

Frequently Asked Questions on Wavelets

Naoki Saito
Department of Mathematics
University of California
Davis, CA 95616 USA
email:saito@math.ucdavis.edu

April 4, 2004

- Q1. What is the wavelet transform?
 - Q2. What is the wavelet packet transform?
 - Q3. What is the local cosine/sine transform?
 - Q4. What are the applications of wavelets?
 - Q5. Where can I get the further information on wavelets?
 - Q6. Which books should I get?
- Appendix: “Wavelet Applications Come to the Fore” by Barry Cipra.

Q1. What is the wavelet transform?

The wavelet transform (in the signal processing context) is a method to decompose an input signal of interest into a set of elementary waveforms, called “wavelets,” and provides a way to analyze the signal by examining the coefficients (or weights) of these wavelets. The Fourier transform also decomposes signals into elementary waveforms; however, in this case, they are global oscillations, i.e., sines and cosines. Thus, when one wants to analyze the local property of the input signals such as edges or transients, the Fourier transform is not an efficient analysis tool.

Operationally, the wavelet transform essentially partitions the frequency axis smoothly and recursively and analyzes each segment (i.e., frequency band) with a resolution matched to its scale. An input signal is first decomposed into low and high frequency bands by convolution-subsampling operations with a pair consisting of a “lowpass” filter $\{h_k\}$ and a “highpass” filter $\{g_k\}$ directly on the discrete time domain. Let H and G be the convolution-subsampling operators using these filters. These are called quadrature mirror filters (QMFs) if they satisfy the following orthogonality (or perfect reconstruction) conditions: $HG^* = GH^* = 0$ and $H^*H + G^*G = I$, where H^* and G^* are the adjoint (i.e., upsampling-anticonvolution) operators and I is the identity operator. Various design criteria (concerning regularity, symmetry etc.) on the lowpass filter coefficients $\{h_k\}$ can be found in [4]. In the wavelet transform, this decomposition (also known as expansion or analysis) process is iterated only on the low frequency bands and each time the high frequency coefficients are retained intact.

Each high frequency subband is spanned by a set of translated versions of a single elementary waveform with a specific scaling parameter. This most elementary waveform is called “mother wavelet.” Each low

frequency subband (subband containing the DC component) is spanned by a set of translated versions of another single elementary waveform called “father wavelet” (or “scaling function”).

The frequency characteristics of the father wavelet essentially determines the quality of the wavelet transform. If we were to partition the frequency axis sharply using the characteristic functions (or boxcar functions), then we would end up with the so-called Shannon (or Littlewood-Paley) wavelets, i.e., the difference of two sinc functions, the corresponding father wavelet is simply the sinc function. Clearly, however, we cannot have a finite-length filter in the time domain in this case. The other extreme is the Haar wavelet which is the simplest and oldest wavelet discovered in 1910. In this case, the father wavelet is just a boxcar function, and the mother wavelet is defined as

$$g(x) = \begin{cases} 1 & \text{for } 0 < x < 1/2, \\ -1 & \text{for } 1/2 \leq x < 1, \\ 0 & \text{otherwise.} \end{cases}$$

This is simply a difference of the two boxcar functions in time. Although the Haar wavelet gives rise to the shortest filter length in the time domain, it partitions the frequency axis quite badly.

The reconstruction (or synthesis) process is also very simple: starting from the lowest frequency components (or coarsest scale coefficients) and the second lowest frequency components, the adjoint operations H^* and G^* are applied respectively and added to obtain the next finer scale coefficients. This process is iterated to reconstruct the original signal. The computational complexity of the decomposition and reconstruction process is in both cases $O(N)$ where N is a number of time samples.

Q2. What is the wavelet packet transform?

For oscillating signals such as acoustic signals, the analysis by the wavelet transform is sometimes inefficient because it only partitions the frequency axis finely toward the low frequency. The wavelet packet transform decomposes even the high frequency bands which are kept intact in the wavelet transform. If we repeat the decomposition process J times, we end up having JN expansion coefficients. Clearly, we have a redundant set of expansion coefficients, in fact, there are more than $2^{2^{(J-1)}}$ possible orthonormal bases. One can choose the most suitable basis out of this large collection of bases. See e.g., [4], [12], [9], [20], The computational cost for both decomposition and reconstruction is $O(N \log N)$.

Q3. What is a local cosine/sine transform?

The local cosine/sine transform is a conjugate or dual of the wavelet packet transform: it partitions the time axis smoothly and recursively, and then does the frequency analysis (the discrete cosine/sine transform) on each time interval. In fact, it is possible to partition the real-line into any disjoint intervals smoothly and construct orthonormal bases on each interval [12], [9], [20]. One of the natural ways of splitting the time interval is to split it into two half intervals of equal length smoothly by a specific tapering (or bell) function, and iterate this split recursively. This process again creates a set of redundant coefficients. One can select the most suitable set of time intervals for one’s need. The space of functions living on each interval is spanned by a set of smooth sinusoidal functions localized around the same time interval. The computational requirement in this transform is at most $O(N(\log N)^2)$.

Q4. What are the applications of wavelets?

Wavelets and their relatives generated a lot of interests in diverse fields ranging from astronomy to geology to biology as well as statistics and computer science. In each of these fields, the wavelets are applied for

- Data Compression
- Noise Removal
- Feature Extraction, Classification, and Regression
- Fast Numerical Analysis

Here is a list of interesting applications which drew my attention:

- FBI's fingerprint image compression
- Enhancement of mammography images
- Analysis of clustering of galaxies

One of the most successful applications so far, however, is data compression. In fact, the new image compression standard called JPEG2000 is fully based on wavelets. See the JPEG (Joint Photographic Experts Group) official website listed below for the details. I would strongly encourage the reader to read an excellent and somewhat advanced review article of Donoho, Vetterli, DeVore, and Daubechies [5] that explains a *deep* relationship between data compression and harmonic analysis.

It is beyond the scope of this FAQ to describe each application. Instead, I would like to refer the reader to the SIAM News article written by Barry Cipra attached in Appendix.

Q5. Where can I get the further information on wavelets?

From WWW

There are many wavelet information one can get from the Web. The following four are particularly useful:

* Amara's Wavelet Page:

<http://www.amara.com/current/wavelet.html>

This page gives also an introductory explanation of the wavelet theory and provides a large amount of links to various wavelet-related home pages.

* The Wavelet Digest:

<http://www.wavelet.org/>

You can get information on wavelets discussed over the moderated news group called "The Wavelet Digest" over the net.

* Wavelet Resources:

<http://www.mathsoft.com/wavelets.html>

This page contains many useful pointers to research papers and software.

* JPEG Official Website:

<http://www.jpeg.org/>

This page provides wealth of information about both the JPEG standard (based on DCT-II) as well as the JPEG2000 standard (based on wavelets) in details.

* Fast Mathematical Algorithms & Hardware Corp.:

<http://www.fmah.com>

FMAH is a company created by Professor Raphy Coifman and Professor Vladimir Rokhlin of Yale University. This company's homepage contains an array of very interesting and practical examples of the use of the wavelets and their relatives.

* WaveLab from Stanford:

<http://www-stat.stanford.edu/~wavelab/>

Professor Dave Donoho and his group are distributing the "WaveLab" package which runs under "matlab." So, if you are a matlab user, you might want to try this package. I highly recommend this to anyone who wants to know "what wavelets really can do for me?" Also, check out his report page

<http://www-stat.stanford.edu/~donoho/Reports/>, which contains very interesting and important papers.

From Journals

There are two journals dedicated to the wavelet-related research:

* Applied and Computational Harmonic Analysis, Academic Press.

* Journal of the Fourier Analysis and Its Applications, Birkhäuser

One often encounters the wavelet-related work in the following journals (just a subset):

* IEEE Trans. on Signal Processing,

* IEEE Trans. on Image Processing,

* IEEE Trans. on Pattern Analysis and Machine Intelligence,

* IEEE Trans. on Information Theory,

* SIAM Journ. on Mathematical Analysis,

* Signal Processing

Q6. Which books should I get?

Here is my favorite list of books on wavelets: [9], [12], [20]. For more advanced topics and details, I would recommend: [4] and [16]. The following are not books, but I highly recommend to read them to understand this field: [5], [11], [13], [15], [18].

For those who want to dig into more mathematical aspects of wavelets (in real and functional analysis setting), I would recommend [21], then more advanced books of Meyer (and Coifman) [14], [17]. Most of the books listed emphasize the orthonormal wavelets and the discrete wavelet transforms. For those who are interested in data interpretation and analysis such as geophysicists and geologists, the continuous wavelet

transforms are also popular tools. The book of Holschneider [8] emphasizes these tools as well as the group theoretical interpretation of wavelet transforms.

Finally, by studying the wavelets, one may want to learn more about the classical harmonic analysis (Fourier analysis) and its applications. In fact, by studying the wavelets, one may gain a fresher view toward the classical harmonic analysis. There are many books on this subject, but I found the following extremely useful. The books of Dym and McKean [6] and Körner [10] are jewels. In fact, I would like to see the same type of books for emphasizing wavelet aspects. The book of Folland [7] is a very good text book, in particular, describes the nice interaction between the Fourier analysis and the Sturm-Liouville theory. Out of the more engineering oriented books, Bracewell [2] is a classic and contains a highly useful pictorial dictionary of the functions and their Fourier transforms. For the discrete Fourier transforms, I found [3] very useful. The more mathematically oriented readers may want to consult [1], [19], [22].

References

- [1] J. M. Ash (ed.), *Studies in Harmonic Analysis*, MAA Studies in Mathematics, vol. 13, Math. Assoc. Amer., 1976.
- [2] R. N. Bracewell, *The Fourier Transform and Its Applications*, second, revised ed., McGraw-Hill, Inc., 1986.
- [3] W. L. Briggs and V. E. Henson, *The DFT: An Owner's Manual for the Discrete Fourier Transform*, SIAM, Philadelphia, PA, 1995.
- [4] I. Daubechies, *Ten Lectures on Wavelets*, CBMS-NSF Regional Conference Series in Applied Mathematics, vol. 61, SIAM, Philadelphia, PA, 1992.
- [5] D. L. Donoho, M. Vetterli, R. A. DeVore, and I. Daubechies, *Data compression and harmonic analysis*, IEEE Trans. Inform. Theory **44** (1998), no. 6, 2435–2476, Invited paper.
- [6] H. Dym and H. P. McKean, *Fourier Series and Integrals*, Academic Press, 1972.
- [7] G. B. Folland, *Fourier Analysis and Its Applications*, Wadsworth & Brooks/Cole, 1992.
- [8] M. Holschneider, *Wavelets: An Analysis Tool*, Oxford Univ. Press, 1995.
- [9] S. Jaffard, Y. Meyer, and R. D. Ryan, *Wavelets: Tools for Science & Technology*, SIAM, Philadelphia, PA, 2001.
- [10] T. W. Körner, *Fourier Analysis*, Cambridge Univ. Press, 1988.
- [11] S. Mallat, *Multifrequency channel decompositions of images and wavelet models*, IEEE Trans. Acoust., Speech, Signal Process. **37** (1989), no. 12, 2091–2110.
- [12] ———, *A Wavelet Tour of Signal Processing*, second ed., Academic Press, San Diego, CA, 1999.
- [13] Y. Meyer, *Book review of “An Introduction to Wavelets” by C. K. Chui, Academic Press, NY, 1992, and “Ten Lectures on Wavelets” by I. Daubechies, SIAM, 1992*, Bull. Amer. Math. Soc. **28** (1993), no. 2, 350–360.

- [14] ———, *Wavelets and Operators*, Cambridge Studies in Advanced Mathematics, vol. 37, Cambridge Univ. Press, New York, 1993, Translated by D. H. Salinger.
- [15] ———, *Wavelets: Their past and their future*, Progress in Wavelet Analysis and Applications (Y. Meyer and S. Roques, eds.), Editions Frontieres, B.P.33, 91192 Gif-sur-Yvette Cedex, France, 1993, pp. 9–18.
- [16] ———, *Oscillating patterns in image processing and nonlinear evolution equations*, University Lecture Series, vol. 22, AMS, Providence, RI, 2001.
- [17] Y. Meyer and R. Coifman, *Wavelets: Calderón-Zygmund and multilinear operators*, Cambridge Studies in Advanced Mathematics, vol. 48, Cambridge University Press, New York, 1997.
- [18] O. Rioul and M. Vetterli, *Wavelets and signal processing*, IEEE SP Magazine **8** (1991), no. 4, 14–38.
- [19] E. M. Stein and G. Weiss, *Introduction to Fourier Analysis on Euclidean Spaces*, Princeton Univ. Press, 1971.
- [20] M. V. Wickerhauser, *Adapted Wavelet Analysis from Theory to Software*, A K Peters, Ltd., Wellesley, MA, 1994, with diskette.
- [21] P. Wojtaszczyk, *A Mathematical Introduction to Wavelets*, London Mathematical Society Student Texts, vol. 37, Cambridge Univ. Press, 1997.
- [22] A. Zygmund, *Trigonometric Series*, third ed., Cambridge Mathematical Library, Cambridge Univ. Press, 2003, Volumes I & II combined.

Appendix: Barry Cipra's Article

=====

Reprinted from SIAM NEWS
Volume 26-7, November 1993
(C) 1993 by Society for Industrial and Applied Mathematics
All rights reserved.

=====

Wavelet Applications Come to the Fore

By Barry A. Cipra

“Stick your finger in the gizmo,” the weary patrolman orders the besotted occupant of his squad car. There’s a gentle electric hum as the Digitron 6000 scans the ridges and whorls on the drunk driver’s right index finger. Moments later the report from a database hundreds of miles away lights up the screen of the squad car’s computer. The message startles the officer back into complete wakefulness. Seconds later he’s speeding jailward, the drunk in the back seat now under arrest on a charge of murder.

That kind of instant electronic police work is still a ways off. But the technology that will make it possible is coming into place. In a key step forward, the Federal Bureau of Investigation recently adopted a wavelet-based standard for digital fingerprint image compression. At compression ratios of about 20:1, the new standard will facilitate the rapid transmission of information that is crucial for effective police work.

The FBI’s fingerprint standard is just one way that wavelets are making an impact. In less than ten years, the new analytic tool has found a litany of applications. Wavelet theory is rapidly becoming the foundation of applied analysis. It is being used for all kinds of signal processing, from image compression to sound enhancement to statistical analysis.

“In the beginning people tried wavelets on everything in sight,” recalls Ingrid Daubechies, one of the key contributors to the theory. “We know by now on what wavelets work and on what they don’t work, better than we did a couple of years ago. And I think there are some very nice success stories.”

New Bases for Analysis

Loosely—very loosely—speaking, wavelet theory can be described as abstract harmonic analysis made easy, or at least intelligible to engineers. The basic idea is to represent a given function—say the electrical activity of an individual’s brain—as a combination of “basis” functions belonging to a specified set, whose analytic properties are readily accessible. With luck, the given function is well approximated by just a few basis functions. In that case, it is easy to work with the function. In particular, its description can be reduced from a painstaking point-by-point report to a handful of nonzero coefficients.

The classic tools for analyzing functions are Fourier series and integrals, which use the trigonometric sine and cosine functions (or, to be fancy, the complex exponential, $e^{i2\pi x}$), to represent arbitrary functions. Fourier analysis is marvelously well suited for a huge number and variety of problems, and it provides the theoretical framework for virtually all analytic techniques, including wavelets. But sines and cosines are not ideal for every problem.

Wavelet theory takes a more eclectic approach. Almost any function can serve as the starting point for a system of wavelets. Other wavelets in the system are constructed by simply translating the given “mother” wavelet and either dilating or contracting it. In this way, wavelets are able to localize behavior in both “time” (via translation) and “frequency” (via dilation/contraction).

Daubechies likens a wavelet transform to a musical score, which tells the musician which note to play at what time. One of the attractive features of wavelets is their “zoom in” property: They are designed to deal with fine details that affect only part of an image or signal—something that always leaves Fourier analysis with a case of the jitters. That is in part why wavelets seem to be so good at data compression for things like fingerprint images.

Megabytes of Ink

One problem with fingerprints is simply that there are so many of them. The FBI has approximately 200 million fingerprint cards. Some come from employment and security checks, but 114.5 million cards belong to some 29 million criminals (bad guys tend to get fingerprinted more than once). According to Peter Higgins, deputy assistant director of the FBI’s Criminal Justice Information Services division, the files occupy an acre of office space.

That’s a lot of black ink.

By digitizing the files and storing them electronically, “we hope to put [them] in something that would fit in a 20 x 20 - foot room,” Higgins says. Putting things in electronic form should speed up the submission process. (Currently, the FBI receives anywhere from 30,000 to 40,000 fingerprint identification requests every day, mostly through the mail. Approximately half pertain to criminal arrests and half to employment checks.) It’s also hard to imagine doing automated fingerprint identification any other way.

But the digitized images have to be of high quality. Faxed fingerprints may be OK for post office reproductions, but not for permanent records. And that’s another problem: At a resolution of 500 pixels per inch with 256 levels of gray-scale, a single inch-and-a-half-square fingerprint block takes up approximately 600 kilobytes, and an entire card weighs in at a hefty 10 megabytes of data. Multiplied by 200 million cards, that’s—well, quite a bit. Moreover, 10 megabytes is a nontrivial amount of data to transmit. At a standard modem rate of 9600 bits per second with 20tie up a phone line for nearly three hours.

That’s where the wavelet compression standard comes in. Developed by Tom Hopper at the FBI and Jonathan Bradley and Christopher Brislawn at Los Alamos National Laboratory, the standard is based on scalar quantization of a 64-subband discrete wavelet transform. Compression takes place in the quantization step, where the coefficients of the transform within each subband are, in effect, assigned to integer-valued “bins.” (The information is further compressed by Huffman coding, which uses strings of variable length to represent data.)

The wavelet/scalar quantization standard is not locked in to a particular wavelet basis. A digitized fingerprint file will contain not only the compressed image, but also tables specifying the wavelet transform, scalar quantizer, and Huffman code. “We allow for a number of different encoders,” Brislawn explains. So far, however, only one system, using a basis of bi-orthogonal wavelets constructed by Cohen, Daubechies, and Feauveau and reported in a 1990 paper, has been approved by the FBI. The system gives reconstructions that are hard to distinguish from originals at compression ratios of about 20:1.

Wavelets, PDEs, and Brahms

Image compression is just one of the uses that have been found for wavelets. Among the many commercial possibilities, researchers at Aware Inc. in Cambridge, Massachusetts, and Fast Mathematical Algorithms & Hardware Corp. in Hamden, Connecticut, are exploring a variety of applications, including video and speech compression. In 1990, Aware introduced the first wavelet transform processor chip. The company recently licensed the chip to ImMix, a start-up company based in Grass Valley, California, for use in VideoCube, a new video editing workstation.

Not all the applications under consideration have immediate market potential. At Aware, for example, John Weiss and Sam Qian have been investigating the use of wavelets in the numerical solution of partial differential equations. “One of the things that they seem to be very useful for are situations where you have strong gradients,” Weiss explains. That includes problems involving shock waves and turbulence, he adds.

“What we’re trying to do is see if you can make the wavelet method as universal as, say, the finite element method but obtain a better rate of convergence.” Weiss, who co-organized a minisymposium on wavelet solutions to PDEs at the SIAM annual meeting last July, envisions the incorporation of wavelet-based solvers into a CAD/CAM package for design engineers who don’t have access to supercomputers: “If you can offer something which is reliable and fairly simple to use, I think there is a market for that.”

At Los Alamos, Bradley and Brislawn are also applying wavelet technology to the solutions of partial differential equations, but in a completely different way: They are using wavelets not to solve PDEs, but to help manage the volumes of data that supercomputers spew out when they run through a global climate or ocean simulation. A typical simulation generates on the order of a terabyte of uncompressed data. Much like their wavelet/scalar quantization approach to fingerprint image compression, Bradley and Brislawn have developed a wavelet/vector quantization method for the multidimensional data sets of climate and ocean models. In this case the purpose is to give researchers a rough and ready look at what the supercomputer is trying to tell them.

In one of the more pleasing applications of the new mathematics, Ronald Coifman and colleagues at Yale University used wavelets to clean up an old recording of Brahms playing his First Hungarian Dance on the piano. The “original” recording—actually a re-recording of a radio broadcast (complete with static) of a 78 record copied from a partially melted wax cylinder—was unrecognizable as music. But using a technique he calls adapted waveform analysis, Coifman, who is a co-founder of Fast Mathematical Algorithms & Hardware, managed to strip out much of the noise. The remaining sound is good enough to give a sense of Brahms’s style of play.

Cleaning Up Statistics

Music is not the only domain in which noise is a problem. Statisticians have long grappled with the problem of noisy data. It appears that wavelets may hold some of the answers. At least that’s the view expressed by David Donoho, a statistician at Stanford University who has led the way in applying wavelet techniques in the theory of statistics. He and colleague Iain Johnstone have developed a “wavelet shrinkage” technique that works wonders on a variety of data sets.

The technique starts with the application of a wavelet transform to the noisy signal or data set. It then “shrinks” each of the wavelet coefficients toward zero, using a soft-threshold nonlinearity, so that suitably small coefficients are set precisely to zero. Finally, the altered coefficients are inverted to produce a “denoised” signal. Donoho and Johnstone, with co-authors Gerard Kerkycharian and Dominique Picard of the Universite de Paris-VII, have shown that denoising by wavelet shrinkage is either optimal or nearly so for a number of technical criteria. For example, Donoho has proved that the reconstructed signal is, with high probability, at least as smooth as the original (true) signal, for a wide variety of smoothness measures.

Whether wavelets will have an impact in specific areas of applied statistics—say clinical trials in medical research—remains to be seen. But there’s no question they’ve changed the landscape of theoretical statistics, says Donoho: “As soon as we were exposed to wavelets, we made the equivalent of about ten years’ progress in months.” One change is likely to be in the questions that are asked about such problems as smoothness. “Part of the reason we were looking at these problems was that we didn’t know how to do them,” Donoho says. “Now that we know how to do them, I think we’re in a better position to say what are the right

questions for statistical theory to focus on.”

(Barry A. Cipra is a mathematician and writer based in Northfield, Minnesota.)