

Mathematics of Image and Data Analysis
Math 5467

The Discrete Wavelet Transform

Instructor: Jeff Calder
Email: jcalder@umn.edu

<http://www-users.math.umn.edu/~jwcalder/5467>

Last time

- The sampling theorem
- Discrete Cosine Transform and image compression

Today

- The Wavelet Transform (1D and 2D Haar Wavelets)

Localization/delocalization of the DFT

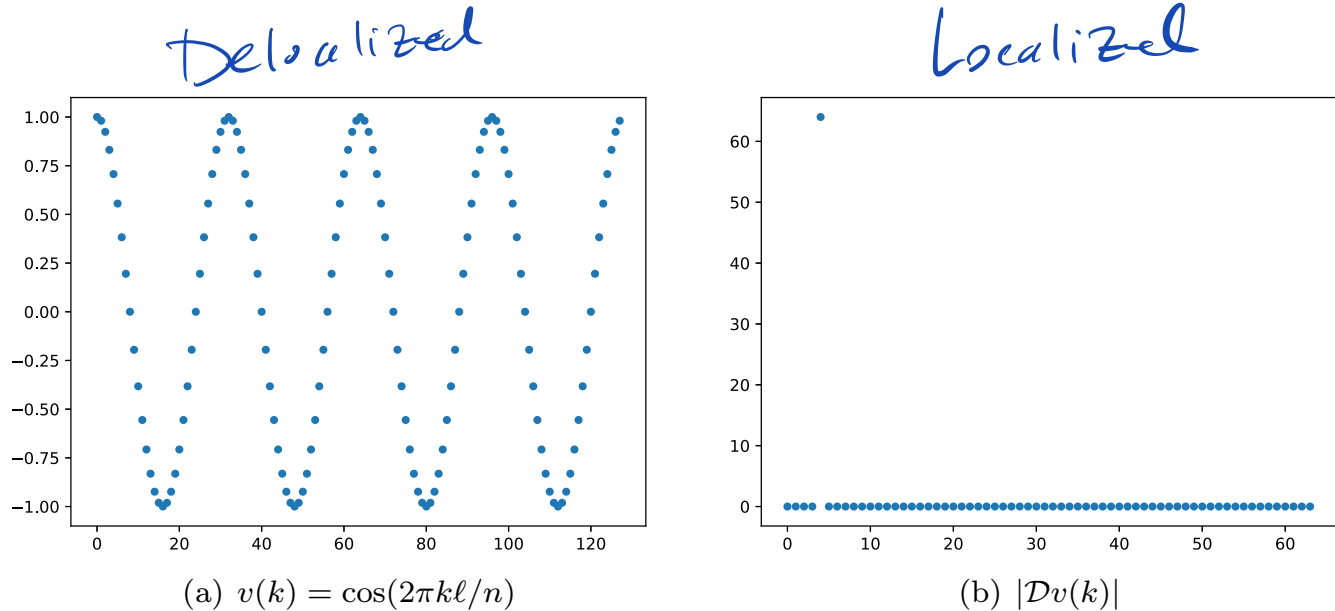


Figure 1: A plot the real part of a Fourier basis function and its Discrete Fourier Transform (DFT). The function v ($\ell = 4, n = 128$) is completely delocalized (most values are nonzero), while its DFT is highly localized (all values vanish except one).

Forcing spatial localization through blocking/windowing






Figure 2: The cameraman image and its Discrete Cosine Transform (DCT) coefficients computed on 8×8 blocks.

Windowing delocalizes the frequency domain




(a) $v(k) = \cos(2\pi k\ell/n)$




(b) $|\mathcal{D}v(k)|$

Figure 3: A plot the real part of a Fourier basis function and its Discrete Fourier Transform (DFT). The function v ($\ell = 4, n = 128$) is completely delocalized (most values are nonzero), while its DFT is highly localized (all values vanish except one).

Windowing delocalizes the frequency domain



(a) Windowed Fourier basis function



(b) Windowed DFT

Figure 4: A plot the real part of a windowed Fourier basis function and its Discrete Fourier Transform (DFT). The windowed function is more localized in space, while its DFT is less localized in the frequency domain, compared to Figure 3.

Uncertainty Principle

In fact, there is a fundamental limit to how much a function and its DFT can both be localized. The *uncertainty principle* states that

$$\|f\|_0 \|\mathcal{D}f\|_0 \geq n,$$

where $\|f\|_0$ is the number of nonzero values of f .

- This says that it is impossible for both f and $\mathcal{D}f$ to both be localized (i.e., have mostly zero entries).
- This bound is saturated by the Fourier basis functions u_ℓ , which satisfy $\|u_\ell\|_0 = n$ and $\|\mathcal{D}u_\ell\|_0 = 1$.

Uncertainty Principle

In fact, there is a fundamental limit to how much a function and its DFT can both be localized. The *uncertainty principle* states that

$$(1) \quad \|f\|_0 \|\mathcal{D}f\|_0 \geq n,$$

where $\|f\|_0$ is the number of nonzero values of f .

- This says that it is impossible for both f and $\mathcal{D}f$ to both be localized (i.e., have mostly zero entries).
- This bound is saturated by the Fourier basis functions u_ℓ , which satisfy $\|u_\ell\|_0 = n$ and $\|\mathcal{D}u_\ell\|_0 = 1$.

Heisenberg Uncertainty Principle:

In quantum mechanics, the probability distributions of the position and momentum of a particle are the Fourier Transforms of each other. In this context, the uncertainty principle (1) says that their distributions cannot both be localized, meaning we cannot determine both the position and momentum of a particle with high precision.

The Wavelet Transformation

The Wavelet Transformation is a principled approach to finding a decomposition of a signal or image into frequency components where the basis functions are localized in time/space and frequency, as much as is possible.

- Multi-scale analysis.
- Many different varieties of wavelets.

In this course, we will cover the Haar Wavelet in depth, and discuss general Wavelet transformations briefly.

The 1D Haar Wavelet Transformation

The Wavelet Transform is based on repeatedly decomposing a signal into a low frequency part, called the *approximation coefficients*, and a high frequency part, called the *detail coefficients*. This is best illustrated at first with an example. Consider the following length $n = 8$ signal

Signal: (7, 5, 6, 3, 2, 5, 4, 1)

AB

$$\text{Approximation} = A + B$$

$$\text{Detail} = B - A$$

App.

Detail

(12, 9, 7, 5)

(-2, -3, 3, -3)

1-Level Haar:

(12, 9, 7, 5, -2, -3, 3, -3)

Approx

Detail

2-level: Apply 1-level to Approx. eseff.

(12, 9, 7, 5, -2, -3, 3, -3)

(21, 12, -3, -2, -2, -3, 3, -3)

level 2 level 2 level 1 detail.
App. Detail

3-level:

(33, -9, -3, -2, -2, -3, 3, -3)

level 3 level 3 level 2 level 1 detail.
Approx detail Detail

Inverse: $L = \text{Approximation} = A + B$

$H = \text{Detail} = B - A$

$$L + H = 2B, \quad B = \frac{L + H}{2}$$

$$L - H = 2A, \quad A = \frac{L - H}{2}$$

Matrix of Haar Wavelet Transformation

We can view the 3-level Haar Wavelet Transformation as multiplication of the signal $f = (7, 5, 6, 3, 2, 5, 4, 1)$ by the matrix

$$(2) \quad W = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ -1 & -1 & -1 & -1 & 1 & 1 & 1 & 1 \\ -1 & -1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & -1 & 1 & 1 \\ -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 \end{bmatrix}$$

-L3 Approx
-L3 Detail
-L2 Detail
-L1 Detail.

Rows are orthogonal.

Haar Transform

In general, we can define a one-level Haar Wavelet Transform of a signal $f \in L^2(\mathbb{Z}_n)$ of length $n = 2^k$ as the signal $\mathcal{W}_1 f$ given by

$$(3) \quad \mathcal{W}_1 f(j) = \begin{cases} f(2j+1) + f(2j), & \text{if } 0 \leq j \leq \frac{n}{2} - 1 \\ f(2j-n+1) - f(2j-n), & \text{if } \frac{n}{2} \leq j \leq n-1. \end{cases}$$

Approx.

Detail

In general, the ℓ^{th} -level Haar Wavelet Transformation is

$$(4) \quad \mathcal{W}_\ell f(j) = \begin{cases} \mathcal{W}_{\ell-1} f(2j+1) + \mathcal{W}_{\ell-1} f(2j), & \text{if } 0 \leq j \leq \frac{n}{2^\ell} - 1 \\ \mathcal{W}_{\ell-1} f(2j - \frac{n}{2^{\ell-1}} + 1) - \mathcal{W}_{\ell-1} f(2j - \frac{n}{2^{\ell-1}}), & \text{if } \frac{n}{2^\ell} \leq j \leq \frac{n}{2^{\ell-1}} - 1 \\ \mathcal{W}_{\ell-1} f(j), & \text{if } \frac{n}{2^{\ell-1}} \leq j \leq n-1 \end{cases}$$

ℓ -level transform

$\ell-1$ detail

Python Code

Algorithm 1 The Haar Wavelet Transformation in Python

```
1  import numpy as np
2
3  def haar_wavelet(f,depth):
4      g = np.zeros_like(f)
5      n2 = len(f)>>1
6      g[:n2] = f[::2] + f[1::2] #Approximation coeff
7      g[n2:] = f[1::2] - f[::2] #Detail coeff
8      if depth >= 2:
9          g[:n2] = haar_wavelet(g[:n2],depth-1)
10     return g
```


Similar recursive structure as the FFT. Since only half of the signal is recursed on, the total complexity is $O(n)$.

The inverse Haar Transform

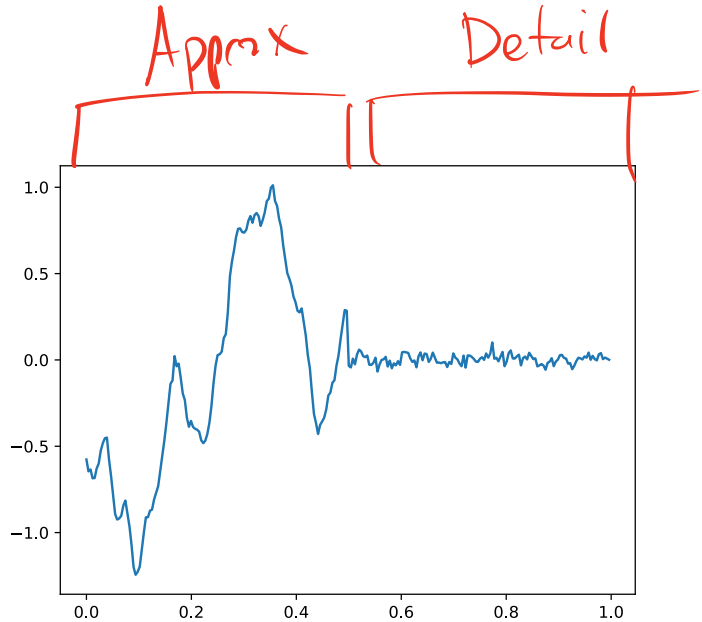
Algorithm 2 The Inverse Haar Wavelet Transformation in Python

```
1  import numpy as np
2
3  def inverse_haar_wavelet(f,depth):
4      if depth == 0:
5          return f
6      else:
7          n2 = len(f)>>1
8          h = inverse_haar_wavelet(f[:n2],depth-1)
9          g = np.zeros_like(f)
10         g[1::2] = (h + f[n2:])/2
11         g[::2] = (h - f[n2:])/2
12         return g
```

Example on a noisy signal




(a) Noisy Signal



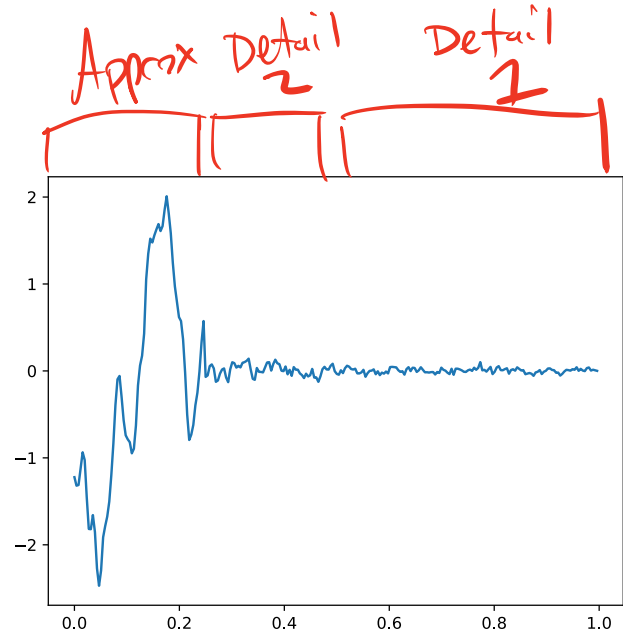
(b) Haar Wavelet Transform (1 level)

Figure 5: One level of the Haar Wavelet Transformation applied to a noisy signal. Notice that much of the noise appears in the detail coefficients. Wavelet based denoising and compression algorithms are based on thresholding the detail coefficients.

Example on a noisy signal




(a) Noisy Signal




(b) Haar Wavelet Transform (2 level)

Figure 6: Two levels of the Haar Wavelet Transformation applied to a noisy signal. Notice that much of the noise appears in the detail coefficients. Wavelet based denoising and compression algorithms are based on thresholding the detail coefficients.

Example on a noisy signal




(a) Noisy Signal




(b) Haar Wavelet Transform (3 level)

Figure 7: Three level of the Haar Wavelet Transformation applied to a noisy signal. Notice that much of the noise appears in the detail coefficients. Wavelet based denoising and compression algorithms are based on thresholding the detail coefficients.

Example on a noisy signal



(a) Noisy Signal



(b) Haar Wavelet Transform (8 level)

Figure 8: Eight level of the Haar Wavelet Transformation applied to a noisy signal. Notice that much of the noise appears in the detail coefficients. Wavelet based denoising and compression algorithms are based on thresholding the detail coefficients.

2D Haar Wavelets

One level of the 2D Haar Wavelet Transformation acts on 2×2 blocks in images:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

The approximation, vertical, horizontal, and diagonal detail coefficients, denoted A, V, H, D , respectively, are given by

$$\begin{aligned} A &= a + b + c + d && - \text{Approx} \\ H &= -a - b + c + d && - \text{Horizontal} \\ V &= -a + b - c + d && - \text{Vertical} \\ D &= a - b - c + d && - \text{Diagonal} \end{aligned}$$

2D Haar Wavelets

The inverse transform is simple to obtain; indeed, we have

$$a = \frac{1}{4}(A - H - V + D)$$

$$b = \frac{1}{4}(A - H + V - D)$$

$$c = \frac{1}{4}(A + H - V - D)$$

$$d = \frac{1}{4}(A + H + V + D).$$

Further levels of the transform are obtained by applying the transform again to the approximation image.

Example of 2D Haar Wavelets

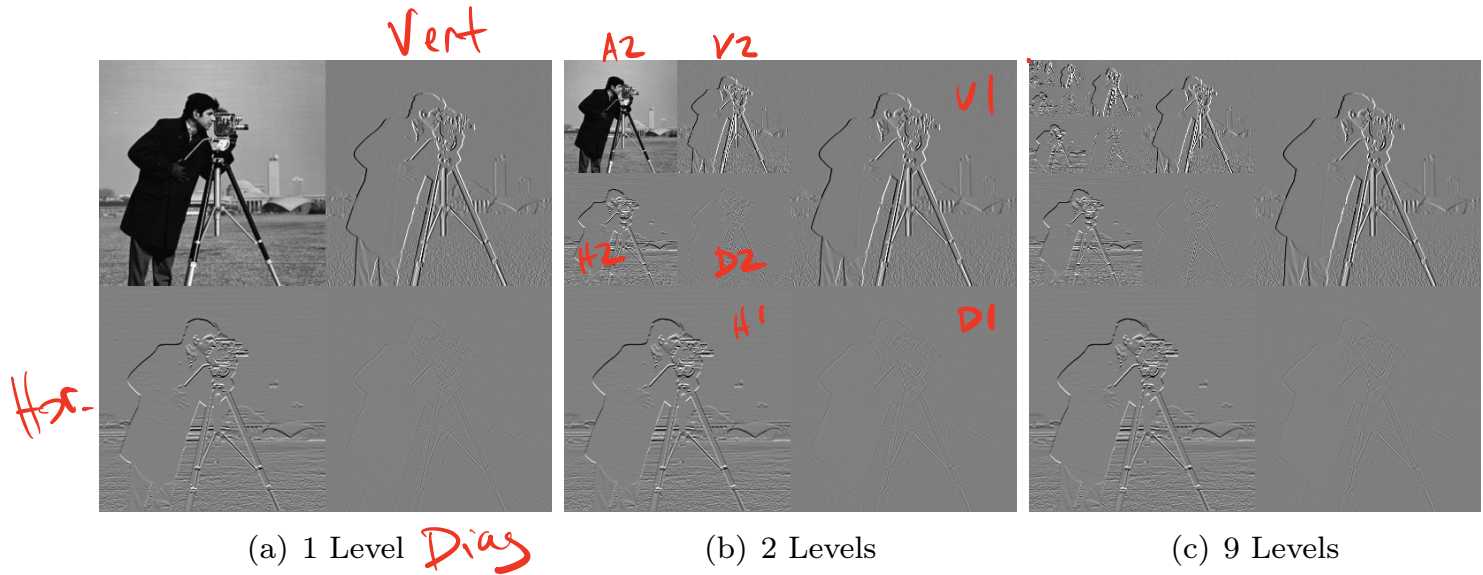


Figure 9: The Haar Wavelet Transformation of levels 1, 2, and 9 on the cameraman image. The approximation image is placed in the upper left corner, the horizontal detail in the lower left, the vertical detail in the upper right, and the diagonal detail in the lower right.

Wavelet compression and denoising (.ipynb)