

How Many Pixels Do We Need?

March 2005

Over the last few years the digital camera market went through a veritable pixel race. This has now calmed down, once the consumer realised that all you need is a 4 MP camera to produce a perfect A4-sized print. It is nice to have some overhead for possible cropping, but 6 MP is really all the average photographer will ever utilise.

For the professional it is a different story. He or she will think of medium format film and the details you can find in a 160 MB scan. This is the kind of MB figure we'd like to have for a 50x76 cm print. Even today's high-resolution medium format camera backs don't come anywhere close to that. However, there is no film grain and digital camera files lend themselves to upsampling. The question is, how many MP or MB do we really need and, if we haven't got enough, how do we upsize our files?

Optimum File Resolution for Printing

Commercial printers routinely ask for 300 dpi files. This figure goes back to the resolution of the first LaserWriters in the mid 1980s and 300 dpi has become something like a magic number for printing (like the 72 dpi for monitor display).

In **offset printing** the print resolution is determined by the halftone screen frequency. The file resolution should be between 1.2 and 2 times the screen frequency. Newspapers are printed with a coarse 80 line screen. A 2x factor will give us a file resolution of 160 dpi. Glossy magazines typically use a 133-150 line screen and high-quality art is printed with a 200 line screen. At this high level of screen frequencies a factor of 1.5 is considered as sufficient. This gives us a maximum file resolution of 225 for glossy magazines and 300 dpi for the best fine art print.

For **contone output** the file resolution simply needs to match the resolution of the output device. A 300 dpi dye-sub printer will give its best with a 300 dpi file. An 8k-film recorder would, ideally, need a 128 MB file (8192x5461 pixels for 35mm film), although this can be halved with not much of a penalty.

Today's high-resolution **inkjet printers** work by putting 720 to 2880 drops of ink per inch onto paper. Every pixel of our file needs to be dithered with many ink drops to simulate thousands of colours. We certainly don't need to (and we shouldn't) feed our printer with 720-2880 dpi files!

Firstly, a print resolution of 1440 dpi is plenty for matt and fine art paper and you will be hard pressed to tell the difference between 1440 and 2880 dpi even on the best photo gloss paper. Secondly, the resolution of the file should be an integral divisor of the printer resolution. Most experts agree that a file resolution of 240 dpi is sufficient for photographic images, which happens to be 1/6 of 1440.

I have done many tests with ultra-sharp images and could only tell the difference between a 240 and a 360 dpi file by critically looking at the print with an 8x loupe.

Without loupe you will be hard pressed to find any fault with a print coming from a 180 dpi file! Because of the viewing distances, poster prints can be done at fairly low file resolutions, typically 150 dpi down to 75. If the file resolution is too low, we start to see individual pixels, most evident in the 'jaggies' on diagonal sharp lines.

After establishing an optimum file resolution of 240 dpi for a high-quality inkjet print, we can work out how big we can print to: e.g. an 8 MP camera will give us a 37x25 cm print. For anything larger we need to upsample the file.

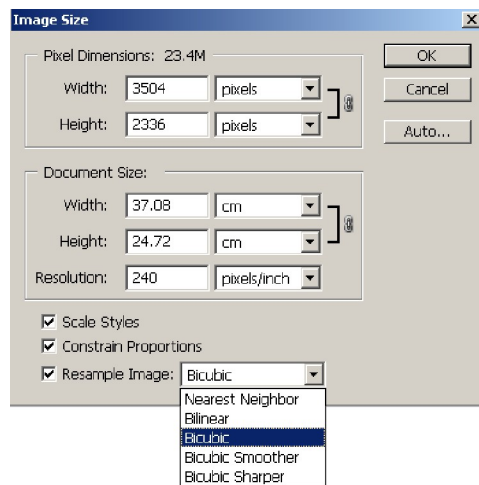
Upsampling of Digital Files

Upsampling means asking the computer to 'invent' new pixels between existing ones. Needless to say, there is no way that we can add details that weren't there in the first place. Adding pixels to an image can reduce aliasing (the 'jaggies'), but in general the print will become soft and blurry. The challenge is to retain sharpness and contrast.

In scanning there is a simple rule: Never upsample. Film grain is usually not resolved, unless the film is scanned at extremely high resolutions. Instead, the film grain and the scanner's sampling grid form interference patterns which tend to make things worse when upsizing. For best quality, scan at the maximum optical resolution and then downsample to the required size.

Digital camera captures, on the other hand, don't have any grain and are much more tolerant of resizing. Upsampling by 200-300% will still get you very acceptable results.

Until recently Photoshop gave us three options for resizing.



'Nearest Neighbour' is the fastest and most basic algorithm. It is only useful for images made of coloured lines and shapes (I used it for the above screen capture).

I cannot think of any reason for using 'Bilinear', but the 'Bicubic' interpolation is the algorithm recommended for photographic images.

The latest version of Photoshop CS gives us two more options: 'Bicubic Smoother' is specifically designed for upsampling. It provides smoother results that respond better to subsequent sharpening than the normal 'Bicubic' resampling.

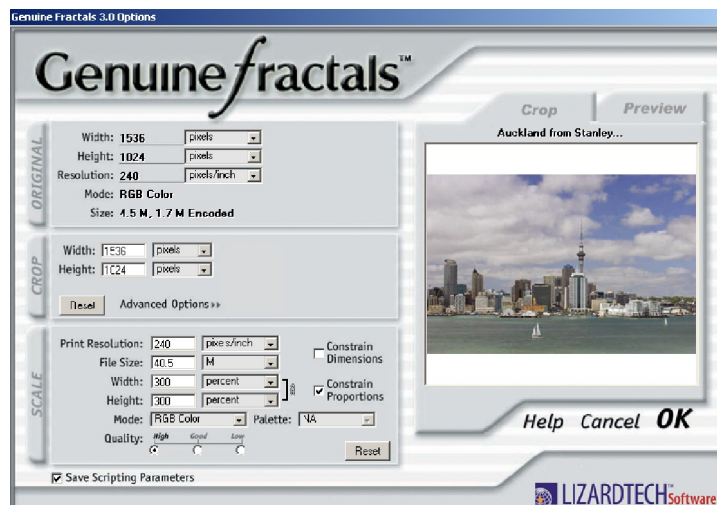
'Bicubic Sharper' is intended for downsampling and does a slightly better job in preserving details than the straight 'Bicubic'.

After upsampling the file needs to be sharpened. I always use the Unsharp Mask filter with an edge mask, often in combination with High Pass sharpening (both these sharpening techniques have been described in The Photographer's Mail Vol 13 No.3). Upsampling does enhance any camera noise and I usually run the file after upsizing, but before sharpening, through a noise reduction programme, such as Neat Image or Noise Ninja.

Photoshop's Camera Raw plug-in gives an option for resizing during RAW conversion. It is not recommended to use this for upsizing, unless you use a digital camera with non-square pixels.

Other Algorithms for Upsizing

There are several programmes which claim to do a better job than Photoshop's Bicubic formula. The best-known one is **Genuine Fractals** from LizardTech (www.lizardtech.com), currently vs 4.0 (as of July 2005). The more expensive PrintPro version also works with CMYK files, but for photographers the lighter version is fine, which can handle 16-bit RGB and 8-bit grayscale files.



The programme works as a plug-in and you need to save the file in the .STN format, which comes up as a new option once Genuine Fractals is installed. When reopening the file, a dialogue box allows you to specify the new image dimensions in percent, pixels, cm, inches or Mbytes.

PhotoZoom (S-Spline) from Shortcut Software (www.shortcut.nl) is a standalone application that offers three interpolation methods: S-Spline, Bicubic and Bilinear.

Whereas the Bicubic algorithm interpolates the new pixels from its surrounding neighbours, S-Spline 'looks' at the image in an intelligent way and tries to adapt its interpolation method from pixel to pixel, depending on how sharp or smooth the image is in that area. For smooth areas, the algorithm will perform much like B-Spline or bicubic interpolation. For sharp areas it uses different interpolation techniques, trying to maintain the original sharpness of the image.

XFile from Human Software (www.humansoftware.com) is a plug-in that is accessed through the filter menu. It offers a number of interpolation options, including Gaussian, Lanczos and the proprietary Xfile. A large preview makes it convenient to choose the best algorithm.

The **FM Stair Interpolation** is an Action created by Fred Miranda (www.fredmiranda) which applies Photoshop's Bicubic interpolation in several small steps, followed by the Unsharp Mask filter.

Lastly I want to mention **Irfanview** (www.irfanview.com). This is a freeware programme for non-commercial applications. It has a number of interpolation options, of which Lanczos is the slowest, but the best.

... And which One is the Winner?


For comparison I chose a shot taken with an 8.2 MP camera of Downtown Auckland. It has all the sharpness and detail of a 50mm lens at f/8. I downsized the file to 9 MB and then upsampled it with the various interpolation methods to the original 23.4 MB size, followed by careful sharpening. Three of the results are shown here together with the original file:




Central Auckland across the Harbour (Canon 20D, 50 mm f1.4 @ F/8)



Bicubic Smoother



Genuine Fractals



Original



Lanczos

I hope that the differences will come out on the sheetfed press; on an inkjet print the winner is the Lanczos algorithm from Irfanview, closely followed by Photoshop's Bicubic Smoother. In a similar exercise with a portrait shot, Genuine Fractals was the clear winner. The best algorithm will depend on the particular image, its content, size and quality.

Genuine Fractals is my favourite and I was quite surprised with the results of the Auckland shot. In defence of GF I must say that the programme gets better the larger the original file is. In other words, if I had started with the original 23.4 MB file instead of the 9 MB file, GF would have done a much better job. But then we would not have been able to compare the results with the original.

In general, the difference between the various interpolations is small and in many cases it will be hardly worthwhile using anything other than Bicubic or, if available, Bicubic Smoother. A great deal will depend on your skill of sharpening the image after upsampling.