

Lecture 17

Wavelets: Motivation and Description

17.1 Learning Objectives

- Recognize the key limitations of the Fourier transform and the STFT, which provides a fixed temporal resolution for all frequency components.
- Understand the multi-resolution logic behind wavelet analysis, which provides finer temporal (spatial) resolution for higher frequency components, and coarser temporal (spatial) resolution for lower frequency components.

17.2 Introduction

Despite the enormous success of Fourier analysis in a variety of applications, one key limitation of Fourier analysis is the infinite support of its basis functions (in other words, the sines and cosines go on forever). This poses some fundamental challenges for representing signals that have a compact support in the spatial (or time) domain, or signals whose properties vary substantially in space. As we have seen in previous lectures, one relatively straightforward approach to overcome this limitation is the Short-Time Fourier Transform (STFT), where the signal is essentially broken up into smaller portions (using a pre-defined window function) and then the Fourier transform is applied to each of the portions. In other words, the STFT enables us to pick the temporal resolution at which we perform our Fourier analysis. Even though the STFT has proved useful, it has limited flexibility since the width and shape of the STFT window need to be set a priori, ie: the temporal resolution is fixed. Instead, it is highly appealing to design signal analysis methods where the temporal resolution (or spatial resolution in the case of imaging applications) is adaptive: high frequency components are analyzed with high resolution (since they can change rapidly in time/space), and low frequency components are analyzed with low resolution (since they change slowly in time/space). This leads us to the study of wavelets, which will occupy our next few lectures.




Figure 17.1: Alfred Haar and Ingrid Daubechies, two of the central contributors to the development of wavelets.

The first to mention the functions that later became known as wavelets was Alfred Haar in 1909, within an Appendix of his PhD thesis. The Haar wavelet is quite simple, defined as:

$$\psi(x) = \begin{cases} 1, & \text{if } 0 \leq x < 1/2 \\ -1, & \text{if } 1/2 \leq x < 1 \\ 0, & \text{otherwise} \end{cases} \quad (17.1)$$

The main idea behind wavelets is that functions such as $\psi(x)$ defined above can be scaled, shifted, and then combined to represent other functions. This is similar to the use of multiple sines and cosines at different frequencies in Fourier analysis, except that wavelets, unlike sinusoids, are well localized in space (or time). This localization is the source of the diminutive 'let' in the word wavelet ("small wave").

Just like in Fourier analysis, wavelet transforms can be divided into continuous and discrete versions. In the next few sections we will describe both the continuous and discrete wavelet transforms. Then, in subsequent lectures we will review one- and multi-dimensional applications of wavelets, with a focus on imaging and image processing.

17.3 Continuous Wavelet Transform

The wavelet transform of a function $f(x)$, at a scale $a > 0$ and translation b , is defined as:

$$F_w(a, b) = \frac{1}{|a|^{1/2}} \int_{-\infty}^{\infty} f(x) \overline{\psi} \left(\frac{x-b}{a} \right) dx \quad (17.2)$$

where $\psi(x)$ is called the mother wavelet, and the overline symbol denotes complex conjugate. In wavelet lingo, the mother wavelet generates the daughter wavelets, ie: translated and scaled versions of the mother wavelet.




Figure 17.2: Multi-resolution approach to signal analysis.

17.4 Discrete Wavelet Transform

A discrete wavelet transform (DWT)¹ can be seen as arising from discrete sampling of a continuous wavelet. The first DWT can be traced back to Alfred Haar. This transform can be understood as progressively calculating finite differences and averages between consecutive pairs of samples, leading to ‘high-pass’ and ‘low-pass’ representations of the signal. This same process can be performed subsequently on the low-pass filtered portion of the signal. The most widespread set of DWTs, formulated by Ingrid Daubechies in the 1980s, applies recurrence relations to generate progressively finer discrete samplings of an implicit mother wavelet function. The first wavelet in Daubechies’ family of wavelets is the Haar wavelet. A very common approach for understanding (and implementing!) discrete wavelet decompositions is to use a filter bank: a pair of low-pass and high-pass filters, where the output of the low-pass filter is fed into another pair of low-pass and high-pass filters, and so on. The output of each filter is downsampled by a factor of two to preserve the total number of coefficients. Figure 17.3 illustrates this filter bank approach to wavelet decompositions.

¹Make sure you have access to Matlab with the Wavelet toolbox for the computational examples




Figure 17.3: Filter bank view of wavelet decomposition.

17.5 Wavelet vs Fourier Analysis

17.5.1 Similarities

- Both the DFT and the DWT can be seen as matrix operations where we multiply our spatial-domain signal (of length M) by a matrix (eg: of size $M \times M$, although it may have different numbers of rows in wavelet transforms), leading to a transform-domain representation of our signal. Of course, the specific matrix is different for DFT and for DWT.
- The inverse DFT and inverse DWT can be obtained by applying the inverse of the transform matrix described in the previous point (for square matrices).
- Both the DFT and the DWT provide information about the frequency content of signals.
- Both the DFT and the DWT can be computed efficiently using fast algorithms. In previous lectures, we have discussed the FFT as a fast algorithm (in fact, a family of fast algorithms) for computing the DFT. Over the next few lectures, we will also mention fast algorithms for computing the DWT.
- Both the DFT and the DWT have broad utility in signal processing and imaging, including image reconstruction, processing (eg: denoising), and signal compression. Indeed, the image compression standard “JPEG” is based on a modified version of the DFT termed the discrete cosine transform (DCT), whereas the more recent standard “JPEG 2000” is based on the DWT.

17.5.2 Differences

- First of all, it is important to note that there is a single DFT, whereas there are infinitely many choices of the DWT. Different wavelet families present different tradeoffs between compactness in space (ability to capture localized features) and smoothness.

- The most relevant difference between the DFT and the DWT is that the basis functions for the DFT cover all of space (they are sines and cosines, or complex exponentials with constant amplitude over the entire domain), whereas the basis functions for the DWT are localized in space. What this means in practice is that wavelets are a lot better than Fourier analysis at efficiently describing signals with localized sharp features (which would result in non-zero Fourier coefficients at all frequencies).
- Although the STFT makes an attempt at providing both temporal (or spatial) and frequency localization, the multi-resolution approach inherent to wavelets is arguably much more elegant and general. In wavelet analysis, low frequencies are analyzed over long spans of time (or space), and high frequencies (features that change fast) are analyzed over short spans of time (or space). In other words, the wavelet transform includes multi-resolution basis functions: short, high-frequency basis functions in addition to long, low-frequency basis functions. See Figure [17.2](#) for an illustration of this critical difference.

