
Image Resampling

Written by Jonathan Sachs
Copyright © 2001 Digital Light & Color

What is Resampling?

Digital images consist of a rectangular grid of evenly spaced pixels. Each pixel can be thought of as a measurement or *sample* of the light from a subject. Commonly, the original samples are obtained using a scanner or digital camera by averaging the amount of red, green and blue light that falls on the sensitive area of each of its CCD sensing elements.

Resampling is the mathematical technique used to create a new version of the image with a different width and/or height in pixels. Increasing the size of an image is called *upsampling*; reducing its size is called *downsampling*.

When images are upsampled, the number of pixels increases, but, with reference to the original subject, new image detail cannot be created that was not already present in the original image. As a result, images normally become softer the more they are enlarged since the amount of information per pixel goes down.

When images are downsampled, information in the original image has to be discarded to make the image smaller. Thus if you downsample and then upsample an image, you will not get all the original image detail back. Downsampling a soft image can make it appear sharper even though it contains less information than the original.

Resampling in Picture Window is performed by the Resize transformation. The Resize dialog box displays the input image width, height and dpi setting and lets you specify the output image width, height, dpi setting and resampling method. Note that changing the dpi setting (resolution) does not affect the image data in any way—it is simply a suggested scale factor that is associated with the image to expedite printing or document the size of the original subject.

Why is Resampling Important?

A few situations where image resampling comes into play are:

- Reducing the size of an image from a scanner or digital camera for emailing or display on the web.
This step is normally followed by saving the file in JPEG format and uploading it to a web site or attaching it to an email message.
- Increasing the size of an image before or during the printing process.
Printing an image almost always requires resampling it as it will usually be too small or too large to make a print of the desired size. This resampling is handled automatically by your Windows printer driver. Most drivers do a good job of resizing images, but you can resample the image yourself prior to printing if you suspect the driver is not doing a good job.
- Creating Windows wallpaper
Wallpaper is a term used to describe an image that is used as a background for the Windows desktop. You can create wallpaper from your own images by cropping them to proportions of 4:3, resizing them to your current screen size (e.g. 1024x768) and then saving them as a (Version 1) BMP file in your Windows folder. The file will then be available in the image list displayed in the Windows Display control panel under Background.

How Resampling Works

It turns out that resampling a 2-dimensional image can be broken down into two one-dimensional resampling passes. In one pass, horizontal resampling is performed producing an image with a different width but the same height. In the next pass, this intermediate image is resampled vertically, changing its height while leaving the width the same. This is computationally much more efficient than trying to combine the work into a single pass.

Upsampling involves interpolating between the existing pixels to obtain an estimate of their values at the new pixel locations.

Downsampling involves computing a weighted average of the original pixels that overlap each new pixel.

Color images are treated like three black and white images which are separated from the original image, individually resampled, and finally recombined to create the final image.

Resampling Methods

Many different resampling schemes are possible. Most techniques work by computing new pixels as a weighted average of the surrounding pixels. The weights depend on the distance between the new pixel location and the neighboring pixels. The simplest methods consider only the immediate neighbors; more advanced methods examine more of the surround pixels to attempt to produce a more accurate result. Picture Window's Resize transformation (starting with Version 3.0d) supports the resampling methods listed below. Each one is illustrated with a small sample image which is magnified by 1000% (i.e. enlarged by a factor of 10). This huge amount of magnification is much greater than you would use under normal circumstances—it is used here to make some of the relatively subtle differences between the different methods more apparent.

Nearest Neighbor

The nearest neighbor resampling method is very simple; it computes new pixels as the value of the nearest pixel in the original image. This produces a blocky result when upsampling and a coarse, grainy effect when downsampling.



Bilinear

Bilinear resampling computes new pixels using linear interpolation. When upsampling, this method operates on the 2 by 2 cell of pixels surrounding each new pixel location. The resulting images are much smoother than those produced by the nearest neighbor method but are not as sharp or as smooth as image computed using bicubic resampling (see below).



Bicubic

Bicubic resampling computes new pixels using a cubic splines. When upsampling, this method operates on the 4 by 4 cell of pixels surrounding each new pixel location. This is the recommended resampling method for most images as it represents a good trade-off between accuracy and speed



This method has an extra parameter you can use to sharpen the image during the resampling process. No sharpening gives the smoothest results; excessive sharpening can produce halos and more jagged diagonals. The default sharpening value of 16.7% is usually a good compromise.


Lanczos 4X4, 6X6 and 8X8

Lanczos resampling computes new pixels using a *Lanczos-windowed sinc kernel*. When upsampling, these methods operate on the 4 by 4, 6 by 6, or 8 by 8 cell of pixels surrounding each new pixel location. The 4 X4 method is nearly identical to the bicubic method with no sharpening; the 6X6 and 8X8 methods produce slightly more accurate results than bicubic but take a little longer to compute.





Resampling Kernels

Each of the resampling methods can be characterized by its *resampling kernel*. The resampling kernel is simply a curve that determines the relative weights of adjacent pixels based on their distance from the new pixel. The various kernels are represented by the following graph:



Genuine Fractals

An alternative to the standard resampling techniques is a Photoshop plug-in called Genuine Fractals from Altamira. This program does not resample images directly, but it does let you save them using their proprietary fractal compression algorithms and a special file format. When you reload the compressed file, you have the option of resizing the image when it is opened and this in effect provides a fractal resampling feature. Genuine Fractals has been hailed by some as a major breakthrough for enlarging low resolution images from digital cameras so they can be printed at large sizes without becoming fuzzy. Actual results seem to depend quite a bit on the subject matter. The fractal images appear to be built up as areas of more or less solid color separated by relatively sharp edges, a little like a watercolor painting. The following examples compare fractal resampling to traditional resampling for 10X magnification



Here are some more samples of fractal vs. traditional resampling with a 10X magnification:



Genuine Fractals






Lanczos 8X8






Fractal Resampling





Lanczos 8X8



Depending on the subject matter, it appears that fractal resampling is sometimes fairly effective, sometimes produces results that are visibly different from traditional resampling but not necessarily better, and sometimes produces results that are very similar. Whether or not you find the type of artifacts fractal resampling introduces more or less pleasing is probably a matter of personal taste. Interestingly, the first image (for which the results seem fairly good) is a tiny piece of a Photo CD scan while the second two images are from digital cameras. The noise in the digital camera images seems to reduce the effectiveness of the fractal resampling scheme somewhat as there are fewer areas of nearly solid color and a lot of false edges.