

A fast and accurate algorithm for spherical harmonic analysis on HEALPix grids with applications to the cosmic microwave background radiation

Kathryn P. Drake*, Grady B. Wright

Department of Mathematics, Boise State University, 1910 University Drive, Boise, ID 83725-1555, United States of America

ARTICLE INFO

Article history:

Received 23 October 2019

Received in revised form 1 May 2020

Accepted 5 May 2020

Available online 14 May 2020

Keywords:

HEALPix

Fast spherical harmonic transform

Non-uniform fast Fourier transform

Double Fourier sphere

Cosmic microwave background radiation

ABSTRACT

The Hierarchical Equal Area isoLatitude Pixelation (HEALPix) scheme is used extensively in astrophysics for data collection and analysis on the sphere. The scheme was originally designed for studying the Cosmic Microwave Background (CMB) radiation, which represents the first light to travel during the early stages of the universe's development and gives the strongest evidence for the Big Bang theory to date. Refined analysis of the CMB angular power spectrum can lead to revolutionary developments in understanding the nature of dark matter and dark energy. In this paper, we present a new method for performing spherical harmonic analysis for HEALPix data, which is a central component to computing and analyzing the angular power spectrum of the massive CMB data sets. The method uses a novel combination of a non-uniform fast Fourier transform, the double Fourier sphere method, and Slevinsky's fast spherical harmonic transform [38]. For a HEALPix grid with N pixels (points), the computational complexity of the method is $\mathcal{O}(N \log^2 N)$, with an initial set-up cost of $\mathcal{O}(N^{3/2} \log N)$. This compares favorably with $\mathcal{O}(N^{3/2})$ runtime complexity of the current methods available in the HEALPix software when multiple maps need to be analyzed at the same time. Using numerical experiments, we demonstrate that the new method also appears to provide better accuracy over the entire angular power spectrum of synthetic data when compared to the current methods, with a convergence rate at least two times higher.

© 2020 Elsevier Inc. All rights reserved.

1. Introduction

About 379,000 years after the universe began, the dense plasma of matter cooled enough for neutral hydrogen to form. During this epoch of recombination, the universe was becoming increasingly transparent to photons, which eventually began to move freely through space. Now faintly glowing at the edge of the observable universe, these photons form the Cosmic Microwave Background (CMB) radiation, which has become the strongest evidence for the Big Bang Theory to date [3]. While the CMB has been deemed “the most perfect black body ever measured in nature” [42], there are temperature and polarization fluctuations that give insight into the primordial universe [28]. These anisotropies are consequences of the initial density distribution of matter, and analyzing them can provide a better understanding of the geometry and composition of the universe [3,15].

* Corresponding author.

E-mail addresses: kathryndrake@u.boisestate.edu (K.P. Drake), gradywright@boisestate.edu (G.B. Wright).




Fig. 1. CMB component map from the Planck mission [30] (a) and corresponding (scaled) angular power spectrum (b).

Using ground-based telescopes, balloons, and satellites which probe the sky in the microwave and infra-red frequencies, scientists have measured the CMB temperature differences at small angular scales. These measurements are quantized and stored as a high resolution sky map of the CMB using the Hierarchical Equal Area isoLatitude Pixelation (HEALPix) scheme [11] for the sphere; see Fig. 1a) for an example sky map. Once a sky map is composed, it can then be analyzed by its angular power spectrum. This quantity measures the amplitude of the CMB temperature fluctuations as a function of angular scale and is used to estimate parameters of the cosmological model for the universe [42]. For example, the confirmation of the first peak in the temperature angular power spectrum affirmed that the universe is spatially flat [17]. The values of the temperature angular spectrum at higher frequencies are crucial to many aspects of modern cosmology, including the density of dark matter and dark energy in the universe. The CMB power spectrum (Fig. 1b) is calculated from the spherical harmonic coefficients, a_ℓ^m , of the sky map as follows:

$$C_\ell = \frac{1}{2\ell + 1} \sum_m |a_\ell^m|^2. \quad (1)$$

The spherical harmonic conventions used in this work are detailed in Appendix A.

The HEALPix scheme [11] and the associated eponymous software [10] have a number of desirable properties for data collection on the sphere. First, each pixel has the same surface area and the pixel centers (points) are quasi-uniformly distributed over the sphere. This is important since any white noise produced by the microwave receivers is exactly integrated into white noise in the pixel area. Second, the pixels produced by the scheme are based on a hierarchical subdivision of the sphere, which allows for data locality in computer memory and fast search procedures. Finally, the pixel centers are isolatitudinal, allowing for a significant reduction in the computational cost of performing discrete spherical harmonic transforms—a central component to computing and analyzing the angular power spectrum of the CMB data sets, which from the Planck mission consist of millions of pixels [30]. These properties have made the HEALPix scheme popular for other applications in astrophysics/astronomy [21,29,35], and to several other disciplines, including geophysics [41], planetary science [25], nuclear engineering [32], and computer vision [16].

In this paper, we focus on an aspect of the HEALPix scheme that has received very little attention in the literature: accuracy and computational complexity improvements of the discrete spherical harmonic transform. We first review the current techniques used in the HEALPix software [10], which are based on equal-weight quadrature, ring-weight quadrature, and pixel-weight quadrature. We then introduce a new algorithm for computing spherical harmonic coefficients for data collected on HEALPix grids. The main motivation for the method is Slevinsky's recently developed fast spherical harmonic transform (FSHT) [38], which converts bivariate Fourier coefficients for data on the sphere to spherical harmonic coefficients of the data with near optimal complexity. By combining the nonuniform fast Fourier transform (NUFFT) [33] and the double Fourier sphere (DFS) [40] methods, we give a fast and accurate method for obtaining the bivariate Fourier coefficients for functions sampled on the HEALPix grid, which we then use with the FSHT to obtain the spherical harmonic coefficients. For a HEALPix grid with N pixels (points), the computational complexity of the method is $\mathcal{O}(N \log^2 N)$, with an initial set-up cost of $\mathcal{O}(N^{3/2} \log N)$, which compares favorably with the complexity of the current methods available in the HEALPix software when multiple maps need to be analyzed at the same time. Using numerical experiments, we demonstrate that the new method also appears to be more accurate than the current methods for synthetic data over the whole spectrum, with a convergence rate at least two times higher. We believe this new scheme will be useful not only for CMB analysis, but also for the many applications of the HEALPix scheme given above that require a spherical harmonic analysis. Additionally, the algorithm presented here has natural generalizations for other “equal-area” isolatitudinal sampling strategies for sphere that do not have a natural way to do fast and accurate spherical harmonic transforms [6,20,22,34].




Fig. 2. HEALPix grid with resolutions, from left to right, $N_{\text{side}} = 1, 2, 4, 8$. The lines indicate the pixel boundaries and the solid dots represent the pixel centers or points.

The remainder of the paper is structured in the following manner. In section 2, we offer supporting information on the HEALPix grid as well as details and analysis of the current methods used in the HEALPix software for computing the spherical harmonic coefficients of CMB maps. We present the new algorithm for fast spherical harmonic analysis of data collected on the HEALPix grid in section 3. Numerical results comparing the presented method with that of the methods in the HEALPix software for calculating the angular power spectrum of functions on the sphere are given in section 4. Finally, in section 5, we give some brief conclusions and remarks on future directions of the work.

2. Background and current approach

2.1. HEALPix scheme

The HEALPix scheme¹ was created to discretize functions on the sphere at high resolutions. In addition to creating an equal area pixelization of the sphere, one of the primary motivations behind the scheme was to allow for more computationally efficient spherical harmonic transforms on increasingly large CMB datasets [11]. While there are many options for discretizing the sphere, there is no known deterministic method that gives a set of quasiuniform points and allows for exact spherical harmonic decompositions of band-limited functions using equal-weight quadrature. While the HEALPix scheme does not offer optimal complexity for spherical harmonic analyses, it does achieve some efficiency gains over existing schemes for discretizing the sphere. This improvement is accomplished primarily by the isolatitudinal distribution of pixels.

The HEALPix grid resolution is defined using the parameter $N_{\text{side}} = 2^t$, $t \in \mathbb{N}$, which creates N_{side}^2 equal area divisions of each base pixel. Fig. 2 illustrates the base resolution grid, $t = 0$, and the increasing levels of refinement $t = 1, 2, 3$, where each base pixel is subdivided further into four equal area pixels. A HEALPix map therefore has $N = 12N_{\text{side}}^2$ equal area (but differently shaped) pixels, with the centers placed on $4N_{\text{side}} - 1$ rings of constant latitude. For any N_{side} , the HEALPix centers, which we henceforth call the HEALPix points, are defined algebraically using three regions of the sphere, two polar ($\mathcal{N}$ and $\mathcal{S}$) and one equatorial ($\mathcal{E}$) [19]. In spherical coordinates, the points in these regions are given as

$$\begin{aligned} \mathcal{N} &:= \left\{ \left(\arccos \left(1 - \frac{j^2}{3N_{\text{side}}^2} \right), \frac{\pi \left(k + \frac{1}{2} \right)}{2j} \right) : j = 1, \dots, N_{\text{side}} - 1, k = 0, \dots, 4j - 1 \right\} \\ \mathcal{E} &:= \left\{ \left(\arccos \left(\frac{2(2N_{\text{side}} - j)}{3N_{\text{side}}} \right), \frac{\pi \left(k + \frac{(j+1) \bmod 2}{2} \right)}{2N_{\text{side}}} \right) : j = N_{\text{side}}, \dots, 3N_{\text{side}}, k = 0, \dots, 4N_{\text{side}} - 1 \right\} \\ \mathcal{S} &:= \left\{ \left(\arccos \left(- \left(1 - \frac{j^2}{3N_{\text{side}}^2} \right) \right), \frac{\pi \left(k + \frac{1}{2} \right)}{2j} \right) : j = 1, \dots, N_{\text{side}} - 1, k = 0, \dots, 4j - 1 \right\}. \end{aligned} \quad (2)$$

The final HEALPix point set is $\mathcal{X} = \mathcal{N} \cup \mathcal{E} \cup \mathcal{S}$. The number of points on each ring varies in the polar regions, with only four points on the rings closest to the north and south poles of the sphere, whereas the rings in the equatorial region have the same number of points. This point distribution is illustrated more clearly in Fig. 3, where the HEALPix points are displayed to a latitude-longitude grid.

The biggest computational advantage for spherical harmonic analysis in the HEALPix scheme lies in the equally-spaced points on each ring of constant latitude. While this aides computation in the longitude direction with FFTs, the misaligned and unequally spaced points in latitude make fast bivariate Fourier analysis impossible without modification. We address this in the new algorithm presented in section 3.

¹ The HEALPix scheme produces a grid consisting of a collection of pixels of different shapes but the same area. However, for our method we do not exploit this fact and simply treat the center of each pixel as a point with the given value of the pixel.




Fig. 3. HEALPix grid on $[0, 2\pi] \times [0, \pi]$, where θ is latitude, and λ is longitude. The point sets in the northern ($\mathcal{N}$), equatorial ($\mathcal{E}$), and southern ($\mathcal{S}$) regions are shown in blue, red, and yellow, respectively. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

2.2. HEALPix software spherical harmonic analysis

The standard method in the HEALPix software [10] for estimating the angular power spectrum (1) of data at the HEALPix points approximates the exact spherical harmonic coefficients ($\tilde{a}_\ell^m$) of the data as

$$a_\ell^m = \frac{4\pi}{N} \sum_{i=1}^N \bar{Y}_\ell^m(\lambda_i, \theta_i) f(\lambda_i, \theta_i), \quad 0 \leq \ell \leq \ell_{\max}, -\ell \leq m \leq \ell, \quad (3)$$

where (λ_i, θ_i) are HEALPix points in longitude-latitude, f is the data, and Y_ℓ^m is a spherical harmonic of degree ℓ and order m (see Appendix A for a discussion of the spherical harmonic conventions used in this paper). While the user can input any band limit $\ell_{\max}$ for this approximation, the software default is $\ell_{\max} = 3N_{\text{side}} - 1$. Due to the isolatitudinal nature of the HEALPix points, this computation is done with $\mathcal{O}(N^{3/2})$ complexity as opposed to $\mathcal{O}(N^2)$ [11]. Note that $N = \mathcal{O}(\ell_{\max}^2)$, so the complexity of the a_ℓ^m computation is equivalent to $\mathcal{O}(\ell_{\max}^3)$. The expression (3) is a low-order approximation to the continuous inner product (23) which defines the coefficients, since it uses a simple equal weight quadrature. To improve this approximation, the software employs an iterative procedure, which is referred to as a “Jacobi iteration” [11]. In order to illustrate how the iterative method converges, we explain it below in the language of linear algebra.

The analysis operation, defined in (3), produces an approximation to the spherical harmonic coefficients from the data f on the sphere, whereas the synthesis operation reconstructs the data given the spherical harmonic coefficients:

$$\hat{f}(\lambda_i, \theta_i) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} a_\ell^m Y_\ell^m(\lambda_i, \theta_i), \quad i = 1, \dots, N \quad (4)$$

Note that we use a hat on f to indicate that computing the spherical harmonic coefficients according to (3) and using them in (4) gives different function values in general. In matrix-vector notation, we denote (3) and (4) as

Analysis: $\mathbf{a} = \mathbf{A}\mathbf{f}$

Synthesis: $\hat{\mathbf{f}} = \mathbf{S}\mathbf{a}$,

where $\mathbf{a}$ is the vector of spherical harmonic coefficients and $\mathbf{f}$ and $\hat{\mathbf{f}}$ are the vectors of data values at the HEALPix points. Using this notation, the iterative procedure in the HEALPix software can be written as

$$\begin{aligned} \mathbf{r}^{(k+1)} &= \mathbf{f} - \mathbf{S}\mathbf{a}^{(k)}, \\ \mathbf{a}^{(k+1)} &= \mathbf{a}^{(k)} + \mathbf{A}\mathbf{r}^{(k+1)}, \end{aligned} \quad (5)$$

where $\mathbf{r}$ is the residual vector and $\mathbf{a}^{(0)} = \mathbf{A}\mathbf{f}$. Substituting the first equation of (5) into the last and using the fact that the analysis matrix satisfies $\mathbf{A} = \frac{4\pi}{N}\mathbf{S}^*$, gives the iteration

$$\mathbf{a}^{(k+1)} = \mathbf{a}^{(k)} + \frac{4\pi}{N}\mathbf{S}^*(\mathbf{f} - \mathbf{S}\mathbf{a}^{(k)}) = \frac{4\pi}{N}\mathbf{S}^*\mathbf{f} + \left(\mathbf{I} - \frac{4\pi}{N}\mathbf{S}^*\mathbf{S}\right)\mathbf{a}^{(k)}. \quad (6)$$

Table 1Spectral radius of the Richardson iteration matrix from (6) for different values of N .

N_{side}	N	$\rho\left(\mathbf{I} - \frac{4\pi}{N}\mathbf{S}^*\mathbf{S}\right)$
2	48	0.1986
4	192	0.0932
8	768	0.0600
16	3072	0.0475
32	12288	0.0421

This is just stationary Richardson iteration (or Gradient Decent) with relaxation parameter $\frac{4\pi}{N}$ applied to the normal equations $\mathbf{S}^*\mathbf{S}\mathbf{a} = \mathbf{S}^*\mathbf{f}$ [4, pp. 44–45]. Thus, the iterative procedure converges to the least squares solution to (3), provided the spectral radius of $\mathbf{I} - \frac{4\pi}{N}\mathbf{S}^*\mathbf{S}$ is less than one. The spectral radius also determines the convergence rate. Since the HEALPix points are equidistributed, we know that (3) converges to the integral (23) as $N \rightarrow \infty$ (in exact arithmetic) [14]. Thus, the spectral radius of $\mathbf{I} - \frac{4\pi}{N}\mathbf{S}^*\mathbf{S}$ converges to 0 as $N \rightarrow \infty$ and we expect the iteration (6) to converge more rapidly as N increases. Table 1 gives evidence of this result by displaying the spectral radius of $\mathbf{I} - \frac{4\pi}{N}\mathbf{S}^*\mathbf{S}$ for increasing values of N .

The default option in the HEALPix software sets the number of iterations of (6) to 3. While this does improve the accuracy of computing the spherical harmonic coefficients, it adds to the cost, as each iteration requires doing an analysis and synthesis ((3) and (4)) at a cost of $\mathcal{O}(\ell_{\text{max}}^3)$ operations each. Since the solution converges to the least squares solution, one could improve the convergence of the Richardson iteration method by using algorithms like LSQR or conjugate gradient on the normal equations [27].

2.2.1. Pixel weights and ring weights

As an alternative to the iterative scheme, the HEALPix software also has the option of using quadrature weights to improve the accuracy of the computation of the spherical harmonic coefficients. In this case, the equal weight quadrature approximation (3) is generalized to

$$a_\ell^m = \sum_{i=1}^N w_i \bar{Y}_\ell^m(\lambda_i, \theta_i) f(\lambda_i, \theta_i), \quad 0 \leq \ell \leq \ell_{\text{max}}, -\ell \leq m \leq \ell, \quad (7)$$

where w_i are the quadrature weights. There are two options for using quadrature weights. The first is “pixel weights”, which uses different weights for *each HEALPix point*. These weights are computed using the positive quadrature weight algorithm from [19], which consists of solving a system involving a Gram matrix containing the spherical harmonics whose size is proportional to N [31]. For large N , the weights are computed once and stored. The second option is to use “ring weights”, which use different weights for *each ring* of the HEALPix point sets. The computation of the ring weights is done using similar ideas to the pixel weights, but a much smaller system has to be solved [31]. The new method introduced in this paper does not use quadrature weights directly, but instead computes the bivariate Fourier coefficients of the HEALPix data and then uses these to obtain the spherical harmonic coefficients.

3. HP2SPH

The algorithm presented here, named HP2SPH, introduces a new way to calculate the spherical harmonic coefficients of data sampled at the HEALPix points (2). The outline for the algorithm is given in Algorithm 1, and each of the pieces are described below.

Algorithm 1 HP2SPH.

Input: Data sampled at the HEALPix point set of size N , $\{f_j\}$, $j = 1, \dots, N$.

Output: Approximate spherical harmonic coefficients, $\{a_\ell^m\}$, $0 \leq \ell \leq \ell_{\text{max}}$, $-\ell \leq m \leq \ell$

1. Transform the data to a tensor product latitude-longitude grid:
 - (i) Upsample the data in longitude from the northern ($\mathcal{N}$) and southern ($\mathcal{S}$) point sets using FFT
 - (ii) Shift (interpolate) the data from the equatorial ($\mathcal{E}$) point set so it is longitudinally aligned
 2. Compute the bivariate Fourier coefficients:
 - (i) Apply the DFS method
 - (ii) Apply the inverse NUFFT-II in latitude
 - (iii) Apply the inverse FFT in longitude
 3. Obtain the spherical harmonic coefficients via the FSHT
-




Fig. 4. (a) HEALPix points with $N_{\text{side}} = 16$ displayed in latitude and longitude and (b) the corresponding upsampled points.

3.1. Step 1: transform the data to a tensor product latitude-longitude grid

As described in Section 2.1, the HEALPix grid has an unequal number of points on the rings in the northern ($\mathcal{N}$) and southern ($\mathcal{S}$) sets (2), and the points on the rings in the equatorial ($\mathcal{E}$) set are shifted on every other ring. This structure leads to the pixels being misaligned in latitude. By upsampling the data on the northern and southern points in longitude so that there are an equivalent samples of the data on each ring and shifting the data at equatorial points in longitude, we can use fast algorithms to obtain the bivariate Fourier coefficients of the data as discussed in the next section. On the two polar point sets, we upsample the data using the trigonometric interpolant of the data on each ring of these sets to the non-shifted equally spaced longitude points on the equatorial rings, i.e.,

$$\lambda_k = \frac{k}{2N_{\text{side}}} \pi, \quad k = 0, \dots, 4N_{\text{side}} - 1. \quad (8)$$

We also interpolate the data on the rings in the equatorial point set with shifted longitude points, to these λ values. Fig. 4(b) illustrates the upsampling procedure leading to a tensor product latitude-longitude grid of data of size $(4N_{\text{side}} - 1) \times 4N_{\text{side}}$.

We describe the interpolation procedure here for the data in the northern point set $\mathcal{N}$. Consider the latitude values for the northern rings, $\theta_j = \arccos\left(1 - \frac{j^2}{3N_{\text{side}}^2}\right)$, $j = 1, \dots, N_{\text{side}}$. We approximate the data in each ring using a trigonometric expansion of the form

$$f(\lambda, \theta_j) =: f_j(\lambda) = \sum_{n=-2j}^{2j} c_n^{(j)} e^{in\lambda}. \quad (9)$$

The coefficients in the expansion are determined by enforcing interpolation of the given data values

$$f\left(\frac{k + \frac{1}{2}}{2j} \pi, \theta_j\right), \quad k = 0, \dots, 4j - 1.$$

With the minor algebraic manipulation of (9),

$$f_j\left(\frac{k + \frac{1}{2}}{2j} \pi\right) = \sum_{n=-2j}^{2j-1} c_n^{(j)} e^{in \frac{k + \frac{1}{2}}{2j} \pi} = \sum_{n=-2j}^{2j-1} c_n^{(j)} e^{in \frac{\pi}{4j}} e^{in \frac{k}{2j} \pi} = \sum_{n=-2j}^{2j-1} \tilde{c}_n^{(j)} e^{in \frac{k}{2j} \pi}, \quad k = 0, \dots, 4j - 1,$$

we see the interpolation conditions yield the system

$$\sum_{n=-2j}^{2j-1} \tilde{c}_n^{(j)} e^{in \frac{k}{2j} \pi} = f\left(\frac{k + \frac{1}{2}}{2j} \pi, \theta_j\right), \quad k = 0, \dots, 4j - 1, \quad (10)$$

which can be computed using the inverse FFT. We can then obtain the Fourier coefficients $c_n^{(j)}$ in (9) for the data at the non-shifted values through simple multiplication.² Finally, we pad the vector containing the coefficients $c_n^{(j)}$ with the appropriate number of zeros to get to a total of $4N_{side}$, so that the expansion in longitude in each ring has the same number of Fourier coefficients. The values of the interpolant on each ring can then be obtained at the upsampled values (8) by applying the FFT on these padded vectors. A similar procedure can be applied to the data on the rings in the southern point set $\mathcal{S}$.

On the rings in the equatorial set $\mathcal{E}$ where the longitude values are shifted by $\pi(k + \frac{1}{2})/(2N_{side})$, we compute the Fourier coefficients of the data using (10) with $j = N_{side}$. We then obtain the coefficients in (9) at the non-shifted points from which the values can be computed using the FFT. No padding or upsampling is needed in this case.

3.2. Step 2: compute bivariate Fourier coefficients

Bivariate Fourier analysis for data on the sphere requires the application of the DFS method to obtain periodicity of the data in latitude and to retain spherical symmetry. When we apply this method to the upsampled HEALPix data, there is an issue that the points in latitude are not equally-spaced, making the standard FFT unsuitable. To bypass this issue we use an NUFFT. Both the DFS technique and NUFFT method we use are discussed below for completeness.

3.2.1. Double Fourier sphere (DFS) method

A natural approach to representing a function on the sphere is to use a latitude-longitude coordinate transform, defined by

$$x(\lambda, \theta) = \cos \lambda \sin \theta, \quad y(\lambda, \theta) = \sin \lambda \sin \theta, \quad z(\lambda, \theta) = \cos \theta, \quad (\lambda, \theta) \in [0, 2\pi] \times [0, \pi], \quad (11)$$

which maps the sphere to a rectangular domain. While this transformation allows for performing computations with the function $f(\lambda, \theta) = f(x(\lambda, \theta), y(\lambda, \theta), z(\lambda, \theta))$, it also introduces artificial boundaries at the north and south poles. In addition, the change of variables does not maintain the symmetry of functions on the sphere. Specifically, the transform described in (11) does not preserve the periodicity in the latitude direction, which is necessary for bivariate Fourier analysis to be applicable and for results to make physical sense. These problems are solved by the DFS method.

Originally introduced by Merilees in [23] (see also [40]) the DFS method transforms a function on the sphere to a rectangular grid while maintaining bi-periodicity. This can be thought of as “doubling up” the function $f(\lambda, \theta)$ to form a new function that preserves periodicity in both the latitude and longitude directions. Algebraically, this new function, $\tilde{f}(\lambda, \theta)$, is defined on $[0, 2\pi] \times [0, 2\pi]$ as follows [40]

$$\tilde{f}(\lambda, \theta) = \begin{cases} g(\lambda, \theta), & (\lambda, \theta) \in [0, \pi] \times [0, \pi], \\ h(\lambda - \pi, \theta), & (\lambda, \theta) \in [\pi, 2\pi] \times [0, \pi], \\ h(\lambda, 2\pi - \theta), & (\lambda, \theta) \in [0, \pi] \times [\pi, 2\pi], \\ g(\lambda - \pi, 2\pi - \theta), & (\lambda, \theta) \in [\pi, 2\pi] \times [\pi, 2\pi], \end{cases} \quad (12)$$

where $g(\lambda, \theta) = f(\lambda, \theta)$ and $h(\lambda, \theta) = f(\lambda + \pi, \theta)$ for $(\lambda, \theta) \in [0, \pi] \times [0, \pi]$. Fig. 5 illustrates the DFS method applied to the surface of the Earth, which shows the preservation of bi-periodicity in (c). We note that the DFS method can also be easily applied to discrete data sampled at a tensor product latitude-longitude grid using (12), which is what we do for the upsampled HEALPix data. In this case, (12) corresponds to flipping and shifting the data matrix appropriately.

Once the DFS method is applied to a function on the sphere, it can be approximated using a 2D bivariate Fourier expansion:

$$\tilde{f}(\lambda, \theta) \approx \sum_{j=-\lfloor \frac{m}{2} \rfloor}^{\lceil \frac{m}{2} \rceil - 1} \sum_{k=-\lfloor \frac{n}{2} \rfloor}^{\lceil \frac{n}{2} \rceil - 1} C_{jk} e^{ij\theta} e^{ik\lambda}, \quad (13)$$

where m and n represent the number of frequencies in (doubled-up) latitude and longitude, respectively.

Note that the HEALPix grid does not include points at the north and south poles. When applying the DFS to the upsampled data from Step 1, this leads to a relatively large gap in the points in latitude over the poles compared to the other points, which can lead to issues with the inverse NUFFT (see below). To bypass this issue, we construct values at the two poles by using a weighted quadratic least squares fit [8] that combines the data from the three rings closest to each pole.

Remark. The Fourier coefficients of the upsampled data in longitude are computed in Step 1. These can be used directly in the DFS procedure by applying (12) in Fourier space in the λ variable, which amounts to appending the (padded) coefficient matrix from Step 1 with a flipped version of itself with all odd wave numbers multiplied by -1 . It then only remains to

² Horner's rule (and Estrin's scheme for higher accuracy at small angles [7]) could also be used to implement the shift, avoiding loss of accuracy due to evaluation of high frequency complex exponentials.




Fig. 5. Illustration of the DFS method: (a) The surface of earth, (b) the surface mapped onto a latitude-longitude grid, and (c) the surface after applying the DFS method. [40].

compute the Fourier coefficients in latitude θ to obtain the full bivariate Fourier coefficients. This is the focus of the next step.

3.2.2. Nonuniform fast Fourier transform (NUFFT)

The use of the nonuniform discrete Fourier transform (NUDFT) in many domain sciences has led to the development of algorithms for computing it efficiently. If these algorithms are quasi-optimal requiring $\mathcal{O}(n \log n)$, then they are referred to as a nonuniform fast Fourier transform (NUFFT). Given a vector $\mathbf{c} \in \mathbb{C}^n$, the one-dimensional NUDFT computes the vector $\mathbf{f} \in \mathbb{C}^n$ defined by

$$f_j = \sum_{k=0}^{n-1} c_k e^{-2\pi i x_j \omega_k}, \quad 0 \leq j \leq n-1, \quad (14)$$

where $x_j \in [0, 1]$ are the samples and $\omega_k \in [0, n]$ are the frequencies. If the samples are equispaced ($x_j = j/n$) and the frequencies are integer ($\omega_k = k$), then the transform is a uniform DFT, which can be computed by the FFT in $\mathcal{O}(n \log n)$ operations [5]. When either the samples are nonequispaced or the frequencies are noninteger, the FFT does not directly apply without some careful manipulation [2].

Ruiz and Townsend [33] contributed to the collection of NUFFT algorithms by utilizing low rank matrix approximations to relate the NUDFT to the uniform DFT. There are three types of NUDFTs and inverse NUDFTs that they account for in their algorithm: NUDFT-I, which has uniform samples but noninteger frequencies; NUDFT-II, which has nonuniform samples and integer frequencies; NUDFT-III, which has both nonuniform samples and nonuniform frequencies [12]. Since our HP2SPH method only uses the one-dimensional inverse NUFFT of the second type, we discuss the NUFFT-II algorithm [33].

Given Fourier coefficients, $\mathbf{c} \in \mathbb{C}^{n \times 1}$, the NUFFT-II attempts to approximate the vector

$$\mathbf{f} = \hat{F}_2 \mathbf{c}, \quad (15)$$

to machine precision in quasi-optimal complexity. Here $(\hat{F}_2)_{jk} = e^{-2\pi i x_j \omega_k}$, x_j are nonuniform samples (restricted to be in $[0, 1]$), and k are integer frequencies for $0 \leq j, k \leq n-1$. Notice that the DFT matrix for uniform samples and integer frequencies is similarly $F_{jk} = e^{-2\pi i jk/n}$. The key to the NUFFT-II algorithm described in [33] is that if the samples are nearly equispaced, then $\hat{F}_2$ can be related to the Hadamard product of F and a low rank matrix. This means that given a rank K approximation which relates $\hat{F}_2$ and F , the NUFFT-II can then be computed using K FFTs with $\mathcal{O}(Kn \log n)$ cost. In practice, machine (double) precision can be achieved with $K = 14$ [33].

In the case of the inverse NUFFT-II, Ruiz and Townsend advocate for the use of the conjugate gradient (CG) method in order to solve the linear system $\hat{F}_2 \mathbf{c} = \mathbf{f}$ for $\mathbf{c}$. Since $\hat{F}_2$ is not hermitian, the CG method is applied to the normal equations:

$$\hat{F}_2^* \hat{F}_2 \mathbf{c} = \hat{F}_2^* \mathbf{f}, \quad (16)$$

where $\widehat{F}_2^* \widehat{F}_2$ is simply a Toeplitz matrix. Therefore, the inverse NUFFT-II can be calculated using the CG method and a fast Toeplitz multiplication [9] in $\mathcal{O}(R_{CG} n \log n)$ operations, where R_{CG} is the number of CG iterations. The following suggestion is placed on the nonuniform function samples to avoid ill-conditioning in the system (16) [33]:

$$\left| x_j - \frac{j}{n} \right| \leq \frac{\gamma}{n}, \quad 0 \leq j \leq n-1, \quad (17)$$

where $0 \leq \gamma < 1/4$. When this condition is satisfied, it has been experimentally observed that $R_{CG} \leq 10$ for a large range of n .

For the method proposed in this paper, we apply the inverse NUFFT-II in latitude to the DFS upsampled HEALPix data from Step 2. Unfortunately, the HEALPix points in latitude direction do not meet the condition (17). To bypass this issue, we instead use a least squares solution to compute fewer coefficients at the higher wave numbers than what the given data may support. We describe this procedure below since it not discussed in [33].

The inverse NUFFT-II method computes first column of the symmetric Toeplitz matrix $\widehat{F}_2^* \widehat{F}_2$ in (16) in the following manner:

$$\widehat{F}_2^* \widehat{F}_2 \mathbf{e}_1 = \widehat{F}_2^* \mathbf{1} = (\mathbf{1}^T \widehat{F}_2)^* = (\widehat{F}_2^T \mathbf{1})^*.$$

The last expression above can be computed efficiently by the NUFFT-I algorithm, since the NUDFT-I matrix is simply the transpose of the NUDFT-II matrix [33]. To compute a least squares solution to (15) with fewer coefficients, we simply truncate the first column obtained from the above method to $m < n$ terms and form the resulting $m \times m$ Toeplitz matrix $\widehat{F}_2^* \widehat{F}_2$. The right hand side for the least squares solution is obtained by similarly computing $\widehat{F}_2^* \mathbf{f}$ and truncating this to m terms.

For the DFS upsampled HEALPix data from Step 2, there are $8N_{side}$ coefficients in latitude, but only $4N_{side}$ coefficients in longitude. To keep the number of Fourier modes in both directions (nearly) the same, we choose $m = 4N_{side} + 1$ as the truncation parameter for the least squares solution for computing the Fourier coefficients in latitude. This is also a convenient choice since the method in step three for converting bivariate Fourier coefficients of data on the sphere to spherical harmonic coefficients requires the number of coefficients in each direction is the same and an odd number (we explain how to convert the coefficients in longitude to $m = 4N_{side} + 1$ in the next section).

Remark. For problems where the data may contain noise (e.g., for the CMB application), there could be an issue with this noise being amplified in steps 1 and 2. For step 1, we should not expect any additional noise to be introduced, since we are simply computing the Fourier coefficients on each ring using the original data and then shifting the coefficients and padding them with zeros. Step 2 has two areas where there could be an issue with noisy data. The first is in constructing values at the poles and the second is in the application of the NUFFT in latitude. However, both of these steps apply a least squares procedure, which provides some smoothing. In our tests on CMB data, we did not observe any amplification of noise that was present in the data. The HP2SPH method has a further benefit of using a backward stable algorithm for computing the spherical harmonic coefficients (as discussed next), which ensures that the resulting uncertainty in the spherical harmonic coefficients has only a low algebraic growth with respect to degree and is always proportional to the norm of the noise in the data.

3.3. Step 3: obtain the spherical harmonic coefficients via the fast spherical harmonic transform (FSHT)

In [38], Slevinsky derives a fast, backward stable method for the transformation between spherical harmonic expansions and their bivariate Fourier series (given in (13)) by viewing it as a change of basis. This relation is defined as a connection problem, and the matrices that arise in the present case are well-conditioned, making them ideal for fast computations. Slevinsky describes the change of basis in two steps: converting expansions in normalized associated Legendre functions to those of only order zero and one, and then re-expressing these in trigonometric form. In other words, it uses spherical harmonic expansions of order zero and one as intermediate expressions between higher-order spherical harmonics and their corresponding bivariate Fourier coefficients.

The first step of the algorithm takes advantage of the fact that the matrix of connection coefficients between the associated Legendre functions of all orders and those of order zero and one can be represented by a product of Givens rotation matrices. This enables the use of the butterfly algorithm, which can be thought of as an abstraction of the algebraic properties of fast Fourier transforms. The term butterfly was introduced in [24], where it was used for analyzing scattering from electrically large surfaces, and then further developed in [26] for use in special function transforms. Slevinsky uses the butterfly algorithm to perform a factorization of the well-conditioned matrices of connection coefficients.

The second step of this method exploits the hierarchical decompositions of the connection coefficient matrices between the associated Legendre functions of order zero and one to the Chebyshev polynomials of the first and second kind, respectively. This step quickly computes the fast orthogonal polynomial transforms using an adaptation of the Fast Multipole Method [13] and low-rank matrix approximations. The total pre-computation time for both steps is $\mathcal{O}(\ell_{max}^3 \log \ell_{max})$, and execution time is asymptotically optimal with $\mathcal{O}(\ell_{max}^2 \log^2 \ell_{max})$ operations. This FSHT is implemented in Julia with the package FastTransforms [37] (as are the NUFFT methods from [33] used in Step 2).

The FSHT in FastTransforms assumes the input function has a bivariate Fourier expansion of the form

$$\tilde{f}(\lambda, \theta) = \sum_{j=0}^p \sum_{k=-p}^p g_j^k \frac{e^{ik\lambda}}{\sqrt{2\pi}} \begin{cases} \cos j\theta, & k \text{ even} \\ \sin(j+1)\theta, & k \text{ odd} \end{cases}. \quad (18)$$

Any function on the sphere is required to have these even/odd conditions on its bivariate Fourier coefficients [23]. At the end of step 2 we have obtained the bivariate Fourier expansion of the data of the form

$$\tilde{f}(\lambda, \theta) = \sum_{j=-p}^p \sum_{k=-p}^{p-1} C_{jk} e^{ij\theta} e^{ik\lambda}, \quad (19)$$

where $p = N_{\text{side}}/2$. Since we are dealing with real-valued data, we can expand Fourier coefficients array in λ to an odd number of points. The expanded array is defined by

$$X_{j,k} = \begin{cases} C_{j,k} & \text{if } -p+1 \leq k \leq p-1 \\ \frac{1}{2}C_{j,-p} & \text{if } k = \pm p \end{cases}, \quad -p \leq j, k \leq p.$$

Using the array X , we can write (19) as

$$\begin{aligned} \tilde{f}(\lambda, \theta) &= \sum_{k=-p}^p e^{ik\lambda} \sum_{j=0}^p ((X_{jk} + \bar{X}_{-jk}) \cos(j\theta)) + \frac{1}{i} (X_{jk} - \bar{X}_{-jk}) \sin(j\theta) \\ &= \sum_{j=0}^p \sum_{k=-p}^p e^{ik\lambda} \begin{cases} ((X_{jk} + \bar{X}_{-jk}) \cos(j\theta)), & k \text{ even} \\ ((X_{jk} - \bar{X}_{-jk}) \sin(j\theta)), & k \text{ odd} \end{cases}, \end{aligned}$$

from which we can obtain the coefficients g_j^k in (18).

The FSHT software takes bivariate Fourier coefficients g_j^k as input in an array organized as follows:

$$\begin{pmatrix} g_0^0 & g_0^{-1} & g_0^1 & \cdots & g_0^{-p} & g_0^p \\ g_1^0 & g_1^{-1} & g_1^1 & \cdots & g_1^{-p} & g_1^p \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ g_{p-1}^0 & g_{p-1}^{-1} & g_{p-1}^1 & \cdots & g_{p-1}^{-p} & g_{p-1}^p \\ g_p^0 & 0 & 0 & \cdots & g_p^{-p} & g_p^p \end{pmatrix}.$$

The output of the software is the approximate spherical harmonic coefficients of the data arranged in an array of the form

$$\begin{pmatrix} a_0^0 & a_1^{-1} & a_1^1 & a_2^{-2} & a_2^2 & \cdots & a_p^{-p} & a_p^p \\ a_1^0 & a_2^{-1} & a_2^1 & a_3^{-2} & a_3^2 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ a_{p-2}^0 & a_{p-1}^{-1} & a_{p-1}^1 & a_p^{-2} & a_p^2 & & \vdots & \vdots \\ a_{p-1}^0 & a_p^{-1} & a_p^1 & 0 & 0 & \cdots & 0 & 0 \\ a_p^0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 \end{pmatrix}.$$

The angular power spectrum (1) can then be computed from this array.

4. Numerical results

In this section we present a few numerical tests comparing the spherical harmonics and angular power spectrum (1) computed by our new method HP2SPH to the values computed by the HEALPix software employing the iterative scheme (6), pixel weights, and ring weights (7). The first test compares the rate at which the two methods converge to the spherical harmonic coefficients by applying them to deterministic (i.e. non-noisy) functions sampled at the HEALPix points with known coefficients. The second test compares the accuracy of the methods using deterministic functions, which have a known power spectrum. In the third test, we compare the methods after calculating the angular power spectrum for a real CMB data map, which contains noise.




Fig. 6. Maximum absolute errors as a function of t for the computed spherical harmonic coefficients of (20) using HP2SPH and (a) HEALPix (3 iterative refinement steps), pixel weights, ring weights and (b) HEALPix with increasing iterative steps. The lines in the figure are the lines of best fit to the data and the convergence rates are determined from the slope of this line (as displayed in the plot legends).

4.1. Convergence of spherical harmonic coefficients

We choose the test function

$$f(\lambda, \theta) = \sum_{j=1}^3 c_j (2 - 2\mathbf{x}(\lambda, \theta) \cdot \mathbf{x}(\lambda_j, \theta_j))^{3/2}, \quad (20)$$

where $\mathbf{x}(\lambda, \theta) = [x(\lambda, \theta) \ y(\lambda, \theta) \ z(\lambda, \theta)]$ from (11) and the parameters, which we picked randomly, are given by

$$\{c_1, c_2, c_3\} = \{5, -3, 8\},$$

$$\{\lambda_1, \lambda_2, \lambda_3\} = \{0.891498158152027, 2.650004294134628, 5.753735997130328\},$$

$$\{\theta_1, \theta_2, \theta_3\} = \{1.232217523107963, 2.059244524372349, 0.537798840821172\}.$$

The function $(2 - 2\mathbf{x}(\lambda, \theta) \cdot \mathbf{x}(\lambda_c, \theta_c))^{3/2}$ is called a *potential spline* of order $3/2$ centered at $\mathbf{x}(\lambda_c, \theta_c)$ and its exact spherical harmonic coefficients are given by [18]

$$\tilde{a}_\ell^m = \frac{18\pi}{(\ell + 5/2)(\ell + 3/2)(\ell + 1/2)(\ell - 1/2)(\ell - 3/2)} Y_\ell^m(\lambda_c, \theta_c). \quad (21)$$

These values are used to compare the convergence rates of the methods to the exact spherical harmonic coefficients of f . We do this by plotting in Fig. 6 the maximum absolute errors of the coefficients against the parameter t , which is used to determine the grid resolution parameter ($N_{\text{side}} = 2^t$). Note that the spherical harmonic coefficients of the CMB decay at a rate between $\mathcal{O}(\ell^{-2})$ and $\mathcal{O}(\ell^{-3})$ [39], which is slower than the decay rate of the coefficients of the test function (20) (since, for all $-\ell \leq m \leq \ell$, $|Y_\ell^m(\lambda_c, \theta_c)| \leq \sqrt{(2\ell + 1)/4\pi}$, and the remaining terms in (21) decay at a rate of $\mathcal{O}(\ell^{-5})$). This means that the test function has more smoothness than an actual CMB data set.

Fig. 6(a) compares the four methods and shows that the HP2SPH method converges to the spherical harmonic coefficients of (20) at a rate at least twice as fast as any of the HEALPix methods. Although the consecutive iterative refinement steps used in the HEALPix method produce progressively better errors, Fig. 6(b) illustrates that this does not improve the convergence rate (as discussed in Section 2.2). It is also important to note that after 8 iterative steps, there are no further improvements in the accuracy, indicating the algorithm has nearly converged to the least squares solution to (3). The results the HEALPix method with pixel weights look pretty erratic with convergence achieved to 8–10 digits around $t = 7$, but no further reductions. This could be because of potential errors in the computed quadrature weights.

Next, we test how the convergence rates of the methods are affected by high frequencies. To do this we add 15 spherical harmonics to (20) with the following randomly chosen degrees and orders:

ℓ	176	190	191	230	248	283	292	303	326	366	388	404	421	446	448
m	56	81	124	40	155	274	27	145	55	343	200	78	420	284	234




Fig. 7. Maximum absolute errors as a function of t for the computed spherical harmonic coefficients of (20) augmented with spherical harmonics using HP2SPH, HEALPix with 3 iterative refinement steps, pixel weights, and ring weights.




Fig. 8. (a) Scaled angular power spectrum of (20) as a function of degree ℓ computed by the HEALPix software with 3 iterative refinement steps, the HP2SPH method, the HEALPix method with ring weights, and the HEALPix method with pixel weights. The exact power spectrum is given by the black o's. Here $N_{\text{side}} = 2^{10}$, which is $N = 12, 582, 912$ total points. (b) Absolute errors in the (scaled) angular power spectrum of the results from (a) as a function of degree ℓ .

The results from this test are displayed in Fig. 7. We note that in order to ensure the calculation of asymptotic convergence rates for all methods, we excluded the errors for $t = 8$ in the lines of best fit. This test shows that the convergence rates of the new method as well as the iterative HEALPix scheme and the HEALPix method with ring weights are not affected by the addition of high degree spherical harmonic terms to (20). The HEALPix method with pixel weights shows a similar erratic behavior to Fig. 6(a), with no steady reductions in the errors after $t = 9$. The new HP2SPH method has the lowest errors of all the methods.

4.2. Errors in the angular power spectrum

In this test, we first explore the accuracy of all the methods for computing the angular power spectrum of (20). These results are compared to the exact spectrum, which is calculated using the exact spherical harmonic coefficients (21) in (1). The angular power spectrum of the four methods are displayed in Fig. 8(a). We see that the algorithms produce similar results for lower degrees ℓ , but the HEALPix method with the iterative refinement scheme (6) diverges for degrees greater than approximately $\ell = 100$, whereas the ring weight and pixel weight quadrature (7) results diverge for degrees greater than $\ell = 200$. To compare the methods more directly, Fig. 8(b) plots the absolute errors in the angular power spectrum for each degree ℓ . While the HEALPix method with ring weights performs the best out of all of the HEALPix methods for $\ell \leq 50$, the errors increase rapidly for higher ℓ . Conversely, the HEALPix method with pixel weights does not perform as well




Fig. 9. Absolute errors in the (scaled) angular power spectrum of (20) augmented with high-degree spherical harmonics computed by the HEALPix software with 3 iterative refinement steps, the HP2SPH method, the HEALPix method with ring weights, and the HEALPix method with pixel weights as a function of degree ℓ . (a) Displays the errors for degrees $\ell = 1, \dots, 2000$, while (b) displays the errors only for $\ell = 450, \dots, 1500$ to better show the how good the methods are at recovering the spectrum at the degrees of the appended spherical harmonics. Here $N_{\text{side}} = 2^{10}$, which is $N = 12, 582, 912$ total points.

for smaller ℓ , yet it performs better than the other HEALPix methods for larger ℓ . The HP2SPH method offers comparable results for low ℓ as the pixel weights method while still maintaining accuracy for high ℓ .

Similar to Test 1, we test the accuracy of the methods when computing the power spectrum of data with high frequencies, as occur in real CMB data maps. As before, we add several spherical harmonic functions of high degree to the function (20). The new test function appends the following degrees and orders:

ℓ	589	633	636	766	829	943	974	1009	1085	1219	1294	1346	1404	1485	1493
m	188	269	414	134	517	912	93	483	183	1143	667	259	1400	946	779

The power spectrum of this function is the same as (20), but with the value at each degree ℓ of appended spherical harmonics increased by $\frac{1}{2\ell+1}$.

Fig. 9(a) displays the errors in the angular power spectrum of this new function for all the methods over all ℓ , while Fig. 9(b) displays the errors only over the range of ℓ that were appended to the base function. We again see that the ring weights provide the highest accuracy for low ℓ , but then the errors increase rapidly, while the HEALPix iterative method has the largest errors for low ℓ , but some of the smallest errors for the ℓ corresponding to the appended spherical harmonics. In contrast, the new HP2SPH method provides small errors over the entire angular power spectrum, clearly showing its benefits over the HEALPix methods for determining the angular power spectrum of deterministic functions on the sphere.

4.3. Application to real CMB map

Our final numerical test compares the methods on the real CMB map shown in Fig. 1. Fig. 10 (a) shows the angular power spectrum for this map computed with the methods, while (b) shows the errors in the three HEALPix methods compared to the new HP2SPH method. We see from the figure that the new method (in blue) produces visually the same results as the HEALPix methods (red, magenta, and cyan), indicating that the new method is not anymore susceptible to noise than the HEALPix methods.

5. Conclusions and remarks

We have presented a new method, HP2SPH, for performing discrete spherical harmonic analysis on data collected using the HEALPix scheme. The method utilizes the FFT, NUFFT, and the FSHT to compute the spherical harmonic transform in near optimal computational complexity ($\mathcal{O}(N \log^2 N)$ complexity for N total HEALPix points). Several numerical tests were presented to demonstrate the improved convergence and accuracy of the new method relative to the various HEALPix approaches for problems involving synthetic data with no noise, except that introduced by roundoff errors. In the case of a real CMB map with additional types of noise, the power spectra computed by the methods show good agreement. The new HP2SPH benefits further from the backward stability of the FSHT, which ensures that the resulting uncertainty in the spherical harmonic coefficients has only a low algebraic growth with respect to degree and is always proportional to the




Fig. 10. (Scaled) angular power spectrum of the CMB data map displayed in Fig. 1 (a) with $N_{\text{side}} = 2^{11}$ for the four methods discussed in the paper (left), and the relative errors of the HEALPix software methods against the HP2SPH method (right).

norm of the noise in the data. We anticipate the new method will be applicable to the many other areas where the HEALPix scheme is used and is naturally generalizable to other equal-area isolatitudinal sampling schemes for the sphere.

For our next steps, we will work to optimize the implementation of the method, which is currently in Julia, to improve its actual run-time. This will include transcribing our code into a lower-level language like C; efforts in this direction are already underway for the FSHT [36]. In addition to this, we will include the ability to perform Fourier synthesis on a CMB map, i.e. given an angular power spectrum, we will return the corresponding CMB map values. For this purpose, our method has another advantage over HEALPix in that we will have the bivariate Fourier coefficients, which will simply make the synthesis an application of the FFT and NUFFT. Finally, we plan to add functionality for analyzing the polarization of CMB maps.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We are grateful to Ann Almgren for her comments and suggestions as well as Richard Mikael Slevinsky for assisting us with using the FastTransforms package. We also thank the entire HEALPix team, especially Kris Górski and Martin Reinecke, for their feedback and assistance on this work. Finally, we thank the anonymous reviewers whose comments helped improve the paper. The first author's work was supported by the SMART Scholarship, which is funded by The Under Secretary of Defense-Research and Engineering, National Defense Education Program/BA-1, Basic Research. The second author's work was supported by National Science Foundation grant CCF 1717556.

Appendix A. Spherical harmonic conventions

We denote a scalar spherical harmonic of degree $\ell \geq 0$ and order $-\ell \leq m \leq \ell$ as $Y_\ell^m(\lambda, \theta)$, where λ is the azimuth angle and θ is the zenith angle. We define these functions as

$$Y_\ell^m(\lambda, \theta) = \sqrt{\frac{2\ell+1}{4\pi}} \sqrt{\frac{(\ell-m)!}{(\ell+m)!}} P_\ell^m(\cos \theta) e^{im\lambda}, \quad m = 0, 1, \dots, \ell, \quad (22)$$

where $Y_\ell^m = (-1)^m Y_\ell^{|m|}$ for $m < 0$ and $P_\ell^m(\cos \theta)$ are the associate Legendre functions. As eigenfunctions of the Laplace-Beltrami operator, spherical harmonics are the natural basis for square integrable functions on the sphere [1]. In other words, any L^2 -integrable function f on the sphere can be uniquely represented as

$$f(\lambda, \theta) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \tilde{a}_\ell^m Y_\ell^m(\lambda, \theta),$$

where the spherical harmonic coefficients, $\tilde{a}_\ell^m$, are found using the usual L^2 -inner product for scalar functions on the sphere:

$$\tilde{a}_\ell^m = \langle f, Y_\ell^m \rangle = \int_0^{2\pi} \int_0^\pi f(\lambda, \theta) \bar{Y}_\ell^m(\lambda, \theta) \sin \theta d\theta d\lambda. \quad (23)$$

References

- [1] K. Atkinson, W. Han, *Spherical Harmonics and Approximations on the Unit Sphere: An Introduction*, Springer, Berlin, Heidelberg, 2012.
- [2] R. Beatson, L. Greengard, A short course on fast multipole methods, in: *Wavelets Multilevel Methods Elliptic PDEs*, Vol. 1, 1997, pp. 1–37.
- [3] C.L. Bennett, D. Larson, J.L. Weiland, N. Jarosik, G. Hinshaw, N. Odegard, K.M. Smith, R.S. Hill, B. Gold, M. Halpern, E. Komatsu, M.R. Nolte, L. Page, D.N. Spergel, E. Wollack, J. Dunkley, A. Kogut, M. Limon, S.S. Meyer, G.S. Tucker, E.L. Wright, Nine-year Wilkinson microwave anisotropy probe (WMAP) observations: final maps and results, *Astrophys. J. Suppl. Ser.* 208 (2013) 20.
- [4] M. Benzi, G.H. Golub, J. Liesen, Numerical solution of saddle point problems, *Acta Numer.* 14 (2005) 1–137.
- [5] J. Cooley, J. Tukey, An algorithm for the machine calculation of complex Fourier series, *Math. Comput.* 19 (1965) 297–301.
- [6] R.G. Crittenden, N.G. Turok, Exactly azimuthal pixelizations of the sky, Tech. Rep. DAMTP-1998-78, University of Cambridge, 1998, arXiv:astro-ph/9806374.
- [7] G. Estrin, Organization of computer systems – the fixed plus variable structure computer, in: *Proc. Western Joint IRE/IEEE-ACM Comput. Conf.*, 1960, pp. 33–40.
- [8] G.E. Fasshauer, *Matrix Approximation Methods with MATLAB*, World Scientific Publishers, Singapore, 2007.
- [9] G. Golub, C.F. Van Loan, *Matrix Computations*, Johns Hopkins University Press, Baltimore, 1996.
- [10] K.M. Górski, B.D. Wandelt, E. Hivon, F.K. Hansen, A.J. Banday, The HEALPix primer, arXiv:astro-ph/9905275, 1999.
- [11] K.M. Górski, B.D. Wandelt, E. Hivon, F.K. Hansen, A.J. Banday, M. Reinecke, M. Bartelmann, HEALPix – a framework for high resolution discretization, and fast analysis of data distributed on the sphere, *Astrophys. J.* 622 (2005) 759–771.
- [12] L. Greengard, J. Lee, Accelerating the nonuniform fast Fourier transform, *SIAM Rev.* 46 (2004) 443–454.
- [13] L. Greengard, V. Rokhlin, A fast algorithm for particle simulations, *J. Comput. Phys.* 73 (1987) 325–348.
- [14] D.P. Hardin, T. Michaels, E.B. Saff, A comparison of popular point configurations on S^2 , *Dolomites Res. Notes Approx.* 9 (2016) 16–49.
- [15] G. Hinshaw, D. Larson, E. Komatsu, D.N. Spergel, C.L. Bennett, J. Dunkley, M.R. Nolte, M. Halpern, R.S. Hill, N. Odegard, L. Page, K.M. Smith, J.L. Weiland, B. Gold, N. Jarosik, A. Kogut, M. Limon, S.S. Meyer, G.S. Tucker, E. Wollack, E.L. Wright, Nine-year Wilkinson microwave anisotropy probe (WMAP) observations: cosmological results, *Astrophys. J. Suppl. Ser.* 208 (2013) 19.
- [16] R.C. Hoover, A.A. Maciejewski, R.G. Roberts, Pose detection of 3-D objects using S^2 -correlated images and discrete spherical harmonic transforms, in: *2008 IEEE International Conference on Robotics and Automation*, 2008, pp. 993–998.
- [17] W. Hu, S. Dodelson, Cosmic microwave background anisotropies, *Annu. Rev. Astron. Astrophys.* 40 (2002) 171–216.
- [18] S. Hubbert, B.J.C. Baxter, Radial basis functions for the sphere, in: *Recent Progress in Multivariate Approximation (Witten-Bommerholz, 2000)*, in: *Internat. Ser. Numer. Math.*, vol. 137, Birkhäuser, Basel, 2001, pp. 33–47.
- [19] J. Keiner, D. Potts, Fast evaluation of quadrature formulae on the sphere, *Math. Comput.* 77 (2008) 397–419.
- [20] P. Leopardi, A partition of the unit sphere into regions of equal area and small diameter, *Electron. Trans. Numer. Anal.* 25 (2006) 309–327.
- [21] C. Majeau, E. Agol, N.B. Cowan, A two-dimensional infrared map of the extrasolar planet HD 189733b, *Astrophys. J.* 747 (2012) L20.
- [22] Z. Malkin, A new equal-area isolatitudinal grid on a spherical surface, *Astron. J.* 158 (2019) 158.
- [23] P.E. Merilees, The pseudospectral approximation applied to the shallow water equations on a sphere, *Atmosphere* 11 (1973) 13–20.
- [24] E. Michielssen, A. Boag, A multilevel matrix decomposition algorithm for analyzing scattering from large structures, *IEEE Trans. Antennas Propag.* 44 (1996) 1086–1093.
- [25] R.S. Miller, G. Nerurkar, D.J. Lawrence, Enhanced hydrogen at the lunar poles: new insights from the detection of epithermal and fast neutron signatures, *J. Geophys. Res., Planets* 117 (2012).
- [26] M. O’Neil, F. Woolfe, V. Rokhlin, An algorithm for the rapid evaluation of special function transforms, *Appl. Comput. Harmon. Anal.* 28 (2010) 203–226.
- [27] C.C. Paige, M.A. Saunders, LSQR: an algorithm for sparse linear equations and sparse least squares, *ACM Trans. Math. Softw.* 8 (1982) 43–71.
- [28] P. Paykari, J.-L. Starck, CMB data analysis, *Adv. Mach. Learning Data Mining Astron.* (2012) 55–87.
- [29] W.J. Percival, S. Cole, D.J. Eisenstein, R.C. Nichol, J.A. Peacock, A.C. Pope, A.S. Szalay, Measuring the baryon acoustic oscillation scale using the sloan digital sky survey and 2dF galaxy redshift survey, *Mon. Not. R. Astron. Soc.* 381 (2007) 1053–1066.
- [30] Planck Collaboration 2005, Planck, the scientific programme, ESA publication ESA-SCI(2005)/01, (2005).
- [31] M. Reinecke, personal communication, June 5, 2019.
- [32] C. Roecker, N.S. Bowden, G. Carosi, M. Heffner, I. Jovanovic, Reconstruction algorithms for directional neutron detection using a time projection chamber, *Nucl. Technol.* 180 (2012) 231–240.
- [33] D. Ruiz-Antolín, A. Townsend, A nonuniform fast Fourier transform based on low rank approximation, *SIAM J. Sci. Comput.* 40 (2018) A529–A547.
- [34] E.B. Saff, A.B.J. Kuijlaars, Distributing many points on a sphere, *Math. Intell.* 19 (1997) 5–11.
- [35] L.P. Singer, L.R. Price, Rapid Bayesian position reconstruction for gravitational-wave transients, *Phys. Rev. D* 93 (2016) 024013.
- [36] R.M. Slevinsky, *FastTransforms*, <https://github.com/MikaelSlevinsky/FastTransforms.GitHub>.
- [37] R.M. Slevinsky, *FastTransforms.jl*, <https://github.com/MikaelSlevinsky/FastTransforms.GitHub>.
- [38] R.M. Slevinsky, Fast and backward stable transforms between spherical harmonic expansions and bivariate fourier series, *Appl. Comput. Harmon. Anal.* 47 (2019) 585–606.
- [39] J.-L. Starck, M.J. Fadili, A. Rassat, Low- ℓ CMB analysis and inpainting, *Astron. Astrophys.* 550 (2013).
- [40] A. Townsend, H. Wilber, G.B. Wright, Computing with functions in spherical and polar geometries I. The sphere, *SIAM J. Sci. Comput.* 38 (2016) C403–C425.
- [41] D.R. Weimer, E.K. Sutton, M.G. Mlynczak, L.A. Hunt, Intercalibration of neutral density measurements for mapping the thermosphere, *J. Geophys. Res. Space Phys.* 121 (2016) 5975–5990.
- [42] M. White, *Anisotropies in the CMB*, Proceedings of Division of Particles and Fields of the American Physical Society, 1999.