

UC Berkeley

UC Berkeley Previously Published Works

Title

Harmonic analysis of discrete tracers of large-scale structure

Permalink

<https://escholarship.org/uc/item/0sm4m3bc>

Journal

Journal of Cosmology and Astroparticle Physics, 2024(05)

ISSN

1475-7516

Authors

Lizancos, Antón Baleato

White, Martin

Publication Date

2024-05-01

DOI

10.1088/1475-7516/2024/05/010

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Harmonic analysis of discrete tracers of large-scale structure

Antón Baleato Lizancos  and Martin White 

Berkeley Center for Cosmological Physics, UC Berkeley, CA 94720, USA

Department of Physics, University of California, Berkeley, CA 94720, USA

Lawrence Berkeley National Laboratory, One Cyclotron Road, Berkeley, CA 94720, USA

E-mail: a.baleatolizancos@berkeley.edu, mwhite@berkeley.edu

Abstract. It is commonplace in cosmology to analyze fields projected onto the celestial sphere, and in particular density fields that are defined by a set of points e.g. galaxies. When performing an harmonic-space analysis of such data (e.g. an angular power spectrum) using a pixelized map one has to deal with aliasing of small-scale power and pixel window functions. We compare and contrast the approaches to this problem taken in the cosmic microwave background and large-scale structure communities, and advocate for a direct approach that avoids pixelization. We describe a method for performing a pseudo-spectrum analysis of a galaxy data set and show that it can be implemented efficiently using well-known algorithms for special functions that are suited to acceleration by graphics processing units (GPUs). The method returns the same spectra as the more traditional map-based approach if in the latter the number of pixels is taken to be sufficiently large and the mask is well sampled. The method is readily generalizable to cross-spectra and higher-order functions. It also provides a convenient route for distributing the information in a galaxy catalog directly in harmonic space, as a complement to releasing the configuration-space positions and weights, and a route to spectral apodization. We make public a code enabling the application of our method to existing and upcoming datasets.

Contents

1	Introduction	1
2	Power spectrum algorithm	3
2.1	Pseudo-spectra	3
2.2	Direct SHT	6
2.3	Algorithm	7
3	Examples	8
3.1	Golden spiral	8
3.2	Lognormal mock catalog	8
4	Conclusions	11
A	Numerical and implementation details	13
A.1	Computing the spherical harmonics	13
A.2	Interpolation	14
B	Wigner 3j symbols	16

1 Introduction

Observational cosmology is entering an exciting era of “big surveys”, each of which will map large fractions of the sky with unprecedented depth and fidelity. Examples include DESI [1], Euclid [2, 3], SPHEREx [4], LSST on the Vera Rubin Observatory [5], Roman [6], the Simons Observatory [7], the South Pole Observatory and CMB-S4 [8]. These surveys probe the Universe in complementary ways, and the correlated signals within the datasets they will produce will offer a wealth of insight into questions of cosmology, astrophysics and fundamental physics [9]. In this paper we are concerned with techniques for performing auto- and cross-correlations of scalar fields on the sphere, which includes photometric galaxy clustering, CMB lensing, the thermal and kinematic Sunyaev-Zeldovich effects [10, 11] and many other probes. Our particular focus will be on how to handle galaxy data, which is provided as a set of points with significant amounts of small-scale power (e.g. shot noise), in a harmonic-space analysis.

The CMB community has developed sets of tools for dealing with pixelized maps of continuous fields defined on the sphere (e.g. CMB temperature or polarization and lensing convergence), with the most common being the Healpix pixelization scheme [12] and the pseudo- C_ℓ method for measuring their (angular) power spectra [13]. These methods work very well for spectra without large amounts of small-scale power¹, and very efficient “community codes” are available for many of the standard operations (e.g. NaMaster [14], PolSpice [15]). When dealing with projected galaxy data most authors have tried to take advantage of these tools, and so as a first step have produced pixelized maps of the (projected) galaxy overdensity

¹CMB experiments observing at microwave frequencies have beams much larger than typical optical surveys and thus smear out much more of the small-scale structures present in the sky. In addition many of the signal components (e.g. primary CMB anisotropies) have intrinsically very red spectra.

(see e.g. ref. [16] for a pedagogical example from the recent literature). Typically (e.g. [16–18]) these maps are created by binning the galaxy positions, with their systematics weights, into a map, and then doing the same for a random catalog that defines the sample in the absence of clustering². The ratio of the two maps then defines the overdensity field, $1 + \delta_{\text{gal}}$. However, this procedure has some drawbacks. First of all, the tails of the distribution of any inaccuracies in the random (or mask) map are enhanced by this division process, which introduces a non-linear function of e.g. the systematics weights. Another issue is that because of shot-noise the galaxy field has significant small-scale power which means that aliasing is a concern. Aliasing could be reduced by interlacing or a higher-order charge assignment scheme [19–21], but this is difficult in curvilinear coordinates and common practice is to assign galaxies and randoms to a single pixel. The finite size of the pixels also means that the pixel window function can depart from unity at the angular scales targeted by analyses; moreover, when these scales are close to the pixel size, the non-trivial shapes of the pixels become important and pixelization corrections become computationally intractable³. For all these reasons the practice has been to employ a high resolution map (e.g. ref. [23] and the discussion in the appendices of refs. [14, 24]). However, fine pixelizations come with a computational cost. Since the density field involves dividing by the random map, a large number of random points is needed in order to prevent stochastic noise in the randoms from significantly distorting the statistics of the overdensity field. And, perhaps more importantly, the spherical harmonic transform algorithm scales with the number of pixels – even in the highly efficient `HealPix` scheme, it goes as $N_{\text{pix}}^{3/2}$ [12].

In this paper we advocate for a different approach motivated by the procedure that is commonly employed in analyzing three dimensional galaxy (redshift) surveys ([25]; hereafter FKP). The “FKP approach” defines the power spectrum to be the square of the Fourier transform of a normalized difference between the data and a scaled set of randoms. Thus, if α is the ratio of the number of data to the number of randoms, then the FKP field is proportional to $n_g(\mathbf{x}) - \alpha n_r(\mathbf{x})$, with the proportionality constant chosen for convenience. Since the randoms define a mean density, \bar{n} , and the data define $\bar{n}(1 + \delta)$, the difference is $\bar{n} \delta$. The square of the Fourier transform of this field has as expectation value the power spectrum convolved with a window function. If we were to follow this approach, the standard definition of the overdensity, $\delta(\mathbf{x}) \propto n_g(\mathbf{x})/n_r(\mathbf{x}) - 1$, a non-linear operation involving division by the randoms or mean density, would be replaced with $\delta(\mathbf{x}) \propto n_g(\mathbf{x}) - \alpha n_r(\mathbf{x})$. Since the latter is linear in data and randoms, the harmonic-space overdensity could then be obtained directly by subtracting the spherical harmonic transforms of data and scaled randoms. In the limit that the objects are pointlike, the spherical harmonic transform of these fields is simply a sum of spherical harmonics evaluated at the positions of the objects, $Y_{\ell m}^*(\theta_i, \phi_i)$. This can be done without recourse to a map or pixelization, thus avoiding the pixel window function and

²Instead of randoms, some authors (e.g., refs. [16, 17]) work with ‘completeness weights’ that define the local mean density as a modulation of the average number of galaxies per pixel by the fraction of each pixel that is rendered unobservable due to non-idealities. The randoms serve the same purpose while providing a convenient avenue to convert between the various pixelizations and coordinate systems involved in the calculation of completeness. Due to this equivalence, we will not address completeness weights explicitly in what follows.

³Pixelized observations result from averaging the underlying signal within each pixel – i.e., convolving the signal with a ‘pixel window function’. On scales much larger than the pixel size, pixels can be assumed to all be the same shape, and the impact of pixelization is well captured by a simple filtering of the angular spectra by a function that depends only on the angular multipole ℓ ; see, e.g., ref. [22]. An analysis that fully captures the non-trivial shapes of the pixels is computationally intractable.

any aliasing. Highly efficient routines for computing $Y_{\ell m}$ exist that make this appealing and allow us to focus computational effort on the particular ℓ s which are of interest. Moreover, the key steps of the computation are well-suited for acceleration on Graphics Processing Units (GPUs). Finally, this approach provides a convenient method for distributing the galaxy information in harmonic space. The coefficients, $a_{\ell m}$, of the harmonic expansion can be released directly (alongside catalogs of positions and weights in configuration space). Rotating coordinates or generating a pixelized map from these coefficients is straightforward using existing tools⁴ [12].

The outline of the paper is as follows. In §2 we define the basic algorithm, focusing on the case of 2-point statistics (i.e. the angular auto-power spectrum) for definiteness though the generalization to cross-spectra and to higher order functions is reasonably straightforward. Then, in §3, we consider two examples that illustrate the benefits of our proposed approach. We conclude with a general discussion in §4. Some implementation details are relegated to the Appendices.

2 Power spectrum algorithm

Let us now explain in detail the proposed procedure, which closely follows the “pseudo- C_ℓ ” method [13] and the FKP method⁵ [25]. A high-level summary is provided in §2.3.

2.1 Pseudo-spectra

Given either a set of weighted randoms or a pixelized map describing the survey “mask” or “window”, which we denote $w(\theta, \phi)$, let $w_{\ell m}$ be this window’s spherical harmonic transform (SHT; we describe how to compute $w_{\ell m}$ later). In large-scale structure studies this would typically be denoted \bar{n} , the mean density in the absence of clustering, but we shall follow the CMB convention where this is a “window”. Define the square of $w_{\ell m}$ to be a window function⁶

$$W_\ell = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} |w_{\ell m}|^2 . \quad (2.1)$$

The masked/window-convolved data minus scaled randoms or normalized mask is $w(1 + \delta) - w = w\delta$ and its SHT is the difference of the SHTs of each term. This product has SHT [26]

$$\{w\delta\}_{\ell m} \equiv (-1)^m \sum_{\ell_1 m_1} \sum_{\ell_2 m_2} G_{-m m_1 m_2}^{\ell \ell_1 \ell_2} w_{\ell_1 m_1} \delta_{\ell_2 m_2} , \quad (2.2)$$

where we have used the definition of the Gaunt integral

$$\begin{aligned} G_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} &\equiv \int d\hat{n} Y_{\ell_1 m_1}(\hat{n}) Y_{\ell_2 m_2}(\hat{n}) Y_{\ell_3 m_3}(\hat{n}) \\ &= \sqrt{\frac{(2\ell_1 + 1)(2\ell_2 + 1)(2\ell_3 + 1)}{4\pi}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix} , \end{aligned} \quad (2.3)$$

⁴e.g. <https://healpy.readthedocs.io/en/latest/>

⁵Ref. [25] additionally describe “optimal weights” – frequently referred to as “FKP weights” – that balance sample variance and shot noise. We shall not consider those here.

⁶If W_ℓ is estimated by a small number of randoms, one should subtract the shot-noise contribution.

with the terms in parentheses indicating 3j symbols [27] (see Appendix B for an algorithm to compute them efficiently). Now average over m to define

$$\widehat{C}_\ell \equiv \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} |\{w\delta\}_{\ell m}|^2 \quad . \quad (2.4)$$

Taking the expectation value

$$\langle |\{w\delta\}_{\ell m}|^2 \rangle = \sum_{\ell_1 \dots \ell_4} \sum_{m_1 \dots m_4} G_{-m m_1 m_2}^{\ell \ell_1 \ell_2} G_{-m m_3 m_4}^{\ell \ell_3 \ell_4} w_{\ell_1 m_1} w_{\ell_3 m_3}^* \langle \delta_{\ell_2 m_2} \delta_{\ell_4 m_4}^* \rangle \quad (2.5)$$

$$= \sum_{\ell_1 \dots \ell_3} \sum_{m_1 \dots m_3} G_{-m m_2 m_1}^{\ell \ell_2 \ell_1} G_{-m m_2 m_3}^{\ell \ell_2 \ell_3} w_{\ell_1 m_1} w_{\ell_3 m_3}^* C_{\ell_2} \quad (2.6)$$

and using [27]

$$\sum_{m_4 m_1} G_{-m_4 m_1 m_2}^{\ell_4 \ell_1 \ell_2} G_{-m_4 m_2 m_3}^{\ell_4 \ell_1 \ell_3} = \delta_{\ell_2 \ell_3}^{(K)} \delta_{m_2 m_3}^{(K)} \frac{(2\ell_4+1)(2\ell_1+1)}{4\pi} \begin{pmatrix} \ell_4 & \ell_1 & \ell_2 \\ 0 & 0 & 0 \end{pmatrix}^2, \quad (2.7)$$

where $\delta^{(K)}$ is the Kronecker delta function, we find [13]

$$\langle \widehat{C}_\ell \rangle = \sum_{\ell'} M_{\ell \ell'} C_{\ell'} \quad \text{with} \quad M_{\ell \ell'} = \frac{2\ell'+1}{4\pi} \sum_{\lambda} (2\lambda+1) \begin{pmatrix} \ell & \ell' & \lambda \\ 0 & 0 & 0 \end{pmatrix}^2 W_\lambda, \quad (2.8)$$

where $C_{\ell'}$ is the ‘true’ angular power spectrum of the overdensity field.

This mirrors the ‘normal’ derivation of the pseudo- C_ℓ calculation and definition of the mode-coupling matrix, $M_{\ell \ell'}$ [13]. In fact, we can see that simply differencing the data and random catalogs achieves the same result as the more usual approach [17] of generating an overdensity field $\delta(\hat{n})$ (by dividing a data catalog by a random catalog and subtracting one) multiplying it by the mask $w(\hat{n})$ (generated from the random catalog and possibly apodized), and transforming the product: $(D/R - 1) \times R = D - R$.

This \widehat{C}_ℓ (likely binned into bandpowers) and $M_{\ell \ell'}$ are sufficient to enable a comparison of theory and data. However, in general, the spectra are not reported as \widehat{C}_ℓ . Within the CMB community the normal procedure [13, 17] is to bin the \widehat{C}_ℓ in bands of ℓ and apply a mode-decoupling. The mode-decoupling step ensures that if the true theory were well approximated by a series of constant bandpowers in the chosen bins then the reported bandpowers would match this, i.e. the mode-decoupling step is a convenience. On the other hand, the FKP procedure, commonly used in large-scale structure, does not apply such a step but simply a normalization⁷.

Let us consider each approach in turn. In the standard pseudo- C_ℓ approach we take a set of bins, B , defined through a matrix θ_b^ℓ with $b \in B$. A common choice would be $\theta_b^\ell = 1/\Delta\ell$ for ℓ in bin b and zero otherwise, where $\Delta\ell$ is the number of ℓ s in bin b . This simply averages the \widehat{C}_ℓ within each bin. Alternatives to weight the ℓ s within a bin differently can also be considered. Now bin \widehat{C}_ℓ into these bins to form $\widehat{C}_b = \sum_{\ell} \theta_b^\ell \widehat{C}_\ell$. Further assume the true C_ℓ is piecewise constant within the same bins, with values C_b . The binned coupling matrix,

⁷Ignoring the ‘FKP weights’, the normalization adopted in ref. [25] divides by the square root of $\int dV \bar{n}^2$.

$M_{bb'} = \sum_{\ell \in b} \sum_{\ell' \in b'} \theta_b^\ell M_{\ell\ell'}$, relates $\langle \widehat{C}_b \rangle$ to $C_{b'}$ and is invertible for sufficiently broad bins. Multiply by the inverse of this matrix so that

$$\widetilde{C}_b = \sum_{b'} M_{bb'}^{-1} \widehat{C}_{b'} = \sum_{b'} M_{bb'}^{-1} \sum_{\ell m} \frac{\theta_{b'}^\ell}{2\ell + 1} \left| a_{\ell m}^{(d)} - w_{\ell m} \right|^2. \quad (2.9)$$

Note that we have defined our overdensity field following FKP, denoting the SHT of data and scaled randoms as $a_{\ell m}^{(d)}$ and $w_{\ell m}$, respectively. The above ‘mode-decoupled’ bandpowers have expectation value

$$\langle \widetilde{C}_b \rangle = \sum_{b'} M_{bb'}^{-1} \langle \widehat{C}_{b'} \rangle = \sum_{b'\ell\ell'} \left(M_{bb'}^{-1} \theta_{b'}^\ell M_{\ell\ell'} \right) C_{\ell'} \equiv \sum_{\ell'} \mathcal{M}_{bb'\ell'} C_{\ell'}. \quad (2.10)$$

Implicitly, the mode-coupling matrix knows about the normalization of the weighted data (which matches that of the weighted randoms by construction), so the mode-decoupled bandpowers are independent of the chosen normalization.

The mode-coupling matrix $M_{\ell\ell'}$ is strictly positive. However, since the mode-decoupling step involves $M_{bb'}^{-1}$, the matrix $\mathcal{M}_{bb'\ell'}$ for any bin b depends upon all of the chosen bins and will typically have negative entries. The reason is that it is trying to correct for changes in the shape of the pseudo-spectrum induced by the harmonic-space convolution with the mask⁸.

An alternative approach, closer to how FKP proposed to process 3D power spectra, would be to simply normalize and bin $M_{\ell\ell'}$ to define a new pseudo-spectrum which would have only positive weights that do not depend on distant bins. One possibility is

$$\widetilde{C}'_b = \mathcal{N} \sum_{\ell} \theta_b^\ell \widehat{C}_\ell \quad \text{with} \quad \langle \widetilde{C}'_b \rangle = \mathcal{N} \sum_{\ell\ell'} \theta_b^\ell M_{\ell\ell'} C_{\ell'} = \sum_{\ell'} \mathcal{M}'_{bb'\ell'} C_{\ell'}. \quad (2.11)$$

While the choice of \mathcal{N} is arbitrary, a sensible choice is to leave the amplitude of a shot-noise spectrum unchanged. Since C_ℓ is constant for shot-noise, using the addition theorem for spherical harmonics and the conventional normalization $P_\ell(1) = 1$ we have

$$\langle \widehat{C}_\ell \rangle = \left[\frac{\int d\Omega w^2(\hat{n})}{4\pi} \right] C_\ell^{\text{shot}} = \left[\frac{W_0 \int d\Omega w^2(\hat{n})}{\left| \int d\Omega w(\hat{n}) \right|^2} \right] C_\ell^{\text{shot}} = \left[\sum_{\ell'} \frac{(2\ell' + 1)}{4\pi} W_{\ell'} \right] C_\ell^{\text{shot}}, \quad (2.12)$$

where the last equality comes from expanding $w(\hat{n})$ in $Y_{\ell m}$ and using the orthonormality of the spherical harmonics. For the θ_b^ℓ described above, choosing \mathcal{N} as the inverse of the term(s) in square brackets leaves the amplitude of shot-noise at high ℓ unchanged.

Note that this normalization does not address the change in the shape of the power spectrum induced by the convolution in Eq. (2.8). The shape of the reported power spectrum will therefore differ from the true one ‘‘on the sky’’, especially at low ℓ where W_ℓ has significant support. However, we emphasize that this is simply a convention for quoting the data: from the perspective of comparing theory and data one multiplies the theory curve by the supplied $N_{\text{bin}} \times N_\ell$ matrix ($\mathcal{M}_{bb'}$ or $\mathcal{M}'_{bb'}$ in our notation above) to compare to the supplied data (\widetilde{C}_b or \widetilde{C}'_b). Any self-consistent convention is equivalent.

The resulting pseudo-spectrum includes stochastic or Poisson noise. This could be subtracted if desired (as e.g. done in ref. [25]), though the approach gaining traction in the

⁸This can lead to odd behavior in bandpowers near the edges of the domain, as the mode-deconvolution algorithm implicitly assumes that the true C_ℓ are zero beyond the multipole range we have chosen to decouple (see also ref. [28] for a related discussion).




Figure 1. A comparison of the two conventions for \tilde{C}_b for the footprint shown in Fig. 4 and linear bins of $\Delta\ell = 25$. (Left) The ‘true’ full-sky spectrum is shown as the solid black line, the mode-decoupled spectrum as blue points and the normalized spectrum as orange points. (Right) The matrices $\mathcal{M}_{b\ell}$ for each convention, shown for every other bandpower.

field is to include it in both the measurement and the model. This then allows freedom for the noise to have a non-Poisson value whose degree of departure from the Poisson value can be handled with priors if so desired.

Finally we mention that if the computation of the pseudo-spectrum is being done as part of a Monte Carlo in which the footprint (and/or random catalog) is not changing, then the computation of the $w_{\ell m}$, W_ℓ and $M_{\ell\ell'}$ does not need to be redone. In addition, the $Y_{\ell m}(\hat{n}_k)$ grid doesn’t need to be recomputed if ℓ_{\max} is unchanged. This can speed up the computation.

2.2 Direct SHT

Note that by following FKP we don’t need to divide by the mean density, and thus the harmonic transforms can be done directly at the level of the data and random fields. If we are dealing with discrete objects, we don’t really need to go through the pixelization step, and in fact there are some advantages to avoiding it. For a set of point-like objects at $\{\theta_i, \phi_i\}$ with weights ω_i , the spherical harmonic coefficients are

$$a_{\ell m} = \sum_i \omega_i Y_{\ell m}^*(\theta_i, \phi_i), \quad (2.13)$$

where as usual

$$Y_{\ell m}(\theta, \phi) = N_{\ell m} P_\ell^m(\cos\theta) e^{im\phi} \quad \text{and} \quad Y_{\ell, -m} = (-1)^m Y_{\ell m}^*, \quad (2.14)$$

with P_ℓ^m denoting the associated Legendre functions [27, 29] and with normalization

$$N_{\ell m} = \sqrt{\frac{(2\ell+1)}{4\pi} \frac{(\ell-m)!}{(\ell+m)!}} \quad \text{for } m \geq 0 \quad (2.15)$$

independent of position. The computation of the exponential factor (sines and cosines) is straightforward, but calculating P_ℓ^m requires more attention. This is handled in Appendix




Figure 2. Speed of the direct SHT, running on four NVIDIA A100 GPU cores of the Perlmutter computer at NERSC. The colored curves with squares show the wall-time taken (in seconds) to compute the $a_{\ell m}$ up to various ℓ_{\max} values for N_{pnt} points on the sphere with the number of spline points fixed at $N_x = 2048$ (see Appendix A.2). The black dashed line shows the measured cost of determining what interval of the interpolation grid $\cos \theta$ falls into (i.e. ‘binning’) and sorting the data in such a way that their corresponding grid values are arranged in ascending order. The sorting scales as $\mathcal{O}(N \log N)$, as expected. We see that this sorting, which only needs to happen once per dataset, dominates the cost for large N_{pnt} . It could be skipped if the data were already sorted, in which case the time would be dominated by binning (the black dotted line).

A.1, where we describe how to compute $Y_{\ell m}(\theta_k, 0)$ and its first derivative for a grid of θ_k , and in Appendix A.2, where we show how to efficiently interpolate from this grid using Hermite splines (which express the interpolated value as a linear combination of the function and its first derivative at the endpoints of the sub-interval). Storage and computation can be trivially distributed across devices, though this may not be necessary given how fast the algorithm is. With a CPU implementation in C it takes ~ 30 seconds to evaluate the SHT of $2^{24} \simeq 17$ M points up to $\ell_{\max} \simeq 10^3$ on the Perlmutter computer at NERSC (using one AMD Milan CPU). However the algorithm is well-suited for acceleration on GPUs, and using the JAX⁹ library on a Perlmutter node with four NVIDIA A100 GPU cores, we are able to obtain the SHT of 10^8 points up to $\ell_{\max} \simeq 10^3$ in approximately the same amount of time (~ 30 seconds; Fig. 2). This establishes the direct SHT as feasible (and competitive with other approaches) and is likely sufficient for many purposes. However, if high efficiency is needed, there are highly optimized libraries¹⁰ that perform even better than this.

2.3 Algorithm

Let us now summarize our proposed algorithm. Suppose we are given a set of data and randoms together with their corresponding weights. Use Eq. (2.13) on the data to compute $a_{\ell m}^{(d)}$ and on the randoms to produce $a_{\ell m}^{(r)}$, normalizing the randoms such that $a_{00}^{(r)} = a_{00}^{(d)}$. Use Eq. (2.1) with $w_{\ell m} = a_{\ell m}^{(r)}$ to compute W_{ℓ} (subtracting shot noise) and Eq. (2.9) to define \tilde{C}_b ,

⁹<https://github.com/google/jax>

¹⁰As one example, the [ducc0 library](#) that is used in ref. [30] has an “adjoint_synthesis_general” method that does the forward SHT.

the binned, mode-decoupled, pseudo- C_ℓ estimate of the power spectrum (or, alternatively, the binned, normalized pseudo- C_ℓ estimate).

In the case that the mask or window is provided in pixelized form, the SHT should be done using whatever routines apply to that pixelization. For example, for HEALPIX this can be done with the `anafast` routine of `healpy` [31]. Use Eq. (2.13) on the data to compute $a_{\ell m}^{(d)}$. Normalize the $a_{\ell m}^{(r)}$ such that $a_{00}^{(r)} = a_{00}^{(d)}$, i.e.

$$\sum_{\text{pix}} \Omega_{\text{pix}} w_{\text{pix}} = \sum_{i=1}^{N_d} \omega_i^{(d)}, \quad (2.16)$$

where Ω_{pix} is the area of the pixel, w_{pix} is the value of the mask in that pixel and $\omega_i^{(d)}$ are the weights for the data points. If the mask is binary, i.e. $w_{\text{pix}} = 0$ or 1 , then the sum is $4\pi f_{\text{sky}}$, where f_{sky} is the observed sky fraction. Proceed as above to define W_ℓ , M , \mathcal{M} and \tilde{C}_b .

3 Examples

3.1 Golden spiral

Now we give two examples of the code and approach. First we give an example with a lot of small-scale power to show how taming aliasing can require very fine pixelization. For this example we consider a set of $N_{\text{pnt}} = 163\,840$ points laid out in a ‘‘golden spiral’’ (a.k.a. Fibonacci spiral) pattern between $\cos\theta = -1/2$ and $+1/2$:

$$\cos\theta_k = \frac{1}{2} \left[1 - \frac{2k+1}{N_{\text{pnt}}} \right], \quad \phi = \frac{2\pi k}{\varphi} \quad \text{for } 0 \leq k \leq N_{\text{pnt}} \quad (3.1)$$

where $\varphi = (1+\sqrt{5})/2$ is the golden ratio. This leads to an approximately uniform distribution of points in the covered region, and we perturb each point randomly by a very small amount so that the spacing isn’t completely regular. The left panel of Fig. 3 shows the $a_{\ell m}$ computed by direct harmonic transform, squared and averaged over m . Note that this point set has significant ‘power’ at high ℓ (analogous to a spike of power at the Nyquist frequency of a regular, Cartesian grid of points). We then generate a HEALPIX map of these points with various N_{side} from 256 to 4096 and use the `healpy` [31] ‘‘`anafast`’’ routine to compute the same average over m of $|a_{\ell m}|^2$. The right panels show the ratio of this estimate to the direct-sum approach (on a fine and expanded scale) up to $\ell = 600$. Note that the large amount of power at very high ℓ aliases into lower ℓ power unless N_{side} is quite large. All cases shown have less than one point per pixel (in the ‘survey region’) on average but still show significant excess power, even at ℓ of a few hundred. If we do not perturb the points in the spiral from their initial positions the aliasing effect is even more extreme and none of the pixelized approaches converge for even moderate ℓ .

3.2 Lognormal mock catalog

The golden spiral example has a much bluer spectrum (i.e., with more small-scale structure) than we commonly consider in cosmology, and thus serves to illustrate the impact of aliasing. Our second example is closer to the most common use-case for this sort of analysis in cosmology – a set of points with an angular power spectrum that is approximately a (red) power-law, plus shot noise. For this example we shall err in the other direction and enhance

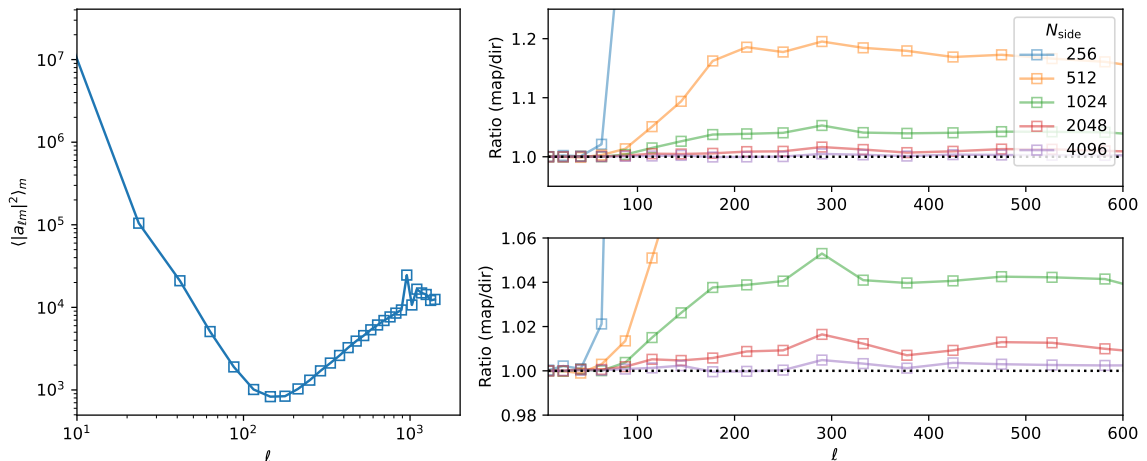


Figure 3. Comparison of direct harmonic transform and transform via pixelization for the Golden Spiral example (see text). Aliasing of small-scale power is particularly strong for this example. (Left) The average over m of $|a_{\ell m}|^2$ computed by direct transform, in bins of $\Delta\ell \propto \sqrt{\ell}$. (Top right) The ratio of the Healpix map-based power to the direct harmonic transform (i.e. the left panel) on an expanded y -axis scale for N_{side} from 256 to 4096. (Lower right) The same ratio on a smaller y -axis scale about unity.

the clustering signal over the shot noise (e.g. as for galaxies in a narrow shell in redshift where the clustering is not washed out by projection) to produce a “red” spectrum with $C_\ell \approx \ell^{-3/2}$ that is sample variance limited to $\ell \approx 10^3$. This should optimize the performance of the map-based approach. To create such a set of points we first generate a Gaussian random field over the whole sphere at $N_{\text{side}}=8192$ using the `synfast` routine of `healpy` [31]. We exponentiate this field and then populate each pixel with probability proportional to the pixel value, perturbing the positions of points away from the pixel centers by a very small amount. We then cut the set to $-0.4 < \cos\theta < 0.5$ and $0.2 \leq \phi < 5$, approximately one third of the sky, as an example of a “survey mask”, and we place 10 ‘masked’ regions of 2.5° radius randomly within the footprint as a mockup of a (very) bright star mask. The resulting data, along with randoms thrown evenly within the same range, forms our mock dataset – see Fig. 4. We analyze this dataset using the direct approach and the more traditional approach of first forming the overdensity field on a HEALPIX grid. In the absence of numerical effects both approaches should give the same pseudo- C_ℓ , mode coupling matrix, etc.

We find that the ‘direct’ approach is relatively robust to our analysis choices, and we get good results if the number of randoms exceeds the number of data by a factor of 20 – 50 \times . The precise number of randoms needed will depend upon the degree of structure in the mask. For the pixelized approach more care needs to be taken, especially when the overdensity field is generated by dividing a data map by a random map. In this case we found it necessary to ensure tens of randoms per pixel on average to get robust answers (since “division by R ” enhances the tails of the distribution). However, for a very large number of random points and sufficiently high resolution, the agreement between the two methods was very good, as illustrated in Fig. 5. We note that if percent level convergence is desired for this sample one needs more stringent requirements than advocated in e.g. ref. [24].




Figure 4. Map of the overdensity field from our lognormal mock catalog, in Mollweide projection, covering approximately 34% of the sky. The catalog contains about 1 million objects within the footprint shown, or 70 deg^{-2} . Given our fiducial clustering model the objects are signal dominated over most of the angular scales of interest to us, so the power spectrum is “red”.




Figure 5. A comparison of pseudo-spectra for the 1M mock galaxies shown in Fig. 4. (Left) The binned and mode-decoupled angular power spectrum produced by direct spherical harmonic transform. (Right) The ratio of the map-based estimates of the power spectrum to the direct (harmonic) spectrum for $N_{\text{side}} = 256, 512$ and 1024 ($N_{\text{pix}} \approx 0.8\text{M}, 3.1\text{M}, 12.5\text{M}$). The grey, horizontal band shows $\pm 1\%$ while the colored, dotted lines show the pixel window function (squared).

4 Conclusions

Harmonic analysis of the clustering of point sources on the sphere can be computationally demanding with the traditional approach of generating pixelized maps that are then transformed. We have presented an alternative, aimed specifically at analysis of cosmology surveys which cover significant fractions of the sky with the information mostly confined to large angular scales (due to observational and modeling limitations and the inherent complexities of galaxy formation that obscure cosmological signals). These include surveys such as DESI [1], Euclid [2, 3], SPHEREx [4], LSST on the Vera Rubin Observatory [5], Roman [6], and their cross-correlations with the Simons Observatory [7], the South Pole Observatory and CMB-S4 [8]. The method directly transforms the point set as a sum of spherical harmonics and we show that this can be efficiently performed numerically and is well-suited to GPU acceleration. In order to avoid having to divide by a mean density (e.g. estimated from a random catalog), we follow the approach more commonly adopted in 3D large-scale structure studies [25] which involves forming $\bar{n} \delta$ by subtracting a scaled random catalog from the data: $n_g - \alpha n_r$. If this field is transformed and squared one obtains the convolution of the angular power spectrum by a window function. By applying a “bandpower decoupling” step, as usually done in CMB analyses [13], one obtains a properly normalized, binned estimate of the angular power spectrum. This method is mathematically equivalent to the more common approach of forming an overdensity map that is then multiplied¹¹ by a mask or window map and transformed, but more straightforward and robust. We also discuss an alternative to “bandpower decoupling” which is again closer to the technique adopted in large-scale structure studies [25], where an overall normalization is applied without attempting to correct for the shape change induced by convolution with the mask. We show that both approaches lead to very similar pseudo-spectra (Fig. 1).

Section 2 introduces the algorithm and shows the connection to the “FKP method” [25] usually employed in redshift surveys and the “pseudo- C_ℓ ” approach [13] that is standard in CMB analyses. We show how the measured pseudo- C_ℓ is related to the ‘true’ spectrum on average and describe how our algorithm can be implemented (with some mathematical details relegated to appendices A.1, A.2 and B). Our GPU implementation is able to handle the SHT of $> 10^8$ points in about 30 seconds, and can be trivially parallelized.

Section 3 presents some examples of the approach and a comparison to the more traditional route. First we show how the aliasing that inevitably follows from pixelization can be challenging to handle if the spectrum under consideration has significant small-scale power. Fig. 3 presents an extreme example in which the points being analyzed are laid out in a “golden spiral” pattern. Fig. 5 shows a comparison of our method to the traditional approach for a spectrum closer to those considered in large-scale structure. In this example we have generated a lognormal mock catalog (Fig. 4) with an enhanced clustering component in order to compare with a “red” spectrum which is most favorable to the pixel or map-based approach. We see that for large enough N_{pix} the map-based approach converges to the direct SHT, though typically this requires $N_{\text{pix}} > N_{\text{pnt}}$ even for modest ℓ . In the oft-used approach where an overdensity field is generated by dividing the data map by the random map, the number of randoms should be $\gg 1$ per pixel so that the statistics of $1/R$ are not affected by the tails. This, combined with the need for $N_{\text{pix}} > N_{\text{pnt}}$, typically requires $N_{\text{rand}} \gg N_{\text{pnt}}$. Handling a large number of randoms and computing their pixel numbers can grow to be a

¹¹If the mask is either 0 or 1 this multiplication step can be skipped, but in the more usual case of an apodized mask it needs to be included.

non-trivial task, in addition to the $\mathcal{O}(N_{\text{pix}}^{3/2})$ scaling of the basic algorithm. Even when the use of randoms is avoided entirely, treating the ‘mask’ or ‘weight map’ linearly rather than dividing by it makes our algorithm more robust.

Finally we mention that giving the spherical harmonic transforms of the data ($a_{\ell m}$) and random ($w_{\ell m}$) fields provides a convenient complement to publishing a catalog of positions and weights. The coefficients provide a compact representation of the large-scale distribution, can be easily transformed between frames and allow simple computation of the 2-point and higher-point auto- and cross-spectra with other fields defined in harmonic space. Handling systematics weights or mode deprojection is also straightforward (see e.g. the discussion in ref. [14]). Harmonic-space apodization of the $a_{\ell m}$, e.g. as an alternative to higher-order charge assignment schemes, can also be applied to produce a band-limited function that can be analyzed in a map-based manner while leaving the large-scale properties of the signal unchanged (or changed in a known and controlled manner).

Our approach has several natural generalizations. First, a ‘flat-sky’ approximation to our method would simply entail replacing the $Y_{\ell m}(\theta, \phi)$ in equation (2.13) with $\exp[i\boldsymbol{\ell} \cdot \boldsymbol{x}]$, where \boldsymbol{x} is a vector living on the plane tangent to the sphere and $\boldsymbol{\ell}$ is the flat-sky analog of (ℓ, m) . Since exponentials can be computed very efficiently in terms of sines and cosines, it may be possible to obtain the Fourier coefficients exactly, even for a large number of points, by evaluating the exponential at all the points rather than interpolating. However, given the accuracy and speed of our full-sky implementation, this approximation is likely unnecessary. In fact, our method provides a straightforward and efficient way to upgrade analyses to the full-sky formalism.

Second, though we have presented the algorithm for describing the (scalar) density field, the generalization to fields with non-zero spin (such as cosmic shear) is relatively straightforward: rather than use $Y_{\ell m}$, one uses spin-spherical harmonics ($_{\pm 2}Y_{\ell m}$), which can also be efficiently computed using recurrence relations. One can also use this approach to implement a “radial-angular” (e.g. spherical Fourier-Bessel) analysis of 3D fields [32–37], such as galaxy redshift surveys. Such an approach would also be particularly appropriate for harmonic analyses of the Ly α forest, where sightlines are sparse but each sightline is well sampled in the radial direction, as an alternative to the 3D flux power spectrum [38, 39]. We choose not to pursue these directions here.

We make our `DirectSHT` code publicly available¹² and provide several tutorial notebooks.

Acknowledgments

We are grateful to Noah Sailer, Julien Carron, Zvonimir Vlah and David Alonso for useful comments on our manuscript. A.B.L. would also like to thank Marcelo Alvarez, Shamik Ghosh, Minas Karamanis, David Nabergoj and Stephen Bailey for discussions regarding the JAX library. M.W. is supported by the DOE. This research has made use of NASA’s Astrophysics Data System, the arXiv preprint server, the Python programming language and packages NUMPY, MATPLOTLIB, SCIPY, ASTROPY, JAX, HEALPY [31]. This research is supported by the Director, Office of Science, Office of High Energy Physics of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, and by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract.

¹²<https://github.com/martinjameswhite/directsht/tree/main>

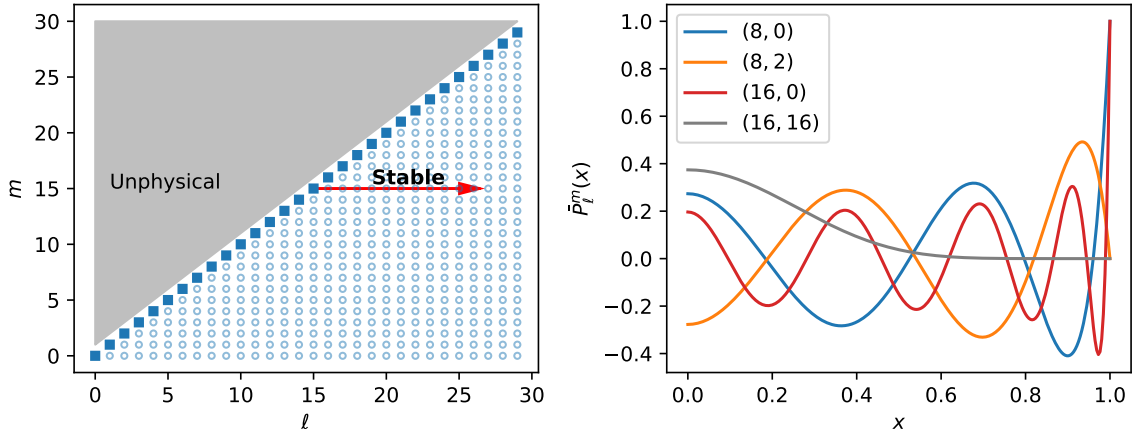


Figure 6. (Left) An illustration of the method for computing the (scaled) Legendre functions, $\bar{P}_\ell^m(x)$ via recurrence. \bar{P} is dominant on the degree (ℓ) so the recurrence relations are stable with increasing ℓ at fixed m . (Right) Some examples of $\bar{P}_\ell^m(x)$ for several (ℓ, m) pairs.

A Numerical and implementation details

A.1 Computing the spherical harmonics

In order to compute our harmonic coefficients we need to be able to (efficiently) compute $Y_{\ell m}(x = \cos \theta, \phi)$. It suffices to evaluate these for $\ell \geq 0$, $0 \leq m \leq \ell$ and $x \in [0, 1)$. The values for $m < 0$ and $x < 0$ can be computed by symmetry since $Y_{\ell, -m} = (-1)^m Y_{\ell m}^*$ while reflection in x incurs a phase $(-1)^{\ell-m}$. Let us decompose $Y_{\ell m}$ as

$$Y_{\ell m}(\theta, \phi) = \sqrt{\frac{(2\ell+1)}{4\pi}} \bar{P}_\ell^m(x = \cos \theta) e^{im\phi} \quad , \quad m \geq 0 \quad (\text{A.1})$$

where

$$\bar{P}_\ell^m(x) \equiv \sqrt{\frac{(\ell-m)!}{(\ell+m)!}} P_\ell^m(x) \quad , \quad m \geq 0 \quad (\text{A.2})$$

is a normalized version of the associated Legendre functions, with the normalization included for numerical convenience (since P_ℓ^m can become very large for $m \approx \ell$). The only computationally challenging piece is P_ℓ^m . However, these functions are smooth so their values and first derivatives can be computed on a grid of $x = \cos \theta$ and then interpolated using a Hermite spline. Both \bar{P} and its derivative can be computed very efficiently using recurrence relations¹³ that can be derived from their generating function [27, 29], so the resulting algorithm is fast and highly vectorizable.

The \bar{P} can be computed by recurrence starting from

$$\bar{P}_m^m = -\sqrt{1-x^2} \sqrt{1-(2m)^{-1}} \bar{P}_{m-1}^{m-1} \quad \text{and} \quad \bar{P}_{m+1}^m = \sqrt{2m+1} x \bar{P}_m^m, \quad (\text{A.3})$$

with the special cases $\bar{P}_0^0(x) = P_0(x) = 1$ and $\bar{P}_1^0(x) = P_1(x) = x$. Now, \bar{P} is dominant on the degree (ℓ) and minimal on the order (m), so for all other ℓ and m

$$\bar{P}_\ell^m = \sqrt{\frac{(\ell-m)}{(\ell+m)}} \left[\frac{2\ell-1}{\ell-m} x \bar{P}_{\ell-1}^m(x) - \sqrt{\frac{(\ell-m-1)}{(\ell+m-1)}} \frac{\ell+m-1}{\ell-m} \bar{P}_{\ell-2}^m(x) \right] \quad (\text{A.4})$$

¹³See e.g. <https://dlmf.nist.gov/14.10>

is stable with m increasing from 0 and ℓ increasing from $m + 2$ for each m (see Fig. 6). If high numerical efficiency is desired, there are reformulations of this recurrence that require fewer floating point operations [40].

Once the \bar{P}_ℓ^m are known, the derivatives can be computed from

$$\frac{d}{dx} \bar{P}_\ell^m(x) = -\frac{\ell x}{1-x^2} \bar{P}_\ell^m(x) + \frac{\ell+m}{1-x^2} \sqrt{\frac{(\ell-m)}{(\ell+m)}} \bar{P}_{\ell-1}^m(x) \quad (\text{A.5})$$

for $\ell > 0$, with $P'_0(x) = 0$ and $P'_1(x) = 1$.

A.2 Interpolation

The sum in Eq. (2.13) must be performed for both the data and random samples. It is not uncommon for this to involve 10^7 or 10^8 points, and while the problem is “embarrassingly parallel”, some care is still required in this evaluation. Since $Y_{\ell m}(\theta, \phi)$ is smooth, it is computationally advantageous to compute $Y_{\ell m}$ only at certain sample points and then interpolate between them. On the other hand, the ϕ -dependence can be included exactly at every ϕ_r position, as it is cheap to obtain. For notational convenience, let us define $y_{\ell m}(x) \equiv Y_{\ell m}^*(\theta, 0) = Y_{\ell m}(\theta, 0)$ with $x = \cos \theta$.

We found that when N_{pnt} is large the interpolation step can become a bottleneck for the calculation unless some care is taken. Thankfully, our analytic scheme to calculate $\bar{P}_{\ell m}(x)$ and its derivatives (appendix A.1) allows for fast and accurate interpolation using cubic Hermite splines, which are the unique piecewise-cubic polynomial with the correct function value and first derivative at the sample points. For a set of N_x points labeled x_j where our function is sampled, and defining $t_j \equiv (x - x_j)/(x_{j+1} - x_j)$, these splines are

$$S^{(\ell, m)}(x) = \sum_{j=0}^{N_x-1} [c_{j0}(x)y_{\ell m}(x_j) + c_{j1}(x)(x_{j+1} - x_j)y'_{\ell m}(x_j) + c_{j2}(x)y_{\ell m}(x_{j+1}) + c_{j3}(x)(x_{j+1} - x_j)y'_{\ell m}(x_{j+1})], \quad (\text{A.6})$$

with

$$c_{j0}(x) \equiv (2t_j+1)(1-t_j)^2, \quad c_{j1}(x) \equiv t_j(1-t_j)^2, \quad c_{j2}(x) \equiv t_j^2(3-2t_j), \quad c_{j3}(x) \equiv t_j^2(t_j-1) \quad (\text{A.7})$$

if $x_j \leq x < x_{j+1}$ and are zero otherwise. Though the splines $S^{(\ell, m)}(x)$ are unique to each (ℓ, m) pair, the dependence enters only through the values of $y_{\ell m}$ and its derivative at the sample points – in which the spline is linear. Both of these can be tabulated in advance following appendix A.1, thus eliminating the need to re-compute the interpolation weights at each (ℓ, m) , which is an expensive requirement of some other schemes. In addition, this being a ‘local’ interpolation method linear in y and y' , we are free to choose our sample points without regard for the ringing artefacts that plague ‘global’ methods or that require careful sampling (e.g. using Chebyshev nodes) to ensure a linear system can be stably inverted. In fact there is no linear system to solve for when computing the Hermite spline.

While the dependence on (ℓ, m) enters the splines only through the values of $y_{\ell m}$ and its derivative at the sample points, the $c_{j\alpha}$ depend upon x . This dichotomy is key to an efficient




Figure 7. The absolute error induced in our interpolation of $y_{\ell m}(x)$ as a function of $x = \cos \theta$. (Left) The absolute error for $\ell = 250$ and 500 with $m = 0$, for different numbers of spline points. (Right) The absolute error for $\ell = 500$ and $m = 125$ and 250 for different numbers of spline points (the error for $m = 500$ is much smaller). The characteristic size of $y_{\ell m}(x)$ is $\mathcal{O}(1)$ while for $N_x = 2\ell_{\max}$ (solid lines) the error is $\mathcal{O}(10^{-4})$ except near $x \approx 1$ where for $m = 0$ it grows rapidly (see text).

implementation. Interpolating with our cubic Hermite spline, we can write

$$\begin{aligned}
 a_{\ell m} &= \sum_{r=0}^{N_{\text{pnt}}} w_r e^{-im\phi_r} y_{\ell m}(x_r) \\
 &= \sum_{r=0}^{N_{\text{pnt}}} w_r e^{-im\phi_r} \sum_{j=0}^{N_x} [c_{j0}(x_r) y_{\ell m}(x_j) + c_{j1}(x_r) y'_{\ell m}(x_j) \\
 &\quad + c_{j2}(x_r) y_{\ell m}(x_{j+1}) + c_{j3}(x_r) y'_{\ell m}(x_{j+1})] \\
 &= \sum_{j=0}^{N_x} [v_{j0}^{(m)} y_{\ell m}(x_j) + v_{j1}^{(m)} y'_{\ell m}(x_j) + v_{j2}^{(m)} y_{\ell m}(x_{j+1}) + v_{j3}^{(m)} y'_{\ell m}(x_{j+1})] \quad (\text{A.8})
 \end{aligned}$$

where we have defined

$$v_{j\alpha}^{(m)} \equiv \sum_{r=0}^{N_{\text{pnt}}} w_r e^{-im\phi_r} c_{j\alpha}(x_r). \quad (\text{A.9})$$

These objects can be tabulated efficiently harnessing the fact that the $c_{j\alpha}(x_r)$ are zero outside of the spline region that x_r falls into: we therefore just need to accumulate contributions within bins bounded by the x_i sample points, rather than actually summing over all $N_{\text{pnt}} \sim 10^6 - 10^8$ elements at each j . Since it needs to happen at every m , this binning is one of the most computationally demanding steps of our pipeline. It is ideally carried on a GPU, the high throughput of which allows for efficient parallelization of the accumulation step.

Fig. 7 shows the interpolation error as a function of $x = \cos \theta$ for several values of (ℓ, m) and numbers of spline points, N_x . It appears that $N_x = N_\ell$ returns fractional errors of approximately 10^{-3} for $\ell = 500$ and $m = 0$ while $N_x = 2N_\ell$ does about an order of magnitude better. However, in practical situations involving many points, the implied error is actually less. For example, we find that the fractional error on W_ℓ computed from the random catalogs associated with the mock galaxy footprint described in §3 and shown in Fig. 4 is less than one percent for $N_x = N_\ell/2$, and almost undetectable for $N_x = N_\ell$. Furthermore, this fractional error occurs at $\ell \approx \ell_{\max}$, where W_ℓ is very small anyway. Thus, choosing $N_x \geq N_\ell$ should

lead to very small errors in practice. Fortunately, the code is sufficiently fast (Fig. 2) that the choice of N_x is not particularly critical.

Our interpolation scheme does have some drawbacks. Near $x \approx 1$, the largest of the $y_{\ell m}$ are those with $m = 0$. These become very steep as $x \rightarrow 1^-$ (Fig. 6). In fact $P'_\ell(1) = \ell(\ell+1)/2$, so the slope grows quadratically with the degree. This renders any interpolation scheme ineffective for large enough ℓ (see e.g. Fig. 7). For our situation we are able to simply limit $|x|$ such that we avoid the poles, and the rest of the function is very well behaved. For a fully general scheme, one approach would be to split off the points where $|x|$ exceeds some tolerance and then explicitly compute $y_{\ell m}(x)$ for those points either by recurrence as in appendix A.1 or via one of the many small-angle approximations, e.g. $P_\ell^m(x) \approx J_m(\ell\theta)$, coupled with the ability to rapidly compute special functions [29]. An alternative is to split the sample into regions near the equator and regions near both poles. The equatorial points are treated as above. The polar regions can be rotated to the equator, the SHT applied, and then the $a_{\ell m}$ rotated back to the original frame using Wigner matrices (which are block-diagonal in ℓ). We utilize the latter approach in our code for any points with $|x| > 3/4$. The Wigner matrices can also be computed from recurrence relations, see e.g. §3.3 of ref. [41].

B Wigner 3j symbols

In computing the mode-coupling matrix (Eq. 2.8) it is necessary to have an efficient means for evaluating the Wigner 3j symbols. This can also be done via recurrence [27] and we include the steps here for completeness. Let the three multipoles be ℓ_1 , ℓ_2 and ℓ_3 . For $m_1 = m_2 = m_3 = 0$, the 3j symbol is invariant under any permutation of the ℓ_i and vanishes if $L = \ell_1 + \ell_2 + \ell_3$ is odd. Non-vanishing values can be obtained from

$$\begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix} = \sqrt{\frac{(L - 2\ell_2 - 1)(L - 2\ell_3 + 2)}{(L - 2\ell_2)(L - 2\ell_3 + 1)}} \begin{pmatrix} \ell_1 & \ell_2 + 1 & \ell_3 - 1 \\ 0 & 0 & 0 \end{pmatrix}, \quad (\text{B.1})$$

with the terminal case

$$\begin{pmatrix} \ell & \ell & 0 \\ 0 & 0 & 0 \end{pmatrix} = \frac{(-1)^\ell}{\sqrt{2\ell + 1}}. \quad (\text{B.2})$$

This allows rapid evaluation of the required 3j symbols, and the calculation can be further optimized if previously computed values are cached.

References

- [1] DESI Collaboration, A. Aghamousa, J. Aguilar, et al., *The DESI Experiment Part I: Science, Targeting, and Survey Design*, *ArXiv e-prints* (Oct., 2016) [[arXiv:1611.00036](#)].
- [2] R. Laureijs, J. Amiaux, S. Arduini, et al., *Euclid definition study report*, 2011.
- [3] L. Amendola, S. Appleby, A. Avgoustidis, et al., *Cosmology and fundamental physics with the Euclid satellite*, *Living Reviews in Relativity* **21** (Apr., 2018) 2, [[arXiv:1606.00180](#)].
- [4] O. Doré, J. Bock, M. Ashby, et al., *Cosmology with the SPHEREX All-Sky Spectral Survey*, *arXiv e-prints* (Dec., 2014) arXiv:1412.4872, [[arXiv:1412.4872](#)].
- [5] **LSST Dark Energy Science** Collaboration, D. Alonso et al., *The LSST Dark Energy Science Collaboration (DESC) Science Requirements Document*, [arXiv:1809.01669](#).
- [6] D. Spergel, N. Gehrels, C. Baltay, et al., *Wide-Field Infrared Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report*, *arXiv e-prints* (Mar., 2015) arXiv:1503.03757, [[arXiv:1503.03757](#)].

- [7] N. Galitzki, A. Ali, K. S. Arnold, et al., *The Simons Observatory: instrument overview*, in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, vol. 10708 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 1070804, July, 2018. [arXiv:1808.04493](#).
- [8] K. N. Abazajian, P. Adshead, Z. Ahmed, et al., *CMB-S4 Science Book, First Edition, ArXiv e-prints* (Oct., 2016) [[arXiv:1610.02743](#)].
- [9] A. S. Chou, M. Soares-Santos, T. M. P. Tait, et al., *Snowmass Cosmic Frontier Report, arXiv e-prints* (Nov., 2022) arXiv:2211.09978, [[arXiv:2211.09978](#)].
- [10] R. A. Sunyaev and I. B. Zeldovich, *Microwave background radiation as a probe of the contemporary structure and history of the universe*, *Ann. Rev. Astron. & Astrophys.* **18** (Jan., 1980) 537–560.
- [11] J. E. Carlstrom, G. P. Holder, and E. D. Reese, *Cosmology with the Sunyaev-Zel’dovich Effect*, *Ann. Rev. Astron. & Astrophys.* **40** (Jan., 2002) 643–680, [[astro-ph/0208192](#)].
- [12] K. M. Górski, E. Hivon, A. J. Banday, et al., *HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere*, *Astrophys. J.* **622** (Apr, 2005) 759–771, [[astro-ph/0409513](#)].
- [13] E. Hivon, K. M. Górski, C. B. Netterfield, et al., *MASTER of the Cosmic Microwave Background Anisotropy Power Spectrum: A Fast Method for Statistical Analysis of Large and Complex Cosmic Microwave Background Data Sets*, *Astrophys. J.* **567** (Mar., 2002) 2–17, [[astro-ph/0105302](#)].
- [14] D. Alonso, J. Sanchez, A. Slosar, and LSST Dark Energy Science Collaboration, *A unified pseudo- C_ℓ framework*, *Mon. Not. R. Astron. Soc.* **484** (Apr., 2019) 4127–4151, [[arXiv:1809.09603](#)].
- [15] G. Chon, A. Challinor, S. Prunet, et al., *Fast estimation of polarization power spectra using correlation functions*, *Mon. Not. R. Astron. Soc.* **350** (May, 2004) 914–926, [[astro-ph/0303414](#)].
- [16] C. García-García, J. Ruiz-Zapatero, D. Alonso, et al., *The growth of density perturbations in the last 10 billion years from tomographic large-scale structure data*, *Journal of Cosmology and Astro-Particle Physics* **2021** (Oct., 2021) 030, [[arXiv:2105.12108](#)].
- [17] A. Nicola, D. Alonso, J. Sánchez, et al., *Tomographic galaxy clustering with the Subaru Hyper Suprime-Cam first year public data release*, *Journal of Cosmology and Astro-Particle Physics* **2020** (Mar., 2020) 044, [[arXiv:1912.08209](#)].
- [18] M. White, R. Zhou, J. DeRose, et al., *Cosmological constraints from the tomographic cross-correlation of DESI Luminous Red Galaxies and Planck CMB lensing*, *arXiv e-prints* (Nov., 2021) arXiv:2111.09898, [[arXiv:2111.09898](#)].
- [19] R. W. Hockney and J. W. Eastwood, *Computer simulation using particles*. Bristol: Hilger, 1988, 1988.
- [20] Y. P. Jing, *Correcting for the alias effect when measuring the power spectrum using FFT*, *Astrophys. J.* **620** (2005) 559–563, [[astro-ph/0409240](#)].
- [21] E. Sefusatti, M. Crocce, R. Scoccimarro, and H. Couchman, *Accurate Estimators of Correlation Functions in Fourier Space*, *Mon. Not. Roy. Astron. Soc.* **460** (2016), no. 4 3624–3636, [[arXiv:1512.07295](#)].
- [22] E. Hivon, “Pixel window functions.”
- [23] A. Krolewski, S. Ferraro, E. F. Schlafly, and M. White, *unWISE tomography of Planck CMB lensing*, *Journal of Cosmology and Astro-Particle Physics* **2020** (May, 2020) 047, [[arXiv:1909.07412](#)].

- [24] B. Leistedt, H. V. Peiris, D. J. Mortlock, et al., *Estimating the large-scale angular power spectrum in the presence of systematics: a case study of Sloan Digital Sky Survey quasars*, *Mon. Not. R. Astron. Soc.* **435** (Nov., 2013) 1857–1873, [[arXiv:1306.0005](#)].
- [25] H. A. Feldman, N. Kaiser, and J. A. Peacock, *Power-Spectrum Analysis of Three-dimensional Redshift Surveys*, *Astrophys. J.* **426** (May, 1994) 23, [[astro-ph/9304022](#)].
- [26] A. Baleato Lizancos and M. White, *The impact of anisotropic redshift distributions on angular clustering*, *J. Cosmol. Astropart. Phys.* **2023** (July, 2023) 044.
- [27] A. Edmonds, *Angular Momentum in Quantum Mechanics*. Investigations in Physics Series. Princeton University Press, 1996.
- [28] S. Singh, *improved Master for the LSS: fast and accurate analysis of the two-point power spectra and correlation functions*, *Mon. Not. Roy. Astron. Soc.* **508** (2021), no. 2 1632–1651, [[arXiv:2105.04548](#)].
- [29] N. Lebedev and R. Silverman, *Special Functions and Their Applications*. Dover Books on Mathematics. Dover Publications, 1972.
- [30] M. Reinecke, S. Belkner, and J. Carron, *Improved cosmic microwave background (de-)lensing using general spherical harmonic transforms*, *Astron. Astrophys.* **678** (Oct., 2023) A165, [[arXiv:2304.10431](#)].
- [31] A. Zonca, L. Singer, D. Lenz, et al., *healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in python*, *Journal of Open Source Software* **4** (Mar., 2019) 1298.
- [32] O. Lahav, K. B. Fisher, Y. Hoffman, et al., *Wiener Reconstruction of All-Sky Galaxy Surveys in Spherical Harmonics*, *Astrophys. J. Lett.* **423** (Mar., 1994) L93, [[astro-ph/9311059](#)].
- [33] K. B. Fisher, C. A. Scharf, and O. Lahav, *A spherical harmonic approach to redshift distortion and a measurement of $\Omega(0)$ from the 1.2-Jy IRAS Redshift Survey*, *Mon. Not. R. Astron. Soc.* **266** (Jan., 1994) 219, [[astro-ph/9309027](#)].
- [34] A. F. Heavens and A. N. Taylor, *A spherical harmonic analysis of redshift space*, *Mon. Not. R. Astron. Soc.* **275** (July, 1995) 483–497, [[astro-ph/9409027](#)].
- [35] E. Castorina and M. White, *The Zeldovich approximation and wide-angle redshift-space distortions*, *Mon. Not. R. Astron. Soc.* **479** (Sept., 2018) 741–752, [[arXiv:1803.08185](#)].
- [36] H. S. Grasshorn Gebhardt and O. Doré, *Fabulous code for spherical Fourier-Bessel decomposition*, *Phys. Rev. D* **104** (Dec., 2021) 123548, [[arXiv:2102.10079](#)].
- [37] Z. Gao, A. Raccañelli, and Z. Vlah, *Asymptotic connection between full- and flat-sky angular correlators*, *Phys. Rev. D* **108** (Aug., 2023) 043503, [[arXiv:2306.02993](#)].
- [38] A. Font-Ribera, P. McDonald, and A. Slosar, *How to estimate the 3D power spectrum of the Lyman- α forest*, *Journal of Cosmology and Astro-Particle Physics* **2018** (Jan., 2018) 003, [[arXiv:1710.11036](#)].
- [39] M. L. A. Karim, E. Armengaud, G. Mention, et al., *Measurement of the small-scale 3D Lyman- α forest power spectrum*, *arXiv e-prints* (Oct., 2023) arXiv:2310.09116, [[arXiv:2310.09116](#)].
- [40] K. ISHIOKA, *A new recurrence formula for efficient computation of spherical harmonic transform*, *Journal of the Meteorological Society of Japan. Ser. II* **96** (2018), no. 2 241–249.
- [41] M. A. Price and J. D. McEwen, *Differentiable and accelerated spherical harmonic and Wigner transforms*, *arXiv e-prints* (Nov., 2023) arXiv:2311.14670, [[arXiv:2311.14670](#)].