

Coming to Terms with DPI, PPI and Size

By Mike Chaney

Used with Permission

Background

I get a lot of questions about DPI, PPI and image size, enough to indicate that a lot of people are having trouble grasping some of these concepts when applied to their images. In this article, we'll try to make the terms and concepts of DPI, PPI, and physical size a little less of a mystery while at the same time offering some simpler methods of dealing with image/file size and resolution.

Understanding the terms

DPI or "dots per inch" is a term that normally describes the resolution of a printer since printers produce colors by placing colored dots on a piece of paper. A printer that has a resolution or DPI rating of 2400 can place 2400 individual dots of color within a one inch span on the page. If the printer is listed as a 2400 x 1200 printer, the printer can place 2400 dots in an inch across the page and 1200 dots in an inch down the page. A 2400 x 1200 printer can therefore place about 3 million dots in a 1 inch x 1 inch square on the paper. Resolutions of 2400 x 1200 DPI, 5760 x 1440 DPI, etc. are typical for inkjet printers.

Dye sub printers are usually between 300 and 500 DPI. Why the difference? Dye sub printers can produce continuous tone color, meaning that every dot on the page can be an arbitrary color of up to 16.7 million possibilities. With an inkjet printer, the printer can only "spit" a few different colors (up to maybe 32 in some cases) for each dot, so much smaller dots are used in combination to simulate continuous color. A dye sub printer might place a single dot of gray on the page for example, while an inkjet printer might lay down 32 intermingled dots (16 black and 16 empty/white) in the same area in an 8x4 pattern to "simulate" gray by using half black dots and half white dots. Let's not get too caught up in printer technology, however, and just remember that DPI refers to a printer's capability to place individual dots on the page.

PPI or "pixels per inch" is a term that most often indicates the resolution of an image file (a JPEG, TIFF, etc.) when that image file is displayed/printed at a certain size. For example, if you take a photo with your 6 megapixel camera that produces a 3000 x 2000 pixel image, that 3000 x 2000 pixel image will be 300 PPI *only* if displayed/printed at a size of 10 x 6.67 inches, 150 PPI if displayed/printed at a size of 20 x 13.33 inches, and so forth. The bigger you display/print that 3000 x 2000 image, the less PPI you have to work with because the same number of pixels (6 million) is being spread over a larger area. As you can see, PPI (resolution) and image size are directly related. With respect to resolution (PPI) and image size, increasing one will always decrease the

other, and vice versa because there are always 6 million pixels in your 3000 x 2000 images from your camera and the number of captured pixels does not change.

One Size Fits All

When you look at an image that was captured using your 6 megapixel (MP) camera, it will most likely be stamped with an arbitrary resolution such as 300 PPI. A 3000 x 2000 pixel image from a Canon 10D camera for example has 300 PPI recorded in the image file as its "resolution". With 3000 x 2000 pixels in the image, 300 PPI only equates to one size: 10 x 6.67 inches. When you open this image in your photo editor, it will tell you that the image is 3000 x 2000 pixels, 300 PPI, and its size is 10 x 6.67 inches. Does that mean that you should print all images from that camera at 10 x 6.67 inches? Does it mean that 10 x 6.67 is the *maximum* size you can print?

Of course, stamping an image file that came from a camera with a physical size is completely arbitrary and an almost "backward" way of thinking because the scene that you photographed with the camera most likely wasn't 10 x 6.67 inches nor is it likely that you will always print all your photos at 10 x 6.67 inches. So why put information in the image file that gives it a physical size of 10 x 6.67 inches? The answer is that 300 PPI is just a general guideline often used in the industry for "minimum resolution for true photo quality". In reality, it is possible to print excellent quality photos far above and below the 300 PPI "photo quality" threshold, so the resolution/size that you see in an image when you open that image in a photo editor is rather arbitrary and should not be considered a magic number.

For most workflows involving editing and printing images, an image should not have a physical size until that image is displayed or printed on a device and you shouldn't need to worry about size or resolution until print time. You can't change the number of captured pixels in the image anyway, so resolution is determined (automatically) by size at print time. Scanners are the exception to this way of thinking since they scan an object with a known size at a chosen resolution and you may wish to reproduce scanned material at the same (scanned) size. Having to deal with image resolution/size up front when working with photographs from a camera, however, is a major point of confusion for many people because typical photo editing software will tell you that your photograph is 10 x 6.67 inches and you begin to wonder what you have to do to print a 4x6, 5x7, 11x14, or some size other than the size indicated in the image file. As a result, a lot of people end up resampling (interpolating) their images to 300 PPI at the new/desired print size as the first step, then editing them, and finally printing. Such a workflow can and often does cause image degradation due to applying the interpolation at the wrong point in the process. Interpolation to a desired print size/resolution should be the last step in the workflow.

Instead of working backwards from an arbitrary size of 10 x 6.67 inches and 300 PPI and then trying to figure out what you have to do to print the size you want, it makes more sense to simply remember that you have 6 million captured or "original" pixels to work with: an image that has 3000 x 2000 pixels. Physical size (inches or cm) is not

something that should be dealt with until you determine what size you need at display/print time. Leaving the image at its original resolution (3000 x 2000) and resizing it only at print time allows you to edit one copy of the image and then create different print sizes from the same original. Just remember that the larger you choose to display or print, the lower the resolution (PPI) that 3000 x 2000 image will be. If you print your 3000 x 2000 image at a size of 6x4, you'll have 500 PPI of data available in the original image because $3000 / 6$ is 500. If you print the same image at 24 x 16 inches, your 3000 x 2000 pixels will have to be stretched across a much larger space and you'll only have 125 PPI from the image available at that size.

If 300 PPI is generally considered good enough resolution to reproduce printed photographs, 500 PPI is more resolution than you need so your 6x4 will look fine. Does 125 PPI mean that you cannot print a 24 x 16 inch print because that is so much lower than the 300 PPI considered "photo quality"? No. Not necessarily. What it means is that the image probably will not be quite as sharp at 24 x 16 as it is at 6x4 when viewed closely. Resampling that 3000 x 2000 image to a higher resolution before printing can make that 125 PPI available in the image go a lot farther. Let's see how we can make use of all this information in the most effective way.

Making the most of it

The key to understanding how to deal with resolution and size is to realize that there is only one thing that is stable and will never change with regard to your photograph that came out of the camera: it has a given resolution and that resolution is static and will not change. If you have a 6 MP camera that outputs a 3000 x 2000 pixel image file, that's the data you have to work with: 6 million pixels. How you choose to "spread out" those pixels on screen or in print is up to you. Print the image at a small size like 6x4, and you'll get a very sharp print. Print it at a larger size such as 10x8, 14x11, or more, and your print will begin to reduce in sharpness and detail, but depending on the methods used to print, you can often get very good photographs down to 150 PPI (20 x 13.3 with a 6 MP image) and even lower. The key is in *how* you stretch your 6 million pixels to cover a larger area.

As an example of "stretching", you could take your 3000 x 2000 image and upsample/interpolate the resolution to 6000 x 4000 before you print at that 20 x 13.3 size and in doing so, increase the image resolution to 300 PPI. Interpolation is a way of "extrapolating" data between pixels to create more pixels than were recorded in the original. While this does not add any real data to the image in the sense that only 6 million pixels were recorded, a good interpolation algorithm can predict what a higher resolution image might look like and can reduce artifacts such as jaggies. Click on [this link](#) to view how various interpolation methods can improve image quality. Note that the "pixel resize" version is what the small image on the upper right looks like displayed at the larger size. By using different interpolation methods, we are able to reduce or eliminate the pixelization or "jaggies" that occur from displaying a low resolution image at a large size. Also note that while the images are improved by interpolation, they are no match for the original image on the top/left that was photographed at the higher

resolution to begin with. Needless to say, there is no substitute for *real* data, but interpolation can improve things quite a bit when printing at large sizes! As such, resampling (interpolating) can improve print results when printing large sizes that cause resolution to drop below 300 PPI.

Given the fact that our original image (photograph) contains all the pixels that the device could record and therefore the maximum amount of data available to you, I would recommend doing all work such as color changes, cloning out blemishes, and even any sharpening needed because the original appears too soft without changing the size or resolution of the image. Remember that interpolated/resampled pixels are based on the original/captured pixels so resampling first and then editing only serves to increase the pixel count of the images you are working with. It makes more sense to edit an image that is 6 million pixels and then interpolate to 12 million pixels than to create 12 million pixels up front and have to edit twice as many pixels.

Once any needed changes have been made to the original image, the image can then be resized or interpolated to the desired resolution for printing. If you are using a photo editor, you can enter the desired print size (say 11x14), enter a resolution (PPI), and check "resample" so that the photo editor will interpolate the image to the chosen size and resolution. Size is often the easy part because you know the size you want to print. For resolution, you want to use a multiple of the actual/physical printer resolution. For Canon/HP printers that would be 300 PPI for typical photos or 600 PPI for optimal quality with the finest details. For Epson printers, 360 PPI for photo quality or 720 PPI for photo quality with the finest details possible. If you use a dye sub printer, always resample to the "native" resolution of the dye sub printer, 314 PPI, 320, 480, etc. Note that depending on how effective your print driver is at stretching (or shrinking) the image to fit on the paper, you may get better results if you always resample to a multiple of the printer resolution, even if you start with a *higher* resolution than needed. For example, when printing to a dye sub printer that runs at 314 DPI, the 6x4 print from that 6 MP camera will likely print better if you *downsample* the 500 PPI image down to 314 PPI before sending it to the printer! In this case, sending an image to the printer at 314 PPI actually produces better results than sending it at the higher (but mismatched) 500 PPI!

Of course, it's always nice to have a tool that will do all these calculations for you so that you never have to worry about anything other than making color, levels, and a few other adjustments on the original, allowing the software to automatically interpolate to the best possible quality at print time. Qimage is such a tool that will allow you to correct the original image (in Qimage or using a photo editor) without ever having to worry about DPI, PPI, or size. Simply make any changes you like on the original without modifying the size or resolution of the image and Qimage will handle all the sizing and PPI calculations when you tell it the size you would like to print, and will automatically resample to the optimal resolution for your printer at print time using advanced interpolation and sharpening methods. While I did write Qimage and take the opportunity to plug it here, it does allow you to work with images in a much simpler and much more "forward" approach where you edit the original pixels and let the software worry about size, PPI, and DPI at the appropriate time: when you are ready to print. It is probably the easiest way to avoid the backward mentality of dealing with images that

have arbitrary size and resolution stamps that can often prompt users to deal with resolution/size at inappropriate times.

In Summary:

I hope this article will assist in the understanding of PPI, DPI, and image size and will offer an approach that is easier for most to work with. Photographic images such as JPEG and TIFF images come with a size/resolution stamp that can be useful when doing things like scanning prints and reproducing them at the same size. When scanning a given size media in a scanner, recording the original size of the media can be useful because it tells us how to reproduce that media at the same size. If you want to reproduce the scanned media at an arbitrary or different size, however, or print a photo that came from a camera, such resolution/size recordings in the image become arbitrary and often confuse the issue. When you are capturing images at the resolution limit of a device and/or you intend to reproduce those images at an (arbitrary) size of your choosing, it is best to just leave the images at their native/captured resolution and only interpolate/resample the images if necessary at print time since that is the only time at which assigning a physical size to an image makes sense.

Resampling images as a first (or early) step and then editing them is a common mistake that can cause problems when working with images that you'd ultimately like to be able to reproduce at more than one size. For example, resampling an image to 300 PPI at 11x14 and then editing color and other aspects of the image can cause trouble if at a later time you would like to print some 16x20 prints from the edited image. The edited image has already been resampled to 300 PPI at 11x14 and resampling it again to 300 PPI at the new 16x20 size means that the image has gone through two resampling/interpolation steps which can degrade the image and not give you the best possible results.