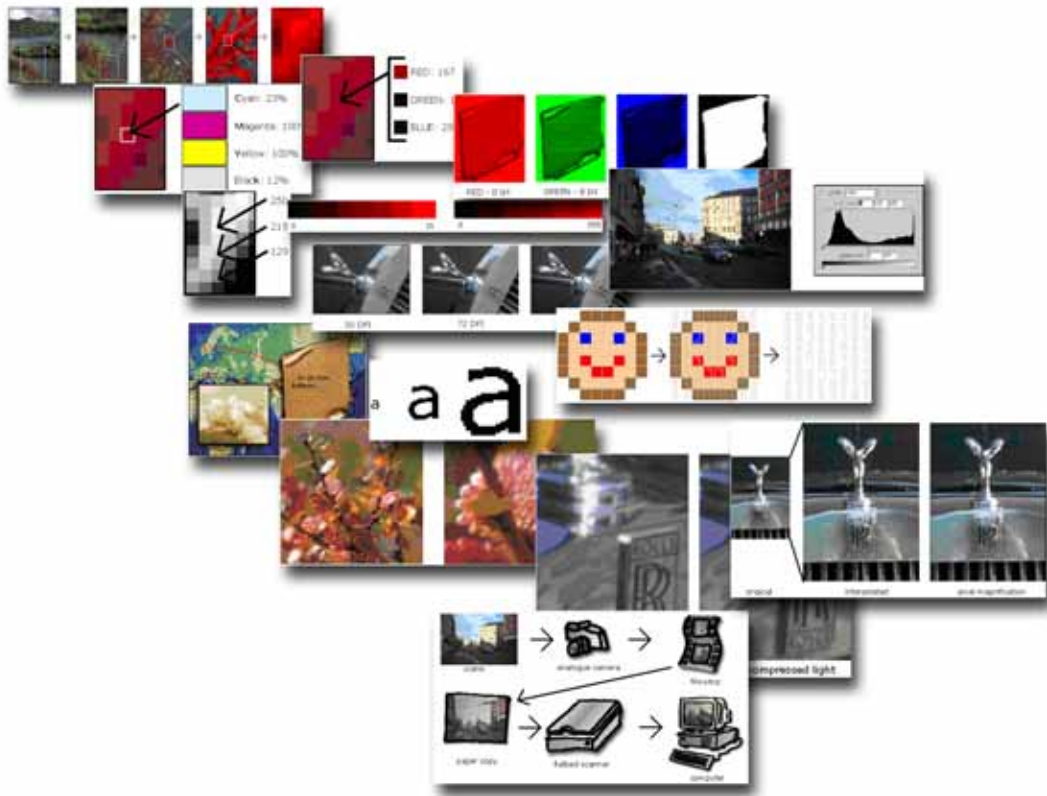


digital images



CC Eirik Solheim 2005
www.eirikso.com

INTRODUCTION TO DIGITAL IMAGES	3
Pixels	3
Description of colour and brightness	3
CMYK	3
RGB	4
Other methods of representation	4
Greyscale and black/white	4
Colour depth	5
Compression	6
Loss-less run length compression	6
Indexed colours	7
Photo compression	8

Introduction to digital images

Pixels

A digital image is an image that is represented by digits (numbers). A photo is a typical analogue representation with huge amounts of different shades and details. To represent a photo digitally it has to be deconstructed into smaller pieces that can be described by numbers. Magnifying a digitised picture several times can reveal the basic parts of it.




Figure 1 - Digital picture magnification

The illustration clearly shows that the picture consists of small rectangular dots. In digital images these dots are called *pixels* – **P**icture **E**lements.

Description of colour and brightness

Each pixel has to be described to be able to store its colour and brightness. There are several different ways to describe these attributes. The most common method is to deconstruct the colour into elements based on primary colours. The primary colours are different depending on the target medium.

CMYK

For print on paper a system called CMYK is the most common. This is based on the primary colours for subtractive blending. CMYK is short for **C**yan, **M**agenta, **Y**ellow and **B**lack (K is used to avoid confusion with B for Blue).




Figure 2 - Pixel deconstructed to CMYK

The CMYK levels are given in percent, from 0 to 100.

RGB

For representation on a screen or projector another system for colour blending is used. This is called RGB and is based on the primary colours for additive blending, **Red**, **Green** and **Blue**.




Figure 3 - Pixel deconstructed to RGB

The RGB levels are given by numeric values.

Other methods of representation

In addition to CMYK and RGB a pixel can also be described by HSB-values. This is short for **H**ue, **S**aturation and **B**rightness. This method is based on the human perception of colour. Hue is given by an angle of 0 to 360 degrees, pointing to a colour in the standard colour wheel. Saturation reflects the strength of the colour, from 0% (grey) to 100% (fully saturated). Brightness is the relative darkness or lightness of the colour. Most often given by a percentage between 0 and 100.

L*a*b is yet another way to describe colour. The LAB system was developed in 1931 to be a system for colour representation independent of medium. A colour represented by the LAB system consists of three components. L for luminance (lightness), A-component (from green to red) and B-component (from blue to yellow).

Greyscale and black/white

A picture can also be stored in greyscale or simple black and white. For a greyscale picture each pixel have only one value, from 0 to 255. A value of 0 is black and a value of 255 is white.




Figure 4 - Greyscale

For true black and white each pixel is described by a single bit, 0 or 1.




Figure 5 – Black and white

Colour depth

The R, G and B colours can be described with different levels of accuracy. For representation on screen an accuracy of 8 bits per colour is the most common. This gives a possibility of describing each colour with a number from 0 to 255. For very high quality prints on paper a greater level of detail is needed. Professional image editing programs usually gives the possibility of a colour depth of up to 16 bits. A 16-bit colour depth gives 65535 different shades of each colour.



Figure 6 - Colour depth: 4 bits (16 shades) and 8 bits (256 shades)

When 8 bits are used for each channel, red, green and blue. A total of 24 bits is needed to describe the picture. When a picture is stored with all this information it is called a 24-bit picture.

Compression

High quality pixel-based pictures contain large amounts of digital information.




Figure 7 - Simple 1-bit picture

The simple picture in Figure 7 will be stored in the computer as a series of zero and one:

```
0,0,1,1,1,1,0,0,0,1,0,0,0,0,1,0,1,0,1,0,0,1,0,1,1,0,0,0,0,0,0,1,1,0,1,0,0,1,0,1,1,0,0,1,1,0,
0,1,0,1,0,0,0,0,1,0,0,0,0,1,1,1,0,0
```

The header of this picture-file would tell the computer that this is a 1-bit picture and that the dimensions are 8 by 8 pixels. After ordering the bits in rows with 8 bits in each row it starts to look like an image:




Figure 8 - From bits to image

In a computer everything is stored as a representation of zero and one. A simplified view could be to compare the zeros and ones to switches. A switch that is off or a switch that is on. The simple picture in Figure 7 needs 64 switches to be described. When sent over a network a total of 64 small signals have to be transmitted. Storage and network communication cost time and money. If it was possible to describe the same information with less amount of switches it would be possible to send it faster over a network or store it on a cheaper hard-disk.

Loss-less run length compression

One of the simplest forms for compression is called run length compression (RLE). This is usually a so-called loss-less compression technique. It reduces the amount of numbers that is needed to describe an image.

If we add colour to the simple image from Figure 7 we need three numbers between 0 and 255 to describe each pixel.




Figure 9 – Three numbers for each pixel (RGB)

The picture contains 64 pixels; each is described with three numbers. This sums up to a total of 196 numbers to describe the picture.

A picture like this one has an amount of pixels that have identical colours. The first line of the picture is described by the following sequence:

(255, 255, 255), (255, 255, 255), (153, 102, 53), (153, 102, 53), (153, 102, 53), (153, 102, 53), (255, 255, 255), (255, 255, 255).

The two first pixels are identical (white), and then there are four identical brown pixels, and two more identical white ones. This information can also be described like this:

2, (255, 255, 255), 4, (153, 102, 53), 2, (255, 255, 255)

For the computer this is less information. If we describe the whole picture like this we end up with a total of 116 numbers to describe exactly the same that we originally needed 196 numbers to describe.

Indexed colours

The 24-bit RGB system for describing colour is flexible and let us describe a picture with substantial accuracy. Still, for simple images this system can be too complicated. The image in Figure 9 contains a total of 5 different colours. Instead of describing each colour with three values each time it appears it is possible to reduce the amount of information needed by using indexing.

If we start by defining a table of colours:

- 0- (255, 255, 255) White
- 1- (153, 102, 53) Dark brown
- 2- (255, 204, 153) Light brown
- 3- (0,0,255) Blue
- 4- (255,0,0) Red

Then we can map these numbers to each pixel:

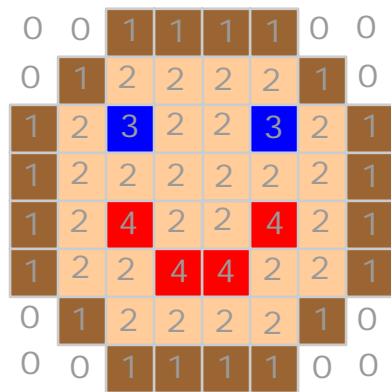


Figure 10 - Indexed colours

The picture now contains 64 numbers. The table of colour itself contain 15 numbers. This gives a total of 79 numbers to describe the picture. If we add the RLE-compression to this we can compress it down to 65 numbers.


The GIF-file format utilizes both RLE-compression and colour indexing. The highest quality GIF-file cannot contain more than 256 different colours. This kind of compression is very effective on simple shapes, logos, diagrams and drawings.

Photo compression

A digital photo is usually much more complex than the image in Figure 10. By combining different forms of compression techniques it is possible to reduce the file size of high quality digital images substantially.

The most common compression technique for photos is JPEG. It uses advanced algorithms to take out information that is not needed to describe the visual impression of the photo. JPEG is not loss-less, which means that a picture that has been compressed with this technique contains less information than the original.

JPEG-compression is variable, and when saving an image the user can select the amount of compression wanted. A high amount of compression will be more visible in the picture than a lower level of compression.



Uncompressed



JPEG-compressed

Figure 11 - JPEG-Compression (JPG)

A high-resolution photo of 3500 by 2700 pixels will have these file-sizes before and after compression:

- 27 870 kBytes - Uncompressed TIF
- 18 705 kBytes - LZW-compressed TIF (Loss-less type of RLE-compression)
- 7 844 kBytes - JPG, low compression
- 222 kBytes - JPG, high compression
- 5 739 kBytes - GIF, 256 colours

JPG with the lowest amount of compression produces a picture that is nearly identical to the original. High compression JPG is extremely effective, but degrades quality. GIF is not very effective on photos, because reducing the amount of colours decreases the image quality.