

A Unifying Parametric Framework for 2D Steerable Wavelet Transforms*

Michael Unser[†] and Nicolas Chenouard[†]

Abstract. We introduce a complete parameterization of the family of two-dimensional steerable wavelets that are polar-separable in the Fourier domain under the constraint of self-reversibility. These wavelets are constructed by multiorder generalized Riesz transformation of a primary isotropic bandpass pyramid. The backbone of the transform (pyramid) is characterized by a radial frequency profile function $h(\omega)$, while the directional wavelet components at each scale are encoded by an $M \times (2N + 1)$ shaping matrix \mathbf{U} , where M is the number of wavelet channels and N the order of the Riesz transform. We provide general conditions on $h(\omega)$ and \mathbf{U} for the underlying wavelet system to form a tight frame of $L_2(\mathbb{R}^2)$ (with a redundancy factor $4/3M$). The proposed framework ensures that the wavelets are steerable and provides new degrees of freedom (shaping matrix \mathbf{U}) that can be exploited for designing specific wavelet systems. It encompasses many known transforms as particular cases: Simoncelli's steerable pyramid, Marr gradient and Hessian wavelets, monogenic wavelets, and N th-order Riesz and circular harmonic wavelets. We take advantage of the framework to construct new generalized spheroidal prolate wavelets, whose angular selectivity is maximized, as well as signal-adapted detectors based on principal component analysis. We also introduce a curvelet-like steerable wavelet system. Finally, we illustrate the advantages of some of the designs for signal denoising, feature extraction, pattern analysis, and source separation.

Key words. wavelets, steerability, Riesz transform, frame

AMS subject classifications. 68U10, 42C40, 42C15, 47B06

DOI. 10.1137/120866014

1. Introduction. Scale and directionality are essential ingredients for visual perception and processing. This has prompted researchers in image processing and applied mathematics to develop representation schemes and function dictionaries that are capable of extracting and quantifying this type of information explicitly.

The fundamental operation underlying the notion of scale is *dilation*, which calls for a wavelet-type dictionary in which dilated sets of basis functions that “live” at different scales coexist. The elegant aspect here is that it is possible to specify wavelet bases of $L_2(\mathbb{R}^2)$ that result in an orthogonal decomposition of an image in terms of its multiresolution components [1]. The multiscale analysis achieved by these “classical” wavelets is well suited for extracting isotropic image features and isolated singularities, but is not quite as efficient for sparsely encoding edges and curvilinear structures.

To capture directionality, the basis functions need to be angularly selective and the original dilation scheme complemented with spatial *rotation*. This can be achieved at the expense of

*Received by the editors February 14, 2012; accepted for publication (in revised form) October 1, 2012; published electronically January 29, 2013. This work was funded by the Center for Biomedical Imaging, the Foundations Leenaards and Louis-Jeannet, and ERC grant ERC-2010-AdG 267439-FUN-SP.

<http://www.siam.org/journals/siims/6-1/86601.html>

[†]Biomedical Imaging Group (BIG), École Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland (michael.unser@epfl.ch, nicolas.chenouard@epfl.ch).

some redundancy, which rules out the use of a basis. The next best option is a tight frame representation, which offers the same type of functionality (i.e., a self-reversible transformation). The first successful example of multiscale, multiorientation decomposition is the steerable pyramid constructed by Simoncelli and coworkers [2, 3]. Remarkably, this transform has the ability to provide a signal representation that is perfectly rotation-invariant; specifically, the wavelets can be *steered* adaptively—that is, rotated along some dominant local orientation $\theta \in [-\pi, \pi]$ —by forming an appropriate linear combination of the M directional filters at the same location. The steerable pyramid has been used extensively in image processing; typical applications include image denoising [4], contour detection [5], texture analysis and synthesis [6, 7], image extrapolation [8], image fusion [9], and solution of inverse problems under sparsity constraints [10, 11]. The directional wavelets in Simoncelli’s design are rotated versions of a single template in an equiangular configuration, but other options for specifying steerable decompositions are available as well. One extreme solution is given by the circular harmonic wavelets which have the remarkable property of being self-steerable [12], but which have no angular selectivity at all. There is also an intimate link between steerability and the Riesz transform, which has been exploited by several teams [13, 14], and which can result in wavelet designs with greater shape diversity [15].

While steerability is attractive conceptually, it is not a strict requirement. Other solutions to the problem of directional multiresolution image representation include the two-dimensional (2D) Gabor transform [16], curvelets [17], contourlets [18], directionlets [19], and shearlets [20]. The key idea behind curvelets, for instance, is to bring (approximate) rotation and dilation invariance by building a set of basis functions from a series of rotated and dilated versions of an anisotropic mother wavelet. Contourlets basically reproduce the same frequency partitioning, but are based on a tree-structured filterbank. The latter is fast and enjoys a greater flexibility, including different subsampling rates [21]. Shearlets exploit the property that the effect of a shear is analogous to that of a rotation; since the former operation is better suited to a discrete grid, there is a natural link between the discrete and continuous versions of the transform which is much harder to obtain for curvelets [22].

The purpose of this work is to present a unifying framework that ties most of these directional analysis methods together, while providing a universal parameterization of steerable wavelet frames of $L_2(\mathbb{R}^2)$. The key idea is to exploit a functional link between a complex version of the Riesz transform, which dates back to the work of Larkin in optics [23], and the polar-phase components of the circular harmonics. In what follows, we shall investigate the following list of topics, which also summarizes the scope of the present contribution:

- The characterization of the functional properties of the complex Riesz transform and its iterates, including the derivation of their impulse responses (section 2).
- The definition of a one-to-many multiorder generalized Riesz transform that is parameterized by a shaping matrix \mathbf{U} and the derivation of the necessary and sufficient conditions for the mapping to be energy-preserving and self-invertible (section 3). Note that this transform is designed specifically to map an isotropic bandpass pyramid into a bona fide steerable wavelet system.
- The specification of generalized steerable wavelets and the investigation of their mathematical properties, including componentwise orthogonality (section 4). We shall see that the

shaping matrix translates into a Fourier-domain characterization of wavelets as weighted sums of circular harmonics.

- The description of some popular wavelet systems (gradient, Simoncelli, monogenic, circular harmonic, etc.) in the proposed parameterization, as well as the specification of new steerable wavelet families such as the prolate spheroidal wavelets and the principal component analysis (PCA) wavelets which are tuned to a particular class of signals (section 5).
- The demonstration of the practical usefulness of the framework for the design, implementation, and optimization of wavelet systems for specific tasks such as image denoising, curvelet-like image analysis, source separation, and texture discrimination (section 6).

2. Mathematical preliminaries.

2.1. The 2D complex Riesz transform. We will rely heavily on a “complexified” version of the Riesz transform that combines the usual x and y components of the 2D Riesz transform into a single complex signal; it was introduced by Larkin in optics under the name of the *spiral phase quadrature* transform [23, 24] and used in our prior work to define the monogenic extension of a wavelet transform [25].

In this work, we adopt the definition given by Larkin, which differs from that used in our prior paper by a factor of $i = \sqrt{-1}$. The Fourier-domain definition of the transform is

$$(2.1) \quad \mathcal{R}f(\mathbf{x}) \xleftrightarrow{\mathcal{F}} \frac{(\omega_x + i\omega_y)}{\|\boldsymbol{\omega}\|} \hat{f}(\boldsymbol{\omega}) = e^{i\theta} \hat{f}_{\text{pol}}(\omega, \theta),$$

where $\hat{f}(\boldsymbol{\omega}) = \int_{\mathbb{R}^2} f(\mathbf{x}) e^{-i\langle \boldsymbol{\omega}, \mathbf{x} \rangle} dx dy$ with $\boldsymbol{\omega} = (\omega_x, \omega_y)$ and $\hat{f}_{\text{pol}}(\omega, \theta) = \hat{f}(\omega \cos \theta, \omega \sin \theta)$ are the Cartesian and polar representations of the 2D Fourier transform of $f \in L_2(\mathbb{R}^2)$, respectively. The advantage of the present definition is that the frequency response $\widehat{R}(\boldsymbol{\omega}) = \widehat{R}_{\text{pol}}(\omega, \theta) = e^{i\theta}$ is a pure complex exponential of the angular frequency variable $\theta = \angle \boldsymbol{\omega}$. This highlights the fact that the transform is a convolution-type operator that acts as an all-pass filter with a phase response that is completely encoded in the orientation.

The complex Riesz transform satisfies the following properties, which can be established in the Fourier domain [23, 25, 26]:

- *Translation invariance:*

$$\forall \mathbf{x}_0 \in \mathbb{R}^2, \quad \mathcal{R}\{f(\cdot - \mathbf{x}_0)\}(\mathbf{x}) = \mathcal{R}\{f(\cdot)\}(\mathbf{x} - \mathbf{x}_0).$$

- *Scale invariance:*

$$\forall a \in \mathbb{R}^+, \quad \mathcal{R}\{f(\cdot/a)\}(\mathbf{x}) = \mathcal{R}\{f(\cdot)\}(\mathbf{x}/a).$$

- *Inner-product preservation:*

$$\forall f, g \in L_2(\mathbb{R}^2), \quad \langle f, g \rangle_{L_2} = \langle \mathcal{R}f, \mathcal{R}g \rangle_{L_2}.$$

- *Steerability of its impulse response:*

$$\mathcal{R}\{\delta\}(\mathbf{R}_{\theta_0} \mathbf{x}) = e^{-i\theta_0} \mathcal{R}\{\delta\}(\mathbf{x}),$$

where $\mathbf{R}_{\theta_0} = \begin{bmatrix} \cos \theta_0 & -\sin \theta_0 \\ \sin \theta_0 & \cos \theta_0 \end{bmatrix}$ is the matrix that implements a 2D spatial rotation by the angle θ_0 .

The adjoint of the complex Riesz operator, which is such that $\langle \mathcal{R}f, g \rangle_{L_2} = \langle f, \mathcal{R}^*g \rangle_{L_2}$, is specified by

$$(2.2) \quad \mathcal{R}^*f(\mathbf{x}) \xleftrightarrow{\mathcal{F}} \frac{(\omega_x - i\omega_y)}{\|\boldsymbol{\omega}\|} \hat{f}(\boldsymbol{\omega}) = e^{-i\theta} \hat{f}_{\text{pol}}(\omega, \theta),$$

with the property that $\mathcal{R}^{-1} = \mathcal{R}^*$.

In this work, we will also consider the n th-order complex Riesz transform \mathcal{R}^n , which represents the n -fold iterate of \mathcal{R} . It is a convolution operator that has the same invariance properties as \mathcal{R} . Its frequency response is a purely polar function given by

$$\hat{R}_{\text{pol}}^n(\omega, \theta) = e^{in\theta}.$$

The definition is extendable to negative orders as well, owing to the fact that $\mathcal{R}^{n*} = \mathcal{R}^{-n}$. The composition rule for these complex Riesz operators is $\mathcal{R}^{n_1}\mathcal{R}^{n_2} = \mathcal{R}^{n_1+n_2}$ for all $n_1, n_2 \in \mathbb{Z}$ with the convention that $\mathcal{R}^0 = \text{Identity}$.

2.2. Advanced functional properties. The inner-product preservation property implies that the complex Riesz transform is a unitary operator in $L_2(\mathbb{R}^2)$. More generally, it is a continuous mapping from $L_p(\mathbb{R}^2)$ into itself for $1 < p < \infty$ [26, 27]; in other words, there exist some constants C_p such that

$$\forall f \in L_p(\mathbb{R}^2), \quad \|\mathcal{R}f\|_{L_p} \leq C_p \|f\|_{L_p}.$$

The equality is achieved for $p = 2$ and $C_2 = 1$. The tricky aspect of the transform is that the cases $p = 1$ and $p = +\infty$ are excluded, meaning that \mathcal{R} is not stable in the classical BIBO sense (bounded input and bounded output). The difficulty stems from the fact that the impulse response of \mathcal{R} is unbounded at the origin and slowly decaying at infinity at the rate $O(\|\mathbf{x}\|^{-2})$. Specifically, if f is a function that is continuous and locally integrable, one can specify its complex Riesz transform via the following convolution integral:

$$(2.3) \quad \mathcal{R}f(\mathbf{x}) = \text{P.V.} \frac{1}{2\pi} \int_{\mathbb{R}^2} \frac{ix' - y'}{\|\mathbf{x}'\|^3} f(\mathbf{x} - \mathbf{x}') dx' dy',$$

where the symbol P.V. denotes Cauchy’s principal value (the latter is required to resolve the singular part of the integral around the origin). An equivalent statement is that $\frac{ix-y}{2\pi\|x^2+y^2\|^{3/2}}$ is the impulse response of \mathcal{R} (in the sense of distributions). Interestingly, we can provide the same kind of explicit characterization for the n th-order Riesz transforms with $n \in \mathbb{Z} \setminus \{0\}$ and show that their responses are polar-separable.

Proposition 2.1. *The impulse responses of the n th-order complex Riesz operators are given by*

$$(2.4) \quad \mathcal{R}^n\{\delta\}(x, y) = ni^n \frac{(x + iy)^n}{2\pi(x^2 + y^2)^{\frac{n+2}{2}}} = ni^n \frac{e^{in\phi}}{2\pi r^2},$$

where $(x, y) = r(\cos \phi, \sin \phi)$. They are tempered distributions whose explicit action on a function $f(x, y)$ involves a P.V. limit as in (2.3) to give a proper meaning to the space-domain integral.

Proof. The proposition can be deduced from general results on singular integrals [26, Chapter IV]. The key is to observe that $(x + iy)^n = P(\mathbf{x})$ is a harmonic polynomial that is homogeneous (with respect to scaling) of degree n . We then apply Theorem 4.5 of Stein and Weiss [26], which states that the generalized Fourier transform of $P(\mathbf{x})/\|\mathbf{x}\|^{2+n}$ (for $d = 2$) is given by

$$i^{-n}\pi \frac{\Gamma(\frac{n}{2})}{\Gamma(\frac{n+2}{2})} \frac{P(\boldsymbol{\omega})}{\|\boldsymbol{\omega}\|^n} = \frac{2\pi}{i^n n} \frac{P(\boldsymbol{\omega})}{\|\boldsymbol{\omega}\|^n},$$

which, upon substitution, yields $\widehat{R}^n(\boldsymbol{\omega}) = e^{in\theta}$ with $e^{j\theta} = \frac{\omega_x + i\omega_y}{\|\boldsymbol{\omega}\|}$. Conversely, the form of the principal value distribution (2.4) is consistent with Theorem 4.7 of Stein and Weiss [26] since $\widehat{R}^n(\boldsymbol{\omega})$ is homogeneous of degree 0 (because it is purely polar) and its restriction to the circle is square-integrable such that $\int_{-\pi}^{\pi} e^{in\theta} d\theta = 0$. ■

The Riesz operator \mathcal{R} is also related to the complex gradient or Wirtinger operator in complex analysis $\partial_{\mathbf{x}} = \frac{\partial}{\partial x} + i\frac{\partial}{\partial y}$. Specifically, we have that

$$(2.5) \quad \partial_{\mathbf{x}} f(\mathbf{x}) = i\mathcal{R}(-\Delta)^{\frac{1}{2}} f(\mathbf{x}),$$

$$(2.6) \quad \mathcal{R}\{f\}(\mathbf{x}) = -i(-\Delta)^{-\frac{1}{2}} \partial_{\mathbf{x}} f(\mathbf{x}),$$

where the first-order differential operator $(-\Delta)^{\frac{1}{2}}$ is the square-root Laplacian whose frequency response is $\|\boldsymbol{\omega}\|$. Its inverse $(-\Delta)^{-\frac{1}{2}}$ is a fractional integrator of order 1 which acts as an isotropic smoothing kernel. More generally, by expanding the frequency-domain expression $(\omega_x + i\omega_y)^n$, we can relate the n th iterate of the complex Riesz operator to the partial derivatives of order n of the signal to which it is applied:

$$(2.7) \quad \mathcal{R}^n\{f\}(x, y) = (-\Delta)^{-\frac{n}{2}} \sum_{n_1=0}^n \binom{n}{n_1} (-i)^{n_1} \partial_x^{n_1} \partial_y^{n-n_1} f(x, y).$$

The operator $(-\Delta)^{-\frac{n}{2}}$ is a fractional integrator of order n which is best specified in the Fourier domain:

$$(2.8) \quad (-\Delta)^{-\frac{n}{2}} \psi(\mathbf{x}) \xleftrightarrow{\mathcal{F}} \hat{\psi}(\boldsymbol{\omega}) \frac{1}{\|\boldsymbol{\omega}\|^n}.$$

A necessary requirement for the above definition to make sense is that the Fourier transform of ψ have a sufficient number of zeros at the origin to compensate for the singularity of the frequency response [28]. This is a condition that is generally met by wavelets that have vanishing moments up to order n . The global effect of $(-\Delta)^{-\frac{n}{2}}$ is that of a lowpass filter with a smoothing strength that increases with n .

Since the impulse response of \mathcal{R}^n is only decaying like $O(\|\mathbf{x}\|^{-2})$ (cf. Proposition 2.1), the Riesz operators will tend to spoil the decay of the functions to which they are applied. Fortunately, this behavior is tempered if the input function has a sufficient number of vanishing moments, as is typically the case with wavelets.

Theorem 2.2. *Let $\psi(\mathbf{x})$ be a function whose moments up to degree $m_0 \geq 0$ are vanishing and that is differentiable with sufficient inverse-polynomial decay; i.e., $\partial_x^{n_1} \partial_y^{n_2} \psi(\mathbf{x}) =$*

$O(1/\|\mathbf{x}\|^{-2-m_0-(n_1+n_2)+\epsilon})$ for some $0 \leq \epsilon < 1$ and for all $n_1 + n_2 \leq 1$. Then, we have

$$(2.9) \quad \mathcal{R}^n \psi(\mathbf{x}) \leq \frac{C}{1 + \|\mathbf{x}\|^{2+m_0-\epsilon'}}$$

for some $0 \leq \epsilon' < 1$.

This follows from the proofs of Theorems 3.2 and 3.4 in [29].

Thus, in order to minimize the delocalization effect of \mathcal{R}^n , it makes good sense to consider wavelets that decay faster than any polynomial and that have an infinite number of vanishing moments. In that case, $\mathcal{R}^n \psi(\mathbf{x})$ will be rapidly decreasing as well, which is the very best one can hope for. On the other hand, applying \mathcal{R}^n to a scaling function φ whose integral is nonzero will necessarily yield a poorly localized output with an asymptotic decay that is no better than $1/\|\mathbf{x}\|^2$.

Interestingly, there is no limitation with wavelet regularity since \mathcal{R}^n preserves L_2 -differentiability (or Sobolev smoothness) as a result of its unitary character (all-pass filter). The complex Riesz transform also has the convenient property of preserving vanishing moments.

3. Multiorder generalized Riesz transforms. Let $\mathbf{U}_{M,N}$ be a (possibly complex-valued) matrix of size $M \times (2N + 1)$ with $M \geq 1$.

Definition 3.1. The multiorder generalized Riesz transform with coefficient matrix $\mathbf{U}_{M,N}$ is the scalar to M -vector signal transformation $\mathcal{R}_{\mathbf{U}_{M,N}} f(\mathbf{x})$ whose m th component is given by

$$(3.1) \quad [\mathcal{R}_{\mathbf{U}_{M,N}} f(\mathbf{x})]_m = \sum_{n=-N}^{+N} u_{m,n} \mathcal{R}^n f(\mathbf{x}).$$

The adjoint transformation $\mathcal{R}_{\mathbf{U}_{M,N}}^*$ maps an M -vector signal $\mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_M(\mathbf{x}))$ back into the scalar signal domain:

$$(3.2) \quad \mathcal{R}_{\mathbf{U}_{M,N}}^* \mathbf{f}(\mathbf{x}) = \sum_{n=-N}^{+N} \mathcal{R}^{-n} \left\{ \sum_{m=1}^M \bar{u}_{m,n} f_m \right\}(\mathbf{x})$$

$$(3.3) \quad = \sum_{n=-N}^{+N} \sum_{m=1}^M \bar{u}_{m,n} \mathcal{R}^{-n} f_m(\mathbf{x}),$$

where $\bar{u}_{m,n}$ is the complex conjugate of $u_{m,n}$.

The matrix weighting $\mathbf{U}_{M,N}$ adds a level of generalization that is crucial for our later purpose while retaining all the desirable invariance properties of the elementary constituent operators \mathcal{R}^n . As we shall see, this is mostly a matter of appropriate factorization. In what follows, we will often use \mathbf{U} to denote a generic matrix of size $M \times (2N + 1)$ to simplify the notation.

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

$$\forall \mathbf{x}_0 \in \mathbb{R}^2, \quad \mathcal{R}_{\mathbf{U}}\{f(\cdot - \mathbf{x}_0)\}(\mathbf{x}) = \mathcal{R}_{\mathbf{U}}\{f(\cdot)\}(\mathbf{x} - \mathbf{x}_0),$$

$$\forall a \in \mathbb{R}^+, \quad \mathcal{R}_{\mathbf{U}}\{f(\cdot/a)\}(\mathbf{x}) = \mathcal{R}_{\mathbf{U}}\{f(\cdot)\}(\mathbf{x}/a).$$

This simply follows from the fact that both translation and dilation invariances are preserved through linear combination.

Property 2 (norm preservation and self-invertibility). *Let \mathbf{U} be a complex $M \times (2N+1)$ matrix such that the product $\mathbf{U}^H \mathbf{U} = \text{diag}(d_N, \dots, d_0, \dots, d_{-N})$ is diagonal with $\sum_{n=-N}^N d_n = 1$. Then, the corresponding multiorder Riesz transform is norm-preserving and self-invertible; that is, for all $f \in L_2(\mathbb{R}^2)$*

$$\begin{aligned}\|\mathcal{R}_{\mathbf{U}} f\|_{L_2^M(\mathbb{R}^2)} &= \|f\|_{L_2(\mathbb{R}^2)}, \\ \mathcal{R}_{\mathbf{U}}^* \mathcal{R}_{\mathbf{U}} f &= f.\end{aligned}$$

Proof. The two key relations, which follow from Definition 3.1, are

$$\begin{aligned}\mathcal{R}_{\mathbf{U}} f(\mathbf{x}) &= \mathbf{U} \mathcal{R}_{\mathbf{I}_{2N+1}} f(\mathbf{x}), \\ \mathcal{R}_{\mathbf{U}}^* \mathbf{f}(\mathbf{x}) &= \mathcal{R}_{\mathbf{I}_{2N+1}}^* \mathbf{U}^H \mathbf{f}(\mathbf{x}),\end{aligned}$$

where \mathbf{I}_{2N+1} is the identity matrix of size $2N+1$ and where \mathbf{U}^H is the Hermitian transpose of \mathbf{U} which maps a vector of \mathbb{C}^M into a vector of \mathbb{C}^{2N+1} . The generalized multiorder Riesz transform $\mathcal{R}_{\mathbf{U}} f$ yields an M -vector signal whose energy is computed as

$$\begin{aligned}\|\mathcal{R}_{\mathbf{U}} f\|_{L_2^M(\mathbb{R}^2)}^2 &= \langle \mathcal{R}_{\mathbf{I}_{2N+1}} f, \mathbf{U}^H \mathbf{U} \mathcal{R}_{\mathbf{I}_{2N+1}} f \rangle_{L_2^{2N+1}(\mathbb{R}^2)} \\ &= \sum_{n=-N}^N d_n \langle \mathcal{R}^n f, \mathcal{R}^n f \rangle_{L_2(\mathbb{R}^2)} \\ &= \sum_{n=-N}^N d_n \|f\|_{L_2(\mathbb{R}^2)}^2 = \|f\|_{L_2(\mathbb{R}^2)}^2,\end{aligned}$$

where we have used the property that \mathcal{R} and all its n -fold iterates are unitary operators. Similarly, we show that

$$\begin{aligned}\mathcal{R}_{\mathbf{U}}^* \mathcal{R}_{\mathbf{U}} f &= \mathcal{R}_{\mathbf{I}_{2N+1}}^* \mathbf{U}^H \mathbf{U} \mathcal{R}_{\mathbf{I}_{2N+1}} f \\ &= \sum_{n=-N}^N d_n \underbrace{\mathcal{R}^{n*} \mathcal{R}^n}_{\text{Identity}} f \\ &= \sum_{n=-N}^N d_n f = f. \quad \blacksquare\end{aligned}$$

The property of the generalized multiorder Riesz transform that is probably the most interesting for image processing is the fact that its action commutes with spatial rotations. This is to say that the spatially rotated versions of the operator are part of the same parametric family.

Property 3 (steerability). *The generalized multiorder 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}_{θ_0} be a 2×2 spatial*

rotation matrix with angle θ_0 . Then, $\mathcal{R}_{\mathbf{U}_{M,N}}\{\delta\}(\mathbf{R}_{\theta_0}\mathbf{x}) = \mathcal{R}_{\mathbf{U}_{M,N}\mathbf{S}_{\theta_0}}\{\delta\}(\mathbf{x})$, where $\mathbf{S}_{\theta_0} = \text{diag}(e^{-iN\theta_0}, \dots, e^{-i\theta_0}, 1, e^{i\theta_0}, \dots, e^{iN\theta_0})$ is the corresponding diagonal steering matrix of size $2N + 1$.

Proof. We recall that the frequency response of \mathcal{R}^n is $e^{in\theta}$, where θ is the angular frequency variable. We then apply the rotation property of the Fourier transform, which gives

$$\begin{aligned} \mathcal{R}^n\{\delta\}(\mathbf{R}_{\theta_0}\mathbf{x}) &= \mathcal{F}_{2D}^{-1}\{e^{in(\theta-\theta_0)}\}(\mathbf{x}) \\ &= e^{-in\theta_0}\mathcal{F}_{2D}^{-1}\{e^{in\theta}\}(\mathbf{x}) \\ &= e^{-in\theta_0}\mathcal{R}^n\{\delta\}(\mathbf{x}). \end{aligned}$$

This shows that the elementary component operators \mathcal{R}^n are self-steerable and yields the desired result. ■

4. Characterization of steerable wavelet frames. In this section, we set the foundation of our approach, which relies on the specification of a primal isotropic wavelet ψ and a matrix \mathbf{U} that determines the shape of the steerable wavelets. After a brief review of isotropic bandlimited wavelet frames, we show how we can generate a large variety of steerable wavelet frames by multiorder generalized Riesz transformation of such primal wavelets. Finally, we investigate the functional properties of generic classes of steerable wavelets. The study of specific examples is deferred to section 5.

4.1. Primal isotropic wavelet frame. The multiresolution backbone of our method is an isotropic tight wavelet frame of $L_2(\mathbb{R}^2)$ whose basis functions are generated by suitable dilation and translation of a single mother wavelet $\psi(\mathbf{x})$. Several such decompositions are available in the literature within the framework of radially bandlimited wavelets [6, 14, 30, 31]. Each brand is uniquely specified in terms of its radial frequency profile.

Proposition 4.1. *Let $h(\omega)$ be a radial frequency profile such that the following hold:*

- Condition (1): $h(\omega) = 0$ for all $\omega > \pi$.
- Condition (2): $\sum_{i \in \mathbb{Z}} |h(2^i\omega)|^2 = 1$.
- Condition (3): $\frac{d^n h(\omega)}{d\omega^n} \Big|_{\omega=0} = 0$ for $n = 0, \dots, N$.

Then, the isotropic mother wavelet ψ whose 2D Fourier transform is given by

$$(4.1) \quad \hat{\psi}(\boldsymbol{\omega}) = h(\|\boldsymbol{\omega}\|)$$

generates a tight wavelet frame of $L_2(\mathbb{R}^2)$ whose basis functions

$$(4.2) \quad \psi_{i,\mathbf{k}}(\mathbf{x}) = \psi_i(\mathbf{x} - 2^i\mathbf{k}) \quad \text{with} \quad \psi_i(\mathbf{x}) = 2^{-2i}\psi(\mathbf{x}/2^i)$$

are isotropic with vanishing moments up to order N .

The tight frame property implies that any finite energy function $f \in L_2(\mathbb{R}^2)$ can be expanded as

$$(4.3) \quad f(\mathbf{x}) = \sum_{i \in \mathbb{Z}} \sum_{\mathbf{k} \in \mathbb{Z}^2} \langle f, \psi_{i,\mathbf{k}} \rangle \psi_{i,\mathbf{k}}(\mathbf{x}).$$

Likewise, we can represent a bandlimited function using a scale-truncated wavelet expansion ($i \in \mathbb{N}$) with an overall redundancy of $1 + \frac{1}{4} + \frac{1}{4^2} + \dots = 4/3$ over the canonical representation in terms of its sampled values. The simplest choice of radial frequency profile that fulfills the conditions in Proposition 4.1 is $h(\omega) = \text{rect}\left(\frac{\omega - 3\pi/4}{\pi/2}\right)$ (Shannon ideal-bandpass wavelet) with

$$\text{rect}(x) = \begin{cases} 1, & -\frac{1}{2} < x \leq \frac{1}{2}, \\ 0 & \text{otherwise,} \end{cases}$$

which yields a Bessel-type wavelet [30]. Another prominent example is the filter that is implemented in the popular version of the steerable pyramid described in [4]:

$$h(\omega) = \begin{cases} \cos\left(\frac{\pi}{2} \log_2\left(\frac{2\omega}{\pi}\right)\right), & \frac{\pi}{4} < |\omega| \leq \pi, \\ 0 & \text{otherwise.} \end{cases}$$

The latter has the advantage of producing a wavelet that is better localized in space; it is the design that is adopted for the experimental part of this paper.

4.2. Construction of steerable wavelets by generalized Riesz transformation. Having selected our primal isotropic wavelet frame, we can now apply the proposed one-to- M functional mapping to convert it into a bona fide steerable wavelet transform.

Proposition 4.2. *The multiorder generalized Riesz transform $\mathcal{R}_{\mathbf{U}_{M,N}}$ maps a primal isotropic wavelet frame of $L_2(\mathbb{R}^2)$, $\{\psi_{i,\mathbf{k}}\}_{i \in \mathbb{Z}, \mathbf{k} \in \mathbb{Z}^2}$, into a steerable wavelet frame $\{\psi_{i,\mathbf{k}}^{(m)}\}_{m=1,\dots,M, i \in \mathbb{Z}, \mathbf{k} \in \mathbb{Z}^2}$ of $L_2(\mathbb{R}^2)$ with*

$$(4.4) \quad \psi_{i,\mathbf{k}}^{(m)} = \sum_{n=-N}^N u_{m,n} \mathcal{R}^n \psi_{i,\mathbf{k}}.$$

Moreover, the frame bounds are conserved if $\mathbf{U}_{M,N}$ satisfies the condition for Property 2.

While the present multiorder extension of the Riesz transform is more general than the N th-order one introduced in our earlier work, it has the same kind of invariances (Properties 1, 2, and 3) so that the proof of [13, Proposition 1] is directly transposable to the present case as well.

Let

$$w_{m,i}[\mathbf{k}] = \langle f, \psi_{i,\mathbf{k}}^{(m)} \rangle$$

denote the corresponding (steerable) wavelet coefficients at scale i and channel m of a signal $f(\mathbf{x}) \in L_2(\mathbb{R}^2)$. Then, Proposition 4.2 implies that $f(\mathbf{x})$ is completely specified by its wavelet coefficients $\{w_{m,i}[\mathbf{k}]\}$ and that it can be reconstructed from its wavelet expansion:

$$(4.5) \quad f(\mathbf{x}) = \sum_{i \in \mathbb{Z}} \sum_{\mathbf{k} \in \mathbb{Z}^2} \sum_{m=1}^M w_{m,i}[\mathbf{k}] \psi_{i,\mathbf{k}}^{(m)}(\mathbf{x}),$$

which is the multichannel counterpart of (4.3). Hence, the wavelets $\{\psi_{i,\mathbf{k}}^{(m)}\}$ form a frame for $L_2(\mathbb{R}^2)$ with a global redundancy of $(4/3)M$. Moreover, this steerable wavelet transform admits a fast filterbank implementation, with a computational complexity that is at most M times that of the primal decomposition.

The bottom line is that the proposed scheme yields a whole family of steerable wavelet transforms which are parameterized by the rectangular matrix $\mathbf{U}_{M,N}$.

4.3. Steerable wavelet properties. By using the fact that the polar representation of the Fourier transform of $\mathcal{R}^n\psi(\mathbf{x})$ is $e^{in\theta}h(\omega)$, we can readily show that the Fourier transform of a generic steerable wavelet is polar-separable; i.e.,

$$(4.6) \quad \psi_{\text{Gen}}(\mathbf{x}) = \sum_{n=-N}^N u_n \mathcal{R}^n\psi(\mathbf{x}) \xleftrightarrow{\mathcal{F}} \hat{\psi}_{\text{Gen}}(\omega) = h(\omega)\hat{u}(\theta),$$

where $\hat{u}(\theta) = \sum_{n=-N}^N u_n e^{in\theta}$ is 2π -periodic. Conversely, we have the guarantee that the proposed representation provides a full parameterization of the wavelets whose Fourier transform is polar-separable because the complex exponentials $\{e^{in\theta}\}_{n \in \mathbb{Z}}$ form a basis of $L_2([-\pi, \pi])$.

The steerable wavelet $\psi_{\text{Gen}}(\mathbf{x})$ is real-valued if and only if its Fourier transform is Hermitian-symmetric (i.e., $\overline{\hat{\psi}_{\text{Gen}}(-\omega)} = \hat{\psi}_{\text{Gen}}(\omega)$). Since the real term $h(\omega)$ can be factored out, this gets translated into the angular condition $\overline{\hat{u}(\theta + \pi)} = \hat{u}(\theta)$. Additionally, we can impose symmetry by considering Fourier series with even or odd harmonic terms.

4.3.1. Even-harmonic-type wavelets. These wavelets are constructed by restricting the sum to even terms only:

$$\psi_{\text{Even}}(\mathbf{x}) = \sum_{n=-\lfloor N/2 \rfloor}^{\lfloor N/2 \rfloor} u_{2n} \mathcal{R}^{2n}\psi(\mathbf{x}).$$

In that case $\hat{u}(\theta + \pi) = \hat{u}(\theta)$. It follows that such wavelets will be real-valued symmetric if and only if their angular Fourier coefficients are Hermitian-symmetric:

$$\psi_{\text{Even}}(\mathbf{x}) = \psi_{\text{Even}}(-\mathbf{x}) \Leftrightarrow u_{2n} = \overline{u_{-2n}}.$$

Conversely, they will be real-valued antisymmetric if and only if their Fourier coefficients are Hermitian-antisymmetric:

$$\psi_{\text{Even}}(\mathbf{x}) = -\psi_{\text{Even}}(-\mathbf{x}) \Leftrightarrow u_{2n} = -\overline{u_{-2n}}.$$

4.3.2. Odd-harmonic-type wavelets. This is the complementary type of wavelets that involves odd-harmonic terms only:

$$\psi_{\text{Odd}}(\mathbf{x}) = \sum_{n=-\lfloor (N+1)/2 \rfloor}^{\lfloor (N-1)/2 \rfloor} u_{2n+1} \mathcal{R}^{2n+1}\psi(\mathbf{x}).$$

In that case $\hat{u}(\theta + \pi) = -\hat{u}(\theta)$. This leads to real-valued wavelet configurations that are transposed versions of the previous ones:

$$\psi_{\text{Odd}}(\mathbf{x}) = -\psi_{\text{Odd}}(-\mathbf{x}) \Leftrightarrow u_{2n+1} = \overline{u_{-2n-1}},$$

$$\psi_{\text{Odd}}(\mathbf{x}) = \psi_{\text{Odd}}(-\mathbf{x}) \Leftrightarrow u_{2n+1} = -\overline{u_{-2n-1}}.$$

4.3.3. Complex-quadrature-type wavelets. The idea here is to constrain the sum to be one-sided:

$$\psi_{\text{Comp}}(\mathbf{x}) = \sum_{n=0}^N u_n \mathcal{R}^n \psi(\mathbf{x}).$$

When the u_n are real-valued (or purely imaginary), this design results in complex wavelets whose real and imaginary components are in angular quadrature. The practical interest is that such wavelets have a natural amplitude/phase interpretation which makes them more robust to local deviations from the main orientation. Imposing an additional even (resp., odd) harmonic constraint defines wavelet components that are symmetric (resp., antisymmetric) in the spatial domain.

4.3.4. Componentwise orthogonality. By using Parseval's relation, the isotropy of the primal wavelet ψ , and the orthogonality of the circular harmonics, we readily compute the inner product between two generic steerable wavelets $\psi^{(m)}$ and $\psi^{(m')}$ as

$$\begin{aligned} \langle \psi^{(m)}, \psi^{(m')} \rangle_{L_2(\mathbb{R}^2)} &= \frac{1}{(2\pi)^2} \int_0^{+\infty} \int_{-\pi}^{\pi} |h(\omega)|^2 \hat{u}_m(\theta) \bar{\hat{u}}_{m'}(\theta) \omega d\omega d\theta \\ &= \frac{1}{2\pi} \int_0^{+\infty} |h(\omega)|^2 \omega d\omega \cdot \frac{1}{2\pi} \int_{-\pi}^{\pi} \hat{u}_m(\theta) \bar{\hat{u}}_{m'}(\theta) d\theta \\ (4.7) \qquad &= \|\psi\|_{L_2(\mathbb{R}^2)}^2 \cdot \sum_{n=-N}^N u_{m,n} \bar{u}_{m',n}. \end{aligned}$$

The implication is that the L_2 -norm of a steerable wavelet is proportional to the ℓ_2 -norm of its circular harmonic coefficients u_n :

$$\|\psi_{\text{Gen}}\|_{L_2(\mathbb{R}^2)} \propto \left(\sum_{n=-N}^N |u_n|^2 \right)^{\frac{1}{2}}.$$

The above formula yields the correct normalization factor for specifying wavelet-domain thresholding functions for image denoising. In what follows, we will refer to $1/\|u\|_{\ell_2}$, which is a crucial algorithmic component (cf. [15]), as the wavelet equalization factor. Equation (4.7) also implies that even-harmonic-type wavelets are necessarily orthogonal to all odd-harmonic-type wavelets. Along the same line of thought, the most general statement that can be made about componentwise wavelet orthogonality is as follows.

Proposition 4.3. *Let $\mathcal{R}_{\mathbf{U}_{M,N}} \psi(\mathbf{x})$ be a set of steerable wavelets obtained by multiorder generalized Riesz transform of a primal isotropic wavelet ψ with $M \leq 2N+1$. Then, the component wavelets are orthogonal if and only if $\mathbf{U}_{M,N} \mathbf{U}_{M,N}^H$ is a diagonal matrix of size M .*

While the latter requirement for orthogonality is reminiscent of the condition for self-reversibility in Property 2, it is generally not equivalent to it unless the underlying matrices are unitary (up to a normalization factor), in which case $M = 2N + 1$ and the diagonal terms are all equal.

One should also keep in mind that the type of orthogonality that is covered by Proposition 4.3 is only valid across the wavelet channels (index m) at a given wavelet-domain location

(\mathbf{k}, i) . Indeed, there is little hope in enforcing full orthogonality across translations and dilations as well, simply because the wavelet frame $\{\psi_{i,\mathbf{k}}^{(m)}\}$ is overcomplete by a factor of $(4/3)M$.

5. Specific examples of steerable wavelet transforms. We will now investigate particular choices of matrices \mathbf{U} leading to the specification of interesting families of wavelets that are both steerable and self-reversible. First, we will consider known constructions and show that these can be formulated as particular cases of the proposed framework. We will then introduce new families of steerable wavelets with optimized correlation and/or localization properties.

5.1. Primary examples of low-order steerable wavelets.

5.1.1. Gradient or Marr-like wavelets. The first nontrivial, real-valued case for $N = 1$ is obtained with

$$\mathbf{U}_G = \begin{pmatrix} \frac{i}{2} & 0 & \frac{i}{2} \\ -\frac{i}{2} & 0 & \frac{i}{2} \end{pmatrix}.$$

The corresponding frequency-domain formulae are $i\frac{\omega_x}{\|\boldsymbol{\omega}\|} = i\cos\theta = \frac{i}{2}(e^{i\theta} + e^{-i\theta})$ and $i\frac{\omega_y}{\|\boldsymbol{\omega}\|} = i\sin\theta = \frac{i}{2}(e^{i\theta} - e^{-i\theta})$. This design, which can be traced back to the early work of Freeman and Simoncelli (see [2, 32]), yields two gradient-like wavelets $(\partial_x\psi_1, \partial_y\psi_1) = \nabla\psi_1$, where $\psi_1 = (-\Delta)^{-\frac{1}{2}}\psi$ is a smoothed version of the primal isotropic wavelet (isotropic fractional integral of order 1) (cf. Figure 1(a)). This particular wavelet configuration is also the one that is used implicitly in the Marr-like pyramid, which involves a nonbandlimited primal wavelet that is the Laplacian of a Gaussian-like smoothing kernel [33]. Observe that the two gradient wavelets are antisymmetric, which is consistent with the fact that they are of odd-harmonic type with Hermitian-symmetric coefficients. Moreover, $\mathbf{U}_G^H\mathbf{U}_G = \text{diag}(\frac{1}{2}, 0, \frac{1}{2})$ and $\mathbf{U}_G\mathbf{U}_G^H = \text{diag}(\frac{1}{2}, \frac{1}{2})$, which implies that the gradient wavelet transform is not only self-reversible (tight frame), but also equalized (and orthogonal) on a componentwise basis. Alternatively, the wavelets may also be encoded using a single complex quadrature-type wavelet $(\partial_x\psi_1)$ with

$$\mathbf{U}_{G,\text{Comp}} = \begin{pmatrix} i & 0 & 0 \end{pmatrix},$$

which is a more concise representation of the same system.

5.1.2. Monogenic wavelets. The next interesting case that yields a full set of real-valued wavelets for $N = 1$ is

$$\mathbf{U}_{\text{Mono}} = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 1 & 0 \\ -\frac{i}{2} & 0 & -\frac{i}{2} \\ \frac{1}{2} & 0 & -\frac{1}{2} \end{pmatrix}.$$

The corresponding wavelets, which are shown in Figure 1(b), actually provide the monogenic signal extension of the primal one: $(\psi, \mathcal{R}_x\psi, \mathcal{R}_y\psi)$, where \mathcal{R}_x and \mathcal{R}_y denote the x and y component operators of the conventional (noncomplex) Riesz transform [25, 34]. We recall that the monogenic signal is the 2D counterpart of the one-dimensional (1D) analytic signal. It gives access to characteristic signal parameters such as the local orientation, phase, and amplitude, which are transposable to the wavelet domain as well. This opens the door to various forms of nonconventional wavelet-domain processing, such as instantaneous frequency estimation, demodulation, tensor-based orientation, and coherence analysis [25]. The monogenic wavelet




Figure 1. Comparative display of first- and second-order steerable wavelets.

transform may be viewed as an augmented version of the gradient one, which offers significant advantages with respect to global shift invariance (thanks to its additional phase parameter). It is built from the concatenation of an odd-harmonic set (the two gradient wavelets) and an even-harmonic set (primal wavelet). It is easy to verify that $\mathbf{U}_{\text{Mono}}^H \mathbf{U}_{\text{Mono}} = \text{diag}(\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$ and $\mathbf{U}_{\text{Mono}} \mathbf{U}_{\text{Mono}}^H = \text{diag}(\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$, which implies that the transform is self-reversible and componentwise orthogonal, but not fully equalized. The latter needs to be taken into account when designing some corresponding wavelet-domain denoising procedure.

5.1.3. Hessian wavelets. With $N = 2$, we get access to second-order spatial derivatives. In particular, the choice

$$\mathbf{U}_{\text{H}} = \mathbf{U}_{\text{Riesz},2} = \begin{pmatrix} -\frac{1}{4} & 0 & -\frac{1}{2} & 0 & -\frac{1}{4} \\ -\frac{i}{2\sqrt{2}} & 0 & 0 & 0 & \frac{i}{2\sqrt{2}} \\ \frac{1}{4} & 0 & -\frac{1}{2} & 0 & \frac{1}{4} \end{pmatrix}$$

leads to the specification of the Hessian wavelets $(\partial_{xx}\psi_2, \sqrt{2}\partial_{xy}\psi_2, \partial_{yy}\psi_2)$, where $\psi_2 = (-\Delta)^{-1}\psi$ is a smoothed version of the primal isotropic wavelet (cf. Figure 1(c)). Here, we find that $\mathbf{U}_{\text{H}}^H \mathbf{U}_{\text{H}} = \text{diag}(\frac{1}{4}, 0, \frac{1}{2}, 0, \frac{1}{4})$, which confirms that the transform is self-reversible (cf. Property 2 and Proposition 4.2). On the other hand, we have that

$$\mathbf{U}_{\text{H}} \mathbf{U}_{\text{H}}^H = \begin{pmatrix} \frac{3}{8} & 0 & \frac{1}{8} \\ 0 & \frac{1}{4} & 0 \\ \frac{1}{8} & 0 & \frac{3}{8} \end{pmatrix},$$

which indicates that the Hessian wavelets are not orthogonal (or equalized) in a componentwise fashion.

5.1.4. Simoncelli's two- and three-component wavelets. The angular components of Simoncelli's filters in the M -channel steerable pyramid are proportional to $\{(j \cos(\theta - \theta_m))^N\}$ with $\theta_m = \frac{\pi(m-1)}{M}$, $m = 1, \dots, M$, and $N = M - 1$. Using the connection with the directional Hilbert transform [13], we can obtain the differential interpretation, $\psi_{\text{Sim}}^{(m)} \propto D_{\theta_m}^N (-\Delta)^{-\frac{N}{2}} \psi$, where $D_{\theta_m}^N$ is the N th directional derivative along the direction θ_m and $(-\Delta)^{-\frac{N}{2}}$ is the isotropic fractional integrator of order N , which is a smoothing operator. For $N = 1$, we end up with

directional derivatives along the two coordinate directions ($\theta_1 = 0$ and $\theta_2 = \pi/2$) so that the 2-channel version of the steerable pyramid is in fact rigorously equivalent to the gradient wavelet transform described above.

The case $N = 2$ is more instructive because it brings out the specificities of the equiangular design, which is distinct from the other solutions considered here. The transform parameters for Simoncelli’s 3-channel solution are given by

$$\mathbf{U}_{\text{Sim},2} = \begin{pmatrix} -\frac{1}{3\sqrt{2}} & 0 & -\frac{\sqrt{2}}{3} & 0 & -\frac{1}{3\sqrt{2}} \\ -\frac{(-1)^{2/3}}{3\sqrt{2}} & 0 & -\frac{\sqrt{2}}{3} & 0 & \frac{\sqrt[3]{-1}}{3\sqrt{2}} \\ \frac{\sqrt[3]{-1}}{3\sqrt{2}} & 0 & -\frac{\sqrt{2}}{3} & 0 & -\frac{(-1)^{2/3}}{3\sqrt{2}} \end{pmatrix}.$$

The corresponding wavelets, which are shown in Figure 1(d), are rotated versions at angles $0, \pi/3, 2\pi/3$ of the second-derivative wavelet $\partial_{xx}\psi_2$ (first component of Hessian-like transform). A direct calculation shows that $\mathbf{U}_{\text{Sim},2}^H \mathbf{U}_{\text{Sim},2} = \text{diag}(\frac{1}{6}, 0, \frac{2}{3}, 0, \frac{1}{6})$, which can be taken as a confirmation that the steerable pyramid is indeed self-reversible. The Gram matrix of the wavelets is given by

$$\mathbf{U}_{\text{Sim},2} \mathbf{U}_{\text{Sim},2}^H = \begin{pmatrix} \frac{1}{3} & \frac{1}{6} & \frac{1}{6} \\ \frac{1}{6} & \frac{1}{3} & \frac{1}{6} \\ \frac{1}{6} & \frac{1}{6} & \frac{1}{3} \end{pmatrix},$$

which is far from diagonal.

The important point that we want to make here is that the Simoncelli-3 and Hessian wavelets, which are both of even-harmonic type (ridge detectors), actually span the same steerable subspaces; yet, they are both fundamentally distinct from the point of view of shape diversity. The present analysis would even suggest that the Hessian wavelet transform could be an attractive substitute because of the natural link it makes with differential geometry and its more favorable correlation properties (smaller off-diagonal terms).

5.2. Circular harmonic wavelets. The circular Harmonic wavelets, which go back to the work of Jacovitti and Neri [12], stand out as the canonical basis of the proposed wavelet parameterization. The associated weighting matrix is proportional to the $(2N + 1) \times (2N + 1)$ identity matrix

$$\mathbf{U}_{2N+1,N} = \frac{1}{\sqrt{2N + 1}} \mathbf{I}_{2N+1}.$$

The corresponding wavelet transform is self-reversible and fully equalized by construction. The circular harmonic wavelets with $n \in \mathbb{Z}$ are best characterized in the frequency domain:

$$(5.1) \quad \hat{\psi}_{\text{circ}}^{(n)}(\omega) = \hat{\psi}(\omega) e^{in\theta} = h(\omega) e^{in\theta},$$

where ω and θ are the polar frequency-domain variables. Note that the above definition does not include the normalization factor which is dependent on N .

The circular Harmonic wavelets satisfy the recursive space-domain formula $\psi_{\text{circ}}^{(n)} = \mathcal{R}\psi_{\text{circ}}^{(n-1)} = \mathcal{R}^n \psi_{\text{circ}}^{(0)}$, where $\psi_{\text{circ}}^{(0)} = \psi$ is the primary isotropic wavelet.

They can also be determined analytically by calculating their inverse Fourier transform.

Proposition 5.1. *The circular harmonic wavelets are given by*

$$(5.2) \quad \psi_{\text{circ}}^{(n)}(r, \phi) = \frac{1}{2\pi} H_n(r) e^{in\phi},$$

where

$$H_n(r) = \int_0^{+\infty} h(\omega) J_n(r\omega) \omega d\omega$$

is the n th-order Hankel transform of the radial frequency profile $h(\omega)$.

Proof. First, we note that $\langle \mathbf{x}, \boldsymbol{\omega} \rangle = r\omega \cos(\phi - \theta)$, where $\mathbf{x} = r(\cos \phi, \sin \phi)$ and $\boldsymbol{\omega} = \omega(\cos \theta, \sin \theta)$. This allows us to write the polar version of the inverse 2D Fourier transform of (5.1) as

$$\begin{aligned} \psi_{\text{circ}}^{(n)}(r, \phi) &= \frac{1}{(2\pi)^2} \int_{-\pi}^{\pi} \int_0^{+\infty} h(\omega) e^{in\theta} e^{j\omega r \cos(\phi-\theta)} \omega d\omega d\theta \\ &= \frac{1}{(2\pi)^2} \int_0^{+\infty} h(\omega) \omega \int_{\alpha_0}^{\alpha_0+2\pi} e^{in(\phi+\alpha)-i\omega r \sin \alpha} d\alpha d\omega \\ &= \frac{1}{2\pi} \int_0^{+\infty} h(\omega) \omega \left(\frac{1}{2\pi} \int_0^{2\pi} e^{in\alpha-i\omega r \sin \alpha} d\alpha \right) d\omega e^{in\phi}, \end{aligned}$$

where we have made the change of variables $\phi - \theta = \alpha + \frac{\pi}{2}$ with $\alpha_0 = \phi - \frac{3\pi}{2}$. Next, we identify the latter inner integral as the n th-order Bessel function of the first kind,


$$(5.3) \quad J_n(x) = \frac{1}{2\pi} \int_0^{2\pi} e^{i(n\alpha - x \sin \alpha)} d\alpha.$$

This allows us to rewrite $\psi_{\text{circ}}^{(n)}(r, \phi)$ as


$$\psi_{\text{circ}}^{(n)}(r, \phi) = \frac{e^{in\phi}}{2\pi} \underbrace{\int_0^{+\infty} h(\omega) J_n(r\omega) \omega d\omega}_{H_n(r)},$$

where the remaining integral is the n th-order Hankel transform of h (cf. [35]). ■


The interest of this result is that the circular harmonic wavelets are polar-separable in the space domain as well. Indeed, there is a nice duality between the polar Fourier and space-domain formulae (5.1) and (5.2) with the radial profiles $h(\omega)$ and $H_n(r)$ being n th-order Hankel transforms of one another. This yields a series of complex wavelets with an aesthetically appealing n -fold circular symmetry (cf. Figure 2(a)). We need to emphasize that the above space-domain separability property is truly specific to the circular harmonic wavelets. The flip side of the coin is that these wavelets completely lack angular selectivity, which happens to be a handicap for most applications.




(a) Circular harmonic wavelets ($N = 5$). Top/middle/bottom line: real part, imaginary part and absolute value, respectively.



(b) Symmetric Riesz wavelets ($N = 4$).



(c) Anti-symmetric Riesz wavelets ($N = 5$).



(d) 5-channel equiangular wavelets ($N = 4$).



(e) 6-channel equiangular wavelets ($N = 5$).

Figure 2. Examples of high-order steerable wavelets.

5.3. Equiangular designs and Simoncelli's wavelets. We are now proposing a generalization of Simoncelli's equiangular design without any restriction on the angular shaping filter $\hat{u}(\theta)$. We will consider two situations: the general M -channel complex case with $M \geq 2N + 1$, and a reduced M' -channel real-valued version with $M' = M/2 \geq N + 1$ which encompasses Simoncelli's solution.

The idea is to pick a first directional wavelet $\psi^{(1)}(\mathbf{x}) = \sum_{n=-N}^N u_{1,n} \mathcal{R}^n \psi(\mathbf{x})$ where the weights in the expansion can be selected arbitrarily up to a normalization factor and to specify the others as the rotated versions of the first in an equiangular configuration. Specifically, we have that $\psi^{(m)}(\mathbf{x}) = \psi^{(1)}(\mathbf{R}_{\theta_m} \mathbf{x})$, where the rotation angles $\theta_m = \frac{2\pi(m-1)}{M}$ with $m = 1, \dots, M$ are equally spaced around the circle.

The weighting matrix entries of the rotated wavelets are then obtained by applying the steering property (Property 3):

$$u_{m,n} = u_{1,n} e^{-in\theta_m}.$$

The remarkable property is that the proposed directional wavelet configuration is self-invertible, irrespective of the actual choice of $\psi^{(1)}$. While this may be established in the linear algebra framework via the factorization of some underlying discrete Fourier matrices, we prefer to approach the problem from a signal processing perspective. Specifically, we will show that the required frame property

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

is automatically satisfied provided that $\sum_{n=1}^N |u_{1,n}|^2 = \frac{1}{M}$.

Theorem 5.2. *Let $U(z) = \sum_{n=-N}^N u_n z^{-n}$, where the u_n are arbitrary complex coefficients, and let M be an integer greater than or equal to $2N + 1$. Then, the equiangular directional wavelets $\{\psi^{(m)}\}_{m=1}^M$ whose Fourier transforms are given by $\hat{\psi}^{(m)}(\boldsymbol{\omega}) = \hat{\psi}(\boldsymbol{\omega}) \hat{u}_m(\theta)$ with $\theta = \angle \boldsymbol{\omega}$ and*

$$(5.4) \quad \hat{u}_m(\theta) = \frac{U(e^{i(\theta - \frac{2\pi(m-1)}{M})})}{\sqrt{M \sum_{n=-N}^N |u_n|^2}}$$

are such that $\sum_{m=1}^M |\hat{\psi}^{(m)}(\boldsymbol{\omega})|^2 = |\hat{\psi}(\boldsymbol{\omega})|^2$.

Proof. The crucial observation is that the coefficients a_n of the product polynomial

$$U(z)U(z^{-1}) = \sum_{n=-2N}^{2N} a_n z^n$$

correspond to the autocorrelation of the sequence u_n (of length $2N + 1$); in particular, this implies that the terms with index $|n| > 2N$ are necessarily zero. Therefore, if we down-sample the sequence a_n by a factor $M > 2N$, we are left with a single nonzero coefficient at the origin: $a_0 = \sum_{n=-N}^N |u_n|^2$. In the frequency domain, this down-sampling operation corresponds to a periodization, leading to the identity

$$(5.5) \quad \frac{1}{M} \sum_{m=1}^M |U(e^{i(\theta - \frac{2\pi(m-1)}{M})})|^2 = a_0,$$

where the right-hand side is the Fourier transform of the remaining impulse. This is equivalent to

$$\sum_{m=1}^M |\hat{u}_m(\theta)|^2 = 1$$

with the circular harmonic filters being specified by (5.4). ■

Next we observe that there are many instances of the above design where the wavelets appear in duplicate form, meaning that the effective number of wavelet channels can be reduced by a factor of two.

Corollary 5.3. *Let $U(e^{i\theta}) = \sum_{n=-N}^N u_n e^{-in\theta}$ be an arbitrary trigonometric polynomial of order N such that $|U(e^{i\theta})| = |U(e^{i(\theta+\pi)})|$. Then, the equiangular directional wavelets $\{\psi^{(m)}\}_{m=1}^{M'}$*

with $M' \geq N + 1$ whose Fourier transforms are given by $\hat{\psi}^{(m)}(\boldsymbol{\omega}) = \hat{\psi}(\boldsymbol{\omega})\hat{u}_m(\theta)$ with $\theta = \angle\boldsymbol{\omega}$ and

$$(5.6) \quad \hat{u}_m(\theta) = \frac{U(e^{i(\theta - \frac{\pi(m-1)}{M'})})}{\sqrt{M' \sum_{n=-N}^N |u_n|^2}}$$

are such that $\sum_{m=1}^{M'} |\hat{\psi}^{(m)}(\boldsymbol{\omega})|^2 = |\hat{\psi}(\boldsymbol{\omega})|^2$.

Proof. This can be seen by examining (5.5) in the particular case where $M = 2M'$ is even. The π -periodicity condition on $|U(e^{i\theta})|$ implies that the sum from 1 to M is twice the sum from 1 to M' . ■

Note that the condition in Corollary 5.3 is automatically satisfied when the first directional wavelet $\psi^{(1)}(\mathbf{x})$ is real-valued or purely imaginary. Corollary 5.3 slightly extends and corrects¹ an earlier result of ours [13, Theorem 2] where the admissibility condition on the filter was incorrectly stated. It also covers Simoncelli’s classical design, which corresponds to the particular choice $U(z) = i(\frac{z+z^{-1}}{2})^N$ with $M' = N + 1$.

5.4. High-order Riesz wavelets and partial derivatives. We have investigated the N th-order Riesz wavelets extensively in our previous work [13]. To show how these fit into the present framework, we write the Fourier transform of the Riesz wavelets of order N as

$$\hat{\psi}_{\text{Riesz}}^{(m)}(\boldsymbol{\omega}) = \hat{\psi}(\boldsymbol{\omega})\hat{u}_m(\theta) \quad \text{with} \quad \hat{u}_m(\theta) = \sqrt{\binom{N}{m}} (-i)^N (\cos \theta)^{N-m} (\sin \theta)^m,$$

with $m = 0, \dots, N$ and $M = N + 1$. In the space domain, this translates into real-valued (resp., purely imaginary) wavelets with an N th-order partial-derivative-like behavior:

$$\psi_{\text{Riesz}}^{(m)}(\mathbf{x}) \propto \partial_x^{N-m} \partial_y^m \psi_N(\mathbf{x}),$$

where the isotropic kernel function $\psi_N(\mathbf{x}) = (-\Delta)^{-\frac{N}{2}} \psi(\mathbf{x})$ is a smoothed version of the primal wavelet ψ . This makes the link with the gradient and Hessian wavelets in section 5.1, which are the Riesz wavelets of orders 1 and 2, respectively.

The thought-after parameterization is obtained by computing the Fourier series coefficients of $\hat{u}_m(\theta)$, which amounts to plugging in the Euler relations $\cos \theta = \frac{e^{i\theta} + e^{-i\theta}}{2}$ and $\sin \theta = \frac{e^{i\theta} - e^{-i\theta}}{2i}$ and performing the polynomial expansion. The end result is an $(N + 1) \times (2N + 1)$ weighting matrix $\mathbf{U}_{\text{Riesz},N} = \mathbf{U}_{N+1,N}$ which automatically meets the condition in Property 2 because the N th-order Riesz transform is self-reversible by construction [13]. While $\mathbf{U}_{\text{Riesz},N}^H \mathbf{U}_{\text{Riesz},N}$ is diagonal, the converse property is not satisfied for $N > 1$, meaning that the higher-order Riesz wavelets are not equalized. Another important observation is that the Riesz wavelets of odd order are of odd-harmonic type (antisymmetric contour detectors), while the Riesz wavelets

¹Addendum to Theorem 2 of [13]: The introductory statement “Let $H(e^{i\theta}) = \sum_{k=-N}^N c[k]e^{ik\theta}$ where the $c[k]$ ’s are arbitrary real-valued (or purely imaginary) coefficients” should be replaced with “Let $H(e^{i\theta}) = \sum_{k=-N}^N c[k]e^{ik\theta}$ be an arbitrary trigonometric polynomial such that $|H(e^{i\theta})| = |H(e^{i(\theta+\pi)})|$ ” and the corresponding statement in the proof on p. 645, 2nd column, line 11 “since the coefficients are real-valued (or purely imaginary)” deleted.

of even order are of even-harmonic type (symmetric, ridge detectors). This is illustrated in Figure 2(b)–(c).

As example of higher-order parameterization, we provide

$$\mathbf{U}_{\text{Riesz},3} = \begin{pmatrix} \frac{i}{8} & 0 & \frac{3i}{8} & 0 & \frac{3i}{8} & 0 & \frac{i}{8} \\ -\frac{\sqrt{3}}{8} & 0 & -\frac{\sqrt{3}}{8} & 0 & \frac{\sqrt{3}}{8} & 0 & \frac{\sqrt{3}}{8} \\ -\frac{i\sqrt{3}}{8} & 0 & \frac{i\sqrt{3}}{8} & 0 & \frac{i\sqrt{3}}{8} & 0 & -\frac{i\sqrt{3}}{8} \\ \frac{1}{8} & 0 & -\frac{3}{8} & 0 & \frac{3}{8} & 0 & -\frac{1}{8} \end{pmatrix}.$$

A direct calculation shows that $\mathbf{U}_{\text{Riesz},3}^H \mathbf{U}_{\text{Riesz},3} = \text{diag}(\frac{1}{8}, 0, \frac{3}{8}, 0, \frac{3}{8}, 0, \frac{1}{8})$, which confirms that the transform is self-reversible. On the other hand, we have that

$$\mathbf{U}_{\text{Riesz},3} \mathbf{U}_{\text{Riesz},3}^H = \begin{pmatrix} \frac{5}{16} & 0 & \frac{\sqrt{3}}{16} & 0 \\ 0 & \frac{3}{16} & 0 & \frac{\sqrt{3}}{16} \\ \frac{\sqrt{3}}{16} & 0 & \frac{3}{16} & 0 \\ 0 & \frac{\sqrt{3}}{16} & 0 & \frac{5}{16} \end{pmatrix},$$

which expresses the fact that the Riesz component wavelets are partially correlated.

5.5. Concatenation of even- and odd-harmonic wavelets. While the above Riesz wavelets provide a family with interesting differential properties, they do not span the full space of steerable functions of order N . This becomes obvious if we recall that they are made up of functions that are all either symmetric or antisymmetric. The same remark is applicable to Simoncelli's equiangular design. The situation can be fixed easily through the concatenation of odd/even wavelet families at successive orders $N - 1$ and N . The advantage of such a construct is two-fold: (1) the resulting matrix is square of size $(2N + 1)$, meaning that the concatenated wavelets provide a basis of the steerable functions of order N , and (2) the correlation structure does not deteriorate because the even wavelets are necessarily orthogonal to the odd ones. The latter property actually ensures that the concatenated transform is self-reversible, albeit not necessarily equalized. A prototypical example is the monogenic wavelet transform in section 5.1.2, which is obtained from the concatenation of the Riesz wavelets of orders 0 and 1.

5.6. Prolate spheroidal wavelets. A convenient way of constructing other families of wavelets is to consider the generic parameterization

$$\mathbf{U}_{2N+1,N} = \frac{1}{\sqrt{2N+1}} \mathbf{U},$$

where \mathbf{U} is some unitary matrix of size $(2N + 1)$. The trivial case $\mathbf{U} = \text{Identity}$ yields the circular harmonic wavelets which have no angular selectivity. Here, we propose exploring the other extreme, which calls for the identification of the “most directional” wavelets with an angular profile $\hat{u}(\theta)$ that is maximally concentrated around some central orientation θ_0 . A possible solution to this design problem is provided by Slepian's discrete prolate spheroidal sequences, which maximize the energy concentration of $\hat{u}(\theta) = \sum_n u_n e^{i\theta}$ in a rectangular

window of relative size B [36, 37]. A variation of Slepian's formulation, which is presented in the appendix and better suited to our problem, is to minimize the angular variance of the profile:

$$\sigma_u^2 = \frac{\int_{-\pi}^{+\pi} |\hat{u}(\theta)|^2 \theta^2 d\theta}{\int_{-\pi}^{+\pi} |\hat{u}(\theta)|^2 d\theta}.$$

More generally, we may replace the quadratic term $w_1(\theta) = \theta^2$ in the numerator by some weighting function $w(\theta) \geq 0$. To account for the fact that the angular profiles corresponding to real-valued wavelets are such that $|\hat{u}(\theta)| = |\hat{u}(\theta + \pi)|$ (Hermitian symmetry of the 2D Fourier transform), we also introduce a periodized version of the above variance measure corresponding to the window function

$$w_2(\theta) = \chi_{[-\pi, -\pi/2]}(\theta) (\theta + \pi)^2 + \chi_{[-\pi/2, \pi/2]}(\theta) \theta^2 + \chi_{[\pi/2, \pi]}(\theta) (\theta - \pi)^2,$$

where $\chi_{[\theta_1, \theta_2]}(\theta)$ denotes the indicator function for the interval $[\theta_1, \theta_2]$. The bottom line is that the specification of a particular window function w gives rise to an eigenvalue problem that involves a symmetric matrix $\mathbf{W}(w)$ of size $2N + 1$, as detailed in the appendix. The eigenvectors of $\mathbf{W}(w)$, which are generalized Slepian sequences, then specify the unitary matrix \mathbf{U} corresponding to an orthogonal set of wavelets with optimal angular localization.

We note that a similar method of optimization was proposed by Simoncelli and Farid for the derivation of the harmonic components of steerable wedge-like feature detectors [38]. The main difference is that these authors restricted their attention to an equiangular design (first eigenvector only) and did not investigate the issue of the reversibility of such a feature extraction process.

5.7. Signal-adapted wavelets. In a recent paper [15], we introduced the steerable PCA wavelets which were constructed by appropriate linear transformation of the N th-order Riesz wavelets of section 5.4. We also found that the application of an equalization step prior to PCA would significantly boost the denoising performance of such signal-adapted transforms. This concept is transposable to the present framework, which brings two advantages. First, performing the training on the circular harmonic wavelets simplifies the process and avoids the need for equalization. This phase involves the estimation and eigenvector decomposition of the scatter matrices of the steered wavelet coefficients of some reference image(s) on a scale-by-scale basis. Second, the fact of considering an enlarged space of steerable wavelets ($2N + 1$ circular harmonic wavelets as compared to the $N + 1$ Riesz wavelets of our initial formulation) gives access to a wider range of wavelet shapes that combines symmetric and antisymmetric feature detectors.

The implementation details are directly transposable from [15, section V.A-B] after substitution of the equalized Riesz wavelets coefficients by the circular harmonic ones. The fundamental ingredient that makes the transform rotation invariant is the steering mechanism that is applied at every wavelet-domain location (i, \mathbf{k}) prior to the evaluation of the wavelet coefficients. This is achieved by using a structure tensor approach, which amounts to the determination of the direction that maximizes the local energy in the first component of the gradient wavelets in section 5.1.1. Some examples of fine-scale PCA wavelets for the *Lena* image and $N = 4, 5$ are shown in Figure 3. Interestingly, they happen to be rather similar to the corresponding sets of generalized Slepian wavelets shown below.

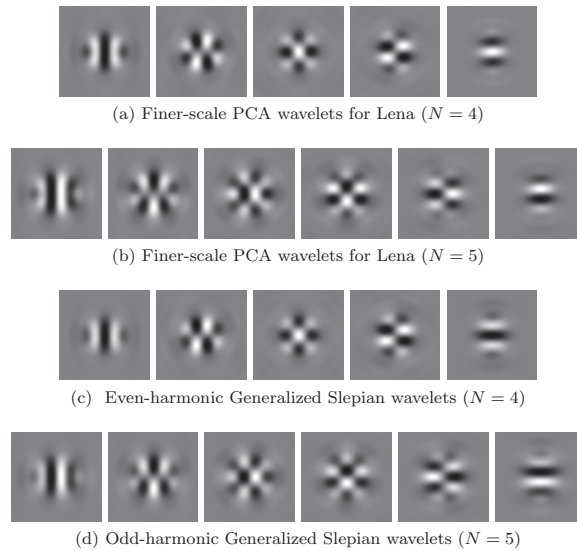


Figure 3. Examples of optimized steerable wavelets.

6. Experimental results. We now illustrate the ability of the proposed framework to reproduce state-of-the-art results in wavelet-based image processing. A significant aspect is that we are actually able to improve upon previous algorithms by optimizing the steerable wavelets for some given task.

6.1. Equiangular design for BLS-GSM denoising. The Bayes least squares Gaussian scale mixture (BLS-GSM) algorithm exploits Simoncelli's pyramid for removing noise in images [4]. It provides state-of-the-art performance among wavelet-based methods. The BLS-GSM relies on local wavelet-domain statistics and uses an elaborate processing to estimate the wavelet coefficients of the signal. While the original version uses Simoncelli's wavelets, the algorithm can be run on other equiangular configurations with $N + 1$ rotated filters equally distributed between 0 and π (as specified in Corollary 5.3 with $M' = N + 1$). To test the influence of the angular filter, we considered three choices for the primary wavelet function:

- $\hat{u}_1(\theta) \propto (-i)^N \cos(\theta)^N$, which corresponds to an N th-order Riesz wavelet and generates the standard Simoncelli pyramid.
- The most directional profile $\hat{u}_1(\theta)$ according to the prolate spheroidal design (cf. section 5.6).
- The first component of the wavelet-domain PCA (cf. section 5.7). Since the energy of the noise is constant across all wavelet channels, this is the filter that maximizes the SNR after proper steering.

We have applied the BLS-GSM to several images corrupted by additive white Gaussian noise with different standard deviation values σ_{noise} using the three equiangular frames with odd harmonics of degree $N = 7$.

The results in Table 1 indicate that the angular profile of the equiangular frame affects denoising performance. While the difference between the different methods varies with the noise level and the type of image, we found that the prolate spheroidal design consistently gave

Table 1

BLS-GSM denoising performance measured as the peak signal-to-noise ratio ($20 \log_{10}(255/\sigma_{\text{error}})$) in dB. The results are averaged over 50 independent noise samples (white and Gaussian) for three conditions specified by the standard deviation σ_{noise} .

Barbara			
$\sigma_{\text{noise}}/\text{PSNR}$	15 / 20.17	25 / 24.61	50 / 14.15
Simoncelli	31.89	29.14	25.47
Prolate	32.05	29.46	25.95
Signal-adapted	31.89	29.30	25.83
Lena			
$\sigma_{\text{noise}}/\text{PSNR}$	15 / 20.17	25 / 24.61	50 / 14.15
Simoncelli	34.02	31.83	28.82
Prolate	34.09	31.90	28.89
Signal-adapted	34.05	31.88	28.86
Texture			
$\sigma_{\text{noise}}/\text{PSNR}$	15 / 20.17	25 / 24.61	50 / 14.15
Simoncelli	28.33	25.55	22.10
Prolate	28.46	25.73	22.36
Signal-adapted	28.35	25.66	22.35

the best results. We believe that this is a consequence of its optimal angular localization, which also minimizes the residual correlation among the channels ($\nu_{\text{Prol}} = \|\mathbf{G} - \text{diag}(\mathbf{G})\|^2 / \|\mathbf{G}\|^2 = 0.163$ with Gram matrix $\mathbf{G} = \mathbf{U}\mathbf{U}^H$), as opposed to the Simoncelli frame channels, which are more correlated ($\nu_{\text{Sim}} = 0.566$). This property is favorable for the BLS-GSM algorithm which processes the wavelet channels independently. More surprising is the finding that the prolate solution also (slightly) outperforms the signal-adapted design (PCA). This may be explained by the learning procedure which uses rotation-invariant coefficients through steering, while only approximate rotation invariance is achieved by the BLS-GSM algorithm which does not accommodate steering. Another observation is that the gap between the different methods tends to grow with the noise energy, except for *Lena*, where the differences are marginal.

To test the influence of the number of channels, we run BLS-GSM denoising on the *Barbara* image corrupted by a white-Gaussian noise of standard deviation 50, while varying the order of the Riesz frame. The SNR results shown in Figure 4 are averaged over 100 trials. While the differences with the basic algorithm (Simoncelli) are negligible at low orders where the degrees of freedom are few ($N < 3$), they become significant as the number of channels increases. The performance eventually reaches a plateau, which happens around $N = 20$ for the Prolate filters. We therefore conclude that this latter design is the most advantageous computationally because it can yield better results with fewer channels. We also note that there is no major difference between using even and odd harmonics, which is somewhat surprising.

6.2. Curvelet-like wavelets and application to pattern separations. Directional systems of functions such as curvelets [17], contourlets [18], and shearlets [20] are often contrasted with conventional wavelets and presented as alternatives. In the following, we draw a parallel between these directional transforms, which we call curvelet-like frames for historical reasons, and equiangular generalized Riesz wavelets which can be judiciously combined to offer the




Figure 4. BLS-GSM denoising performance for the Barbara image, measured as the peak signal-to-noise ratio ($20 \log_{10}(255/\sigma_{\text{error}})$) in dB, as a function of the Riesz order. The results are averaged over 100 independent white-Gaussian noise samples with standard deviation 50.

same functionality. We then demonstrate the use of our wavelets for the separation of image patterns.

6.2.1. Construction of curvelet-like frames. Since curvelets and steerable wavelets share the same notion of multiresolution and directional analysis, we focus here on the task of reproducing the main features of the first ones in the proposed framework. We recall that the continuous-domain curvelet transform of a signal $f \in L_2(\mathbb{R}^2)$ is encoded in a set of coefficients $c(m, l, \mathbf{k})$ which is indexed with respect to scale ($m = 1, \dots, J$), orientation ($\theta_l = 2\pi l \cdot 2^{\lfloor m/2 \rfloor}$ with $l = 0, 1, \dots$ such that $0 \leq \theta_l < 2\pi$), and location ($\mathbf{k} = (k_1, k_2) \in \mathbb{Z}^2$) [39]. Using Parseval's relation, the coefficients are obtained by computing the frequency-domain inner products

$$(6.1) \quad c(m, l, \mathbf{k}) = \frac{1}{(2\pi)^2} \int U_m(\mathbf{R}_{\theta_l} \boldsymbol{\omega}) e^{-i\langle \mathbf{x}_{\mathbf{k}}^{(m,l)}, \boldsymbol{\omega} \rangle} \hat{f}(\boldsymbol{\omega}) d\boldsymbol{\omega},$$

where \mathbf{R}_{θ_l} is the rotation matrix for the angle θ_l , $\mathbf{x}_{\mathbf{k}}^{m,l} = \mathbf{R}_{\theta_l}^{-1}(k_1 \cdot 2^{-m}, k_2 \cdot 2^{-m/2})$ the corresponding sampling location, and U_m a smooth frequency window which has a polar-separable expression. The two implementations proposed in [39] are based on a digital coronization of the frequency plane on a Cartesian grid, which allows for some sampling rate reduction. Ultimately, this results in a discrete transform with a moderate redundancy factor ($\simeq 7.2$). The link with the continuous version of the transform, however, is partly lost (e.g., rotations are replaced by shearing operations). By contrast, shearlets are defined via the continuous-domain translation, dilation, and shearing of a single mother function [40]. The main motivation behind this design is that shearing is easier to discretize than rotation, which results in a more faithful digital implementation [22]. We propose here using the generalized Riesz transform paradigm to obtain a digital version of the continuous-domain curvelet transform at the expense of some redundancy.

Our generalized Riesz wavelets are suitable candidates for approximating curvelets since both the frequency window U_m and the Fourier transform of the wavelets are polar-separable




Figure 5. Example of basis functions for the real curvelet transform and the proposed Riesz-wavelet counterpart. Left: some real curvelet functions at different scales (in separated boxes) for 16 bands at the coarsest scale. Then, from left to right: Riesz-wavelet functions for 16, 12, and 8 bands at the coarsest scale, respectively.

[cf. (4.6)]. Equation (6.1) also suggests taking an equiangular Riesz transform to reproduce the equispaced rotation sequence of the frequency window U_m . Finally, to replicate the parabolic scaling of curvelets, we double the order of the generalized Riesz transform every two scales. For the angular shape of the frequency window, we propose two different configurations:

1. *Prolate spheroidal design with positive harmonics.* This design is closely related to the usual curvelet transform; it yields max-directional complex-valued wavelets whose Fourier transforms are thin, one-sided functions.
2. *Real prolate spheroidal design.* This design uses both positive and negative harmonics to produce real-valued basis functions, which are better suited for certain image processing tasks.

The so-defined frames are automatically tight as a result of the construction. Some examples of basis functions are shown in Figure 5 for the real-valued transforms. The main point is that we are able to closely reproduce the curvelet behavior within the proposed framework with the added benefit that our steerable wavelets have a better angular selectivity (max-directional design). The prolate-Riesz wavelets are typically more elongated in space, which allows us to achieve an equivalent angular discrimination with fewer channels (e.g., 16 bands for the curvelet transform vs. 8 bands for the max-directional wavelet transform). Moreover, reducing the number of channels does not affect rotation invariance since our Riesz wavelets are inherently steerable, unlike curvelet-like frames. Figure 6 illustrates the Fourier-domain partitioning achieved with complex-valued curvelets, real-valued shearlets, and the proposed steerable wavelets. The latter configuration is more favorable for directional analyses since the frequency responses of the curvelet-like wavelets are rotated versions of each other. Note that the shearlets' frequency profiles are not quite as sharp because the underlying basis functions are compactly supported in the space domain [41].

6.2.2. Pattern separation with directional frames. A nice application of wavelets and curvelet-like frames is sparsity-based source separation [42]. The idea is to separate signal components with different morphologies based on the premise that these are compactly represented in terms of distinct families of basis functions (frames). The formulation assumes a linear mixture model where the observed image $\mathbf{f} \in \mathbb{R}^K$ is decomposed as $\mathbf{f} = \mathbf{n} + \sum_{i=1}^I \mathbf{s}_i$, where \mathbf{n} is a disturbance term (noise) and each of the sources $\mathbf{s}_i = \mathbf{F}_i \mathbf{c}_i$ has a “sparse” representation (with coefficients $\mathbf{c}_i \in \mathbb{R}^{L_i}$) in some corresponding frame represented by the $K \times L_i$ matrix \mathbf{F}_i . The morphological component analysis (MCA) algorithm [42] separates the sources

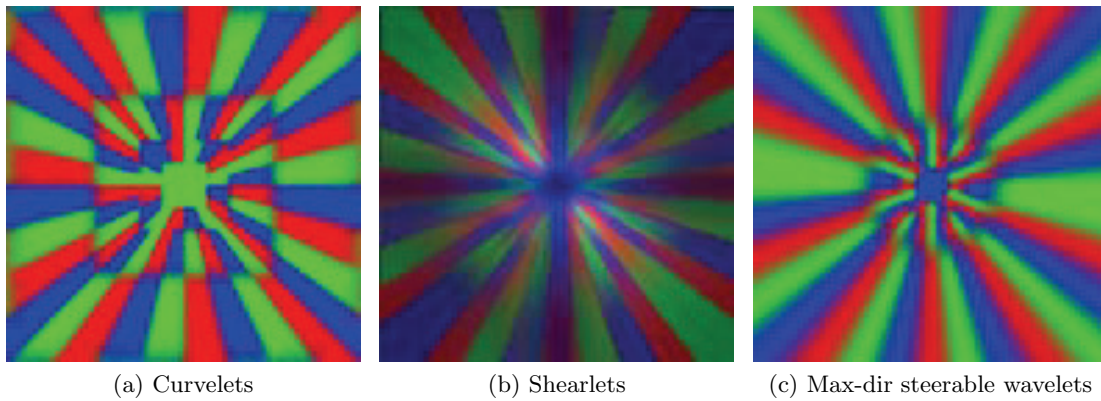


Figure 6. Fourier domain tiling achieved by the complex curvelet, compactly supported shearlets, and max-directional wavelet transforms. The bandpass region of each band has been sequentially colored in red, blue, and green, so as to highlight frequency coverage overlaps between bands.

by maximizing the following cost function:

$$(6.2) \quad \{\mathbf{c}_1^*, \dots, \mathbf{c}_I^*\} = \arg \min_{\mathbf{c}_1, \dots, \mathbf{c}_I} \left\| \mathbf{f} - \sum_{i=1}^I \mathbf{F}_i \mathbf{c}_i \right\|_2^2 + \lambda \sum_{i=1}^I \|\mathbf{c}_i\|_p^p,$$

where the right-hand-side regularization typically involves an ℓ_1 -norm or the pseudo- ℓ_0 -norm of the coefficients. MCA is an iterative coarse-to-fine algorithm. Each iteration t requires the computation of the residuals $\Delta \mathbf{f}_j = \mathbf{f} - \sum_{i=1, i \neq j}^I \mathbf{F}_i \mathbf{c}_i'$ (where the \mathbf{c}_i' are the current source estimates) and a solution update via the evaluation of $\mathbf{c}_j = \arg \min_{\mathbf{c}} \|\Delta \mathbf{f}_j - \mathbf{F}_j \mathbf{c}\|_2^2 + \lambda_t \|\mathbf{c}\|_p^p$ for each dictionary \mathbf{F}_j . When the frame is tight, the latter is achieved in one step by thresholding the projection $\mathbf{F}_j^T \Delta \mathbf{f}_j$ of $\Delta \mathbf{f}_j$ in the current dictionary. The relaxation parameter $\lambda_t \in \mathbb{R}^+$ is decreased over the iterations towards λ , and the sources are finally recovered as $\{\mathbf{s}_i = \mathbf{F}_i \mathbf{c}_i^*\}_{i=1}^I$.

Combining wavelets and curvelets is a typical choice for separating isotropic objects (e.g., stars) from more elongated patterns (e.g., galaxies) [42]. This technique has been applied quite successfully in astronomy and, more recently, in biological imaging [43, 44]. Here, we propose instead using specific combinations of generalized Riesz wavelets. Figure 7 displays an example of isotropic vs. elongated source separation obtained from a fluorescence micrograph of neuronal cells. The first function system is provided by the primary isotropic wavelet pyramid ($N = 0$), while the directional set is given by the curvelet-like transform with $M = 16$ fine-scale directional channels. The optimization was achieved by performing 100 MCA iterations and applying a hard threshold to solve the inner minimization problems ($p = 0$); we also imposed a positivity constraint on the reconstructed sources, which complies with the additive intensity model of fluorescence images. We see that the sources \mathbf{s}_1 and \mathbf{s}_2 contain exclusively isotropic and elongated features, respectively, which correspond to different biological objects (vesicles vs. axons). Automatic image analysis and biological event quantitation (such as particle detection) are therefore facilitated.

In this application, working with generalized Riesz wavelets can also bring design flexibility and computational benefits. For instance, we can neglect between-scale dependencies




Figure 7. Source separation results for a microscopy image using generalized Riesz-wavelet frames. The original fluorescence micrograph shows a mixture of linear (axons) and spotty features (vesicles). These are separated using MCA with a combination of isotropic ($N = 0$) and directional, curvelet-like ($N = 15$) wavelet dictionaries.

and perform the separation one scale at the time since the function systems share the same multiresolution structure/elementary Riesz atoms. Moreover, the curvelet-like transform can be replaced by a less redundant nonequiangular design, such as the full prolate spheroidal or PCA solutions. This is justifiable provided that we properly steer the transform and express the sparsity constraint in the locally oriented wavelet system. There is also the possibility of penalizing certain basis functions more than others. As an example, we show in Figure 8 the separation results obtained with a full PCA frame (trained on elongated features-only images) of lower order (even harmonics, $N = 4$). The quality of the separation is comparable to that shown in Figure 7, or when using the original curvelet transform, while the computational cost is lowered significantly. One promising future direction is to specifically adapt each dictionary to a source with a learning technique such as the proposed PCA-based procedure.

We have also performed experiments on synthetic images in order to compare the separation performance of the proposed wavelet frames with that of curvelets and shearlets. Using the benchmark proposed in [45, Figure 5], we found that steerable wavelets could essentially replicate the performance of compactly supported shearlets and yield better results than curvelets (data not shown). These examples using real and synthetic images are intended to demonstrate that the proposed wavelets constitute an attractive alternative to curvelets and shearlets for image analysis and processing. On the other hand, generalized Riesz wavelets are

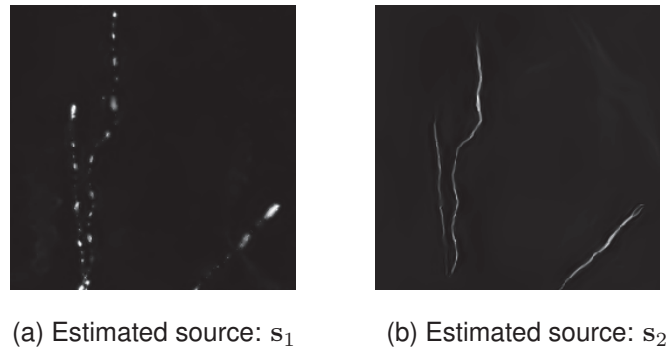


Figure 8. Source separation for the biological image in Figure 7(a) using MCA with a combination of isotropic ($N = 0$) and low-order PCA wavelet frames ($N = 4$). Rotation invariance is achieved by steering the coefficients according to the local orientation.

not as favorable for data compression—or asymptotically optimal for encoding cartoon-like images—because they are more redundant, which is the price to pay for steerability.

6.3. Discriminant frame learning. We have shown above that a combination of generalized Riesz transforms can be used to separate signals with different morphologies. An alternative approach is to design a single transform that discriminates between signal classes. We formulate this principle as the construction of the generalization matrix $\tilde{\mathbf{U}}$ that maximizes the difference of relative energy contribution of the two signals across wavelet channels. The relevant discriminant index is

$$(6.3) \quad \beta_{\mathbf{U}}(W_1, W_2) = \sum_{m=1}^M \frac{(\|W_1 \mathbf{u}_m\|^2 - \|W_2 \mathbf{u}_m\|^2)^2}{\|W_1 \mathbf{u}_m\|^2 \cdot \|W_2 \mathbf{u}_m\|^2} = \sum_{m=1}^M \beta_{\mathbf{u}_m}(W_1, W_2),$$

where W_1 and W_2 are $K \times M$ matrices containing the wavelet coefficients of the two signals for a real primary generalized Riesz transform with M channels at a given scale. It is shown in [46] that for such a problem the optimal linear transform $\tilde{\mathbf{U}} = \arg \max_{\mathbf{U}=[\mathbf{u}_1, \dots, \mathbf{u}_M]} \beta_{\mathbf{U}}(W_1, W_2)$ satisfies

$$(6.4) \quad \mathbf{C}_1 \tilde{\mathbf{u}}_m = \gamma_m \mathbf{C}_2 \tilde{\mathbf{u}}_m \quad \text{for } m = 1, \dots, M,$$

where \mathbf{C}_1 and \mathbf{C}_2 are the $M \times M$ covariances matrices of the coefficients for the two signal classes ($\mathbf{C}_1 = W_1^T W_1$ and $\mathbf{C}_2 = W_2^T W_2$ in our case), and $\gamma_m \in \mathbb{R}^+$. A particular solution of (6.4) is given by a matrix $\tilde{\mathbf{U}}$ that jointly diagonalizes \mathbf{C}_1 and \mathbf{C}_2 . Such a matrix exists for any symmetrical matrices \mathbf{C}_1 and \mathbf{C}_2 ; however, it is not unique and is generally nonorthogonal. In practice we rewrite the joint-diagonalization task as a generalized eigenvalue problem with symmetric-definite matrices, which we solve using Cholesky factorization and Schur decomposition [47].

We have used the proposed frame learning technique for discriminating the two textures from the Brodatz database shown in Figure 9. As the primary Riesz frame, we have used a steerable pyramid with five channels (even harmonics) and four scales. In a first experiment,




Figure 9. Two sample texture images from the Brodatz dataset. We focus on learning generalized Riesz wavelet frames that are able to discriminate them.

we learned a generalization matrix \mathbf{U} without steering the wavelet coefficients. The resulting basis functions for the first two scales are shown in Figure 10 along with the corresponding partial separation indices. At the finest scale, the first texture exhibits a dominant diagonal orientation (45 degrees), while the second is mainly composed of horizontal and vertical edges. The most discriminant filter ($\beta_{\mathbf{u}_1} = 1.83$) is a ridge-like pattern with 45 degree orientation. At the second scale, the second texture is composed mainly of ridge patterns oriented along the two diagonals. The 45 degree ridge pattern is thus no longer discriminating. Figure 10(b) shows that the best filter ($\beta_{\mathbf{u}_3} = 10.20$) is now a ridge-like function with a -45 degree orientation. This demonstrates the ability of the method to adapt to the texture classes at different scales. Table 2 documents the improvement in texture separation that is achieved by this type of learning technique.

There are also applications where one would like to factor out orientation. This can be achieved easily by steering the wavelet coefficients along the preferential local orientation prior to feature extraction, which makes the analysis rotation-invariant. We show the resulting filters in Figure 11. The main difference with Figure 10 is that the orientation of the new discriminating wavelets is no longer correlated with that of the initial pattern. For the first scale, the most discriminant filter ($\beta_{\mathbf{u}_M} = 1.89$) is orthogonal to the dominant local direction (horizontal axis). This helps in separating the two textures: the second contains many cross-like patterns with strong orthogonal components to the main direction, while the first is composed mainly of “pure” ridges. For the second scale, the most discriminating filter ($\beta_{\mathbf{u}_M} = 3.30$) is once again typical of the second texture: a step-like pattern, which is hardly found in texture 1. Finally, the results in Table 2 confirm that the optimized wavelets are better at discriminating the two textures than the standard equiangular design. We also see that a joint optimization is superior to a PCA-type design targeted to either one of the textures.

7. Conclusion. We have presented a general parameterization of 2D steerable wavelet frames. The scheme is interesting both conceptually and computationally. Since the con-

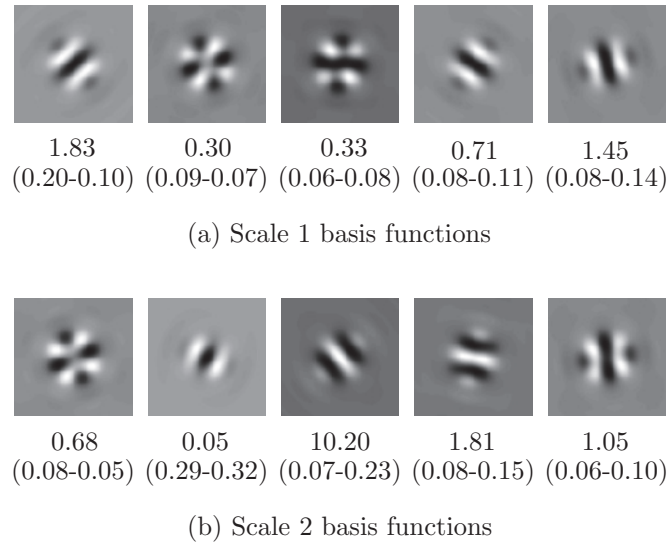


Figure 10. Discriminator filters for the two textures in Figure 9 and nonsteered wavelet coefficients ($N = 4$ with even harmonics). Below each image is the discriminant index $\beta_{\mathbf{u}_m}$ and the wavelet-coefficient standard deviation for the two textures (second line).

Table 2

Comparison of the ability of generalized wavelet families to discriminate the textures in Figure 9. The figure of merit defined by (6.3) is computed for the first two scales. The experiment was repeated twice using conventional vs. steered (rotation-invariant) feature extraction.

Coefficient steering	Frame	$\beta_{\mathbf{U}}(W1, W2)$	
		scale 1	scale 2
Nonsteered	Simoncelli	3.07	9.90
	Optimized	4.61	13.79
Steered	Simoncelli	3.43	6.77
	Signal-adapted (texture 1)	5.43	8.65
	Signal-adapted (texture 2)	4.19	6.52
	Optimized	5.57	8.67

straints on the wavelet shaping matrix \mathbf{U} are minimal, it facilitates the design of steerable wavelets while opening up new possibilities. It also provides a unifying perspective and a better global understanding of the choices and design compromise made in existing transforms, including Simoncelli's steerable pyramid. The fact that the wavelets are bandlimited with simple Fourier-domain expressions also suggests a generic decomposition algorithm where one first expands the signal in terms of circular harmonic wavelets using FFT-based filtering and then extracts the desired wavelet coefficients by simple matrix multiplication with \mathbf{U} . We also note that steering is best done in the circular harmonic domain where it amounts to a simple pointwise (complex) multiplication (self-steerability property). The image reconstruction algorithm applies the same steps in reverse order and amounts to the flow graph transpose of the analysis, thanks to the tight frame property (self-reversibility). Our generic MATLAB software is available publicly at <http://bigwww.epfl.ch/demo/circular-wavelets>.

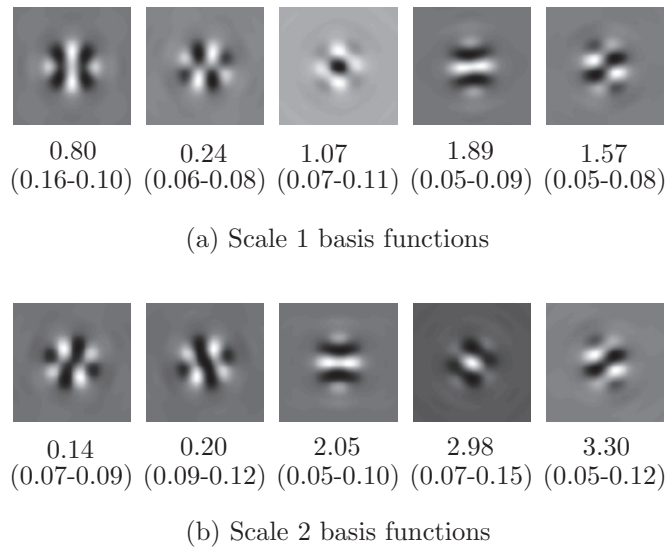


Figure 11. Discriminant filters for the two textures in Figure 9 and steered Riesz wavelet coefficients ($N = 4$ with even harmonics). Below each image is the discriminant index $\beta_{\mathbf{u}_m}$ and the wavelet-coefficient standard deviation for the two textures (second line).

We have also shown that the framework lends itself to the design of wavelets with optimized properties. In particular, we have constructed new prolate spheroidal wavelets whose angular profile is maximally localized. Our experimental results suggest that these are particularly well suited for applications such as denoising and directional feature extraction. It is actually remarkable that a mere change in shaping matrix \mathbf{U} can result in a notable improvement upon the state-of-the-art performance in wavelet processing. At the other extreme, we have observed that the canonical choice $\mathbf{U} = \mathbf{I}$ substantially degrades performance (data not shown), probably due to the fact that the circular harmonics have no angular selectivity at all. We take these as signs that the topic of wavelet design is not closed yet and that there is still room for improvement.

Appendix. Generalized Slepian sequences. Let us consider the 2π -periodic function $\hat{u}(\theta) = \sum_{n \in S} u_n e^{in\theta}$ that is described by its Fourier coefficients u_n over some finite indexing set S (e.g., $S = \{0, 1, \dots, N\}$) with $\text{Card}(S) = M$. We are interested in characterizing the optimal coefficients u_n such that the weighted-energy criterion

$$E_w(u) = \frac{1}{2\pi} \int_{-\pi}^{+\pi} |\hat{u}(\theta)|^2 w(\theta) d\theta$$

with $w(\theta) \geq 0$ is maximized (or minimized) subject to the normalization constraint $\|\hat{u}(\theta)\|_{L_2([-\pi, \pi])}^2 = \sum_{n \in S} |u_n|^2 = 1$. In his classical paper on discrete prolate spheroidal sequences (DPSS) [36], Slepian investigates the localization problem associated with the particular weighting function $w_0(\theta) = \text{rect}(\theta/(2\pi B))$, where $B < 1$ is a relative bandwidth parameter. With the above generalized statement of the problem, it is not difficult to extend Slepian's proof for an arbitrary nonnegative measurable function $w(\theta)$. The key idea is to

rewrite the weighted energy criterion as follows:

$$\begin{aligned}
 E_w(u) &= \frac{1}{2\pi} \int_{-\pi}^{+\pi} \sum_{n \in S} u_n e^{in\theta} \sum_{n' \in S} \bar{u}_{n'} e^{-in'\theta} w(\theta) d\theta \\
 &= \sum_{n \in S} \sum_{n' \in S} u_n \bar{u}_{n'} \frac{1}{2\pi} \int_{-\pi}^{+\pi} e^{-i(n'-n)\theta} w(\theta) d\theta \\
 (A.1) \quad &= \sum_{n' \in S} \sum_{n \in S} \bar{u}_{n'} W(n' - n) u_n = \mathbf{u}^H \mathbf{W} \mathbf{u},
 \end{aligned}$$

where the kernel,

$$W(x) = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-i\theta x} w(\theta) d\theta = \mathcal{F}_\theta \left\{ \frac{1}{2\pi} \text{rect} \left(\frac{\theta}{2\pi} \right) w(\theta) \right\} (x)$$

with $x \in \mathbb{R}$, is proportional to the Fourier transform of the restriction of $w(\theta)$ to the main period $\theta \in [-\pi, \pi]$. The notation for the right-hand side of (A.1) is as follows: \mathbf{u} is the M -dimensional coefficient vector with components $[\mathbf{u}]_n = u_n$, while \mathbf{W} is the $M \times M$ symmetric matrix whose entries are given by

$$[\mathbf{W}]_{n',n} = W(n' - n).$$

Since $w(\theta)$ is a nonnegative Borel measure, the kernel function $W(x)$ is positive-definite by Bochner's theorem [48]. The bottom line is that (A.1) specifies a positive-definite quadratic form (i.e., for all $\mathbf{u} \in \mathbb{C}^M$, $\mathbf{u}^H \mathbf{W} \mathbf{u} \geq 0$), irrespective of the index set S .

The numerical form of the problem is now the following: minimize $E_w(u) = \mathbf{u}^H \mathbf{W} \mathbf{u}$ subject to the condition $\mathbf{u}^H \mathbf{u} = 1$. This is a classical eigenvalue problem whose solution is given by

$$\mathbf{W} \mathbf{u} = \lambda \mathbf{u},$$

with $E_w(\mathbf{u}) = \lambda$. The important point for our purpose is that the corresponding eigenvectors define an orthogonal transformation whose extreme component achieves the best localization as characterized by λ_{\max} (or λ_{\min} , depending on the type of weighting).

The classical case, which yields the Slepian sequences, corresponds to the reproducing kernel:

$$W_0(x) = \frac{\sin(B\pi x)}{\pi x} \quad \xleftrightarrow{\mathcal{F}^{-1}} \quad w_0(\theta) = \text{rect} \left(\frac{\theta}{2\pi B} \right),$$

with $B < 1$.

In the present context of steerable wavelets, we have chosen an alternative variance-based measure of localization which may be extracted by means of the following symmetric kernel:

$$W_1(x) = \frac{(\pi^2 x^2 - 2) \sin(\pi x) + 2\pi x \cos(\pi x)}{\pi x^3} \quad \xleftrightarrow{\mathcal{F}^{-1}} \quad w_1(\theta) = \text{rect} \left(\frac{\theta}{2\pi} \right) \theta^2.$$

The most concentrated angular profile around $\theta_0 = 0$ (minimum variance solution) is the one that minimizes $E_{w_1}(u)$, while the least concentrated one maximizes it (which is the reverse of the classical ordering).

We are also introducing the functions $w_2(\theta)$ and $w_3(\theta)$, which implement a variation of the above quadratic weighting that is compatible with the angular periodicity condition $|\hat{u}(\theta)| = |\hat{u}(\theta + \pi)|$ of real-valued wavelets. $w_2(\theta)$ is described in section 5.6 and is designed to identify profiles that are simultaneously concentrated around zero and π . The third function is targeted towards the identification of angular profiles that are centered around $\pm\pi/2$ (instead of zero). The corresponding kernel is

$$W_3(x) = \frac{(\pi^2 x^2 - 8) \sin(\pi x) + 4\pi x + 4\pi x \cos(\pi x)}{4\pi x^3}$$

$$\xleftrightarrow{\mathcal{F}^{-1}} w_3(\theta) = \left(\theta - \frac{\pi}{2}\right)^2 \operatorname{rect}\left(\frac{\theta - \pi/2}{\pi}\right) + \left(\theta + \frac{\pi}{2}\right)^2 \operatorname{rect}\left(\frac{\theta + \pi/2}{\pi}\right).$$

Observe that the maximization of $E_{w_3}(u)$ will favor profiles that are centered and maximally concentrated around zero and π so that the classical ordering is restored.

REFERENCES

- [1] S. G. MALLAT, *A theory of multiresolution signal decomposition: The wavelet representation*, IEEE Trans. Pattern Anal. Mach. Intell., 11 (1989), pp. 674–693.
- [2] E. SIMONCELLI AND W. FREEMAN, *The steerable pyramid: A flexible architecture for multi-scale derivative computation*, in Proceedings of the International Conference on Image Processing, Vol. 3, 1995, pp. 444–447.
- [3] E. P. SIMONCELLI, W. T. FREEMAN, E. H. ADELSON, AND D. J. HEEGER, *Shiftable multiscale transforms*, IEEE Trans. Inform. Theory, 38 (1992), pp. 587–607.
- [4] 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., 12 (2003), pp. 1338–1351.
- [5] F. DENIS AND A. BASKURT, *Multidirectional curvilinear structures detection using steerable pyramid*, J. Electron. Imag., 13 (2004), pp. 756–765.
- [6] J. PORTILLA AND E. P. SIMONCELLI, *A parametric texture model based on joint statistics of complex wavelet coefficients*, Int. J. Comput. Vis., 40 (2000), pp. 49–70.
- [7] G. TZAGKARAKIS, B. BEFERULL-LOZANO, AND P. TSAKALIDES, *Rotation-invariant texture retrieval with Gaussianized steerable pyramids*, IEEE Trans. Image Process., 15 (2006), pp. 2702–2718.
- [8] C. Y. SU, Y. T. ZHUANG, L. HUANG, AND F. WU, *Steerable pyramid-based face hallucination*, Pattern Recognition, 38 (2005), pp. 813–824.
- [9] Z. LIU, K. TSUKADA, K. HANASAKI, Y. K. HO, AND Y. P. DAI, *Image fusion by using steerable pyramid*, Pattern Recognition Lett., 22 (2001), pp. 929–939.
- [10] I. DAUBECHIES, M. DEFRISE, AND C. DE MOL, *An iterative thresholding algorithm for linear inverse problems with a sparsity constraint*, Comm. Pure Appl. Math., 57 (2004), pp. 1413–1457.
- [11] H. RABBANI, *Image denoising in steerable pyramid domain based on a local Laplace prior*, Pattern Recognition, 42 (2009), pp. 2181–2193.
- [12] G. JACOVITTI AND A. NERI, *Multiresolution circular harmonic decomposition*, IEEE Trans. Signal Process., 48 (2000), pp. 3242–3247.
- [13] M. UNSER AND D. VAN DE VILLE, *Wavelet steerability and the higher-order Riesz transforms*, IEEE Trans. Image Process., 19 (2010), pp. 636–652.
- [14] S. HELD, M. STORAH, P. MASSOPUST, AND B. FORSTER, *Steerable wavelet frames based on the Riesz transform*, IEEE Trans. Image Process., 19 (2010), pp. 653–667.
- [15] M. UNSER, N. CHENOUEARD, AND D. VAN DE VILLE, *Steerable pyramids and tight wavelet frames in $L_2(\mathbb{R}^d)$* , IEEE Trans. Image Process., 20 (2011), pp. 2705–2721.
- [16] J. G. DAUGMAN, *Complete discrete 2-D Gabor transforms by neural networks for image-analysis and compression*, IEEE Trans. Acoust. Speech Signal Process., 36 (1988), pp. 1169–1179.

- [17] E. J. CANDÈS AND D. L. DONOHO, *New tight frames of curvelets and optimal representations of objects with piecewise C^2 singularities*, *Comm. Pure Appl. Math.*, 57 (2004), pp. 219–266.
- [18] M. N. DO AND M. VETTERLI, *The contourlet transform: An efficient directional multiresolution image representation*, *IEEE Trans. Image Process.*, 14 (2005), pp. 2091–2106.
- [19] V. VELISAVLJEVIC, B. BEFERULL-LOZANO, M. VETTERLI, AND P. DRAGOTTI, *Directionlets: Anisotropic multidirectional representation with separable filtering*, *IEEE Trans. Image Process.*, 15 (2006), pp. 1916–1933.
- [20] K. GUO AND D. LABATE, *Optimally sparse multidimensional representation using shearlets*, *SIAM J. Math. Anal.*, 39 (2007), pp. 298–318.
- [21] A. DA CUNHA, J. ZHOU, AND M. DO, *The nonsubsamped contourlet transform: Theory, design, and applications*, *IEEE Trans. Image Process.*, 15 (2006), pp. 3089–3101.
- [22] D. DONOHO, G. KUTYNIOK, M. SHAHRAM, AND X. ZHUANG, *A rational design of a digital shearlet transform*, in *Proceedings of the 9th International Conference on Sampling Theory and Applications (SampTA)*, 2011.
- [23] K. G. LARKIN, D. J. BONE, AND M. A. OLDFIELD, *Natural demodulation of two-dimensional fringe patterns. I. General background of the spiral phase quadrature transform*, *J. Opt. Soc. Amer. A*, 18 (2001), pp. 1862–1870.
- [24] K. G. LARKIN, *Natural demodulation of two-dimensional fringe patterns. II. Stationary phase analysis of the spiral phase quadrature transform*, *J. Opt. Soc. Amer. A*, 18 (2001), pp. 1871–1881.
- [25] M. UNSER, D. SAGE, AND D. VAN DE VILLE, *Multiresolution monogenic signal analysis using the Riesz-Laplace wavelet transforms*, *IEEE Trans. Image Process.*, 18 (2009), pp. 2402–2418.
- [26] E. STEIN AND G. WEISS, *Introduction to Fourier Analysis on Euclidean Spaces*, Princeton University Press, Princeton, NJ, 1971.
- [27] E. STEIN, *Singular Integrals and Differentiability Properties of Functions*, Princeton University Press, Princeton, NJ, 1970.
- [28] Q. SUN AND M. UNSER, *Left-inverses of fractional Laplacian and sparse stochastic processes*, *Adv. Comput. Math.*, 36 (2012), pp. 399–441.
- [29] J. P. WARD, K. N. CHAUDHURY, AND M. UNSER, *Decay properties of Riesz transforms and steerable wavelets*, *SIAM J. Imaging Sci.*, submitted.
- [30] 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.*, 24 (2003), pp. 907–928.
- [31] 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.*, 31 (2009), pp. 283–328.
- [32] W. T. FREEMAN AND E. H. ADELSON, *The design and use of steerable filters*, *IEEE Trans. Pattern Anal. Mach. Intell.*, 13 (1991), pp. 891–906.
- [33] D. VAN DE VILLE AND M. UNSER, *Complex wavelet bases, steerability, and the Marr-like pyramid*, *IEEE Trans. Image Process.*, 17 (2008), pp. 2063–2080.
- [34] M. FELSBERG AND G. SOMMER, *The monogenic signal*, *IEEE Trans. Signal Process.*, 49 (2001), pp. 3136–3144.
- [35] J. W. GOODMAN, *Introduction to Fourier Optics*, Roberts and Company, Greenwood Village, CO, 2005.
- [36] D. SLEPIAN, *Prolate spheroidal wave-functions, Fourier-analysis, and uncertainty—V: The discrete case*, *Bell System Tech. J.*, 57 (1978), pp. 1371–1430.
- [37] D. SLEPIAN, *Some comments on Fourier analysis, uncertainty and modeling*, *SIAM Rev.*, 25 (1983), pp. 379–393.
- [38] E. SIMONCELLI AND H. FARID, *Steerable wedge filters for local orientation analysis*, *IEEE Trans. Image Process.*, 5 (1996), pp. 1377–1382.
- [39] E. CANDÈS, L. DEMANET, D. DONOHO, AND L. YING, *Fast discrete curvelet transforms*, *Multiscale Model. Simul.*, 5 (2006), pp. 861–899.
- [40] C. S. GUO, Y. J. HAN, J. B. XU, AND J. P. DING, *Radial Hilbert transform with Laguerre-Gaussian spatial filters*, *Optic Lett.*, 31 (2006), pp. 1394–1396.
- [41] P. KITTIPOOM, G. KUTYNIOK, AND W.-Q. LIM, *Construction of compactly supported shearlet frames*, *Constr. Approx.*, 35 (2012), pp. 21–72.

- [42] J.-L. STARCK, D. ELAD, AND D. DONOHO, *Redundant multiscale transforms and their application for morphological component separation*, in Advances in Imaging and Electron Physics, Vol. 132, Elsevier, Amsterdam, 2004, pp. 287–348.
- [43] N. CHENOUEARD, S. VERNHETTES, I. BLOCH, AND J.-C. OLIVO-MARIN, *Morphological source separation for particle tracking in complex biological environments*, in Proceedings of the 19th International Conference on Pattern Recognition (ICPR 2008), 2008, pp. 1–4.
- [44] K. GOUSSET, E. SCHIFF, C. LANGEVIN, Z. MARIJANOVIC, A. CAPUTO, D. T. BROWMAN, N. CHENOUEARD, F. DE CHAUMONT, A. MARTINO, J. ENNINGA, J.-C. OLIVO-MARIN, D. MÄNNEL, AND C. ZURZOLO, *Prions hijack tunneling nanotubes for intercellular spread*, Nat. Cell Biol., 11 (2009), pp. 328–336.
- [45] G. KUTYNIOK AND W.-Q. LIM, *Image separation using wavelets and shearlets*, in Curves and Surfaces, Lecture Notes in Comput. Sci. 6920, J.-D. Boissonnat, P. Chenin, A. Cohen, C. Gout, T. Lyche, M.-L. Mazure, and L. Schumaker, eds., Springer, Berlin, Heidelberg, 2012, pp. 416–430.
- [46] M. UNSER, *Local linear transforms for texture measurements*, Signal Process., 11 (1986), pp. 61–79.
- [47] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., The Johns Hopkins University Press, Baltimore, MD, 1996.
- [48] J. STEWART, *Positive definite functions and generalizations, an historical survey*, Rocky Mountain J. Math., 6 (1976), pp. 409–434.