3D Steerable Wavelets in Practice

Nicolas Chenouard, *Member, IEEE*, and Michael Unser, *Fellow, IEEE*

***Abstract*— We introduce a systematic and practical design for steerable wavelet frames in 3D. Our steerable wavelets are obtained by applying a 3D version of the generalized Riesz transform to a primary isotropic wavelet frame. The novel transform is self-reversible (tight frame) and its elementary constituents (Riesz wavelets) can be efficiently rotated in any 3D direction by forming appropriate linear combinations. Moreover, the basis functions at a given location can be linearly combined to design custom (and adaptive) steerable wavelets. The features of the proposed method are illustrated with the processing and analysis of 3D biomedical data. In particular, we show how those wavelets can be used to characterize directional patterns and to detect edges by means of a 3D monogenic analysis. We also propose a new inverse-problem formalism along with an optimization algorithm for reconstructing 3D images from a sparse set of wavelet-domain edges. The scheme results in high-quality image reconstructions which demonstrate the feature-reduction ability of the steerable wavelets as well as their potential for solving inverse problems.**

***Index Terms*— 3D wavelet transform, edge detection, image reconstruction, monogenic signal, riesz transform, steerability.**

I. Introduction

THE MULTISCALE representations offered by three-dimensional (3D) wavelets are especially relevant in the context of biomedical imaging where the data is inherently 3D. In particular, wavelet-based techniques have been developed for MRI reconstruction [1], 3D deconvolution microscopy [2], and medical image coding [3]. Various 3D wavelets have also been designed and used for video processing [4]–[7]. While the use of separable wavelets is the simplest option available to practitioners, these decompositions have the drawback of lacking invariance with respect to shifts and rotations and of providing an uneven angular selectivity. During the past five years, researchers have investigated how a certain amount of redundancy makes possible the construction of wavelet-like representations of 3D data that are better suited to the geometrical features of images. These efforts have resulted in several versions of 3D curvelet transforms [8], [9], the 3D beamlet transform [10], surflets [11], and surfacelets [12]. The purpose of this paper is to present an alternative wavelet representation that uses steerable basis functions that can be rotated in 3D space in a data-adaptive fashion.

Steerable filters have been known for many years and have been successfully applied to a variety of feature-extraction tasks, including in 3D. Simoncelli and co-workers were able to combine steerable filters and multiresolution decomposition in the mid 1990s with their groundbreaking design of the 2D steerable pyramid [13], [14], which has been extremely popular in applications ever since. This representation owes its success to the fact that it is self-inverting (tight-frame property) with steerable basis functions that achieve a wavelet analysis that is truly rotation-invariant. Moreover, the angular selectivity of the wavelets can be controlled through the addition of orientation channels at the expense of a corresponding increase in redundancy. In practice, the design of a steerable wavelet frame has two main requirements: 1) the use of steerable functions as elementary building blocks and 2) self-invertibility to promote the tightness of the frame [15]. The steerable pyramid solution consists in taking steerable functions that are polar-separable in the Fourier domain and rotated versions of one another in an equiangular configuration. While this design ensures self-invertibility in 2D under fairly general conditions, it does not extend to 3D because of the lack of a systematic scheme to uniformly tile the sphere. The 3D steerable-wavelet design that we propose in this paper overcomes this issue by making use of a 3D version of the generalized Riesz transform [16], [17] to map an isotropic wavelet frame of $L_2(\mathbb{R}^3)$ into a new one that is steerable. Extending our previous work on 2D steerable wavelets [17] to 3D is feasible because the Riesz construction does not rely on equiangular basis functions. The proposed steerable transform offers a large palette of possibilities as its rotational components are parameterized by the shaping matrix of the generalized Riesz transform. The amount of redundancy of the whole transform stays reasonably low. It benefits from the low redundancy of the primary transform (8/7 asymptotically) and the ability to limit the number of bands of the Riesz-transform by controlling its order.

The contributions of this paper can be grouped under three primary headings.

1) Functional Specification of Steerable Wavelets: We are providing a general and systematic design of steerable wavelets in 3D with a complete decoupling of the multiresolution and steerable aspects of the transform. We have selected a multiresolution backbone that is the direct 3D counterpart of the radial bandpass filter (with a log-Gabor type profile) used in the steerable pyramid. The steerable part of the transform is obtained by application of the Nth-order generalized Riesz transform which has the remarkable property of preserving self-invertibility. The transform parameters are the Riesz order,

Manuscript received December 16, 2011; revised April 25, 2012; accepted June 6, 2012. Date of publication June 26, 2012; date of current version October 12, 2012. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Charles Creusere.

N. Chenouard is with the Biomedical Imaging Group, École Polytechnique Fédérale de Lausanne, Lausanne CH-1015, Switzerland, and also with the Center for Biomedical Imaging, Université de Lausanne, Lausanne CH-1015, Switzerland (e-mail: nicolas.chenouard@gmail.com).

M. Unser is with the Biomedical Imaging Group, École Polytechnique Fédérale de Lausanne, Lausanne CH-1015, Switzerland (e-mail: michael.unser@epfl.ch).

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

Digital Object Identifier 10.1109/TIP.2012.2206044

1057–7149/$31.00 © 2012 IEEE

which determines the number M of orientation channels, and the $(M \times M)$ wavelet-shaping matrix $\mathbf{U}$. Note that $\mathbf{U}$ is freely adjustable with the only constraint that it needs to be invertible, in the general case, and unitary for the construction of a tight frame.

2) Implementation: Our steerable wavelet transform is specified in the continuous domain. To address the issue of its proper discretization, we have projected the input signal onto the space of bandlimited functions. This results in the specification of a fast digital filterbank algorithm that is perfectly reversible. We are providing a generic FFT-based implementation[1] that can handle arbitrary orders and wavelet shaping matrices.

3) Application to 3D Image Processing: We are illustrating specific advantages of the proposed decomposition for the processing of biomedical images and are proposing novel algorithms. First, we have made use of the connection between the Riesz and the directional Hilbert transforms to define a truly 3D version of a wavelet-domain monogenic analysis, as an extension of prior works of our team and others in 2D [16], [18]. This scheme provides a multiscale description of the local orientation, amplitude, and phase of the 3D signal and turns out to be particularly well-suited to the characterization of directional structures such as sheets and filaments. Second, we exploit the differential interpretation of our wavelets as multiscale derivatives to design a 3D wavelet-domain contour-detection algorithm that is inspired by Canny's edge detector. We have also proposed a novel scheme for reconstructing the image from its multiscale edge map which is encoded by the local extrema of the wavelet transform. In doing so, we are able to demonstrate the pertinence of edges for the characterization of images, in the spirit of Marr's primal sketch [19], and to extend those ideas to 3D. The proposed algorithm also serves as an illustration of the successful application of our framework to the solution of inverse problems.

The paper is organized as follows: We start with a brief introduction on steerable functions and the presentation of the generalized Riesz transform. This constitutes the mathematical foundation of our approach. In Section III, we give a detailed description of the principles and digital implementation of the new 3D steerable wavelet transform. In Section IV, we introduce the 3D monogenic signal analysis and illustrate its capability for a set of real microscope images. Finally, in Section V, we demonstrate the benefits of the Riesz-wavelet transform for image reconstruction from edge coefficients, before discussing our results in the concluding part of the paper.

II. Preliminaries

A. From 2D to 3D Steerable Wavelets

In their seminal work [20], Freeman and Adelson have shown that a 2D function f is steerable if its polar representation can be expanded as

$$f(r,\phi) = \sum_{n=-N}^{N} a_n(r) \mathrm{e}^{\mathrm{j}n\phi} \tag{1}$$

with $0 \leq N < \infty$, $a_n(r) \in \mathbb{C}$, $\mathrm{j}^2 = -1$, and polar coordinates $(r,\phi) \in \mathbb{R}^+ \times]-\pi, \pi]$. Indeed, when (1) holds, it is not hard to show that there exists a set of primary functions $\{g_m\}_{m=1\ldots M}$ and a set of interpolating functions $\{\kappa_m\}_{m=1\ldots M}$ with $M < \infty$ such that

$$f(r,\phi+\phi_0) = \sum_{m=1}^{M} \kappa_m(\phi_0) g_m(r,\phi) \tag{2}$$

for any $\phi_0 \in \mathbb{R}$. In particular, Freeman and Adelson have proposed to use the rotated versions of f in M directions with equispaced angles as the primary functions

$$g_m(r,\phi) = f(r,\phi+\pi(m-1)/M). \tag{3}$$

The key consequence of (2) is that the versions of f rotated by any arbitrary angle can be efficiently computed by linear combination of the primary functions, which need to be computed once only. This property has lead Simoncelli and his colleagues to use steerability as the key concept for the construction of rotation-invariant wavelets in 2D [14]. Their steerable pyramid [13] relies on a polar separable design consisting of two parts: 1) a set of primary wavelet functions which are isotropic and define a tight frame of $L_2(\mathbb{R}^2)$ for the radial part, 2) an equiangular steerable filterbank as in (3) for the polar part. The steerable filterbank is carefully chosen such that self-invertibility of the whole construction holds.

In 3D, the equiangular concept for steerable wavelets is not well-defined and has proved difficult to apply because of the self-invertibility constraints [21], except for the special case $M = 3$ [22]. As as consequence, we have adopted a more general steerable-wavelet design that does not require equiangular functions for $M \geq 3$. To do so, we state our definition of 3D steerability as the direct extension to 3D of (2) given by

$$f(\mathbf{R}\boldsymbol{x}) = \sum_{m=1}^{M} \kappa_m(\mathbf{R}) g_m(\boldsymbol{x}) \tag{4}$$

for any vector of Cartesian coordinates $\boldsymbol{x} = (x_1, x_2, x_3)$ and any rotation matrix $\mathbf{R}$ in three dimensions.

B. Riesz-Wavelet Transform

A flexible alternative to the steerable pyramid construction is to replace the equiangular filterbank by a generalized Riesz transform of high order [17]. We introduce hereafter some useful notations and the fundamentals of this technique, which we extend to 3D in Section III.

1) Notations: In the sequel, we use a bold font for any column vector $\boldsymbol{x} = (x_1, \ldots, x_d) \in \mathbb{R}^d$. The notation $\|\cdot\|$ is used for the ℓ_2 norm of $\mathbb{R}^d$, with $\|\boldsymbol{x}\| = \|\boldsymbol{x}\|_2 = (\sum_{i=1}^{d} |x_i|^2)^{1/2}$. We also consider d-dimensional multiindex vectors of the form $\mathbf{n} = (n_1, \ldots, n_d)$ whose entries n_i are nonnegative integers. We then define the following multiindex operations and operators.

[1]These tools are made available via a dedicated open-source toolbox, available from http://bigwww.epfl.ch/demo/steerable-wavelets-3D/.

1) Sum of components $|\mathbf{n}| = \sum_{i=1}^{d} n_i = N$.
2) Factorial $\mathbf{n}! = n_1! n_2! \cdots n_d!$.
3) Exponentiation $\boldsymbol{x}^{\mathbf{n}} = x_1^{n_1} \cdots x_d^{n_d}$.
4) Higher-order partial derivative of a function $f : \mathbb{R}^d \mapsto \mathbb{R}$

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

5) Fourier transform of $f \in L_2(\mathbb{R}^d)$

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

for $\boldsymbol{\omega} \in \mathbb{R}^d$.

These notations allow us to describe the multinomial theorem concisely as

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

which involves a summation over $\binom{N+d-1}{d-1}$ distinct monomials of the form $\boldsymbol{x}^{\mathbf{n}} = x_1^{n_1} \cdots x_d^{n_d}$ with $n_1 + \cdots + n_d = N$.

2) Generalized Riesz Transforms: The Riesz transform of a function $f(\boldsymbol{x})$ of $L_2(\mathbb{R}^d)$ is a scalar-to-vector transformation $\boldsymbol{\mathcal{R}} f(\boldsymbol{x}) = (\mathcal{R}_1 f(\boldsymbol{x}), \ldots, \mathcal{R}_d f(\boldsymbol{x}))$. The Riesz-operator component $\mathcal{R}_i : \mathbb{R}^d \mapsto \mathbb{R}$ is linear, space-invariant, and characterized by the frequency response

$$\widehat{\mathcal{R}_i f}(\boldsymbol{\omega}) = -\mathrm{j} \frac{\omega_i}{\|\boldsymbol{\omega}\|} \hat{f}(\boldsymbol{\omega}). \tag{5}$$

We readily check that the Riesz transform is self-reversible:

$$\widehat{\boldsymbol{\mathcal{R}}^{\star} \boldsymbol{\mathcal{R}} f}(\boldsymbol{\omega}) = \frac{(-\mathrm{j}\boldsymbol{\omega})^H (-\mathrm{j}\boldsymbol{\omega})}{\|\boldsymbol{\omega}\|^2} \hat{f}(\boldsymbol{\omega}) = \hat{f}(\boldsymbol{\omega})$$

with $(-\mathrm{j}\boldsymbol{\omega})^H = \mathrm{j}\boldsymbol{\omega}^T$ the Hermitian transpose of the complex-valued vector $(-\mathrm{j}\boldsymbol{\omega})$. In [23], a high-order version of the Riesz transform is proposed for $\mathbb{R}^d$ by iterating the Riesz-component operators. The Riesz transform $\boldsymbol{\mathcal{R}}^N$ of order $N \in \mathbb{N}$ is a scalar-to-vector transformation with

$$p(N, d) = \begin{pmatrix} N + d - 1 \\ d - 1 \end{pmatrix}$$

components. The Riesz-operator components are defined as

$$\boldsymbol{\mathcal{R}}^{\boldsymbol{n}} = \sqrt{\frac{N!}{\boldsymbol{n}!}} \mathcal{R}_1^{n_1} \mathcal{R}_2^{n_2} \cdots \mathcal{R}_d^{n_d}$$

for every $\boldsymbol{n} \in \mathbb{N}^d$ such that $|\boldsymbol{n}| = N$.

Using multiindex notations, the frequency response of each component finds the compact expression

$$\widehat{\boldsymbol{\mathcal{R}}^{\boldsymbol{n}}}(\boldsymbol{\omega}) = (-\mathrm{j})^N \sqrt{\frac{N!}{\boldsymbol{n}!}} \frac{\boldsymbol{\omega}^{\boldsymbol{n}}}{\|\boldsymbol{\omega}\|^N}.$$

Due to the appropriate normalization factor for each component, one can now apply the multinomial theorem to prove self-invertibility of the high-order Riesz transform. It shows that the high-order Riesz transform can map any tight frame of $L_2(\mathbb{R}^d)$ into another one [22], [23].

The framework can be extended by parameterizing the high-order Riesz transform with an $(M \times M)$ complex-valued matrix $\mathbf{U}$ (with $M = p(N, d)$) to build the family of generalized Riesz transforms defined as $\boldsymbol{\mathcal{R}}_{\mathbf{U}} = \mathbf{U} \boldsymbol{\mathcal{R}}^N$ [17]. It is easily shown that the self-invertibility property is preserved as long as $\mathbf{U}$ is unitary. The benefits of adapting the generalization matrix $\mathbf{U}$ to a given task was demonstrated in [17] for the image-denoising case. It is worth noting that the generalized Riesz transform is steerable, which means that the impulse responses of its components can be simultaneously oriented in any direction by forming suitable linear combinations [17], [23]. Specifically, let $\mathbf{R}$ be a $d \times d$ spatial rotation matrix. Then, we have that

$$\boldsymbol{\mathcal{R}}_{\mathbf{U}} f(\mathbf{R}\boldsymbol{x}) = \boldsymbol{\mathcal{R}}_{\mathbf{U}\mathbf{S}_{\mathbf{R}}} f(\boldsymbol{x}) \tag{6}$$

with $\mathbf{S}_{\mathbf{R}}$ the appropriate $(M \times M)$ steering matrix for $\mathbf{R}$.

3) 2D Riesz-Wavelet Pyramid: The ability of Riesz transforms to map any tight frame of $L_2(\mathbb{R}^2)$ into another tight frame of $L_2(\mathbb{R}^2)$ has been exploited for the design of 2D Riesz-wavelet pyramids, first with spline-based wavelets for the high-order transform [23], and then with an isotropic and non-separable wavelet pyramid for the generalized version [17]. In the latter case, the isotropy of the wavelet functions yields a perfect steerability of the Riesz-wavelet basis functions. As compared to Simoncelli's pyramid, the Riesz-wavelet transform enjoys a greater degree of flexibility for the polar part of the basis functions, as it does not require the individual wavelets to be rotated versions of the same primary profile. Shaping these functions can be conveniently achieved with the design of the generalization matrix $\mathbf{U}$.

III. Design of 3D Steerable Wavelets

We propose a 3D steerable wavelet transform consisting of the composition of two operators: a 3D isotropic wavelet transform that performs multiscale decomposition and gives vanishing moments, together with a 3D Riesz transform that brings directionality and steerability. We give the continuous-domain characterization of the transform and we make explicit the link with a digital implementation for bandlimited signals of $L_2(\mathbb{R}^3)$.

A. Primal Isotropic 3D Wavelet Frame

The multiresolution backbone of our method is the tight wavelet frame $\{\psi_{i,\boldsymbol{k}}\}_{i \in \mathbb{Z}, \boldsymbol{k} \in \mathbb{Z}^3}$, whose basis functions in $L_2(\mathbb{R}^3)$ are generated by suitable dilations and translations of a single mother wavelet $\psi(\boldsymbol{x})$. The normalized wavelets at scale i and translation parameter $\boldsymbol{k}$ are given by

$$\psi_{i,\boldsymbol{k}}(\boldsymbol{x}) = \psi_i(\boldsymbol{x} - 2^i \boldsymbol{k}) \text{ with } \psi_i(\boldsymbol{x}) = 2^{-3i/2} \psi(\boldsymbol{x}/2^i). \tag{7}$$

The present requirement is that the mother wavelet $\psi(\boldsymbol{x})$ be isotropic. It has been shown by various authors that this can be achieved by selecting a function that is bandlimited radially [22], [24]–[26]. The specific Fourier-domain wavelet-design constraints are given in Proposition 1.

Proposition 1: Let $h(\omega)$ be a radial frequency profile such that

Condition 1): $h(\omega) = 0, \forall \omega > \pi$

Condition 2): $\sum_{i\in\mathbb{Z}} |h(2^i\omega)|^2 = 1$

Condition 3): $\left.\frac{\mathrm{d}^n h(\omega)}{\mathrm{d}\omega^n}\right|_{\omega=0} = 0,$ for $n = 0, \ldots, N.$

Then, the mother wavelet ψ whose 3D Fourier transform is given by

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

generates a tight wavelet frame of $L_2(\mathbb{R}^3)$. The corresponding basis functions, which are specified by (7), are isotropic with vanishing moments up to order N.

It is instructive to have a closer look at the underlying multiresolution decomposition mechanism. Condition 1) ensures that the bandpass signals

$$r_i(\boldsymbol{x}) = \langle f, \psi_i(\cdot - \boldsymbol{x})\rangle = (f * \psi_i)(\boldsymbol{x})$$

which are obtained by performing a continuous wavelet analysis at the scale i, are bandlimited. By Shannon's sampling theorem, r_i is completely specified by its uniform samples with sampling period $T_i = 2^i$, which happen to be the wavelet coefficients of f at scale i computed as

$$w_{i,\boldsymbol{k}} = \langle f, \psi_{i,\boldsymbol{k}}\rangle = r_i(\boldsymbol{x})|_{\boldsymbol{x}=T_i\boldsymbol{k}}\,.$$

This leads to the Shannon reconstruction formula

$$r_i(\boldsymbol{x}) = \sum_{\boldsymbol{k}\in\mathbb{Z}^3} w_{i,\boldsymbol{k}}\, \frac{1}{T_i^3}\mathrm{sinc}\left(\frac{\boldsymbol{x} - T_i\boldsymbol{k}}{T_i}\right).$$

Condition 2) guarantees that the dyadic sequence of wavelets ψ_i provides a full coverage of the 3D frequency domain since

$$\sum_{i\in\mathbb{Z}} |\hat{\psi}_i(\boldsymbol{\omega})|^2 = 1 \tag{8}$$

which is the key to ensuring the tight-frame property [27]. In fact, (8) is equivalent to

$$\sum_{i\in\mathbb{Z}} \big(\psi_i * \psi_i * f\big)(\boldsymbol{x}) = \sum_{i\in\mathbb{Z}} \big(\psi_i * r_i\big)(\boldsymbol{x}) = f(\boldsymbol{x}). \tag{9}$$

Because ψ_i is bandlimited and the rescaled sinc function acts an ideal lowpass filter, let us now write that $\psi_i * \frac{1}{T_i^3}\mathrm{sinc}(\cdot/T_i) = \psi_i$. We find that

$$(\psi_i * r_i)(\boldsymbol{x}) = \sum_{\boldsymbol{k}\in\mathbb{Z}^3} w_{i,\boldsymbol{k}}\, \psi_{i,\boldsymbol{k}}(\boldsymbol{x}).$$

By inserting this result into (9), we end up with

$$f(\boldsymbol{x}) = \sum_{i\in\mathbb{Z}} \sum_{\boldsymbol{k}\in\mathbb{Z}^3} \langle f, \psi_{i,\boldsymbol{k}}\rangle \psi_{i,\boldsymbol{k}}(\boldsymbol{x})$$

which is a wavelet expansion that holds for any finite energy function $f \in L_2(\mathbb{R}^3)$. Likewise, we can represent a bandlimited function using a scale-truncated wavelet expansion that yields $J \in \mathbb{N}$ subsampled bands and a low-frequency residual. In the 3D case, the overall redundancy is $1 + \frac{1}{8} + \cdots + (\frac{1}{8})^J = \frac{8}{7}(1 - (\frac{1}{8})^{J+1}) \xrightarrow{J\to\infty} \frac{8}{7}$ over the canonical representation in terms of its sampled values.

The fact that the wavelet frame is tight is also equivalent to the existence of a constant A such that (cf. [28], [29])

$$A \cdot \|f\|_{L_2(\mathbb{R}^3)} = \sum_i \sum_{\boldsymbol{k}\in\mathbb{Z}^3} |\langle f, \psi_{i,\boldsymbol{k}}\rangle|^2.$$

Finally, Condition 3) implies that

$$\left.\frac{\partial^{|\boldsymbol{n}|}}{\partial^{n_1}\omega_1\partial^{n_2}\omega_2\partial^{n_3}\omega_3}\hat{\psi}(\boldsymbol{\omega})\right|_{\boldsymbol{\omega}=\boldsymbol{0}} = 0$$

and

$$(-\mathrm{j})^{|\boldsymbol{n}|}\int_{\mathbb{R}^3} x_1^{n_1}x_2^{n_2}x_3^{n_3}\psi(\boldsymbol{x})\mathrm{d}\boldsymbol{x} = 0$$

for all multiindices $\boldsymbol{n}$, with $|\boldsymbol{n}| < N$. Since the design is in the frequency domain, it is actually possible to obtain N arbitrarily large ($N \to \infty$) by imposing that $h(\omega) = 0$ for $\omega \in [-\epsilon, \epsilon]$ in a finite interval ($0 < \epsilon < \infty$) around the origin.

The simplest choice of radial frequency profile that fulfills the conditions in Proposition 1 is $h(\omega) = \mathrm{rect}(\frac{\omega - 3\pi/4}{\pi/2})$ (Shannon ideal-bandpass wavelet), which yields a Bessel-type wavelet that can be characterized analytically [25]. Another prominent solution is the filter h that is implemented in the popular version of the steerable pyramid described in [14], with

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

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. We illustrate in Figure 1 the frequency profile of these wavelet functions.

B. High-Order Riesz Transform in 3D

For the polar part of the transform, we rely on the generalized Riesz transforms in 3D. For the sake of simplicity, we investigate the standard Nth order case (*i.e.*, $\mathbf{U}$ is the $(M \times M)$ identity matrix). The generalized version is straightforward to obtain by taking linear combination of the Riesz basis functions. In 3D, the Riesz tranform of order $N \in \mathbb{N}$ is composed of $M = (N + 2)(N + 1)/2$ components which we denote by $\mathcal{R}^{\boldsymbol{n}}$, with $\boldsymbol{n} = (n_1, n_2, n_3) \in \mathbb{N}^3$ such that $|\boldsymbol{n}| = N$. Each of these convolutive operators has the frequency response

$$\widehat{\mathcal{R}^{\boldsymbol{n}}}(\boldsymbol{\omega}) = (-\mathrm{j})^N \sqrt{\frac{N!}{n_1!n_2!n_3!}}\, \frac{\omega_1^{n_1}\omega_2^{n_2}\omega_3^{n_3}}{(\omega_1^2 + \omega_2^2 + \omega_3^2)^{N/2}} \tag{10}$$

with $\boldsymbol{\omega} = (\omega_1, \omega_2, \omega_3) \in \mathbb{R}^3$ the 3D pulsation vector. Expression (10) highlights the link of the 3D Riesz transform with more-standard operators. On one hand, the fractional Laplacian $(-\Delta)^{s/2}$ of order $s \in \mathbb{R}$ is defined in the Fourier domain in the sense of distributions as

$$(-\Delta)^{s/2} f(\boldsymbol{x}) \overset{\mathcal{F}}{\longleftrightarrow} \|\boldsymbol{\omega}\|^s \hat{f}(\boldsymbol{\omega}).$$

On the other hand, the Fourier transform of $\partial^{\boldsymbol{n}} f(\boldsymbol{x})$ is

$$\partial^{\boldsymbol{n}} f(\boldsymbol{x}) \overset{\mathcal{F}}{\longleftrightarrow} (j\boldsymbol{\omega})^{\boldsymbol{n}} \hat{f}(\boldsymbol{\omega}).$$

The frequency response of the 3D Riesz transform thus shows that it can be interpreted as the 3D partial derivative of the integral operator $(\Delta)^{-N/2}$, which turns out to be an isotropic smoothing filter. The 3D Riesz transform may thus be viewed as a high-order gradient operator which is regularized for self-invertibility. We exploit this view to obtain an explicit space-domain expression for the impulse response $\mathcal{R}^{\boldsymbol{n}}\{\delta\}(\boldsymbol{x})$ of each component.

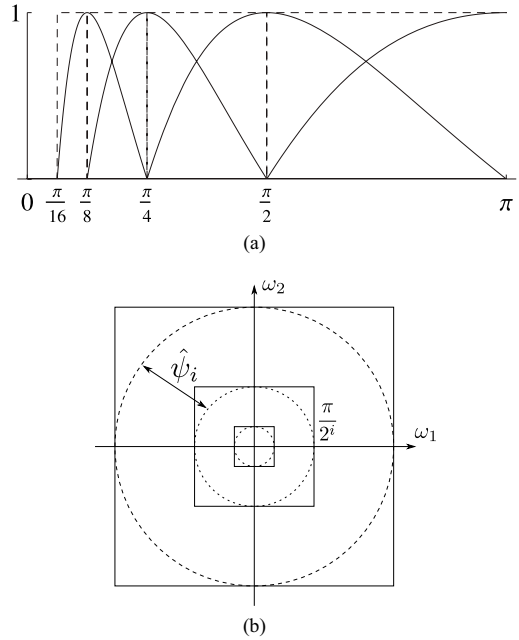

Fig. 1. Frequency-domain partition achieved by the primary isotropic wavelet transform. (a) $h(\omega)$ for Shannon's wavelets (dashed lines) and Simoncelli's wavelets (solid lines). (b) Support of each Shannon's wavelet function in 2D is delineated by dashed lines. The space-domain subsampling operations, which restrict the frequency plane to the support of each wavelet function, are shown with boxes.

The space-domain expression of the impulse response of $(-\Delta)^{s/2}$ is given by

$$\|\boldsymbol{\omega}\|^s \overset{\mathcal{F}}{\longleftrightarrow} \begin{cases} C_{d,s}\|\boldsymbol{x}\|^{-(s+d)}, \\ \quad \text{if } |s| < d \text{ or if } s+d \text{ is odd} \\ C'_{d,s}\|\boldsymbol{x}\|^{-(s+d)} \log \|\boldsymbol{x}\|, \\ \quad \text{otherwise,} \end{cases}$$

where d is the dimension of the space. This is a standard result from the theory of distributions [30]. For $s < 0$, the real constants $C_{d,s}$ and $C'_{d,s}$ are given by

$$C_{d,s} = \frac{2^s \Gamma\left(\frac{s+d}{2}\right)}{\pi^{\frac{d}{2}} \Gamma\left(-\frac{s}{2}\right)}, \qquad C'_{d,s} = \frac{2^{s+1}(-1)^{-\frac{s+d}{2}+1}}{\pi^{\frac{d}{2}} \Gamma\left(-\frac{s}{2}\right) \Gamma\left(-\frac{s+d}{2}+1\right)}.$$

The impulse response of the $\boldsymbol{n}$th order Riesz transform can thus be expressed as

$$\mathcal{R}^{\boldsymbol{n}}\{\delta\}(\boldsymbol{x}) = \begin{cases} (-1)^N C_{d,-N}\partial^{\boldsymbol{n}}\{\|\cdot\|^{-(d-N)}\}(\boldsymbol{x}), \\ \quad \text{if } N < d \text{ or if } N+d \text{ is odd} \\ (-1)^N C'_{d,-N}\partial^{\boldsymbol{n}}\{\|\cdot\|^{-(d-N)} \log \|\cdot\|\}(\boldsymbol{x}), \\ \quad \text{otherwise.} \end{cases}$$

We give in Table I the formulas for the 3D Riesz filters of order 1, 2, and 3. The impulse response is a homogeneous distribution which is not a conventional function around the origin. It is non compactly supported and decays like $\|\boldsymbol{x}\|^{-3}$. We shall see in Section III-C that the transform is more conveniently implemented in the frequency domain.

Using a high-order Riesz transform allows characterizing 3D structures of various morphologies. In Figure 2 we give examples of filters that are adapted to edges (a), sheets (b), filaments (d), and other complex shapes (c).

TABLE I
IMPULSE AND FREQUENCY RESPONSES OF SOME RIESZ TRANSFORMS OF ORDER N IN 3D

N	$\boldsymbol{n}$	$\widehat{\mathcal{R}^{\boldsymbol{n}}}(\boldsymbol{\omega})$	$\mathcal{R}^{\boldsymbol{n}}\{\delta\}(\boldsymbol{x})$
1	(1, 0, 0)	$\frac{-j\omega_1}{\|\boldsymbol{\omega}\|}$	$\frac{x_1}{\pi^2\|\boldsymbol{x}\|^4}$
2	(2, 0, 0)	$\frac{(-j\omega_1)^2}{\|\boldsymbol{\omega}\|^2}$	$-\frac{-2x_1^2+x_2^2+x_3^2}{4\pi\|\boldsymbol{x}\|^5}$
	(1, 1, 0)	$\frac{(-j\omega_1)(-j\omega_2)}{\|\boldsymbol{\omega}\|^2}$	$\frac{3x_1x_2}{4\pi\|\boldsymbol{x}\|^5}$
3	(3, 0, 0)	$\frac{(-j\omega_1)^3}{\|\boldsymbol{\omega}\|^3}$	$\frac{x_1(x_1^2-3x_2^2-3x_3^2)}{\pi^2\|\boldsymbol{x}\|^6}$
	(2, 1, 0)	$\frac{(-j\omega_1)^2(-j\omega_2)}{\|\boldsymbol{\omega}\|^3}$	$-\frac{x_2(-3x_1^2+x_2^2+x_3^2)}{\pi^2\|\boldsymbol{x}\|^6}$
	(1, 1, 1)	$\frac{(-j\omega_1)(-j\omega_2)(-j\omega_3)}{\|\boldsymbol{\omega}\|^3}$	$-\frac{4x_1x_2x_3}{\pi^2\|\boldsymbol{x}\|^6}$

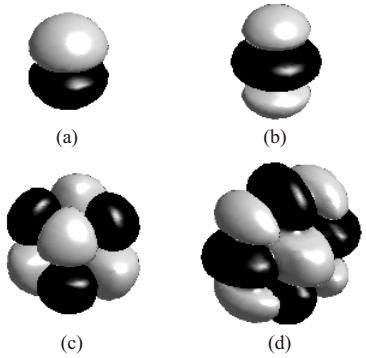

Fig. 2. (a)–(d) Isosurface representation of some Riesz-wavelet components of different orders. The surface is colored white if the component value is positive, and black otherwise. Simoncelli's wavelet [14] is used for the primary wavelet function ψ. (a) $\mathcal{R}^{(0,0,1)}\psi$. (b) $\mathcal{R}^{(0,0,2)}\psi$. (c) $\mathcal{R}^{(1,1,1)}\psi$. (d) $\mathcal{R}^{(0,2,2)}\psi$.

C. 3D Riesz-Wavelet Frames

We combine the 3D Riesz transform with the isotropic wavelet frame presented in Section III-A to obtain steerable wavelet frames in 3D. Remarkably, because the Riesz transform is scale-invariant, the basis functions remain dilated versions of a primary profile after applying the Riesz transform

$$\boldsymbol{\mathcal{R}}\psi_i(\boldsymbol{x}) = 2^{-3i/2}\boldsymbol{\mathcal{R}}\psi(\boldsymbol{x}/2^i).$$

This justifies the *wavelet* frame appelation. The Riesz and wavelet analysis operators commute as they are both convolution operators (*i.e.*, operators that are linear and invariant by translation). Therefore, for any $f \in L_2(\mathbb{R}^3)$, Riesz-wavelet coefficients can be obtained either by applying the Riesz transform to each wavelet band $\psi_{i,\boldsymbol{k}} * f$, or by analyzing each Riesz band $\mathcal{R}^{\boldsymbol{n}} f$ with the wavelet functions. The bandlimitedness and symmetry of ψ_i allows us to

compute the Riesz-wavelet expansion

$$\begin{aligned} f(\boldsymbol{x}) &= \left\{ \sum_{i\in\mathbb{Z}} \psi_i * \sum_{|\boldsymbol{n}|=N} (\mathcal{R}^{\boldsymbol{n}})^*(\mathcal{R}^{\boldsymbol{n}}\psi_i * f) \right\}(\boldsymbol{x}) \\ &= \left\{ \sum_{i\in\mathbb{Z}} \sum_{|\boldsymbol{n}|=N} ((\mathcal{R}^{\boldsymbol{n}})^*\psi_i) * (\mathcal{R}^{\boldsymbol{n}}\psi_i * f) \right\}(\boldsymbol{x}) \\ &= \left\{ \sum_{i\in\mathbb{Z}} \sum_{|\boldsymbol{n}|=N} ((\mathcal{R}^{\boldsymbol{n}})^*\psi_i) \right. \\ &\quad \left. * \sum_{\boldsymbol{k}\in\mathbb{Z}^3} s^{\boldsymbol{n}}_{i,\boldsymbol{k}} \mathrm{sinc}\left(\frac{\cdot - T_i\boldsymbol{k}}{T_i}\right)\frac{1}{T_i^3} \right\}(\boldsymbol{x}) \\ &= \sum_{i\in\mathbb{Z}} \sum_{|\boldsymbol{n}|=N} \sum_{\boldsymbol{k}\in\mathbb{Z}^3} s^{\boldsymbol{n}}_{i,\boldsymbol{k}} \left\{(\mathcal{R}^{\boldsymbol{n}})^*\psi_{i,\boldsymbol{k}}\right\}(\boldsymbol{x}) \end{aligned}$$

where $s^{\boldsymbol{n}}_{i,\boldsymbol{k}} = \{\mathcal{R}^{\boldsymbol{n}}\psi_i * f\}(\boldsymbol{x})|_{\boldsymbol{x}=T_i\boldsymbol{k}}$ is the Riesz-wavelet coefficient at scale i for the $\boldsymbol{n}$th Riesz band at shift $\boldsymbol{k}$.

D. Digital Riesz-Wavelet Transforms

We consider hereafter that any 3D image f corresponds to a set of uniform samples of a function of $L_2(\mathbb{R}^3)$: $f[\boldsymbol{k}] = f(\boldsymbol{k})$, such that f is radially bandlimited ($\hat{f}(\boldsymbol{\omega}) = 0$ if $||\boldsymbol{\omega}|| \geq \pi$). The sampling step is fixed to 1, without loss of generality. We denote B_π to be the set of such functions and V_0 the space of functions that can be written as

$$f(\boldsymbol{x}) = \sum_{\boldsymbol{k}\in\mathbb{Z}^3} f[\boldsymbol{k}]\mathrm{sinc}(\boldsymbol{x} - \boldsymbol{k}) \tag{11}$$

with $f[\cdot] \in \ell_2$, where sinc is the separable and orthonormalized cardinal sine function of $\mathbb{R}^3$. Because $B_\pi \subset V_0$, the Shannon reconstruction formula (11) applies for any $f \in B_\pi$, and the Riesz-wavelet coefficients of f can be computed as

$$\begin{aligned} s^{\boldsymbol{n}}_{i,\boldsymbol{k}} &= \left\{ \mathcal{R}^{\boldsymbol{n}}\psi_i * \sum_{\boldsymbol{k}'\in\mathbb{Z}^3} f[\boldsymbol{k}']\mathrm{sinc}(\cdot - \boldsymbol{k}') \right\}(T_i\boldsymbol{k}) \\ &= \sum_{\boldsymbol{k}'\in\mathbb{Z}^3} f[\boldsymbol{k}']\left\{\mathcal{R}^{\boldsymbol{n}}\psi_i\right\}(T_i\boldsymbol{k} - \boldsymbol{k}'). \end{aligned}$$

The 3D discrete sequence of coefficients $s^{\boldsymbol{n}}_i[\boldsymbol{k}] = s^{\boldsymbol{n}}_{i,\boldsymbol{k}}$ for $\boldsymbol{k} \in \mathbb{Z}^3$ are therefore obtained as the discrete convolution at points $T_i\boldsymbol{k}$ of the 3D sequences $f[\cdot]$ and $\{\mathcal{R}^{\boldsymbol{n}}\psi_i\}[\cdot]$. The latter corresponds to a discrete set of uniform samples of the function $\mathcal{R}^{\boldsymbol{n}}\psi_i$ on the grid $\mathbb{Z}^3$. The set of Riesz-wavelet coefficients can thus be computed in two steps: 1) discrete convolution between the image and the sampled filters, 2) downsampling by a factor T_i.

A different digital implementation consists in applying first the digital wavelet transform to f and then the digital Riesz transform to each wavelet band, as for the steerable pyramid in 2D. Indeed, it is not hard to show that the sequence of wavelet coefficients $w_i[\boldsymbol{k}] = w_{i,\boldsymbol{k}}$ can be obtained by subsampling by a factor T_i the discrete convolution between f and the samples of the wavelet filters. This leads to

$$w_i[\boldsymbol{k}] = \sum_{\boldsymbol{k}'\in\mathbb{Z}^3} f[k']\psi_i(T_i\boldsymbol{k} - \boldsymbol{k}').$$

Then, the set of Riesz-wavelet coefficients $s^{\boldsymbol{n}}_{i,\boldsymbol{k}} = \{\mathcal{R}^{\boldsymbol{n}}\psi_i * f\}(\boldsymbol{x})|_{\boldsymbol{x}=T_i k}$ is computed as

$$\begin{aligned} s^{\boldsymbol{n}}_{i,\boldsymbol{k}} &= \left\{ \mathcal{R}^{\boldsymbol{n}} \sum_{\boldsymbol{k}'\in\mathbb{Z}^3} w_{i,\boldsymbol{k}'} \frac{1}{T_i^3}\mathrm{sinc}\left(\frac{\cdot - T_i\boldsymbol{k}'}{T_i}\right) \right\}(\boldsymbol{x})|_{\boldsymbol{x}=T_i k} \\ &= \sum_{\boldsymbol{k}'\in\mathbb{Z}^3} w_{i,\boldsymbol{k}'}\{\mathcal{R}^{\boldsymbol{n}}\mathrm{sinc}\}\left(\boldsymbol{k} - \boldsymbol{k}'\right) \end{aligned} \tag{12}$$

which corresponds to a discrete convolution product between wavelet coefficients at scale i and a discrete set of uniform samples of $\mathcal{R}^{\boldsymbol{n}}$sinc (a bandlimited version of the impulse response of the Riesz component) on the Cartesian grid. In practice, this digital two-step scheme is efficient as the digital wavelet transform can be computed as a set of cascaded discrete convolutions and subsampling steps [13]. We then apply the Riesz transform to each set of wavelet coefficients (12). The space-domain convolution in (12) is not practical as the Riesz impulse response is not compactly supported and because it is not a conventional function at the origin. Instead, the discrete convolution is made possible in the frequency domain due to the smoothness of the wavelet transform at the origin [Condition (iii)] which annihilates the singularity of the frequency response of the Riesz transform at this point. Inverting the full transform to obtain a 3D digital signal f from the Riesz-wavelet coefficients is achieved by inverting sequentially each step of the forward transform, using a series of discrete adjoint operators.

E. Steerability

As the primary wavelet transform is isotropic, the directional information is fully encoded by the Riesz part of the frame. The proposed Riesz-wavelet transform consequently inherits the steerable property of the high-order Riesz transform as defined in (6) for $\mathbf{R}$ a (3×3) spatial rotation matrix, and we check that the general steerability definition in (4) is satisfied.

Applying Theorem 1 from [17], we have that the entries $s_{\boldsymbol{n},\boldsymbol{m}}$ (in multiindex notation) of $\mathbf{S}_{\mathbf{R}}$ for $\mathbf{R} = [\boldsymbol{r}_1\ \boldsymbol{r}_2\ \boldsymbol{r}_3]$ are computed as

$$\begin{aligned} s_{\boldsymbol{n},\boldsymbol{m}} = \sqrt{\frac{\boldsymbol{m}!}{\boldsymbol{n}!}} \sum_{|\boldsymbol{k}_1|=n_1} \sum_{|\boldsymbol{k}_2|=n_2} \sum_{|\boldsymbol{k}_3|=n_3} \delta_{\boldsymbol{k}_1+\boldsymbol{k}_2+\boldsymbol{k}_3,\boldsymbol{m}} \\ \cdot \frac{\boldsymbol{n}!}{\boldsymbol{k}_1!\boldsymbol{k}_2!\boldsymbol{k}_3!} \boldsymbol{r}_1^{\boldsymbol{k}_1}\boldsymbol{r}_2^{\boldsymbol{k}_2}\boldsymbol{r}_3^{\boldsymbol{k}_3} \end{aligned}$$

where $\delta_{\boldsymbol{k}_1+\boldsymbol{k}_2+\boldsymbol{k}_3,\boldsymbol{m}}$ denotes the Kronecker symbol that selects only multiindex coefficients $\boldsymbol{k}_1$, $\boldsymbol{k}_2$, and $\boldsymbol{k}_3$ such that $\boldsymbol{k}_1 + \boldsymbol{k}_2 + \boldsymbol{k}_3 = \boldsymbol{m}$. Theorem 2 in [17] moreover shows that we can steer back the Riesz-wavelet coefficients to their original values by applying the transpose of the steering matrix $\mathbf{S}^T_{\mathbf{R}}$ since $\mathbf{S}_{\mathbf{R}}\mathbf{S}^T_{\mathbf{R}} = \mathbf{I}_{M\times M}$ (the $(M \times M)$ identity matrix).

In practice, the sampled version of a steerable function is still steerable with the same interpolating functions since the steering and sampling operators commute. Hence, the steering of digital coefficients is performed using the standard matrix $\mathbf{S}_{\mathbf{R}}$ as defined above.

IV. 3D Monogenic Signal Analysis

As first application, we investigate the design of a multiscale analytic-like signal in 3D. The analytic signal was introduced by Gabor [31] as the complex extension of a 1D signal

$$f_{\text{anal}}(x) = f(x) + \mathrm{j}\mathcal{H}f(x) = A(x)\mathrm{e}^{\mathrm{j}\xi(x)}$$

based upon the Hilbert transform $\mathcal{H}$, which is the linear and shift-invariant operator that maps all 1D cosine functions into their corresponding sine functions. An AM/FM signal analysis is achieved with this representation by taking the time-varying amplitude as $A(x) = |f_{\text{anal}}(x)|$ and the instantaneous frequency as $\nu(x) = \mathrm{d}\xi(x)/\mathrm{d}x$. The extension of the analytical signal to 2D proposed in [32] is via the Riesz transform, which can be viewed as a natural extension of the Hilbert transform to multiple dimensions. This so-called *monogenic signal* was recently extended to the analysis of 2D Riesz-wavelet coefficients in [16], and to nD in [22]. We extend to 3D the directional Hilbert interpretation of the monogenic signal that is proposed in [16] as it is simpler and more directly applicable than the Clifford frame framework of [22].

A. 3D Monogenic Signal

We propose to perform in 3D the same type of AM/FM signal analysis as in 1D, but along the particular 3D direction $\boldsymbol{u} \in \mathbb{R}^3$. More precisely, we choose $\boldsymbol{u}$ as the local orientation that maximizes the directional Hilbert-transform response, with $\|\boldsymbol{u}\| = 1$. (We present in the next section a fast and robust way to compute $\boldsymbol{u}$.) When comparing the frequency response $\widehat{\mathcal{H}f}(\omega) = -\mathrm{j}\omega/|\omega|\hat{f}(\boldsymbol{\omega})$ of the Hilbert transform with the Fourier-domain definition of the Riesz components (5), it is apparent that $\mathcal{H}_{\boldsymbol{u}}f$, which is the Hilbert transform of f along the direction $\boldsymbol{u}$, corresponds to the projection of its Riesz transform onto $\boldsymbol{u}$. Formally, we write that $\mathcal{H}_{\boldsymbol{u}}f(\boldsymbol{x}) = \boldsymbol{u}^T\boldsymbol{\mathcal{R}}f(\boldsymbol{x})$, with $\boldsymbol{\mathcal{R}}f(\boldsymbol{x})$ the column-vector of first-order Riesz coefficients for f at position $\boldsymbol{x}$. We thus propose the following 3D extension of the monogenic signal which is based on the 3D Riesz transform:

$$\mathbf{f}_{\text{mono}}(\boldsymbol{x}) = (f(\boldsymbol{x}), \mathcal{R}_1 f(\boldsymbol{x}), \mathcal{R}_2 f(\boldsymbol{x}), \mathcal{R}_3 f(\boldsymbol{x})).$$

We compute the local signal amplitude (AM component) as $A(\boldsymbol{x}) = \|\mathbf{f}_{\text{mono}}(\boldsymbol{x})\|$. The local phase ξ in the direction $\boldsymbol{u}$ is computed as $\xi(\boldsymbol{x}) = \arctan\left(\mathcal{H}_{\boldsymbol{u}}(\boldsymbol{x})/f(\boldsymbol{x})\right) = \arctan\left(\boldsymbol{u}^T\boldsymbol{\mathcal{R}}f(\boldsymbol{x})/f(\boldsymbol{x})\right)$.

B. Tensor-Based Estimation of the Local Orientation

We extend to 3D the local orientation-estimation procedure proposed in [16] that relies on a structure tensor [33]. We select the orientation $\boldsymbol{u}$ that maximizes $\mathcal{H}_{\boldsymbol{u}}(\boldsymbol{x})$, but, instead of doing it pointwise, we optimize the response over a 3D local neighborhood that is specified by the isotropic weighting function $\upsilon : \mathbb{R}^3 \mapsto \mathbb{R}_+$. The function υ is typically a 3D Gaussian window which acts as a regularization function of the orientation map and enhances the consistency and the robustness to noise of the estimated values. The pointwise optimization problem for estimating an orientation is

$$\boldsymbol{u}_\upsilon(\boldsymbol{x}_0) = \arg\max_{\|\boldsymbol{u}\|=1} \int_{\mathbb{R}^3} \upsilon(\boldsymbol{x}-\boldsymbol{x}_0)|\mathcal{H}_{\boldsymbol{u}}f(\boldsymbol{x})|^2\mathrm{d}\boldsymbol{x}.$$

Exploiting the link between the directional Hilbert transform and the Riesz transform, and taking advantage of the steerability of the latter, we write

$$|\mathcal{H}_{\boldsymbol{u}}f(\boldsymbol{x})|^2 = \boldsymbol{u}^T(\boldsymbol{\mathcal{R}}f(\boldsymbol{x}))(\boldsymbol{\mathcal{R}}f(\boldsymbol{x}))^T\boldsymbol{u}.$$

The above formula shows that estimating the local orientation is a quadratic-form maximization problem which boils down to an eigenvalue problem. The vector $\boldsymbol{u}_\upsilon$ is computed pointwise as the eigenvector corresponding to the largest eigenvalue of the tensor matrix $\mathbf{J}(\boldsymbol{x}_0)$, with

$$[\mathbf{J}(\boldsymbol{x}_0)]_{mn} = \int_{\mathbb{R}^3} \upsilon(\boldsymbol{x}-\boldsymbol{x}_0)\mathcal{R}_m f(\boldsymbol{x})\mathcal{R}_n f(\boldsymbol{x})\mathrm{d}\boldsymbol{x}$$

and $m, n \in \{1, 2, 3\}$. The sorted collection of eigenvectors of $\mathbf{J}(\boldsymbol{x}_0)$ defines a rotation matrix $\mathbf{U}_\upsilon$ such that the energy of the first band of the rotated Riesz transform $\mathbf{U}_\upsilon\boldsymbol{\mathcal{R}}f$ is maximized at $\boldsymbol{x}_0$, the first row of $\mathbf{U}_\upsilon$ being $\boldsymbol{u}_\upsilon(\boldsymbol{x}_0)^T$.

The eigenvalue decomposition of $\mathbf{J}(\boldsymbol{x}_0)$ also allows us to quantify the degree of directionality of the neighborhood through the relative weights of the eigenvalues. To do so, we define the coherency of the 3D local neighborhood as

$$\chi(\boldsymbol{x}_0) = \frac{\lambda_1(\boldsymbol{x}_0) - (\lambda_2(\boldsymbol{x}_0) + \lambda_3(\boldsymbol{x}_0))/2}{\lambda_1(\boldsymbol{x}_0) + (\lambda_2(\boldsymbol{x}_0) + \lambda_3(\boldsymbol{x}_0))/2}$$

with $\lambda_i(\boldsymbol{x}_0)$ the ith largest eigenvalue of $\mathbf{J}(\boldsymbol{x}_0)$. The coherency indicator takes value 1 when all the eigenvalues except the first do vanish (the energy is concentrated along the $\boldsymbol{u}_\upsilon(\boldsymbol{x}_0)$ direction), and 0 when the three eigenvalues are equal (isotropic gradient values).

C. 3D Monogenic Wavelet for the Analysis of Biological Images

We apply the proposed 3D monogenic analysis to each scale of the 3D isotropic wavelet pyramid described in Section III, which yields at each scale the monogenic signal

$$\mathbf{f}_{\text{mono},i}(\boldsymbol{x}) = (\psi_i * f(\boldsymbol{x}), \mathcal{R}_1\{\psi_i * f\}(\boldsymbol{x}), \times\ \mathcal{R}_2\{\psi_i * f\}(\boldsymbol{x}), \mathcal{R}_3\{\psi_i * f\}(\boldsymbol{x})).$$

As a result, we estimate locally the orientation and coherency for different scales of the underlying 3D image. It is worth pointing out that the isotropy of the primary wavelet pyramid is crucial here. It prevents the introduction of a bias in the monogenic analysis of the Riesz coefficients at each scale. We show in Figure 3 the local orientations obtained with the monogenic Riesz-wavelet transform for fluorescence images of collagen filaments. The figure is an illustration of the ability of our method to estimate the local orientation at a given scale. The estimation is robust and consistent, which is of major importance for 3D microscope images which are generally corrupted by noise. This ability is exploited in the Section V to guide a multiscale edge-detection algorithm which was found to perform well for a set of 3D medical images.

V. Image Reconstruction from Edges: An Inverse Problem Approach

We demonstrate now the use of the 3D Riesz-wavelet frame for inverse problems. In particular, we investigate the

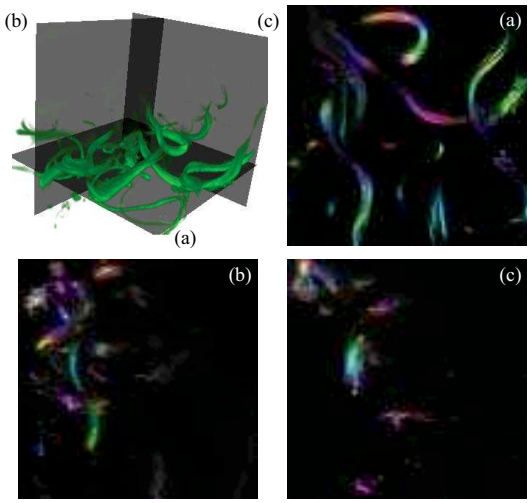

Fig. 3. 3D image stack of collagen filaments (top left). (a)–(c) Color encoded (hue) monogenic direction at scale 2 for the plans A, B, and C, after an orthogonal projection onto the *xy*, *xz*, and *yz* planes, respectively. The saturation indicates the value of the local coherency.

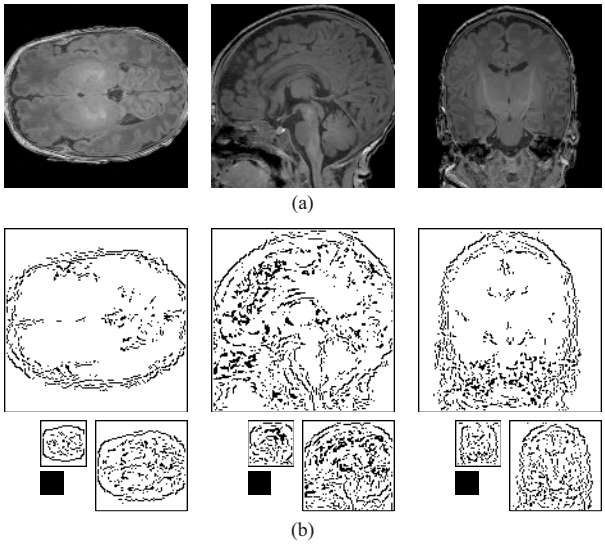

Fig. 4. Multiscale edge detection for a 3D MRI volume. (a) Transverse, sagittal, and coronal slices from the original ($144 \times 144 \times 144$) data. (b) 3D edges detected for each slice in the three analyzed wavelet bands.

reconstruction of an image from a subset of coefficients located on edges. This application leads us to solve constrained and unconstrained optimization problems involving the Riesz-wavelet transform of 3D images.

A. 3D Singularity Detection in the Riesz-Wavelet Domain

We adapt the well-known Canny edge detector [34] to the 3D multiscale case by using the monogenic signal analysis proposed above. For an input signal $f \in V_0$, coefficients corresponding to singularities in 3D (surfaces, ridges, dots) are detected by exploiting the estimated directional monogenic amplitude and pointwise local orientation. For each wavelet scale, we

1) compute the local orientation $\boldsymbol{u}_v(\boldsymbol{x}_0)$ at each point $\boldsymbol{x}_0$;
2) identify $\mathcal{B}_{\boldsymbol{x}_0,\boldsymbol{u}_v(\boldsymbol{x}_0)}$, the set of neighboring voxels of $\boldsymbol{x}_0$ in the direction $\boldsymbol{u}_v(\boldsymbol{x}_0)$;
3) retain $\boldsymbol{x}_0$ as a singular point if $A(\boldsymbol{x}_0) \geq A(\boldsymbol{x})$, $\forall \boldsymbol{x} \in \mathcal{B}_{\boldsymbol{x}_0,\boldsymbol{u}_v(\boldsymbol{x}_0)}$.

For each singular point, thresholding with hysteresis [34] is then performed on the monogenic amplitude values. High amplitude singularities are detected using a high threshold value, and lower-amplitude singularities (exceeding a low threshold value) are kept only if connected (using 26 connectivity in 3D) to a high-amplitude component. This results in a set of singular points which represent in a robust manner both high- and low-intensity 3D features. The coarsest-scale coefficients are also included in the representation in order to preserve the brightness information of the image. We show in Figure 4 the set of edges which are detected for a real 3D MRI image. We check that the procedure of thresholding with hysteresis yields the identification of both strong and weak features with only few false-positive detections.

B. Reconstruction from Edge Coefficients

Let us denote by $\boldsymbol{f} \in \mathbb{R}^N$ the column vector that represents a 3D image with N pixels, and by $\boldsymbol{s}$ the column vector of its stacked Riesz-wavelet coefficients $\{s_{i,\boldsymbol{k}}^{\boldsymbol{n}}\}$. Then, $\boldsymbol{f} = \mathbf{W}\boldsymbol{s}$ with $\mathbf{W}$ the $(N \times M)$ matrix corresponding to the digital Riesz-wavelet synthesis operator. As we are considering tight Riesz-wavelet frames, we also write $\boldsymbol{s} = \mathbf{W}^H \boldsymbol{f}$, where the Hermitian transpose of the matrix $\mathbf{W}$ corresponds to the Riesz-wavelet analysis operator.

Let us assume for the image $\boldsymbol{f}_0 = \mathbf{W}\boldsymbol{s}_0$ that the singular-point detection method presented in Section V-A yields a set of indices $\mathcal{S}$ for the Riesz-wavelet coefficients $\boldsymbol{s}_0$. We wish to build a close approximation $\hat{\boldsymbol{f}}$ of $\boldsymbol{f}_0$ from the knowledge of the subset $\mathcal{S}$ of Riesz-wavelet coefficients only. For the 2D Marr-like pyramid [35], which is inspired from experiments on visual perception and is technically closely related to the monogenic Riesz-wavelet pyramid in 2D, a primal sketch was proposed to build $\hat{\boldsymbol{f}}$ as an image with consistent edges with $\boldsymbol{s}_0$. In practice, $\hat{\boldsymbol{f}}$ was obtained by alternating between 1) the penalization of wrong edges and 2) the projection onto the space of feasible wavelet coefficients, so that the edge-detection procedure applied to $\hat{\boldsymbol{f}}$ yielded the index set $\mathcal{S}$. Despite some convicing results, this approach was suffering from some limitations. First, the set of images with consistent edges is not finite in general and there is no hint that the image towards which the procedure converges is close to the original image $\boldsymbol{f}_0$. Second, the edge-detection procedure is time-consuming in 3D, which hampers the use of methods repeatedly using it. We tackle these issues by proposing a new paradigm for reconstructing images from edge coefficients that exploits a novel inverse-problem formulation of the task.

1) Constrained Optimization With the Riesz-Wavelet Transform: We express the task of reconstructing an image from wavelet-domain singularities as the constrained minimization problem over the image space

$$\hat{\boldsymbol{f}} = \underset{\boldsymbol{f} \in \mathbb{R}^N}{\arg\min} \|\mathbf{W}_\Lambda^H \boldsymbol{f}\|_1 \text{ s.t. } \mathcal{C} \tag{13}$$

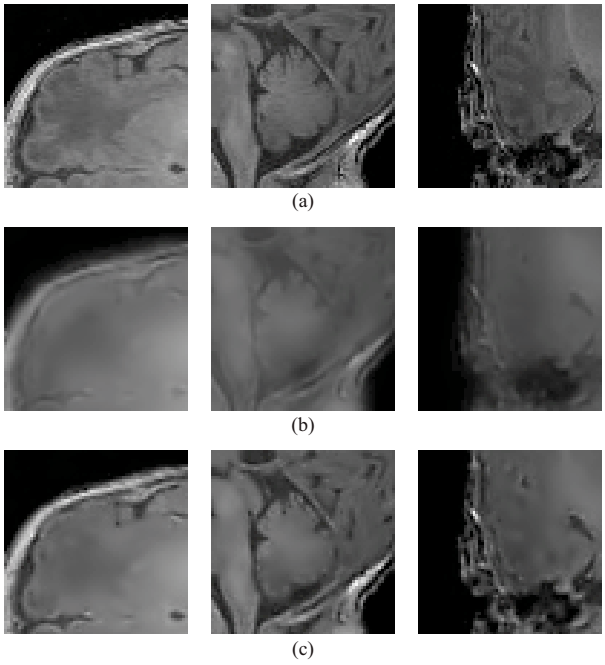

Fig. 5. 3D image reconstruction for the edge set shown in Fig. 4. (a) 2D areas from the original MRI volume. (b) Reconstruction by orthogonal projection of the edge coefficients: 25.81 dB. (c) Reconstructed images by (13) with Algorithm 1 (50 iterations): 32.06 dB. The orthogonal projection reconstruction shown in (b) is used as a starting point for the proposed method shown in (c).

with

$$\mathcal{C} : [\mathbf{W}^H \boldsymbol{f}]_k = [\boldsymbol{s}_0]_k, \ \forall k \in \mathcal{S}.$$

There, $\| \cdot \|_1$ stands for the ℓ_1 norm of $\mathbb{R}^M$. In (13), we take $\mathbf{W}_\Lambda^H = \mathbf{\Lambda}\mathbf{W}^H$ with $\mathbf{\Lambda}$ an $(M \times M)$ diagonal matrix with positive elements which weight the Riesz-wavelet coefficients for the computation of the regularization term. We typically take equal weights for coefficients from the same Riesz-wavelet band and set the diagonal elements of $\mathbf{\Lambda}$ to 0 for the coefficients of the coarsest scale (lowpass residual) that should not be penalized.

With the proposed formulation (13), the set of constraints $\mathcal{C}$ for Riesz-wavelet coefficients enforces the consistency between the original sampled values and the coefficients of the estimated image. The regularization term $\|\mathbf{W}_\Lambda^H \boldsymbol{f}\|_1$ corresponds to a sparse prior for the solution image in the Riesz-wavelet space. The sparsity criterion naturally complies with the edge consistency between $\boldsymbol{f}_0$ and $\hat{\boldsymbol{f}}$ as spurious singularities in the Riesz-wavelet representation of $\hat{\boldsymbol{f}}$ generally yield an increased ℓ_1 penalty term. This intuitive assumption is verified in the sequel with various reconstruction experiments.

It is worth noting that if we take $\mathbf{\Lambda} = \mathbf{I}$, and if an orthonormal frame ($\mathbf{W}^{-1} = \mathbf{W}^H$) is used in place of the Riesz-wavelet frame, then there exists a trivial solution to the constrained problem (13). This solution is given by $\hat{\boldsymbol{f}} = \mathbf{W}\tilde{\boldsymbol{s}}$, with $\tilde{s}_k = [\boldsymbol{s}_0]_k$ if $k \in \mathcal{S}$, and $\tilde{s}_k = 0$ otherwise. Indeed, in this particular case we have that $\tilde{\boldsymbol{s}} = \mathbf{W}^H\mathbf{W}\tilde{\boldsymbol{s}}$. By contrast, because the Riesz-wavelet frame is redundant, there may not exist $\boldsymbol{f} \in \mathbb{R}^N$ such that $\tilde{\boldsymbol{s}} = \mathbf{W}^H \boldsymbol{f}$. In other words, $\tilde{\boldsymbol{s}}$ may not correspond to the Riesz-wavelet coefficients of an image. A naive idea to solve the reconstruction problem is to orthogonally project $\tilde{\boldsymbol{s}}$ onto the space of Riesz-wavelet coefficients. This requires the solution of $\arg\min_{f \in \mathbb{R}^N} ||\tilde{\boldsymbol{s}} - \mathbf{W}^H f||_2$, which can be obtained by exploiting the tight-frame property of the Riesz-wavelet frame ($\mathbf{W}\mathbf{W}^H = \mathbf{I}$). It is easy to show that this solution, which we call the orthogonal projection solution, is $\tilde{\boldsymbol{f}} = \mathbf{W}\tilde{\boldsymbol{s}}$, as for the case of a basis. We however show in Figure 5 for the MRI volume that this procedure is detrimental to the edge consistency. We thus aim at solving (13) to improve the reconstruction quality by jointly imposing edge consistency and exploiting a sparsity prior.

Algorithm 1 Augmented-Lagrangian Optimization Procedure

Require: $\boldsymbol{\mu}^{(0)} \in \mathbb{R}_+^M$, $\boldsymbol{\lambda}^{(0)} \in \mathbb{R}^M$, $\boldsymbol{f}_0^{(0)} \in \mathbb{R}^N$, and a final tolerance $\epsilon \in \mathbb{R}_+$ or a maximal number of iterations T

for $t = 0, 1, 2 \ldots$ **do**

Compute $\boldsymbol{f}^{(t)} = \arg\min_{f \in \mathbb{R}^N} \mathcal{L}(\boldsymbol{f}, \boldsymbol{\lambda}^{(t)}, \boldsymbol{\mu}^{(t)})$ with the starting point $\boldsymbol{f}_0^{(t)}$

if convergence holds ($\|\nabla_f \mathcal{L}(\boldsymbol{f}^{(t)}, \boldsymbol{\lambda}^{(t)}, \boldsymbol{\mu}^{(t)})\| \leq \epsilon$ or $t \geq T$) **then**

Exit

else

for $k \in \mathcal{S}$ **do**

$\lambda_k^{(t+1)} = \lambda_k^{(t)} - \mu_k^{(t)} \left([\mathbf{W}^H \boldsymbol{f}^{(t)}]_k - [\boldsymbol{s}_0]_k \right)$

$\mu_k^{(t+1)} = \beta \mu_k^{(t)}$ with $\beta \in (1, \infty)$

end for

$\boldsymbol{f}_0^{(t+1)} = \boldsymbol{f}^{(t)}$

end if

end for $\hat{\boldsymbol{f}} = \boldsymbol{f}^{(t)}$

Overall, (13) corresponds to the minimization of a convex functional subject to a finite set of linear constraints. We can thus guarantee that there exists a feasible minimum $\hat{\boldsymbol{f}}$, and that we can reach it with appropriate optimization algorithms. We propose to use an augmented-Lagrangian (AL) technique [36] in order to maximize (13) subject to $\mathcal{C}$. For the vectors of multipliers $\boldsymbol{\lambda} \in \mathbb{R}^M$ and $\boldsymbol{\mu} \in \mathbb{R}_+^M$ we write the AL functional as

$$\mathcal{L}(\boldsymbol{f}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \|\mathbf{W}_\Lambda^H \boldsymbol{f}\|_1 - \sum_{k \in S} \lambda_k \left([\mathbf{W}^H \boldsymbol{f}]_k - [\boldsymbol{s}_0]_k \right) + \sum_{k \in S} \frac{\mu_k}{2} \left([\mathbf{W}^H \boldsymbol{f}]_k - [\boldsymbol{s}_0]_k \right)^2.$$

The global scheme for AL minimization is summarized in Algorithm 1. We initialize the multiplier vectors as $\lambda_k^{(0)} = 0$ for $k = 1 \ldots M$, and $\mu_k^{(0)} = 1$ for $k \in \mathcal{S}$ and $\mu_k^{(0)} = 0$ otherwise. We note that, following Algorithm 1, multipliers values λ_k and μ_k for $k \notin \mathcal{S}$ are kept fixed to their initial zero values. Using these initial conditions, we concisely re-write the Lagrangian functional as

$$\mathcal{L}(\boldsymbol{f}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \|\mathbf{W}_\Lambda^H \boldsymbol{f}\|_1 - \boldsymbol{\lambda}^T \left(\mathbf{W}^H \boldsymbol{f} - \boldsymbol{s}_0 \right) + \frac{1}{2} \left(\mathbf{W}^H \boldsymbol{f} - \boldsymbol{s}_0 \right)^T \mathbf{C} \left(\mathbf{W}^H \boldsymbol{f} - \boldsymbol{s}_0 \right) \quad (14)$$

with $\mathbf{C}$ a diagonal $(M \times M)$ matrix such that $\mathbf{C}_{k,k} = \mu_k$ for $k = 1 \ldots M$. In order to initiate the optimization algorithm, we take the orthogonal projection solution $\tilde{\boldsymbol{f}}$ as the first estimate $\boldsymbol{f}_0^{(0)}$.

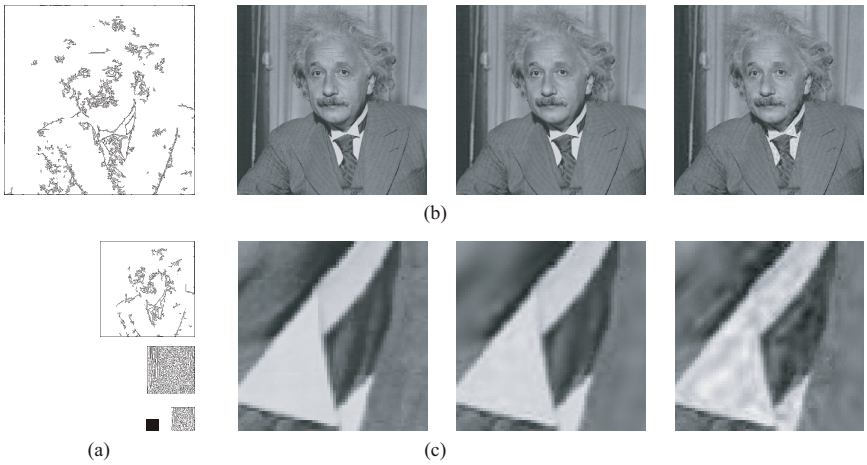

Fig. 6. *Einstein* image reconstructed from edges. (a) Edge map in wavelet bands as obtained from a wavelet-domain Canny-like detector [35]. (b) Full-resolution images. (c) Zoomed-in images of the collar area. The images in (b) and (c) are, from left to right: original image, reconstructed image using the proposed method with 50 iterations (35.37 dB), and image reconstructed with the Marr-like primal sketch [35] with 50 iterations (33.46 dB).

At each step of the AL minimization procedure, we solve the unconstrained minimization problem $\boldsymbol{f}^\star = \arg\min_{\boldsymbol{f}\in\mathbb{R}^N} \mathcal{L}(\boldsymbol{f}, \boldsymbol{\lambda}, \boldsymbol{\mu})$. As the cost function is convex, we use a gradient-descent-like minimization procedure. The gradient vector at point $\boldsymbol{f}$ is computed as

$$\begin{aligned}\mathbf{G}_{\boldsymbol{f}} &\triangleq \nabla_{\boldsymbol{f}}\mathcal{L}(\boldsymbol{f}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \\ &= \mathbf{W}\left(\Lambda \boldsymbol{v}_{\boldsymbol{f}} - \boldsymbol{\lambda} + \mathbf{C}\left(\mathbf{W}^H \boldsymbol{f} - \boldsymbol{s}_0\right)\right)\end{aligned}$$

where $\boldsymbol{v}_{\boldsymbol{f}}$ is the M-dimensional vector with component $[\boldsymbol{v}_{\boldsymbol{f}}]_k = \mathrm{sgn}([\mathbf{W}_\Lambda^H \boldsymbol{f}]_k)$. In practice, we use a Huber loss function to cope for the non-differentiability of the absolute value function at the origin. It is worth pointing out that the computation of $\mathbf{G}_{\boldsymbol{f}}$ is efficient for a tight Riesz-wavelet frame since, in this case, the left-multiplication by $\mathbf{W}$ corresponds to applying the synthesis operator, which admits a fast filterbank implementation. In order to solve (14), we use an iterative procedure in which the tth estimate is computed as $\boldsymbol{f}^{(t)} = \boldsymbol{f}^{(t-1)} - \alpha^{(t)}\mathbf{G}_{\boldsymbol{f}}$. The step length $\alpha^{(t)} \in \mathbb{R}_+$ is set such as to minimize $\mathcal{L}(\boldsymbol{f}^{(t-1)} - \alpha^{(t)}\mathbf{G}_{\boldsymbol{f}}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ for faster convergence. We note that, once $\mathbf{W}^H\mathbf{G}_{\boldsymbol{f}}$ is computed, the optimization of the step length can be performed very efficiently in the transformed domain using a line-search algorithm since we have that $\mathbf{W}^H(\boldsymbol{f}^{(t-1)} - \alpha^{(t)}\mathbf{G}_{\boldsymbol{f}}) = \mathbf{W}^H\boldsymbol{f}^{(t-1)} - \alpha^{(t)}\mathbf{W}^H\mathbf{G}_{\boldsymbol{f}}$. To do so, we keep the estimated image in the Riesz-wavelet domain and perform the update as $\mathbf{W}^H\boldsymbol{f}^{(t)} = \mathbf{W}^H\boldsymbol{f}^{(t-1)} - \alpha^{(t)}\mathbf{W}^H\mathbf{G}_{\boldsymbol{f}}$ because the AL functional involves only transformed-domain coefficients. As a consequence, we need to apply the Riesz-wavelet analysis and synthesis operators only once per iteration of the gradient-descent algorithm in order to estimate $\mathbf{W}^H\mathbf{G}_{\boldsymbol{f}}$. The computation of the Riesz-wavelet forward and backward transforms is the main computational bottleneck as the evaluation of the gradient vectors and the computation of AL functional values both involve only pointwise operations in the transformed domain. The fast filterbank implementation of the transforms is therefore a key advantage of the proposed steerable frames for the reconstruction task.

C. Experiments

1) 2D Image Reconstruction: We first investigate the reconstruction of 2D images from edge coefficients in order to facilitate the comparison of the proposed technique with previous works. We have applied the edge-detection procedure proposed in [35] to the *Einstein* image in order to obtain a mask indicating edges in the wavelet domain. An example of detected edges is shown in Figure 6. We have then applied both the primal-sketch reconstruction with the Marr-like pyramid coefficients and the proposed reconstruction method with the Riesz-wavelet coefficients when using 2D versions of the isotropic wavelet and Riesz transforms. We show in Figure 6 the corresponding reconstructions. We observe that the proposed method yields an image that is very close to the original (PSNR = 35.37 dB), despite the low ratio of coefficients that are retained by the edge detector (7.5%). Edges in the reconstructed image are indeed consistent with those in the original image, although (13) does not explicitly impose consistency. As compared with the state-of-the-art technique, we have found that our reconstructed image contains much fewer oscillating patterns close to the edges. This happens because the ℓ_1 minimization technique we use prevents this kind of artifacts and the creation of new edges. For the reconstruction method based on the primal sketch, oscillating patterns are allowed (despite the extra ℓ_2 penalty term) as long as they are not seen by the edge detector. As a result, the proposed reconstruction significantly outperforms the primal-sketch reconstruction based on the Marr-like pyramid in term of PSNR (+1.91 dB).

We have repeated the reconstruction experiments for different rates of retained coefficients. For the two techniques, the performance measured by the PSNR and the structural similarity (SSIM) index, which better accounts for visual perception [37], are shown in Figure 7. We observe that the proposed method consistently outperforms the primal sketch reconstruction method, for both the PNSR and SSIM criteria. We have measured an SSIM difference as large as 0.049.

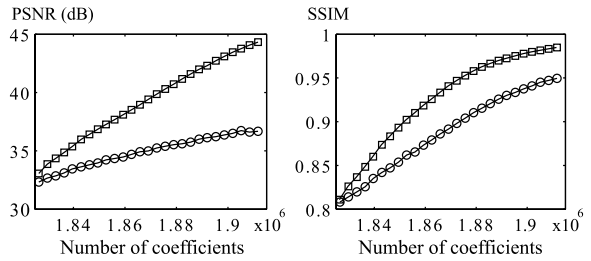

Fig. 7. Performance of reconstruction from edges for the *Einstein* image with an increasing number of coefficients. Results for the proposed reconstruction method in 2D are plotted with boxes. Results for the Marr-like primal sketch [35] are plotted with circles. Fifty iterations have been performed for each reconstruction procedure.

The difference between the measured PSNR values increases with the number of coefficients: In our experiments, the difference monotonically grows from 0.73 dB to 7.63 dB. We note that the PSNR of the primal-sketch tends to saturate for large numbers of coefficients, while that of the proposed method does not. As demonstrated by the very high SSIM values (up to 0.985), our experiments show that the Riesz-wavelet coefficients located on the edges contain the essential visual information of the *Einstein* image and that the proposed reconstruction method is able to fully exploit these data.

2) Reconstruction of 3D Medical Images: We have applied the proposed reconstruction procedure to the set of 3D edges shown in Figure 4 for the MRI data set. In Figure 5, some 2D slices from the initial estimate (orthogonal projection) and from the final reconstruction are shown. It first demonstrates that the orthogonal projection of the edge coefficients onto the space of images does not yield an accurate estimation of the original data. This projection operation tends to blur the edges excessively. By contrast, the proposed method yields consistent and sharpened edges. As a result, the reconstructed image quality is significantly improved as compared to the initial projected images (+6.25 dB for the PSNR). Like for the 2D case, we check that no artifactual oscillatory pattern is created near the edges. Instead, the reconstruction method produces smooth areas in regions where no edge is present. The reconstruction quality which is achieved is high for the analyzed MRI volume: 32.06 dB. It shows, for a real 3D medical data set, that the proposed 3D Riesz-wavelet coefficients contain the essential data information and that our inverse-problem approach for image reconstruction is able to extract and exploit this information.

VI. Conclusion

In this paper, we have introduced the first 3D wavelet frame having the key property of steerability. We have discussed a functional framework involving both non-separable wavelet pyramids and the 3D generalized Riesz transforms. Then, the link between the continuous-domain model and the 3D digital transforms has been established. As a result, fast filterbank transforms for the synthesis and analysis operators have been proposed. We have identified key properties of the 3D Riesz-wavelet frame such as the 3D steerability of the basis functions, its tightness, and the ability to customize it by shaping the parameterization matrix of the 3D generalized Riesz transform.

In order to illustrate the usefulness of our contributions, we have investigated two applications taking advantage of 3D steerable pyramids. First, we have demonstrated that the Riesz-wavelet transform can be used to define a 3D monogenic signal which allows one to estimate AM/FM-like parameters along any 3D direction. We have shown for a set of biological images that one can accurately estimate the local orientation in 3D microscope images. The 3D monogenic signal has been exploited in our second application for detecting singularities at multiple scales of 3D medical images. We have then proposed a new optimization framework for reconstructing images based on the coefficients along these multiscale singularities. Our technique relies on the formulation of the reconstruction problem as a constrained sparsity-promoting optimization of an inverse problem and takes advantage of Riesz-wavelet frames. In practice, experiments have shown that our approach outperforms previous works when considering the 2D case, while its applications to 3D is novel. In both 2D and 3D, the high quality of the reconstructed images shows that the Riesz-wavelet coefficients along singularities contain the essential information of the images, and that the proposed reconstruction method takes full advantage of it. All the 3D experiments in this work have been performed with real biomedical data. This type of data requires extensive use of image analysis, processing, and reconstruction methods.

Acknowledgment

The authors would like to thank D. Van De Ville for providing the code of the Marr pyramid. They would also like to thank M. Guerquin-Kern and J.-C. Baritaux for fruitful discussions regarding optimization techniques.

References

[1] M. Guerquin-Kern, M. Häberlin, K. Pruessmann, and M. Unser, "A fast wavelet-based reconstruction method for magnetic resonance imaging," *IEEE Trans. Med. Imag.*, vol. 30, no. 9, pp. 1649–1660, Sep. 2011.

[2] C. Vonesch and M. Unser, "A fast multilevel algorithm for wavelet-regularized image restoration," *IEEE Trans. Image Process.*, vol. 18, no. 3, pp. 509–523, Mar. 2009.

[3] H. Benoit-Cattin, A. Baskurt, F. Turjman, and R. Prost, "3D medical image coding using separable 3D wavelet decomposition and lattice vector quantization," *Signal Process.*, vol. 59, no. 2, pp. 139–153, Jun. 1997.

[4] J. Kovačević and M. Vetterli, "Nonseparable two- and three-dimensional wavelets," *IEEE Trans. Signal Process.*, vol. 43, no. 5, pp. 1269–1273, May 1995.

[5] A. S. Lewis and G. Knowles, "Video compression using 3D wavelet transforms," *Electron. Lett.*, vol. 26, no. 6, pp. 396–398, Mar. 1990.

[6] C. I. Podilchuk, N. S. Jayant, and N. Farvardin, "Three-dimensional subband coding of video," *IEEE Trans. Image Process.*, vol. 4, no. 2, pp. 125–139, Feb. 1995.

[7] I. W. Selesnick and K. Y. Li, "Video denoising using 2D and 3D dual-tree complex wavelet transforms," *Proc. SPIE*, vol. 5207, no. 1, pp. 607–618, 2003.

[8] L. Ying, L. Demanet, and E. Candes, "3D discrete curvelet transform," *Proc. SPIE*, vol. 5914, no. 1, pp. 351–361, 2005.

[9] A. Woiselle, J. L. Starck, and J. Fadili, "3D curvelet transforms and astronomical data restoration," *Appl. Comput. Harmonic Anal.*, vol. 28, pp. 171–188, Mar. 2010.

[10] D. L. Donoho and O. Levi, "Fast X-ray and beamlet transforms for three-dimensional data," in *Proc. Modern Signal Process.*, 2002, pp. 1–38.

[11] V. Chandrasekaran, M. B. Wakin, D. Baron, and R. G. Baraniuk, "Representation and compression of multidimensional piecewise functions using surflets," *IEEE Trans. Inf. Theory*, vol. 55, no. 1, pp. 374–400, Jan. 2009.
[12] Y. M. Lu and M. N. Do, "Multidimensional directional filter banks and surfacelets," *IEEE Trans. Image Process.*, vol. 16, no. 4, pp. 918–931, Apr. 2007.
[13] E. P. Simoncelli and W. T. Freeman, "The steerable pyramid: A flexible architecture for multi-scale derivative computation," in *Proc. Int. Conf. Image Process.*, vol. 3. Oct. 1995, pp. 444–447.
[14] J. Portilla, V. Strela, M. Wainwright, and E. P. Simoncelli, "Image denoising using scale mixtures of Gaussians in the wavelet domain," *IEEE Trans. Image Process.*, vol. 12, no. 11, pp. 1338–1351, Nov. 2003.
[15] I. Daubechies, A. Grossmann, and Y. Meyer, "Painless nonorthogonal expansions," *J. Math. Phys.*, vol. 27, no. 5, pp. 1271–1283, May 1986.
[16] M. Unser, D. Sage, and D. Van De Ville, "Multiresolution monogenic signal analysis using the Riesz-Laplace wavelet transform," *IEEE Trans. Image Process.*, vol. 18, no. 11, pp. 2402–2418, Nov. 2009.
[17] M. Unser, N. Chenouard, and D. Van De Ville, "Steerable pyramids and tight wavelet frames in $l_2 R^d$," *IEEE Trans. Image Process.*, vol. 20, no. 10, pp. 2705–2721, Oct. 2011.
[18] S. C. Olhede and G. Metikas, "The monogenic wavelet transform," *IEEE Trans. Signal Process.*, vol. 57, no. 9, pp. 3426–3441, Sep. 2009.
[19] D. Marr and E. Hildreth, "Theory of edge detection," *Proc. Royal Soc. London. Series B. Biolog. Sci.*, vol. 207, no. 1167, pp. 187–217, Feb. 1980.
[20] W. T. Freeman and E. H. Adelson, "The design and use of steerable filters," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 9, pp. 891–906, Sep. 1991.
[21] C. Delle Luche, D. Denis, and A. Baskurt, "3D steerable pyramid based on conic filters," *Proc. SPIE*, vol. 5266, no. 2060, pp. 260–268, 2004.
[22] S. Held, M. Storath, P. Massopust, and B. Forster, "Steerable wavelet frames based on the Riesz transform," *IEEE Trans. Image Process.*, vol. 19, no. 3, pp. 653–667, Mar. 2010.
[23] M. Unser and D. Van De Ville, "Wavelet steerability and the higher-order Riesz transform," *IEEE Trans. Image Process.*, vol. 19, no. 3, pp. 636–652, Mar. 2010.
[24] J. Portilla and E. P. Simoncelli, "A parametric texture model based on joint statistics of complex wavelet coefficients," *Int. J. Comput. Vision*, vol. 40, no. 1, pp. 49–70, Oct. 2000.
[25] M. Papadakis, G. Gogoshin, I. A. Kakadiaris, D. J. Kouri, and D. K. Hoffman, "Nonseparable radial frame multiresolution analysis in multidimensions and isotropic fast wavelet algorithms," *Proc. SPIE*, vol. 5207, no. 260, pp. 631–642, 2003.
[26] J. R. Romero, S. K. Alexander, S. Baid, S. Jain, and M. Papadakis, "The geometry and the analytic properties of isotropic multiresolution analysis," *Adv. Comput. Math.*, vol. 31, nos. 1–3, pp. 283–328, Oct. 2009.
[27] I. Daubechies, *Ten Lectures on Wavelets*. Philadelphia, PA: SIAM, 1992.
[28] R. J. Duffin and A. C. Schaeffer, "A class of nonharmonic Fourier series," *Trans. Amer. Math. Soc.*, vol. 72, no. 2, pp. 341–366, Mar. 1952.
[29] O. Christensen, *An Introduction to Frames and Riesz Bases* (Applied and Numerical Harmonic Analysis). Boston, MA: Birkhäuser, 2003.
[30] I. M. Gel'fand, G. Shilov, and E. Saletan, *Generalized Functions*. New York: Academic, 1964.
[31] D. Gabor, "Theory of communications," *J. Inst. Elect. Eng.*, vol. 93, no. 3, pp. 429–457, Nov. 1946.
[32] M. Felsberg and G. Sommer, "The monogenic signal," *IEEE Trans. Signal Process.*, vol. 49, no. 12, pp. 3136–3144, Dec. 2001.
[33] B. Jähne, *Digital Image Processing*. New York: Springer-Verlag, 2005.
[34] J. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 8, no. 6, pp. 679–698, Nov. 1986.
[35] D. Van De Ville and M. Unser, "Complex wavelet bases, steerability, and the Marr-like pyramid," *IEEE Trans. Image Process.*, vol. 17, no. 11, pp. 2063–2080, Nov. 2008.
[36] D. P. Bertsekas, *Constrained Optimization and Lagrange Multiplier Methods*. Belmont, MA: Athena Scientific, 1996.
[37] W. Zhou, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.

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

He was a Post-Doctoral Fellow with the Biomedical Imaging Group, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, and the Center for Biomedical Imaging, Université de Lausanne, Lausanne, from 2010 to 2012. He was engaged in research on the design of steerable wavelet frames in multiple dimensions, learning techniques for adaptive frame design, and bioimaging. Since March 2012, he has been a Post-Doctoral Fellow with the Langone Medical Center, New York University, New York. His current research interests include neuroscience, in particular, synaptic mechanisms by using advanced microscopy and image processing techniques.

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

He was a Scientist with the National Institutes of Health, Bethesda, MD, from 1985 to 1997. He is currently a Full Professor and the Director of the Biomedical Imaging Group, EPFL. He is the author or co-author of 200 journal papers. His current research interests include biomedical image processing, sampling theories, multiresolution algorithms, wavelets, and the use of splines for image processing.

Dr. Unser was an Associate Editor-in-Chief of the IEEE TRANSACTIONS ON MEDICAL IMAGING from 2003 to 2005, and an Associate Editor of the same journal from 1999 to 2002 and from 2006 to 2007, the IEEE TRANSACTIONS ON IMAGE PROCESSING from 1992 to 1995, and the IEEE SIGNAL PROCESSING LETTERS from 1994 to 1998. He is currently an Editorial Board Member of *Foundations and Trends in Signal Processing*, and *Sampling Theory in Signal and Image Processing*. He co-organized the first IEEE International Symposium on Biomedical Imaging in 2002. He was the Founding Chair of the technical committee of the IEEE-SP Society on Bio Imaging and Signal Processing.