Part I

Essentials for Serious Image Editing

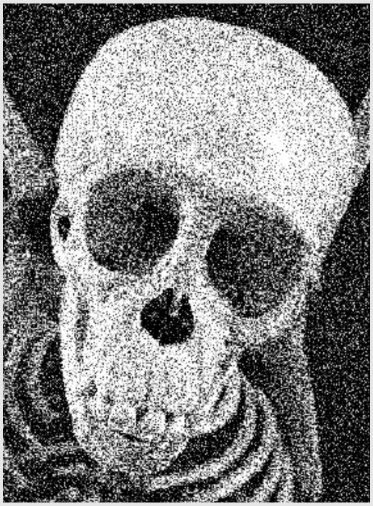

Serious image editing requires preparation and understanding. You have to be prepared with the best source images (the best content, resolution, and color), and you have to understand what corrections you need to make and how best to enhance the images. Problems in your images that you ignore when you are just starting out become more nagging as you gain experience. What you don't want to do when you work on images is waste hours correcting problems that would have taken moments to fix at the time of capture. And you don't want to waste time learning superfluous or redundant tools and techniques. Capturing the best information and getting the best results require understanding the images themselves and how image information is retained and displayed in print or on a monitor. An outline for the tools you should use helps you focus on the right techniques to get the best results with those images from the outset. This part of the book lays the groundwork you'll need for stepping into more advanced concepts later in the book.

Chapter 1

Resolution: The Cornerstone of Image Detail

Your first task in working with images is to always to capture the best image that you can. Starting with the best capture as a photo or scan, instead of one you will just plan to "fix later," will save you time and give you more image to work with. The better the information you start with, the more likely you'll have what you need to make the best result with less effort, in less time.

Images are built around resolution. There is a finite amount of detail in any pixel-based digital image that you get from a digital camera or scanner. *Resolution* is a measure of potential detail; the more resolution you have, the greater amount of potential detail in the image. High resolution suggests that there will be intricate detail; low resolution suggests that detail may be compromised.

While you will want to capture images at high resolution to retain detail, what you really want for images when they are applied is the *correct* resolution. We'll look at how to leverage resolution in this chapter.

In Terms of Resolution

Pixels are the smallest visual unit in an image. Each pixel (short for *picture element*) represents a single color or tone. The pixels are organized in rows and columns like blocks or tiles to map out the details, one tiny image pixel at a time. The multitude of pixels in an image (numbering more often in the millions for consumer digital cameras) blends together in our vision to create the look and feel of an image in tone (light and dark) and color (hue and saturation). The more pixels you have in an image, the greater potential detail.

This chapter has essential information that you need to know at any level of image editing. If you are a novice, read on. If you are an experienced image editor, you may want to read this chapter for review. If you are experienced and don't want the review, skip on to Chapter 2.

"Potential" detail is mentioned several times. While there is potential to retain more detail with higher pixel counts, limitations of your equipment, poor focus, bad lighting, and bad image handling can all lead to situations where detail is compromised no matter how many pixels you have in your images.

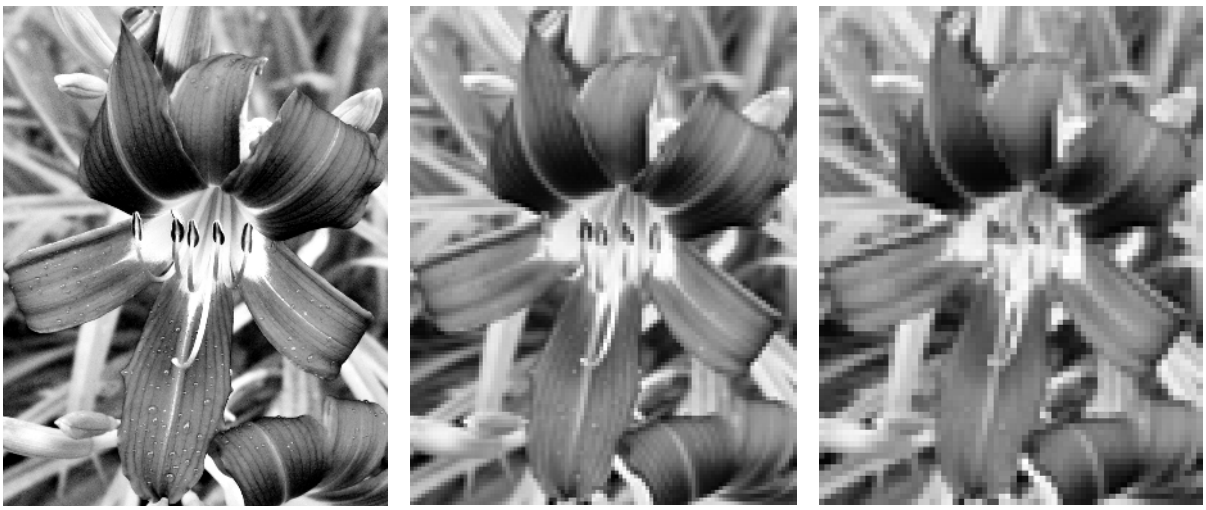

a b c

Figure 1.1

This flower is rendered here with the correct resolution for print in this book (a), one-eighth the proper resolution (b), and one-twelfth the resolution (c). Images with less resolution than necessary will appear with too little detail and become blocky.

Image resolution is usually measured in several ways, including the number of total pixels (e.g., a 5 megapixel camera), the size of the image file (number of bytes, kilobytes, or megabytes), the amount of information per inch (ppi, or pixels per inch), and the physical dimension (pixel dimension or ruler sizing). One way of measuring the image resolution is not necessarily better or worse than another, though some are more intuitive. In fact, the measurements serve different purposes, and none of them is sufficient on its own to really define how much detail is in an image. For example, ppi tells how much of the image pixel information should be applied per inch, but it doesn't tell how much information is in the image or the image size. Without those additional defining parameters, it is impossible to determine how big an image will be when applied.

Total pixels is the actual count of pixels (picture elements) in the image. As suggested, digital cameras are commonly rated in pixel count. Pixel count can also be a dimension of the pixel mapping, such as 2100 × 1500 pixels. An image with a greater pixel dimension or pixel count than another is generally larger in terms of potential detail than an image with fewer pixels. If your camera can shoot images that are 2100 × 1500 pixels, the images will have more detail than a camera that takes images which are 1600 × 1200 pixels. The potential for greater resolution (more detail) is why you will generally want to shoot with a higher resolution setting if your camera offers a choice. Images with higher resolution have more pixels and take more space to store, so you will get fewer of them on your camera's memory card, but the images you do get should have greater detail.

Greater pixel count in different cameras using different lenses and sensors can mean different things as far as sharpness in the final image. For example, using alternative technology such as Foveon X3 chips for image capture can render a sharper image than a standard mosaic Bayer chip that has three or more times the resolution in megapixels. The Foveon chip measures RGB for *each pixel,* while a common CCD (charge-coupled device) or CMOS (complementary metal oxide semiconductor) sensor measures color in an array and interpolates the result. Resolution is a good gauge of sharpness but not an infallible one.

Measuring image resolution in file size is probably the least intuitive or useful approach for most people working in Photoshop Elements. File size is used to describe the quantity of information in the source image, independent of how it is to be applied or stored. If you are storing more than one type of image (e.g., RAW and JPEG), file size becomes all but meaningless because different image types store information in different ways (and with different compression schemes). The color mode of the image (grayscale, RGB, etc.) and bit depth (8 or 16 bits) affect the amount of information per pixel. Image attributes (storage of layers, masks, and other image components) can inflate image size without changing the dimension or visual resolution. So file size is really useful only in comparing resolution of like files—those that are uncompressed, the same color mode, and the same file type—and before processing that will add image elements (layers, vectors, alpha channels, etc.). File size is likely the least-used measure of resolution.

Beyond file size, physical pixel dimension, and ppi, there are some interrelated terms for resolution that you will hear from time to time used both correctly and incorrectly (often in the same sentence). For example, *dpi* is a commonly used term, often used casually as a universal term for resolution. However, it is really an output term specific to printing and the number of dots a printer uses in representing an image. Other measures of resolution include spi, ppi, and lpi. The various resolution-related terms can be tricky to use correctly and consistently, but you should know what they mean and use them properly when you mean something specific.

> To simplify with better accuracy, use *spi* when speaking of capture (scan sampling), *ppi* when discussing digital files, *dpi* when considering output resolution, and *lpi* in the context of halftone dot size.

spi (samples per inch) Capture resolution. The number of scanning and digital capture samples per inch.

ppi (pixels per inch) Digital file resolution. The assigned number of digital pixel elements to be used in printing or displaying an image.

dpi (dots per inch) Printer resolution. The number of bitmap dots (smallest printing component) an output can create per inch.

lpi (lines per inch) A measure of halftone dot size. Halftone dots are made of multiple printer dots. The number of rows of halftone dots per inch.

It would seem that you would always want high resolution in your images if you consider detail important. But that's not always the case. Correct resolution depends not only on what size you want the result to be but also on what display or print medium you will be using and how the image information is applied.

For example, if you don't have enough resolution to meet the needs of print output, images won't look as sharp as they could; they might appear a little soft, fuzzy, or blocky (again, have a good look at Figure 1.1). If you have too much resolution in your image, file sizes are unnecessarily large; processing will take longer than it needs to; you'll take up excess storage; and the results will not improve. On the Web, images without enough resolution will be too small; those with too much resolution will be too large and will take longer to download. You can't just guess how much image information you need when applying an image; you have to know the amount you really need and work within those parameters. Understanding what resolution is and how it is used is the only way to use it correctly.

What Image Resolution to Use

Some people generalize and suggest using 300 ppi as a standard resolution for images going to print. For the Web, it is usually accepted that images should be 72 ppi. While these are pretty good as general-purpose guidelines, they don't tell the entire story of what resolution to use. For example, 300 ppi may be more than is necessary for all home printers, but it may actually be too little for demanding output (such as film recorder output). Because monitor resolutions can vary, your 72 ppi image on a 96 dpi screen would actually be about 75 percent of the intended size. Neither choice is likely to totally ruin your output, in most cases, but it can compromise what you expect.

Because output differs in how it applies image information, there is no one universal magic formula to figure out what resolution to use. Each output type has a target range (minimum and maximum), based on its capability to process and use image information. Once you know the range you need, you simply use that range as a target when working on an image. Know what your printing service or printer manufacturer recommends for output on the devices you use. This may require reading the manuals or calling the printing service to find out. The optimal range is the range where the image will perform the best when applied; it is possible to get acceptable results by going outside the range depending on how you implement the image and the results that you expect.

Table 1.1 shows the approximate resolutions to use for your images, depending on how you want to use them. An image sent to a device that uses a specified output resolution should have a specific target ppi. The table shows some real-world examples of output resolution and workable ppi ranges. Formulas used for the calculations are shown in the Calculation Used column; square brackets in the calculations indicate the range of values used to determine the lowest and highest resolution acceptable in that medium.

Note that these resolutions are suggested and not absolute. Images will still print and display at other resolutions, but the results may not be predictable or efficient.

Table 1.1
Approximate Resolution for Various Media

MEDIUM	OUTPUT RESOLUTION	APPROXIMATE IMAGE FILE RESOLUTION	CALCULATION USED
Web	72–96 dpi (monitor)	72–96 ppi	ppi = dpi
Inkjet (stochastic)	720 dpi	180–234 ppi	[1 to 1.3] × (dpi / 4)
Inkjet 6 color (stochastic)	1440 dpi	240-312 ppi	[1 to 1.3] × (dpi / 6)
Halftone, low resolution	75–100 lpi	116–200 ppi	[1.55 to 2] × lpi
Halftone, normal resolution	133–150 lpi	233–350 ppi	[1.55 to 2] × lpi
Halftone, high resolution	175–200 lpi	271–400 ppi	[1.55 to 2] × lpi
Line art	600–3000 dpi	600–1342 ppi	$(dpi/600)^{1/2} \times 600$
Film recorder	4K (35mm)	2731 × 4096 pixels	Total pixels
Film recorder	8K (6 × 9cm)	5461 × 8192 pixels	Total pixels

Actual resolution needs may be somewhat flexible based on circumstances, such as paper and equipment used, original image quality, expected results, and so forth. Be sure to read manufacturer suggestions, and take most of the advice offered by services that offer printing—they should know how to get the best results from their equipment.

> For now, or until you are sure of what to do, set your camera to the highest resolution, and resize images without interpolation (leave the Resample Image box in the Image Size dialog unchecked) to 240 ppi for printing—at least until you read more and have reason to do otherwise. See Chapter 11, the Appendix, and suggestions for resizing in the next section.

Resizing Images

There are two methods of changing the size of an area that a group of pixels occupies: one causes you to resample an image, actually changing the image content (using Bicubic, Bilinear, or Nearest Neighbor interpolation), and the other changes the resolution to redistribute pixels over a smaller or larger area without actually changing the image content.

> **Redistributing** pixels does nothing to actually change the content (mathematics) of the image information that is stored; it just suggests that the content will be applied over a different area. It is a ppi adjustment.
>
> **Resampling**, on the other hand, actually changes the content of your images and changes it permanently.

The larger the amount of resampling (the greater the percentage increase or decrease), the more it affects the image content. The greater the redistribution, the more it affects image size and efficiency. One of these two things, redistributing or resampling, has to happen each time you either change the size of the whole image (using Image Size, not Canvas Size) or change the size of a selection by stretching or transforming.

When you resample image information (*upsample* or *downsample*), changing the actual count of pixels, Photoshop Elements has to interpret and redistribute tonal and color information, either creating (upsampling) or removing (downsampling) pixels. It does this through *interpolation* (adding image information) or *decimation* (removing image information), which are really fancy names for making an educated guess. Resampling an image to make it larger will never fill in information that is not already there, no matter what you do and which plug-in you use. What resampling will do is estimate and average differences between pixels to make a best guess. Details will tend to soften (upsample) or be lost (downsample). In neither case will it actually increase the captured detail in an image.

That trick you've seen on TV, where a pixilated image gets clearer and clearer as they zoom in, is reverse engineered. You can never enhance image detail that has not been captured. The only thing you can really do to reclaim image detail that you don't already have is reshoot a subject with higher resolution (e.g., using a longer lens, macro setting, or higher pixel dimension to capture more detail) or rescan (assuming that the detail is present in what you are scanning).

Photoshop Elements has three methods of interpolation (methods of figuring out how to insert new pixels or remove existing ones as you change the size of an image) and five interpolation options. Nearest Neighbor, Bilinear, and Bicubic are the methods. Bicubic Smoother and Bicubic Sharper are variations (various levels of sharpening) on the Bicubic method that were added in Elements 3.

Nearest Neighbor When you resize using Nearest Neighbor interpolation, Photoshop Elements adds or removes pixels based on pixel information and color that already exist in the image. Whether upsampling or downsampling, there is no averaging of color and/or tone to create new colors/tones. Nearest Neighbor is useful, for example, for controlled upsampling of screenshots without blurring (quadrupling the pixel count can yield an exact duplication of an image at four times the resolution). Multiply by squares (4, 9, 16, etc.) to achieve controlled upsampling.

Bilinear Bilinear interpolation behaves much like Bicubic and is supposed to be faster, but I've never clocked them. During the sampling, new tones and colors can be introduced between existing colors that are not in the original image. This can blur sampling of hard edges but can provide a smooth transition for tones (Nearest Neighbor might provide a blockier, stepped result). One thing about Bilinear upsampling is that it does simple, true averaging between neighboring tones and adds fewer new qualities to an image than any type of Bicubic resampling. At times these interpolation properties prove to be an advantage in retaining look and feel (when decreasing the image pixel count), and in others they may result in softening (when increasing the image pixel count). Use Bilinear when you want to downsample images.

Bicubic The Bicubic resampling process creates new image information by averaging, like Bilinear, but goes one step further to provide a bit of sharpening to the result. The intensity of the sharpening is stronger depending on the type of Bicubic interpolation selected. The sharpening is intended to counteract the blurring result of averaging when increasing the pixel count. Bicubic resampling changes a greater number of pixels with the same radius setting as Bilinear but may generally give a better visible result in most cases than Bilinear (when upsampling). This type of interpolation is the real workhorse for sizing images. Bicubic Sharper is like Bicubic but with enhanced sharpening; Bicubic Smoother is like Bicubic but with less sharpening. Bicubic resampling can be used for both upsampling and downsampling images.

While making up information and decimating it sound like bad things, each has its purpose. Usually you should avoid upsampling—especially if such options as rescanning or returning to an original camera image exist for gaining more detail. However, images can be upsampled with some success, depending on the desired quality—provided the change isn't huge. Upsampling 10 percent or even 20 percent may not be noticeable if the source image is sharp. Usually you will upsample only to make up small gaps between the resolution you have in an image and what you really need or to adjust borrowed image components (elements you are compositing from other images).

Downsampling, while certainly damaging and compromising to image content, should be less noticeable in your results if you use the right sampling methods. Image information indeed gets averaged or eliminated, but if downsampling is being done for the right reason, any details you lose in resampling would have been lost on output or display anyway. Detail loss is inherent in the process of downsampling, or outputting images at a smaller size. Even if equipment can reproduce detail at a smaller size, eventually details will pass the limit of the human eye's ability to discern them. In other words, at some point you lose the details anyway.

Find even more information on image resizing and interpolation in the Resolution section of the Appendix, under "Interpolation."

Multipurpose Images

Making images that you'll use for more than one purpose (for example, print and Web) can cause a little problem considering the resolution and resizing issues already discussed. Optimally, you'd like to work with images so that you target the result. Doing so ensures that you retain all of the actual image detail rather than relying on interpolation or decimation and your choice of sampling type to interpret detail. However, you can't work on an image at two resolutions or in more than one color mode at the same time. It is a simple fact that an image going to print on a high-resolution printer should have more information than one at the same size used on the Web. This is because of the difference in the way these media use image information. In fact, different printers and printer types will have different optimal utilization of image detail because of their mechanics. You will need to target image information to your output or you will not optimize detail.

You have only two solutions in working with dual-purpose images:

- Create more than one image, each with a specific purpose.
- Create one image and resize it.

Either of these choices poses a trade-off. In creating more than one image, you sacrifice valuable time in repeating processes for correction on different versions of an image. It is often self-defeating to work on two images to produce the same results (even using a detailed script) because the difference in size and volume of information in the image will produce different results with the same application of tools. In creating one image and resizing, you have to allow either interpolation of new image information or decimation, neither of which may be the optimal process. You can't work on small images and resize up because detail will not be present.

The best way to go about working with multipurpose images is usually to work with them at the highest resolution and then resize them smaller. Working at the higher of two or more resolutions retains the details for the higher-resolution presentations and decimates detail that will not be reproducible at lower resolutions. Softening or other ill effects from severe resizing can be countered somewhat by sharpening.

See "(Un)Sharpening and Boosting Contrast" in Chapter 5 for more information on sharpening.

A similar concept in retaining detail holds true when considering color depth. You will want to work in larger color spaces and at greater color depth to retain image detail and then reduce color detail and move to smaller color spaces after making corrections to reduce loss. You will most often use images from your digital camera at full resolution in RGB during corrections before reducing color and resolution for specific purposes.

Can You Have Too Much Resolution?

There are two answers to the question of whether or not you can have too much resolution: yes and no. The answer depends on whether there are other circumstances that make high-resolution images a waste because the information won't ever be used. There does come a point where the amount of image information is simply too much for the purposes of the image as it is being applied, or it gets so fine that more information doesn't really reveal more useful detail.

For the most part, you want all the detail that you can get in your source images from a digital camera. Consumer digital cameras are not so powerful that you will have enormous file sizes that are unwieldy—though you may need to consider alternatives for archiving images and image storage (on camera and off) to make the most of your equipment. Storing your images at high resolution will allow you to return to them for other purposes in the future.

Exceptions to the high-resolution rule happen only when high resolution is absolutely overkill for the purposes that you took the image in the first place. For example, if there are some items that you want to put on eBay that don't require a lot of magnification to see product quality or details, then you may need just enough resolution to show that the item is intact. Taking a full-resolution image may not be necessary, and large image downloads may annoy, rather than attract, potential buyers. If you are taking just a few images of this sort, it may be just as easy to take the high-resolution image and resize the image smaller later (especially if you find you have to look for the camera's manual to figure out how to change the resolution settings). Quality in this case is hardly the issue.

If you think you might ever use the shot for more than one purpose, grab all the resolution that you can with your digital camera.

Images that you obtain using a scanner or scanning service fall into a similar category. You will want to get as much detail as you can from a scan, but at some point you will be examining the grain of the film or paper in the print rather than extracting detail from the existing image. If you are scanning an image from a newspaper, you can use less resolution (samples) than if you scan from a negative, because the resolution of newspaper printing will be far lower (see Figure 1.2). The detail of your source can dictate the resolution as well as the application.

If you are a casual shooter and only send photos to relatives via e-mail, keep in mind that an image will display at about three times the size in a web browser (72 ppi) than it does in print (240 ppi). So even resolution that seems low can be more than you really need depending on what you do.

When going to print, too much image information can slow down processing and can be overkill. If you have a 5″ × 7″ image at 240 ppi, that will be enough to print at that size for many purposes. That is just 1200 × 1680 pixels, roughly what you get from a 2 megapixel camera. If you envision doing larger prints of the same quality, you will usually need more resolution (depending on the output devices). On the other hand, using that same image without resampling for a 2″ × 2.8″ image in a magazine is overboard: you'll have roughly twice the resolution you really need. That additional resolution taxes computer imaging resources and will not improve your result. Imaging equipment ends up just crunching the image information and decimating detail, likely using a simple averaging technique with no sharpening. What that means is you may even get a better result if you resize the image correctly to a lower resolution on your own and sharpen the result. See the chapters in Part Vand the Appendix for more information about printing and resolution.

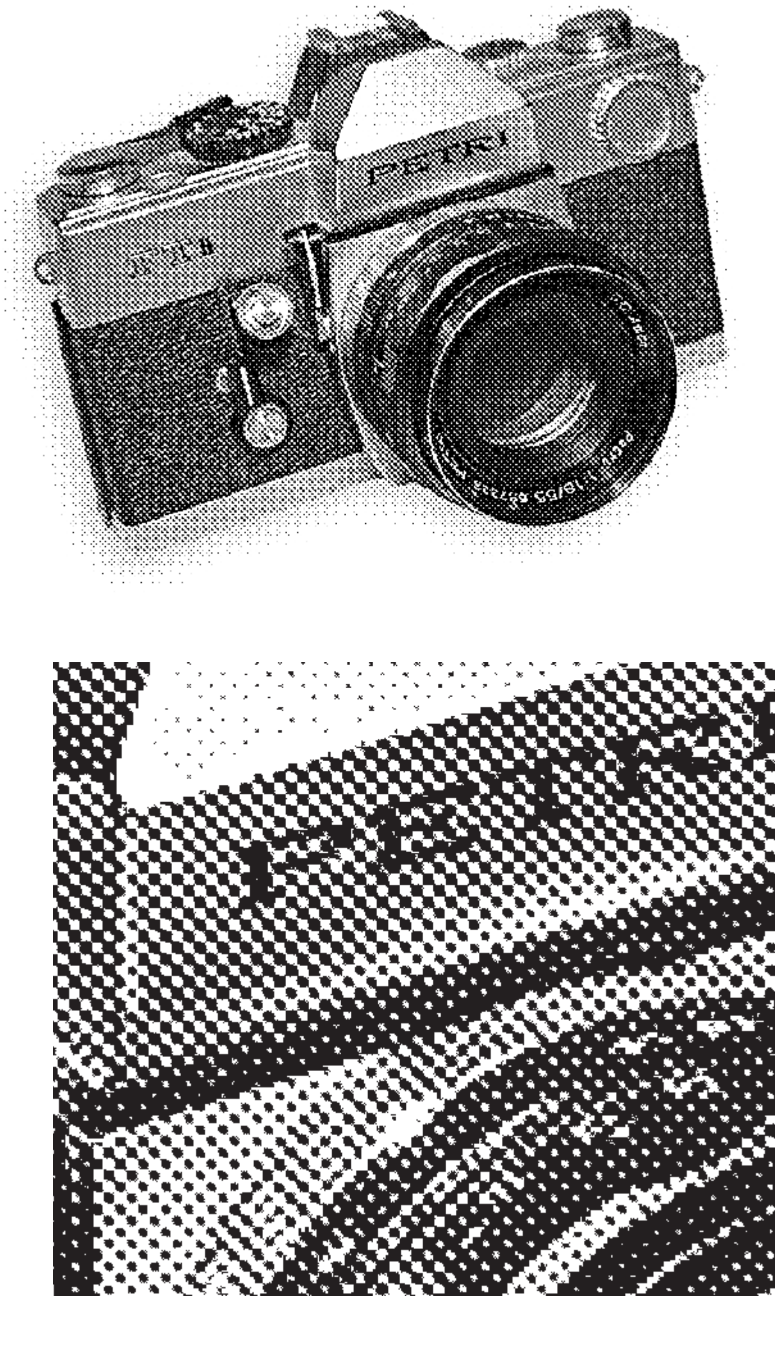

Figure 1.2

If you're scanning or photographing this image from newsprint, the detail soon runs out, making higher-resolution scans or photos superfluous.

You may want to plan a little for the future; not only may technology improve to demand more resolution from images to make the best prints, but your needs may change and additional resolution may leave you with a little leeway to take advantage of future changes. Use resolution to your advantage, rather than just assuming it is correct. So, can you have too much resolution? Yes, but usually only in applying the images (in print and on the Web). Traditionally, the push has been toward more resolution and greater definition. However, with newer consumer digital cameras capturing larger images, understanding how to apply and use image resolution becomes more important. Get what you can for archiving and storing images, but target that resolution to output and display sizes as needed.