

Frame Isotropic Multiresolution Analysis for Micro CT Scans of Coronary Arteries

Bernhard G. Bodmann^a, Manos Papadakis^b, Donald J. Kouri^c, S. David Gertz, MD^{e,g}, Paul Cherukuri^f, Deborah Vela, MD^g, Gregory W. Gladish, MD^h, Dianna D. Cody^h, Ibrahim Aboshady, MD^g, Jodie L. Conyers^g, James T. Willerson, MD^g and S. Ward Casscells, MD^g

^aDepartment of Applied Mathematics, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada

^bDepartment of Mathematics, University of Houston, 651 Phillip G. Hoffman Hall, Houston, TX 77204-3008

^cDepartment of Chemistry, University of Houston, Houston, TX 77204-5003

^eDepartment of Anatomy and Cell Biology, The Hebrew University-Hadassah Medical School, Jerusalem, Israel

^fDepartment of Chemistry, Rice University, P.O. Box 1892, MS-60, Houston, TX 77251-1892

^gUniversity of Texas Health Science Center at Houston and Texas Heart Institute, St. Luke's Episcopal Hospital, 6770 Bertner Avenue MC 2-255 POB 20345, Houston, Texas 77225

^hDivision of Diagnostic Imaging, The University of Texas M.D. Anderson Cancer Center, 1515 Holcombe Blvd., Houston, TX 77030

ABSTRACT

We analyze localized textural consistencies in high-resolution Micro CT scans of coronary arteries to identify the appearance of diagnostically relevant changes in tissue. For the efficient and accurate processing of CT volume data, we use fast algorithms associated with three-dimensional so-called isotropic multiresolution wavelets that implement a redundant, frame-based image encoding without directional preference. Our algorithm identifies textural consistencies by correlating coefficients in the wavelet representation.

Keywords: non-separable multiresolution analysis, wavelets, frames, random fields, segmentation

The general objective of this paper is to introduce texture-based image analysis for the processing of high-resolution CT scans. Our test data was obtained from scanning coronary artery specimens with a General Electric RS-9 Micro CT scanner providing images with cubic voxels of sidelength $27\mu\text{m}$. The motivation for our work is the need to distinguish various types of diagnostically relevant tissue in atherosclerotic plaque found in coronary arteries. It has been suggested that the presence of lipid-rich, lightly calcified lesions close to the arterial lumen is associated with an increased risk of plaque rupture and subsequent acute myocardial infarction.^{1,2} The main concern in this paper is the identification of such lesions within the surrounding fibromuscular tissue. The presence of noise makes this task difficult to achieve with usual threshold-based methods. On the other hand, suppressing noise by smoothing obscures the difference between intensity fluctuations characteristic of speckled calcific deposits embedded in lipid and intensity patterns caused by natural variations in fibromuscular tissue. The application of our algorithm to Micro CT data demonstrates the detection of significant changes in textural consistency that reflect deviations in the structure of tissue.

The concept implemented in our algorithm addresses three main challenges: *1. The image analysis should be applied to the entire, three-dimensional volume of data obtained from CT*, instead of separately treating two-dimensional slices. To achieve the computational efficiency desired for the proposed texture identification, we

Further author information for correspondence: E-mail for B.G.B.: bgb@uwaterloo.ca; E-mail for M.P.: mpapadakis@math.uh.edu; E-mail for D.J.K.: kouri@uh.edu; E-mail for S.D.G.: dgertz@uth.tmc.edu; E-mail for P.C.: cheru@rice.edu; E-mail for D.V.: Deborah.Vela@uth.tmc.edu; E-mail for G.W.G.: G.Gladish@di.mdacc.tmc.edu; E-mail for D.D.C.: dcody@di.mdacc.tmc.edu; E-mail for I.A.: iaboshady@uth.tmc.edu; E-mail for J.L.C.: Jodie.L.Conyers@uth.tmc.edu; E-mail for S.W.C.: s.ward.casscells@uth.tmc.edu

encode the CT images with a so-called fast isotropic wavelet transform. The wavelet transform separates features and textures of different scales in the image, the proportions of which indicate different types of tissue. *2. Features and textures should be identified regardless of their orientation.* Unlike traditional fast wavelet transforms that rely on tensor product constructions, our encoding does not introduce any artificial directional bias in the separation of scales. *3. Information should be extracted in a reliable, unambiguous way.* The encoded images are analyzed for textural consistencies in order to infer the tissue character with a localized version of statistical hypothesis testing.

This paper explains the mathematical structure of our image processing algorithm and demonstrates its application to a Micro CT scan of a human coronary artery. Section 1 contains the details of the Multiresolution Analysis underlying our Fast Isotropic Wavelet Transform. Section 2 describes the statistical analysis of images in the wavelet representation. Finally, Section 3 demonstrates the wavelet-based texture segmentation algorithm applied to Micro CT data and compares an example of image segmentation to tissue characterization by histology.

1. IMAGE ENCODING BY A FAST ISOTROPIC WAVELET TRANSFORM

With the help of wavelet analysis, the information contained in a digital image is separated into features and textures of different scales, i.e. fine-grained vs. coarse-grained levels of detail.

In order to process the large volumes of data generated by CT scans, we use fast algorithms associated with novel isotropic, three-dimensional wavelets. Unlike traditional, directionally biased wavelets, the isotropic wavelets developed by Papadakis and co-workers³ show the same sensitivity for features and textures, regardless of their orientation with respect to any fixed Cartesian coordinate system. The use of a redundant encoding based on frame theory is the key feature allowing us to avoid directional bias.

1.1. Theoretical Background

For the purposes of this paper, an *image* constitutes a square-integrable function on n -dimensional Euclidean space \mathbb{R}^n , or a square-summable set of coefficients indexed by an n -dimensional grid \mathbb{Z}^n . In the latter case, we say the image is in the *digital domain*, and identify the values of the coefficients with those of voxels in n dimensions.

DEFINITION 1.1. *A finite or countable subset $\{f_k : k \in \mathbb{K}\}$ of a Hilbert space H is a frame of H if there exist constants $c_1, c_2 > 0$ such that for every f in H , we have the norm equivalence $c_1 \|f\|^2 \leq \sum_{k \in \mathbb{K}} |\langle f, f_k \rangle|^2 \leq c_2 \|f\|^2$.⁴⁻⁸ If $c_1 = c_2 = 1$ we call $\{f_k : k \in \mathbb{K}\}$ a Parseval frame of H .*

If $\{f_k : k \in \mathbb{K}\}$ is a Parseval frame, then the identity $f = \sum_{k \in \mathbb{K}} \langle f, f_k \rangle f_k$ holds for every $f \in H$. A generalization of this identity is true for every frame. Parseval frames behave very similar to orthonormal bases, but they may not be minimal sets with a dense span in H .

Allowing frames instead of orthonormal bases gives rise to a structure of multiresolution analysis theory that generalizes the classical one.

Let the translations $\{T_{\mathbf{k}}\}_{\mathbf{k} \in \mathbb{Z}^n}$ be given by $T_{\mathbf{k}} f(\mathbf{x}) = f(\mathbf{x} - \mathbf{k})$, and the dyadic dilation operator D by $Df(\mathbf{x}) = 2^{n/2} f(2\mathbf{x})$, with $f \in L^2(\mathbb{R}^n)$.

DEFINITION 1.2. *A GFMRA⁹ of $L^2(\mathbb{R}^n)$ is an increasing family of closed subspaces $\{V_j\}_{j \in \mathbb{Z}}$ satisfying the following properties:*

- *The subspaces observe $V_j = D^j(V_0)$ for every $j \in \mathbb{Z}$.*
- *The intersection of subspaces is $\bigcap_{j \in \mathbb{Z}} V_j = \{0\}$ and the union $\bigcup_{j \in \mathbb{Z}} V_j$ is dense in $L^2(\mathbb{R}^n)$.*
- *There exists a finite or countably infinite set of functions $\{\phi_i, i = 1, 2, \dots\}$ belonging to V_0 , such that the translates $\{T_{\mathbf{k}} \phi_i : \mathbf{k} \in \mathbb{Z}, i = 1, 2, \dots\}$ form a frame of V_0 . We refer to the set $\{\phi_i\}$ as frame scaling functions. A finite or countably infinite set of frame wavelets (or simply wavelets) associated with $\{V_j\}_{j \in \mathbb{Z}}$ is a set of functions $\{\psi_p : p = 1, 2, \dots\}$ such that $\{T_{\mathbf{k}} \psi_p : \mathbf{k} \in \mathbb{Z}^n, p = 1, 2, \dots\}$ is a frame of $W_0 := V_0^\perp \cap V_1$.*

The general structure of GFMRA theory allows frames for the core subspaces to be generated with abelian “translation” groups other than $\{T_k\}$.⁹ Moreover, GFMRA may be constructed with a variety of dilation operators, thereby providing a lot of design freedom. Here, we have exploited this freedom to obtain good symmetry properties for scaling functions and wavelets.

For the separation of detail-levels we use a Fast Wavelet Algorithm based on a new type of Multiresolution Analysis (MRA) called *First Generation Isotropic MRA*,³ which we will show is an example of a GFMRA. Our implementation augments this structure with additional oversampling and asymmetric filters to achieve better localization properties. For simplicity, we first explain the underlying concept without additional oversampling.

DEFINITION 1.3. A First Generation Isotropic MRA of $L^2(\mathbb{R}^n)$ is an increasing family of closed subspaces $\{V_j\}_{j \in \mathbb{Z}}$ such that each V_j contains all functions in $L^2(\mathbb{R}^n)$ whose Fourier transforms vanish outside $2^j\mathbb{B}$, with the closed ball $\mathbb{B} := \{\xi \in \mathbb{R}^n : |\xi| \leq b_0\}$ of radius $1/4 < b_0 < 1/2$. The Fourier transform we use is defined for $f \in L^1(\mathbb{R}^n)$ by $\hat{f}(\xi) := \int_{\mathbb{R}^n} f(x) e^{-2\pi i x \cdot \xi} dx$, $\xi \in \mathbb{R}^n$ and as usual extended to $f \in L^2(\mathbb{R}^n)$. We also adopt the notations $\mathbb{T}^n := [-1/2, 1/2)^n$ and $e_k(\xi) := e^{-2\pi i(\xi \cdot k)}$, $\xi \in \mathbb{R}^n$ and $k \in \mathbb{Z}^n$.

In order to show that this defines an example of a GFMRA, we recall a well-known fact concerning the construction of Parseval frames.

LEMMA 1.4.¹⁰ Let H be a Hilbert space and P be an orthogonal projection defined on H . If $\{\epsilon_i : i \in I\}$ is an orthonormal subset of H , or a Parseval frame of H , then $\{P\epsilon_i : i \in I\}$ is a Parseval frame of $P(H)$.

The linear mapping on $L^2(\mathbb{R}^n)$ defined by multiplication with the characteristic function $\hat{\phi} = \chi_{\mathbb{B}}$ is an orthogonal projection. Then, Lemma 1.4 implies that $\{e_k \chi_{\mathbb{B}} : k \in \mathbb{Z}^n\}$ is a Parseval frame for the closed subspace it generates. Applying the inverse Fourier transform, we obtain that $\{T_k \phi : k \in \mathbb{Z}^n\}$ is a Parseval frame for the closed subspace it generates, which is V_0 . One may then easily check that the sequence $\{V_j\}_{j \in \mathbb{Z}}$ satisfies all the remaining properties of a GFMRA of $L^2(\mathbb{R}^n)$. In particular, the function ϕ is a frame scaling function for $\{V_j\}_{j \in \mathbb{Z}}$. The frame scaling function ϕ is radial, i.e. invariant under all rotations about the origin. We call a function that is a translate of a radial function *isotropic*. A similar symmetry property holds for the spaces $\{V_j\}$ and $\{W_j\}$.

PROPOSITION 1.5. Each V_j is invariant under any rotation about any center $x \in \mathbb{R}^n$. The same is true for each detail space $W_j = V_{j+1} \cap V_j^\perp$, $j \in \mathbb{Z}$.

Proof. To see this, we note that any such rotation is a rotation about the origin, conjugated with a translation that moves x to the origin. Translating the point x to the origin amounts to a modulation in the frequency domain, which leaves the support unchanged and thus V_j invariant. By definition, V_j is also invariant under rotations about the origin. The detail space $W_j = V_j^\perp \cap V_{j+1}$ contains all functions whose Fourier transforms vanish outside the spherical shell $2^j\mathbb{S}$ with inner radius $2^j b_0$ and outer radius $2^{j+1} b_0$, so the same argument applies. \square

Several sets of *frame wavelets* $\{\psi_p : p = 1, 2, \dots\}$ are associated with $\{V_j\}_j$, in the sense that $\{T_k \psi_p : k \in \mathbb{Z}^n, p = 1, 2, \dots\}$ is a frame of W_0 . We are interested in one particular set of isotropic ones, the integer translates of which form a Parseval frame of W_0 :

$$\hat{\psi}_i(\xi) = 2^{-n/2} e_{q_i}(\xi/2) \chi_{\mathbb{S}}(\xi) \quad \xi \in \mathbb{R}^n,$$

where $\{q_i : i = 0, 1, 2, \dots, 2^n - 1\}$ are the representatives of the quotient group $\mathbb{Z}^n/2(\mathbb{Z}^n)$. If, for example, $n = 2$, one may choose $q_0 = (0, 0)$, $q_1 = (1, 0)$, $q_2 = (0, 1)$ and, $q_3 = (1, 1)$.

The use of frames is a major difference to traditional, so-called separable MRA constructions in multidimensions, i.e. those derived from tensor products of one-dimensional MRAs. To remedy the directional bias in those, non-separable MRAs have been proposed, e.g. Refs 11–24 or other approaches where additional directional preference is incorporated, e.g. Refs. 25–33. None of these approaches give MRAs that are invariant under *all* rotations and not all of them lead to fast algorithms similar to those associated with separable MRAs.

Observing that the ball \mathbb{B} is contained in \mathbb{T}^n , a classical sampling theorem applies

THEOREM 1.6.³ Let f be in V_0 , then

$$f = \sum_{k \in \mathbb{Z}^n} f(k) T_k \phi, \tag{1}$$

where the right-hand side of the previous equation converges in the L^2 -norm and uniformly if $f \in V_0$ is continuous.

This theorem shows that the input of the fast wavelet algorithms resulting from the Isotropic MRA $\{V_j\}_{j \in \mathbb{Z}}$ consists of the samples of the input analog signal obtained on the usual cubic sampling grid \mathbb{Z}^n .

1.2. Implementation with Ideal Filters

The next task is to derive the low and high pass filters resulting from the First Generation IMRA. Although the coefficients $\{\alpha_{\mathbf{k}} = \langle D^* \phi, T_{\mathbf{k}} \phi \rangle\}_{\mathbf{k} \in \mathbb{Z}^n}$ are not the unique solution for the two-scale equation

$$D^* \phi = \sum_{\mathbf{k} \in \mathbb{Z}^n} \alpha_{\mathbf{k}} T_{\mathbf{k}} \phi, \quad (2)$$

we adopt the convention that the low pass filter is the \mathbb{Z}^n -periodic function

$$m_0 := \sum_{\mathbf{k} \in \mathbb{Z}^n} \langle D^* \phi, T_{\mathbf{k}} \phi \rangle e_{\mathbf{k}}.$$

Applying the Fourier transform to both sides of (2), we obtain

$$\hat{\phi}(2\xi) = 2^{-n/2} m_0(\xi) \hat{\phi}(\xi) \quad \text{a.e.},$$

as in the case of orthonormal scaling functions. In fact, by our convention the support of m_0 is contained in that of $\hat{\phi}$, and therefore the two-scale equation determines the function m_0 uniquely. Following our convention, $m_0(\xi) = 2^{n/2} \chi_{\mathbb{B}/2}(\xi)$ a.e. in \mathbb{T}^n . The same argument gives the form of the high pass filters $h_i(\xi) = e_{\mathbf{q}_i}(\xi) \chi_{\mathbb{S}/2}(\xi)$, where $\xi \in \mathbb{T}^n$ and $i = 0, 1, \dots, 2^n - 1$.

Since the map $f \mapsto \sum_{\mathbf{k}} \langle f, T_{\mathbf{k}} \phi \rangle e_{\mathbf{k}}$ is an isometry, $T_{2\mathbf{k}} D^* = D^* T_{\mathbf{k}}$, $\{D^* T_{\mathbf{k}} \phi : \mathbf{k} \in \mathbb{Z}^n\}$ is a Parseval frame of V_{-1} , and $\{e_{\mathbf{k}} \chi_{\mathbb{S}/2} : \mathbf{k} \in \mathbb{Z}^n\}$ is a Parseval frame of W_{-1} , we obtain that $\{e_{2\mathbf{k}} m_0 : \mathbf{k} \in \mathbb{Z}^n\} \cup \{e_{\mathbf{k}} \chi_{\mathbb{S}/2} : \mathbf{k} \in \mathbb{Z}^n\}$ is a Parseval frame of the space $\{\mu \in L^2(\mathbb{T}^n) : \mu(\xi) = 0 \text{ a.e. in } \mathbb{B}^c\}$. This observation gives rise to a new fast wavelet transform which we refer to as the *First Generation Fast Isotropic Wavelet Transform*.³⁴ This wavelet transform has the structure of a traditional MRA-based Fast Wavelet Transform (FWT).

REMARK 1.7. We note that the set $\{e_{\mathbf{k}} \chi_{\mathbb{S}/2} : \mathbf{k} \in \mathbb{Z}^n\}$ is equal to $\cup_{i=0}^{2^n-1} \{e_{2\mathbf{k}} h_i : \mathbf{k} \in \mathbb{Z}^n\}$. The latter union would yield 2^n high pass channels with decimation. In our implementation, all these channels are combined to an undecimated one using the filter h .

Accordingly, in our implementation the low-pass filtering is followed by decimation by a factor of two, while the high pass filtered signal stays undecimated.

In contrast to the usual tensor-product constructions, both filters (low and high pass) are isotropic. However, they are discontinuous in the frequency domain, which leads to ringing artifacts and bad localization properties of the decomposition. To rectify this problem we use smooth approximations for the ideal filters m_0 and h combined with so-called radial oversampling.

1.3. Smooth Filters and Radial Oversampling

Let $1/8 < b_2 < b_1 < b_0/2 < 1/4$. We begin with the analysis low-pass filter, a C^∞ -smooth, \mathbb{Z}^n -periodic function \tilde{m}_a satisfying the following three properties:

1. The support of \tilde{m}_a inside \mathbb{T}^n is the set $\mathbb{B}/2$.
2. The restriction $\tilde{m}_a|_{\mathbb{B}}$ is radial.
3. The filter assumes the value $\tilde{m}_a(\xi) = 2^{n/2}$ for all $|\xi| < b_1$.

To approximate the ideal high pass filter $\chi_{\mathbb{S}/2}$ on \mathbb{B} , we define the analysis high pass filter

$$\tilde{h}_a := 1 - \frac{\tilde{m}_a}{2^{n/2}} .$$

By definition, the two analysis filters overlap in the spherical shell with radii b_1 and $b_0/2$.

We use upper case characters to denote each filter in the spatial domain. Using this notation, the analysis filters produce analysis operators \mathcal{M}_a and \mathcal{H}_a given by

$$\mathcal{M}_a \sigma := \mathcal{D}(\tilde{M}_a * \sigma) \quad \text{and} \quad \mathcal{H}_a \sigma = \tilde{H}_a * \sigma, \quad \sigma \in \ell^2(\mathbb{Z}^n),$$

where \mathcal{D} denotes downsampling by a factor of 2,

$$(\mathcal{D}\sigma)^\wedge(\xi) = \frac{1}{2^n} \sum_{l=0}^{2^n-1} \hat{\sigma} \left(\frac{\xi}{2} + \gamma_l \right),$$

and $\{\gamma_l : l = 0, 1, \dots, 2^n - 1\}$ is a set of solutions of the equation $2\xi \equiv 0 \pmod{\mathbb{Z}^n}$ with $\xi \in \mathbb{T}^n$ and $\gamma_0 = \mathbf{0}$. The range space of the operator $\sigma \mapsto \tilde{M}_a * \sigma$ is the subspace of all signals whose Fourier transforms vanish outside the ball $\{\xi : |\xi| \leq b_0/2\}$. Likewise, the range of \tilde{H}_a contains all signals whose Fourier transforms vanish inside the ball $\{\xi : |\xi| < b_1\}$. Both signals $\tilde{M}_a * \sigma$ and $\tilde{H}_a * \sigma$ are oversampled, because the frequency support of the functions in the range of \tilde{M}_a or \tilde{H}_a is not the entire torus \mathbb{T}^n . In analogy with one-dimensional oversampling, we now use synthesis filters for the reconstruction with a bandwidth slightly greater than the bandwidth of the analysis filters. So, we select the synthesis low pass filter \tilde{m}_s as the C^∞ -smooth \mathbb{Z}^n -periodic function with the properties:

1. The support of \tilde{m}_s inside \mathbb{T}^n is the set $\{\xi \in \mathbb{T}^n : |\xi| \leq \frac{1}{4}\}$.
2. The restriction $\tilde{m}_s|_{\mathbb{B}}$ is radial.
3. The value $\tilde{m}_s(\xi) = 2^{n/2}$ is assumed whenever $\xi \in \mathbb{T}^n$ such that $|\xi| \leq b_0/2$.

Similarly, we select \tilde{h}_s to be C^∞ with the properties:

1. It assumes the value $\tilde{h}_s(\xi) = 1$ for all $\xi \in \mathbb{T}^n$ such that $|\xi| \geq b_1$.
2. The restriction $\tilde{h}_s|_{\mathbb{B}}$ is radial.
3. The filter \tilde{h}_s vanishes if $|\xi| < b_2$.

The operators $\mathcal{M}_s \sigma := \tilde{M}_s * (\mathcal{U}\sigma)$ and $\mathcal{H}_s \sigma := \tilde{H}_s * \sigma$ are the low and high pass synthesis operators respectively, where $\sigma \in \ell^2(\mathbb{Z}^n)$ and \mathcal{U} is the upsampling by a factor of two operator, $(\mathcal{U}\sigma)^\wedge(\xi) = \hat{\sigma}(2\xi)$.

THEOREM 1.8. *With the low and high pass filters defined above, any digital signal $\sigma \in \ell^2(\mathbb{Z}^n)$ can be reconstructed exactly,*

$$\sigma = \mathcal{M}_s \mathcal{M}_a \sigma + \mathcal{H}_s \mathcal{H}_a \sigma. \quad (3)$$

Proof. Eq. (3) is true if the operators \mathcal{M}_a , \mathcal{M}_s , \mathcal{H}_a and \mathcal{H}_s are defined with the low and high pass filters m_0 and h only, because $\{e_{2\mathbf{k}} m_0 : \mathbf{k} \in \mathbb{Z}^n\} \cup \{e_{\mathbf{k}} h : \mathbf{k} \in \mathbb{Z}^n\}$ is a Parseval frame of the space $\hat{V}_0 = \{q \in L^2(\mathbb{T}^n) : q(\xi) = 0 \text{ a.e. in } \mathbb{B}^c\}$. The use of the smooth filters \tilde{m}_a , \tilde{h}_a , \tilde{m}_s and \tilde{h}_s is possible by “radial” oversampling.

Let $\sigma \in \ell^2(\mathbb{Z}^n)$. Then $(\mathcal{H}_s \mathcal{H}_a \sigma)^\wedge = \tilde{h}_s \tilde{h}_a \hat{\sigma} = \tilde{h}_a \hat{\sigma}$, because $\tilde{h}_s = 1$ on the entire support of \tilde{h}_a .

From the definition of \mathcal{D} and \mathcal{U} , we infer $(\mathcal{U} \mathcal{M}_a \sigma)^\wedge$ is a $\frac{\mathbb{Z}^n}{2}$ -periodic function. Since \tilde{m}_s and \tilde{m}_a vanish outside the n -torus $\mathbb{T}^n/2$ and the ball $\mathbb{B}/2$ respectively, we obtain $(\mathcal{M}_s \mathcal{M}_a \sigma)^\wedge = \frac{1}{2^n} \tilde{m}_s \tilde{m}_a \hat{\sigma} = \frac{1}{2^{n/2}} \tilde{m}_a \hat{\sigma}$. The exactness of reconstruction is verified once we observe $\frac{\tilde{m}_a}{2^{n/2}} + \tilde{h}_a = 1$. \square

While giving perfect reconstruction, the low and high subband outputs of the above variant of the Fast Isotropic Wavelet Transform only approximately represent the decomposition $V_0 = V_{-1} \oplus W_{-1}$ of the theoretical prototype associated with the filters m_0 and h . We could stay arbitrarily close to the decomposition $V_0 = V_{-1} \oplus W_{-1}$ by choosing the difference $b_0/2 - b_1$ small, but need smoothed filters to avoid ringing and bad localization. For a quantitative discussion of this trade-off between frequency selectivity and filter length, see Ref. 35.

1.4. Specifics of Implementation

In our implementation we do not use the filters $\tilde{m}_{a,s}$ or $\tilde{h}_{a,s}$ described above but instead we use finite-length approximations. In order to guarantee perfect reconstruction, we define the (undecimated) high subband output of the input signal $\sigma \in \ell^2(\mathbb{Z}^n)$ as $\sigma - \mathcal{M}_s \mathcal{M}_a \sigma$. The filter length we used was optimized for a combination of speed and frequency selectivity.³⁵ The analysis low-pass filter has $9 \times 9 \times 9$ filter taps and the synthesis uses a $15 \times 15 \times 15$ filter size.

2. STATISTICAL IMAGE ANALYSIS

It is our experience that various tissue types appear with characteristic, fluctuating components in each resolution level. To recognize the occurrence of these components, we use statistical methods in our tissue classification scheme. Another reason why such methods are necessary is the lack of an a priori method to distinguish intensity fluctuations due to density variations in tissue and those due to noise in our images.

2.1. A statistical model for tissue in CT images

Our model for digital CT images is based on the following hypotheses: There are finitely many types of tissue (e.g.: lumen, calcified, lipid, fibrous, muscle cells) and an image is composed of segments containing these tissues. Each tissue type is represented by a random configuration of intensity values in an image. This randomness may consist of typical density fluctuations in the tissue and additive noise. Unlike more conventional approaches that characterize tissue by an average voxel value, our characterization scheme uses the correlations between all voxels belonging to a given tissue sample. Moreover, our model of tissue includes multiple levels of resolution.

Tissue of one type is modeled by a wide-sense homogeneous, isotropic discrete random field with the property that its first and second moments can be obtained by averaging a tissue sample over all shifts. In the following precise definition of a tissue type, we have suppressed the dependence of random variables on the outcome $\omega \in \Omega$ in our probability space. Accordingly, the convolution of a random field $\sigma : \mathbb{Z}^n \times \Omega \rightarrow \mathbb{R}$ with a digital filter G is denoted by $(G * \sigma)_k := \sum_{l \in \mathbb{Z}^n} G_{k-l} \sigma_l$.

DEFINITION 2.1. A tissue type τ at resolution level j is a family of real-valued random variables $\{\tau_k^{(j)}\}_{k \in \mathbb{Z}^n}$ over some probability space $(\Omega, \mathbb{P}, \mathcal{F})$ with the following properties:

1. For each $k \in \mathbb{Z}^n$, $\tau_k^{(j)} \in L^2(\mathbb{P})$. The expected value $\mathbb{E}[\tau_k^{(j)}] = \bar{\tau}^{(j)}$ is independent of k , and the covariance matrix $C^{(j)}$ with entries $C_{k,k'}^{(j)} := \mathbb{E}[(\tau_k^{(j)} - \bar{\tau}^{(j)})(\tau_{k'}^{(j)} - \bar{\tau}^{(j)})]$ is a bounded operator on $\ell^2(\mathbb{Z}^n)$. Moreover, the entries $C_{k,k'}^{(j)}$ depend only on the distance $|k - k'|$ of $k, k' \in \mathbb{Z}^n$.
2. For each $k \in \mathbb{Z}^n$ and each filter G with (absolutely) summable taps, the spatial average of $(G * \tau^{(j)})_k$ converges almost surely,

$$\lim_{\mathbb{V} \nearrow \mathbb{Z}^n} \frac{1}{|\mathbb{V}|} \sum_{l \in \mathbb{V}} (G * \tau^{(j)})_{k+l} = \bar{\tau}^{(j)} \sum_k G_k$$

as the finite set \mathbb{V} grows and eventually contains any given finite set of indices.

3. For each $k \in \mathbb{Z}^n$ and each finite subset $\mathbb{V} \subset \mathbb{Z}^n$, we abbreviate the local average $\bar{\tau}_k^{(j)}(\mathbb{V}) = \frac{1}{|\mathbb{V}|} \sum_{l \in \mathbb{V}} \tau_{k+l}^{(j)}$ and define an approximate, random covariance matrix $C^{(j)}(\mathbb{V})$ with entries

$$C_{k,k'}^{(j)}(\mathbb{V}) = \frac{1}{|\mathbb{V}|} \sum_{l \in \mathbb{V}} (\tau_{k+l}^{(j)} - \bar{\tau}_k^{(j)}(\mathbb{V}))(\tau_{k'+l}^{(j)} - \bar{\tau}_{k'}^{(j)}(\mathbb{V})).$$

In this matrix notation, we require that for each $\mathbf{k}, \mathbf{k}' \in \mathbb{Z}^n$ and every summable filter G , almost surely

$$\lim_{\mathbb{V} \nearrow \mathbb{Z}^n} (G^* C^{(j)}(\mathbb{V}) G)_{\mathbf{k}, \mathbf{k}'} = (G^* C^{(j)} G)_{\mathbf{k}, \mathbf{k}'}.$$

REMARKS 2.2. The first property states homogeneity and isotropy and ensures that the quantities appearing in the subsequent properties are well-defined. The remaining two properties specify how expectation value and covariance matrix arise from spatial averaging. Birkhoff's ergodic theorem concludes that these two properties are satisfied by all locally square-integrable ergodic random fields.³⁶

A simple example for a tissue type at some resolution level, say $j = 0$, could consist of independent, identically distributed (i.i.d.) square-integrable random variables for the voxel values $\{\tau_{\mathbf{k}}^{(0)}\}$. The averaging properties are then a consequence of the ergodicity of the joint distribution of the i.i.d. variables with respect to the translation group \mathbb{Z}^n .³⁶

The following theorem shows that applying a summable filter preserves the properties required of a tissue type. Consequently, convolving a random field of independent, identically distributed square-integrable voxel values with a summable filter gives a tissue type with voxels that are generally no longer independent. Finally, a convex combination of such random fields with different underlying individual distributions and filters allows to recreate an intuitive picture of tissue, containing objects of a range of intensities and sizes.

THEOREM 2.3. *The properties of a tissue type at resolution level j are preserved when convolving it with a summable filter G .*

Proof. We need to show that under convolution with G , $\sum_{\mathbf{k} \in \mathbb{Z}^n} |G_{\mathbf{k}}| < \infty$, the properties of a tissue type at resolution level j are preserved. The first property is true because filtering and expectation can be interchanged. Next, the covariance matrix of the filtered random field is the previous covariance matrix conjugated with G . Due to its summability, the filter is a bounded operator on $\ell^2(\mathbb{Z}^n)$, the norm of the operator $G^* C^{(j)} G$ is finite. Properties 2 and 3 ensure by definition that they remain valid when the field $\tau^{(j)}$ is replaced with $G * \tau^{(j)}$. \square

Thus, applying a summable low pass filter to a tissue type τ at resolution level j , followed by downsampling, produces a corresponding tissue type at resolution level $j - 1$. This ensures the consistency of our definition with the approximate implementation of a multiresolution structure described in the previous section on smoothed filters.

2.2. Tissue segmentation

Our procedure for the identification of tissue types is based on analyzing the statistics of wavelet coefficients belonging to a given tissue sample. We compare these statistics with reference statistics that are derived from known tissue types. When significant deviations from reference statistics occur (wavelet coefficients violate bounds that are observed with a high probability given a specific tissue type), we infer that we have identified tissue that is different from the reference type.

DEFINITION 2.4. *A tissue segment $\{\tau_{\mathbf{k}}^{(j)}\}_{\mathbf{k} \in \mathbb{V}}$ at resolution level j is specified by selecting a tissue type $\tau^{(j)}$ and an index set $\mathbb{V} \subset \mathbb{Z}^n$. A digital image at a resolution level j that contains various tissues is the union of segments of different tissue types at this resolution level.*

In order to separate tissue types with statistically different behavior, our algorithm proceeds in two steps.

Training step. At first, we pick a segment that is assumed to contain only one type of tissue, henceforth called reference tissue, as a training set.

a) *Customization of prediction filters for each subband.* The idea for this prediction filter is that each voxel gets replaced by a linear prediction based on its neighbors. The filter is chosen so that it best estimates each voxel in the least squares sense when averaged over the segment of reference tissue. The theorem contained in the remainder of this section states how, in the limit of arbitrarily large training sets, this becomes a least squares estimator with respect to the probability measure governing the reference tissue. Selecting a confidence interval for the difference between predicted and actual voxel values is therefore a statistically reliable way to decide whether the values of a voxel and its neighbors belong to tissue of reference type.

For practical purposes, it is enough to consider finite-length prediction filters with bounded filter taps. This is reflected in our definition of prediction filters. In order to avoid introducing an artificial directional bias in the prediction, we require that the filter observes the same isotropy condition as the covariance matrix of the tissue, see Property 1 in Definition 2.1.

DEFINITION 2.5. A finite-length prediction filter $P^{(j)}$ for the resolution level $j \in \mathbb{Z}$ is given by a $2^j \mathbb{Z}^n$ -periodic trigonometric polynomial $p^{(j)}(\xi) = \sum_{\mathbf{k} \neq \mathbf{0}}^{\|\mathbf{k}\| \leq \Lambda} P_{\mathbf{k}}^{(j)} e_{\mathbf{k}}(\xi/2^j)$ in the frequency domain, such that the filter taps $\{P_{\mathbf{k}}^{(j)}\}$ depend only on the magnitude $|\mathbf{k}|$ of the index $\mathbf{k} \in \mathbb{Z}^n$, $|\mathbf{k}| \leq \Lambda$. We denote the space of prediction filters of maximal length Λ , maximal ℓ^2 -norm ρ and resolution level j as $\mathcal{P}_{\Lambda, \rho}^{(j)} := \{P^{(j)} : P_{\mathbf{0}}^{(j)} = 0, P_{\mathbf{k}}^{(j)} = P_{\mathbf{k}'}^{(j)} \text{ whenever } |\mathbf{k}| = |\mathbf{k}'|, \sum_{|\mathbf{k}| \leq \Lambda} |P_{\mathbf{k}}^{(j)}|^2 \leq \rho^2\}$.

DEFINITION 2.6. Given a tissue of type τ at a resolution level j , and a prediction filter $P^{(j)}$, we define its mean-square error to be

$$Q(P^{(j)} - I) := \mathbb{E}[(P^{(j)} - I) * (\tau^{(j)} - \bar{\tau}^{(j)})_{\mathbf{k}}]^2 = ((P^{(j)} - I)^* C^{(j)} (P^{(j)} - I))_{\mathbf{k}, \mathbf{k}},$$

where $\mathbf{k} \in \mathbb{Z}^n$ is arbitrary, and I denotes the digital all-pass filter corresponding to the constant function $\iota(\xi) = 1$ in the frequency domain. We say that the locally averaged square error at $\mathbf{k} \in \mathbb{Z}^n$ is

$$Q_{\mathbb{V}, \mathbf{k}}(P^{(j)} - I) := ((P^{(j)} - I)^* C^{(j)}(\mathbb{V}) (P^{(j)} - I))_{\mathbf{k}, \mathbf{k}}.$$

THEOREM 2.7. Let τ be a tissue at resolution level j , and fix a maximal length $\Lambda \in \mathbb{N}$ and norm $\rho > 0$ for all prediction filters under consideration. Let $P^{(j)}$ be a prediction filter that minimizes $Q(P^{(j)} - I)$ in $\mathcal{P}_{\Lambda, \rho}^{(j)}$. Fix $\mathbf{k} \in \mathbb{Z}^n$ and choose for each $\mathbb{V} \subset \mathbb{Z}^n$ a prediction filter $P^{(j)}(\mathbb{V}) \in \mathcal{P}_{\Lambda, \rho}^{(j)}$ such that $Q_{\mathbb{V}, \mathbf{k}}(P^{(j)}(\mathbb{V}) - I)$ is minimized. Then we obtain almost surely

$$\lim_{\mathbb{V} \nearrow \mathbb{Z}^n} Q(P^{(j)}(\mathbb{V}) - I) = Q(P^{(j)} - I). \quad (4)$$

Proof. To establish this equality, we consider the space of filters of length Λ as a finite-dimensional vector space denoted by \mathcal{V}_{Λ} , which is equipped with the ℓ^2 -norm on the filter taps. The first step to prove the claimed equality is to show uniform convergence of $Q_{\mathbb{V}, \mathbf{k}}$ on compact sets in the space \mathcal{V}_{Λ} . To this end, we note that only finitely many coefficients of either $C^{(j)}$ or $C^{(j)}(\mathbb{V})$ are relevant for evaluating $Q(G)$ or $Q_{\mathbb{V}, \mathbf{k}}(G)$ for $G \in \mathcal{V}_{\Lambda}$. Since the entries of the matrix $C^{(j)}(\mathbb{V})$ converge almost surely to the corresponding entries of $C^{(j)}$, we have uniform convergence of the finite-dimensional restrictions of those matrices. Therefore, for any arbitrarily small $\epsilon > 0$, there exists a sufficiently big index set \mathbb{V}_0 such that almost surely for all $\mathbb{V} \supset \mathbb{V}_0$ and all $G \in \mathcal{V}_{\Lambda}$, we have $|Q(G) - Q_{\mathbb{V}, \mathbf{k}}(G)| \leq \epsilon \|G\|^2$. Moreover, if G is in the compact set $\mathcal{P}_{\Lambda, \rho}^{(j)}$, we conclude

$$|Q_{\mathbb{V}, \mathbf{k}}(G - I) - Q(G - I)| \leq \epsilon(\rho + 1)^2.$$

To begin with the estimates needed for the proof, we observe $Q(P^{(j)} - I) \leq \liminf_{\mathbb{V}} Q(P^{(j)}(\mathbb{V}) - I)$, because $P^{(j)}$ is a minimizer for Q . We can further estimate that for any given $\epsilon > 0$, when \mathbb{V} is sufficiently large, $Q(P^{(j)} - I) \leq \limsup_{\mathbb{V}} Q(P^{(j)}(\mathbb{V}) - I) \leq \limsup_{\mathbb{V}} Q_{\mathbb{V}, \mathbf{k}}(P^{(j)}(\mathbb{V}) - I) + \epsilon(\rho + 1)^2$. Using that $P^{(j)}(\mathbb{V})$ is a minimizer for $Q_{\mathbb{V}, \mathbf{k}}$, we further increase $Q_{\mathbb{V}, \mathbf{k}}$ when replacing $P^{(j)}(\mathbb{V})$ with $P^{(j)}$. By uniform convergence, $Q(P^{(j)} - I) \leq \liminf_{\mathbb{V}} Q(P^{(j)}(\mathbb{V}) - I) \leq \limsup_{\mathbb{V}} Q(P^{(j)}(\mathbb{V}) - I) \leq \lim_{\mathbb{V}} Q_{\mathbb{V}, \mathbf{k}}(P^{(j)} - I) + \epsilon(\rho + 1)^2 = Q(P^{(j)} - I) + \epsilon(\rho + 1)^2$, and since ϵ can be chosen arbitrarily small, the claimed equality follows. \square

COROLLARY 2.8. We can repeat the above theorem by replacing the field variables $\{\tau_{\mathbf{k}}\}$ with $\{(G * \tau)_{\mathbf{k}}\}$, where G is a summable filter such as a smooth filter used to obtain the low or high subband output.

Therefore, we can optimize the prediction filter in different subbands separately. This optimization scheme ignores correlations between different subbands. In our implementation, we have used correlations between “neighboring” subbands by stacking them to random fields on the space \mathbb{Z}^{n+1} , with the extra dimension denoting the subband index. In this case, the scheme has to be modified by allowing different expectation values for the

random field in each n -dimensional layer, and by restricting the homogeneity condition to the components of the covariance matrix within one layer, as well as by requiring an appropriately reduced isotropy of the prediction filters.

b) Calculate tolerance intervals. Again involving the training set, we then calculate *tolerance intervals* so that in each subband W_j , of the given coefficient appearing with reference tissue, only a small fraction $\varpi_j > 0$ is falsely labeled as outliers.

Choosing a small fraction ϖ_j of outliers in the training set guarantees, due to the assumed averaging property of tissue, that *each coefficient in a subband W_j has a probability close to ϖ_j of being rejected. As a consequence, setting all coefficients to zero except for the outliers suppresses the overall intensity of the reference tissue by roughly only a factor $1 - \varpi_j$.* This ensures that reference tissue is minimally affected when we threshold with respect to deviations from predicted values.

Application step. The next step is the application of our algorithm to an image composed of several tissue segments. For the reliability of the tissue segmentation, we require that the effect of surfaces between segments of different tissue types be small.

DEFINITION 2.9. *Given a digital filter G of finite length and a subset $\mathbb{V} \subset \mathbb{Z}^n$, we call $\mathbb{A} = \{\mathbf{k} \in \mathbb{V} : G_{\mathbf{k}-\mathbf{l}} \neq 0 \text{ for some } \mathbf{l} \in \mathbb{Z}^n \setminus \mathbb{V}\}$ the filter-induced surface region of this subset. This set contains all indices in \mathbb{V} that are affected by values of the field beyond \mathbb{V} when the filter G is applied. The complement of \mathbb{A} in \mathbb{V} is called the surface-insensitive interior.*

Given a fixed filter length Λ for the prediction filter in a subband, say W_0 , and a segment of reference tissue $\{\tau_{\mathbf{k}}\}_{\mathbf{k} \in \mathbb{V}}$, then each coefficient belonging to the surface-insensitive interior has a probability close to ϖ_j of being rejected. As a consequence, if the number of voxels in the surface-insensitive interior of tissue segments is large compared to those in surface regions, then for a large fraction of voxels (those contained in the surface-insensitive interior) belonging to reference tissue, the probability that they are classified as outliers converges to ϖ_j as the volume of the training set increases.

In order to identify segments that do not have the statistics of the reference tissue, we use the filters $P^{(j)}(\mathbb{V})$, apply them to each subband of the tissue under consideration, keep only those filtered coefficients that deviate from the original coefficients by more than the tolerance levels in the respective subbands, and reconstruct this “anomalous” part of the image. In other words, we eliminate the part of the signal that has the statistics of the reference tissue.

Based on this comparison method, we may use several reference types for classification purposes. If a segment is rejected by comparing with all types except for one, then it is unambiguously identified as tissue of this type. If it is rejected by comparison with all but two types, we conclude it is either of those, etc.

The idea of exploiting correlations in multiscale representations of images for denoising and texture detection has already been investigated, see e.g. Ref. 37. The novelty of the current work is the use of First Generation Isotropic Multiresolution Analysis for fast image encoding without directional bias. Another difference to previous results is that we use non-parametric methods in our statistical image analysis. Such methods are feasible within our rigid framework of homogeneous random fields with moment averaging properties, whereas for applications in general-purpose image processing one typically has to make the assumption of having parametric distributions of voxel intensities.^{37, 38}

2.3. Specifics of Implementation

The filter lengths for all $P^{(j)}(\mathbb{V})$ were chosen to be one (nearest neighbors only). Including the next coarser and next finer scale in the least-squares optimization of the filter improved performance because of correlations between subbands of tissue types and also because of the use of the finite-length approximations for the low-pass analysis and reconstruction filters used in the Fast Isotropic Wavelet Transform.

The training sets we used for extracting tissue statistics were segments of $30 \times 30 \times 30$ voxels at the highest resolution level, located in positions that by anatomical considerations only contained one tissue type.

3. APPLICATION TO MICRO CT DATA

Our initial studies have found that image processing via these wavelet algorithms permits three-dimensional, high resolution digital discrimination between calcific deposits, lipid-rich tissue that is lightly calcified, and the surrounding lipid-poor (fibrous) tissue.




Figure 1. Comparison of unprocessed and processed data obtained from Micro CT scan of a human coronary artery containing atherosclerotic plaque with a calcium-lipid complex. Left panel shows three-slice view of unprocessed data. Middle panel shows two slices with superimposed 3D rendering of threshold-based processed data. Right panel shows two processed slices with fibromuscular tissue in black and a superimposed 3D rendering of anomalous tissue. Lumen labeled L, arrows indicating calcific deposits with associated lipid-rich tissue. Wavelet-based tissue identification shown here based on statistical deviation of texture from background of fibromuscular tissue.

In Figure 1, we compare the unprocessed data obtained from scanning a coronary artery specimen with results that are based on thresholds or on textural analysis. We show that the wavelet decomposition segments the image into contiguous parts that have statistically distinct textural consistencies. We find that changes in




Figure 2. Comparison of a slice of processed data from Figure 1 with histology. Left panel shows a two-dimensional representation of the isotropic wavelet analysis applied to the Micro CT image. Processing identifies calcific deposits associated with previously undetected lipid rich core (arrows). Right panel shows partially decalcified histologic section through area depicted on left confirming the presence of calcific deposits within the plaque associated with lipid-rich pultaceous debris (H&E stain). Lumen labeled L in both images.

textural consistencies identified in the wavelet representation reflect deviations in the structural components of tissue.

In Figure 2, we validate the tissue identification algorithm by comparison with histological analysis. The region that has been identified as having a textural consistency which is different from that of fibromuscular tissue is verified as lipid-rich, lightly calcified lesion. Our laboratory is currently processing a number of specimens to test the reliability of the correlation between our tissue discrimination algorithm and histology.

ACKNOWLEDGMENTS

This research was partially supported by the following grants: University of Houston TLCC funds, NSF-DMS 0406748, by a subcontract from the University of Texas Health Science Center's "T5"-grant and by the R.A. Welch Foundation.

REFERENCES

1. M. Naghavi, P. Libby, E. Falk, S. Casscells, S. Litovsky, J. Rumberger, J. Badimon, C. Stefanadis, and et al, "From vulnerable plaque to vulnerable patient: a call for new definitions and risk assessment strategies: Part I," *Circulation* **108**(14), pp. 1664–1672, 2003. Full-text available online to subscribers using PubMed's LinkOut feature, <http://www.ncbi.nlm.nih.gov/entrez>.
2. M. Naghavi, P. Libby, E. Falk, S. Casscells, S. Litovsky, J. Rumberger, J. Badimon, C. Stefanadis, and et al, "From vulnerable plaque to vulnerable patient: a call for new definitions and risk assessment strategies: Part II," *Circulation* **18**(15), pp. 1772–1778, 2003.
3. M. Papadakis, G. Gogoshin, I. Kakadiaris, D. Kouri, and D. Hoffman, "Non-separable radial frame multiresolution analysis in multidimensions," *Numer. Function. Anal. Optimization* **24**, pp. 907–928, 2003.
4. J. Benedetto and M. Frazier, *Wavelets: Mathematics and Applications*, CRC Press, Boca Raton, FL, 1994. Chapters 3,7.
5. K. Grochenig, *Foundations of Time-Frequency Analysis*, Applied and Numerical Harmonic Analysis, Birkhauser, Boston, MA, 2001.
6. P. Casazza, "The art of frame theory," *Taiwanese J. Math* **4**, pp. 129–201, 2000.
7. P. Casazza, "Modern tools for Weyl-Heisenberg frames," *Advances in Imaging and Electron Physics* **115**, pp. 1–127, 2001.
8. D. Han and D. Larson, *Frames, Bases and Group Representations*, vol. 147 of *Memoirs*, American Mathematical Society, 2000.
9. M. Papadakis, "Generalized frame multiresolution analysis of abstract Hilbert spaces," in *Sampling, Wavelets, and Tomography*, J. Benedetto and A. Zayed, eds., pp. 179–223, Birkhauser, 2003.
10. A. Aldroubi, "Portraits of frames," *Proceedings of the American Mathematical Society* **123**, pp. 1661–1668, 1995.
11. E. Belogay and Y. Wang, "Arbitrarily smooth orthogonal nonseparable wavelets in R^2 ," *SIAM Journal of Mathematical Analysis* **30**, pp. 678–697, 1999.
12. A. Cohen and I. Daubechies, "Nonseparable bidimensional wavelet bases," *Revista Matematica Iberoamericana* **9**, pp. 51–137, 1993.
13. J. Kovacevic and M. Vetterli, "Nonseparable multidimensional perfect reconstruction filter-banks," *IEEE Transactions on Information Theory* **38**, pp. 533–555, 1992.
14. W. He and M. J. Lai, "Examples of bivariate nonseparable compactly supported orthonormal continuous wavelets," in *Wavelet Applications in Signal and Image Processing IV*, M. U. A. Aldroubi, A. Laine, ed., *Proceedings SPIE* **3169**, pp. 303–314, 1997.
15. W. Lawton and H. Resnikoff, "Multidimensional wavelet bases," February 1991. Techn. Rep. AD910130, Aware Inc., Bedford, MA.
16. K. Grochenig and W. Madych, "Multiresolution analysis, Haar bases and self-similar tilings," *IEEE Transactions on Information Theory* **38**, pp. 558–568, 1992.
17. A. Ayache, E. Belogay, and Y. Wang, "Orthogonal lifting: constructing new (symmetric) orthogonal scaling functions." Personal communication, 2002.

18. A. Ayache, "Construction de bases d'ondelettes orthonormées de $L^2(\mathbb{R}^2)$ non séparables, à support compact et de régularité arbitrairement grande," *Comptes Rendus Académie des Sciences de Paris* **325**, pp. 17–20, 1997.
19. A. Ayache, "Construction of non separable dyadic compactly supported wavelet bases for $L^2(\mathbb{R}^2)$ of arbitrarily high regularity," *Revista Matemática Iberoamericana* **15** **15**(1), pp. 37–58, 1999.
20. A. Ayache, "Some methods for constructing non separable, orthonormal, compactly supported wavelet bases," *Appl. Comput. Harmon. Anal.* **10**, pp. 99–111, 2001.
21. J. Derado, "Nonseparable, compactly supported interpolating refinable functions with arbitrary smoothness," *Appl. Comput. Harmon. Anal.* **10**(2), pp. 113–138, 2001.
22. C. Cabrelli, C. Heil, and U. Molter, *Self-Similarity and Multiwavelets in Higher Dimensions*, vol. 170 of *Memoirs, Amer. Math. Soc.*, 2004.
23. M. Antonini, M. Barlaud, and P. Mathieu, "Image coding using lattice vector quantization of wavelet coefficients," in *IEEE Internat. Conf. Acoust. Signal Speech Process.*, pp. 2273–2276, 1991.
24. B. Forster, T. Blu, and M. Unser, "A new family of complex rotation-covariant multiresolution bases in 2d," in *Wavelets X*, A. Aldroubi, A. Laine, and M. Unser, eds., *Proceedings SPIE* **5207**, pp. 475–479, 2003.
25. E. Candes and D. Donoho, "Ridgelets: a key to higher dimensional intermittency?," *Phil. Trans. R. Soc. London A*, pp. 2495–2509, 1999.
26. D. Donoho and X. Huo, "Beamlets and multiscale image analysis," *Springer Lecture Notes in Computational Science and Engineering* **20**, pp. 149–196, 2002.
27. J. Starck, E. Candes, and D. Donoho, "The curvelet transform for image denoising," *IEEE Transactions Image Processing* **11**(6), pp. 670–684, 2002.
28. A. Aldroubi, C. Cabrelli, and U. Molter, "Wavelets on irregular grids with arbitrary dilation matrices and frame atoms on $L^2(\mathbb{R}^d)$," *Appl. Comput. Harmon. Anal.* **17**(2), 2004. Special issue: Frames in Harmonic Analysis, Part II.
29. N. Kingsbury, "Image processing with complex wavelets," *Phil. Trans. R. Soc. London A* **357**, pp. 2543–2560, 1999.
30. I. Selesnick, "Smooth wavelet tight frames with zero moments," *Appl. Comput. Harmon. Anal.* **10**(2), pp. 163–181, 2001.
31. I. Selesnick and L. Sendur, "Iterated oversampled filter banks and wavelet frames," in *Wavelet Applications in Signal and Image Processing VIII*, M. U. A. Aldroubi, A. Laine, ed., *Proceedings of SPIE* **4119**, 2000.
32. I. Daubechies, B. Han, A. Ron, and Z. Shen, "Framelets: MRA-based constructions of wavelet frames," *Appl. Comput. Harmon. Anal.* **14**(1), pp. 1–46, 2003.
33. L. Shen, M. Papadakis, I. Kakadiaris, I. Konstantinidis, D. Kouri, and D. Hofmann, "Image denoising using tight frames." To appear in *IEEE Trans. Signal Proc.*, 2005.
34. M. Papadakis, G. Gogoshin, I. Kakadiaris, D. Kouri, and D. Hoffman, "Non-separable radial frame multiresolution analysis in multidimensions and the isotropic fast wavelet algorithms," in *Wavelets X*, A. Aldroubi, A. Laine, and M. Unser, eds., *Proceedings SPIE* **5207**, 2003.
35. B. Bodmann, M. Papadakis, and Q. Sun, "An inhomogeneous uncertainty principle for digital low-pass filters." Under review, 2005.
36. L. Breiman, *Probability*, Classics in Applied Mathematics, SIAM, Philadelphia, PA, 1992.
37. J. Portilla, V. Strela, M. Wainwright, and E. Simoncelli, "Image denoising using Gaussian scale mixtures in the wavelet domain," *IEEE Trans. Image Proc.* **12**(11), pp. 1338–1351, 2003. Computer Science techn. rep. nr. TR2002-831, Courant Institute of Mathematical Sciences.
38. A. Srivastava, A. B. Lee, E. P. Simoncelli, and S.-C. Zhu, "On advances in statistical modeling of natural images," *Journal of Mathematical Imaging and Vision* **18**(1), pp. 17–33, 2003.