

digital *multimedia*

nigel chapman and jenny chapman



Graphics and Colour
Video and Animation
Sound
Text and Typography
Hypermedia
Flash and DOM Scripting
Multimedia and Networks

Third
Edition

4

Bitmapped Images

**Based on material from
Digital Multimedia, 3rd edition
published by John Wiley & Sons, 2009
© 2009 Nigel Chapman and Jenny Chapman**

**These lecture slides © 2009
Nigel Chapman and Jenny Chapman**

All figures © MacAvon Media Productions

Resolution

Resolution is a measure of how finely a device approximates continuous images using finite pixels.

The resolution of scanners and printers is usually equated with their pixel density (dots per inch).

The resolution of video frames and computer monitors is usually equated with their pixel dimensions (width \times height).

The resolution of a digital still camera is often quoted as the total number of pixels in the largest image it can record.

The pixel density of monitors and still cameras is assumed to be 72 dpi (“screen resolution”).

A bitmapped image has pixel dimensions but no intrinsic pixel density.

The physical size of an image when it is displayed will depend on the pixel density of the device it is to be displayed on.


72 dpi



115 dpi




600 dpi




Device resolution and image size

Most image formats record a resolution (pixel density) together with the image data; this is usually the resolution of the device on which it originated.

72 dpi, 198 × 149 px



600 dpi, 1654 × 1240 px





Resolution and pixel dimensions

Scaling an image without changing its resolution changes its pixel dimensions in the same way as changing its resolution without changing its size.


Reducing the pixel dimensions is called “downsampling”; increasing them is called “upsampling”. Both can lead to a loss of quality.

Scaling can be done either by applying a transformation to each original pixel or by applying the inverse transformation to each pixel in the scaled image.

Interpolation is needed because of the finite size of pixels.




Scaling pixels



Scaling pixels using a reverse mapping


Nearest-neighbour, bilinear or bicubic interpolation may be used.

Nearest-neighbour interpolation is quickest but produces poor-quality results; bicubic is slowest but produces very good results; bilinear is in between.



Pixel interpolation

*Areas for bilinear
interpolation*



Nearest-neighbour (left), bilinear (middle) and bicubic (right) interpolation

Interpolation causes a loss of information which can never be recovered.

Image Compression

Images can be losslessly compressed using various methods, including run-length encoding (RLE), Huffman encoding and the dictionary-based LZ77, LZ78, LZW and deflate algorithms.

JPEG is the most commonly used lossy compression method for still images.

High-frequency information can be discarded from an image without perceptible loss of quality, because people do not perceive the effects of high frequencies in images very accurately.

The image is mapped into the frequency domain using the Discrete Cosine Transform (DCT).

The Discrete Cosine Transform is applied to 8×8 blocks of pixels.

Applying the DCT does not reduce the size of the data, since the array of frequency coefficients is the same size as the original pixel array.

The coefficients are quantized, according to a quantization matrix which determines the quality. The quantization discards some information.

After quantization there will usually be many zero coefficients. These are RLE-encoded, using a zig-zag sequence to maximize the length of the runs.

The non-zero coefficients are compressed using Huffman encoding.

Decompression is performed by reversing the process, using the Inverse DCT to recover the image from its frequency domain representation.

The decompressed image may exhibit compression artefacts, including blurring and visible edges at the boundaries between the 8×8 pixel blocks, especially at low quality settings.



Original (left) and JPEG (right)

JPEG2000 improves on JPEG in many areas, including image quality at high compression ratios. It can be used losslessly as well as lossily.

For JPEG2000 compression the image is divided into tiles, but these can be any size, up to the entire image.

A Discrete Wavelet Transform (DWT) is applied to the tiles, generating a wavelet decomposition, comprising a coarse (low resolution) version of the image and a set of detail coefficients that can be used to add progressively more detail to the image.



Wavelet decomposition of an image

The DWT may be reversible (lossless) or irreversible (lossy).

The detail coefficients in the wavelet decomposition may be quantized and are then losslessly compressed using arithmetic encoding.




Original (left) and JPEG2000 (right)

File Formats

GIF files use LZW compression and are restricted to 256 colours. One colour may be used to designate transparency.

They are most suitable for simple images with areas of flat colour.



A GIF with transparent areas

PNG was developed to supersede GIF. It uses deflate compression, is not restricted to 256 colours and supports alpha channels for partial transparency.

JPEG data can be stored in several different formats.

JFIF and SPIFF are compatible formats for JPEG images and are widely used on the Web.

Exif can hold either JPEG or TIFF data, together with extensive metadata.

JP2 and JPX formats are defined for storing JPEG2000 data.

TIFF is an extensible format, often used for storing uncompressed digital photographs, and for interchange of images.

BMP is a simple bitmapped image format that is native to Windows, but widely supported. BMP files are often uncompressed.

PDF documents can include bitmapped image data, that may be compressed using JPEG, JPEG2000, LZW, deflate, and others.

Camera raw data is used when complete control over image processing is required, but there is no standard format for camera raw data.

Adobe's DNG (Digital Negative) format is a standard, based on the TIFF format, intended for archiving camera raw images.

Image file formats differ in their support for metadata.


Image Manipulation

Image manipulation software provides high-level operations for systematically altering pixels. Most operations are described by analogy with traditional photographic techniques, such as the use of masks and filters.

Photoshop is the *de facto* industry standard; the Gimp is an Open Source alternative. Image Magick can be used for command-line processing.

Bitmapped images are manipulated to correct technical deficiencies, alter the content or create artificial compositions.

Images are often organized into layers, which are like overlaid sheets that may have transparent areas. Layers are used for compositing or experimenting with different versions of an image.



Compositing layers

Areas may be selected by drawing with marquee and lasso tools or a Bézier pen, or selected on the basis of colour similarity or edges using a magic wand or magnetic lasso.



Magic wand selection



Magnetic lasso selection and mask


Any selection defines a mask – the area that is not selected. Masked areas of the image are protected from changes.

A greyscale mask, which is partially transparent, is an alpha channel.

An alpha channel can be associated with a layer as a layer mask, and used for effects such as knock-outs and vignettes.



Compositing with a layer mask




Constructing a vignette with an alpha channel

In pixel point processing, each pixel's new value depends only on its old value.

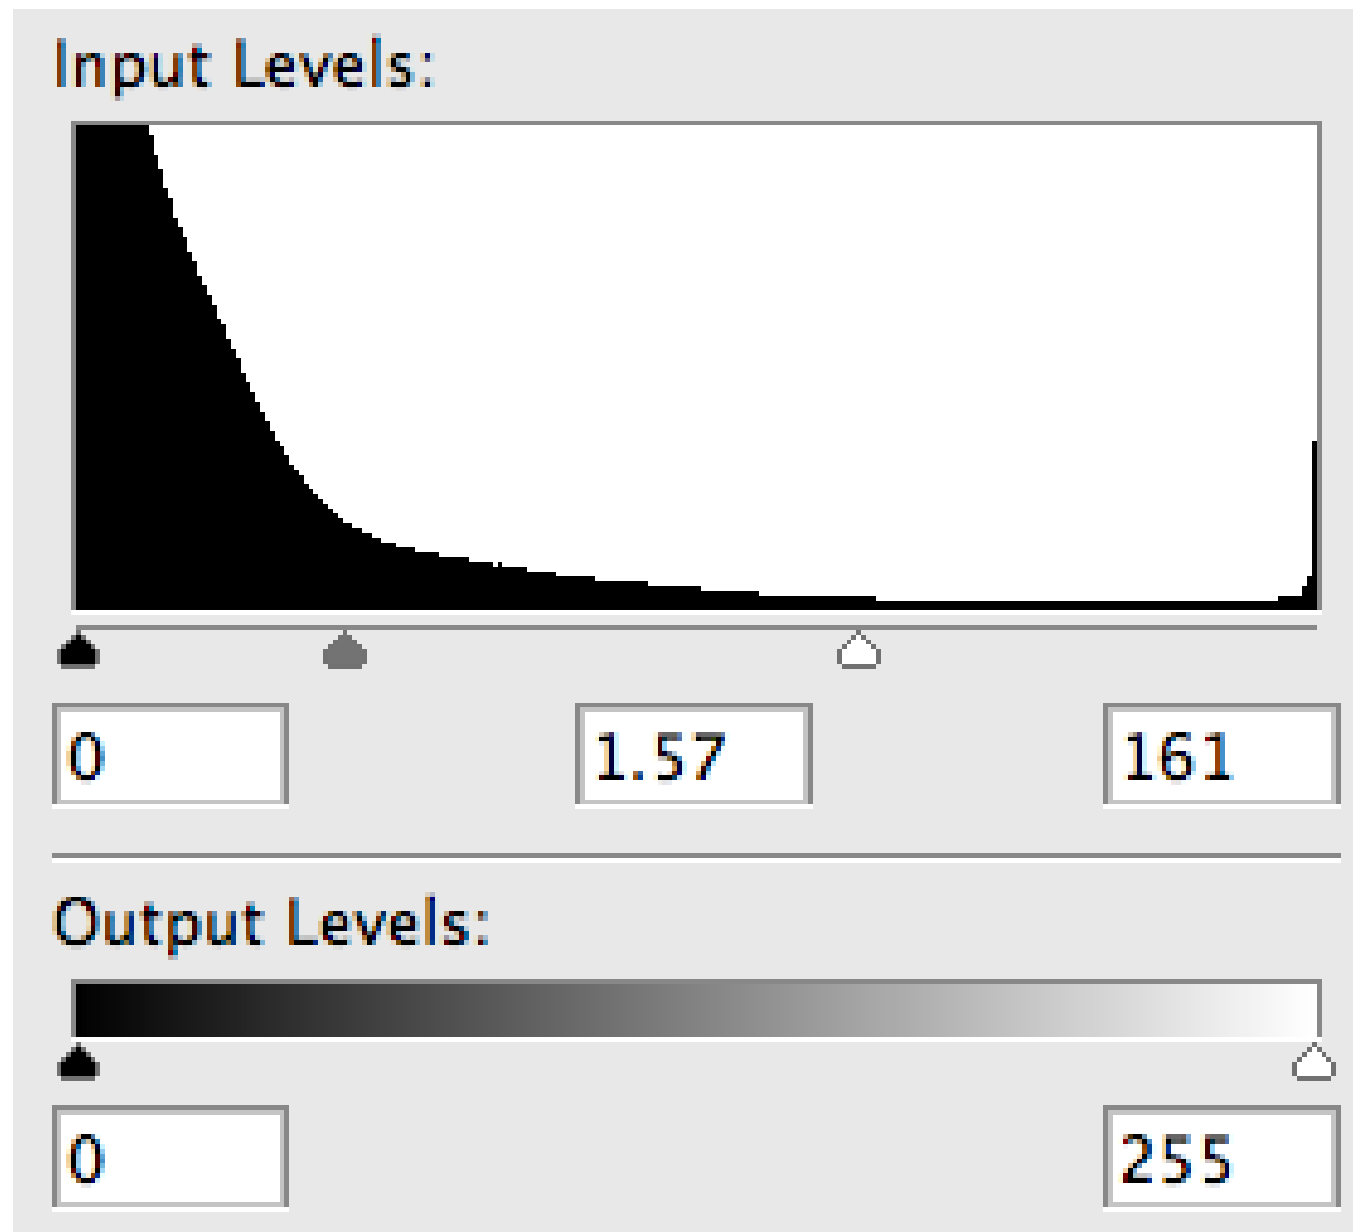
Brightness, contrast and levels are relatively crude pixel point adjustments.



Brightness (top) and contrast (bottom) adjustments




Brightness and contrast mapping functions




An image histogram in Photoshop's Levels dialogue

Curves adjustments provide full control over the relationship between original and new values. A sigmoid curve is often used to enhance contrast.




A sigmoid curve for enhancing contrast



Curve adjustment


Pixel group processing uses the values of neighbouring pixels as well.

The convolution operation in the frequency domain can be implemented as a weighted average in the spatial domain: for each pixel in the filtered image, the pixels of a convolution kernel are combined using a convolution mask.



Filtering with a 3×3 convolution mask

Simple blurring uses a 3×3 mask with equal values but produces crude results. Gaussian blur is preferred, as it produces more natural results.




The Gaussian bell curve




A drop shadow


Sharpening with a 3×3 mask is crude. Unsharp masking – combining an image with a Gaussian blurred copy of itself – produces better-looking results. Over-sharpening should always be avoided.



original image




blurred copy



difference

Unsharp masking



Unsharp masking, showing over-sharpening

The principle of convolution-based filtering can be used to create many special effects, but such filters should be used judiciously.