LETTER TO THE EDITOR

A new approach for the statistical denoising of *Planck* interstellar dust polarization data[★]

Bruno Regaldo-Saint Blancard^{1,2}, Erwan Allys¹, François Boulanger¹, François Levrier¹, and Niall Jeffrey^{1,3}

¹ Laboratoire de Physique de l'École Normale Supérieure, ENS, Université PSL, CNRS, Sorbonne Université, Université de Paris, 75005 Paris, France

e-mail: bruno.regaldo@phys.ens.fr

² Observatoire de Paris, PSL University, Sorbonne Université, LERMA, 75014 Paris, France

³ Department of Physics & Astronomy, University College London, Gower Street, London WC1E 6BT, UK

Received 5 February 2021 / Accepted 13 May 2021

ABSTRACT

Dust emission is the main foreground for cosmic microwave background polarization. Its statistical characterization must be derived from the analysis of observational data because the precision required for a reliable component separation is far greater than what is currently achievable with physical models of the turbulent magnetized interstellar medium. This Letter takes a significant step toward this goal by proposing a method that retrieves non-Gaussian statistical characteristics of dust emission from noisy *Planck* polarization observations at 353 GHz. We devised a statistical denoising method based on wavelet phase harmonics (WPH) statistics, which characterize the coherent structures in non-Gaussian random fields and define a generative model of the data. The method was validated on mock data combining a dust map from a magnetohydrodynamic simulation and *Planck* noise maps. The denoised map reproduces the true power spectrum down to scales where the noise power is an order of magnitude larger than that of the signal. It remains highly correlated to the true emission and retrieves some of its non-Gaussian properties. Applied to *Planck* data, the method provides a new approach to building a generative model of dust polarization that will characterize the full complexity of the dust emission. We also release PyWPH, a public Python package, to perform GPU-accelerated WPH analyses on images.

Key words. dust, extinction – polarization – methods: statistical – cosmic background radiation

1. Introduction

The quest of primordial *B*-modes in the polarization of the cosmic microwave background (CMB) faces a major challenge, namely the accurate characterization of one of its main foregrounds: the polarized emission of interstellar dust from our Galaxy (BICEP2/Keck and Planck Collaborations 2015). To rightfully claim any detection of the cosmological signal, one has to fully take into account the complexity of the Galactic contribution to the sky emission (Hensley & Bull 2018; Clark & Hensley 2019; Pelgrims et al. 2020; Planck Collaboration XI 2020). The statistical characterization of Galactic polarized foregrounds is also important to extract the CMB lensing potential from sky observations (Beck et al. 2020).

The Galactic contribution is sourced by the thermal emission of nonspherical dust grains in the diffuse interstellar medium (ISM). The mixture of dust and gas is described as a turbulent magnetized fluid, and the grains tend to align statistically with their long axes perpendicular to the local magnetic field, thus polarizing the emission (Planck Collaboration XII 2020). The physical and chemical processes regulating and structuring the diffuse ISM are nonlinearly coupled (Draine 2011), leading to the emergence of strong scale couplings, which may be evidenced by the highly non-Gaussian statistics observed for many tracers (Burkhart et al. 2009).

Capturing non-Gaussianity is therefore essential to characterize the dust signal. To tackle this issue, several authors have introduced machine learning algorithms that need to be trained (Aylor et al. 2020; Krachmalnicoff & Puglisi 2021; Petroff et al. 2020; Thorne et al. 2021). Another approach, which does not involve any learning steps, uses statistics defined from nonlinear transforms, namely the wavelet scattering transform and the wavelet phase harmonics (WPH). These describe the couplings between scales in non-Gaussian fields, and they have been applied to ISM data (Allys et al. 2019; Regaldo-Saint Blancard et al. 2020; Saydjari et al. 2021) as well to study the large-scale structure of the Universe (Allys et al. 2020; Cheng et al. 2020). They have been successfully used to characterize simulated and observational data with a high signal-to-noise ratio (S/N), and to build generative models for these data. However, these approaches face difficulties with noisy data, which is the case of *Planck* polarization observations. In this letter, we introduce a statistical denoising method using the WPH statistics to retrieve the statistical properties of the noise-free dust emission.

In Sect. 2 we present our statistical denoising method and show how well it performs on noisy simulated data of the polarized emission from dust. In Sect. 3, we apply it to the Chamaeleon-Musca 353 GHz polarized emission, as observed by the *Planck* satellite, and discuss our results. In Sect. 4 we conclude and suggest a few perspectives for follow-up studies. This Letter also includes two appendices. Appendix A

[★] The public Python package PyWPH is available at <https://github.com/bregaldo/pywph>.

introduces the WPH statistics and the denoising procedure in detail. Appendix B presents a first comparison with other methods. We also provide a public Python package to perform GPU-accelerated WPH analyses on images called PyWPH¹.

2. Method and validation

2.1. Description of the method

We observe a noisy map $\mathbf{d}$ that is modeled as follows:

$$\mathbf{d} = \mathbf{s} + \mathbf{n}, \quad (1)$$

with $\mathbf{s}$ being the target truth map and $\mathbf{n}$ being an additive noise signal. The noise $\mathbf{n}$ is an unknown realization of a random field N . We assume that we know N , meaning that we are able to generate as many independent realizations of it as needed. We call $\{\mathbf{n}_1, \dots, \mathbf{n}_M\}$ M of these.

We introduce a method to retrieve the statistical properties of $\mathbf{s}$ while denoising $\mathbf{d}$. This statistical denoising consists in iteratively deforming $\mathbf{d}$ to build a denoised map $\tilde{\mathbf{s}}$ such that the $\{\tilde{\mathbf{s}} + \mathbf{n}_i\}_i$ maps and the $\mathbf{d}$ map become “close enough” in a given statistical space. It employs the possibility introduced in Zhang & Mallat (2021) to build a generative model from WPH statistics, that is, the ability to generate new random realizations based on statistical constraints. In the following, we call ϕ the operator that computes a set of summary statistics from a given map. The algorithm consists in minimizing the following loss function:

$$\mathcal{L}_M(\mathbf{u}) = \frac{1}{M} \sum_{i=1}^M \|\phi(\mathbf{u} + \mathbf{n}_i) - \phi(\mathbf{d})\|^2, \quad (2)$$

where $\|\cdot\|$ denotes the Euclidean norm. We choose $\mathbf{u}_0 = \mathbf{d}$ to initialize the optimizer. The denoised map $\tilde{\mathbf{s}}$ corresponds to an approximate minimum obtained by performing this optimization in pixel space, using an LBFGS optimizer (Byrd et al. 1995). We note that a limitation of this algorithm relies on the (ideal) assumption that we know N . In practice, any non-modeled statistical property of the noise will be considered to be part of the signal $\mathbf{s}$.

The choice of the operator ϕ is obviously paramount to the quality of the method and must be tailored to the properties of $\mathbf{s}$ and $\mathbf{n}$. In the context of ISM polarization data, we expect $\mathbf{s}$ to be relatively regular and to exhibit non-Gaussian signatures due to the interactions between scales (e.g., filamentary structures), while the noise is expected to be highly irregular and close to a (possibly spatially-varying) Gaussian white noise.

In this work, the ϕ operator computes WPH statistics. These statistics are designed to characterize the coherent structures that appear in non-Gaussian random fields, by quantifying the phase alignment between different scales (Zhang & Mallat 2021; Mallat et al. 2020). Their ability to define a generative model has been studied in particular by Allys et al. (2020), who demonstrate quantitatively that they include most of the information captured by various other statistics, such as the power spectrum, the bispectrum, or the Minkowski functionals. These statistics were applied to real-valued data. Here we extend them to complex-valued polarization maps $Q + iU$. Statistics derived from this complex variable efficiently characterize the full complexity of the polarization signal (Regaldo-Saint Blancard et al. 2020). We refer readers to Appendix A for a presentation of the WPH statistics and a detailed description of the denoising procedure.

2.2. Validation on a simulation

In this section, we assess the performance of our denoising algorithm by applying it to mock data $\mathbf{d} = \mathbf{s} + \mathbf{n}$, emulating a noisy $Q + iU$ Planck polarization signal. We show figures of merit based on power spectra and probability distribution functions (PDFs) of the increments of the maps. These are only shown for Q , but they are similar for U . We postpone figures of merit based on other kinds of statistics, including the WPH statistics, to a future study.

Here $\mathbf{s}$ is a simulated $Q + iU$ map, representing a typical linear polarization signal produced by the thermal emission of dust in the diffuse ISM. It was built from the same magnetohydrodynamic simulation as the one used in Regaldo-Saint Blancard et al. (2020) and following the same procedure. We refer to this Letter for more details on the construction of this simulated map. We simulated a Planck observation of dust polarization using Planck instrumental noise maps introduced in Sect. 3 (Planck Collaboration III 2020). We have a total of 300 realizations of this noise at our disposal. We picked one of these, called $\mathbf{n}$, to build $\mathbf{d}$, and we used the remaining $M = 299$ noise maps, labeled $\{\mathbf{n}_1, \dots, \mathbf{n}_M\}$, for the denoising algorithm.

We define the S/N of $\mathbf{d}$ as the ratio of the standard deviations of $\Re(\mathbf{s})$ and $\Re(\mathbf{n})$ (in the following, the real part $\Re$ of a complex $Q + iU$ map refers to its Q map). We adjusted this S/N by scaling $\mathbf{s}$, so that the impact of the noise “resembles” that on the Planck map presented in Sect. 3. This was not straightforward since the power spectrum of the simulated map $\mathbf{s}$ has a different slope than that of the estimate of the power spectrum of the dust emission from the Planck map (see Figs. 2 and 5). We decided to adjust the S/N so that the scale k_i , at which the power spectra of $\mathbf{n}$ and $\mathbf{s}$ intersect, coincides with the one from the Planck map. Figures 2 and 5 show that $k_i/(2\pi) \approx 0.8 \text{ px}^{-1}$. This procedure leads to $S/N \approx 0.4$ for the simulated map, which is more than two times lower than the S/N of the observational map discussed in Sect. 3.

We show in Fig. 1 (top) the simulated Q maps corresponding to the truth $\mathbf{s}$ (left), the noise $\mathbf{n}$ (middle), and the resulting noisy map $\mathbf{d}$ (right). The map $\Re(\mathbf{s})$ exhibits coherent structures such as filaments and large smooth regions that are characteristic of its non-Gaussian statistical properties. On the other hand, $\Re(\mathbf{n})$ seems close to an inhomogeneous white Gaussian noise. The spatial inhomogeneity appears as variations of the local standard deviation due to the Planck satellite scanning strategy. Finally, in $\Re(\mathbf{d})$, coherent structures at intermediate and small scales are hard or impossible to identify.

We applied our denoising method to the simulated noisy map $\mathbf{d}$ and show, in Fig. 1 (bottom), the resulting denoised Q map $\Re(\tilde{\mathbf{s}})$ (left) next to the difference between the noisy and denoised maps $\Re(\mathbf{d} - \tilde{\mathbf{s}})$ (middle) and the denoising error $\Re(\mathbf{s} - \tilde{\mathbf{s}})$ (right). The map $\Re(\tilde{\mathbf{s}})$ shows that the noise level has been drastically reduced and that we were able to recover the filamentary structure down to a minimum scale. We can identify the smooth regions of $\Re(\mathbf{s})$ even if there still remains a visible noise. The similarities between $\Re(\mathbf{d} - \tilde{\mathbf{s}})$ and $\Re(\mathbf{n})$ are striking. The local variations of the standard deviation of the noise are clearly recovered, demonstrating that the inhomogeneity of the noise is not an issue for our method. Finally, $\Re(\mathbf{s} - \tilde{\mathbf{s}})$ exhibits some remaining structures that match the thinnest filaments appearing in $\Re(\mathbf{s})$, on top of a more diffuse background. This indicates that down to a minimum scale, below which the algorithm struggles to recover features, most of the structures are efficiently reconstructed.

¹ <https://github.com/bregaldo/pywph>

Fig. 1. *Top row:* simulated Q maps corresponding to the truth s (left), the noise n (middle), and the resulting noisy map d (right). *Bottom row:* denoised simulated Q map $\mathfrak{R}(\tilde{s})$ (left) next to the difference noisy-denoised $\mathfrak{R}(d - \tilde{s})$ (middle) and denoising error $\mathfrak{R}(s - \tilde{s})$ (right). Units are arbitrary, but kept consistent between the maps.

Fig. 2. Q maps power spectra for d , s , $\tilde{s}$, n (all in solid lines), $d - \tilde{s}$, and $s - \tilde{s}$, and cross-spectrum between the Q maps of $\tilde{s}$ and s (all in dashed lines). *Bottom panel:* ratio of the power spectra of d and $\tilde{s}$ with that of s .

Figure 2 compares the power spectra of the six maps shown in Fig. 1 plus the cross-spectrum between the denoised and truth maps. These power spectra were estimated computing the mean square moduli of the Fourier components of the maps binned on linearly spaced isotropic wavenumber bins, and the uncertainties correspond to the standard deviation of the mean. The cross-spectrum was computed similarly except for the bins above $k/(2\pi) = 0.14 \text{ px}^{-1}$ that were logarithmically spaced in

order to lower the statistical variance². We first point out that the power spectrum of d coincides with the sum of those of s and n because of the statistical independence between s and n . The power spectra of $\mathfrak{R}(\tilde{s})$ and $\mathfrak{R}(s)$ are in very good agreement with each other up to 0.18 px^{-1} , at which scale the noise power is ten times that of the signal. At smaller scales, where the noise dominates the signal even more, our algorithm is not able to retrieve the true power spectrum but gets closer to it. The power spectrum of $\mathfrak{R}(d - \tilde{s})$ coincides with that of $\mathfrak{R}(n)$ at small scales, and this agreement progressively worsens toward larger scales. This large-scale behavior shows that our algorithm does not remove the already negligible noise, as shown by the superposition of the power spectrum of $\mathfrak{R}(s - \tilde{s})$ and that of $\mathfrak{R}(n)$ at these scales. The cross-spectrum $\tilde{s} \times s$ is slightly below the power spectrum of s at intermediate scales and this discrepancy increases toward the smallest scales. At intermediate scales, where the power spectrum of $\mathfrak{R}(\tilde{s})$ matches that of $\mathfrak{R}(s)$, we suspect this discrepancy to stem from differences between the phases of the Fourier components of s and $\tilde{s}$ that would deserve a further quantification. Nevertheless, the production of a denoised map whose power spectrum coincides with that of s , even though n is ten times more powerful than s , and that retains a significant correlation with s is a striking success of our method.

To better characterize the non-Gaussianity of $\tilde{s}$, we computed the PDFs of the increments of Q for the noisy, denoised, and truth maps, and we plot them in Fig. 3 for three scalar lags. The increment $\delta Q_l(\mathbf{r})$ for a scalar lag l at a position $\mathbf{r}$ is the set of differences $\delta Q_l(\mathbf{r}) = Q(\mathbf{r}) - Q(\mathbf{r} + \mathbf{l})$ with $l \leq |\mathbf{l}| < l + 1$ in pixel units. These statistics are usually computed on velocity maps for the tails of the PDFs that characterize the intermittent dissipation of turbulence in the ISM (Hily-Blant et al. 2008;

² We recall that this cross-spectrum, called $\text{CS}[\mathfrak{R}(s), \mathfrak{R}(\tilde{s})]$, is related to the power spectra by the following relation: $\text{PS}[\mathfrak{R}(s - \tilde{s})] = \text{PS}[\mathfrak{R}(s)] + \text{PS}[\mathfrak{R}(\tilde{s})] - 2\text{CS}[\mathfrak{R}(s), \mathfrak{R}(\tilde{s})]$.

Fig. 3. PDFs of the increments of Q for d (noisy), s (truth), and $\tilde{s}$ (denoised), computed for three logarithmically spaced lags going from 2 to 32 pixels. For each lag, the increments were normalized by the standard deviation σ_l of the Gaussian that fits the core of the PDF of the noisy map. Bottom panels: ratio of the PDFs of $\tilde{s}$ with those of s for each lag.

Fig. 4. Noisy (left) and denoised (middle) Q maps of the Chamaeleon-Musca region as observed by the *Planck* satellite at 353 GHz. The noisy maps were denoised as described in Sect. 2.1. We also show the corresponding GNILC map (right) for reference.

Hily-Blant & Falgarone 2009). Contrary to the case of the noisy map, the distributions of increments for $\mathfrak{K}(s)$ are far from Gaussian for every lag. This is a clear signature of the non-Gaussianity of the data as we expect Gaussian-distributed increments for homogeneous Gaussian data. Our method recovers these statistics and their non-Gaussian tails with limited distortion for each lag, demonstrating its efficiency in retrieving non-Gaussianity in the data. A more thorough characterization of this non-Gaussianity has been postponed for further studies.

3. Application to *Planck* polarization data

3.1. Presentation of the data

We now apply our denoising method to a $Q+iU$ polarization map of the Chamaeleon-Musca region observed at 353 GHz with the *Planck* satellite (PR3 data³, Planck Collaboration III 2020). At this frequency, the CMB is negligible and the dominant components are the dust emission and the noise (see Fig. 34 of Planck Collaboration IV 2020). We consider the Q and U maps corresponding to the full mission, and those corresponding to the two half-missions, as well as the 300 end-to-end simulated

Fig. 5. Q maps power spectra for d , $\tilde{s}$, n (all in solid lines), and $d - \tilde{s}$, and cross-spectrum between d_1 and d_2 (both in dashed lines). This cross-spectrum estimates the power spectrum of the dust emission. Also shown in a solid line is the power spectrum of the GNILC Q map, corresponding to the same Chamaeleon-Musca field.

³ https://wiki.cosmos.esa.int/planck-legacy-archive/index.php/Main_Page

Q and U maps of the noise and the systematics of the instrument for the full mission (the ones used in Sect. 2). We projected all of

Fig. 6. PDFs of the Chamaeleon-Musca Q maps increments for $\mathbf{d}$ (noisy), $\tilde{\mathbf{s}}$ (denoised), and GNILC map, computed for three logarithmically spaced lags ranging from $4.7'$ to $75.2'$. For each lag, the increments were normalized by the standard deviation σ_l of the Gaussian that fits the core of the PDF of the noisy map.

these maps on 512×512 grids with a pixel size of $2.35'$, centered on the region of the Chamaeleon-Musca clouds at Galactic coordinates $(l, b) = (300.26^\circ, -16.77^\circ)$. We used a Gnomonic projection through the HEALPix/healpy⁴ package (Górski et al. 2005; Zonca et al. 2019).

In Fig. 4, we show the projected full-mission Q map that we aim to denoise (left), as well as the denoised map discussed in the next section (middle). The same maps for U are given in Fig. B.2.

3.2. Denoising results

We applied the denoising method presented in Sect. 2 to the full mission $Q + iU$ map, using the corresponding 300 noises. Figures 4 and B.2 (middle) show the resulting denoised Q and U maps, respectively. The overall noise level has been clearly mitigated although subtle residuals of the patterns due to the scanning strategy remain, and we can now discern a more complex layout of structures even in the regions where the signal is weak.

In Fig. 5, we show a power spectrum analysis of the denoising of the Q map, comparing the power spectra of the noisy and denoised maps $\mathcal{R}(\mathbf{d})$ and $\mathcal{R}(\tilde{\mathbf{s}})$, of one noise map $\mathcal{R}(\mathbf{n})$, and of the difference between the noisy and denoised maps $\mathcal{R}(\mathbf{d} - \tilde{\mathbf{s}})$, as well as the cross-spectrum between the two half-missions maps. This cross-spectrum gives an estimate for the power spectrum of the truth signal since the noise signals are independent. It is satisfactory to see that the power spectrum of $\mathcal{R}(\tilde{\mathbf{s}})$ is consistent with this cross-spectrum for scales down to $k/(2\pi) \sim 3.4 \text{ deg}^{-1}$. Also, $\mathcal{R}(\mathbf{d} - \tilde{\mathbf{s}})$ behaves similarly to what we have observed in Sect. 2.2, with a power spectrum consistent with that of $\mathcal{R}(\mathbf{n})$ when the noise has a larger power than the signal, and it falls below what is expected when the emission of the dust begins to dominate. Similar results are obtained for U .

Figure 6 compares the PDFs of the increments of Q for the noisy and denoised maps, and for the same three scalar lags as in Fig. 3. Based on the results of Sect. 2, we expect the PDFs of the denoised map to be good estimates for the statistics of the truth map. These results show important differences between the noisy and denoised maps that are again a signature of the non-Gaussianity of the truth signal. We find similar results for the increments of U .

The literature on image denoising is rich and abundant (see for instance Buades et al. 2010; Dabov et al. 2007; Zhang et al. 2017), and a thorough comparison of our method with other denoising algorithms would be beyond the scope of this Letter. Nevertheless, to give some elements of comparison, we compare our method in Appendix B to the following: (1) Wiener filtering and sampling methods, which are widely used in astrophysics; and (2) the GNILC method (Planck Collaboration Int. XLVIII 2016), which was used on *Planck* polarization data and provides a local smoothing kernel in order to optimally remove the noise. For the Wiener approach, we find that neither the Wiener filtered image nor realizations drawn from the Wiener posterior distribution are able to retrieve the PDFs of the increments, even when the true power spectrum was given as input. The comparison with GNILC shows that our method recovers the true power spectrum and the PDFs of the increments better.

4. Conclusion and prospects

We have introduced a new method for the denoising of *Planck* polarization data based on the generation of random synthetic maps from WPH statistics that characterize the non-Gaussianity of the dust emission. This method takes advantage of the strong statistical differences between the signal of interest (non-Gaussian and regular) and the noise (close to Gaussian and irregular) by performing an optimization that constrains the statistical properties of the denoised map plus noise realizations.

We applied our method to mock $Q + iU$ noisy data designed to emulate typical *Planck* polarization maps of dust in the diffuse ISM. The denoised map has a power spectrum that coincides with the true power spectrum down to a minimum scale where the power of the noise is ten times that of the signal, while being highly correlated with the truth signal. It recovers the PDFs of the increments for various isotropic lags, demonstrating that our method is able to retrieve non-Gaussianity in the data. Finally, we applied this method to a 353 GHz *Planck* observation of the Chamaeleon-Musca field.

Our method has been introduced as a statistical denoising method, but it should be more generally applicable to component separation problems. In particular, we expect this method to be efficient at disentangling dust emission from the CMB, cosmic infrared background, and noise in *Planck* maps, provided that we have relevant models of each of these contaminants at

⁴ <http://healpix.sourceforge.net>

our disposal. Therefore, it could hopefully enhance the scientific outcome of *Planck* data.

One of the main motivations of our work is to build a generative model of the dust polarization signal that takes into account the non-Gaussianity of the data. Such a model derived from the analysis of *Planck* data may be used to simulate the dust polarization sky. Our denoising algorithm constitutes a step toward this modeling, in the sense that WPH statistics of the dust emission, corrected to a first approximation from noise contamination, may be computed from the denoised map. A natural extension of this work would be to quantitatively assess how well we can recover the WPH statistics of the truth map with this approach. This will be the subject of a future paper.

Acknowledgements. We thank the anonymous referee for a very helpful report. This work was undertaken in parallel to a similar project led by J.-M. Delouis on the full sky, and we acknowledge a fruitful interaction with his team. We gratefully acknowledge P. Lesaffre for helping us with the computation of the PDFs of the increments. We also acknowledge fruitful discussions with J.-F. Cardoso, S. Mallat, and T. Marchand. This research was supported by the Agence Nationale de la Recherche (project BxB: ANR-17-CE31-0022). B. Regaldo-Saint Blancard acknowledges support from the Centre National d'Etudes Spatiales (CNES).

References

- Allys, E., Levrier, F., Zhang, S., et al. 2019, *A&A*, **629**, A115
- Allys, E., Marchand, T., Cardoso, J. F., et al. 2020, *Phys. Rev. D*, **102**, 103506
- Alsing, J., Heavens, A., & Jaffe, A. H. 2017, *MNRAS*, **466**, 3272
- Aylor, K., Haq, M., Knox, L., Hezaveh, Y., & Perreault-Levasseur, L. 2020, *MNRAS*, **500**, 3889
- Beck, D., Errard, J., & Stompor, R. 2020, *JCAP*, **2020**, 030
- BICEP2/Keck and Planck Collaborations 2015, *Phys. Rev. Lett.*, **114**, 101301
- Buades, A., Coll, B., & Morel, J. M. 2010, *SIAM Rev.*, **52**, 113
- Burkhart, B., Falceta-Gonçalves, D., Kowal, G., & Lazarian, A. 2009, *ApJ*, **693**, 250
- Byrd, R., Lu, P., Nocedal, J., & Zhu, C. 1995, *SIAM J. Sci. Comput.*, **16**, 1190
- Cheng, S., Ting, Y.-S., Ménard, B., & Bruna, J. 2020, *MNRAS*, **499**, 5902
- Clark, S. E., & Hensley, B. S. 2019, *ApJ*, **887**, 136
- Dabov, K., Foi, A., Katkovnik, V., & Egiazarian, K. 2007, *IEEE Trans. Image Process.*, **16**, 2080
- Draine, B. T. 2011, *Physics of the Interstellar and Intergalactic Medium* (Princeton: Princeton University Press)
- Elsner, F., & Wandelt, B. D. 2013, *A&A*, **549**, A111
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
- Hensley, B. S., & Bull, P. 2018, *ApJ*, **853**, 127
- Hily-Blant, P., & Falgarone, E. 2009, *A&A*, **500**, L29
- Hily-Blant, P., Falgarone, E., & Pety, J. 2008, *A&A*, **481**, 367
- Jasche, J., & Lavaux, G. 2015, *MNRAS*, **447**, 1204
- Jeffrey, N., Abdalla, F. B., Lahav, O., et al. 2018a, *MNRAS*, **479**, 2871
- Jeffrey, N., Heavens, A. F., & Fortio, P. D. 2018b, *Astron. Comput.*, **25**, 230
- Kodi Ramanah, D., Lavaux, G., & Wandelt, B. D. 2019, *MNRAS*, **490**, 947
- Krachmalnicoff, N., & Puglisi, G. 2021, *ApJ*, **911**, 42
- Mallat, S., Zhang, S., & Rochette, G. 2020, *Information and Inference: A Journal of the IMA*, **9**, 721
- Pelgrims, V., Ferrière, K., Boulanger, F., Lallement, R., & Montier, L. 2020, *A&A*, **636**, A17
- Petroff, M. A., Addison, G. E., Bennett, C. L., & Weiland, J. L. 2020, *ApJ*, **903**, 104
- Planck Collaboration III. 2020, *A&A*, **641**, A3
- Planck Collaboration IV. 2020, *A&A*, **641**, A4
- Planck Collaboration XI. 2020, *A&A*, **641**, A11
- Planck Collaboration XII. 2020, *A&A*, **641**, A12
- Planck Collaboration Int. XLVIII. 2016, *A&A*, **596**, A109
- Regaldo-Saint Blancard, B., Levrier, F., Allys, E., Bellomi, E., & Boulanger, F. 2020, *A&A*, **642**, A217
- Saydjari, A. K., Portillo, S. K. N., Slepian, Z., et al. 2021, *ApJ*, **910**, 122
- Thorne, B., Knox, L., & Prabhu, K. 2021, *MNRAS*, **504**, 2603
- Wandelt, B. D., Larson, D. L., & Lakshminarayanan, A. 2004, *Phys. Rev. D*, **70**, 083511
- Wiener, N. 1949, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series* (Cambridge, MA: MIT Press), 7
- Zaroubi, S., Hoffman, Y., Fisher, K. B., & Lahav, O. 1995, *ApJ*, **449**, 446
- Zhang, S., & Mallat, S. 2021, *Appl. Comput. Harmonic Anal.*, **53**, 199
- Zhang, K., Zuo, W., Chen, Y., Meng, D., & Zhang, L. 2017, *IEEE Trans. Image Process.*, **26**, 3142
- Zonca, A., Singer, L., Lenz, D., et al. 2019, *J. Open Sour. Softw.*, **4**, 1298

Appendix A: The WPH statistics

In this appendix, we introduce the WPH statistics used in this Letter, referring to [Allys et al. \(2020\)](#) and [Mallat et al. \(2020\)](#) for additional details. We also describe the detailed denoising procedure introduced in Sect. 2.1. We call X a given random field and $\mathbf{x}$ one of its realizations.

A.1. Bump-steerable wavelets

The WPH statistics rely on a bank of wavelets that are dilations and rotations of a mother bump-steerable wavelet ψ defined in Fourier space as follows ([Mallat et al. 2020](#)):

$$\hat{\psi}(\mathbf{k}) = \exp\left(\frac{-(k - \xi_0)^2}{\xi_0^2 - (k - \xi_0)^2}\right) \cdot \mathbf{1}_{[0, 2\xi_0]}(k) \times \cos^{L-1}(\arg(\mathbf{k})) \cdot \mathbf{1}_{[0, \pi/2]}(|\arg(\mathbf{k})|), \quad (\text{A.1})$$

with $k = \|\mathbf{k}\|$, $\mathbf{1}_A(x)$ being the indicator function that returns 1 if $x \in A$ and 0 otherwise, and $\xi_0 = 0.85\pi$ being the central wavenumber of the mother wavelet. The dilated and rotated wavelets are defined as follows:

$$\psi_{\xi_{j,\theta}}(\mathbf{r}) = 2^{-j}\psi(2^{-j}\text{rot}_{-\theta}[\mathbf{r}]), \quad (\text{A.2})$$

with j being the index of the dilation by a factor 2^j and θ being the angle of rotation. We call $\xi_{j,\theta} = 2^{-j}\xi_0\mathbf{u}_\theta$ with $\mathbf{u}_\theta = \cos\theta\mathbf{u}_x + \sin\theta\mathbf{u}_y$ the central wavevector of the wavelet obtained by such a transform. In practice, we consider J dilation indices j ranging from 0 to $J-1$, and we divide 2π into $2L$ evenly spaced rotation angles θ . The bank of wavelets is thus the set of $2JL$ wavelets $\{\psi_{\xi_{j,\theta}} : j = 0, \dots, J-1, \theta = 0, \frac{\pi}{L}, \dots, \frac{(2L-1)\pi}{L}\}$, and these wavelets cover most of Fourier space with their respective band-passes.

In this Letter, we work with 512×512 maps and choose $J = 8$ and $L = 8$. We show in Fig. A.1 one example wavelet from the bank.

A.2. WPH moments

The WPH moments of X are covariances of the phase harmonics of the wavelet transform of X . With $\{\psi_{\xi_1}, \dots, \psi_{\xi_N}\}$, our bank of wavelets labeled by their central wavevectors ξ_i , the wavelet transform of X corresponds to the set $\{X \star \psi_{\xi_1}, \dots, X \star \psi_{\xi_N}\}$, where $\star$ denotes the convolution operation. These convolutions amount to a local band-pass filtering of X on the scales probed by each of the wavelets. For a homogeneous X , the WPH moments of X are as follows:

$$C_{\xi_1, p_1, \xi_2, p_2}(\tau) = \text{Cov}\left([X \star \psi_{\xi_1}(\mathbf{r})]^{p_1}, [X \star \psi_{\xi_2}(\mathbf{r} + \tau)]^{p_2}\right), \quad (\text{A.3})$$

with $z \mapsto [z]^p = |z| \cdot e^{ip \arg(z)}$ being the phase harmonic operator⁵.

These WPH moments are able to capture interactions between different scales of X thanks to the phase harmonic operator. Indeed, the covariance between $X \star \psi_{\xi_1}$ and $X \star \psi_{\xi_2}$ vanishes when the wavelets ψ_{ξ_1} and ψ_{ξ_2} have nonintersecting band-passes, and it is otherwise a function of the power spectrum of X only. With proper p_i and p_j indices, the phase harmonic operator can make $[X \star \psi_{\xi_1}]^{p_1}$ and $[X \star \psi_{\xi_2}]^{p_2}$ comparable in the sense that they share common spatial frequencies, allowing an extraction of high-order information through their covariance.

⁵ These moments do not depend on the $\mathbf{r}$ variable because of the homogeneity of X .

Fig. A.1. Real part (left) and imaginary part (right) of a 512×512 bump-steerable wavelet $\psi_{6,\pi/4}$. The wavelet is centered in the middle of the map for a better visualization.

We discretize the τ variable similarly to [Allys et al. \(2020\)](#). This introduces a grid of $\tau_{n,\alpha}$ vectors for each wavelet $\psi_{j,\theta}$, defined as

$$\tau_{n,\alpha} = 3n2^j\mathbf{u}_{\theta+\alpha}, \quad (\text{A.4})$$

with $n = 0, \dots, \Delta_n$ and $\alpha \in \{-\pi/4, 0, \pi/4, \pi/2\}$. To avoid a redundancy in the information contained in the coefficients, we discarded the translations for which $n > \min(J-1-j, \Delta_n)$.

[Allys et al. \(2020\)](#) identified a relevant set of WPH moments to build a generative model of real-valued simulated data of the large-scale structure of the Universe, which are highly non-Gaussian at late times and at scales smaller than $100 h^{-1}$ Mpc. In the present work, we have adopted the same corresponding WPH statistics and applied them similarly on complex-valued data, except for the scaling moments discussed below. In practice we estimate a set of moments that can be divided into the following five categories:

- the $S^{(1,1)}$ moments, of the form $C_{\xi_1, \xi_2, 1}$, at every $\tau_{n,\alpha}$ with $\Delta_n = 2$. They capture the power spectrum information of the ξ band-pass.
- the $S^{(0,0)}$ moments, of the form $C_{\xi, 0, \xi, 0}$, at every $\tau_{n,\alpha}$ with $\Delta_n = 2$. They capture information related to the sparsity of the data in the ξ band-pass.
- the $S^{(0,1)}$ moments, of the form $C_{\xi, 0, \xi, 1}$, at $\tau = 0$ only. They capture information related to the couplings between the scales of the same ξ band-pass.
- the $C^{(0,1)}$ moments, of the form $C_{\xi_1, 0, \xi_2, 1}$, considering $0 \leq j_1 < j_2 \leq J-1$ and $|\theta_1 - \theta_2| \leq \frac{\pi}{2}$, at every $\tau_{n,\alpha}$ with $\Delta_n = 2$ when $\theta_1 = \theta_2$ and at $\tau = 0$ when $\theta_1 \neq \theta_2$. They capture information related to the correlation between local levels of oscillation for the scales in the ξ_1 and ξ_2 band-passes.
- the C^{phase} moments, of the form $C_{\xi_1, 1, \xi_2, p_2}$ with $p_2 = \xi_1/\xi_2 > 1$, considering $0 \leq j_1 < j_2 \leq J-1$ and $\theta_1 = \theta_2$, at every $\tau_{n,\alpha}$ with $\Delta_n = 2$. They capture information related to the statistical phase alignment of oscillations between the scales in the ξ_1 and ξ_2 band-passes.

We call $\tilde{C}_{\xi_1, p_1, \xi_2, p_2}(\tau)$ the normalized estimates of the WPH moments $C_{\xi_1, p_1, \xi_2, p_2}(\tau)$. In practice they are defined as follows for a map $\mathbf{x}$:

$$\tilde{C}_{\xi_1, p_1, \xi_2, p_2}(\tau) = \frac{\langle \mathbf{x}^{(\xi_1, p_1)}(\mathbf{r}) \overline{\mathbf{x}^{(\xi_2, p_2)}(\mathbf{r} + \tau)} \rangle}{\sqrt{\langle |\mathbf{x}_0^{(\xi_1, p_1)}|^2 \rangle \langle |\mathbf{x}_0^{(\xi_2, p_2)}|^2 \rangle}}, \quad (\text{A.5})$$

where the brackets stand for a spatial mean on the $\mathbf{r}$ variable, the bar denotes the complex conjugate, and $\mathbf{x}^{(\xi, p)} = [\mathbf{x} \star \psi_\xi]^p - \langle [\mathbf{x}_0 \star \psi_\xi]^p \rangle$. These estimates depend on a

reference map $\mathbf{x}_0$. We found that the choice of $\mathbf{x}_0$ can have a significant impact on the results of our denoising method as explained below.

A.3. Scaling moments

Similarly to [Allys et al. \(2020\)](#), we complement these statistics with a small number of coefficients, the estimates of the so-called scaling moments $L_{j,p}$ that better constrain the largest scales that are not probed by the WPH moments. These moments were constructed from a bank of scaling functions $\{\varphi_j\}_{0 \leq j \leq J-1}$, which correspond to dilations of an isotropic Gaussian filter φ defined in Fourier space by the following:

$$\hat{\varphi}(\mathbf{k}) = \exp\left(-\frac{\|\mathbf{k}\|^2}{2\sigma^2}\right), \quad (\text{A.6})$$

with $\sigma = 0.496 \times 2^{-0.55} \xi_0$. For a real-valued random field $\mathbf{X}$, the scaling moments of $\mathbf{X}$ are as follows:

$$L_{j,0} = \text{Cov} \left[|\mathbf{X} \star \varphi_j|, |\mathbf{X} \star \varphi_j| \right], \quad (\text{A.7})$$

$$L_{j,p} = \text{Cov} \left[\left(\mathbf{X} \star \varphi_j \right)^p, \left(\mathbf{X} \star \varphi_j \right)^p \right] \text{ (for } p > 0). \quad (\text{A.8})$$

In practice, we only estimated the scaling moments with $p \in \{0, 1, 2, 3\}$ and $2 \leq j \leq J-2$. We define the corresponding normalized estimates computed for a map $\mathbf{x}$ by

$$\tilde{L}_{j,p} = \frac{\langle |\mathbf{x}^{(j,p)}|^2 \rangle}{\langle |\mathbf{x}_0^{(j,p)}|^2 \rangle}, \quad (\text{A.9})$$

with

$$\mathbf{x}^{(j,0)} = |(\mathbf{x} - \langle \mathbf{x}_0 \rangle) \star \varphi_j| - \langle |(\mathbf{x} - \langle \mathbf{x}_0 \rangle) \star \varphi_j| \rangle, \quad (\text{A.10})$$

$$\mathbf{x}^{(j,p)} = ((\mathbf{x} - \langle \mathbf{x}_0 \rangle) \star \varphi_j)^p - \langle ((\mathbf{x} - \langle \mathbf{x}_0 \rangle) \star \varphi_j)^p \rangle \text{ (for } p > 0). \quad (\text{A.11})$$

These normalized estimates depend on the same reference map $\mathbf{x}_0$ introduced for the estimates of the WPH moments of $\mathbf{X}$. As we deal with complex maps $\mathbf{x}$ in this work, we computed these estimates for their real and imaginary parts separately.

A.4. Denoising procedure

We applied our denoising method to a simulated noisy map $\mathbf{d}$ in two steps. We performed a first denoising of $\mathbf{d}$ using a ϕ operator that only takes into account the estimates of the $S^{(1,1)}$ moments, related to the power spectrum information, and the estimates of the scaling moments, and by normalizing the estimates choosing $\mathbf{u}_0 = \mathbf{d}$. This first step yields a map $\tilde{\mathbf{s}}_0$ that has a power spectrum that is much more consistent with the truth map $\mathbf{s}$. Then, we performed a second denoising of $\mathbf{d}$ using a ϕ operator that includes the whole set of estimates of WPH moments and scaling moments, but here normalizing these estimates with $\mathbf{u}_0 = \tilde{\mathbf{s}}_0$. We call $\tilde{\mathbf{s}}$ the output of this second step. This particular choice for the normalization of the estimates has proven to be more efficient at retrieving the power spectrum and the PDFs of the increments of the truth map $\mathbf{s}$.

In [Table A.1](#), we show the resulting number of statistical coefficients per class of moments computed for a complex 512×512 pixels map $\mathbf{x}$ and for the chosen parameters. In this Letter, the ϕ operator used for the second step of the procedure computes a set of 11 176 statistical coefficients (of which 40 correspond to the estimates of the scaling moments), which corresponds to approximately 4% of the number of pixels of the input map.

Table A.1. Number of statistical coefficients per class of moments for the parameters of this Letter.

$S^{(1,1)}$	$S^{(0,0)}$	$S^{(0,1)}$	$C^{(0,1)}$	C^{phase}	$L_{j,p}$	Total
960	960	128	6336	2752	40	11 176

Appendix B: Comparison to other methods

B.1. Wiener filtering and posterior sampling

Wiener filtering is a typical approach to signal inference for problems of the form given by [Eq. \(1\)](#). In this case, the Wiener filtered reconstruction $\tilde{\mathbf{s}}_W$ ([Wiener 1949](#); [Zaroubi et al. 1995](#)) is given by

$$\tilde{\mathbf{s}}_W = \mathbf{W}\mathbf{d} = \mathbf{S}(\mathbf{S} + \mathbf{N})^{-1}\mathbf{d}, \quad (\text{B.1})$$

where $\mathbf{S} = \langle \mathbf{s}\mathbf{s}^\dagger \rangle$ and $\mathbf{N} = \langle \mathbf{n}\mathbf{n}^\dagger \rangle$ are the signal and noise covariance matrices, respectively. For brevity we have assumed $\langle \mathbf{s} \rangle = \langle \mathbf{n} \rangle = 0$. This is the linear filter that gives the expected minimum variance of residuals $(\mathbf{W}\mathbf{d} - \mathbf{s})$ for known $\mathbf{S}$ and $\mathbf{N}$.

Furthermore, it is both the maximum a posteriori and mean posterior solution if the noise and signal are drawn from Gaussian distributions. That is, if the real-valued data are distributed with a likelihood

$$p(\mathbf{d}|\mathbf{s}, \mathbf{N}) = \frac{1}{\sqrt{(\det 2\pi\mathbf{N})}} \exp\left[-\frac{1}{2}(\mathbf{d} - \mathbf{s})^\dagger \mathbf{N}^{-1}(\mathbf{d} - \mathbf{s})\right], \quad (\text{B.2})$$

and the prior distribution for the real-valued signal is

$$p(\mathbf{s}|\mathbf{S}) = \frac{1}{\sqrt{(\det 2\pi\mathbf{S})}} \exp\left[-\frac{1}{2}\mathbf{s}^\dagger \mathbf{S}^{-1}\mathbf{s}\right], \quad (\text{B.3})$$

then the Wiener posterior distribution is given by

$$p(\mathbf{s}|\mathbf{S}, \mathbf{N}, \mathbf{d}) \propto p(\mathbf{d}|\mathbf{s}, \mathbf{N})p(\mathbf{s}|\mathbf{S}) \propto \exp\left[-\frac{1}{2}(\mathbf{s} - \mathbf{W}\mathbf{d})^\dagger (\mathbf{S}^{-1} + \mathbf{N}^{-1})(\mathbf{s} - \mathbf{W}\mathbf{d})\right]. \quad (\text{B.4})$$

For the problem described in this work, the noise covariance $\mathbf{N}$ can be estimated from the samples of noise realizations and the signal covariance $\mathbf{S}$ can either be assumed or jointly estimated from the data ([Wandelt et al. 2004](#)).

We avoided inverting the dense $\mathbf{W}$ matrix by implementing a messenger field approach, first described by [Elsner & Wandelt \(2013\)](#) and now widely used in cosmology (e.g., [Jasche & Lavaux 2015](#); [Alsing et al. 2017](#); [Jeffrey et al. 2018a,b](#); [Kodi Ramanah et al. 2019](#)). We assumed that the underlying signal is homogeneous and isotropic, such that the signal covariance $\mathbf{S}$ is diagonal in harmonic space with elements corresponding to the one-dimensional power spectrum. The uncorrelated noise gives a diagonal noise covariance $\mathbf{N}$ in pixel space, so that the messenger field can efficiently iterate between harmonic and pixel space.

As we are concerned with polarization data, we used a spin-2 harmonic transformation between $\{\mathbf{s}_Q, \mathbf{s}_U\}$ and $\{\mathbf{s}_E, \mathbf{s}_B\}$, using the curl-free E -mode and divergence-free B -mode representation. The signal covariance in harmonic space is therefore a concatenation of $\{\mathbf{S}_E, \mathbf{S}_B\}$. This formulation preserves the relevant $Q-U$ correlation.

Even if the Wiener filtered reconstructed signal $\tilde{\mathbf{s}}_W$ is the maximum of the posterior distribution $p(\mathbf{s}|\mathbf{d})$, functions $f(\tilde{\mathbf{s}}_W)$ do not correspond to the maximum $p(f(\mathbf{s})|\mathbf{d})$. For example, the two-point

Fig. B.1. *Left:* Wiener filtered map $\tilde{s}_W$. *Middle:* same as Fig. 2, but for $\tilde{s}_W$. The power spectrum of $\tilde{s}_W$ is suppressed as expected. *Right:* same as Fig. 3, but for $l = 2$ px with the Wiener filter (mean) map and a Wiener posterior sample.

Fig. B.2. Same as Fig. 4, but for U .

statistics (e.g., variance or power spectra) of $\tilde{s}_W$ are not unbiased estimates of the power spectrum of s . Instead, we can draw sample images s_i from the posterior $p(s|\mathbf{d})$ (Eq. (B.4)), so that the transformed samples $f(s_i)$ are correctly distributed according to $p(f(s)|\mathbf{d})$. These realizations are generated by amending the messenger field algorithm (Elsner & Wandelt 2013).

Figure B.1 shows the Wiener filtered reconstruction of the simulated Q map (left panel). The low amplitude power spectrum of the residual map $\mathcal{R}(s - \tilde{s}_W)$ demonstrates that the pixel values are close to the truth, whereas the power spectrum of $\mathcal{R}(\tilde{s}_W)$ is biased low (center panel), which is as expected. To compare the methods, we compared with functions of samples s_i from the Wiener posterior. Though not plotted here, as we input the true power spectrum, which was not done for the WPH denoising method, the realizations s_i did have the correct power spectra with relatively small sample variance.

As a goal of this denoising work is to retain the statistical non-Gaussianity in the signal, which could be represented in the WPH coefficients, we again compared the PDF of increments δQ_l . The right panel of Fig. B.1 shows the comparison between the data $\mathcal{R}(\mathbf{d})$, the signal $\mathcal{R}(s)$, the Wiener filtered (mean) map $\mathcal{R}(\tilde{s}_W)$, and a realization $\mathcal{R}(\tilde{s}_{W,\text{sample}})$ drawn from the Wiener posterior. We see that neither the Wiener mean map $\mathcal{R}(\tilde{s}_W)$ nor the sample $\mathcal{R}(\tilde{s}_{W,\text{sample}})$ capture the non-Gaussianity in terms of δQ_l increments as successfully as our WPH denoising method. Indeed, for $|\delta Q_l|/\sigma_l \gtrsim 1$, there is a clear divergence from the true PDF, showing that the tails of the true statistics are not recovered at all.

Furthermore, inspection of the WPH coefficients of the Wiener posterior samples $\phi(s_i)$ shows them to be noticeably further from those of the truth map than $\phi(\tilde{s})$. These preliminary results are encouraging, but not surprising. The Gaussian prior distribution in the statistical model of the Wiener posterior leads

to a poor recovery of the non-Gaussian structures that are intrinsic to the polarized dust emission.

B.2. GNILC

The GNILC method is a wavelet-based component separation method that is designed to extract the emission of the Galactic foregrounds from the *Planck* full-sky maps. It has been applied to polarization maps Q and U (Planck Collaboration IV 2020; Planck Collaboration XII 2020) to disentangle the thermal dust polarization emission from the CMB polarization and instrumental noise over the entire sky. We show in Figs. 4 and B.2 (right column) the resulting Q and U maps, respectively, for the Chamaeleon-Musca field observed at 353 GHz. We note that at this frequency, the CMB can be safely ignored, so that the main effect of the GNILC algorithm is to denoise the dust emission. In these maps, we clearly see that the smallest scales, most contaminated by the noise, have been filtered out. Compared to GNILC, our denoised maps thus include a wider variety of structures at intermediate and small scales.

A quantitative comparison based on the power spectrum and the PDFs of the increments for the Q maps is shown in Figs. 5 and 6. The power spectrum of the GNILC map plummets at small scales and exhibits a lack of power compared to that of $\tilde{s}$ and the cross-spectrum $\mathbf{d}_1 \times \mathbf{d}_2$. The PDFs of the increments also show important discrepancies, especially at the tails of the PDFs, between the $\tilde{s}$ map and the GNILC map. Assuming that the statistics of $\tilde{s}$ give a relevant order of magnitude of the true statistics, based on the validation on simulated data, this suggests a distortion of these statistics by the GNILC method. This quantitative comparison demonstrates the superiority of our method for recovering the true power spectrum and, a priori, the PDFs of the increments.