

RESOLUTION

An In-The-Trenches Guide To Resolution Basics

BY ANDREW RODNEY

Explaining digital resolution is like trying to explain the metric system to someone brought up on inches and feet. In order to understand resolution, you must start small—at the pixel level.

How big is a pixel? No, it isn't a trick question. When I ask my students this question, I can usually tell if they understand resolution based on their answers. If someone says, "It depends," then I know he understands resolution.

A pixel (short for "picture element") is the smallest unit in a digital file. Digital files consist of a series of pixels running the length and width of the file, thus creating a two-dimensional image. If you would like to see what a pixel looks like up close, open Adobe Photoshop and your favorite file, zoom in as close as you can (1,600 percent), and you'll see individual pixels. When combined and viewed at 100 percent, pixels create a mosaic that resembles a continuous tone image.

Scanner and digital camera resolution is measured in pixels per inch (ppi), and output device resolution is measured in dots per inch (dpi). Many people confuse dpi and ppi, but they are not the

same. Scanners and digital cameras create pixels, not dots. However, the pixels eventually will be output as dots; therefore, a dot will be created for every pixel. For example, if a 300-pixel scan is output on a printer at 300 dpi resolution, the print will be one inch in size, because one dot was created for each pixel. So, perhaps it is better to ask the question, "How big is a dot?" For a printer that is capable of printing at 300 dpi, each dot is 1/300 of an inch.

If you print a file that is 3,000 pixels along the long axis, and the printer makes dots that are 1/300 inch in size, one inch of printed image will appear on the output for each group of 300 pixels in the file. The total output will be 10 inches. If you output the same file to film, and you use a film recorder that has a resolution of 1,000 dpi, each dot is 1/1,000 inch. With 3,000 pixels in the file, the film recorder will produce one inch of image on film for every group of 1,000 pixels. The total output will be three inches. In both cases, there are 3,000 pixels in the file, but one output device yielded 10 inches and the other produced only three inches. So, in this scenario, the film

recorder has higher resolution than the printer, because the dots are smaller (1/1,000 inch vs. 1/300 inch).

Think of resolution in another way. Digital files don't have any size other than the space they take up on a storage device. How big is an RGB file that is 2,000x3,000 pixels? On a hard drive, it takes up about 18MB. How big is it really? That's impossible to answer. It doesn't have size until those 2,000x3,000 pixels are output to a printer. Therefore, the question is, how many dots can the printer lay down per inch and how big are the dots? If the printer can lay down 300 dpi, the file will be a drastically different size than if it can lay down 1,000 dpi. So how can a digital file have "size" other than the space it takes up in storage? With the 2,000x3,000-pixel file, you can create an output that uses dots that are 1/300 inch, 1/1,000 inch, or for that matter, any size.

No Democracy

Not all pixels are created equal. Output quality depends on the quality of the pixels in a file. For example, a 5,000-ppi drum scanner can easily produce a 300-ppi file that

How To Change Image Size

To do this tutorial, open the Resolution Test file on the *PEI* Web site, www.peimag.com.

1. Select Image Size under the Image menu. The dialog will look like the figure at right. The new dialog box in Version 4.0 operates in a fashion similar to 3.0, with a newer interface that is divided into three important sections. The top section, labeled "Pixel Dimension," shows pixels in width and height. Using the pop-up menus, you can also set this to a percentage.

In the second section, labeled "Print Size," you can work with a number of different units, depending on the pop-up menus next to the data entry fields. In this example, I set the unit of measure to inches, but you can also pick percent, centimeters, points, or picas. Notice that the print size is 5.556x6.994 inches because the field labeled "Resolution" is set to 72 dpi.

The last section is the most important. It contains two check boxes that allow you to add or subtract pixels when the Resample Image check box is on. New in Photoshop 4.0 is the ability to control the interpolation method. I recommend you leave "Bicubic" on for best quality, as shown in our example. Nine times out of 10, the Proportions check box will be set to On, or the file can be resampled and distorted.

2. The Resample Image check box is *very* important. Make sure that it is not checked. This forbids Photoshop to add or subtract (resample) the file. Notice that when you do this, the top portion of the dialog box where the number of pixels appears is "locked" and these values cannot be changed. Now we will play around with the width and height pop-up menu and the resolution fields to see how pixels, inches, and resolution work together. When the Resample box is off, we can use this dialog as our "resolution calculator."

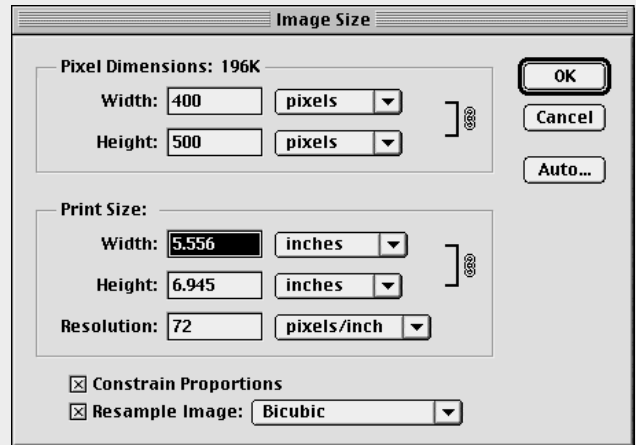
3. Click on both the width and height pop-up menus to make sure they are set to inches. You will notice that the size is 5.556x6.994 inches.

4. Change the figure in the resolution field from 72 to 100 ppi. Notice that instantly the width and height change to 4 and 5, respectively.

5. Change the pull-down menu from inches to percentage. You will see that nothing has changed! You are still at 100% (the original size). Also note that the top portion of the dialog box shows that you still have 400x500 pixels. Set the resolution back to 72 dpi from 100 dpi.

6. Check the Resample Image check box to On. This allows you to change figures in any fields and add or subtract pixels in the files.

7. Enter 1,400 in the Pixel Dimension width field at the top portion of the dialog box. Notice what happens. First, the height field is updated from 500 to 1,750. This is because the Proportions Check Box was on and Photoshop kept those proportions by doing the math for you and adding the correct number of pixels for that dimension. Also notice the new file



size that was updated in the top portion of the dialog box. The file grew to 2.34MB, because you've added a lot more pixels to the file. This is known as resampling up the image.

Also notice that Photoshop shows you what the file size was prior to the resampling (it says, "Was 196K"). Finally, notice that if you kept Width field in the print size in percentage, the new size now states 350%. If you left the Height field in inches, it would now say 19.444 inches. Feel free to toggle the various pull-down menus in this portion of the dialog box to different options. When done, set the pull-down menus back to inches.

8. Examine the Print Size fields in the center portion of the dialog box. Notice the new size? Since we actually added more pixels to the file, it will output larger at 72 dpi. It's now going to be 19.444x24.306 inches.

9. Make sure the Resample check box is off. While you should revert to the original settings, you can also reset the dialog box to the original settings. Hold down the Option key and the "Cancel" button will toggle to the "Reset" button. Now type different figures into the Resolution field and watch as only the print size changes. Notice that when you type in 300 Pixels/Inch, the size shrinks to 1.3x1.6. You can type different output resolution figures into this field and see how the pixels in the file will be mapped to a certain size. You now have a resolution calculator.

Try using this dialog box with different figures until you feel comfortable with these concepts. If you think in terms of pixels along a width and height, then think about the device's actual output resolution, you can figure out how large that file will be output. Or think of it this way: A digital file has no real size until you map out the pixels to a device. A 4,000x5,000-pixel file can be 4x5 inches, 40x50 inches, or any size, depending on how you lay out the pixels. Higher resolution devices create smaller pixels, which means you can lay more of those tiny pixels within an inch, making them more difficult to see individually.

would exactly match the size and resolution of a 300-ppi, \$200 flatbed scanner; however, the difference in quality would be significant. Even if one device is capable of producing more pixels per inch than the other,

it may not produce superior quality. This is especially true of digital cameras. Most people who are in the market for a digital camera use resolution as the primary indication of quality, but there are

many factors that affect quality. The CCD and its noise level, the analog-to-digital converter, the lens, and the amount of compression all play a role in image quality.

Here is a step-by-step tutorial that

further illustrates resolution, as well as Adobe Photoshop's often misunderstood Image Size dialog box.

1. Launch Photoshop.
2. Create a new file by selecting New under the edit menu (Cmd/Ctrl N).
3. In the pop-up window, name the file "Resolution Test."
4. Notice the Width and Height fields. You can specify how you want to work—in pixels, inches, etc. In the menu, select pixels and enter 400 in the width field and 500 in the height field (Figure 1). Set the Resolution Field at 72.



Figure 1

5. Set the Mode in the pull-down menu to grayscale (it's possible to create a file in CMYK, RGB, or grayscale). Later, you will see that in terms of storage, each of these color spaces produces a different size file; however, this is not a change in resolution.

6. The Contents buttons allow you to set the background color in your new document. Leave it on White.

7. Click the OK button or hit the Return Key to create a new file. The newly created file should be 196K in size, as indicated on the small size label in the lower left corner of your file (it should read 196K/0). Option Click on this small label (see Figure 2). Notice that this little pop-up menu shows you information such as the number of pixels in the width and height, the number of channels, and the resolution, which is 72 ppi. What is the actual size in inches? This dialog states that the "size" is




Figure 2

5.556x6.944 inches. If you do the math, 400 divided by 72 equals 5.556, and 500 divided by 72 equals 6.944. Photoshop does the math for you based on the resolution set for the file size (72 dpi). In other words, if you output this 400x500-pixel file to an output device with 72 dpi resolution, the size would be 5.556x6.994 inches.

I Am, Therefore I Have Size

Because files don't have a height or width until they are output, digital files have a size complex. In other words, they need to know what size they *would* be if output. This information is necessary for imaging applications and digital printers to assign size to a file. A file carries a resolution "tag," which is information embedded in the file that tells software and printers how to achieve desired output sizes. Tags are usually placed in a file when it's created by a scanner or digital camera. It's possible to change a tag and, therefore, the size of output without adding or subtracting pixels.

For example, if you place a 4x5-inch transparency on a scanner and want to capture 4,000x5,000 pixels in the file, you must inform the scanning software of the file dimensions as well as the pixel dimension (resolution). The only way to tell the software what the image size should be is to tell it what the resolution should be and how you want to divide the 4,000x5,000 pixels into each inch. This file could be 4x5 inches if the resolution is set to 1,000 ppi (notice I use the term "ppi" because the file hasn't been output yet). However, a 4,000x5,000-pixel file could also be 8x10 inches if the resolution of the file is 500 ppi. For

that matter, if the resolution is set at 100 ppi, the file will be output at a monstrous 40x50 inches.

The combination of resolution and pixel dimensions allows us to do the math to determine a size when output. Many printers will examine the resolution tag in the file and attempt to output the file at 100 percent of the given size. If a digital printer reads that the file is 8x10 inches at 500 ppi, it will do its best to make an 8x10-inch print, regardless of its native resolution. It doesn't really matter that a file has a certain number of pixels that can be divided into hundreds of permutations. The printer only wants to provide the size it "thinks" is correct at 100 percent. Many printers use their native resolutions to internally resample the file up or down in order to produce a 100-percent file based on the resolution tag. Remember, the tag is used only for math purposes. There must be a figure (the resolution) that allows the pixels in the file to be divided per inch to yield a final print size.

Monitors Count, Too

Since you know how to work in pixels and how to use Photoshop's Image Size dialog box, you also know you can turn off the Resample Check box and type 300 dpi into the resolution field. This field will modify the tag from 72 ppi to 300 ppi, so that the file will output at 100 percent at 10.24 x 6.8 inches.

You may be wondering why 72 dpi is a common resolution tag. This is because most monitors have a resolution of 72 dpi, and it is an output device much like a printer. If you have a file that's 216 pixels, it will appear (at 100 percent) in Photoshop as three inches (216/72 = 3). Those who work with scanned images for the Internet know that it commonly uses files that are no more than 640x480 pixels (divide the two dimensions into 72 and you'll see why). This

provides a very nice file for the lowest common denominator, a 14-inch display.

Imagine what would happen if you wanted to view a 2,000x3,000-pixel file in Photoshop. Since the monitor output is 72 dpi, the image you would see on screen at 100 percent would appear so enlarged, that you would see only a fraction of the entire image. That is because every pixel in the file would be displayed by one dot on the screen, making image editing very difficult. Fortunately, Photoshop allows you to zoom out to see the entire image. To do this, however, the software has to do an on-the-fly interpolation for the monitor.

When you zoom out to less than 100 percent, the image you see is sub-sampled and inaccurate. Therefore, it's very important to do some editing tasks at 100 percent. For example, if you apply an Unsharp Mask filter to a file, it should always be done at 100 percent, so that what you see on screen is accurate when you move the sliders back and forth to obtain desired effects. If you zoom in past 100 percent, you start to see the actual pixels. For doing fine work in Photoshop, make sure you're working with a 100 percent view on screen, or you may be in for a shock when the file is output.

Bulking Up On Color

File resolution and physical size are not directly related. True, a 1,000x1,000-pixel file takes up more space than a 100x100-pixel file, but files can be made up of color channels that affect the physical size in cases where the resolution is identical. For example, a 100x100-pixel grayscale file will be a third the physical size of a 100x100-pixel RGB file. That's because the RGB file has three color channels (red, green, blue), one for each color. A CMYK file will be one-fourth larger than the RGB file. By knowing a file's pixel

dimensions, you can always figure out the physical storage size, no matter what the color mode.

Try this exercise: Multiply the pixel dimensions. If you have a file that is 1,000x1,000 pixels, $1,000 \times 1,000 = 1,000,000$. This is the total number of pixels in the file. Multiply that figure by the number of color channels. For an RGB file, that would be: $1,000,000 \times 3 = 3,000,000$. Now, divide 3,000,000 by 1,024 to determine bytes (1,024 is the unit measure of bytes) and you'll get 2,929. Divide again by 1,024 to get megabytes, and you get 2.86MB. So, the formula is: Multiply width by height by the number of channels, then divide by 1,024 to get bytes, and divide again for megabytes. By knowing the pixel dimensions, you can determine the physical size of the file no matter what the color mode.

Buyer Beware

You may have seen ads and spec sheets for flatbed scanners in which a manufacturer entices you into thinking its scanner has a larger resolution capacity than it really has. Have you ever seen the specification 600x1,000 dpi? You've learned that dpi is an incorrect term, but that's a minor point compared to the fact that this scanner can scan at an optical resolution of only 600 ppi. In other words, the second figure is meaningless.

The sensor in most flatbed scanners is a single row of CCDs, which contains the number of pixels within one inch. In this case, the sensor has 600 tiny, square eyes within one inch, capable of creating 600 pixels per inch. It's rather easy for the software controlling the scanner to capture only 300 pixels per inch (we simply "turn off" every other pixel along the CCD sensor).

What does the other figure in the spec mean? It represents the scanner's stepper motor, which moves the CCD up and down the

bed of the scanner as it moves over the artwork. The stepper motor can move 1/1,000 inch. However, the CCD can capture a maximum of only 600 pixels per inch. It is impossible for this scanner to capture 600 pixels in one dimension using the true optical resolution of the CCD and, at the same time, move up the bed 1/1,000 inch. In this case, pixels are square, not rectangular. Therefore, when you see a specification such as 600x1,000 dpi, and you attempt to produce a scan at 1,000 dpi by instructing the scanner to move up the bed using the finest step available, the scanner software must interpolate data on the other axis.

Since the optical resolution is 600 ppi, 400 ppi of data has to be interpolated in order to actually scan at 1,000x1,000 ppi. When you scan at the optical resolution of the scanner, the stepper motor simply matches the resolution as it moves up the bed over the artwork. A scanner that has a rating of 600x1,000 dpi is capable of an optical resolution (true resolution, not interpolated) of only 600x600.

To Interpolate or Not?

If you need to create a larger file than your scanner can produce optically, should you let the scanner interpolate, or should you interpolate in Photoshop? It depends on what kind of interpolation algorithm the scanning software uses. Bicubic interpolation is the best quality algorithm that Photoshop uses. Few scanner manufacturers will tell you if their interpolation is Bicubic, or a faster, lower quality algorithm. I recommend the following test: Scan the file at the optical resolution of the scanner and do your interpolation with Photoshop. Then scan the file with the same resolution by allowing the scanner to interpolate, and examine the two files for quality.

If you scan at the higher resolution given for the stepper motor, interpolation has to take place in the other dimension. But in this case, we are interpolating on only one axis. If the scanner software uses a good interpolation algorithm, you will get better quality scanning at the highest resolution the stepper motor can provide than scanning at the optical resolution and resampling later in Photoshop. Because Photoshop cannot interpolate a file in only one dimension, it must interpolate both horizontally and vertically.

Rather than scanning at 600x600 ppi and interpolating to 1,000x1,000 ppi in Photoshop (which adds 400 pixels to both dimensions), scan at 600x1,000 ppi and let the scanning software add 400 pixels to only one dimension. Many high-end scanners have hardware interpolation boards that interpolate far better than Photoshop, and many expensive output devices have similar capabilities. While you generally achieve superior quality by not interpolating a file, don't dismiss interpolation altogether.

Knowing the optical resolution of your scanner's CCD can help decrease your scan times, because when you scan at an even division of the optical resolution, the scanner software has far less work to do. For example, if you have a flatbed scanner with an optical resolution of 600 ppi, and you scan an image at 300 ppi, you are using half the number of pixels per inch in the CCD. However, if you want to scan at 72 ppi, the scanner software has to do a lot more math to provide an odd number of pixels per inch from the CCD. You will not lose quality, but you'll have to wait a bit longer for the scan.

If you have a fast processor (and a slow scanner), you may want to try scanning at an even division of the CCD, then resampling to the exact size you want in Photoshop. Some very fast film scanners, such

as the line from Polaroid, have the option of scanning in "native resolutions" or any resolution you choose. Should you use a resolution that is not native, the scan time increases. Some scanners use the term "pitch" to aid you in deciding on a native resolution. An even pitch number means you are using a resolution that is an even division of the optical resolution of the CCD.

The smoke and mirrors surrounding resolution and manufacturer specs doesn't stop at scanners. There are digital cameras on the market that use rectangular pixels and then interpolate the data after the image capture to produce a file with square pixels. The only way to know for sure is to find out if the pixels are square or rectangular. Like the flatbed scanner example, if the CCD doesn't have square pixels, the data has to be added somehow, and guess what? It's interpolated. I'm not saying you should always avoid interpolation. Just be aware that when you interpolate, you are bound to experience some data loss. If the loss is OK, or you have no other choice, then go for it.

See Your Nearest Neighbor

Photoshop has three interpolation algorithms: Bicubic, Bilinear, and Nearest Neighbor. Bicubic produces the best quality, but takes the longest time to process. In general, Bicubic interpolation is the algorithm you would want to use most often, particularly for continuous tone images. But if you happen to be interpolating line art, it's better to use Nearest Neighbor. When you interpolate continuous tone images, the interpolation should anti-alias the image to avoid dreaded "jaggies." In other words, you want the edges of the image to be soft. This is certainly not what you want when you interpolate line art. You want the edges to remain sharp and clean.

In order to see the difference between these two Photoshop

interpolation algorithms, make a small file and use the Line tool to make a series of one-pixel-wide lines running up and down the image. Try to make a grid pattern. Be sure to save the file to disk. Now interpolate the file up in size a few hundred percent using Bicubic interpolation first, then after reverting the file, interpolate the same enlargement factor, but use Nearest Neighbor. Zoom into the file and examine the one-pixel black lines. You will see a vast difference between the two interpolation methods. ◀

Andrew Rodney operates Camera & Darkroom Digital in Santa Fe, New Mexico. Camera & Darkroom Digital runs a full-service digital imaging service bureau and training center. He can be reached at rodney@peimag.com.