

On the Construction of Invertible Filter Banks on the 2-Sphere

Boon Thye Thomas Yeo, Wanmei Ou, and Polina Golland

Abstract—The theories of signal sampling, filter banks, wavelets, and “overcomplete wavelets” are well established for the Euclidean spaces and are widely used in the processing and analysis of images. While recent advances have extended some filtering methods to spherical images, many key challenges remain. In this paper, we develop theoretical conditions for the invertibility of filter banks under continuous spherical convolution. Furthermore, we present an analogue of the Papoulis generalized sampling theorem on the 2-Sphere. We use the theoretical results to establish a general framework for the design of invertible filter banks on the sphere and demonstrate the approach with examples of self-invertible spherical wavelets and steerable pyramids. We conclude by examining the use of a self-invertible spherical steerable pyramid in a denoising experiment and discussing the computational complexity of the filtering framework.

Index Terms—Channel bank filters, feature extraction, filtering, frequency response, image sampling, image orientation analysis, spheres, spherical images, wavelet transforms.

I. INTRODUCTION

MULTISCALE filtering methods, such as wavelets [7] and “overcomplete wavelets” [9], [23], [25] have many applications in feature detection, compression, and denoising of planar images. Extending the theories and the methods of filtering to spherical images promises similar benefits in the fields that give rise to such images, including shape analysis in computer vision [5], illumination computation in computer graphics [21], [22], cosmic background radiation analysis in astrophysics [26], [29] and brain cortical surface analysis in medical imaging [32], [33].

We consider a two-stage filtering framework [Fig. 1(a)], conceptually equivalent to the usual Euclidean filtering framework except the planar images and filters are replaced by spherical ones. We can think of the first set of filters as analysis

Manuscript received May 24, 2007; revised October 23, 2007. This work was supported in part by the National Alliance for Medical Image Analysis (NIH NIBIB NAMED U54-EB005149), in part by the Neuroimaging Analysis Center (NIT CRR NAC P41-RR13218), in part by the Morphometry Biomedical Informatics Research Network (NIH NCRR mBIRN U24-RR021382), in part by the NIH Grant NINDS R01-NS051826, and in part by the National Science Foundation (NSF) under Grant CAREER 0642971. B. T. T. Yeo was supported in part by the Agency for Science, Technology, and Research, Singapore. W. Ou was supported by the NSF graduate fellowship. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Ljubisa Stankovic.

The authors are with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: ythomas@csail.mit.edu; wanmei@csail.mit.edu; polina@csail.mit.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2007.915550

filters which project the input image onto the space spanned by the analysis filters. A reconstructed image is then obtained by passing the intermediate outputs through the second layer of filters. Fig. 1(b) shows a modification of the framework in Fig. 1(a) that introduces sampling between the first and the second layers of filters. The sampling is useful if one is interested in processing the outputs of the analysis filters before passing them through the synthesis filters.

In this work, we analyze the relationship between the reconstructed image and the original image, and establish conditions under which the reconstructed image is the same as the original one, i.e., conditions for invertibility. We demonstrate the use of our results on continuous invertibility and generalized sampling for designing filter banks that enable explicit control of both analysis and synthesis filters. We illustrate the framework by creating examples of self-invertible spherical wavelets and steerable pyramids.

Self-invertible filter banks employ identical filters for analysis and synthesis. 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 nonlinear processing of wavelet coefficients could propagate to spatial locations and frequencies other than those which were used to compute the coefficients [23]. To the best of our knowledge, this is the first approach demonstrated on a sphere that enables the design of self-invertible filter banks.

II. RELATED WORK

In the Euclidean domain, convolution is computed efficiently using the fast Fourier transform (FFT) [6]. Once we move to the sphere, FFT must be replaced with an alternative efficient method for computing convolutions. An original algorithm for axisymmetric convolution kernels on the sphere was derived in [11] and was recently extended to arbitrary functions [28], [30]. These results allow us to efficiently compute the outputs of the first set of filters. Unfortunately, they do not apply to the inverse convolution with the second layer of filters because the outputs of the first layer of filters are in general not spherical images as we will see in Section III.

In the past decade, there has been much work on extending the general paradigm of linear filtering to the spherical domain [1], [4], [8], [12], [13], [17], [21], [22], [26], [29]. For example, the lifting scheme in [21] and [22] adopts a nonparametric approach to computing the wavelet decomposition of arbitrary meshes by generalizing the standard two-scale relation of Euclidean wavelets. This method enables a multiscale representa-

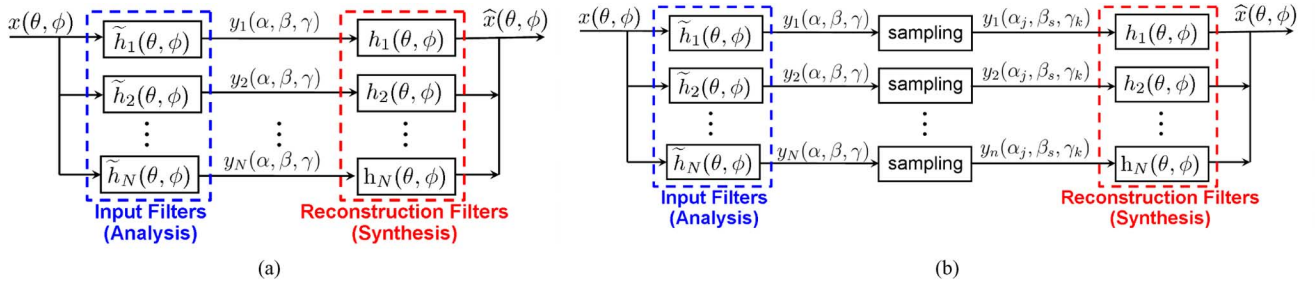


Fig. 1. Continuous and discrete filter bank diagram. (a) Continuous filter bank. (b) Discrete filter bank.

tion of the original mesh (image) with excellent compression and speed 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 and invariance to rotation [33].

A similar problem in the Euclidean domain leads to the invention of “overcomplete wavelets,” such as steerable pyramids [14], [23]. A group theoretic formulation of overcomplete continuous spherical wavelets is proposed in [1]. In particular, it can be shown that 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 [29].

In the group theoretic approach, defining the mother wavelet completely determines the analysis and synthesis filters. However, while the analysis filters are related by stereographic dilation, the synthesis filters are in general not related by dilation. In fact, the support of corresponding analysis and synthesis filters is guaranteed to be the same in frequency domain but not in the spatial domain. Bogdanova *et al.* [4] discretize the group theoretic wavelets, providing a sampling guarantee for the framework of Fig. 1(b) for the restricted class of axisymmetric filters. An axisymmetric spherical function is one which is symmetrical about the north pole. This work is, therefore, the most similar to ours. In contrast, we study general filter banks, without any restriction on the relationships among the cascade of filters. We derive the analogue of the Papoulis generalized sampling theorem [18] on the sphere, applicable to both axisymmetric and nonaxisymmetric filters.

Driscoll and Healy [11] provide the equivalent of the Nyquist–Shannon sampling theorem on the sphere. While the Nyquist–Shannon sampling theorem provides reconstruction guarantees for bandlimited signals in Euclidean space under perfect sampling (convolution with a delta function), the Papoulis generalized sampling theorem provides guarantees for bandlimited signals sampled via convolutions with kernels of sufficient bandwidth.

An earlier version of this work was first presented at the International Conference on Image Processing [31]. In this paper, we include proofs of the invertibility conditions and demonstrate the generation of self-invertible spherical steerable pyramids. In Section III, we introduce the notation used throughout the paper. In Section IV, we present the main theoretical contributions of this paper: continuous invertibility and the generalized sampling theorem. We propose a procedure for generating

self-invertible multiscale filter banks on the sphere in Section V. In Section VI, we illustrate the procedure to design wavelets and steerable pyramids and employ a steerable pyramid in denoising. We conclude with the discussion of future research and outstanding challenges in the proposed framework.

To summarize, our contributions are as follows.

- 1) We present theoretical conditions for the invertibility of axisymmetric and nonaxisymmetric filter banks under continuous spherical convolution.
- 2) We present a generalized sampling theorem of signals on the 2-Sphere for both axisymmetric and nonaxisymmetric filter banks. This generalizes the works of Bogdanova *et al.* [4] and Starck *et al.* [26] to nonaxisymmetric filters and opens a way for nonlinear processing of the wavelet coefficients generated from general filter banks.
- 3) We present a mechanism for generating invertible, as well as self-invertible, wavelets, and steerable pyramids. We provide an analysis of the computational complexity of the filtering framework.

III. DEFINITIONS

Let $x(\theta, \phi) \in L^2(S^2)$ be a square-integrable function on the 2-D unit sphere, where (θ, ϕ) are the spherical coordinates. Suppose $P = (\theta, \phi)$ is a point on the sphere. Then, $\theta \in [0, \pi]$ is the co-latitude, which is the angle between the positive z -axis (north pole) and the vector corresponding to P . $\phi \in [0, 2\pi]$ is the longitude and is taken to be the angle between the positive x -axis and the projection of P onto the $x - y$ plane. ϕ is undefined on the north and south poles.

The spherical harmonics $Y_l^m(\theta, \phi)$ [20] form an orthonormal set of basis functions for $L^2(S^2)$: i.e.,

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

where $x^{l,m}$ is the spherical harmonic coefficient of degree l and order m obtained by projecting the function $x(\theta, \phi)$ onto $Y_l^m(\theta, \phi)$

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

where $d\Omega = \sin \theta d\theta d\phi$ and $*$ denotes complex conjugation. We call $Y_l^m(\theta, \phi)$ a spherical harmonic of degree l and order m . We note that for axisymmetric functions (independent of ϕ), only




Fig. 2. Rotation via euler angles (α, β, γ) . (a) Original spherical image. (b) Rotation by γ about z -axis. (c) Rotation by β about y -axis. (d) Rotation by α about z -axis.

the order 0 harmonics are nonzero. A more detailed background of spherical harmonics is found in Appendix A.

We choose to parameterize rotations on the sphere by the 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 a function by γ about the z -axis [Fig. 2(b)], then by β about the y -axis [Fig. 2(c)], and finally by α about the z -axis [Fig. 2(d)]. 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 spherical 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 [20]. Appendix A provides the explicit expressions for the Wigner-D functions.

A. Continuous Convolution

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

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

is a function of $L^2(SO(3))$ rather than $L^2(S^2)$. This definition of convolution is identical to that in [28], [30], although [30] calls it directional correlation.

By convention, we shall consider the center (origin) of a spherical filter to be at the north pole ($\theta = 0$). Then intuitively, $y(\alpha, \beta, \gamma)$ is the inner product between the re-oriented filter $D(\alpha, \beta, \gamma)\tilde{h}$ [e.g., Fig. 2(d)] and the spherical image. In other words, we obtain $y(\alpha, \beta, \gamma)$ by first re-orienting the spherical filter by a rotation of γ about the z -axis (center still at north pole), then bringing the center of the filter to the point (β, α) of the spherical image, and performing an inner product between the image and filter. Therefore, $y(\alpha, \beta, \gamma)$ is the correlation of the rotated version of \tilde{h} with x , or the projection coefficient of x onto $[D(\alpha, \beta, \gamma)\tilde{h}]$. In the case of the filter shown in Fig. 2, a high value of $y(\alpha, \beta, \gamma)$ would imply a presence of an oriented edge at spherical coordinate (β, α) with orientation γ .

We note that the notion of orientation is inherently local here because a continuous unit-norm vector field does not exist on the

sphere. Therefore, it is not meaningful to claim an existence of an edge of orientation γ at location (β, α) without specifying a local coordinate system. In our case, we can define such a local coordinate system by first specifying one at the north pole. Our choice of parameterizing rotation via the Euler angles then induces a local coordinate system everywhere, except the south pole.

For axisymmetric filters, $\tilde{h}(\theta, \phi) = \tilde{h}(\theta)$, the rotation by γ about the z -axis has no effect, i.e., $y(\alpha, \beta, \gamma) = y(\alpha, \beta)$ is a spherical image parametrized by $\theta = \beta$, $\phi = \alpha$.

To project a function in $L^2(SO(3))$ onto $L^2(S^2)$, we define the inverse convolution of a spherical filter $h(\theta, \phi)$ with $y(\alpha, \beta, \gamma) \in L^2(SO(3))$ to be

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

where the integration is over the Euler angles: $d\rho = \sin\beta d\alpha d\beta d\gamma$. Our definition of inverse convolution generalizes the spherical convolution between two functions on S^2 defined by Driscoll and Healy [11], and is identical to theirs if $y(\alpha, \beta, \gamma) = y(\beta, \alpha)$, i.e., a spherical image. Our inverse convolution operation is similar to the transpose convolution operation defined in [2], [28]. However, the transpose convolution uses the uniform measure $d\alpha d\beta d\gamma$, ignoring the intrinsic nonuniform component $\sin\beta$ of the measure on $SO(3)$.

We can think of the inverse convolution in the following intuitive way. The reconstructed value at a given (θ, ϕ) is obtained by summing (integrating) the contributions of the rotated reconstruction filters h , centered at all possible positions (β, α) and oriented by all possible angles γ , where the weights of the contributions are given by the convolution outputs (projection coefficients on the corresponding input filters).

We distinguish between the forward convolution (3) and the inverse convolution (4) because the forward convolution combines two spherical images to give a function on $SO(3)$, while the inverse convolution combines a spherical image and a function on $SO(3)$ to give a spherical image. Both definitions of convolutions generalize the concept of convolution in Euclidean spaces.

When using a filter bank of N analysis-synthesis filter pairs [Fig. 1(a)], the reconstructed signal is obtained by summing the response of all filter pairs

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

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

B. Discrete Convolution

In the Euclidean case, we typically discretize both the input images and the convolution outputs. When working on the sphere, we 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 [28], [30]. Since no uniform sampling grid exists on the sphere, performing convolution completely by quadrature would be slow. This is because under each rotation of the filter relative to the spherical image, we would need to re-sample (or re-interpolate) the filter or the image.

For $L^2(SO(3))$ (or equivalently, in the spherical wavelet domain), continuous representation is possible through series of complex exponentials [28] or Wigner-D functions [30]. However, both the complex exponentials and the Wigner-D functions have global support. Therefore, in applications where we want to modify the image in the wavelet domain, manipulating the series coefficients would be tantamount to simultaneously altering all the wavelet coefficients, defeating the purpose of the wavelet decomposition, which is to provide localized control in both spatial and frequency domain. To avoid this, we 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 [Fig. 1(b)]. The inverse convolution between the sampled projection coefficients $y_{j,s,k}$ and the continuous reconstruction filters h is then defined as

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

which includes sampling-dependent quadrature weights $w_{j,s,k}$, introduced so that the discrete case converges to the continuous case as the number of samples increases. This definition allows for an easy transfer of continuous filtering theory to its discrete analogue. In Section IV, we show that “good” choices of $w_{j,s,k}$ exist, corresponding to different sampling schemes. In contrast with the Euclidean case, $w_{j,s,k}$ are necessary because of the nonuniform measure on the Euler angles $d\rho = \sin \beta d\alpha d\beta d\gamma$, as we discuss in Section IV.

Similar to the continuous case [cf. (4)], the signal reconstructed through N analysis-synthesis filter pairs is defined as a sum of contributions of all filter pairs

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

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

IV. INVERTIBILITY CONDITIONS

In this section, we present the main theoretical contributions of our work.

1) **Theorem 4.1: (Continuous Frequency Response).** Let $\{\tilde{h}_n, h_n\}_{n=1}^N$ be an analysis-synthesis filter bank. Then for any spherical image $x \in L^2(S^2)$ and its corresponding reconstructed image \hat{x}

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

where $x^{l,m}$ and $\hat{x}^{l,m}$ are the spherical harmonic coefficients of the input and reconstructed images, respectively. $\tilde{h}_n^{l,m'}$ and $h_n^{l,m'}$ are the spherical harmonic coefficients of the n th analysis and synthesis filters, respectively.

Appendix B presents the proof of Theorem 4.1. To draw an analogy with the Euclidean case, we call

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

the frequency response of the analysis-synthesis filter bank. Note that the degree l spherical harmonics coefficients of the reconstructed image are affected only by the degree l spherical harmonic coefficients of the filters. However, the degree l order m spherical harmonic coefficient of the reconstructed signal is affected by all the orders of degree l spherical harmonic coefficients of the filters. In contrast, on the plane, the frequency response is simply the sum of products of the Fourier coefficients of the analysis and the synthesis filters

$$\mathcal{F}\{\hat{x}\}(s_1, s_2) = \mathcal{F}\{x\}(s_1, s_2) \times \sum_{n=1}^N \mathcal{F}\{h_n\}(s_1, s_2)^* \mathcal{F}\{\tilde{h}_n\}(s_1, s_2)^* \quad (10)$$

where $\mathcal{F}\{\hat{x}\}(s_1, s_2)$, $\mathcal{F}\{x\}(s_1, s_2)$, $\mathcal{F}\{\tilde{h}_n\}(s_1, s_2)$ and $\mathcal{F}\{h_n\}(s_1, s_2)$ denote the Fourier transforms of the reconstructed signal, original signal, analysis filters and synthesis filters respectively. We note that under Euclidean space conventions, a “self-invertible” filter bank is one with $\tilde{h}_n(x, y) = h_n(-x, -y)$ and $\sum_{n=1}^N |\mathcal{F}\{h_n\}(s_1, s_2)|^2 = 1$. We see that the effects of s_1 and s_2 are separable, unlike l and m . Furthermore, on the sphere, the frequency response contains an extra modulating factor that decreases with degree l . The following corollary of Theorem 4.1 provides the necessary and sufficient condition for the invertibility of filter banks under continuous convolution.

2) **Corollary 4.2: (Continuous Invertibility).** Let $\{\tilde{h}_n, h_n\}_{n=1}^N$ be an analysis-synthesis filter bank. Then for any spherical image $x \in L^2(S^2)$ and its corresponding reconstructed image \hat{x}

$$\hat{x}^{l,m} = x^{l,m} \text{ for all } (l, m) \text{ iff } \sum_{n=1}^N \sum_{m'=-l}^l [h_n^{l,m'}] [\tilde{h}_n^{l,m'}]^* = \frac{2l+1}{8\pi^2} \text{ for all } l \text{ s.t. } x^{l,m} \neq 0. \quad (11)$$

We note that the corollary is easily satisfied if there is no constraint on the relationships among the cascade of filters: given a set of analysis filters \tilde{h}_n , there are in general multiple sets of synthesis filters that can achieve invertibility. For example, we can define the synthesis filters to be $h_n = L_\psi \tilde{h}_n$, where

$$[L_\psi \tilde{h}_n]^{l,m} = \begin{cases} \frac{1}{H_{\tilde{h},h}(l)} \tilde{h}_n^{l,m}, & \text{for } H_{\tilde{h},h}(l) > 0 \\ 0, & \text{otherwise} \end{cases} \quad (12)$$

and $H_{\tilde{h},h}(l)$ is the frequency response defined in (9). L_ψ is a frequency modulating operator that normalizes the synthesis filters at each degree, such that the combined frequency response of the filter bank is 1 for all l with $H_{\tilde{h},h}(l) > 0$. This filter bank is, therefore, invertible over the frequency range of the support of the filters. This operation is similar to the frame operator in the continuous spherical wavelet transform of [1], where the counterpart of $H_{\tilde{h},h}(l)$ is given by $(8\pi^2/(2l+1)) \sum_{|m| \leq l} \int_0^\infty (1/a^3) |\tilde{h}_a^{l,m}|^2 da$, replacing the summation over n by the integration over the scale a , with measure $1/a^3 da$. For the special case of the analysis filters being dilated versions of each other, this choice of the synthesis filters is a direct discretization of [1], albeit ignoring the measure of a . The complete discretization of the continuous wavelet transform in [1] is actually accomplished in [4]. However, regardless of using L_ψ or the frame operators of [1], [4], the synthesis filters are in general not related by dilation even if the analysis filters are.

We now define \tilde{L}_n (\tilde{O}_n) and L_n (O_n) to be the highest nonzero harmonic degree (order) of \tilde{h}_n and h_n , respectively. The following result relates the spherical harmonics coefficients of the input image $x^{l,m}$ and reconstructed image $\hat{x}^{l,m}$ under the sampling framework of Fig. 1(b).

3) **Theorem 4.3: (Generalized Sampling Theorem).** Let $\{\tilde{h}_n, h_n\}_{n=1}^N$ be a filter bank with \tilde{L}_n (and, thus, \tilde{O}_n) $< \infty$ and L_n (and, thus, O_n) $< \infty$. Suppose the sampling grid and the quadrature weights satisfy:

- $\alpha_{j,n} = 2\pi j / (\tilde{L}_n + L_n + 1)$ for $j = 0, 1, \dots, (\tilde{L}_n + L_n)$;
- $\gamma_{k,n} = 2\pi k / (\tilde{O}_n + O_n + 1)$ for $k = 0, 1, \dots, (\tilde{O}_n + O_n)$;
- $w_{s,n}$ and $\beta_{s,n}$ are the quadrature weights and knots such that

$$\int_0^\pi d_{mm'}^l(\beta) d_{mm'}^{l'}(\beta) \sin(\beta) d\beta = \sum_{s=0}^{S_n-1} w_{s,n} d_{mm'}^l(\beta_{s,n}) d_{mm'}^{l'}(\beta_{s,n}) \quad (13)$$

for $l \leq L_n$, $l' \leq \tilde{L}_n$, where $d_{mm'}^l(\beta)$ and $d_{mm'}^{l'}(\beta)$ are the Wigner-d functions;

- $w_{j,s,k,n} = 4\pi^2 w_{s,n} / ((\tilde{L}_n + L_n + 1)(\tilde{O}_n + O_n + 1))$.

Then

$$\hat{x}^{l,m} = x^{l,m} \frac{8\pi^2}{2l+1} \left\{ \sum_{n=1}^N \sum_{m'=-l}^l [h_n^{l,m'}] [\tilde{h}_n^{l,m'}]^* \right\}. \quad (14)$$

Appendix A provides the definitions and the explicit expressions for the Wigner-d functions. Note that there is an implicit link between the number of samples S_n and the bandwidths \tilde{L}_n

and L_n (see Appendix D). The constraints in this theorem ensure that the number of samples remain finite and is similar to the tradeoff between the number of points in a spherical grid and the maximum bandwidth of signals defined on it [11]. The samples and the weights are picked such that the discrete reconstruction defined in (7) gives the same result as the continuous version in (5). The proof is found in Appendix C. In Appendix D, we demonstrate two sets of quadrature weights and knots that satisfy the conditions of the theorem. We note that the use of quadrature methods as tools for sampling proofs was previously employed in the fast spherical transform [11], [15] and the fast Wigner-D transform (SOFT) [16].

The measures corresponding to α and γ are constants in $SO(3)$, just like in the Euclidean space. We, therefore, use uniform sampling for these parameters in our work. For discrete planar convolution, it is customary to have no weights (or rather, unit weights). On the sphere, however, the nonuniform measure on β , $\sin \beta d\beta$, presents challenges for sampling. If we are simply interested in convergence, then setting $w_{j,s,k} = (2\pi/J)(2\pi/S) \sin(\beta_s)(2\pi/K)$ for uniform samples of α, β, γ corresponds to the Riemann sum of the integral. Theorem 4.3 states that better quadrature schemes exist that allow for faster convergence with finite numbers of samples.

The continuous invertibility corollary and the generalized sampling theorem imply the following:

4) **Corollary 4.4: (Discrete Invertibility).** Consider a filter bank of filters with finite maximum spherical harmonic degrees. There exists quadrature schemes such that the filter bank is invertible over a particular frequency range under the discrete spherical convolution if the filter bank is also invertible over the same frequency range under the continuous spherical convolution.

Because functions in $L^2(S^2)$ have finite energy, their spherical harmonic coefficients must necessarily decay to zero. Therefore, we can reasonably assume that the filters of the filter bank are of finite bandwidth as required by Theorem 4.3 by representing the filters with a finite number of coefficients up to an arbitrary prespecified precision. We will, therefore, focus on constructing invertible filter banks for continuous convolution.

The discrete invertibility corollary opens a way for nonlinear processing of the wavelet coefficients generated from general filter banks. Applications may include compression, denoising and image enhancement. The corollary also enables the perfect reconstruction of an original signal sampled with equipment that introduces blurring during the acquisition process. For example, the first layer of filters \tilde{h}_n could be the blurring kernels of a set of N radio dishes measuring the cosmic background radiation of the sky. We can then hope to recover the true cosmic background radiation signal by passing the recorded signal through the second bank of filters.

V. CONSTRUCTING SELF-INVERTIBLE MULTISCALE FILTER BANKS

In this section, we show how to use the continuous invertibility corollary to generate self-invertible multiscale filter banks. The optimization framework presented here can be easily adapted to design other types of filter banks by altering




Fig. 3. Stereographic dilation. Correspondence between the spherical surface and tangent plane is established via stereographic projection: point P on the sphere is mapped to the point P' on the tangent plane of the north pole. This allows a spherical function $f(\theta, \phi)$ to be mapped onto the tangent plane. The usual dilation is then performed on the tangent plane. The resulting dilated function is mapped back onto the sphere via stereographic dilation: $(1/a)f(2 \tan^{-1}((1/a) \tan(\theta/2)), \phi)$. The normalization factor $(1 + \tan^2(\theta/2))/(1 + ((1/a) \tan(\theta/2))^2)$ is added to conserve the inner product between dilated functions.

the structure of the optimization problem according to an application's needs.

For self-invertible filter banks, $\tilde{h}_n(\theta, \phi)$ is constrained to be the same as $h_n(\theta, \phi)$. Furthermore, in multiscale analysis, the analysis filters are related through dilation and scaling of a particular template $\tilde{h}(\theta, \phi)$, i.e.,

$$\tilde{h}_n(\theta, \phi) = \left(\prod_{k=1}^n b_k \right) D_{a_n} \tilde{h}(\theta, \phi) \quad (15)$$

where $b_k \geq 1$ and D_{a_n} is the nonlinear dilation operator,¹ both of which will be explained shortly. We use the convention that larger n corresponds to smaller a (narrower filters),

In this paper, we adopt the stereographic dilation operator introduced in [1] (Fig. 3), which involves stereographically projecting the function from the sphere onto the tangent plane of the north pole, performing the usual dilation operation on the plane and then projecting the resulting function back onto the sphere²

$$[D_a f](\theta, \phi) = \left(\frac{1 + \tan^2 \frac{\theta}{2}}{1 + \left(\frac{1}{a} \tan \frac{\theta}{2}\right)^2} \right) \times \frac{1}{a} f \left(2 \tan^{-1} \left(\frac{1}{a} \tan \frac{\theta}{2} \right), \phi \right). \quad (16)$$

The normalization factor $(1 + \tan^2(\theta/2))/(1 + ((1/a) \tan(\theta/2))^2)$ ensures the inner product between functions is conserved under stereographic dilation. Stereographic dilation allows for an explicit control of the spatial localization of the wavelets in contrast with previous approaches that define dilation in the frequency domain [4]. Because of the nonlinear nature of stereographic dilation, extreme dilation of a spherical function will eventually lead to high frequencies. Initial

¹We note that the symbol D is overloaded to imply rotation as well as dilation, but the meaning should be clear depending on the context.

²We note that the approach commonly used with planar images of applying a constant filter to a subsampled image fails here because the sphere is periodic and bounded, causing the effective size of the image features (relative to the filter) to stay constant 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, $f(\theta, \phi) \rightarrow f((\theta/a), \phi)$, leads to undesirable "wrap-around" effects.




Fig. 4. Effect on energy due to norm preserving stereographic dilation.

dilation of a function localized at the north pole increases the support of the function and results in lower frequencies. As dilation continues, oscillations will accumulate near the south pole, resulting in high frequencies. A useful analysis of this phenomenon can be found in Bogdanova *et al.* [4]. In practice, we will avoid working in that region, since the dilated filter no longer looks like the original filter.

The b_k 's in (15) are the amplitude scaling parameters that control the tradeoff between self-invertibility and norm-preserving dilation. Corollary 4.2 implies that the sum of squares of the spherical harmonic coefficients of a bank of self-invertible filters must increase linearly with degree. But stretching a function while preserving its norm shifts its spherical harmonic coefficients to the left (spherical harmonic degrees decrease) and magnifies them (Fig. 4).

These extra weights are analogous to the measure of scale $1/a^3 da$ in the group theoretic formulation of spherical wavelets [1] and its discretization in the axisymmetric case [4], resulting in wider filters being assigned smaller weights. On the continuous real line, the measure $1/a^2 da$ nicely cancels out the dilation of the filter (cf. [27, Ch. 5]). On the discrete real line, the convolution outputs of narrower filters are sampled more densely. This suggests two possible approaches: variable sampling of the convolution outputs or variable scaling of the filters. Yet another approach is to sample the scale space unevenly rather than according to a power law. Because the effects of stereographic dilation on the spherical harmonic coefficients of a function are not analytical, none of these approaches leads to a closed-form solution. In this paper, we take the variable scaling approach by finding the appropriate b_k 's as part of the filter design.

Fortunately, stereographic dilation is distributive over addition. Suppose the template \tilde{h} is expressible as a linear combination of the basis functions $B^i(\theta, \phi)$, i.e., $\tilde{h}(\theta, \phi) = \sum_{i=1}^M c_i B^i(\theta, \phi)$. Here, we assume that $B^i(\theta, \phi)$ are spherical harmonics and note that the technique is still applicable if a more suitable basis is found, such as described in [24]. Applying stereographic dilation to \tilde{h}

$$[D_a \tilde{h}]^{l,m} = \left[D_a \sum_{i=1}^M c_i B^i \right]^{l,m} = \sum_{i=1}^M c_i [D_a B^i]^{l,m} \quad (17)$$

yields the spherical harmonic coefficients of the analysis filter at another scale. This is useful since the invertibility condition in Corollary 4.2 is expressed in terms of the spherical harmonic coefficients of the filters. We can, therefore, decide on a set of scales $\{a_n\}_{n=1}^N$ and precompute a table of spherical harmonic coefficients of the dilated basis functions $[D_a B^i]^{l,m}$. Equation (17) allows us to determine the spherical harmonic coefficients

of the dilated filters at each relative scale given $\vec{c} = \{c_i\}_{i=1}^M$ and $\vec{b} = \{b_k\}_{k=1}^M$. This technique can also be applied to other definitions of scale that are distributive over addition.

After fixing the set of basis functions $\{B^i\}$ and the set of scales $\{a_n\}$, we now pose a constrained optimization problem to determine \vec{c} and \vec{b} . Similarly to the filter design in the Euclidean space, the objective function should be application dependent, and could for example be a function of the frequency response. The constraints come from enforcing self-invertibility: we assume that the analysis and synthesis filters are identical and optimize the cost function under the invertibility constraints of Corollary 4.2. Since we cannot have more constraints than variables, self-invertibility cannot be achieved for more degrees than the number of basis functions and scales. We will discuss several specific examples of the objective function in Section VI.

The quadratic penalty method [3] is effective in solving this optimization problem with nonconvex constraints by incorporating the constraints into the objective function and solving the resulting unconstrained optimization problem using nonlinear least squares minimization

$$E(\vec{c}, \vec{b}) = \lambda \sum_l \|H_{\vec{c}, \vec{b}}(l) - 1\|^2 + \sum_{n=1}^N \left\| \vec{\Gamma}_n(\tilde{h}_{a_n}(\vec{c}, \vec{b})) \right\|^2. \quad (18)$$

The first term corresponds to the invertibility conditions by constraining the frequency response $H_{\vec{c}, \vec{b}}(l)$ ((9)) to be 1. The sum in the first term is over the desired invertibility range. Note that \vec{c} and \vec{b} completely determine the first term since they define the filter bank $\{\tilde{h}_{a_n}, h_{a_n}\}$ via (15) and (17). The vector-valued functions $\vec{\Gamma}_n$'s in the second term reflect the (un)desirable properties of the filter bank. For example, $\vec{\Gamma}$ could be the discrete second derivatives of the spherical harmonic coefficients of the filters. Minimizing the norms of $\vec{\Gamma}_n$'s would then lead to smooth filters. We repeat the optimization procedure while increasing the weight λ of the constraints and using the solution corresponding to the previous value of λ as the starting point for the next iteration, until convergence to a local minimum of the original optimization problem. Our implementation uses the Matlab's `lsqnonlin` function to minimize the objective function at each iteration.

VI. EXPERIMENTS

In this section, we demonstrate the optimization procedure formulated in the previous section. We demonstrate the construction of both self-invertible spherical wavelets and spherical steerable pyramids. Similar to the Euclidean domain, we define a spherical wavelet transform to be the decomposition of a spherical signal into component signals at different scales, i.e., employing axisymmetric filter kernels. On the other hand, we reserve the term spherical steerable pyramid transform for the decomposition of a spherical signal into component signals at different scales and orientations, i.e., using nonaxisymmetric filter kernels. We note that in some literature [1], [4], [29], the term ‘‘spherical wavelets’’ includes spherical steerable pyramids.

A. Spherical Wavelets

In designing axisymmetric wavelets, we limit our set of basis functions $\{B^i\}$ to be the first hundred spherical harmonics of

order 0, since the spherical harmonic coefficients of axisymmetric functions are zero for orders other than 0.

We define the set of scales to be $a = \{2^{-n/3}\}$, $n = -6, -5, \dots, 2, 3$, with $a = 1$ ($n = 0$) corresponding to the undilated template. We use S2kit [15]³ to create a table of the spherical harmonic coefficients of $D_a Y_l^0$ for $l = 0, \dots, 99$. As mentioned earlier, extreme stereographic dilation and shrinking of spherical harmonics can result in high frequencies. We find the first 600 order 0 spherical harmonic coefficients of each dilated spherical harmonic (a dilated axisymmetric function remains axisymmetric). We verify that for $a = 4$ ($n = -6$) and $a = 0.5$ ($n = 3$), $|[D_a Y_{99}^0]^{599,0}| / \max_l |[D_a Y_{99}^0]^{l,0}| < 10^{-12}$.

For axisymmetric filters, we can use the fast spherical convolution [11] to compute the forward convolution (3). We quote the results here for completeness

$$[y(\beta, \alpha)]^{l,m} = \sqrt{\frac{4\pi}{(2l+1)}} x^{l,m} \tilde{h}^{l,0*}. \quad (19)$$

Furthermore, as noted before, the inverse convolution (4) between two spherical images is the same as the definition of convolution by Driscoll and Healy [11] (except for a conjugation)

$$[\hat{x}(\theta, \phi)]^{l,m} = 2\pi \sqrt{\frac{4\pi}{(2l+1)}} y^{l,m} h^{l,0}. \quad (20)$$

Because we seek a multiscale decomposition of the original spherical signal, we would like the filter at each scale to act as a bandpass filter. Similar to the Euclidean domain, we require a residual lowpass filter to ensure that the combined wavelet and the lowpass filter bank is invertible up to a particular degree. Since we only use the first 100 spherical harmonics as our basis, the frequency response of $\tilde{h}_{a=1}(\theta, \phi)$ will be zero for all degrees higher than 99. If we also penalize the magnitude of the leading spherical harmonic coefficients of $\tilde{h}_{a=1}(\theta, \phi)$, the frequency response of $\tilde{h}_{a=1}(\theta, \phi)$ will be zeros at both ends, i.e., it will serve as a bandpass filter. To satisfy the self-invertibility conditions, the solution cannot be identically zero, but must rise to a peak somewhere in the middle of the frequency range.

We also penalize the second derivatives of the filters' frequency responses and spherical harmonic coefficients to force the filters to be relatively smooth and to reduce ringing. Note that the second derivatives are discrete since spherical harmonic degrees are discrete. We can induce a sharper cutoff frequency by penalizing the magnitude of the combined frequency response above a cutoff degree L_c . In addition, we fix the amplitude scaling factors b_k 's (15) to be the same. While allowing the b_k 's to take on different values provides the optimization procedure more flexibility in finding a set of desired filters, we find that in practice, having the same b for all the scales results in the frequency responses of the filters at all scales having comparable amplitude. Once again, we note that the energies of the filters at different scales will be different because of Corollary 4.2.

Fig. 5(a) illustrates the frequency response of a ten-scale wavelet filter bank ($a = \{2^{-n/3}\}$, $n = -6, -5, \dots, 2, 3$)

³We use the matlab interface from the publicly available Yet Another Wavelet Toolbox (YAWTb): <http://rhea.tele.ucl.ac.be/yawtb/>

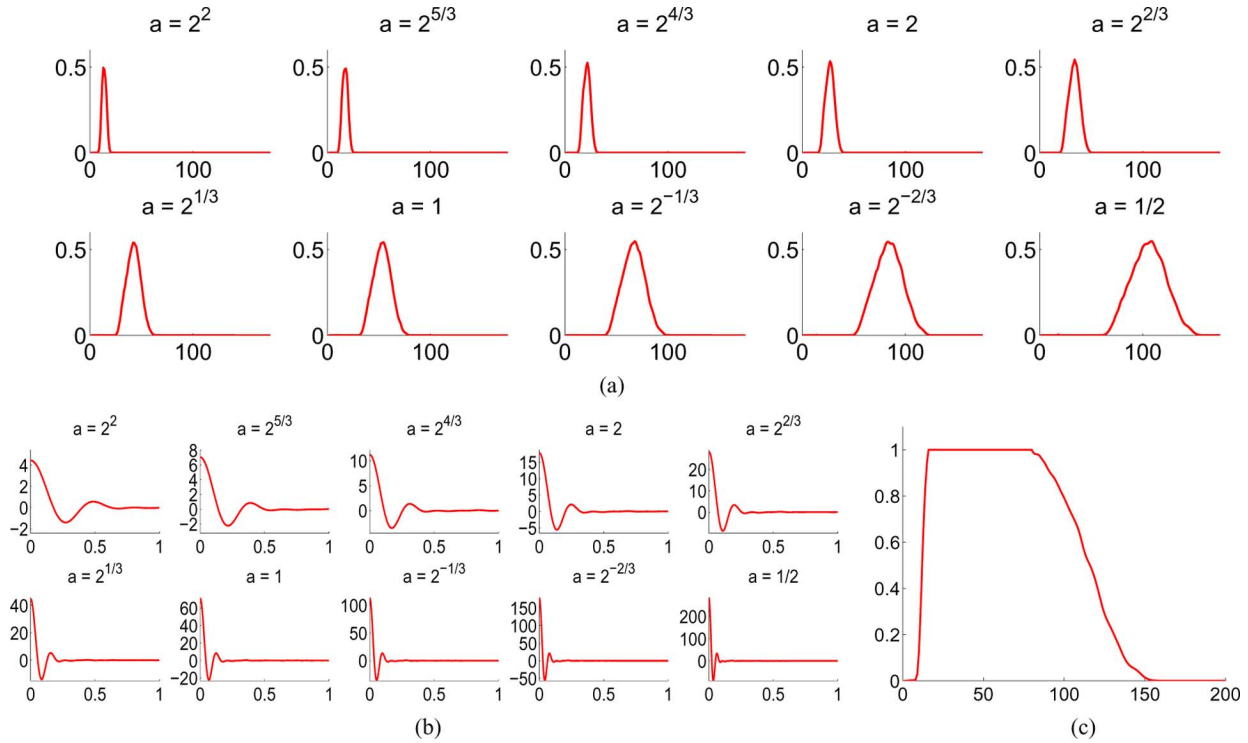


Fig. 5. Ten-scale wavelet filter bank obtained by imposing invertibility from degree 15 to 79 and a combined frequency response cutoff at $L_c = 150$. (a) Frequency response of individual filters. (b) Individual filters in the spatial domain ($0 \leq \theta \leq 1$ radian). The second peak indicates ringing. (c) Combined frequency response, $L_c = 150$.

obtained through our optimization procedure. Invertibility is enforced from degree 15 to 79. Furthermore, we impose a quadratic penalty on the magnitude of the combined frequency response for degrees above $L_c = 150$. The combined frequency response of the filters is shown in Fig. 5(c).

Because the filters are axisymmetric, we can plot the filters in the image domain as a function of θ [Fig. 5(b)]. The existence of a second peak after the peak at $\theta = 0$ (north pole) indicates ringing. When we vary the cutoff frequency penalty, we can trade off the amount of ringing for the slope of the cutoff. For example, Fig. 6(a) shows the combined frequency response of a wavelet filter bank obtained by penalizing the magnitude of the combined frequency response for degrees above $L_c = 100$. Notice the combined frequency response drops rapidly after degree 79. However, this results in increased ringing. If we measure ringing by the ratio of the second maxima to the maxima at the north pole, we can measure the tradeoff between ringing and the cutoff frequency, as shown in Fig. 6(b).

As a verification, we apply the resulting filters to a high resolution world elevation map (size $> 1000 \times 1000$). We bilinearly interpolate the map onto the S2kit [15] grid and convert the spatial image into spherical harmonics using S2kit. This is followed by the inverse transform using S2kit to obtain a final spherical image. The spherical harmonics and the final spherical image become the “ground truth” for measuring invertibility. Lower resolution versions of the elevation map are obtained by truncating the higher degrees spherical harmonics. It is necessary to use the spherical image obtained from inverting the spherical harmonics rather than from the linearly interpolated image because the linearly interpolated image might not be samples




Fig. 6. (a) Combined frequency response of the filters obtained when the combined frequency response cutoff is set to $L_c = 100$. Note the sharper cutoff obtained. However, this is at the expense of ringing. (b) Plot of ringing versus cutoff frequency L_c . Ringing is defined to be the ratio of the second peak to the maximum peak at the lowest scale, $a = 4$. (a) Combined frequency response, $L_c = 100$. (b) Ringing versus L_c .

from a bandlimited signal. In that case, the sampling theorem of Driscoll and Healy [11] does not hold since the quadrature knots and weights were computed assuming samples from a bandlimited signal. This is in contrast to the Euclidean case, where the forward and inverse discrete Fourier transform always return the same results. As a result, the fast spherical harmonic transform of S2kit is not completely invertible for nonbandlimited signals and is instead an approximate fit of the linearly interpolated data. The final spherical image and the spherical harmonics are, however, exact inverses of each other.

We convolve the wavelet filter bank of Fig. 5(a) with the world elevation map [Fig. 7(a)]. The results for four scales are shown in Fig. 7(b)–(e). Upon reconstruction using (20), we find that $\max_{l,m} |(\hat{x}^{l,m} - x^{l,m})/x^{l,m}| = 1.3 \times 10^{-6}$ for degrees between 15 to 79 inclusive. The reconstruction error obtained is




Fig. 7. Outputs of the analysis filter bank of Fig. 5 applied to the world elevation map. Only four scales are shown. (a) Original Image. (b) $a = 4$. (c) $a = 2$. (d) $a = 1$. (e) $a = 0.5$.




Fig. 8. (a) Combined frequency response of a four-scale wavelet filter bank obtained by our optimization procedure. $a = \{4, 2, 1, 0.5\}$. Invertibility is imposed from degree 10 to 89. Combined frequency response cutoff is set to $L_c = 150$. (b) Plot of ringing versus cutoff frequency L_c . (a) Combined frequency response, $L_c = 150$. (b) Ringing versus L_c .

expected since the quadratic penalty optimization method from the previous section is performed until $\sum_l \|H_{\vec{c}, \vec{b}}(l) - 1\|^2$ [in (18)] is in the order of 10^{-6} . The reconstruction errors can be made arbitrarily small by running the quadratic penalty method with increasingly higher weights on the invertibility constraints.

To demonstrate that the optimization procedure is stable across different settings of parameters, we show a second example where we optimize for a four-scale wavelet filter bank ($a = \{4, 2, 1, 0.5\}$). We enforce invertibility from degree 10 to 89 and apply a quadratic penalty on the magnitude of the combined frequency response for degrees above $L_c = 150$. The combined frequency response of the resultant filter bank is shown in Fig. 8(a). Once again, by varying the cutoff frequency threshold, we can obtain a tradeoff between ringing and sharpness of the cutoff [see Fig. 8(b)]. We apply the filter bank to the world elevation map [see Fig. 9] and find that invertibility is indeed obtained for degrees between 10 and 89 inclusive. Notice that there is significantly less ringing artifacts than in Fig. 7 as predicted by our measure of ringing at $L_c = 150$ [compare Figs. 6(b) and 8(b)].

B. Spherical Steerable Pyramid

Just like the Euclidean domain [14], it can be shown that there is a direct tradeoff between angular resolution and steerability of oriented (nonaxisymmetric) filters on the sphere [29], i.e., filters that have higher angular resolving power requires a bigger set of “steering” basis filters. This can also be seen from Theorem 4.3.




Fig. 9. Convolution outputs obtained by applying the analysis filter bank of Fig. 8(a) to the world elevation map of Fig. 7(a). Notice that there is less ringing artifacts than in Fig. 7 because ringing is lower in the four-scale filter bank than in the ten-scale filter bank when L_c is set to 150. (a) $a = 4$. (b) $a = 2$. (c) $a = 1$. (d) $a = 0.5$.

Filters with higher angular resolving power require spherical harmonics of higher orders, translating to more samples needed to satisfy the conditions of the generalized sampling theorem.

In our experiments, we limit our set of basis functions $\{B^i\}$ to be the first two hundred spherical harmonics of order $+1$ and -1 . We note that we can increase our angular power by using higher orders, but this decreases the steerability of our filters. By considering only real filters, we can avoid working directly with the order -1 spherical harmonics, since their coefficients are effectively constrained by those of the order $+1$ spherical harmonics (see Appendix A). For convenience, we further constrain the coefficients of the order $+1$ spherical harmonics to be real.

We define the set of scales to be $a = \{2^{-n/2}\}$, $n = -4, \dots, 4$, with $a = 1$ corresponding to the undilated template. Once again, we use S2kit [15] to create a table of the spherical harmonic coefficients of $D_a Y_l^1$ for $l = 1, \dots, 200$. We find the first 999 order 1 spherical harmonic coefficient of

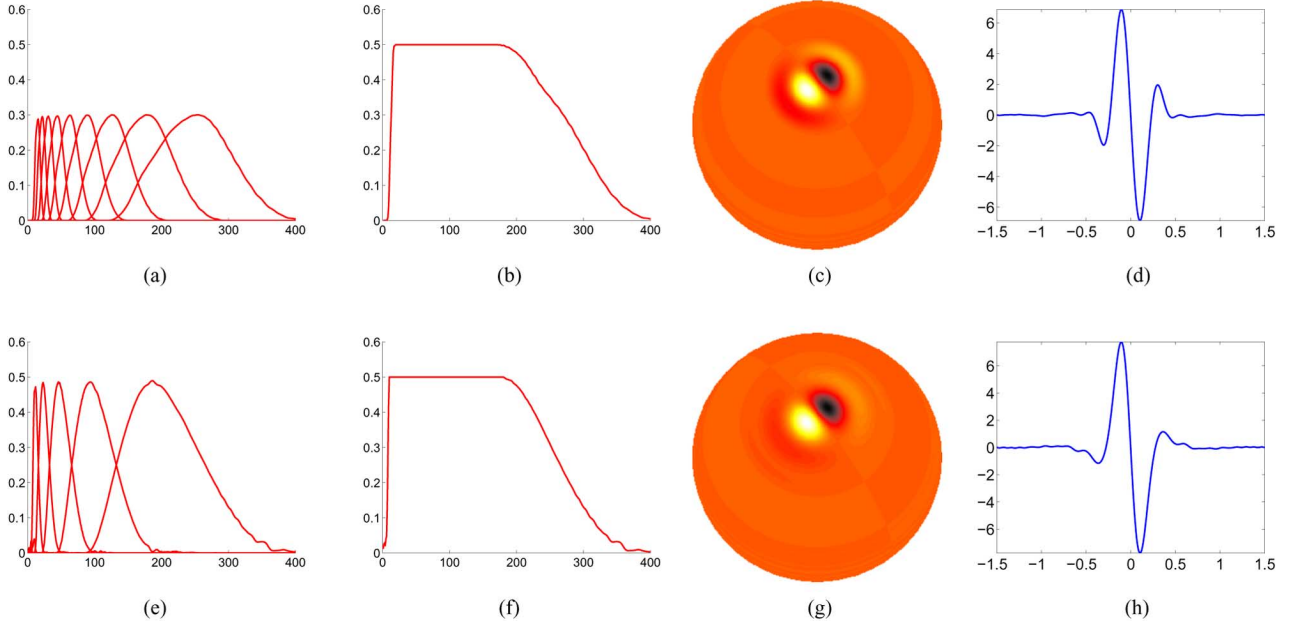


Fig. 10. (a)–(d) Nine-scale steerable pyramid ($a = \{2^{-n/2}\}$, $n = -4, \dots, 4$) obtained by imposing invertibility from degree 20 to 170. Note that the frequency response is equal to 0.5 in the invertibility range because we only plot the frequency response contributed by the order +1 harmonics. (e)–(h) Five-scale steerable pyramid ($a = \{4, 2, 1, 0.5, 0.25\}$) obtained by imposing invertibility from degree 10 to 180. Note that the frequency response is equal to 0.5 in the invertibility range because we only plot the frequency response contributed by the order +1 harmonics. (a) Individual frequency response. (b) Combined frequency response. (c) $h_{a=4}(\theta, \phi)$. (d) Plot of $h_{a=4}(\theta, \phi)$ ($-1.5 \leq \theta \leq 1.5$ radian). (e) Individual frequency response. (f) Combined frequency response. (g) $h_{a=4}(\theta, \phi)$. (h) Plot of $h_{a=4}(\theta, \phi)$ ($-1.5 \leq \theta \leq 1.5$ radian).

each dilated spherical harmonic (the order of a spherical function does not change under dilation). We verify that for $a = 4$ and $a = 0.25$, $|[D_a Y_{200}^1]^{999,1}|/\max_l |[D_a Y_{200}^1]^{l,1}| < 10^{-12}$.

Similar to the previous subsection, we penalize the magnitude of the leading coefficients of $h_{a=1}(\theta, \phi)$. We also penalize the second derivatives of the filters' frequency responses and spherical harmonic coefficients. Finally, we fix the amplitude scaling factors b_k 's at all scales to be the same.

Fig. 10(a) and (b) illustrates the frequency response of a nine-scale steerable pyramid ($a = \{2^{-n/2}\}$, $n = -4, \dots, 4$) obtained through our optimization procedure. Invertibility is enforced from degree 20 to 170. Note that the frequency response is equal to 0.5 in the invertibility range because we only plot the frequency response contributed by the order +1 harmonics. Fig. 10(c) shows $h_{a=4}(\theta, \phi)$ as a spherical image. Note that it looks like a derivative of Gaussian. We can also quantify ringing by plotting $h_{a=4}(\theta, \phi)$ as a function of θ while fixing ϕ to correspond to the great circle passing through the maxima and minima of the filter [Fig. 10(d)].

Similarly, Fig. 10(e)–(f) shows the frequency response of a five-scale steerable pyramid ($a = \{4, 2, 1, 0.5, 0.25\}$) obtained through our optimization procedure. Invertibility is enforced from degree 10 to 180. Fig. 10(g) shows $h_{a=4}(\theta, \phi)$ as a spherical image and Fig. 10(h) is a plot of $h_{a=4}(\theta, \phi)$ as a function of θ by fixing ϕ .

C. Denoising

We now employ the steerable pyramid from Fig. 10(e)–(f) in a denoising experiment. We emphasize that this toy example

only serves to illustrate the use of the self-invertible steerable pyramid and is not meant to be the model on how wavelet or steerable pyramid denoising should be done. We first obtain a full pass steerable pyramid filter bank by computing a residual lowpass and highpass filter bank so that the combined filter bank is invertible up to degree 300. By constraining the analysis and synthesis filters to be identical and axisymmetric, the lowpass and highpass filters are uniquely determined.

We use the world elevation map truncated to a maximum degree of 300. A Gaussian noise map is first produced in the spatial domain on the S2kit grid. The forward and inverse spherical harmonic transform is then performed with S2kit [15]. As noted earlier, the resulting noise image map from the inverse spherical harmonic transform will be different from the original noise map, because the fast spherical harmonic transform [11] is not invertible for nonbandlimited signals. Interestingly, noise near the poles tends to be suppressed under the forward and inverse spherical harmonic transform (not shown), possibly related to the concentration of samples near the poles on a latitude-longitude grid. The noise map obtained from the inverse spherical harmonic transform is then added to the world map. Fig. 11(a)–(d) display world elevation maps with different peak signal-to-noise ratios (PSNR).

We adapt the wavelet shrinkage technique pioneered by Donoho and Johnstone [10] to our denoising problem.

- 1) We perform continuous spherical convolution between the noisy world elevation map and the analysis filters of the steerable pyramid using the algorithm in [28].
- 2) We sample the transform coefficients using the inverse FFT, by using the uniform sampling grid of the second




Fig. 11. (a)–(d) Corrupted world elevation maps. (e)–(h) Corresponding denoised world elevation maps, i.e., (e) shows the denoised version of (a), and so on. (a) PSNR = 27.923 dB. (b) PSNR = 21.943 dB. (c) PSNR = 15.9035 dB. (d) PSNR = 9.8804 dB. (e) PSNR = 31.578 dB. (f) PSNR = 27.74 dB. (g) PSNR = 24.16 dB. (h) PSNR = 21.0649 dB.

quadrature rule (Appendix D) and the fact that the continuous spherical convolution in step 1 is computed by representing functions on $SO(3)$ as sum of complex exponentials.

3) We threshold the sampled transform coefficients

$$y_{njks} \leftarrow \text{sign}(y_{njks}) \max(|y_{njks}| - t\sigma_n, 0) \quad (21)$$

where for brevity, we denote the sampled transform coefficients $y_n(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n})$ by y_{njks} . The soft threshold (shrinkage) sets to zero for y_{njks} below $t\sigma_n$ in absolute value and pulls other data towards the origin by an amount $t\sigma_n$. The noise standard deviation σ_n can be determined empirically or by simulations [26]. In our toy example, we find that σ_n consistently doubles from one steerable pyramid scale to the finer one, except for the scale corresponding to the highpass filter. t is a user-defined parameter. In our experiments, we vary t from 0.5 to 5. We note that in the SureShrink [10] algorithm, the parameter t is found automatically once the noise standard deviation σ_n is known. Unfortunately, naive application of SureShrink gives poor results. One reason is that the coefficients of the overcomplete steerable pyramid is correlated, and, hence, the noise is no longer independently and identically distributed (i.i.d) in the steerable pyramid domain.

4) We utilize (6) and (7) to reconstruct the denoised image.

Details of the computation are found in Appendix E.

Recall that $L_n(O_n)$ is the maximum degree (order) of the n th synthesis filter (and the analysis filter since they are identical). Then, step 1 of the above procedure roughly requires computational complexity of $O(L_n^3 O_n)$. Step 2 requires $O(L_n \log L_n)$ computations. Step 3 requires $O(L_n^2 O_n)$ computations. Step 4 is the bottleneck, requiring $O(L_n^4 O_n^2)$ computations. We are currently working on faster algorithms and quadrature rules for step 4.

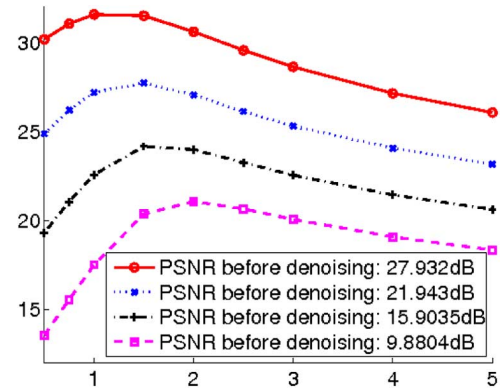


Fig. 12. PSNR of denoised images as t is varied from 0.5 to 5.

Fig. 12 shows a plot of the PSNR of the denoised images as we vary t from 0.5 to 5. In general, the noisier the image before denoising, the greater the improvement. In particular, there is an improvement of more than 10dB for the noisiest map. The best results are obtained for t between 1 and 2. The denoised maps corresponding to the best t 's are shown in Fig. 11(e)–(f). Artifacts become more noticeable for the noisier images. These artifacts are likely the by-products of ringing.

VII. DISCUSSION AND CONCLUSION

In this paper, we present theoretical conditions for the invertibility of filter banks under continuous convolution on the 2-Sphere. We discretize the results using quadrature, thus obtaining a generalized sampling theorem. We propose a general procedure for constructing invertible filter banks and demonstrate the procedure by generating self-invertible spherical wavelets and steerable pyramids. Finally, we provide an analysis of the computational complexity of the filtering framework.

Nonlinear dilation of functions on the sphere remains difficult to work with. While we circumvent the problem by using the distributive property of stereographic dilation, the spherical harmonic coefficients table can take up a substantial amount of space. More efficient methods are, therefore, needed. Other definitions of dilation (such as those defined in the frequency domain) might also fit better into the computational framework. More work is needed to understand the space of invertible and self-invertible filter banks. As seen in our experiments, there is an implicit tradeoff between the sharpness of the frequency response of the filters and ringing. It will be useful to formulate an objective function that directly trades off between ringing and the sharpness of the frequency response. The bandlimited spatially-concentrated eigenfunctions proposed by Simons *et al.* [24] could prove to be a useful set of basis functions for the optimization problem in our framework.

This paper introduces theoretical results on invertibility and sampling, and represents a step towards a general framework for filter design on the 2-Sphere. Just as wavelets and steerable pyramids have been useful for the processing and analysis of planar images, we are optimistic that future work will lead to similar applications on the sphere.

APPENDIX A SPHERICAL HARMONICS BASICS

Here, we review useful facts on the spherical harmonics.

Spherical Harmonics: The spherical harmonics Y_l^m are defined in terms of the associated Legendre polynomials P_l^m . For a given degree $l \geq 0$ and order $|m| \leq l$, $0 \leq \theta \leq \pi$ and $0 \leq \phi \leq 2\pi$

$$Y_l^m(\theta, \phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos \theta) e^{im\phi} \quad (22)$$

where, for $m \geq 0$ and $|x| < 1$

$$P_l^m(x) = \frac{(-1)^m}{2^l l!} (1-x^2)^{m/2} \frac{d^{l+m}}{dx^{l+m}} (x^2-1)^l \quad (23)$$

$$P_l^{-m}(x) = (-1)^m \frac{(l-m)!}{(l+m)!} P_l^m(x). \quad (24)$$

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

$$Y_l^m(\theta, \phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} \frac{(-1)^m}{2^l l!} (1-\cos^2 \theta)^{m/2} \times \left. \frac{d^{l+m}}{dx^{l+m}} (x^2-1)^l \right|_{x=\cos \theta} e^{im\phi} \quad (25)$$

$$Y_l^{-m}(\theta, \phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} \frac{1}{2^l l!} (1-\cos^2 \theta)^{m/2} \times \left. \frac{d^{l+m}}{dx^{l+m}} (x^2-1)^l \right|_{x=\cos \theta} e^{-im\phi} \quad (26)$$

$$= (-1)^m Y_l^{m*}(\theta, \phi). \quad (27)$$

Rotation of Spherical Harmonics on the Sphere: Under rotation, each spherical harmonic of degree l is transformed into

a linear combination of spherical harmonics of the same degree but possibly different orders. In particular, if we parametrize our rotation by the three Euler angles, α, β, γ , and rotate our original function f , it can be shown that

$$[D(\alpha, \beta, \gamma)f]^{l,m} = \sum_{m'=-l}^l D_{mm'}^l(\alpha, \beta, \gamma) f^{l,m'} \quad (28)$$

where $D_{mm'}^l(\alpha, \beta, \gamma)$ is the Wigner-D function [20]. We can further decompose $D_{mm'}^l(\alpha, \beta, \gamma)$ as follows:

$$D_{mm'}^l(\alpha, \beta, \gamma) = e^{-im\alpha} d_{mm'}^l(\beta) e^{-im'\gamma} \quad (29)$$

where $d_{mm'}^l(\beta)$ is the Wigner-d function and is real [20]

$$d_{mm'}^l(\beta) = \sum_j (-1)^{j-m'+m} \times \frac{\sqrt{(l+m')!(l-m')!(l+m)!(l-m)!}}{(l+m'-j)!j!(l-j-m)!(j-m'+m)!} \times \left(\cos \frac{\beta}{2}\right)^{2l-2j+m'-m} \left(\sin \frac{\beta}{2}\right)^{2j-m'+m}. \quad (30)$$

The sum is over all j such that none of the denominator terms with factorials is negative. This reflects the fact that only rotations about the y -axis mixes orders.

By the Peter-Weyl theorem on compact groups [30]

$$\int_{SO(3)} D_{mn}^l(\rho) D_{m'n'}^{l'*}(\rho) d\rho = \frac{8\pi^2}{2l+1} \delta(l-l', m-m', n-n'). \quad (31)$$

By integrating out α and γ , we obtain

$$\int_{\beta} d_{mn}^l(\beta) d_{m'n'}^l(\beta) \sin \beta d\beta = \frac{2}{2l+1} \delta(l-l'). \quad (32)$$

We will also use the identity [28]

$$D(\alpha, \beta, \gamma) = D\left(\alpha + \frac{\pi}{2}, \frac{\pi}{2}, 0\right) D\left(\beta + \pi, \frac{\pi}{2}, \frac{\pi}{2} + \gamma\right). \quad (33)$$

APPENDIX B PROOF OF CONTINUOUS-INVERTIBILITY

Here, we prove Theorem 4.1 on continuous frequency response. We first note that by using Parseval's Theorem and substituting (28), we can re-write the output of the n th analysis filter as

$$y_n(\alpha, \beta, \gamma) = \int_{S^2} [D(\alpha, \beta, \gamma) \tilde{h}_n]^*(\theta, \phi) x(\theta, \phi) d\Omega \quad (34)$$

$$= \sum_{l'=0}^{\infty} \sum_{m'=-l'}^{l'} \left[\sum_{m''=-l'}^{l'} D_{m',m''}^{l'}(\alpha, \beta, \gamma) \tilde{h}_n^{l',m''} \right]^* \times x^{l',m'} \quad (35)$$

and the reconstructed image as

$$\begin{aligned} \hat{x}(\theta, \phi) &= \sum_{n=1}^N \int_{SO(3)} [D(\alpha, \beta, \gamma) h_n](\theta, \phi) y_n(\alpha, \beta, \gamma) d\rho \quad (36) \\ &\stackrel{(35)}{=} \sum_{n=1}^N \int_{SO(3)} [D(\alpha, \beta, \gamma) h_n](\theta, \phi) \\ &\quad \times \sum_{l'=0}^{\infty} \sum_{m'=-l'}^{l'} \left[\sum_{m''=-l'}^{l'} D_{m',m''}^{l'}(\alpha, \beta, \gamma) \tilde{h}_n^{l',m''} \right]^* \\ &\quad \times x^{l',m'} d\rho. \quad (37) \end{aligned}$$

Projecting $\hat{x}(\theta, \phi)$ onto the spherical harmonics basis, we obtain

$$\begin{aligned} \hat{x}^{l,m} &= \int_{S^2} \hat{x}(\theta, \phi) Y_l^{m*}(\theta, \phi) d\Omega \quad (38) \\ &\stackrel{(37)}{=} \sum_{n=1}^N \int_{SO(3)} \left[\int_{S^2} [D(\alpha, \beta, \gamma) h_n](\theta, \phi) Y_l^{m*}(\theta, \phi) d\Omega \right] \\ &\quad \times \sum_{l'=0}^{\infty} \sum_{m'=-l'}^{l'} \left[\sum_{m''=-l'}^{l'} D_{m',m''}^{l'}(\alpha, \beta, \gamma) \tilde{h}_n^{l',m''} \right]^* \\ &\quad \times x^{l',m'} d\rho \quad (39) \\ &\stackrel{(28)}{=} \sum_{n=1}^N \int_{SO(3)} \left[\sum_{m'''=-l}^l h_n^{l,m'''} D_{m,m'''}^l(\alpha, \beta, \gamma) \right] \\ &\quad \times \sum_{l'=0}^{\infty} \sum_{m'=-l'}^{l'} \left[\sum_{m''=-l'}^{l'} D_{m',m''}^{l'}(\alpha, \beta, \gamma) \tilde{h}_n^{l',m''} \right]^* \\ &\quad \times x^{l',m'} d\rho \quad (40) \\ &= \sum_{n=1}^N \sum_{m'''=-l}^l [h_n^{l,m'''}] \sum_{l'=0}^{\infty} \sum_{m'=-l'}^{l'} x^{l',m'} \\ &\quad \times \sum_{m''=-l'}^{l'} [\tilde{h}_n^{l',m''}]^* \int_{SO(3)} D_{m,m'''}^l(\alpha, \beta, \gamma) \\ &\quad \times D_{m',m''}^{l'}(\alpha, \beta, \gamma)^* d\rho \quad (41) \\ &\stackrel{(31)}{=} x^{l,m} \frac{8\pi^2}{2l+1} \left\{ \sum_{n=1}^N \sum_{m'''=-l}^l [h_n^{l,m'''}] [\tilde{h}_n^{l,m'''}]^* \right\}. \quad (42) \end{aligned}$$

APPENDIX C

PROOF OF DISCRETE-INVERTIBILITY

Here, we prove the generalized sampling theorem. Recall from Section III that $y_n(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n})$ are samples of $y_n(\alpha, \beta, \gamma)$ and that \tilde{L}_n (O_n) and L_n (O_n) are the highest nonzero harmonic degree (order) of \tilde{h}_n and h_n , respectively.

Since the output of the n th synthesis filter $\tilde{x}_n(\theta, \phi)$ is a linear combination of rotated versions of $h_n(\theta, \phi)$ (7), for $l > L_n$, $\hat{x}_n^{l,m} = 0$ and $\sum_{m''=-l}^l h_n^{l,m''} \tilde{h}_n^{l,m''*} = 0$ and we obtain the trivial result

$$\hat{x}_n^{l,m} = \frac{8\pi^2}{2l+1} x_n^{l,m} \sum_{m''=-l}^l h_n^{l,m''} \tilde{h}_n^{l,m''*} = 0. \quad (43)$$

We will now show that for $l \leq L_n$

$$\hat{x}_n^{l,m} = \frac{8\pi^2}{2l+1} x_n^{l,m} \sum_{m''=-l}^l h_n^{l,m''} \tilde{h}_n^{l,m''*}. \quad (44)$$

From (35)

$$\begin{aligned} y_n(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) \\ &= \sum_{l'=0}^{\tilde{L}_n} \sum_{m'=-l'}^{l'} \left[\sum_{m''=-l'}^{l'} D_{m',m''}^{l'}(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) \tilde{h}_n^{l',m''} \right]^* \\ &\quad \times x^{l',m'}. \quad (45) \end{aligned}$$

For $w_{j,s,k,n} = 4\pi^2 w_{s,n} / (J_n K_n)$ (as required by the theorem), the output of the n th synthesis filter becomes (46) and (47), shown at the bottom of the page. Projecting $\hat{x}_n(\theta, \phi)$ onto the spherical harmonics, for $l \leq L_n$ (and, thus, $|m| \leq L_n$), we have (48) and (49), shown at the bottom of the next page, where we have arranged the terms so that they look like the setup for Peter–Weyl Theorem, except we have summations instead of integrals. Let Φ be the last part of (49), and, thus, we have

$$\begin{aligned} \hat{x}_n^{l,m} &= \frac{4\pi^2}{J_n K_n} \sum_{m'''=-l}^l h_n^{l,m'''} \sum_{l'=0}^{\tilde{L}_n} \sum_{m'=-l'}^{l'} x^{l',m'} \\ &\quad \times \sum_{m''=-l'}^{l'} [\tilde{h}_n^{l',m''}]^* \Phi. \quad (50) \end{aligned}$$

$$\hat{x}_n(\theta, \phi) \stackrel{(6)}{=} \sum_{j=0}^{J_n-1} \sum_{s=0}^{S_n-1} \sum_{k=0}^{K_n-1} w_{j,s,k,n} [D(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) h_n](\theta, \phi) y_n(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) \quad (46)$$

$$\begin{aligned} &\stackrel{(45)}{=} \sum_{j=0}^{J_n-1} \sum_{s=0}^{S_n-1} \sum_{k=0}^{K_n-1} \frac{4\pi^2 w_{s,n}}{J_n K_n} [D(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) h_n](\theta, \phi) \\ &\quad \times \sum_{l'=0}^{\tilde{L}_n} \sum_{m'=-l'}^{l'} \left[\sum_{m''=-l'}^{l'} D_{m',m''}^{l'}(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) \tilde{h}_n^{l',m''} \right]^* x^{l',m'} \quad (47) \end{aligned}$$

To simplify Φ , we write

$$\Phi = \sum_{j=0}^{J_n-1} \sum_{s=0}^{S_n-1} \sum_{k=0}^{K_n-1} w_{s,n} D_{m,m''}^l(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) \times D_{m',m''}^{l'*}(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) \quad (51)$$

$$\stackrel{(29)}{=} \sum_{j=0}^{J_n-1} \sum_{s=0}^{S_n-1} \sum_{k=0}^{K_n-1} w_{s,n} e^{-im\alpha_{j,n}} d_{m,m''}^l(\beta_{s,n}) \times e^{-im''\gamma_{k,n}} e^{im'\alpha_{j,n}} d_{m',m''}^{l'}(\beta_{s,n}) e^{im''\gamma_{k,n}} \\ = \sum_{s=0}^{S_n-1} w_{s,n} d_{m,m''}^l(\beta_{s,n}) d_{m',m''}^{l'}(\beta_{s,n}) \\ \times \left[\sum_{j=0}^{J_n-1} e^{i(m'-m)\alpha_{j,n}} \right] \left[\sum_{k=0}^{K_n-1} e^{i(m''-m''')\gamma_{k,n}} \right]. \quad (52)$$

We note that $|m'| \leq l' \leq \tilde{L}_n$ and $|m| \leq L_n$, and, therefore, $-\tilde{L}_n - L_n \leq m' - m \leq \tilde{L}_n + L_n$. Since $\alpha_{j,n} = 2\pi j / (\tilde{L}_n + L_n + 1)$, $j = 0, 1, \dots, \tilde{L}_n + L_n$, we can conclude via the geometric series that

$$\sum_{j=0}^{J_n-1} e^{i(m'-m)\alpha_{j,n}} = \sum_{j=0}^{\tilde{L}_n+L_n} e^{i(m'-m)2\pi j / (\tilde{L}_n+L_n+1)} \\ = \begin{cases} \tilde{L}_n + L_n + 1, & \text{if } (m' - m) = 0 \\ 0, & \text{otherwise.} \end{cases} \quad (53)$$

Similarly, from (48), we observe that $|m''| \leq O_{\tilde{h}_n}$ and $|m'''| \leq O_n$. Using the same reasoning, we get

$$\sum_{k=0}^{K_n-1} e^{i(m''-m''')\gamma_{k,n}} \\ = \sum_{k=0}^{O_{\tilde{h}_n} + O_{h_n}} e^{i(m''-m''')2\pi k / (O_{\tilde{h}_n} + O_{h_n} + 1)} \\ = \begin{cases} O_{\tilde{h}_n} + O_{h_n} + 1, & \text{if } (m'' - m''') = 0 \\ 0, & \text{otherwise.} \end{cases} \quad (54)$$

Substituting into (52), we get

$$\Phi = \sum_{s=0}^{S_n-1} w_{s,n} d_{m,m''}^l(\beta_{s,n}) d_{m',m''}^{l'}(\beta_{s,n}) \\ \times \left[\sum_{j=0}^{J_n-1} e^{i(m'-m)\alpha_{j,n}} \right] \left[\sum_{k=0}^{K_n-1} e^{i(m''-m''')\gamma_{k,n}} \right] \quad (55)$$

$$= J_n K_n \sum_{s=0}^{S_n-1} w_{s,n} d_{m,m''}^l(\beta_{s,n}) d_{m',m''}^{l'}(\beta_{s,n}) \\ \times \delta(m - m') \delta(m'' - m''') \quad (56)$$

$$= J_n K_n \sum_{s=0}^{S_n-1} w_{s,n} d_{m,m''}^l(\beta_{s,n}) d_{m,m''}^{l'}(\beta_{s,n}) \\ \times \delta(m - m') \delta(m'' - m''') \quad (57)$$

$$\stackrel{(32)}{=} \frac{2J_n K_n}{2l+1} \delta(l - l') \delta(m - m') \delta(m'' - m''') \quad (58)$$

where in the third equality, $m''' = m''$, $m' = m$ 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 $\int_0^\pi d_{mm''}^{l'}(\beta) d_{mm''}^{l'}(\beta) \sin(\beta) d\beta$. We can now substitute (58) back into (50)

$$\hat{x}_n^{l,m} = \frac{4\pi^2}{J_n K_n} \sum_{m'''=-l}^l h_n^{l,m'''} \\ \times \sum_{l'=0}^{\tilde{L}_n} \sum_{m'=-l'}^{l'} x^{l',m'} \sum_{m''=-l'}^{l'} [\tilde{h}_n^{l',m''}]^* \Phi \quad (59)$$

$$= \frac{4\pi^2}{J_n K_n} \sum_{m'''=-l}^l h_n^{l,m'''} \sum_{l'=0}^{\tilde{L}_n} \\ \times \sum_{m'=-l'}^{l'} x^{l',m'} \sum_{m''=-l'}^{l'} [\tilde{h}_n^{l',m''}]^* \\ \times \frac{2J_n K_n}{2l+1} \delta(l - l') \delta(m - m') \delta(m'' - m''') \quad (60)$$

$$= \frac{8\pi^2}{2l+1} x^{l,m} \sum_{m''=-l}^l h_n^{l,m''} \tilde{h}_n^{l,m''*} \text{ for all } l \leq L_n \quad (61)$$

$$\hat{x}_n^{l,m} = \int_{S^2} \hat{x}_n(\theta, \phi) Y_l^{m*}(\theta, \phi) d\Omega \\ \stackrel{(28)}{=} \sum_{j=0}^{J_n-1} \sum_{s=0}^{S_n-1} \sum_{k=0}^{K_n-1} \frac{4\pi^2 w_{s,n}}{J_n K_n} \left[\sum_{m'''=-l}^l D_{m,m''}^l(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) h_n^{l,m'''} \right] \quad (48)$$

$$\times \sum_{l'=0}^{\tilde{L}_n} \sum_{m'=-l'}^{l'} \left[\sum_{m''=-l'}^{l'} D_{m',m''}^{l'}(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) \tilde{h}_n^{l',m''} \right]^* x^{l',m'} \\ = \frac{4\pi^2}{J_n K_n} \sum_{m'''=-l}^l h_n^{l,m'''} \sum_{l'=0}^{\tilde{L}_n} \sum_{m'=-l'}^{l'} x^{l',m'} \sum_{m''=-l'}^{l'} [\tilde{h}_n^{l',m''}]^* \\ \times \sum_{j=0}^{J_n-1} \sum_{s=0}^{S_n-1} \sum_{k=0}^{K_n-1} w_{s,n} D_{m,m''}^l(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) D_{m',m''}^{l'*}(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) \quad (49)$$

thus proving (44). Together with (43), we have

$$\begin{aligned}\widehat{x}^{l,m} &= \sum_{n=1}^N \widehat{x}_n^{l,m} \\ &= \frac{8\pi^2}{2l+1} x^{l,m} \sum_{n=1}^N \sum_{m''=-l}^l h_n^{l,m''} \widetilde{h}_n^{l,m''*} \text{ for all } l. \quad (62)\end{aligned}$$

APPENDIX D QUADRATURE RULES

In this appendix, we derive two different quadrature rules that satisfy the conditions of Theorem 4.3, namely, we show quadrature weights $w_{s,n}$ and knots $\beta_{s,n}$, such that for $l \leq L_n, l' \leq \widetilde{L}_n$

$$\begin{aligned}\int_0^\pi d_{mm'}^l(\beta) d_{mm'}^{l'}(\beta) \sin(\beta) d\beta \\ = \sum_{s=0}^{S_n-1} w_{s,n} d_{mm'}^l(\beta_{s,n}) d_{mm'}^{l'}(\beta_{s,n}). \quad (63)\end{aligned}$$

Quadrature Rule (1): We denote $f(\beta) = d_{mm'}^l(\beta) d_{mm'}^{l'}(\beta)$ and observe that $f(\beta)$ consists of a linear combination of even powers of $\cos(\beta/2)$ and $\sin(\beta/2)$ (see (30)). Therefore, if we make the substitution $u = \sin(\beta/2)$, and noting that $f(u)$ (where we are overloading f) is now a polynomial with maximum degree, $Q = 2(l+l') \leq 2(L_n + \widetilde{L}_n)$, we get

$$\begin{aligned}\int_0^\pi f(\beta) \sin \beta d\beta &= 2 \int_0^\pi f(\beta) \sin\left(\frac{\beta}{2}\right) \cos\left(\frac{\beta}{2}\right) d\beta \\ &= 4 \int_0^1 f(u) u du. \quad (64)\end{aligned}$$

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

$$\begin{aligned}\int_0^\pi f(\beta) \sin \beta d\beta &= 2 \int_{-1}^1 f\left(\frac{v+1}{2}\right) \frac{v+1}{2} dv \\ &= 2 \sum_{k=0}^{N-1} r_k \frac{v_k+1}{2} f\left(\frac{v_k+1}{2}\right) \quad (65)\end{aligned}$$

where r_k 's are defined to be the weights of the Gauss–Legendre quadrature on the interval $[-1, 1]$, and v_k 's correspond to the sampling knots [19]. The weights and abscissas can be found by standard algorithms (see for example [19]). In general, quadrature integration is exact up to polynomial powers $2N - 1$ where N is the number of samples. Because the integrand's highest polynomial power is $Q+1$, if $N \geq Q/2+1$ (or $N \geq \lceil Q/2 \rceil + 1$), then $2N - 1 \geq Q+1$, and, thus, the quadrature formula is exact.

From the substitution above, we have $\sin \beta_k/2 = v_k = (v_k+1)/2$ or $\beta_k = 2 \sin^{-1}((v_k+1)/2)$. In conclusion, for $N = \lceil Q/2 \rceil + 1$, we have $\int_0^\pi d_{mm'}^l(\beta) d_{mm'}^{l'}(\beta) \sin \beta d\beta = \sum_{s=0}^{N-1} w_s f(\beta_s)$, $w_s = r_s(v_s+1)$ and $\beta_s = 2 \sin^{-1}((v_s+1)/2)$, where r_s and v_s are the quadrature weights and knots of the Gauss–Legendre quadrature.

Quadrature Rule 2: We will derive another rule in this section, using the technique shown in [11]. However, 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

$$SQ(u) = \begin{cases} 1, & -\pi < u < \frac{-\pi}{2} \\ -1, & \frac{-\pi}{2} < u < 0 \\ 1, & 0 < u < \frac{\pi}{2} \\ -1, & \frac{\pi}{2} < u < \pi. \end{cases} \quad (66)$$

Projecting $SQ(u)$ onto the Fourier series basis, we get

$$\begin{aligned}\frac{1}{2\pi} \int_{-\pi}^\pi SQ(u) e^{-iku} du \\ = \frac{1}{2\pi} \left[\int_{-\pi}^{-\pi/2} e^{-iku} du - \int_{-\pi/2}^0 e^{-iku} du \right. \\ \left. + \int_0^{\pi/2} e^{-iku} du - \int_{\pi/2}^\pi e^{-iku} du \right] \quad (67)\end{aligned}$$

$$\begin{aligned}= \frac{i}{2\pi k} \left[e^{-iku} \Big|_{-\pi}^{-\pi/2} - e^{-iku} \Big|_{-\pi/2}^0 \right. \\ \left. + e^{-iku} \Big|_0^{\pi/2} - e^{-iku} \Big|_{\pi/2}^\pi \right] \quad (68)\end{aligned}$$

$$\begin{aligned}= \frac{i}{2\pi k} \left[(e^{ik\pi/2} - e^{ik\pi}) - (1 - e^{ik\pi/2}) \right. \\ \left. + (e^{-ik\pi/2} - 1) - (e^{-ik\pi} - e^{-ik\pi/2}) \right] \quad (69)\end{aligned}$$

$$= \frac{i}{\pi k} \left[e^{ik\pi/2} + e^{-ik\pi/2} - 1 - \frac{1}{2} e^{ik\pi} - \frac{1}{2} e^{-ik\pi} \right] \quad (70)$$

$$= \frac{i}{\pi k} \left[2 \cos\left(k\frac{\pi}{2}\right) - 1 - \cos(k\pi) \right]. \quad (71)$$

If $k = 4n$, we get $2 \cos(2n\pi) - 1 - \cos(4n\pi) = 0$

If $k = 4n + 1$, we get $2 \cos(2n\pi + \pi/2) - 1 - \cos(4n\pi + \pi) = 0$

If $k = 4n + 2$, we get $2 \cos(2n\pi + \pi) - 1 - \cos(4n\pi + 2\pi) = -4$

If $k = 4n + 3$, we get $2 \cos(2n\pi + 3\pi/2) - 1 - \cos(4n\pi + 3\pi) = 0$.

Therefore, the Fourier series for $SQ(u)$ is nonzero for $k = 4n + 2$, and is equal to $-4i/\pi k$, and we have

$$\begin{aligned}SQ(u) &= \sum_{p=-\infty}^{\infty} -\frac{4i}{\pi(4p+2)} e^{i(4p+2)u} \\ &= \sum_{p=-\infty}^{\infty} -\frac{2i}{\pi(2p+1)} e^{i(4p+2)u}. \quad (72)\end{aligned}$$

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, non-negative and bounded (once again, we are overloading f). We can express them as complex exponentials so that $f(\beta) = \sum_q b_q (e^{i\beta/2})^q$ where q can take on negative values, but it is still bounded: $|q| \leq 2(l+l') \leq 2(L_n + \widetilde{L}_n)$. 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

$$\begin{aligned}\int_0^\pi f(\beta) \sin(\beta) d\beta \\ = 2 \int_0^{\pi/2} f(\beta') \sin(2\beta') d\beta' \quad (73)\end{aligned}$$

$$\begin{aligned}
&= \frac{1}{2} \int_{-\pi}^{-\pi/2} f(\beta') \sin(2\beta') d\beta' \\
&\quad - \frac{1}{2} \int_{-\pi/2}^0 f(\beta') \sin(2\beta') d\beta' \\
&\quad + \frac{1}{2} \int_0^{\pi/2} f(\beta') \sin(2\beta') d\beta' \\
&\quad - \frac{1}{2} \int_{\pi/2}^{\pi} f(\beta') \sin(2\beta') d\beta' \quad (74) \\
&= \frac{1}{2} \int_{-\pi}^{\pi} f(\beta') \sin(2\beta') SQ(\beta') d\beta' \quad (75)
\end{aligned}$$

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), thus we only require p such that

$$|4p + 2| \leq Q + 2 \quad (76)$$

$$\Leftrightarrow -Q - 2 \leq 4p + 2 \leq Q + 2 \quad (77)$$

$$\Leftrightarrow -\frac{Q}{4} - 1 \leq p \leq \frac{Q}{4} \quad (78)$$

$$\Leftrightarrow -\left\lfloor \frac{Q}{4} \right\rfloor - 1 \leq p \leq \left\lfloor \frac{Q}{4} \right\rfloor. \quad (79)$$

Defining $\widetilde{SQ}(\beta') = \sum_{p=-\lfloor Q/4 \rfloor - 1}^{\lfloor Q/4 \rfloor} -(2i/(\pi(2p+1)))e^{i(4p+2)\beta'}$, we have

$$\begin{aligned}
&\int_0^{\pi} f(\beta) \sin(\beta) d\beta \\
&= \frac{1}{2} \int_{-\pi}^{\pi} f(\beta') \sin(2\beta') \widetilde{SQ}(\beta') d\beta' \quad (80) \\
&= \frac{1}{2} \sum_{q=-Q}^Q b_q \sum_{p=-\lfloor Q/4 \rfloor - 1}^{\lfloor Q/4 \rfloor} -\frac{2i}{\pi(2p+1)} \\
&\quad \times \int_{-\pi}^{\pi} e^{i\beta'q} \frac{e^{i2\beta'} - e^{-i2\beta'}}{2i} e^{i(4p+2)\beta'} d\beta'. \quad (81)
\end{aligned}$$

Let us define the highest exponential power to be B and notice that $B = Q + 2 + 4\lfloor Q/4 \rfloor + 2 = Q + 4\lfloor Q/4 \rfloor + 4$, while lowest exponential power corresponds to $-Q - 2 - 4\lfloor Q/4 \rfloor - 2 = -Q - 4\lfloor Q/4 \rfloor - 4 = -B$. Let N be the smallest integer such that $B < 4N$, where N is an integer. Therefore, $N = \lfloor Q/4 \rfloor + \lfloor Q/4 \rfloor + 1 + 1 = 2\lfloor Q/4 \rfloor + 2$

It is easy to verify the following identity:

$$\begin{aligned}
\frac{1}{4N} \sum_{k=-2N}^{2N-1} e^{2\pi ikl/(4N)} &= \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{i\beta' l} d\beta' \\
&= \begin{cases} 1, & \text{if } l = 0 \\ 0, & \text{otherwise} \end{cases} \quad \forall |l| < 4N. \quad (82)
\end{aligned}$$

Substituting the identity into (81), we get ($\beta' \rightarrow 2\pi k/(4N)$); see (83)–(87), shown at the bottom of the page, where the second

$$\begin{aligned}
&\frac{1}{2} \sum_{q=-Q}^Q b_q \sum_{p=-\lfloor Q/4 \rfloor - 1}^{\lfloor Q/4 \rfloor} -\frac{2i}{\pi(2p+1)} \frac{\pi}{2N} \\
&\quad \times \sum_{k=-2N}^{2N-1} e^{2\pi ikq/(4N)} \frac{e^{2\pi ik2/(4N)} - e^{-2\pi ik2/(4N)}}{2i} e^{2\pi ik(4p+2)/(4N)} \\
&= \frac{\pi}{4N} \sum_{k=-2N}^{2N-1} \sum_{q=-Q}^Q b_q e^{2\pi ikq/(4N)} \sin\left(\frac{4\pi k}{4N}\right) \\
&\quad \times \sum_{p=-\lfloor Q/4 \rfloor - 1}^{\lfloor Q/4 \rfloor} -\frac{2i}{\pi(2p+1)} e^{2\pi ik(4p+2)/(4N)} \quad (83)
\end{aligned}$$

$$= \frac{\pi}{4N} \sum_{k=-2N}^{2N-1} f\left(\beta'_k = \frac{2\pi k}{4N}\right) \sin\left(\frac{4\pi k}{4N}\right) \widetilde{SQ}\left(\frac{2\pi k}{4N}\right) \quad (84)$$

$$= \frac{\pi}{4N} \sum_{k=-2N}^{2N-1} f\left(\beta'_k = \frac{2\pi k}{4N}\right) \sin\left(\frac{\pi k}{N}\right) \widetilde{SQ}\left(\frac{\pi k}{2N}\right) \quad (85)$$

$$= \frac{\pi}{2N} \sum_{k=0}^{2N-1} f\left(\beta'_k = \frac{2\pi k}{4N}\right) \sin\left(\frac{\pi k}{N}\right) \widetilde{SQ}\left(\frac{\pi k}{2N}\right) \quad (86)$$

$$= \frac{\pi}{N} \sum_{k=0}^{N-1} f\left(\beta'_k = \frac{2\pi k}{4N}\right) \sin\left(\frac{\pi k}{N}\right) \widetilde{SQ}\left(\frac{\pi k}{2N}\right) \quad (87)$$

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 $\beta'_k = \beta_k/2 = 2\pi k/(4N)$, hence $\beta_k = \pi k/N$ for $k = 0, 1, \dots, N - 1$ corresponds to our quadrature knots, with quadrature weights, $w_k = \pi/N \sin(\pi k/N) \widetilde{SQ}(\pi k/(2N))$.

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

- 1) $\beta_s = 2\sin^{-1}((v_s+1)/2)$ and $w_s = r_s(v_s+1)$, $s = 0, 1, \dots, N-1$, where:
 - a) $N = \lceil Q/2 \rceil + 1$;
 - b) Q is the highest power of $d_{mm'}^l(\beta) d_{mm'}^{l'}(\beta)$ when viewed as a polynomial in $\sin(\beta/2)$;
 - c) r_s and v_s are the weights and nodes of the Gaussian-Legendre quadrature on the interval $[-1, 1]$.
- 2) $\beta_s = \pi s/N$ and $w_s = \pi/N \sin(\pi s/N) \widetilde{SQ}(\pi s/(2N))$, $s = 0, 1, \dots, N-1$, where:
 - a) $N = 2\lceil Q/4 \rceil + 2$;
 - b) Q is the highest power of $d_{mm'}^l(\beta) d_{mm'}^{l'}(\beta)$ when viewed as a polynomial in $e^{i\beta/2}$.

Note that $Q = 2(l+l') \leq 2(L_n + \widetilde{L}_n)$. The second set of quadrature knots is similar to that used in the fast Wigner-D transform (SOFT) [16], except that the north pole ($s = 0$) is removed from the SOFT discretization.

APPENDIX E

INVERSE CONVOLUTION WITH NON-AXISYMMETRIC FILTERS

For brevity, we denote $y_n(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n})$ by $y_{nj sk}$. Given $y_{nj sk}$, one can compute $\widehat{x}_n(\theta, \phi)$ by rewriting (6) and computing its spherical harmonics

$$\widehat{x}_n^{l,m} = \sum_{j,s,k} \frac{4\pi^2}{JK} w_{s,n} y_{nj sk} [D(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) h_n]^{l,m}. \quad (88)$$

Since

$$[D(\alpha, \beta, \gamma) h_n]^{l,m} \stackrel{(33;28)}{=} \sum_{M,M'} D_{mM}^l \left(\alpha + \frac{\pi}{2}, \frac{\pi}{2}, 0 \right) \times D_{MM'}^l \left(\beta + \pi, \frac{\pi}{2}, \frac{\pi}{2} + \gamma \right) h_n^{l,M'} \quad (89)$$

for a particular $(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n})$, we can compute $[D(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) h_n]^{l,m}$ for all degrees l and orders m with $O(L^3 O_n)$ computations. This can be done by computing and storing $d_{mm'}^l(\pi/2)$ in advance.

Recall that $L_n(O_n)$ is the maximum degree (order) of the n th synthesis filter (and the analysis filter since they are identical). Since $J_n = O(L_n)$, $S_n = O(L_n)$ and $K_n = O(O_n)$, the brute force method of computing $\widehat{x}_n^{l,m}$ for all degrees and orders yields a $O(L^5 O_n^2)$ algorithm.

However, a faster algorithm can be achieved.

- 1) $x_n^{l,m} \leftarrow 0$ for all l, m .
- 2) For each β_s and γ_k :
 - a) compute $[D(0, \beta_{s,n}, \gamma_{k,n}) h_n]^{l,m}$ for all l, m ;
 - b) for each $\alpha_{j,n}$:

- i) compute $[D(0, 0, \alpha_{j,n}) D(0, \beta_{s,n}, \gamma_{k,n}) h_n]^{l,m}$ for all l, m ;
- ii)

$$x_n^{l,m} \leftarrow x_n^{l,m} + \frac{4\pi^2}{JK} w_{s,n} y_{nj sk} [D(\alpha_{j,n}, \beta_{s,n}, \gamma_{k,n}) h_n]^{l,m}.$$

Since step 2a dominates step 2b, the runtime complexity is $O(L_n^4 O_n^2)$.

ACKNOWLEDGMENT

The authors would like to thank M. Tappen for discussions on optimization procedures, B. Freeman and T. Adelson for discussions on self-invertibility, V. Goyal for discussions on norm preserving dilations in Euclidean wavelets and wavelet denoising, and F. Durand and B. Horn for reading earlier drafts of this paper. The authors would also like to thank the reviewers for their many useful suggestions and pointers to references. T. Yeo would like to thank C.-L. Cheng for discussions on spherical harmonics and Wigner-D functions.

REFERENCES

- [1] J.-P. Antoine and P. Vandergheynst, "Wavelets on the 2-Sphere: A Group-theoretical approach," *Appl. Comput. Harmon. Anal.*, vol. 7, pp. 262–291, 1999.
- [2] C. Armitage and B. D. Wandelt, "Deconvolution map-making for cosmic microwave background observations," *Phys. Rev. D* 70, vol. 123007, pp. 123013–123013, 2004.
- [3] D. P. Bertsekas, *Nonlinear Programming*. Nashua, NH: Athena Scientific, 1998.
- [4] I. Bogdanova, P. Vandergheynst, J.-P. Antoine, L. Jacques, and M. Morvidone, "Stereographic wavelet frames on the sphere," *Appl. Comput. Harmon. Anal.*, vol. 19, pp. 223–252, 2005.
- [5] C. H. Brechbuhler, G. Gerig, and O. Kubler, "Parametrization of closed surfaces for 3-D shape description," *Comput. Vis. Image Understand.*, vol. 61, no. 2, pp. 154–170, 1994.
- [6] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex Fourier series," *Math. Comput.*, vol. 19, no. 90, pp. 297–301, 1965.
- [7] I. Daubechies, "Ten lectures on wavelets," *SIAM*, 1992.
- [8] L. Demanet and P. Vandergheynst, "Gabor wavelets on the sphere," *Proc. SPIE*, vol. 5207, pp. 208–215, 2003.
- [9] M. N. Do and M. Vetterli, "The finite ridgelet transform for image representation," *IEEE Trans. Image Process.*, vol. 12, no. 1, pp. 16–28, Jan. 2003.
- [10] D. L. Donoho and I. M. Johnstone, "Adapting to unknown smoothness via wavelet shrinkage," *J. Amer. Statist. Assoc.*, vol. 90, no. 432, pp. 1200–1224, 1995.
- [11] J. R. Driscoll and D. M. Healy, "Computing Fourier transforms and convolutions on the 2-Sphere," *Adv. Appl. Math.*, vol. 15, pp. 202–250, 1994.
- [12] W. Freeden and U. Windheuser, "Spherical wavelet transform and its discretization," *Adv. Comput. Math.*, vol. 5, pp. 51–94, 1996.
- [13] W. Freeden, T. Gervens, and M. Schneider, *Constructive Approximation on the Sphere: With Applications to Geomathematics*. Oxford, U.K.: Clarendon, 1998.
- [14] W. T. Freeman and E. W. Adelson, "The design and use of steerable filters," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 9, pp. 891–906, Sep. 1991.
- [15] P. J. Kostelec and D. N. Rockmore, A Lite Version of Spharmonic Kit [Online]. Available: <http://www.cs.dartmouth.edu/geelong/sphere/>
- [16] P. J. Kostelec and D. N. Rockmore, "Fits on the rotation group," in *Proc. Santa Fe Institute Working Papers Series Paper 03-11-060*, 2003 [Online]. Available: <http://www.cs.dartmouth.edu/geelong/>
- [17] H. N. Mhaskar, F. J. Narcowich, J. Prestin, and J. D. Ward, "Polynomial frames on the sphere," *Adv. Comput. Math.*, vol. 13, pp. 387–403, 2000.
- [18] A. Papoulis, "Generalized sampling expansion," *IEEE Trans. Circuits Syst.*, vol. 24, no. 11, pp. 652–654, Nov. 1977.
- [19] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press.

- [20] J. J. Sakurai, *Modern Quantum Mechanics*, 2nd ed. Reading, MA: Addison Wesley, 1994.
- [21] P. Schroder and W. Sweldens, "Spherical wavelets: Efficiently representing functions on the sphere," in *Proc. SIGGRAPH*, 1995, pp. 161–172.
- [22] P. Schroder and W. Sweldens, "Spherical wavelets: Texture processing," *Rendering Tech.*, pp. 252–263, Aug. 1995.
- [23] E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, "Shiftable multi-scale transforms," *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 587–607, Feb. 1992.
- [24] F. Simons, F. Dahlen, and M. Wiecek, "Spatiospectral concentration on a sphere," *SIAM Rev.*, vol. 48, no. 3, pp. 504–536, 2006.
- [25] J.-L. Starck, E. J. Candes, and D. L. Donoho, "The curvelet transform for image denoising," *IEEE Trans. Image Process.*, vol. 11, no. 6, pp. 670–684, Jun. 2002.
- [26] J.-L. Starck, Y. Moudden, P. Abrial, and M. K. Nguyen, "Wavelets, ridgelets and curvelets on the sphere," *Astron. Astrophys.*, vol. 446, pp. 1191–1204, 2006.
- [27] M. Vetterli and J. Kovačević, *Wavelets and Subband Coding*. Englewood Cliffs, NJ: Prentice-Hall, 1995.
- [28] B. D. Wandelt and K. M. Gorski, "Fast convolution on the sphere," *Phys. Rev. D* 63, vol. 1230002, pp. 1–6, 2001.
- [29] Y. Wiaux, L. Jacques, and P. Vandergheynst, "Correspondence principle between spherical and Euclidean wavelets," *Astrophys. J.*, vol. 632, pp. 15–28, 2005.
- [30] Y. Wiaux, L. Jacques, and P. Vandergheynst, "Fast directional correlation on the sphere with steerable filters," *Astrophys. J.*, vol. 652, pp. 820–832, 2006.
- [31] B. T. T. Yeo, W. Ou, and P. Golland, "Invertible filter banks on the 2-Sphere," in *Proc. Int. Conf. Image Processing*, 2006, pp. 2161–2164.
- [32] P. Yu, P. E. Grant, Y. Qi, X. Han, F. Segonne, R. Pienaar, E. Busa, J. Pacheco, N. Makris, R. Buckner, P. Golland, and B. Fischl, "Cortical surface shape analysis based on spherical wavelets," *IEEE Trans. Med. Imag.*, vol. 26, no. 4, pp. 582–598, Apr. 2007.
- [33] P. Yu, B. T. T. Yeo, E. Grant, B. Fischl, and P. Golland, "Cortical folding development study based on over-complete spherical wavelets," presented at the IEEE Computer Society Workshop on Mathematical Methods in Biomedical Image Analysis, 2007.



Boon Thye Thomas Yeo received the B.S. and M.S. degrees in electrical engineering from Stanford University, Stanford, CA, in 2002. He is currently pursuing the Ph.D. degree in computer science at the Massachusetts Institute of Technology, Cambridge.

His current research focuses on developing and applying techniques in signal processing, statistics, and differential geometry to the problems of registration, segmentation, and shape analysis in neuroimaging.

Mr. Yeo has received several awards, including the MICCAI Young Scientist Award, A*STAR National Science Scholarship, Frederick E. Terman Engineering Scholastic Award, and the Hewlett-Packard and Agilent Technologies Project Award.



Wanmei Ou received the B.S. degree in electrical engineering from the City College of the City University of New York (Salutatorian) in 2003, the M.S. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, in 2005, where she is currently pursuing the Ph.D. degree.

She is interested in developing mathematical algorithms to improve brain activity detection for functional brain studies, as well as to quantify difference in anatomical structures. Her research

projects include spatial regularization for fMRI, spatio-temporal EEG/MEG inverse problems, neurovascular coupling, wavelets analysis, and joint analysis of fMRI and EEG/MEG data.



Polina Golland received the B.S. and M.S. degrees from the Technion—Israel Institute of Technology, Haifa, and the Ph.D. degree in 2001 from the Massachusetts Institute of Technology (MIT), Cambridge.

She is an Assistant Professor in the Electrical Engineering and Computer Science Department and the Computer Science and Artificial Intelligence Laboratory, MIT. She has worked on various problems in computer vision, including motion and stereo, shape modeling, and representation. Her current research focuses on modeling biological shape and function using images (from MRI to microscopy) as a source of information.

Prof. Golland has served on program committees of several leading conferences and as a referee for international journals in biomedical image analysis and computer vision. She received the NSF CAREER Award in 2007.