

Steerable Pyramids and Tight Wavelet Frames in $L_2(\mathbb{R}^d)$

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

Abstract—We present a functional framework for the design of tight steerable wavelet frames in any number of dimensions. The 2-D version of the method can be viewed as a generalization of Simoncelli’s steerable pyramid that gives access to a larger palette of steerable wavelets via a suitable parametrization. The backbone of our construction is a primal isotropic wavelet frame that provides the multiresolution decomposition of the signal. The steerable wavelets are obtained by applying a one-to-many mapping (N th-order generalized Riesz transform) to the primal ones. The shaping of the steerable wavelets is controlled by an $M \times M$ unitary matrix (where M is the number of wavelet channels) that can be selected arbitrarily; this allows for a much wider range of solutions than the traditional equiangular configuration (steerable pyramid). We give a complete functional description of these generalized wavelet transforms and derive their steering equations. We describe some concrete examples of transforms, including some built around a Mallat-type multiresolution analysis of $L_2(\mathbb{R}^d)$, and provide a fast Fourier transform-based decomposition algorithm. We also propose a principal-component-based method for signal-adapted wavelet design. Finally, we present some illustrative examples together with a comparison of the denoising performance of various brands of steerable transforms. The results are in favor of an optimized wavelet design (equalized principal component analysis), which consistently performs best.

Index Terms—Directional derivatives, multiresolution decomposition, Riesz transform, steerable filters, steerable pyramid, tight frames, wavelet transform.

I. INTRODUCTION

THE steerable pyramid is a multiorientation, multiscale image decomposition that was developed in the 1990s by Simoncelli and others [1]–[4]. It is a wavelet-like representation, whose analysis functions are dilated and rotated versions of a single directional wavelet. *Steerability* refers to the property that the underlying wavelets can be rotated to any orientation by forming suitable linear combinations of a primary set of equiangular directional wavelet components [5]. This provides a powerful mechanism for adapting the transform to the local characteristics of the image by steering

the basis functions in the direction of maximal response. The concept is very valuable for image analysis; in particular, local orientation analysis, contour detection, shape from shading [5], texture retrieval [6], [7], and directional pattern detection [8]. The steerable pyramid is also self-reversible, which translates into the fact that the corresponding wavelets form a *tight frame* of $L_2(\mathbb{R}^2)$ [9]. The combination of these two properties—steerability and self-reversibility—is particularly advantageous for designing wavelet-based algorithms for image denoising [10], image enhancement [11], [12], texture synthesis [13], image extrapolation or “hallucination” [14], image fusion [15], as well as sparsity-constrained image reconstruction [16]–[18], which is presently a very active area of research.

In prior work, we introduced a general operator-based framework for the construction of steerable wavelet frames. The foundation of the approach is a higher order version of the Riesz transform that has the remarkable property of mapping a primal isotropic wavelet frame of $L_2(\mathbb{R}^d)$ into a directional wavelet, whose basis functions are steerable [19]. We did apply the concept to the design of the Riesz–Laplace wavelets, which closely resemble the partial derivatives of a Gaussian. The price to pay for the good localization of these wavelets, however, is that the frame is not tight—although the transform is perfectly reversible.

In this paper, we take the next step focusing on the construction of *tight* steerable wavelet frames. Our motivation is to specify extended families of steerable transforms with a larger diversity of wavelets than what has been considered so far, and to propose concrete design solutions in dimensions greater or equal to two.

Our functional framework is such that it decouples the multiresolution and steerability aspects of the transform, which suggests a two-step approach to the problem. The desired starting point for the construction is a tight primal wavelet frame with the best possible localization and isotropy properties. The design options that are presently available to meet these primary requirements are as follows.

- 1) *Nonseparable quincunx wavelet transforms*: These are the most isotropic wavelet bases among the various families available [20], [21]. The advantage is that there is a single wavelet per scale that can also be specified to be orthogonal. The limitation is that the approach does not generalize well to dimension greater than 2 [22].
- 2) *Laplacian-like pyramid decompositions in the spirit of Burt and Adelson* [23]: Such pyramids can be designed to have good isotropy and energy compaction properties [24], [25]; they can also be specified using orthogonal scaling functions, which automatically yields a tight frame [26].

Manuscript received June 18, 2010; revised February 23, 2011; accepted March 06, 2011. Date of publication April 07, 2011; date of current version September 16, 2011. This work was supported by the Swiss National Science Foundation under Grant 200020-109415 (MU) and Grant PP00P2-123438 (DVD) and by the Center for Biomedical Imaging of the Geneva-Lausanne Universities and the Ecole Polytechnique Fédérale of Lausanne, as well as the foundations Leenaards and Louis-Jeantet. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Hsueh-Ming Hang.

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

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

Digital Object Identifier 10.1109/TIP.2011.2138147

The major difficulty there is that imposing orthogonality is often detrimental to isotropy, especially at higher orders.

- 3) *Bandlimited wavelet frames*: Several authors have found that perfect wavelet isotropy is achievable by imposing a radial bandlimiting constraint [13], [27]–[29]. Interestingly, this is also the solution that leads to an equivalence between the present Riesz-based formulation and Simoncelli's steerable pyramid. The fundamental limitation there is that the basis functions cannot be compactly supported. In fact, Romero *et al.* have proven that it is impossible to construct isotropic refinable functions that are compactly supported [28].

While each of these options imposes some level of compromise, we have chosen to concentrate on the two latter ones, which are better matched to our objectives.

The second aspect of the construction of steerable wavelet frames is the specification of the directional components of the transform. The key observation is that our canonical wavelets, which are given by the N th-order Riesz transform of the primal isotropic one, actually define a whole subspace of steerable functions. The idea is to exploit these degrees of freedom to develop alternatives to Simoncelli's classical equiangular solution. In particular, we will apply principal component analysis (PCA) to optimize the basis functions for image denoising and feature extraction.

The paper is organized as follows. We start with a short introduction to the multiinteger notation and a review of the key properties of the higher order Riesz transform. We then proceed in Section III with the definition of a generalized Riesz transform and the explicit derivation of the corresponding steering relations. The main point is that we can use these generalized transforms to specify a whole family of steerable wavelet frames that are associated to some primal wavelet decomposition. The bottom line is that the sought-after properties of the transform—exactness of the steerability and self-reversibility—are in direct relation with the degree of isotropy of the primal wavelets and the tightness of the initial frame bounds. In Section IV, we introduce a common multiresolution framework for the specification of the primal (orientation-free) wavelet decomposition in accordance with the strategies 2) and 3) outlined earlier. We provide explicit Fourier-domain formulas for isotropic bandlimited wavelets as well as a novel functional characterization of pyramid-based decompositions. The implementation details and quantitative assessment of various brands of steerable wavelet transforms (equiangular, Riesz, and different types of PCA wavelets) are given in Section V together with illustrative examples of image processing.

II. MATHEMATICAL PRELIMINARIES

A. Multiindex Notation: Factorials and Derivatives

We will consider d -dimensional multiindex vectors of the form $\mathbf{n} = (n_1, \dots, n_d)$, whose entries n_i are nonnegative integers. We then define the following multiindex operations and operators:

- 1) Sum of components: $|\mathbf{n}| = \sum_{i=1}^d n_i = N$.
- 2) Factorial: $\mathbf{n}! = n_1!n_2!\cdots n_d!$.

- 3) Exponentiation of a vector $\mathbf{z} = (z_1, \dots, z_d) \in \mathbb{C}^d$: $\mathbf{z}^{\mathbf{n}} = z_1^{n_1} \cdots z_d^{n_d}$.
- 4) Higher order partial derivative of a function $f(\mathbf{x})$, $\mathbf{x} = (x_1, \dots, x_d) \in \mathbb{R}^d$

$$\partial^{\mathbf{n}} f(\mathbf{x}) = \frac{\partial^{|\mathbf{n}|} f(\mathbf{x})}{\partial x_1^{n_1} \cdots \partial x_d^{n_d}}$$

The notation allows for a concise description of the multinomial theorem

$$\left(\sum_{i=1}^d x_i \right)^N = \sum_{|\mathbf{n}|=N} \frac{N!}{\mathbf{n}!} \mathbf{x}^{\mathbf{n}}$$

which involves a summation over $\binom{N+d-1}{d-1}$ distinct monomials of the form $\mathbf{x}^{\mathbf{n}} = x_1^{n_1} \cdots x_d^{n_d}$ with $n_1 + \cdots + n_d = N$. It also yields a compact formula for the N th-order Taylor series expansion of a multidimensional function $f(\mathbf{x})$ with well-defined derivatives up to order $N + 1$

$$f(\mathbf{x} + \mathbf{h}) = \sum_{0 \leq |\mathbf{n}| \leq N} \frac{\partial^{\mathbf{n}} f(\mathbf{x})}{\mathbf{n}!} \mathbf{h}^{\mathbf{n}} + O(\|\mathbf{h}\|^{N+1}). \quad (1)$$

By setting $h = \|\mathbf{h}\|$ and $\mathbf{u} = \mathbf{h}/\|\mathbf{h}\|$, this is equivalent to the univariate N th-order Taylor series of $f(\mathbf{x})$ along the direction \mathbf{u}

$$f(\mathbf{x} + h\mathbf{u}) = \sum_{n=0}^N \frac{D_{\mathbf{u}}^n f(\mathbf{x})}{n!} h^n + O(h^{N+1})$$

where

$$D_{\mathbf{u}}^n f(\mathbf{x}) = \sum_{|\mathbf{n}|=n} \frac{n!}{\mathbf{n}!} \mathbf{u}^{\mathbf{n}} \partial^{\mathbf{n}} f(\mathbf{x})$$

is the n th-order directional derivative of f along the unit vector \mathbf{u} . The Fourier domain counterpart of the latter formula is as follows:

$$\widehat{D_{\mathbf{u}}^n f}(\boldsymbol{\omega}) = (\langle \mathbf{u}, j\boldsymbol{\omega} \rangle)^n \widehat{f}(\boldsymbol{\omega})$$

where $\widehat{f}(\boldsymbol{\omega}) = \int_{\mathbb{R}^d} f(\mathbf{x}) e^{-j\langle \boldsymbol{\omega}, \mathbf{x} \rangle} dx_1, \dots, dx_d$ is the Fourier transform of the input signal $f(\mathbf{x})$.

B. Riesz Transform and Its Higher Order Extension

The Riesz transform of a finite-energy function $f(\mathbf{x})$, $\mathbf{x} \in \mathbb{R}^d$ is the scalar-to-vector signal transformation $\mathcal{R}: L_2(\mathbb{R}^d) \rightarrow L_2^d(\mathbb{R}^d)$

$$\mathcal{R}f(\mathbf{x}) = \begin{pmatrix} \mathcal{R}_1 f(\mathbf{x}) \\ \vdots \\ \mathcal{R}_d f(\mathbf{x}) \end{pmatrix} \quad (2)$$

where the Riesz component operators \mathcal{R}_i are linear, space-invariant and characterized by the frequency responses $\widehat{\mathcal{R}}_i(\boldsymbol{\omega}) = -j\omega_i/\|\boldsymbol{\omega}\|$. Hence, in the frequency domain, we have that

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

The higher order versions of the transform are obtained by iteration. While there are d^N possible ways of forming terms of form $\mathcal{R}_{i_1}\mathcal{R}_{i_2}\dots\mathcal{R}_{i_N}f$ with $i_1, i_2, \dots, i_N \in \{1, \dots, d\}$, there are actually much fewer distinct N th-order Riesz components due to the commutativity and factorization properties of the underlying convolution operators. Specifically, we define the higher order Riesz component operator with multiindex $\mathbf{n} = (n_1, \dots, n_d)$ as follows:

$$\mathcal{R}^{\mathbf{n}} = \sqrt{\frac{|\mathbf{n}|!}{\mathbf{n}!}} \mathcal{R}_1^{n_1} \mathcal{R}_2^{n_2} \dots \mathcal{R}_d^{n_d} \quad (4)$$

which is normalized such that the global transform preserves energy. These are then used to specify the full N th-order Riesz transform of a signal $f(\mathbf{x})$

$$\mathcal{R}^{(N)} f(\mathbf{x}) = \begin{pmatrix} \mathcal{R}^{(N,0,\dots,0)} f(\mathbf{x}) \\ \vdots \\ \mathcal{R}^{(n_1,\dots,n_d)} f(\mathbf{x}) \\ \vdots \\ \mathcal{R}^{(0,\dots,0,N)} f(\mathbf{x}) \end{pmatrix} \quad (5)$$

which involves $M = \binom{N+d-1}{d-1}$ distinct higher order Riesz components indexed by \mathbf{n} with $n_1 + \dots + n_d = N$. The corresponding Fourier-domain description of the \mathbf{n} -component of the transform is as follows:

$$\widehat{\mathcal{R}^{\mathbf{n}} f}(\boldsymbol{\omega}) = \sqrt{\frac{|\mathbf{n}|!}{\mathbf{n}!}} \frac{(-j\boldsymbol{\omega})^{\mathbf{n}}}{\|\boldsymbol{\omega}\|^{|\mathbf{n}|}} \hat{f}(\boldsymbol{\omega}). \quad (6)$$

A fundamental property is that the higher order Riesz transform preserves inner products in $L_2(\mathbb{R}^d)$, which also implies that it is self-invertible:

$$\forall f \in L_2(\mathbb{R}^d), \quad \left(\mathcal{R}^{(N)}\right)^* \mathcal{R}^{(N)} f = \sum_{|\mathbf{n}|=N} (\mathcal{R}^{\mathbf{n}})^* \mathcal{R}^{\mathbf{n}} f = f \quad (7)$$

where the symbol $*$ denotes the adjoint operator. In particular, we have that $(\mathcal{R}^{\mathbf{n}})^* = (-1)^N \mathcal{R}^{\mathbf{n}}$. Note that the Fourier-domain version of this perfect reconstruction property is a direct consequence of the multinomial theorem

$$\begin{aligned} \hat{f}(\boldsymbol{\omega}) &= \frac{\|\boldsymbol{\omega}\|^{2N}}{\|\boldsymbol{\omega}\|^{2N}} \hat{f}(\boldsymbol{\omega}) \\ &= \frac{1}{\|\boldsymbol{\omega}\|^{2N}} \left(\sum_{i=1}^d \omega_i^2 \right)^N \hat{f}(\boldsymbol{\omega}) \\ &= \frac{1}{\|\boldsymbol{\omega}\|^{2N}} \sum_{|\mathbf{n}|=N} \frac{N!}{\mathbf{n}!} \boldsymbol{\omega}^{\mathbf{n}} \boldsymbol{\omega}^{\mathbf{n}} \hat{f}(\boldsymbol{\omega}) \\ &= \sum_{|\mathbf{n}|=N} \sqrt{\frac{N!}{\mathbf{n}!}} \frac{(j\boldsymbol{\omega})^{\mathbf{n}}}{\|\boldsymbol{\omega}\|^N} \cdot \sqrt{\frac{N!}{\mathbf{n}!}} \frac{(-j\boldsymbol{\omega})^{\mathbf{n}}}{\|\boldsymbol{\omega}\|^N} \hat{f}(\boldsymbol{\omega}) \end{aligned}$$

which, thanks to (6), is equivalent to (7).

We can also establish a connection between the Riesz transform and partial derivatives by multiplying (6) by $\|\boldsymbol{\omega}\|^N$ with $N = n_1 + \dots + n_d$ and explicitly writing the corresponding Fourier-transform pair

$$\begin{aligned} \mathcal{R}^{(n_1,\dots,n_d)} (-\Delta)^{N/2} f(\mathbf{x}) \\ \xleftrightarrow{\mathcal{F}} \sqrt{\frac{N!}{\mathbf{n}_1! \dots \mathbf{n}_d!}} (-1)^N (j\omega_1)^{n_1} \dots (j\omega_d)^{n_d} \hat{f}(\boldsymbol{\omega}) \end{aligned}$$

where $(-\Delta)^\gamma$ is the (fractional) Laplace operator of order γ , whose frequency response is $\|\boldsymbol{\omega}\|^{2\gamma}$. This yields the relation

$$\mathcal{R}^{\mathbf{n}} (-\Delta)^{\frac{|\mathbf{n}|}{2}} f(\mathbf{x}) = (-1)^{|\mathbf{n}|} \sqrt{\frac{|\mathbf{n}|!}{\mathbf{n}!}} \partial^{\mathbf{n}} f(\mathbf{x})$$

which holds in the sense of distributions. Since the inverse of $(-\Delta)^{|\mathbf{n}|/2}$ is an isotropic low-pass-filtering operator, the net effect of the higher order Riesz transform is to extract smoothed versions of the derivatives of order N of the signal of interest. Based on the aforementioned relation, we also provide a Riesz version of the Taylor series (1)

$$\begin{aligned} f(\mathbf{x}) &= \sum_{0 \leq |\mathbf{n}| \leq N} \mathcal{R}^{\mathbf{n}} (-\Delta)^{\frac{|\mathbf{n}|}{2}} f(\mathbf{x}) \Big|_{\mathbf{x}=\mathbf{x}_0} a_{\mathbf{n}} (\mathbf{x} - \mathbf{x}_0)^{\mathbf{n}} \\ &\quad + O(\|\mathbf{x} - \mathbf{x}_0\|^{N+1}) \end{aligned}$$

with the weighting factors $a_{\mathbf{n}} = (-1)^{|\mathbf{n}|} / \sqrt{|\mathbf{n}|! \mathbf{n}!}$.

Similar to derivatives, the Riesz transform commutes with translation and scaling operators. What is more remarkable for image processing is that this feature extraction process commutes with spatial rotations, as discussed in the next section.

III. GENERALIZED RIESZ TRANSFORMS AND STEERABLE WAVELET FRAMES

A. Generalized Riesz Transforms

Definition 1: Let \mathbf{U} be a $M \times M$ (possibly complex-valued) nonsingular matrix with $M = \binom{N+d-1}{d-1}$. The generalized Riesz transform of order N and coefficient matrix \mathbf{U} is the scalar to M -vector signal transformation given by

$$\mathcal{R}_{\mathbf{U}} f(\mathbf{x}) = \mathbf{U} \mathcal{R}^{(N)} f(\mathbf{x}).$$

For notational convenience, we will represent the matrix entries of \mathbf{U} using the multiindex notation $u_{\mathbf{m},\mathbf{n}}$; specifically, $[\mathbf{U}]_{\mathbf{m},\mathbf{n}} = u_{\mathbf{m},\mathbf{n}}$, where \mathbf{m} and \mathbf{n} are the multiindices corresponding the m and n components of the N th-order Riesz transform (5), respectively.

The adjoint transformation maps a M -vector signal $\mathbf{f}(\mathbf{x})$ back into the scalar signal domain; it is given by

$$\mathcal{R}_{\mathbf{U}}^* \mathbf{f}(\mathbf{x}) = \left(\mathcal{R}^{(N)}\right)^* \mathbf{U}^H \mathbf{f}(\mathbf{x}) \quad (8)$$

where \mathbf{U}^H is the hermitian transpose of \mathbf{U} .

While the multiplication of the transform with an arbitrary matrix \mathbf{U} adds an interesting level of generalization, it does not fundamentally affect the invariance properties of the Riesz transform.

Property 1: The generalized Riesz transforms are translation- and scale-invariant

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

The translation and dilation invariance properties of the primary Riesz components are preserved, since the matrix multiplication part of the transform can be simply factored out,

Property 2 (Norm Equivalence and Invertibility): The generalized Riesz transform satisfies the following norm equivalence relation for all $f \in L_2(\mathbb{R}^d)$:

$$\lambda_{\min}\|f\|_{L_2(\mathbb{R}^d)}^2 \leq \|\mathcal{R}_{\mathbf{U}}f\|_{L_2^M(\mathbb{R}^d)}^2 \leq \lambda_{\max}\|f\|_{L_2(\mathbb{R}^d)}^2$$

where λ_{\max} and λ_{\min} are the maximum and minimum eigenvalues of $\mathbf{U}^H\mathbf{U}$, respectively. Its (left) inverse is given by the adjoint mapping $\mathcal{R}_{\mathbf{U}^{-1H}}$ and is well defined iff $\lambda_{\min} > 0$. In particular, when \mathbf{U} is a unitary matrix (i.e., $\mathbf{U}^H\mathbf{U} = \mathbf{I}$), we have exact energy and inner-product preservation, which also implies self-reversibility; i.e., $\mathcal{R}_{\mathbf{U}}\mathcal{R}_{\mathbf{U}} = \text{Identity}$.

Proof: The result follows from the energy preservation [19, Property 7] of the N th-order Riesz transform: $\|\mathcal{R}^{(N)}f\|_{L_2^M(\mathbb{R}^d)}^2 = \|f\|_{L_2(\mathbb{R}^d)}^2$. Indeed, we have that

$$\begin{aligned} \|\mathcal{R}_{\mathbf{U}}f\|_{L_2^M(\mathbb{R}^d)}^2 &= \left\langle \mathcal{R}^{(N)}f, \mathbf{U}^H\mathbf{U}\mathcal{R}^{(N)}f \right\rangle_{L_2^M(\mathbb{R}^d)} \\ &\leq \lambda_{\max} \left\langle \mathcal{R}^{(N)}f, \mathcal{R}^{(N)}f \right\rangle_{L_2^M(\mathbb{R}^d)} \\ &= \lambda_{\max}\|f\|_{L_2(\mathbb{R}^d)}^2 \end{aligned}$$

and we obtain the corresponding lower bound in exactly the same fashion.

The fact that the operator is bounded for above and below implies that is invertible. To identify the inverse, we simply observe that $\mathcal{R}_{\mathbf{U}^{-1H}}\mathcal{R}_{\mathbf{U}} = (\mathcal{R}^{(N)})^*(\mathbf{U}^{-1}\mathbf{U})\mathcal{R}^{(N)} = (\mathcal{R}^{(N)})^*\mathcal{R}^{(N)} = \text{Identity}$, as a consequence of (7) and (8). ■

The generalized N th-order Riesz transform is intrinsically rotation-invariant, as expressed by the following theorem.

Theorem 1 (Steerability): The generalized Riesz transform is steerable in the sense that its component impulse responses can be simultaneously rotated to any spatial orientation by forming suitable linear combinations. Specifically, let $\mathbf{R} = (\mathbf{r}_1 \dots \mathbf{r}_d)^T$ be a $d \times d$ spatial rotation matrix, such that $\mathbf{R}^{-1} = \mathbf{R}^T$. Then, $\mathcal{R}_{\mathbf{U}}\{\delta\}(\mathbf{R}\mathbf{x}) = \mathcal{R}_{\mathbf{U}\mathbf{S}_{\mathbf{R}}}\{\delta\}(\mathbf{x})$, where δ is the Dirac impulse and where $\mathbf{S}_{\mathbf{R}}$ is a $M \times M$ steering matrix, whose entries $s_{\mathbf{n},\mathbf{m}}$ (in multiindex notation) are N th-order polynomials of the entries of the rotation matrix \mathbf{R} given by

$$\begin{aligned} s_{\mathbf{n},\mathbf{m}} &= \sqrt{\frac{\mathbf{m}!}{\mathbf{n}!}} \sum_{|\mathbf{k}_1|=n_1} \dots \sum_{|\mathbf{k}_d|=n_d} \delta_{\mathbf{k}_1+\dots+\mathbf{k}_d,\mathbf{m}} \\ &\quad \times \frac{\mathbf{n}!}{\mathbf{k}_1! \dots \mathbf{k}_d!} \mathbf{r}_1^{\mathbf{k}_1} \dots \mathbf{r}_d^{\mathbf{k}_d} \quad (9) \end{aligned}$$

where we are using the Kronecker symbol to exclude the summation terms with $\mathbf{k}_1 + \dots + \mathbf{k}_d \neq \mathbf{m}$.

Proof: We start by writing the frequency response of the \mathbf{m}' -component filter of the generalized Riesz transform:

$$\hat{U}_{\mathbf{m}'}(\boldsymbol{\omega}) = \sum_{|\mathbf{n}|=N} u_{\mathbf{m}',\mathbf{n}} \sqrt{\frac{N}{\mathbf{n}!}} \frac{(-j\boldsymbol{\omega})^{\mathbf{n}}}{\|\boldsymbol{\omega}\|^N}. \quad (10)$$

Due to the rotation property of the Fourier transform, spatially rotating this filter by \mathbf{R} corresponds to the same rotation in the frequency domain

$$\hat{U}_{\mathbf{m}'}(\mathbf{R}\boldsymbol{\omega}) = \sum_{|\mathbf{n}|=N} u_{\mathbf{m}',\mathbf{n}} \sqrt{\frac{N}{\mathbf{n}!}} \frac{(-j\mathbf{R}\boldsymbol{\omega})^{\mathbf{n}}}{\|\boldsymbol{\omega}\|^N}. \quad (11)$$

Next, we expand the right-hand side multimonomial

$$\begin{aligned} (-j\mathbf{R}\boldsymbol{\omega})^{\mathbf{n}} &= \left(-j \sum_{i=1}^d r_{1i}\omega_i \right)^{n_1} \dots \left(-j \sum_{i=1}^d r_{di}\omega_i \right)^{n_d} \\ &= \prod_{i=1}^d \left(\sum_{|\mathbf{k}_i|=n_i} \frac{n_i!}{\mathbf{k}_i!} \mathbf{r}_i^{\mathbf{k}_i} (-j\boldsymbol{\omega})^{\mathbf{k}_i} \right) \\ &= \sum_{|\mathbf{k}_1|=n_1} \dots \sum_{|\mathbf{k}_d|=n_d} \prod_{i=1}^d \left(\sum_{|\mathbf{k}_i|=n_i} \frac{n_i!}{\mathbf{k}_i!} \mathbf{r}_i^{\mathbf{k}_i} (-j\boldsymbol{\omega})^{\mathbf{k}_i} \right) \\ &= \sum_{|\mathbf{k}_1|=n_1} \dots \sum_{|\mathbf{k}_d|=n_d} \frac{\mathbf{n}!}{\mathbf{k}_1! \dots \mathbf{k}_d!} \\ &\quad \times \mathbf{r}_1^{\mathbf{k}_1} \dots \mathbf{r}_d^{\mathbf{k}_d} (-j\boldsymbol{\omega})^{\mathbf{k}_1+\dots+\mathbf{k}_d}. \end{aligned}$$

Regrouping the terms corresponding to the same power $\mathbf{m} = \mathbf{k}_1 + \dots + \mathbf{k}_d$ of $(-j\boldsymbol{\omega})$ with the constraint that $|\mathbf{m}| = n_1 + \dots + n_d = N$, we express the aforementioned expansion in standard polynomial form

$$\sqrt{\frac{N}{\mathbf{n}!}} (-j\mathbf{R}\boldsymbol{\omega})^{\mathbf{n}} = \sum_{|\mathbf{m}|=N} s_{\mathbf{n},\mathbf{m}} \sqrt{\frac{N}{\mathbf{m}!}} (-j\boldsymbol{\omega})^{\mathbf{m}}$$

where $s_{\mathbf{n},\mathbf{m}}$ are the steering coefficients defined by (9). We then plug the aforementioned polynomial in (11) and obtain

$$\begin{aligned} \hat{U}_{\mathbf{m}'}(\mathbf{R}\boldsymbol{\omega}) &= \sum_{|\mathbf{n}|=N} u_{\mathbf{m}',\mathbf{n}} \sum_{|\mathbf{m}|=N} s_{\mathbf{n},\mathbf{m}} \sqrt{\frac{N}{\mathbf{m}!}} \frac{(-j\boldsymbol{\omega})^{\mathbf{m}}}{\|\boldsymbol{\omega}\|^N} \\ &= \sum_{|\mathbf{m}|=N} \underbrace{\left(\sum_{|\mathbf{n}|=N} u_{\mathbf{m}',\mathbf{n}} s_{\mathbf{n},\mathbf{m}} \right)}_{v_{\mathbf{m}',\mathbf{m}}} \sqrt{\frac{N}{\mathbf{m}!}} \frac{(-j\boldsymbol{\omega})^{\mathbf{m}}}{\|\boldsymbol{\omega}\|^N} \end{aligned}$$

which has the form of a generalized Riesz transform with coefficients $v_{\mathbf{m}',\mathbf{m}}$ corresponding to the matrix product $\mathbf{U}\mathbf{S}_{\mathbf{R}}$. ■

Remarkably, the steering matrices $\mathbf{S}_{\mathbf{R}}$ are orthogonal and endowed with the same group structure as the rotation matrices from which they are derived: $\mathbf{S}_{\mathbf{R}_2}\mathbf{S}_{\mathbf{R}_1} = \mathbf{S}_{\mathbf{R}_2\mathbf{R}_1}$.

Theorem 2: If \mathbf{R} is a $d \times d$ spatial rotation matrix, then the steering matrix $\mathbf{S}_{\mathbf{R}}$ with components $s_{\mathbf{n},\mathbf{m}}$ given by (9) is orthogonal and its inverse is $\mathbf{S}_{\mathbf{R}}^{-1} = \mathbf{S}_{\mathbf{R}}^T = \mathbf{S}_{\mathbf{R}^T}$.

Proof: To show that $\mathbf{S}_{\mathbf{R}}^{-1} = \mathbf{S}_{\mathbf{R}^T}$, we write the frequency response of the \mathbf{n} -component Riesz filter in the rotated coordinate system $\boldsymbol{\omega}' = \mathbf{R}\boldsymbol{\omega}$:

$$\sqrt{\frac{N}{\mathbf{n}!}} \frac{(-j\boldsymbol{\omega})^{\mathbf{n}}}{\|\boldsymbol{\omega}\|^N} = \sqrt{\frac{N}{\mathbf{n}!}} \frac{(-j\mathbf{R}'\boldsymbol{\omega}')^{\mathbf{n}}}{\|\boldsymbol{\omega}'\|^N}$$

with $\mathbf{R}' = \mathbf{R}^T$. Next, we apply the result of Theorem 1 to the right hand-side of this expression, which yields

$$\begin{aligned} \sqrt{\frac{N}{\mathbf{n}!}} \frac{(-j\boldsymbol{\omega})^{\mathbf{n}}}{\|\boldsymbol{\omega}\|^N} &= \sum_{|\mathbf{m}|=N} s'_{\mathbf{n},\mathbf{m}} \sqrt{\frac{N}{\mathbf{m}!}} \frac{(-j\boldsymbol{\omega}')^{\mathbf{m}}}{\|\boldsymbol{\omega}'\|^N} \\ &= \sum_{|\mathbf{m}|=N} s'_{\mathbf{n},\mathbf{m}} \sqrt{\frac{N}{\mathbf{m}!}} \frac{(-j\mathbf{R}\boldsymbol{\omega})^{\mathbf{m}}}{\|\boldsymbol{\omega}\|^N} \end{aligned} \quad (12)$$

where the coefficients $s'_{\mathbf{n},\mathbf{m}}$ are the entries of the steering matrix $\mathbf{S}_{\mathbf{R}'}$ = $\mathbf{S}_{\mathbf{R}^T}$, in accordance with (9). Applying Theorem 1 one more time to the rotated version of the frequency response in (12), we end up with an expression that is the composition (matrix multiplication) of the steering matrices $\mathbf{S}_{\mathbf{R}'}$ and $\mathbf{S}_{\mathbf{R}}$, from which we deduce that $\mathbf{S}_{\mathbf{R}'}\mathbf{S}_{\mathbf{R}} = \text{Identity}$, or equivalently, that $\mathbf{S}_{\mathbf{R}}^{-1} = \mathbf{S}_{\mathbf{R}^T}$.

To prove that $\mathbf{S}_{\mathbf{R}^T}$ with $\mathbf{R}^T = [\mathbf{r}'_1 \cdots \mathbf{r}'_d]^T$ is the transpose of $\mathbf{S}_{\mathbf{R}}$, we consider the form of $\mathbf{S}_{\mathbf{R}^T}$ given by (9), where we have explicitly written the sums $\sum_{|\mathbf{k}_j|=n_j} = \sum_{k_{j1}=0}^N \cdots \sum_{k_{jd}=0}^N \delta_{k_{j1}+\cdots+k_{jd},n_j}$ for each multiindex \mathbf{k}_j with $j = 1, \dots, d$

$$\begin{aligned} s'_{\mathbf{n},\mathbf{m}} &= \sqrt{\frac{\mathbf{m}!}{\mathbf{n}!}} \left(\sum_{k_{11}=0}^N \cdots \sum_{k_{dd}=0}^N (\delta_{|\mathbf{k}_1|,n_1} \cdots \delta_{|\mathbf{k}_d|,n_d}) \right. \\ &\quad \left. \cdot (\delta_{|\mathbf{k}_1|,m_1} \cdots \delta_{|\mathbf{k}_d|,m_d}) \frac{\mathbf{n}!}{\mathbf{k}_1! \cdots \mathbf{k}_d!} \mathbf{r}'_1{}^{\mathbf{k}_1} \cdots \mathbf{r}'_d{}^{\mathbf{k}_d} \right). \end{aligned} \quad (13)$$

Further, since $\mathbf{r}'_i{}^{\mathbf{k}_i} = \prod_{m=0}^d (r'_{im})^{k_{im}} = \prod_{m=0}^d (r_{mi})^{k_{im}}$, we rewrite $\mathbf{r}'_1{}^{\mathbf{k}_1} \cdots \mathbf{r}'_d{}^{\mathbf{k}_d} = \mathbf{r}_1{}^{\mathbf{k}'_1} \cdots \mathbf{r}_d{}^{\mathbf{k}'_d}$ in terms of multimonomials of the vector components of \mathbf{R} using the new multiindex vectors $\mathbf{k}'_i = (k_{1i}, \dots, k_{di})$ with $i = 1, \dots, d$. Using the fact that $\mathbf{k}_1! \cdots \mathbf{k}_d! = \mathbf{k}'_1! \cdots \mathbf{k}'_d!$, we then factorize the Kronecker products in (13) with respect to the index vectors \mathbf{k}'_i , which yields

$$\begin{aligned} s'_{\mathbf{n},\mathbf{m}} &= \sqrt{\frac{\mathbf{m}!}{\mathbf{n}!}} \sum_{|\mathbf{k}'_1|=m_1} \cdots \sum_{|\mathbf{k}'_d|=m_d} \delta_{\mathbf{k}'_1+\cdots+\mathbf{k}'_d,\mathbf{n}} \\ &\quad \times \frac{\mathbf{n}!}{\mathbf{k}'_1! \cdots \mathbf{k}'_d!} \mathbf{r}_1{}^{\mathbf{k}'_1} \cdots \mathbf{r}_d{}^{\mathbf{k}'_d}. \end{aligned}$$

By comparing this last equation with (9) and observing that $\mathbf{n}! \sqrt{\mathbf{m}!/\mathbf{n}!} = \mathbf{m}! \sqrt{\mathbf{n}!/\mathbf{m}!}$, we conclude that $s'_{\mathbf{n},\mathbf{m}} = s_{\mathbf{m},\mathbf{n}}$ for any \mathbf{n} and \mathbf{m} , such that $|\mathbf{n}| = N$ and $|\mathbf{m}| = N$, which proves that $\mathbf{S}_{\mathbf{R}}^T = \mathbf{S}_{\mathbf{R}^T} = \mathbf{S}_{\mathbf{R}}^{-1}$. ■

B. Tight Wavelet Frames

The pleasing consequence of Property 2 is that the generalized Riesz transform will automatically map any frame of $L_2(\mathbb{R}^d)$ into another one. In this paper, we are especially interested in tight (or Parseval) wavelet frames of $L_2(\mathbb{R}^d)$, which have the important property of providing self-reversible signal

transformations. These wavelet frames are families of functions $\{\psi_{i,\mathbf{k}}\}_{\mathbf{k} \in \mathbb{Z}^d, i \in \mathbb{Z}}$ that are indexed by a pair (\mathbf{k}, i) of location and scale (or resolution) indices. The wavelet property is that the basis functions at resolution i are rescaled and dilated versions of the ones at resolution 0: $\psi_{\mathbf{k},i}(\mathbf{x}) = \det(\mathbf{D})^{i/2} \psi_{\mathbf{k},0}(\mathbf{D}^{-i}\mathbf{x})$, where \mathbf{D} is an admissible $d \times d$ dilation matrix (i.e., $\det(\mathbf{D})$ must be a positive integer). The standard dyadic configuration corresponds to the choice $\mathbf{D} = 2\mathbf{I}$. A wavelet frame is also endowed with a shift-invariant structure in the sense that the wavelets at any given scale are generated by taking equidistant translates of $K = \det(\mathbf{D})$ mother wavelets with the property that $\psi_{i,\mathbf{k}+K\mathbf{l}}(\mathbf{x}) = \psi_{i,\mathbf{k}}(\mathbf{x} - K\mathbf{D}^i\mathbf{l})$. The K mother wavelets are often distinct, but this is not a strict necessity, e.g., our prior construction of polyharmonic spline wavelets in $L_2(\mathbb{R}^d)$, which involves a single mother wavelet [30], [31].

Definition 2: A family of wavelets $\{\psi_{i,\mathbf{k}}\}_{\mathbf{k} \in \mathbb{Z}^d, i \in \mathbb{Z}}$ is called a Parseval wavelet frame of $L_2(\mathbb{R}^d)$ iff it preserves energy

$$\forall f \in L_2(\mathbb{R}^d), \quad \|f\|_{L_2(\mathbb{R}^d)}^2 = \sum_{i \in \mathbb{Z}} \sum_{\mathbf{k} \in \mathbb{Z}^d} \left| \langle \psi_{i,\mathbf{k}}, f \rangle_{L_2} \right|^2.$$

Note that a Parseval frame is the normalized version of a tight frame and that its frame bounds are $A = B = 1$ [32], [33]. The remarkable aspect of a Parseval wavelet frame is that it satisfies the decomposition/reconstruction formula

$$\forall f \in L_2(\mathbb{R}^d), \quad f = \sum_{i \in \mathbb{Z}} \sum_{\mathbf{k} \in \mathbb{Z}^d} \langle \psi_{i,\mathbf{k}}, f \rangle_{L_2} \psi_{i,\mathbf{k}} \quad (14)$$

which has the same flavor as that associated with an orthonormal basis. The distinguishing feature of the frame generalization is that the wavelet family $\{\psi_{i,\mathbf{k}}\}$ may be redundant.

Based on the fact that the generalized Riesz transform has the same invariance properties as the N th-order Riesz transform considered in our previous work [19], we have the following result, which provides a simple, practical recipe for constructing steerable wavelet frames.

Proposition 1: The generalized N th-order Riesz transform $\mathcal{R}_{\mathbf{U}}$ maps a primal wavelet frame of $L_2(\mathbb{R}^d) \{\psi_{i,\mathbf{k}}\}_{i \in \mathbb{Z}, \mathbf{k} \in \mathbb{Z}^d}$ into another wavelet frame $\{\psi_{i,\mathbf{k}}^{\mathbf{n}}\}_{|\mathbf{n}|=N, i \in \mathbb{Z}, \mathbf{k} \in \mathbb{Z}^d}$ of $L_2(\mathbb{R}^d)$ with

$$\psi_{i,\mathbf{k}}^{\mathbf{n}} = [\mathcal{R}_{\mathbf{U}} \psi_{i,\mathbf{k}}]_{\mathbf{n}} = \sum_{|\mathbf{m}|=N} u_{\mathbf{n},\mathbf{m}} \mathcal{R}^{\mathbf{m}} \psi_{i,\mathbf{k}}. \quad (15)$$

In particular, if the primal transform is a Parseval frame and the generalized Riesz transform is energy preserving (e.g., $\mathbf{U}^H \mathbf{U} = \mathbf{I}$), then the generalized Riesz mapping yields a Parseval wavelet frame that satisfies the decomposition/reconstruction formula

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

The result is a straightforward extension of [19, Proposition 2]. The fundamental ingredient that makes this possible is the fact that the generalized Riesz transform globally preserves the L_2 norm of the signal (cf., Property 2). The other important aspect is the shift and scale invariance of the transform (Property 1), which ensures that the resulting functions are bona-fide wavelets; i.e., the generalized Riesz wavelets are

shifted and dilated versions of a fixed number of generators (mother wavelets). Also note that the redundancy is increased in proportion to the number of Riesz components $M = \binom{N+d-1}{d-1}$.

If, in addition, the mother wavelets in (14) are isotropic, then we end up with a decomposition that is steerable, and therefore, rotation-covariant, due to Theorem 1. We may also consider orientation-free wavelets—i.e., wavelets that do not exhibit any preferential orientation—as an acceptable compromise. We can then rotate the whole set of generalized Riesz wavelets ψ^n to any desired orientation by using the steering relation

$$[\mathcal{R}_{\mathbf{S}_R \mathbf{U}} \psi(\mathbf{x})]_{\mathbf{n}} = \sum_{|\mathbf{m}|=N} s_{\mathbf{n}, \mathbf{m}} \psi^{\mathbf{m}}(\mathbf{x}) \approx \psi^{\mathbf{n}}(\mathbf{R}\mathbf{x}) \quad (16)$$

where the \mathbf{R} is the desired $d \times d$ rotation matrix and \mathbf{S}_R is the steering matrix given by (9). The quality of the right-hand side approximation will depend on the degree of isotropy of the primal wavelet $\psi(\mathbf{x})$; it will be a strict equality iff ψ is such that $\psi(\mathbf{x}) = \psi(\mathbf{R}\mathbf{x})$. The practical interest of this scheme is that it can yield a rich variety of steerable wavelet frames that are parametrized by the matrix \mathbf{U} and associated with a single primal wavelet transform. The remarkable aspect is that the design is essentially constraint-free; the only requirement for a Parseval frame is that \mathbf{U} be unitary.

IV. CONSTRUCTION OF PRIMAL, ORIENTATION-FREE WAVELET FRAMES

The present formulation allows for a wavelet design, where the directional and multiscale aspects are completely decoupled. To obtain a primal wavelet representation with good localization and/or isotropy properties, we consider two options: 1) the radial bandlimited design [13], [27], and 2) a multiresolution-based approach in the spirit of Burt and Adelson that generalizes the L_2 -spline pyramids that we had introduced in early work [25]. To the best of our knowledge, the first approach is the only one that yields wavelets that are perfectly isotropic; the price to pay for this property is a slower decay in space. The second option yields orientation-free wavelets that retain all the desirable features of conventional wavelet bases (good spatial localization, stability, approximation theoretic properties, fast implementation, etc.)—the less favorable aspect there is that the wavelets are only approximately isotropic.

A. Preliminaries: From Multiresolution Riesz Bases to Wavelet Frames

We like to present the two approaches from the common perspective of the transformation of a *Riesz basis*¹ into a frame spanning a “smaller” wavelet subspace. We also rely on Mallat’s multiresolution analysis of $L_2(\mathbb{R}^d)$ to extend the representation to the entire space of finite-energy functions. While the concept is similar to Mallat’s construction of wavelet bases [34], the fundamental difference is that the present multiresolution decompositions, which are slightly redundant, have a simpler,

¹A Riesz basis, which is the nonredundant version of a frame, is a fundamental concept in functional analysis that was introduced by the Hungarian mathematician Frigyes Riesz; it is unrelated to the Riesz transform that is due to Marcel Riesz (Frigyes’ younger brother).

shift-invariant structure with a single-wavelet channel per resolution level. Each level of decomposition is qualitatively equivalent to a bandpass filtering of the input signal.

The first ingredient of our formulation is the following projection mechanism, which converts a Riesz basis of some “larger” space V into a frame of $W \subset V$.

Theorem 3: Let V and W be two closed subspaces of some generic Hilbert space such that $W \subset V$. Furthermore, let $\{\varphi_k\}$ be a Riesz basis of V that is characterized by the decomposition/reconstruction formula

$$\forall f \in V, \quad f = \sum_k \langle \tilde{\varphi}_k, f \rangle \varphi_k \quad (17)$$

where $\{\tilde{\varphi}_k\}$ is the dual basis of $\{\varphi_k\}$; i.e., $\langle \tilde{\varphi}_k, \varphi_l \rangle = \delta_{k-l}$ with $\varphi_k, \tilde{\varphi}_l \in V$. Then, the functions

$$\psi_k = \text{Proj}_W \varphi_k \quad \text{and} \quad \tilde{\psi}_k = \text{Proj}_W \tilde{\varphi}_k$$

where Proj_W is the orthogonal projector operator into W , define a dual set of frames of W with the following dual decomposition/reconstruction formulas:

$$\forall f \in W, \quad f = \sum_k \langle \tilde{\psi}_k, f \rangle \psi_k = \sum_k \langle \psi_k, f \rangle \tilde{\psi}_k.$$

Moreover, the Riesz bounds of $\{\varphi_k\}$ are mapped into the frame bounds of $\{\psi_k\}$ and likewise for their duals.

The proof is given in Appendix A.

The second ingredient is the specification of a multiresolution sequence of spaces $V = V_i$ and their corresponding Riesz bases to be used in Proposition 2. This is done by selecting a valid scaling $\varphi(\mathbf{x})$ (e.g., $\varphi(\mathbf{x}) = \text{sinc}(\mathbf{x})$, if we want the decomposition to be bandlimited). For simplicity, we restrict the discussion to the case of a dyadic dilation matrix $\mathbf{D} = 2\mathbf{I}$. To generate a valid multiresolution of $L_2(\mathbb{R}^d)$, φ must satisfy the three admissibility conditions, which are necessary and sufficient [35]:

- 1) it generates a shift-invariant Riesz basis $\{\varphi(\mathbf{x} - \mathbf{k})\}_{\mathbf{k} \in \mathbb{Z}^d}$, which is equivalent to the condition (cf., [36])

$$0 < A \leq \sum_{\mathbf{n} \in \mathbb{Z}^d} |\hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{n})|^2 \leq B < +\infty;$$

- 2) it satisfies the two-scale relation

$$\varphi(\mathbf{x}/2) = 2^{d/2} \sum_{\mathbf{k} \in \mathbb{Z}^d} h[\mathbf{k}] \varphi(\mathbf{x} - \mathbf{k}); \quad (18)$$

- 3) it yields a partition of unity $\sum_{\mathbf{k} \in \mathbb{Z}^d} \varphi(\mathbf{x} - \mathbf{k}) = 1$. By denoting $\varphi_{i, \mathbf{k}}(\mathbf{x}) = 2^{-id/2} \varphi(\mathbf{x}/2^i - \mathbf{k})$, we specify the multiresolution ladder of approximation spaces

$$V_i = \left\{ f_i(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] \varphi_{i, \mathbf{k}}(\mathbf{x}) : c \in \ell_2(\mathbb{Z}^d) \right\} \quad (19)$$

that are indexed by the scale parameter $i \in \mathbb{Z}$; the step size at resolution i is $a = 2^i$ and the nesting is such that $V_i \subset V_{i-1} \cdots \subset L_2(\mathbb{R}^d)$. Note that the third assumption (partition

of unity) guarantees that $\lim_{i \rightarrow -\infty} V_i$ is dense in $L_2(\mathbb{R}^d)$; it is essential for the L_2 -completeness of the wavelet decomposition.

We now proceed with the construction of orientation-free wavelets by projecting the multiresolution Riesz basis $\{\varphi_{i,\mathbf{k}}\}_{\mathbf{k} \in \mathbb{Z}^d}$ onto some appropriate wavelet subspace $W = W_{i+1} \subset V = V_i$.

B. Isotropic Bandlimited Wavelet Frames

There are a number of constructions in the literature that fall into this category [13], [27]–[29]. Before reviewing them, we apply the aforementioned projection strategy to obtain a straightforward design that is in direct correspondence with Shannon's sampling theorem, and that is the starting point for the specification of Meyer-type wavelets.

1) *Construction of Isotropic, Shannon-Type Wavelets:* By selecting $\varphi(\mathbf{x}) = \text{sinc}(\mathbf{x}) = \prod_{i=1}^d (\sin(\pi x_i)/\pi x_i)$ in (19), we specify the so-called Shannon multiresolution analysis of $L_2(\mathbb{R}^d)$, which consists of a sequence of embedded subspaces V_i that are bandlimited to $\boldsymbol{\omega} \in [-(\pi/2^i), (\pi/2^i)]^d$

$$V_i = \text{span}\{\text{sinc}_{i,\mathbf{k}}\}_{\mathbf{k} \in \mathbb{Z}^d} \\ = \left\{ f \in L_2(\mathbb{R}^d) : \hat{f}(\boldsymbol{\omega}) = 0, \forall \boldsymbol{\omega} \notin \left[-\frac{\pi}{2^i}, \frac{\pi}{2^i}\right]^d \right\}.$$

We then define some corresponding wavelet subspaces of radial bandpass functions

$$W_{i+1} = \left\{ f \in L_2(\mathbb{R}^d) : \hat{f}(\boldsymbol{\omega}) = 0 \text{ for } \|\boldsymbol{\omega}\| < \frac{\pi}{2^{i+1}} \text{ and } \frac{\pi}{2^i} \leq \|\boldsymbol{\omega}\| \right\}. \quad (20)$$

Since W_{i+1} is a closed subspace of V_i , we can apply Proposition 3 to its orthogonal sinc basis to obtain the tight wavelet frame $\{\psi_{i,\mathbf{k}}\}_{\mathbf{k} \in \mathbb{Z}^d}$ of W_{i+1} with

$$\psi_{i,\mathbf{k}} = \text{Proj}_{W_{i+1}} \text{sinc}_{i,\mathbf{k}} = \psi_{i,\mathbf{k}} = g_i * \text{sinc}_{i,\mathbf{k}} \quad (21)$$

where $g_i(\mathbf{x}) = g_i(\|\mathbf{x}\|)$ is the impulse response of the ideal radial bandpass filter, whose frequency response is $\hat{g}_i(\boldsymbol{\omega}) = 1_{\pi/(2^{i+1}) \leq \|\boldsymbol{\omega}\| < (\pi/2^i)}$, i.e., the indicator function corresponding to the spectral support of W_{i+1} . Based on the fact that $\cup_{i \in \mathbb{Z}} W_i$ is dense in $L_2(\mathbb{R}^d)$, the design procedure yields a tight wavelet frame of $L_2(\mathbb{R}^d)$. The energy-preserving condition that ensures that the latter is true is as follows:

$$\sum_{i \in \mathbb{Z}} |\hat{g}_i(\boldsymbol{\omega})|^2 = \sum_{i \in \mathbb{Z}} |\hat{g}_0(\boldsymbol{\omega} 2^i)|^2 = 1. \quad (22)$$

It is automatically fulfilled, since the sequence of ideal radial bandpass filters constitutes a tiling of the frequency domain \mathbb{R}^d , as illustrated in Fig. 1. For a complete characterization of these wavelets, we refer to the work of Papadakis *et al.* [27].

2) *Specification of Meyer-Type Wavelets:* While the aforementioned construction yields a tight isotropic wavelet frame of $L_2(\mathbb{R}^d)$, it has the drawback of producing wavelets with poor spatial decay (e.g., $\psi_{i,\mathbf{k}}(\mathbf{x}) = O(1/\|\mathbf{x}\|)$), due to the sharp cutoff in frequency domain. A remarkable observation, which can be traced back to the early work of Daubechies and Meyer on frames [37], [38], is that this can be corrected via an appropriate adjustment of the radial bandpass filtering functions $\hat{g}_i(\boldsymbol{\omega})$, which need not be indicator functions, as long as they




Fig. 1. Tiling of the 2-D frequency domain using radial-bandpass filters. The shaded area corresponds to the spectral support of the wavelet subspace W_{i+1} ; it is included in the spectral support of V_i (enclosing square).

satisfy (22). This leads to the following extended definition of the wavelet subspaces

$$W'_{i+1} = \{r(\mathbf{x}) = (g_i * f)(\mathbf{x}) : f \in V_i\}$$

which is equivalent to (20) if g_i is the impulse response of the ideal radial bandpass filter. Since r can be written as $r = \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] \text{sinc}_{i,\mathbf{k}}$, there exists a sequence $c[\mathbf{k}]$ such that

$$\forall r \in W'_{i+1}, \quad r = \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] \psi_{i,\mathbf{k}}$$

where the wavelet functions $\psi_{i,\mathbf{k}}$ are still given by (21). This indicates that $\{\psi_{i,\mathbf{k}}\}_{\mathbf{k} \in \mathbb{Z}^d}$ is a frame of W'_{i+1} , albeit not necessarily a tight one. Yet, if Condition (22) is satisfied, then one recovers the tight frame property over $L_2(\mathbb{R}^d)$, which is the union of the wavelet subspaces W'_i , $i \in \mathbb{Z}$. The condition for the wavelet frame to be isotropic is that the restriction of the filtering function h_i over V_i be isotropic, i.e., $\text{Proj}_{V_i} \{g_i\}(\mathbf{x}) = \text{Proj}_{V_i} \{g_i\}(\|\mathbf{x}\|)$.

The aforementioned functional framework accounts for all known constructions of isotropic wavelet frames of $L_2(\mathbb{R}^d)$, which are summarized in Table I. The common feature of the wavelet profile functions in Table I is that they are compactly supported within the frequency interval $(-\pi, +\pi)$ (bandlimited property) and that they satisfy a rescaled version of the partition of unity condition

$$\sum_{i \in \mathbb{Z}} \left| \hat{\psi}(2^i \boldsymbol{\omega}) \right|^2 = 1$$

which is the key to ensuring the tight frame property [38]. Two remarks are in order with respect to Simoncelli's pioneering design which is, by far, the solution most widely used in applications. First, the response of the filter is a warped version of a bump [cf., Fig. 2(a)] and the furthest away from an ideal bandpass filter; this simply reflects the fact that a design objective for the steerable pyramid was to approximate the behavior of a log-Gabor filterbank, which is a well-accepted model of the response properties of cortical cells in the mammalian visual system [39]. Second, Simoncelli and coworkers put a lot of emphasis on finite-impulse response filter design in their initial

TABLE I
RADIAL FREQUENCY RESPONSES OF ISOTROPIC BANDLIMITED WAVELETS

Wavelet type	Wavelet mother function
Shannon	$\hat{\psi}(\omega) = \begin{cases} 1, & \frac{\pi}{2} \leq \omega \leq \pi \\ 0, & \text{otherwise} \end{cases}$
Simoncelli [13]	$\hat{\psi}(\omega) = \begin{cases} \cos\left(\frac{\pi}{2} \log_2\left(\frac{2\omega}{\pi}\right)\right), & \frac{\pi}{4} < \omega \leq \pi \\ 0, & \text{otherwise} \end{cases}$
Meyer [38]	$\hat{\psi}(\omega) = \begin{cases} \sin\left(\frac{\pi}{2} \nu\left(\frac{4\omega}{\pi} - 1\right)\right), & \frac{\pi}{4} < \omega \leq \frac{\pi}{2} \\ \cos\left(\frac{\pi}{2} \nu\left(\frac{2\omega}{\pi} - 1\right)\right), & \frac{\pi}{2} < \omega \leq \pi \\ 0, & \text{otherwise} \end{cases}$
Papadakis [28]	$\hat{\psi}(\omega) = \begin{cases} \frac{\sqrt{1+\sin(5\omega)}}{2}, & \frac{3\pi}{10} < \omega \leq \frac{\pi}{2} \\ 1, & \frac{\pi}{2} < \omega \leq \frac{3\pi}{5} \\ \frac{\sqrt{1-\sin(5\omega)}}{2}, & \frac{3\pi}{5} < \omega \leq \pi \\ 0, & \text{otherwise} \end{cases}$
Held [29]	$\hat{\psi}(\omega) = \begin{cases} \cos\left(2\pi q\left(\frac{\omega}{2\pi}\right)\right), & \frac{\pi}{4} < \omega \leq \frac{\pi}{2} \\ \sin\left(2\pi q\left(\frac{\omega}{4\pi}\right)\right), & \frac{\pi}{2} < \omega \leq \pi \\ 0, & \text{otherwise} \end{cases}$

work [1], [4]; later on they adopted the Fourier-based construction described here, which is implemented in the current software release of the steerable pyramid, as documented in [10] and [13]. The other constructions in Table I are more mathematically oriented and motivated by the desire to soften the cutoff transition of the Shannon-type wavelet to improve the spatial decay of the wavelets. The constraint there is to produce an ideal-filter-like frequency response with the highest possible order of differentiation [38]. For the Meyer wavelet of order N , the auxiliary function $\nu(t)$ is a polynomial of degree $2N + 1$, which is chosen such that: $\nu(t) = 0$ if $t \leq 0$, $\nu(t) = 0$ if $t \geq 1$, and $\nu(t) + \nu(1-t) = 1$, while $\nu \in C^N([0, 1])$. For instance, the auxiliary function that achieves a frequency response with $N = 3$ continuous derivatives is $\nu(t) = t^4(35 - 84t + 70t^2 - 20t^3)$ [cf., Fig. 2(b)]. For the Held wavelet of order N [29], the function $q(t)$ is a polynomial of degree N . For instance, the polynomial $q(t) = 256t^3 - 144t^2 + 24t - 1$ leads to $C^3([0, 1])$. One can also push the scheme up to $C^\infty([0, 1])$ by choosing $q(t) = \lambda(4 - 16t)(4(\lambda(4 - 16t) + \lambda(-2 + 16t)))^{-1}$, with $\lambda(t) = \exp(-t^{-2})$ if $t > 0$, and $\lambda(t) = 0$ otherwise, which yields a primal isotropic wavelet with exponential decay.

Each of these wavelets specifies a tight isotropic wavelet frame that is especially convenient to implement in the frequency domain using the fast Fourier transform (FFT) algorithm. The analysis part of the algorithm involves the creation of a bandlimited pyramid by successive low-pass filtering and downsampling by a factor of two. The signal at each level of the pyramid is then processed with an isotropic high-pass filter (cf., wavelet profile in Table I) to yield the corresponding wavelet coefficients. Note that the filtering and downsampling operations are implemented in the Fourier domain, while the wavelet




Fig. 2. Radial frequency responses of two sequences of bandpass wavelet filters. (a) Simoncelli wavelets. (b) Meyer-type wavelets.

coefficients at a given scale are recovered by performing an inverse FFT of the appropriate size. The main difference with a standard multidimensional wavelet transform is that there is a single mother wavelet at each scale and that there is no downsampling in the high-pass branches of the algorithm.

C. Pyramid-Like Wavelet Frames

The next method is generic and applicable to any multidimensional scaling function φ without any requirement for separability. The multiresolution subspaces V_i are specified by (19) and the corresponding wavelet subspaces are the ones considered in conventional wavelet theory: W_{i+1} is the orthogonal complement of V_{i+1} with respect to V_i , so that $V_i = V_{i+1} \oplus W_{i+1}$ with $W_{i+1} \perp V_{i+1}$.

The principle of our wavelet construction is then straightforward: we take $V = V_i = V_{i+1} \oplus W_{i+1}$ with $\{\varphi_{i,\mathbf{k}}\}_{\mathbf{k} \in \mathbb{Z}^d}$ as its Riesz basis, which we then project onto $W = W_{i+1}$ to obtain a corresponding wavelet frame according to the procedure outlined in Proposition 2. Specifically, we select our analysis/synthesis wavelets as follows:

$$\psi_{i,\mathbf{k}} = \text{Proj}_{W_{i+1}} \varphi_{i,\mathbf{k}} = \varphi_{i,\mathbf{k}} - \text{Proj}_{V_{i+1}} \varphi_{i,\mathbf{k}} \quad (23)$$

$$\tilde{\psi}_{i,\mathbf{k}} = \text{Proj}_{W_{i+1}} \tilde{\varphi}_{i,\mathbf{k}} = \tilde{\varphi}_{i,\mathbf{k}} - \text{Proj}_{V_{i+1}} \tilde{\varphi}_{i,\mathbf{k}} \quad (24)$$

where $\tilde{\varphi}_{i,\mathbf{k}} \in V_i$ is the dual basis of $\varphi_{i,\mathbf{k}}$. The corresponding primary decomposition/reconstruction formula is as follows:

$$\forall g \in W_{i+1}, \quad g = \sum_{\mathbf{k} \in \mathbb{Z}^d} \langle \psi_{i,\mathbf{k}}, g \rangle \tilde{\psi}_{i,\mathbf{k}}.$$




Fig. 3. Two basic building blocks of the wavelet pyramid algorithm. (a) Reduce module. (b) Expand module.




Fig. 4. Schematic description of the (a) analysis and (b) synthesis procedures for the decomposition and reconstruction of a signal in a pyramid-based wavelet frame.

Moreover, because $\psi_{i,\mathbf{k}}, \tilde{\psi}_{i,\mathbf{k}} \in W_{i+1}$, we can write the orthogonal projection of any function $f \in L_2(\mathbb{R}^d)$ into W_{i+1} as $r_{i+1} = \text{Proj}_{W_{i+1}} f = \sum_{\mathbf{k} \in \mathbb{Z}^d} \langle \psi_{i,\mathbf{k}}, f \rangle \psi_{i,\mathbf{k}}$. Finally, by summing up these orthogonal wavelet components, we obtain a wavelet frame decomposition that has the same form as (14), and whose frame bounds are the Riesz bounds of the generating function φ . In particular, the procedure yields a Parseval wavelet frame whenever φ is orthogonal, i.e., when $A = B = 1$, which is equivalent to $\sum_{\mathbf{n} \in \mathbb{Z}^d} |\hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{n})|^2 = 1$ almost everywhere.

While the scheme is reminiscent of a standard decomposition into a wavelet basis, it is not quite equivalent because the wavelet functions at each scale are slightly redundant [by a factor $2^d/(2^d - 1)$].

The design part of the process is to select a scaling function φ —for instance, a B-spline or an orthonormalized version of it—such that the wavelets have good space-domain localization and isotropy properties. Alternatively, one may also specify the basis functions implicitly via the choice of the refinement filter $h[\mathbf{k}]$.

The remarkable aspect of the proposed wavelet pyramid decomposition is that the wavelet coefficients $w_i[\mathbf{k}] = \langle \psi_{i,\mathbf{k}}, f \rangle$ are identical to the expansion coefficients of $r_{i+1} = \text{Proj}_{W_{i+1}} f = \text{Proj}_{V_i} f - \text{Proj}_{V_{i+1}} f$ in the dual basis $\{\tilde{\varphi}_{i,\mathbf{k}}\}_{\mathbf{k} \in \mathbb{Z}^d}$ that also spans the finer scale approximation space V_i . In Appendix B, we exploit this property to derive the filterbank implementation of the transform and to make the connection with the approach of Do and Vetterli [26], which is computationally equivalent. The final algorithm involves the dual pair of refinement filters h and \tilde{h} and is summarized in Figs. 3 and 4. While the latter constitutes a well-known extension of Burt and Adelson’s algorithm [23], [26], the functional

characterization of the underlying wavelet functions given by (23) and (24) is new to the best of our knowledge—it is concise and “filterbank-free”.

V. EXPERIMENTAL RESULTS

In the sequel, we discuss some practical issues and present experimental results that illustrate the potential of the proposed steerable wavelet transforms for basic image processing tasks, such as denoising and feature extraction.

A. Steering the Wavelet Transform

Since the proposed wavelet transforms are steerable, it makes good sense for best feature extraction to reorient the generalized Riesz wavelets $\psi_{i,\mathbf{k}}^{\mathbf{n}}$ at index (i, \mathbf{k}) (as defined by Proposition 1) in a data-adaptive fashion by applying some spatial rotation matrix $\mathbf{R}_{i,\mathbf{k}}$. The result of this process is an adaptive generalized Riesz wavelet frame $\{\mathcal{R}_{\mathbf{U}_{i,\mathbf{k}}} \psi_{i,\mathbf{k}}\}_{i \in \mathbb{Z}, \mathbf{k} \in \mathbb{Z}^d}$ with $\mathbf{U}_{i,\mathbf{k}} = \mathbf{U} \mathbf{S}_{\mathbf{R}_{i,\mathbf{k}}}$, where $\mathbf{S}_{\mathbf{R}_{i,\mathbf{k}}}$ is the steering matrix at wavelet location (i, \mathbf{k}) (cf., Theorem 1). The remarkable feature is that $\mathbf{R}_{i,\mathbf{k}}$ can vary locally without affecting the frame bounds and the perfect reconstruction properties of the underlying generalized Riesz wavelet transform, as a direct consequence of Theorem 2.

We did experiment with different performance criteria for steering the wavelet transform, such as the maximization of the ℓ_∞ norm of the wavelet coefficients at location (i, \mathbf{k}) to favor strong wavelet responses (Max steering), or the minimization of their ℓ_1 -norm with the aim of producing a sparse decomposition. While both approaches gave adequate results, we finally settled for a first-order tensor-based orientation determination, which proved to be more robust and also cheaper computationally.

The proposed solution is to select the local orientation that maximizes the directional Hilbert transform of the primal wavelet decomposition within a neighborhood specified by a Gaussian-like window v . We recall that the directional Hilbert transform of a signal $f(\mathbf{x})$ along some unit vector $\mathbf{u} \in \mathbb{R}^d$ is given by (cf., [19])

$$\mathcal{H}_{\mathbf{u}} f(\mathbf{x}) = \langle \mathbf{u}, \mathcal{R} f(\mathbf{x}) \rangle$$

where $\mathcal{R} f(\mathbf{x})$ is the d -component Riesz transform of f , the aforementioned is actually the Riesz counterpart of the directional derivative $D_{\mathbf{u}} f(\mathbf{x}) = \langle \mathbf{u}, \nabla f(\mathbf{x}) \rangle$. Then, for a given primal isotropic analysis wavelet $\psi_i(\mathbf{x}) = \psi(\mathbf{x}/2^i)$ at resolution i , we search for the orientation such that the energy of $\mathcal{H}_{\mathbf{u}}(\psi_i * f)(\mathbf{x})$ is maximized within a local neighborhood of $\mathbf{x}_0 = 2^i \mathbf{k}_0$. We can readily show (cf., [31, Sec. III-B]) that the corresponding rotation matrix is $\mathbf{R}_{i,\mathbf{k}_0} = [\mathbf{u}_1 \cdots \mathbf{u}_d]^T$, where the \mathbf{u}_n are the eigenvectors of the local $d \times d$ wavelet structure matrix

$$\mathbf{J}_{i,\mathbf{k}_0} = \sum_{\mathbf{k} \in \mathbb{Z}^d} v[\mathbf{k} - \mathbf{k}_0] \mathbf{v}_i[\mathbf{k}] \mathbf{v}_i^T[\mathbf{k}] \quad (25)$$

where $v[\mathbf{k} - \mathbf{k}_0]$ is a nonnegative, symmetric, and compactly supported weighting window centered at \mathbf{k}_0 and where $\mathbf{w}_i[\mathbf{k}] = \langle f, \mathcal{R} \psi_{i,\mathbf{k}} \rangle_{L_2(\mathbb{R}^d)}$ are the (first-order) Riesz wavelet coefficient of f at index location (i, \mathbf{k}) . In our experiments, we used a fixed Gaussian window v of standard deviation 2.

B. Signal-Adapted Wavelets and PCA

Having at our disposal a robust mechanism for locally orienting the wavelets in the direction of the predominant image patterns, we can take advantage of the degrees of freedom offered by the generalized Riesz transform to construct wavelet templates that are best matched to a particular class of images.

To design such optimized wavelets, we developed an algorithm that is based on PCA in a rotated wavelet frame. Given an N th-order generalized Riesz transform $\mathcal{R}_{\mathbf{U}}$, the training phase consists in computing and steering the wavelet coefficients $\mathbf{w}_i[\mathbf{k}] = \langle f, \mathcal{R}_{\mathbf{U}_i, \mathbf{k}} \psi_{i, \mathbf{k}} \rangle$ at every location (i, \mathbf{k}) and in building up some corresponding scatter matrices $\mathbf{C}_i = \sum_{\mathbf{k}} \mathbf{w}_i[\mathbf{k}] \mathbf{w}_i[\mathbf{k}]^T$ at each scale using the test images at our disposal. The corresponding scale-adaptive PCA wavelet transform is then specified by the series of generalized Riesz transforms $\mathcal{R}_{\mathbf{U}_i, \mathbf{U}}$ with $i = 1, \dots, I$, where the transformation matrices \mathbf{U}_i are given by the eigenvectors of the scatter matrices \mathbf{C}_i (without subtraction of the mean). This also means that the PCA wavelet transform can be implemented straightforwardly by performing the primary generalized Riesz wavelet analysis, steering the wavelet coefficients according to the procedure outlined in Section IV-A, and applying the linear transformation matrices \mathbf{U}_i on a scale-by-scale basis. The reconstruction part of the algorithm follows the reverse path: matrix multiplication with $\mathbf{U}_i^T = \mathbf{U}_i^{-1}$, inverse of the steering matrix $\mathbf{S}_{\mathbf{R}_{i, \mathbf{k}}}^T$, and signal recomposition using the synthesis branch of the initial generalized Riesz transform.

Since the PCA is applied at each wavelet location (i, \mathbf{k}) , we found it beneficial to perform our training on an equalized version of the Riesz transform, which involves modified Riesz wavelets that are transformed to be locally orthonormal. These equalized wavelets are defined by considering the generalized Riesz transform $\mathcal{R}_{\mathbf{G}^{-1/2}}$, where $\mathbf{G}^{-1/2}$ is the square root inverse of the componentwise Gram matrix \mathbf{G} of the Riesz wavelets, whose multiindex entries are given by

$$\begin{aligned} [\mathbf{G}]_{\mathbf{m}, \mathbf{n}} &= \langle \psi^{\mathbf{m}}, \psi^{\mathbf{n}} \rangle_{L_2(\mathbb{R}^d)} \\ &= \frac{1}{(2\pi)^d} \frac{N}{\sqrt{\mathbf{m}! \cdot \mathbf{n}!}} \int_{\mathbb{R}^d} |\hat{\psi}(\boldsymbol{\omega})|^2 \frac{\boldsymbol{\omega}^{\mathbf{m}} \boldsymbol{\omega}^{\mathbf{n}}}{\|\boldsymbol{\omega}\|^{2N}} d\boldsymbol{\omega} \end{aligned}$$

with $|\mathbf{m}| = |\mathbf{n}| = N$ and $\hat{\psi}(\boldsymbol{\omega})$ is the Fourier transform of the primary isotropic wavelet $\psi(\mathbf{x})$. The Gram matrix may be computed analytically by switching to polar coordinates—it does actually not depend on the wavelet if the latter is normalized and truly isotropic. A comparative display of nonequalized and equalized Riesz wavelets is given in Fig. 5.

In the sequel, we will refer to the rotation-covariant PCA wavelet analysis performed on the equalized Riesz wavelet coefficients as “equalized PCA” to differentiate it from “Riesz PCA,” which corresponds to PCA applied to the canonical N th-order Riesz transform, i.e., $\mathcal{R}^{(N)} = \mathcal{R}_{\mathbf{I}}$ where \mathbf{I} is the identity matrix. The former is advantageous for denoising because it ensures that the noise contribution is independent and identically distributed across wavelet channels. This justifies componentwise processing with a common thresholding function. Some examples of equalized PCA wavelets are shown in Fig. 6. In all cases, the first wavelet, which is typically the most




Fig. 5. Examples of steerable wavelets (image-domain representation). (a) Riesz wavelets of order 4. (b) Equalized Riesz wavelets of order 4.

directional one within its set, accounts for a large proportion of the total energy. This is consistent with the chosen steering mechanism. Despite the fact that the overall diversity of the wavelets in the upper two lower order examples ($N = 4$) is limited, there are some notable differences in energy distribution and shape. For instance, it is interesting to observe that the second and third wavelets in Fig. 6(b) exhibit a central wave-like blob that is reminiscent of patterns present in the texture fragment displayed on the right. The PCA wavelet shapes are much more elongated in the higher order example (Barbara). While the first wavelet is clearly a vertical line detector, there is a tendency in the next functions to oscillate more and to become less directional; this probably helps encoding slight deviations from the reference in straightness and thickness.

Interestingly, the first basis function of “equalized PCA” is the steerable wavelet that achieves the maximum possible signal-to-noise ratio for the class of signals under consideration. We should note, however, that the “equalized PCA” Riesz wavelet transform loses the tight frame property because it involves a transformation that is the composition of an orthogonal matrix and a diagonal scaling matrix. In practice, this is not a problem; indeed, any generalized Riesz wavelet transform can be inverted straightforwardly by the application of the appropriate matrix inverse: \mathbf{U}^{-1} .

C. Simoncelli’s Equiangular Wavelets

As pointed out in our previous work, equiangular wavelet filter configurations are only feasible in two dimensions. A remarkable property is that the choice of any (normalized) filter that is in the span of the Riesz wavelets will automatically yield a tight wavelet frame (cf., [19, Theorem 2]). The standard implementation of Simoncelli’s steerable pyramid, which is documented in [13], corresponds to choosing the primary directional wavelet $\psi_1(\mathbf{x}) = \mathcal{R}^{(N,0)} \psi(\mathbf{x})$, where $\psi(\mathbf{x})$ is the isotropic bandlimited wavelet, whose radial frequency response is specified by the second equation in Table I. The corresponding polar-separable frequency response is $\hat{\psi}_1(\boldsymbol{\omega}) = \hat{\psi}(\boldsymbol{\omega}) (-j \cos \theta)^N$, where $\boldsymbol{\omega}$ and θ are the radial and angular frequency variable, respectively. The directional pyramid decomposition at level i is implemented via an $(N + 1)$ -channel filterbank, where the individual wavelet filters are rotated versions of the primary one along the directions $\xi_n = (\pi(n - 1))/(N + 1)$ with $n = 1, \dots, N + 1$. This equiangular configuration happens to be a special instance of the generalized N th-order Riesz wavelet




Fig. 6. Examples of steerable signal-adapted wavelets (equalized PCA) at the finer scale of resolution. (a) Complete set of wavelets for noise-free Lena with $N = 4$. (b) Complete set of wavelets for noise-free Brodatz texture D104 with $N = 4$. (c) Seven most significant wavelets for noise-free Barbara with $N = 14$. The percentage below the pictures represents the signal energy contribution as measured by the corresponding eigenvalue.

transform specified in Proposition 1. The underlying transformation matrix $\mathbf{U}_{\text{Simon}}$ can be determined by considering the first component [i.e., $\mathbf{n} = (N, 0)$] of the general steering relation (9) with $\mathbf{u} = (\cos \xi_n, \sin \xi_n)$. Specifically, the entries of the $(N + 1) \times (N + 1)$ matrix that defines Simoncelli's steerable pyramid of order N are given by

$$[\mathbf{U}_{\text{Simon}}]_{m+1, n+1} = \sqrt{\binom{N}{m}} \cos\left(\frac{\pi n}{N+1}\right)^m \times \sin\left(\frac{\pi n}{N+1}\right)^{N-m}$$

with $m + 1, n + 1 \in \{1, \dots, N + 1\}$ (standard row-column matrix indexing).

D. Generic Implementation of Steerable Wavelet Transforms

The steerable wavelet transforms described earlier are all special instances of the generalized Riesz wavelet transform in Proposition 1. This motivated us to develop a generic steerable wavelet decomposition algorithm via the suitable composition of three basic modules: 1) N th-order Riesz transform; 2) primal wavelet decomposition; and 3) linear transformation and steering in the wavelet domain. The key idea here is that the Riesz transform can be factored out of the wavelet decomposition because it commutes with translation and dilation (cf., Property 1). Specifically, we have that

$$\langle f, \psi_{i, \mathbf{k}}^{\mathbf{n}} \rangle = \langle f, \mathcal{R}^{\mathbf{n}} \psi_{i, \mathbf{k}} \rangle = \langle \mathcal{R}^{\mathbf{n}*} f, \psi_{i, \mathbf{k}} \rangle$$

where $\mathcal{R}^{\mathbf{n}*}$ is the adjoint of the Riesz component operator $\mathcal{R}^{\mathbf{n}}$ defined by (4) and where $\psi_{i, \mathbf{k}}$ are the primal, orientation-free wavelets. This means that performing the N th-order Riesz wavelet transform of the signal f is equivalent to applying a sequence of primal wavelet transforms to the (adjoint) Riesz components of f : $\mathcal{R}^{\mathbf{n}*} f$ with $|\mathbf{n}| = N$. In practice, we are assuming that the input signal is bandlimited and we are performing all the filtering operation (Riesz transform and wavelet

TABLE II
ASSESSMENT OF THE EFFECT OF STEERING: SIMONCELLI'S EQUIANGULAR DESIGN VERSUS RIESZ WAVELETS

Equi-Angular						
Order	2	3	4	5	6	7
No steering	25.22	25.33	25.48	25.51	25.58	25.55
Max steering	25.36	25.39	25.51	25.47	25.54	25.51
Tensor steering	25.43	25.43	25.54	25.50	25.58	25.54
Riesz						
Order	2	3	4	5	6	7
No steering	25.25	25.33	25.50	25.49	25.58	25.56
Max steering	25.38	25.42	25.56	25.52	25.59	25.55
Tensor steering	25.44	25.48	25.65	25.61	25.70	25.66

Experimental conditions: cameraman image corrupted by white Gaussian noise ($\sigma = 30$). Initial PSNR = 18.55.

decomposition) in the frequency domain for convenience and greater efficiency. The sequence of operations for the decomposition part of the algorithm is therefore to take the FFT of the input signal, to evaluate its (adjoint) Riesz transform by multiplication in the frequency domain, to compute the wavelet decomposition (by filtering and downsampling) of each component and to get back the wavelet coefficients at each scale by performing a series of inverse FFTs of the appropriate size. The final step is to steer the transform and compute the generalized Riesz wavelet transform by suitable matrix multiplication according to the procedure outlined in Theorem 1. The reconstruction part follows the same steps but in reverse order (flow graph transpose of the decomposition algorithm). We note that the Riesz part of the algorithm and the steering is completely generic, so that the leading part of the effort is the implementation of the primal wavelet decomposition. The reader is referred to Appendix B and Figs. 3 and 4 for the computational aspects of pyramid-based wavelets.

While the description of the isotropic bandlimited wavelets in Section IV-B is self-contained, there is a practical modification that needs to be introduced in the frequency response of the finer scale wavelet filter to ensure exact reconstruction from sampled data. Instead of falling off after $\omega > \pi/2$, the response needs to

TABLE III
COMPARISON OF THE DENOISING PERFORMANCE OF DIFFERENT TYPES OF STEERABLE WAVELET TRANSFORMS

σ	order	Barbara			Lena			zoneplate		
		10	20	30	10	20	30	10	20	30
initial PSNR		28.11	22.10	18.59	28.13	22.11	18.59	28.14	22.12	18.59
Equi-Angular	2	31.43	27.33	25.22	33.63	30.24	28.47	34.03	29.02	26.08
	3	31.57	27.47	25.38	33.68	30.30	28.54	34.56	29.54	26.58
	4	31.72	27.60	25.50	33.77	30.37	28.58	34.92	29.86	26.91
	5	31.81	27.69	25.57	33.76	30.38	28.58	35.08	30.10	27.14
Riesz	2	31.68	27.44	25.29	33.68	30.27	28.47	35.06	29.83	26.47
	3	31.86	27.67	25.48	33.76	30.34	28.53	35.44	30.22	26.78
	4	32.03	27.86	25.66	33.89	30.47	28.64	35.79	30.55	27.16
	5	32.09	27.95	25.74	33.88	30.46	28.63	35.94	30.72	27.49
Equalized Riesz	2	30.85	26.63	24.58	33.14	29.57	27.83	32.79	27.51	24.83
	3	31.09	26.94	24.64	33.25	30.06	28.19	33.18	27.76	24.74
	4	31.02	26.83	24.76	33.28	29.71	27.95	32.95	27.69	25.07
	5	31.06	26.97	24.70	33.37	30.12	28.23	33.19	27.95	24.86
PCA	2	31.65	27.33	25.14	33.59	30.14	28.33	35.00	29.70	26.62
	3	31.75	27.44	25.19	33.58	30.11	28.30	35.27	29.96	26.86
	4	31.86	27.53	25.27	33.64	30.15	28.32	35.46	30.12	27.03
	5	31.87	27.55	25.25	33.59	30.10	28.28	35.54	30.20	27.11
Max SNR Equi-Angular	2	31.72	27.48	25.28	33.63	30.19	28.40	35.09	29.94	26.97
	3	31.85	27.69	25.53	33.72	30.30	28.50	35.68	30.45	27.57
	4	31.93	27.81	25.66	33.80	30.38	28.55	35.93	30.79	27.79
	5	31.96	27.86	25.72	33.78	30.37	28.57	35.78	30.96	27.81
Equalized PCA	2	31.80	27.61	25.41	33.75	30.32	28.50	35.32	30.07	27.03
	3	32.05	27.92	25.70	33.84	30.40	28.56	35.87	30.64	27.58
	4	32.23	28.17	25.94	33.99	30.53	28.67	36.25	31.10	28.01
	5	32.31	28.29	26.04	33.98	30.52	28.66	36.40	31.33	28.29

be maintained to one in order to retain the full information in the upper diagonal parts of the Nyquist region. In essence, this corresponds to a truncation of the frequency tilling in Fig. 1, where the outer frontier of the bandpass region of the signal, which corresponds to V_0 , is a square rather than a circle. Alternatively, one could have chosen to radially bandlimit the input signal prior to wavelet decomposition, which has the disadvantage of throwing part of the information away. We have opted for the first solution, which is in the spirit of the pyramid-based wavelets. The Matlab implementation of the present family of steerable wavelets is available at: <http://bigwww.epfl.ch/demo/steerable-wavelets>.

E. Comparisons of Steerable Transforms

To assess the performance of the various brands of steerable wavelet transforms, we used a basic denoising benchmark with standard test images corrupted by additive white Gaussian noise of variance σ^2 . To highlight the influence of the wavelet transform, we implemented the simplest possible algorithm, which consists in applying a pointwise nonlinearity (soft-threshold) in the steered-wavelet domain prior to reconstruction. In each experiment, the noisy image was decomposed with a fixed number of scales ($I = 4$) and a common optimized threshold was applied in all wavelet channels, while the low-pass components (scaling coefficients) were left unchanged. Unless stated otherwise (cf., Table II), the wavelets were steered according to the procedure described in Section V-A, the idea being to select the 2-D rotation matrix that maximizes the directional Hilbert transform over a local wavelet-domain neighborhood (tensor steering). To account for the fact that the noise does not get distributed evenly across Riesz components, we set the threshold proportional to the theoretical standard deviation

$\sigma_{i,\mathbf{n}} = E\{\langle \psi_{i,\mathbf{k}}^{\mathbf{n}}, \text{noise} \rangle\}^{1/2} = \sigma \|\psi_i^{\mathbf{n}}\|_{L_2}$, where $\psi_i^{\mathbf{n}}$ is the analysis wavelet in channel \mathbf{n} and scale i . Note that the weights are constant by design in the equiangular and equalized Riesz configurations. The proportionality factor (threshold value) was kept fixed across all wavelet channels and optimized for best performance using the noise-free image (oracle) as our reference.

We conducted a series of experiments using a wide range of images, noise levels, and combination of steerable wavelet transforms. Our performance index is the signal-to-noise ratio in decibel with respect to the reference image. By default, the choice of primal wavelet decomposition is the one associated with Simoncelli's steerable pyramid, unless stated otherwise (see Table IV).

1) *Effect of Steering*: We found that wavelet steering systematically improves performance, as documented in Table II. The "Max steering" option refers to the standard procedure, which is to maximize the intensity in the first channel of the filterbank at each wavelet location [1]. In the case of the equiangular filterbank (Simoncelli's steerable pyramid), the gain becomes less as the order increases, which is understandable because the wavelet filters span more and more directions making steering less advantageous. This is not so for the Riesz wavelets, which need to be oriented properly because of their greater shape diversity. The proposed tensor-based steering mechanism based on the eigenvalue decomposition of $\mathbf{J}_{i,\mathbf{k}_0}$ in (25) is consistently superior to the standard strategy (Max steering). This is an interesting finding that also gives a simple way of boosting the performance of Simoncelli's steerable pyramid. It is the steering method that is applied by default for the remainder of the experiments.

2) *Effect of Wavelet Shaping*: The relevant experiments are summarized in Table III. Barbara and Lena are standard




Fig. 7. Evolution of the denoising performance as the order increases. Experimental conditions: 512×512 Barbara image corrupted by white Gaussian noise ($\sigma = 30$). Initial PSNR = 18.60.

“natural” images, while “zoneplate” is synthetic; it is a circular chirp-like wave, whose frequency increases linearly as a function of the distance to the center of the image (cf. [31, Fig 7]). The performance of the canonical Riesz wavelets and the equiangular ones (Simoncelli) are more or less comparable with a slight advantage to the former, especially when the underlying patterns are directional (Barbara, zoneplate). Yet, getting there was not immediate and required use of norm-adapted thresholds, as documented earlier. The worst results are obtained using equalized Riesz wavelets; we suspect that the primary reason is their lack of angular selectivity, as is apparent in Fig. 5. To allow for a fair comparison, the PCA wavelets were trained on the noisy image to which they were applied. The straight (unequalized) Riesz PCA wavelets did not offer any significant advantage over the basic ones. This is not so in the case of the equalized Riesz transform, where the situation is reversed dramatically. Here, the gain of PCA is always quite substantial (a boost of several decibels over equalized Riesz) and the denoising results are consistently the best, both visually and quantitatively, over the complete set of experiments. One can also observe that learning is most beneficial when the number of channels increases (more degrees of freedom) and/or when the noise decreases, in which case the training is more accurate. The case of “equiangular Max SNR” is interesting because it demonstrates that it is possible to outperform Simoncelli’s classical design (as well as Riesz wavelets) by sticking to an equiangular configuration using a basis filter that is optimized for maximum signal-to-noise ratio (i.e., the first component filter of equalized PCA). The ranking of the various transforms does not appear to be affected by the actual noise level; it could also be replicated with the many other images on which we tested the algorithms. We observed that the extent of performance improvement of signal-adapted wavelets (equalized PCA) was dependent upon the type of image and generally better in the presence of texture or patterned areas (e.g., the pants of Barbara or zoneplate). It is less significant with Lena for which Simoncelli’s steerable pyramid performs remarkably well, especially at higher noise levels.

3) *Influence of Order:* For this series of experiments as well as the last one, we restricted our comparison to our most promising solution (“equalized PCA”) and the classical one (Simoncelli’s equiangular design). The general trend is that the performance improves in both cases as the order increases with a tendency to saturation above $N = 10$ (cf., Fig. 7). Also, the

benefit of PCA is most significant at intermediate orders [gain of 0.5 dB for Barbara, and +0.25 dB for Lena (data not show)], most likely because there are enough degrees of freedom to exploit.

4) *Influence of Primal Multiresolution Decomposition:* Among the category of pyramid-based wavelets (cf., Section IV-C), we investigated the class of tight frames associated with the polynomial spline of degree n . The underlying scaling functions are the orthogonalized version of the B-splines (Battle-Lemarié functions), which have played a significant role in the construction of the very first wavelet basis of $L_2(\mathbb{R})$, as well as $L_2(\mathbb{R}^2)$ using tensor products [34]. To quantify the degree of isotropy of these functions in 2-D, we used the performance index

$$\rho = \frac{1}{\pi} \int_0^\pi \frac{\langle \varphi, \varphi(\mathbf{R}_\theta \cdot) \rangle_{L_2(\mathbb{R}^2)}}{\|\varphi\|_{L_2(\mathbb{R}^2)}^2} d\theta$$

where \mathbf{R}_θ is the 2×2 spatial rotation matrix with angle θ . Clearly, $\varphi(\mathbf{R}_\theta \mathbf{x}) = \varphi(\mathbf{x})$ for all θ if and only if φ is perfectly isotropic, which corresponds to a value of $\rho = 1$; otherwise, $\rho < 1$ as a consequence of the Cauchy–Schwarz inequality. Note that the criterion may also be evaluated in the frequency domain (using Parseval’s formula), which turns out to be especially advantageous for splines because of the availability of closed-form formulas in the Fourier domain. This allowed us to identify the case $n = 1$ (piecewise-linear splines) with $\rho = 0.95$ as the most isotropic one in this category. For reference, we are providing the Fourier-domain formula of the corresponding orthogonal scaling function: $\hat{\varphi}_{\text{lin}}(\omega_1, \omega_2) = \hat{\phi}_1(\omega_1)\hat{\phi}_1(\omega_2)$ with

$$\hat{\phi}_1(\omega) = \frac{\text{sinc}(\omega/2\pi)^2}{\sqrt{2/3 + (1/3)\cos\omega}}$$

In addition to this, we undertook a systematic evaluation of the bandlimited wavelets that are specified in Table I and which have the advantage over the former to be truly isotropic $\rho = 1$. The results of these comparisons are summarized in Table IV.

The key observation is that the performance of the various transforms are quite comparable with a slight preference to linear splines and Papadakis. The good performance of the Papadakis and Simoncelli radial filters did subsist over a wide variety of images, while this was not always the case for linear splines, probably due to the fact that they are not perfectly isotropic. On the other hand, we found the Shannon wavelets, which also have the poorest spatial decay, to be consistently the worst. While the slight superiority of Papadakis is suggestive of the fact that there may still be room for improvement in this area, it is probably not significant enough to justify replacing the Simoncelli filters, which are very popular in applications.

F. Discussion

For our performance assessment, we have intentionally kept the processing simple and reproducible (oracle-based soft-thresholding) to really focus on the effect of the wavelet transform and to separate out other factors. The clear message that emerges is that not all wavelets spanning the same steerable subspace are equal: the equalized Riesz wavelets, which essentially have no directionality, are by far the worst, while the

TABLE IV
ASSESSMENT OF THE EFFECT OF THE PRIMAL DECOMPOSITION

Barbara							
method	order	Simoncelli	Meyer	Shannon	Papadakis	Held	Spline
Equalized PCA	2	25.40	25.37	25.30	25.45	25.41	25.32
Equi-Angular	2	25.20	25.18	25.13	25.26	25.22	25.23
Equalized PCA	3	25.68	25.65	25.57	25.73	25.69	25.77
Equi-Angular	3	25.36	25.33	25.29	25.42	25.38	25.55
Equalized PCA	4	25.90	25.87	25.77	25.96	25.91	25.93
Equi-Angular	4	25.47	25.45	25.40	25.53	25.48	25.67
Equalized PCA	5	26.01	25.98	25.89	26.07	26.02	26.14
Equi-Angular	5	25.55	25.52	25.48	25.61	25.56	25.85
Equalized PCA	6	26.15	26.12	26.02	26.21	26.16	26.21
Equi-Angular	6	25.63	25.60	25.55	25.69	25.64	25.93
Equalized PCA	7	26.19	26.16	26.08	26.25	26.21	26.34
Equi-Angular	7	25.68	25.66	25.60	25.75	25.70	26.04

Experimental conditions: 512×512 Barbara image corrupted by white Gaussian noise ($\sigma = 30$). Initial PSNR = 18.60.

equalized PCA ones are systematically the best. Simoncelli's equiangular design, which has good angular selectivity, generally performs well, especially on piecewise-smooth images, such as Lena. Yet, it is not necessarily the best solution, even for an equiangular configuration (e.g., max SNR solution).

While these findings are very promising, more work is required to demonstrate the ability of these transforms to improve upon the state-of-the-art in wavelet-based image processing. There is also good hope that the proposed methodology can be pushed further by addressing the following issues.

- 1) *Refinement of the denoising algorithm*: Here, it is relatively easy to improve upon soft-thresholding by applying more sophisticated methods, such as BLS-GSM [10] or SURELET [40], which presently constitute the state-of-the-art in wavelet-based denoising.
- 2) *Improved angular wavelet design/learning algorithms*: While the results obtained using PCA are encouraging, one can easily conceive of applying more sophisticated algorithms, such as independent component analysis [41] or sparse PCA [42]. The challenges there are twofold: 1) defining suitable task-dependent performance criteria; and 2) introducing a two-way coupling between learning and steering, since the latter greatly influences the former while good learning could probably improve steering.
- 3) *Optimization of radial frequency profile*: The search for a better primal decomposition is still an open research issue. The difficulty is that it not yet clear what the important design parameters are: spectral shaping, space-frequency localization, or asymptotic rate of decay of the wavelet in space.

VI. CONCLUSION

We have presented a general framework for the design of tight steerable wavelets frames in multiple dimensions. The approach allows for a separate handling of the multiresolution and rotational aspects of the transform. The main extension over Simoncelli's steerable pyramid is that the mother wavelets are not necessarily rotated versions of one another, which provides greater design flexibility. The convenient feature of the proposed decomposition is that the shape of the M steerable wavelets at a given scale i is parametrized by an $M \times M$ unitary matrix \mathbf{U}_i

that can be freely selected, without impacting upon the perfect reconstruction (tight frame) properties.

The practical interest is that these generalized steerable transforms admit a fast decomposition algorithm and that their multiresolution basis functions can be optimized using standard learning algorithms, such as PCA. Our results suggest that the combination of steerability and signal-adaptivity holds good promises for improving wavelet-based image processing. In particular, we have introduced a novel family of steerable "equalized PCA" wavelets that systematically outperforms the traditional steerable pyramid in a basic denoising benchmark. Due to their tunability and intrinsic rotation-covariance, these new wavelet transforms may be useful for developing better algorithms for denoising, image filtering/enhancement, texture analysis/synthesis, and solving inverse problems under sparsity constraints.

APPENDIX A PROOF OF THEOREM 3

To prove the result, we first write the dual counterpart of (17) for any function $f \in V$, and more specifically for $g \in W$

$$\forall g \in W, \quad g = \sum_k \langle \varphi_k, g \rangle \tilde{\varphi}_k$$

where we are using the property that the dual basis $\{\tilde{\varphi}_k\}$ is necessarily a Riesz basis of V as well. The next observation is that $\psi_k = \varphi_k - \text{Proj}_{W^\perp} \varphi_k$, where W^\perp is the orthogonal complement of W with respect to V . This implies that

$$\forall g \in W, \quad \langle \psi_k, g \rangle = \langle \varphi_k, g \rangle - \underbrace{\langle \text{Proj}_{W^\perp} \varphi_k, g \rangle}_{=0} \quad (26)$$

because $g \perp W^\perp$. Moreover, we obviously have that

$$\forall g \in W, \quad g = \text{Proj}_W g = \sum_k \langle \psi_k, g \rangle \underbrace{\text{Proj}_W \tilde{\varphi}_k}_{=\tilde{\psi}_k}$$

which proves the first decomposition/reconstruction part of the theorem; the dual version is obtained simply by interchanging the role of $\tilde{\varphi}_k$ and φ_k . As for the bounds, we recall that a Riesz

basis and its dual are characterized by the following norm equivalences:

$$\begin{aligned} A \sum_k |c_k|^2 &\leq \left\| \sum_k c_k \varphi_k \right\|^2 \leq B \sum_k |c_k|^2 \\ \frac{1}{B} \sum_k |\tilde{c}_k|^2 &\leq \left\| \sum_k \tilde{c}_k \tilde{\varphi}_k \right\|^2 \leq \frac{1}{A} \sum_k |\tilde{c}_k|^2 \end{aligned}$$

with $f = \sum_k c_k \varphi_k = \sum_k \tilde{c}_k \tilde{\varphi}_k$ in accordance with (17) and its dual equivalent. The constants $0 < A$ and $B < \infty$ are the so-called Riesz bound of φ_k . By making the substitutions $\tilde{c}_k = \langle \varphi_k, f \rangle$ and $c_k = \langle \tilde{\varphi}_k, f \rangle$, we derive the two frame inequalities

$$\begin{aligned} \forall f \in V, \quad A \|f\|^2 &\leq \sum_k |\langle \varphi_k, f \rangle|^2 \leq B \|f\|^2 \\ \forall f \in V, \quad \frac{1}{B} \|f\|^2 &\leq \sum_k |\langle \tilde{\varphi}_k, f \rangle|^2 \leq \frac{1}{A} \|f\|^2 \end{aligned} \quad (27)$$

which reflect the property that a Riesz basis is a particular case of a frame. Finally, by writing the specialized version of (27) for $g \in W$ and using (26), we obtain

$$\forall g \in W, \quad A \|g\|^2 \leq \sum_k |\langle \psi_k, g \rangle|^2 \leq B \|g\|^2.$$

APPENDIX B

IMPLEMENTATION OF PYRAMID-BASED TRANSFORMS

To derive the corresponding digital-filtering algorithm, we first specify the dual function $\tilde{\varphi} \in V_0$, whose Fourier-domain expression is as follows:

$$\hat{\tilde{\varphi}}(\omega) = \frac{\hat{\varphi}(\omega)}{A_\varphi(e^{j\omega})}$$

where the denominator represents to the so-called autocorrelation filter

$$\begin{aligned} A_\varphi(e^{j\omega}) &= \sum_{\mathbf{k} \in \mathbb{Z}^d} \langle \varphi, \varphi(\cdot - \mathbf{k}) \rangle_{L_2} e^{-j(\omega, \mathbf{k})} \\ &= \sum_{\mathbf{n} \in \mathbb{Z}^d} |\hat{\varphi}(\omega + 2\pi \mathbf{n})|^2 \end{aligned}$$

which is bounded away from zero as a consequence of the Riesz basis condition 1). More generally, the fact that φ is a valid scaling function implies that the same must be true for $\tilde{\varphi}$. In particular, $\tilde{\varphi}$ satisfies the two-scale relation

$$\tilde{\varphi}(\mathbf{x}/2) = 2^{d/2} \sum_{\mathbf{k} \in \mathbb{Z}^d} \tilde{h}[\mathbf{k}] \tilde{\varphi}(\mathbf{x} - \mathbf{k}) \quad (28)$$

where the dual refinement filter \tilde{h} is conveniently described by its d -dimensional z -transform

$$\tilde{H}(z) = H(z) \frac{A_\varphi(z)}{A_\varphi(z^2)}$$

this is easily established using (18) and (28). By considering the dual basis functions $\tilde{\varphi}_{i, \mathbf{k}}(\mathbf{x}) = 2^{-id/2} \tilde{\varphi}(\mathbf{x}/2^i - \mathbf{k})$, we can

now explicitly specify the orthogonal projection of a function $f \in L_2(\mathbb{R}^d)$ into V_i

$$f_i = \text{Proj}_{V_i} f = \sum_{\mathbf{k} \in \mathbb{Z}^d} c_i[\mathbf{k}] \tilde{\varphi}_{i, \mathbf{k}}$$

where $c_i[\mathbf{k}] = \langle \varphi_{i, \mathbf{k}}, f \rangle$. By exploiting the property that the multiresolution spaces are nested (i.e., $V_i \subset V_{i-1}$), one can compute the expansion coefficients iteratively by filtering and downsampling by a factor of two [cf., reduce module in Fig. 3(a)]

$$c_i[\mathbf{k}] = \langle \varphi_{i, \mathbf{k}}, f \rangle = \langle \varphi_{i, \mathbf{k}}, \text{Proj}_{V_{i-1}} f \rangle = (h * c_{i-1})[2\mathbf{k}].$$

This algorithm needs to be initialized by specifying the finer scale coefficients $c_0[\mathbf{k}]$. In practice, these are obtained by interpolating the discrete input samples $f[\mathbf{k}]$ using some appropriate interpolation kernel [e.g., $\varphi_{\text{int}}(\mathbf{x}) = \text{sinc}(\mathbf{x})$] and by projecting the resulting continuous-space signal $f(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} f[\mathbf{k}] \varphi_{\text{int}}(\mathbf{x} - \mathbf{k})$ onto V_0 . This process is implemented efficiently by prefiltering the input data with a digital filter, whose impulse response is $h_0[\mathbf{k}] = \langle \varphi_{\text{int}}, \varphi(\mathbf{k} - \cdot) \rangle$.

The next step is the expansion part of the algorithm that uses the two-scale relation (28) to compute the representation of f_i at the finer scale $i - 1$

$$f_i = \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{c}_{i-1}[\mathbf{k}] \varphi_{i-1, \mathbf{k}}$$

where

$$\hat{c}_{i-1}[\mathbf{k}] = \sum_{\mathbf{k}' \in \mathbb{Z}^d} h[\mathbf{k} - 2\mathbf{k}'] c_i[\mathbf{k}']$$

which amounts to an upsampling by 2 followed by a postfiltering with the refinement filter \tilde{h} [cf., Fig. 3(b)].

Finally, we obtain the residual (or wavelet) signals $f_i - f_{i+1} \in V_i$ by subtracting the two representations

$$r_{i+1} = f_i - f_{i+1} = \sum_{\mathbf{k} \in \mathbb{Z}^d} (c_i[\mathbf{k}] - \hat{c}_i[\mathbf{k}]) \tilde{\varphi}_{i, \mathbf{k}} = \sum_{\mathbf{k} \in \mathbb{Z}^d} w_i[\mathbf{k}] \tilde{\psi}_{i, \mathbf{k}}, \quad (29)$$

from which we conclude that

$$w_i[\mathbf{k}] = \langle \psi_{i, \mathbf{k}}, f \rangle = c_i[\mathbf{k}] - \hat{c}_i[\mathbf{k}]$$

which corresponds to the block diagram in Fig. 4(a).

The reconstruction procedure, which is based on the right-hand side of (29), follows a dual path and corresponds to the flow graph transpose of the decomposition algorithm. We skip the derivation and refer the reader to the full algorithm description in Fig. 4.

Although the wavelet component representations $r_{i+1} = f_i - f_{i+1} = \sum_{\mathbf{k} \in \mathbb{Z}^d} w_i[\mathbf{k}] \tilde{\varphi}_{i, \mathbf{k}}$ and $r_{i+1} = \sum_{\mathbf{k} \in \mathbb{Z}^d} w_i[\mathbf{k}] \tilde{\psi}_{i, \mathbf{k}}$ are mathematically equivalent, it is interesting to note that they do not necessarily lead to the same reconstruction algorithm. In the first case, we may apply the original method of Burt and Adelson [23], and reconstruct $f_0 = \sum_{i=1}^{i_{\text{max}}} r_i$ by adding up the residual progressively from coarse-to-fine and expanding

the intermediate results to the next-finer scale using the expand module. In effect, this amounts for removing the reduce block in the synthesis procedure in Fig. 4(b). While this is more efficient algorithmically, it is less favorable when the wavelet coefficients $w_i[\mathbf{k}]$ are altered through processing. The fundamental difference between the two representations is that $\sum_{\mathbf{k} \in Z^d} \tilde{w}_i[\mathbf{k}] \tilde{\psi}_{i,\mathbf{k}}$, where $\tilde{w}_i[\mathbf{k}]$ are arbitrary coefficients is always included in W_{i-1} , while it is not necessarily so when the wavelets are substituted by the basis functions $\tilde{\varphi}_{i,\mathbf{k}}$. In other words, using the wavelet reconstruction functions implicitly reprojects the subband component onto W_{i+1} , leading to a reconstruction that is more robust to noise. Mathematically, this corresponds to using the dual frame reconstruction operator. This is an idea that was first proposed by Do and Vetterli and analyzed by taking a perfect-reconstruction-filterbank point of view, which offers some further insights; in particular, the polyphase-domain analysis [26].

ACKNOWLEDGMENT

The authors would like to thank I. Bayram and P. Bhatia for their participation in the project.

REFERENCES

- [1] E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, "Shiftable multiscale transforms," *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 587–607, Mar. 1992.
- [2] H. Greenspan, S. Belongie, R. Goodman, P. Perona, S. Rakshit, and C. Anderson, "Overcomplete steerable pyramid filters and rotation invariance," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, June 21–23, 1994, pp. 222–228.
- [3] A. Karasiris and E. Simoncelli, "A filter design technique for steerable pyramid image transforms," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, May 7–10, 1996, vol. 4, pp. 2387–2390.
- [4] E. Simoncelli and W. Freeman, "The steerable pyramid: A flexible architecture for multi-scale derivative computation," in *Proc. Int. Conf. Image Process.*, Oct. 23–26, 1995, vol. 3, pp. 444–447.
- [5] W. T. Freeman and E. H. Adelson, "The design and use of steerable filters," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 9, pp. 891–906, Sep. 1991.
- [6] M. Do and M. Vetterli, "Rotation invariant texture characterization and retrieval using steerable wavelet-domain hidden Markov models," *IEEE Trans. Multimedia*, vol. 4, no. 4, pp. 517–527, Dec. 2002.
- [7] G. Tzagkarakis, B. Beferull-Lozano, and P. Tsakalides, "Rotation-invariant texture retrieval with gaussianized steerable pyramids," *IEEE Trans. Image Process.*, vol. 15, no. 9, pp. 2702–2718, Sep. 2006.
- [8] F. Denis and A. Baskurt, "Multidirectional curvilinear structures detection using steerable pyramid," *J. Electron. Imag.*, vol. 13, no. 4, pp. 756–765, 2004.
- [9] S. Mallat, *A Wavelet Tour of Signal Processing*. San Diego, CA: Academic, 1998.
- [10] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, "Image denoising using scale mixtures of Gaussians in the wavelet domain," *IEEE Trans. Image Process.*, vol. 12, no. 11, pp. 1338–1351, Nov. 2003.
- [11] A. Laine, J. A. Fan, and W. H. Yang, "Wavelets for contrast enhancement of digital mammography," *IEEE Eng. Med. Biol. Mag.*, vol. 14, no. 5, pp. 536–550, Sep./Oct. 1995.
- [12] A. Laine, A. Laine, S. Schuler, J. Fan, and W. Huda, "Mammographic feature enhancement by multiscale analysis," *IEEE Trans. Med. Imag.*, vol. 13, no. 4, pp. 725–740, Dec. 1994.
- [13] J. Portilla and E. P. Simoncelli, "A parametric texture model based on joint statistics of complex wavelet coefficients," *Int. J. Comput. Vis.*, vol. 40, no. 1, pp. 49–70, 2000.
- [14] C. Y. Su, Y. T. Zhuang, L. Huang, and F. Wu, "Steerable pyramid-based face hallucination," *Pattern Recognit.*, vol. 38, no. 6, pp. 813–824, 2005.
- [15] Z. Liu, K. Tsukada, K. Hanasaki, Y. K. Ho, and Y. P. Dai, "Image fusion by using steerable pyramid," *Pattern Recognit. Lett.*, vol. 22, no. 9, pp. 929–939, 2001.
- [16] I. Daubechies, M. Defrise, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Commun. Pure Appl. Math.*, vol. 57, no. 11, pp. 1413–1457, 2004.
- [17] I. Daubechies, G. Teschke, and L. Vese, "Iteratively solving linear inverse problems under general convex constraints," *Inverse Problems Imag.*, vol. 1, no. 1, pp. 29–46, 2007.
- [18] H. Rabbani, "Image denoising in steerable pyramid domain based on a local Laplace prior," *Pattern Recognit.*, vol. 42, no. 9, pp. 2181–2193, 2009.
- [19] M. Unser and D. Van De Ville, "Wavelet steerability and the higher-order Riesz transforms," *IEEE Trans. Image Process.*, vol. 19, no. 3, pp. 636–652, Mar. 2010.
- [20] J. Kovacevic and M. Vetterli, "Nonseparable 2-dimensional and 3-dimensional wavelets," *IEEE Trans. Signal Process.*, vol. 43, no. 5, pp. 1269–1273, May 1995.
- [21] M. Feilner, D. Van De Ville, and M. Unser, "An orthogonal family of quincunx wavelets with continuously adjustable order," *IEEE Trans. Image Process.*, vol. 14, no. 4, pp. 499–510, Apr. 2005.
- [22] D. Van De Ville, T. Blu, and M. Unser, "On the multidimensional extension of the quincunx subsampling matrix," *IEEE Signal Process. Lett.*, vol. 12, no. 2, pp. 112–115, Feb. 2005.
- [23] P. Burt and E. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Trans. Commun.*, vol. 31, no. 4, pp. 532–540, Apr. 1983.
- [24] M. Unser, "An improved least squares Laplacian pyramid for image compression," *Signal Process.*, vol. 27, no. 2, pp. 187–203, May 1992.
- [25] M. Unser, A. Aldroubi, and M. Eden, "The L_2 -polynomial spline pyramid," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 15, no. 4, pp. 364–379, Apr. 1993.
- [26] M. Do and M. Vetterli, "Framing pyramids," *IEEE Trans. Signal Process.*, vol. 51, no. 9, pp. 2329–2342, Sep. 2003.
- [27] M. Papadakis, G. Gogoshin, I. A. Kakadiaris, D. J. Kouri, and D. K. Hoffman, "Nonseparable radial frame multiresolution analysis in multidimensions," *Numer. Funct. Anal. Optim.*, vol. 24, no. 7–8, pp. 907–928, 2003.
- [28] J. R. Romero, S. K. Alexander, S. Baid, S. Jain, and M. Papadakis, "The geometry and the analytic properties of isotropic multiresolution analysis," *Adv. Comput. Math.*, vol. 31, no. 1–3, pp. 283–328, 2009.
- [29] S. Held, M. Storch, P. Massopust, and B. Forster, "Steerable wavelet frames based on the Riesz transform," *IEEE Trans. Image Process.*, vol. 19, no. 3, pp. 653–667, Mar. 2010.
- [30] D. Van De Ville, T. Blu, and M. Unser, "Isotropic polyharmonic B-splines: Scaling functions and wavelets," *IEEE Trans. Image Process.*, vol. 14, no. 11, pp. 1798–1813, Nov. 2005.
- [31] M. Unser, D. Sage, and D. Van De Ville, "Multiresolution monogenic signal analysis using the Riesz-Laplace wavelet transforms," *IEEE Trans. Image Process.*, vol. 18, no. 11, pp. 2402–2418, Nov. 2009.
- [32] R. J. Duffin and A. C. Schaeffer, "A class of nonharmonic Fourier series," *Trans. Amer. Math. Soc.*, vol. 72, no. 2, pp. 341–366, Mar. 1952.
- [33] O. Christensen, *An Introduction to Frames and Riesz Bases*. Cambridge, MA: Birkhauser, 2003.
- [34] S. G. Mallat, "A theory of multiresolution signal decomposition: The wavelet representation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 11, no. 7, pp. 674–693, Jul. 1989.
- [35] M. Unser and T. Blu, "Wavelet theory demystified," *IEEE Trans. Signal Process.*, vol. 51, no. 2, pp. 470–483, Feb. 2003.
- [36] A. Aldroubi and M. Unser, "Sampling procedures in function spaces and asymptotic equivalence with Shannon's sampling theory," *Numer. Funct. Anal. Optim.*, vol. 15, no. 1–2, pp. 1–21, 1994.
- [37] I. Daubechies, A. Grossmann, and Y. Meyer, "Painless nonorthogonal expansions," *J. Math. Phys.*, vol. 27, no. 5, pp. 1271–1283, 1986.
- [38] I. Daubechies, *Ten Lectures on Wavelets*. Philadelphia, PA: Soc. Ind. Appl. Math., 1992.
- [39] D. J. Field, "Relations between the statistics of natural images and the response properties of cortical cells," *J. Opt. Soc. Amer. A-Opt. Image Sci. Vis.*, vol. 4, no. 12, pp. 2379–2394, 1987.
- [40] T. Blu and F. Luisier, "The SURE-LET approach to image denoising," *IEEE Trans. Image Process.*, vol. 16, no. 11, pp. 2778–2786, Nov. 2007.
- [41] P. Comon, "Independent component analysis: A new concept," *Signal Process.*, vol. 36, no. 3, pp. 287–314, 1994.
- [42] B. A. Olshausen and D. J. Field, "Natural image statistics and efficient coding," *Netw.-Comput. Neural Syst.*, vol. 7, no. 2, pp. 333–339, 1996.



Michael Unser (M'89–SM'94–F'99) received the M.S. (*summa cum laude*) and Ph.D. degrees in electrical engineering from the École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, in 1981 and 1984, respectively.

From 1985 to 1997, he was a Scientist with the National Institutes of Health, Bethesda, MD. He is currently a Full Professor and the Director of the Biomedical Imaging Group, EPFL. He is the author or coauthor of 200 journal papers of various topics, and is one of ISI's highly cited authors in engineering

(<http://isihighlycited.com>). His research interests include biomedical image processing, sampling theories, multiresolution algorithms, wavelets, and the use of splines for image processing.

Dr. Unser was an Associate Editor-in-Chief for the IEEE TRANSACTIONS ON MEDICAL IMAGING (2003–2005) and an Associate Editor for the same journal (1999–2002 and 2006–2007), the IEEE TRANSACTIONS ON IMAGE PROCESSING (1992–1995), and the IEEE SIGNAL PROCESSING LETTERS (1994–1998). He is currently a member of the editorial boards of *Foundations and Trends in Signal Processing* and *Sampling Theory in Signal and Image Processing*. He was the Coorganizer of the first IEEE International Symposium on Biomedical Imaging (2002) and was the founding chair of the technical committee of the IEEE-SP Society on Bio Imaging and Signal Processing. He was the recipient of the 1995 and 2003 Best Paper Awards, the 2000 Magazine Award, and the 2008 Technical Achievement Award from the IEEE Signal Processing Society. He is an European Association for Signal Processing (EURASIP) Fellow and a member of the Swiss Academy of Engineering Sciences.



Nicolas Chenouard (M'08) received the M.S. degree in engineering and bioinformatics from the Institut National des Sciences Appliquées, Lyon, France, and the M.Sc. degree in artificial intelligence and data mining from the University Pierre et Marie Curie (Paris VI), Paris, France, both in 2006, and the Ph.D. degree in image and signal processing from the Institut Pasteur and Telecom ParisTech, Paris, in 2010.

Since 2010, he has been a Postdoctoral Fellow with the Biomedical Imaging Group and the Center for Biomedical Imaging, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland. His current research interests include steerable frames, frame learning, 3-D image processing, and bioimaging applications.



Dimitri Van De Ville (M'02) received the M.S. degree in engineering and computer sciences, and the Ph.D. degree from Ghent University, Ghent, Belgium, in 1998, and 2002, respectively.

In 2002, he joined Prof. M. Unser's Biomedical Imaging Group, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland. In December 2005, he became responsible for the Signal Processing Unit, University Hospital of Geneva, Geneva, Switzerland, as a part of the Centre d'Imagerie Biomédicale. In 2009, he received a

Swiss National Science Foundation professorship and currently holds a joint position at the University of Geneva, Geneva, and the EPFL. His research interests include wavelets, sparsity, pattern recognition, and their applications in biomedical imaging, such as functional magnetic resonance imaging.

Dr. Van De Ville was an Associate Editor for the IEEE TRANSACTIONS ON IMAGE PROCESSING from 2006 to 2009, and the IEEE SIGNAL PROCESSING LETTERS from 2004 to 2006. He is the Vice-Chair of the Bio Imaging and Signal Processing Technical Committee of the IEEE Signal Processing Society. Since 2003, he has also been an Editor and Webmaster of *The Wavelet Digest*. He is a Guest Coeditor of the 2011 Special Issue on Brain Decoding in Pattern Recognition. He is Co-Chair of the 2007, 2009, and 2011 Wavelets series conferences, together with V. Goyal and M. Papadakis. He was also the recipient of a Fellowship from the Fund for Scientific Research Flanders Belgium.