Invertible Filter Banks on the 2-Sphere

B.T. Thomas Yeo Wanmei Ou Polina Golland*

Abstract

Multiscale filtering methods, such as wavelets and steerable pyramids, have been widely used in processing and analysis of planar images and promise similar benefits in application to spherical images. While recent advances have extended some filtering methods to the sphere, many key challenges remain. This paper focuses on the self-invertibility property of filter banks, particularly desirable if images are modified in the wavelet domain. More specifically, we develop conditions for invertibility of spherical filter banks for both continuous and discrete convolution and illustrate how such conditions can be incorporated into the design of multiscale axis-symmetric wavelets.

1 Introduction

The theories of filter banks, wavelets and overcomplete wavelets, such as steerable pyramids, are well-established for the Euclidean spaces [5, 15] and have many applications in feature detection, compression and denoising of images. Extending the theory and methods of filtering to spherical images promises similar benefits in the fields that give rise to such images, including computer vision [4], computer graphics [13], astrophysics [19], and geophysics [17]. Our motivation for this work comes from the representation of brain cortical surfaces as functions on the sphere [7, 8]. Neurobiologists believe that most of higher cognitive abilities originate from the cerebral cortex, and that neurological growth or diseases significantly alter the structure of the cortex. Wavelet analysis therefore promises to help the detection and characterization of important cortical features and how they develop over time.

Similarly to the Euclidean case, filtering in the spherical domain involves decomposing the spherical image into correlation coefficients via convolution with a bank of analysis filters, resulting in the convolved outputs, as illustrated in Fig. 1 and 2. Once we move to the sphere, Fast Fourier Transform must be replaced with an alternative efficient method for computing convolutions. An original algorithm for axis-symmetric convolution kernels on the sphere was derived in [6], and was recently extended to arbitrary functions [18, 20]. The reconstructed image is obtained by adding the inverse convolutions of the filtering outputs with the synthesis filters. The filters' shapes and the relationship among the filters determine various properties of the filter bank. For example, in the Euclidean wavelets, the analysis filters are parameterized by dilation, while the steerable pyramids add parametrization through rotation. Invertible filter banks enable perfect reconstruction of the original signal and therefore provide an equivalent image representation in the wavelet domain. In self-invertible filter banks, the analysis and the corresponding synthesis filters are identical [1]. Self-invertibility is desirable for image manipulation in the wavelet domain, leading to an intuitive notion that a convolution coefficient corresponds to the contribution of the corresponding filter to the reconstructed signal. Without self-invertibility, the effects of

*This work is funded in part by the NIH grants R01-NS051826 and 1U54 EB005149. B.T. Thomas Yeo is funded by Agency for Science, Technology and Research, Singapore. Wanmei Ou is supported by the NSF graduate fellowship.

[1] In the Euclidean space, the definition of self-invertibility and convolution requires the synthesis filters be the reflection of the analysis filters [15], but the standard definition of convolution on the sphere leads to analysis filters that are identical to the synthesis filters.

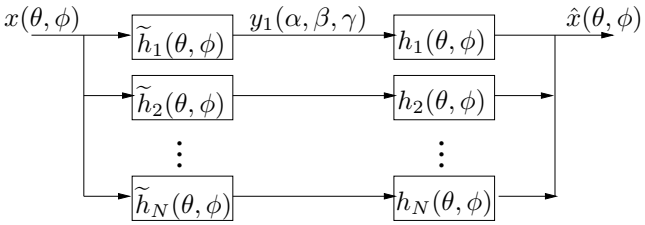

Figure 1: Continuous analysis and synthesis filter bank diagram

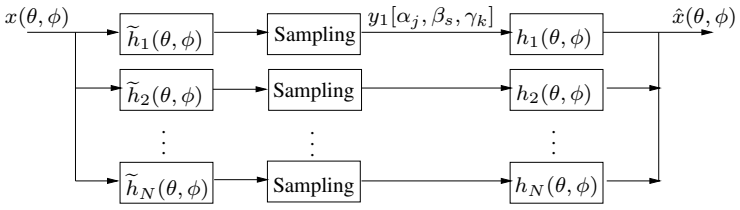

Figure 2: Discrete analysis and synthesis filter bank diagram

nonlinear processing of wavelet coefficients will propagate to locations and frequencies other than those which were used to compute the coefficients [15]. Our work extends the notion of self-invertibility to the sphere.

Recently, the general paradigm of linear filtering has been extended to the spherical domain. For example, the lifting scheme in [13, 14] adopts a non-parametric approach to computing wavelet decomposition of arbitrary meshes by generalizing the standard 2-scale relation of Euclidean wavelets, enabling a multi-scale representation of the original mesh (image) with excellent compression performance. However, the lifting wavelets are not overcomplete, i.e., exactly one wavelet coefficient is created per sample point, causing difficulties in designing filters for oriented feature detection. A similar problem in the Euclidean domain led to the invention of overcomplete wavelets, such as steerable pyramids [9, 15]. The group theoretic formulation of overcomplete spherical wavelets [1] is therefore a step forward. In particular, the stereographic projection of an admissible planar wavelet to the sphere is also admissible under the group theoretic framework, providing a straightforward framework for the design of analysis filters for specific features of interest, such as oriented edges [19]. Unfortunately, the synthesis filters are fully determined by the shapes of the analysis filters, which in general does not lead to self-invertibility. In contrast, we explicitly derive the conditions for self-invertibility and incorporate them into the filter design.

In the next section, we introduce the notation used throughout the paper. We present the conditions for the invertibility of filter banks for continuous and discrete spherical convolution in section 3. In section 4, we further specialize the invertibility conditions and present a procedure for generating self-invertible, multiscale filter banks on the sphere. We then illustrate the procedure for the case of axis-symmetric wavelets in section 5 and conclude with discussion of future research and outstanding challenges in the proposed framework.

2 Spherical Filtering. Notations and Definitions

Let $x(\theta,\phi) \in L^2(S^2)$ be a square-integrable function on the two-dimensional unit sphere. $\theta \in [0,\pi]$ is the co-latitude, which is the angle between the positive z-axis (north pole) and the vector corresponding to the point. $\phi \in [0,2\pi]$ is the longitude and is taken to be the angle between the positive x-axis and the

projection of the point onto the x-y plane. ϕ is undefined on the north and south poles.
The spherical harmonics $Y_l^m(\theta, \phi)$ form an orthonormal set of basis functions for $L^2(S^2)$: i.e.

$$x(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{|m| \leq l} x^{l,m} Y_l^m(\theta, \phi) \tag{1}$$

where $x^{l,m}$'s are the spherical harmonic coefficients of degree l and order m obtained by projecting x onto $Y_l^m(\theta, \phi)$:

$$x^{l,m} = \int_{S^2} x(\theta, \phi) Y_l^{m*}(\theta, \phi) d\Omega \tag{2}$$

where $d\Omega = \sin(\theta) d\theta d\phi$. We call $Y_l^m(\theta, \phi)$ a spherical harmonic of degree l and order m. For more details about spherical harmonics, see Appendix A.

We choose to parameterize rotations on the sphere by the 3-Euler angles, α, β, γ ($\alpha \in [0, 2\pi]$, $\beta \in [0, \pi]$, $\gamma \in [0, 2\pi]$). The rotation operator $D(\alpha, \beta, \gamma)$ first rotates the function by γ about the z-axis, then by β about the y-axis and finally by α about the z-axis. The direction of positive rotation follows the right-hand screw rule. The three angles specify an element of the rotation group $SO(3)$ and provide a natural parametrization of convolution on the sphere. The effects of rotation on the harmonic coefficients of a function is expressible in terms of the so called Wigner-D functions. The Wigner-D functions form an irreducible representation of the rotation group (see for example [20]).

On the plane, convolution is defined in terms of the inner product between two functions translated relative to each other, and is parameterized by the amount of translation. On the sphere, it is more natural to talk about rotation rather than translation, and therefore spherical convolution is parameterized by rotation. Given a spherical image, $x(\theta, \phi)$ and a spherical filter, $\tilde{h}(\theta, \phi)$, their spherical convolution

$$y(\alpha, \beta, \gamma) = \int_{S^2} [D(\alpha, \beta, \gamma)\tilde{h}]^*(\theta, \phi) x(\theta, \phi) d\Omega \tag{3}$$

is a function of $L^2(SO(3))$, rather than $L^2(S^2)$. If we consider the center (origin) of a spherical filter to be initially at the north pole ($\theta = 0$), then intuitively, $y(\alpha, \beta, \gamma)$ is obtained by first re-orienting the spherical filter by a rotation of γ about the z-axis (center still at north pole) and then bringing the center of the filter to the point (β, α) of the spherical image, and then performing an inner product between the image and filter. Therefore $y(\alpha, \beta, \gamma)$ is the inner product (correlation) of the rotated version of $\tilde{h}$ with x, or the projection coefficient of x onto $[D(\alpha, \beta, \gamma)\tilde{h}]$. For axis-symmetric filters $\tilde{h}(\theta, \phi) = \tilde{h}(\theta)$, the rotation by γ about z-axis has no effect, i.e., $y(\alpha, \beta, \gamma) = y(\alpha, \beta)$ is a spherical image parametrized by $\theta = \beta, \phi = \alpha$. Our definition of convolution is identical to that in [18, 20], although [20] calls it directional correlation. In [6], γ is integrated out, resulting in a spherical image.

The inverse convolution of a spherical filter $h(\theta, \phi)$ with $y(\alpha, \beta, \gamma) \in L^2(SO(3))$ produces a spherical image:

$$\widehat{x}_h(\theta, \phi) = \int_{SO(3)} [D(\alpha, \beta, \gamma)h](\theta, \phi) y(\alpha, \beta, \gamma) d\rho \tag{4}$$

where the integration is over the Euler angles: $d\rho = \sin(\beta) d\alpha \, d\beta \, d\gamma$. We can think of the inverse convolution the following way. The reconstructed value at a given (θ, ϕ) is obtained by summing (i.e. integrating) the contributions of inverse convolution filters, h, centered at (β, α) and oriented by γ, where the weights of the contributions are given by the convolution outputs (projection coefficients).

When using a filter bank of N analysis-synthesis filter pairs (Fig. 1) the reconstructed signal is obtained from the convolved outputs of the N analysis filters through the inverse convolution with the corresponding synthesis filters:

$$\widehat{x}(\theta, \phi) = \sum_{n=1}^{N} \int_{SO(3)} [D(\alpha, \beta, \gamma)h_n](\theta, \phi) y_n(\alpha, \beta, \gamma) d\rho \tag{5}$$

which is similar to the definition in [1], with integration over scale replaced by summation over the filter index.

In the Euclidean case, we typically discretize both the input images and the convolution outputs. When working on the sphere, we discretize the convolution outputs, but choose to keep the image domain continuous by working with spherical harmonic coefficients rather than sample values, because this allows us to exploit efficient algorithms for spherical convolution [18, 20]. Since no uniform sampling grid exists on the sphere, performing convolution completely by quadrature would be slow, because under each rotation of the filter relative to the spherical image, we will need to re-sample (or even re-interpolate) the filter or the image values. We note that continuous representation in the wavelet domain is possible through series of complex exponentials [18] or Wigner-D functions [20], but manipulating the series coefficients would be tantamount to simultaneously altering all the wavelet coefficients, defeating the purpose of the wavelet decomposition.

We therefore sample the output of the continuous convolution $y(\alpha, \beta, \gamma)$ to create its discrete counterpart $y(\alpha_j, \beta_s, \gamma_k)$, where $\{\alpha_j, \beta_s, \gamma_k\}$ define a particular sampling grid. The inverse discrete convolution definition

$$\widehat{x}_h(\theta, \phi) = \sum_{j=0}^{J-1} \sum_{s=0}^{S-1} \sum_{k=0}^{K-1} w_{j,s,k} [D(\alpha_j, \beta_s, \gamma_k)h](\theta, \phi) y(\alpha_j, \beta_s, \gamma_k) \tag{6}$$

includes sampling-dependent quadrature weights $w_{j,s,k}$, introduced so that the discrete inverse convolution converges to the continuous inverse convolution as the number of samples increases. This definition allows for an easy transfer of continuous filtering theory to its discrete analogue. In contrast with the planar case, $w_{j,s,k}$ are necessary because of the non-uniform measure on the Euler angle β.

The reconstruction of a signal from the sampled convolved outputs of N analysis filters through the inverse discrete convolution with N synthesis filters is naturally defined as

$$\widehat{x}(\theta, \phi) = \sum_{n=1}^{N} \sum_{j=0}^{J_n-1} \sum_{s=0}^{S_n-1} \sum_{k=0}^{K_n-1} w_{j,s,k,n} [D(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n})h_n](\theta, \phi) y_n(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) \tag{7}$$

The sampling grid and the quadrature weights now depend on n since different filters in the bank might use different sampling schemes.

3 Invertibility Conditions

In this section, we do not assume any relationship among the analysis and synthesis filters and provide general conditions for the invertibility of filter banks. The proofs are provided in the appendices.

Theorem 3.1 (Continuous Invertibility). *Let $\{\widetilde{h}_n, h_n\}_{n=1}^N$ be an analysis-synthesis filter bank. Then for any spherical image $x \in L^2(S^2)$ and its corresponding reconstructed image $\widehat{x}$,*

$$\hat{x}^{l,m} = x^{l,m} \text{ for all } (l,m) \quad \text{iff} \quad \sum_{n=1}^{N} \sum_{m'=-l}^{l} \left[h_n^{l,m'}\right] \left[\widetilde{h}_n^{l,m'}\right]^* = \frac{2l+1}{8\pi^2} \text{ for all } l \text{ s.t. } x^{l,m} \neq 0 \tag{8}$$

where $x^{l,m}$ and $\widehat{x}^{l,m}$ are the spherical harmonic coefficients of the input and reconstructed signals respectively. $\widetilde{h}_n^{l,m'}$ and $h_n^{l,m'}$ are the spherical harmonic coefficients of the n-th analysis and synthesis filters respectively. The proof is in Appendix B.

This theorem provides the necessary and sufficient condition for the invertibility of filter banks under continuous convolution. To draw analogies with the Euclidean case, we call

$$H_{\widetilde{h},h}(l) \quad = \quad \frac{8\pi^2}{2l+1} \sum_{n=1}^{N} \sum_{m'=-l}^{l} \left[h_n^{l,m'}\right] \left[\widetilde{h}_n^{l,m'}\right]^* \tag{9}$$

the frequency response of the analysis-synthesis filter bank. Theorem 3.1 implies that to guarantee perfect reconstructions of signals of bandwidth (maximal degree) L, the frequency response of the filter bank must be equal to 1 for all degrees up to L. On the plane, the frequency response is simply the

sum of products of the Fourier coefficients of the analysis and the synthesis filters. On the sphere, the frequency response contains an extra modulating factor that decreases with degree l.

We now define $L_{\widetilde{h}_n}$ ($O_{\widetilde{h}_n}$) and $L_{h_n}(O_{h_n})$ to be the highest non-zero harmonic degree (order) of $\widetilde{h}_n$ and h_n respectively. The following theorem specifies the sufficient, but not necessary, conditions for the invertibility of filter banks under discrete spherical convolution.

Theorem 3.2 (Discrete Invertibility). *Let $\{\widetilde{h}_n, h_n\}_{n=1}^N$ be a filter bank whose frequency response is equal to 1 up to degree $L < \infty$ and $O_{\widetilde{h}_n} < \infty$ and $O_{h_n} < \infty$. Let $\widetilde{L}_n = \min(L, L_{\widetilde{h}_n})$ and the sampling grid and the quadrature weights satisfy*

- $\alpha_{j,n} = \frac{2\pi j}{\widetilde{L}_n + L + 1}$ *for* $j = 0, 1, \ldots, (\widetilde{L}_n + L)$
- $\gamma_{k,n} = \frac{2\pi k}{O_{\widetilde{h}_n} + O_{h_n} + 1}$ *for* $k = 0, 1, \ldots, (O_{\widetilde{h}_n} + O_{h_n})$
- $w_{s,n}$ *and* $\beta_{s,n}$ *are the quadrature weights and knots such that*

$$\int_0^\pi d^l_{mm'}(\beta) d^{l'}_{mm'}(\beta) \sin(\beta) d\beta = \sum_{s=0}^{S_n - 1} w_{s,n} d^l_{mm'}(\beta_{s,n}) d^{l'}_{mm'}(\beta_{s,n}) \tag{10}$$

  for $l \leq L, l' \leq \widetilde{L}_n$, where $d^l_{mm'}(\beta)$ and $d^{l'}_{mm'}(\beta)$ are the Wigner-d functions.

- $w_{j,s,k,n} = \frac{4\pi^2 w_{s,n}}{(\widetilde{L}_n + L + 1)(O_{\widetilde{h}_n} + O_{h_n} + 1)}$

Then the filter bank is invertible for any spherical image $x \in L^2(S^2)$ with bandwidth L (i.e., $\widehat{x}^{l,m} = x^{l,m}$ for all $0 \leq l \leq L$) under the discrete convolution.

We note that invertibility is only guaranteed for degrees up to L rather than for all degrees as in Theorem 3.1. Additional constraints in this theorem ensure that the number of samples remain finite. The samples and the weights are picked such that the reconstruction obtained in Eq. (7) is the same as that in Eq. (5) up to degree L. The proof is found in Appendix C. In Appendix D, we also demonstrate two sets of quadrature weights and knots that satisfy the conditions of the theorem. The theorem is sufficient rather than necessary because there can exist other quadrature schemes that enable perfect reconstruction.

The measures corresponding to α and γ are constant, just like in the Euclidean space. We therefore assume uniform sampling for these parameters in our work. For discrete planar convolution, it is customary to have no weights (or rather, the weights are constants). On the sphere, however, the non-uniform measure on β, $\sin(\beta)d\beta$, makes sampling tricky. If we are simply interested in convergence, then setting $w_{j,s,k} = \frac{2\pi}{J}\frac{2\pi}{S}\sin(\beta_s)\frac{2\pi}{K}$ for uniform samples of α, β, γ corresponds to the Riemann sum of the integral. Theorem 3.2 states that better quadrature schemes exist that guarantee exact reconstruction up to a certain bandwidth.

The two theorems imply that if a filter bank with a finite maximal spherical harmonic order is invertible up to degree L under the continuous spherical convolution, it is also invertible up to degree L under the discrete spherical convolution. Since spherical harmonic coefficients of functions in $L^2(S^2)$ must necessarily decay to zero, we can reasonably assume that the constraint is satisfied if we represent the filter bank with a finite number of coefficients up to an arbitrary pre-specified precision. We will therefore focus on developing techniques for constructing invertible filter banks for continuous convolution.

Finally, we note that given a set of analysis filters, $\widetilde{h}_n$, there are in general multiple sets of synthesis filters that can achieve invertibility. A simple way is by defining the synthesis filters to be $h_n = L_\psi \widetilde{h}_n$, where,

$$[L_\psi \widetilde{h}_n]^{l,m} = \begin{cases} \frac{1}{H_{\widetilde{h},\widetilde{h}}(l)} \widetilde{h}^{l,m}_n & \text{for } H_{\widetilde{h},\widetilde{h}}(l) > 0 \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

$H_{\widetilde{h},\widetilde{h}}(l)$ is the frequency response defined in Eq. (9). $L_{\widetilde{h}}$ is a frequency modulating operator that normalizes the synthesis filters at each degree, such that the combined frequency response of the filter bank is one for all l with $H_{\widetilde{h},\widetilde{h}}(l) > 0$. The filter bank is therefore invertible for signals with non-zero spherical

harmonic coefficients corresponding to the non-zero values of $H_{\widetilde{h},\widetilde{h}}(l)$. This operation is similar to the frame operator in the continuous spherical wavelet transform of [1], where the counterpart of $H_{\widetilde{h},\widetilde{h}}(l)$ is given by $\frac{8\pi^2}{2l+1}\sum_{|m|\leq l}\int_0^\infty \frac{1}{a^3}\widetilde{h}_a^{l,m}da$, replacing the summation over n by the integration over the scale a, with measure $\frac{1}{a^3}da$. For the special case of the analysis filters being dilated versions of each other, our choice of the synthesis filters is a direct discretization of [1], albeit ignoring the measure of a. The discretization of the continuous wavelet transform in [1] was actually accomplished in [3]. In general, the synthesis filters are not related by dilation.

4 Self-Invertible Multi-scale Filter Bank

In multi-scale analysis, we construct the analysis filters through dilation and scaling of a particular template $\widetilde{h}(\theta,\phi)$, i.e.,

$$\widetilde{h}_k(\theta,\phi) = (\prod_{n=1}^{k} b_n) D_{a_k}\widetilde{h}(\theta,\phi) \tag{12}$$

where $b_n \geq 1$ and D_{a_k} is the nonlinear dilation operator, with larger n corresponding to smaller a (narrower filters).

The magnifying factors b_n's are necessary because of the conflict between self-invertibility and dilation while preserving norm. Theorem 3.1 implies that the sum of squares of the harmonic coefficients of a bank of self-invertible filters must increase linearly with degree. But stretching a function while preserving its norm shifts its harmonic coefficients to the left (harmonic degrees decrease) and magnifies them. These extra weights are less surprising once we note that the measure of scale is $\frac{1}{a^3}da$, i.e., wider filters are assigned smaller weights [1]. In the Euclidean case, with continuous image domain and discrete wavelet domain, enforcing orthonormality sidesteps this problem since convolution outputs of narrower filters are sampled more densely.

In this work, we adopt the stereographic dilation operator defined in [1, 3, 19], which involves stereographically projecting the function from the sphere onto the plane, performing the usual dilation operation on the plane and then projecting the resulting function back onto the sphere. This definition of the stereographic projection and dilation includes a normalization factor such that the inner product between functions is conserved:

$$[D_a f](\theta,\phi) = \frac{1}{a}\left(\frac{1+\tan^2\frac{\theta}{2}}{1+(\frac{1}{a}\tan\frac{\theta}{2})^2}\right) f(2\tan^{-1}(\frac{1}{a}\tan\frac{\theta}{2}),\phi) \tag{13}$$

Because of the nonlinear nature of stereographic dilation, extreme dilation of a spherical function will eventually lead to high frequencies. In practice, we will avoid working in that region, since the dilated filter no longer looks like the original filter. We also note that our techniques can easily handle other definitions of scale.

The approach commonly used with planar images of applying a constant filter to a subsampled image fails here because the sphere is periodic and compact, causing the effective size of the features to stay constant (relative to the filter) with subsampling. We also note that nonlinear dilation is necessary since the sphere is compact, hence dilating a spherical function by naively scaling the radial component of the spherical function, $f(\theta,\phi) \rightarrow f(\frac{\theta}{a},\phi)$, leads to undesired "wrap-around" effects.

Unlike the Fourier transform in the Euclidean space, there is no simple closed-form connection between the spherical harmonic coefficients of a function before and after stereographic dilation. Fortunately, like any reasonable dilation operator, stereographic dilation is distributive over addition. Suppose the template $\widetilde{h}$ is expressible as a linear combination of the basis functions $B^i(\theta,\phi)$, i.e., $\widetilde{h}(\theta,\phi) = \sum_{i=1}^{M} c_i B^i(\theta,\phi)$ (for the purpose of this paper, we will assume that $B^i(\theta,\phi)$ are spherical harmonics

and note that the technique is still applicable if a more suitable basis is found). Therefore,

$$[D_a\widetilde{h}]^{l,m} = \left[D_a \sum_{i=1}^{M} c_i B^i\right]^{l,m} = \sum_{i=1}^{M} c_i [D_a B^i]^{l,m} \tag{14}$$

yielding the harmonic coefficients of the analysis filter at another scale (before accounting for b_n's). This is useful since the invertibility condition in Eq. (8) was expressed in terms of the harmonic coefficients of the filters. We can therefore decide on a set of scales $\{a_n\}_{n=1}^{N}$ and create a table of spherical harmonic coefficients of the dilated basis functions. Eq. (14) allows us to determine the spherical harmonic coefficients of the dilated filters at each relative scale given c_i's and b_n's.

After fixing the set of basis functions $\{B^i\}$ and the set of scales $\{a_n\}$, we now pose an optimization problem to determine c_i's and b_n's. Similarly to the filter design in Euclidean space, the objective function should be application dependent, and could for example be a function of the frequency response. To guarantee self-invertibility, we assume that the analysis and synthesis filters are identical and optimize the cost function under the invertibility constraints of Eq. (8). Since we cannot have more constraints than variables, self-invertibility cannot be achieved for more degrees than the number of basis functions and scales.

In the experiments reported in the next section, we find the quadratic penalty method effective in solving this optimization problem with non-convex constraints. This involves incorporating the constraints in the objective function and solving the resulting unconstrained optimization problem using Newton's method. The procedure is repeated while increasing the weights of the constraints and using the solution corresponding to the previous weights as the starting point, until convergence to a local minimum of the original cost function.

5 Experiments

In this section, we perform the optimization problem formulated in the previous section. The experiments involve axis-symmetric filters only. Throughout this section, we will limit our set of basis functions to be the first hundred spherical harmonics of order 0, since spherical harmonic coefficients of axis-symmetric functions are zero for orders not equal 0.

We use the free software, S2kit [10], to create a table of the spherical harmonic coefficients of $D_a Y_l^0$ for $l = 0, \cdots, 99$. We find the first 600 order 0 harmonics of each dilated filter (a dilated axis-symmetric function remains axis-symmetric). As mentioned before, extreme dilation and shrinking of spherical harmonics can result in high frequencies. We find that for $a = 4$ and $a = 0.25$, $[D_a Y_{99}^0]^{599,0} < 10^{-7}$. We sample uniformly (on the logarithmic scale) the intermediate scales, resulting in a table for $a_n = 2^{-\frac{n}{3}}$ for $n = -6, -5, \cdots, 5, 6$, with $a = 1$ corresponding to the undilated spherical harmonics. We note that we might not be using the entire table for the experiments below.

Since our filters are axis-symmetric, we can use the fast spherical convolution [6] to compute forward convolution. We quote the results here for completeness:

$$[y(\alpha,\beta)]^{l,m} = \sqrt{\frac{4\pi}{(2l+1)}}\, x^{l,m} \widetilde{h}^{l,0*} \tag{15}$$

There is an offset of 2π from [6] because their definition of forward convolution includes integrating out γ. We will show in Appendix E that we can use almost the same formula to calculate the inverse convolution of $y(\alpha,\beta)$ with an axis-symmetric filter $h(\theta,\phi)$:

$$[\widehat{x}(\theta,\phi)]^{l,m} = 2\pi \sqrt{\frac{4\pi}{(2l+1)}}\, y^{l,m} h^{l,0} \tag{16}$$

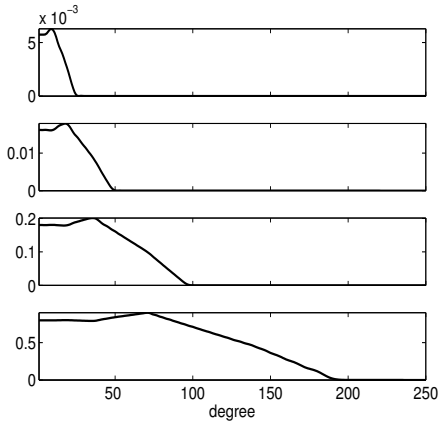

(a) Experiment 1

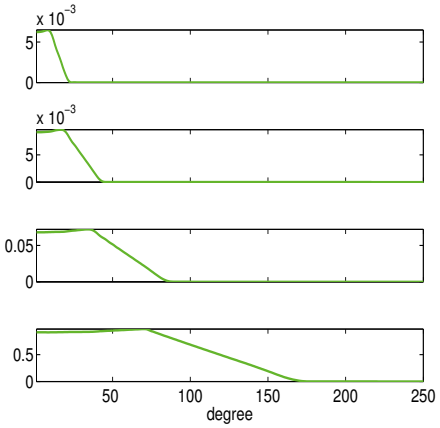

(b) Experiment 2

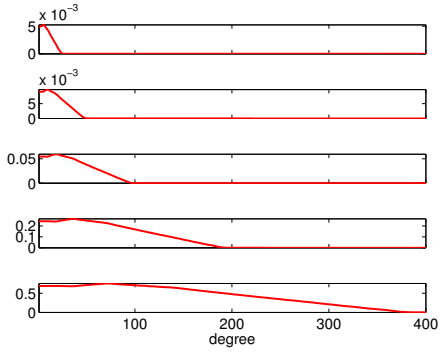

(c) Experiment 3

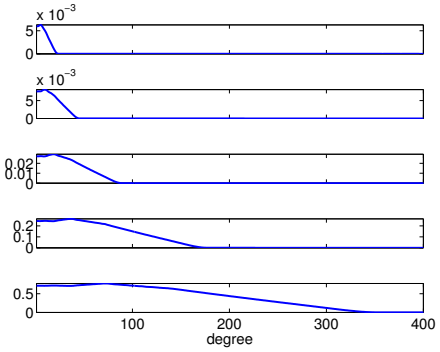

(d) Experiment 4

Figure 3: Frequency responses of lowpass filters found by optimization procedure in four experiments of section 5.1. Self-invertibility enforced from degree 0 to 69 for all four experiments. Objective function minimizes first derivatives of frequency responses at all scales. (a) Scales $a = \{4, 2, 1, 0.5\}$. (b) Scales $a = \{4, 2, 1, 0.5\}$, objective function also minimizes frequency response for degree 80 to 99 at $a = 1$. (c) Scales $a = \{4, 2, 1, 0.5, 0.25\}$. (d) Scales $a = \{4, 2, 1, 0.5, 0.25\}$, objective function also minimizes frequency response for degree 80 to 99 at $a = 1$. Note the different horizontal and vertical scales in the graphs

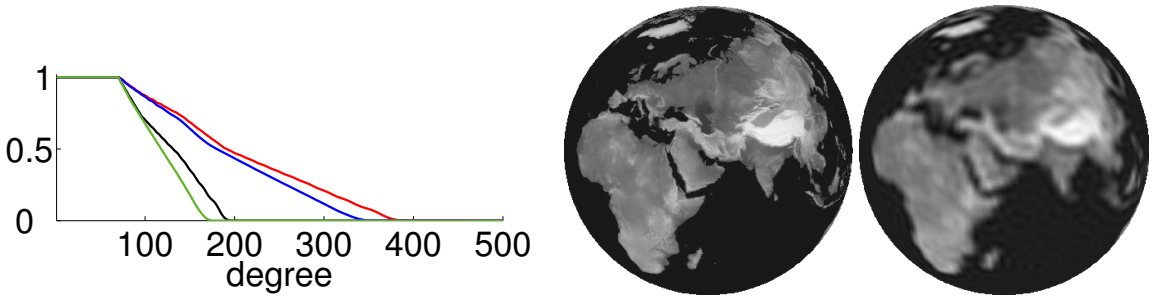

(a) Lowpass Combined Responses (b) BW = 540 (c) BW = 70

Figure 4: (a) Combined frequency responses of the lowpass filters found in the four experiments of section 5.1. Note that the colors correspond to those in Fig. 3 (b) World elevation map, bandwidth = 540. (c) Truncated world elevation map, bandwidth = 70.

5.1 Self-Invertible Lowpass Filters

In this subsection, we would like to derive an objective function that determines a set of lowpass filters. Ideally, we would like the frequency responses of the lowpass filters to be flat up to a certain degree and then smoothly drop down to zero. Since we only use the first 100 harmonics as our basis, the frequency response of $\widetilde{h}_{a=1}(\theta, \phi)$ will be zero for all degrees higher than 99. Setting our objective function to penalize first derivatives of the frequency responses of all the filters will thus result in an almost flat response followed by a gentle slope down to 0 at degree 100 for $a = 1$. The solution will not be degenerate (i.e. completely zero) because we enforce self-invertibility from degree 0 to 69 (70 constraints from Equation (8) for $l = 0, 1, \cdots, 69$).

We performed four experiments, using different sets of scales ($a = \{4, 2, 1, 1/2\}$ and $a = \{4, 2, 1, 1/2, 1/4\}$) and different penalty functions on the frequency responses of the filter bank for degrees between 80 and 100 (from not penalizing at all to imposing quadratic penalty on non-zero frequency responses). The frequency responses of the individual filters found in each experiment are shown in Fig. 3. Fig. 4(a) shows the combined frequency responses of the filter banks from each experiment. Note that an additional penalty term onside the invertibility range leads to a shaper cutoff (green and blue). The combined frequency responses of the filters from the 5-scales experiments also cover a much wider range of harmonics than the self-invertibility range (blue and red). Incorporating smaller scales a leads to frequency responses that are further beyond the region of invertibility. Including narrower filters thus only makes sense if we increase the range of invertibility (by increasing the number of basis functions).

We prefer the filters of the experiment whose combined response is plotted in green due to its sharpest cutoff. Fig. 5 displays the filters as spherical functions. The bright spot corresponds to the north pole. Since the filters are axis-symmetric, we can also plot the filters as functions of θ (see Fig. 6).

Fig. 7 illustrates application of these filters to a spherical image of an elevation map of the world of bandwidth 540 (Fig. 4(b)). The reconstruction performed using this filter bank is an accurate reconstruction of the original world elevation map truncated at bandwidth 70 (Fig. 4(c)), with the maximum absolute difference on the order of 10^{-7}. Similarly to the residual lowpass branch in the canonical wavelet analysis, in this experiment, we require a residual highpass filter that complements the frequency response for degrees 70 and above.

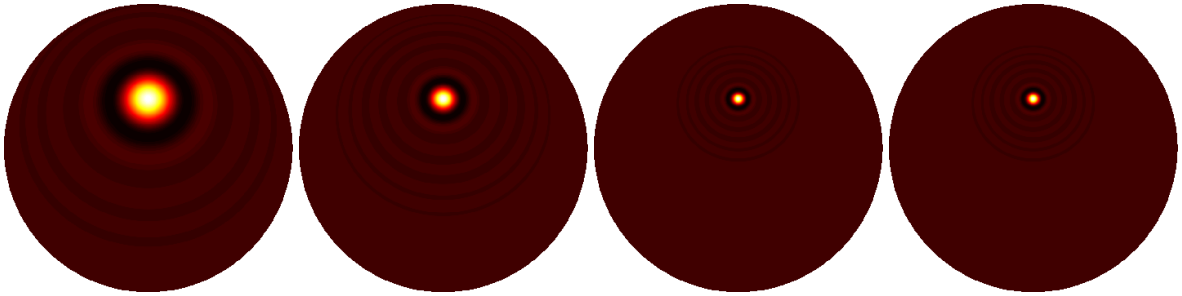

(a) a = 4 (b) a = 2 (c) a = 1 (d) a = 0.5

Figure 5: Plot of the lowpass filters found in Experiment 2 of section 5.1 (Fig. 3(b)) as spherical images.

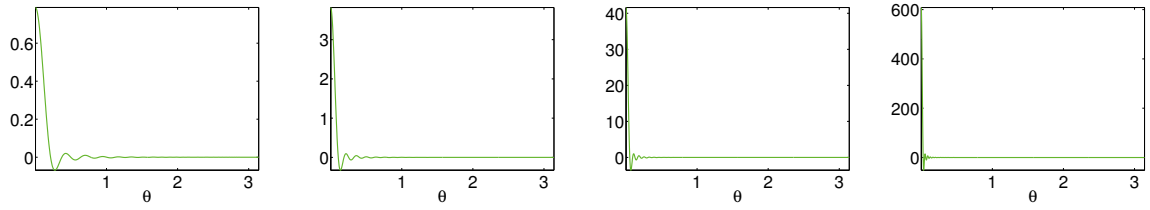

(a) a = 4 (b) a = 2 (c) a = 1 (d) a = 0.5

Figure 6: Plot of the lowpass filters found in Experiment 2 of section 5.1 (Fig. 3(b)) as functions of θ.

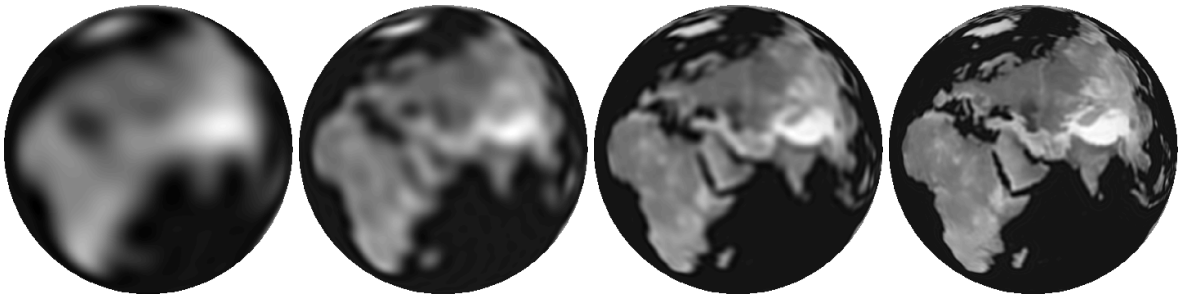

(a) a = 4 (b) a = 2 (c) a = 1 (d) a = 0.5

Figure 7: Convolution outputs between the set of lowpass filters in experiment 2 of section 5.1 (Fig. 3(b)) and world elevation map of bandwidth 540 degrees.

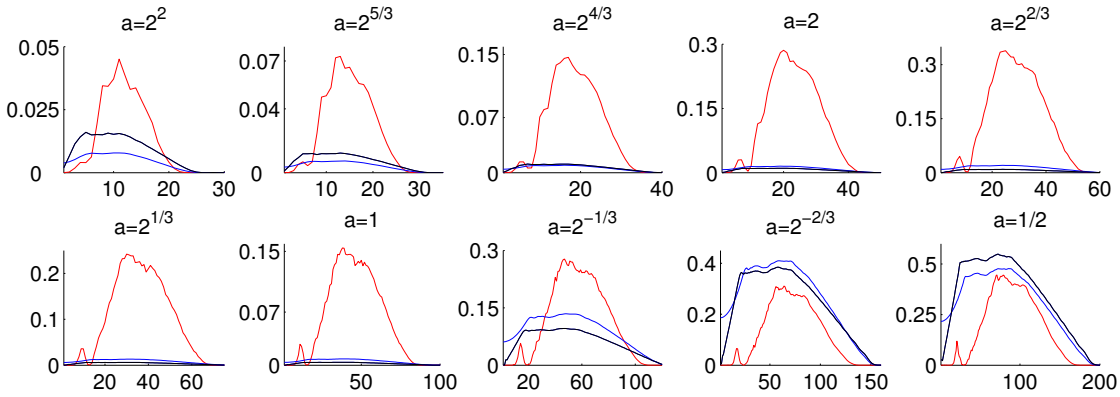

Figure 8: Frequency responses of the bandpass filters found by our optimization procedure in the three experiments of section 5.2. We consider the scales $a = \{2^{-\frac{k}{3}}\}$, for $k = -6, -5, \cdots, 2, 3$. Red: penalize frequency response that deviates from 0 for degree 80 to 99 and degree 0 to 19 at $a = 1$, while enforcing invertibility for degree 20 to 69. Blue: penalize frequency response that deviates from 0 for degree 80 to 99 and degree 0 to 19 at $a = 1$, while enforcing invertibility for degree 30 to 69. Black: penalize frequency response that deviates from 0 for degree 80 to 99 and degree 0 to 24 at $a = 1$, while enforcing invertibility from degree 24 to 69. Note the different horizontal and vertical scales in the graphs.

5.2 Self-Invertible Bandpass Filters

To achieve self-invertible bandpass filters, we use the same objective function as before. As noted previously, the frequency response at $a = 1$ will be zero for degree at least 100. If we also penalize the first few coefficients of the filter corresponding to $a = 1$, this means that the frequency response of the filter at $a = 1$ will be zeros at both ends. To satisfy the self-invertibility conditions, the solution will not be degenerate, i.e., it will not be completely zero, but must instead rise to a peak somewhere in the middle. The original conditions of minimizing first derivatives at all scales force the filters to be relatively smooth and reduce ringing. In addition, we relax invertibility for the first few degrees of the filter bank to make the optimization problem easier because of the additional penalties on the first few coefficients in the objective function.

We perform three experiments. In all three experiments, we consider the scales $a = \{2^{-\frac{k}{3}}\}$, for $k = -6, -5, \cdots, 2, 3$. We stop at 3 because incorporating narrower filters only lead to frequency responses that are way outside the region of invertibility. In the first experiment, we enforce invertibility for degree 20 to 69, while penalizing frequency response that deviates from 0 for degree 80 to 99 and degree 0 to 19 at $a = 1$. In the second experiment, we enforce invertibility for degree 30 to 69, while penalizing frequency response that deviates from 0 for degree 80 to 99 and degree 0 to 19 at $a = 1$. Finally, in the third experiment, we enforce invertibility from degree 24 to 69, while penalizing frequency response that deviates from 0 for degree 80 to 99 and degree 0 to 24 at $a = 1$.

The frequency responses of the filters found by the optimization procedures are shown in Fig. 8. Red corresponds to the first experiment, blue corresponds to the second experiment and black corresponds to the third experiment. As shown in the graphs, filters become narrower as we enforce invertibility over a wider range of degrees and as we increase the number of leading coefficients of $a = 1$ being penalized for not being 0. We prefer the filters from the first experiment (red) because their responses are the sharpest. Fig. 9 shows the results of applying the bandpass filters from the first experiment to the world elevation map of bandwidth 540. Only the convolution outputs corresponding to the scales, $a = \{4, 2, 1, 0.5\}$, are shown.

In this case, we require a residual lowpass filter in addition to a residual highpass filter to ensure that

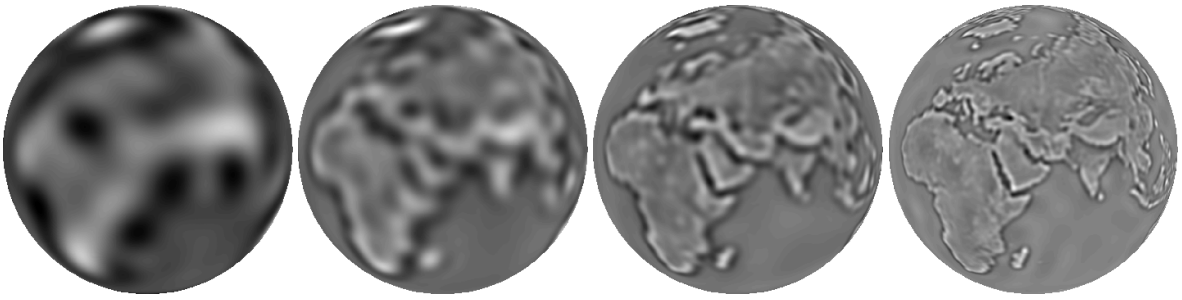

(a) a = 4 (b) a = 2 (c) a = 1 (d) a = 0.5

Figure 9: Convolution outputs between the bandpass filters (red lines in figure 8, first experiment of section 5.2) and world elevation map of bandwidth 540 degrees. Only scales $a = \{4, 2, 1, 0.5\}$ are shown.

the resultant response of the combined bandpass, lowpass and highpass filters is self-invertible up to a certain degree.

6 Conclusion

In this paper, we derive the necessary and sufficient conditions for a bank of filters to achieve perfect reconstruction in the continuous case. We then discretize the results using quadrature, but the conditions are now only sufficient rather than necessary. We present a procedure for obtaining self-invertible wavelets. We demonstrate the procedure with experiments, obtaining self-invertible lowpass and bandpass axis-symmetric wavelets.

Although stereographic dilation has many advantages [1], nonlinear dilation of functions on the sphere remains hard to work with. While we circumvent the problem by using the distributive property of stereographic dilation, the harmonic coefficients table can take up a substantial amount of space. More efficient methods are therefore needed. Perhaps it is also possible to formulate other definitions of dilation that fit better into the analytical framework.

So far, we have only demonstrated the use of our procedures for axis-symmetric wavelets. We plan to extend the approach to axis-asymmetric filters that detect oriented features at multiple scales. A promising approach is to develop an appropriate objective function for use in our optimization procedure. An alternate approach is to leverage upon the bandpass axis-symmetric wavelets which we have already found. Finally, we would like to reformulate the magnifying factors b_n's in a more natural fashion that reflects the measure of scale.

This paper introduces theoretic results on invertibility and represents a step towards creation of general self-invertible multiscale filter banks on the sphere. Steerable pyramids have been useful for feature detection and characterization in planar images, and we are optimistic that future work will lead to similar applications on the sphere.

7 Acknowledgements

The authors would like to thank Marshall Tappen for discussion on optimization procedures, Bill Freeman and Ted Adelson for discussion on self-invertibility, and Vivek Goyal for discussion on norm preserving dilations in Euclidean wavelets.

A Spherical Harmonics Basics

A.1 Associated Legendre Polynomials

For $m \geq 0$ and $|x| < 1$,

$$P_l^m(x) = \frac{(-1)^m}{2^l l!}(1-x^2)^{m/2}\frac{d^{l+m}}{dx^{l+m}}(x^2-1)^l \quad (17)$$

$$P_l^{-m}(x) = (-1)^m\frac{(l-m)!}{(l+m)!}P_l^m(x) \quad (18)$$

A.2 Spherical Harmonics

For a given degree $l \geq 0$ and order $|m| \leq l$, $0 \leq \theta \leq \pi$ and $0 \leq \phi \leq 2\pi$

$$Y_l^m(\theta,\phi) = \sqrt{\frac{2l+1}{4\pi}\frac{(l-m)!}{(l+m)!}}P_l^m(\cos\theta)e^{im\phi} \quad (19)$$

Therefore, for $l \geq m \geq 0$, we have,

$$Y_l^m(\theta,\phi) = \sqrt{\frac{2l+1}{4\pi}\frac{(l-m)!}{(l+m)!}}\frac{(-1)^m}{2^l l!}(1-\cos^2\theta)^{m/2}\frac{d^{l+m}}{dx^{l+m}}(x^2-1)^l\Big|_{x=\cos\theta}e^{im\phi} \quad (20)$$

$$Y_l^{-m}(\theta,\phi) = \sqrt{\frac{2l+1}{4\pi}\frac{(l-m)!}{(l+m)!}}\frac{1}{2^l l!}(1-\cos^2\theta)^{m/2}\frac{d^{l+m}}{dx^{l+m}}(x^2-1)^l\Big|_{x=\cos\theta}e^{-im\phi} \quad (21)$$

$$= (-1)^m\overline{Y_l^m(\theta,\phi)} \quad (22)$$

A.3 Rotation of Spherical Harmonics on the Sphere

Under a rotation, each spherical harmonic of degree l is transformed into a linear combination of only those spherical harmonics, Y_l^m, $-l \leq m \leq l$, of the same degree. In particular if we parametrize our rotation via the three euler angles, α, β, γ, and rotate our original function f, we have:

$$[D(\alpha,\beta,\gamma)f]^{l,m} = \sum_{m'=-l}^{l} D^l_{mm'}(\alpha,\beta,\gamma)f^{l,m'} \quad (23)$$

where $D^l_{mm'}(\alpha,\beta,\gamma)$ is the wigner-D function and is the amplitude of rotated state in m when the original unrotated state is m' (using quantum mechanics terminologies). In fact, only rotation about the y-axis mixes states (or orders). Hence we can further decompose $D^l_{mm'}(\alpha,\beta,\gamma)$ into

$$D^l_{mm'}(\alpha,\beta,\gamma) = e^{-im\alpha}d^l_{mm'}(\beta)e^{-im'\gamma} \quad (24)$$

where, $d^l_{mm'}(\beta)$ is a real function, the wigner-d function, whose formula is included here for completeness [12]:

$$d^l_{mm'}(\beta) = \sum_j(-1)^{j-m'+m}\frac{\sqrt{(l+m')!(l-m')!(l+m)!(l-m)!}}{(l+m'-j)!j!(l-j-m)!(j-m'+m)!} (\cos\frac{\beta}{2})^{2l-2j+m'-m}(\sin\frac{\beta}{2})^{2j-m'+m} \quad (25)$$

The sum is over all j such that none of the denominator terms with factorials are negative.

By the Peter-Weyl theorem on compact groups [20]:

$$\int_{SO(3)} D^l_{mn}(\rho)\overline{D^{l'}_{m'n'}}(\rho)d\rho = \frac{8\pi^2}{2l+1}\delta(l-l', m-m', n-n') \quad (26)$$

Another helpful identity is this:

$$Y_l^m(\beta,\alpha) = \sqrt{\frac{2l+1}{4\pi}}\left[D^l_{m0}(\alpha,\beta,\gamma)\right]^* = \sqrt{\frac{2l+1}{4\pi}}\left[e^{-im\alpha}d^l_{m0}(\beta)\right]^* \tag{27}$$

A.4 Axis-symmetric Functions

For axis-symmetric functions (independent of ϕ), we note that only the order 0 harmonics are non-zero.

B Proof of Continuous-Invertibility

In this appendix, we will prove Theorem 3.1 on continuous invertibility. We first note that using Parseval's Theorem, we can write

$$\begin{aligned} y_n(\alpha,\beta,\gamma) &= \int_{S^2}[D(\alpha,\beta,\gamma)\widetilde{h}_n]^*(\theta,\phi)x(\theta,\phi)d\Omega & (28)\\ &= \sum_{l'=0}^{\infty}\sum_{m'=-l'}^{l'}\left[\sum_{m''=-l'}^{l'} D^{l'}_{m',m''}(\alpha,\beta,\gamma)\widetilde{h}_n^{l',m''}\right]^* x^{l',m'} & (29)\end{aligned}$$

Therefore, in reconstruction,

$$\begin{aligned} \widehat{x}(\theta,\phi) &= \sum_{n=1}^{N}\int_{SO(3)}[D(\alpha,\beta,\gamma)h_n](\theta,\phi)y_n(\alpha,\beta,\gamma)d\rho & (30)\\ &= \sum_{n=1}^{N}\int_{SO(3)}[D(\alpha,\beta,\gamma)h_n](\theta,\phi)\sum_{l'=0}^{\infty}\sum_{m'=-l'}^{l'}\left[\sum_{m''=-l'}^{l'} D^{l'}_{m',m''}(\alpha,\beta,\gamma)\widetilde{h}_n^{l',m''}\right]^* x^{l',m'}d\rho & (31)\end{aligned}$$

Projecting $\widehat{x}(\theta,\phi)$ onto the spherical harmonics basis and letting $\widehat{x}^{l,m}$ be the spherical harmonic coefficient of degree l and order m of $\widehat{x}$, we have

$$\begin{aligned} \widehat{x}^{l,m} &= \int_{S^2}\widehat{x}(\theta,\phi)\overline{Y}_l^m(\theta,\phi)d\Omega & (32)\\ &= \sum_{n=1}^{N}\int_{SO(3)}\left[\int_{S^2}[D(\alpha,\beta,\gamma)h_n](\theta,\phi)\overline{Y}_l^m(\theta,\phi)d\Omega\right]\sum_{l'=0}^{\infty}\sum_{m'=-l'}^{l'}\left[\sum_{m''=-l'}^{l'} D^{l'}_{m',m''}(\alpha,\beta,\gamma)\widetilde{h}_n^{l',m''}\right]^* x^{l',m'}d\rho & (33)\\ &= \sum_{n=1}^{N}\int_{SO(3)}\left[\sum_{m'''=-l}^{l} h_n^{l,m'''}D^l_{m,m'''}(\alpha,\beta,\gamma)\right]\sum_{l'=0}^{\infty}\sum_{m'=-l'}^{l'}\left[\sum_{m''=-l'}^{l'} D^{l'}_{m',m''}(\alpha,\beta,\gamma)\widetilde{h}_n^{l',m''}\right]^* x^{l',m'}d\rho & (34)\\ &= \sum_{n=1}^{N}\sum_{m'''=-l}^{l}\left[h_n^{l,m'''}\right]\sum_{l'=0}^{\infty}\sum_{m'=-l'}^{l'}x^{l',m'}\sum_{m''=-l'}^{l'}\left[\widetilde{h}_n^{l',m''}\right]^*\int_{SO(3)}D^l_{m,m'''}(\alpha,\beta,\gamma)D^{l'}_{m',m''}(\alpha,\beta,\gamma)^*d\rho & (35)\\ &= x^{l,m}\frac{8\pi^2}{2l+1}\sum_{n=1}^{N}\sum_{m'''=-l}^{l}\left[h_n^{l,m'''}\right]\left[\widetilde{h}_n^{l,m'''}\right]^* & (36)\end{aligned}$$

where we have used the Peter-Weyl theorem on compact group (Eq. (26)) in the last equality.

Hence, if $x^{l,m}\neq 0$ and $\sum_{n=1}^{N}\sum_{m'''=-l}^{l}\left[h_n^{l,m'''}\right]\left[\widetilde{h}_n^{l,m'''}\right]^* = \frac{2l+1}{8\pi^2}$, then $\widehat{x}^{l,m} = x^{l,m}$. If $x^{l,m}=0$, then $\widehat{x}^{l,m}$ is also zero, and hence we also get $\widehat{x}^{l,m}=x^{l,m}$. Because the equations above go both ways, we have thus obtained our necessary and sufficient conditions for invertibility.

C Proof of Discrete-Invertibility

In this appendix, we will prove Theorem 3.2 on discrete invertibility. Remember from our definition of discrete convolution that $y_n(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n})$ are samples of $y_n(\alpha, \beta, \gamma)$. Hence, from Eq. (29), and noting that the maximum degree of x to be L and finite, we get,

$$y_n(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) = \sum_{l'=0}^{\widetilde{L}_n} \sum_{m'=-l'}^{l'} \left[\sum_{m''=-l'}^{l'} D^{l'}_{m',m''}(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) \widetilde{h}_n^{l',m''} \right]^* x^{l',m'} \quad (37)$$

Also, letting $w_{j,s,k,n} = \frac{4\pi^2 w_{s,n}}{J_n K_n}$ (as required by the theorem), we have

$$\widehat{x}_n(\theta,\phi) = \sum_{j=0}^{J_n-1} \sum_{s=0}^{S_n-1} \sum_{k=0}^{K_n-1} w_{j,s,k,n} [D(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) h_n](\theta,\phi) y_n(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) \quad (38)$$

$$= \sum_{j=0}^{J_n-1} \sum_{s=0}^{S_n-1} \sum_{k=0}^{K_n-1} \frac{4\pi^2 w_{s,n}}{J_n K_n} [D(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) h_n](\theta,\phi) \quad (39)$$

$$\sum_{l'=0}^{\widetilde{L}_n} \sum_{m'=-l'}^{l'} \left[\sum_{m''=-l'}^{l'} D^{l'}_{m',m''}(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) \widetilde{h}_n^{l',m''} \right]^* x^{l',m'}$$

Projecting $\widehat{x}_n(\theta,\phi)$ onto the spherical harmonics, for $l \leq L$ (and hence $|m| \leq L$), we have

$$\widehat{x}_n^{l,m} = \int_{S^2} \widehat{x}_n(\theta,\phi) \overline{Y}_l^m(\theta,\phi) d\Omega$$

$$= \sum_{j=0}^{J_n-1} \sum_{s=0}^{S_n-1} \sum_{k=0}^{K_n-1} \frac{4\pi^2 w_{s,n}}{J_n K_n} \left[\sum_{m'''=-l}^{l} D^{l}_{m,m'''}(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) h_n^{l,m'''} \right] \quad (40)$$

$$\sum_{l'=0}^{\widetilde{L}_n} \sum_{m'=-l'}^{l'} \left[\sum_{m''=-l'}^{l'} \widetilde{h}_n^{l',m''} D^{l'}_{m',m''}(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) \right]^* x^{l',m'}$$

$$= \frac{4\pi^2}{J_n K_n} \sum_{m'''=-l}^{l} h_n^{l,m'''} \sum_{l'=0}^{\widetilde{L}_n} \sum_{m'=-l'}^{l'} x^{l',m'} \sum_{m''=-l'}^{l'} \left[\widetilde{h}_n^{l',m''} \right]^* \quad (41)$$

$$\sum_{j=0}^{J_n-1} \sum_{s=0}^{S_n-1} \sum_{k=0}^{K_n-1} w_{s,n} D^{l}_{m,m'''}(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) D^{l'*}_{m',m''}(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n})$$

where the second equality is obtained by performing the projection, and in the last equality we arrange the terms so that they look like the setup for Peter-Weyl Theorem, except we have summations instead of integrals. Noting that we can write $D^l_{mm'}(\alpha,\beta,\gamma) = e^{-im\alpha} d^l_{mm'}(\beta) e^{-im'\gamma}$ (Eq. (24)), the last part of Eq. (41) becomes

$$\sum_{j=0}^{J_n-1} \sum_{s=0}^{S_n-1} \sum_{k=0}^{K_n-1} w_{s,n} D^{l}_{m,m'''}(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) D^{l'*}_{m',m''}(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) \quad (42)$$

$$= \sum_{j=0}^{J_n-1} \sum_{s=0}^{S_n-1} \sum_{k=0}^{K_n-1} w_{s,n} e^{-im\alpha_{j,n}} d^{l}_{m,m'''}(\beta_{s,n}) e^{-im'''\gamma_{k,n}} e^{im'\alpha_{j,n}} d^{l'}_{m',m''}(\beta_{s,n}) e^{im''\gamma_{k,n}} \quad (43)$$

$$= \sum_{s=0}^{S_n-1} w_{s,n} d^{l}_{m,m'''}(\beta_{s,n}) d^{l'}_{m',m''}(\beta_{s,n}) \left[\sum_{j=0}^{J_n-1} e^{i(m'-m)\alpha_{j,n}} \right] \left[\sum_{k=0}^{K_n-1} e^{i(m''-m''')\gamma_{k,n}} \right] \quad (44)$$

We note that $|m'| \leq l' \leq \widetilde{L}_n$ and $|m| \leq L$, and hence $m' - m$ has the finite range $[-\widetilde{L}_n - L, \widetilde{L}_n + L]$. Noting the range of $m' - m$ and since $\alpha_{j,n} = \frac{2\pi j}{\widetilde{L}_n + L + 1}$, where $j = 0, 1, \cdots, \widetilde{L}_n + L$, we can conclude that

$$\sum_{j=0}^{J_n-1} e^{i(m'-m)\alpha_{j,n}} = \sum_{j=0}^{\widetilde{L}_n+L} e^{i(m'-m)\frac{2\pi j}{\widetilde{L}_n+L+1}} = \begin{cases} \widetilde{L}_n + L + 1 & \text{if } (m'-m) = 0 \\ 0 & \text{otherwise} \end{cases} \tag{45}$$

Similarly, looking at where m'' and m''' appear in Eq. (40), we can conclude that $|m''| \leq O_{\tilde{h}_n}$ and $|m'''| \leq O_{h_n}$. Hence, using the same reasoning, we get

$$\sum_{k=0}^{K_n-1} e^{i(m''-m''')\gamma_{k,n}} = \sum_{k=0}^{O_{\widetilde{h}_n}+O_{h_n}} e^{i(m''-m''')\frac{2\pi k}{O_{\widetilde{h}_n}+O_{h_n}+1}} = \begin{cases} O_{\widetilde{h}_n} + O_{h_n} + 1 & \text{if } (m''-m''') = 0 \\ 0 & \text{otherwise} \end{cases} \tag{46}$$

Therefore Eq. (44) becomes

$$\sum_{s=0}^{S_n-1} w_{s,n} d^l_{m,m''}(\beta_{s,n}) d^{l'}_{m',m''}(\beta_{s,n}) \left[\sum_{j=0}^{J_n-1} e^{i(m'-m)\alpha_{j,n}}\right] \left[\sum_{k=0}^{K_n-1} e^{i(m''-m''')\gamma_{k,n}}\right] \tag{47}$$

$$= J_n K_n \sum_{s=0}^{S_n-1} w_{s,n} d^l_{m,m''}(\beta_{s,n}) d^{l'}_{m',m''}(\beta_{s,n}) \delta(m-m') \delta(m''-m''') \tag{48}$$

$$= J_n K_n \sum_{s=0}^{S_n-1} w_{s,n} d^l_{m,m''}(\beta_{s,n}) d^{l'}_{m,m''}(\beta_{s,n}) \delta(m-m') \delta(m''-m''') \tag{49}$$

$$= \frac{2J_nK_n}{2l+1} \delta(l-l') \delta(m-m') \delta(m''-m''') \tag{50}$$

where in the second equality, $m''' \to m''$, $m' \to m$ because of the delta functions, and the last equality was obtained using the assumption that $w_{s,n}$ and $\beta_{s,n}$ are the quadrature weights and knots of the integral $\int_0^\pi d^l_{mm''}(\beta) d^{l'}_{mm''}(\beta) \sin(\beta) d\beta$ and hence we can use the Peter-Weyl Theorem. We can now substitute equation (50) back into equation (41), and we get

$$\widehat{x}^{l,m}_n$$

$$= \frac{4\pi^2}{J_nK_n} \sum_{m'''=-l}^{l} h^{l,m'''}_n \sum_{l'=0}^{\widetilde{L}_n} \sum_{m'=-l'}^{l'} x^{l',m'} \sum_{m''=-l'}^{l'} \left[\widetilde{h}^{l',m''}_n\right]^* \frac{2J_nK_n}{2l+1} \delta(l-l') \delta(m-m') \delta(m''-m''') \tag{51}$$

$$= \frac{8\pi^2}{2l+1} x^{l,m} \sum_{m''=-l}^{l} h^{l,m''}_n \widetilde{h}^{l,m''*}_n \tag{52}$$

Noting that $\widehat{x}(\theta, \phi) = \sum_{n=1}^N \widehat{x}_n(\theta, \phi)$, we have

$$\widehat{x}^{l,m} = \sum_{n=1}^N \widehat{x}^{l,m}_n = \frac{8\pi^2}{2l+1} x^{l,m} \sum_{n=1}^N \sum_{m''=-l}^{l} h^{l,m''}_n \widetilde{h}^{l,m''*}_n = x^{l,m} \quad \text{for all } 0 \leq l \leq L \tag{53}$$

We note that subtle variations of the theorem can be obtained, for example by increasing the maximum degree of x to be greater than L or increasing the number of samples on α and β or both.

D Quadrature Rules

In this appendix, we will derive two different quadrature rules that satisfy the following: $w_{s,n}$ and $\beta_{s,n}$ are the quadrature weights and knots such that for $l \leq L, l' \leq \widetilde{L}_n$,

$$\int_0^\pi f(\beta) \sin\beta d\beta = \int_0^\pi d^l_{mm'}(\beta) d^{l'}_{mm'}(\beta) \sin(\beta) d\beta = \sum_{s=0}^{S_n-1} w_{s,n} d^l_{mm'}(\beta_{s,n}) d^{l'}_{mm'}(\beta_{s,n}) \tag{54}$$

D.1 Quadrature Rule (1)

We first note that $f(\beta)$ consists of a linear combination of even powers of $\cos(\beta/2)$ and $\sin(\beta/2)$ (see Eq. (25)). Hence, if we make the substitution $u = \sin(\beta/2)$, and noting that $f(u)$ will now be a polynomial with maximum degree, $Q = 2(l + l') \leq 2(L + \widetilde{L}_n)$, we get

$$\int_0^\pi f(\beta) \sin\beta d\beta = 2\int_0^\pi f(\beta)\sin(\beta/2)\cos(\beta/2)d\beta = 4\int_0^1 f(u)udu \tag{55}$$

Now, making the substitution, $v = 2u - 1$, we have

$$\int_0^\pi f(\beta) \sin\beta d\beta = 2\int_{-1}^1 f(\frac{v+1}{2})\frac{v+1}{2}dv = 2\sum_{k=0}^{N-1} r_k \frac{v_k+1}{2} f(\frac{v_k+1}{2}) \tag{56}$$

where r_k are defined to be the weights of the Gauss-Legendre quadrature on the interval $[-1, 1]$, and v_k corresponds to the sampling knots. The weights and abscissas can be found by standard algorithms (see for example [11]). Note that because the integrand has highest polynomial power $Q + 1$, if $N \geq \frac{Q}{2} + 1$ or $N \geq \text{roof}(\frac{Q}{2}) + 1$, then $2N - 1 \geq Q + 1$ and the quadrature formula is exact.

From the substitution above, we have $\sin\frac{\beta_k}{2} = u_k = \frac{v_k+1}{2}$ or $\beta_k = 2\sin^{-1}(\frac{v_k+1}{2})$. Note that we can assume the range of $\sin^{-1}$ to be $[0, \frac{\pi}{2}]$ since $\cos$ can be expressed in terms of $\sin$ and the even powers of $\cos$ in f will take care of the rest. In conclusion, for $N = \text{roof}(\frac{Q}{2}) + 1$, we have $\int_0^\pi d^l_{mm'}(\beta)d^{l'}_{mm'}(\beta)\sin\beta d\beta = \sum_{s=0}^{N-1} w_s f(\beta_s)$, where $w_s = r_s(v_s + 1)$ and $\beta_s = 2\sin^{-1}(\frac{v_s+1}{2})$.

D.2 Quadrature Rule 2

We will derive another rule in this section, using the technique shown in [6]. But first we need to obtain the fourier series formula for the square wave, $SQ(u)$, which is defined to be periodic from $-\pi$ to π and is $+1$ from $-\pi$ to $-\pi/2$ and 0 to $\pi/2$ and -1 otherwise.

$$\begin{aligned}
\frac{1}{2\pi}\int_{-\pi}^{\pi} SQ(u)e^{-iku}du &= \frac{1}{2\pi}\left[\int_{-\pi}^{-\pi/2} e^{-iku}du - \int_{-\pi/2}^{0} e^{-iku}du + \int_0^{\pi/2} e^{-iku}du - \int_{\pi/2}^{\pi} e^{-iku}du\right] && (57)\\
&= \frac{i}{2\pi k}\left[e^{-iku}\Big|_{-\pi}^{-\pi/2} - e^{-iku}\Big|_{-\pi/2}^{0} + e^{-iku}\Big|_{0}^{\pi/2} - e^{-iku}\Big|_{\pi/2}^{\pi}\right] && (58)\\
&= \frac{i}{2\pi k}\left[(e^{ik\frac{\pi}{2}} - e^{ik\pi}) - (1 - e^{ik\frac{\pi}{2}}) + (e^{-ik\frac{\pi}{2}} - 1) - (e^{-ik\pi} - e^{-ik\frac{\pi}{2}})\right] && (59)\\
&= \frac{i}{\pi k}\left[e^{ik\frac{\pi}{2}} + e^{-ik\frac{\pi}{2}} - 1 - \frac{1}{2}e^{ik\pi} - \frac{1}{2}e^{-ik\pi}\right] && (60)\\
&= \frac{i}{\pi k}\left[2\cos(k\frac{\pi}{2}) - 1 - \cos(k\pi)\right] && (61)
\end{aligned}$$

If $k = 4n$, we get $2\cos(2n\pi) - 1 - \cos(4n\pi) = 0$
If $k = 4n + 1$, we get $2\cos(2n\pi + \frac{\pi}{2}) - 1 - \cos(4n\pi + \pi) = 0$
If $k = 4n + 2$, we get $2\cos(2n\pi + \pi) - 1 - \cos(4n\pi + 2\pi) = -4$
If $k = 4n + 3$, we get $2\cos(2n\pi + \frac{3\pi}{2}) - 1 - \cos(4n\pi + 3\pi) = 0$
Hence, the fourier series for $SQ(u)$ is non-zero for $k = 4n + 2$, and is equal to $-4\frac{i}{\pi k}$, and we have

$$SQ(u) = \sum_{p=-\infty}^{\infty} -\frac{4i}{\pi(4p+2)}e^{i(4p+2)u} = \sum_{p=-\infty}^{\infty} -\frac{2i}{\pi(2p+1)}e^{i(4p+2)u} \tag{62}$$

Now, we can continue with the derivation of the quadrature. We note that $f(\beta) = \sum_{k,k'} a_k \cos(\beta/2)^k \sin(\beta/2)^{k'}$, where k and k' are always even, at least zero and bounded. We can re-express them as complex exponentials so that $f(\beta) = \sum_q b_q(e^{i\beta/2})^q$ where now, q can take on negative values, but it is still bounded

on top and below, $|q| \leq 2(l+l') \leq 2(L+\widetilde{L}_n)$. Note that if we make the substitution $\beta' = \beta/2$, we now have $f(\beta') = \sum_{k,k'} a_k \cos(\beta')^k \sin(\beta')^{k'} = \sum_q b_q (e^{i\beta'})^q$. Therefore,

$$\begin{aligned}
& \int_0^{\pi} f(\beta)\sin(\beta)d\beta \\
= \;& 2\int_0^{\pi/2} f(\beta')\sin(2\beta')d\beta' && (63)\\
= \;& \frac{1}{2}\int_{-\pi}^{-\pi/2} f(\beta')\sin(2\beta')d\beta' - \frac{1}{2}\int_{-\pi/2}^{0} f(\beta')\sin(2\beta')d\beta' && (64)\\
& \qquad + \frac{1}{2}\int_0^{\pi/2} f(\beta')\sin(2\beta')d\beta' - \frac{1}{2}\int_{\pi/2}^{\pi} f(\beta')\sin(2\beta')d\beta' \\
= \;& \frac{1}{2}\int_{-\pi}^{\pi} f(\beta')\sin(2\beta')SQ(\beta')d\beta' && (65)
\end{aligned}$$

where the middle equality was obtained using symmetry arguments since f is a linear combination of even positive powers of sin and cos.

Since $|q| \leq Q$ implies that the term $f(\beta')\sin(2\beta')$ has exponential powers $\leq Q+2$, we can eliminate terms in the fourier series of $SQ(\beta')$ that falls out of the range (by orthonormality of the exponentials), hence we only require p such that

$$\begin{aligned}
& |4p+2| \leq Q+2 && (66)\\
\Leftrightarrow \;& -Q-2 \leq 4p+2 \leq Q+2 && (67)\\
\Leftrightarrow \;& -\frac{Q}{4} - 1 \leq p \leq \frac{Q}{4} && (68)\\
\Leftrightarrow \;& -\text{floor}(\frac{Q}{4}) - 1 \leq p \leq \text{floor}(\frac{Q}{4}) && (69)
\end{aligned}$$

Hence defining $\widetilde{SQ}(\beta') = \sum_{p=-\text{floor}(\frac{Q}{4})-1}^{\text{floor}(\frac{Q}{4})} -\frac{2i}{\pi(2p+1)} e^{i(4p+2)\beta'}$, we have

$$\begin{aligned}
\int_0^{\pi} f(\beta)\sin(\beta)d\beta &= \frac{1}{2}\int_{-\pi}^{\pi} f(\beta')\sin(2\beta')\widetilde{SQ}(\beta')d\beta' && (70)\\
&= \frac{1}{2}\sum_{q=-Q}^{Q} b_q \sum_{p=-\text{floor}(\frac{Q}{4})-1}^{\text{floor}(\frac{Q}{4})} -\frac{2i}{\pi(2p+1)}\int_{-\pi}^{\pi} e^{i\beta' q}\frac{e^{i2\beta'} - e^{-i2\beta'}}{2i} e^{i(4p+2)\beta'}d\beta' && (71)
\end{aligned}$$

Let us define the highest exponential power to be B and notice that $B = Q + 2 + 4\text{floor}(\frac{Q}{4}) + 2 = Q + 4\text{floor}(Q/4) + 4$, while lowest exponential power corresponds to $-Q - 2 - 4\text{floor}(\frac{Q}{4}) - 2 = -Q - 4\text{floor}(\frac{Q}{4}) - 4 = -B$. Let N be the smallest integer such that $B < 4N$, where N is an integer. Hence $N = \text{floor}(\frac{Q}{4}) + \text{floor}(\frac{Q}{4}) + 1 + 1 = 2\text{floor}(\frac{Q}{4}) + 2$

It is easy to verify the following identity:

$$\frac{1}{4N}\sum_{k=-2N}^{2N-1} e^{\frac{2\pi ikl}{4N}} = \frac{1}{2\pi}\int_{-\pi}^{\pi} e^{il\beta'}d\beta' \qquad \forall |l| < 4N \tag{72}$$

Substituting the identity into the equation above, we get

$$\frac{1}{2}\sum_{q=-Q}^{Q} b_q \sum_{p=-\text{floor}(\frac{Q}{4})-1}^{\text{floor}(\frac{Q}{4})} -\frac{2i}{\pi(2p+1)}\frac{\pi}{2N}\sum_{k=-2N}^{2N-1} e^{\frac{2\pi ikq}{4N}}\frac{e^{\frac{2\pi ik2}{4N}}-e^{\frac{-2\pi ik2}{4N}}}{2i}e^{\frac{2\pi ik(4p+2)}{4N}} \quad (73)$$

$$= \frac{\pi}{4N}\sum_{k=-2N}^{2N-1}\sum_{q=-Q}^{Q} b_q e^{\frac{2\pi ikq}{4N}}\sin(\frac{4\pi k}{4N})\sum_{p=-\text{floor}(\frac{Q}{4})-1}^{\text{floor}(\frac{Q}{4})} -\frac{2i}{\pi(2p+1)}e^{\frac{2\pi ik(4p+2)}{4N}} \quad (74)$$

$$= \frac{\pi}{4N}\sum_{k=-2N}^{2N-1} f(\beta'_k=\frac{2\pi k}{4N})\sin(\frac{4\pi k}{4N})\widetilde{SQ}(\frac{2\pi k}{4N}) \quad (75)$$

$$= \frac{\pi}{4N}\sum_{k=-2N}^{2N-1} f(\beta'_k=\frac{2\pi k}{4N})\sin(\frac{\pi k}{N})\widetilde{SQ}(\frac{\pi k}{2N}) \quad (76)$$

$$= \frac{\pi}{2N}\sum_{k=0}^{2N-1} f(\beta'_k=\frac{2\pi k}{4N})\sin(\frac{\pi k}{N})\widetilde{SQ}(\frac{\pi k}{2N}) \quad (77)$$

$$= \frac{\pi}{N}\sum_{k=0}^{N-1} f(\beta'_k=\frac{2\pi k}{4N})\sin(\frac{\pi k}{N})\widetilde{SQ}(\frac{\pi k}{2N}) \quad (78)$$

where the second last equality uses the fact that $\sin(-2\pi)=0$ and the function f is even, and the last equality uses the fact that $\sin\pi=0$ and $f(\beta')$ is even about $\beta'=\pi$. Because $\beta'_k=\beta_k/2=\frac{2\pi k}{4N}$, hence $\beta_k=\frac{\pi k}{N}$ for $k=0,1,\cdots N-1$ corresponds to our quadrature knots, with quadrature weights, $w_k=\frac{\pi}{N}\sin(\frac{\pi k}{N})\widetilde{SQ}(\frac{\pi k}{2N})$.

D.3 Quadrature Conclusion

We have formulated two possible ways of obtaining an exact quadrature of $\int_0^\pi d^l_{mm'}(\beta)d^{l'}_{mm'}(\beta)d\beta$ and the integral is equal to $\sum_{s=0}^{N-1} w_s d^l_{mm'}(\beta_s)d^{l'}_{mm'}(\beta_s)=\frac{2}{2l+1}\delta(l-l')$ if we pick the correct samples and corresponding weights. In particular, this is true for

1. $\beta_s=2\sin^{-1}(\frac{v_s+1}{2})$ for $s=0,1,\cdots N-1$ with $N=\text{roof}(\frac{Q}{2})+1$, where Q is the highest polynomial power of $f(u=\sin\frac{\beta}{2})$ and $w_s=r_s(v_s+1)$, where r_s and v_s are the weights and nodes of the gaussian-lengendre quadrature on the interval [-1,1]
2. $\beta_s=\frac{\pi s}{N}$ for $s=0,1,\cdots N-1$ with $N=2\text{floor}(\frac{Q}{4})+2$, where Q is the highest complex exponential power of $d^l_{mm'}(\beta'=\frac{\beta}{2})d^{l'}_{mm'}(\beta'=\frac{\beta}{2})$ and $w_s=\frac{\pi}{N}\sin(\frac{\pi s}{N})\widetilde{SQ}(\frac{\pi s}{2N})$

Note that $Q=2(l+l')\leq 2(L+\widetilde{L}_n)$

E Inverse Convolution with Axis-symmetric Filter

In this section, we illustrate the computation of the inverse convolution of $y(\alpha,\beta)$ with an axis-symmetric filter $h(\theta,\phi)$. Starting with the definition of inverse convolution, using Parseval's Theorem and Eq. (27),

we get

$$\widehat{x}(\theta,\phi) = \int_{SO(3)} [D(\alpha,\beta,\gamma)h](\theta,\phi)y(\alpha,\beta)d\rho \quad (79)$$

$$= \int_{SO(3)} \left[\sum_{l,m} \left(D^l_{m0}(\alpha,\beta,\gamma)h^{l,0}\right) Y^m_l(\theta,\phi)\right] y(\alpha,\beta)d\rho \quad (80)$$

$$= \int_{SO(3)} \left[\sum_{l,m} \left(D^l_{m0}(\alpha,\beta,\gamma)h^{l,0}\right) \sqrt{\frac{2l+1}{4\pi}} D^{l*}_{m0}(\phi,\theta,\gamma')\right] y(\alpha,\beta)d\rho \quad (81)$$

Noting that γ' can take on any value without affecting the equation, we get

$$\widehat{x}(\theta,\phi) = \int_{SO(3)} \left[\sum_{l,m} \left(D^{l*}_{m0}(\phi,\theta,\gamma')h^{l,0}\right) Y^{m*}_l(\beta,\alpha)\right] y(\alpha,\beta)d\rho \quad (82)$$

$$= \int_{SO(3)} \left[\sum_{l,m} \left(D^l_{m0}(\phi,\theta,\gamma')h^{l,0*}\right) Y^m_l(\beta,\alpha)\right]^* y(\alpha,\beta)d\rho \quad (83)$$

$$= \int_{SO(3)} \left[D(\phi,\theta,\gamma')h^*\right]^* (\beta,\alpha)y(\alpha,\beta)d\rho \quad (84)$$

$$= \int_0^{2\pi}\int_0^{\pi}\int_0^{2\pi} \left[D(\phi,\theta,\gamma')h^*\right]^* (\beta,\alpha)y(\alpha,\beta)\sin\beta d\alpha d\beta d\gamma \quad (85)$$

$$= 2\pi\int_0^{\pi}\int_0^{2\pi} \left[D(\phi,\theta,\gamma')h^*\right]^* (\beta,\alpha)y(\alpha,\beta)\sin\beta d\alpha d\beta \quad (86)$$

Eq. (86) without the 2π is simply a forward convolution between the spherical image $y(\alpha,\beta)$ (where we remind the readers that α is taking the role of ϕ and β is taking the role of θ) and the filter $h^*(\beta,\alpha)$. Hence using Eq. (15), we get

$$[\widehat{x}(\theta,\phi)]^{l,m} = 2\pi\sqrt{\frac{4\pi}{(2l+1)}} y^{l,m}h^{l,0} \quad (87)$$

References

[1] Antoine, J-P and Vandergheynst, P. Wavelets on the 2-Sphere: a Group-Theoretical Approach. *Applied and Computational Harmonic Analysis*, 7:262–291, 1999.

[2] Bertsekas, D.P. Nonlinear Programming. *Athena Scientific*, 1998.

[3] Bogdanova I., Vandergheynst, P., Antoine, J-P, Jacques, L., and Morvidone, M. Stereographic Wavelet Frames on the Sphere. *Applied and Computational Harmonic Analysis*, 19:223–252, 2005.

[4] Brechbuhler, C.H., Gerig, G., and Kubler O. Parametrization of Closed Surfaces for 3-D Shape Description. *Computer Vision and Image Understanding*, 61(2):154–170, 1994.

[5] Daubechies, I. Ten Lectures on Wavelets. *Society for Industrial and Applied Mathematics (SIAM)*.

[6] Driscoll, J.R. and Healy, D.M. Computing Fourier Transforms and Convolutions on the 2-Sphere. *Advances in Applied Mathematics*, 15:202–250, 1994.

[7] Fischl, B., Sereno M.I., and Dale, A.M. Cortical Surface-Based Analysis: II: Inflation, Flattening, and a Surface-Based Coordinate System. *NeuroImage*, 9(2):195–207, 1999.

[8] Fischl, B., Sereno M.I., Totell, B.H., and Dale, A.M. Cortical Surface-Based Analysis: High-resolution Inter-Subject Averaging and a Coordinate System for the Cortical Surface. *Human Brain Mapping*, 8(4):272–284, 1999.

[9] Freeman, W.T. and Adelson, E.W. The design and use of steerable filters. *IEEE Transaction Pattern Analysis and Machine Intelligence*, 13(9):891–906, 1991.

[10] Kostelec, P.J. and Rockmore, D.N. A lite version of spharmonic kit. *http://www.cs.dartmouth.edu/geelong/sphere/*.

[11] Press, W.H., Teukolsky, S.A., Vetterling, W.T. and Flannery, B.P. Numerical Recipes in C: The Art of Scientific Computing. *Second Edition, Cambridge University Press*.

[12] Sakurai, J.J. Modern Quantum Mechanics *Second Edition, Addison Wesley*, 1994.

[13] Schroder, Peter and Sweldens, Wim Spherical Wavelets: Efficiently Representing Functions on the Sphere. *Computer Graphics Proceedings (SIGGRAPH)*, 161–172, 1995.

[14] Schroder, Peter and Sweldens, Wim Spherical Wavelets: Texture Processing. *Rendering Techniques*, 252–263, Aug 1995.

[15] Simoncelli, E.P., Freeman, W.T., Adelson, E.H., and Heeger, D.J. Shiftable multi-scale transforms. *IEEE Transaction Information Theory*, 38(2):587–607, 1992.

[16] Staib, Lawrence H. and Duncan, James S. Model-Based Deformable Surface Finding for Medical Images. *IEEE Transactions on Medical Imaging*, 15(5):720–731, 1996.

[17] Van$\acute{i}\check{c}$ek, P., Jan$\acute{a}$k, J., and Featherstone,W.E. Truncation of spherical convolution integration with an isotropic kernel. *Studia Geophysica et Geodaetica*, 47(3):455–465, 2003.

[18] Wandelt, B.D. and Gorski, K.M. Fast Convolution on the sphere. *Physics Review D 63*, 1230002:1-6, 2001.

[19] Wiaux, Y., Jacques, L., and Vandergheynst. P. Correspondence principle between spherical and euclidean wavelets. *Astrophys. J.*, 15:632, 2005.

[20] Wiaux, Y., Jacques, L., and Vandergheynst. P. Fast directional correlation on the sphere with steerable filters. *Astrophys. J.*, submitted, 2005.