

INVERTIBLE FILTER BANKS ON THE 2-SPHERE

B.T. Thomas Yeo

Wanmei Ou

Polina Golland

Computer Science and Artificial Intelligence Lab, MIT, Cambridge, MA, USA

ABSTRACT

Multiscale filtering methods, such as wavelets and steerable pyramids, are 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. In this paper, we develop conditions for the invertibility of spherical filter banks for both continuous and discrete convolution and illustrate how such conditions can be incorporated into the design of self-invertible axis-symmetric wavelets. Self-invertibility is particularly desirable when modifying images in the wavelet domain.

Index Terms— spheres, frequency response, wavelet transforms, filtering, channel bank filters, image sampling, image orientation analysis, feature extraction

1. INTRODUCTION

The theories of filter banks, wavelets and overcomplete wavelets, such as steerable pyramids, are well-established for the Euclidean spaces and have many applications in feature detection, compression and denoising. Extending the theory and the methods of filtering to spherical images promises similar benefits in the fields that give rise to such images, including computer vision [3], computer graphics [8], astrophysics [13, 14, 15], and geophysics [11]. Our motivation comes from the study of brain cortical surfaces, often represented and manipulated on the sphere as spherical images [5]. These cortical surfaces are characterized by patterns of oriented edges. Detecting features of interest (e.g. consistent edge patterns or patterns of folding differences between populations) requires extending the filtering theory and algorithms to the sphere.

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. 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 [4], and was recently extended to arbitrary functions [13, 15]. The reconstructed image is obtained by adding the inverse convolutions of the filtering outputs with the synthesis filters. The filter shape 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. Self-invertibility is desirable for image manipulation in the




Fig. 1. Continuous analysis and synthesis filter bank diagram

wavelet domain, leading to an intuitive notion that the magnitude of a convolution coefficient signify the contribution of the corresponding filter to the reconstructed signal. Without self-invertibility, the effects of nonlinear processing of wavelet coefficients will propagate to locations and frequencies other than those which were used to compute the coefficients [9]. 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 [8, 10] 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 [6, 9]. The group-theoretic formulation of overcomplete spherical wavelets [1, 2] extends the overcomplete wavelet theory to the sphere through the operation of stereographic projection onto the tangent plane. This approach provides a straightforward framework for the design of analysis filters for specific features of interest, such as oriented edges [14]. Unfortunately, the synthesis filters are fully determined by the shape 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 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. NOTATION AND DEFINITIONS

Let $x(\theta, \phi) \in L^2(S^2)$ be a square-integrable function on the two-dimensional unit sphere, where $\theta \in [0, \pi]$ is the co-latitude and $\phi \in [0, 2\pi]$ is the longitude. The spherical harmonics $Y_l^m(\theta, \phi)$ form an orthonormal basis for $L^2(S^2)$: $x(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{|m| \leq l} x^{l,m} Y_l^m(\theta, \phi)$ for any $x \in L^2(S^2)$, where $x^{l,m}$ are the spherical harmonic coefficients of degree l and order m obtained by projecting x onto Y_l^m , i.e.

This work is funded in part by the NIH NIBIB NAMIC (U54-EB005149), the NIH NCRR mBIRN (U24-RR021382), the NIH NCRR NAC (P41-RR13218), and the NIH NINDS R01-NS051826 grant. B.T. Thomas Yeo is funded by Agency for Science, Technology and Research, Singapore. Wanmei Ou is supported by the NSF graduate fellowship.

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

We 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 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 (see for example [15]).

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 use relative rotations. Given spherical image x and filter \tilde{h} , their spherical convolution $y(\alpha, \beta, \gamma) = \int_{S^2} [D(\alpha, \beta, \gamma)\tilde{h}]^*(\theta, \phi) x(\theta, \phi) \sin(\theta) d\theta d\phi$ is a function in $L^2(SO(3))$. $y(\alpha, \beta, \gamma)$ is the inner product 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 $\theta = \beta, \phi = \alpha$. This symmetry is sometimes achieved for axis-asymmetric filters by integrating the resulting convolution over γ [4]. Our definition of convolution is identical to that in [13, 15], although [15] calls it directional correlation.

The inverse convolution of a spherical filter h with $y \in L^2(SO(3))$ produces a spherical image $\hat{x}_h(\theta, \phi) = \int_{SO(3)} [D(\alpha, \beta, \gamma)h](\theta, \phi) y(\alpha, \beta, \gamma) \sin(\beta) d\alpha d\beta d\gamma$. $\hat{x}_h(\theta, \phi)$ 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:

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

where $d\rho = \sin(\beta) d\alpha d\beta d\gamma$. Our definition is similar to the continuous wavelet transform in the group theoretic formulation [1]. The difference comes from their treatment of scale as a continuous parameter, rather than a discrete one, resulting in an integration over scale rather than a summation over 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, allowing us to exploit efficient algorithms for spherical convolution [13, 15]. We note that continuous representation in the wavelet domain is possible through series of complex exponentials [13] or Wigner-D functions [15], but manipulating the series coefficients would be tantamount to simultaneously altering all the wavelet coefficients, defeating the purpose of the wavelet decomposition. We leave the formal treatment of discrete convolution to the technical report [16].

3. INVERTIBILITY CONDITIONS

In this section, we present general conditions for invertibility of filter banks. The proofs are provided in the technical report [16].

Theorem 3.1 (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}

$$\hat{x}^{l,m} = x^{l,m} \text{ for all } (l, m) \text{ iff}$$

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

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

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

$$H_{\tilde{h},h}(l) = \frac{8\pi^2}{2l+1} \sum_{n=1}^N \sum_{m'=-l}^l [h_n^{l,m'}] [\tilde{h}_n^{l,m'}]^* \quad (3)$$

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 linearly with degree l .

A second theorem concerning discrete invertibility is presented in technical report [16], using quadratures and imposing limits on the bandwidth of the image and filters. Together, the two theorems on continuous and discrete invertibility imply that under very reasonable and mild conditions, if a filter bank is invertible up to degree L under the continuous spherical convolution, it is also invertible up to degree L under the discrete spherical convolution. 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 \tilde{h}_n , there are in general multiple sets of synthesis filters that can achieve invertibility. A simple solution is to define the synthesis filters to be equal to the analysis filters divided by $H_{\tilde{h},\tilde{h}}(l)$ as defined in Eq. (3). This is similar to the frame operator in the continuous spherical wavelet transform under the group theoretic formulation [1], where the counterpart of $H_{\tilde{h},\tilde{h}}(l)$ is given by $\frac{8\pi^2}{2l+1} \sum_{|m| \leq l} \int_0^\infty \frac{1}{a^3} |\tilde{h}_a^{l,m}|^2 da$, 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 . For completeness, we note that the discretization of the group theoretic formulation is achieved in [2]. However, even if the analysis filters are related by dilations, the synthesis filters discussed in all the above approaches will in general not have this relationship. In the next section, we specialize the invertibility conditions to self-invertible multiscale filter banks.

4. SELF-INVERTIBLE MULTI-SCALE FILTER BANKS

The results from the previous section are general and apply to any analysis-synthesis filter bank. We shall now constrain the relationships among the analysis and synthesis filters. In multi-scale analysis, we construct the analysis filters through dilation and amplitude scaling of a particular template $\tilde{h}(\theta, \phi)$, i.e., $\tilde{h}_k = (\prod_{n=1}^k b_n) D_{a_k} \tilde{h}(\theta, \phi)$, $b_n \geq 1$. D_{a_k} is the nonlinear dilation operator. In this work, we adopt the stereographic dilation operator defined in [1, 2], where larger k corresponds to smaller a (narrower filters).




Fig. 2. (a) Combined frequency responses $H_{\tilde{h},h}(l) = \frac{8\pi^2}{2l+1} \sum_{n=1}^N \sum_{m'=-l}^l [h_n^{l,m'}][\tilde{h}_n^{l,m'}]^*$ of the lowpass filter pyramids found in the four experiments, using different scale sets and penalty functions. (b-e) Frequency responses of individual lowpass filters $(\prod_{n=1}^k b_n)D_{a_k}\tilde{h}(\theta, \phi)$ whose combined frequency response is shown in green in (a). Note the different vertical scales in the graphs. See text for details.

The b_n 's are the amplitude scaling parameters that control the tradeoff between self-invertibility and norm-preserving dilation. Theorem 3.1 and Eq. (3) imply 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.

These extra weights are analogous to the measure of scale $\frac{1}{a^3}da$ in the group theoretic formulation of wavelets [1], which results in wider filters being assigned smaller weights. On the continuous real line, the measure $\frac{1}{a^2}da$ nicely cancels out the dilation of the filter (c.f. [12], chapter 5). On the discrete real line, the convolution outputs of narrower filters are sampled more densely by enforcing orthonormality. This suggests two possible approaches: variable sampling or weights. Because the effects of stereographic dilation on the spherical harmonic coefficients of a function is not analytical, neither approach leads to a closed-form solution¹. In this paper, we will take the variable weights approach, i.e. find the appropriate b_n '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 , i.e., $\tilde{h}(\theta, \phi) = \sum_{i=1}^M c_i B^i(\theta, \phi)$ (We will assume that B^i are spherical harmonics and note that the technique is still applicable if a more suitable basis is found). Then,

$$[D_a \tilde{h}]^{l,m} = \left[D_a \sum_{i=1}^M c_i B^i \right]^{l,m} = \sum_{i=1}^M c_i [D_a B^i]^{l,m} \quad (4)$$

yielding the spherical harmonic coefficients of the analysis filter at another scale. This is useful since the invertibility condition in Eq. (2) was 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 create a table of spherical harmonic coefficients of the dilated basis functions. Eq. (4) 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 a constrained optimization problem to determine c_i 's and b_n 's. Similarly to the filter design in Euclidean space, the optimization objective should be application dependent, and could depend on 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 Eq. (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.

¹We 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 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.




Fig. 3. Lowpass filters from Fig. 2(b-e), shown in the spherical domain. The bright spot corresponds to the north pole.

The quadratic penalty method is effective in solving this optimization problem with non-convex constraints by incorporating the constraints into 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 demonstrate the use of the methods discussed in the previous section to construct axis-symmetric lowpass filter pyramids. Our implementation is based on the fast spherical convolution [4]. Due to space constraints, we leave the discussion of axis-symmetric bandpass pyramids to the technical report [16].

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 axis-symmetric functions are zero for orders not equal to 0. We use the S2kit [7] software package to find the first 600 order 0 spherical harmonic coefficients of each dilated spherical harmonic $D_a Y_l^0$, $l = 0, \dots, 99$ (a dilated axis-symmetric function remains axis-symmetric). Due to non-linearity of the stereographic dilation, extreme dilation and shrinking can result in high frequencies [2]. We find that for $a = 4$ and $a = 1/4$, $[D_a Y_{99}^0]^{599,0} < 10^{-7}$, implying that 600 spherical harmonics are sufficient at these scales. We sample the intermediate scales, resulting in a table for $a_n = \{4, 2, 1, 1/2, 1/4\}$, with $a = 1$ corresponding to the undilated spherical harmonics.

We would like to derive an objective function that determines a lowpass filter pyramid. 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 spherical harmonics as our basis, the frequency response of $\tilde{h}_{a=1}(\theta, \phi)$ will be zero for all degrees higher than 99. Setting our objective function to simply penalize the first order 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 will also enforce self-invertibility for degrees lower than 70 (Eq. (2), $l = 0, \dots, 69$). In addition, we can make the cutoff sharper by adding extra terms in the objective function to penalize non-zero frequency responses for degrees above some threshold.

We experimented with 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




Fig. 4. (a) World elevation map, $BW = 539$. (b) World elevation map, truncated at $BW = 69$. (c-f) Convolution outputs between the set of lowpass filters shown in Figs. 2-3 and the world elevation map shown in (a).

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). Fig. 2(a) shows the combined frequency responses of the filter banks from each experiment. We observe that the extra quadratic penalty term leads to a sharper cutoff (green and blue versus black and red). The combined frequency responses of the filters from the 5-scales experiments also cover a much wider range of spherical harmonics than the self-invertibility range (blue and red versus green and black). 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. Their individual responses are shown in Fig. 2(b-e). Notice the energies of the narrower filters are much higher than those of the wider ones. Such discrepancies in energies are less pronounced (but still present) in bandpass filters [16]. Fig. 3 displays the filters as spherical functions. The bright spot corresponds to the north pole.

To verify that the lowpass filter pyramid is indeed self-invertible, we apply the filters to the global elevation map of the Earth. The original elevation data $y(\theta, \phi)$ is sampled on a 1080×1080 equi-angular grid. Using S2kit [7], we project the data onto the spherical harmonics of degrees at most 539. Fig. 4(a) and (b) show the elevation maps up to bandwidth 539 and 69 respectively, obtained by performing an inverse spherical harmonic transform.

Fig. 4(c-f) illustrate the application of these filters to the high-resolution spherical image in Fig. 4(a). The reconstructed image (not shown) obtained by combining the deconvolution of the filter outputs and the filter pyramid (Eq. 1) is an accurate reconstruction of the lower-resolution world map in Fig. 4(b), with 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 at least 70 if we want perfect reconstruction up to degrees 539.

We can also apply the filters directly to the lower-resolution image ($BW = 69$). However, judging from the frequency responses of the filters corresponding to $a = 1$ and $a = 1/2$ in Fig. 2(d-e) and their convolved outputs in Fig. 4(e-f), the convolution outputs of $a = 1$ and $a = 1/2$ will be quite similar since the resolution of the narrower $a = 1/2$ filter will be limited by that of the input image.

6. DISCUSSION AND CONCLUSIONS

In this paper, we derive the necessary and sufficient conditions for a bank of filters to achieve perfect reconstruction in the continuous case and discuss its applications to discrete filter design. We present a procedure for obtaining self-invertible wavelets and generated axis-symmetric self-invertible lowpass wavelets.

Although stereographic dilation has many advantages, 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 spherical 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.

Here we have only demonstrated the use of our procedures for axis-symmetric wavelets that are useful for image representation. We plan to expand the basis set to axis-asymmetric spherical harmonics and derive a useful objective function for axis-asymmetric filter pyramid so as to enable oriented feature detection at different scales.

This paper introduces theoretical results on invertibility and represents a step towards the general theory of 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.

8. REFERENCES

- [1] Antoine, J-P & Vandergheynst, P. Wavelets on the 2-Sphere: a Group-Theoretical Approach. *Appl. Comput. Harmon. Anal.*, 7:262–291, 1999.
- [2] Bogdanova I., Vandergheynst, P., Antoine, J-P, Jacques, L. & Morvidone, M. Stereographic Wavelet Frames on the Sphere. *Appl. Comput. Harmon. Anal.*, 19:223–252, 2005.
- [3] Brechbuhler, C.H., Gerig, G. & Kubler O. Parametrization of Closed Surfaces for 3-D Shape Description. *Comput. Vision and Image Understanding*, 61(2):154–170, 1994.
- [4] Driscoll, J.R. & Healy, D.M. Computing Fourier Transforms & Convolutions on the 2-Sphere. *Adv. in Appl. Math.*, 15:202–250, 1994.
- [5] Fischl, B., Sereno M.I. & Dale, A.M. Cortical Surface-Based Analysis: II: Inflation, Flattening, and a Surface-Based Coordinate System. *NeuroImage*, 9(2):195–207, 1999.
- [6] Freeman, W.T. & Adelson, E.W. The Design and Use of Steerable Filters. *IEEE Trans. Pattern Anal. and Machine Intell.*, 13(9):891–906, 1991.
- [7] Kostelec, P.J. & Rockmore, D.N. A lite version of spharmonic kit. <http://www.cs.dartmouth.edu/geelong/sphere/>.
- [8] Schroder, P. & Sweldens, W. Spherical Wavelets: Efficiently Representing Functions on the Sphere. *Comput. Graphics Proc. (SIGGRAPH)*, 161–172, 1995.
- [9] Simoncelli, E.P., Freeman, W.T., Adelson, E.H. & Heeger, D.J. Shiftable Multiscale Transforms. *IEEE Trans. Inform. Theory*, 38(2):587–607, 1992.
- [10] Sweldens, W. The lifting scheme: A construction of second generation wavelets. *Siam J. Math. Analysis*, 29(2):511–546, 1997.
- [11] Vaníček, P., Janák, J. & Featherstone, W.E. Truncation of Spherical Convolution Integration with an Isotropic Kernel. *Studia Geophysica et Geodaetica*, 47(3):455–465, 2003.
- [12] Vetterli, M. & Kováčević, J. Wavelets and Subband Coding. Prentice-Hall, 1995.
- [13] Wandelt, B.D. & Gorski, K.M. Fast Convolution on the Sphere. *Physics Review D* 63, 1230002:1–6, 2001.
- [14] Wiaux, Y., Jacques, L. & Vandergheynst, P. Correspondence Principle Between Spherical and Euclidean Wavelets. *Astrophys. J.*, 15:632, 2005.
- [15] Wiaux, Y., Jacques, L. & Vandergheynst, P. Fast Directional Correlation on the Sphere with Steerable Filters. *Astrophys. J.*, submitted, 2005.
- [16] Yeo, B.T.T., Ou, W. & Golland, P. Invertible Filter Banks on the 2-Sphere. CSAIL Memo, 2006. <http://people.csail.mit.edu/ythomas/IFB2006.pdf>.