

8. JPEG and JPEG 2000 PICTURE CODING

General idea: **Joint Photographic Experts Group (JPEG)**, formally known as ISO/IEC JTC1/SG29/WG1 standard, the work started in 1983 and draft international standard was formulated in 1991. It is widely used for image exchange, storing imagery, www, and digital photography. All processing operations are performed after segmenting images into 8x8 blocks and performing discrete cosine transform. The transformed coefficients are then quantized and entropy coded either by an arithmetic coder (QM-coder with binary decision tree) or by Huffman coding. Motion JPEG (MJPEG) is uniformly used for digital video editing. The performance of the technique is very much dependent on the content of the imagery at hand and the constraints the user sets. With the introduction of JPEG2000, the quality has significantly improved.

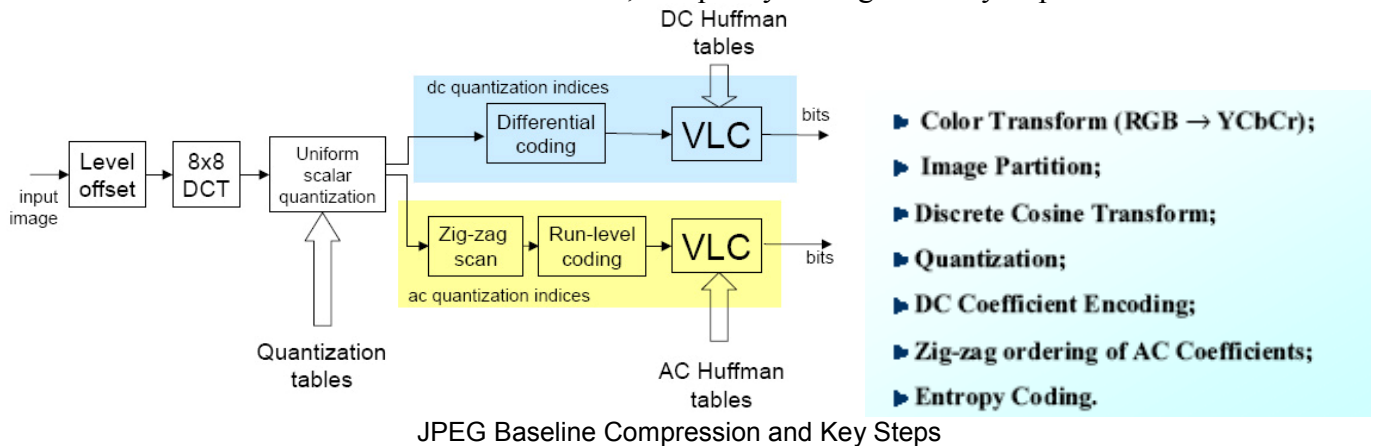
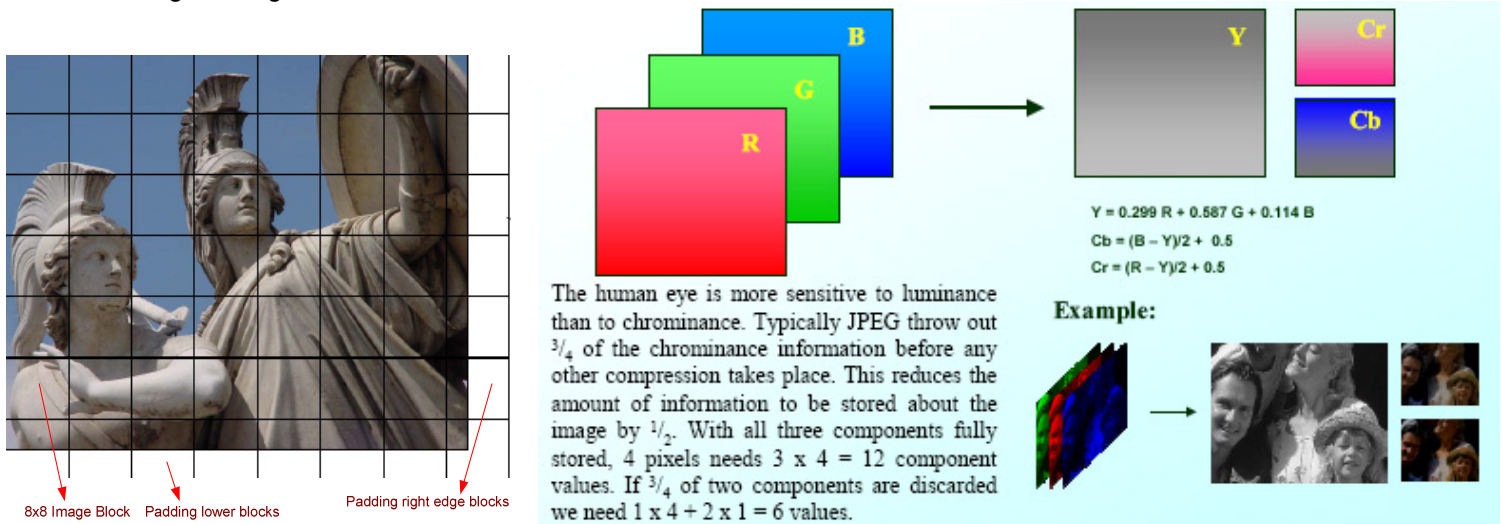


Image arrangement and color transformation:



Baseline Algorithm: Images are padded with extra pixels to conform nice 8x8 blocks. Color transformation for better mimicking the visual perception is handled through the following mapping:

$$Y = 0.3 * R + 0.6 * G + 0.1 * B; \quad C_r = \frac{B - Y}{2} + 0.5; \quad C_b = \frac{R - Y}{1.6} + 0.5 \quad (8.1)$$

1. **DCT Computation:** Over 8x8 blocks, a constant intensity-level of 128 is subtracted from each pixel of 8 bits resolution. Then 2-D DCT is performed.

2. **Quantization Matrix:** DCT coefficients are threshold quantized using a quantization matrix (QM), and then reordered using a zig-zag scanning.
3. **Variable-rate Code (VLC) assignment:** DC coefficient is coded by a DPCM coder relative to the DC of the previous block. Non-zero AC terms are Huffman coded.
4. Chrominance channels are sub-sampled by a factor of 2 in both directions depending on the selection of the case:

- Subsampling
 - 4:4:4 (no subsampling)
 - 4:2:2 (Cr, Cb horizontal subsampling)
 - 4:2:0 (Cr, Cb horizontal + vertical subsampling)

Y1	Y2	Y3	Y4
Y5	Y6	Y7	Y8
Y9	Y10	Y11	Y12
Y13	Y14	Y15	Y16

Cr1	Cr2
Cr3	Cr4

Cb1	Cb2
Cb3	Cb4

5. Pixels or color images are either non-interleaved (three scans) or interleaved to require a single scan.
6. Non-interleaved and interleaved orderings:
 - Scan 1: Y1, Y2, Y3, ..., Y16*
 - Scan 2: Cr1, Cr2, Cr3, Cr4*
 - Scan 3: Cb1, Cb2, Cb3, Cb4*

whereas the interleaved ordering is structured as:

Y1, Y2, Y3, Y4, Cr1, Cb1, Y5, Y6, Y7, Y8, Cr2, Cb2, ...

7. Bit Allocation / Quantization:

It represents which coefficients must be retained and how coarsely each retained coefficient has to quantized via *zonal* or *threshold* coding.

- **Zonal coding:** Locations of coefficients with the K largest variances are indicated by a zonal mask. Same for all blocks.
- **Bit Assignment:** Suppose a total of B bits/frame is the rate and there are M retained coefficients with variances: $\sigma_i^2; i = 1, \dots, M$. The number of bits allocated for each coefficient is given by:

$$b_i = \frac{B}{M} + \frac{1}{2} \log_2 \sigma_i^2 - \frac{1}{2M} \sum_{i=1}^M \log_2 \sigma_i^2$$

- **Threshold Coding:** It is an adaptive method, where only those coefficients whose magnitudes are above a threshold are retained within each block. This operation is expressed in terms of an Quantization Matrix (QM).

$$\hat{S}(k_1, k_2) = NINT\left[\frac{S(k_1, k_2)}{T(k_1, k_2)}\right]$$

where $\hat{S}(k_1, k_2)$ is a limited by a threshold mechanism and quantized approximation of $S(k_1, k_2)$ and $T(k_1, k_2)$ is the corresponding element of the quantization matrix (QM). A coefficient value at location (k_1, k_2) is retained if $\hat{S}(k_1, k_2) \neq 0$. The elements of QM are 8-bit integers which determine the quantization step size for each location. The choice of the QM depends on the source noise level and the viewing conditions. In general, coarser the quantization, larger the weights used for higher-frequency coefficients.

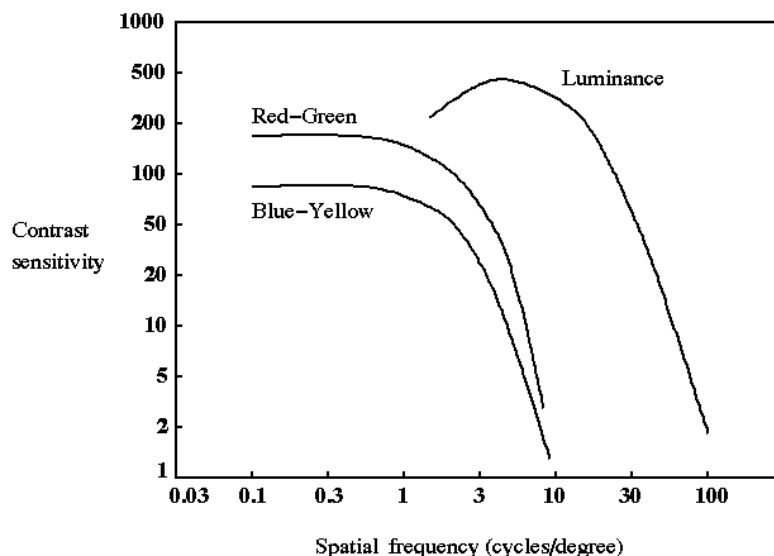
Table: Possible quantization tables

Uniform quantization	More accurate quantization	Less accurate quantization
16 16 16 16 16 16 16 16	1 2 2 4 4 8 16 16	8 64 64 128 256 256 256 256
16 16 16 16 16 16 16 16	2 4 8 8 8 16 16 32	64 128 128 128 256 256 256 256
16 16 16 16 16 16 16 16	4 4 8 16 16 16 32 32	128 256 256 256 256 256 256 256
16 16 16 16 16 16 16 16	4 8 16 16 16 32 32 32	256 256 256 256 256 256 256 256
16 16 16 16 16 16 16 16	8 16 16 32 32 32 32 32	256 256 256 256 256 256 256 256
16 16 16 16 16 16 16 16	8 16 16 32 32 32 64 64	256 256 256 256 256 256 256 256
16 16 16 16 16 16 16 16	16 16 32 32 32 32 64 64	256 256 256 256 256 256 256 256
16 16 16 16 16 16 16 16	16 16 32 32 32 64 64 64	256 256 256 256 256 256 256 256

Table : Default JPEG quantization tables

Luminance	Chrominance
16 11 10 16 24 40 51 61	17 18 24 47 99 99 99 99
12 12 14 19 26 58 60 55	18 21 26 66 99 99 99 99
14 13 16 24 40 57 69 56	24 26 56 99 99 99 99 99
14 17 22 29 51 87 80 62	47 66 99 99 99 99 99 99
18 22 37 56 68 109 103 77	99 99 99 99 99 99 99 99
24 35 55 64 81 104 113 92	99 99 99 99 99 99 99 99
49 64 78 87 103 121 120 101	99 99 99 99 99 99 99 99
72 92 95 98 112 100 103 99	99 99 99 99 99 99 99 99

Sensitivity of the Eye to Luminance and Chrominance Variations



Example 8. 1: Perform JPEG baseline coding for the following 8x8 DCT blocks:

Level shifted 8 × 8 original image

Forward DCT Values


Note: block is shifted by 128 to make 0 mean

-76	-73	-67	-62	-58	-67	-64	-55	-415	-29	-62	25	55	-20	-1	3
-65	-69	-62	-38	-19	-43	-59	-56	7	-21	-62	9	11	-7	-6	6
-66	-69	-60	-15	16	-24	-62	-55	-46	8	77	-25	-30	10	7	-5
-65	-70	-57	-6	26	-22	-58	-59	-50	13	35	-15	-9	6	0	3
-61	-67	-60	-24	-2	-40	-60	-58	11	-8	-13	-2	-1	1	-4	1
-49	-63	-68	-58	-51	-65	-70	-53	-10	1	3	-3	-1	0	2	-1
-43	-57	-64	-69	-73	-67	-63	-45	-4	-1	2	-1	2	-3	1	-2
-41	-49	-59	-60	-63	-52	-50	-34	-1	-1	-1	-2	-1	-1	0	-1

JPEG recommended quantization matrix from step 7

Quantized coefficient array

16	11	10	16	24	40	51	61	-26	-3	-6	2	2	0	0	0
12	12	14	19	26	58	60	55	1	-2	-4	0	0	0	0	0
14	13	16	24	40	57	69	56	-3	1	5	-1	-1	0	0	0
14	17	22	29	51	87	80	62	-4	1	2	-1	0	0	0	0
18	22	37	56	68	109	103	77	1	0	0	0	0	0	0	0
24	35	55	64	81	104	113	92	0	0	0	0	0	0	0	0
49	64	78	87	103	121	120	101	0	0	0	0	0	0	0	0
72	92	95	98	112	100	103	99	0	0	0	0	0	0	0	0





- 1-D coefficient sequence after zig-zag scanning:

-26 -3 1 -3 -2 -6 2 -4 1 -4 1 1 5 0 2 0 0 -1 2 0 0 0 0 0 -1 -1 EOB

where EOB denotes the end of the block.

- Coding the DC coefficient: Encode the difference between DC coefficients of the current and previous blocks.



- Actual bit assignment is done according to categories for both DC and AC coefficients:

Range	DC Difference Category	AC Category
0	0	N/A
-1, 1	1	1
-3, -2, 2, 3	2	2
-7, ..., -4, 4, ..., 7	3	3
-15, ..., -8, 8, ..., 15	4	4
-31, ..., -16, 16, ..., 31	5	5
-63, ..., -32, 32, ..., 63	6	6
-127, ..., -64, 64, ..., 127	7	7
-255, ..., -128, 128, ..., 255	8	8
-511, ..., -256, 256, ..., 511	9	9
-1023, ..., -512, 512, ..., 1023	A	A
-2047, ..., -1024, 1024, ..., 2047	B	B
-4095, ..., -2048, 2048, ..., 4095	C	C
-8191, ..., -4096, 4096, ..., 8191	D	D
-16383, ..., -8192, 8192, ..., 16383	E	E
-32767, ..., -16384, 16384, ..., 32767	F	N/A

- Default DC codes:

Category	Base Code	Length	Category	Base Code	Length
0	010	3	6	1110	10
1	011	4	7	11110	12
2	100	5	8	111110	14
3	00	5	9	1111110	16
4	101	7	A	11111110	18
5	110	8	B	111111110	20

- Coding the AC coefficients: Define (RUN,LEVEL) as symbols: (1,-3); (1,1); (1,-3); ..., (2,-1)
- These symbols are VLC (Huffman or arithmetic) coded according to the run/category.

Run/ Category	Base Code	Length	Run/ Category	Base Code	Length
0/0	1010 (= EOB)	4			
0/1	00	3	8/1	11111010	9
0/2	01	4	8/2	11111111000000	17
0/3	100	6	8/3	111111110110111	19
0/4	1011	8	8/4	111111110111000	20
0/5	11010	10	8/5	111111110111001	21
0/6	111000	12	8/6	111111110111010	22
0/7	1111000	14	8/7	111111110111011	23
0/8	111110110	18	8/8	111111110111100	24
0/9	111111110000010	25	8/9	111111110111101	25
0/A	111111110000011	26	8/A	111111110111110	26
1/1	1100	5	9/1	111111000	10
1/2	111001	8	9/2	111111110111111	18
1/3	1111001	10	9/3	111111111000000	19
1/4	11110110	13	9/4	111111111000001	20
1/5	1111110110	16	9/5	111111111000010	21
1/6	111111110000100	22	9/6	111111111000011	22
1/7	111111110000101	23	9/7	111111111000100	23
1/8	111111110000110	24	9/8	111111111000101	24
1/9	111111110000111	25	9/9	111111111000110	25
1/A	111111110001000	26	9/A	111111111000111	26
2/1	11011	6	A/1	111111001	10
2/2	11111000	10	A/2	111111111001000	18
2/3	111110111	13	A/3	111111111001001	19
2/4	111111110001001	20	A/4	111111111001010	20
2/5	111111110001010	21	A/5	111111111001011	21
2/6	111111110001011	22	A/6	111111111001100	22
2/7	111111110001100	23	A/7	111111111001101	23
2/8	111111110001101	24	A/8	111111111001110	24
2/9	111111110001110	25	A/9	111111111001111	25
2/A	111111110001111	26	A/A	111111111010000	26
3/1	111010	7	B/1	111111010	10
3/2	111110111	11	B/2	111111111010001	18
3/3	11111110111	14	B/3	111111111010010	19
3/4	111111110010000	20	B/4	111111111010011	20
3/5	111111110010001	21	B/5	111111111010100	21
3/6	111111110010010	22	B/6	111111111010101	22
3/7	111111110010011	23	B/7	111111111010110	23
3/8	111111110010100	24	B/8	111111111010111	24
3/9	111111110010101	25	B/9	111111111011000	25
3/A	111111110010110	26	B/A	111111111011001	26
4/1	111011	7	C/1	1111111010	11
4/2	111111000	12	C/2	111111111011010	18
4/3	111111110010111	19	C/3	111111111011011	19
4/4	111111110011000	20	C/4	111111111011100	20
4/5	111111110011001	21	C/5	111111111011101	21
4/6	111111110011010	22	C/6	111111111011110	22
4/7	111111110011011	23	C/7	111111111011111	23
4/8	111111110011100	24	C/8	111111111100000	24
4/9	111111110011101	25	C/9	111111111100001	25
4/A	111111110011110	26	C/A	111111111100010	26

Example 8. 2: Performance of JPEG decoder, which implements the inverse operations and the following is obtained: The reconstruction errors varies in $(-25,+25)$. This is considered a reasonable level of JPEG compression.



Lena: Original (bpp = 8.0; mse = 0.00)



JPEG (bpp = 1.00; mse = 17.26)



JPEG: (bpp = 0.50; mse = 33.08)



JPEG: (bpp = 0.25; mse = 79.11)



Original: (bpp = 8.00; mse = 0.00)



JPEG: (bpp = 1.00; mse = 17.26)



JPEG: (bpp = 0.50; mse = 33.08)



JPEG: (bpp = 0.25; mse = 79.11)

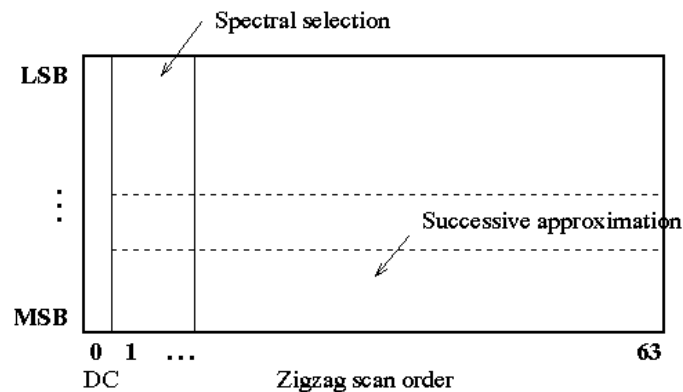
Magnifications of Lena compressed by JPEG.

Example 8. 3: Explore JPEG encoding using the following two demos:

1. VCDemo from Delft University: <http://www-ict.its.tudelft.nl/vcdemo>.
2. JPEG Color Imagery Compression Demo from Simon Fraser University in Burnaby, B.C., Canada: <http://www.cs.sfu.ca/CC/365/mark/interactive-jpeg/Ijpeg.html>

JPEG Modes:


- Sequential/Baseline mode
- Lossless mode
- Progressive mode
- Hierarchical mode
- Motion JPEG: Sequential JPEG is applied to each image frame in a video.



JPEG – Progressive Level:

- The progressive mode consists of a sequence of "scans" each of which codes a part of the quantized DCT coefficients.
- Spectral selection: The DCT coefficients are grouped into spectral bands. The lower frequency bands are usually coded (sent) first.
- Successive approximation: The information is first sent with lower precision, and then refined in later scans. Two processes may be combined to provide a graceful progression.

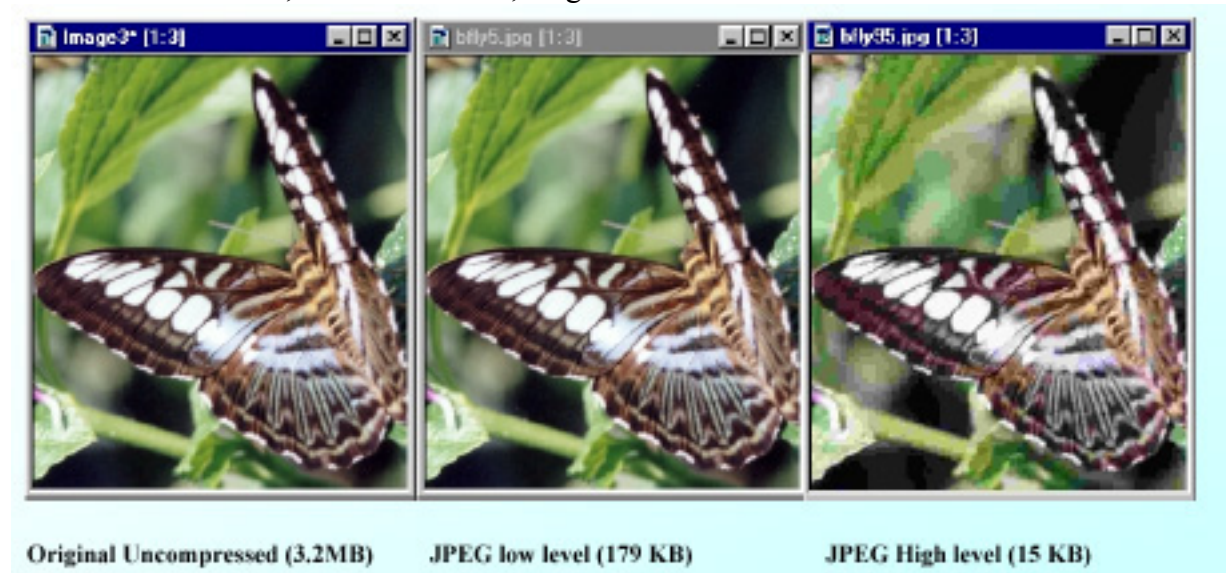
JPEG – Hierarchical Level:



- The first stage (lowest resolution) is coded using one of the sequential or progressive JPEG modes. The output of each hierarchical stage is then up-sampled (interpolated) and used as the prediction for the next stage.
- The image quality at extremely low bit rates surpasses any of the other JPEG modes, but this is achieved at the expense of a higher bit rate at the completion of the progression.

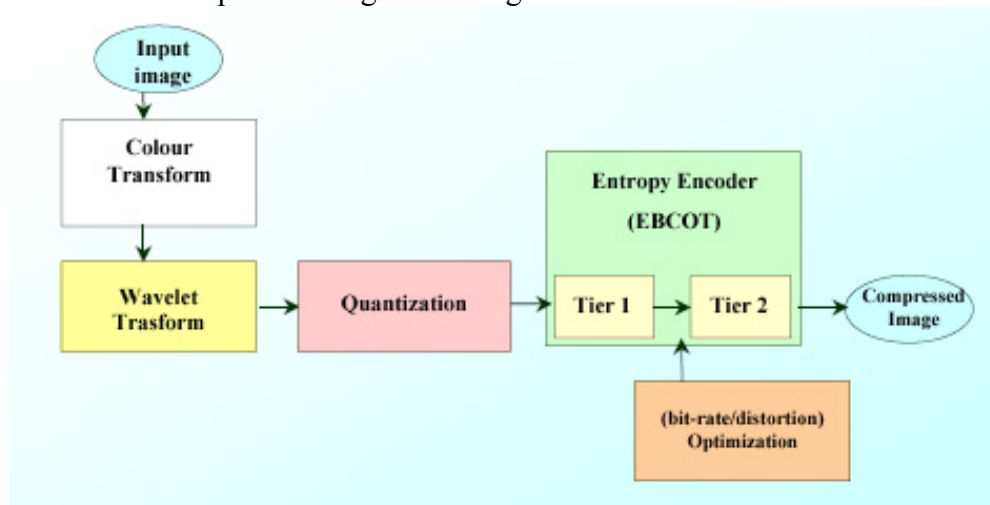
JPEG - Adaptive Quantization

- Allows spatially adaptive quantization, where the quantization matrix can be scaled block-to-block.
 - e.g., separate between high-activity (edges), medium activity (texture) and uniform blocks based on a measure of the intensity variance.
 - provides up to 30% better performance as compared to non-adaptive quantization.
- ISO DIS 10918-3, JPEG Extension, August 1994.



JPEG2000

Block diagram of the JPEG2000 standard and its impact on compression can be clearly seen from a small demo at 0.125 bits/pixel coding of an image.



• 0.125 bpp



JPEG



JPEG2000

JPEG2000 is a Wavelet-based compression and as in the JPEG and other DCT-based compression systems it consists of the following stages:

- Filtering
- Quantizer
- Entropy coding
- Arithmetic coding
- Bit allocation

1. **Aim** of the JPEG2000 was to develop a new still image coding standard for different types of still images (bi-level, gray-level, color, multi-component, hyper-component), with different characteristics (natural, scientific, medical, remote sensing, text, rendered graphics, compound, etc.), allowing different imaging models (client/server, real-time transmission, image library archival, limited buffer and bandwidth resources, etc.) preferably within a unified and integrated system.

2. Objectives

- Superior low bit-rate performance both in lossless and lossy modes
- Continuous-tone and bi-level compression
- Progressive transmission by pixel accuracy and resolution
- Robustness to bit-errors
- Open architecture
- Sequential build-up capability (real time coding)
- Backward compatibility with JPEG
- Content-based description
- Protective image security

3. JPEG2000 Markets and Applications

- Internet, Mobile, and E-Commerce Applications
- Printing, Scanning, Digital Photography
- Remote Sensing
- Facsimile
- Medical
- Digital Libraries

4. JPEG2000 Features

- High compression efficiency in one algorithm both lossless and lossy
- Lossless color transformations
- Progressive by resolution and quality
- Static and dynamic Region-of-Interest
- Error resilience
- Multiple component images
- Block and line based transforms

5. JPEG2000 Requirements:

- Higher compression efficiency than current JPEG
- Backward compatibility with current JPEG
- Progressive coding (by accuracy and by resolution)
- ROI coding (static and dynamic)
- Error resilience capabilities
- Object oriented functionalities (coding, information embedding, ...)

Embedded Block Coding with Optimized Truncation (EBCOT):

- Each sub-band is partitioned into a set of blocks
- All blocks within a sub-band have the same size (possible exception for the blocks at the image boundaries)
- Blocks are encoded independently
- Post-processing operation determines the extent to which each block's bit stream should be truncated
- Final bit stream is composed of a collection of "layers."

Why block coding?

- exploit local variations in the statistics of the image from block to block
- provide support for applications requiring random access to the image


- reduce memory consumption in hardware implementations of the compression or decompression engine
- Allow for parallel implementation

Types of Coding Operations:

- Zero coding (ZC)
- Run-Length coding (RLC)
- Sign coding (SC)
- Magnitude refinement (MR)
 - Arithmetic coding is used
- Reduced complexity in “lazy coding mode”

Zero Coding (ZC):

- Use of 1 of 9 different context states to code the value of the symbol, depending upon the significance state variables of:
 - Immediate horizontal neighbors (h)
 - Immediate vertical neighbors (v)
 - Immediate diagonal neighbors (d)
 - Non-immediate neighbors (f)



Run-Length Coding (RLC):


- . used in conjunction with the ZC primitive, in order to reduce the average number of binary symbols which must be encoded using the arithmetic coding engine
- . Sign coding (SC)
- . used at most once for each sample in the block immediately a previously insignificant symbol is found to be significant during a Zero Coding or Run-Length Coding operation
- . Magnitude Refinement (MR)
- . used to encode an already significant sample

If the sample is non yet significant, a combination of the "Zero Coding" (ZC) and "Run-Length Coding" (RLC) primitives is used to encode whether or not the symbol is significant in the current bit-plane

If so, the "Sign Coding" (SC) primitive must also be invoked to send the sign

- If the sample is already significant, the "Magnitude Refinement" primitive is used to encode the new bit position

JPEG2000 vs JPEG Performance:



Example 8. 4: Explore JPEG2000 compression using VcDemo.

1. Summary of JPEG Image Compression Standards:

ITU-T G3/G4	Binary images (non-adaptive) - for fax transmission of documents (1980)
ISO JBIG	Binary and bit-plane encoding - to handle progressive transmission & halftones (1994)
ISO JPEG	<i>Still frame gray scale and color image</i> - Block based DCT (1993)
ISO JPEG-LS	New Lossless Coding Standard -Nonlinear prediction, context-based, Rice-Goulomb coding (1997)
ISO JPEG 2000	Wavelet Coding Standard (finalized 2002)

2. Summary on Performance criteria in image compression

The aim of image compression is to transform an image into compressed form so that the information content is preserved as much as possible. Two main criteria of measuring the performance of an image compression algorithm are:

- *Compression Efficiency*,
- *Distortion* caused by the compression algorithm.
- *Speed* of the compression and decompression process.

The standard way to measure them is to fix a certain bit rate and then compare the distortion caused by different methods. In on-line applications the waiting times of the user are often critical factors. In the extreme case, a compression algorithm is useless if its processing time causes an intolerable delay in the image processing application. In an image archiving system one can tolerate longer compression times if the compression can be done as a background task. However, fast decompression is usually desired.

Among other interesting features of the compression techniques we may mention the *robustness* against *transmission errors*, and *memory requirements* of the algorithm. The compressed image file is normally an object of a data transmission operation. The transmission is in the simplest form between internal memory and secondary storage but it can as well be between two remote sites via transmission lines. The data transmission systems commonly contain fault tolerant internal data formats so that this property is not always obligatory. The memory requirements are often of secondary importance, however, they may be a crucial factor in hardware implementations.

From the practical point of view the last but often not the least feature is *complexity of the algorithm itself*, i.e. *the ease of implementation*. Reliability of the software often highly depends on the complexity of the algorithm. Let us next examine how these criteria can be measured.

Compression efficiency: Most obvious measure of the compression efficiency is the *bit rate*, which gives the average number of bits per stored pixel of the image:

$$\text{bit rate} = \frac{\text{size of the compressed file}}{\text{pixels in the image}} = \frac{C}{N} \quad (\text{bits per pixel})$$

where C is the number of bits in the compressed file, and $N (=X \cdot Y)$ is the number of pixels in the image. If the bit rate is very low, *compression ratio* might be a more practical measure:

$$\text{Compression Ratio} = \frac{\text{size of the original file}}{\text{size of the compressed file}} = \frac{N \cdot k}{C}$$

where k is the number of bits per pixel in the original image. The overhead information (header) of the files is ignored here.

Distortion: Distortion measures can be divided into two categories: subjective and objective measures. A distortion measure is said to be *subjective*, if the quality is evaluated by humans. The use of human analysts, however, is quite impractical and therefore rarely used. The weakest point of this method is the subjectivity at the first place. It is impossible to establish a single group of humans (preferably experts in the field) that everyone could consult to get a quality evaluation of their pictures. Moreover, the definition of distortion highly depends on the application, i.e. the best quality evaluation is not always made by people at all.

For objective measures distortion is calculated as the *difference* between the original and the reconstructed image by a predefined function. It is assumed that original image is perfect. All changes are considered as occurrences of distortion, no matter how they appear to a human observer. The quantitative distortion of the reconstructed image is commonly measured by the **mean absolute error (MAE)**, **mean square error (MSE)**, and **peak-to-peak signal to noise ratio (PSNR)**:

$$\begin{aligned} \text{MAE} &= \frac{1}{N} \sum_{i=1}^N |y_i - x_i| \\ \text{MSE} &= \frac{1}{N} \sum_{i=1}^N (y_i - x_i)^2 \\ \text{PSNR} &= 10 \cdot \log_{10} \left[255^2 / \text{MSE} \right], \quad \text{assuming } k=8. \end{aligned}$$

These measures are widely used in the literature. Unfortunately these measures do not always coincide with the evaluations of a human expert. The human eye, for example, does not observe small changes of intensity between individual pixels, but is sensitive to the changes in the average value and contrast in larger regions. Thus, one approach would be to calculate the *mean values* and *variances* of some small regions in the image, and then compare them between the original and the reconstructed image. Another deficiency of these distortion functions is that they measure only local, pixel-by-pixel differences, and do not consider global artifacts, like *blockiness*, *blurring*, or the *jaggedness of the edges*.