Phase and Magnitude Perceptual Sensitivities in Nonredundant Complex Wavelet Representations

Michael B. Wakin, Michael T. Orchard, Richard G. Baraniuk, Venkat Chandrasekaran
Department of Electrical and Computer Engineering, Rice University

***Abstract*— The recent development of a nonredundant complex wavelet transform allows a novel framework for image analysis. Work on this representation has recognized that the phase and magnitude of complex coefficients can be related to important geometric properties in images. Existing work on human visual system (HVS) sensitivity offers little guidance in understanding the relative importance of noise (e.g., introduced by lossy coding) in phase components and magnitude components. The distinct geometric significance of the two components would suggest that their respective errors relate to different types of image structure, and thus each would have its own unique HVS sensitivity. In this paper, we extend the study of just-noticeable-differences (JND) to magnitude/phase sensitivities in complex wavelet representations and outline and report on preliminary experiments characterizing them.**

I. Introduction

Human perception of an image can be analyzed in a variety of frameworks and contexts. Research of the mammalian visual cortex, for example, has revealed distinct neural responses for certain localized primitive image characteristics such as orientation and spatial frequency [1]. In addition, Olshausen and Field [2] indicate that natural image components are well described by atoms of various scales and orientations. These observations suggest that wavelet-like representations are appropriate for images. Indeed, wavelets have the additional benefit of a filter-bank implementation and have had a great deal of practical success in image processing [3]. Consequently, psychophysical studies have examined the perceptual effects of quantization in wavelet coefficients [4, 5], and these results have influenced the development of wavelet-based image coders [6].

A. A Cognitive Theory of Perceptual Discrimination

Psychophysical studies are limited by the choice of representation — any given representation reflects certain image characteristics more clearly than others. Consequently, errors[1] in the coefficients will introduce certain kinds of changes in the image, and these are the types of perceptual discrimination that may be readily studied. In the case of the wavelet representation, quantization of wavelet coefficients introduces changes that manifest themselves as ringing artifacts around edges. At a cognitive level, however, there is no natural meaning associated with such artifacts. A study of these effects may lead to practical benefits (such as a new image coder [6], for example), but may also fail to illuminate certain aspects of the human perception of image information.

In this paper, we extend the study of perceptual discrimination in the context of a *cognitive* theory of image information. In [7], Marr writes that, "vision is the *process* of discovering from images what is present in the world, and where it is." We propose a simple model in which the human understanding of image information can be classified into two distinct concepts:

- ***What*** information describes the contents and features contained in the image. Example: "The image contains a shiny red car with a dent."
- ***Where*** information describes the locations and arrangement of the *What* information. Example: "The car is on the left side of the image, facing right, and the dent is in the fender."

Although it would be meaningless to completely decouple these pieces of information, and the nuances of such a model are open to debate, we believe this framework provides a reasonable basis for a novel study of human perception.

Section II provides a further discussion regarding our interpretation of this model and reveals that within a wavelet representation, *What* and *Where* information are strongly coupled. This prevents any reasonable wavelet-based psychophysical study from revealing valuable information about the human discrimination of these two effects.

B. The Role of a Nonredundant Complex Wavelet Representation

In this paper, we propose that a complex wavelet representation serves to reasonably decouple these two functionalities. The potential benefit of a complex-valued signal representation is widely accepted due to the Fourier transform. Fourier coefficient magnitudes relate to signal intensity; coefficient phases relate to locations, or offsets. With the Fourier transform, however, such properties pertain to global signal characteristics. The wavelet transform provides a similar expansion using localized real-valued sinusoids; unfortunately, the concepts of phase and magnitude are lost. As discussed in Section II, these components can be recovered through the concept of the analytic signal. The result is a complex wavelet transform that distinguishes localized *What* and *Where* information using the coefficient magnitudes and phases, respectively.

This work was supported by the NSF, ONR, AFOSR, DARPA, and the Texas Instruments Leadership University Program.

Email: {wakin, orchard, richb, venkatc}@rice.edu. Web: dsp.rice.edu

[1] We consider only simple kinds of distortion in the coefficients, since it is possible to obtain any kind of distortion through a proper (but complicated) adjustment of wavelet coefficients.

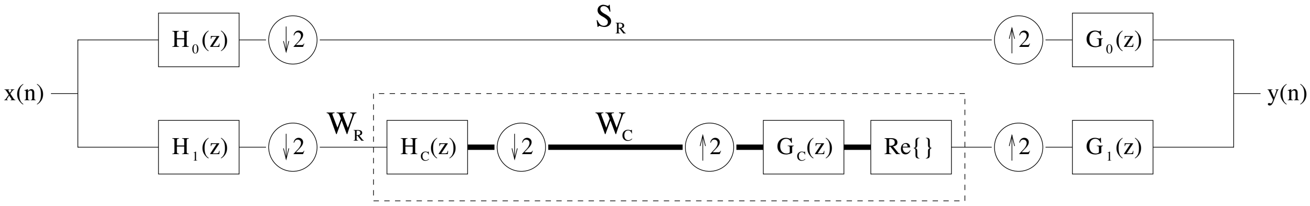

Fig. 1. *Single stage 1D nonredundant complex wavelet transform, consisting of real wavelet filters H_0 and H_1, with an additional complex filter stage H_C. Complex signals are shown in bold. Real scaling coefficients are denoted S_R; complex wavelet coefficients are denoted W_C. The complex stage (boxed) may be removed to leave a standard real DWT with coefficients W_R.*

The distinct geometric significance of the magnitude and phase components motivates a psychophysical study to better understand human visual system (HVS) sensitivity to errors in *What* and *Where* information. This study will be facilitated through the use of a *nonredundant* complex wavelet transform. Although recent years have seen the development of several overcomplete (redundant) complex wavelet transforms [8, 9], the redundancy of such representations may complicate controlled studies. For example, there is no unique coefficient expansion for a given image, and certain coefficient distortions may offset each other. These problems are avoided with a nonredundant complex wavelet representation. Several such designs exist [10–13]; we choose and discuss one [13] in this paper.

This paper is organized as follows. Section II develops the nonredundant complex wavelet transform. In Section III, we propose preliminary methods for extending the study of just-noticeable-differences (JND) to magnitude/phase sensitivities. We discuss the unique and novel issues which arise in developing the appropriate psychophysical experiments. Section IV presents some preliminary experimental results, indicating that the HVS may have a unique sensitivity for errors in phase and magnitude information, and that the relative importance may depend on the spatial frequency.

II. Nonredundant Complex Wavelets

The Fourier transform (decomposition of a signal into complex sinusoids) gives rise to the concepts of amplitude and phase. These relate to distinct, global signal properties: amplitude defines the energy of each sinusoid, and phase defines the offset, or location. The wavelet transform provides a similar expansion using localized real-valued sinusoids; unfortunately, the concepts of phase and magnitude are lost.

A. Analytic Signal

An oscillatory real-valued signal can be viewed as the real component of a complex-valued *analytic* signal, whose phase and magnitude reveal fundamental local signal characteristics that are difficult to discern from the real component alone. Given a real 1D signal $f(t)$, the analytic signal $f_A(t)$ is given by $f_A(t) = f(t) + jf_{\mathcal{H}}(t)$, where $f_{\mathcal{H}}(t)$ represents the Hilbert transform of $f(t)$ [14]. In the frequency domain,

$$F_A(\omega) = F(\omega) \cdot \mathbf{1}_{\omega \geq 0} = \begin{cases} F(\omega), & \omega \geq 0 \\ 0, & \omega < 0. \end{cases}$$

As discussed in [14], the magnitude envelope of the analytic signal defines the amplitude of a local sinusoid, and the phase relates to small local shifts.

B. Nonredundant Complex Wavelet Transform

In general, a real signal is composed of a combination of different frequencies. The real wavelet transform provides a bandpass filter bank that decomposes the signal into oscillatory components. By constructing the corresponding analytic signals, we obtain a complex wavelet transform. Along these lines, a number of redundant [8, 9] and nonredundant [10–13] complex wavelet designs have been developed in recent years.

In [13], for example, Orchard and Ates define a complex filter pair $H_C(z), G_C(z)$ according to the following criteria:

- The idealized frequency response of H_C is $H_C(\omega) = \mathbf{1}_{\omega \geq 0}$. Thus, for a real input f, the convolution $f * H_C$ provides an approximate analytic signal.
- The complex signal $f * H_C$ may be downsampled by a factor of 2. The filter G_C allows perfect reconstruction of the real input f.

These filters are used to define a nonredundant complex wavelet transform, as shown in Fig. 1. After passing a real input through the standard real wavelet filters $H_0(z)$ (lowpass) and $H_1(z)$ (highpass), the real highpass wavelet coefficients are filtered with $H_C(z)$ and downsampled. The frequency-domain effects are shown in Fig. 2. If we start with a 1D length-N signal, one stage of the nonredundant complex wavelet transform produces $N/2$ real scaling coefficients and $N/4$ complex wavelet coefficients. To implement multiple stages, the real wavelet filters continue to iterate as usual on the real signals; the complex filters are the final operation at each scale.

C. 2D Implementation for Images

As shown in Fig. 3, construction of a 2D nonredundant complex wavelet transform follows directly from the above procedure. For example, on the real vertical wavelet subband (generated by a highpass row filter and a lowpass column filter), we apply the complex filter $H_C(z)$ to each row. Similarly, we apply $H_C(z)$ to each column of the horizontal subband.[2] Thus, if we start with a 2D $N \times N$ signal, one stage

[2]For the real diagonal subband, we are unable to apply $H_C(z)$ to both the rows and columns because the H_C, G_C pair is designed for real inputs. For our preliminary investigation, we apply no complex filters to the diagonal subband and restrict our attention to the others.

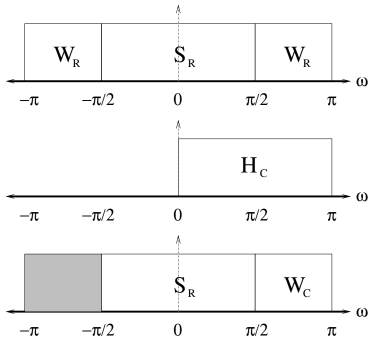

Fig. 2. *Idealized frequency-domain responses. Top: 1D real wavelet transform. Middle: complex filter $H_C(z)$. Bottom: resulting 1D nonredundant complex wavelet transform. The shaded region can be reconstructed without loss of information.*

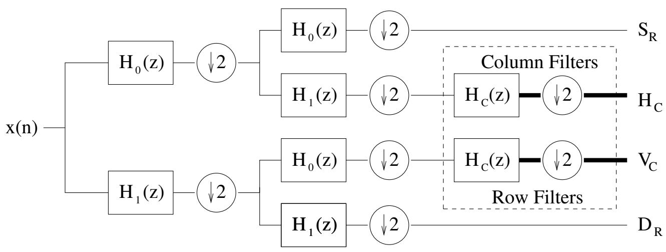

Fig. 3. *Single stage 2D nonredundant complex wavelet decomposition, resulting in a real scaling subband (S_R), complex horizontal (H_C) and vertical (V_C) subbands, and a real diagonal subband (D_R).*

of the nonredundant complex wavelet transform produces an $N/2 \times N/2$ real scaling subband, an $N/2 \times N/2$ real diagonal wavelet subband, an $N/2 \times N/4$ complex vertical wavelet subband, and an $N/4 \times N/2$ complex horizontal wavelet subband.

D. *Separation of* What *and* Where *Information*

By introducing the concepts of magnitude and phase to a wavelet representation, the nonredundant complex wavelet representation decouples *What* and *Where* information that are intertwined within a real wavelet representation. As a brief illustration, we present the following example. In Fig. 4(a), we construct a sinusoidal grating image (with energy concentrated primarily within one wavelet subband). We quantize the real wavelet coefficients (see Fig. 4(b)), then restore *either* the actual complex phase or the actual complex magnitude. (The resulting images contain errors only in the complex magnitude or phase, respectively.) The effects are profound:

- Magnitude errors (Fig. 4(c)) appear primarily as local brightness changes of the grating lines.
- Phase errors (Fig. 4(d)) appear primarily as local shifts (bends) of the grating lines, with few noticeable changes in brightness.
- The original (real wavelet) errors (Fig. 4(b)) appear as a combination of the above effects.

This illustrates the mixing of *What* and *Where* information within a real wavelet representation.

III. Perceptual Experiment Design

As demonstrated above, the phase and magnitude components of the nonredundant complex wavelet coefficients represent distinct image characteristics. The distinct geometric significance of the two components would suggest that each would have its own unique HVS sensitivity. Following the demonstration of Oppenheim and Lim [15], for example, where Fourier transforms are reconstructed using only magnitude or phase information, a commonly held belief is that phase information is much more critical to human perception of an image. In the nonredundant complex wavelet transform, phase and magnitude play similar, but more localized, roles. In this section, we discuss a preliminary investigation into the relative importance of local phase and magnitude information, as we extend the study of just-noticeable-differences (JND) to nonredundant phase and magnitude sensitivities. We present preliminary results in Section IV.

A. *Quantization*

A traditional psychophysical study of JND operates by introducing errors to an image, and determining the noise level at which these errors become noticeable. Recent investigations involving real wavelets produce these errors errors as a result of scalar quantization on individual wavelet subbands [4, 5]. Quantization provides a generic model for the typical effects of lossy compression. Other sources of distortion (such as additive noise in the coefficients [16]), may produce different artifacts but may be less relevant in practice. The differences of these effects using real wavelets are most pronounced in smooth image regions, where quantization may set large patches of coefficients to zero.

The design of a JND experiment for nonredundant complex wavelets must consider a number of factors. Some of these result from the non-orthogonality of the magnitude and phase representation. For example, the maximum phase error is bounded by π — beyond that point the phase error begins to decrease. Also, the maximum magnitude error is bounded in one direction (toward the origin) — moving beyond that point would effectively change the phase by π.

As a preliminary investigation, we propose as a source of error the process of scalar quantization on complex phases and magnitudes. We believe this is a reasonable model for the operation of a lossy coder, and it can be designed to operate appropriately on the phase and magnitude separately. We quantize all coefficients with a uniform quantization stepsize. To be precise, let $z = re^{j\theta}$ be a complex coefficient with $\theta \in [0, 2\pi)$, and let Δ be the desired quantization stepsize. After magnitude quantization, the coefficient is given by $\widehat{r}e^{j\theta}$ where $\widehat{r} = \Delta \cdot \mathrm{round}(r/\Delta)$. For phase quantization, we divide the circle into an integral number of segments, each with arclength approximately Δ. After phase quantization, the coefficient is given by $re^{j\widehat{\theta}}$ where

$$\begin{aligned} \Delta' &= \frac{2\pi}{\lceil 2\pi r/\Delta \rceil} \\ \widehat{\theta} &= (\lceil \theta/\Delta' \rceil - 0.5)\Delta'. \end{aligned}$$

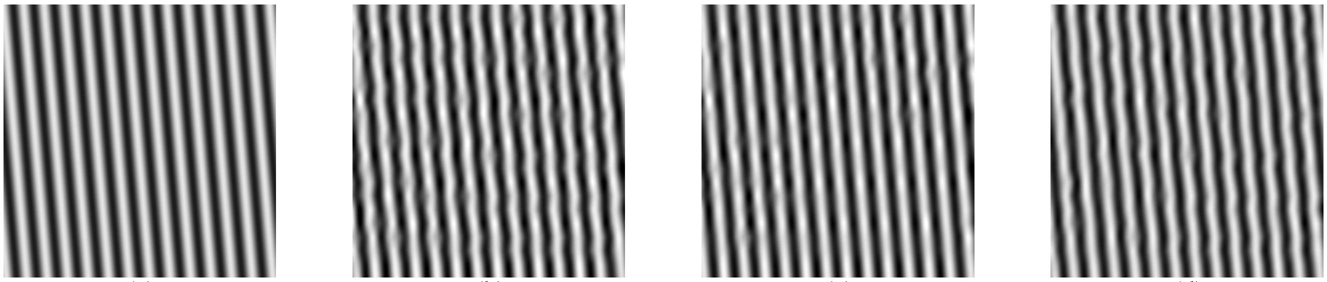

(a) (b) (c) (d)

Fig. 4. *(a) Original grating image. (b) Subject to quantization of real wavelet coefficients, PSNR = 26.2dB. (c) Image from (b) with complex phases restored, PSNR = 29.4dB; note the local brightness changes on grating lines. (d) Image from (b) with complex magnitudes restored, PSNR = 28.7dB; note the local bends in grating lines.*

We consider quantization of one subband (at one orientation) at a time. Future studies may consider more advanced sources of error.

B. Images

In addition to the source of error, the choice of image may impact the nature of the observed distortions. On natural images, for example, quantization effects of phase or magnitude within a single subband are not easily distinguishable (compared to the demonstration in Section II-D). Most artifacts are observed near edges, where each type resembles the common ringing artifacts of real wavelet quantization. Indeed, limited perceptual tests with one author (MW) indicate that the JND levels for magnitude and phase quantization are nearly equal to the JND level for real wavelet quantization. These effects are likely a consequence of the fact that natural edges contain energy across the frequency spectrum. Adjusting wavelet coefficients at a single scale fails to exploit their multiscale coherency. Future studies may examine the impact of coherent, multiscale quantization of complex wavelet coefficients. For the purposes of this paper, we propose a perceptual study involving sinusoidal grating images. These images can be tuned to maintain frequency content within the octave of interest. As shown in Section II-D, the resulting effects of phase and magnitude errors are distinct.

In a standard grating image, we observe within the complex wavelet subband of interest that most coefficient magnitudes are large and nearly equal. Roughly speaking, then, magnitude quantization will treat them uniformly (taking them all to the same new value). The visual impact is a uniform change in the grating contrast (the sinusoidal amplitude). This is contrary to the magnitude effects observed in Section II-D, and it illustrates the fact that the magnitude field is low-dimensional — there are strong constraints on the magnitudes that are not observed in the randomly spread real wavelet coefficients. Indeed, similar low-dimensional constraints hold for the *differences* between neighboring phase values. The complex representation brings forth the simple structure of the grating image.

As a slightly more complicated grating image with frequency in the appropriate octave, we propose a *chirp grating* as shown in Fig. 5(a). Such a grating can be be constructed with sinusoidal frequency ranging from the low to high end of the octave. As shown in Fig. 5(c), the complex magnitudes of such a grating exhibit sufficient variation to allow for spatially varying quantization artifacts.

C. Experimental Procedure

We use a three-alternative forced-choice (3AFC) procedure for determining JND thresholds. In each trial, the test subject is shown three stimuli. Two stimuli contain only the original image (mask only), while one stimulus contains the original image (mask) subject to quantization distortions (target). The amount of error is varied according to a QUEST staircase using the Psychophysics Toolbox software [17]. The JND threshold is estimated as the point where the subject correctly identifies 75% of target images. Many of the experimental parameters (number of trials, spatial frequency, etc.) follow the procedure carefully outlined in [5].

Following similar experiments, our metric for the amount of error introduced in the visible image is the RMS contrast [5]. Let X be the original image containing N pixels, and let X^q be the image containing quantization noise. (The image X^q is presented to the subject.) Define L_{mean} to be the mean luminance of X, and let the distortion image $X^d = X^q - X + X^\ell$, where X^ℓ is a constant image with luminance L_{mean}. The RMS contrast is then given by

$$C_{\text{RMS}}(X, X^q) = \frac{1}{L_{\text{mean}}}\left[\frac{1}{N}\sum_{i=1}^{N}(L_i - L_{\text{mean}})^2\right]^{1/2},$$

where L_i is the luminance at pixel i of X^d.

We run separate, randomly ordered trial blocks to determine the subject's sensitivity to real wavelet quantization, complex phase quantization, and complex magnitude quantization. Throughout a given trial block, quantization is performed on a single vertical subband of the appropriate wavelet transform.

IV. Preliminary Experimental Results

One author (VC) has participated in experiments measuring real, phase, and magnitude quantization thresholds for the chirp grating images. Results are shown in Fig. 6. The horizontal axis plots spatial frequency (for 5 levels of the wavelet transform). The vertical axis plots the RMS constrast

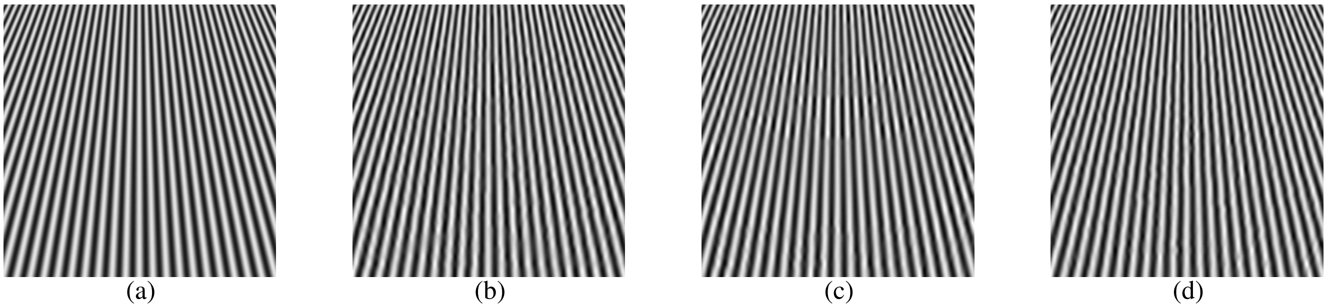

Fig. 5. *(a) Chirp grating image used for perceptual experiments. (b) Real wavelet quantization. (c) Complex magnitude quantization, producing local changes in brightness. (d) Complex phase quantization, producing local bending in grating lines. All quantized images have PSNR ≈ 30dB.*

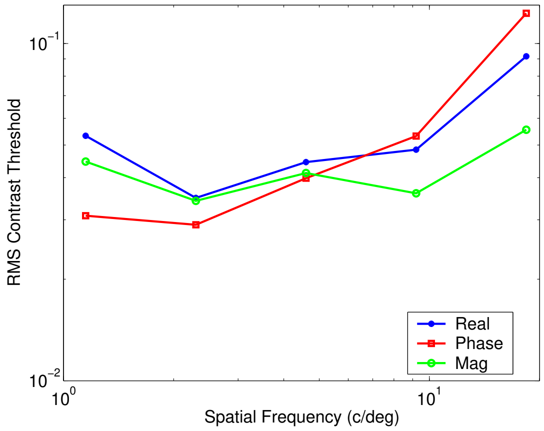

Fig. 6. *RMS contrast thresholds for real, phase, and magnitude quantization of chirp grating images (subject VC).*

threshold — higher values indicate lower sensitivity to the target distortion.

At low frequencies, we observe a greater sensitivity to phase quantization; this sensitivity decreases as spatial frequency is increased. At the highest spatial frequency, we observe that phase quantization has the least sensitivity, while magnitude quantization has the greatest sensitivity. In most cases, sensitivity to real quantization is near the lower sensitivity of phase or magnitude quantization. We note that the order of these differences is somewhat small compared to other common variations in RMS contrast threshold (see [5], for example).

These results serve as a preliminary indication that the human visual system may have a unique sensitivity for errors in phase and magnitude information, and that the relative importance may depend on the spatial frequency. Also, in some cases perception of errors due to real wavelet quantization may be due primarily to either the phase or the magnitude error component.

V. Conclusion

The development of a nonredundant complex wavelet transform has allowed for a novel investigation into the HVS. By focusing on complex magnitudes and phases, we have taken the first steps toward a study of the human perception of *What* and *Where* information. For a more thorough understanding of the relative importance of local phase and magnitude information, future work should extend the JND perceptual experiments using various images and sources of error. As discussed in Section III-B, it will be important to consider the multiscale coherency among complex wavelet coefficients. In addition, an intelligent image coder should exploit the low-dimensionality of magnitude and phase fields (when relevant), and the JND experiments should be updated as appropriate.

References

[1] R. M. Everson, A. K. Prashanth, M. Gabbay, B. W. Knight, L. Sirovich, and E. Kaplan, "Representation of spatial frequency and orientation in the visual cortex," *Proc. Natl. Acad. Sci.*, vol. 95, pp. 8334–8338, 1998.
[2] B. A. Olshausen and D. J. Field, "Sparse coding with an overcomplete basis set: A strategy employed by V1?" *Vision Res.*, vol. 37, pp. 3311–3325, 1997.
[3] M. Vetterli and J. Kovacevic, *Wavelets and Subband Coding*. Prentice Hall, 1995.
[4] A. B. Watson, G. Y. Yang, J. A. Solomon, and J. Villasenor, "Visibility of wavelet quantization noise," *IEEE Trans. Image Proc.*, vol. 6, no. 8, August 1997.
[5] D. M. Chandler and S. S. Hemami, "Effects of natural images on the detectability of simple and compound wavelet subband quantization distortions," *J. Opt. Soc. Am. A*, vol. 20, no. 7, July 2003.
[6] ——, "Contrast-based quantization and rate control for wavelet-coded images," in *IEEE Int. Conf. on Image Proc. — ICIP '02*, Rochester, New York, Oct. 2002.
[7] D. Marr, *Vision*. San Francisco: W. H. Freeman and Company, 1982.
[8] N. Kingsbury, "Image processing with complex wavelets," *Phil. Trans. R. Soc. Lond. A*, vol. 357, pp. 2543–2560, September 1999.
[9] I. W. Selesnick, "The design of approximate Hilbert transform pairs of wavelet bases," *IEEE Trans. Signal Proc.*, vol. 50, no. 5, May 2002.
[10] F. C. A. Fernandes, R. L. C. van Spaendonck, and C. S. Burrus, "A new framework for complex wavelet transforms," *IEEE Trans. Signal Proc.*, vol. 51, no. 7, July 2003.
[11] R. L. van Spaendonck, T. Blu, R. G. Baraniuk, and M. Vetterli, "Orthogonal Hilbert transform filter banks and wavelets," in *Proc., IEEE Int. Conf. Acoust., Speech, Signal Proc. — ICASSP '03*, Hong Kong, April 2003.
[12] F. C. A. Fernandes, M. B. Wakin, and R. G. Baraniuk, "Non-redundant, linear-phase, semi-orthogonal, directional complex wavelets," in *Proc., IEEE Int. Conf. Acoust., Speech, Signal Proc. — ICASSP '04*, 2004, submitted.
[13] M. T. Orchard and H. Ates, "Equiripple design of real and complex filter banks," Rice University, Tech. Rep., 2003.
[14] T. Bulow, "Hypercomplex spectral signal representations for the processing and analysis of images," Ph.D. dissertation, University of Kiel, 1999.
[15] A. V. Oppenheim and J. S. Lim, "The importance of phase in signals," *Proc. IEEE*, vol. 69, no. 5, 1981.
[16] N. Jayant, J. Johnston, and R. Safranek, "Signal compression based on models of human perception," *Proc. IEEE*, vol. 81, no. 10, 1993.
[17] D. H. Brainard, "The Psychophysics Toolbox," *Spatial Vision*, vol. 10, pp. 433–436, 1997.