

# Print Resolution: How Many Pixels per Inch?

*Phil Service*  
*Flagstaff, Arizona, USA*  
*8 September 2014*

## Summary

Matched pictures were taken with a Nikon D90 and a Nikon D7100. Both images were upsampled for printing at 20 x 13.3 inches (51 x 34 cm) and 360 pixels per inch (ppi) on an Epson Stylus Pro 3880. PhotoZoom Pro 6 (BenVista, Ltd.) was used for upsampling. The resulting prints from the two cameras were nearly indistinguishable with respect to amount and quality of fine detail. As a “control”, prints were also made with upsampling left in the background, under the control of the printer driver or Lightroom. The resulting prints were quite good, but noticeably lower-quality than those made from images upsampled with PhotoZoom Pro. The “native” resolution of the D90 image at 20 x 13.3 inches is only 214 ppi. These results argue that the frequently cited 300-pixel-per-inch rule for high quality ink jet printing is overly conservative when using current software and hardware.

**Key words:** print resolution, resampling, Nikon D90, Nikon D7100, Lightroom, Photoshop, PhotoZoom Pro, S-Spline Max, Epson Stylus Pro 3880, Eric Chan

## 1. Introduction

Conventional wisdom says that linear resolution divided by 300 gives the largest (in inches) good or excellent quality print that can be produced from a digital image. For example, the largest high-quality print that can be made from a 6000 x 4000 pixel image is 20 x 13.3 inches (51 x 34 cm)<sup>1</sup> I will refer to this as the “300-pixels-per -inch rule”. As is the case for most “rules”, not all choose to follow — some feel that considerably lower resolutions are sufficient. And, as is also often the case, the reasoning behind this bit of conventional wisdom is seldom given.

My goal in this paper is to suggest at least a plausible basis for the idea that 300 ppi resolution is a necessary requirement for high quality digital printing. The short answer, I think, is that it represents a compromise between maximizing image *sharpness* and preserving *detail*. The longer answer will require us to delve into the concept of sharpness, to review what is known about the limits of human visual resolution, to discuss image resampling, and to consider the details and limits of current ink-jet printing technology. By the end, and with an example, I

---

<sup>1</sup> The obvious corollary is that if you want to make 20-inch prints, you need to use a camera with sensor that is at least 6000 photosites long. According to this rule, the largest high quality print size for a 12MP APS-C camera would be 14 in (36 cm).

hope to convince you that the 300-pixel-per-inch rule is overly conservative in many, if not most, cases.

## 2. Sharpness, Detail, Human Vision, and Image Size

### 2.1. Smaller is Sharper

We have probably all had the experience of coming across an interesting, apparently sharp, image on the web — and then clicked to view it at larger size only to discover that it has become disappointingly soft. The lesson is that apparent sharpness depends on the physical size of an image — that is, either the size at which it is printed or the size at which it is displayed on a computer screen.<sup>2</sup>

To appreciate the importance of physical size on apparent sharpness, let's look at a not-very-sharp image. Fig. 1 shows portions of an image taken with an iPhone 4S under less than optimal conditions. The right side displays the image at 100% (109 ppi on my display). The left side is a crop of the image after it was resampled in Photoshop to simulate a resolution of 450 ppi. No additional processing was done, although the resampling necessarily involved some sharpening.<sup>3</sup> Clearly, smaller is sharper.

In this context, sharpness is principally a quality of edges, or borders, of objects in an image. We perceive an image as sharp if the edges of “hard” objects are clearly (dare I say “sharply”) defined. The blurred profile in the right-hand image of Fig. 1 becomes acceptably sharp in the smaller left-hand image. In essence, reducing the image size also reduces the width of blurry edges. Eventually the blur becomes too narrow for us to see, or is eliminated altogether in the resampling that reduces pixel dimensions and, hence, displayed size.<sup>4</sup>

---

<sup>2</sup> I will try to be consistent here. I will use **size** to mean the physical dimensions of an image, usually expressed in inches or centimeters. The pixel dimensions of an image are its **resolution**, although it is very common to use “size” in this context as well. Even this meaning of “resolution” is controversial. Some would argue that resolution, whether of images or camera sensors, should include a unit of length — pixels per inch or photosites per mm, for example. I will use “resolution” both to mean the total number of pixels in an image and also the number of pixels per inch of a displayed image. It is worth remembering in this context that an image pixel, as opposed to a sensor photosite, has no intrinsic size. Therefore, a digital image has no physical size until it is displayed on a monitor or printed.

<sup>3</sup> Note that both halves of the image are actually being displayed at the same resolution — the resolution of whatever monitor you are using. The “reduced” image was created in Photoshop by reducing its resolution (i.e., number of pixels) so that its physical size would be smaller when viewed on a computer display.

<sup>4</sup> If you are worried that some sharpening trickery is involved in resampling, the effect of reducing image size without resampling can be simulated by viewing an image at greater distance. Again, blur will become less perceptible as we move farther away and the image appears smaller.



Fig. 1. Right side: a portion of an image displayed at 100%. Left side: the same image reduced to approximately 1/4 “native” size.

Typical sharpening methods are consistent with this notion of sharpness as an edge phenomenon. They increase the definition of edges by slightly lightening pixels on the lighter side of a border and/or slightly darkening pixels on the darker side. The effect is to make edges more clearly defined by increasing *contrast* across the edge. That is, transitions become more abrupt, or less fuzzy.

## 2.2. Image Size and Detail

As images are displayed at smaller and smaller sizes, details become more difficult to discern. Image reduction for display on computer monitors necessarily involves loss of information. The fixed resolution of the monitor requires us to reduce the pixel dimensions of the image in order to show it at smaller size. The un-cropped “parent” for the right-hand image in Fig. 1 had a resolution of 1973 x 2821 pixels. The left-hand image is a crop from an image with 478 x 673 pixels. We “discarded” fully 94% of the pixels in the original in order to make the reduced image. Of course, the resampling method doesn’t really “discard” pixels — it creates a reduced number of entirely new pixels in a presumably “intelligent” way.

When we *print* images at smaller sizes, we don’t necessarily need to reduce the pixel dimensions.<sup>5</sup> However at some point, details will either become so small that they will require a magnifying loupe in order to be seen, or they will become too small to be printed.

To summarize the discussion thus far, images appear sharper when they are displayed at smaller sizes. However, there is a tension between sharpness and image detail, because detail tends to be “lost” as images become smaller.

## 2.3. Sharpness and Detail

Details, to the extent that they are large enough to have meaningful edges, can also be more or less sharp in the same way as larger objects in an image. When describing sharp details, we often use adjectives such as “crisp” or “clear”. In other words, the edges of the detail are well-defined. Even a single pixel can represent a detail, albeit without an edge. The “sharpness” of a single-pixel detail is enhanced by increasing the contrast between that pixel and the surrounding pixels. Because we are altering local contrast, the procedure is similar to sharpening of edges.

## 2.4. Human Visual Acuity

At a close viewing distance of 25 cm (about 10 in), a person with normal vision can resolve black and white line pairs at a frequency of about 6.88/mm.<sup>6</sup> Assuming that each line *pair* is two pixels wide (*i.e.*, the minimum possible), that resolution corresponds to 350 ppi. My Epson 3880 can print a line pair pattern at 7 lp/mm on glossy paper. The pattern is very clear

---

<sup>5</sup> When an image is displayed at 100% on a monitor, there is a one-to-one correspondence between image pixels and display pixels. That one-to-one correspondence still holds when the image is displayed at a smaller size. It’s just that the number of image pixels has been temporarily reduced on the fly, in the background, by an application or by the operating system. The situation with printers is more complex, and will be discussed later. In general however, printing at smaller sizes can require either decreasing or *increasing* the number of pixels in an image as an intermediate step (typically done in the background by the printing software, or the application used for printing, such as Lightroom).

<sup>6</sup> Koren, Norman. <http://www.normankoren.com/Tutorials/MTF.html>

when examined with an 8x loupe.<sup>7</sup> With naked eye, good lighting, close viewing and some effort, I can tell that there is a pattern of alternating dark and light lines. Thus, I am willing to take Koren's figure of 6.88 lp/mm as essentially correct.

## 2.5. Synthesis

We can now see why the 300-pixel-per-inch rule represents a compromise between maximizing image sharpness and preserving detail. 300 ppi corresponds to 5.9 lp/mm. This is just slightly above the threshold for visible detail. If we print at a significantly higher resolution, say 420 ppi, the image will be smaller and it might be sharper, but very fine detail will be lost to view (although it might still be present in the print). On the other hand if we could print at a lower resolution, such as 240 ppi, details will be easier to see in the larger image, but the image may also appear less sharp.<sup>8</sup>

## 3. How Ink Jet Printers Work

Much of the material in this section is adapted from information provided by Eric Chan.<sup>9</sup> Chan addresses in considerable detail the workings of the Epson 3800 printer. I will assume that most of what he describes about the 3800 applies to the Epson 3880 that I use, and is also generally applicable to most, if not all, current inkjet printers.<sup>10</sup>

Manufacturers like to tout the resolution of their printers in terms of dots per inch (dpi). The Epson 3880, for example, has maximum horizontal and vertical resolutions of 2880 and 1440 dpi, respectively. However, the most important information for photographers is the printer's "native" resolution in image pixels per inch (ppi). Unfortunately, that information is

---

<sup>7</sup> In fact, the pattern is clear only when the lines are parallel to the direction of print head movement. I refer to this as vertical resolution because the line pairs would be counted vertically (in the direction of paper movement through the printer). When line pairs run at a right angle to the direction of head movement, the pattern is much less clear, although discernible with a loupe. In that case, the line pairs are counted horizontally and resolution is horizontal. Thus, if we are counting line pairs, vertical resolution is superior to horizontal resolution. *This is just the opposite of the the customary statement by printer manufacturers that horizontal resolution (in dpi) is superior to vertical resolution.*

<sup>8</sup> Contrary to some suggestions about the necessity for 300 ppi resolution, I do not think that "pixellation" is likely to be an issue even at resolutions well below 240 ppi, for the reason that image pixels are not simply enlarged when printed at low resolution. As will be explained below, printer software will upsample a low-resolution image prior to printing — that is, increase its number of pixels.

I am well aware of the usual arguments about printing larger images at lower resolution under the assumption that they will be viewed at greater distance, making them appear sharper and making details more difficult to see. I make prints to hang on the wall of my home. I assume that they will be closely scrutinized, regardless of their size.

<sup>9</sup> <http://people.csail.mit.edu/ericchan/>

<sup>10</sup> See also this article by Jeff Schewe. <http://www.digitalphotopro.com/technique/workflow/the-right-resolution.html?tmpl=component&print=1#.VBXReEue939>

more difficult to come by. According to Chan, the Epson 3800 has two native resolutions: 360 and 720 ppi.<sup>11</sup> Here is the fundamental point:

*No matter what resolution you use for sending your image to the Epson 3800, it will be converted to 360 (or 720) ppi before printing. The significance of this cannot be overstated. Unless you are willing to print your image at whatever physical size happens to correspond to the native resolution of the printer, the printing workflow will **inescapably involve image resampling**. You can do the resampling yourself and so control it and view its effects before printing, or you can trust the printer driver (or Photoshop) to do it for you in the background. Printers from other manufacturers appear to work the same way.*

*Resampling*, also known as interpolation, means increasing or decreasing the *number of pixels* in an image.

## 4. Image Resampling

### 4.1. A Sample Calculation

Imagine that we have a 6000 × 4000 pixel image file, and wish to make a 20 × 13.3 inch (51 × 34 cm) print. The “native” resolution of the image printed at this size would be 6000 / 20 = 300 ppi. That seems pretty good, but in fact the printer’s native resolution is 360 ppi. Therefore, the image will be upsampled before printing to 7200 × 4800 pixels. It is common practice in resampling software to refer to this as a 120% upsampling. That is, the linear pixel dimensions of the resampled image are 120% larger than the dimensions of the original. The total number of pixels in the upsampled image would be 144% greater than in the original — 34.56MP vs 24MP.<sup>12</sup> This is a relatively modest amount of upsampling, and presents no great problems. In order to avoid resampling, it would have been necessary to print the image at 16.7 × 11.1 in (42 × 28 cm) — or to have used a printer with a native resolution of 300 ppi.

In this example, we started with a 24MP image file and found that the necessary upsampling was relatively minor. What about a 4288 × 2848 image from a Nikon D90? Its native resolution when printed at 20 × 13.3 in would be only 214.4 ppi. Upsampling to 360 ppi for printing at that size represents a 168% resample. The total number of pixels in the image would almost *triple*, increasing from 12.2 to 34.56 million. One might well wonder about the quality of the upsampled image. I hope to convince you below with an example that it can be very good.

### 4.2. Software for Resampling and Workflow

The details of various resampling methods — such as bicubic, Lanczos, S-spline, etc. — are far beyond the scope of my expertise and this paper. There are a number of resources on the

---

<sup>11</sup> In light of what we know about human visual acuity (Section 2.4), the choice of 360 ppi for printer resolution is clearly not arbitrary. Canon and Hewlett-Packard printers apparently have native resolutions of 300 (and 600) ppi.

<sup>12</sup> I am using the terminology as it is commonly encountered in resampling software. Less ambiguously, an increase from 6000 to 7200 pixels is a 20% increase:  $(7200 - 6000) / 6000 = 0.20$ .

web. My only comment here is that no one method will work best for all images. Also, there are quite a few software choices above and beyond the resampling features of Photoshop. A non-exhaustive, OS X-centered list includes Blow Up 3, SizeFixer, PhotoZoom Pro 6, and Perfect Resize.<sup>13</sup> All are available as free trials. Some work as plug-ins, some as stand-alone applications, or as both. I currently use Photoshop CS6/CC and PhotoZoom Pro 6. The latter, unfortunately, appears to be one of the more expensive options. My advice would be to try several before buying. Note also that Photoshop CC has somewhat more capable resampling features than CS6. I like Lightroom for initial development of raw files. However, the only resampling capability in Lightroom appears to be in the Export dialog, for example exporting images to TIFF, and there is no choice of method.

From the point of view of workflow, the first issue will be whether to resample yourself, or to leave it as a background printing process. I would guess that most people who are serious about fine art printing would want to resample prior to printing. The obvious reason is that one can see and adjust the results beforehand. If resampling is an explicit step, where should it occur in the workflow? As far as I can tell, the only hard and fast rule is that **no, or very minimal, sharpening should be applied before resampling**. The reason is that sharpening can introduce artifacts, and we wish to avoid resampling those artifacts. A simplified workflow for me would start with raw development in Lightroom. Adjust exposure, color balance, highlights, shadows, and tonality; remove distortion, chromatic aberrations, and noise, being careful with the latter. Set sharpening to amount to 0. Export the image in its native resolution as a 16-bit TIFF file in ProPhoto color space. Resample the TIFF image to the resolution necessary to print at the desired size at 360 ppi. Resampling is done in PhotoZoom Pro or Photoshop. Either way, the resampled image will receive final editing in Photoshop or Lightroom. Final editing will include sharpening and possibly some other (hopefully) minor adjustments, such as additional noise removal. Printing is done from either Lightroom or Photoshop.

If you choose to leave resampling as a background printing process, it will be done either by the printer software or the application that you print from. In the case of Lightroom 5, Lightroom will resample if the Print Resolution box is checked in the print module. If one takes that option, the obvious setting would be the printer's native resolution.

### 4.3. An Example and Comparison: Nikon D90 vs. Nikon D7100

For a project not related to resampling, I took pictures with a Nikon D90 and D7100. The images were made a few minutes apart, from the same tripod location, and with the same lens, a Nikkor AF-S 50mm f/1.8G, which was moved between cameras. EXIF data are shown in Table 1. The uncropped images are shown side-by-side in [Fig. 2](#), downsampled in Photoshop with Bicubic Sharper to fit on a 2560 × 1440 pixel monitor at the same size. (In order to show

---

<sup>13</sup> Blow Up 3 is from Alien Skin Software, <http://www.alienskin.com/blowup/> ; SizeFixer is from Fixer Labs, <http://www.fixerlabs.com/EN/upsizer/sizefixer.htm> ; Perfect Resize is from onOne Software, <http://www.ononesoftware.com/products/resize9/> ; PhotoZoom Pro is from BenVista, <http://www.benvista.com/photozoompro> .

Table 1. Test Image EXIF Data		
	Nikon D90	Nikon D7100
Shooting mode	A	A
Metering	"Pattern"	"Pattern"
Exposure compensation (EV)	-1.0	-1.0
Aperture	f/4	f/4
Shutter	1/400	1/200
ISO	200	100
Time	8:18:24	8:14:31
Date (yyyy-mm-dd)	2014-06-22	2014-06-22
Image quality	12-bit NEF	14-bit NEF (lossless compressed)

images at necessary sizes, they are provided as separate files that can be opened by opening the links associated with “Fig. 2”, “Fig. 3”, etc., in the text.) The images in [Fig. 2](#) are unsharpened and at the point at which they were exported from Lightroom prior to upsampling. The scene is challenging: the EV range is high, and there is a great deal of fine detail — for example in the bark and exposed, weathered wood of the juniper, in the rocks at the lower left, in the forest floor litter in the sunlit lower right, and in the small bare twigs to the right of the tree trunk. The sun angle changed slightly between frames, which complicates the comparison. Unresampled detail from each image is shown at actual pixels in [Figs. 3, 4, and 5](#). These are crops from the full-resolution images made at the same point in the workflow as the images in [Fig. 2](#). Most importantly, that means that they are *unsharpened*.

Both images were exported as native-resolution TIFFs from Lightroom, and were upsampled in PhotoZoom Pro 6. The chosen method was S-Spline Max. For the D7100 image, I accepted the default sharpening settings for “Photo – Extra Detailed”. For the D90 image, I used the same method and settings, except for slightly increased sharpening. Both upsampled images were 7200 pixels high (= 360 ppi x 20 in). The upsampled D7100 image was 4800 pixels wide; the D90 image only 4783 pixels because the proportions of the D90 NEF files are not precisely 3:2. The upsampled images were opened in Photoshop CS6 for final sharpening and other minor adjustments.

A series of full-resolution crops of the sharpened, upsampled images are shown in [Figs. 6, 7, and 8](#). These crops correspond to those in [Figs. 3, 4, and 5](#), although they necessarily show smaller portions of the entire images. Brief inspection will reveal that the upsampled D7100 crops look better than the D90 crops when viewed at 100%. That result is hardly surprising given the finer grain of the original D7100 image and the fact that its total pixel count was increased by a relatively much smaller amount. However, the comparison that counts most, in fact the only one that counts at all, is the comparison of the 20-inch prints. The upsampled

images when viewed at 100% on a typical 109 ppi computer display would be 66 in, or 5.5 ft (1.68 m) tall, if viewed in their entirety. As we have already seen, smaller is sharper, and the images will be much smaller when printed to 20 in high. I cannot show the actual prints here. But the prints can be *simulated* by showing the images at 20 in high on a computer monitor. Both upsampled images — at the stage at which they were originally exported from PhotoZoom Pro — were downsampled, also in PhotoZoom Pro. The downsample resolution was set to 109 ppi, which reduced the pixel dimensions to 2180 x 1453 (or 1449), and the same method was used for both images: S-Spline Max with Photo - Detailed default settings. Crops of the simulated prints are shown in Figs. [9](#) and [10](#). I would argue that there is very little difference to be seen between the D90 and D7100 simulated prints. Certainly, the D7100 image is not obviously superior.<sup>14</sup>

The acid test, of course, is the actual prints. Printing was done from Lightroom after a few final adjustments to tonality. Images were printed uni-directionally at 360 ppi and 2880 dpi on an Epson Stylus Pro 3880, using Epson Premium Semi-Gloss Photo Paper. Print sharpening was disabled in the Lightroom print module. If for no other reason than a reality check, I also printed the un-upsampled D90 image after sharpening in Lightroom. Two prints were made: one in which the upsampling was ostensibly left completely to the Epson printer software and print sharpening was set to Standard in the Lightroom print module; and one in which print resolution was set to 360 ppi in the Lightroom print module and print sharpening was set to Low. Upsampling was presumably done by Lightroom in the latter case.

#### 4.4. Evaluation of the Prints

You will have to take my word for this. First, the easy part. Both prints made from the D90 image that I did not upsample were surprisingly good. As far as I could tell, they are as near identical to each other as is possible. In isolation, I think that they would receive very favorable evaluations. On careful inspection, the print made from the D90 image that was upsampled in PhotoZoom Pro is superior. The very fine detail is just better — particularly in areas of bare, weathered wood; and on the rocks along the bottom of the image. That said, the difference between my upsampling and the background upsampling is not obvious without careful examination. Most people, I think, would be very happy with the prints made with background upsampling.

Compared to the differences between the D90 prints made with different upsampling methods, the differences between the D7100 and D90 prints that were both processed through PhotoZoom Pro are much, much more subtle. For lack of a better word, I would say that the very finest detail in the D7100 print is rendered a bit more *delicately*, or smoothly. I am thinking here about the fine grain in the weathered wood of the trunk, and the smoothness of individual pine needles on the ground. The burned region at the base of the juniper trunk is about the only macroscopic region that looks better in the D7100 image, and that may be the only difference that can be seen between the simulated prints ([Fig. 9](#)). Otherwise, the superiority of the D7100 image can only be seen if prints are examined at very close distance and with strong light.

---

<sup>14</sup> When comparing images, either in print or on a computer, it is important to compare regions that are similarly lit. The portions of each image that received direct sunlight have reduced detail.

## 5. Conclusion and Summary

First, it's time to discard the 300 ppi rule. It may have been useful in the past, with older-generation printers and less sophisticated drivers and software. But, the example here shows that very good prints of highly detailed subjects can be made from images whose “native” resolution is as little as 214 ppi, even without the use of dedicated resampling software. I am not the first to suggest that lower resolutions, such as 240 ppi, may be quite adequate for high-quality prints. However, I think I have provided good support for that conclusion. Second, prints of the same size made from 12 and 24MP cameras can be almost indistinguishable if processed carefully. That is not to say that 12MP is “just as good” as 24MP. For one thing, the 24MP image could withstand greater enlargement. But for the example given here — a 20-inch tall print — the difference is virtually negligible.

### 5.1. Summary of Major Points

- The 300-pixel-per-inch rule is overly conservative for many, perhaps most, images.
- Printers have native pixel resolutions. Printers intended for high quality (fine art) printing have native resolutions that coincide with the approximate limits of human visual acuity. Individual image pixels are effectively imperceptible, unless viewed with a magnifier.
- Almost all images will be resampled for printing. The only exception is when the chosen print size results in an image pixel resolution that is the same as the native pixel resolution of the printer
- Resampling can be done by the printer driver, the application used for printing (such as Photoshop or Lightroom), or by special resampling software.
- Dedicated resampling software, such as PhotoZoom Pro, SizeFixer, etc., provides the largest number of resampling options and the greatest control over the results.
- Background resampling, for example by the printer driver, can produce very good results in many cases.
- It is not clear where the practical limits on upsampling lie. The general rule should be that “less is more”. Different images will respond differently to upsampling, and different resampling methods will yield different results. The example in this paper, with the Nikon D90 image, shows that almost tripling the total pixel count can still yield excellent results at 360 ppi print resolution.
- All other things — noise, dynamic range, etc. — being equal, more image pixels are better if upsampling is required for printing. That said, the particular example used in this paper shows that the improvement in print quality given by a 24MP camera over a 12MP camera is almost too little to be seen in a 20-inch tall print.
- In order to fully exploit the implied advantage of 24MP over 12MP, it would be necessary to print at larger sizes. The D7100 image, if upsampled to the same degree as the D90 image, would make a 28-inch tall print.