

DIRECTIONAL HYPERCOMPLEX WAVELETS FOR MULTIDIMENSIONAL SIGNAL ANALYSIS AND PROCESSING

Wai Lam Chan, Hyeokho Choi, Richard G. Baraniuk

Department of Electrical and Computer Engineering, Rice University, Houston, Texas, USA

ABSTRACT

We extend the wavelet transform to handle multidimensional signals that are smooth save for singularities along lower-dimensional manifolds. We first generalize the complex wavelet transform to higher dimensions using a multidimensional Hilbert transform. Then, using the resulting *hypercomplex wavelet transform* (HWT) as a building block, we construct new classes of nearly shift-invariant wavelet frames that are oriented along lower-dimensional subspaces. The HWT can be computed conveniently using a 1-D dual-tree complex wavelet transform along each signal axis. A series of denoising experiments demonstrates that the HWT significantly outperforms both standard tensor product wavelets and complex wavelets when the signals of interest sport a local manifold structure.

1. INTRODUCTION

The multiscale analysis of the discrete wavelet transform (DWT) has proved natural and powerful for dealing with singularity-rich one-dimensional (1-D) signals. In contrast to the sinusoidal basis functions underlying Fourier analysis, wavelet basis functions are localized in both time and frequency and thus yield very sparse and structured representations of *piecewise smooth signals* (signals that are smooth except for a finite number of jump discontinuities). This sparsity and structure boost the performance of algorithms such as denoising by shrinkage [1] and compression [2].

In this paper, we study how to best extend wavelet representations to handle piecewise smooth *multidimensional signals* (2-D, 3-D, and beyond). Such signals are smooth almost everywhere over their n -D domain but feature singularities on *manifold* structures of lower dimension $m \leq n$.¹ For example, an edge discontinuity in a 2-D image lies along a 1-D curve; a bee or aircraft flying through 3-D space tracks out a smooth 1-D curve in 4-D space-time; a video of a smooth object moving smoothly in time lies on a 2-D manifold in 3-D space-time; and points in space can be viewed as lying on 0-D manifolds in 3-D space.

Two fundamental problems have hindered the application of wavelets to multidimensional piecewise smooth signals. First, manifold singularities have a preferred *orientation*: they are smooth in the “direction” of the manifold (along the m -D tangent space) and contain a singularity (rapid change) in the “direction” of the normal (along the $(n-m)$ -D normal space). Thus, to obtain a sparse representation, wavelets intersecting the manifold should

be smooth (lowpass) in the direction of the manifold and oscillatory (bandpass) in the direction of the normal. Unfortunately, standard separable wavelets constructed from tensor products of 1-D wavelets have poor orientation selectivity. In 2-D, for example, one of the three tensor wavelets is simultaneously oscillatory at *both* $+45^\circ$ and -45° in the image plane. Figure 1(a) illustrates that the situation is even more grim in 3-D; this tensor wavelet is simultaneously oscillatory along the 4 diagonal directions within the cube. In general, tensor wavelets are only appropriate for representing 0-D structures (points) in n -D space.

Second, unlike the complex Fourier representation, the standard real-valued wavelet transform in any dimension is highly *shift varying*, meaning that a small shift in the signal drastically changes the values of the wavelet coefficients. This complicates wavelet coefficient modeling and processing considerably.

Recently, a promising approach for generating directional wavelets has emerged via the *complex wavelet transform* (CWT) [4, 5, 6]. Kingsbury’s 1-D *dual-tree* CWT expansion employs two independent wavelet filterbanks corresponding to a “real” wavelet and an “imaginary” wavelet that form an approximate Hilbert transform pair [5]. The resulting $2\times$ redundant tight frame enables a magnitude/phase coherent representation akin to that of Fourier analysis. Moreover, in 2-D a clever combination of the independent 1-D CWT filterbanks creates a directional set of six complex wavelets oriented at multiples of 30° ; these wavelets are very natural for representing 1-D edge structure in images [7]. As an added bonus, the smooth magnitude envelope of the complex wavelet yields a nearly shift-invariant transform.

In [6], Selesnick and Li extended this 2-D construction to a 3-D CWT that is natural for studying video sequences and other data with $(n = 3, m = 2)$. As we see in Fig. 1(b), each 3-D CWT wavelet oscillates along a 1-D line (its orientation) and is smooth in the orthogonal 2-D plane. Extended further to n -D, Selesnick and Li’s CWT is natural for representing signals always lying on $m = n - 1$ dimensional manifolds.

But what about signals with singularities on $m < n - 1$ dimensional manifolds (the 1-D trajectory of a bee in 3-D, for example)? In this paper, we develop a general theory for directional wavelets that are matched to any signal dimension n and manifold dimension $m \leq n$. We take a two-step approach. First, we introduce a new *hypercomplex wavelet transform* (HWT) that generalizes the CWT via an analogy to the hypercomplex Fourier transform and its associated generalized n -D Hilbert transform [8]. The HWT is conveniently obtained by a simple non-separable combination of 1-D dual-tree CWTs. Second, we show that the HWT serves as a basic building block for generating new classes of *directional hypercomplex wavelet transforms* (DHWTs) for handling different manifold dimensions m ; we merely form a simple linear combi-

Email: {wailam, choi, richb}@rice.edu. Web: dsp.rice.edu. This work was supported by ONR grant N00014-02-1-0353, AFOSR grant F49620-01-1-0378, DARPA, and the Texas Instruments Leadership University Program.

¹A manifold is a topological space that is locally Euclidean [3].

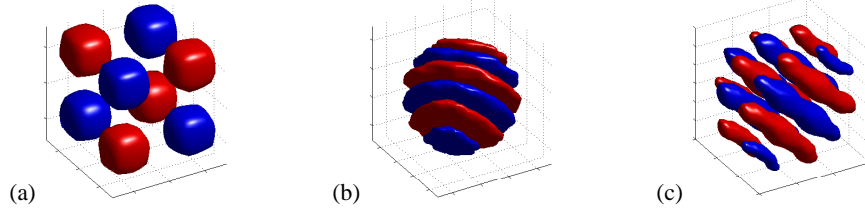


Fig. 1. Isosurfaces (points taking the same function value) of wavelet bases in 3-d space ($n = 3$). Red is large positive; blue is large negative. (a) Standard DWT tensor-product wavelet lacks any strong directionality; (b) $m = 2$ DHWT wavelet (real part) is bandpass along a 1-D line and lowpass along the orthogonal 2-D plane [6]; (c) $m = 1$ DHWT wavelet (real part) is bandpass on a 2-D plane and lowpass on the orthogonal 1-D line.

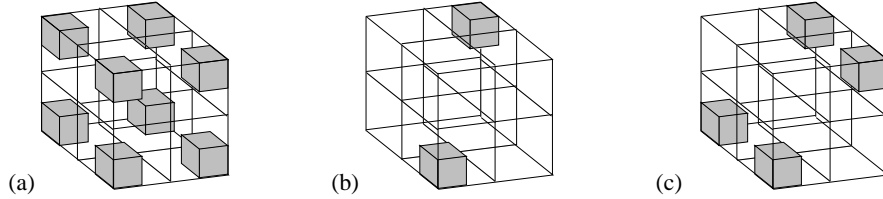


Fig. 2. Fourier-domain support of the 3-D wavelets from Fig. 1 ($n = 3$). Part (a) also represents the Fourier support of the corresponding HWT wavelet.

nations of the HWT components (orthogonal projections) to obtain each new directional transform. As an added bonus, the magnitude envelopes of the HWT and DHWT wavelets are smooth bumps, which makes the transforms nearly shift-invariant.

For different choices of n and m the HWT corresponds to some interesting known and new wavelets, including dual-tree complex wavelets ($n = 1$) [6], new *quaternion wavelets* ($n = 2$), new *octonion wavelets* ($n = 3$), and so on. In 3-D, octonion wavelets are the basic building blocks for constructing both Slesnick and Li's CWT for $m = 2$ (see Fig. 1(b)) and a new DHWT for $m = 1$ (see Fig. 1(c)).

This paper is organized as follows. In Sections 2 and 3 we briefly review the DWT and CWT. Section 4 introduces the HWT and Section 5 develops the idea of orthogonal projections to obtain DHWTs. Section 6 presents the results of a number of denoising experiments that demonstrate significant gains using a DHWT with directionality matched to the manifold dimension at hand. In the following, we will mainly focus on the cases $n = 2$ and $n = 3$ for clarity; however, all results generalize to arbitrary n and m .

2. DISCRETE WAVELET TRANSFORM (DWT)

The discrete wavelet transform (DWT) represents a 1-D real signal $f(t)$ in terms of shifted versions of a *scaling function* $\phi(t)$ and shifted and scaled versions of a *wavelet function* $\psi(t)$ [2]. When $\phi_{L,p}(t) = 2^L \phi(2^L t - p)$ and $\psi_{\ell,p}(t) = 2^\ell \psi(2^\ell t - p)$ form an orthonormal basis, we can represent any $f(t) \in L_2$ as

$$f(t) = \sum_{p \in \mathbb{Z}} c_{L,p} \phi_{L,p}(t) + \sum_{\ell \geq L, p \in \mathbb{Z}} d_{\ell,p} \psi_{\ell,p}(t) \quad (1)$$

where $c_{L,p} = \int f(t) \phi_{L,p}(t) dt$ and $d_{\ell,p} = \int f(t) \psi_{\ell,p}(t) dt$ are the scaling and wavelet coefficients. L sets the coarsest scale space that is spanned by $\phi_{L,p}(t)$. Behind each wavelet transform is a filterbank based on a lowpass filter; we will use the notation $\phi_h(t)$, $\psi_h(t)$ to denote the scaling and wavelet functions corresponding to a particular filter h .

DWTs in higher dimensions are typically obtained using tensor products of 1-D DWTs over each dimension. In 2-D, for example, we use the scaling function $\phi(x)\phi(y)$ and three wavelets

$\psi(x)\psi(y)$, $\phi(x)\psi(y)$, and $\psi(x)\phi(y)$ that are oriented in the diagonal, horizontal, and vertical directions, respectively [6].

3. COMPLEX WAVELET TRANSFORM (CWT)

The 1-D dual-tree CWT expands a real signal in terms of *two* sets of wavelets and scaling functions obtained from two independent filterbanks based on real filters h and g [5, 6]. $\psi_h(t)$ and $\psi_g(t)$ play the roles of the real and imaginary parts of a complex analytic wavelet $\psi^c(t) = \psi_h(t) + j\psi_g(t)$ that is supported only on positive frequencies in the Fourier domain. Equivalently, the imaginary wavelet is the *Hilbert transform* of the real wavelet. The combined system is a $2 \times$ redundant frame that, by virtue of the fact that $|\psi^c(t)|$ is a smooth bump, is shift-invariant.²

It is useful to recall that the Fourier transforms of the Hilbert pair of wavelets are related by $\Psi_g(f) = -j \operatorname{sgn}(f) \Psi_h(f)$, where $\operatorname{sgn}(f)$ equals 1 when $f > 0$ and -1 when $f < 0$. Thus, summing $\Psi_h(f) + j\Psi_g(f) = \Psi^c(f)$ cancels the negative frequency part of the complex wavelet, rendering it analytic.

This notion of using 1-D Hilbert transforms to cancel out negative frequencies has been used to create directional wavelets in 2-D [5], 3-D, and beyond [6]. However, as we will see, these CWTs are matched only to manifolds of intrinsic dimension $m = n - 1$. To best analyze manifolds of lower intrinsic dimension, we must turn to the theory of multidimensional Hilbert transforms.

4. HYPERCOMPLEX WAVELET TRANSFORM (HWT)

There are no unique definitions of Hilbert transform and analytic signal in 2-D and higher. To illustrate two possibilities, consider a 2-D wavelet Fourier transform $\Psi(f_x, f_y)$. The straightforward extension of the 1-D definition sets $\Psi(f_x, f_y) = 0$ on all but a half-plane ($f_x > 0$, for example) and uses standard complex algebra for manipulation. Another natural definition sets $\Psi(f_x, f_y) = 0$

²In practice, in order to have finite-length wavelets, the Hilbert transform is only approximately satisfied, $\psi^c(t)$ is only approximately analytic, and the CWT is only approximately shift-invariant [5, 6].

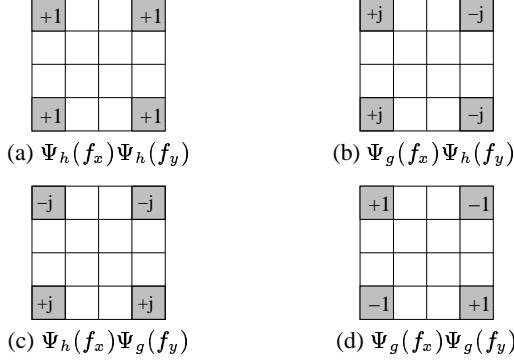


Fig. 3. Fourier-domain relationships among the four components of a 2-D HWT wavelet $\Psi(f_x, f_y)$ in the diagonal subband.

on all but a quadrant ($f_x, f_y > 0$, for example) and uses *quaternion* algebra for manipulation [9].³

Adopting the latter rather than the former definition, we will now introduce a new multiscale signal representation: the dual-tree *hypercomplex wavelet transform* (HWT). For clarity, we will stick to 2-D, but the HWT extends easily to any dimension n .

Each 2-D hypercomplex wavelet consists of a standard DWT tensor wavelet plus three additional real wavelets obtained from the three combinations of 1-D Hilbert transforms along the two coordinates. More specifically, denote the 1-D Hilbert transform operators along the x and y coordinates as \mathcal{H}_x and \mathcal{H}_y , respectively. Then given the diagonal tensor product wavelet $\psi_h(x)\psi_h(y)$ from Section 2, we complement it with

$$\mathcal{H}_x\{\psi_h(x)\psi_h(y)\} = \psi_g(x)\psi_h(y), \quad (2)$$

$$\mathcal{H}_y\{\psi_h(x)\psi_h(y)\} = \psi_h(x)\psi_g(y), \quad (3)$$

$$\mathcal{H}_x\mathcal{H}_y\{\psi_h(x)\psi_h(y)\} = \psi_g(x)\psi_g(y). \quad (4)$$

Conveniently, each component can be computed as a combination of 1-D dual-tree complex wavelets. These three components closely resemble $\psi_h(x)\psi_h(y)$ but are phase-shifted by 90° in the horizontal, vertical, and diagonal directions, respectively; Fig. 3 interprets these relationships in the Fourier domain. Using quaternion algebra, we can compactly summarize the above process in a *quaternion wavelet* $\psi^q(x, y) = \psi_h(x)\psi_h(y) - j_1\psi_g(x)\psi_h(y) - j_2\psi_h(x)\psi_g(y) + j_3\psi_g(x)\psi_g(y)$. The construction of the other two quaternion wavelets based on $\phi_h(x)\psi_h(y)$ and $\psi_h(x)\phi_h(y)$ is identical. The resulting HWT, which we term the dual-tree *quaternion wavelet transform* (QWT) in this special 2-D case, is a $4 \times$ redundant tight frame.

Generalization to 3-D and higher is straightforward. In n -D there are $2^n - 1$ different hypercomplex wavelets. Each wavelet consists of a standard DWT tensor wavelet plus $2^n - 1$ additional real wavelets obtained from the $2^n - 1$ combinations of 1-D Hilbert transforms along the n coordinates. The total redundancy is $2^n \times$.

Since we construct the HWT basis by Hilbert transforming a standard DWT basis, the HWT basis has no better directional selectivity than the DWT basis. However, just as the 1-D Hilbert transform enables us to zero-out the negative frequency axis to form an analytic signal, the n -D hypercomplex Hilbert transform allows us zero-out all but one “ n -tant”⁴ in the Fourier domain. For example, to obtain a single diagonally oriented 2-D wavelet

³The set of quaternions $\mathbb{H} = \{a + j_1b + j_2c + j_3d | a, b, c, d \in \mathbb{R}\}$ with multiplication rules $j_1j_2 = -j_2j_1 = j_3$ and $j_1^2 = j_2^2 = j_3^2 = -1$ [10].

⁴The natural n -D generalization of a quadrant in 2-D and octant in 3-D.

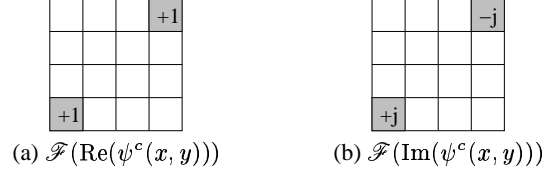


Fig. 4. Fourier-domain relationships between the real and imaginary parts of a 2-D diagonally oriented complex wavelet. $\mathcal{F}(\cdot)$ is the Fourier Transform.

whose frequency support is focused in the second quadrant of the Fourier plane, we combine the four quaternion components as $(\psi_h(x)\psi_h(y) + \psi_g(x)\psi_g(y)) + j(\psi_h(x)\psi_g(y) - \psi_g(x)\psi_h(y))$. (See Fig. 3, and also note that this is an orthogonal projection.) This capability will prove extremely useful for constructing directional wavelets in the next section.

5. DIRECTIONAL HWT (DHWT)

The HWT provides a convenient set of building blocks for constructing new directional wavelet transforms. The idea is simple: we can generate a new wavelet with localized Fourier support, and hence directionality [6], by carefully recombining the 2^n components of an n -D HWT wavelet. The idea is also powerful: with the HWT we can generate n -D wavelets that are bandpass in subspaces of arbitrary dimension m .

5.1. DHWT in 2-D

We illustrate the procedure by first re-deriving the 2-D CWT of [5, 6] from the QWT of last section. The Fourier-domain support of each 2-D QWT wavelet consists of four dyadic squares symmetrically placed with respect to the origin. This limits its directionality. For example, the Fourier support of the diagonally oriented QWT wavelet lies at the four corners of a square in the Fourier plane (see Fig. 3(a)); hence this wavelet is localized simultaneously along two lines of slope $\pm 45^\circ$. To obtain a wavelet directed along a single line (at either $+45^\circ$ or -45°) we need to partition the Fourier plane in half to isolate pairs of squares that lie on a straight line through the origin. Since we can isolate individual quadrants by recombining the QWT components (recall the end of the Section 4), we can generate a wavelet with two-quadrant Fourier support as in Fig. 4(a). In the signal domain, this wavelet will be lowpass along the $+45^\circ$ direction and bandpass along the orthogonal -45° direction [5, 6].

Repeating the same procedure for all three QWT wavelets with two directions per subband, we obtain six directional basis functions, each of which is a complex analytic signal. This is precisely the 2-D CWT [5, 6]. Since the basis functions are analytic, the CWT is shift-invariant.

5.2. DHWT in 3-D

The 3-D case is more interesting because there exist two entirely different DHWTs corresponding to $m = 1, 2$, and one is new. Let us focus on the “diagonal” HWT wavelet whose Fourier support is located in the eight corner cubes in 3-D Fourier space (see Fig. 2(a)). This dual-tree HWT basis function has eight *octonion* components:

$$\begin{aligned} \psi_{1,1} &= \psi_h(x)\psi_h(y)\psi_h(z), & \psi_{2,1} &= \psi_g(x)\psi_g(y)\psi_g(z), \\ \psi_{1,2} &= \psi_g(x)\psi_g(y)\psi_h(z), & \psi_{2,2} &= \psi_h(x)\psi_h(y)\psi_g(z), \\ \psi_{1,3} &= \psi_h(x)\psi_g(y)\psi_g(z), & \psi_{2,3} &= \psi_g(x)\psi_h(y)\psi_h(z), \\ \psi_{1,4} &= \psi_g(x)\psi_h(y)\psi_g(z), & \psi_{2,4} &= \psi_h(x)\psi_g(y)\psi_h(z). \end{aligned} \quad (5)$$

Table 1. 3-D denoising results in log-MSE.

test signal	CWT ($m = 2$)	DHWT ($m = 1$)	DWT ($m = 0$)
Stars	60.96 dB	53.82 dB	62.84 dB
Rods	62.77 dB	60.24 dB	64.97 dB
Helix	63.68 dB	61.02 dB	64.98 dB
Sheet	59.35 dB	52.35 dB	61.67 dB

As in 2-D, we can linearly combine these eight components to obtain a wavelet with spectral support in a single octant. We now have two choices, depending on whether $m = 1$ or 2.

Case $m = 2$: To form each DHWT wavelet, we group a pair of cubes located on a 1-D line running through the origin (see Fig. 2(b)). The resulting wavelet is complex and bandpass along the 1-D line⁵ and lowpass in the orthogonal direction, which is a 2-D plane through the origin (see Fig. 1(b)). This coincides with the 3-D CWT of Selesnick and Li [6].

Case $m = 1$: Alternatively, we group sets of four cubes located on a 2-D plane passing through the origin (see Fig. 2(c)). The resulting wavelet is bandpass along the 2-D plane and lowpass in the orthogonal direction, which is a 1-D line through the origin (see Fig. 1(c)).

Due to the way the four cubes are combined – for example as $\psi_{1,1} \pm \psi_{1,3}$ and $\psi_{2,1} \pm \psi_{2,3}$ – the wavelets behave as a quaternion in the direction parallel to the 2-D plane. This yields a smooth magnitude envelope and shift-invariant transform. A similar recipe works for the other directions.

Starting from the HWT wavelet in Fig. 1(a), there are six ways to group four corner cubes on 2-D planes through the origin. This result holds for each of the three HWT wavelets. Therefore, the $m = 1$ DHWT basis that covers all possible 3-D directions has 18 such directions and thus three times the redundancy of the $m = 2$ DHWT (24× redundant in total).

The same construction applies for arbitrary n and m .

6. DENOISING EXPERIMENTS

To evaluate the performance of the DHWT, we performed a number of denoising experiments in 3-D using data of different intrinsic manifold dimension and wavelets matched to $m = 0, 1, 2$.

The first three test signals have $m = 1$: *Stars* is a synthetic video that pans across a field of stars and galaxies (size: $64 \times 64 \times 64$); *Rods* contains solid rods of various radii and intensities (size: $64 \times 64 \times 64$); *Helix* maps out a smooth 1-D helical curve (size: $64 \times 64 \times 64$). The final test signal has $m = 2$: *Sheet* is a thin tilted plane (size: $64 \times 64 \times 64$).⁶

We added white Gaussian noise of $\sigma = 15$ to each data set and applied optimal hard thresholding following the program in [6] (see taco.poly.edu/WaveletSoftware). We used the dual-tree wavelet transform from [6] with length-10 wavelets. We measured the log-mean-square (LMS) error of the denoised data; LMS errors before denoising were 65.58dB, 76.43dB, 69.67dB, and 77.71dB, respectively.

From Table 1 we observe the superiority of the $m = 1$ DHWT over the $m = 2$ CWT and the usual tensor product DWT. We are surprised that the $m = 2$ CWT does not dominate on the sheet

data, since it should theoretically be better matched to the intrinsic manifold dimension. Perhaps the difference is due to the extra averaging available in the $3\times$ more redundant $m = 1$ transform.

7. CONCLUSIONS

We have introduced a new family of multiscale signal representations, the directional hypercomplex wavelet transforms, which show considerable promise for multidimensional signal processing. DHWTs are especially designed for signals with a local manifold structure, precisely where standard tensor product wavelets have often disappointed. Their superior performance on a number of 3-D denoising experiments suggests that they have excellent sparsification capabilities. Perhaps the main challenge to both CWT and DHWT methods will be their redundancy, since multidimensional data processing problems are already often computationally challenging. Certainly controlling redundancy while preserving directionality is a fertile area for future research.

8. REFERENCES

- [1] D. L. Donoho, “De-noising by soft-thresholding,” *IEEE Trans. Info. Theory*, vol. 41, no. 3, pp. 613–627, 1995.
- [2] M. Vetterli and J. Kovacevic, *Wavelets and Subband Coding*, Prentice Hall, Englewood Cliffs, NJ, USA, 1995.
- [3] W. Boothby, *An Introduction to Differentiable Manifolds and Riemannian Geometry*, Academic, San Diego, CA, USA, 2003.
- [4] J. M. Lina and M. Mayrand, “Complex Daubechies wavelets,” *ACHA*, vol. 2, pp. 219–229, 1995.
- [5] N. G. Kingsbury, “Complex wavelets for shift invariant analysis and filtering of signals,” *ACHA*, vol. 10, no. 3, pp. 234–253, May 2002.
- [6] I. W. Selesnick and K. Y. Li, “Video denoising using 2d and 3d dual-tree complex wavelet transforms,” in *SPIE Wavelets X*, San Diego, CA, August 4-8 2003, vol. 76.
- [7] J. Romberg, M. Wakin, H. Choi, and R. Baraniuk, “A geometric hidden Markov tree wavelet model,” in *SPIE Wavelets X*, San Diego, CA, August 2003, vol. 10.
- [8] T. Bülöw and G. Sommer, “Hypercomplex Signals – A novel extension of the analytic signal to the multidimensional case,” *IEEE Trans. Signal Proc.*, vol. 49, no. 11, pp. 2844–2852, 2001.
- [9] T. Bülöw, M. Felsberg, and G. Sommer, *Geometric Computing with Clifford Algebra*, chapter 11, Springer-Verlag, Berlin, 2001.
- [10] I. L. Kantor and A. S. Solodovnikov, *Hypercomplex Numbers*, Springer-Verlag, 1989.

⁵Due to the rotation property of the Fourier transform, the wavelet will actually be bandpass along a 1-D line in the space domain that is rotated 90° with respect to the Fourier-domain line.

⁶Test signals are available from dsp.rice.edu/Software/HWT.