





RESOLVING RESOLUTION

The Interlude (Part One and a Half)

by Richard Gross



The Pixel Pot



In the last issue, I began a discussion of resolution in digital imaging and what a stained glass artist needs to know about it to make resolution work for him instead of against him. I thought I'd done a pretty good job with that and was ready in this issue to move into telling what the resolution of a given image is and how to use that knowledge to your benefit.

Then I got a phone call from a reader who complained that he didn't understand the article at all because he wasn't really sure what a pixel is in the first place. It's probably good that it was a phone call, because, had it been face-to-face, I most likely would have looked at him like he'd just sprouted a second head that began quoting "Much Ado About Nothing." It just wasn't expected at all.

I guess I figured people are just born knowing what a pixel is. No; I suppose that's not a reasonable assumption. But there it is. So, in the interest of helping those who are pixelly challenged, I decided to back up a bit add this explanation of what a pixel really is. Again, this is a simplified explanation of many of these concepts. It must necessarily be so, because to go into the detail they deserve would (and has) fill books. The point is, if this is very interesting to you, then there is much more detailed information available. If it's not so interesting, then this should be enough to get by. So, without further ado...

The Pixel Pot

The Pixel Pot, shown opposite page, has no existence in the real world; it is strictly a digital image of a three-dimensional computer-generated model. It is nothing more than a 1500x1500 grid of dots of color.

On the far left is the final image; even though this does not represent a real object and it has only a virtual existence, our minds can examine this grid of colors and see Egyptian pottery.

On the right are screen shots of this pixel grid as it is examined in Photoshop. In the first image, the pixel grid is examined at 41.5%; we can clearly see that it represents a pot. In the second image, the pixel grid is viewed at 100%, which means that each pixel on the grid occupies one and only one pixel on the screen. The eye still clearly sees a pot.

The third image is a view at 800%. In this image, a light gray grid is overlaid on the pixel grid to show individual pixels. Remember that each pixel is a single, individual dot of color. The eye can still make out some detail and even see that this image is a magnification of one of the round decorative motifs on the Pixel Pot. In this view, each individual image pixel would have occupied an eight-by-eight grid of screen pixels.

In the last view, we see clear representations of each individual image pixel. Here, the magnification is at 3200%, so each image pixel is represented by a grid of screen pixels that measures 32x32 and all display the same color.

A pixel is, quite simply, the smallest unit of color in a digital image. There are technical definitions that are far more complex, but they really are of very little value in this context. A digital image is made up of many hundreds – even thousands or tens of thousands – of rows of pixels and each row is hundreds, thousands, or tens of thousands of pixels long. So, basically, a digital image is made up of a two-dimensional grid of pixels, each pixel being a single dot of color.

Different digital devices can use different names for the same concept of “pixel.” For example, a printer’s ad copy will talk about resolution in terms of “dots per inch” (dpi) while a monitor might refer to “pixels per inch” (ppi). There are technical differences between the two, but essentially each refers to the number of individual points within a given linear inch that can be addressed. More of either is always higher resolution.

In the early days of the Windows operating system, a standard screen resolution was 640x480 pixels. That means that screens of that resolution would show an image made up of 480 rows of pixels, each of which was 640 pixels long. That means there would be a grand total of 307,200 pixels in that image, or what we would today call just shy of one-third of a megapixel (a million pixels). As can be seen, we know an awful lot about that image; however, we don’t know what physical size it is.

To know what size the image is, you have to know how many pixels per inch the device displaying it can handle. Standard monitor resolutions are historically 72 dpi (Macintosh) and 96 dpi (Windows). That means, on a 72 dpi monitor the 640x480 pixel image would be almost nine inches wide by six-and-two-thirds inches tall. On a 96 dpi monitor, the same image would drop to six-and-two-thirds inches wide by five inches tall. (This next statement is certain to

confuse the issue even more; however, in the interest of being somewhat accurate – even realizing that most of this is already leaving out a huge amount of important information – the concept of “dots” or “pixels” per inch in a monitor is sort of a false idea. Strictly speaking, a monitor doesn’t care about a “per inch” measurement; it works solely with the overall dimensions in pixels of the image. The whole idea of a “per inch” measurement was based on a logical inch, not a physical one, and relates to the way text is displayed on a screen.)

Monitors in use today generally display a much higher resolution than did those early GUI monitors; it’s rare to see a monitor today that is lower than 1024x768 pixels. However, 640x480 pixels is still a resolution you might use for a photograph to be displayed on a website. As can be seen from the numbers above, it would be a reasonably sized image on a 1024x768 monitor when viewing a website.

However, the image size changes dramatically when it is printed; suddenly the concept of a per-inch measurement becomes very physical instead of abstract. If the above-mentioned image is 640x480 pixels and your inkjet printer prints 150 pixels per inch, you can see that the image will print out at four-and-a-quarter inches by three-and-a-quarter (roughly) inches. Suddenly, your image that is nine logical inches wide is printing out at four.

It gets worse on the printing press, which needs 300 dots per inch of resolution. Now, the nine-logical-inch image has dropped to just over two physical, magazine inches wide – less than one column, and entirely too small to be of any use.

In terms of image quality, the more pixels you start with, the better off you are. However, the trade-off comes in terms of file size. The more pixels that

Continued on Page 61.

Resolving Resolution

Continued from Page 49

are contained in an image, the larger the file size of that image. It is not at all unusual for images to be in the forty- to fifty-megabyte range when they leave my office to go to the printing press; however, if you had a web page with a dozen or so fifty-megabyte images, no one would ever see your page because no one would ever wait for all of the images to load.

It is easy to take a large image and make it smaller; it is much more difficult to take a small image and make it larger while still retaining quality. To make an image smaller, the software looks at groups of pixels – there are different models that are used for selecting groups – and averages that group into a single pixel. To simplify somewhat, software might look at a grid three pixels by three pixels (with each pixel, remember, being an individual point of color) and average those nine distinct colors into a single pixel of color. It would continue stepping through the entire image in three-by-three grids and doing this until the entire image had been greatly reduced in size. This is, generally, speaking, an effective way of making an image smaller.

However, it does not work as well when you go the other direction. This is because when making an image larger, the computer is forced to create new pixels based on the averages of surrounding pixels in the area where the new pixel is being created. Instead of working with actual image data, the computer is creating what seems mathematically probable based on the existing image data. Subtle gradations in color are quickly lost, as is detail and image sharpness.

Best practice is to take digital photographs on the highest resolution setting that a given camera offers and then to store those original images (perhaps on a CD or DVD) and work with copies of the originals. This way, if you want to put a given image on your website, you can size the duplicate appropriately for that application while retaining the original at a much higher resolution should you later want that same image to include in printed material.

My personal practice is to save all digital images that I take on external hard drives (a one-terabyte drive now costs under \$100) organized in folders named so that the state abbreviation comes first, followed by the city, followed by the studio name, followed by the installation name. For example, I recently photographed some windows by the Pickel Studio that will be a part of the stained glass tour in Kansas City at the SGAA's Annual Summer Conference this year. Those pictures are in a folder named /KS-lenexa-pickel_studio-holy_family/.

This works for me because the geographic location is generally the most important factor when searching for an image, followed by studio name, followed by installation name. If most of

STAINED GLASS

QUARTERLY OF THE STAINED GLASS ASSOCIATION OF AMERICA

The Stained Glass Quarterly, the Stained Glass Association of America, and the SGAA Stained Glass School are pleased to announce a whole new generation in the Association's online presence. Not simply an evolution, this is a revolution and a radical step forward in the way the SGAA serves the stained, decorative and architectural art glass communities both now and in the future.

This plan will be implemented over the course of the coming year. To glimpse the future now, visit The Stained Glass Quarterly at its new website today. You'll find us online at:

StainedGlassQuarterly.com

the images in your image collection are of your own work, then you might want to start with the installation name because that would likely be far more important to you when searching for a given image. Another option would be to start with the date; for example, /20100219-holy_family/ would give a date that the computer would organize in order chronologically no matter how many folders were in a given directory, as long as you always follow the naming scheme of a four-digit year followed by a two-digit month followed by a two-digit day compressed into a single eight-digit number – use the biggest time measurement to smallest.

Remember: a pixel is an individual point of color. Digital images love diversity; the more distinct individual points of different color they have, the better they are. Really, we could learn a lot from pixels.

Next time, we'll look at how you can tell what the resolution of a given image is and how you can put that knowledge to work for you, plus we'll look at some of the dangers (and advantages) of emailing photographs at any resolution. I really mean it this time.