

Introduction to Digital Resolution

2011

Copyright Les Walkling 2011

Adobe Photoshop screen shots reprinted with permission
from Adobe Systems Incorporated.

Version 2011:02


CONTENTS

• Pixels of Resolution and Levels of Tonality	3
• Digital Resolution: How many Pixels do I need?	4
• Bit Depth: How many Levels do I need?	5
• Levels of Brightness	6
• Camera Resolution	7
• Image Resolution	8
• Monitor Resolution	9
• Colour Gamuts and RGB Profiles	10
• Histograms	11

Pixels of Resolution and Levels of Tonality

Our everyday perception of reality leads us to interact with the world as though it were continuous. A magnifying lens for example reveals more detail than can be readily seen with the unaided eye. Our experience of traditional analogue photography has reinforced this perception, and is why photography is known as a 'continuous tone' medium as opposed to a 'half tone' medium such as the offset printing of images in books, magazines and newspapers that rearrange the original photograph's 'continuousness' into a sea of half tone dots of ink on paper.

While not the same as a printing press, digital photography is also a half tone medium because it reduces the scene and its 'continuous resolution' into steps of resolution that we call Pixels and its 'continuous tone' into steps of brightness called Levels.



The smaller the intervals of resolution (pixelation) and brightness (posterization) that an image is digitised into, the more its appearance will maintain the illusion of continuousness. The world can be digitised by photographing it directly with a digital camera, or by scanning photographic negatives, slides or prints. But once digitised the image is pixelated and posterized and it will remain that way. Therefore choosing the correct image pixelation and posterization is critical for maintaining the illusion of analogue photography.

This raises two important questions:

How many Pixels do I need?

How many Levels do I need?

The type of digital camera and its settings control these output values.

Digital Resolution: How many Pixels do I need?

The sensor in a digital camera is rated in terms of the number of pixels it divides the image into. A one megapixel (mpx) camera therefore will capture approximately one million pixels which when processed into an RGB file will produce a 3 megabyte (mb) image (1mb Red + 1mb Green + 1mb Blue = 3mb RGB). Canon's EOS DSLR range includes:

1Ds Mark III	21.1	Megapixels	5.0	fps	5616 x 3744 px	63.3 mb RGB
1D Mark III	10.1	Megapixels	10.0	fps	3888 x 2592 px	30.3 mb RGB
5D Mark II	21.0	Megapixels	3.9	fps	5616 x 3744 px	63.0 mb RGB
50D	15.1	Megapixels	6.3	fps	4752 x 3168 px	45.3 mb RGB
450D	12.2	Megapixels	3.5	fps	4272 x 2848 px	36.6 mb RGB

Internet/Monitor Distribution

The internet only requires very low resolution, 72 pixels per inch (ppi) images, because that is the typical resolution of the low resolution monitors that make up most of the internet. For example, a 6" x 4" (15cm x 10cm) image @ 72ppi = 364.5kb which is approximately a third of a megabyte (mb).

Print Distribution

When printing digital images different file sizes are required for optimum quality on various printers. Typical examples are listed below. Some printers, such as the Lambda digital photographic printer found in many professional labs will accept images at either 200 ppi or 400 ppi, the later only required for the highest resolution printing of black text on white backgrounds.

Noritsu Photo MiniLab

6" x 4" @ 400ppi print = 11.0MB RGB
8" x 10" @ 400ppi print = 36.6MB RGB

Epson R2400 Inkjet Printer

6" x 4" @ 360ppi print = 8.9MB RGB
8" x 10" @ 360ppi print = 29.7MB RGB
16" x 20" @ 360ppi print = 118.7MB RGB

Kodak Pegasus LED Photo Printer

8" x 10" @ 250ppi print = 14.3MB RGB
20" x 32" @ 250ppi print = 114.4MB RGB

Lambda Digital Photo Printer


8" x 10" @ 200ppi print = 9.16MB RGB
40" x 60" @ 200ppi print = 274.7MB RGB

More pixels will sometimes be required than a camera can produce. In this case the files will need to be resized and/or printed at a lower than optimum resolution. Most digital camera files can usually be doubled in file size with minimal loss of quality, especially in RAW mode. Greater enlargements must be made with care and/or the resulting images need to be viewed at a greater distance to avoid an obvious decrease in image quality, such as lowered sharpness (not as crisp) and rougher tones (not as smooth and continuous). In practice a digital file can often be considerably enlarged and still appear acceptable.


Bit Depth: How many Levels do I need?

Digital photography requires, at a minimum an 8 bit per channel environment where 256 discrete levels or steps of tone ($2^8 = 256$) describe the entire tonal range from black to white. If the steps of tone are small enough increments we don't notice them as discrete steps or jumps of tone and the digital image will then preserve the illusion of smooth continuous colour, and under normal viewing conditions will be indistinguishable from an original continuous tone 'photograph'.

A 24-bit colour image = three 8-bit channels. One 8-bit channel describes the image's RED values, the second describes its GREEN values, and the third describes its BLUE values.



Most monitors/printers are also designed to work with 8-bits per channel. In order to achieve an optimum 24-bit RGB file, digital cameras record the scene in more than 8 bits per channel. The higher the camera's bit depth the greater number of levels of tone the image is captured in. A higher bit depth results in smoother tonal transitions and less obvious posterization. A Canon 1Ds Mk III for example captures in 14-bits (16384 Levels) but when processed in-camera reduces the bit depth to produce an optimum 8-bits per channel. The aim of 'normal processing' is an image whose posterization will not be visible when viewed on a monitor or in a print. For example:



4 bits per channel (16 levels)



8 bits per channel (256 levels)

$2^{24} = 256 \text{ levels RED} \times 256 \text{ levels GREEN} \times 256 \text{ levels BLUE} = 16.77 \text{ million colours.}$

Levels of Brightness

The simplest digital image is a 1 bit image (2^1) which divides the original image into only two levels of brightness; black and white. The bit depth determines the maximum number of levels or steps of brightness from black to white in the image. For example:

1 bit image = $2^1 = 2$ levels of brightness (= black or white)

2 bit image = $2^2 = 4$ levels of brightness

3 bit image = $2^3 = 8$ levels of brightness


4 bit image = $2^4 = 16$ levels of brightness

5 bit image = $2^5 = 32$ levels of brightness

6 bit image = $2^6 = 64$ levels of brightness

7 bit image = $2^7 = 128$ levels of brightness

8 bit image = $2^8 = 256$ levels of brightness (= the illusion of continuous tone)




Increasing levels of brightness decreases levels of visible posterization

A digital image is a posterized image. Therefore for the highest possible image quality, in terms of maintaining the illusion of continuous tone and colour, you want to end with, rather than start with, 8-bits per channel. To do this you have to digitise (scanner/camera) your images into more than 8-bits per channel, so that after image editing you still have enough bit depth left to maintain this illusion. For example a 16-bit workflow (2^{16}) results in 65,535 levels of tone per channel which can support extensive editing without visible posterization. Photoshop offers extensive support for 16-bit images, though its controls still only edit with 8-bit resolution and its 16-bit Info Palette only displays truncated 16-bit values (0 to 32,768). Another issue is that 16-bit resolution doubles the file size which increases storage requirements and significantly slows down all pixel editing operations.

Camera Resolution

Digital Single Lens Reflex cameras replace film with a grid or array of photo-sites (pixels). The scene projected by the lens is digitised according to the number of photo-sites in the array. For example a 6 megapixel (mpx) camera will divide the image into 6 million pixels per channel. A 6 mpx camera will therefore result in an 18 megabyte (mb) RGB image.


Each photo-site can only record the amount (or level) of light, not its colour. In order to produce a full colour image, each pixel in the array must represent its corresponding point in the scene according to its original red, green and blue brightness.




Bayer Filter Array

But each photo-site is covered by only a red or green or blue filter that measures only the red or green or blue brightness of the light falling on it. A typical array also has twice as many green as red or blue filtered photo-sites. While different manufacturers arrange the pixel array and filtration in different configurations according to their propriety design, the captured file always has to be processed into an RGB image.

Processing the Raw file includes interpolating each photo-site's 'missing' GB or RB or RG levels from the surrounding photo-sites. Unfortunately this 'demosaicing' also leads to digital artifacts and aberrations that may need to be corrected. Higher megapixel chips can be enlarged further before the artifacts become visible, but all chips result in artifacts. Dark current noise and JPEG compression artifacts further reduce image resolution.



Dark current noise




JPEG compression artefacts

The final picture quality depends on many factors including the array's sensitivity to light, its signal to noise ratio, the heat it generates and/or how well it is cooled, the camera's propriety image processing algorithms, and the number of photo-sites (megapixels).


Image Resolution

The optimum pixel resolution of an image is dependent on its subject matter and required optical resolution, output medium, and the intended print viewing distance.


Subject Matter: Low contrast and/or diffuse subjects, such as mist, clouds and out of focus areas can be satisfactorily rendered at lower resolutions than sharply focused highly detailed images. The following image was rendered at the same physical size but different pixel resolutions. Photoshop's type tool was used to indicate each file's resolution. All images were printed at the printer's highest quality and the same size on glossy paper.




Scaled to 50ppi and printed at 2880dpi



Scaled to 100ppi and printed at 2880dpi



Scaled to 200ppi and printed at 2880dpi




Scaled to 360ppi and printed at 2880dpi

Pixels per Inch (ppi): Refers to the resolution of the file. Crisp black text on a white background is a good test for the highest required ppi resolution for an image. Rendering and printing text at higher ppi resolutions will reveal how many pixels are required for clean sharp looking text. The rendering of pictorial subjects like skin, sky, and clouds will normally not be improved by rendering at a higher ppi resolution. However some printer drivers expect to receive a precise ppi value, and they will automatically resize images to that required ppi. This can result in decreased optical resolution despite higher ppi resolution, and can only be verified by critical printing tests.

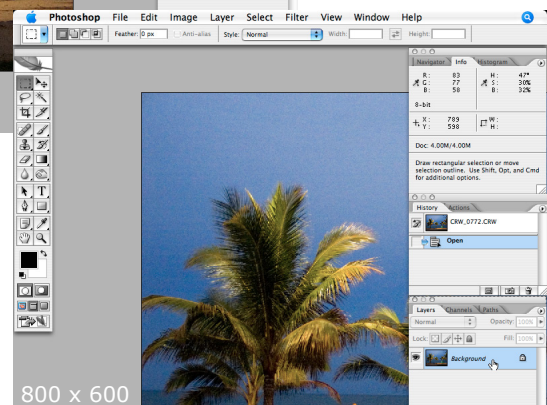
Monitor Resolution

The monitor's maximum resolution equals its physical size multiplied by its pixel density.

Pixel Dimensions: Fixed size objects (such as Photoshop palettes) will appear much larger (approximately four times larger) on a small 640 x 480 pixel desktop compared to a larger 1280 x 1024 pixel desktop. Larger pixel dimensions, not a physically larger screen provides more desktop space for your image relative to Photoshop's palettes, menus and tool bars. Smaller pixel dimensions result in a crowded and less efficient desktop/workspace.



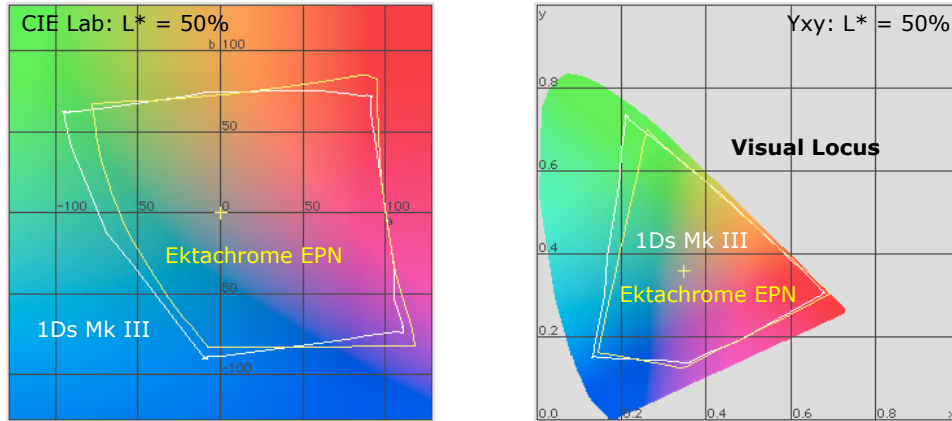
A wide screen 1440 X 900 pixel resolution provides ample room for Photoshop palettes without encroaching on the image. At 800 x 600 pixels the screen is crowded. A resolution of 1920 x 1200 pixels is often considered optimum for image editing on large (23" to 30") screens.



The computer's Display Control Panel (PC) or Displays Preference Pane (Mac) will indicate the range of monitor resolutions its video card can support. Set the monitor to a minimum of 24 bit colour and the highest pixel dimensions that still allows easy navigation of the desktop icons, tools and menus. While many LCD panels operate best at their 'native' or maximum pixel resolution, very large screens (30") may render menu items and dialogue boxes too small (at their native resolution) to read when seated at a comfortable working distance. Sitting closer to more easily read and interact with dialogue boxes and menu items may result in frequent head turning and Occupational Health and Safety issues. These concerns can be reduced on screens larger than 24" by working on images that are smaller than the entire screen, while using the surrounding areas for palettes.


Colour Gamuts and RGB Profiles

Digital SLR cameras are capable of recording very large colour gamuts, often equal to or greater than that captured by film. For example the gamut of a Canon EOS 1Ds Mk III is very similar to the gamut of Kodak Ektachrome EPN (ISO 100).



These are very large gamuts capable of accurately capturing many of the most saturated colours we can see (within the visual locus). They also substantially exceed the colour gamut of all monitors and printers. Therefore it is tempting to save camera files in a colour space (gamut) substantially smaller than its original gamut. This will potentially result in many saturated colours becoming clipped and lost. For example, selecting sRGB as the in-camera colour space, a colour space roughly equivalent to an uncalibrated PC monitor, is fine if your output is destined for the internet, but it would not be recommended for images full of saturated colours that are to be printed on a wide gamut inkjet printer.

The danger of compressing a colourful image into sRGB is that important colour distinctions might be lost. It is often better to save the image in a large colour space like Pro Photo RGB.




All the original colour differences will then be faithfully preserved and can be converted into printable colours at the time of printing. Unfortunately at this time no digital camera supports wide gamut spaces like Pro Photo RGB. However capturing RAW data and processing in a program like Adobe Camera Raw™ will allow the use of Pro Photo RGB.

Histograms

A histogram is a 'picture of a picture' that graphically reveals the distribution of tones in the image. It maps the two constituents of the digital image, 'pixels of resolution,' against the 'levels of brightness' represented by each pixel. Learning to 'read' the shape of a histogram is an important skill, because it indicates the nature of the image, and aids its diagnostic and aesthetic evaluation.



Midtones




Reading Histograms

The horizontal axis represents the levels of brightness, where the shadows (darker values) are to the left, and the highlights (lighter values) are to the right. The vertical axis represents the relative percentage of pixels at each level of brightness.



Shadows and Highlights



A histogram typically consists of at least 256 screen pixel wide columns that each represent a discrete level of brightness in 8-bit notation. Gaps in a histogram (empty columns) represent a loss of contiguous levels of tone (< 256). As the gaps increase, so does the posterization in the image. Smooth toned areas like sky and skin will first show this increased posterization as rougher or courser tonal transitions, and at worst, as 'banding' (distinct bands of brightness across the image).

Understanding Unsharp Masking

Sharpening (and blurring) effects are proportional to the degree of enlargement of the final print. For example sharpening that appears perfectly balanced on a 100 ppi screen will become excessive when enlarged in a 200 ppi Lambda print. Local effects are proportional to the distance (enlargement) over which they are projected. Therefore it is probably wise to apply these sharpening and blurring effects only once the final image size has been established and a test print (not the screen image) has been carefully evaluated.

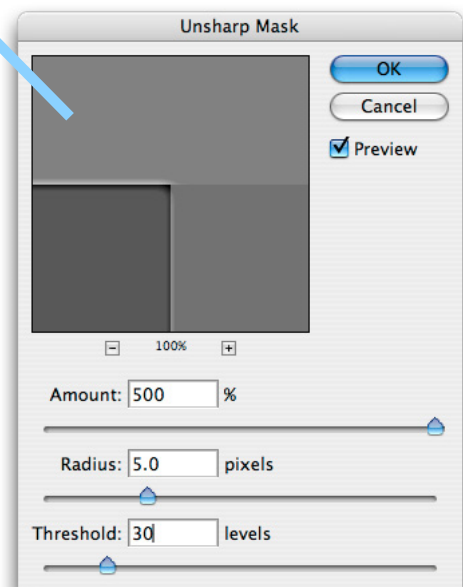


Unsharp Mask: 500% 1.0 pixels, 30 Levels

AMOUNT is the strength of the filter. In the Unsharp Mask filter, 500% results in much darker and lighter outlines (halos) than 100%.

RADIUS is the width of the outline, in pixels either side of the boundary being sharpened.

THRESHOLD is the number of levels of tonal difference that must first exist across a boundary for that 'amount and radius' to be applied. Begin adjusting the Adobe Photoshop Unsharp Mask (USM) filter at a low Radius setting around 0.5 to 2.0 pixels and 0 Threshold with a high Amount (100% to 500%).



Gradually increase the Radius until the edges have been adequately 'outlined', then slowly increase the Threshold setting until the effect just begins to wane. Finally reduce the Amount to suit the image's character. This will give you the basic sharpening setting for that image.

Always inspect the effects of the Unsharp Mask filter at 100% screen magnification. To limit unwanted sharpening artifacts such as colour fringing, initially apply the filter to a duplicate image layer then blend on Luminosity. The duplicate layer will also allow the sharpening to be locally inserted via a layer mask. The Adobe Photoshop sharpen and blur filters can also be applied in 16-bit colour depth which will further reduce unwanted image artifacts and promote smoother transitions in the final print.

Sharpen Layers can also be blended on Darken and Lighten to separate the dark and light outlines, which can then be applied at different opacities that best suit the image.