Image Compression

- Image compression means the reduction of the amount of data required to represent a digital image by removing the redundant data. It involves reducing the size of image data files, while retaining necessary information.
- Mathematically, this means transforming a 2D pixel array (i.e. image) into a statistically uncorrelated data set. The transformation is applied prior to storage or transmission of the image. At later time, the compressed image is decompressed to reconstruct the original (uncompressed) image or an approximation of it.
- The ratio of the original (uncompressed) image to the compressed image is referred to as the *Compression Ratio* C_R:

$$C_R = \frac{Uncompressed\ Image\ Size}{Compressed\ Image\ Size} = \frac{Usize}{Csize}$$

where

$$Usize = M \times N \times k$$

$$Csize = size\ of\ compressed\ image\ file\ stored\ in\ a\ disk$$

Example:

Consider an 8-bit image of 256×256 pixels. After compression, the image size is 6,554 bytes. Find the compression ratio.

Solution:

Usize = (256 × 256 × 8) / 8 = 65,536 bytes

Compression Ratio = 65536 / 6554 = 9.999 ≈ 10 (also written 10:1)

This means that the original image has 10 bytes for every 1 byte in the compressed image.

Image Data Redundancies

There are three basic data redundancies that can be exploited by image compression techniques:

- Coding redundancy: occurs when the data used to represent the image are not utilized in an optimal manner. For example, we have an 8-bit image that allows 256 gray levels, but the actual image contains only 16 gray levels (i.e. only 4-bits are needed).
- Interpixel redundancy: occurs because adjacent pixels tend to be highly correlated. In most images, the gray levels do not change rapidly, but change gradually so that adjacent pixel values tend to be relatively close to each other in value.
- Psychovisual redundancy: means that some information is less important to the human visual system than other types of information. This information is said to be *psychovisually redundant* and can be eliminated without impairing the image quality.

Image compression is achieved when one or more of these redundancies are reduced or eliminated.

Fidelity Criteria

These criteria are used to assess (measure) image fidelity. They quantify the nature and extent of information loss in image compression. Fidelity criteria can be divided into classes:

1. Objective fidelity criteria
2. Subjective fidelity criteria

Objective fidelity criteria

These are metrics that can be used to measure the amount of information loss (i.e. error) in the reconstructed (decompressed) image.

Commonly used objective fidelity criteria include:

- **root-mean-square error**, $\boldsymbol{e_{rms}}$, between an input image *f(x,y)* and output image $\hat{f}(x,y)$:

$$e_{rms} = \sqrt{\frac{1}{MN}\sum_{x=0}^{M-1}\sum_{y=0}^{N-1}\left[\hat{f}(x,y) - f(x,y)\right]^2}$$

where the images are of size *M* × *N*. The smaller the value of e_{rms}, the better the compressed image represents the original image.

- **mean-square signal-to-noise ratio, *SNR_{ms}***:

$$SNR_{ms} = \frac{\sum_{x=0}^{M-1}\sum_{y=0}^{N-1}\hat{f}(x,y)^2}{\sum_{x=0}^{M-1}\sum_{y=0}^{N-1}\left[\hat{f}(x,y) - f(x,y)\right]^2}$$

- **Peak signal-to-noise ratio, $\boldsymbol{SNR_{peak}}$**:

$$SNR_{peak} = 10\log_{10}\frac{(L-1)^2}{\frac{1}{MN}\sum_{x=0}^{M-1}\sum_{y=0}^{N-1}\left[\hat{f}(x,y) - f(x,y)\right]^2}$$

where *L* is the number of gray levels.

A larger number of *SNR* implies a better image.

Subjective fidelity criteria

These criteria measure image quality by the subjective evaluations of a human observer. This can be done by showing a decompressed image to a group of viewers and averaging their evaluations. The evaluations may be made using an absolute rating scale, for example {Excellent, Fine, Passable, Marginal, Inferior, and Unusable}.

Image Compression System

As shown in the figure below, the image compression system consists of two distinct structural blocks: an encoder and a decoder.

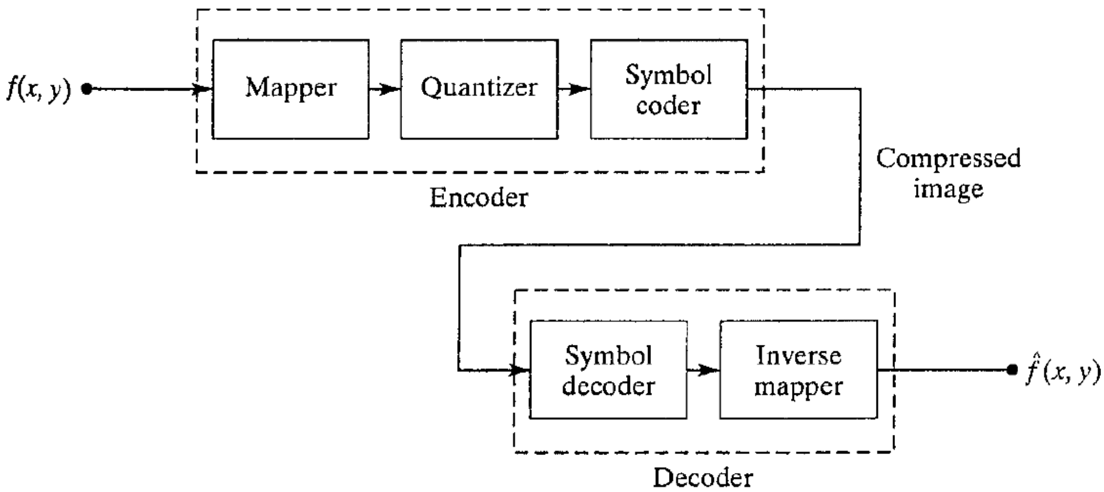

Figure 14.1 A general image compression system

The encoder is responsible for reducing or eliminating any coding, interpixel, or psychovisual redundancies in the input image. It consists of:

- Mapper: it transforms the input image into a nonvisual format designed to reduce interpixel redundancies in the input image. This operation is reversible and may or may not reduce directly the amount of data required.

- Quantizer: it reduces the psychovisual redundancies of the input image. This operation is irreversible.
- Symbol coder: it creates a fixed- or variable-length code to represent the quantizer output. In a variable-length code, the shortest code words are assigned to the most frequently occurring output values, and thus reduce coding redundancy. This operation is reversible.

The decoder contains only two components: symbol decoder and an inverse mapper.

Image Compression Types

Compression techniques are classified into two primary types:

- Lossless compression
- Lossy compression

Lossless compression

- It allows an image to be compressed and decompressed without losing information (i.e. the original image can be recreated exactly from the compressed image).
- This is useful in image archiving (as in the storage of legal or medical records).
- For complex images, the compression ratio is limited (2:1 to 3:1). For simple images (e.g. text-only images) lossless methods may achieve much higher compression.
- An example of lossless compression techniques is Huffman coding.

Huffman Coding

is a popular technique for removing coding redundancy. The result after Huffman coding is variable length code, where the code words are unequal length. Huffman coding yields the smallest possible number of bits per gray level value.

Example:

Consider the 8-bit gray image shown below. Use Huffman coding technique for eliminating coding redundancy in this image.

119	123	168	119
123	119	168	168
119	119	107	119
107	107	119	119

Solution:

Gray level	Histogram	Probability
119	8	0.5
168	3	0.1875
107	3	0.1875
123	2	0.125

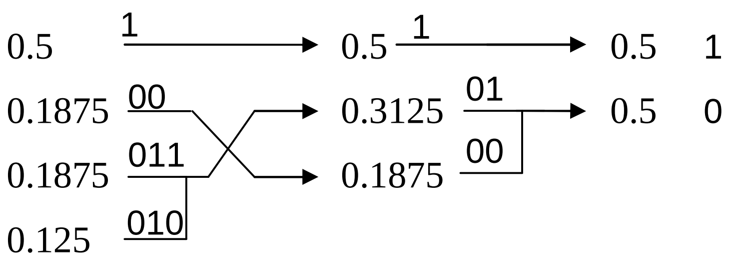

We build a lookup table:

Lookup table:

Gray level	Probability	Code
119	0.5	1
168	0.1875	00
107	0.1875	011
123	0.125	010

We use this code to represent the gray level values of the compressed image:

1	010	00	1
010	1	00	00
1	1	011	1
011	011	1	1

Hence, the total number of bits required to represent the gray levels of the compressed image is 29-bit: 10101011010110110000011110011.

Whereas the original (uncompressed) image requires 4*4*8 = 128 bits.

Compression ratio = 128 / 29 ≈ 4.4

Lossy compression

- It allows a loss in the actual image data, so the original uncompressed image cannot be recreated exactly from the compressed image).
- Lossy compression techniques provide higher levels of data reduction but result in a less than perfect reproduction of the original image.
- This is useful in applications such as broadcast television and videoconferencing. These techniques can achieve compression ratios of 10 or 20 for complex images, and 100 to 200 for simple images.
- An example of lossy compression techniques is JPEG compression and JPEG2000 compression.