HIGHER-ORDER RIESZ TRANSFORMS AND STEERABLE WAVELET FRAMES

Michael Unser and Dimitri Van De Ville

Biomedical Imaging Group,
Ecole Polytechnique Fédérale de Lausanne (EPFL),
CH-1015, Lausanne, Switzerland

ABSTRACT

We introduce an Nth-order extension of the Riesz transform in d dimensions. We prove that this generalized transform has the following remarkable properties: shift-invariance, scale-invariance, inner-product preservation, and steerability. The pleasing consequence is that the transform maps any primary wavelet frame (or basis) of $L_2(\mathbb{R}^d)$ into another "steerable" wavelet frame, while preserving the frame bounds. The concept provides a rigorous functional counterpart to Simoncelli's steerable pyramid whose construction was entirely based on digital filter design. The proposed mechanism allows for the specification of wavelets with any order of steerability in any number of dimensions; it also yields a perfect reconstruction filterbank algorithm. We illustrate the method using a Mexican-hat-like polyharmonic spline wavelet transform as our primary frame.

Index Terms— wavelet transform, steerable filters, frames, multiresolution analysis

1. INTRODUCTION

Scale and directionality are key considerations for visual processing and perception. Researchers in image processing and applied mathematics have worked hard on formalizing these notions and on developping some corresponding feature extraction algorithms and/or signal representation schemes. The concept of scale is admirably captured in Mallat's multiresolution theory of the wavelet transform, while the idea of steerable filterbanks provides an elegant approach to the extraction of directional image features [1]. The steerable pyramid of Simoncelli represents the best effort to date to merge these two approaches [2].

The steerable pyramid has had a strong impact on the field and has led to many successful applications. Yet, it remains that it is a purely discrete scheme with the disadvantage of not providing the type of continuous-domain interpretation that is so central to wavelet theory. The other point is that Simoncelli's construction is somewhat empirical and difficult to generalize.

In this paper, we introduce a general operator framework that provides a mathematical bridge between the continuous-domain theory of the wavelet transform and steerable filterbanks. The main intent is to facilitate the design of steerable wavelet transforms and to extend the concept for dimensions higher than two. The mathematical proofs are available from the authors and will be published elsewhere.

This work was supported by the Swiss National Science Foundation under Grant 200020-109415, the Center for Biomedical Imaging (CIBM) and the foundations Leenaards and Louis-Jeantet.

2. RIESZ TRANSFORMS AND PROPERTIES

2.1. The Riesz transform

The Riesz transform is the natural multidimensional extension of the Hilbert transform [3]. It is the scalar-to-vector signal transformation $\boldsymbol{\mathcal{R}}$ whose frequency-domain definition is

$$\widehat{\boldsymbol{\mathcal{R}}f}(\boldsymbol{\omega}) = -j\frac{\boldsymbol{\omega}}{\|\boldsymbol{\omega}\|}\hat{f}(\boldsymbol{\omega}) \tag{1}$$

where $\hat{f}(\boldsymbol{\omega})$ is the Fourier transform of the d-dimensional input signal $f(\boldsymbol{x})$. This corresponds to a multidimensional d-channel filterbank whose space-domain description is

$$\boldsymbol{\mathcal{R}}f(\boldsymbol{x}) = \begin{pmatrix} \mathcal{R}_1 f(\boldsymbol{x}) \\ \vdots \\ \mathcal{R}_d f(\boldsymbol{x}) \end{pmatrix} = \begin{pmatrix} (h_1 * f)(\boldsymbol{x}) \\ \vdots \\ (h_d * f)(\boldsymbol{x}) \end{pmatrix} \tag{2}$$

where the filters $(h_n)_{n=1}^d$ are characterized by their frequency responses $\hat{h}_n(\boldsymbol{\omega}) = -j\omega_n/\|\boldsymbol{\omega}\|$. For $d = 1$, one recovers the classical Hilbert transform whose frequency response is $\hat{h}(\omega) = -\frac{j\omega}{|\omega|}$.

2.2. Steerability and directional Hilbert transform

Similar to the definition of directional derivatives, we introduce a directional version of the Hilbert transform along any given unit vector $\boldsymbol{u} = (u_1, \ldots, u_d)$:

$$\mathcal{H}_{\boldsymbol{u}} f(\boldsymbol{x}) = \sum_{n=1}^{d} u_n \mathcal{R}_n f(\boldsymbol{x}) = \langle \boldsymbol{u}, \boldsymbol{\mathcal{R}}f(\boldsymbol{x}) \rangle \tag{3}$$

Denoting the canonical basis in $\mathbb{R}^d$ by $\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_d\}$, we have that $\mathcal{H}_{\boldsymbol{e}_n} f(\boldsymbol{x}) = (h_n * f)(\boldsymbol{x})$ which corresponds to the nth component of the Riesz transform. Moreover, by using the rotation property of the Fourier transform, we can show that the impulse response of the directional Hilbert transform corresponds to the rotated version of $h_1(\boldsymbol{x})$ (first component filter) in the direction specified by $\boldsymbol{u}$; i.e., $\mathcal{H}_{\boldsymbol{u}}\delta(\boldsymbol{x}) = h_1(\mathbf{R}_{\boldsymbol{u}}\boldsymbol{x})$ where $\delta(\boldsymbol{x})$ is the multidimensional dirac distribution and $\mathbf{R}_{\boldsymbol{u}}$ is a rotation matrix such that $\boldsymbol{u} = \mathbf{R}_{\boldsymbol{u}}\boldsymbol{e}_1$.

The bottom line is that the Riesz component filters in (2) are 90°-rotated versions of each other and that the filterbank response can be steered in any direction according to (3). Another way to put it is that the Riesz transform commutes with spatial rotations: $\boldsymbol{\mathcal{R}}\{\delta\}(\mathbf{R}\boldsymbol{x}) = \mathbf{R}\boldsymbol{\mathcal{R}}\{\delta\}(\boldsymbol{x})$ where $\mathbf{R}$ is an arbitrary $d \times d$ rotation matrix.

2.3. Higher-order Riesz transforms

It is possible to define higher-order Riesz transforms by considering individual signal components of the form $\mathcal{R}_{i_1}\mathcal{R}_{i_2}\cdots\mathcal{R}_{i_N}f$ with $i_1, i_2, \ldots, i_N \in \{1, \cdots, d\}$. While there are d^N possible ways of forming such Nth order terms, there are actually much less distinct Riesz components due to the commutativity and factorization properties of the underlying convolution operators. Our definition of higher-order transform removes this intrinsic redundancy and is based on the following operator identity.

Theorem 1. *The Nth-order Riesz transform achieves the following decomposition of the identity*

$$\sum_{\substack{n_1,\ldots,n_d\geq 0\\ n_1+\cdots+n_d=N}} \binom{N}{n_1,\ldots,n_d} \left(\mathcal{R}_1^{n1}\cdots\mathcal{R}_d^{n_d}\right)^* \left(\mathcal{R}_1^{n1}\cdots\mathcal{R}_d^{n_d}\right) = \mathrm{Id}$$

where the weighting factors are the multinomial coefficients

$$\binom{N}{n_1,n_2,\ldots,n_d} = \frac{N!}{n_1!\,n_2!\cdots n_d!}.$$

Proof. : By applying the above left-hand side decomposition formula to an input signal f, we obtain a series of components $g_{n_1,\ldots,n_d}(\boldsymbol{x}) = \left(\mathcal{R}_1^{n1}\cdots\mathcal{R}_d^{n_d}\right)^*\left(\mathcal{R}_1^{n1}\cdots\mathcal{R}_d^{n_d}\right)f(\boldsymbol{x})$, whose Fourier-domain equivalent is

$$\hat{g}_{n_1,\ldots,n_d}(\boldsymbol{\omega}) = \frac{|\omega_1|^{2n_1}\cdots|\omega_d|^{2n_d}}{\|\boldsymbol{\omega}\|^{2N}}\hat{f}(\boldsymbol{\omega}). \tag{4}$$

We also have the identity

$$\hat{f}(\boldsymbol{\omega}) = \frac{\|\boldsymbol{\omega}\|^{2N}}{\|\boldsymbol{\omega}\|^{2N}}\hat{f}(\boldsymbol{\omega}) = \frac{1}{\|\boldsymbol{\omega}\|^{2N}}\left(|\omega_1|^2+\cdots+|\omega_d|^2\right)^N\hat{f}(\boldsymbol{\omega}).$$

Next, we apply the multinomial theorem to expand the central factor

$$\begin{aligned}\hat{f}(\boldsymbol{\omega}) &= \sum_{\substack{n_1,n_2,\ldots,n_d\geq 0\\ n_1+n_2+\cdots+n_d=N}} \binom{N}{n_1,\ldots,n_d}\frac{|\omega_1|^{2n_1}\cdots|\omega_d|^{2n_d}}{\|\boldsymbol{\omega}\|^{2N}}\hat{f}(\boldsymbol{\omega})\\ &= \sum_{\substack{n_1,\ldots,n_d\geq 0\\ n_1+n_2+\cdots+n_d=N}} \binom{N}{n_1,n_2,\ldots,n_d}\hat{g}_{n_1,\cdots,n_d}(\boldsymbol{\omega})\end{aligned}$$

where we have identified the relevant higher-order reconstruction components using (4). □

To simplify the identification/indexing of higher-order Riesz components in Theorem 1, we define the set

$$P_N^d = \left\{(n_1,n_2,\ldots,n_d)\in\mathbb{N}^d \mid \sum_{i=1}^d n_i = N\right\}$$

with the property that $p(N,d) = \mathrm{card}(P_N^d)$. We then specify the Nth order Riesz transform of the signal f as the $p(N,d)$-vector signal transformation whose components are given by

$$f_{n_1,\ldots,n_d} = \mathcal{R}^{(n_1,\ldots,n_d)}f = \sqrt{\frac{N!}{n_1!\,n_2!\cdots n_d!}}\mathcal{R}_1^{n_1}\cdots\mathcal{R}_d^{n_d}f$$

with $(n_1,n_2,\ldots,n_d)\in P_N^d$. The crucial point in this definition is that we have included proper normalization factors to ensure energy preservation (cf. Property 2 below). In some instances, we will use the more concise vector notation $\boldsymbol{\mathcal{R}}^{(N)}f$, keeping in mind that each component is associated with its own normalized, Nth-order Riesz operator $\mathcal{R}^{(n_1,\ldots,n_d)}$, as defined above.

2.4. Properties of the Riesz transform and its extensions

Property 1. *The Riesz transform and its higher-order extensions are translation- and scale-invariant:*

$$\begin{aligned}\forall \boldsymbol{x}_0\in\mathbb{R}^d,\quad \boldsymbol{\mathcal{R}}^{(N)}\{f(\cdot-\boldsymbol{x}_0)\}(\boldsymbol{x}) &= \boldsymbol{\mathcal{R}}^{(N)}\{f(\cdot)\}(\boldsymbol{x}-\boldsymbol{x}_0)\\ \forall a\in\mathbb{R}^+,\quad \boldsymbol{\mathcal{R}}^{(N)}\{f(\cdot/a)\}(\boldsymbol{x}) &= \boldsymbol{\mathcal{R}}^{(N)}\{f(\cdot)\}(\boldsymbol{x}/a)\end{aligned}$$

The translation invariance follows from the definition, while the scale invariance of each component is verified in the Fourier domain.

Property 2. *The Riesz transform and its Nth-order extension satisfy the following Parseval-like identities*

$$\begin{aligned}\forall f,\phi\in L_2(\mathbb{R}^d),\quad \langle\boldsymbol{\mathcal{R}}f,\boldsymbol{\mathcal{R}}\phi\rangle_{L_2^d} &= \sum_{n=1}^d\langle\mathcal{R}_n f,\mathcal{R}_n\phi\rangle_{L_2}\\ &= \langle f,\phi\rangle_{L_2}.\end{aligned}$$

$$\forall f,\phi\in L_2(\mathbb{R}^d),\quad \langle\boldsymbol{\mathcal{R}}^{(N)}f,\boldsymbol{\mathcal{R}}^{(N)}\phi\rangle_{L_2^{p(N,d)}} = \sum_{(n_1,\ldots,n_d)\in P_N^d}\langle\mathcal{R}^{(n_1,\ldots,n_d)}f,\mathcal{R}^{(n_1,\ldots,n_d)}\phi\rangle_{L_2} = \langle f,\phi\rangle_{L_2}.$$

The latter is a consequence of Theorem 2.

Theorem 2 (steerability)**.** *The Nth order Riesz transform provides the complete information for computing the Nth order directional Hilbert transform along any direction $\boldsymbol{u}=(u_1,\ldots,u_d)$*

$$\mathcal{H}_{\boldsymbol{u}}^N f(\boldsymbol{x}) = \sum_{\substack{n_1,\ldots,n_d\geq 0\\ n_1+\cdots+n_d=N}} c_{n_1,\ldots,n_d}(\boldsymbol{u})\,f_{n_1,\ldots,n_d}(\boldsymbol{x})$$

where the "steering" coefficients are given by

$$c_{n_1,\ldots,n_d}(\boldsymbol{u}) = \sqrt{\binom{N}{n_1,n_2,\ldots,n_d}}\frac{u_1^{n_1}\cdots u_d^{n_d}}{\|\boldsymbol{u}\|}. \tag{5}$$

Because of the operator identity $\mathcal{H}_{\boldsymbol{e}_1}^N = \mathcal{R}^{(N,0,\ldots,0)} = \mathcal{R}_1^N$, this is equivalent to steering the first component filter of the corresponding filterbank. This steerability property extends for any generalized Riesz operator of the form $\boldsymbol{a}^T\boldsymbol{\mathcal{R}}^{(N)}$ where $\boldsymbol{a}$ is a vector of weights.

This result is established in the Fourier domain using the multinomial expansion once more.

3. CONSTRUCTION OF STEERABLE WAVELET FRAMES

3.1. A general method for constructing steerable frames

The pleasing consequence of Property 2 is that the Riesz transform will automatically map any frame of $L_2(\mathbb{R}^d)$ into another one. Before establishing this result, we start by recalling a few basic facts about frames [4].

Definition 1. *A family of functions $\{\phi_{\boldsymbol{k}}\}_{\boldsymbol{k}\in\mathbb{Z}^d}$ is called a frame of $L_2(\mathbb{R}^d)$ iff. there exists two strictly positive constants A and $B<\infty$ such that*

$$\forall f\in L_2(\mathbb{R}^d),\quad A\|f\|_{L_2(R^d)}^2 \leq \sum_{\boldsymbol{k}\in\mathbb{Z}^d}|\langle\phi_{\boldsymbol{k}},f\rangle_{L_2}|^2 \leq B\|f\|_{L_2(R^d)}^2.$$

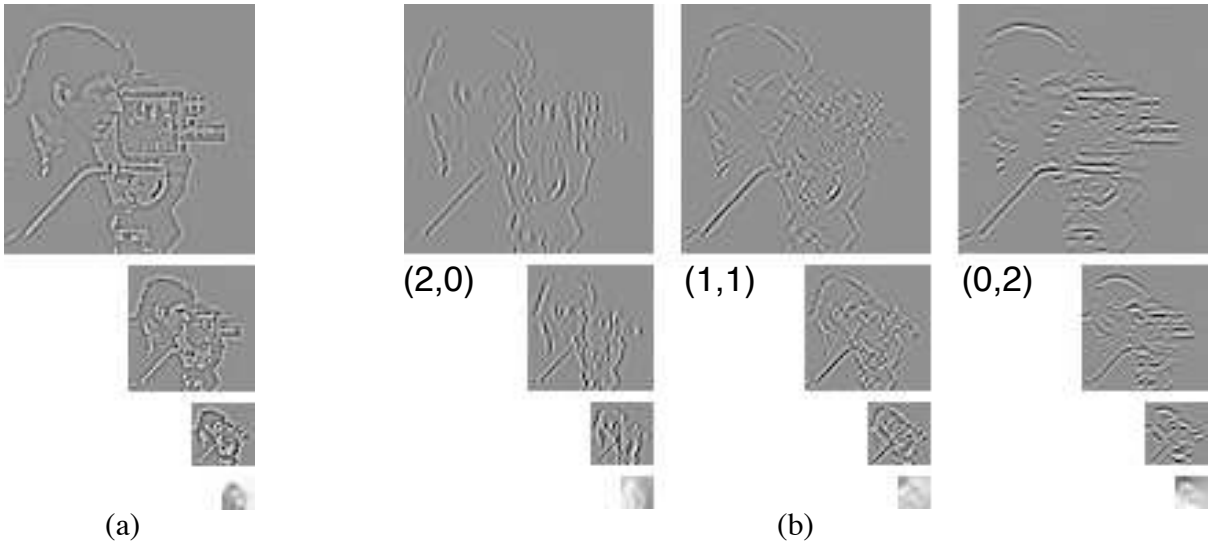

Fig. 1. Reversible wavelet decompositions of Cameraman image (zoom on a 64×64 subimage). (a) Primary Mexican-hat polyharmonic spline transform with $\gamma = 3$ (cf. [5]). (b) Corresponding 2nd order steerable wavelet transform.

The frame called is *tight* iff. the frame bounds can be chosen such that $A = B$. The fundamental property of a frame is that there exist a set of dual synthesis functions $(\tilde{\phi}_{\boldsymbol{k}})$ (not necessarily unique) such that the following decomposition/reconstruction formula holds

$$\forall f \in L_2(\mathbb{R}^d), \quad f = \sum_{\boldsymbol{k} \in \mathbb{Z}^d} \langle \phi_{\boldsymbol{k}}, f \rangle_{L_2} \, \tilde{\phi}_{\boldsymbol{k}}.$$

While the form of this expansion is the same as that associated with a biorthogonal basis, the fundamental aspect of the frame generalization is that the family $(\phi_{\boldsymbol{k}})$ may be redundant.

We are now ready to state the fundamental frame preservation property of the Riesz transform and its higher-order extensions.

Theorem 3. *The Nth order Riesz transform $\boldsymbol{\mathcal{R}}^{(N)}$ maps the frame $\{\phi_{\boldsymbol{k}}\}_{\boldsymbol{k} \in \mathbb{Z}^d}$ into another frame $\{\mathcal{R}^{(n_1,\ldots,n_d)} \phi_{\boldsymbol{k}}\}_{(n_1,\ldots,n_d) \in P_N^d, \boldsymbol{k} \in \mathbb{Z}^d}$ of $L_2(\mathbb{R}^d)$.*

The key argument for this proof is "energy conservation" which comes as a direct consequence of Property 2. Note that the redundancy is increased in proportion to the number of Riesz components.

3.2. Steerable wavelets in $L_2(\mathbb{R}^d)$

What is most remarkable is that the Riesz operator will transform any wavelet frame into another one, thanks to its translation- and scale-invariance properties. Specifically, let us consider the following wavelet decomposition of a finite energy signal:

$$\forall f \in L_2(\mathbb{R}^d), \quad f(\boldsymbol{x}) = \sum_{i \in \mathbb{Z}} \sum_{\boldsymbol{k} \in \mathbb{Z}^d} \langle f, \psi_{i,\boldsymbol{k}} \rangle_{L_2} \, \tilde{\psi}_{i,\boldsymbol{k}}(\boldsymbol{x}), \qquad (6)$$

where the wavelets at scale i are dilated versions of the ones at the reference scale $i = 0$: $\psi_{i,\boldsymbol{k}}(\boldsymbol{x}) = 2^{-id/2} \psi_{0,\boldsymbol{k}}(\boldsymbol{x}/2^i)$. General wavelet schemes also have an inherent translation-invariant structure in that the $\psi_{0,\boldsymbol{k}}$'s are constructed from the integer shifts (index $\boldsymbol{k}$) of up to 2^d distinct generators (mother wavelets).

Proposition 1. *Let $\{\psi_{i,\boldsymbol{k}}\}$ be a primal wavelet frame of $L_2(\mathbb{R}^d)$ associated with the reconstruction formula* (6)*. Then, $\{\mathcal{R}^{(n_1,\ldots,n_d)} \psi_{i,\boldsymbol{k}} \, | (n_1,\ldots,n_d) \in P_N^d\}$ and $\{\mathcal{R}^{(n_1,\ldots,n_d)} \tilde{\psi}_{i,\boldsymbol{k}} \, | (n_1,\ldots,n_d) \in P_N^d\}$ form a dual set of wavelet frames. They admit the same type of decomposition/reconstruction formula as* (6) *with an additional summation over the Riesz components.*

If, in addition, the primary mother wavelets in (6) are isotropic or at least orientation-free, then we end up with a decomposition that is steerable and essentially rotation-invariant, thanks to Theorem 2.

Finally, if the wavelet decomposition (6) is built around a Mallat-type multiresolution analysis and admits a fast filterbank algorithm then there is a corresponding perfect reconstruction filterbank implementation for its Nth-order Riesz counterpart in Proposition 1.

4. ILLUSTRATIVE EXAMPLE

While the construction method outlined in §3.2 is applicable to any primary wavelet transform, it will provide the most satisfactory results when the underlying primary wavelets are very nearly isotropic. We will therefore illustrate the concept using a primary Mexican-hat-like decomposition of $L_2(\mathbb{R}^d)$ [5].

4.1. Primary quasi-isotropy polyharmonic spline wavelets

Polyharmonic splines are the natural spline functions associated with the Laplace operator; they can be used to specify a family of multiresolution analyses $L_2(\mathbb{R}^d)$ which is indexed by the order of the Laplacian [6]. In prior work, we have introduced a pyramid-like representation of the underlying wavelet subspaces that involves a single quasi-isotropic analysis wavelet (cf [5, Section IV.B]):

$$\psi_{\text{iso}}(\boldsymbol{x}) = (-\Delta)^{\gamma/2} \beta_{2\gamma}(2\boldsymbol{x})$$

where $\gamma > 1$ is the order (possibly fractional) of the wavelet and $(-\Delta)^{\alpha} \stackrel{\mathcal{F}}{\longleftrightarrow} \|\boldsymbol{\omega}\|^{2\alpha}$ denotes the αth iterate of the Laplace operator. The function $\beta_{2\gamma}(\boldsymbol{x})$ is the "most-isotropic" polyharmonic B-spline of order 2γ; it is a smoothing kernel that converges to a Gaussian as γ increases [6, Proposition 2].

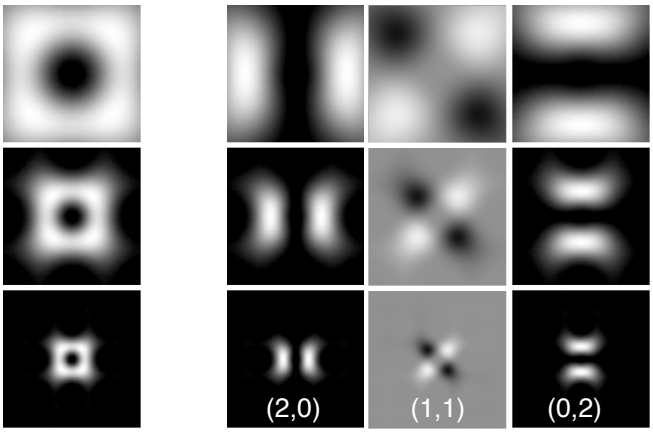

(a) (b)

Fig. 2. Frequency responses of the wavelet analysis filters in Fig. 1. (a) Primary Mexican-hat polyharmonic spline wavelet $|\widehat{\psi}_{\mathrm{iso},i}(\boldsymbol{\omega})|$ with $\gamma = 3$. (b) 2nd-order Riesz-Laplace wavelets $|\widehat{\psi}_i^{(2,0)}(\boldsymbol{\omega})|$, $\mathrm{Re}\Big(\widehat{\psi}_i^{(1,1)}(\boldsymbol{\omega})\Big)$, and $|\widehat{\psi}_i^{(0,2)}(\boldsymbol{\omega})|$.

The corresponding analysis wavelets in (6) are given by $\psi_{i,\boldsymbol{k}}(\boldsymbol{x}) = 2^{-id/2}\psi_{\mathrm{iso}}(\boldsymbol{x}/2^i - \boldsymbol{k})$ and the wavelet transform has a fast reversible filterbank implementation, as described in [5]. An example of such a pyramid-like representation is shown in Fig. 1a. The decomposition is a frame with a small redundancy factor ($R = 4/3$ in 2-D and $R = 6/5$ for $d = 3$). The corresponding 2-D wavelet frequency responses are shown in Fig. 2a. Note the bandpass behavior of these filters and their good isotropy properties, especially over the lower frequency range.

4.2. Steerable higher-order Riesz-Laplace wavelets

The associated Nth-order steerable Riesz-Laplace wavelet transform is obtained through the direct application of Proposition 1. The explicit form of the underlying directional analysis wavelets is

$$\psi^{(n_1,\ldots,n_d)}(\boldsymbol{x}) = \sqrt{\frac{N!}{n_1!\,n_2!\cdots n_d!}}\mathcal{R}_1^{n_1}\cdots\mathcal{R}_d^{n_d}(-\Delta)^{\gamma/2}\beta_{2\gamma}(2\boldsymbol{x})$$

with $(n_1,\ldots,n_d) \in P_N^d$. The enlightening aspect in this formula is that the directional behavior of these wavelets is entirely encoded in the differential operators that are acting on the B-spline smoothing kernel: the Riesz transform factor is steerable while the Laplacian part, which is responsible for the vanishing moments, is perfectly isotropic.

Interestingly, the simplest case $d = 2$, $N = 1$, and $\gamma = 3$ yields two gradient-like wavelets that happen to be mathematically equivalent to the Marr wavelets described in [7]. The next case $d = 2$ and $N = 2$ is new and it yields three Hessian-like wavelets. An example of such a three-channel directional wavelet transform is shown in Fig. 1b for a fragment of the Cameraman image to facilitate the visualization; the frequency responses of the corresponding wavelet filters are displayed in Fig. 2b.

5. DISCUSSION

The proposed operator-based framework is very general: it can yield a large variety of multidimensional wavelets with arbitrary orders of steerability. We have implemented such reversible transforms for orders up to 11, but we are only showing results for $N = 2$ because of space limitations.

The goal that we are pursuing here is quite similar to that achieved by Simoncelli with his steerable pyramid [2]. However, there are fundamental differences. First, Simoncelli's approach is purely numerical and based on filterbank design. Ours, on the other hand, is analytical and rooted in functional analysis; the distinction is similar to that between Mallat's multiresolution theory, which provides a continuous-domain interpretation of wavelets, and the design of critically-sampled perfect reconstruction filterbanks that actually predate wavelets but contain the algorithmic essence of this type of decomposition. The second difference is that for $N > d$, our directional Riesz wavelets are no longer rotated version of a single template, as is the case with traditional steerable filters. This is quite apparent in Fig. 2b. The "à la Simoncelli" counterpart of such a transform would involve three rotated versions of the first filter uniformly spread over the circle (e.g., at 0°, 120°, and 240°). In fact, our wavelets are much closer to derivatives than to rotated filterbanks; we believe that this could eventually be an advantage for analysis purposes. A fundamental problem with rotated filterbanks is that they do not generalize well for $d > 2$.

Our approach is not limited to polyharmonic splines, but it works best—i.e., provides exact steerability—when the primary wavelets are isotropic. A convenient aspect is that it will yield a tight frame whenever the primary decomposition is tight.

The proposed framework decouples the directional and multiscale aspects of wavelet design. While this is valuable conceptually, it also suggest some new design challenges; i.e., the specification of tight wavelet frames with improved isotropy properties.

To the best of our knowledge, this is the first scheme that can provide steerable wavelets in dimensions greater than 2. Such transforms could turn out to be valuable for the analysis and processing of 3-D biomedical data sets (e.g., X-ray tomograms, MRI, and confocal micrographs).

6. REFERENCES

[1] W. T. Freeman and E. H. Adelson, "The design and use of steerable filters," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 9, pp. 891–906, 1991.

[2] E.P. Simoncelli and W.T. Freeman, "The steerable pyramid: a flexible architecture for multi-scale derivative computation," in *Proc. International Conference on Image Processing*, 23–26 Oct. 1995, vol. 3, pp. 444–447.

[3] E.M. Stein and G. Weiss, *Introduction to Fourier Analysis on Euclidean Spaces*, Princeton Univ. Press, Princeton, NJ, 1971.

[4] A. Aldroubi, "Portraits of frames," *Proceedings of the American Mathematical Society*, vol. 123, no. 6, pp. 1661–1668, 1995.

[5] M. Unser and D. Van De Ville, "The pairing of a wavelet basis with a mildly redundant analysis via subband regression," *IEEE Trans. Image Processing*, vol. 17, no. 11, pp. 2040–2052, November 2008.

[6] D. Van De Ville, T. Blu, and M. Unser, "Isotropic polyharmonic B-splines: Scaling functions and wavelets," *IEEE Trans. Image Processing*, vol. 14, no. 11, pp. 1798–1813, November 2005.

[7] D. Van De Ville and M. Unser, "Complex wavelet bases, steerability, and the Marr-like pyramid," *IEEE Trans. Image Processing*, vol. 17, no. 11, pp. 2063–2080, November 2008.