

Point Source Detection and False Discovery Rate Control on CMB Maps

Carrón Duque, Javier^{a,*}, Buzzelli, Alessandro^{a,e}, Fantaye, Yabebal^b, Marinucci, Domenico^c, Schwartzman, Armin^d, Vittorio, Nicola^{a,e}

^a*Dipartimento di Fisica, Università di Roma Tor Vergata, via della Ricerca Scientifica 1, I-00133, Roma, Italy*

^b*African Institute for Mathematical Sciences and Department of Mathematics, University of Stellenbosch, South Africa*

^c*Dipartimento di Matematica, Università di Roma Tor Vergata, via della Ricerca Scientifica 1, I-00133, Roma, Italy*

^d*Division of Biostatistics, University of California, San Diego, USA*

^e*Sezione INFN Roma 2, via della Ricerca Scientifica 1, I-00133, Roma, Italy*

Abstract

We discuss the *STEM* (Smoothing and Testing Multiple hypotheses) procedure to search for point sources in Cosmic Microwave background maps; in particular, we aim at controlling the so-called False Discovery Rate, which is defined as the expected value of false discoveries among pixels which are labelled as contaminated by point sources. *STEM* is based on the following four steps: 1) needlet filtering of the observed CMB maps, to improve the signal to noise ratio; 2) selection of candidate peaks, i.e., the local maxima of filtered maps; 3) computation of p -values for local maxima; 4) implementation of the multiple testing procedure, by means of the so-called Benjamini-Hochberg method. These procedures are also implemented on the latest release of Planck CMB maps.

Keywords: False Discovery Rate, Multiple Testing, Point Source Detection, Needlets, Cosmic Microwave Background
2010 MSC: Primary 62M40; Secondary 62M30, 62M15, 60G60, 42C40

1. Introduction

In the last decades, observations and characterization of Cosmic Microwave Background (CMB) temperature anisotropy established the foundations for the Λ CDM concordance model (see, e.g., BOOMERanG: MacTavish et al. 2006; WMAP: Hinshaw et al. 2013; Planck: Planck Collaboration XIII 2016). According to this scenario, the Universe is described by a flat Euclidean geometry, with cosmic structures originated by an almost scale-invariant spectrum of adiabatic Gaussian primordial fluctuations, and the cosmic energy density is in the form of barionic matter ($\sim 5\%$), cold Dark Matter ($\sim 26.5\%$) and Dark Energy ($\sim 68.5\%$).

However, an accurate analysis of the CMB polarization pattern is required in order to break the degeneracy among some parameters and to constrain critical aspects of the Early Universe. CMB polarization is usually decomposed into a gradient and a curl component, so-called E and B modes, respectively (Kamionkowski et al., 1997). E-modes have been widely detected (see, e.g., Kovac et al., 2002; Planck Collaboration XI, 2016), while primordial B-modes have escaped observation so far (BICEP2/Keck and

Planck Collaborations, 2015). The detection of the B-mode polarization represents nowadays the new frontier of observational cosmology, as B-modes would provide an ultimate confirmation to the existence of a stochastic primordial background of gravitational waves, as predicted by inflationary models (Lyth and Riotto, 1999).

The primary concern in CMB observations is the presence of foreground contamination, due to both diffuse Galactic emission (e.g. Planck Collaboration X 2016) and extragalactic (point) sources (e.g. Planck Collaboration XXVI 2016). In intensity, Galactic emission consists in a mixture of many processes, such as free-free emission, anomalous microwave emission eventually due to non-thermal (spinning) dust and synchrotron emission at low frequency, and thermal dust emission at high frequency. Extragalactic sources, as well, have different astrophysical origin. In particular, the two main classes of sources are radio galaxies and dusty galaxies that are present mainly at low and high frequencies, respectively.

To disentangle the CMB signal from foreground contamination, many approaches have been followed (Planck Collaboration IV, 2018). The latest released Planck cleaned-CMB maps have been produced by the SEVEM template fitting technique, the NILC and SMICA non-parametric procedures and the COMMANDER parametric method. Despite the difference in their approaches, these methods yield largely compatible output maps.

In this work we investigate the application of a general method for the localization of peaks on the sphere, under isotropic Gaussian noise, to detect point sources out

*Corresponding author

Email addresses: javier.carron@uniroma1.it (Carrón Duque, Javier), alessandro.buzzelli@roma2.infn.it (Buzzelli, Alessandro), yabi@aims.ac.za (Fantaye, Yabebal), marinucc@mat.uniroma2.it (Marinucci, Domenico), armins@ucsd.edu (Schwartzman, Armin), nicola.vittorio@uniroma2.it (Vittorio, Nicola)

of CMB maps. On one hand, this identification may allow a better cleaning of the CMB maps from the source contaminants; on the other hand, these sources are of astrophysical interest themselves. Here, we focus on intensity CMB maps and leave the extension to polarization to a forthcoming paper. Our method is based on the so-called STEM procedure (see Cheng et al., 2016), which consists of the following four steps: i) filtering the CMB maps with (Mexican) needlets, in order to increase the signal-to-noise ratio (see Marinucci et al., 2008); ii) detection of local maxima in the filtered maps; iii) computation of p -values for each of the local maxima; iv) implementation of the multiple testing scheme by applying the so-called Benjamini-Hochberg (BH) procedure. The algorithm allows to effectively control the False Discovery Rate, i.e. the expected number of false discoveries among the critical points identified as point sources. We validate this procedure on realistic simulations of CMB maps and, then, apply the algorithm to the latest release of Planck cleaned-CMB maps; in particular, we compare the performance of the different component separation methods in the removal of the extragalactic sources. We find some candidate point sources in most of the maps; however, we conclude that the inpainting procedure adopted for the 2018 release seems to have produced maps which are much closer to being purely Gaussian, with a number of local maxima largely consistent with theoretical predictions.

This paper is organized as follows: in Section 2 we provide a formal description of STEM procedure; in Section 3 we present the numerical implementation of the algorithm and discuss about its performance on simulations; in Section 4 we show our results from Planck cleaned-CMB maps; finally, in Section 5 we draw our conclusions and outline some future perspectives for the application of the algorithm.

2. Methodology: The STEM Procedure and FDR Control

The procedure that we shall exploit in this paper can be viewed as an extension to the sphere of the STEM algorithm, which was introduced in a 1-dimensional Euclidean framework by Schwartzman et al. (2011) and further extended to 2 and 3 dimensions by Cheng et al. (2017); for the sphere, the mathematical background for our proposal and some theoretical results which we shall present below have been discussed in Cheng et al. (2016).

In short, the algorithm can be summarized in the following four steps (STEM stands for Smoothing and TEsting Multiple hypothesis):

- In the first step, the map is smoothed to enhance the signal-to-noise ratio of possible sources, and (equivalently) to get rid of as much Cosmic Variance as possible. The proper implementation of this smoothing step is one of the most delicate parts of our algorithm, and is achieved by means of the (Mexican)

needlet transform, which we shall describe extensively in Subsection 2.1

- In the second step, candidate point sources are selected by numerically computing local maxima of the filtered maps. The algorithm to detect the maxima is described in Subsection 2.2 and has been extensively validated.
- The third step requires the computation of p -values for each of the computed local maxima. Needless to say, even for a purely Gaussian map the distribution of local maxima is not Gaussian, and can be derived analytically along the lines of Bardeen et al. (1985); Cheng and Schwartzman (2015) and Cheng et al. (2016) (see also Cammarota et al. (2016) for related results in the case of single multipole fields/spherical harmonic components). More remarkably, it can be shown that in the high-frequency limit the sample distribution on filtered maps converges to the theoretical expectation, thus making a principled statistical analysis doable. These results are discussed below in Subsection 2.3.
- In the fourth and final step, the multiple testing procedure is implemented. Here we are resorting to the control of the False Discovery Rate, an approach which has become classical in the statistical community over the last decade or so (see Benjamini and Hochberg, 1995, for the pioneering contribution). Heuristically, the idea is to control the proportion of detected point sources that can turn out to be false, as opposed to the control of each of them individually. The advantage of this joint approach to testing have been now very widely recognized in the statistical/mathematical literature, but their impact in Cosmology and Astrophysics has so far been rather limited. The details of our methods, in particular the so-called Benjamini-Hochberg procedure, are discussed below in Subsection 2.4.

In the subsections to follow, we describe in greater details each of the four steps in our algorithm.

2.1. (Mexican) Needlets Filtering

The first step in the procedure is the proper filtering of CMB maps in order to enhance the signal-to-noise ratio. Heuristically, point sources are clearly confined to small scales / high frequencies, hence all the features related to the smallest values of the multipoles ℓ should be considered as “noise” and hence discarded. We are hence looking for a high-pass filter with optimal characteristics.

Needlets are a form of spherical wavelets which were introduced in Cosmology roughly one decade ago and have hence been shown to enjoy a number of very important features. Let us denote by $j, j = 1, 2, \dots$ a set of integer-valued frequencies, and by $P_\ell(\langle x, y \rangle)$ the family of Legendre polynomials, which for $x, y \in S^2$ satisfies the identity




Figure 1: Transversal cut of a needlet function. In the x-axis, it is represented the angular distance between the center of the given needlet (ξ) and an arbitrary point x , in minutes. The amplitude of the needlet is found on the y-axis. The width of the desired needlet can be selected through the parameters j and B . In this work we will fix $B = 1.2$ and use the values of j found in the image: 38, 39, and 40.

$\frac{2\ell+1}{4\pi} P_\ell(\langle x, y \rangle) = \sum_{m=-\ell}^{\ell} Y_{\ell m}(x) \bar{Y}_{\ell m}(y)$, where the bar denotes complex conjugation and $Y_{\ell m}$, as usual, the standard basis of spherical harmonics. The needlet filter is then defined to be

$$\psi_j(x, \xi) := \sum_{\ell} b\left(\frac{\ell}{Bj}\right) \frac{2\ell+1}{4\pi} P_\ell(\langle x, \xi \rangle). \quad (1)$$

An example of this needlet filter can be seen in Fig.1, represented in pixel space, with a fixed value of $B = 1.2$ and j taking the values 38, 39, 40, as we will do throughout this work.

Loosely speaking, the needlet filter is hence nothing more than a weighted average of the usual Legendre polynomial, the latter projecting a spherical map into its distinct multipole components ℓ ; the standard needlet construction was introduced in mathematics by Narcowich et al. (2006), and then in statistics/cosmology by Baldi et al. (2009), Marinucci et al. (2008), Pietrobon et al. (2006). The key ingredient in the construction is then the choice of the weight function $b(\cdot)$, which allows the optimal tradeoff between localization properties in the real domain and those in multipole space/frequency domain. In particular, in the standard needlet construction the function $b(\cdot)$ is supposed to be infinitely differentiable, compactly supported, and satisfying the partition of unity property, which entails $\sum_j b^2(\frac{\ell}{Bj}) = 1$; here B is a user-chosen parameter, whose value will be discussed later (see also Marinucci and Peccati, 2011). These properties ensure, in particular, that each needlet component is finitely supported in multipole space, and hence very-well localized in frequency space. More than that, it has been possible to show (see Narcowich et al., 2006) that the needlet fil-

ter enjoys very good localization properties in real space, the tail of the filter decaying “nearly-exponentially”, i.e., faster than any polynomial, as the frequency j increases; more precisely, one has that, for all integers M there exists a constant C_M such that

$$\psi(x, \xi) \leq \frac{C_M \times B^{2j}}{(1 + B^j \times d_{S^2}(x, \xi))^M}, \quad (2)$$

where $d_{S^2}(x, \xi)$ is the usual geodesic distance on the sphere. The needlet components are then given by

$$\beta_j(\xi) := \int_{S^2} T(x) \psi_j(x, \xi) dx = \sum_{\ell} b\left(\frac{\ell}{Bj}\right) a_{\ell m} Y_{\ell m}(\xi), \quad (3)$$

so that on one hand the filtered field is only supported on the multipoles where the function $b(\cdot)$ takes a non-zero value, while on the other hand the value of the filtered field at any given location ξ is only influenced by the points nearby in the original map $T(x)$. Filtered maps enjoy further very useful (and rather remarkable) properties from the statistical point of view; indeed, under standard assumptions it can be shown that for any two different locations in the sky x and y , the values of $\beta_j(x)$ and $\beta_j(y)$ are asymptotically uncorrelated (and hence independent, under Gaussianity, see Baldi et al., 2009; Marinucci et al., 2008) as the frequency j diverges to infinity. This property will play a crucial role in the determination of the statistical properties of the STEM algorithm.

The filter that is actually going to be implemented is indeed a modification of the original needlet idea, which was introduced soon after by Geller and Mayeli (2009a,b) and then in the statistical/cosmological literature by Lan et al. (2008), Mayeli (2010) and Scodeller et al. (2011). The idea is simply to replace the compactly supported function $b(\cdot)$ by means of the Gaussian-related weight

$$b_p\left(\frac{\ell}{Bj}\right) = \frac{1}{\sqrt{2\pi}} \left(\frac{\ell}{Bj}\right)^{2p} \exp\left(-\frac{1}{2} \left(\frac{\ell}{Bj}\right)^2\right), \quad (4)$$

where p is a parameter, that we have fixed at $p = 1$ in what follows, see Fig. 2.

There are two motivations for the choice of Mexican needlets:

a) the Mexican needlet filter is not compactly supported; however, in the harmonic domain the decay of the filter function has Gaussian tails, hence for all practical purposes the contribution of multipoles which are not concentrated around the centering value is utterly negligible. Moreover, there is a huge theoretical and numerical evidence to show that Mexican needlets optimize the trade-off between concentration in the real and harmonic domain, see the discussion in Geller and Mayeli (2009a,b), Mayeli (2010), Lan et al. (2008) and Scodeller et al. (2011). Heuristically, this is due to the fact that a Gaussian distribution saturates Heisenberg inequality on the concentration properties of a function and its Fourier transform;

b) Mexican needlets have a very explicit form in both the real and harmonic domain, and this makes the investigation of their power properties simpler. In particular,




Figure 2: Representation of the filter function b for a Mexican needlet with $B = 1.2$ and $j = 38, 39, 40$. The value of j determines the multipole region that the filtering will extract.

it is simpler to control the oscillation of their lobes, which is more pronounced in the standard needlet case (with the risk to introduce spurious maxima).

2.2. Selection of Candidate Point Sources

Candidate point sources are selected by simply collecting the local maxima in the filtered maps, i.e., the points where the gradient is zero and the (covariant) Hessian matrix is negative definite. We write $G(\beta_j)$ for the set of detected peaks and $M(\beta_j)$ for their total number, that is to say that a point $x \in S^2$ belongs to $G(\beta_j)$ if and only if

$$\{x \in G(\beta_j)\} \Leftrightarrow \{x : \nabla \beta_j(x) = 0 \text{ and } \nabla^2 \beta_j(x) \prec 0\}, \quad (5)$$

$A \prec 0$ denoting a negative definite matrix.

In practice, candidate peaks are simply identified by the routine *Hotspot* in HEALPix, see <https://healpix.jpl.nasa.gov/html/facilitiesnode8.htm>.

2.3. Distribution of Local Maxima and p-values

The key technical step in this algorithm is based on the possibility to evaluate the exact distribution of local maxima for the CMB maps, under the (null) assumption that no point source is present. Computing the density of local maxima, or equivalently the expected number of maxima for a given Gaussian field, is a topic which has drawn a lot of work in Cosmology, starting from the seminal paper Bardeen et al. (1985) in the eighties. Under the circumstances of the present paper, the density of these maxima can be shown to be given by (see Cheng et al., 2016, and the references therein)

$$f_j(x) = \frac{2\sqrt{3+\eta_j^2}}{2+\eta_j^2\sqrt{3+\eta_j^2}} [A + B + C] \quad (6)$$

where we defined

$$\begin{aligned} A &= [\eta_j^2 + \kappa_j^2(x^2 - 1)] \phi(x) \Phi \left(\frac{\kappa_j x}{\sqrt{2 + \eta_j^2 - \kappa_j^2}} \right) \\ B &= \frac{\kappa_j \sqrt{(2 + \eta_j^2 - \kappa_j^2)}}{2\pi} x e^{-\frac{(2 + \eta_j^2)x^2}{2(2 + \eta_j^2 - \kappa_j^2)}} \\ C &= \frac{\sqrt{2}}{\sqrt{\pi(3 + \eta_j^2 - \kappa_j^2)}} e^{-\frac{(3 + \eta_j^2)x^2}{2(3 + \eta_j^2 - \kappa_j^2)}} \\ &\times \Phi \left(\frac{\kappa_j x}{\sqrt{(2 + \eta_j^2 - \kappa_j^2)(3 + \eta_j^2 - \kappa_j^2)}} \right). \end{aligned}$$

The notation here requires some clarification. Indeed, ϕ, Φ denote, respectively, the standard density and cumulative distribution function for a Gaussian random variable; on the other hand, κ_j and η_j are constants which can be explicitly computed from the angular power spectrum of the original (unfiltered) CMB map; more precisely, they are given by

$$\eta_j = \frac{\sqrt{\Gamma'_j(C_\ell)}}{\sqrt{\Gamma''_j(C_\ell)}}, \quad \kappa_j = \frac{\Gamma'_j(C_\ell)}{\sqrt{\Gamma''_j(C_\ell)}}, \quad (7)$$

where

$$\Gamma'_j(C_\ell) := \sum_{\ell} \frac{2\ell+1}{4\pi} b_p\left(\frac{\ell}{B_j}\right) P'_\ell(1), \quad (8)$$

$$\Gamma''_j(C_\ell) := \sum_{\ell} \frac{2\ell+1}{4\pi} b_p\left(\frac{\ell}{B_j}\right) P''_\ell(1), \quad (9)$$

and the derivatives of Legendre polynomials evaluated at 1 are given by

$$P'_\ell(1) = \frac{\ell(\ell+1)}{2} \quad \text{and} \quad P''_\ell(1) = \frac{\ell(\ell-1)(\ell+1)(\ell+2)}{8}. \quad (10)$$

In the sequel, it should be kept in mind that all our computations will be carried over components which have been normalized to have unit variance.

Of course, once the density of local maxima is known, it is immediate to compute the p -value of any one of them taking value u (say), which is indeed given by

$$\pi_j(u) := \int_u^\infty f_j(t) dt \quad (11)$$

A further, more remarkable result was also established in Cheng et al. (2016); it was indeed shown that at high frequencies, the realized (observed) distributions of critical points on CMB maps converges to the theoretical density f_j given above. In other words, even on a single realization of CMB maps the observed distribution of local maxima for large values of j is going to track closely the theoretical prediction, under the assumptions of Gaussianity and

isotropy. This result is grounded on the capacity of needlet filtered maps to control Cosmic Variance, and it is clearly the foundation for the statistical multiple testing procedure which we shall describe in the next subsection; see Fig.5 later for numerical evidence supporting these claims.

2.4. The Multiple Testing Procedure

As a multiple testing procedure in step (4) we apply the Benjamini-Hochberg (BH) approach Benjamini and Hochberg (1995). This algorithm is now very popular among statisticians; in this paper it is implemented directly as follows. Recall that we write $G(\beta_j)$ for the set of detected peaks and $M(\beta_j)$ for their total number; let us arrange their p -values in increasing order and fix a significance level $\alpha \in (0, 1)$ whose role will be made clearer below.

Now let k be the largest index for which the i th smallest p -value is less than $i\alpha/M(\beta_j)$; then the null hypothesis that there is no point source at a given local maximum location $x \in G(\beta_j)$ is rejected if

$$p_j(x) < \frac{k\alpha}{M(\beta_j)}, \quad (12)$$

where we assume without loss of generality that at least a peak is detected in the map (otherwise the test is clearly unnecessary). In words, the algorithm can be explained as follows: we draw a line starting from the origin with slope α , and on the same plot we represent in ascending order the p -values corresponding to the detected peaks: local maxima are considered significant (and hence identified with point sources) if and only if they are small enough to fall below the line. The procedure is illustrated in Fig. 3.

This procedure has a multiple testing nature in an obvious sense: for any local maxima, not only its value is considered, but also its ranking in the full map (a single source at 3σ , say, may not be significant in a map with million pixels, but one thousand such maxima have definitely another meaning). More rigorously, this algorithm guarantees control of the so called *False Discovery Rate*, which is defined as the expected number of false discoveries out of the critical points which are identified as point sources (see Cheng et al. (2016) for more discussion and details):

$$FDR_j = \langle \frac{W_j}{S_j} \rangle, \quad (13)$$

where S_j, W_j denote, respectively, the number of local maxima which are identified as sources and those for which this identification is actually wrong, and $\langle \cdot \rangle$ is as usual the ensemble expected value. Under some standard assumptions, it is indeed possible to show that the STEM algorithm that we introduced on one hand allows to control the False Discovery Rate by the user chosen parameter α , in the high-frequency limit; i.e.,

$$\lim_{j \rightarrow \infty} FDR_j \leq \alpha; \quad (14)$$




Figure 3: Example of the Benjamini-Hochberg procedure for multiple testing. In the x-axis, we have the rank i of the p -values, sorted from lower (less likely) to higher (more likely). The value of the p -values times the total number of maxima is plotted for each maximum as blue points. This is an estimate of the number of maxima expected to be at least as intense as a given point in a purely Gaussian CMB map. The three colored lines represent the threshold for 3 different values of α ; the highest value still below the line determines the number of points to be classified as candidate point sources. In dashed black line, corresponding to $\alpha = 1$, we have the expected behaviour for a purely Gaussian map.

moreover, in a suitable sense the procedure has statistical power growing to unity at the largest scales, i.e., it is able to recover a proportion growing to 100% of the existing sources. These results are clearly of a theoretical nature, but as we shall show in the Sections to follow they do provide a very good guidance on the actual performance of the STEM algorithm under realistic experimental conditions and on Planck 2018 data.

3. Numerical Implementation and Simulations

The algorithm described in the previous section has been implemented on Planck-like simulations including point sources, as described below.

3.1. Implementation

We implement the algorithm by creating a software on Python 3.5, exploiting in particular `numpy` (Oliphant, 2006) and `scipy` (Jones et al., 2001). We use `astropy` (The Astropy Collaboration, 2018) to import FITS images as the ones provided by the Planck Collaboration, and `HEALPix` (Gorski et al., 2005) to deal with full-sky spherical images, as this is the standard adopted by the Planck Collaboration. The Python package concerning the needlet treatment and the multiple testing approach has been programmed entirely by us, following the algorithm explained in the previous section. Finally, we use `pandas` (McKinney, 2010) and `Matplotlib` (Hunter, 2007) to analyze the results.

The software is structured into three parts: i) The main algorithm, which takes a CMB map and extract a list of points believed to be point sources, according to the STEM Procedure explained in Section 2; ii) The simulations, which create a number of possible CMB realizations and introduces artificial point sources according to specifications, in order to test the validity of the procedure; and iii) The Planck maps analysis, which helps to perform the study of the maps provided by the Planck Collaboration.

Here we explain the main steps of the implementation for each of these parts.

3.1.1. Implementation of the Main Algorithm

We follow the four steps described in Section 2: filtering the map, selecting candidates, computing their p -values, and applying the multiple testing procedure.

We start by filtering the input CMB map; this is done in harmonic space, using Eq. 3. First, we extract the spherical harmonics coefficients $\alpha_{\ell m}$ using the `map2alm` routine on `HEALPix`. Then, these quantities are multiplied by the filter $b(\frac{\ell}{Bj})$ of the corresponding Mexican needlelet with fixed parameters B and j : of course, this filtering factor depends on ℓ but not m (needlelets are isotropic); this procedure is significantly more efficient than filtering in pixel space. The result is then used to reconstruct the now-filtered map $\beta_j(\xi)$ with the `alm2map` routine on `HEALPix`. Going through the spherical harmonic coefficients may be suspected to affect the properties of needlelet-filtered random fields, in the presence of masked regions; however it has been shown both numerically and mathematically that this effect is negligible at high frequencies j , as those on which we focus for point source detection, see i.e., Scodeller et al. (2011).

An example of a CMB map before and after the needlelet filtering is in Fig. 4. The angular size of this image is $5^\circ \times 5^\circ$; of course, the small scale clumps in the temperature map are enhanced in the β -map, especially at the scale corresponding to the selected needlelet (around $\sim 10'$).

Once the filtered map $\beta_j(\xi)$ is computed, we extract its local maxima with the `hotspot` routine in `HEALPix`. It checks the intensity of the neighboring pixels for every pixel in the map and delivers both the location and the intensity of the maxima. At this point, the maxima distribution can be compared with the theoretical distribution $f_j(x)$.

At this stage, the p -value is computed for each maximum, following Eq. 11: for each maximum of intensity x , we integrate $f_j(x)$ numerically from x to infinity, obtaining the p -value $p_j(x)$. However, it has to be noted that every map has tens of thousands of maxima, with similar values for their intensities. Because of this, performing the integration for every point is highly inefficient. Instead, we take the highest and lowest values for the intensity and integrate in this interval with a small step of $\Delta x = 0.05$; the error in this procedure is lower than the error in the numerical integration, which is $\frac{\Delta p}{p} < 10^{-4}$. Nonetheless,




Figure 4: Comparison between a CMB temperature map (on the left) and the β -map resulting of filtering this map with a needlelet $B = 1.2$, $j = 39$ (on the right). The field of view for both images corresponds to $5^\circ \times 5^\circ$.

we make the exact computation only for the 1000 most intense maxima.

The last step is to apply the Benjamini-Hochberg algorithm to identify the possible point sources among the whole maxima population. In order to do that, we proceed as described in Section 2.4: first we sort the p -values from lowest (less likely) to highest (more likely); and then we extract the highest p -value that satisfies Eq. 12. This routine can be implemented for any value of the confidence parameter $\alpha \in (0, 1)$, see Fig. 3.

All the steps described in this section have been implemented in a Python package named `MTNeedlet`¹. This package consists of a series of flexible functions that can be used individually or together. The package is fully documented and available through the standard procedure “`pip install mtneedlet`”.

3.1.2. Implementation of the simulations

We are now in the position to apply the algorithm to simulated maps. We generate a number of realizations of the CMB and a set of artificial point sources, according to the input parameters. After applying the STEM algorithm, this routine calculates the true and false detections obtained by the algorithm. The complete Python code described in this section can be found online² as a Jupyter Notebook.

More precisely, we generate an angular power spectrum C_ℓ from the Planck cosmological parameters, using `CAMB` (see Lewis, 2011). We use this angular power spectrum to generate a series of compatible CMB maps; we choose to use the same resolution of the Planck maps, $\text{nside} = 2048$. According to the `HEALPix` standard, the total number of pixels is $12 \cdot \text{nside}^2$. The maps are then convolved with a Gaussian profile of $fwhm = 5'$, similar to the nominal resolution of Planck data.

¹<https://javiccarron.github.io/mtneedlet/>

²<https://javiccarron.com/code/>

A population of artificial point sources is then introduced. In order to do that, a blank map of higher resolution is created ($n_{\text{side}} = 4096$) and a set of pixels is selected at random; the intensities of these points are selected following a uniform distribution between the chosen limits. These artificial point sources are now convolved with a Gaussian profile equal to the one used for the map before. This map of artificial point sources is then added to the different CMB realizations.

The main algorithm is run on each map (CMB + point sources) to obtain the list of points selected as point sources. We will claim that a detection is successful if it is less than $\rho = 3$ pixels away from the artificial source that had been injected into the map, as in Cheng et al. (2016). We find that increasing this tolerance to any reasonable extent does not have a significant impact on the results, while of course excluding this tolerance factor decreases the number of detections.

The values chosen for the parameters and the results of these simulations will be explained in Section 3.2.

3.1.3. Implementation for the Planck data

The last part of the numerical implementation concerns a pipeline that loads the target maps (in this case Planck CMB observations), applies the algorithm, and gather the most important results (which we present in the tables below). The complete Python code described in this section can be found online³ as a Jupyter Notebook.

First, we extract some information concerning each map: number of pixels, map-making algorithm (and its version), and whether it is the inpainted version provided by Planck or not. Then, the user can choose a set of values for the parameters of the main algorithm: in particular, the needlet parameters B and j , and the confidence level α .

After the main algorithm is implemented on each map, the procedure checks the location of the point sources selected by the algorithm. In order to do that, it uses the two confidence masks provided by the Planck Collaboration. The first one is the common confidence mask, obtained through the combination of the masks for the individual algorithms; it covers around $\sim 22\%$ of the sky. The second is the inpainting mask, covering the areas with apparent contamination that are to be inpainted to obtain realistic maps; it covers around $\sim 2\%$ of the sky (see Section 4.2 in Planck Collaboration IV (2018) for more information about the masks). In our implementation, the number of point sources detected inside and outside each mask is reported; the Planck Catalogue of point sources is also explored to count how many of the detections in each region are matched by known sources.

We then produce a `pandas` table with the results; for every map we can find the information about the parameters of the algorithm and the number of sources reported




Figure 5: Maxima distribution in filtered CMB maps. On the top, the comparison between the observed distribution (mean of 200 realizations), in continuous black line; and the theoretical prediction from Eq. 6, in dashed red line. On the bottom, the residual: $\text{measured} - \text{predicted}$. Both cases include the regions where the 68% and the 90% of the maps lay.

(total number, number inside each mask and number in the catalogues).

3.2. Numerical Validation

Let us now describe the validating simulations. In the case of Mexican needlets, the only free parameter to be chosen is B^j , see Eq. 4; in particular, we select $B = 1.2$ at frequencies $j = 38, 39, 40$, meaning multipole regions around $B^j \approx 1020, 1225, 1470$. The reasons of this choice will be explained in Section 3.2.2. In any case, selecting much lower multipoles (larger scales) makes the point source detection less efficient, while selecting much higher multipoles (smaller scales) makes the algorithm more sensitive to noise and pixelization effects.

We start by generating 200 CMB realizations without artificial point sources but with realistic noise (estimated from the semidifference of odd and even observational rings in Planck data). In these simulations, we observe that the maxima population follows the theoretical distribution from Eq. 6, as it can be seen in Fig. 5. We note that the residual is not exactly 0, but it presents a characteristic pattern, possibly due to pixelization effects. However the case may be, the maximum of the residual is less than 0.01; this is aligned with the result by Cheng et al. (2016), which showed that at high frequency, the maxima distribution of CMB maps converge to the theoretical prediction.

We use again the sample of 200 CMB realizations in order to test the algorithm on maps without artificial sources.

³<https://javiercarron.com/code/>

α	$j = 38$	$j = 39$	$j = 40$
0.05	4.0 %	6.5 %	99.5 %
0.01	1.5 %	2.5 %	98.5 %
0.002	0.5 %	0.5 %	93.5 %
0.05	0.045	0.075	4.595
0.01	0.015	0.030	3.215
0.002	0.005	0.005	2.270

Table 1: Percentage of maps with incorrect detections and average number of detections per map. Both are calculated for a total run of 200 CMB maps without artificial sources but with noise. As expected, very few candidates are reported for $j \leq 39$, but the noise becomes significant at larger multipoles.

It is applied with different values of the confidence parameter $\alpha = 0.05, 0.01, 0.002$. Since no artificial source is added, we expect the amount of reported point sources to be due to noise.

In Table 1 we can see the percentage of maps that report at least one (incorrect) detection and the average number of detections per map, for different combinations of parameters. As expected, the number of detections that are reported will increase as α increases. Additionally, we observe that the number of incorrect detections is very low for $j \leq 39$, almost unaffected by noise. For the time being, we will do the analysis with the parameters $j = 39$, $\alpha = 0.01$, although a further discussion of this choice can be found in Section 3.2.2.

3.2.1. Sensitivity of the algorithm

In order to test the detection power of the algorithm, we introduce a total of 200 artificial point sources in each of the simulated maps; they are produced with a peak intensity between 0 and 7σ , where σ is the standard deviation of the convolved maps before filtering (calculated from the theoretical angular power spectrum). We are interested in knowing the minimum intensity for which the algorithm is able to recover most point sources: this is the lower limit of intensity for a complete detection. Of course, we are also interested in the number and intensities of false detections.

The results of the simulations for $\alpha = 0.01$, $B = 1.2$, and $j = 39$ can be seen in Fig. 6, which provides the number of artificial sources, the correct detections, and the false detections. The intensity of the maxima is measured with respect to the CMB temperature map before the needlet filtering, in units of its standard deviation σ . It can be seen that the algorithm detects essentially every artificial source over 4σ ; indeed, a 90% completeness level is reached at 3.66σ . Additionally, the number of false detections is very small and almost negligible over 2σ . The average number of false detections in these simulations is 1.05 per map, meaning a False Discovery Rate of 0.009.

These results are reported with the intensity referring to the CMB temperature map, before the filtering is applied. Since the detection occurs on the filtered maps β , it is interesting to see the intensity of the point sources after filtering; this is given in Fig. 7. The first difference




Figure 6: Sensitivity of the algorithm. In the x-axis, we have the intensity of the point sources, referred to the CMB temperature maps before filtering, in units of its standard deviation. In the y-axis, we have the number of detection of the maxima in each bin. In green, we have the artificial sources; 10 sources have been introduced in each bin, between 0 and 7σ . In blue, we have the true detections: number of artificial sources correctly detected. In orange, we have the false detections: reported candidates that do not correspond to any artificial source. For both true and false detections, we plot the average and the confidence regions which include 70 and 90% of the maps. The intensity for the artificial sources is the one with which they were generated, the one for the false detections is measured on the map. The algorithm has been applied 200 times, with 200 sources at a time, a needlet filter of $B = 1.2$ and $j = 39$, and $\alpha = 0.01$.

is a clear enhancement of the signal: while point sources are generated with intensities between 0 and $7\sigma_{CMB}$, the intensity of these points in the β -maps is up to $15\sigma_\beta$. The signal to noise is increased up to a factor $\lesssim 2$ for most point sources.

Another difference here is the noisy profile of the detections, resulting in wider confidence regions. This effect is expected, it is a consequence of measuring the intensity on filtered maps after adding the CMB realization: the same source can be located on a minimum of the CMB signal in one map, but on a maximum in another.

3.2.2. Choice of parameters

The implementation of the algorithm requires to select the values of two parameters. The first one is the characteristic scale of the needlet, i.e., the multipole region to be filtered. The second one is the control level for FDR (α) in the Benjamini-Hochberg procedure. The choice of these parameters entails a trade-off between the power of detection and the ratio of false discoveries. For our results below, we shall set the limit $FDR \leq 0.01$: in words, we require that at most 1% of the reported point sources can be false detections in the simulations. Within this limit, we will choose the parameters which maximize the detection power.

A priori, the characteristic scale of the needlet used to filter the map should be comparable to the typical size




Figure 7: Sensitivity of the algorithm on the filtered CMB maps (β -maps). In the x-axis, we have the intensity of the point sources, measured directly on the filtered CMB maps. In the y-axis, we have the number of point sources. Detections that correspond to artificial point sources are plotted in blue, while detections that do not, are represented in orange. Artificial point sources that are not detected are represented in red. The average and the confidence regions are plotted, where 70 and 90% of the maps lay. As before, the algorithm has been applied 200 times, with 200 sources at a time, a needlet filter of $B = 1.2$ and $j = 39$, and $\alpha = 0.01$.

α	$j = 38$	$j = 39$	$j = 40$
0.05	0.042	0.040	0.082
0.01	0.009	0.009	0.046
0.002	0.0013	0.0018	0.0325

Table 2: False Discovery Rate for a simulation of 200 CMB maps with 200 added artificial sources between 0 and $7\sigma_{CMB}$ and realistic noise.

of a point source. Larger scales are expected to be less efficient, as the signals of the point sources are diluted into a larger area; and smaller scales are expected to be more sensitive to instrumental noise and pixelization effects. We fix $B = 1.2$ and select $j = 38, 39, 40$ to probe these hypotheses and select the optimal size (peaking at $\ell \sim 1020, 1225, 1470$, respectively). As before, we simulate 200 CMB realizations with 200 artificial sources between 0 and 7σ and realistic noise. The algorithm is then applied to all the maps with the different values for j and α .

Let us first discuss the effect of the parameters on the False Discovery Rate (FDR). We report the values of the FDR for our simulation in Table 2, and the average number of false discoveries per map in Table 3. We observe that $j = 38$ and $j = 39$ yield similar realized values for FDR ; the actual values are lower than (but close to) α in both cases, as predicted by Eq. 14 for high j . However, this is not the case for $j = 40$, where the instrumental noise produce a significant distortion. Indeed, this disparity is not observed if noise is not added to the maps. The required limit $FDR \leq 0.01$ leads hence to $j \leq 39$.

α	$j = 38$	$j = 39$	$j = 40$
0.05	3.65	5.01	12.67
0.01	0.66	1.05	6.74
0.002	0.09	0.20	4.53

Table 3: Average number of incorrect detections per map for a simulation of 200 CMB maps with 200 added artificial sources and noise.

α	$j = 38$	$j = 39$	$j = 40$
0.05	5.08	3.50	2.50
0.01	5.36	3.66	2.66
0.002	5.74	3.86	2.74

Table 4: Level of intensity for which an artificial point source is recovered at least 90% of the times. The intensities are measured on the peak of the signal before the needlet filtering, in units of the standard deviation of the CMB. These results have been obtained from a run of 200 CMB maps with 200 added artificial sources between 0 and $7\sigma_{CMB}$ and realistic noise.

Let us now investigate detection power: in order to achieve this goal, the sensitivity curve is derived for each case, as in Fig. 6. We are interested in the intensity of the points that the algorithm is able to detect with frequency 90% or larger, this is the 90% completeness limit. The values of this limit for different parameters are shown in Table 4. The variation between different values of j is significant: $\sim 1.5\sigma$ between 38 and 39; and $\sim 1\sigma$ between 39 and 40. As expected, a higher value of j yields a stronger detection power.

To sum up, under the constraint that $FDR \leq 0.01$, detection power is maximized by $j = 39$; this is therefore our choice for the investigations to follow.

Lastly, we can also compare the results of the algorithm in the case where no artificial source is introduced, i.e., we observe only CMB and noise; the results were shown before in Table 1. We see that the algorithm with this choice of parameters performs satisfactorily: it only reports a detection in 1 in 20 maps and, even in those cases no more than 2 candidates are reported. On the other hand, as before the case of $j = 40$ is significantly contaminated by noise and the algorithm reports several detections in the majority of the maps.

3.2.3. More comments on noise and masking

One of the advantages of needlet filtering is that it focuses only on the contribution of a band-limited multipole region; this is also true in practice for Mexican needlets, as the tail of the filter has a Gaussian decay. The noise contribution is more important at high ℓ , so we can minimize its effect by choosing lower values for ℓ , while these multipoles can be chosen large enough to be still suitable to detect point sources. In Fig. 8, we can see the estimated contribution of the noise and the signal to the angular power spectrum C_ℓ , together with the shape of the needlet filter $b(\frac{\ell}{Bj})$ for the values considered in this work $B = 1.2$, $j = 38, 39, 40$ (peaking at $\ell \sim 1020, 1225, 1470$,




Figure 8: Estimated angular power spectrum C_ℓ of the signal and the noise. Overlapped in dashed line, we can see the filter $b\left(\frac{\ell}{B^j}\right)$ for needlets with $B = 1.2$ and $j = 38, 39, 40$. As expected, low multipoles are dominated by signal, while high multipoles are dominated by noise. The needlet scales are chosen in order to correspond to the size of point sources, while staying in the signal-dominated region.

respectively). It can be seen that the needlet filtering will mostly reflect the signal-dominated scales, while the noise starts to become significant (but not dominant) for $j = 40$, as shown in the previous section.

To understand better the role played by noise, we repeated the simulations in the previous subsection without adding a realistic noise component to the maps. In the case of $j = 38$ and $j = 39$, all results are very similar to those reported in (Tables 1 to 4), within statistical error; this seems to suggest that the noise has a minimal impact in the algorithm for this multipole region. On the other hand, the addition of a noise component drastically affects the results for $j = 40$: without noise, the False Discovery Rate is similar to what we obtained for the other values of j , that is $FDR \lesssim \alpha$. For $\alpha = 0.01$, $j = 40$ the False Discovery Rate increases by a factor 5 when noise is introduced.

It is also worth mentioning the effect of masking, i.e., setting to 0 the pixels of areas considered to be contaminated by the Galaxy or known point sources. We recall that one of the main advantages of using needlets is their extremely good localization properties; therefore, we may expect that masking would remove the areas with known contamination, while not affecting the algorithm on the rest on the map. There is, however, some technical difficulties that may arise: since masking a region makes its value equal to 0, some masked sources in cold (negative) regions can still appear as very intense point sources; additionally, in some cases masking generates artificial detections near the border of the masked area. We avoid these technical problems by using the masks only to check the location of the detections (i.e., inside or outside the confidence masks); we do not mask the map before applying

the algorithm.

4. Results on Planck 2018 Maps

In this section we apply the algorithm to the temperature maps produced by the four algorithms of the last Planck data release: COMMANDER, SMICA, SEVEM, and NILC. All the maps are processed at their native resolution of $n_{\text{side}} = 2048$. We analyze the last two data releases: 2015 and 2018 (see Planck Collaboration IX, 2016; Planck Collaboration IV, 2018); for the latter, we apply the algorithm both in the inpainted and not inpainted cases. These maps are expected to be very close outside the confidence masks, but the results will be different since the algorithm considers the maxima population as a whole in the entire map. Using inpainting is a way to exclude the bright and known sources from this analysis.

The algorithm is applied with the parameters discussed in Section 3.2.2: $B = 1.2$, $j = 39$, and confidence parameter $\alpha = 0.01$. Several results are stored for each map: first, the total number of detections reported (their location and intensity are also stored but will not be shown here for brevity's sake). We then calculate the amount of the detections outside the inpainting mask (98% of the sky) and outside the common confidence mask (78% of the sky). Lastly, we check how many detections match the Planck catalogues of point sources. These catalogues are reported for each frequency plus additional effects (such as Sunyaev-Zel'dovich); the procedure checks all the catalogues individually and all together.

The file containing the complete table with all the results can be found in <http://javiercarron.com/code>, along with the code used to generate it. This table includes the results for other values of the parameters j and α . It consists of 144 rows and 18 columns, including the location and intensity of every point reported by the algorithm and, therefore, is too large to be reproduced in this article. In Table 5 we can see a part of these results for $j = 39$ and $\alpha = 0.01$, on which we are going to focus.

We note that the algorithm reports a significant number of detections for most of the maps. Some of them are sources already present in the Planck catalogues that, apparently, could not be completely removed for the final maps. However, a fraction of them are sources that are not intense enough to be present in these catalogues. We note that, as expected, the algorithm does not detect point sources in the high- z catalogue or in the Sunyaev-Zel'dovich catalogue.

Before analyzing the results for the different algorithms, we are going to focus on the similarities and the general behavior of the procedure. The effect of the frequency filtered (j) and the confidence parameter α is similar for all maps. We can see an example of this effect for the not-inpainted 2018 SMICA map in Table 6.

A lower value of $j = 38$ ($\ell \sim 1020$) means filtering the map with wider needlets; this means that the signal will

Map	version	inpainted	detections	mask C	mask I	cat	cat C
commander	2.01	False	55	0	2	53	0
commander	3.00	False	6334	37	716	2927	11
commander	3.00	True	70	1	12	8	0
nilc	2.01	False	1222	6	10	753	1
nilc	3.00	False	1240	8	15	755	1
nilc	3.00	True	0	0	0	0	0
sevem	2.01	False	7063	39	690	4056	13
sevem	3.00	False	5894	31	261	3424	10
sevem	3.00	True	3	0	1	0	0
smica	2.01	False	558	2	10	69	0
smica	3.00	False	318	2	16	197	0
smica	3.00	True	11	1	3	7	0

Table 5: Part of the results table. Here, the three versions for the four algorithms have been filtered with a needlet $B = 1.2$ and $j = 39$. Also, α is fixed to 0.01. The columns are as follows. Map: algorithm used to extract the cleaned temperature map. Version: data release used, 2.01 corresponding to 2015 and 3.00 to 2018. Detections: number of candidates reported. Mask C: number of candidates *outside* the common confidence mask (78% of the sky). Mask I: number of candidates *outside* the inpainting mask (98% of the sky). Cat: number of candidates present in a catalogue. Cat C: number of candidates *outside* the common confidence mask that are present in a catalogue.

Map	version	inpainted	j	alpha	detections	mask C	mask I
smica	3.00	False	38	0.01	169	1	6
smica	3.00	False	39	0.01	318	2	16
smica	3.00	False	40	0.01	744	5	79
smica	3.00	False	39	0.002	268	1	8
smica	3.00	False	39	0.01	318	2	16
smica	3.00	False	39	0.05	403	11	49

Table 6: Example of the behaviour of the results with j and α for the not-inpainted SMICA maps from the last data release. On the top, varying j with fixed α . On the bottom, varying α with fixed j .

be more diluted, making it more difficult to detect sources but more robust against noise. On the other hand, a higher value of $j = 40$ ($\ell \sim 1470$) means filtering with narrower needlets, which are similar in size to the point sources; this makes this value more sensitive to point sources but also much more vulnerable to noise. This explains the observed result: the number of detections grows with j .

Similarly, the number of detections also grows with the confidence parameter α . As explained before, a higher value of α means that the confidence required to report a detection is lower; therefore, more detections will be reported. With these values for α we are able to detect a significant number of point sources introducing a low count of false detection, as explored with simulations in Section 3.2.

As explained in Section 3.2.2, the choice of parameters $j = 39$, $\alpha = 0.01$ is the one that maximizes the power of detection while maintaining $FDR \leq 0.01$

4.1. Comparison between the different algorithms

We are now going to analyze the differences between the results for the different algorithms used to extract the CMB temperature map; we are also going to discuss the version, i.e., data release of the map and whether it is inpainted or not. As discussed before, all values refers to the results for the parameters $j = 39$ and $\alpha = 0.01$.

SEVEM maps present the highest amount of point sources for the 2015 version (7064). This number is reduced for the 2018 version of the map, with 5894 detections. In these two cases the Galactic component is strong and produces a large number of detections in the Galactic plane; indeed, only a small number of them (39 and 31, respectively) are detected outside the confidence mask. As mentioned before, we also analyze the inpainted map in order to limit the influence of the known contamination, especially from Galactic origin: in this case, the overwhelming majority of the detections are removed, only 3 remaining. None of these detections is outside of the confidence mask: this means that the maxima population of the SEVEM inpainted map is compatible with the purely Gaussian case outside the confidence mask.

NILC maps present a lower amount of detections overall: 1222 and 1240 for the 2015 and 2018 versions, respectively: again, only a very small part (6 and 8) are outside the confidence mask. More interesting here is the case of the inpainted map: this is the only case where the algorithm does not report any detection, not even inside the masks. Once the strongest point sources are removed through inpainting, the maxima population of the NILC map is perfectly compatible with the Gaussian case.

SMICA maps present the lowest number of detections: 558 and 318 for the 2015 and 2018 versions, respectively. Of these detections, only 2 are outside the confidence mask

in both cases. The inpainted SMICA map present 11 detections, a higher quantity than the previous algorithms; of these, only one of them is located outside the confidence mask. This source can be seen in Fig. 9.

COMMANDER maps present the highest number of detections in the 2018 version (6335), although it was the lowest for the 2015 version (55). This is probably due to the fact that, in the last version, this is the only algorithm that does not preprocess the frequency maps to remove possible point sources. This produces a more robust procedure in exchange of a higher level of contamination. Indeed, if we look at the inpainted map, where most of these points are removed a posteriori, the number of detections is reduced to 70, only 1 of them located outside the confidence mask. This point is the same that was reported in the SMICA map, which can be seen in Fig. 9.

In general, we observe that the inpainting procedures have worked extremely efficiently, and all four algorithms seem basically without or with very few spurious maxima (i.e., undetected point sources) after inpainting has been implemented.

4.2. Reported point sources

We will close this section by studying the point sources detected by the algorithm. It is interesting to know the intensities of the detections: on the filtered (β) maps, they have a signal of at least $5.4\sigma_\beta$. The signal of these points before filtering, directly on the CMB maps, can be seen in Fig. 10; there is not a clear cut on the intensity, since the filtering is able to exploit their shape in order to boost the detecting power.

As we mentioned before, there is a point source outside of the confidence mask that has been reported by the algorithm both in the SMICA and COMMANDER maps (see Fig. 9; it also corresponds to the green point in Fig. 10). It can be observed that its shape and intensity are similar to the ones expected from an astrophysical point source. This point is located at galactic longitude $l = 103^\circ 18'$ and galactic latitude $b = -27^\circ 17'$ and, as the majority of the detections, we have to note that this point does not correspond to a source present in any of the Planck Catalogues of point sources. Identifying this possible source is not straightforward, but looking at this region in an infrared survey such as 2MASS (see Skrutskie et al., 2006), we find a bright triplet of point sources within the pixel where the maximum is detected, as it can be seen in 11.

5. Conclusions

We have implemented the STEM algorithm to extract candidates point sources from a map, controlling the False Discovery Rate. In order to do that, we have written a code from scratch in Python 3, trying to make our pipeline as flexible as possible, accepting a wide variety of parameters as input. The code could be used to analyze maps even outside the CMB framework. We have extensively tested

the code on CMB simulations, generated using HEALPix and CAMB; we have then the ability to control the proportion of False Detections and the power of the algorithm: we recover most of the sources at intensities higher than 3 to 4 times the standard deviation of the map. Using simulations again, we have tested the effect of noise and masks. We have concluded that the effect of noise is largely negligible, while masks can introduce some false maxima near the boundary.

We have run the algorithm on a set of foreground cleaned Planck CMB maps, including second and final data release (2015 and 2018, respectively, see i.e., Planck Collaboration IX 2016; Planck Collaboration IV 2018). For the final data release, we have included both inpainted and not inpainted maps (in all cases we focused on temperature anisotropy maps). We have observed that the inpainting procedure adopted for the 2018 release seems to have produced maps which are much closer to being purely Gaussian, with a number of local maxima largely consistent with theoretical predictions. This was not necessarily the case with some of the earlier releases.

In this paper, we have only focused on CMB temperature maps. However, these techniques can be readily applied to any kind of spherical map that we suspect to be contaminated by point sources. In particular, we plan to apply this algorithm to polarization maps, both for the E and B modes. Likewise, this algorithm could be eventually applied to frequency maps, before they are combined to obtain the temperature (or polarization) maps. In this case, we have the additional problem of diffuse Galactic radiation contaminating the background, plus all the population of physical point sources. Our aim is also to apply the algorithm to different frequency bands and then combine the results to improve sensitivity; indeed, in this paper we did not make any use of the spectral information of the CMB observations.

We made the code publicly available for the whole scientific community and we plan to expand it for future work.

Acknowledgements

We acknowledge financial support by ASI grant 2016-24-H.0; the research by DM was also supported by the MIUR Excellence Department Project awarded to the Department of Mathematics, University of Rome Tor Vergata, CUP E83C18000100006. Data analysis was based on observations obtained with Planck (<http://www.esa.int/Planck>), an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA, and Canada. Some of the results in this paper have been derived using the HEALPix (Gorski et al., 2005) package.

References

- Adler, R.J., Taylor, J.E., 2009. Random Fields and Geometry. Springer Science & Business Media.




Figure 9: Point source detected by the algorithm in the SMICA map. This point is located at galactic longitude $l = 103^\circ 18'$ and galactic latitude $b = -27^\circ 17'$. The signal to noise ratio at the center is 3.01 in the CMB map and 5.84 in the filtered map. It has also been reported in the COMMANDER map, where it has a signal to noise ratio at the center of 2.94 in the CMB map and 5.69 in the filtered map. The image corresponds to a $3^\circ \times 3^\circ$ patch of the sky.




Figure 10: Signal of the detections reported by the algorithm of the inpainted maps, measured on the (unfiltered) CMB maps, for each of the four methods. In green, we have the detections outside the confidence masks; in blue, detections outside the inpainting mask but inside the common confidence mask; and in salmon, detections inside the inpainting mask.



Figure 11: Region of the sky where the point source is detected, in infrared (2MASS). The width of the image is $16'$, while a pixel in Planck CMB maps is around $1.7'$ wide. The center of the Planck pixel where the maximum is detected is marked with a cyan cross.

- Argüeso, F., Salerno, E., Herranz, D., Sanz, J.L., Kuruoğlu, E.E., Kayabol, K., 2011. A bayesian technique for the detection of point sources in cosmic microwave background maps. *Monthly Notices of the Royal Astronomical Society* 414, 410–417.
- Axelsson, M., Ihle, H.T., Scodeller, S., Hansen, F.K., 2015. Testing for foreground residuals in the planck foreground cleaned maps: A new method for designing confidence masks. *Astronomy and Astrophysics* 578, A44.
- Baldi, P., Kerkycharian, G., Marinucci, D., Picard, D., 2009. Asymptotics for spherical needlets. *The Annals of Statistics* 37, 1150–1171.
- Bardeen, J.M., Szalay, A., Kaiser, N., Bond, J., 1985. The statistics of peaks of gaussian random fields. *Astrophysical Journal* 304, 15–61.
- Benjamini, Y., Hochberg, Y., 1995. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society. Series B (Methodological)* 57, 289–300.
- BICEP2/Keck and Planck Collaborations, 2015. Joint analysis of bicep2/keck array and planck data. *Physical Review Letters* 114, 101301.
- Bobin, J., Sureau, F., Starck, J.L., Rassat, A., Paykari, P., 2014. Joint planck and wmap cmb map reconstruction. *Astronomy & Astrophysics* 563, A105.
- Cammarota, V., Marinucci, D., Wigman, I., 2016. On the distribution of the critical values of random spherical harmonics. *The Journal of Geometric Analysis* 26, 3252–3324.
- Cheng, D., Cammarota, V., Fantaye, Y., Marinucci, D., Schwartzman, A., 2016. Multiple testing of local maxima for detection of peaks on the (celestial) sphere. *Bernoulli*, in press [arXiv:1602.08296](https://arxiv.org/abs/1602.08296).
- Cheng, D., Schwartzman, A., 2015. Distribution of the Height of Local Maxima of Gaussian Random Fields. *Extremes* 18, 213–240.
- Cheng, D., Schwartzman, A., et al., 2017. Multiple testing of local maxima for detection of peaks in random fields. *The Annals of Statistics* 45, 529–556.
- Delabrouille, J., Betoule, M., Melin, J.B., Miville-Deschênes, M.A., Gonzalez-Nuevo, J., Jeune, M.L., Castex, G., de Zotti, G., Basak, S., Ashdown, M., Aumont, J., Baccigalupi, C., Banday, A., Bernard, J.P., Bouchet, F.R., Clements, D.L., da Silva, A., Dickinson, C., Dodu, F., Dolag, K., Elsner, F., Fauvet, L., Fa, G., Giardino, G., Leach, S., Lesgourgues, J., Liguori, M., Macias-Perez, J.F., Massardi, M., Matarrese, S., Mazzotta, P., Montier, L., Mottet, S., Paladini, R., Partridge, B., Piffaretti, R., Prezeau, G., Prunet, S., Ricciardi, S., Roman, M., Schaefer, B., Toffolatti, L., 2013. The pre-launch Planck Sky Model: A model of sky emission at submillimetre to centimetre wavelengths. *Astronomy & Astrophysics* 553, A96.
- Durastanti, C., 2017. Tail behaviour of mexican needlets. *Journal of Mathematical Analysis and Applications* 447, 716–735.
- Geller, D., Mayeli, A., 2009a. Continuous wavelets on compact manifolds. *Mathematische Zeitschrift* 262, 895.
- Geller, D., Mayeli, A., 2009b. Nearly tight frames and space-frequency analysis on compact manifolds. *Mathematische Zeitschrift* 263, 235–264.
- Gorski, K.M., Hivon, E., Banday, A.J., Wandelt, B.D., Hansen, F.K., Reinecke, M., Bartelman, M., 2005. HEALPix – a Framework for High Resolution Discretization, and Fast Analysis of Data Distributed on the Sphere. *The Astrophysical Journal* 622, 759–771.
- Hinshaw, G., Larson, D., Komatsu, E., Spergel, D., Bennett, C., Dunkley, J., Nolte, M., Halpern, M., Hill, R., Odegard, N., et al., 2013. Nine-year wilkinson microwave anisotropy probe (wmap) observations: cosmological parameter results. *The Astrophysical Journal Supplement Series* 208, 19.
- Hunter, J.D., 2007. Matplotlib: A 2D Graphics Environment. *Computing in Science & Engineering* 9, 90–95.
- Jones, E., Oliphant, T., Peterson, P., others, 2001. SciPy: Open source scientific tools for Python.
- Kamionkowski, M., Kosowsky, A., Stebbins, A., 1997. A probe of primordial gravity waves and vorticity. *Physical Review Letters* 78, 2058.
- Kovac, J.M., Leitch, E., Pryke, C., Carlstrom, J., Halverson, N., Holzapfel, W., 2002. Detection of polarization in the cosmic microwave background using dasi. *Nature* 420, 772.
- Lan, X., Marinucci, D., 2009. On the dependence structure of wavelet coefficients for spherical random fields. *Stochastic Processes and their Applications* 119, 3749–3766.
- Lan, X., Marinucci, D., et al., 2008. The needlets bispectrum. *Electronic Journal of Statistics* 2, 332–367.
- Lewis, A., 2011. CAMB notes. URL: <https://cosmologist.info/notes/CAMB.pdf>.
- Lyth, D.H., Riotto, A., 1999. Particle physics models of inflation and the cosmological density perturbation. *Physics Reports* 314, 1–146.
- MacTavish, C., Ade, P., Bock, J., Bond, J., Borrill, J., Boscaleri, A., Cabella, P., Contaldi, C., Crill, B., De Bernardis, P., et al., 2006. Cosmological parameters from the 2003 flight of boomerang. *The Astrophysical Journal* 647, 799.
- Marinucci, D., Peccati, G., 2011. Random fields on the sphere: representation, limit theorems and cosmological applications. volume 389. Cambridge University Press.
- Marinucci, D., Pietrobon, D., Balbi, A., Baldi, P., Cabella, P., Kerkycharian, G., Natoli, P., Picard, D., Vittorio, N., 2008. Spherical needlets for cosmic microwave background data analysis. *Monthly Notices of the Royal Astronomical Society* 383, 539–545.
- Mayeli, A., 2010. Asymptotic uncorrelation for mexican needlets. *Journal of Mathematical Analysis and Applications* 363, 336–344.
- McKinney, W., 2010. Data Structures for Statistical Computing in Python, pp. 51–56.
- Narcowich, F., Petrushev, P., Ward, J., 2006. Localized Tight Frames on Spheres. *SIAM Journal on Mathematical Analysis* 38, 574–594.
- Oliphant, T.E., 2006. A Guide to NumPy. volume 1. Trelgol Publishing USA.
- Pietrobon, D., Balbi, A., Marinucci, D., 2006. Integrated sachs-wolfe effect from the cross correlation of wmap 3year and the nrao vla sky survey data: New results and constraints on dark energy. *Physical Review D* 74, 043524.
- Planck Collaboration IV, 2018. Planck 2018 results. IV. Diffuse component separation [arXiv:1807.06208](https://arxiv.org/abs/1807.06208).
- Planck Collaboration IX, 2016. Planck 2015 results. IX. Diffuse component separation: CMB maps. *Astronomy and Astrophysics* 594, A9.
- Planck Collaboration X, 2016. Planck 2015 results. X. Diffuse component separation: Foreground maps. *Astronomy and Astrophysics* 594, A10.
- Planck Collaboration XI, 2016. Planck 2015 results. xi. cmb power spectra, likelihoods, and robustness of parameters. *Astronomy & Astrophysics* 594, A11.
- Planck Collaboration XIII, 2016. Planck 2015 results. xiii. cosmological parameters. *Astronomy & Astrophysics* 594, A13.
- Planck Collaboration XXVI, 2016. Planck 2015 results. XXVI. The Second Planck Catalogue of Compact Sources. *Astronomy and Astrophysics* 594, A26.
- Schwartzman, A., Gavrilov, Y., Adler, R.J., 2011. Multiple testing of local maxima for detections of peaks in 1D. *Annals of statistics* 39, 3290–3319.
- Scodeller, S., Hansen, F.K., Marinucci, D., 2012. Detection of new point sources in wmap 7 year data using internal templates and needlets. *Astrophysical Journal* 753, 27.
- Scodeller, S., Rudjord, C., Hansen, F.K., Marinucci, D., Geller, D., Mayeli, A., 2011. Introducing Mexican Needlets for CMB Analysis: Issues for Practical Applications and Comparison with Standard Needlets. *The Astrophysical Journal* 733, 121.
- Skrutskie, M.F., Cutri, R.M., Stiening, R., Weinberg, M.D., Schneider, S., Carpenter, J.M., Beichman, C., Capps, R., Chester, T., Elias, J., Huchra, J., Liebert, J., Lonsdale, C., Monet, D.G., Price, S., Seitzer, P., Jarrett, T., Kirkpatrick, J.D., Gizis, J.E., Howard, E., Evans, T., Fowler, J., Fullmer, L., Hurt, R., Light, R., Kopan, E.L., Marsh, K.A., McCallon, H.L., Tam, R., Dyk, S.V., Wheelock, S., 2006. The two micron all sky survey (2mass). *The As-*

- tronomical Journal 131, 1163–1183.
- Starck, J.L., Murtagh, F., Fadili, J.M., 2010. Sparse image and signal processing: Wavelets, curvelets, morphological diversity. *Journal of Electronic Imaging* 19, 049901–049901–2.
- The Astropy Collaboration, 2018. The Astropy Project: Building an inclusive, open-science project and status of the v2.0 core package [arXiv:1801.02634](#).