras/stw1678

Panopoulou, G. V., Tassis, K., Skalidis, R., et al. 2019, *Astrophys. J.*, 872, 56, doi: 10.3847/1538-4357/aafdb2

Pattle, K., Fissel, L., Tahani, M., Liu, T., & Ntormousi, E. 2023, in *Astronomical Society of the Pacific Conference Series*, Vol. 534, *Protostars and Planets VII*, ed. S. Inutsuka, Y. Aikawa, T. Muto, K. Tomida, & M. Tamura, 193, doi: 10.48550/arXiv.2203.11179

Pearson, S., Clark, S. E., Demirjian, A. J., et al. 2022, *Astrophys. J.*, 926, 166, doi: 10.3847/1538-4357/ac4496

Peek, J. E. G., Babler, B. L., Zheng, Y., et al. 2018, *Astrophys. J. Suppl. Ser.*, 234, 2, doi: 10.3847/1538-4365/aa91d3

Pelgrims, V., Clark, S. E., Hensley, B. S., et al. 2021a, *Astron. Astrophys.*, 647, A16, doi: 10.1051/0004-6361/202040218

—. 2021b, *Astron. Astrophys.*, 647, A16, doi: 10.1051/0004-6361/202040218

Pelgrims, V., Ferrière, K., Boulanger, F., Lallement, R., & Montier, L. 2020, *Astron. Astrophys.*, 636, A17, doi: 10.1051/0004-6361/201937157

Pelgrims, V., Panopoulou, G. V., Tassis, K., et al. 2023, *Astron. Astrophys.*, 670, A164, doi: 10.1051/0004-6361/202244625

Pelgrims, V., Mandarakas, N., Skalidis, R., et al. 2024, arXiv e-prints, arXiv:2404.10821, doi: 10.48550/arXiv.2404.10821

Pelkonen, V. M., Juvela, M., & Padoan, P. 2007, *Astron. Astrophys.*, 461, 551, doi: 10.1051/0004-6361:20065838

Penzias, A. A., & Wilson, R. W. 1965, *Astrophys. J.*, 142, 419, doi: 10.1086/148307

Planck Collaboration, Abergel, A., Ade, P. A. R., et al. 2011a, *Astron. Astrophys.*, 536, A24, doi: 10.1051/0004-6361/201116485

—. 2011b, *Astron. Astrophys.*, 536, A24, doi: 10.1051/0004-6361/201116485

—. 2014a, *Astron. Astrophys.*, 571, A11, doi: 10.1051/0004-6361/201323195

—. 2014b, *Astron. Astrophys.*, 571, A11, doi: 10.1051/0004-6361/201323195

Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2015, *Astron. Astrophys.*, 576, A104, doi: 10.1051/0004-6361/201424082

—. 2016a, *Astron. Astrophys.*, 586, A141, doi: 10.1051/0004-6361/201526506

Planck Collaboration, Adam, R., Ade, P. A. R., et al. 2016b, *Astron. Astrophys.*, 596, A103, doi: 10.1051/0004-6361/201528033

Planck Collaboration, Aghanim, N., Ashdown, M., et al. 2016c, *Astron. Astrophys.*, 596, A109, doi: 10.1051/0004-6361/201629022

Planck Collaboration, Adam, R., Ade, P. A. R., et al. 2016d, *Astron. Astrophys.*, 594, A10, doi: 10.1051/0004-6361/201525967

—. 2016e, *Astron. Astrophys.*, 594, A10, doi: 10.1051/0004-6361/201525967

—. 2016f, *Astron. Astrophys.*, 586, A135, doi: 10.1051/0004-6361/201425044

Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020a, *Astron. Astrophys.*, 641, A12, doi: 10.1051/0004-6361/201833885

Planck Collaboration, Akrami, Y., Ashdown, M., et al. 2020b, *Astron. Astrophys.*, 641, A11, doi: 10.1051/0004-6361/201832618

Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020c, *Astron. Astrophys.*, 641, A6, doi: 10.1051/0004-6361/201833910

—. 2020d, *Astron. Astrophys.*, 641, A1, doi: 10.1051/0004-6361/201833880

Planck Collaboration, Akrami, Y., Andersen, K. J., et al. 2020e, *Astronomy & Astrophysics*, 643, A42, doi: 10.1051/0004-6361/202038073

Planck Collaboration, Akrami, Y., Ashdown, M., et al. 2020f, *A&A*, 641, A4, doi: 10.1051/0004-6361/201833881

Planck Collaboration, Akrami, Y., Ashdown, M., et al. 2020g, *Astron. Astrophys.*, 641, A11, doi: 10.1051/0004-6361/201832618

Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020h, *Astron. Astrophys.*, 641, A8, doi: 10.1051/0004-6361/201833886

Plaszczynski, S., Montier, L., Levrier, F., & Tristram, M. 2014, *Mon. Not. R. Astron. Soc.*, 439, 4048, doi: 10.1093/mnras/stu270

POLARBEAR Collaboration, Adachi, S., Adkins, T., et al. 2022, *Astrophys. J.*, 931, 101, doi: 10.3847/1538-4357/ac6809

Polychroni, D., Schisano, E., Elia, D., et al. 2013, *Astrophys. J. Lett.*, 777, L33, doi: 10.1088/2041-8205/777/2/L33

Popescu, C. C., & Tuffs, R. J. 2002, *Mon. Not. R. Astron. Soc.*, 335, L41, doi: 10.1046/j.1365-8711.2002.05881.x

Prézeau, G., & Reinecke, M. 2010, *Astrophys. J. Suppl. Ser.*, 190, 267, doi: 10.1088/0067-0049/190/2/267

Purcell, E. M. 1975, *Interstellar grains as pinwheels.*, 155–167

—. 1979, *Astrophys. J.*, 231, 404, doi: 10.1086/157204

Puspitarini, L., & Lallement, R. 2012, *Astron. Astrophys.*, 545, A21, doi: 10.1051/0004-6361/201219284

Putman, M., Peek, J., & Joung, M. 2012, *Annual Review of Astronomy and Astrophysics*, 50, 491, doi: 10.1146/annurev-astro-081811-125612

Raymond, J. C., Slavin, J. D., Blair, W. P., et al. 2020, *Astrophys. J.*, 903, 2, doi: 10.3847/1538-4357/abb821

Reach, W. T., Wall, W. F., & Odegard, N. 1998, *Astrophys. J.*, 507, 507, doi: 10.1086/306357

Reichardt, C. L. 2016, in *Astrophysics and Space Science Library*, Vol. 423, *Understanding the Epoch of Cosmic Reionization: Challenges and Progress*, ed. A. Mesinger, 227, doi: 10.1007/978-3-319-21957-8_8

Remazeilles, M., Delabrouille, J., & Cardoso, J.-F. 2011, *Mon. Not. R. Astron. Soc.*, 418, 467, doi: 10.1111/j.1365-2966.2011.19497.x

Reynolds, R. J., & Ogden, P. M. 1979, *Astrophys. J.*, 229, 942, doi: 10.1086/157028

Rezaei Kh., S., Bailer-Jones, C. A. L., Hanson, R. J., & Fouesneau, M. 2017, *Astron. Astrophys.*, 598, A125, doi: 10.1051/0004-6361/201628885

Rezaei Kh., S., Bailer-Jones, C. A. L., Schlafly, E. F., & Fouesneau, M. 2018, *Astron. Astrophys.*, 616, A44, doi: 10.1051/0004-6361/201732503

Rezaei Kh., S., Bailer-Jones, C. A. L., Soler, J. D., & Zari, E. 2020, *Astron. Astrophys.*, 643, A151, doi: 10.1051/0004-6361/202038708

Rezaei Kh., S., & Kainulainen, J. 2022, *Astrophys. J. Lett.*, 930, L22, doi: 10.3847/2041-8213/ac67db

Richter, P. 2000, *Astron. Astrophys.*, 359, 1111, doi: 10.48550/arXiv.astro-ph/0005266

Richter, P., & De Boer, K. S. 2005, *The Coldest Phase in Halo High-velocity Gas: Dust and Molecules* (Dordrecht: Springer Netherlands), 183–194, doi: 10.1007/1-4020-2579-3_9

Ritacco, A., Boulanger, F., Guillet, V., et al. 2022, arXiv e-prints, arXiv:2206.07671.
<https://arxiv.org/abs/2206.07671>

Rix, H.-W., Chandra, V., Andrae, R., et al. 2022, *Astrophys. J.*, 941, 45, doi: 10.3847/1538-4357/ac9e01

Rosolowsky, E. W., Pineda, J. E., Kauffmann, J., & Goodman, A. A. 2008, *Astrophys. J.*, 679, 1338, doi: 10.1086/587685

Salvetti, F., Mazzia, V., Khaliq, A., & Chiaberge, M. 2020, *Remote Sensing*, 12, 2207

Sandstrom, K. M., Koch, E. W., Leroy, A. K., et al. 2023, *Astrophys. J. Lett.*, 944, L8, doi: 10.3847/2041-8213/aca972

Santos, F. P., Corradi, W., & Reis, W. 2011, *Astrophys. J.*, 728, 104, doi: 10.1088/0004-637X/728/2/104

Schlafly, E. F., Meisner, A. M., & Green, G. M. 2019, *Astrophys. J. Suppl. Ser.*, 240, 30, doi: 10.3847/1538-4365/aafbea

Schulreich, M. M., Breitschwerdt, D., Feige, J., & Dettbarn, C. 2017, *Astron. Astrophys.*, 604, A81, doi: 10.1051/0004-6361/201629837

Schulreich, M. M., Feige, J., & Breitschwerdt, D. 2023, *Astron. Astrophys.*, 680, A39, doi: 10.1051/0004-6361/202347532

Seljak, U. 1997, *The Astrophysical Journal*, 482, 6–16, doi: 10.1086/304123

Seljak, U., & Zaldarriaga, M. 1997, *Physical Review Letters*, 78, 2054–2057, doi: 10.1103/physrevlett.78.2054

Sfeir, D. M., Lallement, R., Crifo, F., & Welsh, B. Y. 1999, *Astron. Astrophys.*, 346, 785

Simard-Normandin, M., & Kronberg, P. P. 1980, *Astrophys. J.*, 242, 74, doi: 10.1086/158445

Skalidis, R., & Pelgrims, V. 2019, *Astron. Astrophys.*, 631, L11, doi: 10.1051/0004-6361/201936547

Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, *Astron. J.*, 131, 1163, doi: 10.1086/498708

Smith, R. K., & Cox, D. P. 2001, *Astrophys. J. Suppl. Ser.*, 134, 283, doi: 10.1086/320850

Smoot, G. F., Bennett, C. L., Kogut, A., et al. 1992, *Astrophys. J. Lett.*, 396, L1, doi: 10.1086/186504

Snowden, S. L., Egger, R., Finkbeiner, D. P., Freyberg, M. J., & Plucinsky, P. P. 1998, *Astrophys. J.*, 493, 715, doi: 10.1086/305135

Snowden, S. L., Egger, R., Freyberg, M. J., et al. 1997, *Astrophys. J.*, 485, 125, doi: 10.1086/304399

Sofue, Y., & Rubin, V. 2001, *Annu. Rev. Astron. Astrophys.*, 39, 137, doi: 10.1146/annurev.astro.39.1.137

Soler, J. D., Bracco, A., & Pon, A. 2018, *Astron. Astrophys.*, 609, L3, doi: 10.1051/0004-6361/201732203

Soler, J. D., Zucker, C., Peek, J. E. G., et al. 2023, *Astron. Astrophys.*, 675, A206, doi: 10.1051/0004-6361/202346241

SPIDER Collaboration, Ade, P. A. R., Amiri, M., et al. 2022, *Astrophys. J.*, 927, 174, doi: 10.3847/1538-4357/ac20df

SPT Collaboration, Balkenhol, L., Dutcher, D., et al. 2023, *Phys. Rev. D*, 108, 023510, doi: 10.1103/PhysRevD.108.023510

Stompor, R., Leach, S., Stivoli, F., & Baccigalupi, C. 2009, *Mon. Not. R. Astron. Soc.*, 392, 216, doi: 10.1111/j.1365-2966.2008.14023.x

Sturch, C. 1969, *Astron. J.*, 74, 82, doi: 10.1086/110779

Sullivan, C. H., Fissel, L. M., King, P. K., et al. 2021, *Mon. Not. R. Astron. Soc.*, 503, 5006, doi: 10.1093/mnras/stab596

Tahani, M. 2022, *Frontiers in Astronomy and Space Sciences*, 9, 940027, doi: 10.3389/fspas.2022.940027

Tahani, M., Plume, R., Brown, J. C., & Kainulainen, J. 2018, *Astron. Astrophys.*, 614, A100, doi: 10.1051/0004-6361/201732219

Tahani, M., Glover, J., Lupyniciw, W., et al. 2022a, *Astron. Astrophys.*, 660, L7, doi: 10.1051/0004-6361/202243322

Tahani, M., Lupyniciw, W., Glover, J., et al. 2022b, *Astron. Astrophys.*, 660, A97, doi: 10.1051/0004-6361/202141170

Tahani, M., Bastien, P., Furuya, R. S., et al. 2023, *Astrophys. J.*, 944, 139, doi: 10.3847/1538-4357/acac81

Tassis, K., & Pavlidou, V. 2015, *Mon. Not. R. Astron. Soc.*, 451, L90, doi: 10.1093/mnrasl/slv077

Tassis, K., Ramaprakash, A. N., Readhead, A. C. S., et al. 2018, *arXiv e-prints*, arXiv:1810.05652, doi: 10.48550/arXiv.1810.05652

Thorne, B., Dunkley, J., Alonso, D., & Naess, S. 2017, *Monthly Notices of the Royal Astronomical Society*, 469, 2821

Torrado, J., & Lewis, A. 2019, Cobaya: Bayesian analysis in cosmology, *Astrophysics Source Code Library*, record ascl:1910.019. <http://ascl.net/1910.019>

—. 2021, *JCAP*, 2021, 057, doi: 10.1088/1475-7516/2021/05/057

Tram, L. N., & Hoang, T. 2022, *Frontiers in Astronomy and Space Sciences*, 9, 923927, doi: 10.3389/fspas.2022.923927

Tumlinson, J., Shull, J. M., Rachford, B. L., et al. 2002, *The Astrophysical Journal*, 566, 857, doi: 10.1086/338112

Unger, M., & Farrar, G. R. 2023, arXiv e-prints, arXiv:2311.12120, doi: 10.48550/arXiv.2311.12120

Vacher, L., Aumont, J., Boulanger, F., et al. 2022, Frequency dependence of the thermal dust E/B ratio and EB correlation: insights from the spin-moment expansion, arXiv, doi: 10.48550/ARXIV.2210.14768

Valsesia, D., & Magli, E. 2021, *IEEE Transactions on Geoscience and Remote Sensing*, 60, 1

Van Eck, C. L., Haverkorn, M., Alves, M. I. R., et al. 2017, *Astron. Astrophys.*, 597, A98, doi: 10.1051/0004-6361/201629707

Van Eck, C. L., Gaensler, B. M., Hutschenreuter, S., et al. 2023, *Astrophys. J. Suppl. Ser.*, 267, 28, doi: 10.3847/1538-4365/acda24

Vaswani, A., Shazeer, N., Parmar, N., et al. 2017, arXiv e-prints, arXiv:1706.03762, doi: 10.48550/arXiv.1706.03762

Vergely, J. L., Lallement, R., & Cox, N. L. J. 2022a, *Astron. Astrophys.*, 664, A174, doi: 10.1051/0004-6361/202243319

—. 2022b, *Astron. Astrophys.*, 664, A174, doi: 10.1051/0004-6361/202243319

Verschuur, G. L. 1968, *Phys. Rev. Lett.*, 21, 775, doi: 10.1103/PhysRevLett.21.775

Wakker, B. P. 1991, *Astron. Astrophys.*, 250, 499

Wakker, B. P. 2001, *The Astrophysical Journal Supplement Series*, 136, 463–535, doi: 10.1086/321783

Wakker, B. P., & Boulanger, F. 1986, *Astron. Astrophys.*, 170, 84

Wandelt, B. D., & Górski, K. M. 2001, *Physical Review D*, 63, doi: 10.1103/physrevd.63.123002

Watkins, E. J., Barnes, A. T., Henny, K., et al. 2023, *Astrophys. J. Lett.*, 944, L24, doi: 10.3847/2041-8213/aca6e4

Welsh, B. Y. 1991, *Astrophys. J.*, 373, 556, doi: 10.1086/170074

Welsh, B. Y., & Lallement, R. 2010, *PASP*, 122, 1320, doi: 10.1086/657323

Welsh, B. Y., Sallmen, S., & Lallement, R. 2004, *Astron. Astrophys.*, 414, 261, doi: 10.1051/0004-6361:20034367

Welsh, B. Y., Sfeir, D. M., Sirk, M. M., & Lallement, R. 1999, *Astron. Astrophys.*, 352, 308

West, J. L., Safi-Harb, S., Jaffe, T., et al. 2016, *Astron. Astrophys.*, 587, A148, doi: 10.1051/0004-6361/201527001

Westmeier, T. 2018, *Mon. Not. R. Astron. Soc.*, 474, 289, doi: 10.1093/mnras/stx2757

Winkel, B., Kerp, J., Flöer, L., et al. 2016, *Astron. Astrophys.*, 585, A41, doi: 10.1051/0004-6361/201527007

Wolfire, M. G., Hollenbach, D., McKee, C. F., Tielens, A. G. G. M., & Bakes, E. L. O. 1995, *Astrophys. J.*, 443, 152, doi: 10.1086/175510

Wolfire, M. G., McKee, C. F., Hollenbach, D., & Tielens, A. G. G. M. 2003, *Astrophys. J.*, 587, 278, doi: 10.1086/368016

Yao, J., Krachmalnicoff, N., Foschi, M., Puglisi, G., & Baccigalupi, C. 2024, *Astron. Astrophys.*, 686, A290, doi: 10.1051/0004-6361/202449827

Zaldarriaga, M. 2001, *Phys. Rev. D*, 64, 103001, doi: 10.1103/PhysRevD.64.103001

Zhang, X., Green, G. M., & Rix, H.-W. 2023, *Mon. Not. R. Astron. Soc.*, 524, 1855, doi: 10.1093/mnras/stad1941

Zonca, A., Thorne, B., Krachmalnicoff, N., & Borrill, J. 2021, The Journal of Open Source Software, 6, 3783, doi: 10.21105/joss.03783

Zucker, C., Speagle, J. S., Schlafly, E. F., et al. 2019, *Astrophys. J.*, 879, 125, doi: 10.3847/1538-4357/ab2388

Zucker, C., Goodman, A., Alves, J., et al. 2021, *Astrophys. J.*, 919, 35, doi: 10.3847/1538-4357/ac1f96

Zucker, C., Goodman, A. A., Alves, J., et al. 2022, *Nature*, 601, 334, doi: 10.1038/s41586-021-04286-5

Zweibel, E. G. 2005, in *Fluid Dynamics and Dynamos in Astrophysics and Geophysics*, ed. A. M. Soward, C. A. Jones, D. W. Hughes, & N. O. Weiss, 115, doi: 10.1201/9780203017692.ch5