



An Introduction to Computer Images

13 September 2005

M. J. Dallwitz

Pixels and resolution

A 'pixel' (picture element) is the smallest component of a computer image. When an image is being manipulated by the computer, a pixel is regarded as a dot of a single colour (although when the image is displayed, each pixel may be made up of even smaller dots). An image is made up of pixels located on a rectangular grid. The resolution of the image may be defined as the number of points in the grid, e.g. 640x480 (640 pixels wide by 480 pixels high); or as the spacing of the grid points, e.g. 75dpi (75 dots per inch). When discussing the internal storage and processing of images, it is better to specify the number of grid points. The grid spacing is meaningful only when the image is outside the computer, that is, when it is being scanned or displayed.

Vector images

'Vector' images store mathematical descriptions of shapes. For example, a straight line could be represented as the coordinates of its start and finish; a circle as the coordinates of its centre, and its radius; and a text symbol as the coordinates of a set of points around its outline. The colour and thickness of lines can be specified, and closed curves can be filled with simple patterns. However, vectors cannot easily represent complex 'continuous-tone' images, such as a typical photograph.

Vector images are easy to modify. Each object in the image retains a separate identity, and can be altered, moved, or deleted, and placed over or under other objects. The mathematical objects are converted to pixels only when the image is to be displayed. Thus, output quality can be improved almost indefinitely by using high-resolution output devices such as laser printers or typesetters. Even on computer displays of comparatively low resolution, fine detail can be seen by zooming in on a part of the image.

These properties make vector images very suitable for drawings containing lines and simple shading. Programs for creating and editing vector drawings have been successfully used for biological illustration by several organizations.

Bitmap images

'Bitmap' ('raster') images store the colour value of each pixel. They are therefore suitable for storing photographic images, which can be converted to bitmap form by means of a scanner (see below). The resolution at which an image is scanned sets a permanent limit on the amount of detail that can be seen.

Bitmap images tend to be difficult to modify. In most formats, pixels are not explicitly grouped into larger structures or objects. If part of the image is moved or deleted, it is difficult or impossible to correctly fill in the area thus exposed, unless it is uniform or almost uniform. However, some proprietary formats associated with particular image editors do allow the grouping of parts of the image into objects or layers.

'Unmapped' ('true-colour' or 'grey-scale') images store the colour value for each pixel directly. Unmapped colour images usually use 3 bytes for each pixel — 1 for each of the primary colours (1 byte = 8 bits). This allows 256 intensity levels for each primary colour, and therefore 16777216 different colours. (Some formats use only 5 bits per primary colour, giving 32 intensity levels and 32768 different colours. This is rather unsatisfactory, as the eye can readily discern the difference between adjacent

levels.) Grey-scale images (i.e. continuous-tone 'black-and-white' photographs) require 1 byte per pixel, and true black-and-white images require only 1 bit per pixel.

'Mapped' ('paletted') images store colour values indirectly. The colours in a mapped image are restricted to a numbered set of colours, called the 'palette'. Each colour in the palette is defined by three numbers giving the intensities of the primary colours. The colour of a pixel is defined by specifying the number of its colour in the palette. The advantage of this scheme is a reduction in the amount of memory required to store the image. (For example, an image using a palette of 256 3-byte colours requires 1 byte per pixel, plus 768 bytes for the palette.) The disadvantage is that it may not be possible to represent the colours with sufficient accuracy. A typical colour image may contain several thousand different colours. A process called 'colour reduction' is used to choose a palette for an image, and to replace each colour by the closest colour from the palette. A palette of 256 colours can give reasonably good results for most images, but a careful examination of a mapped image will usually reveal colour 'contours', where the jump from one palette element to another takes place. This effect is particularly noticeable in large areas of gradually changing colour, such as a sky. A process called 'dithering', which may be carried out during the conversion of a true-colour image to a mapped image, reduces the effect by interspersing dots of different colours. However, dithering can substantially increase the size of an image file in compressed formats such as GIF (see below).

Image formats and compression

There is a large number of formats for storing digital images. Most were developed for use with a particular program, but some of these formats have become de facto standards, and are used by many programs. A few formats were developed specifically for interchanging files between different programs and computers.

Most bitmap formats have provision for reducing the file size by 'compressing' the image data. The amount of compression depends on both the compression method and the nature of the image. Most compression methods work well on images that have large areas of uniform colour. For example, a simple line drawing could usually be compressed to less than 1/10th of its original size.

The following are some of the formats most commonly used.

- GIF (CompuServe Graphics Interchange Format). A mapped bitmap format developed for file interchange. A good compression method is built into the format.
- TIF/TIFF (Tagged Image File Format). A very flexible bitmap format, supporting mapped and unmapped images. Several compression options are available.
- JPG/JPEG (Joint Photographic Experts Group). A bitmap format for unmapped images, developed as a standard by ISO and CCITT. The main feature of this format is extremely good compression of continuous-tone images. Some information is lost in the compression, and there is a trade off between amount of compression (which may be specified by the user) and image quality. For most images, 20:1 compression may be achieved without noticeable degradation. (By comparison, GIF format, and TIFF format with LZW compression, give about 2:1 compression with continuous-tone images. Because GIF images are mapped, they effectively have an additional compression of 3:1 for colour images, but information is lost in the colour reduction, as explained above.)
- PCD (PhotoCD). A bitmap format used by Kodak for their Photo CD disks. Each file contains 5 or 6 versions of the image, at different resolutions.
- PCX. A mapped bitmap format developed by Zsoft for their paint programs.
- CGM (Computer Graphics Metafile). A vector format developed for file interchange.

Creating and editing images

Bitmap images can be produced from existing drawings or photographs by 'scanners', which measure the brightness and/or colour of the original at each point of a grid. 'Flatbed' scanners scan large, usually opaque, material, such as the page of a book. The maximum resolution is usually 600–1200dpi. 'Slide' scanners scan small transparencies, such as 35mm slides or negatives. The maximum resolution is usually 2000–5000dpi.

Bitmap images are edited with ‘image editing’ programs, such as Adobe Photoshop. Image editing programs usually have provision for scaling, and conversion between different formats, but these facilities may be rather limited. Image Alchemy (Handmade Software, <http://www.handmade.com>) is a comprehensive program for format conversion, colour reduction, and scaling of bitmap images in Windows, MacOS X, and Unix.

Vector images are created and edited with ‘draw’ programs, such as Adobe Illustrator or CorelDraw. These programs often allow the use of a bitmap image as a guide for producing a vector image. A bitmap image containing well-defined edges, such as a scan of a line drawing, can be traced automatically to produce a first draft of the vector image. Bitmap images produced from continuous-tone photographs can be traced manually. The programs can also convert vector images to bitmap images.

Software for editing bitmap images makes heavy demands on memory and processing power. Large hard disks and backup systems may be necessary, particularly for bitmap images.

Recommendations for scanning and storing images

Images should be scanned and stored in such a way that they will be able to take advantage of improvements in display hardware. The resolution of the master copies should be as high as possible (preferably at least 1024x768). Lower-resolution versions can easily be made for display on currently available hardware.

Colour images should be scanned with at least 24 bits per pixel, and line drawings with 8 bits per pixel (i.e. 256 grey levels, NOT as ‘line art’ with 1 bit per pixel).

The scaling method built into scanning software is often very poor. It is best to do no scaling at this stage (‘100%’ scaling), or restrict it to integer factors (25%, 50%, 200%; not 33%). Note that there may be an interaction with the dots-per-inch (dpi) setting, which should be set to the ‘native’ setting of the scanner (usually 300dpi or 600dpi). If you want to check the quality of scaling, use line drawings with fine lines; poor scaling usually shows up as discontinuities in the lines.

Line drawings can be stored compactly as GIF or compressed TIFF files with 16 or 256 grey levels. Using 16 or 256 grey levels gives smoother lines than 2 levels (black and white). Using 256 levels does not give noticeably better results than 16 levels, and produces larger files. However, it is best to use 256 levels until any touching up, such as cleaning up the background and increasing the contrast, has been done.

Line drawings tend to become greyer when scaled to reduce the image size, particularly when the lines are fine. The appearance of the scaled versions can be improved by darkening the lines. This can be done by changing the ‘gamma’ of the image (provided the background of the original image has been made pure white, which should be done routinely). For example, the following Image Alchemy command will convert all GIF files in the current directory to a maximum size of 640x480 and 16 grey levels, with moderate darkening, and place the output files in the subdirectory ‘scaled’.

```
alchemy -g -d -b -c16 -+ -Xb640 -Yb480 -Gi1.0 -Go0.7 -o *.gif scaled
```

Lower values of ‘Go’ (‘gamma output’) give more darkening.

Continuous-tone colour images should be stored as TIFF or JPEG, and grey-scale images as GIF, TIFF, or JPEG. They can be converted to other formats as necessary. Mapped formats such as GIF should not be used for master copies of colour images, because of the sometimes-severe artifacts and information loss. If JPEG format is used for master copies, a high quality factor should be used. Even then the files will contain artifacts, which, though not perceptible, may have adverse effects on later processing of the image, particularly sharpening.