Wavelets

Chris Perkins, Tobin Fricke
Department of Electrical Engineering
University of California at Berkeley

December 1, 2000

1 Introduction

Convential time- and frequency- domain analysis have met with enormous success in modern science and engineering. However, the scientific community is becoming increasingly aware that signals found in nature often exhibit complex, sometimes self-similar, even *fractal* characteristics. Wavelets are a broad class of new tools developed with the specific intent of being better able to analyse these properties of real signals.

2 Background

2.1 The Fourier Transform

In the early 1800's, Joseph Fourier discovered that nearly any function can be expressed as a (possibly infinite) sum of sines and cosines, or, equivalently, complex exponentials (equivalent because of Euler's identity: $e^{j\omega t} = j \sin \omega + \cos \omega$, where $j = \sqrt{-1}$).

A function $f(t)$ is traditionally represented by defining its value for each value of t in the domain of the function. In the terms of Linear Algebra, we can say that the function is represented as a linear combination (a weighted sum) of delta functions. The discrete-time (Krönecker) delta function is defined as follows:

$$\delta(x) = \begin{cases} 1 & \text{if } x = 0, \\ 0 & \text{if } x \neq 0. \end{cases} \tag{1}$$

Thus, for instance, we might define a simple sampled function implementing the mapping $\{0, 1, 2, 3\} \longmapsto \{2, 7, 1, 8\}$ as a linear combination of delta functions:

$$f(x) = 2 \cdot \delta(x-0) + 7 \cdot \delta(x-1) + 1 \cdot \delta(x-2) + 8 \cdot \delta(x-3) + \cdots \tag{2}$$

The Fourier Transform allows us to find a new representation for the function, given by the coefficients of complex exponentials of integer multiples of some fundamental frequency. In other words, the Fourier Transform is simply a change of basis from the basis of delta functions $\{\delta(x)\}$ to the basis of complex sinusoids $\{e^{j\omega n}\}$. The basis of delta functions is usually called the time (or space) domain, while the basis of complex exponentials is called the frequency domain. The continuous-time transform to the frequency domain is:

$$f(\omega) = F\{f(t)\} = \int_{-\infty}^{\infty} f(t) e^{-j\omega t} dt \tag{3}$$

To continue our example, we can use the Discrete Fourier Transform (not shown) to transform from the basis of delta functions to the basis of complex sinusoids. The resulting linear combination implements the same function, just via a different representation:

$$f(x) = \frac{9}{2} \cdot e^{0 \cdot j\pi n/2} + \frac{1+j}{4} \cdot e^{1 \cdot j\pi n/2} + (-3) \cdot e^{2 \cdot j\pi n/2} + \frac{1-j}{4} \cdot e^{3 \cdot j\pi n/2} \tag{4}$$

The frequency domain has one serious shortcoming, however: it has no time resolution. True, the frequency domain will tell us exactly what frequency components are present in a signal, but it tells us nothing about the locality of those frequency components in time. For practical purposes, the Fourier Transform is only useful on steady-state, or "stationary" signals.

2.2 A Bit of Progress: The Windowed Fourier Transform

The basis functions of the frequency domain extend in time from negative infinity to positive infinity, oscillating forever in both directions. We say that these functions have "global extent". In an attempt to add some locality of time to the Fourier Transform, we can choose a new basis, one of sines and cosines modulated by the Gaussian.

We now have a transform that takes a one-dimensional function and returns a two dimensional function: time on one axis, frequency on the other. This allows us to construct "spectrograms", showing the relative frequency content of a signal versus time.

While this can be extremely useful, it too has a serious shortcoming: the uncertainty principle. There is a fundamental trade-off between locality in the frequency domain and locality in the time domain. As we gain one, we lose the other. For instance, the frequency content localized at a single point in a function is completely meaningless.

3 The Wavelet Transform

The Wavelet transform was inspired by the idea that we could vary the scale of the basis functions instead of their frequency: a subtle yet powerful modification. In the words of Amara Graps, "The fundamental idea behind wavelets is to analyze according to scale." [1] Instead of representing a function as a sum of weighted delta functions (as in the time domain), or as a sum of weighted sinusoids (as in the frequency domain), we represent the function as a sum of time-shifted (translated) and scaled (dilated) representations of *some arbitrary function*, which as we shall see is called a *wavelet*. The power of this idea is that it examines the structure of information on all *scales*. The wavelet transform first compares the entire function to the wavelet, then compares smaller pieces of the function to the wavelet. This process is completed on successively smaller and smaller scales. This process forms a representation of the original function as a sum of wavelets of various scales and positions in time, acheiving a balance between locality in time and locality in frequency/scale.

In figure 1, it can be seen that when the frequency of the windowed fourier transform basis function is varied, the width of the pulse remains the same. When the frequency of the wavelet transform basis function is changed, however, the width of the wavelet also changes. For higher frequency wavelets, a narrower envelope results. This allows wavelet analysis to look at different frequency components on different time scales.

We will first discuss the selection of a wavelet $\psi(t)$ and then we shall introduce the continuous and discrete wavelet transforms.

3.1 Wavelets

The Wavelet transform introduces a degree of freedom not present in time or frequency domain analysis: we get to choose the wavelet that we're going to use in the analysis.

What qualities would we like to see in a wavelet? It turns out that there are some properties a wavelet *must* have, and others that are very desirable. The must-have properties are grouped together in what is called

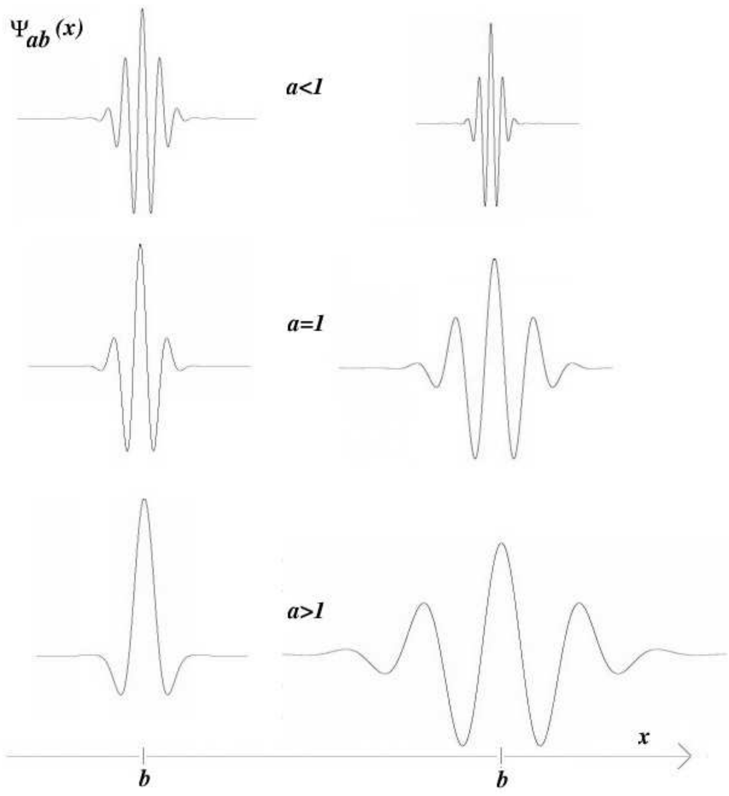

Figure 1: A comparison of the Windowed Fourier Transform (left) and Wavelet Transform (right) basis functions.

the admissibility criteria. The primary requirement is that the integral of the wavelet over all t be zero; it must spend equal time above and below the axis:

$$\int_{-\infty}^{\infty} \psi(t)dt = 0 \tag{5}$$

Furthermore, to be useful, a wavelet must have *local extent*. In other words, it must be localized in time (or space); it must be nonzero only for a finite interval. Furthermore, we prefer wavelet bases that are *orthonormal*. Two vectors are orthogonal if the projection of one on the other has zero length; a set of vectors is orthonormal if all pairs of vectors in the set are orthogonal and all vectors in the set are normal, eg, have length 1. Intuitively, if the basis functions are orthogonal, then the coefficients needed to represent a linear combination will all represent independent information. With an orthonormal wavelet basis, it is possible that more information will be compressed into fewer coefficients.

The first wavelet with these properties was discovered (or invented, depending on your *weltanschauung*) in 1910 by Alfred Haar [2], a Hungarian mathematician. The Haar wavelet is very simple; it just just a step function. Nonetheless it forms an orthonormal wavelet basis, and due to its simplicity and place in history it has also become the canonical example used in introducing wavelets.

$$\psi(x) = \begin{cases} 1 & \text{for } 0 \leq x < \frac{1}{2} \\ -1 & \text{for } \frac{1}{2} \leq x < 1 \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

The Haar wavelet is simple, but it is not ideal. It produces a basis set of predominantly flat functions, large numbers of which are required in linear superposition to approximate a non-flat function. A better, smoother wavelet is needed. Yet, we would like to preserve the orthonormality of the wavelet basis. We need a mother wavelet that is localized in time and in frequency, is reasonably smooth, *and* ideally produces orthonormal translations and dilations.

Only recently was such a wavelet found, by Ingrid Daubechies, now a professor at Princeton, in 1988 at AT&T Bell Laboratories. This function is now known as the Daubechies Wavelet after its discoverer and has become the canonical "real" wavelet.

To understand the Daubechies wavelet, we must first introduce the *scaling equation*. The scaling equation is a recursive expression that is used to generate wavelets, so in some sense it is a "grandmother wavelet". In the emphatic words of Gilbert Strang, "*All good wavelet calculations use recursion.*" The scaling equation is given as [3]:

$$\phi_j(t) = \sum c_k \phi_{j-1}(2x - k) \tag{7}$$

To use the scaling equation to generate a mother wavelet, we must specify several parameters. First, we must specify the the base case of the recursion, the function $\phi_0(t)$. Second, we must specify coefficients $\{c_0, c_1, \cdots, c_N\}$. For ϕ_0 we use the box function, which is zero everywhere except the interval $[0, 1]$:

$$\phi_0(t) = \begin{cases} 1 & \text{for } 0 \leq x \leq 1 \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

The summation is over k: one term in the sum for each of the coefficients. It remains to specify what these "magic" coefficients are. The special values found by Daubechies in 1998 are [3]:

$$\left\{ c_0 = \frac{1}{4}(1+\sqrt{3}), c_1 = \frac{1}{4}(3+\sqrt{3}), c_2 = \frac{1}{4}(3-\sqrt{3}), c_3 = \frac{1}{4}(1-\sqrt{3}) \right\} \quad (9)$$

Plugging these coefficients into the scaling equation and iterating, we get the Daubechies Wavlet (the Daubechies Wavlet is actually a *class* of wavelets; this is the member that is generated from four coefficients.). It should not come as a suprise that the Daubechies Wavelet, generated in this way, has fractal properties! Fractals arise in systems involving iteration or recursion, and the generation of the Daubechies Wavelet very much follows the recipe for generating fractals. In fact, the Daubechies Wavlet has an amazing property of self similarity: it actually contains a scaled down version of itself! And this scaled down version in turn contains a scaled down version, *ad infinitim.* [4]

It is very easy to write a short computer program to demonstrate this, as we did. The source code is attached in the appendics.

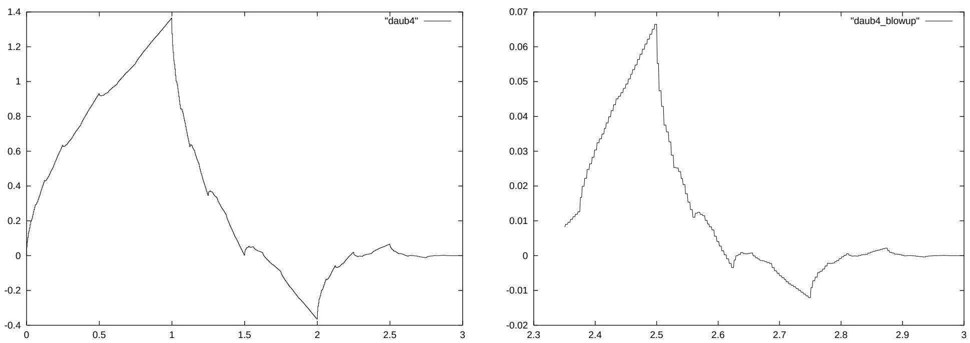

Figure 2: The Daubechies-4 wavelet. On the left, the complete Daubechies wavelet is shown. The self-similarity of this wavelet is easily seen when looking more closely at the region between 2.4 and 3.0, as shown on the right.

When choosing a wavelet to use, it is advantageous to use a wavelet that resembles the signal to be analyzed. If such a wavelet is chosen, the signal can be represented using significantly fewer coefficients than if a non-similar wavelet were used. In addition, choosing a wavelet that possesses self-similarity is also useful.

3.2 The Continuous Wavelet Transform

To use the Continuous Wavelet Transform (CWT), we first select a "mother wavelet" function, $\psi(t)$. This could be the Haar wavelet, the Daubechies wavelet, or any other wavelet. We then form the vector projections of the function under analysis $f(t)$ with dilations and translations of the mother wavelet:

$$W\{x(t)\} = \int_{-\infty}^{\infty} f(t)\psi_{\nu}^{\mu}(t)dt \qquad \forall \nu, \mu \tag{10}$$

where ψ_{ν}^{μ} represents a specific dilation (μ) and translation (ν) of the mother wavelet, $\psi(t)$:

$$\psi_{\nu}^{\mu}(t) = \sqrt{\mu}\psi\left(\frac{t-\nu}{\mu}\right) \tag{11}$$

Although the continuous wavelet transform is very useful for theoretical purposes, in practice most signals are sampled, forming discrete data series. For these series we must use the Discrete Wavelet Transform.

3.3 The Discrete Wavelet Transform

The result of the Continuous Wavelet Transform of a one-dimensional signal (say, a time series for example) is a two-dimensional signal, with independent axes of scale, and of time. Thus the CWT has introduced *redundancy* into the representation of the signal. In the discrete case, the mother wavelet is only scaled and dilated in discrete steps. "Dyadic" scalings are usually chosen: one translation on the largest scale, two translations on a scale half as large, four translations on a scale one quarter as large, etc. Thus the translation and dilation parameters in the wavelet transform become:

$$\begin{cases} \mu = 2^{-m} \\ \nu = n2^{-m} \end{cases} \tag{12}$$

For a data series of length $N = N_0 2^M$ samples, the transform is evaluated for the following scales and translations:

$$\begin{cases} m = 1, 2, ..., M \\ n = 0, 1, ..., N_0 2^{m-1} - 1 \end{cases} \tag{13}$$

Notice that m chooses the scale (dilation) of the wavelet, and n chooses the location (translation) of the wavelet. The value $m = 1$ corresponds to the largest, most general scale, corresponding to the general shape of $f(t)$

over all t. Thus, we only have to evaluate one translation of the wavelet. For $m = 2$, we've halved the scale of our analysis, so it is necessary to evalue two translations. Hopefully this conveys an intuitive view of the dyadic scaling.

Using these relationships, we can form a set of orthogonal basis functions from the mother wavelet $\psi(t)$ and the relationships given above as:

$$\psi_n^m(t) = 2^{m/2}\psi(2^m t - n) \tag{14}$$

We then perform the wavelet transform as given in the continuous case, but with the integral reduced to a summation over a finite number of terms [5].

4 Applications of Wavelet Analysis

Wavelet Analysis has found applications in numerous fields.

Perhaps the most widespread use of Wavelets is in data compression in general and in image compression in particular. The poster child of this application cannot be anything other than the FBI fingerprint database, and the success of Wavelet analysis can also be seen in its adoption as the compression technology to be used in the upcoming JPEG-2000 image compression standard. We will discuss image compression in some detail, and then summarize other applications of Wavelet technology.

4.1 Image Compression

The wavelet representation is extremely well suited for image compression because it does not require the image to be broken down into sub-blocks for processing. Furthermore, with Wavelet compression, an image can be transmitted as a data stream allowing progressive display of an image. The first coefficients contain information about the image on the broadest scale; successive coefficients contain information about successively finer details. A user will have a good idea of what an image looks like after only a few coefficeints have been received, and may stop the transmission at any time.

In Wavelet-based image compression, the source image is first transformed to the Wavelet domain using the Discrete Wavelet Transform. The source is already of finite size, discretely sampled (into pixels), and quantized into a discrete set of colors (say, 256 levels of gray, red, green, or blue). Thus, the transformation into the wavelet domain is only a change-of-basis; no information is gained nor lost, and the resulting data takes up exactly the same amount of space. We now quantize the data. To each wavelet coefficient, we allocate a number of bits proportional to that coefficient's importance in reconstruction. Low valued coefficients, by definition, contribute little to the reconstructed image. It is in this quantization step that some information is lost. Finally, the resulting quantized, Wavelet-domain representation of the image is compressed using conventional, loss-less techniques, usually Huffman compression. The wavelet transform reärranges the information in an image so that the variations on differing scales are grouped together. Usually the high-numbered coefficients, corresponding to variation on the smallest scales, are very small: very often zero. Conventional compression is highly effective in compressing these regions of little variation. [6]

4.2 FBI Fingerprint Database

The American FBI has been collecting fingerprint cards since 1924 and in this time has accumulated more than 200 million cards. The FBI digitizes these cards for electronic storage, at a resolution of 500 dots per inch with 256 levels of gray. A single fingerprint card results in about 10 megabytes of data; thus the entire collection would require 3000 terabytes, a truly huge amount of data. Furthermore, about 40,000 new cards are now being added *per day.* Add to this the needs to search for and retrieve cards from the database on a daily basis, and it's clear there's a problem.

The FBI investigated the use of JPEG as a compressed image format for fingerprint data, but after some consultations with their friends at the UK Home Office Police Research Group, the conclusion was reached that JPEG, and even custom JPEG variants, are unsuitable for compression of fingerprints at ratios greater than about 10:1. The blocking artifacts produced by JPEG at high levels of compression inhibited the interpretation of fingerprint data.

Eventually the FBI settled on a Wavelet based compression scheme developed at Los Alamos National Laboratory, called Wavelet Scalar Quantization. WSQ is a fairly typical wavelet image compression algorithm and has proven to be enormously successful in storing fingerprint information. With the adoption of WSQ by the FBI as a standard data format for fingerprint information, it has become a de-facto standard worldwide as well. [7] [1] [8]

4.3 JPEG 2000

With the upcoming JPEG 2000 standard set to use wavelets as the foundation of its image compression mechanism, wavelets will soon enter the homes of millions of people worldwide.

JPEG 2000 will offer many improvements over the current generation of JPEG compression. It will offer both lossy and lossless compression. It will fully exploit the desirable properties of the wavelet domain. Users will be able to view images at different resolutions corresponding to their needs and capabilities. If only a preview of an image is needed, the wavelet datastream can be terminated very early, but when a higher quality image is needed, more of the wavelet datastream can be decoded. [9]

4.4 Audio Compression

It will be interesting to see what fields Wavelet Analysis shall next infiltrate. MP3, a Fourier-based audio compression technology, coupled with relatively wide access to high-bandwidth internet connectivity, is poised to cause a complete reorganization of the music industry as well as American copyright law. Could wavelets be applied to audio for even greater compression or less degradation in quality?

Unfortunately one of the qualities that makes Wavelets so incredibly well suited for image compression is a drawback when audio compression is considered. When we view images, we take in the entire image at once. First we perceive general features of the image, and then we look at the finer details. The wavelet transform allows images to be displayed progressively in a way very suited to the human visual system. The more coefficients are received, the higher the quality of the resulting image. Audio, on the other hand, the human brain processes linearly. When listening to a song, we do not comprehend the entire song in its entirety all at once; we comprehend the song moment-by-moment, and how it develops over time is important. Thus the wavelet approach – where the "general idea" of the song would be available first, followed by the details, would not be appropriate for progressive, real-time transmission of audio, for example over the internet. We would have to wait for a sufficient number of coefficentsto be received before starting playback.

This practical detail aside, we speculate that audio is well suited to wavelet compression. Music traditionally exhibits a high degree of self-similarity, and, moreover, self-similarity on a multitude of scales. These features are exactly those that the wavelet transform exploits, so it would be interesting to examine the application of wavelet analysis to audio in another project at another time.

4.5 Other Applications

Wavelet analysis has permeated nearly every region of modern science. In any field where data needs to be analyzed, wavelet analysis is often able to give insights that more traditional data analysis techniques cannot provide. In the field of astronomy and cosmology, wavelet analysis is useful in analyzing stellar and solar data and has provided information about the sun and other stellar processes that methods such as Fourier analysis have not been able to provide. Wavelet analysis is also used in these fields to compress telescopic image data (such as that from the Hubble Space Telescope). Wavelet analysis is also useful in analyzing seismic data. Algorithms are needed to determine when an earthquake begins and to study various properties of the motion. Wavelet analysis is a natural candidate to analyze such data. Wavelet analysis has also provided new insights in the study of turbulence. Wavelets are also used extensively in computer graphics to perform edge-detection, texture analysis, and denoising of image data. [10]

Wavelet analysis has also been incorporated into so-called "Fractal Modulation", a new scheme for modulating data for transmission over a wire. Officially known as DWMT, or Discrete Wavelet Multi-Tone Modulation, and developed by Aware, Inc., this technology is one of several being investigated for use in next-generation Digital Subscriber Line technology. Digital Subscriber Line (DSL) is a means of sending and receiving high-speed digital data to and from residences over the same copper wires that are used for regular telephone service – indeed, the system operates *simultaneously* with regular telephone service. With greater and greater needs for high-bandwidth internet connectivity at home, being able to squeeze every last bit per second out of existing infrastructure is also increasing in importance, and Wavelets may have a significant influence in this field. [11]

One reason that wavelet analysis is so useful in such a wide range of physical applications is that physical signals usually have different frequency components on different spatial scales. There are usually very short, high frequency signals mixed with longer, low frequency signals. Wavelet analysis lends itself perfectly to such signals because of its ability to look at different spatial regions using different scales.

5 An Experiment In Image Compression

5.1 Motivation

There is a vast amount of literature available on the subject of wavelets, however most consists of mathematical text with little information on practical application of wavelet technology. Outrageous claims are made regarding the capabilities of wavelet-based techniques, yet it is difficult to find elegant, compelling examples of the success of wavelet analysis.

We set out to verify whether or not wavelets could do everything they promised, specifically in the area of image compression.

5.2 Procedure

We first selected a test image. We could not resist the temptation to use the professor's portrait as it appears on his webpage, instead of a more traditional IEEE test image of some sort. We first converted this image from JPEG format to the much-simpler Portable Greymap (PGM) image file format, which is essentially a list of pixel values in plain, human-readable text. To do this we used a free software program developed at UC Berkeley's eXperimental Computing Facility (the XCF), the GNU Image Manipulation Program (more affectionately known as The GIMP). The next step was to implement the two-dimensional Discrete Wavelet Transform. We were pleased to find an implementation available in the popular reference, *Numerical Recipes In C* [12]. We wrote a C program that reads in the PGM format image of the professor into a row-major linear array. Next, we use the Numerical Recipes routines to compute the discrete wavelet transform along two dimensions. Now, we make a copy of the resulting array, and we use the C standard library routine $qsort()$ to sort it. This allows us to find the value of the nth highest coefficient. We then set to zero some percentage of the lowest-valued coefficients. Finally, we take the inverse discrete wavelet transform resulting in a regular image once again. The program outputs the pixel values of the image to a file, which we then load in Matlab and compare to the original. To summarize, the procedure is as follows:

1. Load the source image data from a file into an array
2. Compute the Discrete Wavelet Transform of the data
3. Remove (set to zero) all coefficients whose value is below a threshold. (This is the compression step.)

4. Reconstruct the image by computing the Inverse Discrete Wavelet Transform
5. Compare the resulting reconstruction of the compressed image to the original image

The C source code to this software is listed in the appendix.

5.3 Results

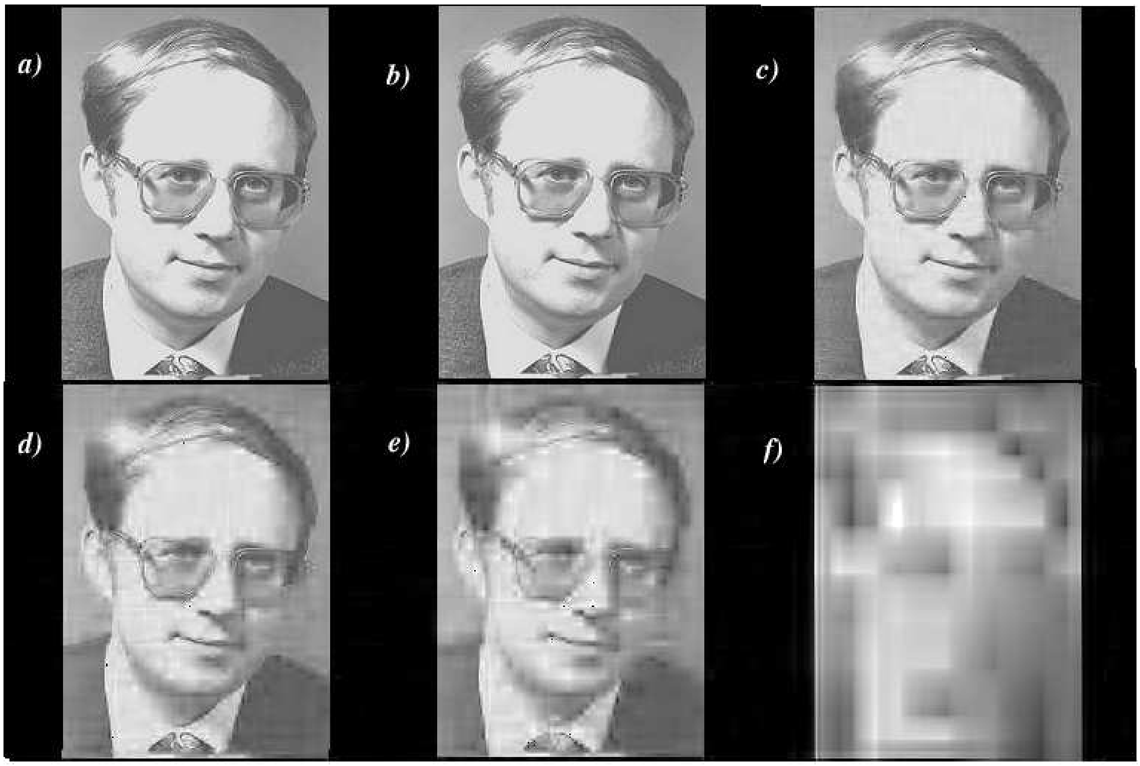

Figure 3: Various levels of wavelet compression. Picture a) shows the original image. Picture b) has had 45% of the wavelet coefficients removed. Picture c) has had 95% removed. Picture d) has had 98% removed. Picture e) has had 99% removed. Picture f) has had 99.9% of the coefficients removed.

For comparison purposes, we compressed the original image using JPEG, selecting a very high compression ratio. JPEG is the current de-facto standard for compressing photorealistic images (as opposed to diagrams and line drawings) for distribution over the internet. The resulting image, at an approximate compression ratio of 10:1, appears very "blocky". JPEG operates by first separating an image into 8x8 pixel blocks and then performing the discrete cosine transformation (DCT), a form of the Fourier transform, on each 8x8 pixel block. The resulting coefficients are truncated and stored in the resulting JPEG file. For higher compression ratios, fewer coefficients are stored, and/or the coefficients are stored at a lower numeric precision. In the limiting case where only one coefficient is stored, the only information that is available is the *average value* of the pixels in each block; each 8x8 pixel block assumes this color, hence the "blocky" look to the image.

Figure 4: A comparison of images using 10:1 JPEG compression (left) and approximately 100:1 wavelet compression using the Daubechies-4 wavelet (right).

On the other hand, the wavelet transform operates on the entire image at once, so there are never any block-like artifacts. We were astounded to find that we could remove (set to zero) 98% of the wavelet coefficients and still be able to reconstruct a decent-looking image. This translates to a compression ratio of nearly 100:1, with visual quality far better than JPEG compression at a ratio of only 10:1. The results using various wavelet compression ratios can be seen in figure 3.

Another interesting comparison is the location of error in the JPEG and wavelet compressed images. In images compressed using JPEG, the error usually is most visually apparent in areas of high spatial frequencies. In contrast, the wavelet compressed image tends to smooth out those areas of an image that are lacking details, or are already smooth. As can be seen in figure 3, even with a compression ratio of approximately 100:1, high resolution regions can still be seen, such as the areas around the glasses, mouth, necktie, and hair, while lower frequency regions, such as the suit and other facial regions have been smoothed out. This mix of sharp and smooth regions demonstrates how the wavelet transform is able to analyze different spatial regions using different scales.

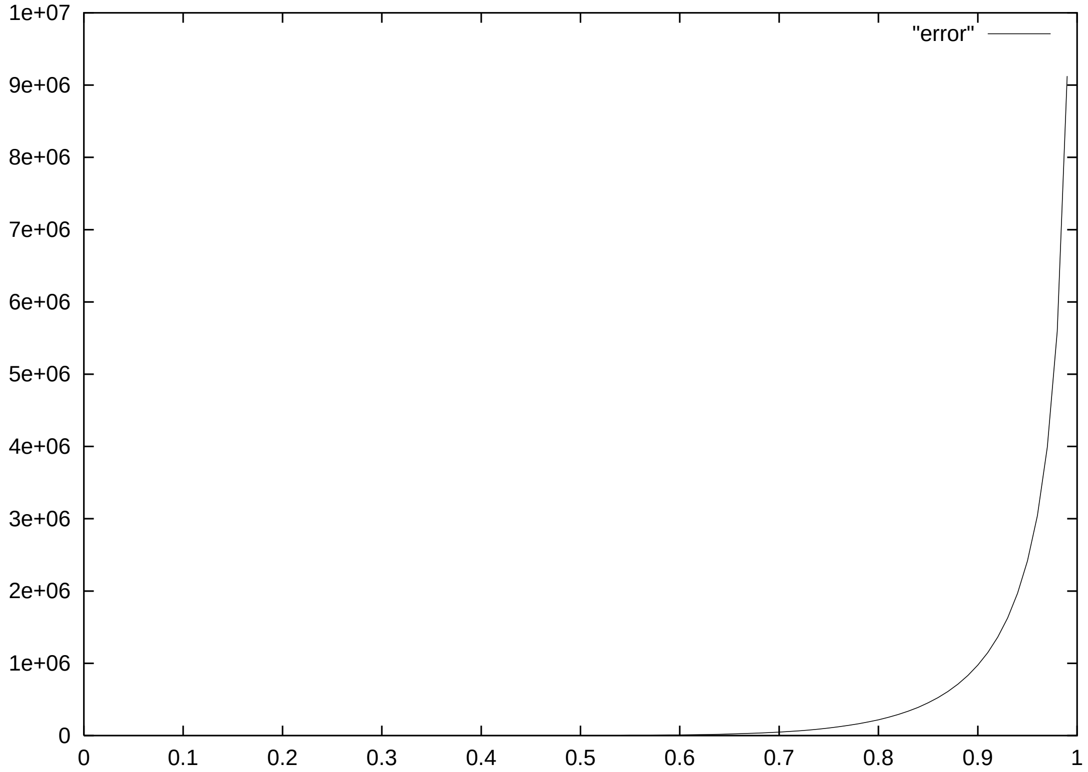

Figure 5: Mean square error as a function of wavelet coefficients removed.

A plot of the mean square error as a function of the percentage of wavelet coefficients removed can be seen in figure 5. Very little error is introduced in removing wavelet coefficients until a certain point is reached, in this case when around 70% of the coefficients are removed, at which time the increase in error becomes exponential. The shape of this curve supports the qualitative claim that a large number of wavelet coefficients can be removed while still rendering an accurate image.

We were interested to find that in the wavelet domain, the vast majority of coefficients are very small. Indeed, the highest coefficient in our test case exceeded the lowest coefficient by five orders of magnitude!

6 Conclusion

Wavelet analysis is able to transform a time- or space- domain signal to a new basis, the wavelet domain, that takes advantage of the character of real signals. Most real-life data cannot be concisely modeled using only time series or fourier series, but the wavelet transform has been remarkably successful in exploiting the fractal characteristics of these signals. Because of this, wavelet-based compression is phenominally effective, and as time goes on will become more and more pervasive. Furthermore, knowledge gained through wavelet analysis in diverse fields may lead to significant new knowledge of the world. Wavelet theory is maturing rapidly and is becoming more and more widely known. This will likely lead to even more applications of wavelet analysis in years to come.

References

[1] Amara Graps. "An Introduction to Wavelets". *IEEE Computational Science and Engineering*, **vol. 2, num. 2** (Summer 1995).

[2] Gilbert Strang. "Wavelets". *American Scientist*, **82** (1994) 250-255.

[3] Gilbert Strang. "Wavelets and Dilation Equations: A Brief Introduction". *Siam Review*, **31** (1989) 613-627.

[4] Ingrid Daubechies. *"Ten Lectures On Wavelets"*. Society for Industrial and Applied Mathematics, (1992).

[5] Gregory W. Wornell. *"Signal Processing With Fractals: A Wavelet-Based Approach"*. Prentice Hall, (1996).

[6] Subhasis Saha. "Image Compression - from DCT to Wavelets". *ACM Crossroads*. `http://www.acm.org/crossroads/xrds6-3/sahaimgcoding.html`.

[7] Chris Brislawn. The FBI Fingerprint Image Compression Standard. `http://www.c3.lanl.gov/~brislawn/FBI/FBI.html`.

[8] Chris Brislawn. Fingerprints Go Digital, **vol. 42 no. 11** (Nov. 1995) 1278-1283.

[9] Michael J. Gormish. "JPEG 2000: Worth the Wait?".

[10] J.C. van den Berg. *"Wavelets in Physics"*. Cambridge University Press, (1999).

[11] Stuart D. Sandberg and Michael A. Tzannes. "Overlapped Discrete Multitone Modulations for High Speed Copper Wire Communications". *Aware, Inc.*, (1995).

[12] William H. Press *et al.* *"Numerical Recipes in C"*. Cambridge University Press, (1992).