

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/40766988>

Wavelet Steerability and the Higher-Order Riesz Transform

Article in IEEE Transactions on Image Processing · March 2010

DOI: 10.1109/TIP.2009.2038832 · Source: PubMed

CITATIONS

91

READS

126

2 authors:



Michael Unser

École Polytechnique Fédérale de Lausanne

969 PUBLICATIONS 41,091 CITATIONS

SEE PROFILE



Dimitri Van De Ville

École Polytechnique Fédérale de Lausanne

678 PUBLICATIONS 10,658 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Image Reconstruction Algorithms [View project](#)



Resting-state fMRI/essential tremor/thalamotomy/radiosurgery/Gamma Knife [View project](#)

Wavelet steerability and the higher-order Riesz transform

Michael Unser, *Fellow, IEEE*, and Dimitri Van De Ville, *Member, IEEE*

Abstract—Our main goal in this paper is to set the foundations of a general continuous-domain framework for designing steerable, reversible signal transformations (a.k.a. frames) in multiple dimensions ($d \geq 2$). To that end, we introduce a self-reversible, N th-order extension of the Riesz transform. 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 functional counterpart to Simoncelli’s steerable pyramid whose construction was primarily based on filterbank 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 with the design of a novel family of multi-dimensional Riesz-Laplace wavelets that essentially behave like the N th-order partial derivatives of an isotropic Gaussian kernel.

I. 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 developing some corresponding feature extraction algorithms and/or signal representation schemes.

The proper handling of the notion of scale in image processing owes much to the work of Mallat who set the foundation of the multiresolution theory of the wavelet transform which involves continuously-defined functions in $L_2(\mathbb{R}^d)$ [1]. The remarkable aspect of this theory is that it yields a reversible one-to-one decomposition of a signal across scale; unfortunately, the underlying wavelet bases (which are typically separable) have uneven angular responses; they are strongly biased towards the vertical and horizontal directions.

The steerable filters of Freeman and Adelson constitute another elegant framework that is more specifically tuned to orientation [2]. The primary focus there is on the image analysis aspect of the problem and the search for an efficient computational scheme for the filtering of an image with a reference template that can be oriented arbitrarily. Several solutions have been suggested for the design of optimized filtering templates [3]–[5]. Researchers have also proposed a unifying Lie-group formulation of steerability that allows for the extension of the concept to others classes of transformations (e.g. translation and dilation) [6]–[8].

It is possible—and highly desirable conceptually—to reconcile the two types of approaches by allowing for wavelet decompositions with a certain level of redundancy [9]. In

particular, Simoncelli and co-workers were able to design a complete family of directional multi-scale transforms which is indexed by the number of orientation bands k [9]–[11]. Their steerable pyramid is over-complete by a factor $4k/3 + 1$ and has the important property of being self-inverting. Steerable wavelets have also been proposed in the redundant framework of the continuous wavelet transform, including some constructions on the sphere [12]. While steerability is attractive conceptually, it is obviously not a strict requirement; other popular solutions to the problem of directional multiresolution image representations include the 2-D Gabor transform [13], curvelets [14], the dual-tree complex wavelet transform [15], [16], directional wavelet frames [17], contourlets [18], bandelets [19], and directionlets [20].

The steerable pyramid has been used quite extensively in image processing. Typical examples of applications include local orientation analysis, adaptive angular filtering, deformable templates [4], contour detection, shape from shading [2], [9], texture analysis and retrieval [21]–[23], directional pattern detection [24], feature extraction for motion tracking [25] and, of course, image denoising [26], [27]. Mathematically, the steerable pyramid is a tight frame of $\ell_2(\mathbb{Z}^2)$, due to the self-reversibility of the underlying discrete filterbank. The link with wavelet theory would be complete if there was a continuous-domain formulation of the transform, which has not been worked out explicitly so far. The other point is that Simoncelli’s construction is quite specific and difficult to generalize to higher dimensions.

The purpose of this paper is to revisit the construction of steerable wavelets and pyramids using a continuous-domain operator-based formalism. The cornerstone of our approach is the Riesz transform which constitutes the natural multi-dimensional generalization of the Hilbert transform. The Riesz transform has a long tradition in mathematics [28] and has been studied extensively in the context of Calderón-Sygmund’s theory of singular integral operators [29]. Its introduction to signal processing is more recent. Felsberg used the transform to define the “monogenic signal” as a 2-D generalization of the analytical signal representation [30], [31]. This pioneering work inspired Metikas and Olhede to specify a monogenic version of the continuous wavelet transform which is a fully redundant signal representation [32]. At about the same time as Felsberg, Larkin independently introduced a complexified version of the 2-D Riesz transform in optics—calling it the spiral quadrature phase transform—and applied it to the demodulation of interferograms and the analysis of fringe patterns [33], [34]. In recent work, we took advantage of the unitary property of this latter transform to specify a proper, reversible wavelet-domain monogenic signal analysis with a

The authors are with the Biomedical Imaging Group (BIG), École Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland.

minimal amount of redundancy [35].

Our present focus on the Riesz transform is motivated by its invariance with respect to the primary coordinate transformations: translation, dilation and rotation (cf. Subsection II.D). The first two invariance properties are ideally matched to the framework of wavelets, while the latter provides a connection with steerable filters and gradient-like signal analyses [35], [36]. The formal extension of this link calls for a higher-order version of the transform, which we introduce in Subsection II.E. Thanks to this operator, we are able to take any multiresolution decomposition of $L_2(\mathbb{R}^d)$ and to map it into a series of directional counterparts. The direct benefits of the proposed methodology are:

- A general framework for constructing families of steerable wavelet frames and pyramids of $L_2(\mathbb{R}^d)$ with arbitrary orders¹ in any number of dimensions $d \geq 2$.
- A clear decoupling between the multiresolution and the steerability aspects of the representation. In the present work, we adopt a less constraining definition of a steerable transform which does not require the basis functions to be rotated versions of one another—the basic requirement is that the component subspaces be rotation-invariant [7]. This offers design flexibility and also carries over naturally to dimensions greater than two.
- A continuous-domain formulation that formalizes and extends the technique proposed by Simoncelli et al.; in essence, it represents the multiresolution theory of $L_2(\mathbb{R}^d)$ counterpart of their original filterbank design. The perfect reconstruction property and closedness in $L_2(\mathbb{R}^d)$ is automatically guaranteed.
- A full control over the regularity (Sobolev smoothness) and approximation theoretic properties (vanishing moments, order of approximation, etc.) of the basis functions.
- The perspective of a purely analytical design of novel families of bona-fide steerable wavelets—in particular, higher-order Riesz-Laplace wavelets which admit explicit closed-form formulas in any number of dimensions.

The paper is organized as follows. In Section II, we review the mathematical properties of the Riesz transform; we then proceed with the specification of a higher-order transform that has the fundamental property of conserving energy. Another crucial property is that the components of our higher-order Riesz transform define a subspace that is rotation-invariant, which is the key to steerability. In Section III, we prove that the Riesz transform and its higher-order extension have the property of mapping a wavelet frame of $L_2(\mathbb{R}^d)$ onto another one, while preserving the frame bounds. We thereby establish a natural correspondence between a (quasi-isotropic) wavelet frame and its directional Riesz counterparts. In Section IV, we explicitly make the connection between the 2-D version of our generalized Riesz transform and Simoncelli's construction. In doing so, we also introduce an extended angular parametrization of a rotating filterbank that ensures self-reversibility. In Section V, we illustrate our concept with the specification of an extended family of high-order steerable

Riesz-Laplace wavelets. The distinguishing features of these wavelets is that they correspond to the N th-order derivatives of a multi-dimensional B-spline kernel. The Gaussian shape of the underlying smoothing kernel, which is common to all wavelets, makes the transform attractive for multi-dimensional feature extraction. We conclude the paper by summarizing the key features of our approach and by discussing the similarities and differences with existing directional transforms.

II. THE RIESZ TRANSFORM AND ITS HIGHER-ORDER EXTENSIONS

The purpose of this section is to lay the mathematical foundations for our approach in a fashion that is self-contained and geared towards a signal processing audience. We start by reviewing the primary features of the Riesz transform with a special emphasis on its invariance properties. We then introduce a novel higher-order extension of the transform that is specifically designed to preserve inner-products (subsection E). The important aspect that we want to bring forth is the link between these generalized Riesz operators and steerable filterbanks, as introduced by Freeman and Adelson.

A. Notations

We consider a generic d -dimensional signal f indexed by the continuous-domain space variable $\mathbf{x} = (x_1, \dots, x_d) \in \mathbb{R}^d$. The multidimensional Fourier transform of $f \in L_1(\mathbb{R}^d)$ is specified as

$$\hat{f}(\boldsymbol{\omega}) = \int_{\mathbb{R}^d} f(\mathbf{x}) e^{-j\langle \boldsymbol{\omega}, \mathbf{x} \rangle} d\mathbf{x}_1 \cdots d\mathbf{x}_d,$$

with $\boldsymbol{\omega} = (\omega_1, \dots, \omega_d) \in \mathbb{R}^d$. This definition is extended appropriately for Lebesgue's space of finite energy functions $L_2(\mathbb{R}^d)$, as well as for Schwartz's class of tempered distributions $\mathcal{S}'(\mathbb{R}^d)$. We recall that the squared L_2 -norm of f is given by

$$\begin{aligned} \|f\|_{L_2(\mathbb{R}^d)}^2 &= \int_{\mathbb{R}^d} |f(\mathbf{x})|^2 d\mathbf{x}_1 \cdots d\mathbf{x}_d \\ &= \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} |\hat{f}(\boldsymbol{\omega})|^2 d\omega_1 \cdots d\omega_d \end{aligned}$$

where the right-hand side follows from Parseval's identity. The adjoint of a linear operator $L : L_2(\mathbb{R}^d) \rightarrow L_2(\mathbb{R}^d)$ is denoted by L^* ; it is such that $\langle f, Lg \rangle_{L_2} = \langle L^*f, g \rangle_{L_2}, \forall f, g \in L_2(\mathbb{R}^d)$.

The fractional Laplacian $(-\Delta)^{s/2}$ with $s \in \mathbb{R}^+$ is the isotropic differential operator of order s whose Fourier-domain definition in the sense of distributions is

$$(-\Delta)^{\frac{s}{2}} f(\mathbf{x}) \xleftrightarrow{\mathcal{F}} \|\boldsymbol{\omega}\|^s \hat{f}(\boldsymbol{\omega}). \quad (1)$$

For $s/2 = n \in \mathbb{N}$, it is a purely local operator that is proportional to the classical n -fold Laplacian. Otherwise, its effect is non-local and described by the following distributional impulse response [37]

$$\frac{C_{s,d}}{\|\mathbf{x}\|^{s+d}} \xleftrightarrow{\mathcal{F}} \|\boldsymbol{\omega}\|^s, \quad (2)$$

¹In this work, we do equate the notions of order of steerability and the order of the Riesz transform.

where $C_{s,d} = \frac{2^s \Gamma(\frac{d+s}{2})}{\pi^{d/2} \Gamma(-s/2)}$ is an appropriate normalization constant and $\Gamma(r) = (r-1)!$ is Euler's gamma function. The definition (1) of the fractional Laplacian can be extended to negative orders as well, which leads to a fractional integral-like behavior. The so-defined family of fractional operators satisfies the composition rule $(-\Delta)^{\alpha_1}(-\Delta)^{\alpha_2} = (-\Delta)^{\alpha_1+\alpha_2}$ which follows directly from their Fourier-domain definition, with the convention that $(-\Delta)^0 = \text{Identity}$. Note, however, that the fractional integral $(-\Delta)^{-s/2}f$ with $s > 0$ is well-defined only for functions (or distributions) f whose moments up to order $\lceil s \rceil$ are zero, due to the singularity of the frequency response at the origin.

We will use the multi-index notation to state some of our results concisely: specifically, $\mathbf{n} = (n_1, \dots, n_d)$ is a d -dimensional multi-index vector with non-negative integers entries; $|\mathbf{n}| = \sum_{i=1}^d n_i$; $\mathbf{z}^{\mathbf{n}} = z_1^{n_1} \dots z_d^{n_d}$ for any $\mathbf{z} = (z_1, \dots, z_d) \in \mathbb{C}^d$; and $\mathbf{n}! = n_1! n_2! \dots n_d!$. With these conventions, the multinomial theorem reads

$$(z_1 + \dots + z_d)^N = \sum_{|\mathbf{n}|=N} \frac{N!}{\mathbf{n}!} \mathbf{z}^{\mathbf{n}},$$

where the summation is over all multi-indices \mathbf{n} satisfying the condition $|\mathbf{n}| = N$.

B. The 1-D Hilbert transform

The Hilbert transform plays a central role in the analytical signal formalism which is a powerful tool for AM/FM analysis [38]. It is a 1-D linear, shift-invariant operator whose defining property is to map all cosine functions into their corresponding sine functions, without affecting their amplitude. The Hilbert operator, denoted by \mathcal{H} , therefore acts as an allpass filter; it is characterized by the transfer function $\hat{h}(\omega) = -j \text{sign}(\omega) = -j\omega/|\omega|$. Its impulse response (in the sense of distributions) is $h(x) = 1/(\pi x)$, which is slowly decaying, expressing the fact that \mathcal{H} is a non-local operator. A potential mathematical difficulty is that $h(x) \notin L_1(\mathbb{R})$ which means that the corresponding filter is not BIBO-stable (bounded input/bounded output), even though it is well-behaved in the L_2 -sense (e.g., bounded and perfectly reversible). A remarkable property in this latter respect is that $\mathcal{H}^{-1} = -\mathcal{H} = \mathcal{H}^*$, indicating that \mathcal{H} is a unitary operator; this follows directly from the Fourier-domain definition.

C. The Riesz transform

The Riesz transform is the natural multidimensional extension of the Hilbert transform [29]. It is the scalar-to-vector signal transformation \mathcal{R} whose frequency-domain definition is

$$\widehat{\mathcal{R}f}(\boldsymbol{\omega}) = -j \frac{\boldsymbol{\omega}}{\|\boldsymbol{\omega}\|} \hat{f}(\boldsymbol{\omega}) \quad (3)$$

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

$$\mathcal{R}f(\mathbf{x}) = \begin{pmatrix} \mathcal{R}_1 f(\mathbf{x}) \\ \vdots \\ \mathcal{R}_d f(\mathbf{x}) \end{pmatrix} = \begin{pmatrix} (h_1 * f)(\mathbf{x}) \\ \vdots \\ (h_d * f)(\mathbf{x}) \end{pmatrix} \quad (4)$$

where the filters $(h_n)_{n=1}^d$ are characterized by their frequency responses $\hat{h}_n(\boldsymbol{\omega}) = -j\omega_n/\|\boldsymbol{\omega}\|$. To obtain explicit space-domain expressions, it is helpful to view these filters as partial derivatives of the function $g(\mathbf{x}) = \mathcal{F}^{-1}\{\|\boldsymbol{\omega}\|^{-1}\}(\mathbf{x})$ (the impulse response of the isotropic integral operator $(-\Delta)^{-\frac{1}{2}}$). Specifically, by making use of the Fourier-transform relation (2) (which is also valid for negative s provided that $s+d \notin -2\mathbb{N}$ [37]), we get

$$g(\mathbf{x}) = \frac{C_{-1,d}}{\|\mathbf{x}\|^{d-1}} = \frac{\frac{1}{2}\Gamma(\frac{d-1}{2})}{\pi^{\frac{d+1}{2}}} \cdot \frac{1}{\|\mathbf{x}\|^{d-1}},$$

which we then differentiate to obtain

$$h_n(x) = -\frac{\partial}{\partial x_n} g(\mathbf{x}) = \frac{(d-1)C_{-1,d}}{\pi^{\frac{d+1}{2}}} \cdot \frac{x_n}{\|\mathbf{x}\|^{d+1}}.$$

The case $d=1$ is excluded from the above argument; in that particular instance, we have $g(x) = \frac{\log|x|}{\pi}$ whose derivative yields $h(x) = -\frac{dg(x)}{dx} = 1/(\pi x)$. The explicit form of these various operators for $d=1, 2, 3, 4$ is recapitulated in Table 1 for further reference.

Note that the impulse responses of the Riesz transform are all anti-symmetric: $\mathcal{R}_n\{\delta\}(\mathbf{x}) = h_n(\mathbf{x}) = -h_n(-\mathbf{x})$ and essentially decaying like $1/\|\mathbf{x}\|^d$ where d is the number of dimensions. While the Riesz transform remains a non-local operator, it is interesting to observe that the delocalization becomes less pronounced in higher dimensions. It is also known from functional analysis that the Riesz transform is a bounded operator for the primary Lebesgue spaces; i.e.,

$$\forall f \in L_p(\mathbb{R}^d), \quad \|\mathcal{R}f\|_{L_p^d} \leq A_{d,p} \cdot \|f\|_{L_p}, \quad 1 < p < \infty \quad (5)$$

where the $A_{d,p}$ are some suitable constants [29]. The singular aspect of the transform is that the extreme cases $p=1, +\infty$ are excluded; this also reflects the fact that the Riesz transform is not stable in the ordinary BIBO sense. Fortunately for us, the situation is particularly favorable for $p=2$, as will be seen next (cf. Property 3).

D. Properties of the multidimensional Riesz transform

We now present and discuss the key properties of the operator \mathcal{R} .

Property 1 (Invariances): The Riesz transform is translation- and scale-invariant:

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

The translation invariance directly follows from the definition, while the scale invariance is easily verified in the Fourier domain.

Remarkably, the Riesz transform is also rotation-invariant. To formalize this property, we consider the group of rotation matrices in \mathbb{R}^d . Specifically, let \mathbf{R}_u denote a $d \times d$ rotation matrix whose first row is the unit vector $\mathbf{u} = (u_1, \dots, u_d)$; i.e., $[\mathbf{R}_u]_{1,n} = [\mathbf{u}]_n = u_n$ (\mathbf{R}_u is such that it rotates the first coordinate vector \mathbf{e}_1 into \mathbf{u}).

TABLE I
RIESZ TRANSFORM COMPONENTS AND RELATED DIFFERENTIAL OPERATORS FOR $d = 1, 2, 3, 4$

Operator	$(-\Delta)^{-\frac{1}{2}}$	$\mathcal{R}_n = -\frac{\partial}{\partial x_n}(-\Delta)^{-\frac{1}{2}}$	$(-\Delta)^{\frac{1}{2}}$
Frequency response	$\frac{1}{\ \boldsymbol{\omega}\ }$	$\frac{-j\omega_n}{\ \boldsymbol{\omega}\ }$	$\ \boldsymbol{\omega}\ $
Impulse responses			
$d = 1$	$\frac{\log x }{\pi}$	$\frac{1}{\pi x}$	$\frac{-1}{\pi x ^2}$
$d = 2$	$\frac{1}{2\pi\ \boldsymbol{x}\ }$	$\frac{x_n}{2\pi\ \boldsymbol{x}\ ^3}$	$\frac{-1}{2\pi\ \boldsymbol{x}\ ^3}$
$d = 3$	$\frac{1}{2\pi^2\ \boldsymbol{x}\ ^2}$	$\frac{x_n}{\pi^2\ \boldsymbol{x}\ ^4}$	$\frac{-1}{\pi^2\ \boldsymbol{x}\ ^4}$
$d = 4$	$\frac{1}{4\pi^2\ \boldsymbol{x}\ ^3}$	$\frac{3x_n}{4\pi^2\ \boldsymbol{x}\ ^5}$	$\frac{-3}{4\pi^2\ \boldsymbol{x}\ ^5}$

Property 2 (Steerability): The Riesz transform filterbank defined by (4) is steerable in the sense that

$$h_1(\mathbf{R}_u \mathbf{x}) = \langle \mathbf{u}, \mathcal{R}\delta(\mathbf{x}) \rangle = \sum_{n=1}^d u_n h_n(\mathbf{x}),$$

while the component filters are 90° rotated versions² of each other; i.e., $h_n(\mathbf{x}) = h_1(\mathbf{R}_{\mathbf{e}_n} \mathbf{x})$ where $[\mathbf{e}_n]_k = \delta_k$.

This is shown by using the rotation property of the Fourier transform

$$\begin{aligned} h_1(\mathbf{R}_u \mathbf{x}) \xleftrightarrow{\mathcal{F}} \hat{h}_1(\mathbf{R}_u \boldsymbol{\omega}) &= -j \frac{[\mathbf{R}_u \boldsymbol{\omega}]_1}{\|\mathbf{R}_u \boldsymbol{\omega}\|} = \sum_{n=1}^d -j \frac{u_n \omega_n}{\|\boldsymbol{\omega}\|} \\ &= \sum_{n=1}^d u_n \hat{h}_n(\boldsymbol{\omega}). \end{aligned}$$

The above result implies that the Riesz transform commutes with rotations, as expressed by the following property of its impulse response

$$\mathcal{R}\{\delta\}(\mathbf{R}_u \mathbf{x}) = \mathbf{R}_u \mathcal{R}\{\delta\}(\mathbf{x})$$

where $\delta(\mathbf{x}) = \prod_{n=1}^d \delta(x_n)$ is the multidimensional Dirac distribution.

Property 2 suggests linking the Riesz transform to the following directional version of the Hilbert transform

$$\mathcal{H}_u f(\mathbf{x}) = \sum_{n=1}^d u_n \mathcal{R}_n f(\mathbf{x}) = \langle \mathbf{u}, \mathcal{R}f(\mathbf{x}) \rangle \quad (6)$$

whose impulse response $h_u(\mathbf{x}) = h_1(\mathbf{R}_u \mathbf{x})$ simply corresponds to the rotated version of $h_1(\mathbf{x})$. Note that the association between the Riesz and the directional Hilbert transforms is essentially the same as the link between the gradient and the directional derivative (more on this later).

²In general, a single vector (or direction) is not sufficient to specify a rotation matrix in d dimensions. In the present case, however, the functions $h_i(\mathbf{x})$ are isotropic within the hyperplane perpendicular to the corresponding coordinate vector \mathbf{e}_i so that our “geometrical” statements implicitly refer to a whole equivalence class of rotation matrices.

It is instructive to determine the Riesz transform of a pure cosine wave with frequency (or wave number) $\boldsymbol{\omega}_0$. A simple Fourier calculation yields

$$\mathcal{R}\{\cos(\boldsymbol{\omega}_0^T \mathbf{x})\}(\mathbf{x}) = \frac{\boldsymbol{\omega}_0}{\|\boldsymbol{\omega}_0\|} \sin(\boldsymbol{\omega}_0^T \mathbf{x}).$$

We can then apply (6) and recover the corresponding sine wave by steering the directional Hilbert transform in the appropriate direction $\mathbf{u}_0 = \frac{\boldsymbol{\omega}_0}{\|\boldsymbol{\omega}_0\|}$:

$$\mathcal{H}_{\mathbf{u}_0}\{\cos(\boldsymbol{\omega}_0^T \mathbf{x})\}(\mathbf{x}) = \sin(\boldsymbol{\omega}_0^T \mathbf{x}).$$

This desirable behavior follows from the fact that the central cut of the frequency response of \mathcal{H}_u along the direction \mathbf{u} perfectly replicates the behavior the 1-D Hilbert transform: $\mathcal{H}_u(\boldsymbol{\omega} \mathbf{u}) = -j \text{sign}(\omega)$. The only limitation of the technique is that the ideal Hilbert-transform-like behavior falls off like $\langle \mathbf{u}, \mathbf{u}_0 \rangle$ (the cosine of the angle), which calls for a precise adjustment of the analysis direction \mathbf{u}_0 .

Property 3 (Inner-product preservation): The Riesz transform satisfies the following Parseval-like identity

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

This is established by using Parseval’s relation for the Fourier transform

$$\begin{aligned} &\sum_{n=1}^d \langle \mathcal{R}_n f, \mathcal{R}_n \phi \rangle_{L_2} \\ &= \frac{1}{(2\pi)^d} \sum_{n=1}^d \int_{\mathbb{R}^d} -j \frac{\omega_n}{\|\boldsymbol{\omega}\|} \hat{f}(\boldsymbol{\omega}) j \frac{\omega_n}{\|\boldsymbol{\omega}\|} \hat{\phi}^*(\boldsymbol{\omega}) d\omega_1 \cdots d\omega_d \\ &= \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \frac{\sum_{n=1}^d \omega_n^2}{\|\boldsymbol{\omega}\|^2} \hat{f}(\boldsymbol{\omega}) \hat{\phi}^*(\boldsymbol{\omega}) d\omega_1 \cdots d\omega_d \\ &= \frac{1}{(2\pi)^d} \langle \hat{f}, \hat{\phi} \rangle_{L_2} = \langle f, \phi \rangle_{L_2} \end{aligned}$$

A similar manipulation leads to the identification of the adjoint operator, which is specified by

$$\mathcal{R}^* \mathbf{r}(\mathbf{x}) = \mathcal{R}_1^* r_1(\mathbf{x}) + \dots + \mathcal{R}_d^* r_d(\mathbf{x}) \xleftarrow{\mathcal{F}} j \frac{\boldsymbol{\omega}^T}{\|\boldsymbol{\omega}\|} \hat{\mathbf{r}}(\boldsymbol{\omega}), \quad (7)$$

where $\mathbf{r}(\mathbf{x}) = (r_1(\mathbf{x}), \dots, r_d(\mathbf{x}))$ is a generic finite-energy d -vector signal and $\mathcal{R}_n^* = -\mathcal{R}_n$. It also corresponds to the left inverse of the transform $\mathcal{R}^{-1} = \mathcal{R}^*$, owing to the fact that $\left(j \frac{\boldsymbol{\omega}^T}{\|\boldsymbol{\omega}\|}\right) \times \left(-j \frac{\boldsymbol{\omega}^T}{\|\boldsymbol{\omega}\|}\right) = 1$. A direct consequence is the energy conservation property

$$\|\mathcal{R}f\|_{L_2^d}^2 = \sum_{n=1}^d \|\mathcal{R}_n f\|_{L_2}^2 = \|f\|_{L_2}^2, \quad (8)$$

which emphasizes the unitary character of the transform.

We conclude this section by listing a few additional functional properties of the components of the Riesz transform for specific classes of signals or basis functions.

Property 4 (Orthogonality): Let $\psi(\mathbf{x}) \in L_2(\mathbb{R}^d)$ be a real-valued function whose Fourier transform $\hat{\psi}$ satisfies the coordinate-wise symmetry relations: $|\hat{\psi}(\omega_1, \dots, \omega_d)|^2 = |\hat{\psi}(\pm\omega_1, \dots, \pm\omega_d)|^2$. Then,

$$\begin{aligned} \langle \mathcal{R}_n \psi, \psi \rangle_{L_2} &= 0, \quad \text{for } n = 1, \dots, d \\ \langle \mathcal{R}_n \psi, \mathcal{R}_m \psi \rangle_{L_2} &= 0, \quad \text{when } m \neq n. \end{aligned}$$

Proof: The direct application of Parseval's relation yields

$$\begin{aligned} &\langle \mathcal{R}_n \psi, \mathcal{R}_m \psi \rangle_{L_2} \\ &= \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} -j \frac{\omega_n}{\|\boldsymbol{\omega}\|} \hat{\psi}(\boldsymbol{\omega}) j \frac{\omega_m}{\|\boldsymbol{\omega}\|} \hat{\psi}^*(\boldsymbol{\omega}) d\omega_1 \dots d\omega_d \\ &= \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \omega_n \omega_m \frac{|\hat{\psi}(\boldsymbol{\omega})|^2}{\|\boldsymbol{\omega}\|^2} d\omega_1 \dots d\omega_d \\ &= \frac{1}{(2\pi)^d} \int_{\mathbb{R}} \omega_n \underbrace{\int_{\mathbb{R}^{d-1}} \omega_m \frac{|\hat{\psi}(\boldsymbol{\omega})|^2}{\|\boldsymbol{\omega}\|^2} d\omega_1 \dots d\omega_d}_{\phi(\omega_n)} d\omega_n = 0 \end{aligned}$$

The partial integration with respect to all frequency variables except ω_n yields the auxiliary function $\phi(\omega_n)$ which is symmetric, as a consequence of our hypotheses. The product function $\omega_n \phi(\omega_n)$ is therefore anti-symmetric, which necessarily leads to a vanishing outer integral. ■

While the assumptions for Property 4 do not hold for arbitrary signals, they are met by most basis functions commonly used for image processing. Indeed, the required symmetry relations are satisfied by all functions that are either separable (because of the Hermitian symmetry of the 1-D Fourier transform), isotropic, or coordinate-wise symmetric/anti-symmetric in the space domain (up to some arbitrary translation).

The smoothness of a function is often specified by its inclusion in the Sobolev space W_2^s (the class of functions that are “ s -times” differentiable in the L_2 -sense) [39]. A nice feature of the Riesz transform is that it fully preserves Sobolev regularity.

Property 5 (Sobolev smoothness): $f \in W_2^s(\mathbb{R}^d) \Leftrightarrow \mathcal{R}_n f \in W_2^s(\mathbb{R}^d)$ for $n = 1, \dots, d$.

Proof: The explicit definition of Sobolev's space of order s is:

$$W_2^s(\mathbb{R}^d) = \left\{ f(\mathbf{x}), \mathbf{x} \in \mathbb{R}^d : \|f\|_{L_2}^2 + \|(-\Delta)^{s/2} f\|_{L_2}^2 < +\infty \right\}$$

where $(-\Delta)^{s/2}$ is the isotropic derivative operator of order s [cf. (1)]. The Sobolev smoothness equivalence then directly follows from the energy preservation property (8) and the translation invariance of \mathcal{R} , which implies that $\|f\|_{L_2}^2 = \sum_{n=1}^d \|\mathcal{R}_n f\|_{L_2}^2 > \|\mathcal{R}_n f\|_{L_2}^2$, and $\|(-\Delta)^{s/2} f\|_{L_2}^2 = \sum_{n=1}^d \|\mathcal{R}_n (-\Delta)^{s/2} f\|_{L_2}^2 > \|(-\Delta)^{s/2} \mathcal{R}_n f\|_{L_2}^2$. ■

In this work, we will use the Riesz transform to map a set of primary wavelets into another augmented one. Property 1 and 3 are essential for that purpose because they ensure that the mapping preserves the scale- and shift-invariant structure as well as the L_2 -stability of the representation. Note that the latter goes hand-in-hand with the reversibility of the transform and its ability to perfectly represent all finite-energy signals. Thanks to Property 5, we also have the guarantee that the Riesz wavelets are functionally well-behaved in that they retain the same regularity (degree of differentiability) as the primary templates from which they are derived. Property 4 is relevant as well because it enforces a certain level of coordinate-wise orthogonality which is potentially favorable for local feature extraction and signal analysis. Finally, we will be able to take advantage of Property 2 to design wavelet transforms whose basis functions are steerable.

E. Higher-order Riesz transform

It is also possible to define higher-order Riesz transforms by considering individual signal components of the form $\mathcal{R}_{i_1} \mathcal{R}_{i_2} \dots \mathcal{R}_{i_N} f$ with $i_1, i_2, \dots, i_N \in \{1, \dots, d\}$. While there are d^N possible ways of forming such N th 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 N th-order Riesz transform achieves the following decomposition of the identity

$$\sum_{|\mathbf{n}|=N} \frac{N!}{\mathbf{n}!} (\mathcal{R}_1^{n_1} \dots \mathcal{R}_d^{n_d})^* (\mathcal{R}_1^{n_1} \dots \mathcal{R}_d^{n_d}) = \text{Id}$$

using the multi-index vector $\mathbf{n} = (n_1, \dots, n_d)$.

Here, \mathcal{R}_n^k denotes the k -fold iterate of \mathcal{R}_n , while the weighting factors in the expansion are the multinomial coefficients of order $|\mathbf{n}| = N$

$$\frac{|\mathbf{n}|!}{\mathbf{n}!} = \binom{N}{n_1, \dots, n_d} = \frac{N!}{n_1! n_2! \dots n_d!}.$$

Proof: By applying the Theorem's left-hand side decomposition to an input signal f , we obtain a series of components of the form

$$g_{n_1, \dots, n_d}(\mathbf{x}) = (\mathcal{R}_1^{n_1} \dots \mathcal{R}_d^{n_d})^* (\mathcal{R}_1^{n_1} \dots \mathcal{R}_d^{n_d}) f(\mathbf{x}).$$

This corresponds to a convolution whose Fourier-domain equivalent is

$$\begin{aligned}
\hat{g}_{n_1, \dots, n_d}(\boldsymbol{\omega}) &= \left(\frac{j\boldsymbol{\omega}}{\|\boldsymbol{\omega}\|} \right)^{\mathbf{n}} \left(\frac{-j\boldsymbol{\omega}}{\|\boldsymbol{\omega}\|} \right)^{\mathbf{n}} \hat{f}(\boldsymbol{\omega}) \\
&= \left(\frac{|\omega_1|^2}{\|\boldsymbol{\omega}\|^2} \right)^{n_1} \cdots \left(\frac{|\omega_d|^2}{\|\boldsymbol{\omega}\|^2} \right)^{n_d} \hat{f}(\boldsymbol{\omega}) \\
&= \frac{|\omega_1|^{2n_1} \cdots |\omega_d|^{2n_d}}{\|\boldsymbol{\omega}\|^{2N}} \hat{f}(\boldsymbol{\omega}). \tag{9}
\end{aligned}$$

Now, we obviously have the identity

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

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

$$\begin{aligned}
\hat{f}(\boldsymbol{\omega}) &\stackrel{\text{N!}}{\sum_{|\mathbf{n}|=N}} \frac{|\omega_1|^{2n_1} \cdots |\omega_d|^{2n_d}}{\|\boldsymbol{\omega}\|^{2N}} \hat{f}(\boldsymbol{\omega}) \\
&\stackrel{\text{N!}}{\sum_{|\mathbf{n}|=N}} \frac{N!}{\mathbf{n}!} \hat{g}_{n_1, \dots, n_d}(\boldsymbol{\omega})
\end{aligned}$$

where we have identified the relevant higher-order reconstruction components using (9). ■

From the above, it is clear that an N th order Riesz transform decomposes a signal into $p(N, d)$ distinct components, where $p(N, d) = \binom{N+d-1}{d-1}$ is the cardinality of the indexing set $\{\mathbf{n} \in \mathbb{N}^d \text{ s.t. } |\mathbf{n}| = N\}$. We also want to ensure that the transform preserves energy which calls for an appropriate normalization. We therefore specify the N th order Riesz transform of the signal f as the $p(N, d)$ -vector signal transformation whose components are given by

$$f^{\mathbf{n}} = \mathcal{R}^{\mathbf{n}} f = \sqrt{\frac{N!}{n_1! \cdots n_d!}} \mathcal{R}_1^{n_1} \cdots \mathcal{R}_d^{n_d} f. \tag{10}$$

In some instances, we will use the more concise vector notation $\mathcal{R}^{(N)} f$, keeping in mind that each component is associated with its own normalized, N th-order Riesz operator $\mathcal{R}^{(n_1, \dots, n_d)}$, as defined above.

Note that the number of higher-order components roughly increases like the $(d-1)$ th power of the dimension. For instance, $p(N, 1) = 1$, $p(N, 2) = N+1$, and $p(N, 3) = (N+2)(N+1)/2$. The good news is that the increase is moderate in dimension 2 (the context of image processing) and still manageable in dimension 3 (volumetric imaging). The case $d=1$ is uninteresting since the higher-order transform cycles back and forth between the identity and the conventional Hilbert transform.

Not too surprisingly, the higher-order Riesz transform inherits all the fundamental properties of the operator \mathcal{R} , as these are preserved through the iteration mechanism.

Property 6: The higher-order Riesz transforms $\mathcal{R}^{(N)}$ are translation and scale-invariant.

What is more remarkable is the higher-order inner-product preservation property—it comes as a corollary of the decomposition of the identity given by Theorem 1.

Property 7: The N th-order Riesz transform satisfies the following Parseval-like identity

$$\begin{aligned}
\forall f, \phi \in L_2(\mathbb{R}^d), \langle \mathcal{R}^{(N)} f, \mathcal{R}^{(N)} \phi \rangle_{L_2^{p(N,d)}} &= \\
\sum_{|\mathbf{n}|=N} \langle \mathcal{R}^{\mathbf{n}} f, \mathcal{R}^{\mathbf{n}} \phi \rangle_{L_2} &= \langle f, \phi \rangle_{L_2}.
\end{aligned}$$

In particular, by setting $f = \phi$ in the above property, we observe that the higher-order Riesz transform perfectly preserves the L_2 -norm of the signal. The generalized version of this result is the boundedness of the higher-order Riesz transform for the classical Lebesgue spaces (with the notable exception of $p = 1, +\infty$).

Property 8: The higher-order Riesz transform satisfies the following norm inequalities:

$$\forall f \in L_p(\mathbb{R}^d), \|\mathcal{R}^{(N)} f\|_{L_p^{p(N,d)}} \leq A_{d,p,N} \|f\|_{L_p} \quad 1 < p < \infty$$

where the $A_{d,p,N}$ are appropriate constants.

Since the higher-order Riesz transform can be decomposed into a succession of first-order transforms, the result is a consequence of the L_p boundedness of all elementary component transformations (cf. (5)).

The final important feature is rotation invariance; it is intimately linked to the concept of steerability which is discussed in the next section.

F. Higher-order steerability

The most interesting aspect of the higher-order Riesz transform is its steerability which also comes hand-in-hand with improved angular selectivity. To best explain this behavior, we relate $\mathcal{R}^{(N)}$ to the N -fold version of the directional Hilbert transform $\mathcal{H}_{\mathbf{u}}$ along the direction specified by the unit vector \mathbf{u} . To that end, we write the Fourier-domain equivalent of (6)

$$\widehat{\mathcal{H}_{\mathbf{u}} f}(\boldsymbol{\omega}) = \sum_{i=1}^d u_i \frac{-j\omega_i}{\|\boldsymbol{\omega}\|} \hat{f}(\boldsymbol{\omega}). \tag{11}$$

We then iterate the operator N times and expand the right-hand side using the multinomial theorem

$$\begin{aligned}
\widehat{\mathcal{H}_{\mathbf{u}}^N f}(\boldsymbol{\omega}) &= \left(-j \frac{\langle \mathbf{u}, \boldsymbol{\omega} \rangle}{\|\boldsymbol{\omega}\|} \right)^N \hat{f}(\boldsymbol{\omega}) \\
&= \left(\sum_{i=1}^d u_i \frac{-j\omega_i}{\|\boldsymbol{\omega}\|} \right)^N \hat{f}(\boldsymbol{\omega}) \\
&= \sum_{|\mathbf{n}|=N} \frac{N!}{\mathbf{n}!} \mathbf{u}^{\mathbf{n}} \left(\frac{-j\boldsymbol{\omega}}{\|\boldsymbol{\omega}\|} \right)^{\mathbf{n}} \hat{f}(\boldsymbol{\omega}).
\end{aligned} \tag{12}$$

Finally, we go back to the space domain, which yields

$$\begin{aligned}
\mathcal{H}_{\mathbf{u}}^N f(\mathbf{x}) &= \sum_{|\mathbf{n}|=N} \frac{N!}{\mathbf{n}!} u_1^{n_1} \mathcal{R}_1^{n_1} \cdots u_d^{n_d} \mathcal{R}_d^{n_d} f(\mathbf{x}) \\
&= \sum_{|\mathbf{n}|=N} \sqrt{\frac{N!}{\mathbf{n}!}} \mathbf{u}^{\mathbf{n}} \mathcal{R}^{\mathbf{n}} f(\mathbf{x}) \\
&= \sum_{|\mathbf{n}|=N} c_{n_1, \dots, n_d}(\mathbf{u}) f^{(n_1, \dots, n_d)}(\mathbf{x}) \tag{13}
\end{aligned}$$




Fig. 1. Steerable filterbank implementation of the N th order directional Hilbert transform in multiple dimensions. The output signal is formed by taking an appropriate linear combination of the components of the N th order Riesz transform of the signal. The scheme is adaptive in the sense the direction of analysis \mathbf{u} may vary spatially.

where the “steering” coefficients are given by

$$c_{n_1, \dots, n_d}(\mathbf{u}) = \sqrt{\frac{N!}{n_1! \dots n_d!}} (u_1^{n_1} \dots u_d^{n_d}). \quad (14)$$

The signal processing transcription of (13) leads to the block diagram in Fig. 1. The corresponding system falls within the general framework of steerable filterbanks of Freeman and Adelson [2], although these authors did not pursue their concept for dimensions larger than 3. The main point is that we can apply a $p(N, d)$ -channel filterbank with impulse responses $\{h_{\mathbf{n}} = \mathcal{R}^{\mathbf{n}} \delta\}_{(|\mathbf{n}|=N)}$ to the signal f and then compute $\mathcal{H}_{\mathbf{u}}^N f(\mathbf{x})$ along any direction \mathbf{u} by forming an appropriate linear combination with the weights given by (14). While some authors define the order of steerability as the dimension of the invariant subspace $k = p(N, d)$,

The argument also extends for generalized Riesz operators of the form $\mathbf{a}^T \mathcal{R}^{(N)}$ where \mathbf{a} is a given $p(N, d)$ -vector of weights. This implies that the function space $V_{N,d} = \text{span}\{h_{\mathbf{n}}(\mathbf{x})\}_{(|\mathbf{n}|=N)}$ is rotation-invariant and that there is a simple weighting mechanism for steering generalized Riesz operators in arbitrary directions. Following this line of thought, it is also possible to define generalized higher-order Riesz operators of the form $\mathbf{A} \mathcal{R}^{(N)}$ where \mathbf{A} is a $p(N, d) \times p(N, d)$ non-singular complex matrix. This generalized operator will share all the fundamental properties of $\mathcal{R}^{(N)}$ (shift-invariance, scale-invariance and steerability) and will preserve inner-products provided that $\mathbf{A} = \mathbf{U}$ is a unitary matrix. Indeed, such a representation leads to the following decomposition of the identity which comes as a corollary of Theorem 1:

$$\begin{aligned} (\mathbf{U} \mathcal{R}^{(N)})^* (\mathbf{U} \mathcal{R}^{(N)}) &= (\mathcal{R}^{(N)})^* \underbrace{\mathbf{U}^H \mathbf{U}}_{\mathbf{I}} \mathcal{R}^{(N)} \\ &= (\mathcal{R}^{(N)})^* \mathcal{R}^{(N)} = \text{Id} \end{aligned} \quad (15)$$

This simple mechanism offers additional degrees of freedom for shaping higher-order impulse responses and designing steerable, reversible filterbanks in arbitrary dimensions.

It is also instructive to investigate the directional behavior of the operator $\mathcal{H}_{\mathbf{u}}^N$ which corresponds to the \mathbf{u} -directional version of the operator \mathcal{R}_1^N . The important point is that its

frequency response $\widehat{\mathcal{H}_{\mathbf{u}}^N}(\boldsymbol{\omega}) = \left(-j \frac{\langle \mathbf{u}, \boldsymbol{\omega} \rangle}{\|\boldsymbol{\omega}\|}\right)^N = (-j \cos \theta)^N$ does not depend upon the radial frequency $\omega = \|\boldsymbol{\omega}\|$, but only upon the angle θ between $\boldsymbol{\omega}$ and the direction vector \mathbf{u} . It is maximum along the direction $\theta = 0$; along that ray, $\widehat{\mathcal{H}_{\mathbf{u}}^N}(\mathbf{u}\omega) = (-j \text{sgn}(\omega))^N$ which coincides (up to a sign change) with the 1-D identity when N is even (symmetric operator) and the 1-D Hilbert transform otherwise (anti-symmetric operator). As θ increases from 0 to $\pi/2$, the response falls off like $\cos(\theta)^N$ —a clear indication that angular selectivity improves with N .

G. Connection with the gradient and partial-derivatives

There is a direct connection between the Riesz transform and the gradient, which is easily seen in the frequency domain. Specifically, we have that

$$\nabla f(\mathbf{x}) = (-1)(-\Delta)^{\frac{1}{2}} \mathcal{R} f(\mathbf{x}) \quad (16)$$

where $(-\Delta)^{\frac{1}{2}}$ is the square-root-Laplace operator whose frequency is $\|\boldsymbol{\omega}\|$ [cf. Table 1]. The converse relation is

$$\mathcal{R} f(\mathbf{x}) = (-1)(-\Delta)^{-\frac{1}{2}} \nabla f(\mathbf{x}), \quad (17)$$

which should be interpreted as a smoothed version of the gradient (up to a sign change); the integral operator $(-\Delta)^{-\frac{1}{2}}$ acts on all derivative components and has an isotropic smoothing effect.

The directional counterparts of this last equation is:

$$\mathcal{H}_{\mathbf{u}} f(\mathbf{x}) = (-1)(-\Delta)^{-\frac{1}{2}} \mathbf{D}_{\mathbf{u}} f(\mathbf{x}) \quad (18)$$

where $\mathcal{H}_{\mathbf{u}}$ is defined by (6) and where $\mathbf{D}_{\mathbf{u}} f = \langle \nabla f, \mathbf{u} \rangle$ is the derivative of f along the direction \mathbf{u} .

By iteration, we obtain the link with the partial derivatives of order $N = n_1 + \dots + n_d$:

$$\mathcal{R}_1^{n_1} \dots \mathcal{R}_d^{n_d} f(\mathbf{x}) = (-1)^N (-\Delta)^{-\frac{N}{2}} \frac{\partial^N f(\mathbf{x})}{\partial^{n_1} x_1 \dots \partial^{n_d} x_d} \quad (19)$$

and the higher-order directional derivatives:

$$\mathcal{H}_{\mathbf{u}}^N f(\mathbf{x}) = (-1)^N (-\Delta)^{-\frac{N}{2}} \mathbf{D}_{\mathbf{u}}^N f(\mathbf{x}) \quad (20)$$

The crucial element here is the fractional integral operator of order N : $(-\Delta)^{-\frac{N}{2}} \xleftrightarrow{\mathcal{F}} 1/\|\boldsymbol{\omega}\|^N$. It is an isotropic lowpass filter whose smoothing strength increases with N . The global interpretation of these relations is that there is a one-to-one connection between Riesz components and spatial derivatives with the former being smoothed versions of the latter. The Riesz transform therefore captures the same directional information as derivatives but it has the advantage of being much better conditioned since there is no amplification of high frequencies. We will now see how we can exploit its invariance and reversibility properties to construct new steerable wavelet-type decompositions.

III. GENERAL CONSTRUCTION OF STEERABLE WAVELET FRAMES

The pleasing consequence of Property 3 is that the Riesz transform will automatically map any frame of $L_2(\mathbb{R}^d)$ into

another one. This fundamental observation was made independently by Held et al. [36] in full generality, and by us in the more constrained 2-D setting, considering a complex version of the Riesz transform that defines a unitary mapping from $L_2(\mathbb{R}^2)$ into itself [35]. Here, we will present a proper N th-order extension of these ideas. Before establishing the results, we start by recalling a few basic facts about frames [40]–[42].

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

$$\forall f \in L_2(\mathbb{R}^d), \quad A \|f\|_{L_2}^2 \leq \sum_{\mathbf{k} \in \mathbb{Z}^d} |\langle \phi_{\mathbf{k}}, f \rangle_{L_2}|^2 \leq B \|f\|_{L_2}^2. \quad (21)$$

The frame is called *tight* iff. the frame bounds can be chosen such that $A = B$. If $A = B = 1$, we have a *Parseval frame* that satisfies the remarkable decomposition/reconstruction formula

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

While the form of this expansion is the same as that associated with an orthonormal basis, the fundamental aspect of the frame generalization is that the family $(\phi_{\mathbf{k}})$ may be redundant. More generally, we have that

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

where $(\tilde{\phi}_{\mathbf{k}})$ is a so-called dual frame. While the dual frame is not unique in the redundant case, there is one—the so-called canonical one—that is used preferentially since it maps into the minimum-norm inverse [42].

We are now ready to state the fundamental frame preservation property of the N th order Riesz transform, which generalizes the first-order result of Held et al.

Proposition 1: The N th order Riesz transform maps the frame $\{\phi_{\mathbf{k}}\}_{\mathbf{k} \in \mathbb{Z}^d}$ of $L_2(\mathbb{R}^d)$ into another frame $\{\mathcal{R}^{\mathbf{n}} \phi_{\mathbf{k}}\}_{(|\mathbf{n}|=N, \mathbf{k} \in \mathbb{Z}^d)}$ that has the same frame bounds and a redundancy increased by $p(N, d) = \binom{N+d-1}{d-1}$.

Proof: First we observe that $\|\mathcal{R}^{\mathbf{n}} f\|_{L_2} = \|\mathcal{R}^{\mathbf{n}*} f\|_{L_2} \leq \|f\|_{L_2}$ where $\mathcal{R}^{\mathbf{n}*} = (-1)^N \mathcal{R}^{\mathbf{n}}$ is the adjoint of the \mathbf{n} -component Riesz operator $\mathcal{R}^{\mathbf{n}}$ defined by (10). We can therefore apply the frame inequality (21) to each (adjoint) component of the N th-order Riesz transform:

$$A \|\mathcal{R}^{\mathbf{n}*} f\|_{L_2}^2 \leq \sum_{\mathbf{k} \in \mathbb{Z}^d} |\langle \phi_{\mathbf{k}}, \mathcal{R}^{\mathbf{n}*} f \rangle_{L_2}|^2 \leq B \|\mathcal{R}^{\mathbf{n}*} f\|_{L_2}^2$$

Next, we sum up these inequalities over the multi-index components $|\mathbf{n}| = N$, which yields

$$A \|f\|_{L_2}^2 \leq \sum_{\mathbf{k} \in \mathbb{Z}^d} \sum_{|\mathbf{n}|=N} |\langle \phi_{\mathbf{k}}, \mathcal{R}^{\mathbf{n}*} f \rangle_{L_2}|^2 \leq B \|f\|_{L_2}^2$$

where we have made use of the energy preservation feature of the Riesz transform which is implied by Property 7. The proposition then simply follows from the adjoint-defining relation $\langle \phi_{\mathbf{k}}, \mathcal{R}^{\mathbf{n}*} f \rangle_{L_2} = \langle \mathcal{R}^{\mathbf{n}} \phi_{\mathbf{k}}, f \rangle_{L_2}$. ■

What is of even greater significance to us is that the N th-order Riesz transform 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:

$$\forall f \in L_2(\mathbb{R}^d), \quad f(\mathbf{x}) = \sum_{i \in \mathbb{Z}} \sum_{\mathbf{k} \in \mathbb{Z}^d} \langle f, \psi_{i,\mathbf{k}} \rangle_{L_2} \tilde{\psi}_{i,\mathbf{k}}(\mathbf{x}), \quad (22)$$

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

For convenience, we now use a multi-index representation of the corresponding N th-order Riesz wavelets, which are defined as follows

$$\psi_{i,\mathbf{k}}^{\mathbf{n}}(\mathbf{x}) = \mathcal{R}^{\mathbf{n}} \psi_{i,\mathbf{k}} = \sqrt{\frac{N!}{n_1! \cdots n_d!}} \mathcal{R}_1^{n_1} \cdots \mathcal{R}_d^{n_d} \psi_{i,\mathbf{k}}, \quad (23)$$

with $\mathbf{n} = (n_1, \dots, n_d)$ and $|\mathbf{n}| = \sum_{i=1}^d n_i = N$.

Proposition 2: Let $\{\psi_{i,\mathbf{k}}\}$ be a primal wavelet frame of $L_2(\mathbb{R}^d)$ associated with the reconstruction formula (22). Then, $\{\psi_{i,\mathbf{k}}^{\mathbf{n}}\}_{(|\mathbf{n}|=N, \mathbf{k} \in \mathbb{Z}^d)}$ and $\{\tilde{\psi}_{i,\mathbf{k}}^{\mathbf{n}}\}_{(|\mathbf{n}|=N, \mathbf{k} \in \mathbb{Z}^d)}$ form a dual set of wavelet frames satisfying the general decomposition/reconstruction formula

$$\forall f \in L_2(\mathbb{R}^d), \quad f(\mathbf{x}) = \sum_{i \in \mathbb{Z}} \sum_{\mathbf{k} \in \mathbb{Z}^d} \sum_{|\mathbf{n}|=N} \langle f, \psi_{i,\mathbf{k}}^{\mathbf{n}} \rangle_{L_2} \tilde{\psi}_{i,\mathbf{k}}^{\mathbf{n}}(\mathbf{x}).$$

Proof: Theorem 1 is equivalent to the following decompositions of the identity operator:

$$\sum_{|\mathbf{n}|=N} \mathcal{R}^{\mathbf{n}*} \mathcal{R}^{\mathbf{n}} = \sum_{|\mathbf{n}|=N} \mathcal{R}^{\mathbf{n}} \mathcal{R}^{\mathbf{n}*} = \text{Id}, \quad (24)$$

which allows us to interchange the role $\mathcal{R}^{\mathbf{n}}$ and $\mathcal{R}^{\mathbf{n}*}$ when analyzing (resp. synthesizing) a signal. First, we apply the wavelet decomposition (22) to the transformed signals $\mathcal{R}^{\mathbf{n}*} f \in L_2(\mathbb{R}^d)$ for all \mathbf{n} such that $|\mathbf{n}| = N$:

$$\begin{aligned} \mathcal{R}^{\mathbf{n}*} f(\mathbf{x}) &= \sum_{i \in \mathbb{Z}} \sum_{\mathbf{k} \in \mathbb{Z}^d} \langle \mathcal{R}^{\mathbf{n}*} f, \psi_{i,\mathbf{k}} \rangle_{L_2} \tilde{\psi}_{i,\mathbf{k}}(\mathbf{x}) \\ &= \sum_{i \in \mathbb{Z}} \sum_{\mathbf{k} \in \mathbb{Z}^d} \langle f, \mathcal{R}^{\mathbf{n}} \psi_{i,\mathbf{k}} \rangle_{L_2} \tilde{\psi}_{i,\mathbf{k}}(\mathbf{x}) \quad (\text{by duality}). \end{aligned}$$

We then make use of (24) to resynthesize f :

$$f(\mathbf{x}) = \sum_{|\mathbf{n}|=N} \mathcal{R}^{\mathbf{n}} \mathcal{R}^{\mathbf{n}*} f(\mathbf{x}),$$

which yields the desired reconstruction formula by linearity. The final ingredient for this construction is Property 6, which ensures that the wavelet structure is conserved; i.e., $\mathcal{R}^{\mathbf{n}} \psi_{i,\mathbf{k}}(\mathbf{x}) = 2^{-id/2} \mathcal{R}^{\mathbf{n}} \psi_{0,\mathbf{k}}(\mathbf{x}/2^i)$, as well as the basic shift-invariant structure that involves a finite number of generators which are the higher-order Riesz transforms of the primary mother wavelets. ■

If, in addition, the mother wavelets in (22) are (quasi-)isotropic, then we end up with a decomposition that is steerable (cf. Section II.F). For instance, we can use the steering equation (13) to rotate the first component wavelet $\psi_{i,\mathbf{k}}^{(N,0,\dots,0)} = \mathcal{R}^{(N,0,\dots,0)} \psi_{i,\mathbf{k}}$ to any desired orientation:

$$\mathcal{H}_{\mathbf{u}}^N \psi_{i,\mathbf{k}}(\mathbf{x}) = \sum_{|\mathbf{n}|=N} c_{\mathbf{n}}(\mathbf{u}) \psi_{i,\mathbf{k}}^{\mathbf{n}} = \psi_{i,\mathbf{k}}^{(N,0,\dots,0)}(\mathbf{R}_{\mathbf{u}} \mathbf{x}), \quad (25)$$

where the right-hand side of the equality is valid provided that the primal wavelet $\psi(\mathbf{x}) = \psi(\|\mathbf{x}\|)$ is isotropic. Similar steering equations can be derived for all other Riesz wavelet components. Concretely, this means that the global wavelet analysis is rotation-covariant.

Finally, if the wavelet decomposition (22) is build around a Mallat-type multiresolution analysis and admits a fast filterbank algorithm, then there is a corresponding perfect reconstruction filterbank implementation for its N th-order Riesz counterpart in Proposition 2.

IV. THE GENERALIZED RIESZ TRANSFORM AND 2-D STEERABLE FILTERBANKS

In 2-D, there is an interesting connection between the generalized Riesz transform (as defined in §II.E), steerable filterbanks that are polar-separable, and the steerable pyramid of Simoncelli et al. [11].

The N th-order Riesz transform for $d = 2$ has $N + 1$ distinct components. The corresponding frequency responses span the subspace $\hat{V}_{N,2} = \text{span} \left\{ \hat{h}_{n_1, N-n_1}(\boldsymbol{\omega}) \right\}_{n_1=0}^N$, which is rotation-invariant—this is simply the frequency-domain transposition of what has already been stated for the space domain in §II.F. The specific form of these filters is

$$\begin{aligned} \hat{h}_{n_1, N-n_1}(\boldsymbol{\omega}) &= \sqrt{\binom{N}{n_1}} \left(\frac{-j\omega_1}{\|\boldsymbol{\omega}\|} \right)^{n_1} \left(\frac{-j\omega_2}{\|\boldsymbol{\omega}\|} \right)^{N-n_1} \\ &= (-j)^N \sqrt{\binom{N}{n_1}} (\cos \theta)^{n_1} (\sin \theta)^{N-n_1} \\ &= H_{n_1, N-n_1}(e^{j\theta}) \end{aligned}$$

where we have made the polar-coordinate substitution $\cos \theta = \frac{\omega_1}{\sqrt{\omega_1^2 + \omega_2^2}}$ and $\sin \theta = \frac{\omega_2}{\sqrt{\omega_1^2 + \omega_2^2}}$. The main point is that these basis functions are purely polar and that they are completely characterized by the 2π -periodic radial profile functions

$$H_{m,n}(z) = (-j)^{m+n} \sqrt{\binom{m+n}{m}} \left(\frac{z+z^{-1}}{2} \right)^m \left(\frac{z-z^{-1}}{2j} \right)^n. \quad (26)$$

with $z = e^{j\theta}$. The complete set of 2-D basis functions for $N = 2, 3, 4$ are displayed in Fig. 2.

For any given order, we can associate the filters in pairs $(\hat{h}_{n_1, n_2}(\boldsymbol{\omega}), \hat{h}_{n_2, n_1}(\boldsymbol{\omega}))$ that are rotated with respect to each other by 90° . As the order increases, the frequency patterns become more intricate and distinct from those observed in conventional steerable filterbanks where all components are rotated versions of a single template. The case $N = 2$ calls for a comment: even though the central pattern looks visually similar to the two others, the $(1, 1)$ filter is not a rotated version of the two others because one also has to take the range of the display into consideration; e.g., (black=-1, white=0) for $(n_1, n_2) \in \{(2, 0), (0, 2)\}$ and (black=-1, white=1) otherwise.

As far as steerability is concerned, the important property is that the profile functions are angularly “bandlimited” [2]; that is, that their Fourier series have a limited number of terms. In fact, they all have the form of a N -term truncated




Fig. 2. Frequency responses (real or imaginary part) of the components filters of the first, second and third order Riesz transform. The origin of the 2-D frequency domain is in the center and the intensity scale is stretched linearly for maximum contrast. The plots are labeled by (n_1, n_2) with $n_2 = N - n_1$.




Fig. 3. Frequency responses of the angular filters of Simoncelli’s third (a) and fourth (b) order steerable pyramids. The origin of the 2-D frequency domain is in the center and the intensity scale is stretched linearly for maximum contrast.

Fourier series $H_{n_1, n_2}(e^{j\theta}) = \sum_{k=-N}^N c_{n_1, n_2}[k] e^{jk\theta}$ with $c_{n_1, n_2}[\pm N] \neq 0$, as is readily checked by expanding (26).

The idea of the generalized Riesz transform is to construct new basis functions by taking suitable linear combinations of the primary ones. Each of these basis functions will be polar as well and characterized by a generic profile function

$$H(e^{j\theta}) = \sum_{k=-N}^N c[k] e^{jk\theta}. \quad (27)$$

A truly remarkable property is that we can use *any* profile function to define a self-inverting generalized Riesz transform that has the structure of a steerable filterbank corresponding to the block diagram in Fig. 4.

Theorem 2: Let $H(e^{j\theta}) = \sum_{k=-N}^N c[k] e^{jk\theta}$ where the $c[k]$ ’s are arbitrary real-valued (or purely imaginary) coefficients. Then, the $(N + 1)$ -channel filterbank whose filters for




Fig. 4. 2-D self-invertible steerable filterbank. The filters are rotated versions of each other with an equi-angular spacing. Their frequency responses are purely polar and angularly bandlimited. They span the same space as the components of the N th order Riesz transform.

$n = 0, \dots, N$ are given by

$$\hat{h}_n(\omega \cos \theta, \omega \sin \theta) = \frac{H(e^{j(\theta + \frac{\pi n}{N+1})})}{\sqrt{(N+1) \sum_{k=-N}^N |c[k]|^2}}$$

is steerable and self-inverting.

Proof: We make the substitution $z = e^{j\theta}$ and consider the Laurent polynomial $H(z) = \sum_{k=-N}^N c[k]z^k$. The crucial observation is that the coefficients $a[k]$ of the product polynomial

$$H(z)H(z^{-1}) = \sum_{k=-2N}^{2N} a[k]z^k$$

correspond to the autocorrelation of the sequence $c[k]$ (of length $2N+1$); in particular, this implies that the terms with index $|k| > 2N$ are necessarily zero. Therefore, if we down-sample the sequence $a[k]$ by a factor $(2N+2)$, we are left with a single non-zero coefficient at the origin: $a[0] = \sum_{k=-N}^N |c[k]|^2$. In the frequency domain, this down-sampling operation corresponds to a periodization, leading to the identity

$$\frac{1}{2N+2} \sum_{n=0}^{2N+1} |H(e^{j(\theta + \frac{2\pi n}{2N+2})})|^2 = a[0],$$

where the right-hand side is the Fourier transform of the remaining impulse. Next we note that $|H(e^{j\theta})|^2$ is π -periodic since the coefficients are real-valued (or purely imaginary). We can therefore divide the sum in two equal parts, which leads to the partition of unity formula

$$\frac{1}{a[0](N+1)} \sum_{n=0}^N |H(e^{j(\theta + \frac{\pi n}{N+1})})|^2 = 1.$$

This is equivalent to

$$\sum_{n=0}^N |\hat{h}_n(\omega)|^2 = 1,$$

which ensures that the filterbank in Fig. 4 has the perfect reconstruction property. ■

While the proof relies on rather basic signal processing concepts, this result does not appear to have been reported before; at least not in such general terms. It allows for generalizations of the design of Simoncelli et al. who used the profile function




Fig. 5. Original 256×256 cameraman with subimage.

$H_N(e^{j\theta}) = (\cos \theta)^N$ to construct their steerable pyramid, having noted that this function would satisfy the required self-reversibility condition (up to some unspecified normalization constant). The autocorrelation filter in that particular case is

$$H_N(z)H_N(z^{-1}) = \left(\frac{z+z^{-1}}{2}\right)^{2N} = \frac{1}{2^{2N}} \sum_{k=-N}^N \binom{2N}{k+N} z^{-2k},$$

allowing for the identification of $a_N[0] = \frac{1}{2^{2N}} \binom{2N}{N}$. The application of Theorem 2 then yields the normalized version of Simoncelli's original reproduction formula,

$$\frac{(2^N N!)^2}{(2N)!(N+1)} \cdot \sum_{n=0}^N \cos^{2N} \left(\theta + \frac{\pi n}{N+1}\right) = 1,$$

which is given here for reference. This particular choice of polar profile corresponds to a steerable implementation of the directional Hilbert transform operator $\mathcal{H}_{(\cos \theta, \sin \theta)}^N$, which clearly falls within the framework of the generalized Riesz transform. The corresponding frequency responses for $N+1 = 3$ and $N+1 = 4$ are given in Fig. 3a and 3b, respectively, and should be compared to the second and third row in Fig. 2. The filterbanks are identical for $N = 1$, but not otherwise. Yet, they span the same spaces and define a self-invertible transform in either formulation.

We would like to emphasize that the results in this section are also directly relevant for the design of general 2-D polar-separable steerable filters whose polar frequency response is of the form $\hat{h}(\omega, \theta) = \hat{w}(\omega)H(e^{j\theta})$ where $\hat{w}(\omega)$ is a given radial frequency window. This factorization allows for a complete decoupling of the radial and angular parts of the design. Since the self-reversibility of the angular part is automatically guaranteed (thanks to Theorems 1 and 2), the designer can concentrate on the problem of specifying a radial (or quasi-isotropic) multi-resolution decomposition that ensures a perfect (or near perfect) reconstruction, which is essentially the approach that was taken by Simoncelli et al. in [11]. An important point, though, is that the present framework gives access to a larger variety of angular profiles and steerable basis functions than what has been considered so far.

V. HIGH-ORDER STEERABLE RIESZ-LAPLACE WAVELETS

While the construction method outlined in §III is applicable to any primary wavelet transform, it will give the most satisfactory results when the underlying primary wavelets are nearly isotropic. We will therefore illustrate the concept using a primary Laplacian-like (or Mexican hat) decomposition of $L_2(\mathbb{R}^d)$ [43].




Fig. 8. Reversible wavelet decompositions of Cameraman image (zoom on a 64×64 subimage). (a) Primary Laplace wavelet pyramid with $\gamma = 5$. (b) 3rd order Riesz-Laplace wavelet transform. (c) Rotated filterbank spanning the same steerable wavelet subspaces as (b).




Fig. 9. Frequency responses (real or imaginary part) of the wavelet analysis filters associated with Fig. 8. (a) Primary Laplace wavelets $\widehat{\psi}_{\text{iso},i}(\omega)$ with $\gamma = 5$. (b) 3rd-order Riesz-Laplace wavelets $\widehat{\psi}_i^{(n_1, n_2)}(\omega)$ with $n_1 + n_2 = 3$. (c) Corresponding “à la Simoncelli” directional filterbank.




Fig. 6. Reversible wavelet decompositions of Cameraman image (zoom on a 64×64 subimage). (a)-(b): Primary Laplace wavelet pyramid with $\gamma = 3$ and its corresponding 1st order Riesz-Laplace (or Marr) wavelet transform. (c)-(d): Primary Laplace wavelet pyramid with $\gamma = 4$ and its corresponding 2nd order Riesz-Laplace wavelet transform.

A. Primary Laplacian-like 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 [44]. In prior work, we have introduced a pyramid-like representation of the underlying wavelet subspaces that involves a single quasi-isotropic analysis wavelet (cf [43, Section IV.B]):

$$\psi_{\text{iso}}(\mathbf{x}) = (-\Delta)^{\gamma/2} \beta_{2\gamma}(2\mathbf{x}) \quad (28)$$

where $\gamma > \frac{d}{2}$ is the order (possibly fractional) of the wavelet. The function $\beta_{2\gamma}(\mathbf{x})$ is the “most-isotropic” polyharmonic B-spline of order 2γ ; it is a smoothing kernel that converges to a Gaussian as γ increases [44, Proposition 2].

The corresponding “Laplacian” analysis wavelets in (22) are given by $\psi_{i,\mathbf{k}}(\mathbf{x}) = 2^{-id/2} \psi_{\text{iso}}(\mathbf{x}/2^i - \mathbf{k})$ and the wavelet transform has a fast reversible filterbank implementation, as described in [43]. The resulting wavelet decomposition has a pyramid-like structure that is similar to the Laplacian pyramid of Burt and Adelson [45]. The fundamental difference, however, is that the proposed transform involves a single analysis wavelet (28) that is specifically tied to the operator $(-\Delta)^{\gamma/2}$. Two examples of wavelet decompositions with $\gamma = 3, 4$ of Cameraman (cf. Fig. 5) are displayed in Fig. 6a and 6c; we have zoomed on a fragment of the image to facilitate the visualization. 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. 7a and 7c. Note the bandpass behavior of these filters




Fig. 7. Frequency responses (real or imaginary part) of the wavelet analysis filters associated with Fig. 6. (a)-(b) Primary Laplace wavelets $\hat{\psi}_{\text{iso},i}(\boldsymbol{\omega})$ with $\gamma = 3$ and 1st-order Riesz-Laplace wavelets $\hat{\psi}_i^{(1,0)}(\boldsymbol{\omega})$ and $\hat{\psi}_i^{(0,1)}(\boldsymbol{\omega})$. (c)-(d) Primary Laplace wavelets $\hat{\psi}_{\text{iso},i}(\boldsymbol{\omega})$ with $\gamma = 4$ and 2nd-order Riesz-Laplace wavelets $\hat{\psi}_i^{(2,0)}(\boldsymbol{\omega})$, $\hat{\psi}_i^{(1,1)}(\boldsymbol{\omega})$, and $\hat{\psi}_i^{(0,2)}(\boldsymbol{\omega})$. The scale i is increasing from $i = 0$ upto 2 as one moves down.

and their good isotropy properties, especially over the lower frequency range.

B. Steerable Riesz-Laplace wavelets

The associated N th-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

$$\psi_{\gamma}^{(n_1, \dots, n_d)}(\mathbf{x}) = \sqrt{\frac{N!}{n_1! \dots n_d!}} \mathcal{R}_1^{n_1} \dots \mathcal{R}_d^{n_d} (-\Delta)^{\frac{\gamma}{2}} \beta_{2\gamma}(2\mathbf{x}) \quad (29)$$

for all multi-indices $\mathbf{n} = (n_1, \dots, n_d)$ with $|\mathbf{n}| = N$. 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. The order parameter $\gamma > \frac{d}{2}$ is a degree of freedom of the transform that can be tuned to the spectral characteristics of the signal; this is especially handy if we are dealing with

fractal processes because of the whitening properties of the fractional Laplacian [46]. In the sequel, we will set it to $\gamma = N + 2m$ with integer $m \geq 0$ to establish a link with conventional (non-fractional) derivative operators.

The proposed wavelets appear to be well suited for feature extraction and analysis because they provide multiscale estimates of the N th order derivatives of the image. Indeed, by using Eq. (19) that links the Riesz transform and partial derivatives, we get

$$\psi_{N+2m}^{(n_1, \dots, n_d)}(\mathbf{x}) = (-1)^N \sqrt{\frac{N!}{n_1! \dots n_d!}} \frac{\partial^N \xi(\mathbf{x})}{\partial^{n_1} x_1 \dots \partial^{n_d} x_d} \quad (30)$$

where the smoothing kernel $\xi(\mathbf{x}) = -\Delta^m \beta_{2N+4m}(2\mathbf{x})$ is isotropic-bandpass (m th Laplacian of a Gaussian-like function). Note that for $N > 1$, we can select $m = 0$ to impose a Gaussian-like smoothing kernel instead of a Mexican hat. If we now consider the analysis of the signal $f(\mathbf{x})$ at scale i with the rescaled wavelet $\psi_{N+2m,i}^{(n_1, \dots, n_d)}(\mathbf{x}) = 2^{-id/2} \psi_{N+2m}^{(n_1, \dots, n_d)}(\mathbf{x}/2^i)$ at location \mathbf{x}_0 , we obtain that

$$\langle f, \psi_{N+2m,i}^{(n_1, \dots, n_d)}(\cdot - \mathbf{x}_0) \rangle_{L_2} \propto \left. \frac{\partial^N (f * \xi_i)(\mathbf{x})}{\partial^{n_1} x_1 \dots \partial^{n_d} x_d} \right|_{\mathbf{x}=\mathbf{x}_0} \quad (31)$$

where we have made use of the adjoint and scale-covariance properties of the partial derivative operators. The interpretation of this relation is that the Riesz wavelet coefficients $w_i^{(n_1, \dots, n_d)}[\mathbf{k}] = \langle f, \mathcal{R}^{(n_1, \dots, n_d)} \psi_{i,\mathbf{k}} \rangle$ correspond to the partial derivatives of order N (sampled at location $\mathbf{x}_0 = 2^i \mathbf{k}$) of the bandpass-filtered version ($f * \xi_i$) of the signal with $\xi_i(\mathbf{x}) = 2^{-id/2} \xi(\mathbf{x}/2^i)$. The non-trivial aspect of these partial-derivative wavelets is that the corresponding signal decomposition is fully reversible (frame property) and that the transform admits an efficient perfect-reconstruction filterbank implementation.

Interestingly, the simplest case $(d, N, m) = (2, 1, 1)$ yields two gradient-like wavelets that happen to be mathematically equivalent to the Marr wavelets described in [47]. An example of such a wavelet decomposition is shown in Fig. 6(b). In our earlier work, we have argued that such a transform could be interpreted as a multiscale edge detector and that it could yield a compact description of an image in terms of its primal wavelet sketch [47]. The interest of the present formulation is that it gives the 3-D generalization of this complex representation without resorting to any sophisticated algebra (e.g., quaternions).

The steerable Riesz-Laplace transforms for $N > 1$ are all new. For instance, for $(d, N, m) = (2, 2, 1)$, we obtain three parallel wavelet channels that provide a multiscale Hessian analysis. The corresponding three-channel directional wavelet transform is shown in Fig. 6(b). The frequency responses of the 1st and 2nd order wavelet filters are displayed in Fig. 7. It is also interesting to consider the 3-D versions of these wavelets (i.e., $(d, N, m) = (3, n, 1)$ with $n = 1, 2$), which are represented in Fig. 10. The first-order wavelet may serve as a contour/surface detector, while the second order wavelet is more suitable for detecting lines or ridges.

The scheme can be pushed further to obtain multiscale image representations with higher orders of steerability. Two examples of 3rd order decompositions with $\gamma = 5$ are shown

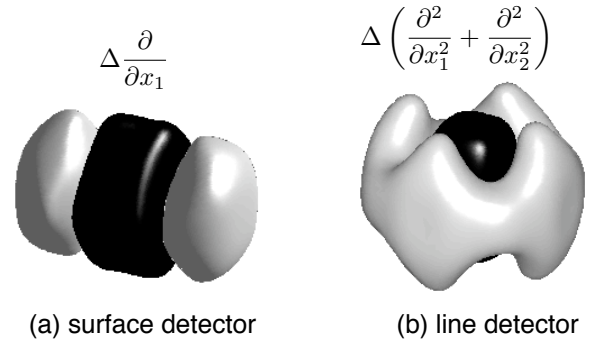


Fig. 10. Isosurface representation of the first and second-order 3-D Riesz-Laplace wavelets $\psi_{n+2}^{(n,0,0)}(\mathbf{x})$ (white: positive threshold, black : negative threshold). (a) : surface or contour detector with $n = 1$. (b) : line or ridge detector with $n = 2$.

in Fig. 8. The first involves the Riesz-Laplace wavelets which produces a local description of image features in terms of third-order derivatives [cf. (31)]. The second corresponds to an equivalent directional filterbank with a basic detector (first Riesz wavelet) that is rotated in the four principal directions $\theta_0 = \pi/2, \pi/4, 0, -\pi/4$ using the steering equations (13) and (14) with $\mathbf{u} = (\cos \theta_0, \sin \theta_0)$. This decomposition represents the “à la Simoncelli” counterpart of the basic Riesz-Laplace transform; it is obtained through the application of Theorem 2 with $H(e^{j\theta}) = (\cos \theta)^N$ with $N = 3$. These filters also have a differential interpretation: they all correspond to third-order directional derivatives, in accordance with (20) and (25). The latter decomposition obviously achieves a better decomposition of the image in terms of directional components, while the former has a greater pattern diversity since it involves two distinct templates $(3, 0)$ and $(2, 1)$ which can, in principle, also be steered in arbitrary directions.

The comparison of the frequency responses of these two systems is given in Fig 9. The directional character of the “rotated-filterbank” decomposition is quite apparent in Fig 9(c). The steering is not perfect in the higher portion of the spectrum due to the “square-shaped” outer frequency profile of the primary wavelet decomposition, which is only approximately isotropic [cf. Fig 9(a)]. While this constitutes a limitation of the current Laplacian-based design, the distorting effect is far less noticeable in the space domain because the underlying polyharmonic B-spline smoothing kernel $\beta_{2\gamma}$ in (29) is guaranteed to converge to an isotropic Gaussian as γ increases.

VI. DISCUSSION

The proposed operator-based framework is very general: it can yield a large variety of multidimensional wavelets with arbitrary orders of steerability. The goal that we are pursuing here is quite similar to that achieved by Simoncelli with his steerable pyramid [11]. However, there are two fundamental differences. First, Simoncelli’s approach is primarily based on digital 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 $N > 1$, the prototypical Riesz wavelets are no longer rotated version of a single template, as is the case with traditional steerable filters. This is quite apparent in Figures 8 and 9. The “à la Simoncelli” counterpart of the 3rd-order Riesz-Laplace transform involves four rotated versions of the first filter uniformly spread over the circle (e.g., at $0^\circ, 45^\circ, 90^\circ$ and 135°). In fact, the Riesz-Laplace wavelets that we introduced in Section V to illustrate our concept are much closer to derivatives than to rotated filterbanks; we believe that this could be an advantage for analysis purposes.

Our approach is obviously not limited to polyharmonic splines. It works best—i.e., provides exact steerability—when the primary wavelets are isotropic; it also automatically yields a tight frame whenever the primary decomposition is tight to start with. Other relevant considerations—including topics for future research—are as follows:

1) *Decoupling between multiresolution and multiorientation properties*: The proposed framework decouples the directional and multi-scale aspects of wavelet design. While this is valuable conceptually, it also suggests some new design challenges; i.e., the specification of wavelet frames with improved isotropy properties. The Riesz-Laplace wavelets that we have presented are very nearly isotropic, but they are not tight. Alternatively, we could have considered a primary orthogonal wavelet transformation, such a quincunx decomposition with good isotropy properties [48], which would have automatically yielded a tight Riesz frame. The practical difficulty is to optimally integrate the requirements of self-reversibility and isotropy (or radial symmetry) into a given computational structure; this is a research subject in its own right which we plan to investigate in future work.

2) *Steerable wavelets diversity*: Our extended subspace-based definition of steerability opens up the perspective of a much larger variety of steerable wavelets than what has been used so far. We are proposing two possible design strategies, both of which preserve self-reversibility. The first is to consider a generalized Riesz transform of the form $\mathbf{U}\mathcal{R}^{(N)}$ where \mathbf{U} is any user-specified unitary matrix; for a justification, see (15). An attractive possibility is to derive \mathbf{U} using some kind of principal component analysis. The second option, which is limited to 2-D, is to take advantage of Theorem 2 to specify a rotating filterbank profile that is optimized for a specific task. Since there is no constraint on the selection of the coefficients of the filter profile (27), one can optimize these for better angular selectivity or tune them to the characteristics of a given type of pattern. For instance, we can optimize the 2nd-order (or Hessian-like) wavelets for the task of ridge detection using a Canny-like criterion [5]. The fact that such a weight adjustment can improve the central elongation of the wavelet is illustrated in Fig. 11; the first detector corresponds to the conventional choice $(\cos\theta)^2$, while the second uses the optimal weights given in [5, Table 2]. The remarkable feature is that the two proposed design methodologies are essentially constraint-free. The extent to which this increased flexibility can improve performance is an aspect that remains




Fig. 11. Two examples of steerable 2nd-order wavelets. (a): Hessian-like wavelet $\psi_4^{(2,0)}$. (b) improved ridge detector: $\psi_4^{(2,0)} - \frac{1}{3\sqrt{2}}\psi_4^{(1,1)}$.

to be investigated.

3) *Simplicity of implementation*: The Riesz components/wavelet coefficients are straightforward to compute in the Fourier domain using a FFT-based version of the fast wavelet transform algorithm. They can then be recombined linearly as desired to generate a large variety of steerable wavelet transforms, including rotated filterbanks. This computational approach simplifies the implementation and frees us from the requirement of using finitely-supported filters. For the present work, we did implement the 2-D Riesz-Laplace transforms for orders up to 11, even though we did only show results for $N = 1, 2, 3$ (Figs. 6b-c, 8b). The “à la Simoncelli” version of these transforms (Fig. 8(c)) were obtained by performing the adequate linear combinations of the Riesz wavelet coefficients; we found this to be a convenient, cost-effective solution.

4) *Exact reconstruction property*: The exact reconstruction property comes as a direct consequence of Proposition 3 which ensures that the generalized Riesz transform preserves the frame property. In other words, we have a general mechanism for mapping any reversible multi-resolution decomposition into its directional counterpart which is automatically guaranteed to be reversible as well. This is not necessarily so in the case of Simoncelli’s steerable pyramid. Indeed the structure of the underlying filterbank is such that it will only yield perfect reconstruction when the aliasing is completely suppressed; that is, when the radial lowpass/highpass filters are bandlimited. The implication is that the reconstruction with finite length filters will only be approximate. There is no such constraint in the present scheme because the computational structure of the primary “isotropic” multiresolution decomposition is essentially irrelevant, as long as it satisfies the frame property.

5) *Preservation of approximation theoretic properties*: Because the Riesz transform and its higher order extension are well-defined convolution operators in $L_2(\mathbb{R}^d)$, they will preserve all important approximation theoretic properties of the primary wavelet transform that is used for the construction. These include:

- the order of approximation γ , which means that the approximation error of a scale-truncated expansion decays like the γ th power of that scale.
- the initial number of vanishing moments; this implies that the Riesz wavelet coefficients are essentially zero in smooth image areas where the image is well represented by its lower-order Taylor series.
- the Sobolev degree of smoothness of the basis functions.

Another important aspect is that, even though the Riesz transform is a non-local operator, it will minimally impact the decay of the wavelet functions. This is because its Fourier-domain singularity at the origin is tempered by the zeros (vanishing moments) of the primary wavelets.

6) *Generalization to higher dimensions:* 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). By contrast, the extension of Simoncelli's steerable pyramid, which uses rotated filters that are angularly uniformly distributed, to 3-D does not seem feasible because of the lack of a systematic mechanism for uniformly tiling the sphere.

VII. CONCLUSION

We introduced a general operator framework that provides a mathematical bridge between the continuous-domain theory of the wavelet transform and steerable filterbanks. In effect, we are taking advantage of the invariance properties (with respect to translation, scaling and rotation) of the generalized Riesz transform, together with its self-reversibility, to decouple the directional and multi-scale aspects of wavelet design. The primary theoretical contribution is a rigorous and straightforward formalism for the specification of steerable wavelet transforms in the continuous-domain—i.e., as frames of $L_2(\mathbb{R}^d)$ —while ensuring that they are endowed with all important approximation-theoretic wavelet properties: order of approximation, vanishing moments, regularity, etc.

The present work is also practically relevant because of the new perspective that it offers for image processing. Interesting topics for future research include the design of novel families of steerable wavelet transforms (in particular, tight frames), the development of steerable wavelet algorithms for the processing of volumetric data, and the investigation of the concept of steerable wavelet diversity—in particular, signal-adapted design—which holds good promises for data processing (e.g., denoising) and feature extraction.

ACKNOWLEDGEMENTS

This work was supported by the Swiss National Science Foundation (MU) under Grant 200020-109415 and by the Center for Biomedical Imaging (DVIDV) of the Geneva-Lausanne Universities and the EPFL, as well as the foundations Leenaards and Louis-Jeantet.

REFERENCES

- [1] S. G. Mallat, "A theory of multiresolution signal decomposition: The wavelet representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 7, pp. 674–693, 1989.
- [2] 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.
- [3] E. Simoncelli and H. Farid, "Steerable wedge filters for local orientation analysis," *IEEE Transactions on Image Processing*, vol. 5, no. 9, pp. 1377–1382, Sept. 1996.
- [4] P. Perona, "Deformable kernels for early vision," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 17, no. 5, pp. 488–499, May 1995.
- [5] M. Jacob and M. Unser, "Design of steerable filters for feature detection using Canny-like criteria," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 8, pp. 1007–1019, August 2004.
- [6] M. Michaelis and G. Sommer, "A Lie group approach to steerable filters," *Pattern Recognition Letters*, vol. 16, no. 11, pp. 1165–1174, 1995.
- [7] Y. Hel-Or and P. C. Teo, "Canonical decomposition of steerable functions," *Journal of Mathematical Imaging and Vision*, vol. 9, no. 1, pp. 83–95, 1998.
- [8] P. C. Teo and Y. Hel-Or, "Lie generators for computing steerable functions," *Pattern Recognition Letters*, vol. 19, no. 1, pp. 7–17, 1998.
- [9] E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, "Shiftable multiscale transforms," *IEEE Transaction on Information Theory*, vol. 38, no. 2, pp. 587–607, 1992.
- [10] A. Karasiris and E. Simoncelli, "A filter design technique for steerable pyramid image transforms," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing ICASSP-96*, vol. 4, 7–10 May 1996, pp. 2387–2390.
- [11] E. Simoncelli and W. Freeman, "The steerable pyramid: a flexible architecture for multi-scale derivative computation," in *Proc. International Conference on Image Processing*, vol. 3, 23–26 Oct. 1995, pp. 444–447.
- [12] Y. Wiaux, L. Jacques, and P. Vandergheynst, "Correspondence principle between spherical and Euclidean wavelets," *Astrophys. J.*, vol. 632, p. 15, 2005.
- [13] J. G. Daugman, "Complete discrete 2-D Gabor transforms by neural networks for image-analysis and compression," *IEEE Transactions on Acoustics Speech and Signal Processing*, vol. 36, no. 7, pp. 1169–1179, 1988.
- [14] E. J. Candès and D. L. Donoho, "Ridgelets: A key to higher-dimensional intermittency?" *Phil. Trans. R. Soc. Lond. A.*, pp. 2495–2509, 1999.
- [15] N. Kingsbury, "Image processing with complex wavelets," *Philos. Trans. Royal Soc. London: Series A—Math. Phys. Eng. Sci.*, vol. 357, no. 1760, pp. 2543–2560, 1999.
- [16] I. W. Selesnick, R. G. Baraniuk, and N. C. Kingsbury, "The dual-tree complex wavelet transform," *IEEE Sig. Proc. Magazine*, vol. 22, no. 6, pp. 123–151, 2005.
- [17] P. Vandergheynst and J.-F. Gobbers, "Directional dyadic wavelet transforms: Design and algorithms," *IEEE Transactions on Image Processing*, vol. 11, no. 4, pp. 363–372, Apr. 2002.
- [18] M. N. Do and M. Vetterli, "The contourlet transform: An efficient directional multiresolution image representation," *IEEE Transactions on Image Processing*, vol. 14, no. 12, pp. 2091–2106, 2005.
- [19] E. Le Pennec and S. Mallat, "Sparse geometric image representations with bandelets," *IEEE Transactions on Image Processing*, vol. 14, no. 4, pp. 423–438, 2005.
- [20] V. Velisavljevic, B. Beferull-Lozano, M. Vetterli, and P. Dragotti, "Directionlets: Anisotropic multidirectional representation with separable filtering," *IEEE Transactions on Image Processing*, vol. 15, no. 7, pp. 1916–1933, 2006.
- [21] H. Greenspan, S. Belongie, R. Goodman, P. Perona, S. Rakshit, and C. Anderson, "Overcomplete steerable pyramid filters and rotation invariance," in *Proc. CVPR '94. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 21–23 June 1994, pp. 222–228.
- [22] M. Do and M. Vetterli, "Rotation invariant texture characterization and retrieval using steerable wavelet-domain hidden markov models," *IEEE Transactions on Multimedia*, vol. 4, no. 4, pp. 517–527, 2002.
- [23] G. Tzagarakis, B. Beferull-Lozano, and P. Tsakalides, "Rotation-invariant texture retrieval with gaussianized steerable pyramids," *IEEE Transactions on Image Processing*, vol. 15, no. 9, pp. 2702–2718, 2006.
- [24] F. Denis and A. Baskurt, "Multidirectional curvilinear structures detection using steerable pyramid," *Journal of Electronic Imaging*, vol. 13, no. 4, pp. 756–765, 2004.
- [25] A. Jepsen, D. Fleet, and T. El-Maraghi, "Robust online appearance models for visual tracking," in *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition CVPR 2001*, vol. 1, 2001, pp. 1–415–1–422 vol.1.
- [26] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, "Image denoising using scale mixtures of Gaussians in the wavelet domain," *IEEE Transactions on Image Processing*, vol. 12, no. 11, pp. 1338–1351, November 2003.
- [27] A. A. Bharath and J. Ng, "A steerable complex wavelet construction and its application to image denoising," *IEEE Transactions on Image Processing*, vol. 14, no. 7, pp. 948–959, July 2005.
- [28] M. Riesz, "Sur les fonctions conjuguées," *Math. Z.*, vol. 27, pp. 218–244, 1927.

- [29] E. Stein and G. Weiss, *Introduction to Fourier Analysis on Euclidean Spaces*. Princeton, NJ: Princeton Univ. Press, 1971.
- [30] M. Felsberg and G. Sommer, "The monogenic signal," *IEEE Trans. Signal Processing*, vol. 49, no. 12, pp. 3136–3144, 2001.
- [31] —, "The monogenic scale-space: A unifying approach to phase-based image processing in scale-space," *J. Math. Imaging and Vision*, vol. 21, no. 1, pp. 5–26, 2004.
- [32] G. Metikas and S. C. Olhede, "Multiple multidimensional Morse wavelets," *IEEE Trans. Signal Processing*, vol. 55, no. 3, pp. 921–936, 2007.
- [33] K. G. Larkin, "Natural demodulation of two-dimensional fringe patterns. II. Stationary phase analysis of the spiral phase quadrature transform," *Journal of the Optical Society of America A*, vol. 18, no. 8, pp. 1871–1881, 2001.
- [34] K. G. Larkin, D. J. Bone, and M. A. Oldfield, "Natural demodulation of two-dimensional fringe patterns. I. General background of the spiral phase quadrature transform," *Journal of the Optical Society of America A*, vol. 18, no. 8, pp. 1862–1870, 2001.
- [35] M. Unser, D. Sage, and D. Van De Ville, "Multiresolution monogenic signal analysis using the Riesz-Laplace wavelet transform," *IEEE Trans. Image Processing*, in press.
- [36] S. Held, M. Stroh, P. Massopust, and B. Forster, "Steerable wavelet frames based on the Riesz transform," *IEEE Transactions on Image Processing*, under review.
- [37] I. Gelfand and G. Shilov, *Generalized functions. Vol. 1*. New York, USA: Academic press, 1964.
- [38] S. L. Hahn, *Hilbert Transforms in Signal Processing*. Artech House, Boston, 1996.
- [39] R. A. Adams, *Sobolev Spaces*. New York: Academic Press, 1975.
- [40] R. J. Duffin and A. C. Schaeffer, "A class of nonharmonic Fourier series," *Transactions of the American Mathematical Society*, vol. 72, no. MAR, pp. 341–366, 1952.
- [41] A. Aldroubi, "Portraits of frames," *Proceedings of the American Mathematical Society*, vol. 123, no. 6, pp. 1661–1668, 1995.
- [42] O. Christensen, *An Introduction to Frames and Riesz Bases*. Birkhauser, 2003.
- [43] 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.
- [44] 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.
- [45] P. Burt and E. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Transactions on Communications*, vol. 31, no. 4, pp. 532–540, Apr 1983.
- [46] P. Tafti, D. Van De Ville, and M. Unser, "Invariances, Laplacian-like wavelet bases, and the whitening of fractal processes," *IEEE Transactions on Image Processing*, vol. 18, no. 4, pp. 689–702, April 2009.
- [47] 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.
- [48] M. Feilner, D. Van De Ville, and M. Unser, "An orthogonal family of quincunx wavelets with continuously adjustable order," *IEEE Transactions on Image Processing*, vol. 14, no. 4, pp. 499–510, April 2005.