

# 4F8 Image Coding Course

Nick Kingsbury

February 9, 2016

## Contents

<b>1</b>	<b>Vision and Image Characteristics useful for Compression</b>	<b>3</b>
1.1	Introduction . . . . .	3
1.2	Human Vision . . . . .	3
1.3	Image Characteristics . . . . .	8
<b>2</b>	<b>A Basic Image Compression Example</b>	<b>11</b>
2.1	The Haar transform . . . . .	12
2.2	Entropy . . . . .	14
2.3	The Multi-level Haar Transform . . . . .	17
2.4	Use of Laplacian PDFs in Image Compression . . . . .	23
2.5	Practical Entropy Coding Techniques . . . . .	27
<b>3</b>	<b>The DCT and the JPEG Standard</b>	<b>33</b>
3.1	The Discrete Cosine Transform (DCT) . . . . .	33
3.2	Fast algorithms for the DCT . . . . .	36
3.3	The 2-dimensional DCT . . . . .	37
3.4	Quantisation of DCT Coefficients . . . . .	43
3.5	JPEG Entropy Coding . . . . .	46
3.6	Sync and Headers . . . . .	53

<b>4</b>	<b>Filter Banks and Wavelets</b>	<b>55</b>
4.1	The 2-band Filter Bank . . . . .	55
4.2	Perfect Reconstruction (PR) . . . . .	57
4.3	The binary filter tree . . . . .	58
4.4	Wavelets . . . . .	62
4.5	Good filters / wavelets . . . . .	63
4.6	The 2-D DWT . . . . .	74
4.7	Compression properties of wavelets . . . . .	78
<b>5</b>	<b>Video Compression and Motion Processing</b>	<b>83</b>
5.1	Motion-Compensated Predictive Coding . . . . .	83
5.2	Motion Estimation . . . . .	85
5.3	The MPEG Standard . . . . .	86

### Recommended Textbooks

General Image Processing (with chapters on transforms and image compression):

- A K Jain, [Fundamentals of Digital Image Processing](#), Prentice-Hall, 1989.
- R C Gonzalez and R E Woods, [Digital Image Processing](#), Addison Wesley, 1992.
- M Petrou and P Bosdogianni, [Image Processing: The Fundamentals](#), John Wiley, 1999.

Specific to Image Compression or Wavelets:

- W B Pennebaker and J L Mitchell, [The JPEG Book](#), Van Nostrand Reinhold, 1993.
- M Ghanbari, [Video Coding: An Introduction to Standard Codecs](#), IEE, 1999.
- M Vetterli and J Kovacevic, [Wavelets and Subband Coding](#), Prentice Hall PTR, 1995.
- S Mallat, [A Wavelet Tour of Signal Processing](#), Academic Press, 1998 (2nd ed: 2000). (Rather mathematical.)
- Iain E G Richardson, [H.264 and MPEG-4 Video Compression](#), John Wiley, 2003. (Advanced – only for the really keen who want to know about standards for video beyond MPEG-2.)
- C Tu, S Srinivasan, G Sullivan, S Regunathan, and H Malvar, [Low-complexity hierarchical lapped transform for lossy-to-lossless image coding in JPEG XR / HD Photo](#), Proc SPIE, Vol 7073, 70730C, (2008), Applications of Digital Image Processing XXXI. (Good paper on JPEG XR transforms and *Lifting*.)

# 1 Vision and Image Characteristics useful for Compression

## 1.1 Introduction

Image Coding (often more correctly known as Image Compression) is the art / science of representing images with the **least information** (no. of bits) consistent with achieving an **acceptable image quality / usefulness**.

In order to do this, we try to take advantage of:

1. **Physiological characteristics** of human vision;
2. **Statistical characteristics** of typical images;
3. Efficient binary **source coding** methods.

We shall consider these in turn.

## 1.2 Human Vision

### Colours

The human vision system perceives images in colour using receptors on the retina of the eye which respond to three relatively broad colour bands in the regions of red, green and blue (RGB) in the colour spectrum (red, orange, yellow, green, blue, indigo, violet).

Colours in between these are perceived as different linear combinations of RGB. Hence colour TVs and monitors can form almost any perceivable colour by controlling the relative intensities of R, G and B light sources. Thus most colour images which exist in electronic form are fundamentally represented by 3 intensities (R, G and B) at each picture element (pixel or pel) position.

The numerical values used for these intensities are usually chosen such that equal increments in value result in approximately equal apparent increases in brightness. In practice this means that the numerical value is approximately proportional to the log of the true light intensity (energy of the wave) – this is Weber's Law. Throughout this course, we shall refer to these numerical values as intensities, since for compression it is most convenient to use a subjectively linear scale.

## The YUV Colour Space

The eye is much more sensitive to overall intensity (luminance) changes than to colour changes. Usually most of the information about a scene is contained in its luminance rather than its colour (chrominance).

This is why black-and-white (monochrome) reproduction was acceptable for photography and TV for many years until technology provided colour reproduction at a sufficiently cheap price to make its modest advantages worth having.

The luminance ( $Y$ ) of a pel may be obtained from its RGB components as:

$$Y = 0.3R + 0.6G + 0.1B \quad (1.1)$$

These coefficients are only approximate, and are the values defined in the JPEG Book. In other places values of 0.3, 0.59 and 0.11 are used.

RGB representations of images are normally defined so that if  $R = G = B$ , the pel is always some shade of grey, and if  $Y = R = G = B$  in these cases, the 3 coefficients in (1.1) must sum to unity (which they do).

When  $Y$  defines the luminance of a pel, its chrominance is usually defined by  $U$  and  $V$  such that:

$$\begin{aligned} U &= 0.5(B - Y) \\ V &= 0.625(R - Y) \end{aligned} \quad (1.2)$$

Note that grey pels will always have  $U = V = 0$ .

The transformation between RGB and YUV colour spaces is linear and may be achieved by a  $3 \times 3$  matrix  $\mathbf{C}$  and its inverse:

$$\begin{bmatrix} Y \\ U \\ V \end{bmatrix} = \mathbf{C} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \quad \text{where} \quad \mathbf{C} = \begin{bmatrix} 0.3 & 0.6 & 0.1 \\ -0.15 & -0.3 & 0.45 \\ 0.4375 & -0.3750 & -0.0625 \end{bmatrix} \quad (1.3)$$

and

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \mathbf{C}^{-1} \begin{bmatrix} Y \\ U \\ V \end{bmatrix} \quad \text{where} \quad \mathbf{C}^{-1} = \begin{bmatrix} 1 & 0 & 1.6 \\ 1 & -0.3333 & -0.8 \\ 1 & 2 & 0 \end{bmatrix} \quad (1.4)$$

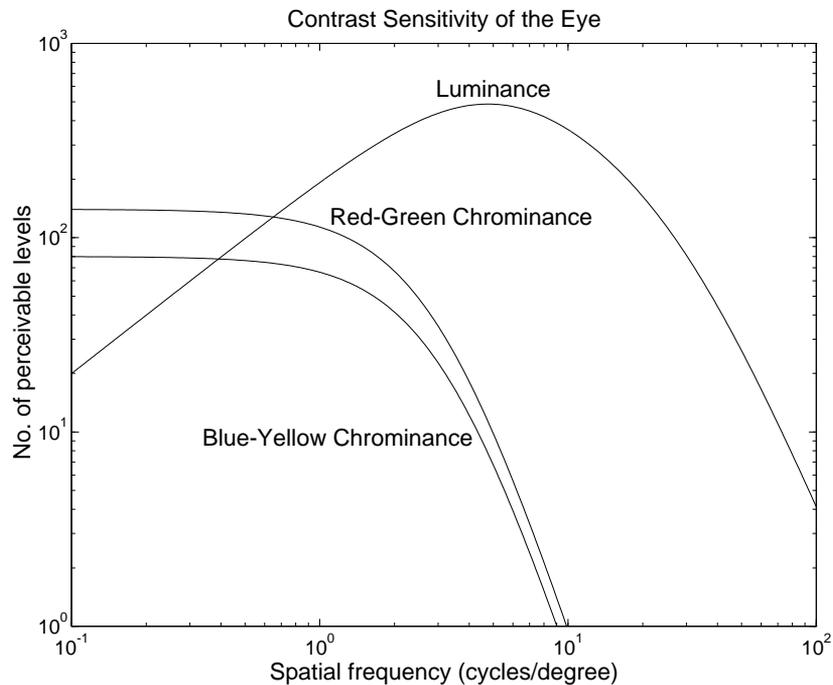


Fig 1.1: Sensitivity of the eye to luminance and chrominance intensity changes.

### Visual sensitivity

Fig 1.1 shows the sensitivity of the eye to luminance (Y) and chrominance (U,V) components of images. The horizontal scale is spatial frequency, and represents the frequency of an alternating pattern of parallel stripes with sinusoidally varying intensity. The vertical scale is the contrast sensitivity of human vision, which is the ratio of the maximum visible range of intensities to the minimum discernable peak-to-peak intensity variation at the specified frequency.

In fig 1.1 we see that:

- The maximum sensitivity to Y occurs for spatial frequencies around 5 cycles / degree, which corresponds to striped patterns with a half-period (stripe width) of 1.8 mm at a distance of 1 m ( $\sim$  arm's length).
- The eye has very little response above 100 cycles / degree, which corresponds to a stripe width of about 0.1 mm at 1 m. A modern flat-panel 1080p High-Definition computer monitor with a width of 500mm and  $1920 \times 1080$  pixels will have a pixel size of  $500/1920 = 0.26\text{mm}$ ; so these displays still fall a little short of the ideal for people with good eyesight. However despite the fact that modern flat-panel displays have resolutions that are slightly coarser than was possible with the older CRT technology, in general they are more pleasing and less tiring to view because the pixel edges are

so sharp and there is no flicker. Note that ultra-high-definition (UHD) displays are now available with 4K pixels across and these get quite close to the limit of human vision unless the displays are very large (as many UHD televisions are!).

- The sensitivity to luminance drops off at low spatial frequencies, showing that we are not very good at estimating absolute luminance levels **as long as they do not change with time** – the luminance sensitivity to temporal fluctuations (flicker) does not fall off at low spatial frequencies.
- The maximum chrominance sensitivity is much lower than the maximum luminance sensitivity, with blue-yellow (U) sensitivity being about half of red-green (V) sensitivity and about  $\frac{1}{6}$  of the maximum luminance sensitivity.
- The chrominance sensitivities fall off above 1 cycle / degree, requiring a much lower spatial bandwidth than luminance.

We can now see why it is better to convert to the YUV domain before attempting image compression. The U and V components may be sampled at a lower rate than Y (due to narrower bandwidth) and may be quantised more coarsely (due to lower contrast sensitivity).

A colour demonstration on the computer will show this effect.

## Colour compression strategy

The 3 RGB samples at each pel are transformed into 3 YUV samples using (1.3).

Most image compression systems then subsample the U and V information by 2:1 horizontally and vertically so that there is one U and one V pel for each  $2 \times 2$  block of Y pels. The subsampled U and V pels are obtained by averaging the four U and V samples, from (1.3). The quarter-size U and V subimages are then compressed using the same techniques as the full-size Y image, except that coarser quantisation may be used for U and V, so the total cost of adding colour may only be about 25% increase in bit rate. Sometimes U and V are subsampled 4:1 each way (16:1 total), giving an even lower cost of colour.

From now on we will mostly be considering compression of the monochrome Y image, and assume that similar techniques will be used for the smaller U and V subimages.

## Activity masking

A final feature of human vision, which is useful for compression, is that the contrast sensitivity to a given pattern is reduced in the presence of other patterns (activity) in the same region. This is known as activity masking.

It is a complicated subject as it depends on the similarity between the given pattern and the background activity. However in general, the higher the variance of the pels in a given region (typically  $\sim 8$  to 16 pels across), the lower is the contrast sensitivity.

Hence compression schemes which adapt the quantisation to local image activity tend to perform better than those which use uniform quantisation.

A computer demonstration will show the effect of reduced sensitivity to quantisation effects when noise is added to an image.



Fig 1.2:  $256 \times 256$  pel 'Lenna' image and its Fourier log power spectrum.

### 1.3 Image Characteristics

We now consider statistical characteristics of typical images which can permit compression. If all images comprised dots with uncorrelated random intensities, then each pel would need to be coded independently and we could not achieve any useful gains. However typical images are very different from random dot patterns and significant compression gains are possible.

Some compression can be achieved even if no additional distortion is permitted (**lossless coding**) but much greater compression is possible if some additional distortion is allowed (**lossy coding**). Lossy coding is the main topic of this course but we try to keep the added distortions near or below the human visual sensitivity thresholds discussed above.

Statistical characteristics of signals can often be most readily appreciated by frequency domain analysis since the power spectrum is the Fourier transform of the autocorrelation function. The 2-D FFT is a convenient tool for analysing images. Fig 1.2 shows the  $256 \times 256$  pel 'Lenna' image and its Fourier log power spectrum. Zero frequency is at the centre of the spectrum image and the log scale shows the lower spectral components much more clearly.

The bright region near the centre of the spectrum shows that the main concentration of image energy is at low frequencies, which implies strong correlation between nearby pels and is typical of real-world images. The diagonal line of spectral energy at about  $-30^\circ$  is due to the strong diagonal edges of the hat normal to this direction. Similarly the near-horizontal spectral line comes from the strong near-vertical stripe of hair to the right of the face. Any other features are difficult to distinguish in this global spectrum.

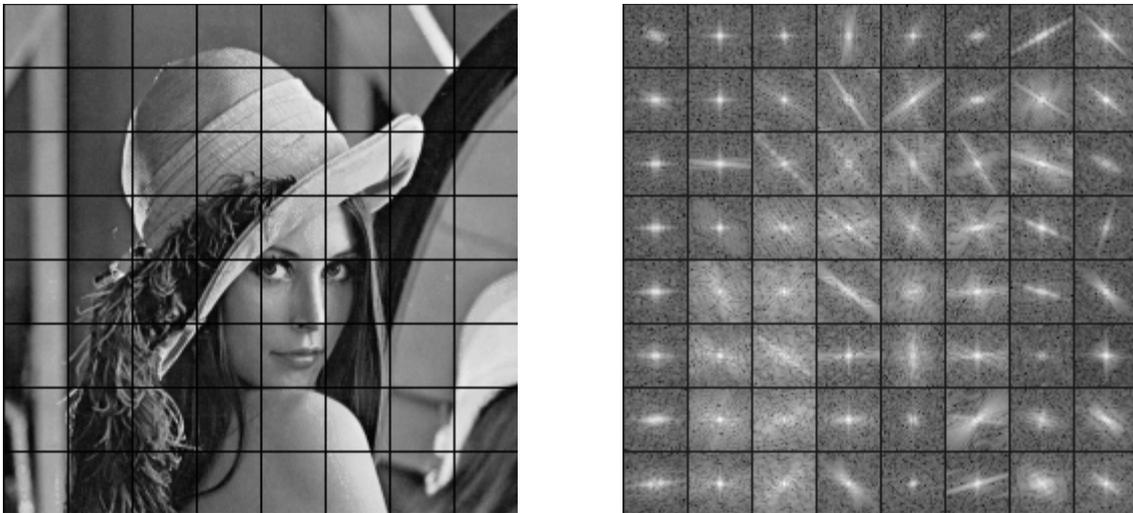


Fig 1.3: Fourier log power spectra of 'Lenna' image split into 64 blocks of  $32 \times 32$  pels.

A key property of real-world images is that their statistics are not stationary over the image. Fig 1.3 demonstrates this by splitting the 'Lenna' image into 64 blocks of  $32 \times 32$  pels, and calculating the Fourier log power spectrum of each block. The wide variation in spectra is clearly seen. Blocks with dominant edge directions produce spectra with lines normal to the edges, and those containing the feathers of the hat generate a broad spread of energy at all frequencies. However a bright centre, indicating dominant low frequency components, is common to all blocks.

We conclude that in many regions of a typical image, most of the signal energy is contained in a relatively small number of spectral components, many of which are at low frequencies. However, between regions, the location of the main components changes significantly.

**The concentration of spectral energy is the key to compression. For example, if a signal can be reconstructed from its Fourier transform, and many of the transform coefficients are very small, then a close approximation to the original can be reconstructed from just the larger transform coefficients, so only these coefficients need be transmitted.**

In practice, the Fourier transform is not very suitable for compression because it generates complex coefficients and it is badly affected by discontinuities at block boundaries (half-sine windowing was used in figs 1.2 and 1.3 to reduce boundary effects but this would prevent proper reconstruction of the image). In the next section, we demonstrate the principles of image compression using the Haar transform, perhaps the simplest of all transforms.



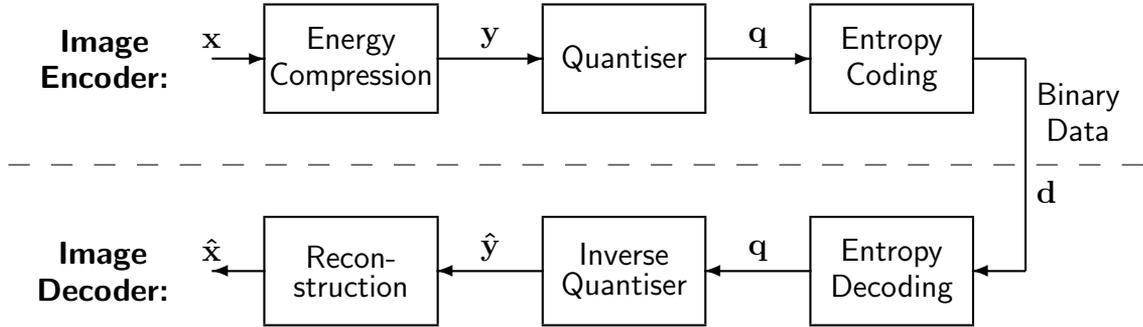


Fig 2.1: The basic block diagram of an image coding system.

## 2 A Basic Image Compression Example

We shall represent a monochrome (luminance) image by a matrix  $\mathbf{x}$  whose elements are  $x(\mathbf{n})$ , where  $\mathbf{n} = [n_1, n_2]$  is the integer vector of row and column indices. The energy of  $\mathbf{x}$  is defined as:

$$\text{Energy of } \mathbf{x} = \sum_{\mathbf{n}} x(\mathbf{n})^2 \quad (2.1)$$

where the sum is performed over all  $\mathbf{n}$  in  $\mathbf{x}$ .

Fig 2.1 shows the main blocks in any image coding system. The decoder is the inverse of the encoder. The three encoder blocks perform the following tasks:

**Energy compression** – This is usually a transformation or filtering process which aims to concentrate a high proportion of the energy of the image  $\mathbf{x}$  into as few samples (coefficients) of  $\mathbf{y}$  as possible while preserving the total energy of  $\mathbf{x}$  in  $\mathbf{y}$ . This minimises the number of non-zero samples of  $\mathbf{y}$  which need to be transmitted for a given level of distortion in the reconstructed image  $\hat{\mathbf{x}}$ .

**Quantisation** – This represents the samples of  $\mathbf{y}$  to a given level of accuracy in the integer matrix  $\mathbf{q}$ . The quantiser step size controls the tradeoff between distortion and bit rate and may be adapted to take account of human visual sensitivities. The inverse quantiser reconstructs  $\hat{\mathbf{y}}$ , the best estimate of  $\mathbf{y}$  from  $\mathbf{q}$ .

**Entropy coding** – This encodes the integers in  $\mathbf{q}$  into a serial bit stream  $\mathbf{d}$ , using variable-length entropy codes which attempt to minimise the total number of bits in  $\mathbf{d}$ , based on the statistics (PDFs) of various classes of samples in  $\mathbf{q}$ .

The energy compression / reconstruction and the entropy coding / decoding processes are normally all lossless. Only the quantiser introduces loss and distortion:  $\hat{\mathbf{y}}$  is a distorted version of  $\mathbf{y}$ , and hence  $\hat{\mathbf{x}}$  is a distorted version of  $\mathbf{x}$ . In the absence of quantisation, if  $\hat{\mathbf{y}} = \mathbf{y}$ , then  $\hat{\mathbf{x}} = \mathbf{x}$ .



Fig 2.2: Original (a) and the Level 1 Haar transform (b) of the 'Lenna' image.

## 2.1 The Haar transform

Probably the simplest useful energy compression process is the Haar transform. In 1-dimension, this transforms a 2-element vector  $[x(1), x(2)]^T$  into  $[y(1), y(2)]^T$  using:

$$\begin{bmatrix} y(1) \\ y(2) \end{bmatrix} = \mathbf{T} \begin{bmatrix} x(1) \\ x(2) \end{bmatrix} \quad \text{where} \quad \mathbf{T} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \quad (2.2)$$

Thus  $y(1)$  and  $y(2)$  are simply the sum and difference of  $x(1)$  and  $x(2)$ , scaled by  $1/\sqrt{2}$  to preserve energy.

Note that  $\mathbf{T}$  is an orthonormal matrix because its rows are orthogonal to each other (their dot products are zero) and they are normalised to unit magnitude. Therefore  $\mathbf{T}^{-1} = \mathbf{T}^T$ . (In this case  $\mathbf{T}$  is symmetric so  $\mathbf{T}^T = \mathbf{T}$ .) Hence we may recover  $\mathbf{x}$  from  $\mathbf{y}$  using:

$$\begin{bmatrix} x(1) \\ x(2) \end{bmatrix} = \mathbf{T}^T \begin{bmatrix} y(1) \\ y(2) \end{bmatrix} \quad (2.3)$$

An important property of orthonormal transformations is that they preserve the energy of  $\mathbf{x}$ , when it is transformed into  $\mathbf{y} = \mathbf{T}\mathbf{x}$ . We see this because  $\mathbf{T}^T\mathbf{T} = \mathbf{I}$  and so:

$$(\text{Energy of } \mathbf{y}) = \mathbf{y}^T \mathbf{y} = \mathbf{x}^T \mathbf{T}^T \mathbf{T} \mathbf{x} = \mathbf{x}^T \mathbf{I} \mathbf{x} = \mathbf{x}^T \mathbf{x} = (\text{Energy of } \mathbf{x}) \quad (2.4)$$

In 2-dimensions  $\mathbf{x}$  and  $\mathbf{y}$  become  $2 \times 2$  matrices. We may transform first the columns of  $\mathbf{x}$ , by premultiplying by  $\mathbf{T}$ , and then the rows of the result by postmultiplying by  $\mathbf{T}^T$ . Hence:

$$\mathbf{y} = \mathbf{T} \mathbf{x} \mathbf{T}^T \quad \text{and to invert:} \quad \mathbf{x} = \mathbf{T}^T \mathbf{y} \mathbf{T} \quad (2.5)$$

To show more clearly what is happening:

$$\text{If } \mathbf{x} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \quad \text{then } \mathbf{y} = \frac{1}{2} \begin{bmatrix} a + b + c + d & a - b + c - d \\ a + b - c - d & a - b - c + d \end{bmatrix}$$

These operations correspond to the following filtering processes:

**Top left:**  $a + b + c + d =$  4-point average or 2-D lowpass (Lo-Lo) filter.

**Top right:**  $a - b + c - d =$  Average horizontal gradient or horizontal highpass and vertical lowpass (Hi-Lo) filter.

**Lower left:**  $a + b - c - d =$  Average vertical gradient or horizontal lowpass and vertical highpass (Lo-Hi) filter.

**Lower right:**  $a - b - c + d =$  Diagonal curvature or 2-D highpass (Hi-Hi) filter.

To apply this transform to a complete image, we group the pels into  $2 \times 2$  blocks and apply (2.5) to each block. The result (after reordering) is shown in fig 2.2b. To view the result sensibly, we have grouped all the top left components of the  $2 \times 2$  blocks in  $\mathbf{y}$  together to form the top left subimage in fig 2.2b, and done the same for the components in the other 3 positions to form the corresponding other 3 subimages.

It is clear from fig 2.2b that most of the energy is contained in the top left (Lo-Lo) subimage and the least energy is in the lower right (Hi-Hi) subimage. Note how the top right (Hi-Lo) subimage contains the near-vertical edges and the lower left (Lo-Hi) subimage contains the near-horizontal edges.

Energy is conserved by the Haar transform (eq. (2.4)), and the energies of the subimages and their percentages of the total input energy are:

<b>Lo-Lo:</b>	<b>Hi-Lo:</b>
201.73 . $10^6$	4.56 . $10^6$
96.5%	2.2%
<b>Lo-Hi:</b>	<b>Hi-Hi:</b>
1.89 . $10^6$	0.82 . $10^6$
0.9%	0.4%

Total energy in figs 2.2a and 2.2b = 208.99 .  $10^6$

We see that a significant compression of energy into the Lo-Lo subimage has been achieved. However the energy measurements do not tell us directly how much data compression this gives.

## 2.2 Entropy

A much more useful measure than energy is the **entropy** of the subimages after a given amount of quantisation. This gives the minimum number of bits per pel needed to represent the quantised data for each subimage, to a given accuracy, assuming that we use an ideal entropy code. By comparing the total entropy of the 4 subimages with that of the original image, we can estimate the compression that one level of the Haar transform can provide.

Entropy of source information was discussed in the third-year 3F1 Information Theory course. For an image  $\mathbf{x}$ , quantised to  $M$  levels, the entropy  $H_{\mathbf{x}}$  is defined as:

$$H_{\mathbf{x}} = \sum_{i=0}^{M-1} p_i \log_2 \left( \frac{1}{p_i} \right) = - \sum_{i=0}^{M-1} p_i \log_2(p_i) \quad (2.6)$$

where  $p_i$ ,  $i = 0$  to  $M - 1$ , is the probability of the  $i^{\text{th}}$  quantiser level being used (often obtained from a histogram of the pel intensities).

$H_{\mathbf{x}}$  represents the mean number of bits per pel with which the quantised image  $\mathbf{x}$  can be represented using an ideal variable-length entropy code. A Huffman code usually approximates this bit-rate quite closely.

To obtain the number of bits to code an image (or subimage)  $\mathbf{x}$  containing  $N$  pels:

- A histogram of  $\mathbf{x}$  is measured using  $M$  bins corresponding to the  $M$  quantiser levels.
- The  $M$  histogram counts are each divided by  $N$  to give the probabilities  $p_i$ , which are then converted into entropies  $h_i = -p_i \log_2(p_i)$ . This conversion law is illustrated in fig 2.3 and shows that probabilities close to zero or one produce low entropy and intermediate values produce entropies near 0.5.
- The entropies  $h_i$  of the separate quantiser levels are summed to give the total entropy  $H_{\mathbf{x}}$  for the subimage.
- Multiplying  $H_{\mathbf{x}}$  by  $N$  gives the estimated total number of bits needed to code  $\mathbf{x}$ , assuming an ideal entropy code is available which is matched to the histogram of  $\mathbf{x}$ .

Fig 2.4 shows the probabilities  $p_i$  and entropies  $h_i$  for the original Lenna image and fig 2.5 shows these for each of the four subimages in fig 2.2, assuming a uniform quantiser with a step-size  $Q_{step} = 15$  in each case. The original Lenna image contained pel values from 3 to 238 and a mean level of 120 was subtracted from each pel value before the image was analysed or transformed in order that all samples would be approximately evenly distributed about zero (a natural feature of highpass subimages).

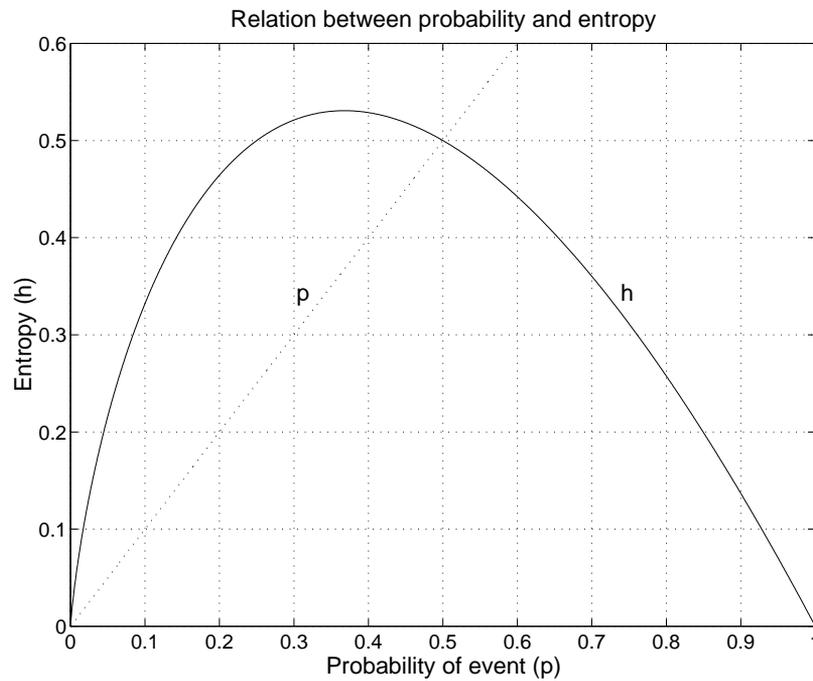


Fig 2.3: Conversion from probability  $p_i$  to entropy  $h_i = -p_i \log_2(p_i)$ .

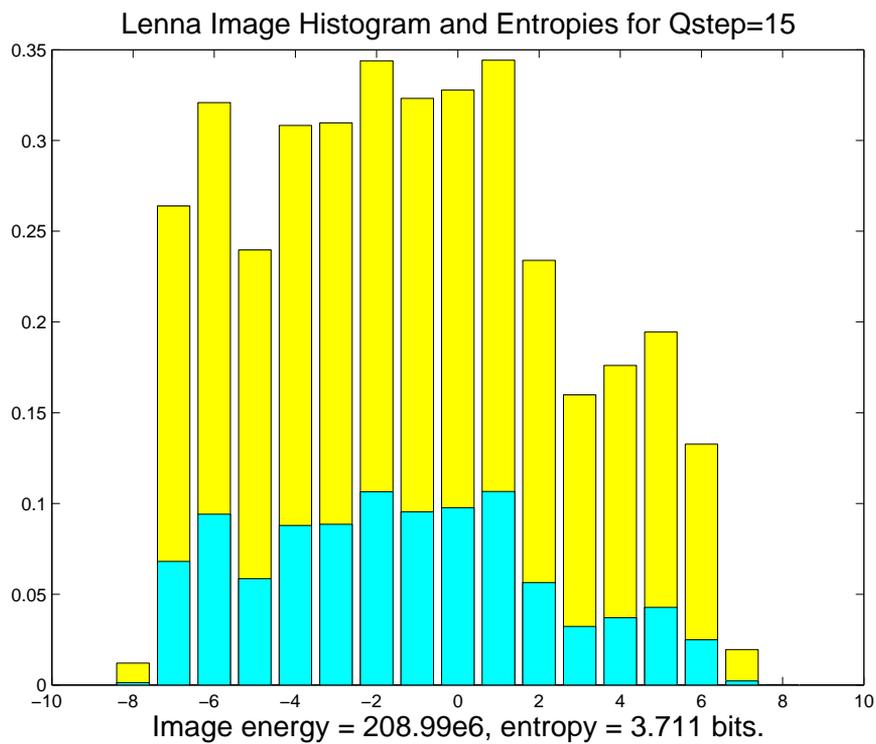


Fig 2.4: Probability histogram (dashed) and entropies (solid) of the Lenna image (fig 2.2a).

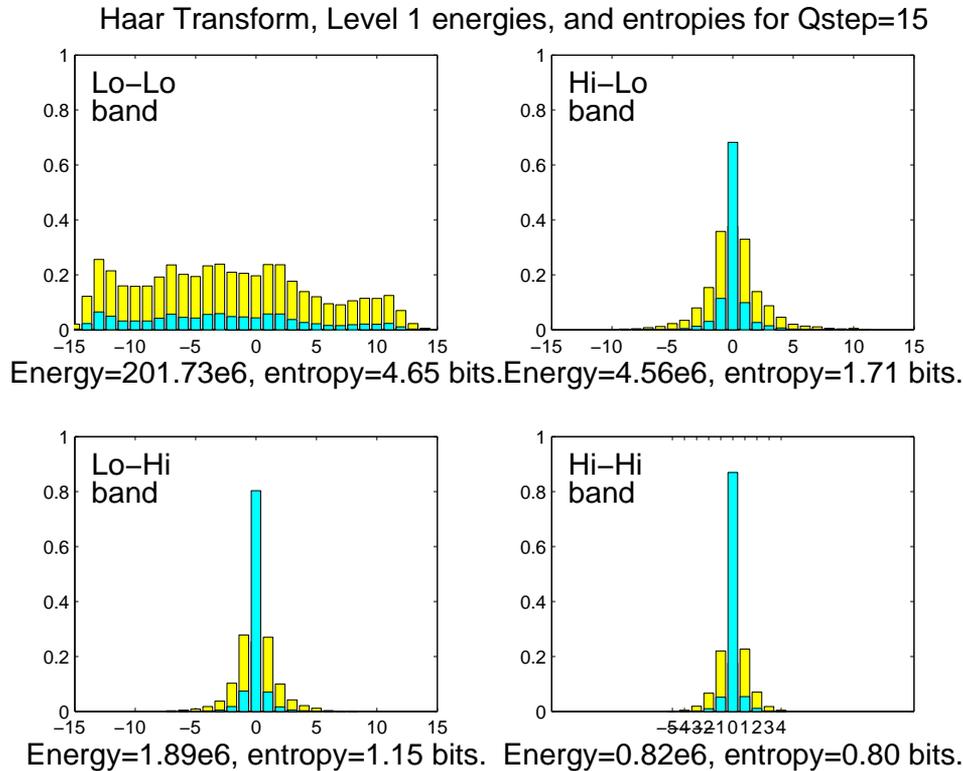


Fig 2.5: Probability histograms (dashed) and entropies (solid) of the four subimages of the Level 1 Haar transform of Lenna (fig 2.2b).

The Haar transform preserves energy and so the expected distortion energy from quantising the transformed image  $\mathbf{y}$  with a given step size  $Q_{step}$  will be approximately the same as that from quantising the input image  $\mathbf{x}$  with the same step size. This is because quantising errors can usually be modelled as independent random processes with variance (energy)  $= Q_{step}^2/12$  and the total squared quantising error (distortion) will tend to the sum of the variances over all pels. This applies whether the error energies are summed before or after the inverse transform (reconstruction) in the decoder.

**Hence equal quantiser step sizes before and after an energy-preserving transformation should generate equivalent quantising distortions and provide a fair estimate of the compression achieved by the transformation.**

The first two columns of Fig 2.6 (**original** and **level 1**) compare the entropy (mean bit rate) per pel for the original image (3.71 bit / pel) with that of the Haar transformed image of fig 2.2 (2.08 bit / pel), using  $Q_{step} = 15$ . Notice that the entropy of the original image is almost as great as the 4 bit / pel that would be needed to code the 16 levels using a simple fixed-length code, because the histogram is relatively uniform.

The level 1 column of fig 2.6 shows the contributions of each of the four subimages of fig 2.2b to the total entropy per pel (the entropies from fig 2.5 have been divided by 4 since each

subimage has one quarter of the total number of pels). The Lo-Lo subimage contributes 56% to the total entropy (bit rate) and has similar spatial correlations to the original image. Hence it is a logical step to apply the Haar transform again to this subimage.

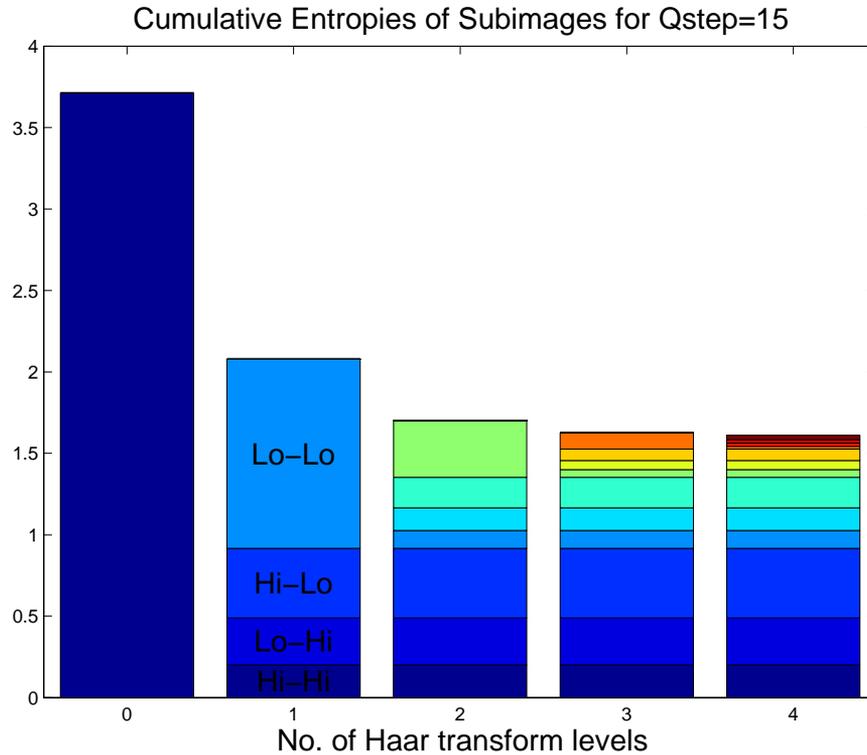


Fig 2.6: Mean bit rate for the original Lenna image and for the Haar transforms of the image after 1 to 4 levels, using a quantiser step size  $Q_{step} = 15$ .

## 2.3 The Multi-level Haar Transform

Fig 2.7a shows the result of applying the Haar transform to the Lo-Lo subimage of fig 2.2b and fig 2.8 shows the probabilities  $p_i$  and entropies  $h_i$  for the 4 new subimages.

The level 2 column of fig 2.6 shows how the total bit rate can be reduced by transforming the level 1 Lo-Lo subimage into four level 2 subimages. The process can be repeated by transforming the final Lo-Lo subimage again and again, giving the subimages in figs 2.7b and 2.7c and the histograms in figs 2.9 and 2.10. The levels 3 and 4 columns of fig 2.6 show that little is gained by transforming to more than 4 levels.

**However a total compression ratio of 4 bit/pel : 1.61 bit/pel = 2.45 : 1 has been achieved (in theory).**

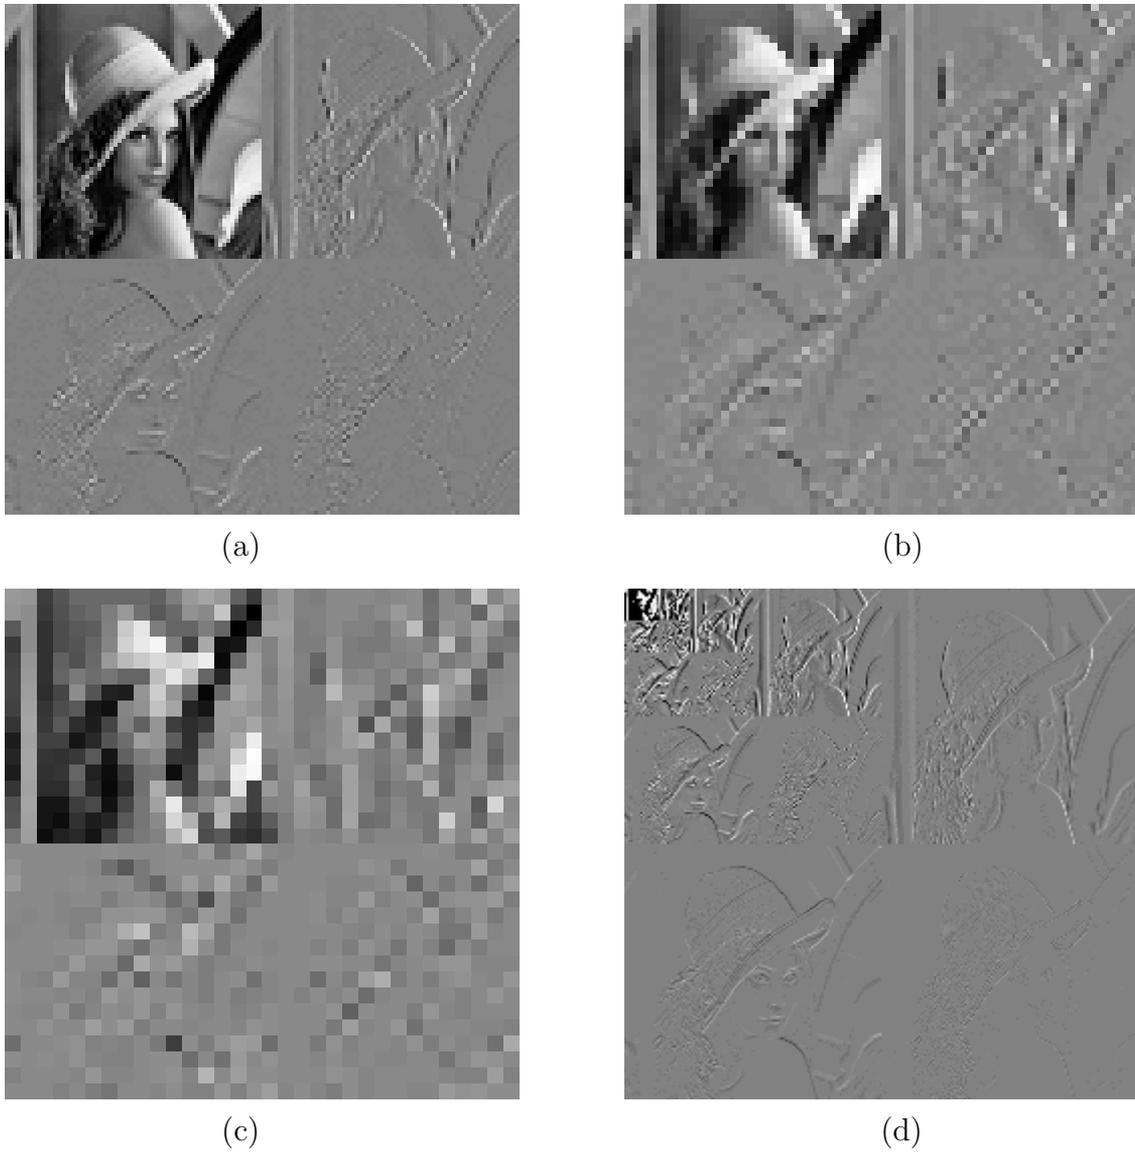


Fig 2.7: Levels 2 (a), 3 (b), and 4 (c) Haar transforms of Lenna;  
and at all of levels 1 to 4 (d).

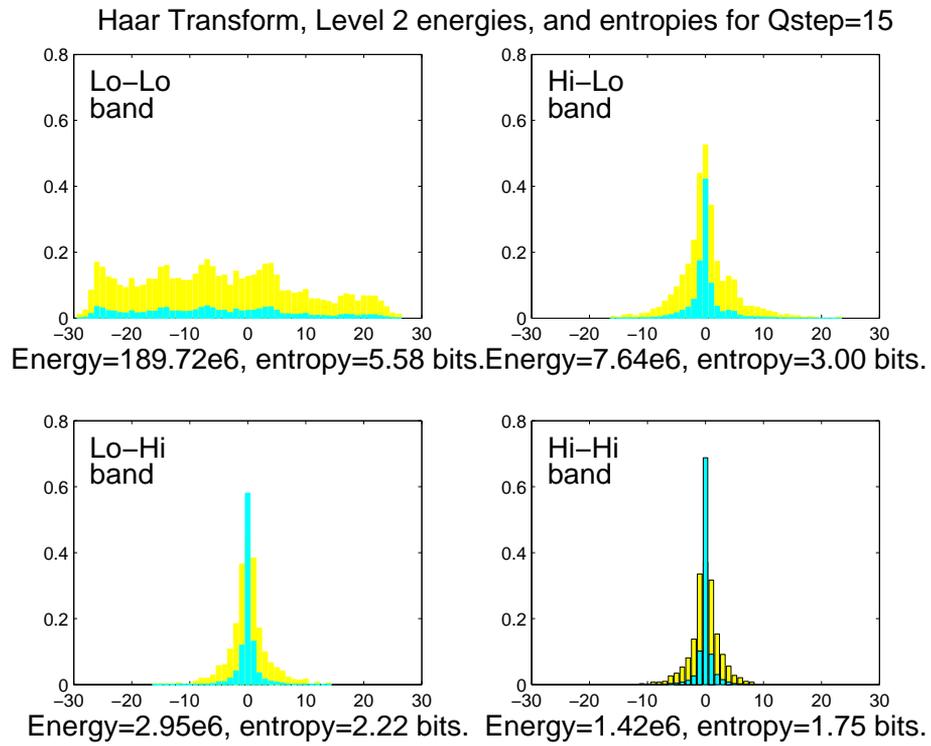


Fig 2.8: The probabilities  $p_i$  and entropies  $h_i$  for the 4 subimages at level 2.

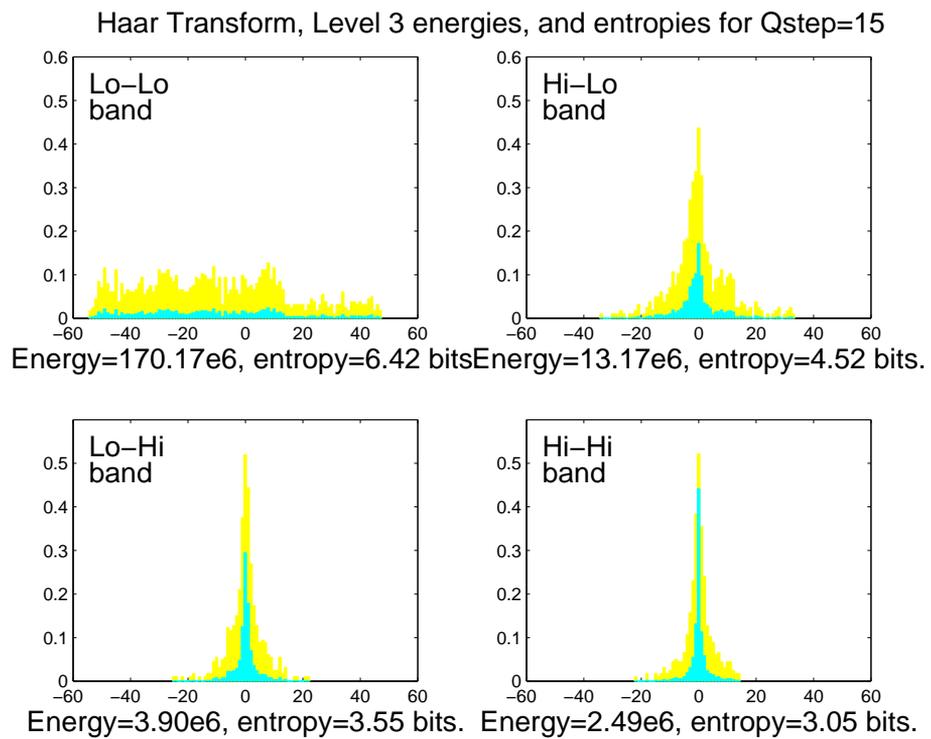


Fig 2.9: The probabilities  $p_i$  and entropies  $h_i$  for the 4 subimages at level 3.

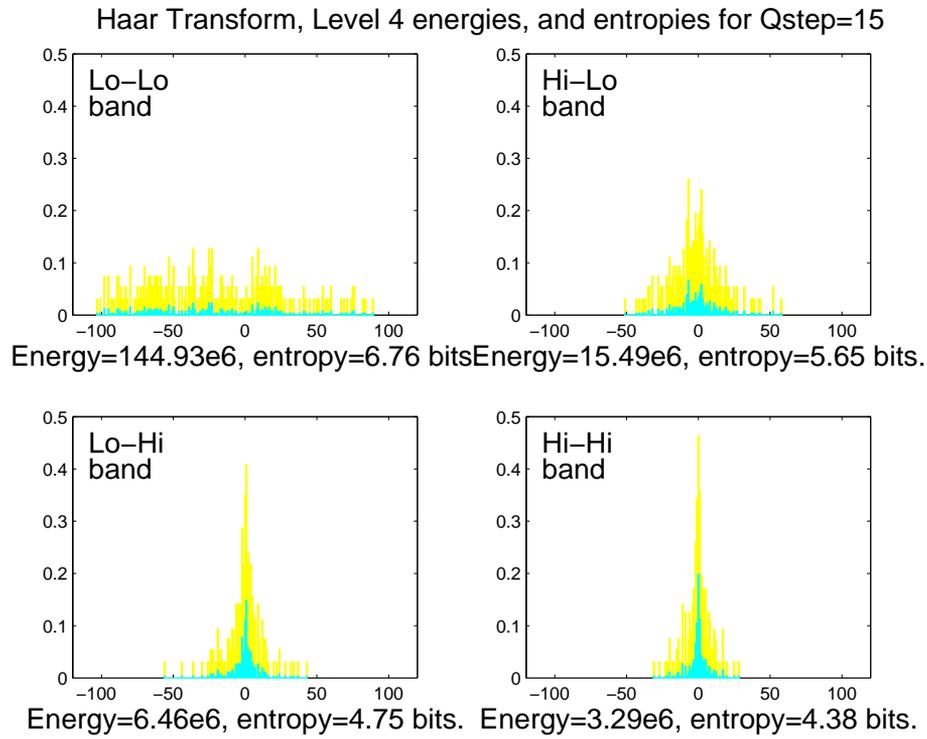


Fig 2.10: The probabilities  $p_i$  and entropies  $h_i$  for the 4 subimages at level 4.



Fig 2.11: Images reconstructed from (a) the original Lenna, and (b) the 4-level Haar transform, each quantised with  $Q_{step} = 15$ . The rms error of (a) = 4.3513, and of (b) = 3.5343.

Note the following features of the 4-level Haar transform:

- Fig 2.7d shows the subimages from all 4 levels of the transform and illustrates the transform's **multi-scale** nature. It also shows that all the subimages occupy the same total area as the original and hence that the total number of transform output samples (coefficients) equals the number of input pels – there is **no redundancy**.
- From the Lo-Lo subimage histograms of figs 2.5, 2.8, 2.9 and 2.10, we see the magnitudes of the Lo-Lo subimage samples increasing with transform level. This is because energy is being conserved and most of it is being concentrated in fewer and fewer Lo-Lo samples. (The DC gain of the Lo-Lo filter of equation (2.5) is 2.)
- We may reconstruct the image from the transform samples (fig 2.7d), quantised to  $Q_{step} = 15$ , by inverting the transform, using the right-hand part of (2.5). We then get the image in fig 2.11b. Contrast this with fig 2.11a, obtained by quantising the pels of the original directly to  $Q_{step} = 15$ , in which contour artifacts are much more visible. Thus the transform provides improved subjective quality as well as significant data compression. The improved quality arises mainly from the high amplitude of the low frequency transform samples, which means that they are quantised to many more levels than the basic pels would be for a given  $Q_{step}$ .
- If  $Q_{step}$  is doubled to 30, then the entropies of all the subimages are reduced as shown in fig 2.12 (compare this with fig 2.6 in which  $Q_{step} = 15$ ). The mean bit rate with the 4-level Haar transform drops from 1.61 to 0.97 bit/pel. However the reconstructed image quality drops to that shown in fig 2.13b. For comparison, fig 2.13a shows the quality if  $\mathbf{x}$  is directly quantised with  $Q_{step} = 30$ .

Multi-level transforms are a key feature of many image compression methods, and later in this course we shall see how they are used in both Cosine and Wavelet based systems.

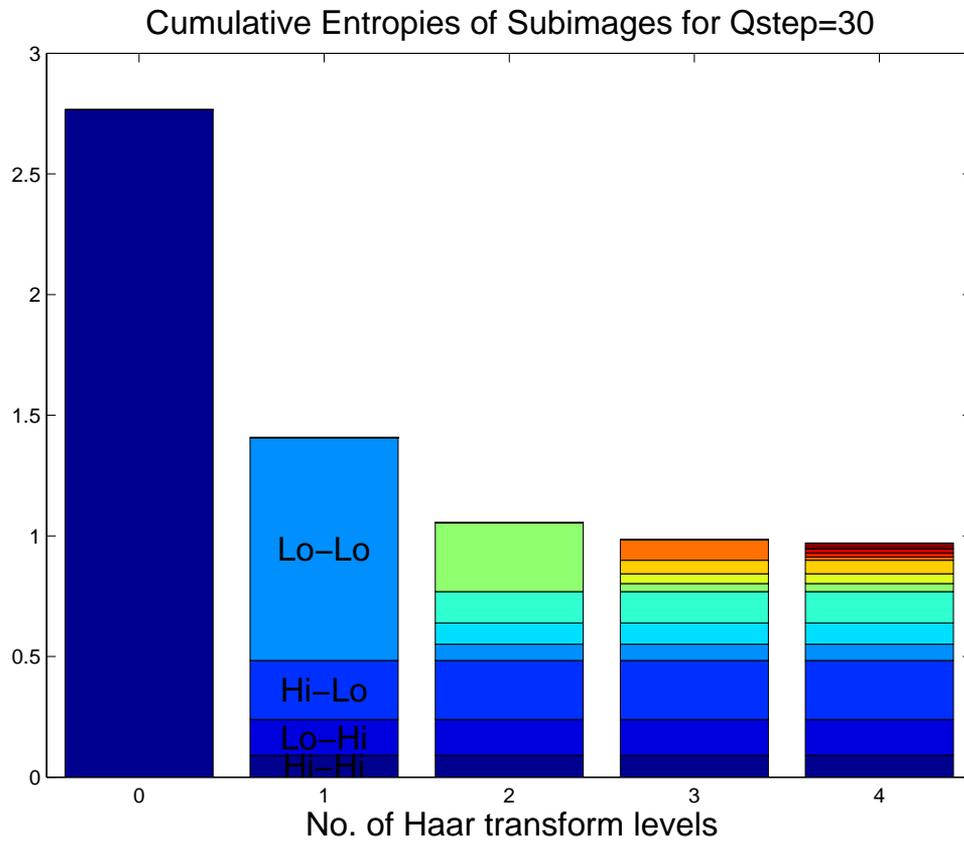


Fig 2.12: Mean bit rate for the original Lenna image and for the Haar transforms of the image after 1 to 4 levels, using a quantiser step size  $Q_{step} = 30$ .



Fig 2.13: Images reconstructed from (a) the original Lenna, and (b) the 4-level Haar transform, each quantised with  $Q_{step} = 30$ . Rms errors: (a) 8.6219; (b) 5.8781.

## 2.4 Use of Laplacian PDFs in Image Compression

It is found to be appropriate and convenient to model the distribution of many types of transformed image coefficients by Laplacian distributions. It is appropriate because much real data is approximately modelled by the Laplacian probability density function (PDF), and it is convenient because the mathematical form of the Laplacian PDF is simple enough to allow some useful analytical results to be derived.

A Laplacian PDF is a back-to-back pair of exponential decays and is given by:

$$p(x) = \frac{1}{2x_0} e^{-|x|/x_0} \quad (2.7)$$

where  $x_0$  is the equivalent of a **time constant** which defines the **width** of the PDF from the centre to the  $1/e$  points. The initial scaling factor ensures that the area under  $p(x)$  is unity, so that it is a valid PDF. Fig 2.14 shows the shape of  $p(x)$ .

The mean of this PDF is zero and the **variance** is given by:

$$v(x_0) = \int_{-\infty}^{\infty} x^2 p(x) dx = 2 \int_0^{\infty} \frac{x^2}{2x_0} e^{-x/x_0} dx = 2x_0^2 \quad (2.8)$$

(using integration by parts twice).

Hence the **standard deviation** is:

$$\sigma(x_0) = \sqrt{v(x_0)} = \sqrt{2} x_0 \quad (2.9)$$

Given the variance (power) of a subimage of transformed pels, we may calculate  $x_0$  and hence determine the PDF of the subimage, assuming a Laplacian shape. We now show that, if we quantise the subimage using a uniform quantiser with step size  $Q$ , we can calculate the entropy of the quantised samples and thus estimate the bit rate needed to encode the subimage in bits/pel. This is a powerful analytical tool as it shows how the compressed bit rate relates directly to the energy of a subimage. The vertical dashed lines in fig 2.14 show the decision thresholds for a typical quantiser for the case when  $Q = 2x_0$ .

**First we analyse the probability of a pel being quantised to each step of the quantiser.** This is given by the area under  $p(x)$  between each adjacent pair of quantiser thresholds.

Probability of being at step 0,  $p_0 = \text{prob}(-\frac{1}{2}Q < x < \frac{1}{2}Q) = 2 \text{prob}(0 < x < \frac{1}{2}Q)$

Probability of being at step  $k$ ,  $p_k = \text{prob}((k - \frac{1}{2})Q < x < (k + \frac{1}{2})Q)$

First, for  $x_2 \geq x_1 \geq 0$ , we calculate:

$$\text{prob}(x_1 < x < x_2) = \int_{x_1}^{x_2} p(x) dx = \left[ -\frac{1}{2} e^{-x/x_0} \right]_{x_1}^{x_2} = \frac{1}{2}(e^{-x_1/x_0} - e^{-x_2/x_0})$$

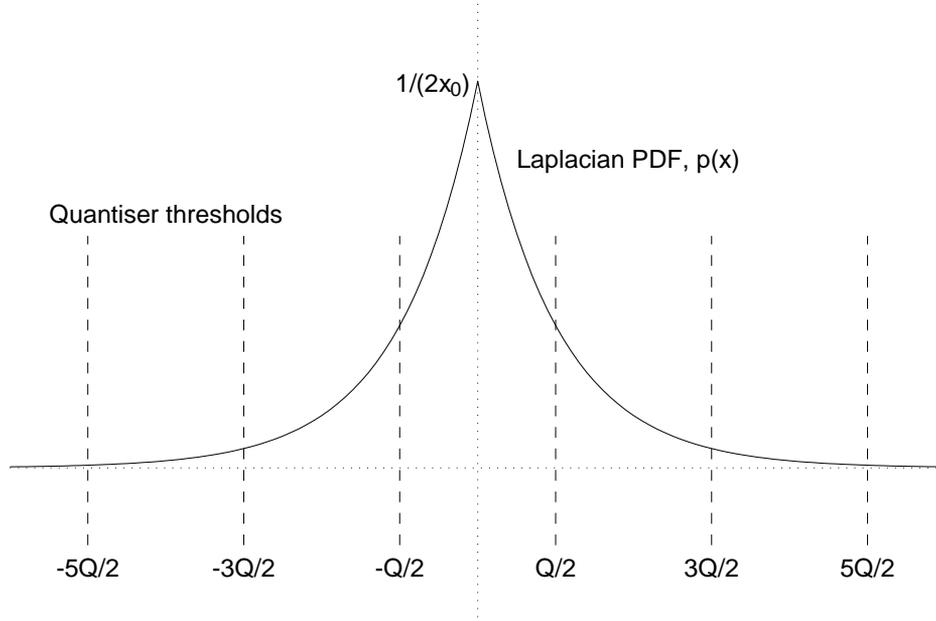


Fig 2.14: Laplacian PDF,  $p(x)$ , and typical quantiser decision thresholds, shown for the case when the quantiser step size  $Q = 2x_0$ .

$$\therefore p_0 = (1 - e^{-Q/(2x_0)}) \quad (2.10)$$

$$\text{and, for } k \geq 1, p_k = \frac{1}{2}(e^{-(k-\frac{1}{2})Q/x_0} - e^{-(k+\frac{1}{2})Q/x_0}) = \sinh\left(\frac{Q}{2x_0}\right) e^{-kQ/x_0} \quad (2.11)$$

By symmetry, if  $k$  is nonzero,  $p_{-k} = p_k = \sinh\left(\frac{Q}{2x_0}\right) e^{-|k|Q/x_0}$

Now we can **calculate the entropy** of the subimage:

$$H = - \sum_{k=-\infty}^{\infty} p_k \log_2 p_k = -p_0 \log_2 p_0 - 2 \sum_{k=1}^{\infty} p_k \log_2 p_k \quad (2.12)$$

To make the evaluation of the summation easier when we substitute for  $p_k$ , we let

$$\begin{aligned} p_k &= \alpha r^k \quad \text{where } \alpha = \sinh\left(\frac{Q}{2x_0}\right) \quad \text{and } r = e^{-Q/x_0} \\ \therefore \sum_{k=1}^{\infty} p_k \log_2 p_k &= \sum_{k=1}^{\infty} \alpha r^k \log_2(\alpha r^k) = \sum_{k=1}^{\infty} \alpha r^k (\log_2 \alpha + k \log_2 r) \\ &= \alpha (\log_2 \alpha) \sum_{k=1}^{\infty} r^k + \alpha (\log_2 r) \sum_{k=1}^{\infty} k r^k \end{aligned}$$

$$\text{Now } \sum_{k=1}^{\infty} r^k = \frac{r}{1-r} \quad \text{and, differentiating by } r : \sum_{k=1}^{\infty} k r^{k-1} = \frac{1}{(1-r)^2}$$

$$\begin{aligned} \therefore \sum_{k=1}^{\infty} p_k \log_2 p_k &= \alpha (\log_2 \alpha) \frac{r}{1-r} + \alpha (\log_2 r) \frac{r}{(1-r)^2} \\ &= \frac{\alpha r}{1-r} \left[ \log_2 \alpha + \frac{\log_2 r}{1-r} \right] \\ \text{and } p_0 \log_2 p_0 &= (1 - \sqrt{r}) \log_2(1 - \sqrt{r}) \end{aligned}$$

Hence the entropy is given by:

$$H = -(1 - \sqrt{r}) \log_2(1 - \sqrt{r}) - \frac{2\alpha r}{1-r} \left[ \log_2 \alpha + \frac{\log_2 r}{1-r} \right] \quad (2.13)$$

Because both  $\alpha$  and  $r$  are functions of  $Q/x_0$ , then  $H$  is a function of just  $Q/x_0$  too. We expect that, for constant  $Q$ , as the energy of the subimage increases, the entropy will also increase approximately logarithmically, so we plot  $H$  against  $x_0/Q$  in dB in fig 2.15. This shows that our expectations are born out.

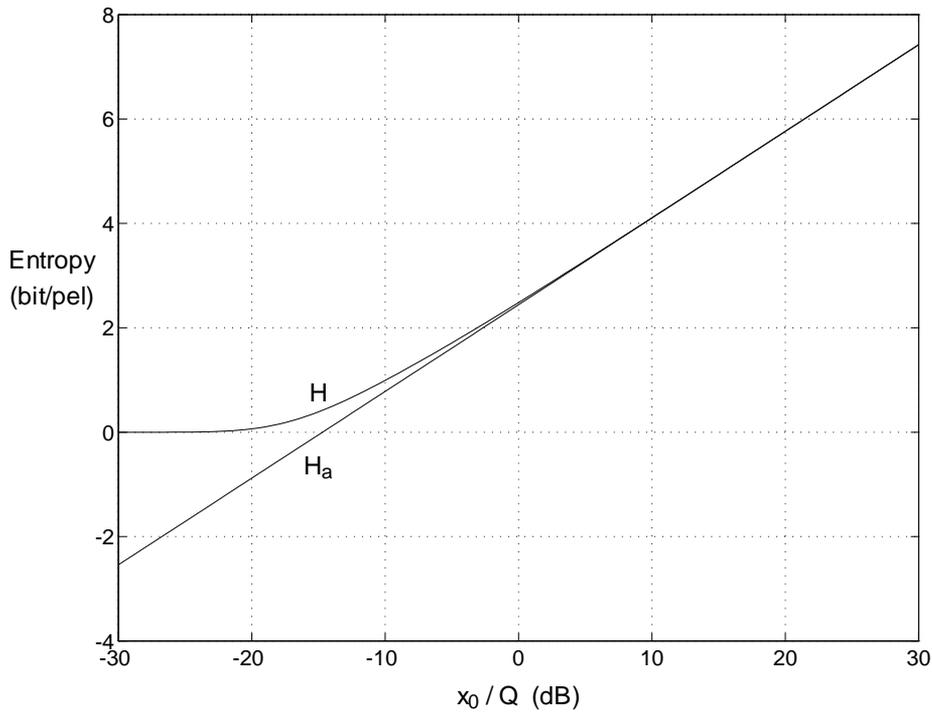


Fig 2.15: Entropy  $H$  and approximate entropy  $H_a$  of a quantised subimage with Laplacian PDF, as a function of  $x_0/Q$  in dB.

We can show this in theory by considering the case when  $x_0/Q \gg 1$ , when we find that:

$$\alpha \approx \frac{Q}{2x_0} \quad r \approx \left(1 - \frac{Q}{x_0}\right) \approx 1 - 2\alpha \quad \sqrt{r} \approx 1 - \alpha$$

Using the approximation  $\log_2(1-\epsilon) \approx -\epsilon/\log_e 2$  for small  $\epsilon$ , it is then fairly straightforward to show that

$$H \approx -\log_2 \alpha + \frac{1}{\log_e 2} \approx \log_2 \left( \frac{2e x_0}{Q} \right)$$

We denote this approximation as  $H_a$  in fig 2.15, which shows how close to  $H$  the approximation is, for  $x_0 > Q$  (i.e. for  $x_0/Q > 0$  dB).

**We can compare the entropies calculated using equation (2.13) with those that were calculated from the bandpass subimage histograms**, as given in figs 2.5, 2.8, 2.9 and 2.10. (The Lo-Lo subimages have PDFs which are more uniform and do not fit the Laplacian model well.) The values of  $x_0$  are calculated from:

$$x_0 = \frac{\text{std. dev.}}{\sqrt{2}} = \sqrt{\frac{\text{subimage energy}}{2 \text{ (no of pels in subimage)}}}$$

The following table shows this comparison:

Transform level	Subimage type	Energy ( $\times 10^6$ )	No of pels	$x_0$	Laplacian entropy	Measured entropy
1	Hi-Lo	4.56	16384	11.80	2.16	1.71
1	Lo-Hi	1.89	16384	7.59	1.58	1.15
1	Hi-Hi	0.82	16384	5.09	1.08	0.80
2	Hi-Lo	7.64	4096	30.54	3.48	3.00
2	Lo-Hi	2.95	4096	18.98	2.81	2.22
2	Hi-Hi	1.42	4096	13.17	2.31	1.75
3	Hi-Lo	13.17	1024	80.19	4.86	4.52
3	Lo-Hi	3.90	1024	43.64	3.99	3.55
3	Hi-Hi	2.49	1024	34.87	3.67	3.05
4	Hi-Lo	15.49	256	173.9	5.98	5.65
4	Lo-Hi	6.46	256	112.3	5.35	4.75
4	Hi-Hi	3.29	256	80.2	4.86	4.38

We see that the entropies calculated from the energy via the Laplacian PDF method (second column from the right) are approximately 0.5 bit/pel greater than the entropies measured from the Lenna subimage histograms. This is due to the heavier tails of the actual PDFs compared with the Laplacian exponentially decreasing tails. More accurate entropies can be obtained if  $x_0$  is obtained from the mean absolute values of the pels in each subimage. For a Laplacian PDF we can show that

$$\text{Mean absolute value} = \int_{-\infty}^{\infty} |x| p(x) dx = 2 \int_0^{\infty} \frac{x}{2x_0} e^{-x/x_0} dx = x_0 \quad (2.14)$$

Estimating  $x_0$  from measurements of the mean absolute values of the subband pels gives values of  $x_0$  that are about 20% lower than those calculated from the energies of the pels, and the calculated entropies are then within approximately 0.2 bit/pel of the measured entropies.

## 2.5 Practical Entropy Coding Techniques

In the previous sections we have assumed that ideal entropy coding has been used in order to calculate the bit rates for the coded data. In practise we must use real codes and we shall now see how this affects the compression performance.

There are three main techniques for achieving entropy coding:

**Huffman Coding** – one of the simplest variable length coding schemes.

**Run-length Coding (RLC)** – very useful for binary data containing long runs of ones or zeros.

**Arithmetic Coding** – a relatively new variable length coding scheme that can combine the best features of Huffman and run-length coding, and also adapt to data with non-stationary statistics.

We shall concentrate on the Huffman and RLC methods for simplicity. Interested readers may find out more about Arithmetic Coding in chapters 12 and 13 of the JPEG Book.

First we consider the change in compression performance if simple Huffman Coding is used to code the subimages of the 4-level Haar transform.

The calculation of entropy in equation (2.6) assumed that each message with probability  $p_i$  could be represented by a word of length  $w_i = -\log_2(p_i)$  bits. Huffman codes require the  $w_i$  to be integers and assume that the  $p_i$  are adjusted to become:

$$\hat{p}_i = 2^{-w_i} \quad (2.15)$$

where the  $w_i$  are integers, chosen subject to the constraint that  $\sum \hat{p}_i \leq 1$  (to guarantee that sufficient uniquely decodable code words are available) and such that the mean Huffman word length (Huffman entropy),  $\hat{H} = \sum p_i w_i$ , is minimised.

We can use the probability histograms which generated the entropy plots of figs 2.5, 2.8, 2.9 and 2.10 to calculate the Huffman entropies  $\hat{H}$  for each subimage and compare these with the true entropies to see the loss in performance caused by using real Huffman codes.

An algorithm for finding the optimum codesizes  $w_i$  is recommended in the JPEG specification [The JPEG Book, Appendix A, Annex K.2, fig K.1]; and a Matlab M-file to implement it is given in fig 2.19.

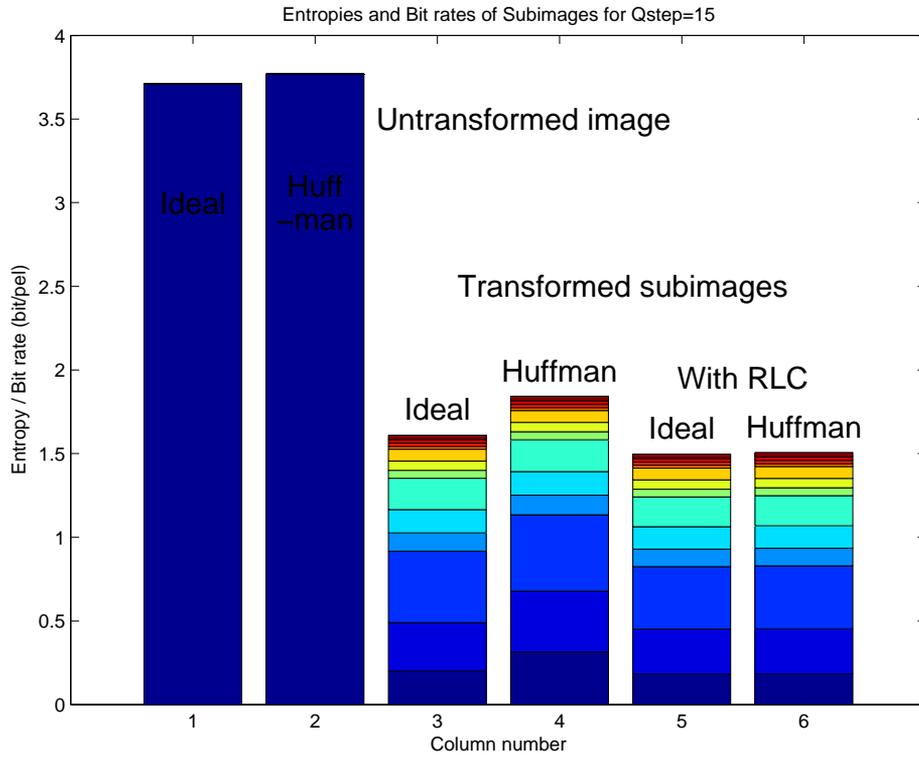


Fig 2.16: Comparison of entropies (columns 1,3,5) and Huffman coded bit rates (columns 2,4,6) for the original (columns 1 and 2) and transformed (columns 3 to 6) Lenna images. In columns 5 and 6, the zero amplitude state is run-length encoded to produce many states with probabilities < 0.5.

Column:	1	2	3	4	5	6	
			0.0264	0.0265	0.0264	0.0266	-
			0.0220	0.0222	0.0221	0.0221	Level 4
			0.0186	0.0187	0.0185	0.0186	
			0.0171	0.0172	0.0171	0.0173	-
			0.0706	0.0713	0.0701	0.0705	
			0.0556	0.0561	0.0557	0.0560	Level 3
	3.7106	3.7676	0.0476	0.0482	0.0466	0.0471	-
			0.1872	0.1897	0.1785	0.1796	
			0.1389	0.1413	0.1340	0.1353	Level 2
			0.1096	0.1170	0.1038	0.1048	-
			0.4269	0.4566	0.3739	0.3762	
			0.2886	0.3634	0.2691	0.2702	Level 1
			0.2012	0.3143	0.1819	0.1828	-
Totals:	3.7106	3.7676	1.6103	1.8425	1.4977	1.5071	

Fig 2.17: Numerical results used in fig 2.16.

Fig 2.16 shows the results of applying this algorithm to the probability histograms and fig 2.17 lists the same results numerically for ease of analysis. Columns 1 and 2 compare the ideal entropy with the mean word length or bit rate from using a Huffman code (the Huffman entropy) for the case of the untransformed image where the original pels are quantised with  $Q_{step} = 15$ . We see that the increase in bit rate from using the real code is:

$$\frac{3.7676}{3.7106} - 1 = 1.5\%$$

**But** when we do the same for the 4-level transformed subimages, we get columns 3 and 4. Here we see that real Huffman codes require an increase in bit rate of:

$$\frac{1.8425}{1.6103} - 1 = 14.4\%$$

Comparing the results for each subimage in columns 3 and 4, we see that most of the increase in bit rate arises in the three level-1 subimages at the bottom of the columns. This is because each of the probability histograms for these subimages (fig 2.5) contain one probability that is greater than 0.5. Huffman codes cannot allocate a word length of less than 1 bit to a given event, and so they start to lose efficiency rapidly when  $-\log_2(p_i)$  becomes less than 1, ie when  $p_i > 0.5$ .

**Run-length codes (RLCs) are a simple and effective way of improving the efficiency of Huffman coding when one event is much more probable than all of the others combined.** They operate as follows:

- The pels of the subimage are scanned sequentially (usually in columns or rows) to form a long 1-dimensional vector.
- Each run of consecutive zero samples (the most probable events) in the vector is coded as a single event.
- Each non-zero sample is coded as a single event in the normal way.
- The two types of event (runs-of-zeros and non-zero samples) are allocated separate sets of codewords in the same Huffman code, which may be designed from a histogram showing the frequencies of all events.
- To limit the number of run events, the maximum run length may be limited to a certain value (we have used 128) and runs longer than this may be represented by two or more run codes in sequence, with negligible loss of efficiency.

Hence RLC may be added before Huffman coding as an extra processing step, which converts the most probable event into many separate events, each of which has  $p_i < 0.5$  and

may therefore be coded efficiently. Fig 2.18 shows the new probability histograms and entropies for level 1 of the Haar transform when RLC is applied to the zero event of the three bandpass subimages. Comparing this with fig 2.5, note the absence of the high probability zero events and the new states to the right of the original histograms corresponding to the run lengths.

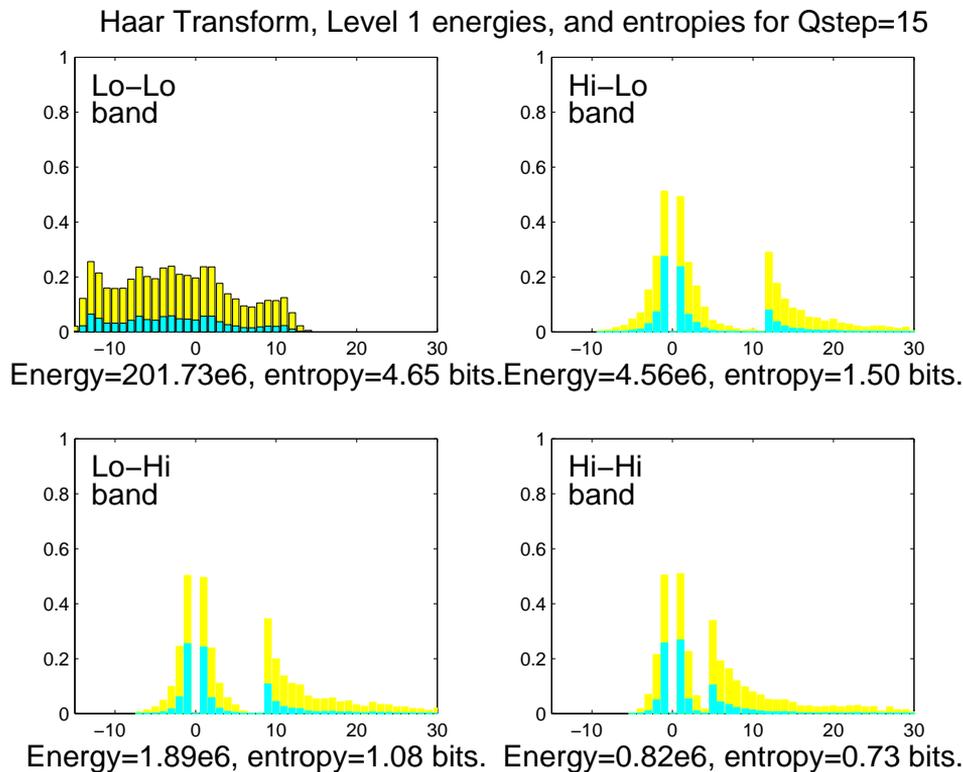


Fig 2.18: Probability histograms (dashed) and entropies (solid) of the four subimages of the Level 1 Haar transform of Lenna (fig 2.2b) after RLC.

The total entropy **per event** for an RLC subimage is calculated as before from the entropy histogram. However to get the entropy **per pel** we scale the entropy by the ratio of the number of events (runs and non-zero samples) in the subimage to the number of pels in the subimage (note that with RLC this ratio will no longer equal one – it will hopefully be much less).

Fig 2.18 gives the entropies per pel after RLC for each subimage, which are now **less** than the entropies in fig 2.5. This is because RLC takes advantage of **spatial clustering** of the zero samples in a subimage, rather than just depending on the histogram of amplitudes.

Clearly if all the zeros were clustered into a single run, this could be coded much more efficiently than if they are distributed into many runs. On the other hand, the entropy of the zero event just tells us the mean number of bits to code each zero pel **if the zero pels**

**are distributed randomly**, ie if the probability of a given pel being zero does not depend on the amplitudes of any nearby pels.

In typical bandpass subimages, non-zero samples tend to be clustered around key features such as object boundaries and areas of high texture. Hence RLC usually reduces the entropy of the data to be coded. There are many other ways to take advantage of clustering (correlation) of the data – RLC is just one of the simplest.

In fig 2.16, comparing column 5 with column 3, we see the modest (7%) reduction in entropy per pel achieved by RLC, due clustering in the Lenna image. The main advantage of RLC is apparent in column 6, which shows the mean bit rate per pel when we use a real Huffman code on the RLC histograms of fig 2.18. The increase in bit rate over the RLC entropy is only

$$\frac{1.5071}{1.4977} - 1 = 0.63\%$$

compared with 14.4% when RLC is not used (columns 3 and 4).

Finally, comparing column 6 with column 3, we see that, relative to the simple entropy measure, combined RLC and Huffman coding can **reduce** the bit rate by

$$1 - \frac{1.5071}{1.6103} = 6.4\%$$

The closeness of this ratio to unity justifies our use of simple entropy as a tool for assessing the information compression properties of the Haar transform – and of other energy compression techniques as we meet them.

```

% Find Huffman code sizes: JPEG fig K.1, procedure Code_size.
% huffhist contains the histogram of event counts (frequencies).
freq = huffhist(:);
codesize = zeros(size(freq));
others = -ones(size(freq)); % Pointers to next symbols in code tree.

% Find non-zero entries in freq, and loop until only 1 entry left.
nz = find(freq > 0);
while length(nz) > 1,
% Find v1 for least value of freq(v1) > 0.
    [y,i] = min(freq(nz));
    v1 = nz(i);
% Find v2 for next least value of freq(v2) > 0.
    nz = nz([1:(i-1) (i+1):length(nz)]); % Remove v1 from nz.
    [y,i] = min(freq(nz));
    v2 = nz(i);
% Combine frequency values.
    freq(v1) = freq(v1) + freq(v2);
    freq(v2) = 0;
    codesize(v1) = codesize(v1) + 1;
% Increment code sizes for all codewords in this tree branch.
    while others(v1) > -1,
        v1 = others(v1);
        codesize(v1) = codesize(v1) + 1;
    end
    others(v1) = v2;
    codesize(v2) = codesize(v2) + 1;
    while others(v2) > -1,
        v2 = others(v2);
        codesize(v2) = codesize(v2) + 1;
    end
    nz = find(freq > 0);
end

% Generate Huffman entropies by multiplying probabilities by code sizes.
huffent = (huffhist(:)/sum(huffhist(:))) .* codesize;

```

Fig 2.19: Listing of the M-file to calculate the Huffman entropy from a given histogram.

### 3 The DCT and the JPEG Standard

The main standard for image compression in current use is the JPEG (Joint Picture Experts Group) standard, devised and refined over the period 1985 to 1993. It is formally known as ISO Draft International Standard 10981-1 and CCITT Recommendation T.81, and is described in depth in *The JPEG Book* by W B Pennebaker and J L Mitchell, Van Nostrand Reinhold 1993. There are also two newer standards: JPEG 2000 (JP2k), which was finalised in 2002; and JPEG XR (JPXR – eXtended Range), finalised in 2009. JP2k uses a wavelet transform at the front end, while JPXR uses two levels of cosine transform; and both employ some pretty sophisticated context-based entropy coding methods. We shall describe cosine transforms in this section and wavelets in section 4, but unfortunately we do not have time on this course to go into details of JP2k or JPXR entropy coding methods, as they are rather too complicated.

Take-up of JP2k has been very slow, because some intellectual property (IP) ownership and rights problems surfaced shortly after the standard was published. This meant that Microsoft and other large players were reluctant to bundle JP2k encoders and decoders in their systems for fear of damaging lawsuits. As a result in 2006 JPXR was proposed by Microsoft, and numerous agreements (see JPEG XR on Wikipedia and the *Microsoft Community Promise*) are now in place to allow it to be freely implemented by other suppliers, which should encourage more widespread take-up. It is also more efficient to implement than JP2k. However at present we still see basic JPEG nearly everywhere.

In this section of the course we shall mainly describe the standard baseline version of JPEG, which is currently in extensive use for most images in cameras, phones and on the internet. Also the same techniques form the basis of the MPEG standards for video compression, which we discuss in section 5. First we consider the JPEG energy compression technique – the discrete cosine transform or DCT.

#### 3.1 The Discrete Cosine Transform (DCT)

In equation (2.2) we met the 2-point Haar transform and in (2.3) we saw that it can be easily inverted if the transform matrix  $\mathbf{T}$  is orthonormal so that  $\mathbf{T}^{-1} = \mathbf{T}^T$ .

If  $\mathbf{T}$  is of size  $n \times n$ , where  $n = 2^m$ , then we may easily generate larger orthonormal matrices, which lead to definitions of larger transforms.

An  $n$ -point transform is defined as:

$$\begin{bmatrix} y(1) \\ \vdots \\ y(n) \end{bmatrix} = \mathbf{T} \begin{bmatrix} x(1) \\ \vdots \\ x(n) \end{bmatrix} \quad \text{where} \quad \mathbf{T} = \begin{bmatrix} t_{11} & \cdots & t_{1n} \\ \vdots & & \vdots \\ t_{n1} & \cdots & t_{nn} \end{bmatrix} \quad (3.1)$$

A 4-point orthonormal transform matrix, equivalent to 2 levels of the Haar transform, is:

$$\mathbf{T} = \underbrace{\frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 0 & 1 & 0 \\ 1 & 0 & -1 & 0 \\ 0 & \sqrt{2} & 0 & 0 \\ 0 & 0 & 0 & \sqrt{2} \end{bmatrix}}_{\text{Haar level 2}} \underbrace{\frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & -1 \end{bmatrix}}_{\text{Haar level 1}} = \frac{1}{2} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ \sqrt{2} & -\sqrt{2} & 0 & 0 \\ 0 & 0 & \sqrt{2} & -\sqrt{2} \end{bmatrix} \quad (3.2)$$

Similarly 3 and 4 level Haar transforms may be expressed using 8 and 16 point transform matrices respectively.

However for  $n > 2$ , there are better matrices than those based on the Haar transform, where **better** means **with improved energy compression properties for typical images**.

Discrete Cosine Transforms (DCTs) have some of these improved properties and are also simple to define and implement. The  $n$  rows of an  $n$ -point DCT matrix  $\mathbf{T}$  are defined by:

$$\begin{aligned} t_{1i} &= \sqrt{\frac{1}{n}} && \text{for } i = 1 \rightarrow n, \\ t_{ki} &= \sqrt{\frac{2}{n}} \cos\left(\frac{\pi(i - \frac{1}{2})(k - 1)}{n}\right) && \text{for } i = 1 \rightarrow n, \quad k = 2 \rightarrow n. \end{aligned} \quad (3.3)$$

It is straightforward to show that this matrix is orthonormal for  $n$  even, since the norm of each row is unity and the dot product of any pair of rows is zero (the product terms may be expressed as the sum of a pair of cosine functions, which are each zero mean).

The 8-point DCT matrix ( $n = 8$ ) is:

$$\mathbf{T} = \begin{bmatrix} 0.3536 & 0.3536 & 0.3536 & 0.3536 & 0.3536 & 0.3536 & 0.3536 & 0.3536 \\ 0.4904 & 0.4157 & 0.2778 & 0.0975 & -0.0975 & -0.2778 & -0.4157 & -0.4904 \\ 0.4619 & 0.1913 & -0.1913 & -0.4619 & -0.4619 & -0.1913 & 0.1913 & 0.4619 \\ 0.4157 & -0.0975 & -0.4904 & -0.2778 & 0.2778 & 0.4904 & 0.0975 & -0.4157 \\ 0.3536 & -0.3536 & -0.3536 & 0.3536 & 0.3536 & -0.3536 & -0.3536 & 0.3536 \\ 0.2778 & -0.4904 & 0.0975 & 0.4157 & -0.4157 & -0.0975 & 0.4904 & -0.2778 \\ 0.1913 & -0.4619 & 0.4619 & -0.1913 & -0.1913 & 0.4619 & -0.4619 & 0.1913 \\ 0.0975 & -0.2778 & 0.4157 & -0.4904 & 0.4904 & -0.4157 & 0.2778 & -0.0975 \end{bmatrix} \quad (3.4)$$

The rows of  $\mathbf{T}$ , known as **basis functions**, are plotted as asterisks in fig 3.1. The asterisks are superimposed on the underlying continuous cosine functions, used in all sizes of DCT. Only the amplitude scaling and the maximum frequency vary with the size  $n$ .

When we take the transform of an  $n$ -point vector using  $\mathbf{y} = \mathbf{T}\mathbf{x}$ ,  $\mathbf{x}$  is decomposed into a linear combination of the basis functions (rows) of  $\mathbf{T}$ , whose coefficients are the samples of  $\mathbf{y}$ , because  $\mathbf{x} = \mathbf{T}^T\mathbf{y}$ .

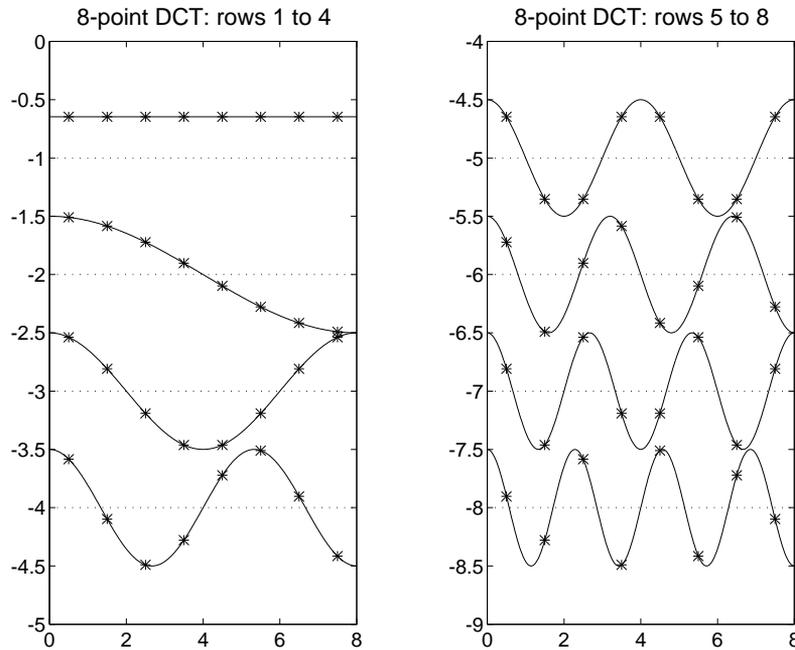


Fig 3.1: The 8-point DCT basis functions (\*) and their underlying continuous cosine waves.

The basis functions may also be viewed as the **impulse responses of FIR filters**, being applied to the data  $\mathbf{x}$ .

The DCT is closely related to the discrete Fourier transform (DFT). It represents the result of applying the  $2n$ -point DFT to a vector:

$$\mathbf{x}_{2n} = \begin{bmatrix} \mathbf{x} \\ \mathbf{x}_{rev} \end{bmatrix} \quad \text{where } \mathbf{x}_{rev} = \begin{bmatrix} x(n) \\ \vdots \\ x(1) \end{bmatrix}$$

$\mathbf{x}_{2n}$  is symmetric about its centre and so the  $2n$  Fourier coefficients are all purely real and symmetric about zero frequency. The  $n$  DCT coefficients are then the first  $n$  Fourier coefficients.

**Note:** the DFT must be defined with a half sample period offset on the indexing of the input samples for the above to be strictly true – see  $(i - \frac{1}{2})$  in eq(3.3).

## Standards

The **8-point DCT** is the basis of the JPEG standard, as well as several other standards such as MPEG-1 and MPEG-2 (for TV and video) and H.263 (for video-phones). Hence we shall concentrate on it as our main example, but bear in mind that DCTs may be defined for a wide range of sizes  $n$ .

### 3.2 Fast algorithms for the DCT

The basic  $n$ -point DCT requires  $n^2$  multiplications and  $n(n - 1)$  additions to calculate  $\mathbf{y} = \mathbf{T}\mathbf{x}$  (64 mults and 56 adds for  $n = 8$ ).

From fig 3.1, it is clear that **symmetries** exist in the DCT basis functions. These can be exploited to **reduce the computation load** of the DCT.

All the odd rows of  $\mathbf{T}$  in (3.4) possess even symmetry about their centres and all the even rows possess odd symmetry. Hence we may form:

$$u(i) = x(i) + x(9 - i) \quad \text{and} \quad v(i) = x(i) - x(9 - i) \quad \text{for } i = 1 \rightarrow 4 \quad (3.5)$$

and then form the odd and even terms in  $\mathbf{y}$  from two  $4 \times 4$  transforms:

$$\begin{bmatrix} y(1) \\ y(3) \\ y(5) \\ y(7) \end{bmatrix} = \mathbf{T}_{left,odd} \mathbf{u} \quad \text{and} \quad \begin{bmatrix} y(2) \\ y(4) \\ y(6) \\ y(8) \end{bmatrix} = \mathbf{T}_{left,even} \mathbf{v} \quad (3.6)$$

where  $\mathbf{T}_{left,odd}$  and  $\mathbf{T}_{left,even}$  are the  $4 \times 4$  matrices formed by the left halves of the odd and even rows of  $\mathbf{T}$ .

This reduces the computation to 8 add/subtract operations for equations (3.5) and  $2 \times 16$  mults and  $2 \times 12$  adds for equations (3.6) – almost halving the total computation load.

The matrix  $\mathbf{T}_{left,even}$  cannot easily be simplified much further, but  $\mathbf{T}_{left,odd}$  can, as it possesses the same symmetries as  $\mathbf{T}$  (it is equivalent to a 4-point DCT matrix). Hence we may use the same technique on this matrix to reduce the 16 mults and 12 adds for this product to 4 add/subtract operations followed by a pair of  $2 \times 2$  matrix products, requiring  $2 \times 4$  mults and  $2 \times 2$  adds. Finally two of these mults may be eliminated since one of the  $2 \times 2$  matrices is just a scaled add/subtract matrix (like the Haar transform).

The total computation load for the  $8 \times 8$  DCT then becomes:

- $8 + 12 + (4 + 2 + 2) = 28$  add/subtract operations;
- $0 + 16 + (0 + 4 + 2) = 22$  multiply operations.

More complicated algorithms exist (JPEG Book, sections 4.3.2 to 4.3.5) which reduce the number of multiplies further. However these all require more intermediate results to be stored. In modern DSP chips this can cost more CPU cycles than the extra multiplications which can often be done simultaneously with additions. Hence the simple approach given above is frequently optimal.

### 3.3 The 2-dimensional DCT

In equation (2.5):

$$\mathbf{y} = \mathbf{T} \mathbf{x} \mathbf{T}^T \quad \text{and its inverse:} \quad \mathbf{x} = \mathbf{T}^T \mathbf{y} \mathbf{T}$$

we saw how a 1-D transform could be extended to 2-D by pre- and post-multiplication of a square matrix  $\mathbf{x}$  to give a matrix result  $\mathbf{y}$ . In that example we used  $2 \times 2$  matrices, but this technique may be applied to square matrices of any size.

Hence the **DCT may be extended into 2-D** by this method.

E.g. an  $8 \times 8$  DCT in 2-D transforms a subimage of  $8 \times 8$  pels into a matrix of  $8 \times 8$  DCT coefficients.

The **2-D basis functions**, from which  $\mathbf{x}$  may be reconstructed, are given by the  $n^2$  separate **products of the columns of  $\mathbf{T}^T$  with the rows of  $\mathbf{T}$** , obtained by putting a single '1' in turn in each of the  $n^2$  elements of  $\mathbf{y}$  (with the rest being zero) and looking at the resulting  $\mathbf{x} = \mathbf{T}^T \mathbf{y} \mathbf{T}$ . These are shown for  $n = 8$  in fig 3.2a as 64 subimages of size  $8 \times 8$  pels.

The result of applying the  $8 \times 8$  DCT to the Lenna image is shown in fig 3.2b. Here each  $8 \times 8$  block of pels  $\mathbf{x}$  is replaced by the  $8 \times 8$  block of DCT coefficients  $\mathbf{y}$ . This shows the  $8 \times 8$  block structure clearly but is not very meaningful otherwise.

Fig 3.2c shows the same data, reordered into 64 subimages of  $32 \times 32$  coefficients each so that each subimage contains all the coefficients of a given type – e.g: the top left subimage contains all the coefficients for the top left basis function from fig 3.2a. The other subimages and basis functions correspond in the same way.

We see the **major energy concentration** to the subimages in the top left corner. Fig 3.2d is an enlargement of the top left 4 subimages of fig 3.2c and bears a strong similarity to the group of third level Haar subimages in fig 2.7b. To emphasise this the histograms and entropies of these 4 subimages are shown in fig 3.3.

Comparing fig 3.3 with fig 2.9, the Haar transform equivalent on page 19, we see that the Lo-Lo bands have identical energies and entropies. This is because the basis functions are identical flat surfaces in both cases. Comparing the other 3 bands, we see that the DCT bands contain more energy and entropy than their Haar equivalents, which means **less energy** (and so hopefully **less entropy**) in the higher DCT bands (not shown) because the total energy is fixed (the transforms all preserve total energy because they are orthonormal).

The mean entropy for all 64 subimages of the DCT is **1.3621 bit/pel**, which compares favourably with the **1.6103 bit/pel** for the 4-level Haar transformed subimages using the same  $Q_{step} = 15$ .

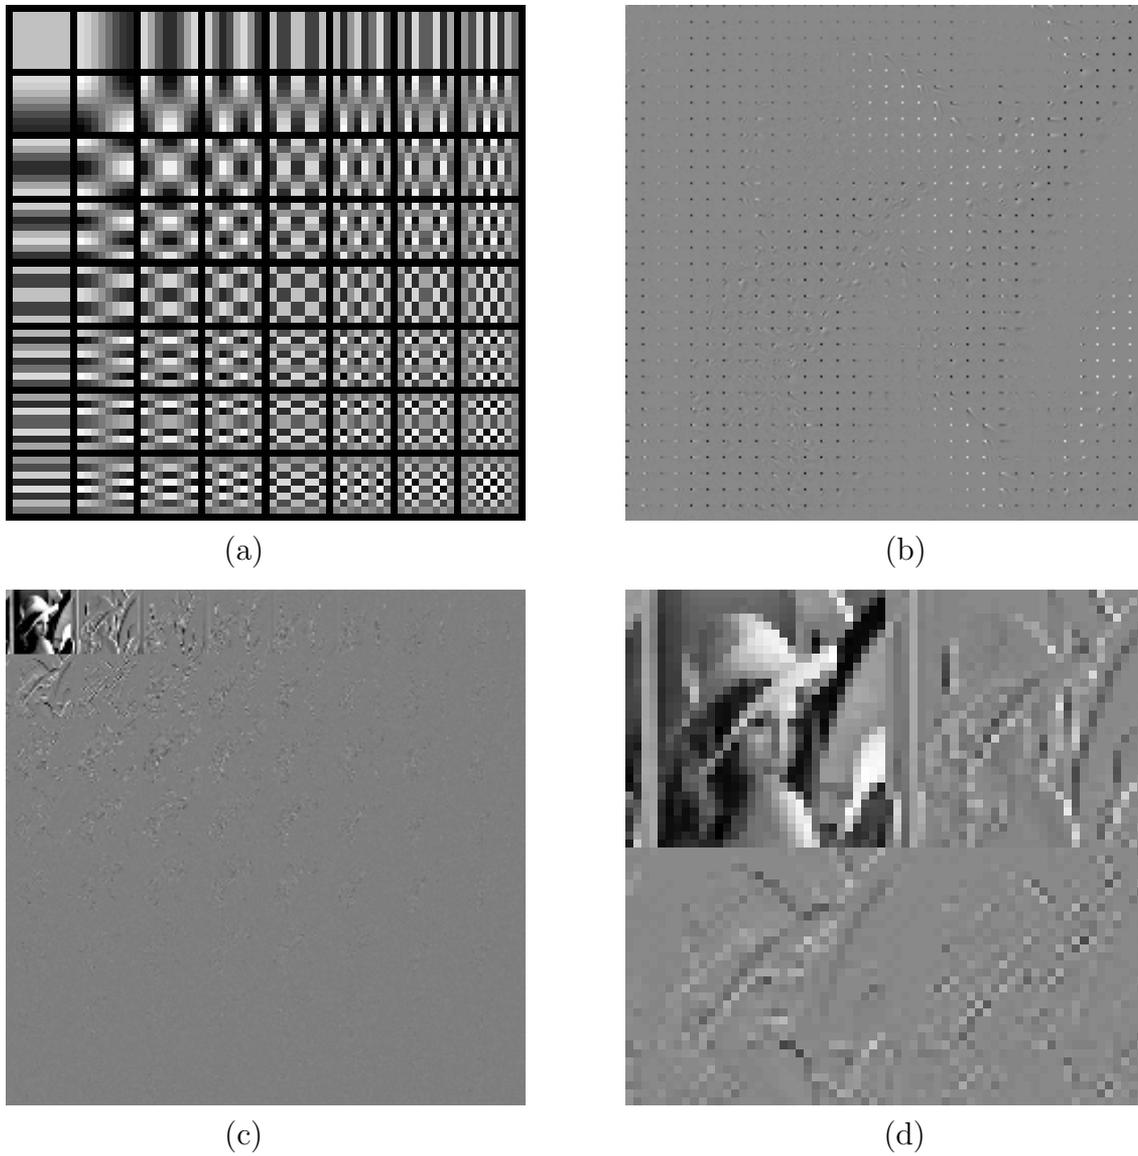


Fig 3.2: (a) Basis functions of the  $8 \times 8$  DCT; (b) Lenna transformed by the  $8 \times 8$  DCT; (c) reordered into subimages grouped by coefficient type; (d) top left 4 subimages from (c).

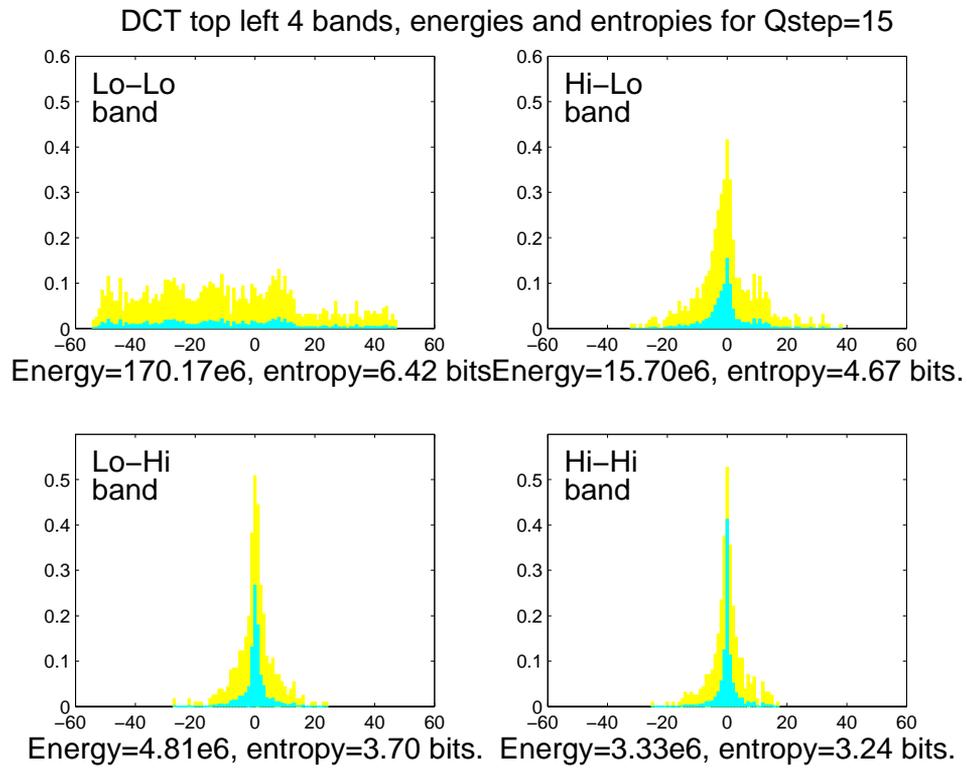


Fig 3.3: The probabilities  $p_i$  and entropies  $h_i$  for the 4 subimages from the top left of the  $8 \times 8$  DCT (fig 3.2d).

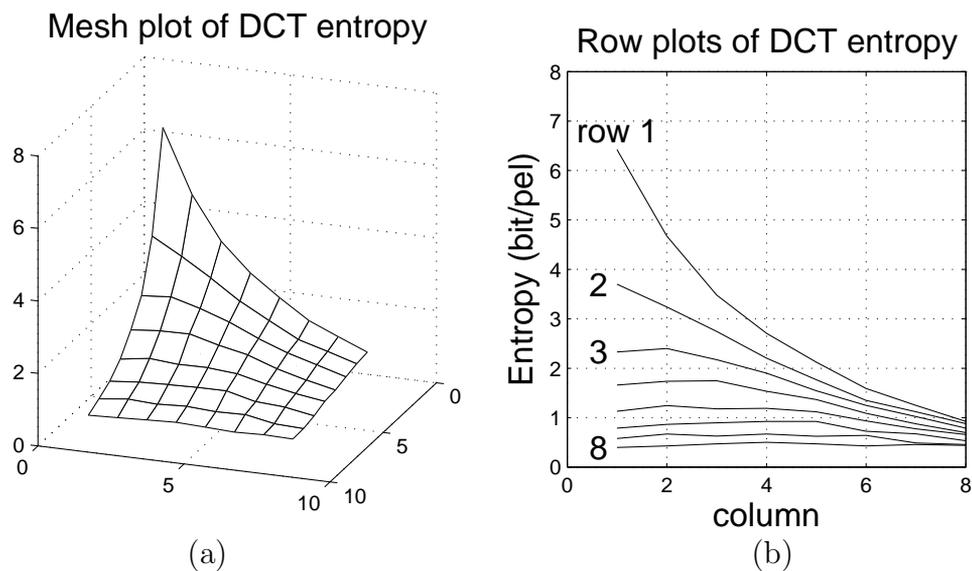


Fig 3.4: (a) Mesh and (b) row plots of the entropies of the subimages of fig 3.2c.

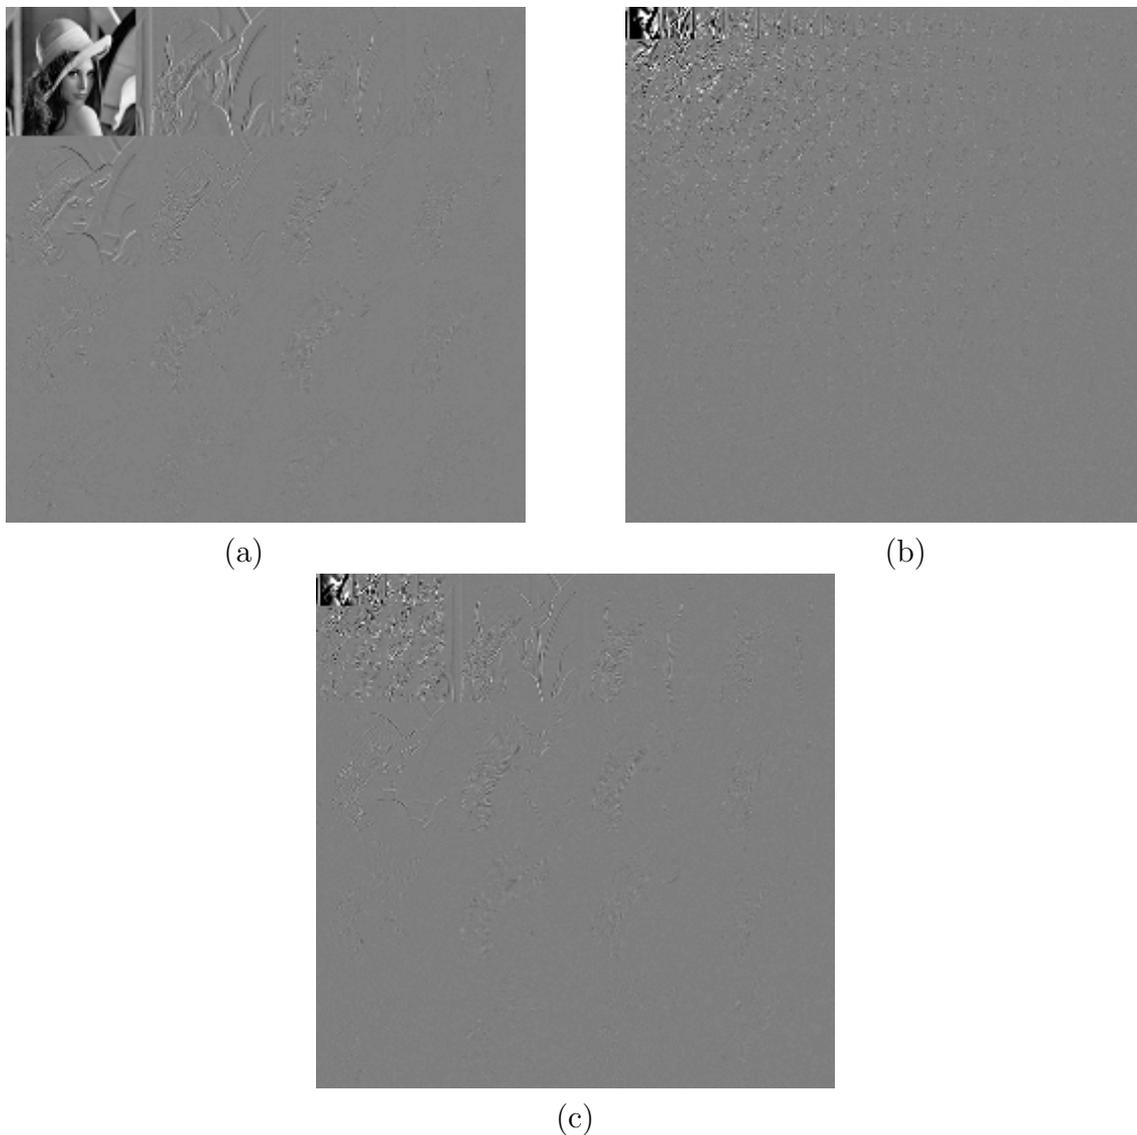


Fig 3.5: Lenna transformed by the  $4 \times 4$  DCT (a),  $16 \times 16$  DCT (b), and 2-level  $4 \times 4$  DCT with prefiltering (c).

### What is the optimum DCT size?

This is a similar question to: What is the optimum number of levels for the Haar transform?

We have analysed Lenna using DCT sizes from  $2 \times 2$  to  $16 \times 16$  to investigate this. Figs 3.5a and 3.5b show the  $4 \times 4$  and  $16 \times 16$  sets of DCT subimages. The  $2 \times 2$  DCT is identical to the level 1 Haar transform (so see fig 2.2b) and the  $8 \times 8$  set is in fig 3.2c.

Fig 3.6 shows the mesh plots of the entropies of the subimages in fig 3.5.

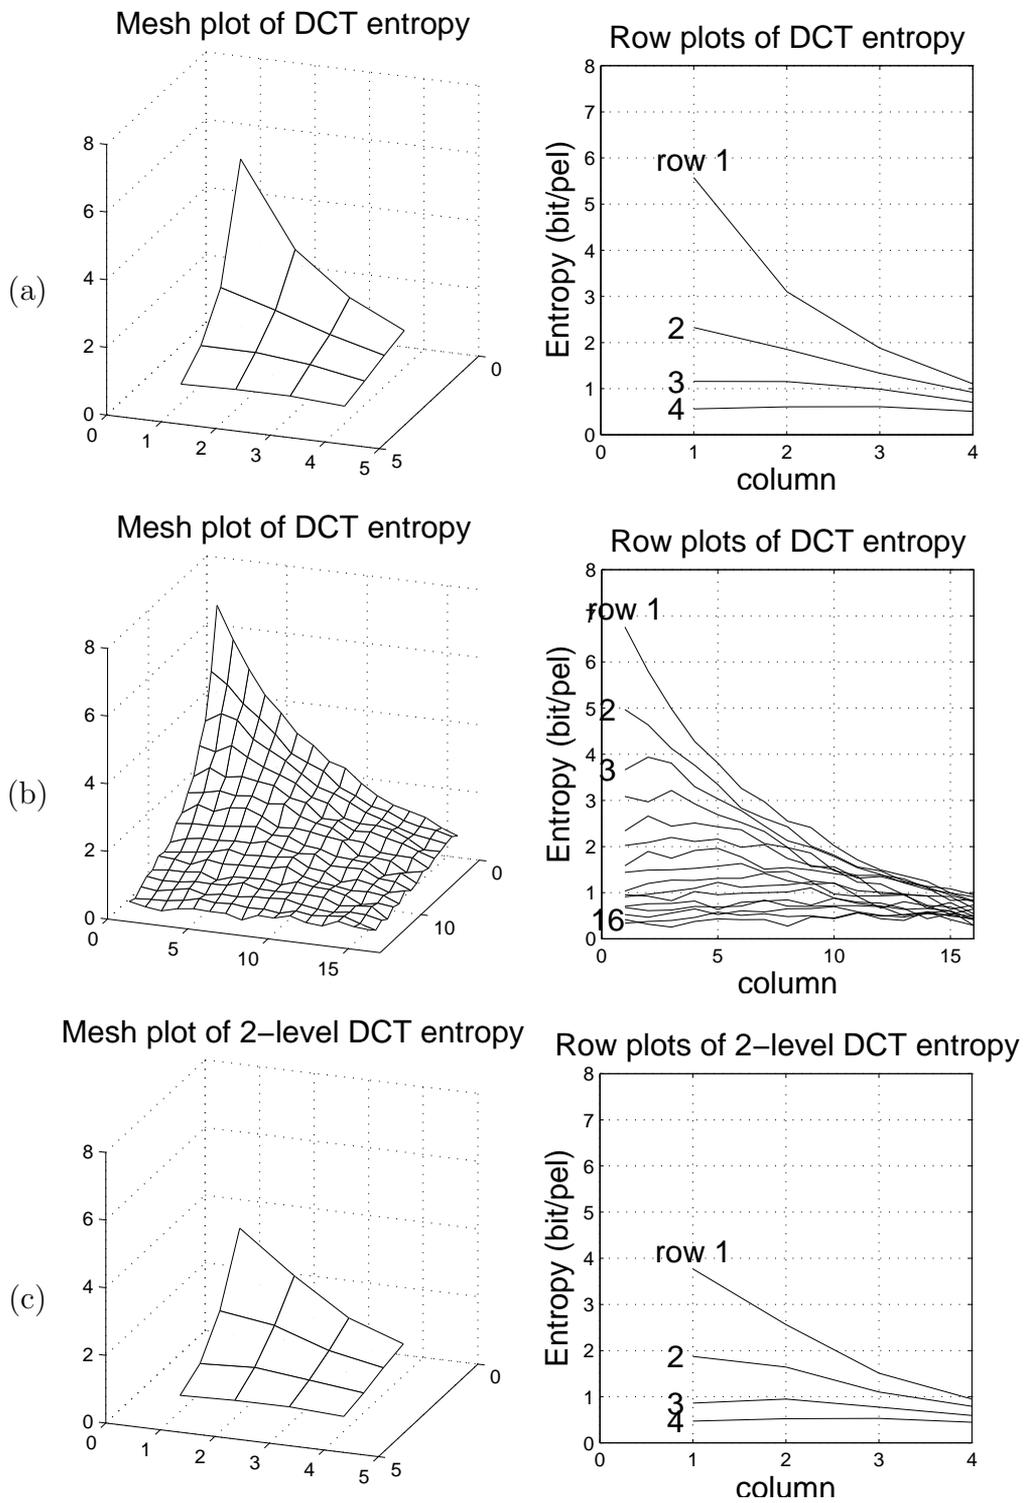


Fig 3.6: Mesh and row plots of the entropies of the  $4 \times 4$  DCT in fig 3.5a, the  $16 \times 16$  DCT in fig 3.5b, and the 2-level  $4 \times 4$  DCT with prefilter in fig 3.5c.

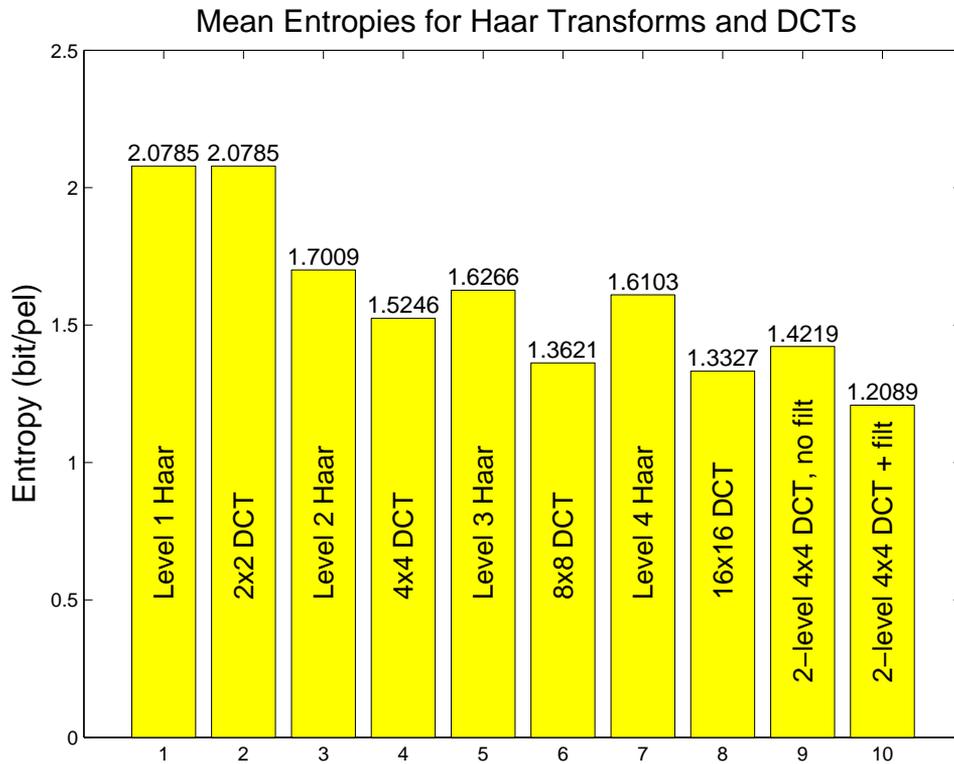


Fig 3.7: Comparison of the mean entropies of the Haar transform of Lenna at levels 1 to 4, and of the DCT for sizes from  $2 \times 2$  to  $16 \times 16$  pels with  $Q_{step} = 15$ . Columns 9 and 10 show the performance of the transform in JPXR.

In fig 3.7, columns 1 to 8 compare the total entropy per pel for the 4 DCT sizes with the equivalent 4 Haar transform sizes. We see that the DCT is significantly better than the rather simpler Haar transform.

As regards the optimum DCT size, from fig 3.7, the  $16 \times 16$  DCT seems to be marginally better than the  $8 \times 8$  DCT, but subjectively this is not the case since quantisation artefacts become more visible as the block size increases. In practice, for a wide range of images and viewing conditions,  **$8 \times 8$  has been found to be a good DCT block size** and is specified in many current coding standards.

We note here that JPEG XR uses a pair of  $4 \times 4$  transform stages to make a 2-level DCT system (like a 2-level Haar system). The second-level DCT is applied only to the lowest frequency coefficients from the level-1 DCT (top left subimage in fig 3.5a). In addition, there is an optional pre-filtering stage that can be applied to the pixels before each DCT stage. The resulting subimages and entropy plots, with the prefilters, are shown in figs 3.5c and 3.6c. The final two columns of fig 3.7 show the entropy of the JPXR 2-level DCT, both without and with the prefilters. We see that the prefilters make a significant difference and allow the JPXR system to outperform the  $8 \times 8$  DCT of JPEG. There are also important subjective advantages to the JPXR system which we discuss later. The paper by Tu et al. in the booklist gives a lot of information on these transforms and filters.

### 3.4 Quantisation of DCT Coefficients

For the previous discussion we assumed a quantiser step size of 15 to allow direct comparison of entropies with the Haar transform. But what step size do we really need?

Figs 3.8a and 3.8b show images reconstructed from the  $8 \times 8$  DCT of Lenna (fig 3.2c), when all the DCT coefficients are quantised with step sizes of 15 and 30 respectively. It is difficult to see quantising artefacts in fig 3.8a ( $Q_{step} = 15$ ) but they are quite noticeable in fig 3.8b ( $Q_{step} = 30$ ).

The visibility of the  $8 \times 8$  DCT basis functions of fig 3.2a has been measured (for a  $720 \times 576$  image viewed from 6 times the image width) and the minimum quantiser steps have been determined which will give artefacts **just at the threshold of visibility**. The matrices (JPEG Book, p37) for the luminance and chrominance threshold step sizes are:

$$\mathbf{Q}_{lum} = \begin{bmatrix} 16 & 11 & 10 & 16 & 24 & 40 & 51 & 61 \\ 12 & 12 & 14 & 19 & 26 & 58 & 60 & 55 \\ 14 & 13 & 16 & 24 & 40 & 57 & 69 & 56 \\ 14 & 17 & 22 & 29 & 51 & 87 & 80 & 62 \\ 18 & 22 & 37 & 56 & 68 & 109 & 103 & 77 \\ 24 & 35 & 55 & 64 & 81 & 104 & 113 & 92 \\ 49 & 64 & 78 & 87 & 103 & 121 & 120 & 101 \\ 72 & 92 & 95 & 98 & 112 & 100 & 103 & 99 \end{bmatrix} \quad (3.7)$$

$$\mathbf{Q}_{chr} = \begin{bmatrix} 17 & 18 & 24 & 47 & 99 & 99 & 99 & 99 \\ 18 & 21 & 26 & 66 & 99 & 99 & 99 & 99 \\ 24 & 26 & 56 & 99 & 99 & 99 & 99 & 99 \\ 47 & 66 & 99 & 99 & 99 & 99 & 99 & 99 \\ 99 & 99 & 99 & 99 & 99 & 99 & 99 & 99 \\ 99 & 99 & 99 & 99 & 99 & 99 & 99 & 99 \\ 99 & 99 & 99 & 99 & 99 & 99 & 99 & 99 \\ 99 & 99 & 99 & 99 & 99 & 99 & 99 & 99 \end{bmatrix} \quad (3.8)$$

Fig 3.8c shows the reconstructed image when each subimage of fig 3.2c is quantised using the corresponding step size from  $\mathbf{Q}_{lum}$ . It is certainly difficult to detect any quantising artefacts, even though many of the step sizes are greater than  $Q_{step} = 30$ , used in fig 3.2b. Fig 3.8d is the reconstructed image using step sizes of  $2 \times \mathbf{Q}_{lum}$  and the artefacts are still quite low.



Fig 3.8: Images reconstructed using the  $8 \times 8$  DCT with (a)  $Q_{step} = 15$ , (b)  $Q_{step} = 30$ , (c)  $Q_{step} = Q_{lum}$ , the JPEG luminance matrix, and (d)  $Q_{step} = 2 \times Q_{lum}$ .

Fig 3.9 shows the entropies of the 64 quantised subimages used to reconstruct each of the four images in fig 3.8. Also given on each plot is the mean entropy (giving the bits/pel for the image) and the rms quantising error between the quantised image and the original.

We see that image (c) has about the same mean entropy and rms error as image (b), but that its quantising artefacts are much less visible. Image (d) has similar visibility of artefacts to (b), but has significantly lower entropy and hence **greater compression** (similarly for images (c) versus (a)).

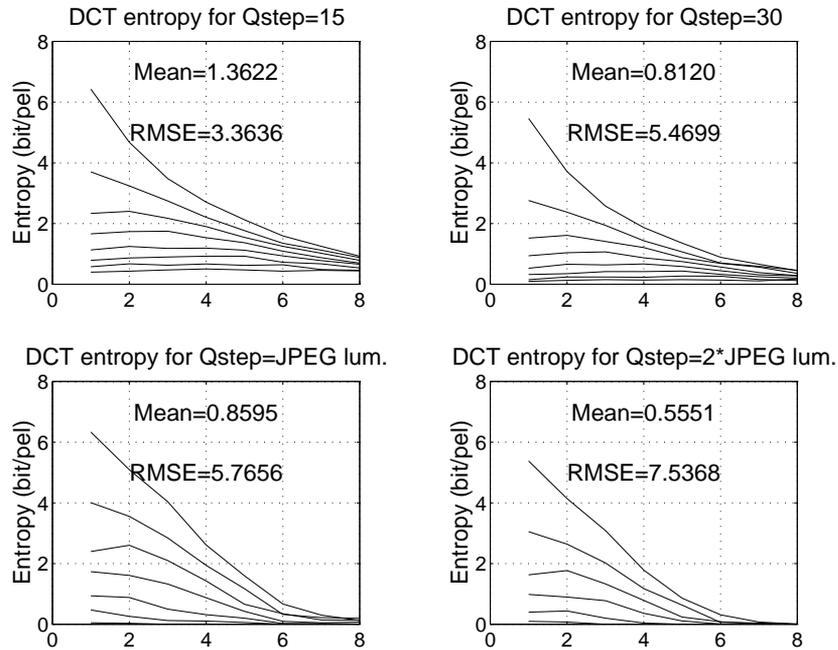


Fig 3.9: Plots of the entropies of the  $8 \times 8$  DCT quantised subimages for the four reconstructed images of fig 3.8.

This shows the distinct advantages of **subjectively weighted quantisation**, and also that it is unwise to rely too much on the rms error as a measure of image quality. After we have introduced wavelets in section 4, we will discuss subjective effects some more and will show how smoothness of the basis functions of a transform is critical to good performance. In fact this is also how JPXR achieves very good performance when its prefilters are enabled, because the filters result in smooth overlapping basis functions, rather than the block-like basis functions of the DCT.

The JPXR results, with prefilters, for the same tests as in fig 3.9 are:

$Q_{step} = 15:$	Mean ent = 1.2089 rms error = 3.2647	$Q_{step} = 30:$	Mean ent = 0.6977 rms error = 5.0711
$Q_{step} = Q_{lum}:$	Mean ent = 0.7323 rms error = 5.7096	$Q_{step} = 2Q_{lum}:$	Mean ent = 0.4552 rms error = 7.5473

These show a reduction of around 15% in entropy for almost no change in rms error, compared with standard JPEG.

### 3.5 JPEG Entropy Coding

The entropy plots of the last section show the theoretical entropies of each DCT sub-band. In practise this would be a poor way to code the data because:

- 64 separate entropy codes would be required (each requiring many extra states to represent run-length coding of zeros).
- The statistics for each code are likely to vary significantly from image to image.
- To transmit the code table for each sub-band as header information would involve a large coding overhead (many extra bits).
- Coding the sub-bands separately does not take account of the correlations which exist between the positions of the non-zero coefs in one sub-band with those of nearby sub-bands (see figs 3.2c and 3.2d).

JPEG uses a clever alternative method of coding, based on **combining run-length and amplitude information into a single Huffman code for the whole of the image** (except the DC sub-band which is coded separately because its statistics are so different).

The code is applied to each block of  $8 \times 8$  quantised DCT coefs from a single  $8 \times 8$  pel region. The blocks are the coefs **before reordering** as shown in fig 3.2b and comprise one coef from each of the 64 sub-bands.

Each block of  $8 \times 8$  quantised coefs is formed into a 1-D vector by zig-zag scanning in the sequence:

0	1	5	6	14	15	27	28
2	4	7	13	16	26	29	42
3	8	12	17	25	30	41	43
9	11	18	24	31	40	44	53
10	19	23	32	39	45	52	54
20	22	33	38	46	51	55	60
21	34	37	47	50	56	59	61
35	36	48	49	57	58	62	63

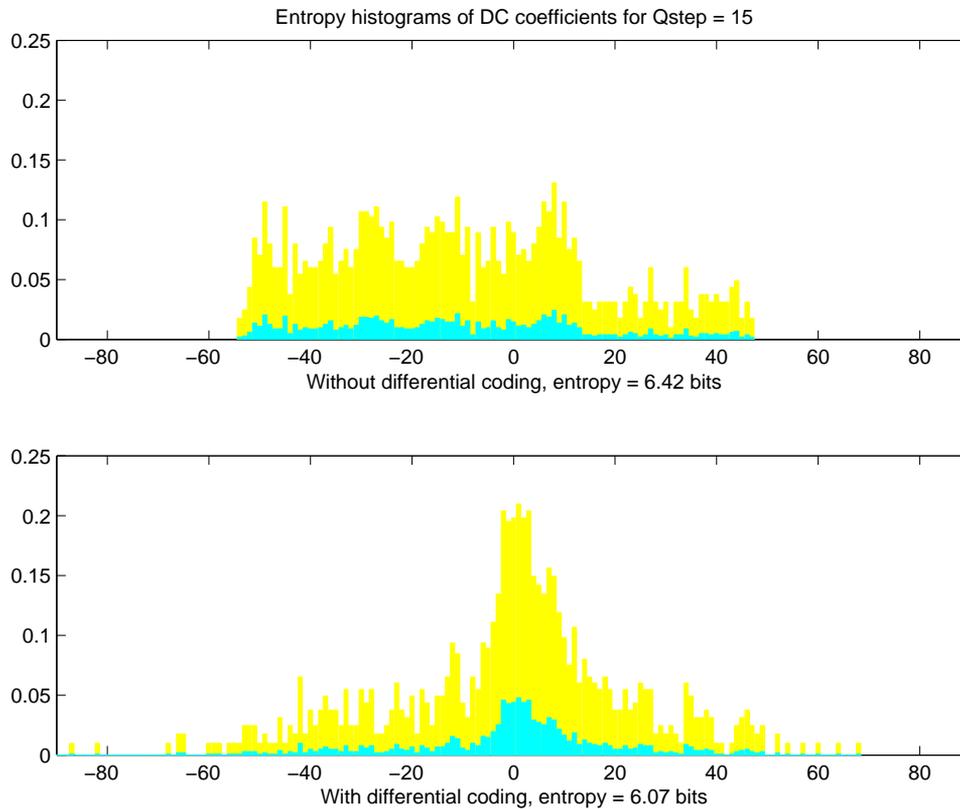


Fig 3.10: Histograms of the DC coefficients from the  $8 \times 8$  DCT of Lenna, showing the entropy reduction with differential coding.

### The JPEG Code for DC coeffs

The first coefficient (0) of each block (vector) is the DC coef, which represents the mean value of the pels in the block (see the top left basis function in fig 3.2a).

The DC coeffs still exhibit significant local correlations (top left of fig 3.2d), so **differential coding** is used in which the value to be coded is the **difference between the current DC coef and that of the previous block** – the blocks are scanned from left to right, row by row. The first block in each row is coded with respect to zero.

The histogram of entropies of the DC coef differences is compared in fig 3.10 with that of the raw DC coeffs from fig 3.3. We note the histogram peak around zero and see that the entropy is reduced from 6.42 bits to 6.07 bits.

The size of the differences can in theory be up to  $\pm 255 \times 8 = \pm 2040$  if the input pels occupy the range  $-128$  to  $+127$  (because the DCT has a gain of 8 at very low frequencies). Hence the Huffman code table would have to be quite large. JPEG adopts a much smaller code by using a form of floating-point representation, where **Size** is the base-2 exponent and **Additional Bits** are used to code the polarity and precise amplitude as follows:

DC Coef Difference	Size	Typical Huffman codes for Size	Additional Bits (in binary)
0	0	00	—
-1, 1	1	010	0, 1
-3, -2, 2, 3	2	011	00, 01, 10, 11
-7, ..., -4, 4, ..., 7	3	100	000, ..., 011, 100, ..., 111
-15, ..., -8, 8, ..., 15	4	101	0000, ..., 0111, 1000, ..., 1111
⋮	⋮	⋮	
-1023, ..., -512, 512, ..., 1023	10	1111 1110	00 0000 0000, ..., 11 1111 1111
-2047, ..., -1024, 1024, ..., 2047	11	1 1111 1110	000 0000 0000, ..., 111 1111 1111

Only Size needs to be Huffman coded in the above scheme, since, within a given Size, all the input values have sufficiently similar probabilities for there to be little gain from entropy coding the Additional Bits (hence they are coded in simple binary as listed). Each coded Size is followed by the appropriate number of Additional Bits (equal to Size) to define the sign and magnitude of the coefficient difference exactly.

There are **only 12 values for Size** to be Huffman coded, so specifying the code table can be **very simple** and require relatively few bits in the header.

In JPEG all Huffman code tables are defined in the image header. Each table requires  $16 + n$  bytes, where  $n$  is the number of codewords in the table.

The first 16 bytes list the number of codewords of each length from 1 to 16 bits (codewords longer than 16 bits are forbidden). The remaining  $n$  bytes list the decoded output values of the  $n$  codewords in ascending codeword order ( $n < 256$ ).

Hence  $16 + 12 = 28$  bytes are needed to specify the code table for DC coefficients.

## The JPEG Run-Amplitude Code

The remaining 63 coefs (the AC coefs) of each 64-element vector usually contain many zeros and so are coded with a **combined run-amplitude Huffman code**.

The codeword represents the run-length of zeros before a non-zero coef **and** the Size of that coef. This is then followed by the Additional Bits which define the coef amplitude and sign precisely. Size and Additional Bits are defined just as for DC coefs.

This 2-dimensional Huffman code (Run, Size) is efficient because there is a **strong correlation** between the Size of a coef and the expected Run of zeros which precedes it – small coefs usually follow long runs; larger coefs tend to follow shorter runs. No single 2-D event is so probable that the Huffman code becomes inefficient.

In order to keep the code table size  $n$  below 256, only the following Run and Size values are coded:

$$\text{Run} = 0 \rightarrow 15 \quad \text{Size} = 1 \rightarrow 10$$

These require 160 codes. Two extra codes, corresponding to (Run,Size) = (0,0) and (15,0) are used for EOB (End-of-block) and ZRL (Zero run length).

EOB is transmitted after the last non-zero coef in a 64-vector. It is the most efficient way of coding the final run of zeros. It is omitted in the rare case that the final element of the vector is non-zero.

ZRL is transmitted whenever  $\text{Run} > 15$ , and represents a run of 16 zeros (15 zeros and a zero amplitude coef) which can be part of a longer run of any length. Hence a run of 20 zeros followed by -5 would be coded as

$$(\text{ZRL}) \ (4,3) \ 010$$

When the code tables are defined in the image header, each codeword is assigned to a given (Run,Size) pair by making the decoded output byte **Code Byte** equal to  $(16 \times \text{Run} + \text{Size})$ .

The default JPEG code for (Run,Size) of AC luminance DCT coefficients is summarised below in order of decreasing code probability:

(Run,Size)	Code Byte (hex)	Code Word (binary)	(Run,Size)	Code Byte (hex)	Code Word (binary)
(0, 1)	01	00	(0, 6)	06	1111000
(0, 2)	02	01	(1, 3)	13	1111001
(0, 3)	03	100	(5, 1)	51	1111010
(EOB)	00	1010	(6, 1)	61	1111011
(0, 4)	04	1011	(0, 7)	07	11111000
(1, 1)	11	1100	(2, 2)	22	11111001
(0, 5)	05	11010	(7, 1)	71	11111010
(1, 2)	12	11011	(1, 4)	14	111110110
(2, 1)	21	11100		⋮	
(3, 1)	31	111010	(ZRL)	F0	1111111001
(4, 1)	41	111011		⋮	

As an example, let us code the following  $8 \times 8$  block:

-13	-3	2	0	0	0	1	0
6	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
-1	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0

Converting this to (DC Size) or (Run,Size) and values for the Additional Bits gives:

(4)	-13	(0, 2)	-3	(0, 3)	6	(2, 2)	2	(3, 1)	-1	(ZRL)	(1, 1)	1	(EOB)
101	0010	01	00	100	110	11111001	10	111010	0	1111111001	1100	1	1010

The compressed bitstream for this block is listed on the lower line, assuming that the default Huffman code tables, given above, are used.

Fig 3.11 shows the histogram of probabilities for the (Run,Size) codewords used to code Lenna using the  $\mathbf{Q}_{lum}$  quantisation matrix. The bin number represents the decoded byte value.

Fig 3.12 shows the equivalent histogram when the quantisation matrix is  $2\mathbf{Q}_{lum}$ .

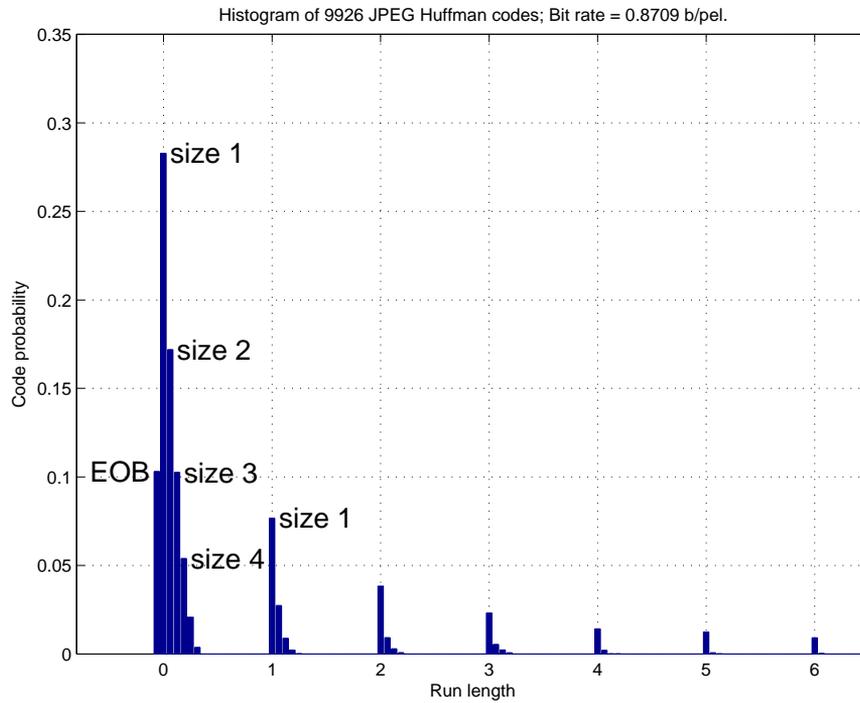


Fig 3.11: Histogram of the (Run,Size) codewords for the DCT of Lenna, quantised using  $Q_{lum}$ .

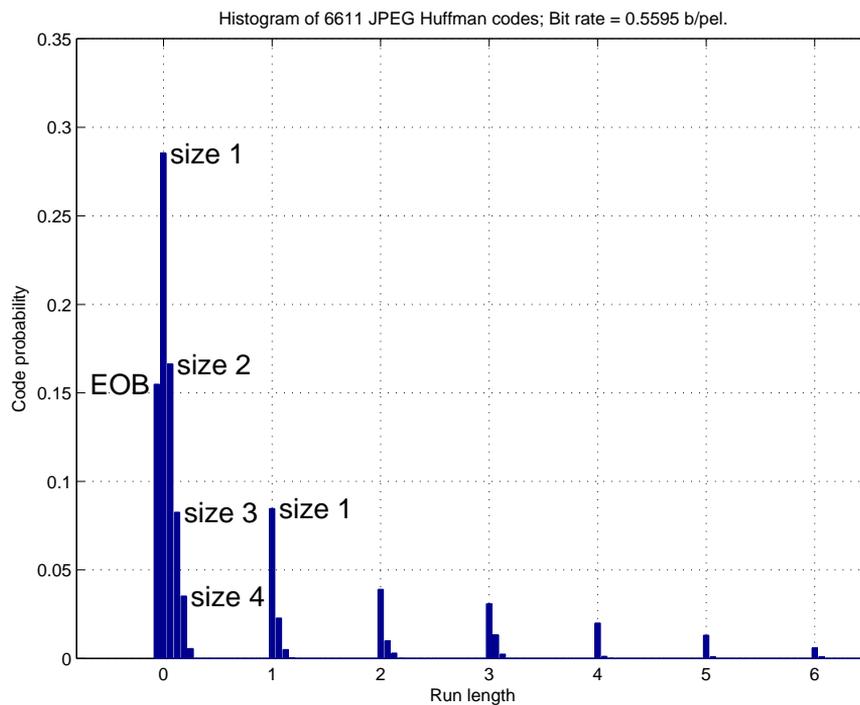


Fig 3.12: Histogram of the (Run,Size) codewords for the DCT of Lenna, quantised using  $2Q_{lum}$ .

Note the strong similarity between these histograms, despite the fact that fig 3.12 represents only  $\frac{2}{3}$  as many events. Only the EOB probability changes significantly, because its probability goes up as the number of events (non-zero coefs) per block goes down.

It turns out that the (Run,Size) histogram remains relatively constant over a wide range of image material and across different regions of each image. This is because of the strong correlation between the run lengths and expected coef sizes. The number of events per block varies considerably depending on the local activity in the image, but the probability distribution of those events (except for EOB) changes much less.

Figs 3.11 and 3.12 also give the mean bit rates to code Lenna for the two quantisation matrices. Comparing these with the theoretical entropies from fig 3.9 (lower row) we get:

Q matrix	Mean Entropy b/pel	JPEG Bit Rate b/pel	JPEG efficiency
$Q_{lum}$	0.8595	0.8709	98.7%
$2Q_{lum}$	0.5551	0.5595	99.21%

Hence we see the **high efficiency of the (Run,Size) code** at two quite different compression factors. This tends to apply over a wide range of images and compression factors and is an impressive achievement.

There is even very little efficiency lost if a single code table is used for many images, which can avoid the need to transmit the  $16 + n$  bytes (168 bytes) of code definition in the header of each image. Using the recommended JPEG default luminance tables (Annex K.3.3) the above efficiencies drop to 97.35% and 95.74% respectively.

### 3.6 Sync and Headers

We have described how individual  $8 \times 8$  blocks of DCT coefficients are coded. Now we shall briefly look at the **sync codes and header information** that are needed in order to complete the coding process.

JPEG is rather complex in this aspect, so we shall just give an overview of the basic principles (see the JPEG Book, chapter 7 for the full picture).

JPEG data is divided into **segments**, each of which starts with a **2-byte marker**.

All markers are byte-aligned – they start on the byte boundaries of the transmission/storage medium. Any variable-length data which precedes a marker is padded with extra ones to achieve this.

The first byte of each marker is  $FF_H$ . The second byte defines the type of marker.

To allow for **recovery in the presence of errors**, it must be possible to detect markers without decoding all of the intervening data. Hence markers must be unique. To achieve this, if an  $FF_H$  byte occurs in the middle of a segment, an extra  $00_H$  **stuffed byte** is inserted after it and  $00_H$  is never used as the second byte of a marker.

Some important markers in the order they are often used are:

Name	Code (hex)	Purpose
SOI	$FFD8$	Start of image.
COM	$FFFE$ $L_{seg}, <\text{Text comments}>$	Comment (segment ignored by decoder).
DQT	$FFDB$ $L_{seg}, < \mathbf{Q}_{lum}, \mathbf{Q}_{chr} \dots >$	Define quantisation table(s).
SOF <sub>0</sub>	$FFC0$ $L_{seg}, <\text{Frame size, no. of components (colours), sub-sampling factors, Q-table selectors}>$	Start of Baseline DCT frame.
DHT	$FFC4$ $L_{seg}, <\text{DC Size and AC (Run,Size) tables for each component}>$	Define Huffman table(s).
SOS	$FFDA$ $L_{seg}, <\text{Huffman table selectors for each component}>$ $<\text{Entropy coded DCT blocks}>$	Start of scan.
EOI	$FFD9$	End of image.

In this table the data which follows each marker is shown between  $\langle \rangle$  brackets. The first 2-byte word of most segments is the length (in bytes) of the segment,  $L_{seg}$ . The length of  $\langle$ Entropy coded DCT blocks $\rangle$ , which forms the main bulk of the compressed data, is not specified explicitly, since it may be determined by decoding the entropy codes. This also allows the data to be transmitted with minimal delay, since it is not necessary to determine the total length of the compressed data before any of the DCT block data can be sent.

Long blocks of entropy-coded data are rather prone to being corrupted by transmission errors. To mitigate the worst aspects of this, **Restart Markers** ( $FFD0 \rightarrow FFD7$ ) may be included at regular intervals (say at the start of each row of DCT blocks in the image) so that separate parts of the entropy coded stream may be decoded independently of errors in other parts. The restart interval, if required, is defined by a DRI ( $FFDD$ ) marker segment. There are 8 restart markers, which are used in sequence, so that if one (or more) is corrupted by errors, its absence may be easily detected.

The use of **multiple scans** within each image frame and **multiple frames** within a given image allows many variations on the ordering and interleaving of the compressed data. For example:

- Chrominance and luminance components may be sent in separate scans or interleaved into a single scan.
- Lower frequency DCT coefs may be sent in one or more scans before higher frequency coefs.
- Coarsely quantised coefs may be sent in one or more scans before finer (refinement) coefs.
- A coarsely sampled frame of the image may be sent initially and then the detail may be progressively improved by adding differentially-coded correction frames of increasing resolution.

JPXR includes even more options than this, including the following:

- Lossless coding is available as an alternative to lossy coding for critical applications such as medical imaging. This is achieved by clever use of reversible *Integer-based Lifting* in all transforms and stages of filtering.
- Large images may be *tiled* for greater ease of processing and reduction in processor memory requirements.
- Arbitrary numbers of colours or transparency maps may be coded at a wide range of amplitude resolution levels, to handle high-dynamic-range imagery of all forms.
- The overlapped prefilters may be enabled / disabled for different image regions as required by the source material.

## 4 Filter Banks and Wavelets

**Digital filter banks** have been actively studied since the 1960s, whereas **Wavelet theory** is a new subject area that was developed in the 1980s, principally by French and Belgian mathematicians, notably Y Meyer, I Daubechies, and S Mallat. The two topics are now firmly linked and of great importance for signal analysis and compression. They are becoming increasingly used in computer vision systems too.

### 4.1 The 2-band Filter Bank

Recall the 1-D Haar transform from section 2.1.

$$\begin{bmatrix} y(1) \\ y(2) \end{bmatrix} = \mathbf{T} \begin{bmatrix} x(1) \\ x(2) \end{bmatrix} \quad \text{where} \quad \mathbf{T} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \quad (4.1)$$

We can write this in expanded form as:

$$\begin{aligned} y(1) &= \frac{1}{\sqrt{2}}x(1) + \frac{1}{\sqrt{2}}x(2) \\ y(2) &= \frac{1}{\sqrt{2}}x(1) - \frac{1}{\sqrt{2}}x(2) \end{aligned} \quad (4.2)$$

More generally if  $\mathbf{x}$  is a longer sequence and the results are placed in two separate sequences  $\mathbf{y}_0$  and  $\mathbf{y}_1$ , we define the process as:

$$\begin{aligned} y_0(n) &= \frac{1}{\sqrt{2}}x(n-1) + \frac{1}{\sqrt{2}}x(n) \\ y_1(n) &= \frac{1}{\sqrt{2}}x(n-1) - \frac{1}{\sqrt{2}}x(n) \end{aligned} \quad (4.3)$$

These can be expressed as 2 FIR filters with tap vectors  $\mathbf{h}_0 = [\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}]$  and  $\mathbf{h}_1 = [\frac{1}{\sqrt{2}}, \frac{-1}{\sqrt{2}}]$ .

Hence as z-transforms, equations (4.3) become:

$$\begin{aligned} Y_0(z) &= H_0(z) X(z) \quad \text{where} \quad H_0(z) = \frac{1}{\sqrt{2}}(z^{-1} + 1) \\ Y_1(z) &= H_1(z) X(z) \quad \text{where} \quad H_1(z) = \frac{1}{\sqrt{2}}(z^{-1} - 1) \end{aligned} \quad (4.4)$$

(We shall later extend these filters to be more complicated.)

In practice, we only calculate  $y_0(n)$  and  $y_1(n)$  at alternate (say even) values of  $n$  so that the total number of samples in  $\mathbf{y}_0$  and  $\mathbf{y}_1$  is the same as in  $\mathbf{x}$ .

We may thus represent the Haar transform operation by a pair of filters followed by down-sampling by 2, as shown in fig 4.1a. This is known as a **2-band analysis filter bank**.

In equation (2.3), to reconstruct  $\mathbf{x}$  from  $\mathbf{y}$  we calculated  $\mathbf{x} = \mathbf{T}^T \mathbf{y}$ . For long sequences this may be written:

$$\begin{aligned} x(n-1) &= \frac{1}{\sqrt{2}}y_0(n) + \frac{1}{\sqrt{2}}y_1(n) \\ x(n) &= \frac{1}{\sqrt{2}}y_0(n) - \frac{1}{\sqrt{2}}y_1(n) \end{aligned} \quad \text{for } n \text{ even.} \quad (4.5)$$

Since  $y_0(n)$  and  $y_1(n)$  are only calculated at even values of  $n$ , we may assume that they are zero at odd values of  $n$ . We may then combine equations (4.5) into a single expression for  $x(n)$ , valid for all  $n$ :

$$x(n) = \frac{1}{\sqrt{2}}[y_0(n+1) + y_0(n)] + \frac{1}{\sqrt{2}}[y_1(n+1) - y_1(n)] \quad (4.6)$$

or as z-transforms:

$$X(z) = G_0(z) Y_0(z) + G_1(z) Y_1(z) \quad (4.7)$$

where

$$G_0(z) = \frac{1}{\sqrt{2}}(z+1) \quad \text{and} \quad G_1(z) = \frac{1}{\sqrt{2}}(z-1) \quad (4.8)$$

In (4.7) the signals  $Y_0(z)$  and  $Y_1(z)$  are not really the same as  $Y_0(z)$  and  $Y_1(z)$  in (4.4) because those in (4.4) have not (yet) had alternate samples set to zero. Also, in (4.7)  $X(z)$  is the reconstructed output whereas in (4.4) it is the input signal.

To avoid confusion we shall use  $\hat{X}$ ,  $\hat{Y}_0$  and  $\hat{Y}_1$  for the signals in (4.7) so it becomes:

$$\hat{X}(z) = G_0(z) \hat{Y}_0(z) + G_1(z) \hat{Y}_1(z) \quad (4.9)$$

We may show this reconstruction operation as upsampling followed by 2 filters, as in fig 4.1b, forming a **2-band reconstruction filter bank**.

If  $\hat{Y}_0$  and  $\hat{Y}_1$  are not the same as  $Y_0$  and  $Y_1$ , how do they relate to each other?

Now

$$\hat{y}_0(n) = y_0(n) \quad \text{for } n \text{ even,} \quad \hat{y}_0(n) = 0 \quad \text{for } n \text{ odd.} \quad (4.10)$$

Therefore  $\hat{Y}_0(z)$  is a polynomial in  $z$ , comprising **only** the terms in even powers of  $z$  from  $Y_0(z)$ . This may be written as:

$$\hat{Y}_0(z) = \sum_{\text{even } n} y_0(n) z^{-n} = \sum_{\text{all } n} \frac{1}{2}[y_0(n) z^{-n} + y_0(n) (-z)^{-n}] = \frac{1}{2}[Y_0(z) + Y_0(-z)] \quad (4.11)$$

Similarly

$$\hat{Y}_1(z) = \frac{1}{2}[Y_1(z) + Y_1(-z)] \quad (4.12)$$

This is our general model for downsampling by 2, followed by upsampling by 2 as defined in equation (4.10).

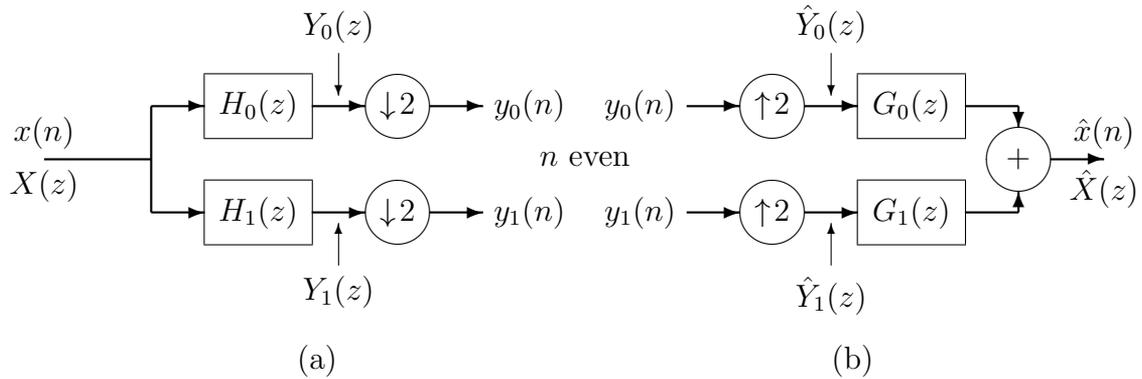


Fig 4.1: Two-band filter banks for analysis (a) and reconstruction (b).

## 4.2 Perfect Reconstruction (PR)

We are now able to **generalise our analysis for arbitrary filters**  $H_0$ ,  $H_1$ ,  $G_0$  and  $G_1$ . Substituting (4.11) and (4.12) into (4.9) and then using (4.4), we get:

$$\begin{aligned}
 \hat{X}(z) &= \frac{1}{2}G_0(z)[Y_0(z) + Y_0(-z)] + \frac{1}{2}G_1(z)[Y_1(z) + Y_1(-z)] \\
 &= \frac{1}{2}G_0(z)H_0(z)X(z) + \frac{1}{2}G_0(z)H_0(-z)X(-z) \\
 &\quad + \frac{1}{2}G_1(z)H_1(z)X(z) + \frac{1}{2}G_1(z)H_1(-z)X(-z) \\
 &= \frac{1}{2}X(z)[G_0(z)H_0(z) + G_1(z)H_1(z)] \\
 &\quad + \frac{1}{2}X(-z)[G_0(z)H_0(-z) + G_1(z)H_1(-z)]
 \end{aligned} \tag{4.13}$$

If we require  $\hat{X}(z) \equiv X(z)$  — **the Perfect Reconstruction (PR) condition** — then:

$$G_0(z)H_0(z) + G_1(z)H_1(z) \equiv 2 \tag{4.14}$$

and

$$G_0(z)H_0(-z) + G_1(z)H_1(-z) \equiv 0 \tag{4.15}$$

Identity (4.15) is known as the **anti-aliasing condition** because the term in  $X(-z)$  in (4.13) is the unwanted aliasing term caused by down-sampling  $\mathbf{y}_0$  and  $\mathbf{y}_1$  by 2.

It is straightforward to show that the expressions for  $H_0$ ,  $H_1$ ,  $G_0$  and  $G_1$ , given in (4.4) and (4.8) for the filters based on the Haar transform, satisfy (4.14) and (4.15). They are the simplest set of filters which do.

Before we look at more complicated PR filters, we examine how the filter structures of fig 4.1 may be extended to form a binary filter tree (and the discrete wavelet transform).

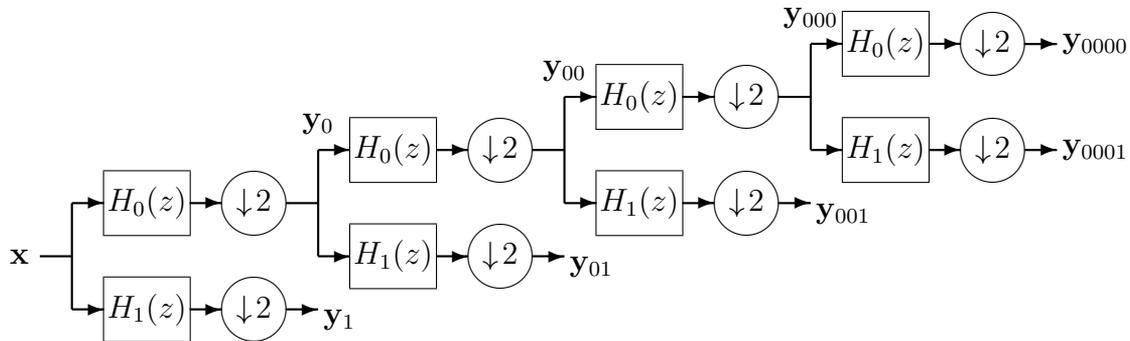


Fig 4.2: Extension of the 2-band filter bank into a binary filter tree.

### 4.3 The binary filter tree

Recall that for image compression, the purpose of the 2-band filter bank in the Haar transform is to compress most of the signal energy into the low-frequency band.

We may achieve **greater compression if the low band is further split into two**. This may be repeated a number of times to give the binary filter tree, shown with 4 levels in fig 4.2.

In 1-D, this is analogous to the way the 2-D Haar transform of section 2.1 was extended to the multi-level Haar transform in section 2.3.

For an  $N$ -sample input vector  $\mathbf{x}$ , the sizes and bandwidths of the signals of the 4-level filter tree are:

Signal	No. of samples	Approximate pass band
$\mathbf{x}$	$N$	$0 \rightarrow \frac{1}{2} f_s$
$\mathbf{y}_1$	$N/2$	$\frac{1}{4} \rightarrow \frac{1}{2} f_s$
$\mathbf{y}_{01}$	$N/4$	$\frac{1}{8} \rightarrow \frac{1}{4} f_s$
$\mathbf{y}_{001}$	$N/8$	$\frac{1}{16} \rightarrow \frac{1}{8} f_s$
$\mathbf{y}_{0001}$	$N/16$	$\frac{1}{32} \rightarrow \frac{1}{16} f_s$
$\mathbf{y}_{0000}$	$N/16$	$0 \rightarrow \frac{1}{32} f_s$

Because of the downsampling by 2 at each level, the **total number of output samples equals  $N$** , regardless of the number of levels in the tree; so the process is **non-redundant**.

The  $H_0$  filter is normally designed to be a **lowpass** filter with a passband from 0 to approximately  $\frac{1}{4}$  of the input sampling frequency for that stage; and  $H_1$  is a **highpass** (bandpass) filter with a pass band approximately from  $\frac{1}{4}$  to  $\frac{1}{2}$  of that same frequency.

When formed into a 4-level tree, the filter outputs have the approximate pass bands given in the above table. The final output  $\mathbf{y}_{0000}$  is a lowpass signal, while the other outputs are all bandpass signals, each covering a band of approximately one octave.

An **inverse tree**, mirroring fig 4.2, may be constructed using filters  $G_0$  and  $G_1$  instead of  $H_0$  and  $H_1$ , as shown for just one level in fig 4.1b. If the PR conditions of (4.14) and (4.15) are satisfied, then the output of each level will be identical to the input of the equivalent level in fig 4.2, and the final output will be a perfect reconstruction of the input signal.

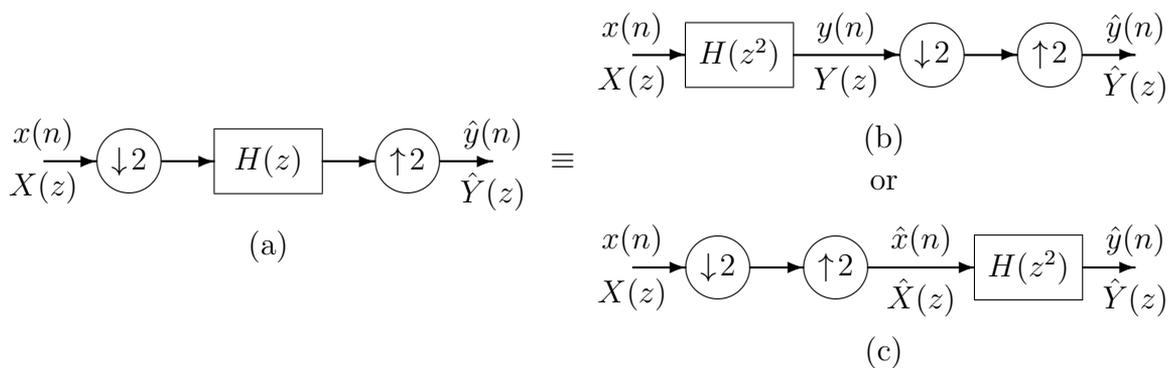


Fig 4.3: Multi-rate filtering — the result of shifting a filter ahead of a downsampling operation or after an upsampling operation.

### Multi-rate filtering theorem

To calculate the impulse and frequency responses for a multistage network with downsampling at each stage, as in fig 4.2, we must first derive an important theorem for multi-rate filters.

#### Theorem:

The downsample–filter–upsample operation of fig 4.3a is equivalent to either the filter–downsample–upsample operation of fig 4.3b or the downsample–upsample–filter operation of fig 4.3c, if the filter is changed from  $H(z)$  to  $H(z^2)$ .

**Proof:**

From fig 4.3a:

$$\begin{aligned}\hat{y}(n) &= \sum_i x(n-2i) h(i) \quad \text{for } n \text{ even} \\ &= 0 \quad \text{for } n \text{ odd}\end{aligned}\tag{4.16}$$

Take z-transforms:

$$\hat{Y}(z) = \sum_n \hat{y}(n) z^{-n} = \sum_{\text{even } n} \sum_i x(n-2i) h(i) z^{-n}\tag{4.17}$$

Reverse the order of summation and let  $m = n - 2i$ :

$$\begin{aligned}\therefore \hat{Y}(z) &= \sum_i h(i) \left( \sum_{\text{even } m} x(m) z^{-m} z^{-2i} \right) \\ &= \left( \sum_i h(i) z^{-2i} \right) \left( \sum_{\text{even } m} x(m) z^{-m} \right) \\ &= H(z^2) \frac{1}{2}[X(z) + X(-z)] \\ &= \frac{1}{2}[H(z^2) X(z) + H((-z)^2) X(-z)] \\ &= \frac{1}{2}[Y(z) + Y(-z)] \quad \text{where } Y(z) = H(z^2) X(z)\end{aligned}\tag{4.18}$$

This describes the operations of fig 4.3b. Hence the first result is proved.

The result from line 3 above

$$\hat{Y}(z) = \frac{1}{2}[X(z) + X(-z)] H(z^2) = \hat{X}(z) H(z^2)\tag{4.19}$$

shows that the filter  $H(z^2)$  may be placed after the down/up-sampler as in fig 4.3c, which proves the second result.

**General results for M:1 subsampling**

It can be shown that:

- $H(z)$  becomes  $H(z^M)$  if shifted ahead of an  $M : 1$  downsampler or following an  $M : 1$  upsampler.
- $M : 1$  down/up-sampling of a signal  $X(z)$  produces:

$$\hat{X}(z) = \frac{1}{M} \sum_{m=0}^{M-1} X(z e^{j2\pi m/M})\tag{4.20}$$

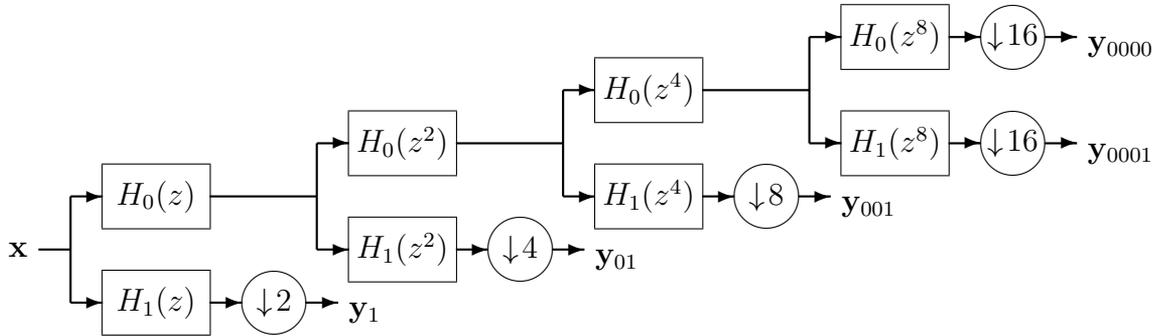


Fig 4.4: Binary filter tree, transformed so that all downsampling operations occur at the outputs.

### Transformation of the filter tree

Using the result of equation (4.18), fig 4.2 can be redrawn as in fig 4.4 with **all downsamplers moved to the outputs**. (Note that fig 4.4 requires **much more computation** than fig 4.2 and so would not be used in practice, but nevertheless it is an important analysis aid.) We can now calculate the transfer function to each output (before the downsamplers) as:

$$\begin{aligned}
 H_{01}(z) &= H_0(z) H_1(z^2) \\
 H_{001}(z) &= H_0(z) H_0(z^2) H_1(z^4) \\
 H_{0001}(z) &= H_0(z) H_0(z^2) H_0(z^4) H_1(z^8) \\
 H_{0000}(z) &= H_0(z) H_0(z^2) H_0(z^4) H_0(z^8)
 \end{aligned} \tag{4.21}$$

In general the transfer functions to the two outputs at level  $k$  of the tree are given by:

$$\begin{aligned}
 H_{k,1} &= \prod_{i=0}^{k-2} H_0(z^{2^i}) H_1(z^{2^{k-1}}) \\
 H_{k,0} &= \prod_{i=0}^{k-1} H_0(z^{2^i})
 \end{aligned} \tag{4.22}$$

For the Haar filters of equation (4.4), the transfer functions to the outputs of the 4-level tree become:

$$\begin{aligned}
 H_{01}(z) &= \frac{1}{2} [ (z^{-3} + z^{-2}) - (z^{-1} + 1) ] \\
 H_{001}(z) &= \frac{1}{2\sqrt{2}} [ (z^{-7} + z^{-6} + z^{-5} + z^{-4}) - (z^{-3} + z^{-2} + z^{-1} + 1) ] \\
 H_{0001}(z) &= \frac{1}{4} [ (z^{-15} + z^{-14} + z^{-13} + z^{-12} + z^{-11} + z^{-10} + z^{-9} + z^{-8}) \\
 &\quad - (z^{-7} + z^{-6} + z^{-5} + z^{-4} + z^{-3} + z^{-2} + z^{-1} + 1) ] \\
 H_{0000}(z) &= \frac{1}{4} [ z^{-15} + z^{-14} + z^{-13} + z^{-12} + z^{-11} + z^{-10} + z^{-9} + z^{-8} \\
 &\quad + z^{-7} + z^{-6} + z^{-5} + z^{-4} + z^{-3} + z^{-2} + z^{-1} + 1 ]
 \end{aligned} \tag{4.23}$$

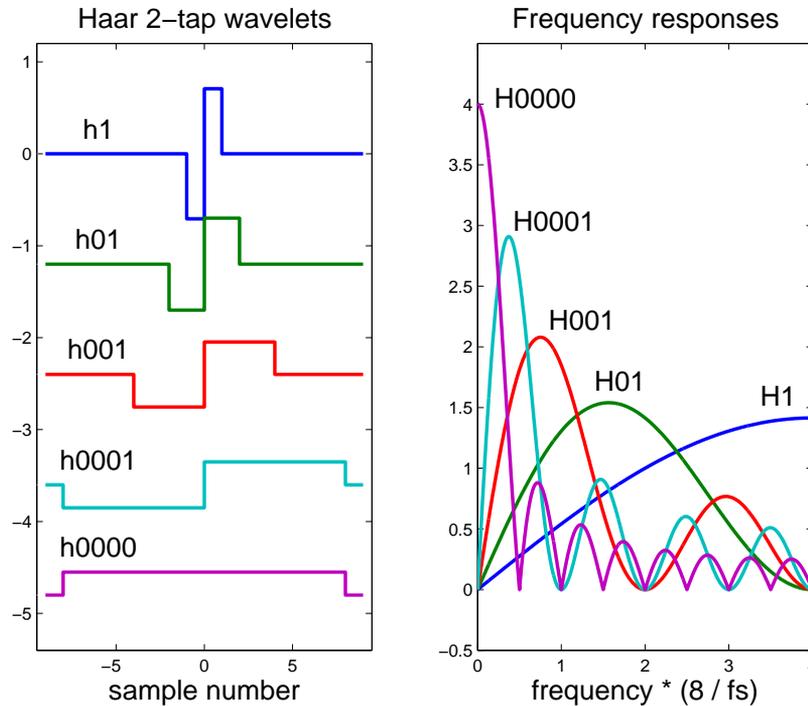


Fig 4.5: Impulse responses and frequency responses of the 4-level tree of Haar filters.

## 4.4 Wavelets

The process of creating the outputs  $\mathbf{y}_1$  to  $\mathbf{y}_{0000}$  from  $\mathbf{x}$  is known as the **discrete wavelet transform (DWT)**; and the reconstruction process is the inverse DWT.

The word **wavelet** refers to the impulse response of the cascade of filters which leads to a given bandpass output. The frequency response of the wavelet at level  $k$  is obtained by substituting  $z = e^{j\omega T_s}$  in the z-transfer function  $H_{k,1}$  from equation (4.22).  $T_s$  is the sampling period at the input to the filter tree.

Since the frequency responses of the bandpass bands are scaled down by 2:1 at each level, their impulse responses become longer by the same factor at each level, BUT their shapes remain very similar. The basic impulse response wave shape is almost independent of scale and is known as the **mother wavelet**.

The impulse response to a lowpass output  $H_{k,0}$  is called the **scaling function** at level  $k$ .

Fig 4.5 shows these effects using the impulse responses and frequency responses for the five outputs of the 4-level tree of Haar filters, based on the z-transforms given in equations (4.23). Notice the abrupt transitions in the middle and at the ends of the Haar wavelets. These result in noticeable **blocking artefacts** in decompressed images (as in fig 2.13b).

## 4.5 Good filters / wavelets

Our main aim now is to search for better filters / wavelets which result in **compression performance that rivals or beats the DCT**.

We assume that perfect reconstruction is a prime requirement, so that the only image degradations are caused by coefficient quantisation, and may be made as small as we wish by increasing bit rate.

We start our search with the two PR identities from section 4.2, equations (4.14) and (4.15), which we repeat here:

$$G_0(z)H_0(z) + G_1(z)H_1(z) \equiv 2 \quad (4.24)$$

and

$$G_0(z)H_0(-z) + G_1(z)H_1(-z) \equiv 0 \quad (4.25)$$

The usual way of satisfying the anti-aliasing condition (4.25), while permitting  $H_0$  and  $G_0$  to have lowpass responses (passband where  $Re[z] > 0$ ) and  $H_1$  and  $G_1$  to have highpass responses (passband where  $Re[z] < 0$ ), is with the following relations:

$$H_1(z) = z^{-k}G_0(-z) \quad \text{and} \quad G_1(z) = z^kH_0(-z) \quad (4.26)$$

where  $k$  must be odd so that:

$$G_0(z)H_0(-z) + G_1(z)H_1(-z) = G_0(z)H_0(-z) + z^kH_0(-z) (-z)^{-k}G_0(z) = 0$$

Now define the lowpass product filter:

$$P(z) = H_0(z)G_0(z) \quad (4.27)$$

and substitute relations (4.26) into identity (4.24) to get:

$$G_0(z)H_0(z) + G_1(z)H_1(z) = G_0(z)H_0(z) + H_0(-z)G_0(-z) = P(z) + P(-z) = 2 \quad (4.28)$$

This requires **all  $P(z)$  terms in even powers of  $z$  to be zero**, except the  $z^0$  term which should be 1. The  $P(z)$  terms in odd powers of  $z$  may take any desired values since they cancel out in (4.28).

A further constraint on  $P(z)$  is that it should be **zero phase**, in order to minimise the visibility of any distortions due to the high-band being quantised to zero. Hence  $P(z)$  should be symmetrical about the  $z^0$  term and be of the form:

$$P(z) = \dots + p_5z^5 + p_3z^3 + p_1z + 1 + p_1z^{-1} + p_3z^{-3} + p_5z^{-5} + \dots \quad (4.29)$$

The design of a set of PR filters  $H_0, H_1$  and  $G_0, G_1$  can now be summarised as:

1. Choose a set of coefficients  $p_1, p_3, p_5 \dots$  to give a zero-phase lowpass product filter  $P(z)$  with desirable characteristics. (This is non-trivial and is discussed below.)
2. Factorize  $P(z)$  into  $H_0(z)$  and  $G_0(z)$ , preferably so that these two filters have similar lowpass frequency responses.
3. Calculate  $H_1(z)$  and  $G_1(z)$  from  $H_0(z)$  and  $G_0(z)$ , using equations (4.26).

It can help to simplify the tasks of choosing  $P(z)$  and factorising it if, based on the zero-phase requirement, we transform  $P(z)$  into  $P_t(Z)$  such that:

$$P(z) = P_t(Z) = 1 + p_{t,1}Z + p_{t,3}Z^3 + p_{t,5}Z^5 + \dots \quad \text{where } Z = \frac{1}{2}(z + z^{-1}) \quad (4.30)$$

We note here that, with this substitution, a polynomial in odd powers of  $Z$  will only produce a polynomial in odd powers of  $z$ , with coefficients that are symmetric about the  $z^0$  term as required, and that the  $z^0$  term is unaffected by the transformation.

To calculate the frequency response of  $P_t$ , let  $z = e^{j\omega T_s}$ :

$$\therefore Z = \frac{1}{2}(e^{j\omega T_s} + e^{-j\omega T_s}) = \cos(\omega T_s) \quad (4.31)$$

This is a purely real function of  $\omega$ , varying from 1 at  $\omega = 0$  to  $-1$  at  $\omega T_s = \pi$  (half the sampling frequency).

The Belgian mathematician, Ingrid Daubechies, did much pioneering work on wavelets in the 1980s. She discovered that, to achieve **smooth wavelets** after many levels of the binary tree, the lowpass filters  $H_0(z)$  and  $G_0(z)$  must both have a number of **zeros at half the sampling frequency** (i.e. at  $z = -1$ ). These will also be zeros of  $P(z)$  at  $z = -1$ , and so  $P_t(Z)$  will have half as many zeros at  $Z = -1$ .

The simplest case is a single zero at  $Z = -1$ , so that  $P_t(Z) = 1 + Z$ .

$$\therefore P(z) = \frac{1}{2}(z + 2 + z^{-1}) = \frac{1}{2}(z + 1)(1 + z^{-1}) = G_0(z) H_0(z)$$

which gives the familiar Haar filters.

As we have seen, the Haar wavelets have significant discontinuities so we need to add more zeros at  $Z = -1$ . However to maintain PR, we must also ensure that all terms in even powers of  $Z$  are zero, so the next more complicated  $P_t$  must be of the form:

$$\begin{aligned} P_t(Z) &= (1 + Z)^2 (1 + aZ) = 1 + (2 + a)Z + (1 + 2a)Z^2 + aZ^3 \\ &= 1 + \frac{3}{2}Z - \frac{1}{2}Z^3 \quad \text{where } a = -\frac{1}{2} \text{ to suppress the term in } Z^2 \end{aligned} \quad (4.32)$$

If we allocate the factors of  $P_t$  such that  $(1 + Z)$  gives  $H_0$  and  $(1 + Z)(1 + aZ)$  gives  $G_0$ , we get:

$$\begin{aligned} H_0(z) &= \frac{1}{2}(z + 2 + z^{-1}) \\ G_0(z) &= \frac{1}{8}(z + 2 + z^{-1})(-z + 4 - z^{-1}) = \frac{1}{8}(-z^2 + 2z + 6 + 2z^{-1} - z^{-2}) \end{aligned} \quad (4.33)$$

Using (4.26) with  $k = 1$ , the corresponding highpass filters then become:

$$\begin{aligned} G_1(z) &= z H_0(-z) = \frac{1}{2}z(-z + 2 - z^{-1}) \\ H_1(z) &= z^{-1} G_0(-z) = \frac{1}{8}z^{-1}(-z^2 - 2z + 6 - 2z^{-1} - z^{-2}) \end{aligned} \quad (4.34)$$

This is often known as the **LeGall 3,5-tap filter set**, since it was first published in the context of 2-band filter banks by Didier LeGall in 1988.

The wavelets of the LeGall 3,5-tap filters,  $H_0$  and  $H_1$  above, and their frequency responses are shown in fig 4.6. The scaling function (bottom left) **converges to a pure triangular pulse** and the wavelets are the superposition of two triangular pulses.

The triangular scaling function produces **linear interpolation** between consecutive low-band coefficients and also causes the wavelets to be linear interpolations of the coefficients of the  $H_1$  filter,  $-1, -2, 6, -2, -1$  (scaled appropriately).

These wavelets have quite desirable properties for image compression (note the absence of waveform discontinuities and the much lower sidelobes of the frequency responses), and they represent probably the simplest useful wavelet design. Unfortunately there is one drawback — **the inverse wavelets are not very good**. These are formed from the LeGall 5,3-tap filter pair,  $G_0$  and  $G_1$  above, whose wavelets and frequency responses are shown in fig 4.7.

The main problem is that the wavelets do not converge after many levels to a smooth function and hence the frequency responses have large unwanted sidelobes. The jaggedness of the scaling function and wavelets causes highly visible coding artefacts if these filters are used for reconstruction of a compressed image.

However the allocation of the factors of  $P_t(Z)$  to  $H_0$  and  $G_0$  is a free design choice, so we may **swap the factors** (and hence swap  $G$  and  $H$ ) in order that the smoother 3,5-tap filters become  $G_0, G_1$  and are used for reconstruction. We shall show later that this leads to a good low-complexity solution for image compression and that the jaggedness of the analysis filters is not critical.

Unbalance between analysis and reconstruction filters / wavelets is nevertheless often regarded as being undesirable, particularly as it prevents the filtering process from being represented as an orthonormal transformation of the input signal (since an orthonormally transformed signal may be reconstructed simply by transposing the transform matrix). An unbalanced PR filter system is often termed a **bi-orthogonal** transformation.

We now consider ways to reduce this unbalance.

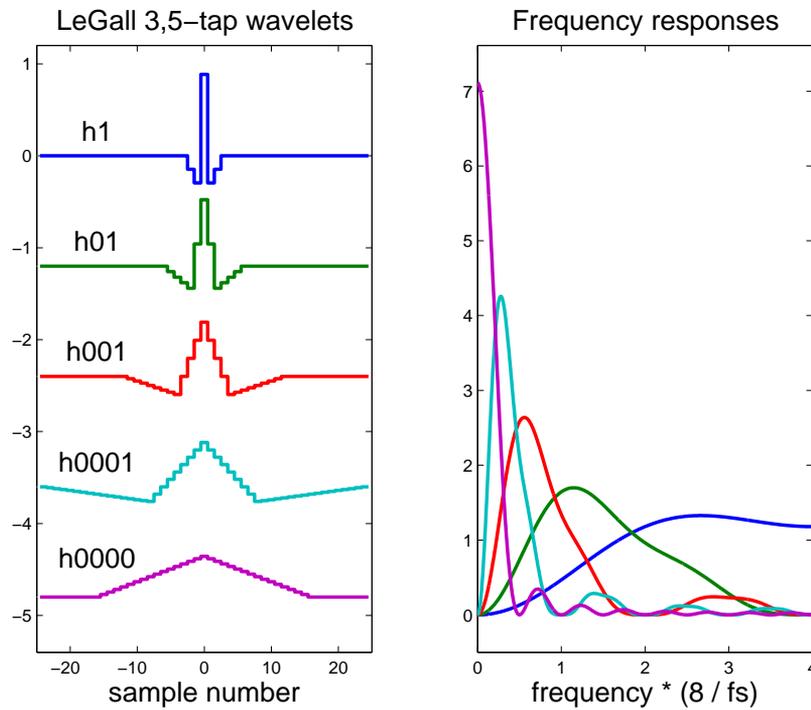


Fig 4.6: Impulse responses and frequency responses of the 4-level tree of LeGall 3,5-tap filters.

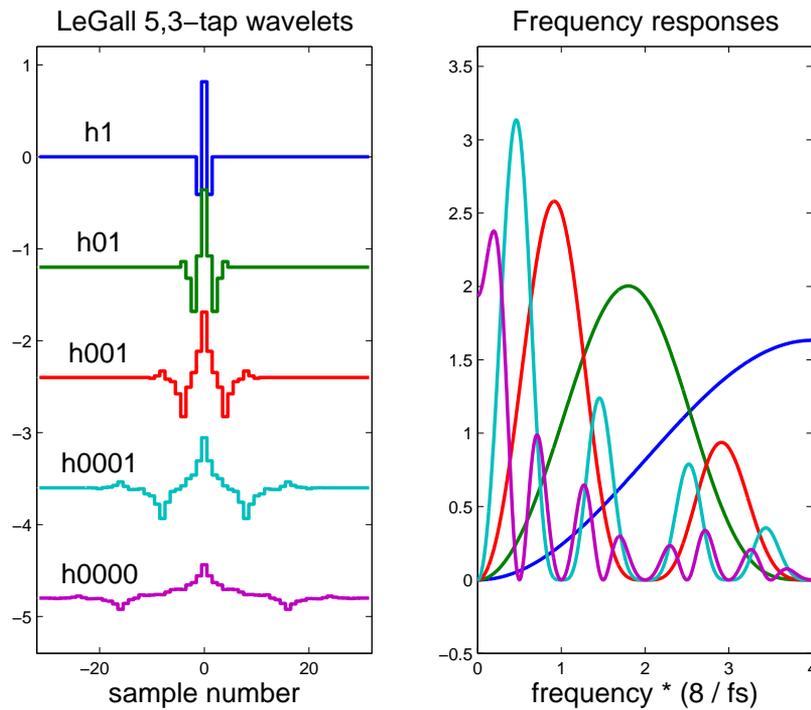


Fig 4.7: Impulse responses and frequency responses of the 4-level tree of LeGall 5,3-tap filters.

### Filters with balanced $H$ and $G$ frequency responses (but non-linear phase responses) — Daubechies wavelets:

In the above analysis, we used the factorisation of  $P_t(Z)$  to give us  $H_0$  and  $G_0$ . This always gives unbalanced factors if terms of  $P_t$  in even powers of  $Z$  are zero.

However each of these factors in  $Z$  may itself be factorised into a pair of factors in  $z$ , since:

$$\begin{aligned} (\alpha z + 1)(1 + \alpha z^{-1}) &= \alpha z + (1 + \alpha^2) + \alpha z^{-1} = (1 + \alpha^2) + 2\alpha Z \\ &= (1 + \alpha^2)(1 + \beta Z) \quad \text{where } \beta = \frac{2\alpha}{1 + \alpha^2} \end{aligned} \quad (4.35)$$

For each factor of  $P_t(Z)$ , we may allocate one of its  $z$  subfactors to  $H_0(z)$  and the other to  $G_0(z)$ . Where roots of  $P_t(Z)$  are complex, the subfactors must be allocated in conjugate pairs so that  $H_0$  and  $G_0$  remain purely real.

Since the subfactors occur in reciprocal pairs (roots at  $z = \alpha$  and  $\alpha^{-1}$ ), we find that

$$G_0(z) = H_0(z^{-1}) \quad (4.36)$$

which means that the impulse response of  $G_0$  is the time-reverse of  $H_0$ .

Therefore the frequency responses are related by  $G_0(e^{j\omega T_s}) = H_0(e^{-j\omega T_s})$ .

Hence the magnitudes of the frequency responses are the same, and their phases are opposite. It may be shown that this is sufficient to obtain **orthogonal** wavelets, but unfortunately the separate filters are no longer zero (or linear) phase. (Linear phase is zero phase with an arbitrary delay  $z^{-k}$ .)

Daubechies wavelets may be generated in this way, with the added constraint that the maximum number of zeros of  $P_t(Z)$  are placed at  $Z = -1$  (producing pairs of zeros of  $P(z)$  at  $z = -1$ ), consistent with terms in even powers of  $Z$  being zero.

If  $P_t(Z)$  is of order  $2K - 1$ , then it may have  $K$  zeros at  $Z = -1$  such that

$$P_t(Z) = (1 + Z)^K R_t(Z) \quad (4.37)$$

where  $R_t(Z)$  is of order  $K - 1$  and its  $K - 1$  roots may be chosen such that terms of  $P_t(Z)$  in the  $K - 1$  even powers of  $Z$  are zero.

Equation (4.32) is the  $K = 2$  solution to (4.37).

$\therefore R_t(Z) = 1 - \frac{1}{2}Z$  so  $\beta = -\frac{1}{2}$  and, from (4.35), the factors of  $R(z)$  are

$$R(z) = \frac{(\alpha z + 1)(1 + \alpha z^{-1})}{1 + \alpha^2} \quad \text{where } \alpha = \sqrt{3} - 2$$

Also

$$(1 + Z)^2 = \frac{1}{2}(z + 1)^2 \frac{1}{2}(1 + z^{-1})^2$$

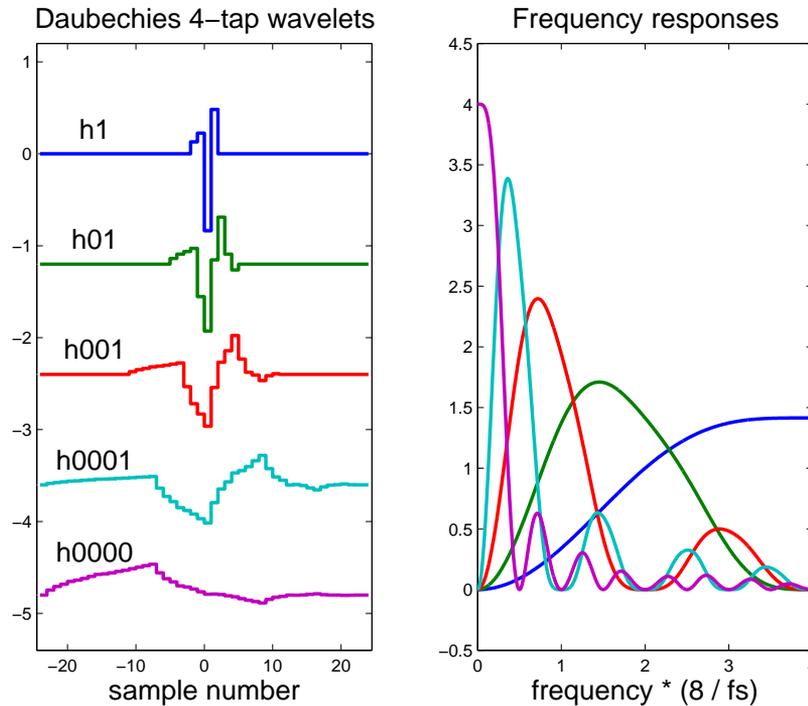


Fig 4.8: Impulse responses and frequency responses of the 4-level tree of Daubechies 4-tap filters.

Hence

$$H_0(z) = \frac{1}{2\sqrt{1+\alpha^2}}(1+z^{-1})^2(1+\alpha z^{-1}) = 0.4830 + 0.8365z^{-1} + 0.2241z^{-2} - 0.1294z^{-3}$$

and

$$H_1(z) = z^{-3}G_0(-z) = z^{-3}H_0(-z^{-1}) = 0.1294 + 0.2241z^{-1} - 0.8365z^{-2} + 0.4830z^{-3}$$

The wavelets and frequency responses for these filters are shown in fig 4.8. It is clear that the wavelets and scaling function are no longer linear phase and are less smooth than those for the LeGall 3,5-tap filters. The frequency responses also show worse sidelobes. The  $G_0, G_1$  filters give the time reverse of these wavelets and identical frequency responses.

Higher order Daubechies filters achieve smooth wavelets but they still suffer from non-linear phase. This tends to result in more visible coding artefacts than linear phase filters, which distribute any artefacts equally on either side of sharp edges in the image.

Linear phase filters also allow an elegant technique, known as symmetric extension, to be used at the outer edges of images, where wavelet filters would otherwise require the size of the transformed image to be increased to allow for convolution with the filters. Symmetric extension assumes that the image is reflected by mirrors at each edge, so that

an infinitely tessellated plane of reflected images is generated. Reflections avoid unwanted edge discontinuities. If the filters are linear phase, then the DWT coefficients also form reflections and no increase in size of the transformed image is necessary to accommodate convolution effects.

### Filters with linear phase and nearly balanced frequency responses:

To ensure that the filters  $H_0, H_1$  and  $G_0, G_1$  are linear phase, the factors in  $Z$  must be allocated to  $H_0$  or  $G_0$  as a whole and not be split, as was done for the Daubechies filters. In this way the symmetry between  $z$  and  $z^{-1}$  is preserved in all filters.

Perfect balance of frequency responses between  $H_0$  and  $G_0$  is then not possible, if PR is preserved, but we have found a factorisation of  $P_t(Z)$  which achieves near balance of the responses.

This is:

$$P_t(Z) = (1 + Z)(1 + aZ + bZ^2) (1 + Z)(1 + cZ) \quad (4.38)$$

This is a 5<sup>th</sup> order polynomial, and if the terms in  $Z^2$  and  $Z^4$  are to be zero, there are two constraints on the 3 unknowns  $[a, b, c]$  so that one of them (say  $c$ ) may be regarded as a free parameter. These constraints require that:

$$a = -\frac{(1 + 2c)^2}{2(1 + c)^2} \quad \text{and} \quad b = \frac{c(1 + 2c)}{2(1 + c)^2} \quad (4.39)$$

$c$  may then be adjusted to give maximum similarity between the left and right pairs of factors in (4.38) as  $Z$  varies from 1 to  $-1$  ( $\omega T_s$  varies from 0 to  $\pi$ ).

It turns out that  $c = -\frac{2}{7}$  gives good similarity and when substituted into (4.39) and (4.38) gives:

$$P_t(Z) = \frac{1}{50}(50 + 41Z - 15Z^2 - 6Z^3) \frac{1}{7}(7 + 5Z - 2Z^2) \quad (4.40)$$

We get  $G_0(z)$  and  $H_0(z)$  by substituting  $Z = \frac{1}{2}(z + z^{-1})$  into these two polynomial factors. This results in 5,7-tap filters whose wavelets and frequency responses are shown in fig 4.9.

The near balance of the responses may be judged from fig 4.10 which shows the alternative 7,5-tap versions (i.e. with  $H$  and  $G$  swapped). The balance is sufficiently good that it is quite hard to see the minor differences between figs 4.9 and 4.10.

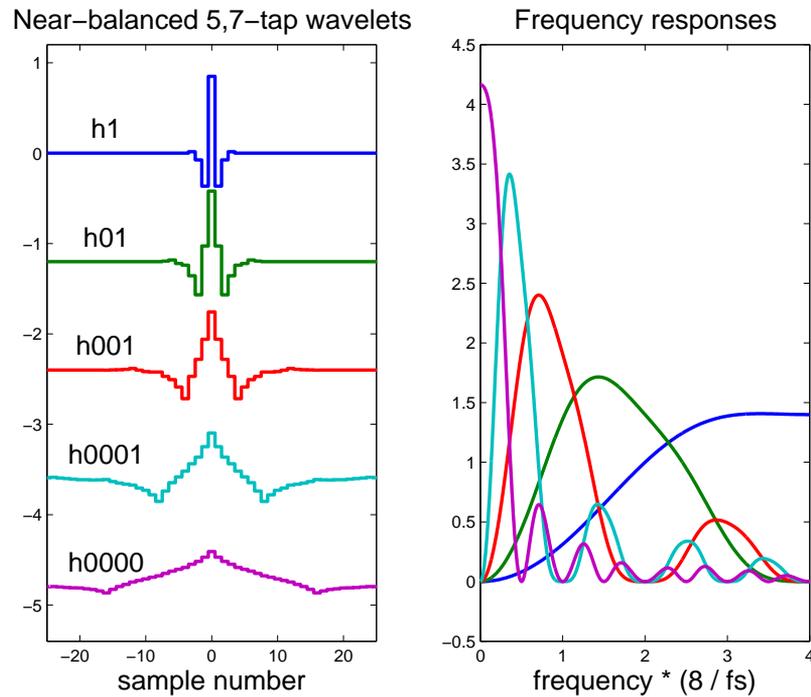


Fig 4.9: Impulse responses and frequency responses of the 4-level tree of near-balanced 5,7-tap filters.

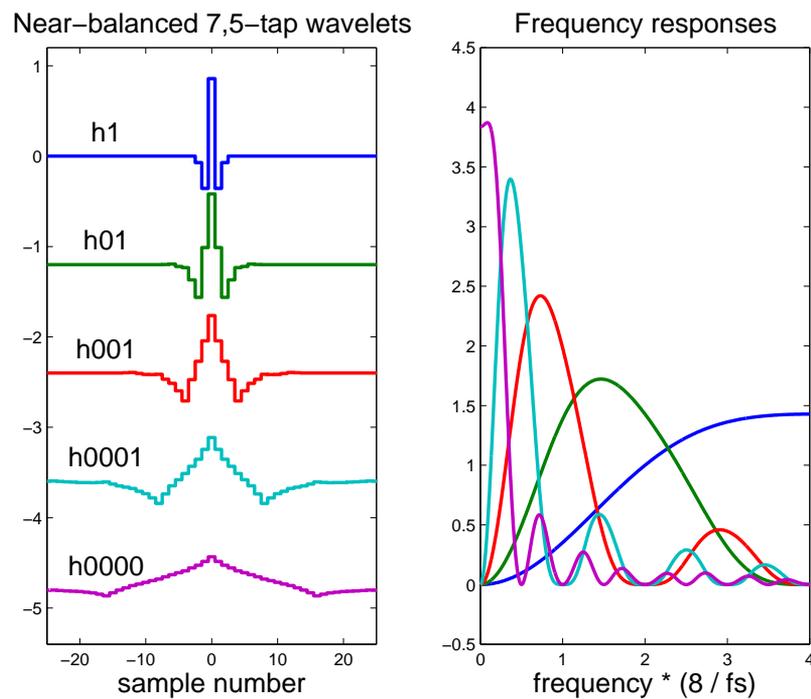


Fig 4.10: Impulse responses and frequency responses of the 4-level tree of near-balanced 7,5-tap filters.

### Smoother Wavelets

In all of the above designs we have used the substitution  $Z = \frac{1}{2}(z + z^{-1})$ . However other substitutions may be used to create improved wavelets. To preserve PR, the substitution should contain only odd powers of  $z$  (so that odd powers of  $Z$  should produce only odd powers of  $z$ ), and to obtain zero phase for  $P(z)$ , the coefficients of the substitution should be symmetric about  $z^0$ .

A substitution, which can give much greater flatness near  $z = \pm 1$  while still satisfying  $Z = \pm 1$  when  $z = \pm 1$ , is:

$$Z = pz^3 + (\frac{1}{2} - p)(z + z^{-1}) + pz^{-3} \quad (4.41)$$

$Z$  then becomes the following function of frequency when  $z = e^{j\omega T_s}$ :

$$Z = (1 - 2p) \cos(\omega T_s) + 2p \cos(3\omega T_s) \quad (4.42)$$

A high degree of flatness (with some ripple) is achieved near  $\omega T_s = 0$  and  $\pi$ , if  $p = -\frac{3}{32}$ . This results in 4 zeros near  $z = -1$  for each  $(Z + 1)$  factor, compared with only 2 zeros per factor when  $Z = \frac{1}{2}(z + z^{-1})$  is used.

The  $2^{nd}$ -order factor in  $P_t(Z)$  now produces terms from  $z^6$  to  $z^{-6}$  and the  $3^{rd}$ -order factor produces terms from  $z^9$  to  $z^{-9}$ . Hence the filters become 13 and 19 tap filters, although 2 taps of each are zero and the outer two taps of the 19-tap filter are very small ( $\sim 10^{-4}$ ).

Fig 4.11a shows the wavelets and frequency responses of the 13,19-tap filters, obtained by substituting (4.41) into (4.40). Note the smoother wavelets and scaling function and the much lower sidelobes in the frequency responses of these higher-order filters.

Fig 4.11b demonstrates that the near-balanced properties of (4.40) are preserved in the resulting higher-order filters.

There are many other types of wavelets with varying features and complexities, including a popular 7,9-tap bi-orthogonal linear-phase filter pair based on equation (4.37) with  $K = 4$ , but we have found the examples given to be near optimum for image compression.

The 2-level 4-point DCT system employed in JPEG XR can also be regarded as a 4-band wavelet system, and the analogy is particularly appropriate when the overlapped prefilters are enabled in that system. Fig 4.12a shows the resulting reconstruction 'wavelets' when the prefilters are disabled, and fig 4.12b shows them when the prefilters are enabled. Note the much smoother waveforms and lower sidelobes in the frequency domain in fig 4.12b. Hence it turns out that the JPXR system has compression performance that closely matches that of the smooth wavelets of fig 4.11.

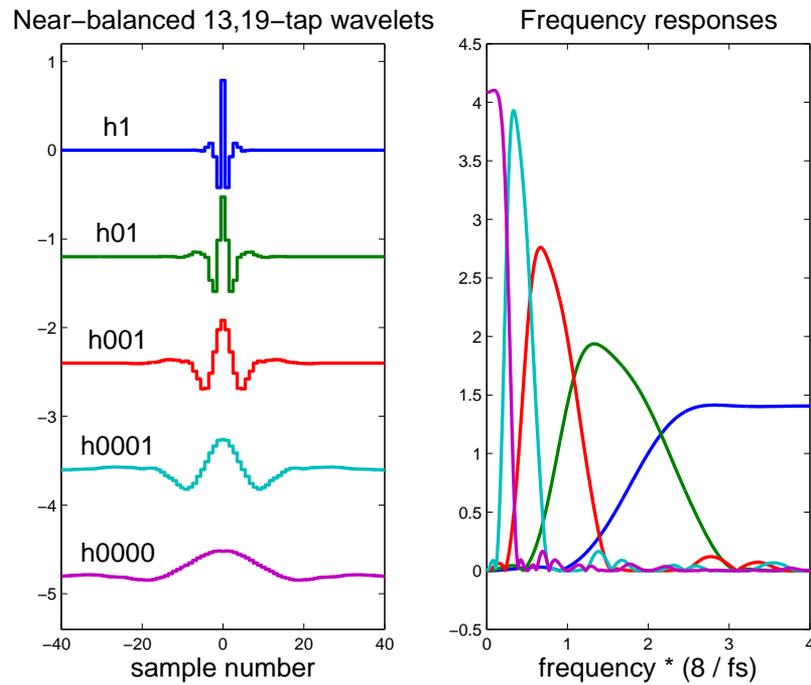


Fig 4.11a: Impulse responses and frequency responses of the 4-level tree of near-balanced 13,19-tap filters.

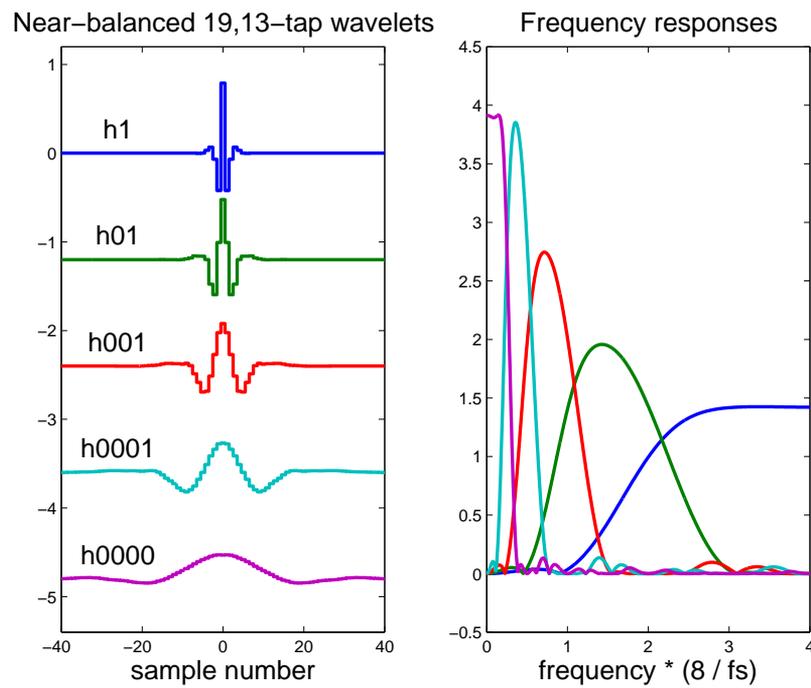


Fig 4.11b: Impulse responses and frequency responses of the 4-level tree of near-balanced 19,13-tap filters.

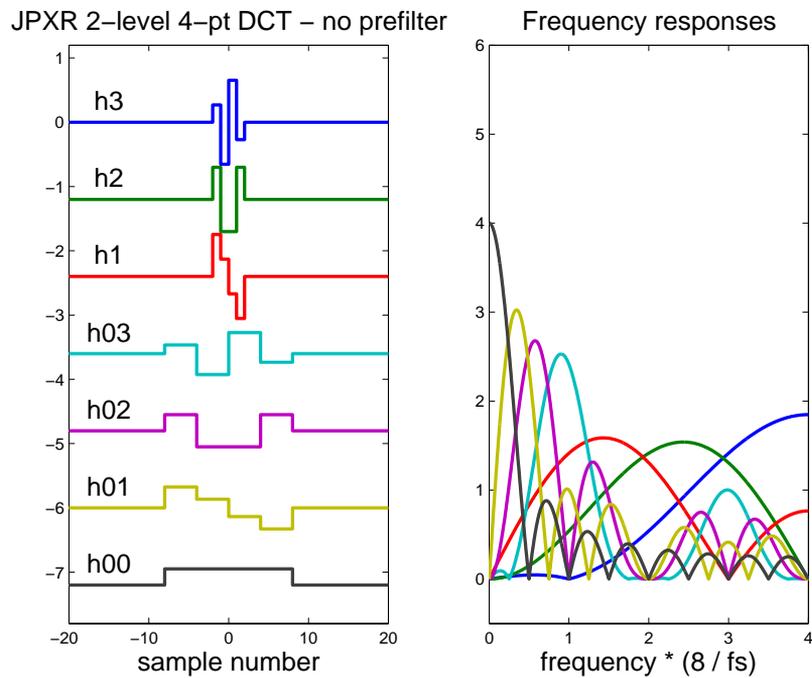


Fig 4.12a: Impulse responses and frequency responses of the JPXR 2-level tree of 4-pt DCT basis functions without prefiltering.

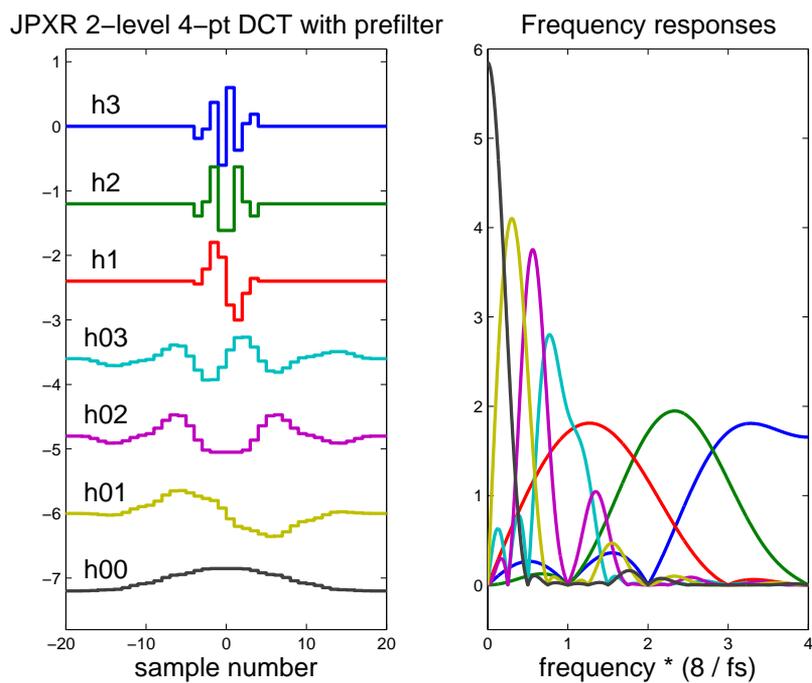


Fig 4.12b: Impulse responses and frequency responses of the JPXR 2-level tree of 4-pt DCT basis functions with prefiltering.

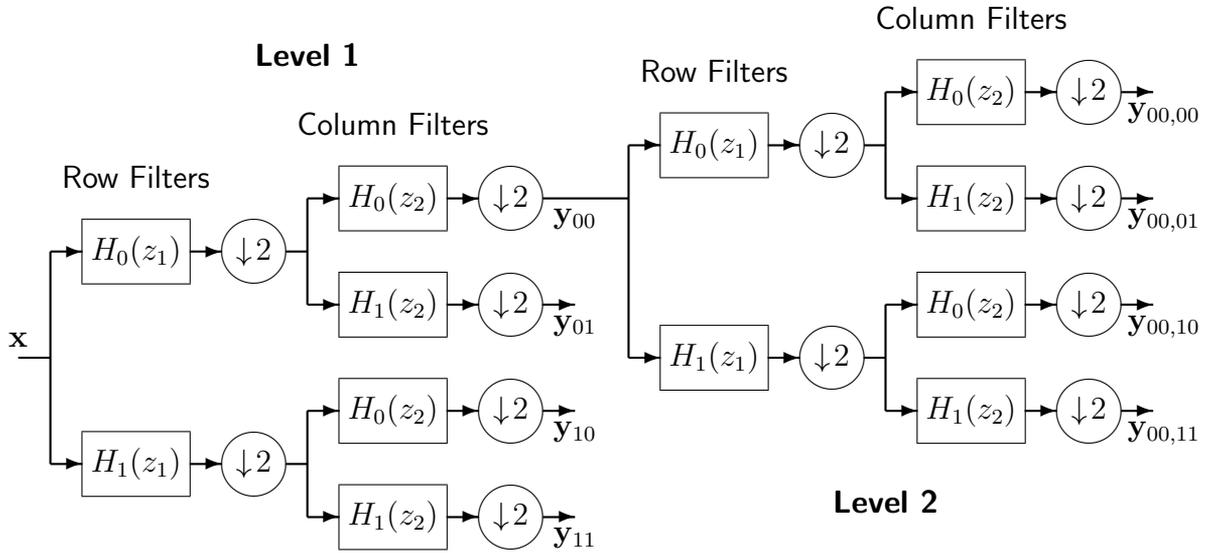


Fig 4.13: Two levels of a 2-D filter tree, formed from 1-D lowpass ( $H_0$ ) and highpass ( $H_1$ ) filters.

## 4.6 The 2-D DWT

We have already seen in section 2.1 how the 1-D Haar transform (or wavelet) could be extended to 2-D by filtering the rows and columns of an image separably.

All 1-D 2-band wavelet filter banks can be extended in a similar way. Fig 4.13 shows two levels of a 2-D filter tree. The input image at each level is split into 4 bands (Lo-Lo =  $\mathbf{y}_{00}$ , Lo-Hi =  $\mathbf{y}_{01}$ , Hi-Lo =  $\mathbf{y}_{10}$ , and Hi-Hi =  $\mathbf{y}_{11}$ ) using the lowpass and highpass wavelet filters on the rows and columns in turn. The Lo-Lo band subimage  $\mathbf{y}_{00}$  is then used as the input image to the next level. Typically 4 levels are used, as for the Haar transform.

Filtering of the rows of an image by  $H_a(z_1)$  and of the columns by  $H_b(z_2)$ , where  $a, b = 0$  or 1, is equivalent to filtering by the 2-D filter:

$$H_{ab}(z_1, z_2) = H_a(z_1) H_b(z_2) \quad (4.43)$$

In the spatial domain, this is equivalent to convolving the image matrix with the 2-D impulse response matrix

$$\mathbf{h}_{ab} = \mathbf{h}_a \mathbf{h}_b^T \quad (4.44)$$

where  $\mathbf{h}_a$  and  $\mathbf{h}_b$  are column vectors of the 1-D filter impulse responses. However note that performing the filtering separably (i.e. as separate 1-D convolutions along the rows and columns) is much more computationally efficient than the equivalent 2-D convolution.

To obtain the impulse responses of the four 2-D filters at each level of the 2-D DWT, we form  $\mathbf{h}_{ab}$  from  $\mathbf{h}_0$  and  $\mathbf{h}_1$  using equation (4.44) with  $ab = 00, 01, 10$  and  $11$  in turn.

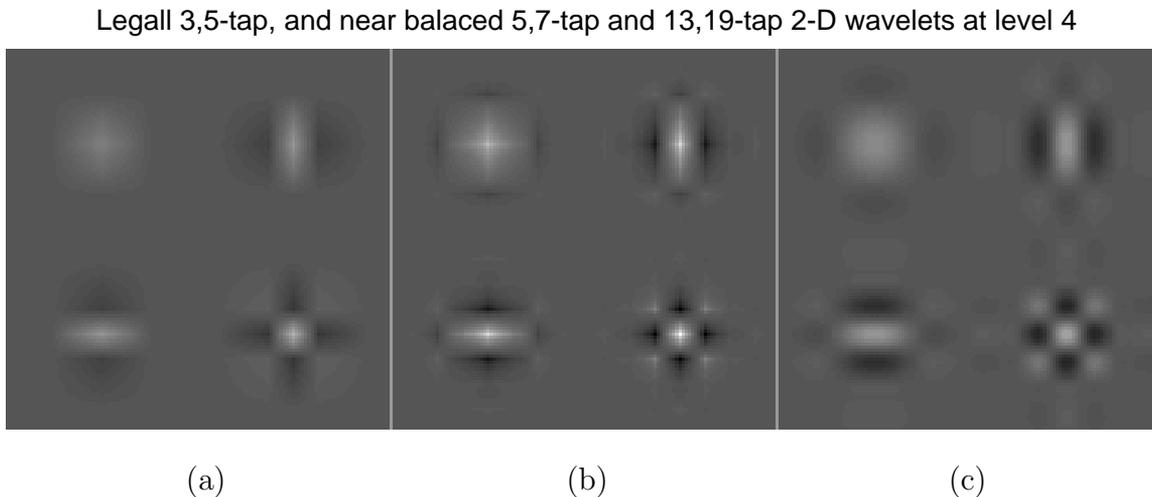


Fig 4.14: 2-D impulse responses of the level-4 wavelets and scaling functions derived from the LeGall 3,5-tap filters (a), and the near-balanced 5,7-tap (b) and 13,19-tap (c) filters.

Fig 4.14 shows the impulse responses at level 4 as images for three 2-D wavelet filter sets, formed from the following 1-D wavelet filter sets:

1. The LeGall 3,5-tap filters:  $H_0$  and  $H_1$  from equations (4.33) and (4.34).
2. The near-balanced 5,7-tap filters: substituting  $Z = \frac{1}{2}(z + z^{-1})$  into equation (4.40).
3. The near-balanced 13,19-tap filters: substituting equation (4.41) into equation (4.40).

Note the sharp points in fig 4.14b, produced by the sharp peaks in the 1-D wavelets of fig 4.9. These result in noticeable artefacts in reconstructed images when these wavelets are used. The smoother wavelets of fig 4.14c are much better in this respect.

The 2-D frequency responses of the level 1 filters, derived from the LeGall 3,5-tap filters, are shown in figs 4.15 (in mesh form) and 4.16 (in contour form). These are obtained by substituting  $z_1 = e^{j\omega_1}$  and  $z_2 = e^{j\omega_2}$  into equation (4.43). This equation demonstrates that the 2-D frequency response is just the product of the responses of the relevant 1-D filters.

Figs 4.17 and 4.18 are the equivalent plots for the 2-D filters derived from the near-balanced 13,19-tap filters. We see the much sharper cut-offs and better defined pass and stop bands of these filters. The high-band filters no longer exhibit gain peaks, which are rather undesirable features of the LeGall 5-tap filters.

Frequency responses of 2-D LeGall 3,5-tap wavelets

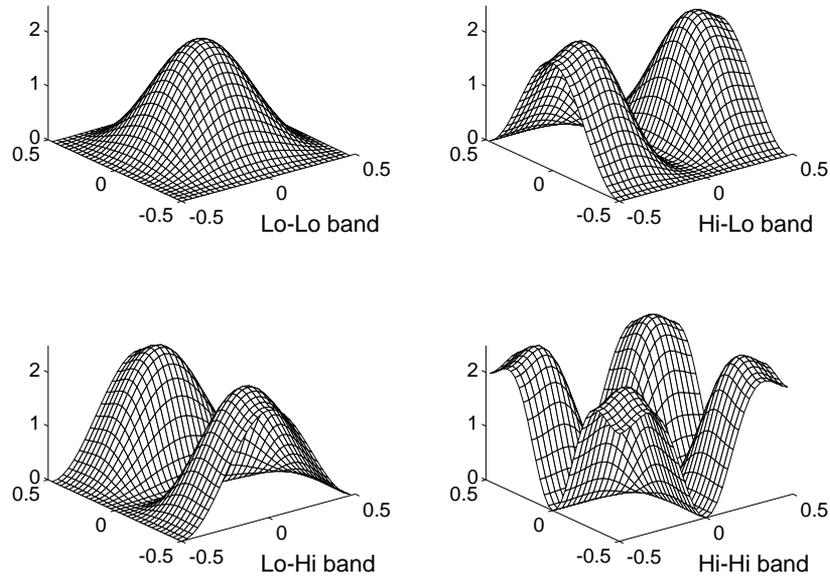


Fig 4.15: Mesh frequency response plots of the 2-D level 1 filters, derived from the LeGall 3,5-tap filters.

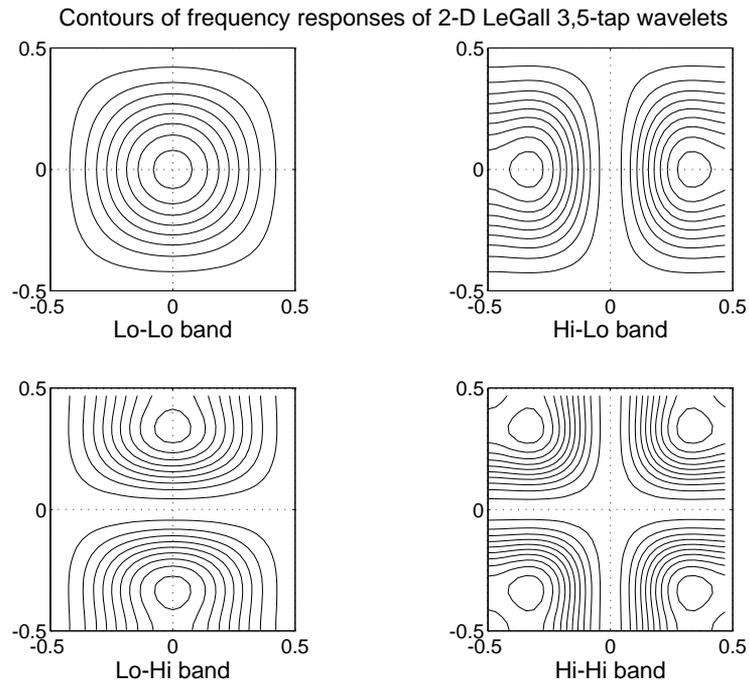


Fig 4.16: Contour plots of the frequency responses of fig 4.15.

Frequency responses of 2-D near-balanced 13,19-tap wavelets

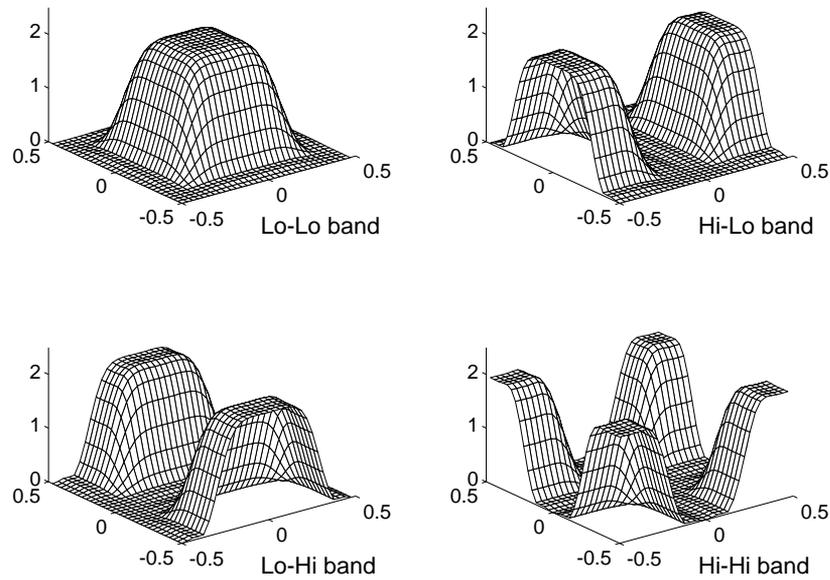


Fig 4.17: Mesh frequency response plots of the 2-D level 1 filters, derived from the near-balanced 13,19-tap filters.

Contours of frequency responses of 2-D near-balanced 13,19-tap wavelets

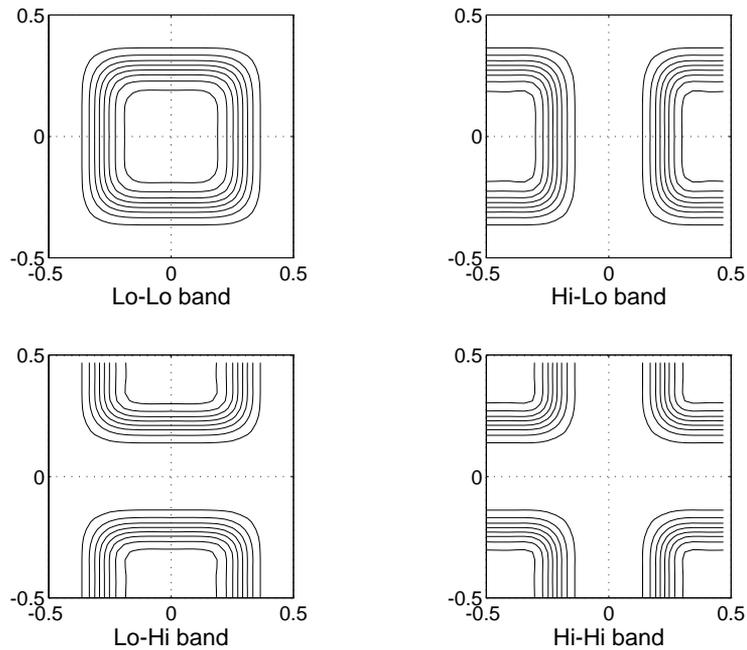


Fig 4.18: Contour plots of the frequency responses of fig 4.17.

## 4.7 Compression properties of wavelets

We now look at how well the various wavelet filters perform in practice. We have used them in place of the Haar transform of section 2, and have measured the entropies and reconstructed the images from quantised coefficients.

In order to allow a fair comparison with the JPEG DCT results, we have modified the DWT quantising strategy to take advantage of the reduced visibility of the higher frequency wavelets. This approximately matches the effects achieved by the JPEG  $\mathbf{Q}_{lum}$  matrix of equation (3.7). To achieve a high degree of compression we have used the following allocation of quantiser step sizes to the 4-level DWT bands:

$$\begin{aligned}
 \text{All bands at levels 3 and 4:} & \quad Q_{step} = 50 \\
 \text{Hi-Lo and Lo-Hi bands at level 2:} & \quad Q_{step} = 50 \\
 \text{Hi-Hi band at level 2:} & \quad Q_{step} = 100 \\
 \text{Hi-Lo and Lo-Hi bands at level 1:} & \quad Q_{step} = 100 \\
 \text{Hi-Hi band at level 1:} & \quad Q_{step} = 200
 \end{aligned} \tag{4.45}$$

A similar compressed bit rate is produced by the  $8 \times 8$  DCT when  $\mathbf{Q}_{step} = 5\mathbf{Q}_{lum}$ .

For reference, figs 4.19a and 4.19b compare the DCT and Haar transforms using these two quantisers. The rms errors between the reconstructed images and the original are virtually the same at 10.49 and 10.61 respectively, but the DCT entropy of 0.2910 bit/pel is significantly lower than the Haar entropy of 0.3820 bit/pel. Both images display significant **blocking artefacts** at this compression level. Figs 4.19c and 4.19d show how the 2-level DCT system of JPXR compares with these under similar quantiser conditions. When the overlapping prefilters are not used (fig 4.19c), we see that the blocking artefacts are similar to the images above it, but when the prefilters are enabled (fig 4.19d) the blocking is much less noticeable, the rms error is reduced from 9.28 to 8.93, and the entropy is reduced substantially, from 0.4084 to 0.3253 bit/pel.

Fig 4.20 shows the reconstructed images for the following four DWTs using the quantiser of equations (4.45):

1. The LeGall 3,5-tap filters:  $H_0, H_1$  and  $G_0, G_1$  from equations (4.33) and (4.34).
2. The inverse-LeGall 5,3-tap filters: equations (4.33) and (4.34) with  $H_0, H_1$  and  $G_0, G_1$  swapped.
3. The near-balanced 5,7-tap filters: substituting  $Z = \frac{1}{2}(z + z^{-1})$  into equation (4.40).
4. The near-balanced 13,19-tap filters: substituting equation (4.41) into equation (4.40).



(a)



(b)



(c)



(d)

Fig 4.19: Reconstructions after coding using: (a) the Haar transform with  $\mathbf{Q}_{step}$  from equation (4.45); (b) the  $8 \times 8$  DCT with  $\mathbf{Q}_{step} = 5\mathbf{Q}_{lum}$ ; (c) the JPXR 2-level  $4 \times 4$  DCTs with  $\mathbf{Q}_{step} \simeq 3\mathbf{Q}_{lum}$ , but no prefilters; (d) JPXR with prefilters and the same  $\mathbf{Q}_{step}$ .

We see that the LeGall 3,5-tap filters (fig 4.20a) produce a poor image, whereas the other three images are all significantly better. The poor image is caused by the roughness of the LeGall 5,3-tap filters (shown in fig 4.7) which are used for reconstructing the image when the 3,5-tap filters are used for analysing the image. When these filters are swapped, so that the reconstruction filters are the 3,5-tap ones of fig 4.6, the quality is greatly improved (fig 4.20b).



Fig 4.20: Reconstructions after coding using  $\mathbf{Q}_{step}$  from equations (4.45) with (a) the LeGall 3,5-tap filters, (b) the inverse-Legall 5,3-tap filters, (c) the near-balanced 5,7-tap filters, and (d) the near-balanced 13,19-tap filters.

The near-balanced 5,7-tap filters (fig 4.20c) produce a relatively good image but there are still a few bright or dark point-artefacts produced by the sharp peaks in the wavelets (shown in fig 4.9). The smoother 13,19-tap wavelets (fig 4.10) eliminate these, but their longer impulse responses tend to cause the image to have a slightly **blotchy** or **mottled** appearance.

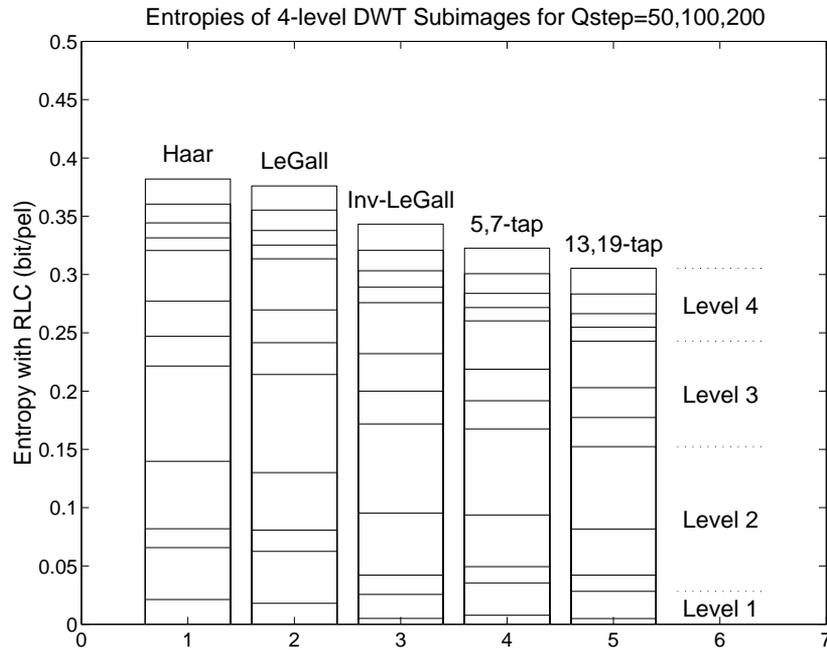


Fig 4.21: Entropies of 4-level DWT subimages using  $Q_{step}$  defined by equations (4.45), for five different wavelet filter pairs.

Fig 4.21 shows the entropies (with RLC) of the separate subimages of the 4-level DWT for the Haar filter set and the four filter sets of fig 4.20.  $Q_{step}$  is defined by equations (4.45) and it is particularly noticeable how the higher step sizes at levels 1 and 2 substantially reduce the entropy required to code these levels (compare with fig 2.12). In fact the Hi-Hi band at level 1 is not coded at all! The reduction of entropy with increasing filter smoothness is also apparent.

**We see that we have now been able to reduce the bit rate to around 0.3 bit/pel.**

However measurement of entropy is not the whole story, as it is the tradeoff of entropy vs quantising error which is important. Fig 4.22 attempts to show this trade-off by plotting rms quantising error (obtained by subtracting the reconstructed image from the original) versus the entropy for the  $8 \times 8$  DCT, the JPXR 2-level 4 DCT with prefilters, and the five DWTs. To show the slope of the curves, the measurements are repeated with a 20% lower quantiser step-size, giving lower rms errors and higher entropies. The pair of points for each configuration are joined by a line which indicates the slope of the rate-distortion curve (the 20% lower stepsize point is to the right in each pair).

Measurements at many more step sizes can be taken in order to give more complete rate-distortion curves if required.

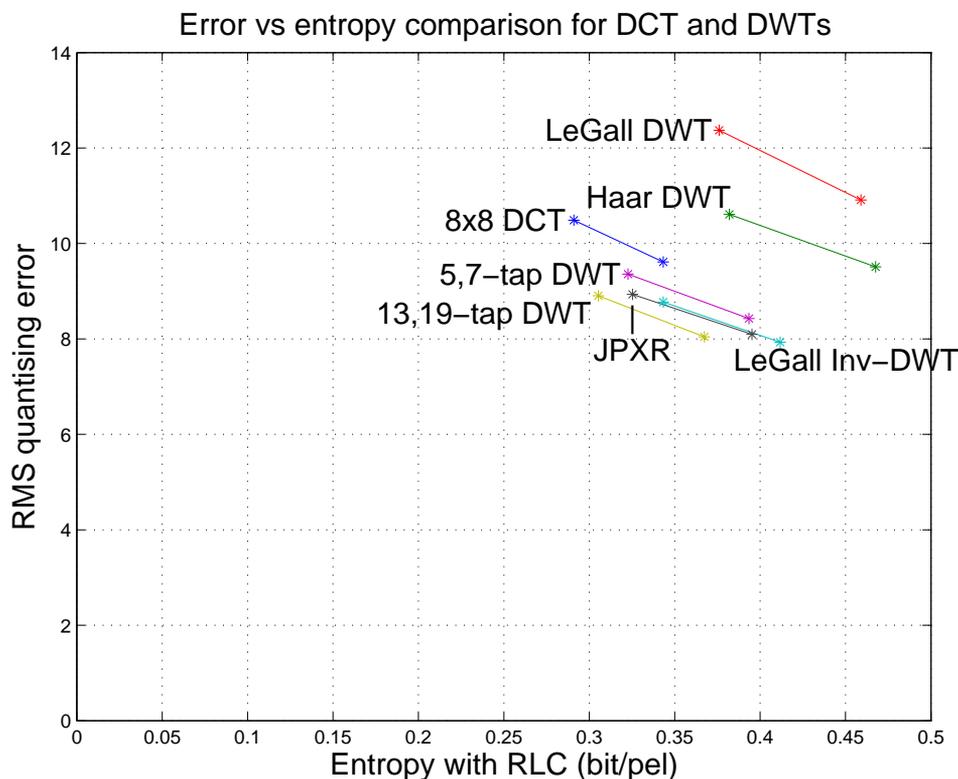


Fig 4.22: RMS error vs. entropy for the  $8 \times 8$  DCT, JPXR transforms with prefilters, and five wavelet filter pairs. For the DCT,  $Q_{step} = 5Q_{lum}$  and  $4Q_{lum}$ ; for the DWTs,  $Q_{step}$  is 1.0 and 0.8 of the values in (4.45). For JPXR, the quantizer levels approximately match those used for the DWTs.

The good performance of the 13,19-tap filters is clear, but the inverse-LeGall filters do surprisingly well — showing that the poor smoothness of the analysis filters does not seem to matter. Correct ways to characterise unbalanced filter sets to account properly for this phenomenon are still the subject of some debate.

**What is clear is that when filters are unbalanced between analysis and reconstruction, the ones which give smoother wavelets *must* be used for reconstruction.**

This is further supported by the good performance of the JPXR system with filters, where the filters produce quite smooth basis functions, as seen in fig 4.12b. Without these filters, the JPXR system would perform somewhere between the  $8 \times 8$  DCT and the Haar DWT.

Finally, in these tests, we find that the assessments of subjective image quality approximately match the assessments based on rms errors. However this is not always true and one must be careful to backup any conclusions from rms error measurements with at least some subjective tests.

## 5 Video Compression and Motion Processing

Video compression is concerned with coding image sequences at low bit rates. In an image sequence, there are typically high correlations between consecutive frames of the sequence, in addition to the spatial correlations which exist naturally within each frame.

Video coders aim to take maximum advantage of **interframe** temporal correlations (between frames) as well as **intraframe** spatial correlations (within frames).

### 5.1 Motion-Compensated Predictive Coding

Motion-compensated predictive coding (MCPC) is the technique that has been found to be most successful for exploiting interframe correlations.

Fig 5.1 shows the basic block diagram of a MCPC video encoder.

The transform, quantise, and entropy encode functions are basically the same as those employed for still image coding. The first frame in a sequence is coded in the normal way for a still image by switching the prediction frame to zero.

For subsequent frames, the input to the transform stage is the difference between the **input** frame and the **prediction** frame, based on the previous decoded frame. This difference frame is usually known as the **prediction error** frame.

The purpose of employing prediction is to reduce the energy of the prediction error frames so that they have lower entropy after transformation and can therefore be coded with a lower bit rate.

If there is motion in the sequence, the prediction error energy may be significantly reduced by **motion compensation**. This allows regions in the prediction frame to be generated from **shifted** regions from the previous decoded frame. Each shift is defined by a **motion vector** which is transmitted to the decoder in addition to the coded transform coefficients. The motion vectors are usually entropy coded to minimise the extra bit rate needed to do this.

The multiplexer combines the various types of coded information into a single serial bit stream, and the buffer smooths out the fluctuations in bit rate caused by varying motion within the sequence and by scene changes. The controller adjusts coding parameters (such as the quantiser step size) in order to maintain the buffer at approximately half-full, and hence it keeps the mean bit rate of the encoder equal to that of the channel.

Decoded frames are produced in the encoder, which are identical to those generated in the decoder. The decoder comprises a buffer, de-multiplexer, and entropy decoder to invert the operations of the equivalent encoder blocks, and then the decoded frames are produced by the part of the encoder loop comprising the inverse quantiser, inverse transform, adder,

frame store, motion compensator and switch.

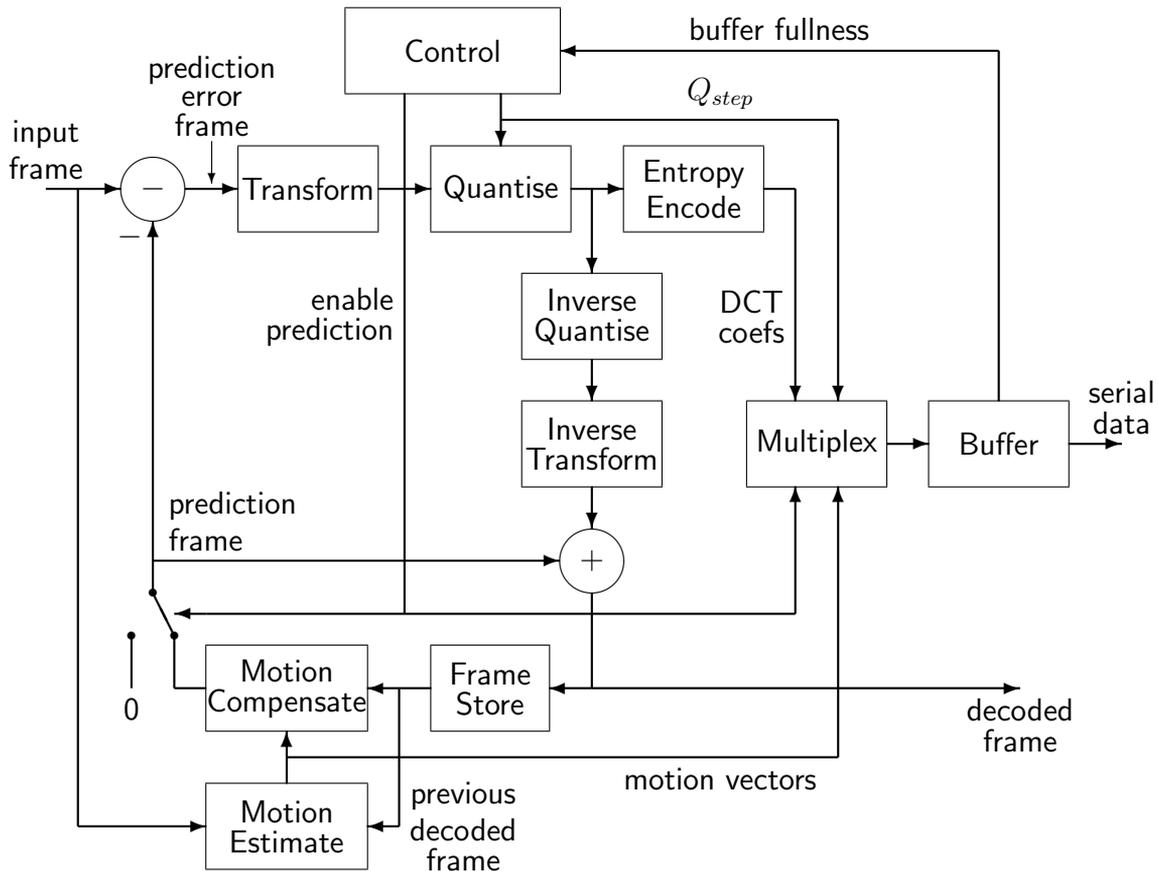


Fig 5.1: Motion compensated prediction coding (MCPC) video encoder.

H.261 is a CCITT standard for video encoding for video-phone and video conferencing applications. Video is much more important in a multi-speaker conferencing environment than in simple one-to-one conversations. H.261 employs coders of the form shown in fig 5.1 to achieve reasonable quality head-and-shoulder images at rates down to 64 kb/s (one ISDN channel). A demonstration of H.261 coding at 64 and 32 kb/s will be shown.

A development of this, H.263, is now in current use, to allow the bit rate to be reduced down to about 20 kb/s, without too much loss of quality, for modems and mobile channels. This uses some of the more advanced motion methods from MPEG (see later), and it also forms the basis for baseline MPEG-4 coders.

## 5.2 Motion Estimation

Motion estimation is the process which generates the motion vectors that determine how each motion compensated prediction frame is created from the previous frame.

**Block Matching** (BM) is the most common method of motion estimation. Typically each macroblock ( $16 \times 16$  pels) in the new frame is compared with shifted regions of the same size from the previous decoded frame, and the shift which results in the minimum error is selected as the best motion vector for that macroblock. The motion compensated prediction frame is then formed from all of the shifted regions from the previous decoded frame.

BM can be very computationally demanding if all shifts of each macroblock are analysed. For example, to analyse shifts of up to  $\pm 15$  pels in the horizontal **and** vertical directions requires  $31 \times 31 = 961$  shifts, each of which involves  $16 \times 16 = 256$  pel difference computations for a given macroblock. This is known as **exhaustive search BM**.

Significant savings can be made with **heirarchical BM**, in which an approximate motion estimate is obtained from exhaustive search using a lowpass subsampled pair of images, and then the estimate is refined by a small local search using the full resolution images. Subsampling 2:1 in each direction reduces the number of macroblock pels **and** the number of shifts both by 4:1, producing a computational saving of 16:1!

There are many other approaches to motion estimation, some using the frequency or wavelet domains, and designers have considerable scope to invent new methods since this process does not need to be specified in coding standards. The standards need only specify how the motion vectors should be interpreted by the decoder (a much simpler process). Unfortunately, we do not have time to discuss these other approaches here.

### 5.3 The MPEG Standard

As a sequel to the JPEG standards committee, the Moving Picture Experts Group (MPEG) was set up in the mid 1980s to agree standards for video sequence compression.

Their first standard was MPEG-1, designed for CD-ROM applications at 1.5Mb/s, and their more recent standard, MPEG-2, is aimed at broadcast quality TV signals at 4 to 10 Mb/s and is also suitable for high-definition TV (HDTV) at 20 Mb/s. We shall not go into the detailed differences between these standards, but simply describe some of their important features. MPEG-2 is used for digital TV and DVD in the UK and throughout the world.

MPEG coders all use the MCPC structure of fig 5.1, and employ the  $8 \times 8$  DCT as the basic transform process. So in many respects they are similar to H.261 coders, except that they operate with higher resolution frames and higher bit rates.

The main difference from H.261 is the concept of a Group of Pictures (GOP) Layer in the coding heirarchy, shown in fig 5.2. However we describe the other layers first:

- The Sequence Layer contains a complete image sequence, possibly hundreds or thousands of frames.
- The Picture Layer contains the code for a single frame, which may either be coded in absolute form or coded as the difference from a predicted frame.
- The Slice Layer contains one row of macroblocks ( $16 \times 16$  pels) from a frame. (48 macroblocks give a row 768 pels wide.)
- The Macroblock Layer contains a single macroblock — usually 4 blocks of luminance, 2 blocks of chrominance and a motion vector.
- The Block Layer contains the DCT coefficients for a single  $8 \times 8$  block of pels, coded almost as in JPEG using zig-zag scanning and run-amplitude Huffman codes.

The GOP Layer contains a small number of frames (typically 12) coded so that they can be decoded completely as a unit, without reference to frames outside of the group. There are three types of frame:

- I** — **Intra** coded frames, which are coded as single frames as in JPEG, without reference to any other frames.
- P** — **Predictive** coded frames, which are coded as the difference from a motion compensated prediction frame, generated from an earlier I or P frame in the GOP.
- B** — **Bi-directional** coded frames, which are coded as the difference from a bi-directionally interpolated frame, generated from earlier and later I or P frames in the sequence (with motion compensation).

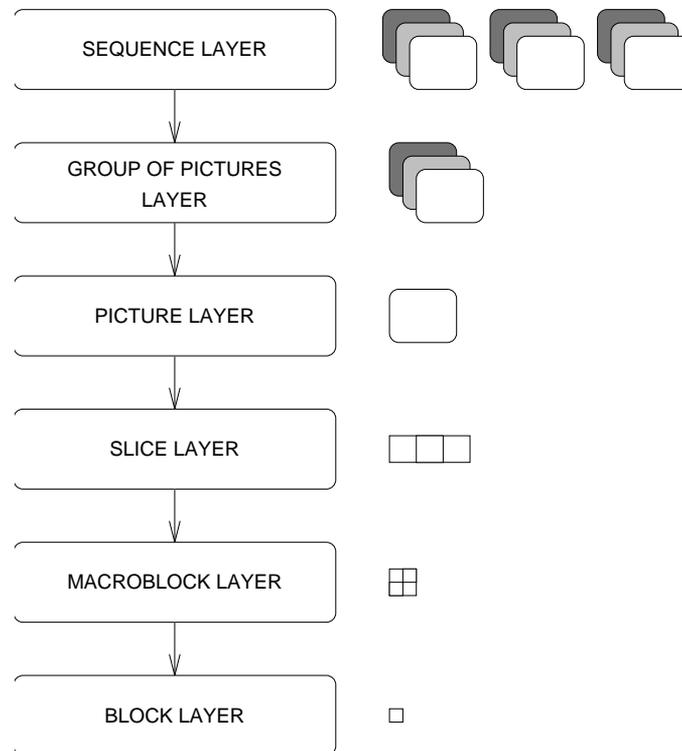


Fig 5.2: MPEG Layers.

The main purpose of the GOP is to allow editing and splicing of video material from different sources and to allow rapid forward or reverse searching through sequences. A GOP usually represents about half a second of the image sequence.

Fig 5.3 shows a typical GOP and how the coded frames depend on each other. The first frame of the GOP is always an I frame, which may be decoded without needing data from any other frame. At regular intervals through the GOP, there are P frames, which are coded relative to a prediction from the I frame or previous P frame in the GOP. Between each pair of I / P frames are one or more B frames.

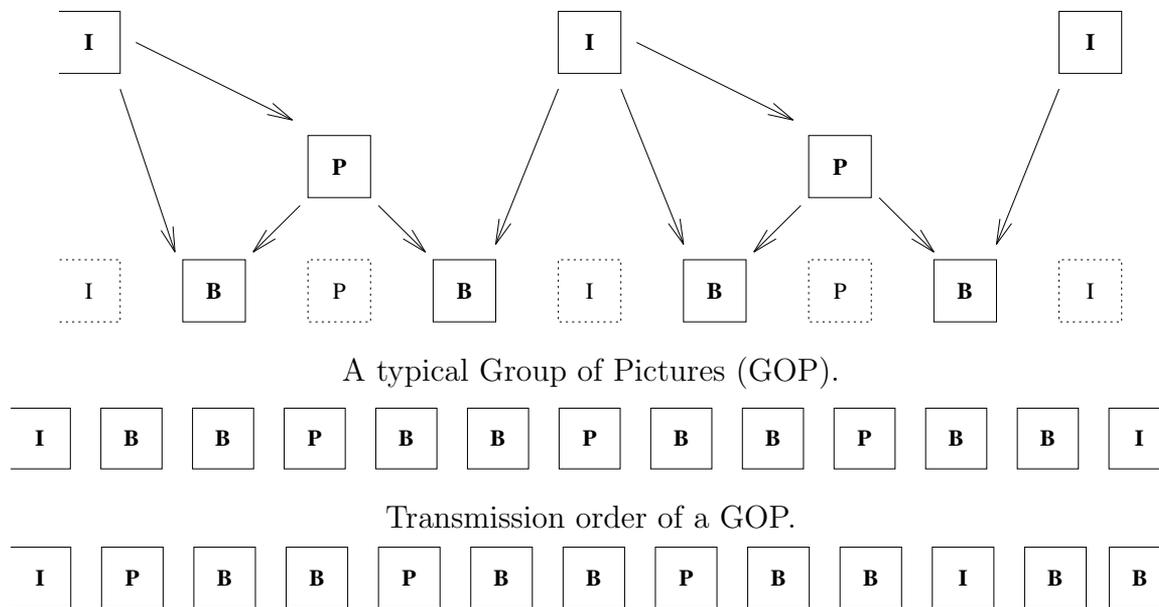


Fig 5.3: GOP Layer — Intra (I), Predicted (P) and Bi-directional (B) frames.

The I frame in each GOP requires the most bits per frame and provides the initial reference for all other frames in the GOP. Each P frame typically requires about one third of the bits of an I frame, and there may be 3 of these per GOP. Each B frame requires about half the bits of a P frame and there may be 8 of these per GOP. Hence the coded bits are split about evenly between the three frame types.

B frames require fewer bits than P frames mainly because bi-directional prediction allows uncovered background areas to be predicted from a subsequent frame. The motion-compensated prediction in a B frame may be forward, backward, or a combination of the two (selected in the macroblock layer). Since no other frames are predicted from them, B frames may be coarsely quantised in areas of high motion and comprise mainly motion prediction information elsewhere.

In order to keep all frames in the coded bit stream causal, B frames are always transmitted **after** the I / P frames to which they refer, as shown at the bottom of fig 5.3.

One of the main ways that the H.263 (enhanced H.261) standard is able to code at very low bit rates is the incorporation of the B frame concept.

Considerable research work at present is being directed towards more sophisticated motion models, which are based more on the outlines of objects rather than on simple blocks. These will form the basis of extensions to the new low bit-rate video standard, MPEG-4 (MPEG-3 is an **audio** coding standard).