digital basics

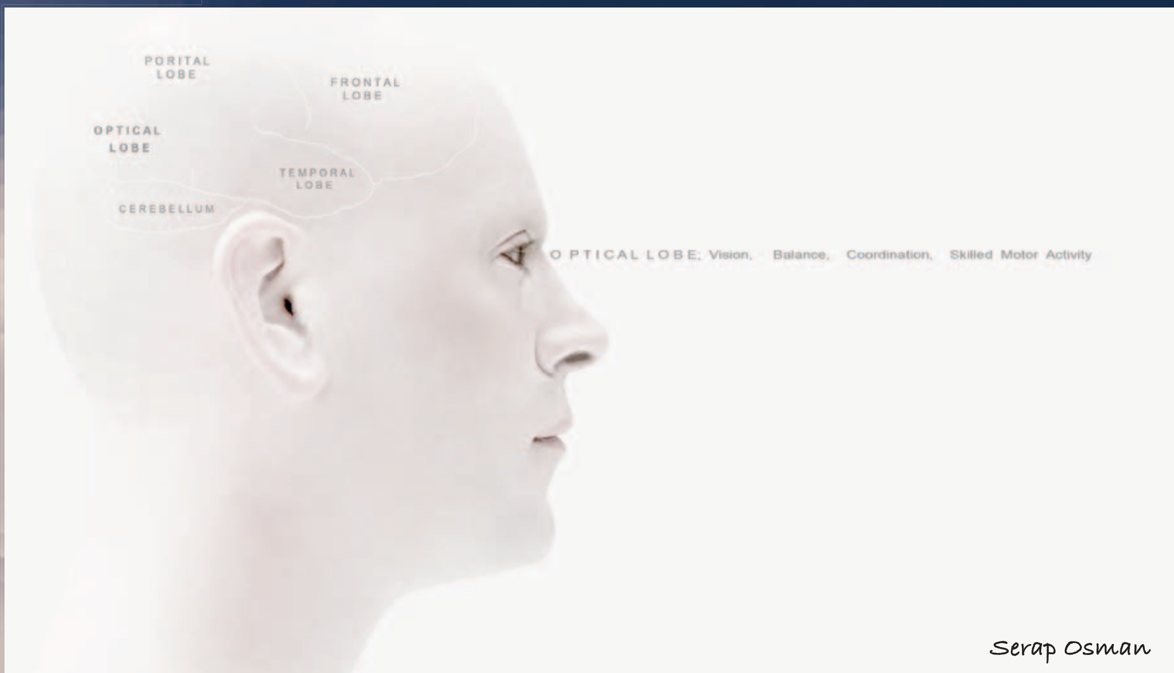

essential skills

- ~ Gain a working knowledge of digital image structure.
- ~ Understand file size, bit depth, image modes, channels, file format and resolution.
- ~ Understand color theory and color perception.

Introduction

Digital imaging is now revolutionizing not only the process of photography but also the way we view photography as a visual communications medium. This new photographic medium affords the individual greater scope for creative expression, image enhancement and manipulation. Before we rush into making changes to our digital files in order to create great art, or turn a warty old frog into a handsome prince, it makes sense to slow down and take time out to understand the structure of the digital image file. In this way the technical terms used to identify, quantify and specify the digital file as a whole, or the component parts of the digital file, serve to clarify rather than bamboozle our overloaded gray matter.

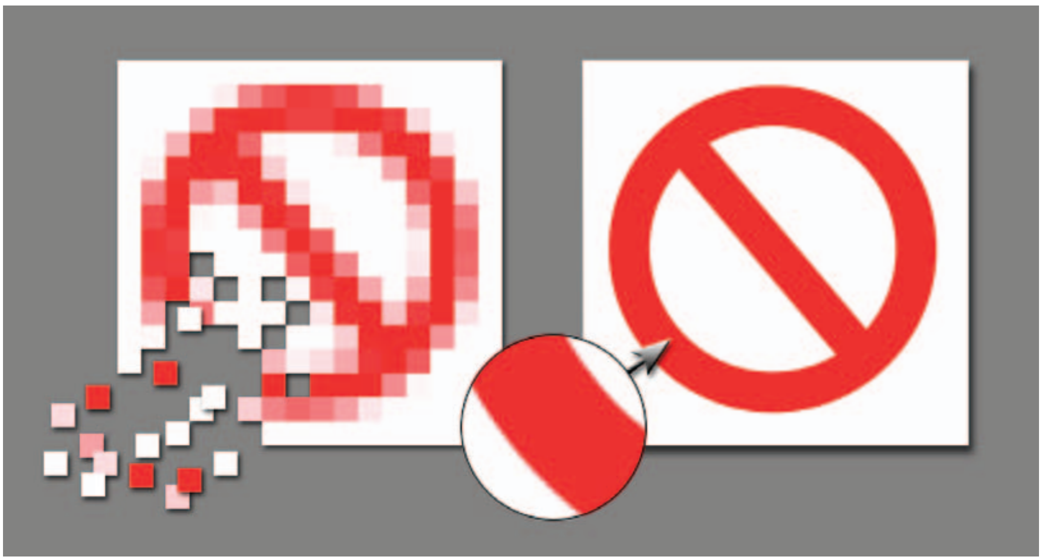

Anti-aliasing and small pixels ensure that a staircase of pixels is rendered as a smooth arc

Pixels

The basic building block of the digital image is the humble pixel (picture element). Pixels for digital imaging are square and positioned in rows horizontally and vertically to form a grid or mosaic. Each pixel in the grid is the same size and is uniform in color and brightness, i.e. the color does not vary from one side of the pixel to the other. If we fully zoom in on the pixels of a digital image, using image-editing software, we will see how smooth flowing shapes can be convincingly constructed out of rectangular building blocks (with not a curved pixel in sight). There are two processes used to create the illusion of curved lines in our photographs. The first is a process called anti-aliasing, where some of the edge pixels adopt a transitional (in-between) color to help create a smoother join between two different adjacent colors or tones. This process helps camouflage noticeable 'staircase' or 'shark's teeth' pixels. The most convincing way to render a smooth flowing line, however, is to simply display the pixels so small that we cannot make them out to be square using the naked eye.

Channels and modes

All the colors of the rainbow when mixed together create white light (a prism is often used to split white light into its component colors to demonstrate the connection between light and color). All the colors of the rainbow can be created by mixing just three of these colors – **Red**, **Green** and **Blue** light (called the **primary** colors of light) – in differing amounts. Using these simple scientific principles all the variations of color in our multicolored world can be captured and stored in three separate component parts of our digital image file. These component parts are called the Red, Green and Blue '**Channels**'. An image that uses this process to store the color data is called an RGB image. RGB is the most common type of **'Image Mode**'.

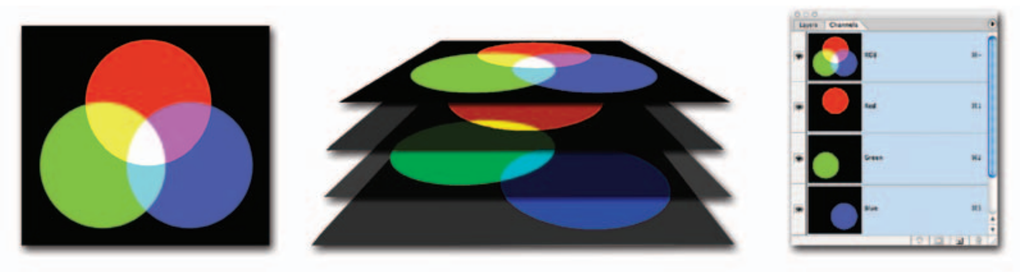

The primary colors of light (stored in three separate channels) create the secondary colors when mixed

In Adobe Photoshop the colored RGB channels are the powerful backbone of Photoshop, working behind the scenes to create multicolored images by providing three sets of information regarding color, i.e. the amount of red, green and blue present in each pixel location. When the color from only one channel is present, a primary color is created in the image window. When information from two channels is present a secondary color is displayed. These secondary colors (created by mixing two primaries) are called Cyan, Magenta and Yellow (CMY). When there is an absence of any color from the three RGB channels the pixel location appears black (no illumination). Mixing all three RGB channels together creates white light or gray if the brightness value from each of the three channels is lowered (see 'Levels'). Color information about the image can also be stored using the secondary colors (mixing two secondary colors creates a primary) plus black (K). Images using this system or **Mode** are called **CMYK** images. Photoshop users can view the information stored in the component channels by clicking on the Channels palette tab.

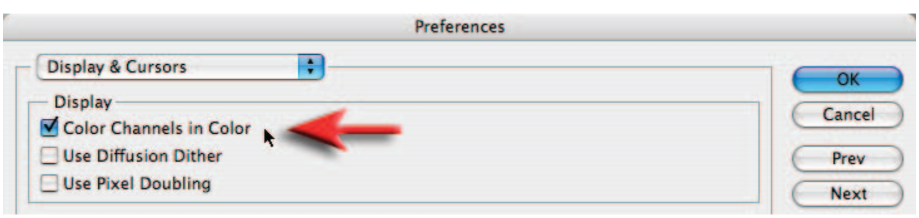

Note > You can view the information in each channel with or without color (Edit/ Photoshop > Preferences > Display & Cursors). It is usually beneficial to view the information in the channels without color when conducting advanced post-production editing but for the purposes of understanding what is actually happening, color is a distinct advantage.

Levels

We have seen how mixing primary colors of light can create the secondary colors. In the previous illustration, where three colored circles were overlapped, the color in each of the three channels is either 'on' or 'off'. In this way six colors are created from three RGB channels. The three channels can, however, house a greater range of information about color than simply 'yes' (fully on) or 'no' (fully off) in any one pixel location. Capture devices are capable of measuring '**how much**' color is present in any one given location. In a standard RGB image, 256 different levels of color can be assigned to each pixel location. The channels operate very much like a mixing desk, mixing varying amounts of color from each of the three color channels to create the full color spectrum.

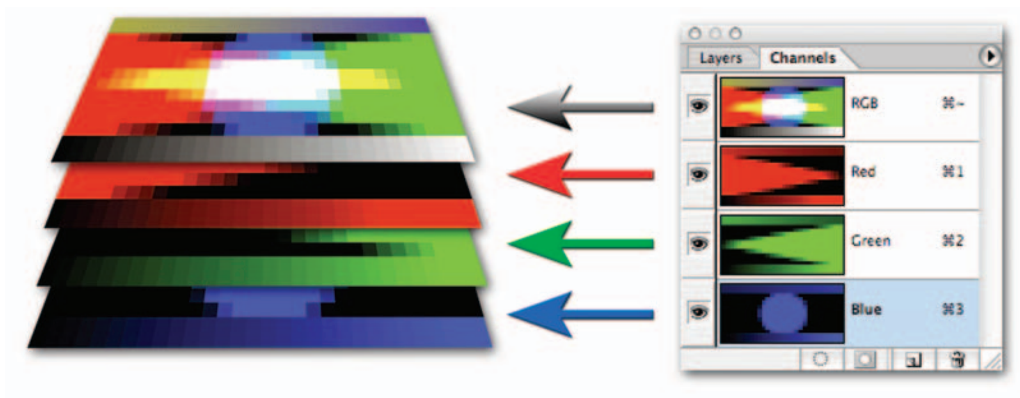

If the three channels are mixed in equal proportions what we see is a series of tonal steps from black (0 in all three channels) to white (255 in all channels).

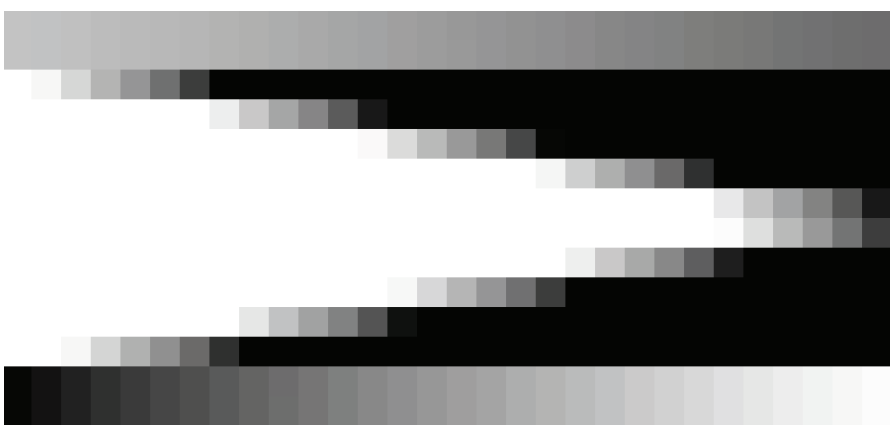

256 levels of tone are reduced to 30 so the steps can be clearly seen

256 separate tones are sufficient to create a smooth transition from dark to light with no visible steps. If the pixels are sufficiently small when printed out, the viewer of the image cannot see either the individual pixels or the steps in tone, and the illusion of '**continuous tone**' or '**photographic quality**' is achieved.

Hue, Saturation and Brightness

Equally high levels in each of the three color channels creates not only a bright or light toned pixel but may also indicate a bright level of illumination in the scene that has been captured. This can be attributed to the pixel being a record of a bright light, or the reflected light off a brightly illuminated subject, but it may alternatively be due to possible overexposure during the capture process, i.e. the sensor being exposed to the light source for too long.

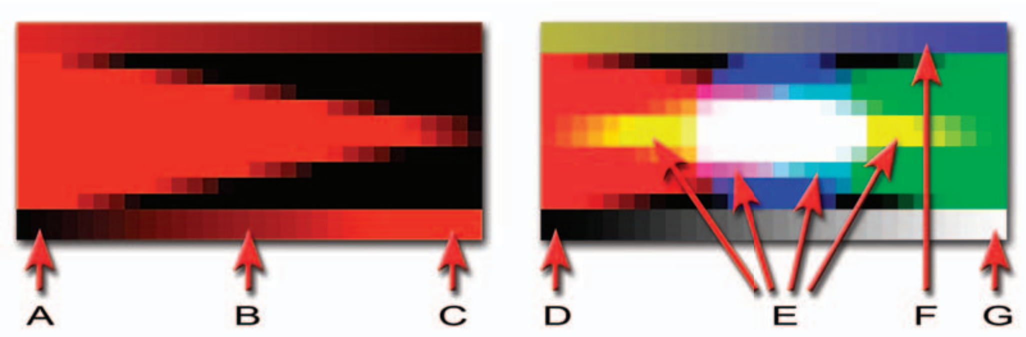

The illustration above shows the enormous variety of tones that can be achieved by combining 256 levels of information from the three color channels. Given that each channel can vary its level of information, independently of the other two, it is possible to create a single pixel with one of a possible 16.7 million different values (256 × 256 × 256). To describe the nature of a particular color value without resorting to numbers, Adobe has adopted a system where the characteristics of the color can be described in three ways. These descriptive categories are:

Hue – as dictated by the dominant primary or secondary color, e.g. red, yellow, blue, etc.
Saturation – the strength of the color, e.g. when one or two of the channels registers 0 the resulting color is fully saturated, i.e. no level of gray or white is weakening the purity of the color.
Brightness – from 0 (black – all channels 0) to 255 (at least one channel registering 255).

Using a common language of Hue, Saturation and Brightness (HSB) we can identify the colors indicated by letters in the illustration above, in terms that can be readily understood by the broader community that are neither mathematicians nor Photoshop nerds.

A to C are levels of brightness from black to fully saturated bright red. The levels from the other two RGB channels are not influencing the overall color or brightness of any of the pixels (all Green and Blue values are set to level 0).

D and G are levels of brightness from black to white. When all RGB channels read the same level the resulting tones are fully desaturated.

E indicates fully saturated secondary colors created by mixing two primary RGB channels at level 255.

F indicates colors of lower saturation as information from the three RGB channels is unequal (therefore creating a gray component to the color's characteristic).

Color Picker

As we have discussed previously it is essential when describing and analyzing color in the digital domain to use the appropriate terminology. A greater understanding of the characteristics of Hue, Saturation and Brightness (HSB) can be gained by viewing colors in the Adobe Color Picker.

To open the Color Picker, click on the foreground swatch in the Tools palette.

> **Hue** – Click on any color in the vertical color bar (in the center of the Color Picker) to view the range of colors associated with that particular hue. Note the number in the top field next to the 'H' radio button. Each of the six primary and secondary colors are positioned 60° apart, e.g. Red at 0°, Yellow at 60°, Green at 120°, etc. Each hue is assigned a number or 'degree' between 0° and 360°. Saturation and Brightness values are not effected by changes in this bar.
>
> **Saturation and Brightness** – Click in the large square box to the left of the bar to choose a Saturation and Brightness value for the selected Hue. Saturation increases when the selection circle is moved to the right side of the box and decreases when moved to the left. Brightness increases towards the top of the box and decreases towards the bottom of the box. Hue is not affected by these changes.

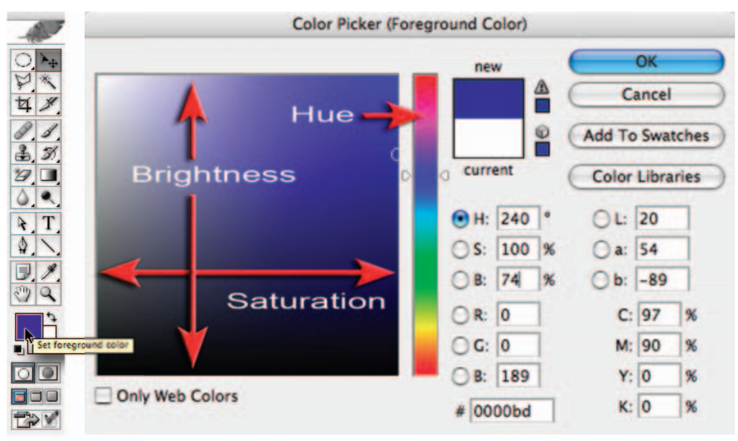

The 'foreground swatch' and 'Color Picker'

Creating and sampling color from an image

If the mouse cursor is moved out of the Color Picker dialog box and into the image window the mouse cursor icon turns into the Eyedropper icon, regardless of what tool was selected in the Tools palette when the Color Picker was opened.

1. With the Color Picker open and an image open behind the picker, press the Caps Lock key on the keyboard to turn this icon into a target for precise selection of a color or tone.
2. Click in an area of the image window to sample the color and reveal its characteristics in the Color Picker dialog box.
3. Change the sample size from a single pixel to a '3 by 3 Average' or '5 by 5 Average' by right-clicking (PC) or Control-clicking (Mac) in the image window to reveal the context menu for the Eyedropper Tool.

Color and light overview

Additive color

The additive primary colors of light are Red, Green and Blue or RGB. Mixing any two of these primary colors creates one of the three secondary colors Magenta, Cyan or Yellow.

Note > Mixing all three primary colors of light in equal proportions creates white light.

RGB – additive color

Subtractive color

The three subtractive secondary colors are Cyan, Magenta and Yellow or CMY. Mixing any two of these secondary colors creates one of the three primary colors Red, Green or Blue. Mixing all three secondary colors in equal proportions in a CMYK file creates black or an absence of light.

CMY – subtractive color

Activity 1– Channels and Info

1. Open the files 'RGB' and 'CMY' from the supporting DVD. When opening these files choose 'Leave as is (don't color manage)' in the Missing Profile dialog box.
2. Open the Channels palette to see how these six colors plus white and black were created using information from three (RGB) or four (CMYK) channels. Use the Eyedropper Tool and the Info palette (Window > Info) to measure the color values.

Hue, Saturation and Brightness

Although most of the digital images are captured in RGB it is sometimes a difficult or awkward color model for some aspects of color editing. Photoshop allows the color information of a digital image to be edited using the HSB model.

Hue, Saturation and Brightness or HSB is an alternative model for image editing which allows the user to edit either the Hue, Saturation or Brightness independently of the other two.

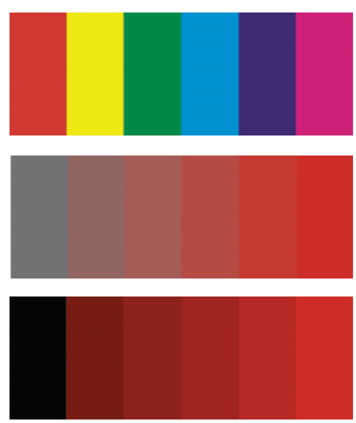

HSB – Hue, Saturation and Brightness

Activity 2 – Color Picker

1. Open the HSB image from the supporting DVD.
2. Click on the foreground color swatch in the Tools palette to open the Color Picker.
3. Move the cursor into the image window and click on each color in turn to review the HSB color values.

Color perception

Our perception of color changes and is dependent on many factors. We perceive color differently when viewing conditions change. Depending on the tones that surround the tone we are assessing, we may see it darker or lighter. Our perception of a particular hue is also dependent on both the lighting conditions and the colors or tones that are adjacent to the color we are assessing.

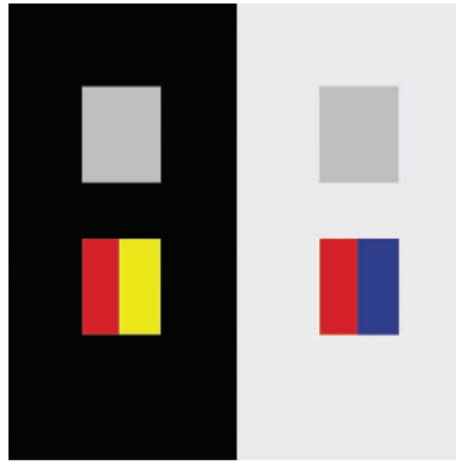

Color perception

Activity 3

Evaluate the tones and colors in the image opposite. Describe the gray squares at the top of the image in terms of tonality. Describe the red bars at the bottom of the image in terms of hue, saturation and brightness. Measure the actual values. If they are the same why do they appear different?

Color gamut

Color gamut (or range) varies, depending on the quality of paper and colorants used (inks, toners and dyes, etc.). Printed images have a smaller color gamut than transparency film or monitors and this needs to be considered when printing. In the image opposite the out of gamut colors are masked by a gray tone. These colors are not able to be printed using the default Photoshop CMYK ink values.

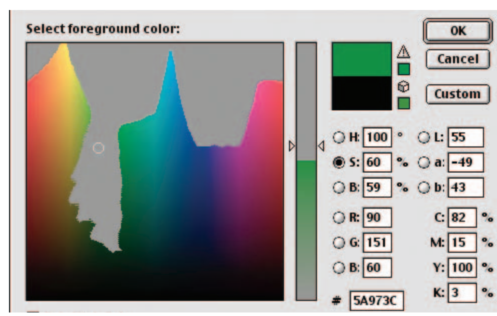

Out of gamut colors

Color management issues

The issue of obtaining consistent color – from original, to its display on a monitor, through to its reproduction in print – is considerable. The variety of devices and materials used to capture, display and reproduce color in print all have a profound effect on the end result.

Color management ensures consistent colors

Bit depth and mode

As discussed earlier in 'Levels' each pixel in a single channel of a standard RGB image is described in one of a possible 256 tones or levels. The computer memory required to calculate and store this color data is '8 bits', a bit (binary digit) being the basic unit of the computer's memory. The amount of bits dedicated to describing and recording tonal or color variations is called the '**bit depth**'. If only tonal information is required (no color) a single channel 8-bit image is sufficient to create a good quality black and white image, reproducing all of the tonal variations needed to produce 'continuous tone'. An 8-bit image that handles only tonal variations is more commonly referred to as a **Grayscale** image.

When 8 bits are needed for each of the three channels of an RGB image this results in what is often referred to as a 24-bit image (3 × 8). Photoshop, however, does not refer to an RGB image as a 24-bit image but rather as an RGB Color and lists the bit depth of each channel rather than the entire image, e.g. 8 Bits/Channel. Images with a higher '**bit depth**' have a greater potential for color or tonal accuracy although this sometimes cannot be viewed because of the limitations of the output device. Images with a higher bit depth, however, require more data or memory to be stored in the image file (Grayscale images are a third of the size of RGB images with the same pixel dimensions and print size). Photoshop offers support for 16 Bits/Channel and 32 Bits/Channel images.

RGB image, 256 levels per channel (24-bit)

256 levels (8-bit)

Capturing and editing at bit depths exceeding 8 bits per channel

Sophisticated 'prosumer' (point and shoot) digital cameras and digital SLRs (DSLRs) are able to export files in the 'RAW' format in bit depths higher than 8 bits per channel to the computer. Higher quality scanners are able to scan and export files at 16 bits per channel (48-bit). In Photoshop it is possible to edit an image using 32, 16 or 8 bits per channel. The size of the file (megabytes rather than pixel dimensions) doubles each time the bit depth is doubled. 16 and 32 Bits/Channel image editing is used by professionals for high-quality or specialized image editing. When extensive tonal or color corrections are required it is recommended to work in 16 bits per channel whenever possible. It is, however, important to note that not all of the Adobe editing tools function in 16-bit mode (even less in 32 Bits/Channel).

File size

Digital images are data hungry (this data being required to record the extensive variations in color and/or tone of the original image or subject). The simple binary language of computers and the visual complexities of a photographic image lead to large '**file sizes**'. This data can require large amounts of computer memory to display, print or store the image. The text file for this entire book would only be a small fraction of the memory required for the cover image (10 megabytes).

Units of memory

8 bits	=	1 byte
1024 bytes	=	1 kilobyte
1024 kilobytes	=	1 megabyte
1024 megabytes	=	1 gigabyte

Storage capacity of disks and drives

Flash/USB pen	=	64 megabytes–2 gigabyte
CD	=	700–800 megabytes
DVD	=	4–9 gigabytes
iPod	=	512 megabytes to 80 gigabytes

Fortunately files can be '**compressed**' (reduced in memory size) when closing the file for storage or uploading over the Internet. Portable hard drives (such as Apple's 'iPod' or the smaller 'USB' or 'Flash' drives) are now commonly used for storing and transferring large image files conveniently and quickly. A 10-megapixel digital image can be saved as a 15-megabyte RAW file or a 1-megabyte JPEG file using a high-quality compression setting. The same file opens up to a 27.2-megabyte file in Photoshop. When talking about file size it helps to know whether you are talking about an open or closed file and whether any image compression has been used.

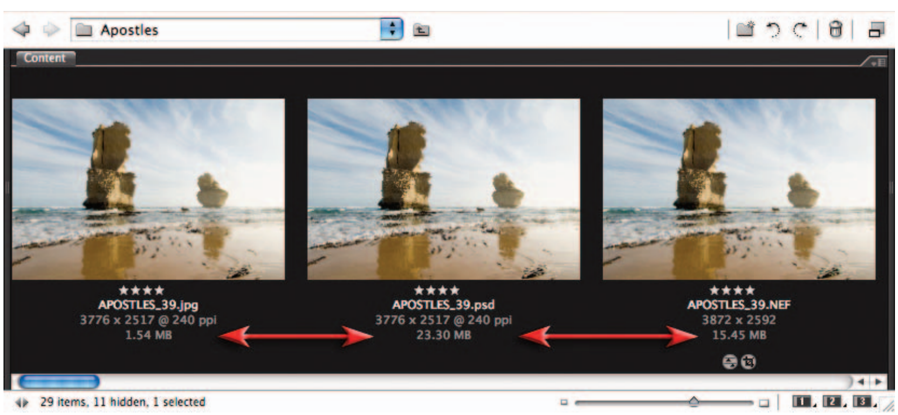

Same image – different file formats (viewed in Bridge with file size preference checked)

Note > If you are using Bridge (see the following chapter) it is possible to gain information about image size (megabytes) and pixel dimensions from files that have not been opened, either directly underneath the image thumbnail or in the 'Metadata' tab. When an image that has been compressed is opened in Photoshop the file size in megabytes will be larger but the pixel dimensions will remain the same.

File formats

When an image is captured by a camera or scanning device it has to be '**saved**' or memorized in a '**file format**'. If the binary information is seen as the communication, the file format can be likened to the language or vehicle for this communication. The information can only be read and understood if the software recognizes the format. Images can be stored in numerous different formats. The four dominant formats in most common usage are:

- Raw (.dng) – Camera Raw and Digital Negative
- JPEG (.jpg, jpf and jpx) – Joint Photographic Experts Group
- TIFF (.tif) – Tagged Image File Format
- Photoshop (.psd) – Photoshop Document

Camera RAW and Digital Negatives – Unlike the other file formats, RAW is not an acronym for a much longer name. Selecting the RAW format in the camera instead of JPEG or TIFF stops the camera from processing the color data collected from the sensor. The RAW data is what the sensor 'saw' before the camera processes the image, and many photographers have started to refer to this file as the 'digital negative'. The unprocessed RAW has to be converted into a usable image file format using image-editing software supplied by the camera manufacturer or built into software packages such as Adobe Photoshop and can be compressed and archived as a 'Digital Negative'.

A close-up detail of an image file that has been compressed using maximum image quality in the JPEG options box

A close-up detail of an image that has been compressed using low image quality in the JPEG options box. Notice the artifacts that appear as blocks

JPEG (Joint Photographic Experts Group) – Industry standard for compressing continuous tone photographic images destined for the World Wide Web (www) or for storage when space is limited. JPEG compression uses a 'lossy compression' (image data and quality are sacrificed for smaller file sizes when the image files are closed). The user is able to control the amount of compression. A high level of compression leads to a lower quality image and a smaller file size. A low level of compression results in a higher quality image but a larger file size.

Note > It is recommended that you only use the JPEG file format after you have completed your image editing and always keep a master Photoshop document for archival purposes.

Format	Compression	Color modes	Layers	Transparency	Uses
RAW	No	Unprocessed	No	No	Master file
JPEG	Yes	RGB, CMYK, Grayscale	No	No	Internet and camera format (compressed)
JPEG2000	Yes	RGB, CMYK, Grayscale	No	No	Internet and archival
Photoshop	No	RGB, CMYK, Grayscale, Indexed color	Yes	Yes	Master file (modified)
TIFF	Yes	RGB, CMYK, Grayscale	Yes	Yes	Commercial printing and generic camera format (lossless)
GIF	Yes	Indexed color	No	Yes	Internet graphics and animations
DNG	Yes	Unprocessed	No	No	Archival format for storing original RAW and metadata

JPEG2000 – This version of the JPEG format supports 16 Bits/Channel and alpha channels and produces less image artifacts than the standard JPEG compression but uses a more complex list of saving options than the standard JPEG format. Photoshop CS3 supports the file format but it is not available as part of the 'Save for Web' options.

PSD (Photoshop Document) – This is the default format used by the Adobe image-editing software. A Photoshop document is usually kept as the master file from which all other files are produced depending on the requirements of the output device. The PSB format is another version of PSD and is designed specifically for creating documents larger than 2GB.

TIFF (Tagged Image File Format) – This has been the industry standard for images destined for publishing (magazines and books, etc.). TIFF uses a 'lossless' compression (no loss of image data or quality) called '**LZW compression**'. Although preserving the quality of the image, LZW compression is only capable of compressing images by a small amount. TIFF files now support layers and transparency that can be read by other Adobe software products such as InDesign.

GIF (Graphics Interchange Format) – This format is used for logos and images with a small number of colors and is very popular with web professionals. It is capable of storing up to 256 colors, animation and areas of transparency. It is not generally used for photographic images.

DNG (Digital Negative Format) – The DNG format is a new archival file format that stores both the RAW picture data as well as the metadata saved by the camera at the time of shooting.

TOP TIPS for cross-platform saving

Many work and education environments contain a mix of Windows and Macintosh machines. Though both systems are far better at reading each other's files than they used to be, there are still occasions when you will have trouble when sharing files between the two platforms. Use these tips to ensure that work that you save is available for use in both environments.

1. Make sure that you always append your file names.
This means add the three-letter abbreviation of the file format you are using after the name. So if you were saving a file named 'Image1' as a TIFF the saved file would be 'Image1.tif', a JPEG version would be 'Image1.jpg' and a Photoshop file would be 'Image1.psd'. Macintosh Photoshop users can force the program to 'Always Append' by selecting this option in the 'Saving Files' section of Preferences.

2. Save TIFF files in the IBM version.
When saving TIFF files you are prompted to choose which platform you prefer to work with; choose IBM if you want to share files. Macintosh machines can generally read IBM (Windows) TIFFs, but the same is not true the other way around.

3. Macintosh users save images to be shared on Windows formatted disks.
If you are sharing images on a portable storage disk such as a Zip drive always use media that are formatted for Windows. Macintosh drives can usually read the Windows disks but Windows machines can't read the Macintosh versions.

4. Try to keep file names to eight characters or less and don't use spaces.
Older Windows machines and some web servers have difficulty reading file names longer than eight characters. So just in case you happen to be trying to share with a cantankerous old machine get into the habit of using short names – and always appended of course. Use a hyphen or underscore instead of a space and use lower case characters only (no capitals) if the images are destined for a web server that likes to say 'no'.

Image compression

Imaging files are huge. This is especially noticeable when you compare them with other digital files such as those used for word processing. A text document that is 100 pages long can easily be less than 1% the size of a file that contains a single 8 × 10 inch digital photograph. With files this large it soon became obvious to the industry that some form of compression was needed to help alleviate the need for us photographers to be continuously buying bigger and bigger hard drives.

What emerged was two different ways to compress pictures. Each enables you to squeeze large image files into smaller spaces but one system does this with no loss of picture quality – *lossless* compression – whereas the other enables greater space savings with the price of losing some of your image's detail – *lossy* compression.

What is compression?

All digital picture files store information about the color, brightness and position of the pixels that make up the image. Compression systems reorder and rationalize the way in which this information is stored. The result is a file that is optimized and therefore reduced in size. Large space savings can be made by identifying patterns of color, texture and brightness within images and storing these patterns once, and then simply referencing them for the rest of the image. This pattern recognition and file optimization is known as compression.

The compression and decompression process, or CODEC, contains three stages:

1. The original image is compressed using an algorithm to optimize the file.
2. This version of the file becomes the one that is stored on your hard drive or web site.
3. The compressed file is decompressed ready for viewing or editing.

If the decompressed file is exactly the same as the original after decompression, then the process is called 'lossless'. If some image information is lost along the way then it is said to be 'lossy'. Lossless systems typically can reduce files to about 60% of their original size, whereas lossy compression can reduce images to less than 1%.

There is no doubt that if you want to save space and maintain the absolute quality of the image then the only choice is the lossless system. A good example of this would be photographers, or illustrators, archiving original pictures. The integrity of the image in this circumstance is more important than the extra space it takes to store it.

On the other hand (no matter how much it goes against the grain), sometimes the circumstances dictate the need for smaller file sizes even if some image quality is lost along the way. Initially you might think that any system that degrades the image is not worth using, and in most circumstances, I would have to agree with you. But sometimes the image quality and the file size have to be balanced. In the case of images on the web they need to be incredibly small so that they can be transmitted quickly over slow telephone lines. Here some loss in quality is preferable to images that take 4 or 5 minutes to appear on the page. This said, I always store images in a lossless format on my own computer and only use a lossy format when it is absolutely crucial to do so.

Format	Compression amount	Original file size	Compressed file size	Lossy/ Lossless
JPEG	Minimum		0.14MB	Lossy
JPEG	Maximum		2.86MB	Lossy
JPEG2000	Minimum		0.07MB	Lossy
JPEG2000	Maximum		2.54MB	Lossy
JPEG2000	–	20.0MB (PSD file)	5.40MB	Lossless
TIFF	LZW		10.30MB	Lossless
TIFF	Zip		10.10MB	Lossless
TIFF	None		27.30MB	Lossless
PSD	–		20.00MB	Lossless

Comparing the compression abilities of different file types and settings using the same 20MB base file

How lossy is lossy?

The term lossy means that some of the image's quality is lost in the compression process. The amount and type of compression used determines the look of the end result. Standard JPEG and JPEG2000 display different types of 'artifacts' or areas where compression is apparent. The level of acceptable artifacts and practical file sizes will depend on the required outcome for the picture.

To help ensure that you have the best balance of file size and image quality make sure that you:

- Use the Save for Web or Save As > JPEG2000 features – both contain a post-compression preview option.
- Always examine the compressed image at a magnification of 100% or greater so that unacceptable artifacts will be obvious.

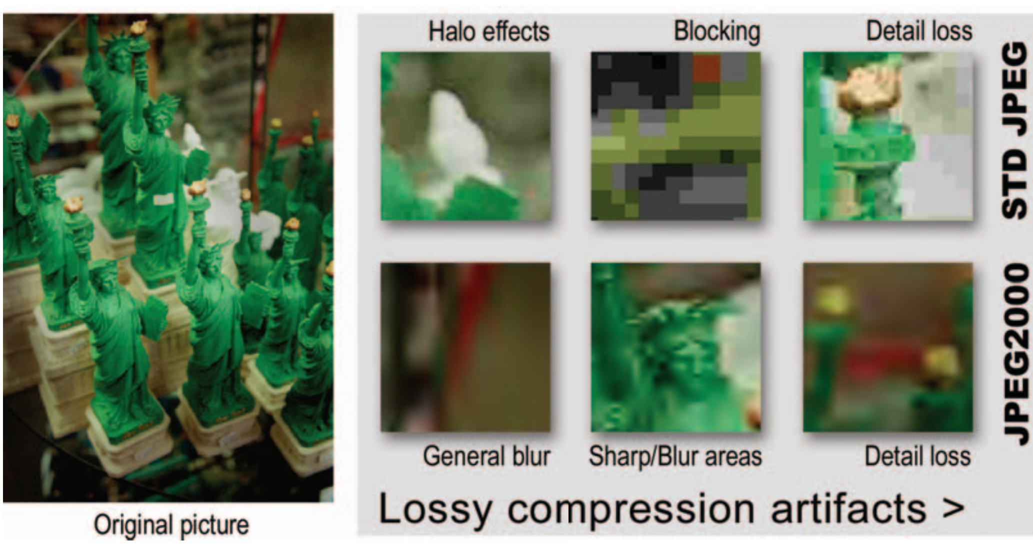

Typical artifacts resulting from applying too much compression to a photo

Resolution

Resolution is a term that is used to specify the size of a pixel, a dot of colored light on a monitor or a dot of ink on the printed page. There are usually two resolutions at play at any one time – the resolution of the digital file and that of the output device. We can talk about capture size, image resolution, monitor resolution and printer resolution. They are all different, but they all come into play when handling a single digital image that is to be printed. Various resolutions can be quoted as we move through the chain of processes involved in creating a digital print (in the example below the total number of pixels remains constant throughout the chain of events).

Summary

An image captured at a resolution greater than 3000ppi is displayed at 100ppi on a high-resolution monitor. Using Photoshop, the image resolution is lowered to 256ppi (the pixel dimensions remain the same). The image is then printed using an inkjet printer with a printer resolution of 1440dpi (dots per inch). The different resolutions associated with this chain of events are:

Capture size > Display resolution > Image resolution > Output device resolution

Image sensor

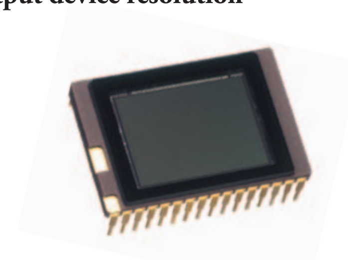

The sensor to the right creates an image file with 5 million pixels or 5 megapixels (2560 × 1920 pixels). The resolution assigned to the image file by the capture device may be a print or monitor resolution. Either way it has no bearing on the file size, which is determined by the total number of pixels.

Digital file displayed on screen

The monitor resolution (the size of its display pixels, e.g. 1024 × 768) is defined by its resolution setting (approximately 100 pixels for every linear inch or 10,000 pixels for every square inch in a high definition TFT display). The image pixels (different than the display pixels) can be viewed in a variety of sizes by zooming in and out of the image using image-editing software.

Digital file adjusted in Photoshop

The resolution of the digital file is adjusted to 256 pixels per inch (ppi). Each pixel is allocated a size of 1/256th of an inch. Because the digital file is 2560 pixels wide this will create a print that is 10 inches wide if printed (256 × 10 = 2560).

Note > Increasing the document size further will start to lower the resolution below an acceptable level (the pixels will become large enough to see with the naked eye).

The image is printed

The image is printed on an inkjet printer using a printing resolution of 1440 dots per inch. Many colored dots of ink are used to render a single image pixel (see 'Dpi and ppi').

Understanding resolution

Resolution is perhaps the most important, and the most confusing, subject in digital imaging. It is important because it is linked to quality. It is confusing because the term '**resolution**' is used to describe at what quality the image is captured, displayed or output through various devices.

10 pixels per inch

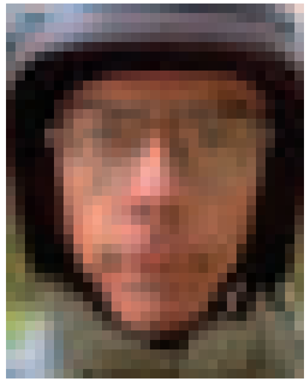

20 pixels per inch

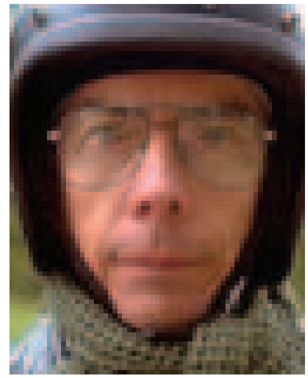

40 pixels per inch

Resolution determines image quality and size

Increasing the total number of pixels in an image at the capture or scanning stage increases both the quality of the image and its file size. It is '**resolution**' that determines how large or small the pixels appear in the final printed image. The greater the image resolution the smaller the pixels, and the greater the apparent sharpness of the final image. Resolution is stated in '**pixels per inch**' or '**ppi**'.

Note > With the USA dominating digital photography, measurements in inches rather than centimeters are commonly used – 1 inch equals exactly 2.54 centimeters.

The images to the right have the same pixel dimensions (300 × 300) but different resolutions. The large image has a resolution half that of the small one. A digital image can be made to appear bigger or smaller without changing the total number of pixels, e.g. a small print or a big poster. This is because a pixel has no fixed size. The pixel size can be modified by the image-editing software to change the document size. Increasing the resolution of the image file decreases the size of the pixels and therefore the output size of the file.

Note > When talking about the 'size' of a digital image it is important to clarify whether it is the pixel dimensions or the document size (measured in inches or centimeters) that are being referred to.

Dpi and ppi

If manufacturers of software and hardware were to agree that dots were round and pixels were square it might help users differentiate between the various resolutions that are often quoted. If this was the case the resolution of a digital image file would always be quoted in 'pixels per inch', but this is not the case.

At the scanning stage some manufacturers use the term dpi instead of 'ppi'. When scanning, 'ppi' and 'dpi' are essentially the same and the terms are interchangeable, e.g. if you scan at 300dpi you get an image that is 300ppi.

When working in Photoshop image resolution is always stated in ppi. You will usually only encounter dpi again when discussing the monitor or printer resolution. The resolutions used to capture, display or print the image are usually different to the image resolution itself.

Note > Just in case you thought this differentiation between ppi and dpi is entirely logical – it isn't. The industry uses the two terms to describe resolution, 'pixels per inch' (ppi) and 'dots per inch' (dpi), indiscriminately. Sometimes even the manufacturers of the software and hardware can't make up their minds which of the two they should be using, e.g. Adobe refer to image resolution as ppi in Photoshop and dpi in InDesign – such is the non-standardized nomenclature that remains in digital imaging.

File size and resolution

When we use the measurement 'ppi' or 'pixels per inch' we are referring to a linear inch, not a square inch (ignore the surface area and look at the length).

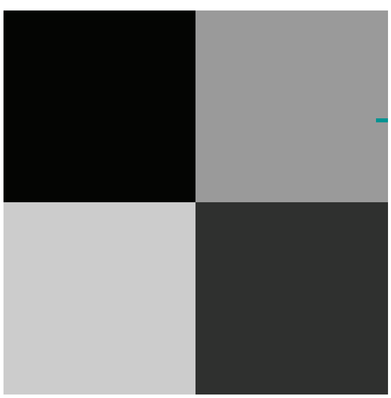

2 inch × 2inch file @ 1ppi = 4 pixels

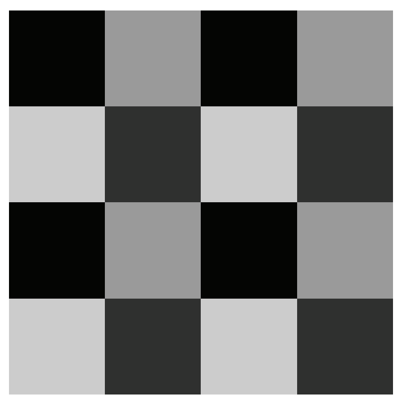

2 inch × 2 inch file @ 2ppi = 16 pixels

File size, however, is directly linked to the **total** number of pixels covering the entire surface area of the digital image. Doubling the image output dimensions or image resolution quadruples the total number of pixels and the associated file size in kilobytes or megabytes.

Note > Handling files with excessive pixel dimensions for your output needs will slow down every stage of your digital image process, including scanning, saving, opening, editing and printing. Extra pixels above and beyond what your output device needs will not lead to extra quality. Quality is limited or 'capped' by the capability of the output device.

Calculating a suitable file size and scanning resolution

Scanning resolution is rarely the same as the resolution you require to print out your image. If you are going to create a print larger than the original you are scanning, the scanning resolution will be greater than the output resolution, e.g. a 35mm negative would have to be scanned at 1200ppi if a 6 × 4 inch commerical print is required. If the print you require is smaller than the original, the scanning resolution will be smaller than the output resolution.

The smaller the original the higher the scanning resolution.

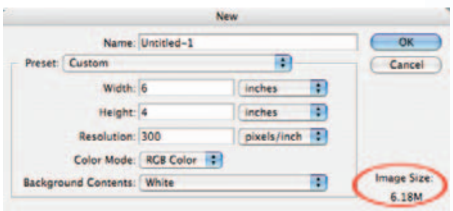

Magnification × output resolution = scanning resolution
scanning resolution = 4 × 300ppi = 1200ppi

To calculate the correct file size and scanning resolution for the job in hand you can:

Either: Go to 'File > New' in Photoshop, type in the document size, resolution and mode you require and then make a note of the number of megabytes you require from the scanning process. Then adjust the scanning software resolution until the required number of megabytes is captured.

Or: Multiply the magnification factor (original size to output size) by the output resolution (as dictated by the output device) to find the scanning resolution (not so difficult as it sounds!).

Size and mode	Output device resolution		
	100ppi screen	240ppi inkjet	300ppi commercial
8 × 10 RGB	2.29MB	13.20MB	20.60MB
8 × 10 Grayscale	781K	4.39MB	6.87MB
5 × 7 RGB	1.00MB	5.77MB	9.01MB
5 × 7 Grayscale	342K	1.92MB	3.00MB
4 × 6 RGB	703K	3.96MB	6.18MB
4 × 6 Grayscale	234K	1.32MB	2.06MB

File size

Image size

Before you adjust the size of the image you have to know how to determine the size you need. Six- and eight-megapixel digital cameras are currently the affordable 'end' of professional digital capture. The image resolution produced by these digital cameras is not directly comparable to 35mm film capture but the images produced can satisfy most of the requirements associated with professional 35mm image capture. DSLRs using full frame sensors can match medium format film cameras for quality.

Six megapixel cameras capture images with pixel dimensions of around 3000 × 2000 (6 million pixels or 6 'megapixels'). The resulting file size of around 17MB (1 megapixel translates to nearly 3 megabytes of data) is suitable for an image in a commercial magazine that would nearly fill the page. Ten-megapixel cameras are capable of producing files that can be used to illustrate double-page spreads in magazines with just a small amount of resampling (see Interpolation).

Useful specifications to remember

- Typical standard-resolution monitor: 1024 × 768 pixels
- Typical full-page magazine illustration: 3400 × 2500 (8.5 million pixels)
- High-resolution TFT monitor: 100ppi
- High quality inkjet print: 240ppi
- Magazine quality printing requirements: 300ppi
- Full-screen image: 2.25MB (1024 × 768)
- Postcard-sized inkjet print: 4MB
- 10 × 8 inch inkjet print: 13.2MB
- Full-page magazine image at commercial resolution: 20MB

Note > Remember to double the above file sizes if you intend to edit in 16-bit per channel mode.

A 20MB file will usually suffice if you are not sure of the intended use of the digital file. Thirty-five-millimeter film scanned with a scanning resolution of **2300** will produce a 20.3MB file (2173 pixels × 3260 pixels).

Pixel dimensions, document size and resolution

Before retouching and enhancement takes place, determine if the '**image size**' needs to be scaled for the intended output (the capture resolution will probably require changing to output resolution). This will ensure that optimum image quality and computer operating speed are maintained. To control image size go to '**Image > Image Size**' in Photoshop.
Image size is described in three ways:

- Pixel dimensions (the number of pixels determines the file size in terms of kilobytes).
- Print size (output dimensions in inches or centimeters).
- Resolution (measured in pixels per inch or ppi).

If one is altered it will affect or impact on one or both of the others, e.g. increasing the print size must either lower the resolution or increase the pixel dimensions and file size. The image size is usually changed for the following reasons:

- Resolution is changed to match the requirements of the print output device.
- Print output dimensions are changed to match display requirements.

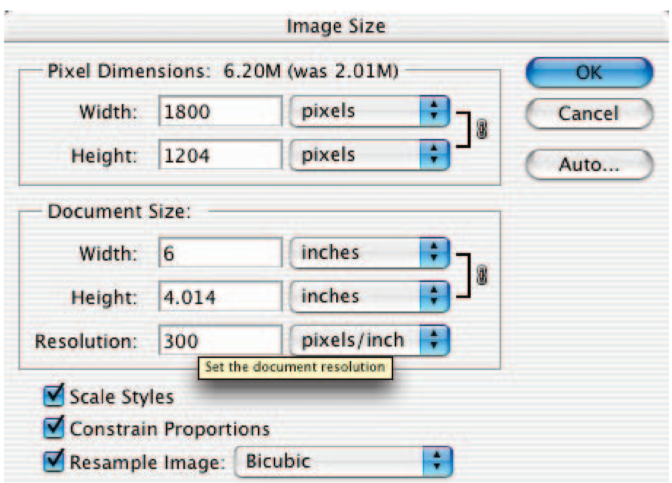

Image size options

When changing an image's size a decision can be made to retain the proportions of the image and/or the pixel dimensions. These are controlled by the following:

- If '**Constrain Proportions**' is selected the proportional dimensions between image width and image height are linked. If either one is altered the other is adjusted automatically and proportionally. If this is not selected the width or height can be adjusted independently of the other and may lead to a distorted image.

- If '**Resample Image**' is selected (use with caution, see 'Resampling') adjusting the dimensions or resolution of the image will allow the file size to be increased or decreased to accommodate the changes. Pixels are either removed or added. If deselected the print size and resolution are linked to prevent resampling. Changing width, height or resolution will change the other two. Pixel dimensions and file size remain constant.

Resampling

An image is ‘resampled’ when its pixel dimensions (and resulting file size) are changed. It is possible to change the output size or resolution without affecting the pixel dimensions (see ‘Understanding resolution’). Resampling usually takes place when the pixel dimensions of the original capture or scan do not precisely match the requirements for output (size and resolution). Downsampling decreases the number of pixels and information is deleted from the image. Increasing the total number of pixels or resampling up requires new pixels to be added to the file. The new pixels use information based on color values of the existing pixels in the image file.

Image scanned at correct resolution

Effects of excessive resampling up

Excessive resampling up can result in poor image quality (the image will start to appear blurry). Avoid the need for resampling to enlarge the file size, if at all possible, by capturing at a high enough resolution or by limiting the output size. If you have to resample due to the limitations of your capture device (not enough megapixels) then files that resample best are those that have been captured from a digital camera using a low ISO setting and have not previously been sharpened.

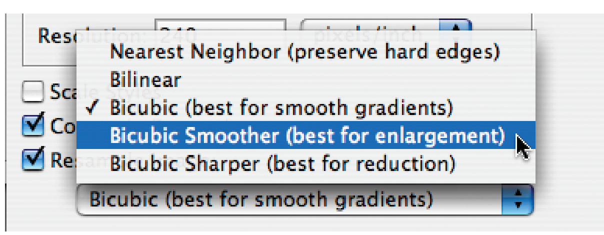

When resampling up an image to create a larger file, choose the ‘Bicubic Smoother’ from the Interpolation options in the Image Size dialog box. Use ‘Bicubic Sharper’ when decreasing the size of the file. Bilinear and Nearest Neighbor are used for hard-edged graphics and are rarely used for the interpolation of photographs. If you have already sharpened your image prior to resampling (best avoided) you will need to reapply the Unsharp Mask or Smart Sharpen to regain the sharp quality of the image.

Interpolation

Well it seems that in recent years a small revolution in refinement has been happening in the area of interpolation technologies. The algorithms and processes used to apply them have been continuously increasing in quality until now they are at such a point that the old adages such as

Sensor dimension/output resolution = maximum print size

don't always apply. Using either software or hardware versions of the latest algorithms it is now possible to take comparatively small files and produce truly large prints of good quality.

Resampling techniques

Bicubic – All resampling techniques in Photoshop use the best interpolation settings of Bicubic, Bicubic Sharper or Bicubic Smoother in conjunction with the Image > Image Size feature. The standard approach uses a 4 × 4 sampling scheme of the original pixels as a way of generating new image data. With the Resample and Constrain Properties options selected, the new picture dimensions are entered into the 'width' and 'height' areas of the dialog. Clicking OK will then increase the number of pixels in the original.
Use the Smart Sharpen or Unsharp Mask filter after resampling rather than before and restrict the amount of resampling that is performed on a single image. If the software allows the user to crop, resize and rotate the image at the same time, this function should be utilized whenever possible.

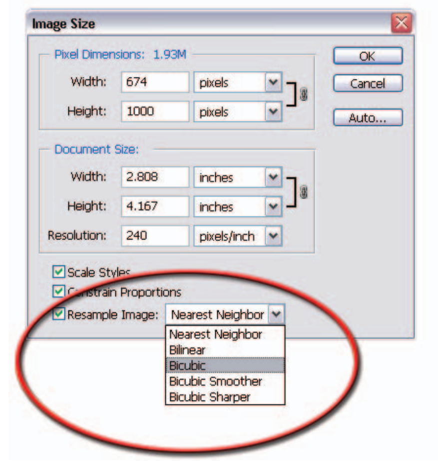

Bicubic via LAB – In this technique the mode of the picture is changed from RGB to LAB first using the Image > Mode > Lab command. The two color channels (A and B) are then hidden in the channels palette and the bicubic interpolation is applied to just the lightness (L) channel. The color channels are then switched back on. The theory behind this approach is that by only interpolating the lightness channel the enlarged image will suffer less deterioration overall.

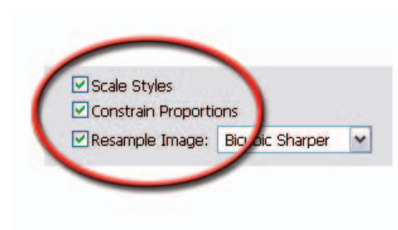

Stair interpolation – There is a growing school of thought that states that increasing the size of an image by several smaller steps will produce a sharper, more detailed result than making the upscale enlargement in a single jump. Most professionals who use this approach increase the size of their images by 10% each time until they reach the desired pixel dimensions.

The advances in the algorithms and procedures used to create large images have dramatically improved over the last few years. I still cringe saying it, but it is now possible to break the 'I must never interpolate my images rule' in order to produce more print area for the pixels I have available.

Resampling guidelines:

When resampling keep in mind the following guidelines for ensuring the best results:

1. Images captured with the correct number of pixels for the required print job will always produce better results than those that have been interpolated.
2. The softening effect that results from high levels of interpolation is less noticeable in general, landscape or portrait images and more apparent in images with sharp-edged elements.
3. The more detail and pixels in the original file the better the interpolated results will be.
4. A well-exposed, sharply focused original file that is saved in a lossless format such as TIFF is the best candidate image for upsizing.

What is interpolation anyway?

Interpolation is a process by which extra pixels are generated in a low-resolution image so that it can be printed at a larger size than its original dimensions would normally allow. Interpolation, or as it is sometimes called, upsizing, can be implemented via software products such as the Image Size > Resample option in Photoshop or by using the resize options in the printer's hardware.

Both approaches work by sampling a group of pixels (a 4 × 4 matrix in the case of bicubic interpolation) and using this information together with a special algorithm as a basis for generating the values for newly created pixels that will be added to the image. The sophistication of the algorithm and the size of the sampling set determine the quality of the interpolated results.

Printer-based resampling

An alternative approach to using Photoshop to resample your picture is to make use of the 'scaling' options in your printer driver dialog. Most desktop and high-end laboratory digital printers are capable of the interpolation necessary to produce big prints.

When outsourcing to a professional printer the file is kept to its original pixel dimensions and the resolution is reduced so that the print size will equal the required print dimensions. The digital printer is then instructed to interpolate the file as it was printing to its optimum resolution of the machine. On the desktop a similar process is used with the digital photographer selecting a scale value, or an option to match the image size to the paper dimensions, in the print or print driver dialog. Letting the output device perform the interpolation of the image has the following advantages: the process does not require the photographer to change the original file in any way and it removes an extra processing step from the file preparation process.

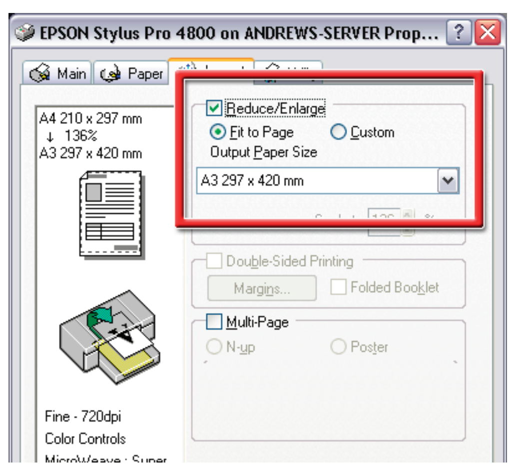

5.5 inches x 8 inches @ 300ppi – CMYK file – 15.1 megabytes (converted from an 11.3M RGB file)

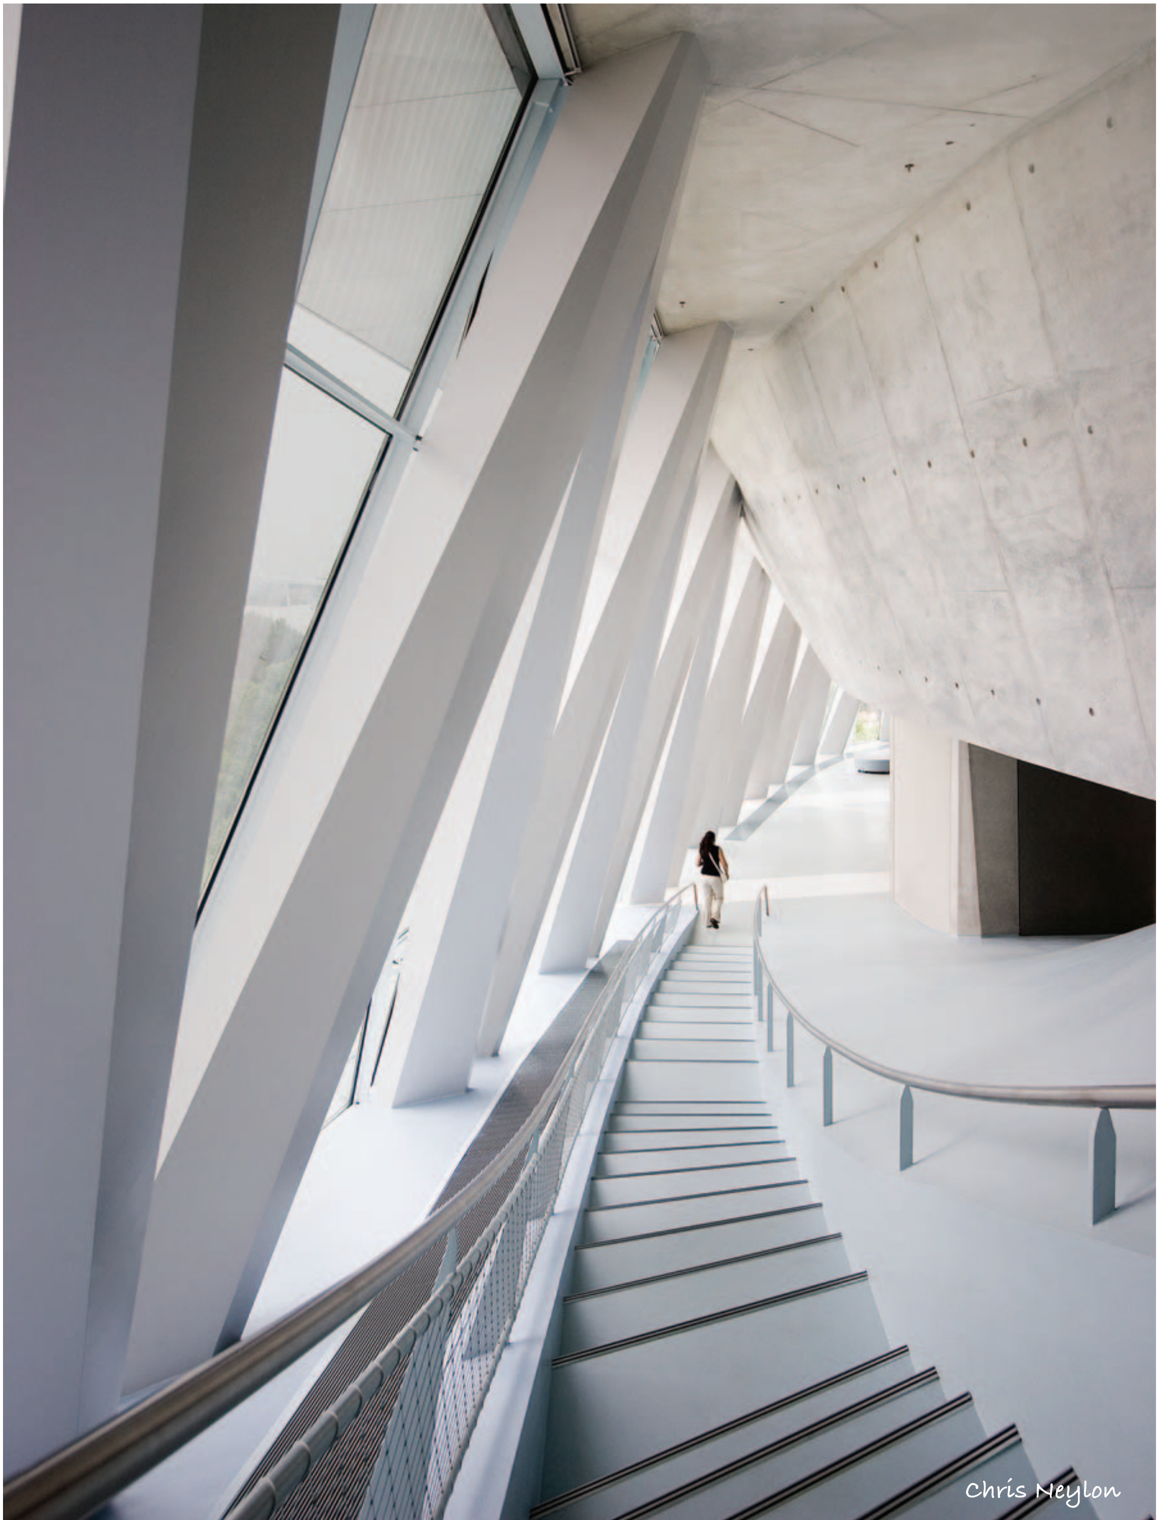

6 inches x 8 inches @ 300ppi - CMYK file - 16.5 Megabytes (converted from an 12.4M RGB file)