

# Multi-resolution and wavelet representations for identifying signatures of disease

Paul Sajda<sup>a,\*</sup>, Andrew Laine<sup>a</sup> and Yehoshua Zeevi<sup>b</sup>

<sup>a</sup>*Department of Biomedical Engineering, Columbia University, New York NY, USA*

<sup>b</sup>*Department of Electrical Engineering, Technion – Israel Institute of Technology, Haifa, Israel*

**Abstract.** Identifying physiological and anatomical signatures of disease in signals and images is one of the fundamental challenges in biomedical engineering. The challenge is most apparent given that such signatures must be identified in spite of tremendous inter and intra-subject variability and noise. Crucial for uncovering these signatures has been the development of methods that exploit general statistical properties of natural signals. The signal processing and applied mathematics communities have developed, in recent years, signal representations which take advantage of Gabor-type and wavelet-type functions that localize signal energy in a joint time-frequency and/or space-frequency domain. These techniques can be expressed as multi-resolution transformations, of which perhaps the best known is the wavelet transform. In this paper we review wavelets, and other related multi-resolution transforms, within the context of identifying signatures for disease. These transforms construct a general representation of signals which can be used in detection, diagnosis and treatment monitoring. We present several examples where these transforms are applied to biomedical signal and imaging processing. These include computer-aided diagnosis in mammography, real-time mosaicking of ophthalmic slit-lamp imagery, characterization of heart disease via ultrasound, predicting epileptic seizures and signature analysis of the electroencephalogram, and reconstruction of positron emission tomography data.

## 1. Introduction

It has long been recognized that the localized structure of natural signals (and images) calls for the development of new theoretical and computational tools to enhance the processing and detection of specific signal attributes and features – i.e. signatures. This is no more true than in biomedical image and signal processing, where measurement noise and inter and in-patient variability obscure physiological and anatomical signatures important for the detection, diagnosis and monitoring of disease. Computer assisted detection and diagnosis systems have been, and are currently being developed to uncover and exploit such signatures. A major challenge has been to develop an “optimal” representation of the data so that signatures are easily uncovered and differentiated.

Classical representation and processing techniques, such as the widely used Fourier transform, consider the global, periodic properties of signals and images. In contrast, in medical imaging and/or electrophysiological measurement, one is looking for small, or short-lived conspicuous segments of the signal that may indicate diseased tissue or short episodes of dysfunction. Thus, even if the image of the normal tissue or the signal of the healthy organ may be considered to a good approximation to be stationary, (i.e. with statistical properties that do not vary over time or space) any localized abnormal structure or function induces, by its very nature, non-stationarity. Such signals and images call for the development and implementation of “localized” time (space)-frequency representations.

It is instructive to consider an example of the use of localized representations in the domain of Telecommunications. Communication engineers had already known in the early part of the 20th century that for speech, the periodic components of phoneme structure, known as formants, vary from phoneme to phoneme. Consequently, an efficient/optimal representation of

---

\*Corresponding author. Department of Biomedical Engineering, Columbia University, 351 Engineering Terrace, MC 8904, New York, NY 10027, USA. E-mail: ps629@columbia.edu.

speech calls for the application of a technique, or a transform, that incorporates both time and frequency. From such considerations came the development of the spectrogram and the Gabor transform [1], developed by Dennis Gabor, who was awarded the Nobel Prize for his invention of holography. Unlike the original time domain signal that identifies effectively the time of production of a sound or, alternatively, the Fourier transform that indicates its spectral components, the Gabor transform is “localized” in both time (or space) and frequency. This representational framework revolutionized the field of Telecommunications and is at the core of much of today’s modern image and signal processing.

In this paper we provide a review of the use of localized representations, including wavelets, for identifying signatures of disease in biomedical signals and images. We begin by providing a short tutorial on the mathematics behind the wavelet and other multi-resolution transforms. We then consider several clinically significant examples and case studies where such representations have been successfully applied. We conclude by discussing the limitations of current approaches, technical challenges, and opportunities that lie ahead for using these techniques for improving the detection, diagnosis and monitoring of disease.

## 2. A tutorial on multi-resolution representations and wavelets

There are many comprehensive references on wavelets and multi-resolution methods for signal and image processing. In this article we present only a brief tutorial on the subject, with additional details and more advanced topics deferred to the appendices. Readers interested in a more thorough and in-depth treatment of the mathematics and signal processing are referred to Akansu and Haddad [2], Vetterli and Kovacevic [3] and Strang and Nguyen [4]. Additional topics and connections with biomedical applications can be found in the edited works of Akay [5], Aldroubi and Unser [6] and Petrosian and Meyer [7], as well as the review article by Unser and Aldroubi [8]. Mathematically oriented studies on both Gabor and Wavelet-type representations in combined spaces can be found in Zeevi and Coifman [9].

### 2.1. Basic properties of wavelets and localized frequency transforms

A fundamental problem in signal processing and pattern recognition is identifying a signal representation which is well-matched to some desired task. For example, we may wish to identify and remove noise sources, compress data with minimal perceptual loss, or identify a signature of a particular disease process. All of these processing tasks are not necessarily best accomplished using the traditional signal space (often the time or space domain). Instead alternative representations often provide more insight and flexibility for processing. Wavelets are one such representation.

As with the Fourier transform, the wavelet transform is a linear transformation. Any function can be represented as a linear combination of wavelet functions. Consider a 1-D time domain signal  $f(t)$ .<sup>1</sup> We can represent  $f(t)$  as a weighted sum of wavelet functions,  $\psi_{j,k}(t)$ ,

$$f(t) = \sum_{j \in Z} \sum_{k \in Z} c_{j,k} \psi_{j,k}(t), \quad (1)$$

where the summation is over non-negative integer values  $j$  and  $k$ , indicated by  $Z$ . The wavelet functions are constructed via scaling and translation of a template function, often called a “mother” wavelet function  $\psi(\cdot)$

$$\psi_{j,k}(t) = 2^{j/2} \psi(2^j t - k), \quad \text{for } j, k \in Z, \quad (2)$$

where  $j$  indicates the discrete scaling number and  $k$  the discrete translation number.

One of the important properties of the wavelet functions is that they localize the representation functions of the signal (e.g. signatures) in both time and frequency.<sup>2</sup> For wavelets to be localized in this joint representation requires that they be of a finite support, that is short-lived, and have an oscillatory structure that constrains their power spectrum to be localized – i.e. 1)  $\psi \in L^1$ , 2)  $\int \psi dt = 0$  and 3) the Fourier transform of  $\psi$  has its power concentrated in the angular frequency range of  $[\pi, 2\pi]$  (and  $[-2\pi, -\pi]$  by symmetry of the Fourier transform). Localization of a wavelet function can be depicted by an appropriate Heisenberg box that indicates within the time-frequency space (also known as the *phase space*, shown in Fig. 1), the effective spread

<sup>1</sup>We present this framework for 1-D temporal signals, nevertheless it generalizes to multi-dimensional signals and images.

<sup>2</sup>Note that, unlike the Gabor representation, the explicit variable in the case of wavelets is scale, with the frequency band being inversely proportional to the scale.

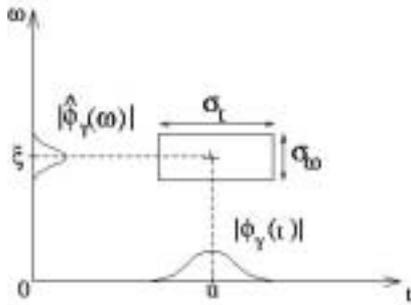


Fig. 1. Localization of a signal in the time-frequency plane is limited by the uncertainty theorem of Heisenberg. The localization of a signal is defined by a Heisenberg box, with dimensions  $\sigma_\omega \times \sigma_t$ .

of the function. There is a limit on the localization accuracy that a wavelet function, or any other linear function, can achieve. This is specified by the Heisenberg uncertainty principle (analogous to the limit on measurement precision of position and momentum that we know from physics) and is given as,

$$\sigma_\omega \sigma_t \geq 1/2 \tag{3}$$

where  $\sigma_\omega$  and  $\sigma_t$  define the widths of the Heisenberg box in frequency and time, respectively. Intuitively, the wavelet and other multi-resolution representations can be thought of as tiling the combined time-frequency plane with Heisenberg boxes. The Heisenberg boxes tile the plane with the effective support of the set of wavelet functions, obtained through a combination of translations and dilations of the “mother” wavelet. In Eq. (2), dilation is indexed by  $j$ , and represents the shrinking or stretching of the wavelet function. In particular, dilation modifies the frequency concentration of the wavelet, shrinking or stretching it from  $[\pi, 2\pi]$  to  $[2^j\pi, 2^{j+1}\pi]$ . In Eq. (2), translation is indexed by  $k$ , and can be seen as localizing the signal in time, with  $t = 2^{-j}k$ .

Figure 2 shows 1-D examples of some of the more popular wavelet families. Note that the shape of the wavelet can vary dramatically, and it is desirable to match the shape of the wavelet to the structure of the signal. Using such wavelets we can completely tile the time-frequency plane. An example of such a wavelet tiling is shown in Fig. 3D. Also shown is the original time domain signal, a Fourier tiling and a windowed Fourier tiling. It is worth mentioning the windowed Fourier tiling, also called the Gabor tiling, since it is widely used in applications such as speech and texture analysis, as well as being the basis for the *spectrogram*. The Gaborian decomposition can be written as,

$$f(t) = \sum_{m \in \mathbb{Z}} \sum_{n \in \mathbb{Z}} a_{m,n} g_{m,n}(t), \tag{4}$$

where  $g_{m,n}(t) = e^{-im\omega_0 t} g(t_0 - nt_0)$ . This essentially represents the superposition of a complex exponential, modulated and translated by an envelope function. Gabor showed that by choosing  $g(x)$  to be a Gaussian minimizes the effective spread of the transform in the phase space and thus, in some sense, optimally localizes the signal energy. These functions have been subsequently termed Gaborian functions (see Appendix 6.2 for additional details about Gaborian functions).

The Gaborian functions tile the time-frequency plane using Heisenberg boxes of equal size, and thus do not constitute a multi-resolution transform. The multi-resolution property of wavelets is useful in that it enables one to “zoom in” for capturing signatures in natural signals. For many types of natural signals, high frequency components have short correlation times and low frequency components have long correlation times. Given that one would like a sparse transform for the signal (i.e. the signal to be localized to only a few non-zero coefficients) the wavelet transform is a near optimal representation.

## 2.2. Computing the wavelet transform

Given a wavelet family,  $\psi_{j,k}$ , and a signal,  $f(t)$ , we must now consider how to compute the wavelet transform – i.e. how to compute the coefficients. Assume that we have chosen the wavelet family  $\psi_{j,k}$  to be orthonormal. The coefficients,  $c_{j,k}$  are computed as a scalar product of the signal and wavelets,

$$c_{i,j} = \langle f, \psi_{j,k} \rangle = \int f(t) \psi_{j,k}(t) dt. \tag{5}$$

With a reasonably large signal (large  $t$ ) and significant number of wavelets, this integration can be computationally expensive. As with the Fourier transform, there exist fast techniques for computing the wavelet coefficients. These methods exploit the fact that the wavelets constitute a multi-resolution analysis. To illustrate this fast transform, we begin first by substituting Eq. (5) into Eq. (1) and consider constructing an approximation to  $f(t)$ , which we call  $\hat{f}_J(t)$

$$\hat{f}_J(t) = \sum_{j < J} \sum_{k \in \mathbb{Z}} \langle f, \psi_{j,k} \rangle \psi_{j,k}(t). \tag{6}$$

The latter is a partial sum up to  $J$  scales. The approximation improves as we increase  $J$ ,

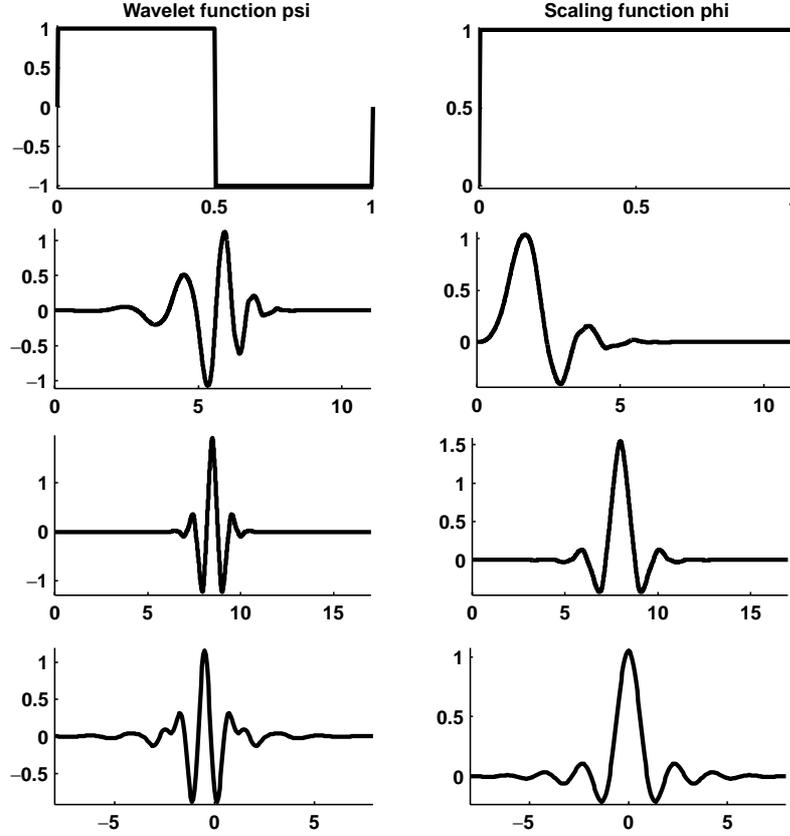


Fig. 2. Examples of “mother” wavelets and scaling functions. (first row) Haar (second row) Daubechie 6 (third row) Biorthogonal 2.8 (fourth row) Meyer 8.

$$\hat{f}_{J+1}(t) = \hat{f}_J(t) + \sum_{j < J} \sum_{k \in Z} \langle f, \psi_{j,k} \rangle \psi_{j,k}(t) \quad (7)$$

with the two right-hand side terms being orthogonal. This begins to hint at a multi-resolution or *pyramid* structure for computing the wavelet coefficients, which we describe in detail below.

In computing the wavelet coefficients, it is often convenient to begin with a *scaling function*,  $\phi$ , analogous to Eq. (2), from which the other lowpass scaling functions are derived,

$$\phi_{j,k}(t) = 2^{j/2} \phi(2^j t - k), \text{ for } j, k \in Z \quad (8)$$

Similarly we can also define a scaling coefficient,

$$s_{j,k} = \int f(t) \phi_{j,k}(t) dt = \langle f, \phi_{j,k} \rangle \quad (9)$$

Examples of scaling functions are shown in Fig. 2. The scaling and wavelet coefficients at coarse scale are computed from the coefficients of the scaling function at fine scale,

$$\begin{aligned} \langle f, \psi_{j,k} \rangle &= \sum_n g_{n-2k} \langle f, \phi_{j+1,n} \rangle \\ \langle f, \phi_{j,k} \rangle &= \sum_n h_{n-2k} \langle f, \phi_{j+1,n} \rangle, \end{aligned} \quad (10)$$

where  $g_{n-2k}$  and  $h_{n-2k}$  are a set of function (i.e. filters) that we need to define. Since we began our discussion by saying that we require the wavelets to be orthogonal, we will need to set  $g_{n-2k}$  and  $h_{n-2k}$  to be orthogonal. One natural choice is to choose them to be high-pass and low-pass filters, which when normalized properly, yield  $\sum_n h_n = 0$  and  $\sum_n g_n = 1$ . This provides an intuitive characterization of the scaling function and wavelets:

1. *The scaling function computes local averages for capturing the gross structure of the signal at some level J.*
2. *Wavelets compute local differences to capture the detailed structure at some level J.*

We would now like to represent  $f(t)$  exactly as an expansion of scaling functions and wavelets. Returning

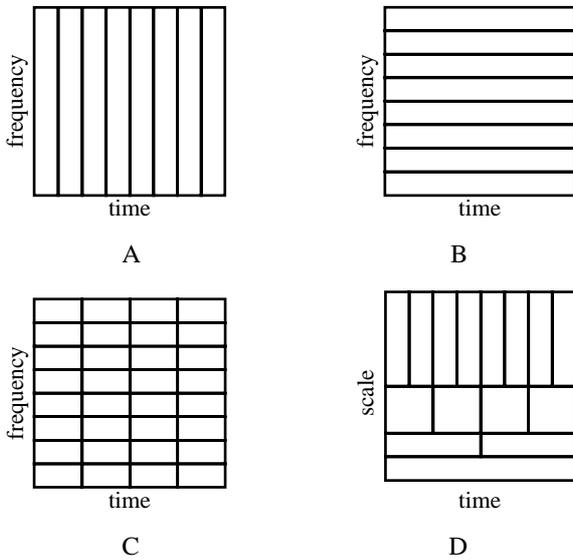


Fig. 3. Tilings of the time-frequency plane. (A) Time-domain sampled signal representation, (B) Fourier representation, (C) Short-time or windowed Fourier tiling (i.e. Gaborian), and (D) Wavelet tiling.

to our partial sums notation, we see that  $f(t) = \hat{f}(t) + \Delta$ , where  $\Delta$  is the difference between our function and its approximation partial sum. We will require that difference be captured by our scaling function at the coarsest resolution (defined as  $J = 0$ ),

$$\begin{aligned}
 f(t) &= \hat{f}(t) + \Delta \\
 &= \sum_{j=0}^{J-1} \sum_{k \in \mathbb{Z}} c_{j,k} \psi_{j,k}(t) \\
 &\quad + \sum_{k \in \mathbb{Z}} s_{0,k}(t) \phi_{0,k}(t).
 \end{aligned} \tag{11}$$

This is the *decomposition* or *analysis* equation for wavelets. The pyramid-based structure of the decomposition is depicted in Fig. 4A.

One of the nice properties of the wavelet transform is that since we require the transform to be orthonormal, computing the inverse (i.e. putting the signal back together) is simply done using the transpose,

$$s_{j+1,n} = \sum_k [h_{n-2l} s_{j,l} + g_{n-2l} c_{j,l}] \tag{12}$$

Equation 12 is called the *reconstruction* or *synthesis* equation. It also has a pyramid-based representation, depicted in Fig. 4(B).

### 2.3. Beyond simple wavelets

Much research has gone into further developing and extending the basic theory of wavelets. In this section

we briefly describe some of the extensions which are most relevant to biomedical image and signal processing.

#### 2.3.1. Biorthogonal wavelets and symmetry

On inspection of the orthogonal Harr and Daubachie wavelets shown in Fig. 2 it becomes apparent that the basic wavelets lack symmetry. However symmetry is a desirable property, particularly when dealing with images, where there is no preferred direction to the signal. We can generalize the wavelet transform and incorporate symmetry if we relax the orthonormality of the wavelet bases and instead require them to be *biorthogonal* to another set of synthesis functions. Biorthogonality can be understood in the context of linear algebra. For example, another way of writing the decomposition equation of the wavelet transform is as a linear transformation of a signal  $\mathbf{f}$  through a basis  $\mathbf{W}$  to obtain the wavelet coefficients  $\mathbf{c}$ ,

$$\mathbf{c} = \mathbf{W}\mathbf{f}. \tag{13}$$

The reconstruction equation is written similarly as,

$$\mathbf{f} = \mathbf{V}\mathbf{c}. \tag{14}$$

If we design  $\mathbf{W}$  to be an orthonormal basis then  $\mathbf{V} = \mathbf{W}^T$ , since the inverse of an orthogonal matrix is just its transpose. If  $\mathbf{W}$  is not an orthogonal basis we can still construct a basis which has many of the same properties of an orthogonal wavelet transform by requiring that the rows of  $\mathbf{W}$  be orthogonal to the columns of  $\mathbf{V}$ ,

$$\mathbf{W}_{i,\cdot} \cdot \mathbf{V}_{\cdot,j} = \delta_{i-j}. \tag{15}$$

In terms of our filter bank representation this implies that the decomposition filters ( $H, G$ ) and reconstruction filters ( $\overline{H}, \overline{G}$ ) are not identical.

#### 2.3.2. Translation invariance and over-complete bases

The wavelet bases we have discussed thus far have the problem that they are not invariant under translation – i.e. a shift in the signal will result in a different set of wavelet coefficients. One way to see why this is the case is by considering the downsampling operation, for example shown in Fig. 4. When we downsample by a factor of two, such as the case in the *dyadic* scheme, we “throw-away” all the odd-number indexed entries at the next scale. If there is a shift at a higher scale that forces coefficients into these indices, then the results after the subsampling will be very different from when there is not a shift. One obvious way to intro-

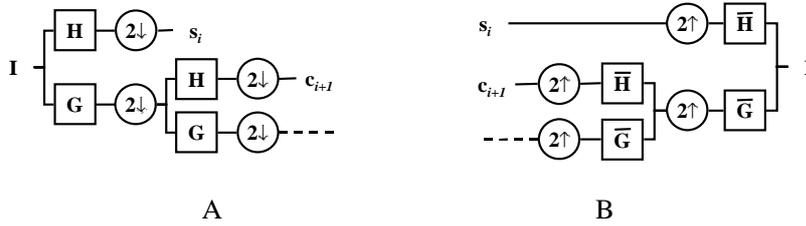


Fig. 4. Tree representation of the filter-bank pyramid framework for analysis and synthesis of a signal/image  $I$ . (A) The wavelet transform shown as a recursive filtering and sub-sampling down the low-pass branch of the tree. ( $H$  is high-pass analysis filter,  $G$  is low-pass analysis filter and  $\downarrow 2$  indicates down sampling by a factor of two, dashed-line indicates recursion down the branch of the tree). (B) Tree representation of wavelet synthesis. ( $\bar{H}$  is the high-pass synthesis filter,  $\bar{G}$  is low-pass synthesis filter and  $\uparrow 2$  is up-sampling by a factor of 2).

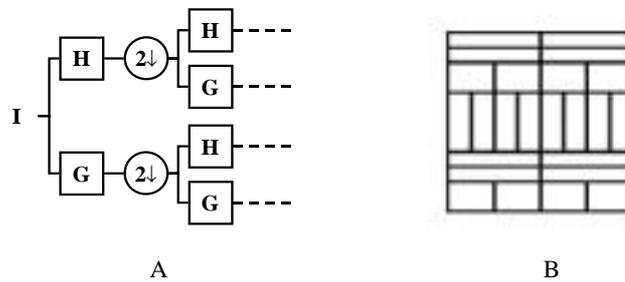


Fig. 5. (A) Wavelet packet tree, where analysis can go along both the high-pass and low-pass branches. (B) An example of a tiling of the time-frequency plane using a wavelet packet.

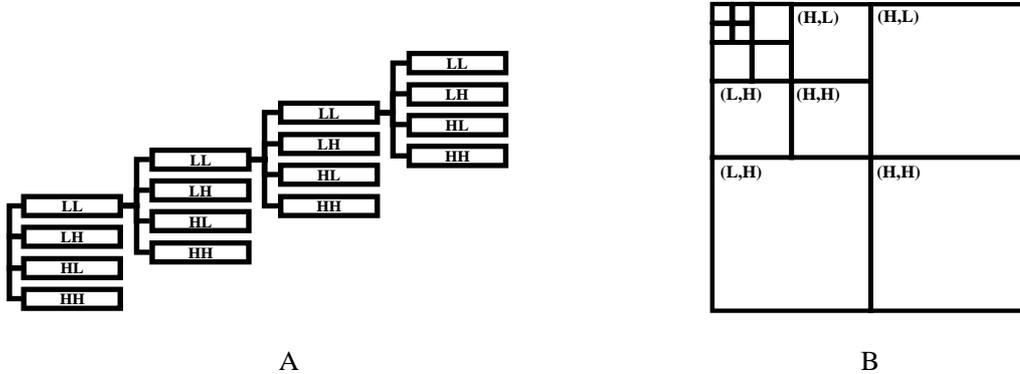


Fig. 6. Extension of wavelets to images (A) Tree structure for computing the 2D transform. (B) Conventional pyramid representation showing the image decomposed at multiple subbands.

duce translation invariance is not to subsample at every level. This comes at the cost of the transform being a redundant or having an *over-complete* basis – i.e. there are more coefficients in the transform than the number of elements in the original signal. With over-complete bases the process of reconstruction is more involved. However in many applications, for example in orientation analysis, one is willing to accept redundancy in order to obtain a representation of the signal which is best matched to the particular analysis. In addition, over-complete representations are often more robust.

2.3.3. Wavelet packets and best basis selection

Depending on the characteristics of the signal and the particular application, one might wish to have additional flexibility in the construction of the wavelet transform. We have represented the wavelet transform using a tree-based filter-bank construction, as shown in Fig. 4. In this case the wavelet decomposition always proceeds down the low-pass branch of the tree. In contrast, we can also construct a transform where we decompose down the high-pass branch as well, as shown in Fig. 5. This decomposition, called a *wavelet packet*,

results in a more flexible tiling of the time-frequency plane. Since one is often interested in a sparse representation, a wavelet packet decomposition can be useful. Coifman and Wickerhauser [10] have proposed the technique of *best basis selection* whereby the sparsest representation of the signal is chosen from a “library” of wavelet packet transforms.

### 3. Extensions to images

One way to extend the wavelet approach to multi-dimensional signals, in particular images, is by using separable wavelets. In this case we can construct a 2D wavelet representation using the same pyramid-based filter bank scheme used for 1D signals. We begin with high and low pass filters and apply them to either the horizontal or vertical dimensions of the image. For the 2D case, this means that each level in the pyramid results in four branches (*HH, HL, LH, LL*), since this is the number of combinations in which we can apply the two filters. Figure 6A shows the pyramid scheme extended to 2D. Figure 6B is the conventional method for depicting a wavelet transform for images.

A multi-resolution transform that has played a major role in computer vision and image processing, particularly real-time processing, is the Laplacian Pyramid, developed by Burt and Adelson [11]. The Laplacian Pyramid can be thought of as a predecessor to the wavelet transform. To understand the Laplacian pyramid it first makes sense to start with an even simpler transform, the Gaussian pyramid. The Gaussian pyramid, shown in Fig. 7, is constructed by iteratively low-pass filtering (i.e. convolving with a Gaussian kernel) and subsampling the image, resulting in a reduced resolution self-similar sequence of images. The Laplacian pyramid can be constructed from the Gaussian pyramid by computing the difference between a Gaussian pyramid image at level  $L$  and the coarser level  $L + 1$ , after it has been expanded (interpolated) to the size of level  $L$ . This construction process is shown in Fig. 7. The Laplacian pyramid therefore represents differences between consecutive resolution levels of the Gaussian pyramid, similar to how the wavelet coefficients represent differences.

There are two characteristics that differentiate the Laplacian pyramid from the wavelet transforms we have focused on thus far. One is that the Laplacian pyramid is overcomplete in that it has  $4/3$  as many coefficients as samples in the image (a critically complete wavelet transform has the same number of coeffi-

cients as samples in the image). This means that there is redundancy in the representation. In addition, the Laplacian pyramid localizes signals in space, but not (oriented) frequency. This is best seen by comparing the Laplacian pyramid to the wavelet transform in 2D frequency space, as shown in Fig. 8.

The utility of the multi-scale representation of the Gaussian and Laplacian pyramids becomes apparent if we consider the problem of pattern matching. Assume we have a signature that we are trying to match in the image. If the image is  $N \times N$  pixels, then we may have to search  $N^2$  points, which if  $N$  is large, can be computationally expensive. However, assuming that we can identify candidate regions by searching at low resolution, then the coarse resolution can be used to constrain the search at subsequently higher resolutions. This *coarse-to-fine* search strategy can be shown to dramatically reduce the computational cost of the search [12]. The coarse-to-fine framework can also be used for estimation, for example estimating velocities in image sequences [13]. The complementary *fine-to-coarse* processing can also prove useful, for instance as a method of propagating local information globally across the scene for surface interpolation [14]. The Laplacian pyramid has shown particular utility in image processing since it is selective to edges in images, which tend to be important features, and is invariant to illumination differences.

There are many extensions and generalizations of wavelets and the multi-resolution approach which are outside the scope of this paper. For example orientation analysis, via steerable filters [15] and steerable pyramids [16], has been used to compute signatures useful for pattern classification and image registration. These transforms tend to be over-complete, sacrificing redundancy for improved representation of features/structures of interest. Multi-scale representations, via construction of a scale-space, have also proven useful in image segmentation (see Appendix 6.3 for a short discussion of scale space).

### 4. Cases studies and example biomedical applications

In this section we present several examples and case studies that implement wavelets and multi-resolution processing techniques for detecting characteristic signatures of disease in medical images and biomedical signals. This is by no means a complete set of examples, given the broad utility of the multi-resolution approach. Our focus will be primarily the application of wavelets to medical image analysis.

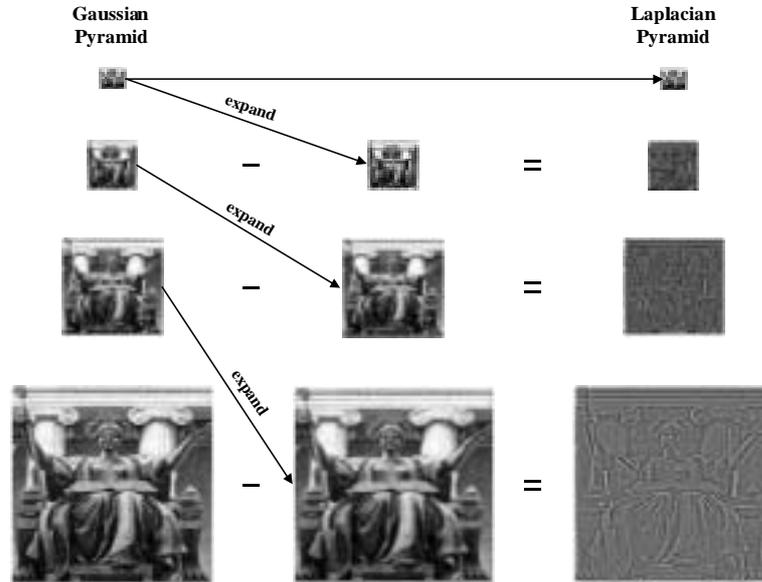


Fig. 7. Gaussian and Laplacian pyramids. The Laplacian pyramid can be constructed by computing the difference between the Gaussian pyramid image at level  $L$  and the interpolated or expanded Gaussian pyramid image at a coarser level  $L + 1$ .

#### 4.1. Cancer detection in x-ray mammography

One of the most clinically successful applications of wavelet and multi-resolution techniques has been Computer-Aided Diagnosis (CAD) systems for mammography. Mammographic CAD can be defined as a diagnosis made by a radiologist which incorporates the results of computer analyzes of the mammograms [17]. The goal of CAD is to improve radiologists' performance by indicating the sites of potential abnormalities, to reduce the number of missed lesions and/or by providing quantitative analysis of specific regions in an image to improve diagnosis. CAD systems typically operate as an automated "second-look" or "double-reading" systems that indicate lesion location and/or type by detecting signatures of disease. Since individual human observers overlook different findings, it has been shown that double reading i.e. (the review of a study by more than one observer) increases the detection rate of breast cancers by 5–15% [18–20]. Double reading, if not done efficiently, can significantly increase the cost of screening. Methods to provide improved detection with little increase in costs will have significant impact on the benefits of screening. Automated CAD systems are a promising approach for low-cost double-reading. Several CAD systems have been in development and the first have been approved by the FDA.

Mammographic CAD systems usually consist of two distinct subsystems, one designed to detect microcalci-

fications and one to directly detect masses [21]. Microcalcifications are indicators of cancerous processes or a pre-malignant stage in that they represent the metabolic bi-products of proliferative cellular activity. Features related to the shape, location, and clustering of microcalcifications are often used by radiologist to detect and diagnose cancer. Masses are the direct manifestation of a cancerous lesion, though they can also be due to a benign process, for example a fluid filled cyst. Features related to the shape of the mass (i.e. the shape of its border and the presence/absence of spiculation) are all cues to whether or not the mass is malignant. The challenge in CAD is to identify methods for optimally representing features related to microcalcifications and masses, and exploiting them so as to maximize detection and/or diagnostic performance.

Several groups have investigated using multi-resolution and wavelet approaches for feature extraction in mammographic mass and microcalcification detection. Brzakovic and Neskovic [22] use a fuzzy pyramid approach for detecting masses while Ng and Bischof [23] apply a template at several scales. One advantage of these approaches is that performance is independent of mass size, as further demonstrated by Miller and Ramsey [24]. Li et al. [25] use an oriented wavelet transform to construct features selective to spicules, thereby creating multi-scale signatures for mass detection. Similarly, Netsch and Peitgen [26] use a Laplacian pyramid to develop scale-space signatures for detecting individ-

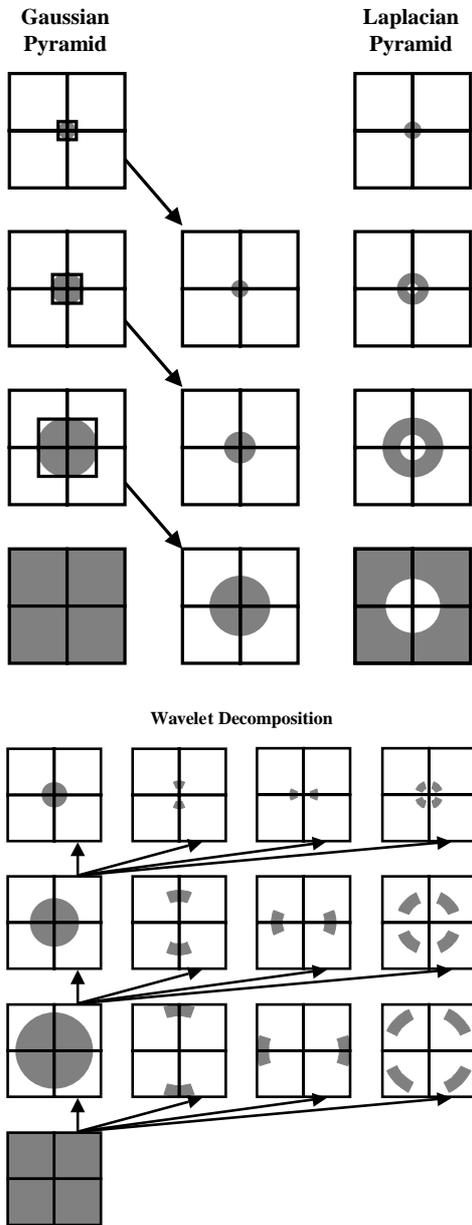


Fig. 8. (top) Laplacian pyramid in frequency domain. (bottom) Wavelet decomposition in the frequency domain. The different quadrants represent different oriented frequency bands. Note that the wavelet transform localizes structure in oriented frequency (as indicated by the wedges of the annuli) while the Laplacian transform does not.

ual microcalcifications. Though these authors show that the signatures are useful for detecting individual calcifications, the signatures do not fully exploit the coarse-scale information that may be indicative of contextual relationships, critical for determining clinical significance (e.g. calcification clustering).

One of the most promising applications of multi-resolution techniques for CAD has been through the integration of such representations with neural networks. In this case, the multi-resolution transformation constructs a good feature space in which the neural network can integrate information for detecting clinically significant structures. For example, te Brake and Karssemeijer [27] compare several multi-scale feature extraction methods against a single-scale method for detecting mammographic masses. They use a single neural network (5 hidden units, trained to minimize the root-mean-squared (RMS) error using backpropagation), to simultaneously integrate the multi-scale features. Aghdasi [28] uses a neural network to learn the optimal set of weights for integrating wavelet coefficients as features for microcalcification detection, an approach similar to that proposed by Yoshida et al. [29]. In all cases, integration is done using a single network.

Sajda et al. [30–32] have developed a multi-resolution neural network called the Hierarchical Pyramid Neural Network (HPNN). The HPNN typically is applied in one of two architectures, shown in Fig. 9, which are characterized as coarse-to-fine and fine-to-coarse. The architectures are designed to automatically learn and exploit contextual information through integration of signatures across scale and space. In mammographic image analysis context is exploited by radiologists and mammographers for detecting and identifying breast abnormalities. The clustering of calcifications, their proximity to ductal tissue, the architectural distortion surrounding potential lesions, are all contextual cues used by radiologists and mammographers [33]. The predictive value, as determined by radiologists, of both local and contextual (global) features for calcification and mass detection has been previously reported [34]. Contextual relationships can be integrated into mammographic CAD systems, being made explicit, given known pathology, through incorporation of preset rules and/or feature detectors tuned to capture the context. Alternatively, contextual relationships can be learned from the data, allowing for more complicated and less obvious contextual cues to be uncovered by the pattern recognition system.

The HPNN integrates contextual information through multi-scale decomposition of an image, via pyramid transforms [12,35] and the subsequent integration of multi-scale image features by a hierarchy of neural networks. These two architectures detect signatures by exploiting coarse-scale (low resolution) and fine-scale (high-resolution) information associated with the signature. For example, in the coarse-to-fine HPNN, net-

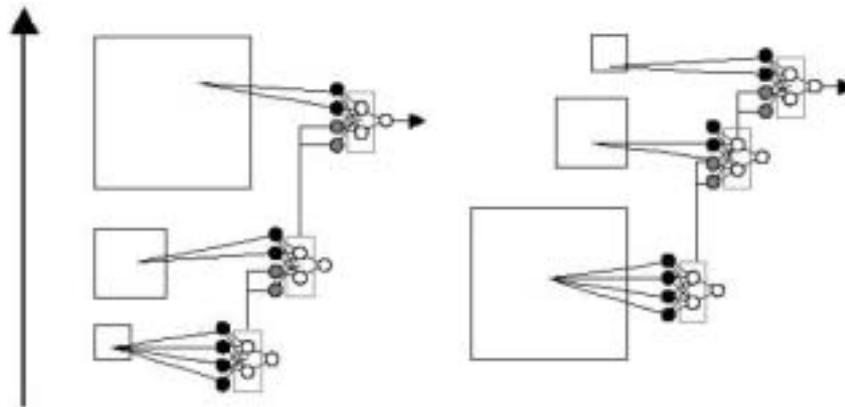


Fig. 9. Coarse-to-fine (left) and fine-to-coarse (right) HPNN architectures. Large arrow on left shows the direction of training and processing.

works operating at low resolution learn contextual features that are passed to networks operating at high resolution and integrated to detect the object of interest (i.e. the contextual inputs condition the probability of target present). For the fine-to-coarse HPNN architecture networks extract detail structure at fine resolutions of the image and then pass this detail information to networks operating at coarser scales (see Fig. 9 right). For many types of signatures, information about the fine detail structure is important for discrimination between different classes, i.e., fine resolution structure occurring within the context of the coarse resolution structure is indicative of an object class.

A key issue in CAD is to minimize false positives, generated by the computer, while maintaining high sensitivity. Results on several test datasets indicate that the HPNN is able to reduce the false positive rate of the computer by approximately 50%, for both microcalcification and mass detection, while still maintaining a 100% sensitivity on the test data. This significant reduction can be attributed to the integration of the multi-resolution features. Figure 10 shows the representation of hidden units from neural networks operating at coarse and fine scales. The hidden unit representations can be thought of as complex, scale dependent, features constructed via the multi-resolution input and network hierarchy. Figure 10B shows a hidden unit representation for a network operating at high resolution and Fig. 10C a network at low resolution. Note that the high resolution hidden units represent structures that appear to resemble microcalcifications. However, the low resolution units represent structure with fundamentally different characteristics. When correlated with the original mammogram, it is clear that the structure captured by the hidden unit at this scale is consistent with

anatomical features such as vascular and the ductal networks. It is well-known that breast cancers begin in the mammary ductal system and tend to be highly vascularized – i.e. those both can serve as contextual signatures for the presence of disease. The HPNN is not given explicit information of these contextual signatures. Instead, through its multi-resolution architecture, it is able to learn relationships and/or associations between structure at various scales, which maximize the probability of correctly detecting disease. Further evidence of the HPNN's ability to learn contextual cues is shown in Fig. 11. In this case, analysis was done on the false positives generated by the University of Chicago (UofC) CAD system which were eliminated by the HPNN. One characteristic of these false positives is that they contained significant linear structure, which when examined as a 1-D cross section, had the appearance of a "mountain range". The peaks of this signature were detected as microcalcifications by the UofC CAD system. The HPNN, on the other hand, could presumably integrate these peaks with the coarse-scale linear structure, so as to determine the region of interest was a false positive. The combined UofC and HPNN system has been tested in a clinical reader study, with results showing significant improvement in mammographer sensitivity when using the system [32].

Another class of mammographic CAD is based on selective image enhancement. Such techniques have been widely used in the field of radiology, where the subjective quality of images is important for human interpretation of signatures/features indicative of disease (e.g. border roughness, homogeneity of enhancement, etc.). Contrast is an important factor in any subjective evaluation of image quality. Many algorithms for accomplishing contrast enhancement have been de-

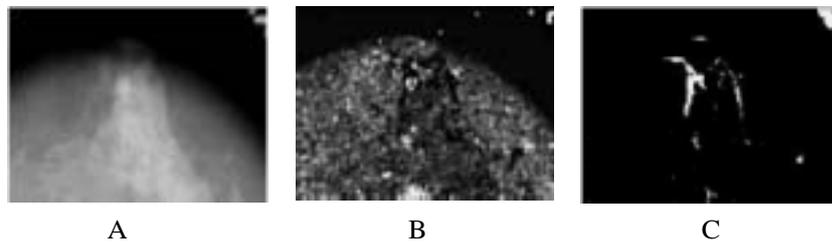


Fig. 10. (A) Original mammogram, (B) hidden unit representation for networks operating at high resolution. (C) hidden unit representations for networks operating at low-resolutions. Radiologists have noted that some of the structure in C appears to correlate with specific anatomy in the breast (ducts and/or blood vessels) indicating that these hidden units may represent contextual information.

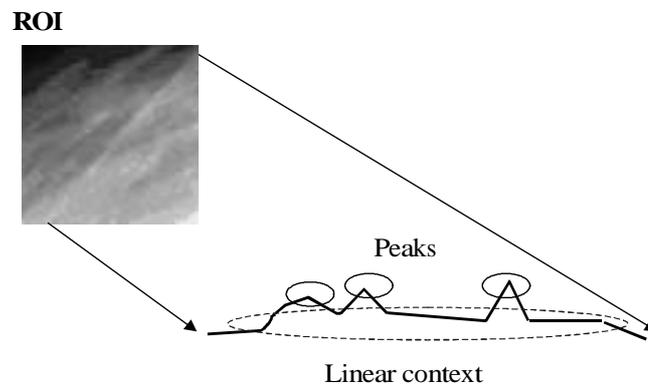


Fig. 11. Typical negative ROI that was eliminated by the coarse-to-fine HPNN. The HPNN learns to associate the fine-scale intensity peaks, which in isolation are interpreted as microcalcifications, with the coarse-scale linear structure, the result being that such false positives are eliminated.

veloped and applied to problems in medical imaging. A comprehensive survey of several methods has been published by Wang et al. [36]. Among the various techniques published, histogram modification and edge enhancement techniques have been most commonly used along with traditional methods of image processing. Histogram modification techniques [37,38] are attractive due to their simplicity and speed, and have achieved acceptable results for some applications. The transformation used is derived from a desired histogram and the histogram of an input image. In general, the transformation function is nonlinear. For continuous functions, a lossless transformation may be achieved. However, for digital images with a finite number of gray levels, such a transformation results in information loss, due to quantization. For example, a subtle edge may be merged with its neighboring pixels and disappear. Attempts to incorporate local context into the transformation process have achieved limited success. For example, simple adaptive histogram equalization [39] supported by fixed contextual regions cannot adapt to features of different size.

Most edge enhancement algorithms implicitly share a common strategy – detection followed by local “edge

sharpening”. The technique of unsharp masking is significant in that it has become a popular enhancement algorithm to assist radiologist in diagnosis [40, 41]. Unsharp masking sharpens edges by subtracting a portion of a filtered component from an original image. Theoretically, this technique was justified as an approximation of a deblurring process by Rosenfield and Kak [42]. Loo et al. [43] studied an extension of this technique in the context of digital radiographs. Another refinement based on Laplacian filtering was proposed by Neycenssac [44]. However, techniques of unsharp masking remain limited by their linear and single-scale properties, and are less effective for images containing a wide range of salient features as typically found in digital mammography. In an attempt to overcome these limitations, a local contrast measure and nonlinear transform functions were introduced in Gordon and Rangayyan, and later refined by Beghdadi and Negrata [45]. Unfortunately, limitations remained in these nonlinear methods as well: (1) no explicit noise suppression stage was included (in fact noise could be amplified), and (2) ad-hoc nonlinear transform functions were introduced without a rigorous mathemati-

cal analysis of their enhancement mechanisms or the possible introduction of artifacts.

Recent advancement of wavelet theory has sparked researchers' interest in the application of image contrast enhancement [46–52]. Wavelet-based feature enhancement traditionally proceeds along the following steps,

1. Construct the forward wavelet decomposition of the image.
2. Linearly or nonlinearly transform: thresholding, scaling or adaptively weight the wavelet coefficients.
3. Reconstructing the image using the modified wavelet coefficients.

Laine [43] achieves a scale-dependent enhancement of mammograms by selectively weighting and scaling the details computed using a first derivative of a Gaussian wavelet. In later work [44] a dyadic wavelet transform method was shown to be equivalent to unsharp masking at multiple scales. A continuous scale representation [53] provides more flexibility on multi-scale analysis over dyadic wavelets. Lu et al. [45] use a multi-scale edge representation to achieve the contrast enhancement. Other wavelets also have been used in related research, such as Complex Daubechies wavelets [46], and wavelet packets [47]. An example of using wavelets for global enhancement of a mammogram is shown in Fig. 12. Notice that the borders of the lesion are more clearly visible and thus are more easily assessed in terms of potential malignancy.

#### 4.2. Identification of retinal disease in slit-lamp fundus imagery

A general problem in medical image analysis is image registration. To detect clinically significant signatures in imagery, clinicians are often forced to fuse or register multiple images in their “minds-eye”. For example, in mammography, serial or temporal change detection (i.e. change between mammograms taken several months or years apart) is a common method for detecting breast cancer [33]. Radiologists will often look at both sets of mammograms and determine, through iterative examination of each set, whether there has been clinically significant change. This requires the radiologist to identify similar features in each mammogram in order to determine a common coordinate system for comparing the structure. A similar problem exists if a physician wants to fuse multiple data types, for example computed tomography (CT), positron emission

tomography (PET) and/or magnetic resonance imaging (MRI). In this case, the image types may be very different and therefore the features that are common to each may be sparse and difficult to identify. Wavelets and multi-resolution approaches have proven valuable techniques for image registration, in that they can exploit structure at multi-scales as well as a coarse-to-fine search strategy [12] for aligning images and volumes.

A promising, and relatively new, clinical application area for exploiting multi-resolution image registration methods is real-time ophthalmology [54,55]. As with other clinical disciplines, change detection and feature identification are fundamental strategies used by ophthalmologists for detection, diagnosis and monitoring of retinal disease. For example, detection and monitoring of drusen, small patches of abnormal metabolic build-up on the fundus, is critical for early detection and differential diagnosis of age-related macular degeneration (AMD) [56].

Imaging plays an important role in ophthalmology, with the workhorse being the slit-lamp biomicroscope. Slit-lamp biomicroscopes are used at the earliest stages of examination and are the imaging modality for ophthalmic screening. Slit-lamp images, examples of which are shown in Fig. 13A, have a narrow field of view (NFOV) because of the need to limit light intensity impinging on the retina. Ophthalmologists typically sweep the field of view across the fundus to construct a wide field of view (WFOV) image in their “minds-eye”. If warranted, a fundus photograph can be taken, which will provide a WFOV. Such photographs can be used to examine the features of the fundus in a broad context, as well as be compared to other types of imagery, for example fluorescein angiograms. However a disadvantage of a fundus photograph is that it requires a special camera and photographer and it often must be taken in another location. It is therefore rarely used at the screening stage.

One approach for providing a WFOV image at the screening stage is to mosaic slit-lamp images collected via a video CCD camera. The challenge is to accurately register the individual video frames and seamlessly blend the individual images. A particularly successful approach has been to use a multi-resolution coarse-to-fine search strategy for estimating the registration parameters [57]. Figure 14 shows the procedure for coarse-to-fine parameter estimation. Frames are first aligned at low resolution, and rough estimates of the registration parameters are computed. These estimates are used to seed the estimation process at the next higher resolution, where the alignment is further

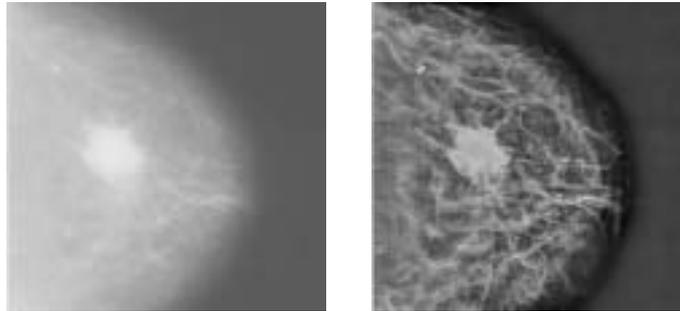


Fig. 12. (left) Original mammogram, (right) wavelet enhanced.

improved. This iterative method proceeds until the highest resolution where the images are registered, often with sub-pixel accuracy [12]. After registration, a Laplacian pyramid can be used to blend the individual frames to eliminate edge artifacts. An example of a resultant image is shown in Fig. 13B. The image provides the same basic type of information as might be provided by a fundus photograph, however without the additional expense and time required for specialized equipment and personnel.

#### 4.3. Analysis of EEG and MEG

Many research efforts have focused on the application of wavelets for detecting signatures of clinical significance in time-domain biomedical signals. For example, wavelets have been used as a means to detect heart rate variability/fluctuation [58–63]. A time-domain biomedical signal that is well-suited to wavelet analysis is the EEG, given that it is highly non-stationary. For example, great advances have been made in epileptic seizure prediction through multi-scale and wavelet-type analysis. Predicting the onset of epileptic seizures can have tremendous clinical utility, given evidence that temporally and spatially precise electrical stimulation preceding a seizure can reduce its severity and/or stop its occurrence [64]. Important has been the observation that discrete signatures, at multiple time scales, precede a seizure event [65]. These signatures can be seen in surface EEG, but they are most clear in the electrocorticogram (ECoG). In the ECoG, for example, one can see chirp-like activity, localized in frequency and time.

An important property of the epileptic signature is that the signal is non-stationary. It is therefore advantageous to track the signal in time-frequency space and watch for the onset of an epileptic signature. These characteristic properties of the seizure precursor have led to several wavelet based methods for signature detection.

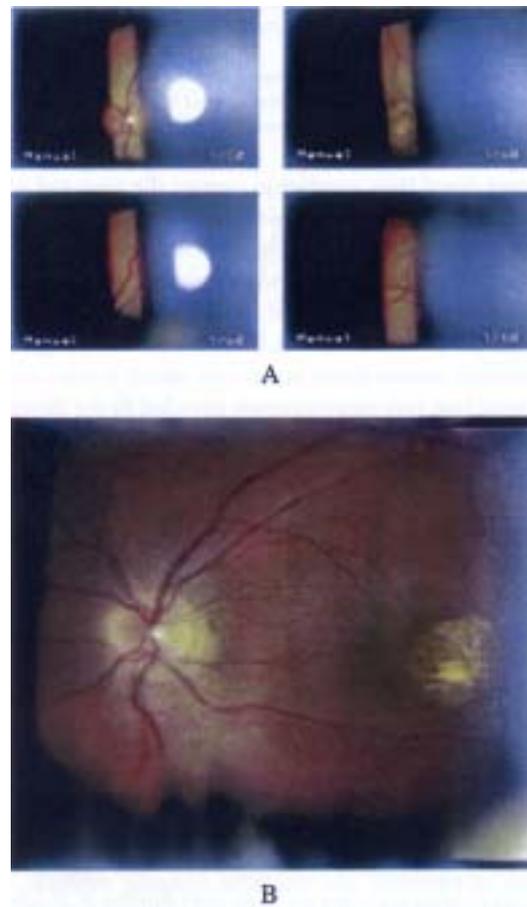


Fig. 13. Real-time mosaicking of slit-lamp biomicroscope images using a multi-resolution framework. (A) Several frames taken from a video sequence. (B) Mosaicked image constructed using a multi-resolution decomposition to align and blend the slit-lamp frames.

For example, Franaszczuk et al. [66] use the matched pursuit methods developed by Mallat and Zhang [67] to track the changing time-frequency signature seen in an epileptic seizure. Benedetto and Pfander [68] use a

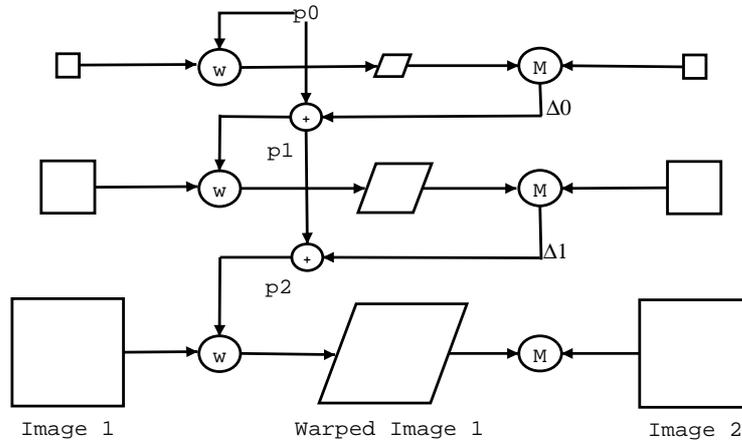


Fig. 14. Coarse-to-fine estimation of registration parameters using a multi-resolution framework. Sets of images are aligned beginning at lowest resolution. This rough alignment is used to subsequently initialize the estimation process at higher resolutions. The framework allows for extremely accurate sub-pixel resolution registration, with applications including change detection, image fusion and mosaicking.

generalized Haar wavelet to capture the time and scale periodicity seen in the wavelet transform – a sharp transition in the signal characteristics indeed calls for the application of Harr wavelets rather than smooth wavelets. Cranstoun et al. [69] use smooth localized exponential functions (SLEX) to detect signatures of epileptic seizure hours prior to the actual seizure event. These fast, real-time methods have led to the development of “brain pacemakers” [70], implantable devices aimed at predicting, in real-time and on-line, seizure onset in order to prevent a seizure through delivery of appropriate electrical stimulation to the cortex.

Wavelet techniques also have been used in the analysis of evoked potentials in the EEG. Advances in imaging techniques and processing of the acquired data have had a particularly important impact on analysis of brain function and dysfunction. The application of modalities such as PET and fMRI provide outstanding means to observe the brain “at work”, in a completely non-invasive manner, with reasonably good spatial resolution. The temporal resolution of these technologies is, however, less satisfactory and does not come close to that of electrical and magnetic imaging methods, i.e. EEG and MEG, which can potentially reveal the dynamics of sensory and/or higher brain functions.

Consider, as an example, the application of wavelet-type sparse representations for extracting single-trial evoked responses from MEG multichannel data [71]. One can formulate the extraction of the evoked response as a problem of Blind Source Separation (BSS) [72], where an  $n$ -channel observed signal,  $\mathbf{x}(t)$  (in this example case 122-channel MEG data) reflects the mixtures

of a source  $s(t)$ , embedded in an  $n$ -channel background signal  $\xi(t)$ ,

$$\mathbf{x}(t) = \mathbf{a} \cdot s(t) + \xi(t), \quad (16)$$

where  $\mathbf{a}$  is an unknown  $n$ -dimensional vector of weights. It is reasonable to assume that a rough template of the event-related evoked response is known *a priori*. Such a template can be obtained by averaging techniques.

BSS-type problems often have been approached using Independent Component Analysis (ICA), where the underlying assumption is linear independence [73, 74]. However, as Zibulevsky et al. have recently stressed [75,71] an alternative assumption, in the case of natural signals, is sparsity. Whereas some natural sources are sparse in their native signal domain, most natural signals and images must be projected onto a proper space to construct a sparse representation. Such sparse BSS techniques often yield better results than those obtained by traditional ICA approaches. As mentioned earlier, a sparse representation of a given signal  $s(t)$  refers to the fact that only a small number of the decomposition coefficients,  $c_k$ , corresponding to the set of representation functions  $\psi_k$  differs significantly from zero. The source signal can be therefore approximated by a linear combination of this subset of  $K$  functions,

$$s(t) = \sum_{k=1}^K c_k \psi_k(t). \quad (17)$$

Indeed, it is a well-known fact that wavelet-type transformations, such as those obtained by wavelet

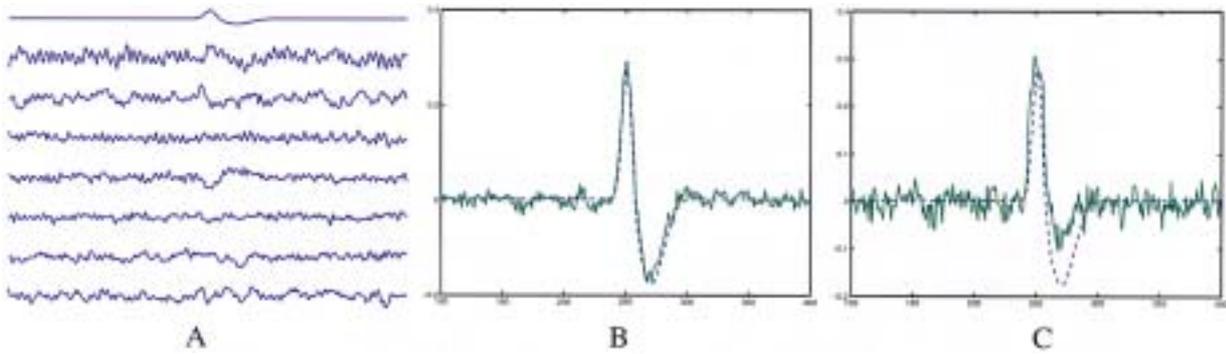


Fig. 15. (A) Subset of MEG channels with embedded evoked response. (B) Recovery of evoked response source using wavelet-based sparsity. (C) Recovery of evoked response source using correlation with a template. Actual embedded evoked response is shown as dashed curve.

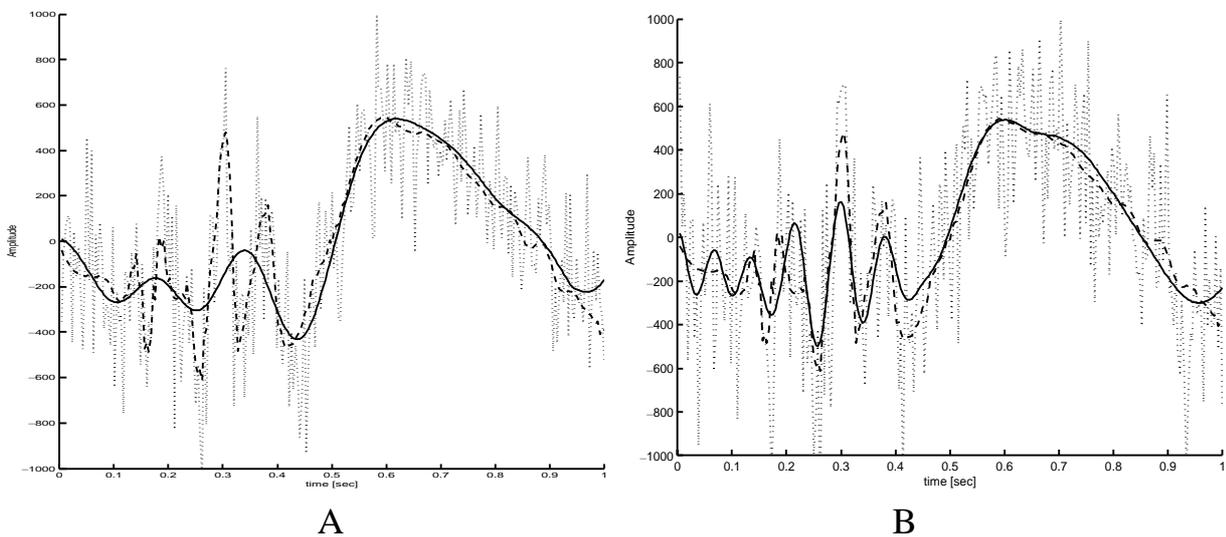


Fig. 16. Estimation with the KL (A) and the local KL (LKL) (B) transforms applied to simulated (synthetic) EP signal. The noise-free, noisy (with white Gaussian noise; SNR = 0 dB) and estimated signals are denoted by dashed, dotted and continuous lines, respectively.

packets [76,72], or Gabor-type transforms [77], when applied to natural signals, result in sparse representations. Quantitatively, the sparsity of the representation depends on the characteristics of the transformation and on its matching to the properties of the signal.

In [71] the authors combine prior knowledge about the sparsity of a source representation with the information regarding the relevant template, into one optimization criterion. Simulations with a synthetic evoked response, superimposed on 122-channel MEG data (Fig. 15), result in good recovery of the desired source (see Fig. 15B). These results are considerably better than those obtained using maximum correlation with a template (Fig. 15C). The template of the source simulated in this example is relatively smooth. This

indicates that a wavelet basis generated by a smooth mother-wavelet is most appropriate.

It should be noted that the number of samples used in this example is much smaller (512) than the number of free parameters associated with the separation of a  $122 \times 122$  matrix (i.e. 122 channels). Consequently, traditional ICA cannot provide a meaningful separation. Indeed, estimation using a traditional ICA technique does not lead to meaningful results.

In another study of processing of evoked potentials (EP), a local Karhunen-Loeve (KL) transform was used in the decomposition of an EEG signal space into complementary subspaces: one for the evoked potentials and the other for the background EEG interference-and-noise [78]. The local KL (LKL) basis was derived from the local autocorrelation of the EEG signals.

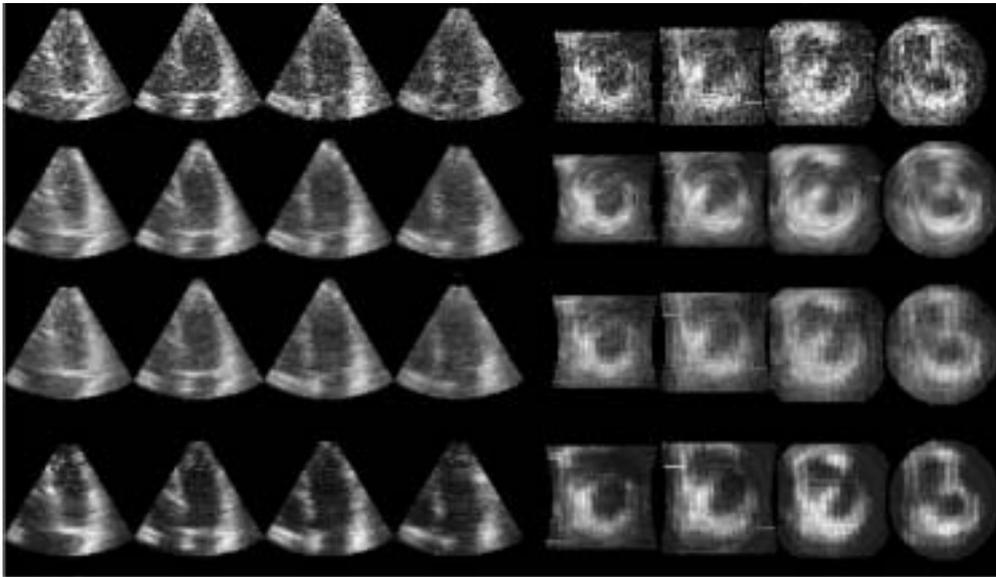


Fig. 17. Qualitative comparison of denoising and enhancement performance on real-time clinical echocardiography data. Denoising/enhancement results are presented for four long-axis and four short-axis views from a single clinical data set. (row 1) Original data, (row 2) denoising with Wiener filtering adapted to Rayleigh speckle noise model, (row 3) denoising with 3-D brushlet analysis, and (row 4) denoising with 4-D (3-D+Time) brushlet analysis.

The KL transform is known to provide the most efficient coordinate system for signal representation, according to the minimum mean-squared error (MMSE) and minimum entropy criteria. Nevertheless, features localized in the combined time-frequency space, as is the case with short-lived EPs, are not well represented by the KL, due to the global nature of its eigenfunctions. This is indeed what one is looking for in localized wavelet-type or Gabor-type functions. An alternative approach is to derive a LKL transform [78]. Examples of simulations with synthetic EP signals are shown in Fig. 16. The results indicate that the LKL can indeed recover the fine structure of a single-event EP, lost in the case of processing with global KL, without resorting to averaging.

#### 4.4. Analysis of the 4D (3D + time) Echocardiogram

A diversity of modalities, such as CT, tagged MRI, SPECT and ultrasound, allow for the acquisition of dynamic sequences of cardiac volumes. Echocardiography is the fastest, least expensive, and least invasive method for imaging the heart. The simplest and most useful clinical parameter used to assess cardiac function is ejection fraction (EF), calculated as the difference between end diastolic and end systolic left ventricular volumes. However, accurate calculations of ventricular volume from standard echocardiographic data

are tedious and costly to employ clinically. This is because existing methods require time to digitize endocardial borders on a series of two-dimensional images, then register the image set and reconstruct each cavity volume. Real-time acquisition via three-dimensional ultrasound obviates the need for slice registration and reconstruction, leaving segmentation as the remaining barrier to an automated, rapid, and therefore clinically applicable calculation of accurate left ventricular cavity volumes and ejection fraction. Because it provides such a rich description of the temporal and spatial environment of any area of interest, three-dimensional ultrasound also offers the potential for increased sensitivity in detecting subtle wall motion abnormality indicative of ischemia (for example during an exercise stress test), compared to fast MRI techniques.

The majority of the volume extraction methods which are used to estimate EF are based on prior models of the entire heart or of the left ventricle only. The parametrization of the model generally uses Finite Element models where the volume is constructed after deformation of the model following physics based constraints for equilibrium. Movement of the cardiac wall extracted from temporal data requires some parametrization of the model. Duncan et al. used contour shape descriptors in [79], Ayache, Cohen et al. [80] have used superquadratics. The nature of the constraints varies between models and can take a

wide range of properties, including differential constraints [81], displacement and velocity constraints [82] as well as other constraints allowing for non-rigid movements.

The 4D ultrasound image is dominated by speckle noise. This is problematic in that most of the model fitting procedures described above are very sensitive to noise in the data. Important therefore are methods for denoising the spatiotemporal ultrasound data so that one can obtain an accurate estimate of EF. Angelini et al. [83] describe a multi-resolution technique that relies on the expansion of temporal volume data on a family of multi-resolution basis functions called “brushlets”, introduced by Meyer and Coifman [84]. These basis functions offer a decomposition of a signal into distinct patterns of oriented texture. In 2D, depending on the tiling chosen prior to analysis, the projected coefficients are associated with distinct “brush strokes” of a particular size (width) and orientation (direction). Final denoising is achieved with the construction of gradient maps, thresholding of selected coefficients and reconstruction of an “enhanced/denoised” volume. The reconstructed data serves as an initial guess for volume extraction and EF estimation. Figure 17 illustrates the improved image quality after denoising with brushlets, with comparisons to traditional denoising methods such as Wiener filtering.

#### 4.5. Using the wavelet transform for maximum likelihood reconstruction of PET data

Positron Emission Tomography (PET) imaging quantifies the distribution of a radioactive marker that is injected into the body and interacts with pathological tissue. The radioactive tracers emit photons that are detected by pairs of detectors corresponding to bins. The projections of the tracer distribution are estimated by counting the number of photons detected in the various bins.

A classical technique for reconstruction of two- and three-dimensional PET images from a set of projections is based on the maximum likelihood approach [85,86]. Exploiting the properties of the Poisson process of photon emission, leads to the Expectation Maximization (EM) algorithm for emission tomography [86].

With the imaged tissue divided into pixels (i.e. image elements, in the 2D case) or voxels (i.e. volumetric elements in the 3D case), the number of photons generated by the radioactive tracer within each voxel is given by  $n(v)$ . Let  $y(b)$  be the number of photons detected in each bin  $b = 1, 2, \dots, B$ , where  $B$  is the total number

of bins. A Poisson process for the generation of photons in each voxel,  $v$ , is characterized by the expected value of the photons,  $\lambda(v)$ , that depends on the tracer distribution and, in turn, on the tissue properties. Further,  $y(b)$  are independent variables of a Poisson distribution with expected value  $\lambda(b)$ . Accordingly, the *log-likelihood* function of the measurements  $y(b)$  is given by

$$L(\lambda) = \sum_v \lambda(v)p(v) + \sum_b y(b) \log \sum_v \lambda(v)p(v, b), \quad (18)$$

where  $\lambda = \{\lambda(v), v = 1, \dots, V\}$ , with  $V$  being the total number of voxels.  $\lambda$  is the set of  $V$  unknown Poisson parameters,  $p(v, b)$  denotes the probability of an event associated with a photon emitted from voxel  $v$  being detected in bin  $b$ , and  $p(v)$  is the probability that an emission from  $v$  is being detected and is given by,

$$p(v) = \sum_b p(v, b) \quad (19)$$

The EM algorithm leads to an iterative approximation of  $\lambda(v)$  that maximizes  $L(\lambda)$ .

In cases where there is some prior information regarding possible pathology, utilization of such information lends itself to the development of new, more powerful, algorithms. Such localized structural information on smoothness and/or edge distribution can be obtained, for example, from another modality or from prior knowledge about specific organ metabolism (i.e. prior knowledge regarding the photon emission means  $\lambda$ ). This leads to a new iterative algorithm that incorporate the local and multi-resolution properties of the wavelet transform into the structure of the EM algorithm.

Let  $\{h_m\}, m = 1, \dots, M$ , be an orthonormal wavelet basis of an  $M$ -dimensional space, and let the discrete image  $\lambda(v)$ , belonging to this space, be represented by this wavelet basis as follows,

$$\lambda(v) = \sum_{m=1}^M c_m \psi_m(v). \quad (20)$$

Estimating the values of  $\lambda$ 's from  $y$ 's is equivalent to estimating  $c_m$ 's from  $y$ 's. One can therefore transform the log-likelihood function into the wavelet transform domain and obtain an expression similar to the one in the original domain, but with representation coefficients  $c_m$ 's instead of the  $\lambda$ 's. Since the transformed function is analogous and similar in structure to the

original one, the optimal  $c_m$  coefficients can be obtained via a wavelet domain iterative formula, similar to the one used in the image space. Since both the log-likelihood functions and the iterative algorithm are similar in both domains, what are the advantages of operating in the wavelet transform domain? Here we reiterate the fundamental empirical fact that for natural images there exists a wavelet-type transform that projects the image onto a sparse representation [72]. This approach can be further refined in the context of PET imaging, by designing sets of wavelet functions that are better matched to the structure of the PET signal.

PET data are very noisy due to their short acquisition time and scatter effects. In solving the optimization problem associated with the log-likelihood function, “penalty” functions reflecting, for example, smoothness of the noise-free image or, as mentioned, other prior information are incorporated into the optimization process. Such *regularization* is common practice in computer vision and image processing. In the context of the PET reconstruction problem, the log-likelihood function thus becomes,

$$L_\pi(\lambda) = L(\lambda) + \mu\Pi(\lambda), \quad (21)$$

where  $\Pi(\lambda)$  is a so-called penalty function, and  $\mu$  is its weighting coefficient that is determined empirically to achieve a desired signal-to-noise ratio. Under the assumption that we wish to construct a sparse representation, the sum of the absolute values of the coefficients provides a reasonably good choice for a penalty function. Since edges “live” in the high resolution subspace of the wavelet representation, the penalty function can be more effective if only the high resolution subset of the coefficients is incorporated [87]. This type of refinement can be carried one step further by considering a more PET-image-specific wavelet packet representation, where it is possible to penalize only a subset (or subsets) of the high resolution coefficients [72].

Kisilev et al. [87] implemented their wavelet transform-based regularization algorithm using the Shepp-Logan phantom [86]. Experiments such as those depicted in Fig. 18 illustrate that the wavelet transform based regularization can improve the signal-to-noise ratio and contrast of the reconstruction. The same approach can be incorporated into other modalities, given knowledge of the specific priors.

Kalifa et al. [88] have developed a new family of regularization methods for PET reconstruction, based on a thresholding procedure using wavelet and wavelet packet decompositions. This approach is based on the fact that the decompositions provide a near-

diagonalization of the inverse Radon transform and of prior information in medical images. A wavelet packet decomposition is adaptively chosen for the specific image to be restored. Corresponding algorithms have been developed for both 2-D and full 3-D reconstruction. These procedures are fast, non-iterative, flexible and outperform, for example, Filtered Back-Projection (see Fig. 19).

## 5. Discussion and conclusions

As discussed in this paper, progress has been made applying wavelet and multi-resolution techniques, originally developed for other signal and image processing domains, to specific problems in biomedical signal and image analysis. However, by many standards, the development and adoption of techniques by the medical community has been slow. This is in part due to the complexities and clinical requirements for dealing with medical data. As mentioned, one faces a tremendous intra and inter-subject variability in data related to physiological measurements and imaging, which can be a challenge for application of the basic wavelet techniques. Expertise in both the clinical domain and signal processing is therefore required to tune and refine the techniques to match the properties of the biomedical signals of interest. Further, due to the physics of acquisition systems and medical standards, one must be very cautious with preprocessing, filtering and approximating medical data. As an example of cooperation that is starting to be seen between the clinical and signal processing communities, we have discussed in this review Computer-Aided Diagnostic (CAD) approaches, where the system provides a ‘second look’. Indeed, such CAD systems have already been introduced into clinical service.<sup>3</sup>

Of the various signal and image processing techniques that have emerged in recent years, we have attempted to convince readers, who are concerned with the outstanding problem of identifying markers of disease, that wavelet-type and Gabor-type multi-resolution analysis techniques have tremendous potential. The examples presented in this review are by no means exhaustive, nor do they claim to be the best. We have selected several examples that represent a wide spectrum of applications with which we are most fa-

<sup>3</sup>R2 Technology M1000 Image Checker, CADx Medical Second-Look, and Intelligent Systems Software MammoReader.

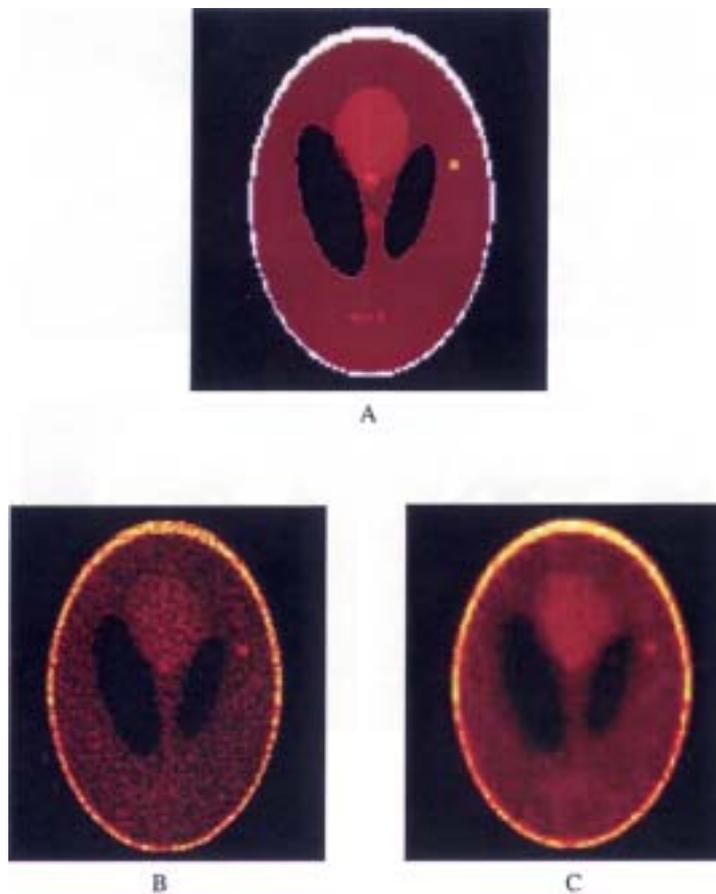


Fig. 18. Comparison of EM and wavelet transform-regularized-EM reconstruction for a Shepp-Logan phantom data. The discrete data of the phantom (A) was Radon-transformed and subjected to a Poisson noise process. The resultant data was then reconstructed using either a standard EM (B) or a wavelet transform regularized-EM procedure (C).

miliar. These and other published examples demonstrate that the proposed techniques are at least as suitable for the analysis of medical data, as they have been successful in other applications.

This short review cannot do justice to the broad range of examples previously reported in the literature, nor to the broad range of derivatives of these generic approaches for representing signals and images. The extension of the one-dimensional approach to two dimensions and higher, as is required in dealing with images and volumetric data, has been dealt with in most cases by extending the one-dimensional case to higher dimensions – in the simplest way by taking a tensor product of one-dimensional wavelet functions. The resulting scheme is characterized in this case by the splitting of the low-pass image into oriented subbands (as was illustrated in Fig. 6). It is often assumed that this is the only approach for constructing higher dimensional representations. However, as we have men-

tioned, two-dimensional non-separable wavelets are available, where many more degrees of freedom exist. This subject is, however, more technically involved and we have refrained from delving into it in depth, other than providing references and a few comments in the Appendices. Fine tuning of a scheme for a specific medical application, including markers of disease, justifies the extra effort of moving into the much richer space of non-separable wavelets.

One research area that has received significant attention in recent years is the use of wavelets as a representation or “feature space” for building generative probabilistic models of natural signals. The advantage of a generative probabilistic approach is that it offers a general framework for image analysis (e.g. classification, segmentation, fusion, compression, etc.) and a principled approach for including prior information and estimating “uncertainty”. For example, Crouse et. al. developed the Hidden Markov Tree (*HMT*) model [89]

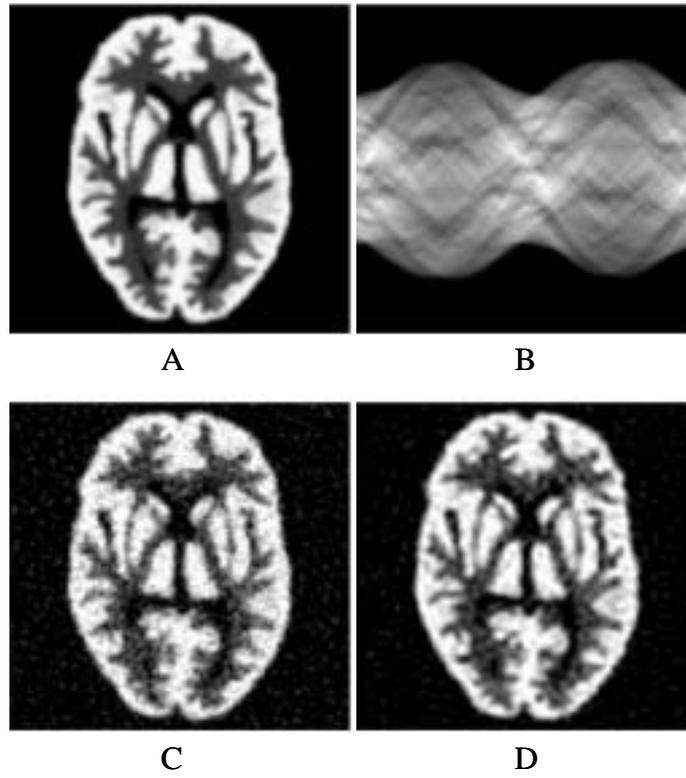


Fig. 19. (A) A reference image of a 256-by-256 Brain Phantom. (B) Sinograms simulated with 192 angular positions and 192 detectors, as in standard PET devices. (C) Image reconstructed with an optimized regularized filtered backprojection (RFBP) (signal to noise ratio (SNR) after reconstruction = 27 db). (D) Image reconstructed with a wavelet packet thresholding (SNR = 29.1 db).

for signals and images. A primary motivation of this model is to capture the tendency for wavelet coefficients to group into two classes, one with large and the other with small coefficient magnitudes. These models have been successfully applied to several problems, especially image enhancement and texture segmentation [90,91]. Cheng and Bouman [92] have applied a similar model for segmentation. Sajda et al. [93,94] have developed an extension of the basic HMT, which they term a hierarchical image probability (*HIP*) model. The *HIP* model allows for the construction of a generative probabilistic model over a wavelet packet representation, and has a more flexible architecture for capturing complex and structured dependencies between the transform coefficients. The *HIP* model has been evaluated for mammographic CAD as well as medical image compression, and has been shown to yield superior results compared to alternative state-of-the-art approaches (e.g. neural networks for CAD and JPEG and JPEG-2000 for image compression).

As often is the case in engineering solutions, the “devil is in the details” of implementation and design.

For example, applying wavelets to the problem of blind source separation of signals, one can achieve good results using the basic wavelets we have described. However the efficiency of the approach depends on the proper choice of wavelets. Indeed, the signatures of disease must be taken into account as well as the physical characteristics of sensing modality, in order to best match the multi-scale representations. In ultrasound, for example, the structure of the texture, as well structures within a lump (in the case of a cyst) are indicative of possible malignant tissue, whereas in PET, radioactive markers must be detected in the smallest possible quantities (minimal dose) by using optimal detection and estimation techniques. As discussed in this review, and as has been true in other applications, the performance of wavelet-based systems can be further improved by introducing tissue-specific and modality-specific prior information into the analysis and synthesis algorithms.

In conclusion, we submit that the potential remains for the development of more sophisticated analysis methods. For example, fruitful would be the devel-

opment of adaptive non-separable wavelet representations, which integrate more detailed and physically meaningful priori information, including specifics of tissue properties, function and modality. This task will challenge the next generation of medical imaging scientists, requiring a multi-disciplinary effort. We would expect the result to be more powerful diagnostic tools enabling minimally invasive identification of signatures of disease.

## 6. Appendices

### 6.1. Wavelet families

The generation of a family of orthogonal functions by dilations and translations traces back to the turn of the 20th century, when the Hungarian Mathematician Alfred Haar proposed a theory for a system of orthogonal functions [95]. The set of orthogonal Haar functions are identical in structure to the set of dyadic wavelets, but they are discontinuous – i.e. the mother wavelet (and so the entire set of wavelets) is constructed of unitary positive and negative pulses. The Haar set of functions is therefore ill-suited for the representation and analysis of medical signals and images and indeed for most natural signals. For many years it was not clear whether one could generate a sequence of smooth, scaled and translated orthogonal functions having finite support. This is, in fact, the important contribution of Ingrid Daubachies [96], who showed that one can generate such wavelets, where the degree of smoothness can be determined by the number of vanishing moments. Yet, it is often desirable to have additional properties that cannot be accomplished by one-dimensional orthogonal wavelets. For example, linear phase, a simple linear relationship between phase delay and frequency, cannot be achieved in the case of dyadic wavelets, except for the case of Haar wavelets that, as noted, are discontinuous. However, linear phase and other properties can be incorporated into a two-dimensional representation by using non-separable two-dimensional wavelets [97–99].

### 6.2. Gaborian transforms

Gabor [1] showed that the choice of a Gaussian for  $g(t)$  (Eq. (4)) minimizes the effective spread in the combined time (position)-frequency plane, compared with the so-called “joint entropy” achieved by any other window function. This result has been extended in

a straightforward manner to two-dimensional signals (e.g. images) that are projected into a four-dimensional combined position-frequency space [100]. Note, however, that even in the case of a one-dimensional signal there is a degree of freedom in choosing the width of the Gaussian window, since there is, according to the uncertainty principle, a trade-off between the spread in time (or space) and in frequency. In fact, the effective spread of one of these canonical variables is inversely proportional to the other. Further, there is also an additional degree of freedom in the tessellation of the combined space. In order to span the representation space, i.e. the space of signals or images to be represented, the set of elementary functions must be complete (or over-complete). This implies that at least a critical sampling density must be accomplished. The requirement of critical sampling density permits, however, a decrease in sampling rate along the time (position) coordinate as long as it is compensated by the corresponding increase in the sampling rate along the complementary frequency axis [77]. However, one must pay attention to the important issue of stability, which is beyond the scope of this paper. Note also that the Gaborian functions are not orthogonal. This necessitates the generation of a complementary, bi-orthogonal, set of functions that are required for obtaining the inverse.

The Gaborian approach to image representation in a combined spatial and spectral four-dimensional space is in fact very suitable for dealing with the complex textures of medical images, since the oriented Gaborian elementary functions are most suitable for the representation of textures [100,101]. Any localized pathology is then identified by deviation from the localized periodic structure of the texture. As such, it has been applied in the analysis of ultrasound images and in a computer vision approach to mass screening of melanoma lesions.

In the case of the original Gaborian approach to signal and image representation, the Gaussian window is fixed in its width, and all the localized functions are of the same (effective) finite support. It has been realized, however, that a large number of classes of natural images (as well as one-dimensional signals), have a peculiar property of self-similarity across scales. Mandelbrot [102] was the first to highlight the fact that this phenomenon is encountered everywhere in nature. This observation has motivated the development of the Laplacian [11] and Gaborian [100] pyramidal representations, and subsequently development of wavelet multi-resolution approach to signal and image analysis [103,104]. Alternatively, the scale can be incorpo-

rated into the Gaborian set of representation function by using a scaled set of multiwindows [77]. Each subset of the Gaborian functions having a fixed frequency depicts in this case the properties of wavelet functions, i.e. a set of localized scaled and translated functions. The specific set of scaled multiwindow Gaborian functions therefore constitutes a wavelet-type representation. In fact, since it incorporates the properties of both wavelets and Gabor representations, it provides a rich set of representation functions suitable for the analysis of medical and other data.

### 6.3. Diffusion processes in scale-space

Another concept closely related to multi-resolution processing and the continuous wavelet transform is the scale-space [105–107]. This concept is directly related to the application of diffusion-type partial differential equations (PDE) in image processing.

We have already mentioned the Laplacian pyramid of Burt and Adelson, and how, in fact, it generates scaled and translated (non-orthogonal) wavelets. The application of a linear diffusion process as an operator acting on an image  $I_0(x, y)$ , results in a dynamic evolution of a low-pass filtered image,  $I(x, y, t)$ , that is a spatio-temporal solution of the partial differential equation:

$$I_t = c\nabla^2 I. \quad (22)$$

$I_0$  is the initial condition of the process,  $c$  is the diffusion coefficient that determines the rate of the filtering process,  $I_t$  denotes the partial derivative of the image with respect to time, and  $\nabla^2$  denotes the Laplacian (i.e. the sum of the second partial derivatives with respect to the two spatial independent variables  $x$  and  $y$ ).

A slice in time of the resultant three-dimensional (i.e. two spatial and one temporal axes) data cube depicts a Gaussian filtered (smoothed) image. In fact, the result of the action of the diffusion process on the image can be obtained at any given time by the convolution of the image with the fundamental Gaussian solution,  $g(x, y, t)$ , of the normalized (i.e.  $c = 1$ ) diffusion equation,

$$g(x, y, t) = \frac{1}{\sqrt{4\pi t}} \exp \frac{-x^2}{4t}. \quad (23)$$

Clearly, the longer the processing time the smoother the image, since the Gaussian kernel of the diffusion equation is characterized by an effective width, i.e. standard deviation  $\sigma = \sqrt{2t}$ . This simple linear image evolution process of low-pass filtering generates a Gaussian scale-space.

While the action of the linear diffusion low-pass filters the noise, it also degrades the edges. Thus nonlinear adaptive diffusion must be applied in order to resolve the conflicting requirements of denoising, i.e. low-pass filtering, and edge enhancement. This can be done effectively by the application of a nonlinear adaptive diffusion equation where the diffusion coefficient is a function of the local gradient. By decreasing the value of the diffusion coefficient to zero at positions of large gradients, edges can be preserved while noise is filtered out over relatively smooth areas [108]. Edges may be enhanced by adapting the value of the diffusion coefficient at points of large gradients. Although such combined “forward-and-backward” diffusion is ill-posed, one can show that it is stable enough to be applied (with care) to images, and results in enhancement of edges [109]. Obviously, this is important in many cases where the task of demarcating a small lump that may depict a tumor cannot be accomplished by simple linear processing techniques because of low signal-to-noise ratio (SNR). The scale-space approach to image processing has been shown to be very effective in segmentation of images [107] and, in particular, medical images.

### Acknowledgments

The authors would like to acknowledge support from NSF CAREER Award BES-0133804 (P.S.), the DoD Multidisciplinary University Research Initiative (MURI) program administered by the Office of Naval Research under N00014-01-0625 (P.S and Y.Z), the Department of Biomedical Engineering Imaging Group of Columbia University, the Ollendorff Center of the Technion and the Israeli Ministry of Science (Y.Z).

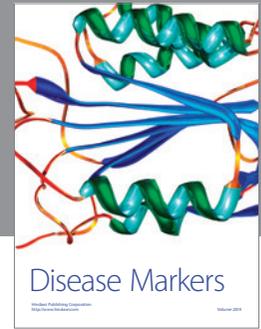
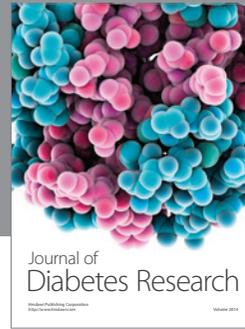
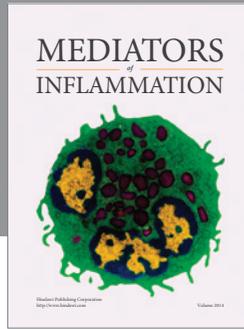
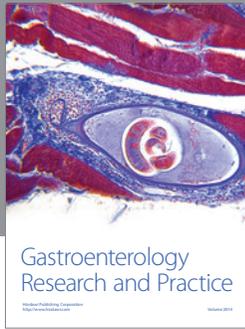
### References

- [1] D. Gabor, Theory of communication, *J.I.E.E.* **93** (1976), 429–459.
- [2] A.N. Akansu and R.A. Haddad, *Multiresolution Signal Decomposition*, Academic Press, New York, 1992.
- [3] M. Vetterli and J. Kovacevic, *Wavelets and Subband Coding*, Prentice Hall PTR, Englewood Cliffs, NJ, 1995.
- [4] G. Strang and T. Nguyen, *Wavelets and Filter Banks*, Wellesley-Cambridge, Boston, 1996.
- [5] M. Akay, *Time Frequency and Wavelets in Biomedical Signal Processing*, IEEE Press, New York, 1999.
- [6] A. Aldroubi and M. Unser, *Wavelets in Medicine and Biology*, CRC Press, Boca Raton, FL., 1996.
- [7] A.A. Petrosian and F.G. Meyer, *Wavelets in Signal and Image Analysis: From Theory to Practice*, Kluwer Academic Publishers, Boston, MA, 2001.

- [8] M. Unser and A. Aldroubi, A review of wavelets in biomedical applications, *Proceedings of the IEEE* **84** (1996), 626–638.
- [9] Y.Y. Zeevi and R. Coifman, *Signal and Image Representations in Combined Spaces*, Academic Press, San Diego, CA, 1997.
- [10] R. Coifman and M. Wickerhauser, Entropy-based algorithms for best basis selection, *IEEE Transactions on Information Theory* **38**(2) (1992), 713–718.
- [11] P.J. Burt and E.A. Adelson, The Laplacian pyramid as a compact image code, *IEEE Trans. Commun.* **31** (1983), 532–540.
- [12] P. Burt, Smart sensing within a pyramid vision machine, *Proceedings IEEE* **76**(8) (1988), 1006–1015. Also in *Neuro-Vision Systems*, Gupta and Knopf, eds, 1994.
- [13] P. Burt, J. Bergen, R. Hingorani, R. Kolczynski, W. Lee, A. Leung, J. Lubin and H. Shvaytser, Object tracking with a moving camera, an application of dynamic motion analysis, in *Proc. IEEE. Workshop on Visual Motion*, IEEE, Irvine, CA, 1988, pp. 2–13.
- [14] P. Burt, Moment images, polynomial fit filters, and the problem of surface interpolation, in *Proc. Computer Vision and Pattern Recognition*, IEEE, Ann Arbor, MI, 1988, pp. 144–152.
- [15] W.T. Freeman and E.H. Adelson, The design and use of steerable filters, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **13**(9) (1991), 891–906.
- [16] E.P. Simoncelli, W.T. Freeman, E.H. Adelson and D.J. Heeger, Shiftable multiscale transforms, *IEEE Trans. Inform. Theory* **38**(2) (1992), 587–607.
- [17] K. Doi, M.L. Giger, R.M. Nishikawa, K. Hoffmann, H. MacMahon, R.A. Schmidt and K.G. Chua, Digital radiography: A useful clinical tool for computer-aided diagnosis by quantitative analysis of radiographic images, *Acta Radiologica* **34** (1993), 426–439.
- [18] R.E. Bird, Professional quality assurance for mammography screening programs, *Radiology* **177** (1990), 8–10.
- [19] C.E. Metz and J.H. Shen, Gains in accuracy from replicated readings of diagnostic images: prediction and assessment in terms of ROC analysis, *Medical Decision Making* **12** (1992), 60–75.
- [20] E.L. Thurffjell, K.A. Lernevall and A.S. Taube, Benefit of independent double reading in a population-based mammography screening program, *Radiology* **191** (1994), 241–244.
- [21] M.L. Giger, Z. Huo, M.A. Kupinski and C.J. Vyborny, Computer-aided diagnosis in mammography, in: *Handbook of Medical Imaging: Volume 2. Medical Image Processing and Analysis*, M. Sonka and J.M. Fitzpatrick, eds, SPIE Press, 2000, pp. 917–986.
- [22] D. Brzakovic and M. Neskovic, Mammogram screening using multiresolution-based image segmentation, in: *State of the Art in Digital Mammographic Image Analysis (Series in Machine Perception and Artificial Intelligence)*, (Vol. 9), K.W. Bowyer and S.M. Astley, eds, New York: World Scientific, 1994, pp. 103–127.
- [23] S.L. Ng and W.F. Bischof, Automated detection and classification of breast tumors, *Computers and Biomedical Research* **25** (1992), 218–237.
- [24] L. Miller and N. Ramsey, The detection of malignant mass by nonlinear multiscale analysis, in: *Digital Mammography*, K. Doi, M.L. Giger, R.M. Nishikawa and R.A. Schmidt, eds, Amsterdam, The Netherlands: Elsevier, 1996, pp. 335–340.
- [25] L. Li, W. Qian, and L.P. Clarke, Digital mammography: Computer-assisted diagnosis method for mass detection with multiorientation and multiresolution wavelet transforms, *Academic Radiology* **11**(4) (1997), 724–731.
- [26] T. Netsch and H.O. Peitgen, Scale-space signatures for the detection of clustered microcalcifications in digital mammograms, *IEEE Transactions on Medical Imaging* **18**(9) (1999), 774–786.
- [27] G.M. te Brake and N. Karssemeijer, Single and multiscale detection of masses in digital mammograms, *IEEE Transactions on Medical Imaging* **18**(7) (1999), 628–639.
- [28] F. Aghdasi, Application of neural network adaptive wavelets for signal representation and classification in digital mammography, in: *Digital Mammography*, K. Doi, M.L. Giger, R.M. Nishikawa and R.A. Schmidt, eds, Amsterdam, The Netherlands: Elsevier, 1996, pp. 307–310.
- [29] H. Yoshida, W. Zhang, W. Cai, K. Doi, R.M. Nishikawa and M.L. Giger, Optimizing wavelet transform based on supervised learning for detection of microcalcifications in digital mammograms, in: *International Conference in Image Processing*, (Vol. 3), 1995, pp. 152–155.
- [30] P. Sajda, C.D. Spence, J.C. Pearson and R.M. Nishikawa, Exploiting context in mammograms: A hierarchical neural network for detecting microcalcifications, in: *Medical Imaging 1996 — Image Processing*, (Vol. 2710), M.H. Loew and K.M. Hanson, eds, P.O. Box 10, Bellingham WA 98227-0010, SPIE, 1996, pp. 733–742.
- [31] C.D. Spence and P. Sajda, Applications of multi-resolution neural networks to mammography, in: *Advances in Neural Information Processing Systems 11*, M.S. Kearns, S.A. Solla and D.A. Cohn, eds, Massachusetts Institute of Technology, MIT Press, Cambridge, MA 02142, 1999, pp. 938–944.
- [32] P. Sajda, C. Spence and J. Pearson, Learning contextual relationship in mammograms using a hierarchical pyramid neural network, *IEEE Transactions on Medical Imaging* **21** **3** (2002), 239–250.
- [33] D. Kopans, *Breast Imaging*, Lippincott, Philadelphia, PA, 1989.
- [34] P. Whatmough, A.G. Gale and A.R.M. Wilson, Do radiologists agree on the importance of mammographic features? in: *Digital Mammography*, K. Doi, M.L. Giger, R.M. Nishikawa and R.A. Schmidt, eds, Elsevier, Amsterdam, The Netherlands, 1999, pp. 111–116.
- [35] P. Sajda, C.D. Spence, S. Hsu and J.C. Pearson, Integrating neural networks with image pyramids to learn target context, *Neural Networks* **8**(7/8) (1995), 1143–1152.
- [36] D.C. Wang, A.H. Vagnucci and C.C. Li, Digital image enhancement, *Computer Vision, Graphics, and Image Processing* **24** (1983), 363–381.
- [37] R. Hummel, Histogram modification techniques, *Computer Graphics and Image Processing* **4** (1975), 209–224.
- [38] W. Frei, Image enhancement by histogram hyperbolization, *Computer Graphics and Image Processing* **6** (1977), 286–294.
- [39] S.M. Pizer, E.P. Amburn, J.D. Austin, R. Cromartie, A. Geselowitz, T. Greer, B.H. Romeny, J.B. Zimmerman and K. Zuiderveld, Adaptive histogram equalization and its variations, *Computer Vision, Graphics, and Image Processing* **39** (1987), 355–368.
- [40] J.W. Oestmann, R. Greene, J.R. Rubens, E. Pile-Spellman, D. Hall, C. Robertson, H.J. Llewellyn, K.A. McCarthy, M. Potsaid and G. White, High frequency edge enhancement in the detection of fine pulmonary lines: parity between storage phosphor digital images and conventional chest radiograph, *Invest. Radiol.* **24** (1989), 643–646.

- [41] M. Prokop, M. Galanski, J.W. Oestmann, U. Von Falkenhäusen, H. Rosenthal, P. Reimer, J. Nischelsky and J. Reichelt, Storage of phosphor versus screenfilm radiography: Effect of varying exposure parameters and unsharp mask filtering on the detectability of cortical bone defects, *Radiology* **177**(1) (1990), 109–113.
- [42] A. Rosenfeld and A.C. Kak, *Digital Picture Processing*, (2nd ed.), New York, NY, 1982.
- [43] L.D. Loo, K. Doi and C. E. Metz, Investigation of basic imaging properties in digital radiography: Effect of unsharp masking on the detectability of simple patterns, *Med. Phys.* **12** (1985), 209–214.
- [44] F. Neyceussac, Contrast enhancement using the Laplacian-of-a-Gaussian filter, *CVGIP: Graphical Models and Image Processing* **55** (1993), 447–463.
- [45] A. Beghdadi and A.L. Negrata, Contrast enhancement technique based on local detection of edges, *Comput. Vision Graphics Image Process.* **46** (1989), 162–174.
- [46] A.F. Laine, Multiscale wavelet representations for mammographic feature analysis, *Proceedings of the National Cancer Institute Breast Imaging Workshop: State-of-the-Art and New Technologies, Image Enhancement Techniques: Computer Science*, Bethesda, MD, 1991.
- [47] A.F. Laine and S. Song, Multiscale wavelet representations for mammographic feature analysis, *Mathematical Methods in Medical Imaging, Proceedings of SPIE*, (Vol. 1768), San Diego, CA, 1992, pp. 306–316.
- [48] A.F. Laine and S. Song, Wavelet processing techniques for digital mammography, in *Proceedings of SPIE: Conference on Visualization in Biomedical Computing*, (Vol. 1808), Chapel Hill, NC, 1992, pp. 610–624.
- [49] A. Laine, S. Song and J. Fan, Adaptive multiscale processing for contrast enhancement, *Proceedings of SPIE: Conference on Biomedical Imaging and Biomedical Visualization*, (Vol. 1905), San Jose, CA, 1993, pp. 521–532.
- [50] B.D. Jawerth, M.L. Hilton and T.L. Huntsberger, Local enhancement of compressed images, *J. Mathematical Imaging and Vision* **3** (1993), 39–49.
- [51] A.F. Laine, S. Schuler, J. Fan and W. Huda, Mammographic feature enhancement by multiscale analysis, *IEEE Transactions on Medical Imaging* **13** (1994), 725–740.
- [52] J. Lu and D.M.H. Jr., Contrast enhancement of medical images using multiscale edge representation, *Proceedings of SPIE: Wavelet applications*, (Vol. 2242), Orlando, FL, 1994, pp. 711–719.
- [53] R. Gordon and R.M. Rangayyan, Feature enhancement of film mammograms using fixed and adaptive neighborhoods, *Applied Optics* **23** (1984), 560–564.
- [54] A. Can, C.V. Stewart and B. Roysam, Robust hierarchical algorithm for constructing a mosaic from images of the curved human retina, in *Proceeding of IEEE Computer Society: CVPR*, 1999, pp. 286–293.
- [55] J.W. Berger and D.S. Shin, Computer-vision enabled augmented reality fundus biomicroscopy, *Ophthalmology* **106** (1999), 1935–1941.
- [56] S.L. Fine, M.G. Maguire and J.W. Berger, *Age-related Macular Degeneration*, Mosby-Year Book, New York, 1999.
- [57] J. Asmuth, B. Madjarov, P. Sajda and J.W. Berger, Mosaicking and enhancement of slit lamp biomicroscopic fundus images, *British Journal of Ophthalmology* **85** (2001), 563–565.
- [58] M. Akay, G. Landesberg, W. Welkowitz, Y.M. Akay and D. Sapoznikov, Carotid-cardiac interaction: Heart rate variability during the unblocking of the carotid artery, *Adv. Exp. Med. Biol.* **72** (1993), 346–365.
- [59] U. Wiklund, Analysis of heart rate variability in patients with severe autonomic dysfunction, in *Proc. 10th Nordic-Baltic Conf. Biomed Eng.*, (Vol. 34), 1996, pp. 383–384.
- [60] V. Pichot, J.M. Gaspoz, S. Molliex, A. Antoniadis, T. Busso, F. Roche, F. Costes, L. Quintin, J.R. Lacour and J.C. Barthelemy, Wavelet transform to quantify heart rate variability and to assess its instantaneous changes, *J. Appl. Physiol.* **86** (1999), 1081–1091.
- [61] U. Wiklund, M. Akay and U. Niklasson, Short-term analysis of heart-rate variability by adapted wavelet transforms, *IEEE Eng. Med. Biol. Mag.* **16**(5) (1997), 113–118.
- [62] L.G. Gamero, M. Risk, J.F. Sobh, A.J. Ramirez and J.P. Saul, Heart rate variability analysis using wavelet transform, *Proc. Computers in Cardiology* (1996), 177–180.
- [63] M. Akay, S. Morrison, U. Niklasson and U. Wiklund, Wavelet decomposition of cardiovascular signals for baroreceptor function tests in pigs, *IEEE Transactions on Biomedical Engineering* **49**(8) (2001), 651–661.
- [64] B. Litt, R. Esteller, J. Echaz, M. D'Alessandro, M. Dichter, R. Shor, T. Henry, P. Pennell, C. Epstein, R. Bakay and G. Vachtsevanos, Epileptic seizures may begin hours in advance of clinical onset: A report of five patients, *Neuron* **30** (2001), 51–64.
- [65] B. Litt and K. Lehnertz, Seizure prediction and the pre-seizure period, *Current Opinion in Neurology* **15**(2) (2002), 173–177.
- [66] P.J. Franaszczuk, G.K. Bergey, P.J. Durka and H.M. Eisenberg, Time-frequency analysis using matching pursuit algorithm applied to seizures originating from the mesial temporal lobe, *Electroenceph. Clin. Neurophysiol.* **106** (1998), 513–521.
- [67] S.G. Mallat and Z. Zhang, Matching pursuit with time-frequency dictionaries, *IEEE Transactions on Signal Processing* **41** (1993), 2297–3415.
- [68] J.J. Benedetto and G.E. Pfander, Periodic wavelet transforms and periodicity detection, *SIAP* **62**(4) (2002), 1329–1368.
- [69] S.D. Cranstoun, B. Litt, H.C. Ombao, W. Guo and R. von Sachs, Time-frequency spectral estimation of multichannel EEG using the auto-SLEX method, *IEEE Transactions on Biomedical Engineering* **49**(9) (2002), 988–996.
- [70] S. Hall, Brain pacemakers, *MIT Technology Review*, September 2001.
- [71] M. Zibulevsky and Y.Y. Zeevi, Extraction of a source from multichannel data using sparse decomposition, *Neurocomputing* **49** (2002), 163–173.
- [72] M. Zibulevsky, P. Kisilev, Y.Y. Zeevi and B.A. Pearlmutter, Blind source separation via multinode sparse representation, in *Proc. NIPS-2001*, 2002, pp. 185–191.
- [73] A.J. Bell and T.J. Sejnowski, An information-maximization approach to blind separation and blind deconvolution, *Neural Computation* **7** (1995), 1129–1159.
- [74] A. Hyvarinen, Survey on independent component analysis, *Neural Computational Surveys* **2** (1999), 99–128.
- [75] M. Zibulevsky and B.A. Pearlmutter, Blind source separation by sparse decomposition in a signal dictionary, *Neural Computation* **13** (2001), 862–882.
- [76] S. Mallat, Multifrequency channel decompositions of images and wavelet models, *IEEE Transactions on Acoustics, Speech, and Signal Processing* **37** (1989), 2091–2110.
- [77] M. Zibulski and Y.Y. Zeevi, Analysis of multiwindow Gabor-type schemes by frame methods, *Applied Comput. Harmonic Anal.* **1997** (1997), 188–221.

- [78] P. Kisilev, Y.Y. Zeevi and H. Pratt, Processing of single-trial evoked potentials: A non-parametric approach, Tech. Rep. 1180, Technion EE Pub., November 1998.
- [79] J. Duncan, R. Owen, L. Staib and P. Anandan, Measurement of non-rigid motion using contour shape descriptors, in *IEEE*, 1991, pp. 318–324.
- [80] E. Bardinet, L.D. Cohen and N. Ayache, Tracking and motion analysis of the left ventricle with deformable superquadrics, *Medical Image Analysis* **1** (1996), 129–149.
- [81] S. Benayoum, C. Nastar and N. Ayache, Dense non-rigid motion estimation in sequences of 3D images using differential constraints, in *Conference on Computer Vision, Virtual Reality and Robotics in Medicine*, 1995, pp. 687.
- [82] P. Shi, G. Robinson, C. Constable, A. Sinusas and J. Duncan, A model-based integrated approach to track myocardial deformation using displacement and velocity constraints, in *IEEE International Conference on Computer Vision*, Cambridge, MA, 1995, pp. 687–692.
- [83] E. Angelini, A. Laine, S. Takuma, J. Holmes and S. Homma, LV volume quantification via spatiotemporal analysis of real-time 3-d echocardiography, *IEEE Transactions on Medical Imaging* **20**(6) (2001), 457–469.
- [84] F. Meyer and R.R. Coifman, Brushlets: A tool for directional image analysis and image compression, *Applied and Computational Harmonic Analysis* **4** (1997), 147–187.
- [85] A. Dempster, M. Laird and D. Rubin, Maximum likelihood estimation from incomplete data via the EM algorithm, *J. Royal Stat. Soc.* **39** (1977), 1–38.
- [86] L. Shepp and Y. Vardi, Maximum likelihood reconstruction for emission tomography, *IEEE Trans. Medical Imaging* **1** (1982), 113–122.
- [87] P. Kisilev, M. Zibulevsky and Y.Y. Zeevi, Wavelet representation and total variation regularization in emission tomography, in *Proc. IEEE Int'l Conf. Image Processing*, Salonicki, 2001, pp. 702–705.
- [88] J. Kalifa, A. Laine and P.D. Essery, Regularization in tomographic reconstruction using thresholding estimators, *IEEE Transactions on Medical Imaging* **22**(3) (2003), 351–359.
- [89] M.S. Crouse, R.D. Nowak and R.G. Baraniuk, Wavelet-based statistical signal processing using hidden Markov models, *IEEE Transactions on Signal Processing* **46**(4) (1998), 886–902.
- [90] J.K. Romberg, H. Coi and R.G. Baraniuk, Bayesian tree-structured image modeling using wavelet domain hidden Markov models, *IEEE Transactions on Image Processing* **10**(7) (2001), 1056–1068.
- [91] H. Coi and R.G. Baraniuk, Multiscale image segmentation using wavelet-domain hidden Markov models, *IEEE Transactions on Image Processing* **10**(9) (2001), 1309–1321.
- [92] H. Cheng and C.A. Bouman, Multiscale Bayesian segmentation using a trainable context model, *IEEE Transactions on Image Processing* **10**(4) (April 2001), 511–525.
- [93] P. Sajda, L. Parra and C.D. Spence, Detection, synthesis and compression in mammographic image analysis using a hierarchical image probability model, in *Mathematical Methods in Biomedical Image Analysis*, M. Staib, ed., IEEE, 2001, pp. 3–10.
- [94] P. Sajda, C. Spence and L. Parra, A multi-scale probabilistic network model for detection, synthesis and compression in mammographic image analysis, *Medical Image Analysis* **1** (2003), 187–204.
- [95] A. Haar, Zur theorie der orthogonalen functionensysteme, *Math. Ann.* **28** (1910), 288–305.
- [96] I. Daubechies, *Ten Lectures on Wavelets*, SIAM, Philadelphia, PA, 1992.
- [97] J. Kovacevic and M. Vetterli, Nonseparable two and three-dimensional wavelets, *IEEE Trans. Signal Processing* **43** (1995), 1269–1273.
- [98] D. Stanhill and Y.Y. Zeevi, Two-dimensional orthogonal wavelets with vanishing moments, *IEEE Trans. Signal Processing* **44** (1996), 2579–2590.
- [99] D. Stanhill and Y.Y. Zeevi, Two-dimensional orthogonal filter banks and wavelets with linear phase, *IEEE Trans. Signal Processing* **46** (1998), 183–190.
- [100] M. Porat and Y.Y. Zeevi, The generalized Gabor scheme of image representation in biological and machine vision, *IEEE Trans. Pattern Anal. Machine Intell.* **10**(4) (1988), 452–468.
- [101] M. Porat and Y.Y. Zeevi, Localized processing of texture in vision, *IEEE Trans. Biomed. Eng.* **36** (1989), 115–129.
- [102] B.B. Mandelbrot, *The Fractal Geometry of Nature*, Freeman and Co., San Francisco, CA, 1982.
- [103] S. Mallat, A theory for multiresolution signal decomposition: the wavelet representation, *IEEE Trans. Pattern Anal. Machine Intell.* **11** (1989), 674–693.
- [104] S. Mallat, *Wavelet Tour of Signal Processing*, Academic Press, San Diego, CA, 1998.
- [105] B.M. ter Haar Romney, *Geometry Driven Diffusion in Computer Vision*, Kluwer Academic Publishers, Netherlands, 1994.
- [106] M. Kerckhove, ed., *Scale-Space and morphology in computer-vision*, Vol. LNCS-2106 of *Lecture Notes in Computer Science*, Berlin, 2001.
- [107] I. Polack, Segmentation and restoration in nonlinear multi-scale filtering, *IEEE Signal Processing Magazine* **19** (2002), 26–36.
- [108] P. Perona and J. Malik, Scale-space and edge detection using anisotropic diffusion, *IEEE Trans. Pattern Anal. Machine Intell.* **12** (1990), 629–639.
- [109] G. Gilboa, N. Sochen and Y.Y. Zeevi, A forward-and-backward diffusion process for adaptive image enhancement and denoising, *IEEE Trans. Image Processing* **11** (2002), 689–703.



**Hindawi**  
Submit your manuscripts at  
<http://www.hindawi.com>

