

Wavelets, a modern tool for signal processing

Ivan W. Selesnick

The simplest wavelet transform is based on sums and differences of nearby points in a function. A variety of sophisticated variations allow signals and images to be effectively compressed, cleaned, and analyzed.

Ivan Selesnick is an associate professor in the department of electrical and computer engineering at Polytechnic University in Brooklyn, New York.

The Federal Bureau of Investigation uses wavelet transforms to compress digitally scanned fingerprints. NASA's Mars rovers use them to compress images acquired by their 18 cameras. The new JPEG 2000 format is based on wavelet transforms; the smaller files allow one to store images using less memory and to transmit those images faster and more reliably. Researchers also use wavelet transforms to clean signals and images—that is, to reduce unwanted noise and blurring. As illustrated in the figure, some algorithms for processing astronomical images are based on wavelets and wavelet-like transforms.

A wavelet transform, like a Fourier transform, involves integrating a product of a signal and an oscillating function. But unlike the everlasting sines and cosines of Fourier analysis, the oscillating functions in a wavelet transform are usually stretched and translated versions of a single oscillating function of short duration; indeed that localized function is the “wavelet.” This tutorial describes an elementary wavelet transform, illustrates why it is effective for compression and noise reduction, and briefly describes how the basic wavelet and noise reduction methods can be improved. The focus is on wavelet transforms used for image compression and reduction of noise and blur; such transforms must be invertible. But a second type of wavelet transform is worth noting briefly. It is designed for signal analysis—for example, to detect faults in machinery from sensor measurements, to study electroencephalograms or other biomedical signals, or to determine how the frequency content of a signal evolves over time. In those cases, the wavelet transform need not be invertible.

A Haar's breadth

The most basic wavelet transform is the Haar transform, introduced by Alfred Haar in 1910. It is easiest to describe for a function of a small number of points, so consider the eight-point signal $x(n)$



which may be written as the eight-component “vector” $x(n) = (3, 4, 5, 5, 7, 6, 4, 2)$. The basis of the Haar transform is

to express $x(n)$ in terms of an average of signal-value pairs, $c(n) = (3.5, 5, 6.5, 3)$, and a signal that encompasses differences, $d(n) = (-0.5, 0, 0.5, 1)$.

The average and difference functions may be written in the form

$$c(n) = 0.5 x(2n - 1) + 0.5 x(2n),$$


$$d(n) = 0.5 x(2n - 1) - 0.5 x(2n),$$

and represented by the block diagram



The decomposition into average and difference signals can be reversed, and so the original function can be reconstituted from the shorter signals $c(n)$ and $d(n)$.

The Haar transform is obtained by sequentially repeating the average and difference process on each output average function. For the eight-point signal $x(n)$, the process may be repeated up to three times, as expressed in the block diagram



The Haar wavelet representation of $x(n)$ is thus the set of four output signals:

$$c_3 = (4.5)$$


$$d_3 = (-0.25)$$

$$d_2 = (-0.75, 1.75)$$

$$d = (-0.5, 0, 0.5, 1).$$

Note that the single value c_3 is simply the average of the eight values in the original function x .

To be useful for compression and noise reduction, a wavelet representation must contain many small values; in the language of the field, it must be sparse. The first step in most signal- and image-compression algorithms is to obtain such a representation. The preponderance of zeros—in practice, small non-zero values—in sparse representations means that fewer bits need to be stored in memory. As discussed below, sparseness also facilitates noise reduction.



Galaxy NGC 2997 thrice viewed. (a) This image gives the conventional view. The galaxy's image has been separated into its stars (b) and vapor (c). Jean-Luc Starck and colleagues implemented the decomposition with a combination of wavelets and their variations. (Images courtesy of Jean-Luc Starck.)

The Haar transform provides a relatively sparse wavelet representation for signals that are approximately piecewise constant. Many commonly encountered signals are not piecewise constant, but they are composed of smooth bits separated by jumps. An exemplar is a photographic scan line—that is, a single row taken from a photographic image. For signals such as scan lines, one must use the newer transforms to obtain sparse wavelet representations.

For example, in 1988 Ingrid Daubechies constructed a family of easily invertible transforms that, like the Haar transform, are readily implemented as a succession of decompositions. Daubechies' method allows for the construction of sparse wavelet representations for signals that are piecewise polynomial. By now, mathematicians and engineers have developed a great number of decompositions from which useful wavelet transforms can be built.

Wavelets are particularly valuable in that they preserve the jumps in scan lines and similar signals. For smooth signals without jumps, highly oscillatory signals, and other signals that confound attempts to build sparse wavelet representations, Fourier transforms or conventional low-pass filters may lead to superior compression and noise abatement.

Noises off

In addition to allowing for efficient compression, sparse representations can also be effective for noise reduction. Suppose you have a noise-free signal and have obtained a sparse wavelet representation. Now imagine adding noise to the original signal in the form of a zero-mean Gaussian random variable of standard deviation σ . The wavelet representation of the noisy signal would itself be noisy, but it is possible to implement the wavelet transform so that the transformed noise is also a zero-mean Gaussian with standard deviation σ .

The simplest wavelet-based method for reducing noise is just to zero out all terms in the wavelet representations whose absolute values are less than some threshold, say 2σ or 3σ . After that modification, one could invert the wavelet transform. On comparing the inverted transform with the original noise-free signal, one typically observes that much of the added noise has indeed been removed, though some noise spikes remain due to the few values in the noisy wavelet representation that exceeded the threshold. One way to reduce those unwanted noise spikes is to choose a larger threshold and so set more values to zero. Setting the thresh-

old too high, however, will cause distortions in the processed signal. Another option is to use a different threshold rule—for example, a “soft-threshold” rule that not only sets small values to zero but also shrinks the remaining values toward zero. With the soft-threshold approach, values in the noisy wavelet representation that slightly exceed the threshold are reduced in magnitude. The signal obtained by inverting the wavelet transform suffers less from spurious noise spikes. But the soft-threshold process can compromise the detail of the original signal that one is trying to reproduce.

Scientists, engineers, and statisticians have actively been developing new transforms and algorithms for processing, analyzing, and storing signals and images. One avenue of research centers on so-called data-adaptive transforms that can automatically find a good transform for a specific signal. Some new decompositions generate redundant data and can offer improved results. For many images and data, a different kind of transform is needed because the conventional wavelet transform provides an optimally sparse representation only for pointlike discontinuities. Many images, though, have discontinuities along curves—the edges of objects. The curvelet transform, introduced by Emmanuel Candes and David Donoho, is designed to provide sparse representations specifically in that case. For some images, it is difficult to obtain a sparse representation with any single transform. In those cases, as illustrated in the figure, it may be useful to use several transforms simultaneously.

Additional resources

- ▶ C. S. Burrus, R. A. Gopinath, H. Guo, *Introduction to Wavelets and Wavelet Transforms: A Primer*, Prentice Hall, Upper Saddle River, NJ (1997).
- ▶ I. Daubechies, *Ten Lectures on Wavelets*, Society for Industrial and Applied Mathematics, Philadelphia (1992).
- ▶ S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, San Diego, CA (1999).
- ▶ I. W. Selesnick, R. G. Baraniuk, N. G. Kingsbury, *IEEE Signal Process. Mag.* **22**(6), 123 (November, 2005).
- ▶ J.-L. Starck, M. Elad, D. L. Donoho, *Adv. Imaging Electron Phys.* **132**, 287 (2004). ■

The online version of this Quick Study includes a discussion of the Daubechies transforms and a detailed example of how wavelets can be used to clean up a noisy signal.