15-853:Algorithms in the Real World

Data Compression 4

Compression Outline

Introduction: Lossy vs. Lossless, Benchmarks, …

Information Theory: Entropy, etc.

Probability Coding: Huffman + Arithmetic Coding

Applications of Probability Coding: PPM + others

Lempel-Ziv Algorithms: LZ77, gzip, compress, …

Other Lossless Algorithms: Burrows-Wheeler

Lossy algorithms for images: JPEG, MPEG, ...

- Scalar and vector quantization
- JPEG and MPEG

Compressing graphs and meshes: BBK

Scalar Quantization

Quantize regions of values into a single value:

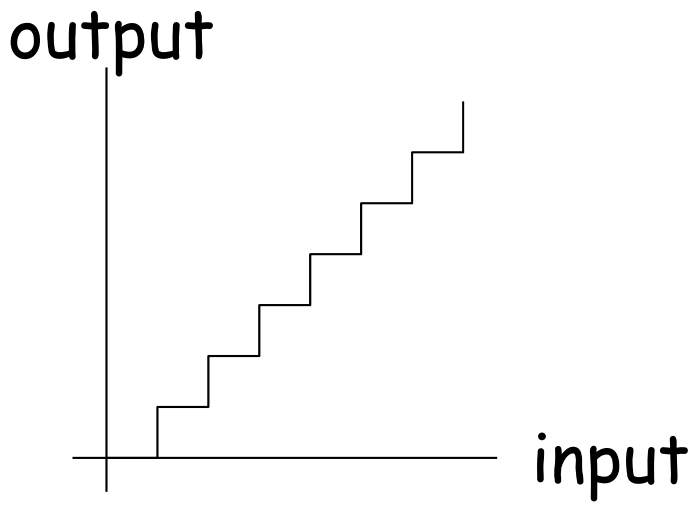

uniform

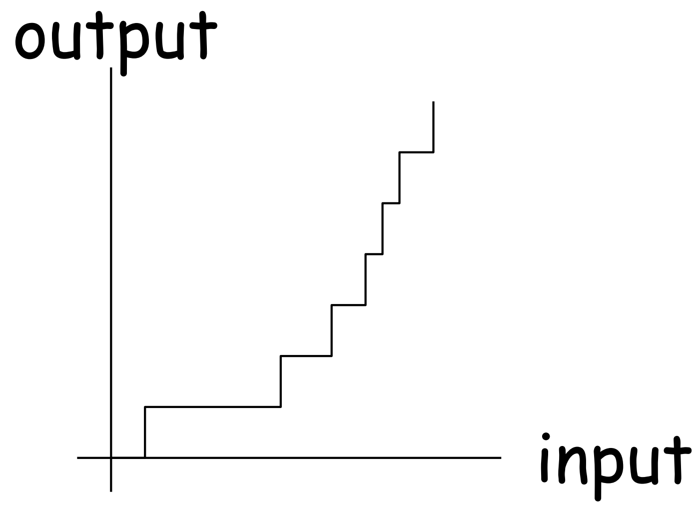

non uniform

Can be used to reduce # of bits for a pixel

Vector Quantization

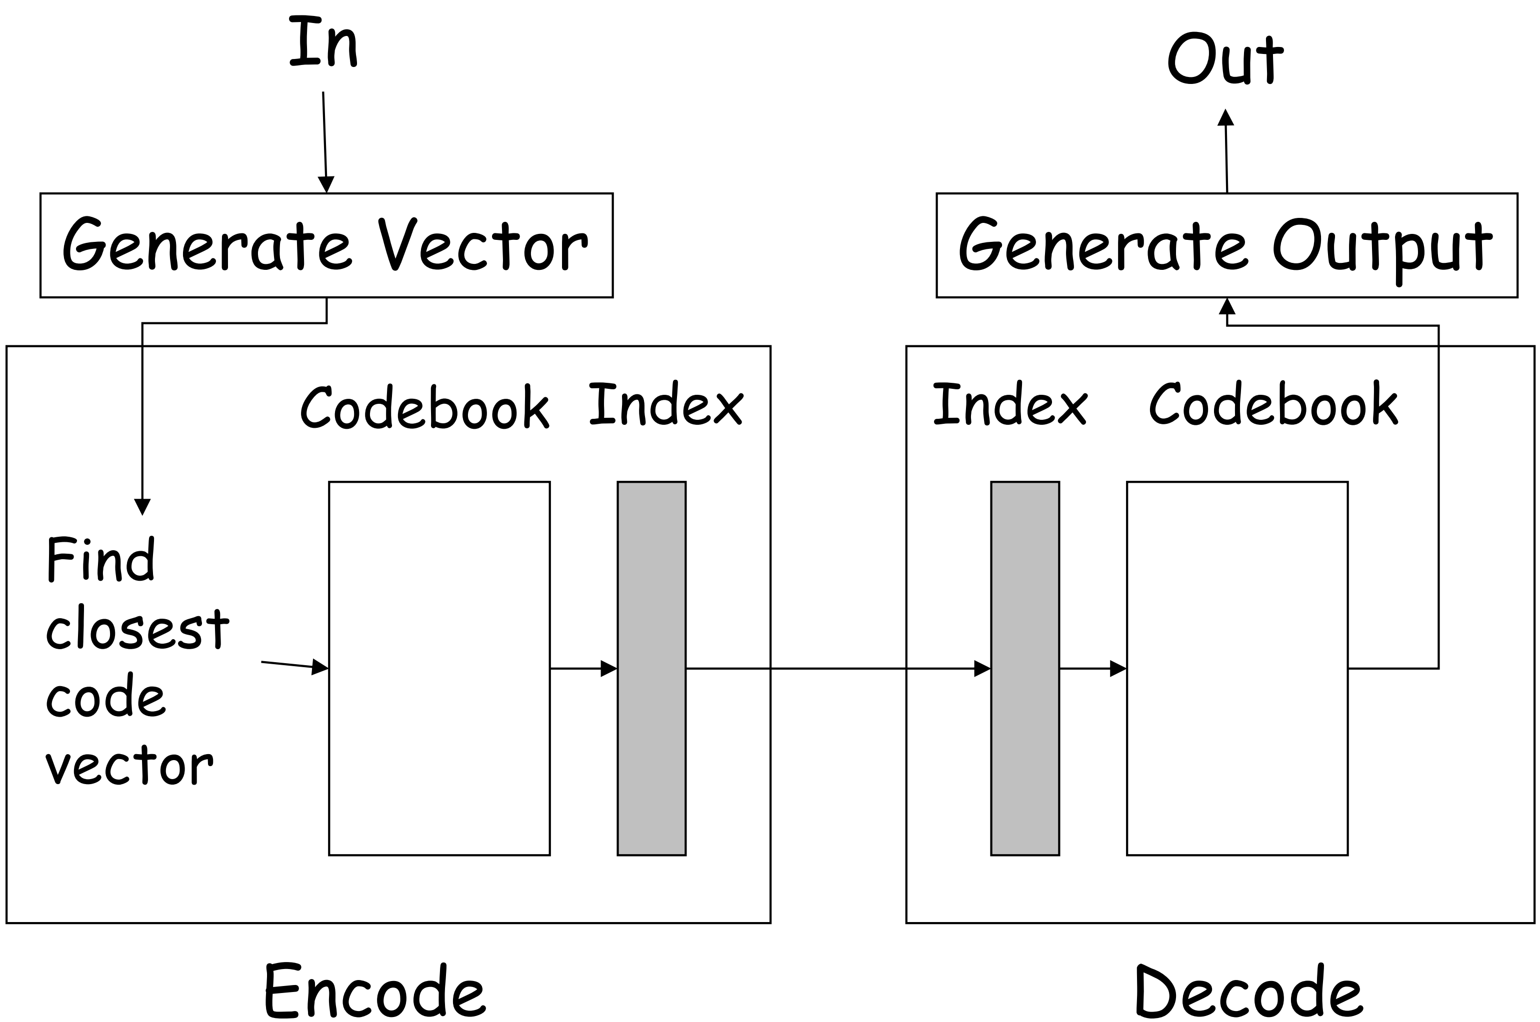

Vector Quantization

What do we use as vectors?

- Color (Red, Green, Blue)
  - Can be used, for example to reduce 24bits/pixel to 8bits/pixel
  - Used in some terminals to reduce data rate from the CPU (colormaps)
- K consecutive samples in audio
- Block of K pixels in an image

How do we decide on a codebook

- Typically done with **clustering**

Vector Quantization: Example

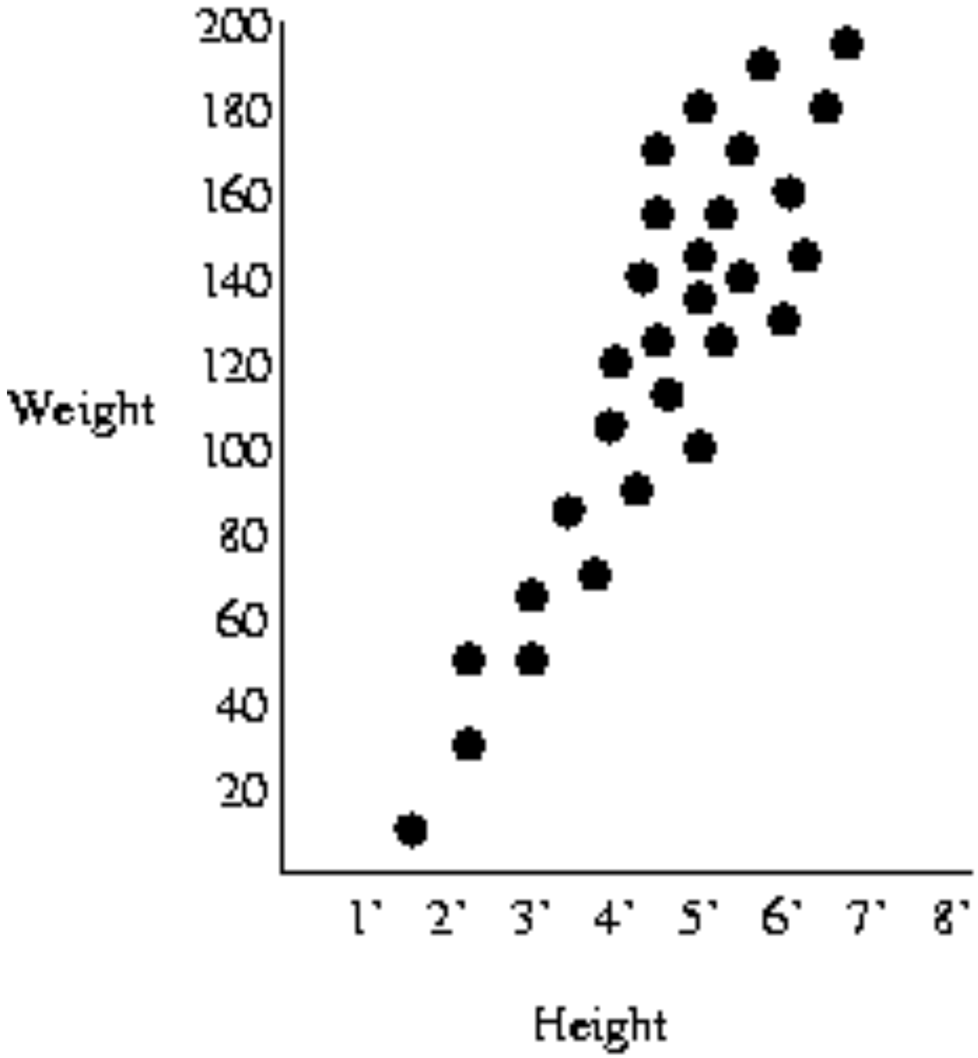

Linear Transform Coding

Want to encode values over a region of time or space

- Typically used for images or audio

Select a set of linear basis functions ϕ_i that span the space

- sin, cos, spherical harmonics, wavelets, …
- Defined at discrete points

Linear Transform Coding

Coefficients: $\Theta_i = \sum_j x_j \phi_i(j) = \sum_j x_j a_{ij}$

Θ_i = i^{th} resulting coefficient

x_j = j^{th} input value

a_{ij} = ij^{th} transform coefficient = $\phi_i(j)$

In matrix notation:

$$\Theta = Ax$$

$$x = A^{-1}\Theta$$

Where A is an n x n matrix, and each row defines a basis function

Example: Cosine Transform

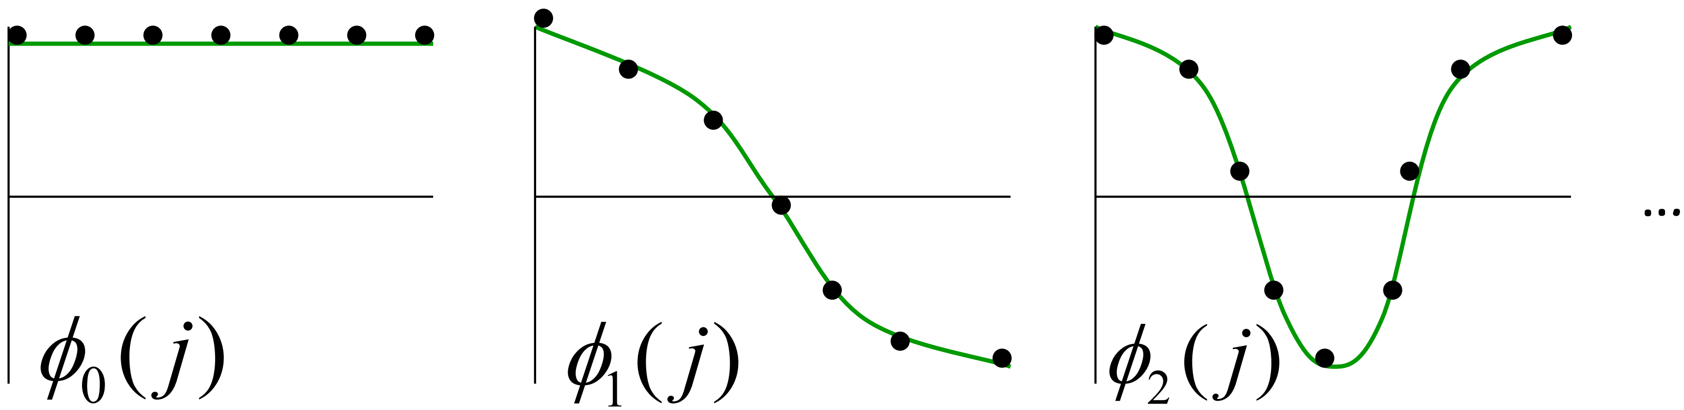

$$\Theta_i = \sum_j x_j \phi_i(j)$$

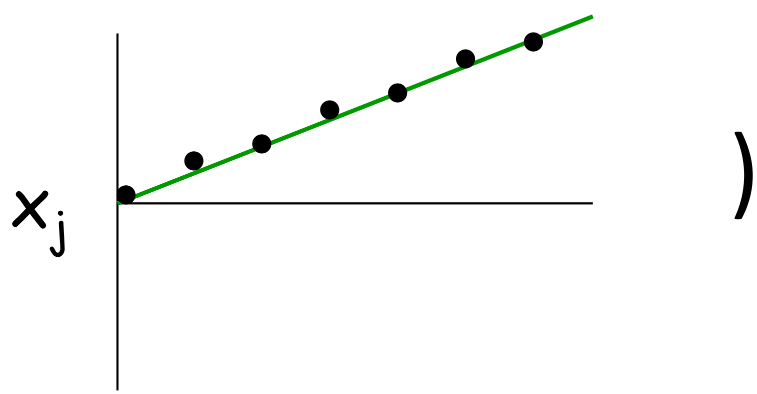

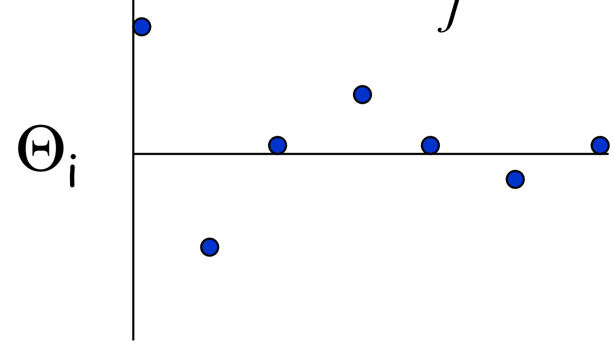

Other Transforms

Polynomial:

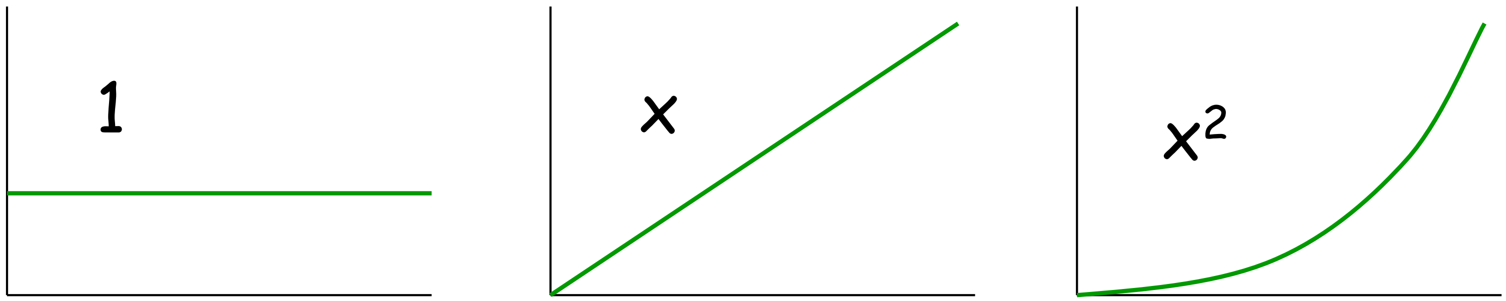

Wavelet (Haar):

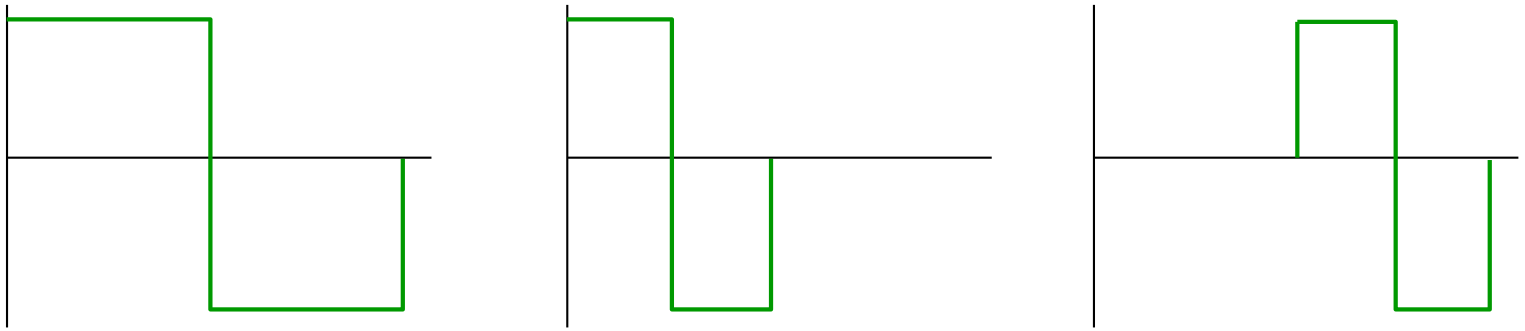

How to Pick a Transform

Goals:

- Decorrelate
- Low coefficients for many terms
- Basis functions that can be ignored by perception

Why is using a Cosine of Fourier transform across a whole image bad?

How might we fix this?

Usefulness of Transform

Typically transforms A are **orthonormal**

Properties of orthonormal transforms:

$\sum x^2 = \sum \Theta^2$ (energy conservation, Parseval's theorem)

Would like to compact energy into as few coefficients as possible

$$G_{TC} = \frac{\frac{1}{n}\sum \sigma_i^2}{\left(\prod \sigma_i^2\right)^{1/n}}$$

(the **transform coding gain**)
arithmetic mean/geometric mean

$\sigma_i = (\Theta_i - \Theta_{av})$

The higher the gain, the better the compression

Case Study: JPEG

A nice example since it uses many techniques:

- Transform coding (Cosine transform)
- Scalar quantization
- Difference coding
- Run-length coding
- Huffman or arithmetic coding

JPEG (Joint Photographic Experts Group) was designed in **1991** for **lossy** and **lossless** compression of **color** or **grayscale images**. The lossless version is rarely used.

Can be adjusted for compression ratio (typically 10:1)

JPEG in a Nutshell

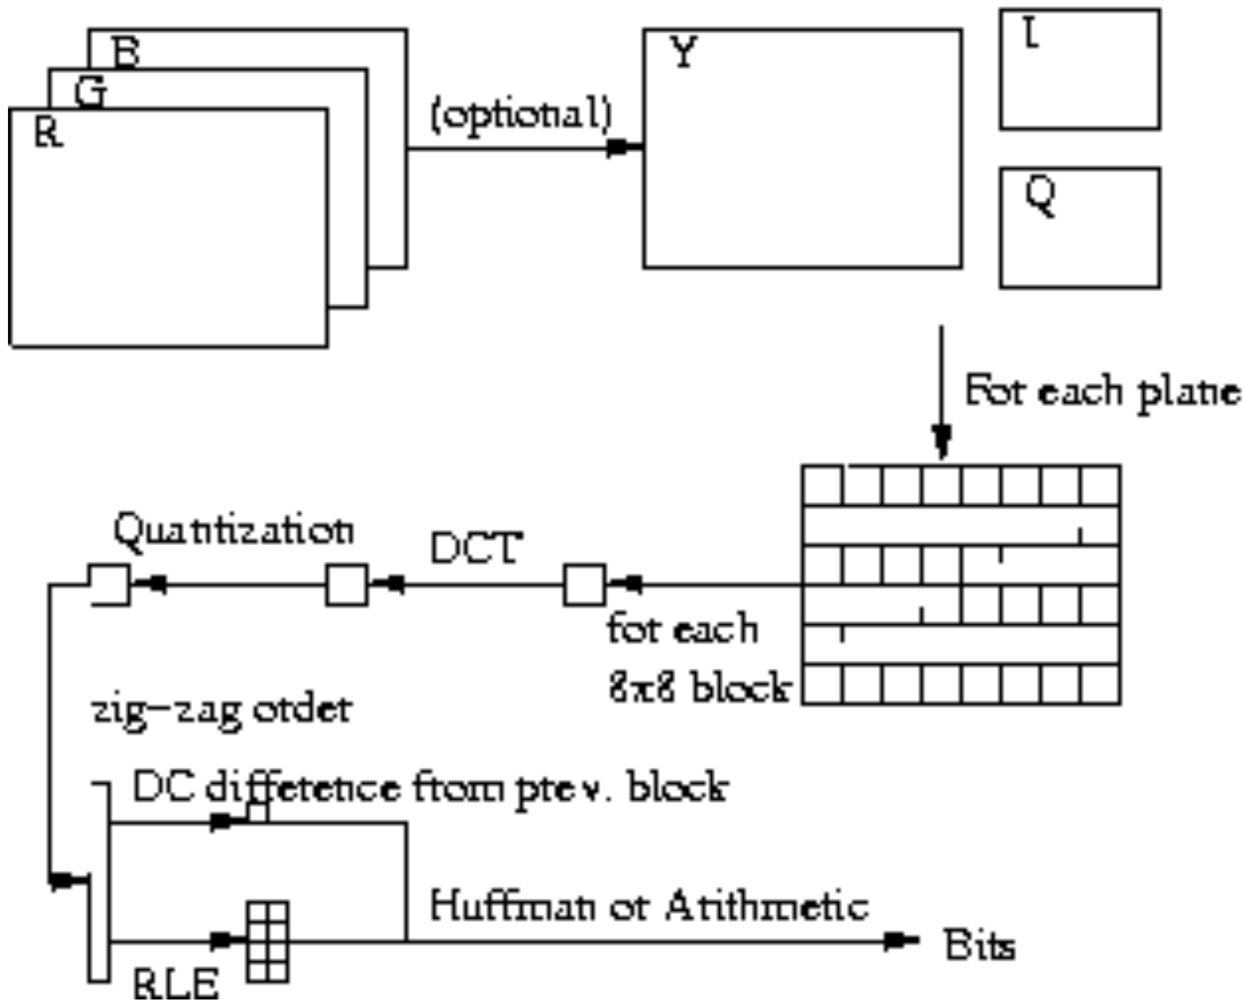

JPEG: Quantization Table

16	11	10	16	24	40	51	61
12	12	14	19	26	58	60	55
14	13	16	24	40	57	69	56
14	17	22	29	51	87	80	62
18	22	37	56	68	109	103	77
24	35	55	64	81	104	113	92
49	64	78	87	103	121	120	101
72	92	95	98	112	100	103	99

Also divided through uniformaly by a quality factor which is under control.

JPEG: Block scanning order

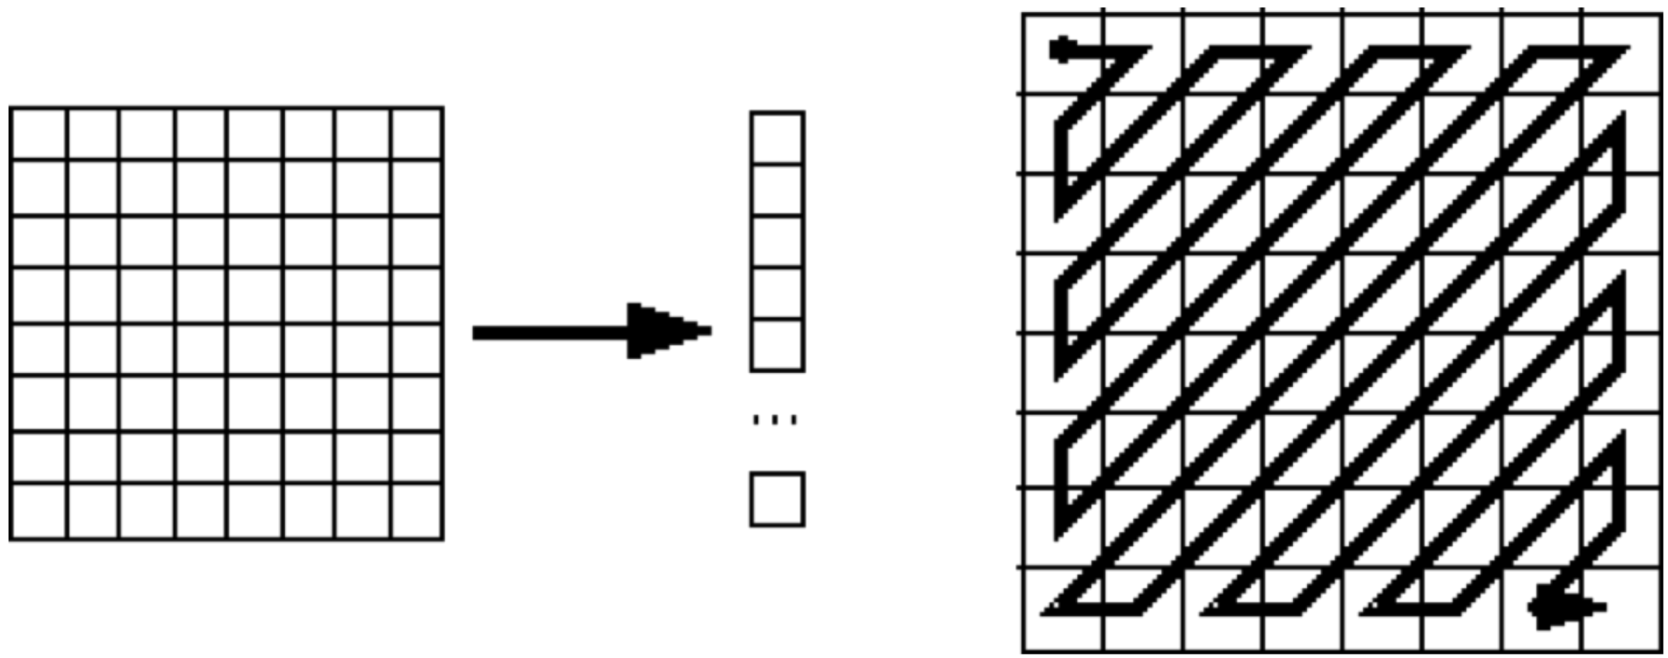

Uses run-length coding for sequences of zeros

JPEG: example

.125 bits/pixel (factor of 200)

Case Study: MPEG

Pretty much JPEG with **interframe coding**

Three types of frames

- **I** = intra frame (aprox. JPEG) anchors
- P = predictive coded frames
- B = bidirectionally predictive coded frames

Example:

Type:	I	B	B	P	B	B	P	B	B	P	B	B	I
Order:	1	3	4	2	6	7	5	9	10	8	12	13	11

I frames are used for random access.

MPEG matching between frames

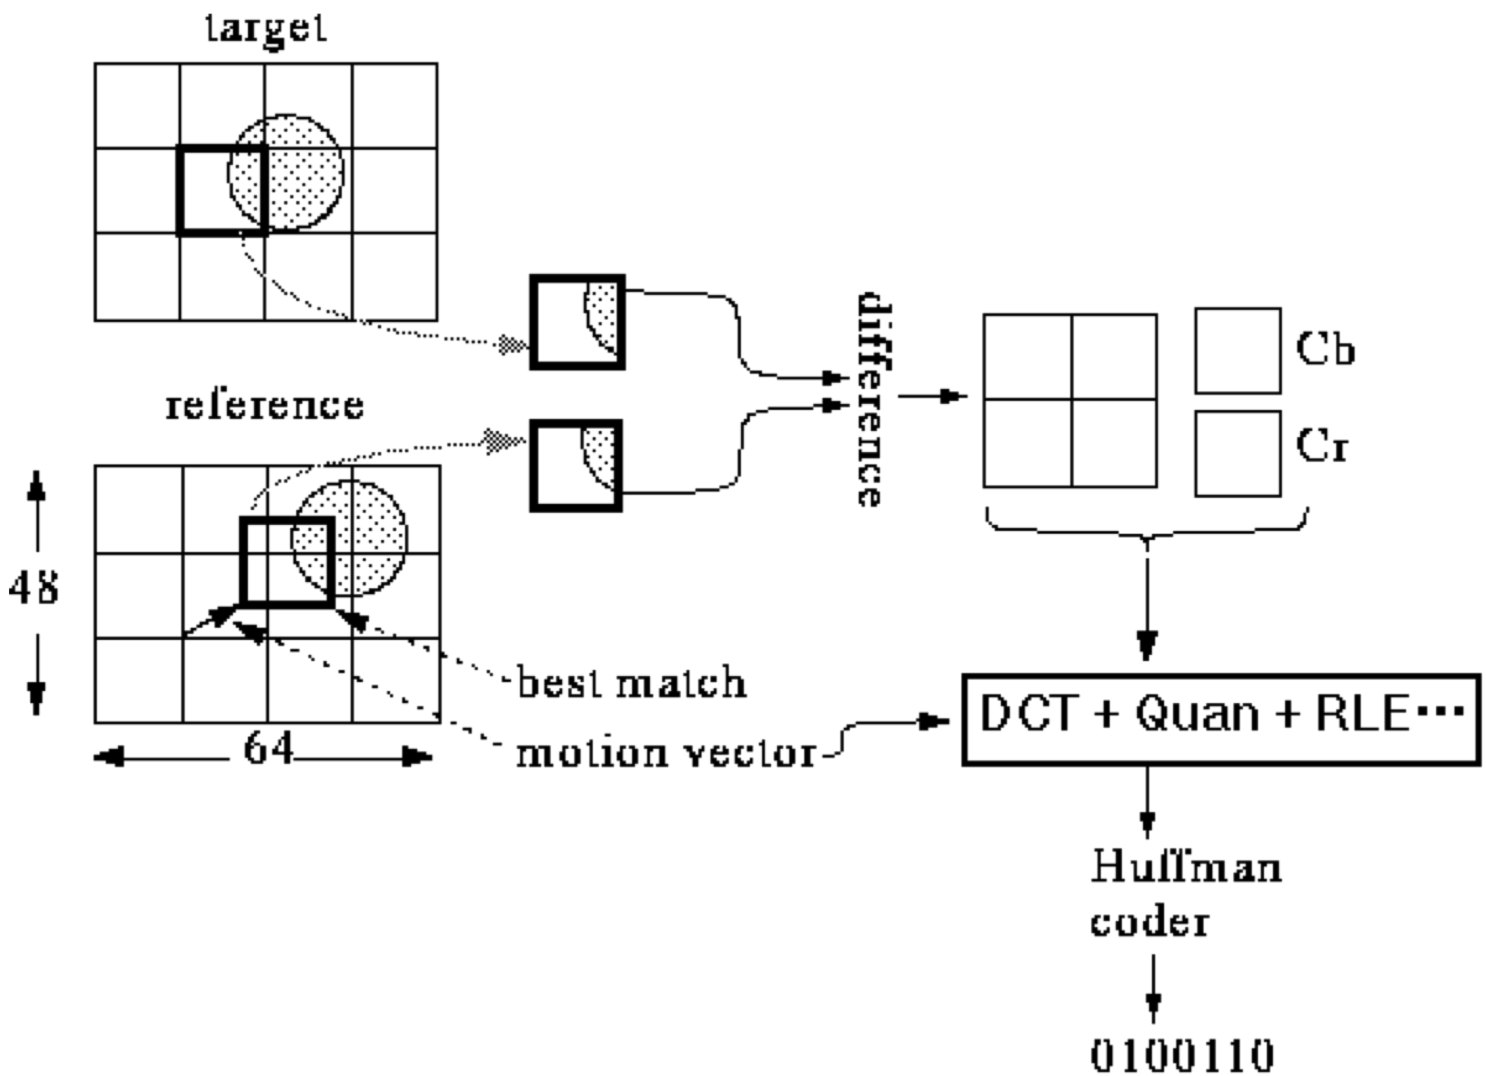

MPEG: Compression Ratio

356 x 240 image

Type	Size	Compression
I	18KB	7/1
P	6KB	20/1
B	2.5KB	50/1
Average	4.8KB	27/1

30 frames/sec x 4.8KB/frame x 8 bits/byte
= 1.2 Mbits/sec + .25 Mbits/sec (stereo audio)

HDTV has 15x more pixels
= 18 Mbits/sec

MPEG in the "real world"

- DVDs
  - Adds "encryption" and error correcting codes
- Direct broadcast satellite
- HDTV standard
  - Adds error correcting code on top
- Storage Tech "Media Vault"
  - Stores 25,000 movies

Encoding is much more expensive than encoding.

Still requires special purpose hardware for high resolution and good compression. Now available on some processors or using GPUs.

H.264 (video)

Uses similar ideas, but

- 4x4 transform that is an approximation of the cosine transform
- More sophisticated prediction of blocks based on their neighbors

Wavelet Compression

- A set of localized basis functions
- Avoids the need to block

"**mother function**" $\varphi(x)$

$\varphi_{sl}(x) = \varphi(2^s x - l)$

s = scale   l = location

Requirements

$$\int_{-\infty}^{\infty} \varphi(x)dx = 0 \quad \text{and} \quad \int_{-\infty}^{\infty} |\varphi(\mathrm{x})|^2 dx < \infty$$

Many **mother** functions have been suggested.

Haar Wavelets

Most described, least used.

$$\varphi(x) = \begin{cases} 1 & 0 \le x < 1/2 \\ -1 & 1/2 \le x < 1 \\ 0 & \text{otherwise} \end{cases}$$

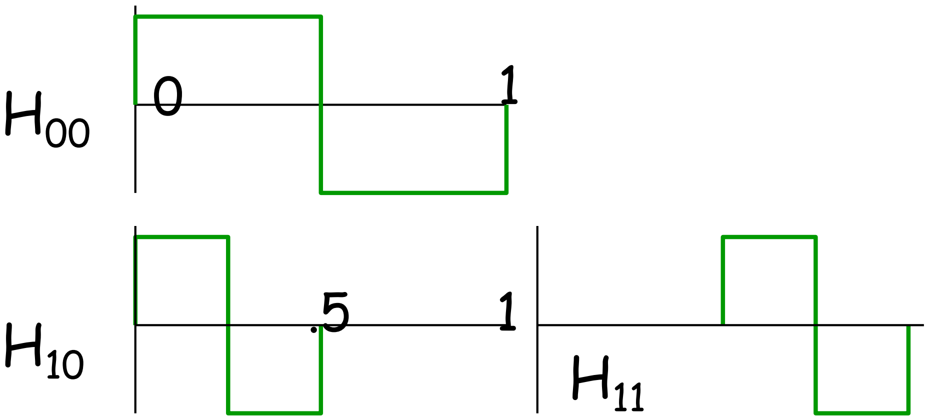

$H_{sl}(x) = \varphi(2^s x - l)$

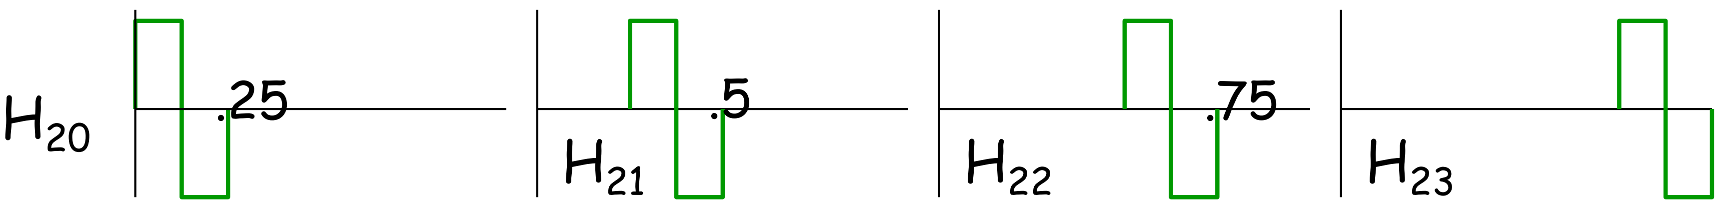

H_{k0} …

+ DC component = 2^{k+1} components

Haar Wavelet in 2d

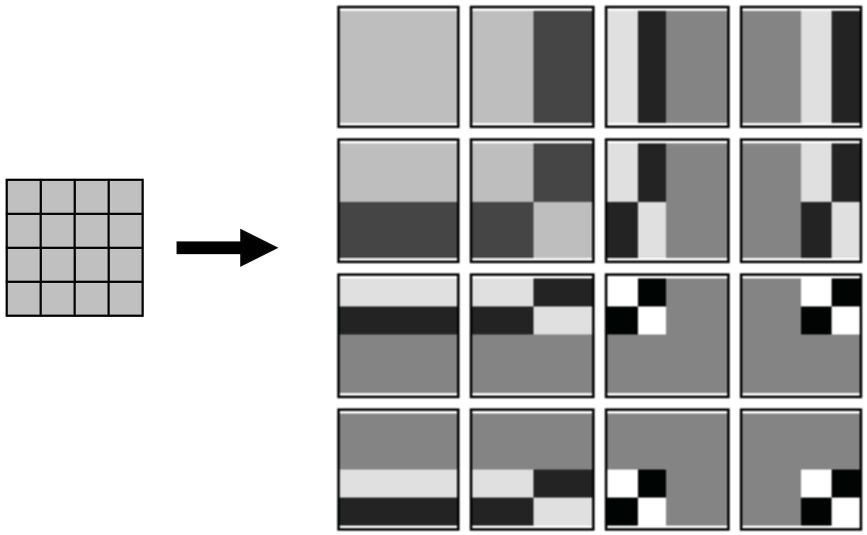

Discrete Haar Wavelet Transform

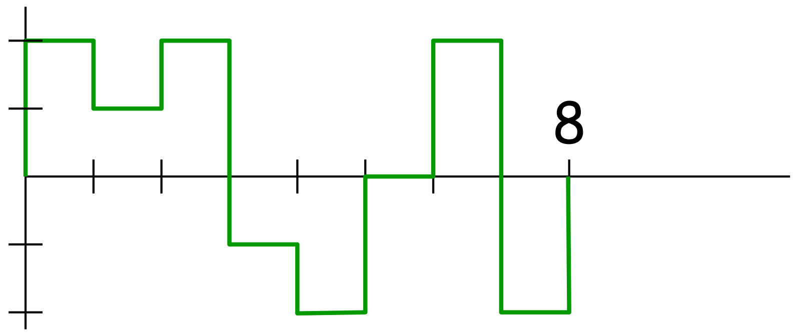

How do we convert this to the wavelet coefficients?

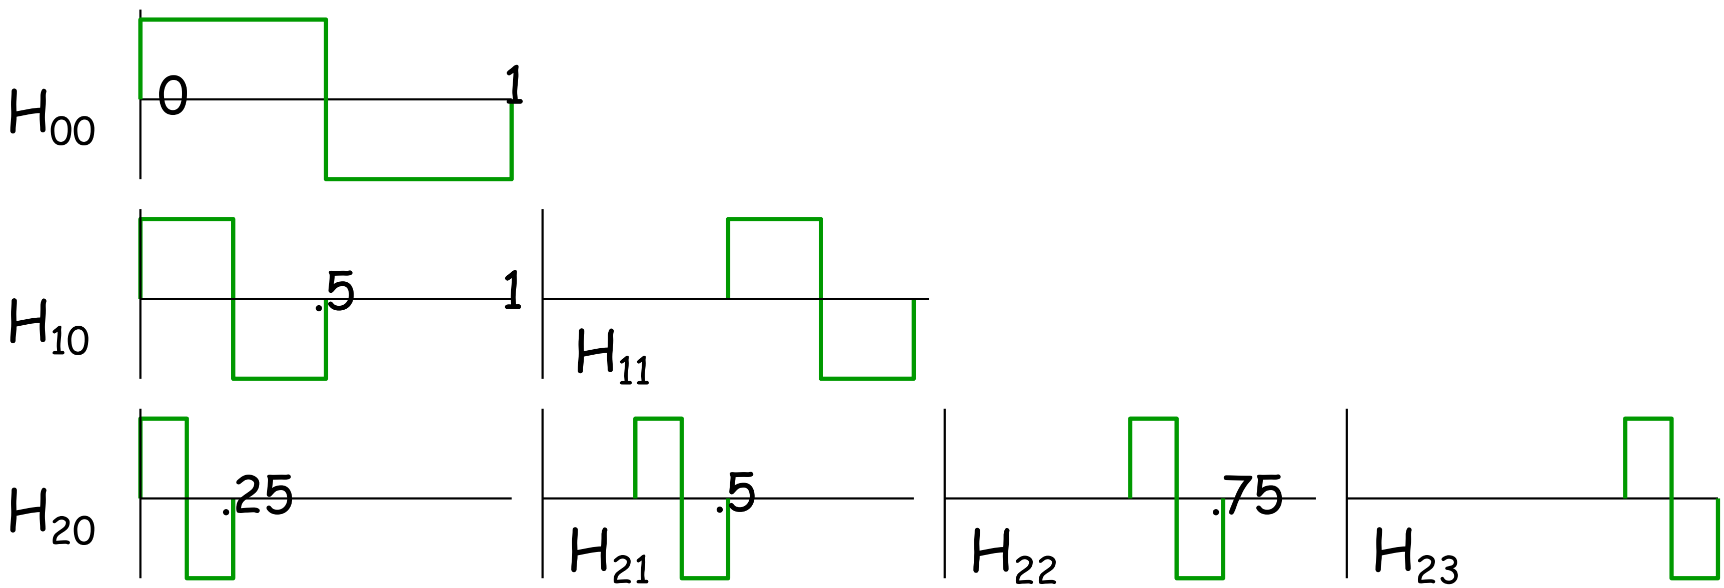

Discrete Haar Wavelet Transform

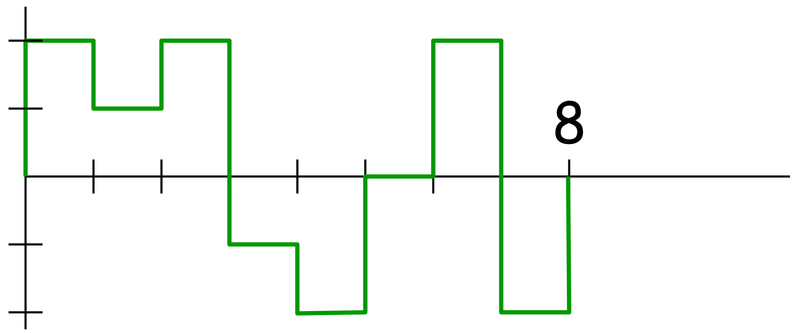

How do we convert this to the wavelet coefficients?

```
for (j = n/2; j >= 1; j = j/2) {
  for (i = 1; i < j; i++) {
    b[i] = (a[2i-1] + a[2i])/2;
    b[j+i] = (a[2i-1] - a[2i])/2; }
  a[1..2*j] = b[1..2*j]; }
```

Averages

Differences

Linear time!

Haar Wavelet Transform: example

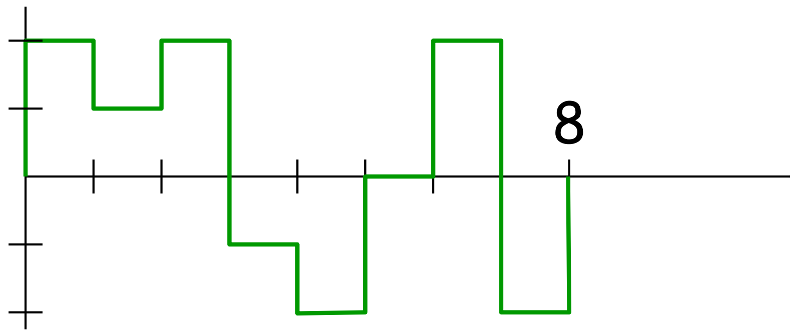

a	=	2	1	2	-1	-2	0	2	-2
	=	1.5	.5	-1	0	.5	1.5	-1	2
	=	1	-.5	.5	-.5				
	=	.25	.75						
a	=	.25	.75	.5	.5	.5	1.5	-1	2

Wavelet decomposition

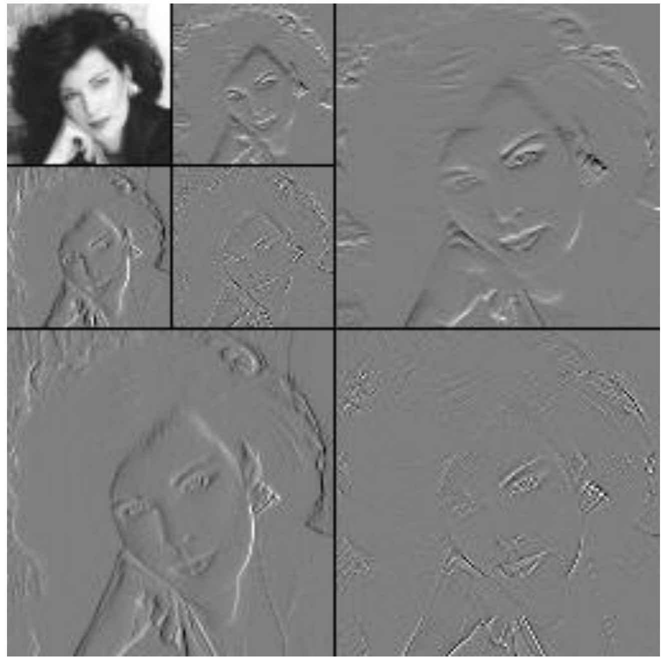

Morlet Wavelet

$\phi(x)$ = Gaussian x Cosine = $e^{-\left(x^2/2\right)}\cos(5x)$

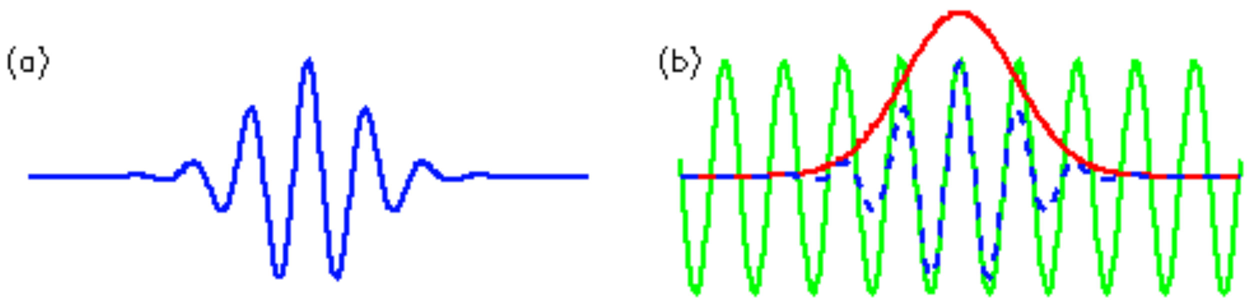

Corresponds to wavepackets in physics.

Daubechies Wavelet

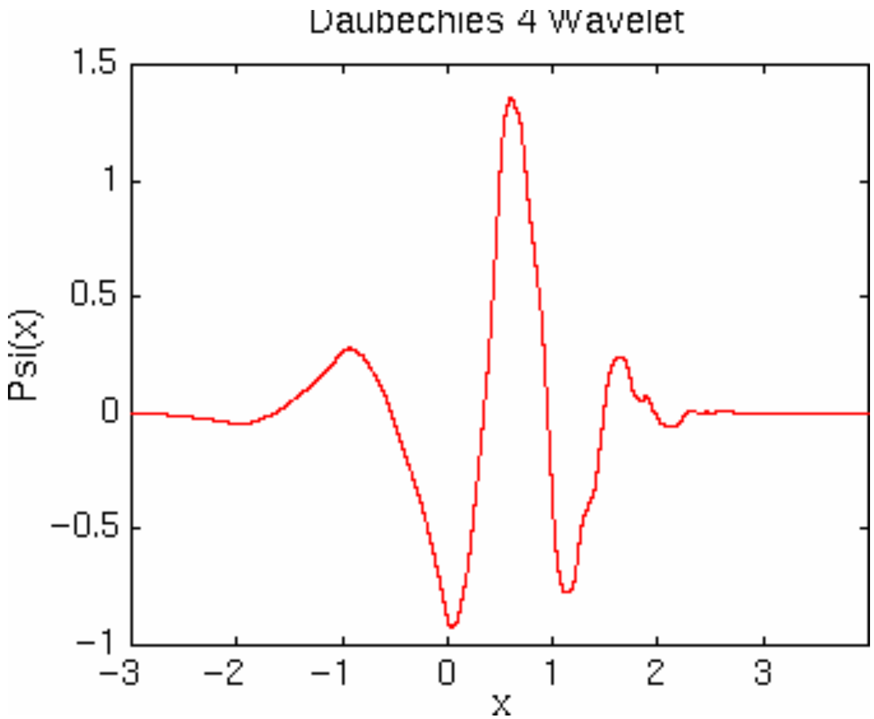

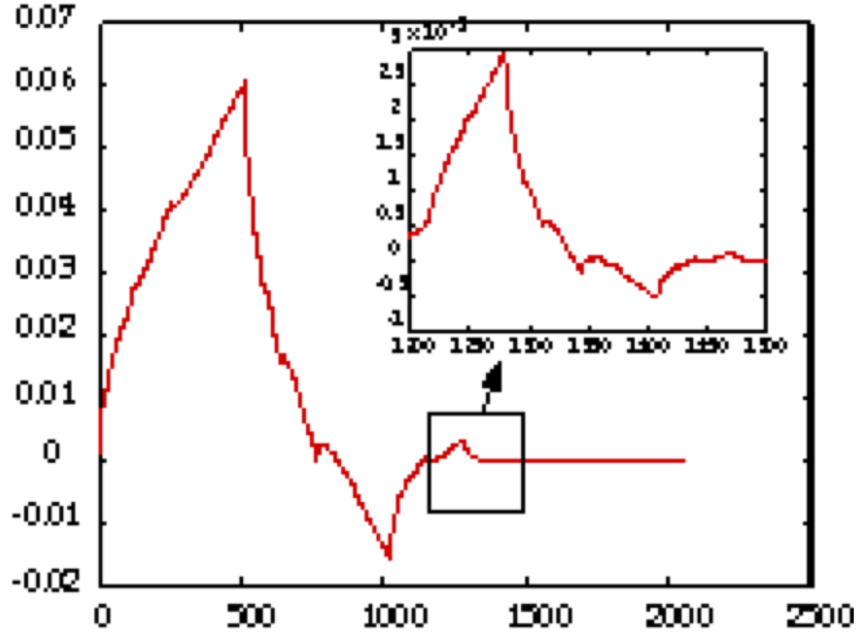

JPEG2000

Overall Goals:

- High compression efficiency with good quality at compression ratios of .25bpp
- Handle large images (up to $2^{32} \times 2^{32}$)
- Progressive image transmission
  - Quality, resolution or region of interest
- Fast access to various points in compressed stream
- Pan and Zoom while only decompressing parts
- Error resilience

Also used in Dirac video compression

JPEG2000: Outline

Main similarities with JPEG

- Separates into Y, I, Q color planes, and can downsample the I and Q planes
- Transform coding

Main differences with JPEG

- Wavelet transform
  - Daubechies 9-tap/7-tap (irreversible)
  - Daubechies 5-tap/3-tap (reversible)
- Many levels of hierarchy (resolution and spatial)
- Only arithmetic coding

JPEG2000: 5-tap/3-tap

```
h[i] = a[2i-1] - (a[2i] + a[2i-2])/2;
l[i] = a[2i] + (h[i-1] + h[i] + 2)/2;
```

`h[i]`: is the "high pass" filter, ie, the **differences**
it depends on 3 values from a (3-tap)

`l[i]`: is the "low pass" filter, ie, the **averages**
it depends on 5 values from a (5-tap)

Need to deal with boundary effects.

This is reversible: assignment

JPEG 2000: Outline

A spatial and resolution hierarchy

- **Tiles:** Makes it easy to decode sections of an image. For our purposes we can imagine the whole image as one tile.
- **Resolution Levels:** These are based on the wavelet transform. High-detail vs. Low detail.
- **Precinct Partitions:** Used within each resolution level to represent a region of space.
- **Code Blocks:** blocks within a precinct
- **Bit Planes:** ordering of significance of the bits

JPEG2000: Precincts

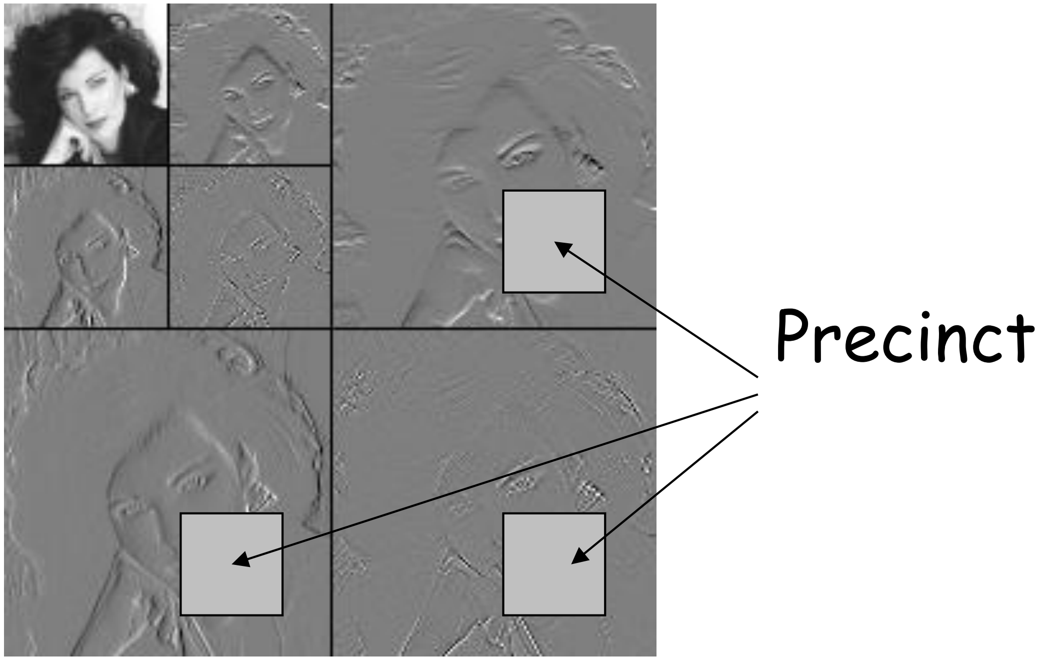

JPEG vs. JPEG2000

JPEG: .125bpp

JPEG2000: .125bpp

Compression Summary

Compression is all about **probabilities**

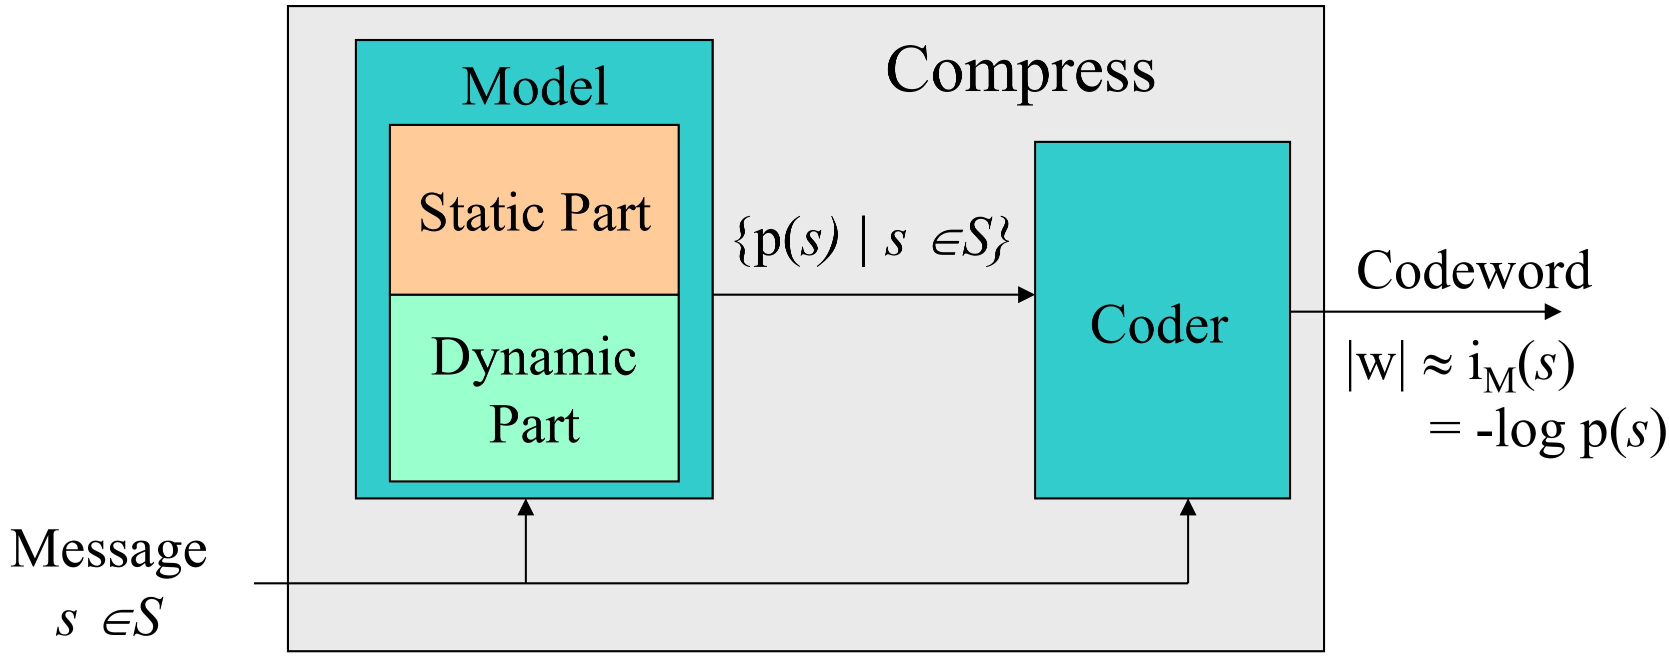

We want the model to skew the probabilities as much as possible (*i.e.*, decrease the **entropy**)

Compression Summary

How do we figure out the probabilities

- Transformations that skew them
  - Guess value and code difference
  - Move to front for temporal locality
  - Run-length
  - Linear transforms (Cosine, Wavelet)
  - Renumber (graph compression)
- Conditional probabilities
  - Neighboring context

In practice one almost always uses a combination of techniques