M.E. Gruchalla, P.E.
Member: Editorial Review Board
4816 Palo Duro, NE
Albuquerque, New Mexico 87110

IMAGE SCANNING

How much resolution is enough?

Digital photography and scanning are among the fastest-growing photographic technologies, and many amateur radio operators like to share memories of DX trips and pictures of their stations via the Internet. The cost of equipment for digital image processing has dropped drastically in the past few years, as the resolution has increased.

What resolution do we need? It's said that there's no need to scan an image at much greater resolution than the output device is capable of printing. However, if our purpose is to digitally archive our images rather than simply scan and print them, we may wish to retain all the information in an image—independent of whether all that information can be reproduced with one of today's output technologies. In this paper, I examine the scanning and printing process primarily from an engineering rather than photographic perspective, to determine—at least to a reasonable estimate—the maximum usable resolution in scanning and printing of amateur negatives and prints to capture effectively all the information in the image.

This is not intended to be a rigorous analysis of the digital-imaging processes. The primary purpose is to provide insight into some of the basic elements of digital-photography from a technical perspective, and to explore the important digital and photographic parameters.

The photographic art

In the photographic art, the needed optical resolution is a function of a number of parameters, including such subjective parameters as viewing distance and the nature of the subject. For example, a much higher resolution may be required for the presentation of an industrial photograph of a refinery than would be required in a soft-toned portrait. Photography as a true art form relies on the subjective artistic ability of the photographer to control such photographic parameters as contrast, color, and focus to produce a pleasing result for the viewer.

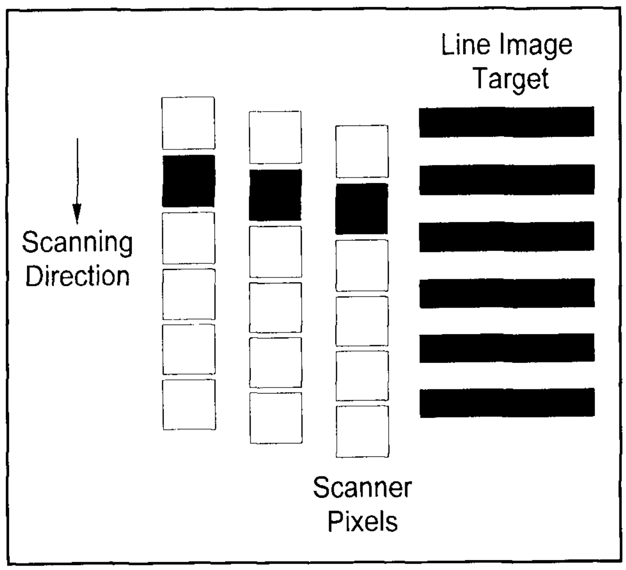

Figure 1. Scanner pixel size equal to target resolution (line pitch).

Photography as a science

Photography as a science is not subjective. Parameters of the various elements in photography as a science are specified analytically. Film resolution is precisely documented by such analytic tools as the Modulation Transfer Function. Film response to optical exposure is precisely provided by the sensitometric curve, also called the gamma curve, of the film.

In the science of photography, the intent is to eliminate virtually all subjective parameters by analytically measuring as precisely as one can as many of the photographic parameters as pos-

Page 9, Communications Quarterly, July 1999, published by American Radio Relay League, Inc.

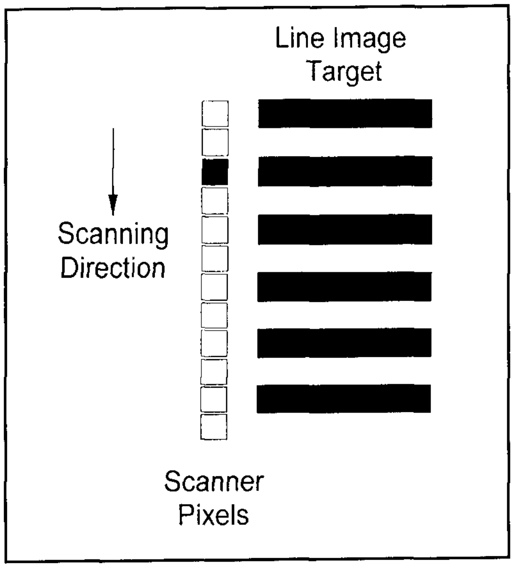

Figure 2. Scanner pixel size equal to target line width (one-half target resolution) pixel centered on black lines.

sible. The objective of this paper is to examine two fundamental analog materials of the photographic art—photographic film and photographic paper—and to determine from an objective technical perspective the equivalent limiting spatial digital resolution and optical digital resolution.

In general, the magnitudes of the values of these parameters from the technical perspective are expected to be much larger than those typically used in photography as an art. Objectively, the resolution is a measure of the limit of information storage of the analog medium. Theoretically, if an image were digitally recorded at spatial and optical resolutions equal to the limiting capability of film, the digital image would be virtually indistinguishable from an analog film recording (the statistics of the noise would be different, but the information content would be similar), regardless of such subjective parameters as subject content and viewing distance.

Digital photography and image scanning are precisely the same signal-processing process as digitizing an analog electrical signal. In processing typical electrical signals, care must be taken to assure that the sampling frequency is adequately above the highest frequencies in the signal to be digitized. If the sampling frequency is too low for the frequencies in the data signal, peculiar aliasing artifacts will result. Often the signal is prefiltered to assure that any signal content above some maximum frequency is sufficiently low to avoid unwanted aliasing. These same concerns apply in digital imaging.

We typically digitize an electrical signal in two ways, or two degrees of freedom or dimensions: We digitize the signal amplitude as a function of time, and we digitize time by digitizing the signal only at discrete time intervals. The digitizing of an image, whether the image of a scene being photographed with a digital camera or a photograph scanned with a scanner, is precisely the same signal-processing process as digitizing a temporal signal. We are simply digitizing in the spatial rather than the time domain. Also, in the spatial domain, we are digitizing in five dimensions rather than the typical two for electrical signals. As with an electrical signal, the actual signal parameter of interest is the magnitude of the signal—the intensity of the image at various points. In order to identify where each intensity measurement is taken, both the X an Y coordinate positions must be recorded with each sample. It is convenient that both the X and Y position of each sample be digitized.

Generally, an electrical signal has only a single feature of interest: magnitude (there may also be interest in phase, but this is related to the timing in the sampling system with respect to some timing reference). When we digitize the "magnitude" of an element of an image, we don't want just a single magnitude, we want the magnitude of each individual color in the image element. Fortunately, we need only record the magnitude of three primary colors, red/green/blue or cyan/magenta/yellow, to record the full color of the image element. Therefore, in addition to digitizing the X and Y coordinate positions of an image element, we also digitize the magnitude of each of the primary colors in the image element. Therefore, we need five digital parameters to completely describe each individual element of an image.

Typically the magnitudes of these parameters are stored in a particular sequence, so the actual X and Y coordinate values associated with each element need not be recorded. For example, when an electrical signal is digitized, we do not actually record the time information. We simply record each digitized sample in order. Because we know the sampling frequency, we know the time spacing between samples and, in turn, the actual relative time that each sample was taken without actually having to record each time value. The same is true in spatial sampling. Because we know the sampling spacing in the X and Y directions, we need only record the magnitude information for each image element in some known sequence in order to keep track of where each piece of digitized magnitude information belongs in the image.

The position sampling or "spatial sampling" of an image is exactly the same as "temporal sampling" of an electrical signal. However, unlike electrical signals, we can't typically "filter" the incoming image signal. After all, how does one filter the "bandwidth" of the image of a scene that is "input" to a digital-camera lens? If we aren't careful in selecting the "spatial sampling frequency" at which we sample our image, we could end up with the same aliasing artifacts in our digital image that can result in digitizing an electrical signal if the sampling frequency is too low for the frequency content of the signal.

To eliminate aliasing, we must assure that the sampling frequency of our digital camera or scanner is sufficiently higher than the maximum spatial frequency content of our image. The frequency of an electrical signal is specified as the number of magnitude cycles per unit time; i.e., cycles per second or Hertz. In the spatial domain, the spatial frequency is specified as cycles of intensity per unit length; i.e., cycles per millimeter, cycles per inch, etc. The image frequency in the X direction will generally be quite different from that in the Y direction. Also, the sampling frequency in the X direction may be different from the sampling frequency in the Y direction; consider a 300-dpi x 600-dpi scanner, for example. But, sampling in the spatial domain and the electrical domain are precisely the same, and all the same mathematics apply. Just the number of dimensions and the units are different.

Resolutions

There are two different resolutions that must be considered in digital imaging: spatial and intensity. Spatial resolution is a measure of how many discrete spatial positions the image is divided into per inch, millimeter, etc. in the horizontal and vertical dimensions. The spatial resolution is typically specified in dots per inch, or dpi. The higher the spatial resolution, the higher the number of positions that can be resolved.

The other type of resolution is optical-intensity. Because the intensity of the digital image is digitized, the original continuous intensity of the image is "quantized" into some number of discrete levels as in any digitizing process. The intensity resolution is typically specified in bits, referring to the precision of the digitizer used in the digitizing process. So, if the intensity of a photographic element is digitized to eight bits, a total of 256 intensity levels (255 levels plus zero level) are resolved. Unfortunately, as noted below, both the spatial resolution and the intensity resolution are confused a bit by the way manufacturers specify their equipment.

Scanning and printing

In general, the scanning and printing processes are considerably different due to the basic

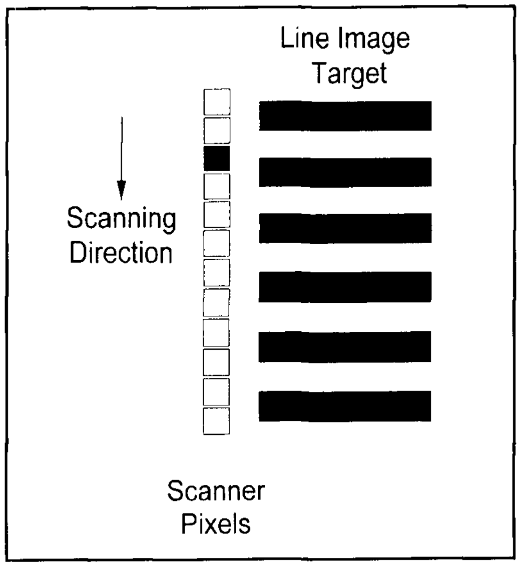

Figure 3. Scanner pixel size equal to target line width (one-half target resolution) pixel centered between black and white lines.

operation of the typical scanner and the typical "photo-quality" printer. The resolution of both scanners and printers are usually given in terms of dots per inch in horizontal and vertical directions, but this specification has a slightly different meaning for a scanner compared to a printer. Each dot a scanner digitizes is essentially a "scanner pixel." The pixel is the smallest picture element that can be resolved.

When the scanner digitizes a pixel, it digitizes the average intensity of each of the primary colors in that pixel. If the scanner is a "24-bit" type, it digitizes each color to eight bits of intensity resolution. The three eight-bit digitizing resolutions corresponding to the three colors are added to obtain the 24-bit specification. Specifying the unit as an "eight-bit" color scanner, where it is understood that each color level is digitized to eight bits, would provide the same information, but the 24-bit number is a much more impressive marketing figure. Similarly, a 30-bit unit digitizes to 10 bits and a 36-bit unit to 12 bits. But anything over 8 bits is not quite what it seems—more about that later.

So, the scanner "sees" a pixel as a dot, more or less, and digitizes its color intensities. A 24-bit scanner can see 256 different levels of each color intensity in a single pixel "dot." However, a printer like an ink-jet type can only print or not print a dot. It can't print part of a dot, or dots of different optical densities. Some printers do provide "image enhancements," and others such as dye sublimation printers can provide better continuous tone in each dot. But, because the ink-jet type printers are the most common, these are the only printers considered here.

To print an image that has a density somewhere between white (assuming a white paper

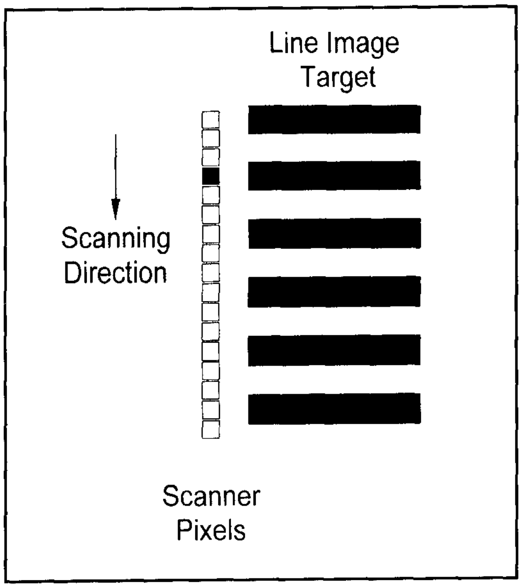

Figure 4. Scanner pixel size equal to one-third target resolution pixel centered on black lines.

printing medium) and maximum density, the printer must print some dots and leave out some dots in a printed area to achieve a visual effect of a midrange intensity density. This is the basic principle used in half-tone printing. The result is that a "printer pixel" must have many dots in it to allow the printer to print pixel densities from total white to maximum density. I'll go into more printer detail later.

The Modulation Transfer Function

One of the critical characteristics of an optical device or material is its ability to resolve the

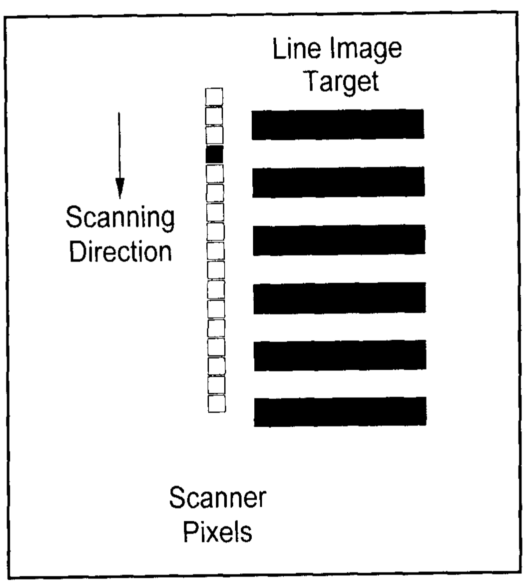

Figure 5. Scanner pixel size equal to one-third target resolution pixel centered on white lines.

fine detail of a subject. As the detail of a subject becomes more fine, optical systems lose the ability to resolve the contrast of the subject. The Modulation Transfer Function, or MTF, is a parameter used to specify the spatial resolving ability of an optical element as a function of frequency. The MTF is a type of "photographic frequency-response" parameter.

Photographic resolution is typically specified in line pairs per millimeter. A line pair consists of a black line adjacent to a white line, where the black and white lines are the same width. When a number of such line pairs are placed together side by side, we see a series of parallel black lines separated by white spaces. The white spaces are actually the white lines of the line pairs. Because a black line plus a white line constitutes a line pair, the spacing or period or pitch of the line pairs is measured from one line of a line pair to the same line in the next pair. For example, the distance from the center of the black line in one pair to the center of the black line in an adjacent pair is the line-pair period. The unit of the line-pair period is millimeters per line pair (mm/lp). The reciprocal of this period is the resolution or frequency in lp/mm.

A test target is typically used to determine the photographic (as opposed to digital) resolution of photographic equipment and materials. Such test targets include a number of different standard patterns, including a number of line pairs of different periods. The test target is photographed and the processed image is examined to determine the line pairs with the smallest period that can be seen distinctly. This may be applied to both negatives and prints, and even to lenses where an aerial image is examined with a microscope.

"What can be seen distinctly" is a somewhat subjective measurement. What one person may be able to see, another may not. Or, if given an entire image, one may be able to "just see" more detail than if only given a very small part of the image containing fine detail. I will use a type of "half-power point" as my measure of the resolving ability of photographic materials.

The density of an image element on film is proportional to the intensity of the image and the time of the exposure. Because the optical intensity is power per unit area and the image element is an area, the image density is proportional to power per unit area times the image-element area times the exposure time. The product is energy in watt-seconds. When we photograph a subject, the exposure time of the entire frame is the same for all points in the frame. If the exposure time is constant for all points on the film, and only the image intensity at each point varies, the image density at each point on the film is proportional to the "input" optical power at each point.

I'll use the half-power point of the developed image as the limit of the useful resolution of photographic film and paper. The MTF is a type of measure of the power bandwidth of photographic materials, and my half-power point is simply the "optical frequency" at which the MTF drops to a value of 0.5. In an electrical circuit, the half-power point is the frequency at which the response drops 3 dB from the mid-band value. However, the response is not zero beyond the half-power point; it simply drops with increasing frequency. So, there is still useful information beyond the half-power point. The half-power point or its equivalent is simply a convenient reference point for many physical phenomena. My choice of an MTF value of 0.5 as the limiting value is also simply a convenient reference point and one consistent with the electronic art. There's certainly resolution information at spatial frequencies above this point, but the density ratio becomes frequency dependent above that point—just as gain in an electrical system becomes frequency dependent above the upper –3 dB frequency.

The MTF is normalized to the contrast of the actual subject being photographed. Consider a line-pair test target containing line pairs of different frequencies used as the subject of a test phonograph. First, the contrast ratio of the subject is determined; for example, a particular line-pair target to be examined. This is called the Modulation M of the subject. Optical modulation is defined as the ratio of the difference of the maximum intensity minus the minimum intensity divided by the sum of the maximum and minimum intensities. The modulation of the subject pattern, the object of the photograph, is called M_O. The test object is photographed and the film processed. The image of the object is examined and the modulation of the image, M_I, is computed in the same manner.

$$\text{Object Modulation} = M_O = \frac{I_{O,MAX} - I_{O,MIN}}{I_{O,MAX} + I_{O,MIN}} \quad (1)$$

$$\text{Image Modulation} = M_I = \frac{I_{I,MAX} - I_{I,MIN}}{I_{I,MAX} + I_{I,MIN}} \quad (2)$$

The MTF is the ratio of M_I to M_O. Typically the object for MTF measurements is a test target designed with a constant modulation, perhaps 35 percent, and with a varying frequency. In other words, the target contains numerous line pairs of different periods all with the same contrast ratio.

$$MTF = {}^{M_I}\!/_{M_O} \quad (3)$$

One particularly useful test target is a sequence of line pairs where the period of the lines decreases across the target. In other words, the frequency increases across the target. By using such a target, the MTF as a function of image frequency may be very easily measured. This then provides the frequency response of the photographic element being tested.

A very good professional lens may have an aerial resolution of 300 lp/mm. A fine-grained film optimally processed may provide as high as 100 to 200 lp/mm resolution.[1] These are pretty much the upper end of what is available with standard techniques.

On a more practical level, a color negative produced with a camera of modest quality and competently processed by a bulk processor will provide a resolution of about 50 lp/mm or perhaps slightly higher.[2,3] In a color print, about 10 lp/mm or a bit higher is typical. Here, I consider only color photographs, as color is far more popular than black and white. Therefore, looking at the camera, lens, and film as a single system, I'll use a resolution on the order of 50 lp/mm in a color negative or slide as about the upper limit of what's achievable in the consumer market. This corresponds to 1,270 line pairs per inch (lp/in). For prints, I'll use 10 lp/mm, or about 250 lp/in, as the maximum typical resolution.

Scanning

First we'll look at scanning in a little more detail. What resolution is needed in a scanner? That depends on the object to be scanned. A photographic negative will require much higher resolution than a photographic print, and a magazine picture much less than a real photograph. Also, it is very likely that we may wish to "enlarge" a scanned image of a negative just as we would enlarge the image in a negative for a photographic print, so we'll need a higher resolution in the scanned negative image than in the final printed image. The best we can do is to capture all of the information available in whatever we are scanning. Let's look at the negative first. In general, this applies to slides as well.

Suppose we photograph a line-pair test target such that the image of the target on our negative is 1270 lp/in as noted above. Now, suppose that this target image is just visually resolvable in our processed negative; i.e., it's at the half-power point of the MTF of the film. This photographed image will now be used as our scanner test target. For convenience, I'll call the maximum-density lines of the target black lines "black" and the minimum-density lines of the target white lines "white." If we

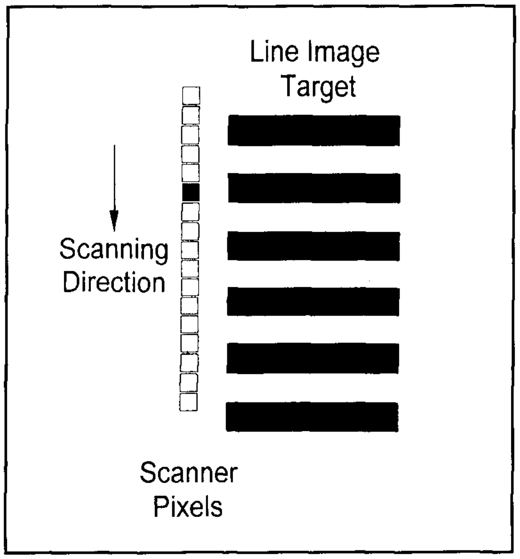

Figure 6. Scanner pixel size equal to one-third target resolution pixel positioned at the edge of black lines.

were to scan this target, what spatial scanning resolution would be required to just resolve this target image?

Test target 1

It might seem at first that 1270 dpi would be adequate; but, suppose we have a scanner with a spatial resolution of 1270 dpi. The pixel size is then 1/1270 pixels per inch or about 0.0008 inch. Then suppose we place the target in the scanner with the test-pattern lines running exactly at right angles to the mechanical scanning direction, so with each step of the stepper motor the scanner advances down the line pattern from one line to the next. We could orient the pattern at right angles to the detector array as well, but it's easier to visualize the stepper stepping along the line pattern. Finally, we carefully adjust the target so the center of each scanning pixel of the 1270-dpi scanner coincides exactly with the center of each black line in the 1270-lp/in test target of the negative.

Now, recall that a line pair consists of a black line and a white line side by side. For the 1270-lp/in target, the spacing between black lines is 0.0008 inch, so the actual black and white lines are each 0.0004 inch wide. However, the scanner pixel is 0.0008 inch; therefore the pixel spans precisely two entire lines, one black and one white. No matter how we adjust the target with respect to the scanner, as long as we do not rotate it, a pixel will always span exactly one black line and one white line (see **Figure 1**). The active pixel in each of the figures is shown filled black. The three columns of pixels in **Figure 1** represent three positions of the scanner with respect to the line target. Notice that no matter how we slide the pixels with respect to the target lines, each pixel always spans exactly a black and a white line.

For example, for the first positioning (column of pixels on the left in **Figure 1**) the active pixel spans a white line, a black line, and another white line, so this pixel sees one whole black line and one whole white line. For the second positioning, the active pixel spans one whole black line and one whole white line. Finally, at the third position, the active pixel spans a black line, a white line, and a black line, so again the pixel sees a complete black line and a complete white line.

By pixel definition, the scanner cannot resolve any detail within the pixel. It can only resolve the average intensity of the entire pixel. So, if the scanner pixel spans both a black line, which is a minimum-intensity line, and a white line, which is a maximum-intensity line, the average intensity of the pixel is half maximum intensity, or 50 percent gray. And, if no matter how we adjust the target along the scanning direction a pixel always spans both a black and a white line, the entire image will always be scanned as 50 percent gray. We won't be able to see the individual lines of the target in the scanned image. Consequently, a 1270-dpi scanner resolution can't resolve a 1270-lp/in target. Instead of just resolvable black and white lines, a uniform gray image will be the result of the scan.

Test target 2

Now suppose a scanning resolution of 2540 dpi is used. In this case the pixel size is exactly the same as the width of the lines, or about 0.0004 inch. In general sampling theory, a periodic signal must be sampled more than twice per signal cycle to just be able to resolve the signal. If the scanner pixel size is exactly the width of a line, then exactly two samples per cycle of the line-pair target will be provided. Again, we position the target so a scanner pixel coincides exactly with the lines of the target. This is shown in **Figure 2**. Now, one pixel falls exactly on a black line and the adjacent pixels on either side fall on white lines. Therefore, as the scanner steps down the target, it will scan alternate black and white lines. It appears that a 2540-dpi scanning resolution is adequate to resolve a 1270-lp/in image, but this is not quite true.

Suppose we now move the target down along the scanning direction a distance of exactly 1/2 pixel, as shown in **Figure 3**. Now each scanner pixel will always span one half a black line and one half a white line. In **Figure 3**, the active pixel spans one-half a black line and one-half a white line. On the next step, the scanner will move one entire pixel, so the pixel will now span one half a white line and one half a black line. In each case, the average pixel intensity is gray. Here, with the target adjusted so the scanner pixel falls halfway between black and white lines, the entire image will again be scanned as continuous gray. So, a scanning resolution of exactly twice the image resolution isn't quite good enough. This example (though perhaps somewhat simplistic) demonstrates why sampling a periodic signal more than twice per signal cycle is required to just resolve the signal.

Test target 3

Finally, suppose we have a scanning resolution that's a factor of three finer than the image resolution. This 3X scanning-resolution factor provides three pixels in the period of each line pair of the test target. For our 1270-lp/in test target, this is a 3810-dpi scanning resolution—a pixel size of about 0.00026 inch. Theoretically, a scanning resolution somewhat more than twice the image resolution could be used, but a factor of three is very easy to demonstrate here.

Again, we adjust the target so the scanner pixel falls exactly on the center of a black line as shown in **Figure 4**. In this case, the pixel is smaller than the target line so we could actually miss it a little and still be totally on the line. At this point, the pixel is digitized as totally black. When the scanner steps one step, the pixel will fall between the black and white lines with one third of the pixel on the black line and two thirds on the adjacent white line. The average pixel intensity will be digitized as about 33 percent gray.

When the scanner steps a second step, the pixel will fall two thirds on the white line and one third on the next black line, so again the average pixel intensity will be 33 percent gray. And, with a third step, the scanner pixel again falls totally on the next black line. So, again, we can resolve the individual black and white lines of the test target if we carefully adjust the position of the target with respect to the scanner. Of course, we should have expected this based on the previous example.

The scanned representation of our image is a series of black and gray lines. At the starting position the pixel is totally black, at the first step it is about 33 percent gray, at the second step it is also 33 percent gray, and at the third step is back to black again. Our scanned image appears as a series of black/light gray/black lines. The scanned image isn't quite the same as the original object, but what we are looking for is the ability to "just resolve" the individual black and white lines of the line pairs.

Now move the test target down one pixel width as before, or about 0.00013 inch, as

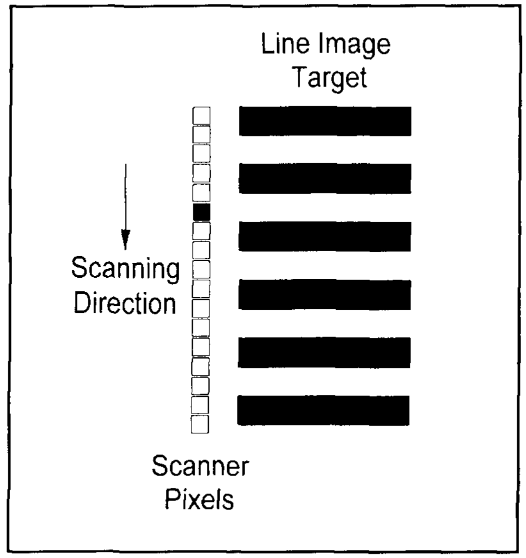

Figure 7. Scanner pixel size equal to one-third target resolution pixel positioned at the edge of white lines.

shown in **Figure 5**. Now the scanner pixel will alternately fall exactly on a white line and then part way between lines. Again scan three steps, but start on a white line. At the starting position the pixel is white, at the first step it's 66 percent gray, at the second it is also 66 percent gray, and the third step is white again. Now our scanned image appears as a series of white/dark gray/white lines. We can still resolve the image, but this image is different from the one obtained with the target adjusted for the scanner pixel initially falling exactly on a black line. Is this what we expected? Is it good enough?

This phenomena is an optical contrast reversal of the image caused by a type of aliasing in the signal processing of the digitized image (the image is discrete black and white lines) and the digital positioning of the scanner (the scanner can only position in discrete steps). This contrast reversal in the image of an object, such as a line-pair target, is an objective point of reference. In other words, it is an optical phenomenon that anyone can see with proper training, and there's little argument as to whether the contrast is reversed or not. Therefore, this point of contrast reversal can be used as a well-defined reference limit of image resolution expressed in lp/mm, or lp/in if you prefer.

Now, suppose we take the last example above with a scanning resolution of 3810 dpi, and position the target so the scanner pixel is totally on a black line but right at the edge of the line in the scanning direction, as shown in **Figure 6**. The pixel is black at this initial position. At the first step, the scanner pixel is then totally on the adjacent white line, so the pixel is totally white. With a second step, the pixel is exactly between lines so it's one-half black and one-half white, and the average pixel density is 50 percent gray. With a third step, the pixel lies totally on the next black line, so the pixel is black. The scanned representation is a series of lines with a black/white/50-percent gray/black sequence. With this positioning of the target, we can still just resolve the target image.

If we move the target so the initial pixel is just on the edge of a white line, as shown in **Figure 7**, the pixel at this initial position will be white. At the first step it will be black, at the second step it will be 50 percent gray, and finally at a third step it will be totally white again. The scanned image will be a sequence of lines white/black/50 percent gray/white. As a result, there is still a contrast reversal because, in the first case, the sequence is black/white/gray/black and in the second it is white/black/gray/white. We could also view the first as gray/black/white/gray and the second as gray/white/black/gray. Such a contrast reversal is also called "pseudo resolution," as it produces apparent additional lines between the real lines.[4,5] In this example, we can easily see that the pseudo resolution has provided a third virtual line in the scanned image. There's no gray line in the actual target between any of the black and white lines.

Test results

In the three test cases examined here, it appears that a 3:1 scanning-to-image resolution is adequate to provide a scanned image that can just resolve the source image, but will provide a contrast reversal. At this scanning resolution, the ability to just resolve the image is not dependent on the positioning of the target in the scanner. Of course, the scanned image is different for different positions, but we can always just resolve the target image. In general sampling theory, the Nyquist Frequency is a factor of two above the frequency of the signal being sampled, whether it is an electrical signal or a photographic line-pair test target, and the sampling frequency must be greater than the Nyquist Frequency to preserve the signal information and to prevent aliasing.[6]

Perhaps a scanning resolution a factor of three greater than our negative resolution is what we need at a minimum. But, even at that resolution our line target is not really faithfully reproduced because, although we can just visually resolve the target on the film image, we will see a contrast reversal and pseudo resolution as a function of how we position the target in the scanner. This may result in Moire patterns in images that contain very fine detail consisting of parallel lines. Also, I have used the half-power point of the MTF as a point of reference. There is certainly additional image information beyond the half-power point. A scanning resolution that's nominally a factor of three finer than the maximum image spatial frequency will only provide about one bit of information about the image at the maximum image frequency—whether anything is there or not.

Generally, a photographic resolution of about a factor of five higher than the maximum spatial frequencies found in the subject to be photographed is considered the typical resolution necessary to capture virtually all the detail of that subject. This is also pretty much true for scanned images. It's equivalent to providing an electrical bandpass in a system that's substantially higher than some electrical signal frequency of interest in order to prevent the bandpass of the system from altering the signal. But, as I'll show later, the higher the scanning resolution, the greater the memory we'll need to store the image—the memory required goes up by the square of the resolution. The storage requirement increases by a factor of four if we simply double the scanning resolution. There is a serious compromise between how precisely we digitize an image and where we store it.

Examined on a microscopic scale, a photograph is composed of a sort of dot structure called "grain." Grain refers roughly to the individual silver-halide crystals that form the photographic image (this is not quite accurate for color materials, but the effect is the same). The grain size and spacing is a function of film speed as well as other parameters, such as storage and processing of the film. A medium-speed film may have a grain size and spacing on the order of a micron or so, or about 500 to 1000 grain elements per mm, about 25,000 grain elements per inch.[7] So, it's virtually impossible to truly capture with a scanner the actual granular information of the film. But such fine detail isn't actually necessary because the film itself is not capable of providing spatial image resolution equivalent to the grain spacing. As noted above, the MTF of a typical color film limits the spatial resolution to about 50 lp/mm, or about 1270 lp/in.

Fortunately, nature is kind to us in typical photographs. Rarely do we wish to photograph resolution charts, such as line charts or pie charts, to proudly display in our home gallery. I focus in this paper primarily on how much scanning resolution is needed to faithfully archive the image information of a typical amateur photograph or negative. In general, these photographs will be of people, pets, flowers, landscapes, and other generally random images—images that tend to have random intensity characteristics as opposed to very structured line characteristics, such as a line-pair test target.

These amateur photographs aren't necessarily of lower quality than professional photographs. They simply won't typically have the same subject content. Neither will these images typically have the equivalent of line-pair resolution test targets. For these typical types of images, a scanning resolution of about a factor of two greater than the image resolution is usually adequate. This isn't necessarily true of "professional" photographs where industrial products are often the subject; for example, a microphotograph of the chip of a Pentium II processor which will have very highly structured image characteristics very much like a resolution test target.

To reliably capture the full image information in a professional industrial photograph in a scanned image, a higher resolution of about a factor of five higher than the resolution of the photographic medium is required. The equipment involved will be much more costly than that required to capture the "useful" information in a typical amateur photograph. Of course, if you feel your photographs are of professional quality and of the types of subjects that demand the highest-possible scanning resolution, you may want to invest in professional-quality equipment, along with some means to store and print the resulting images. I am limiting my analyses to typical amateur photographs and negatives of random-density images.

Two film scanners available on the consumer home market are the Hewlett Packard PhotoSmart scanner and the Nikon CoolScan scanner. In negative/transparency mode, the scanning resolution of the HP unit is 2400 dpi x 2400 dpi and that of the Nikon unit is 2700 dpi x 2700 dpi.[8,9] This is pretty close to the estimated 2540-dpi resolution computed above and is required as a minimum to just resolve a typical amateur negative. This is a good compromise based on the available components, such as CCD imaging devices, positioning technology and the final digital-file storage requirements, and price.

Basic requirements

Recapping so far, a typical negative can provide a resolution of about 50 lp/mm or about 1270 lp/in. A scanning resolution about a factor of three greater than the film resolution, or about 3810 lp/in, is needed to reach the point of contrast reversal in a 1270-lp/in test target. And, if the full detail of the image is to be preserved, a scanning resolution of about a factor of five higher than the film resolution is needed, or about 6350 lp/in. But, because typical subjects of amateur photographs tend not to have highly structured detail expected in professional industrial photographs, a scanning resolution only about a factor of two higher than the film resolution is usually adequate to capture virtually all the image information. The 2400 dpi provided by scanners like the Hewlett Packard PhotoSmart scanner is just at this minimum requirement. Scanners like the Nikon CoolScan provide 2700 dpi, which is just a bit above the minimum requirements. We'll look at this in a little more detail below in the review of storage requirements.

If a resolution of about 2500 dpi to 2700 dpi as provided in film scanners is just adequate to scan a photographic negative if all of the spatial detail is to be captured, what resolution is needed to scan a print? That parameter can be determined in the same manner as for the negative.

The resolution limit for a reasonable-quality, bulk processed print is on the order of 10 lp/mm, or about 250 lp/in. Using the same reasoning as above, we must scan such a print at a resolution of at least 500 dpi (2X) to capture all the useful spatial data in the print image—750 dpi (3X) if we are to reach the contrast-reversal resolution, and perhaps 1250 dpi (5X) to retain virtually all spatial information. A typical magazine photo is printed at about 133 dpi, so a scanning resolution slightly greater than about 266 dpi (or about 300 dpi) is needed to capture a printed magazine image.

Low-end scanners provide an "optical" scanning resolution (as opposed to interpolated) of approximately 300 dpi x 300 dpi. This is quite adequate to capture all of the spatial information in a typical printed image in a magazine or newspaper, but it's not quite adequate for true photographic prints. This does not mean that these 300-dpi scanners are useless or in any way substandard. We must decide what we need, or perhaps more appropriately what we want.

For the most part, magazine pictures look pretty good. If we are scanning prints for articles, letters, or other general interest applications, 300 dpi is quite adequate. But if we are scanning photographic prints to make electronic archives of them for digital preservation, or if we are scanning photographs that will be cropped and enlarged, we probably want to capture as faithfully as practical all the information available in the original print. To do so, we'll need at least a 500-dpi scanner, preferably 750 dpi. And, if we want to preserve essentially all the spatial information, we may need 1250 dpi. The medium-range scanners are 300 dpi x 600 dpi units, which is a reasonable compromise but still a little lacking in one dimension. The 600 x 600 dpi units are getting into the medium to high-end market, but the prices are coming down all the time. There are several 1200 dpi x 1200 dpi units available for about $300 to $500 for the basic scanner.

The unequal scanning-resolution specification for some scanners, such as 600 dpi x 1200 dpi, may seem a bit peculiar, but it is a result of how the scanners operate. In one dimension, typically across the narrower scanning dimension, the resolution is determined by the optical sensor, generally a CCD line array. This is simply a number of photodetectors side by side forming a single line of optical detectors.

The spatial resolution at the detector is determined by the spacing between the individual optical elements. The total width of the scanning area is imaged onto this line array. The scanning resolution is determined by the total number of detectors in the array divided by the width of the scanning area. For example, if there are 2500 elements in the detector array and the scanning width is 8.5 inches, the scanning resolution is approximately 300 dpi. An array slightly larger than 5000 elements would be required for 600 dpi. Because the number of elements in the optical-sensor array of the scanner directly relates to the cost of the optical array, the higher the resolution in the direction of the optical array, the higher the price of the scanner.

To scan the length of the image we move the scanning carriage along its length with a standard stepper motor assembly similar to that used in printers. The length of each step the carriage is moved is a function of the stepper-motor resolution (steps per revolution) and the effective gearing ratio used to couple the motor to the carriage assembly. This motion could also be provided by a linear motor, but it doesn't appear that this technology is as yet used in scanners.

Theoretically, the steps could be made as fine as desired, provided the stability of the mechanical system were adequate. For example, mechanical stepping is also used in computer disk drives to step between tracks (cylinders). Densities of almost 15,000 tracks per inch are used in modern hard drives—the Maxtor Diamond Max Plus Family of drives incorporates track densities as high as 14,522 tracks per inch for example.[10] So, generally, positioning technology is adequate to provide 1200 steps per inch in a scanner.

Initially, the scanner mechanical scanning resolution, that is the number of steps per inch, was made equal to the resolution determined by the optical detector array; a 300 dpi x 300 dpi scanner, for example. Then, for some reason, like marketing or improvements in the mechanical positioning technology of the scanning mechanism, twice the resolution in the mechanical scanning dimension became popular, such as the typical 300 dpi x 600 dpi units common now. This could have been a "no-cost" improvement provided by improved stepper-motor technology, tighter process controls, or other cost-insensitive manufacturing or materials changes. If the doubling of resolution in the mechanical scanning direction could be provided at little increase in cost, it could be offered as an "improvement feature" at little additional price to make the product more attractive over competitive products.

If the pixel size in both the X and Y scanning directions are the same, however—that is the photodetector pixel is square—it's not clear that scanning in one direction at one half the pixel size is particularly useful. It's a type of "interpolation," that creates a synthetic pixel between each real pixel. This provides a type of spatial filtering of the scanned image. It is not clear that having twice as much resolution in one dimension provides much improvement in image quality. It's likely as much a marketing tool as a usable feature. And not all manufacturers are jumping to provide this feature; for instance, the Hewlett Packard 4C scanner is simply a 600 dpi x 600 dpi unit.

Interpolation

In most scanner specifications, we'll see both an "optical" resolution and an interpolated or "maximum" resolution. Sometimes the optical resolution is omitted because it's the smaller (but most important) number. The optical resolution is the actual physical resolution the scanner is capable of providing. The interpolated resolution is a figure that results from the scanning system (software and/or hardware) "creating" additional data points between the actual scanned real data points. This is not real data. It is simply data computed by the scanning system based on a mathematical algorithm.

For example, consider a scanned image that results in alternating black and white pixels and in turn alternating black and white lines. If the scanning system is allowed to "interpolate," it computes several additional pixels in between the true pixels. In other words, it makes the true scanned pixels smaller and adds a few more in between that aren't true pixels. With these added pixels, a smoother transition in density is provided, whether it was in the original image or not. Consider that the system provides a 4X interpolation in both directions, say 300 dpi optical to 1200 dpi interpolated. For each real pixel, it will compute three new ones along both scanning axes so there are 16 times as many pixels in the interpolated image—the original pixel plus the three computed ones in each direction.

Say that a simple linear interpolation between adjacent real pixels is used. For a black and white line target starting at a real black pixel, the next pixel computed will be about 66 percent gray; the next, also computed, about 50 percent gray; the next, computed as well, about 33 percent gray; and finally the fourth pixel will be a real white pixel. So instead of a series of black and white lines corresponding to the original image scanned, the interpolated result is a series of more lines with a type of gray-scale transition between each real black line and each real white line. Clearly this does not add any information to the original image. It just makes the scanned image "look better," or not.

Basically what interpolation does is reduce the apparent pixel size, and this reduces the "pixelization" of the scanned image. For example, if we scan an image at 300 dpi and enlarge it enough, we will begin to see individual pixels. If the scanner interpolated by a factor of four, the pixel size would be one fourth as large; but the added pixels are not true image information. These pixels are simply estimates computed by a set of mathematical rules in the scanner system (software and hardware). As noted above, the grain size of typical film is much finer than the resolution that the film is actually capable of providing. So, at a microscopic level, because there are many more grain elements in the film per inch than resolvable lines, the grain itself provides a type of interpolation in the actual film.

Interpolation apparently is most useful where a small portion of an image is scanned and then enlarged. Here the finer pixelization provided by the interpolation reduces the jaggedness of the image, but the image will still be fuzzy because no actual image information is added by interpolation. It is not clear that interpolation is a useful feature in all applications. But we will get it "for free" whether we want it or not, although we do not have to use it. You will have to determine whether your images "look better" with or with out it.

Interpolation of an image does not necessarily increase the storage memory requirements of the digitized image even though the number of pixels may be increased by as much as a factor of 16 (4X interpolation in both axes) or even more. The interpolated information is computed from the real image data. Therefore, it may be computed at any time using image-processing software if we have the original digital image data available. Specifically, it does not have to be computed at the time of scanning. It may be just as accurately computed at the time of printing. Therefore, the basic scanned image may be stored without any interpolation, and at the time of printing the interpolation algorithms may be applied to provide an interpolated result in the printed image.

Not only does storing only the basic digitized image reduce the storage memory requirements, but it also allows us to apply whatever post processing, such as interpolation, we may desire at the actual time of printing to achieve a specific visual effect. For example, if we are printing the full-frame image in a large format to be viewed at a distance, we may not need any interpolation, but if we are cropping out a small part of the original image to be enlarged we may need a very high degree of interpolation to minimize pixelization in the printed image.

The basic scanned image is in effect the "raw data" of the original image. By storing this raw data, we avoid any additional corruption of the image (other than the fact that we digitized it spatially and digitized the intensity of its three primary colors) by such additional post processing as interpolation. If we always have the raw data available, we can then always apply new and wonderful post-processing techniques to that data to produce interesting printed subjects. However, once we apply some type of post processing to the raw data, generally the original raw data cannot be recovered from the post-processed result.

So, for archival purposes, it is most prudent to store the basic image from the scanner as faithfully as possible without any post processing, and do all post-processing image manipulation at the time of printing. If we wish to be able to reproduce a specific printed image at a later time, we need only record what we did to the raw data to obtain that print. This is precisely what a photographer does. When a negative is printed, the photographer does not in any way modify the actual negative. The negative is the photographer's raw data. The photographer will typically crop the enlarger image, not the negative, to the specific subject of interest and then "post process" on the image, for example by "burning in" dense areas of the image and "dodging" thin areas, to achieve the final desired visual results in the print. To be able to reproduce the print, the photographer simply makes written notes recording how the image was manipulated to achieve the final print, but never actually modifies the negative, the true "raw data."

Digitizing the image intensity

Now that we have a good idea of the spatial precision needed in the digitizing process, we can move on to the digitizing of the intensity information of the image. An individual can visually resolve about a 4 percent difference in the density of side-by-side samples. That is about one part in twenty-five. The dynamic range of film is measured in terms of Optical Density, or OD. The OD ratio is 10log of the ratio of the darkest image the film can produce to the lightest image that the film can provide.[7] The lightest the image can be in a film image is the film base, and the darkest is the density of the emulsion at maximum exposure.

A typical film such as Kodak Royal Gold 100 provides an optical-density range of about OD 2 to 2.5 for each color.[2] An OD 2 is a density range of 100:1 and OD 2.5 is about 316:1. At either end of the exposure curve (sensitometric or gamma curve), the response of the film becomes nonlinear, so the actual usable range of the film is slightly less than the full OD range.[7] An OD of about 2.4 is a convenient usable density since it corresponds to a density range of almost exactly 256:1, or eight bits. Even though the eye can only resolve about a 4-percent intensity variation, typical film is capable of resolving about one part in 256. Therefore, digitizing each color intensity of a pixel to eight bits captures effectively all of the useful intensity information in the film image. As noted above, a scanner that digitizes each color to eight bits is designated a 24-bit scanner (3 colors x 8 bits per color = 3D 24 = bits).

Even though an OD of 2.4 may be the maximum range typically usable in a typical film image, there may be some additional information in the nonlinear areas of an image. For example, in over-exposed areas of the negative or in very thin underexposed areas, there may be some useful detail. As noted below, we would not typically be able to see these details when the negative is normally printed. Nevertheless, the information may be available in the negative. So, a higher digitizing resolution of the intensity would be useful. But if the eye can only see about one part in 25, why would we want to digitize a negative even to 256:1 much less even higher?

The reason is to capture all the information that is in the image, not just that which we may see in a print. If we archive all the information, we can use the additional information to optimize the output image. For example, when printing a negative by hand, one can burn in dense areas and dodge thin areas to bring out the details in those regions of the negative in the print. So if we capture all that detail in a digitized image, we will have the same optimizing opportunities in the digital image when we "print" the digital image with a digital printer.

Eight bits of intensity information is very near the limit of what the film can actually provide. This is a very convenient resolution for storage with eight-bit digital words. However, it is possible to digitize the image intensity to 10 bits or even 12 bits. There is a reasonable compromise among available intensity-digitizing resolution, useful intensity range of the negative, and limitations in digital storage. Some scanner manufacturers provide 10 and 12-bit scanning (specified as 30 bit and 36 bit, of course), but the scanned image is post processed after scanning, and only 8 bits per color are actually stored.

By applying a mathematical function to the digitized image (this has also been referred to as a gamma function), information in the dense areas may be lightened and information in the thin areas may be darkened. The gamma of a film, γ, is the slope of the plot of optical density as a function of the log of the exposure. This mathematical function may be used to effectively modify the gamma characteristics of the scanned image. It is typically directly manipulated by the user to "optimize" the scanned image. The additional digitized bits of information are used as additional data to make modifications to the scanned image. Effectively, the 10 or 12-bit intensity information of each color provided by the scanner is compressed by the mathematical function into an 8-bit intensity word for each color for final storage.

Printing

Printing of a digitized image is quite different from the digitizing process. Because the ink-jet-type printers are by far the most common and least expensive, I will consider only those in this article. However, you should be aware that there are quite a number of other printing technologies. Since typical ink-jet printers can print only fixed-sized dots of fixed density, continuous-tone information must be printed as half-tone fields of dots with different numbers of dots per unit area (different dot densities). Therefore, one could define a print pixel as being made up of some number of possible dot positions. If a dot is printed in each possible position in a print pixel, that pixel will be of maximum density. If no dots are printed in the pixel, the pixel is obviously of minimum density.

We must determine what range of density is needed in typical printed images. For the purposes of this paper, we will assume that the three individual dots of the three primary colors are printed at the same point. Each color of a print pixel is individually half toned by the dot density of that color in the printed pixel. Be sure to keep in mind that a scanner pixel and a printer pixel are different; the scanner pixel is the smallest dot that the scanner can resolve, but the printer pixel is typically made up of many printer dots.

One convenient place to start is to examine reproducing all eight bits of intensity information recorded for each color. This would require 256 dots per print pixel, which can be provided in an 16 x 16 dot array. Earlier, I stated that a typical photographic print resolution is about 250 1p/in. If we have a printer that can print 250 dpi, it can exactly place 250 dots per inch in each printing direction. If the width and length of the dots are somewhat smaller than the dot and line spacing, a 250-dpi printer can just barely print 250 1p/in—each dot is part of a black line and the space between dots is a white line. Generally, the dots overlap somewhat so the printer is capable of printing solid colors without a discernible dot pattern. If adjacent dots just overlap, a slightly higher printer resolution will be required to reproduce the 250-1p/in line image. For a 250 1p/in line image, about a 300-dpi printer should be just capable of reproducing a 250 1p/in image.

It may seem that we do not need as high a resolution in a printer to print an image as we need in a scanner to capture an image. How-ever, the printer cannot print continuous-tone dots. So, although we may only need a printer that can print 300 1p/in, each printer pixel must be made up of many dots. For a 16 x 16 array, each pixel along a printed line must be broken into 16 dots, and there must be 16 lines per print pixel. Therefore, to produce a 300-1p/in image resolution with a 16 x 16 dot array for each print pixel required to reproduce an 8-bit gray scale, a printer with a 4800 dpi x 4800 dpi resolution is required. This is well beyond the present state of the art in typical ink-jet printing devices.

In the case of scanning, the goal was to record for archival purposes all the image information available in the negative or print being scanned. This required a scanning resolution consistent with the resolution of the negative or print. However, when we print an image, our primary purpose, if not only purpose, is to view the image. Therefore, we need only print at a minimum the information the eye can resolve. As noted earlier, the eye can resolve a density difference of about 4 percent, or about one part in 25. Therefore, if a 5 x 5 dot array is used for each print pixel, a 26:1 half-tone gray scale can be reproduced (one dot to 25 dots, plus no dots). This would require each pixel along a printed line to be broken into five dots, and five lines would be required for each print pixel. So, to print 300 1p/in with a 4-percent gray-scale resolution, a printer resolution of about 1500 dpi x 1500 dpi is needed.

This is just at the state of the art in ink-jet-type printers. Several printers are available that provide in excess of 1400 dpi x 1400 dpi color resolution. These printers should be capable of printing images that are virtually visually equal to a photograph in printed resolution. However, this only is true if these images are viewed as normal photographs are viewed. If viewed in close detail with a loupe or microscope, the individual dots forming the half-tone scale will likely be easily seen. A photograph is a continuous tone image, almost. So, on a microscopic scale, a 1500-dpi printed image will be very much more course than a photographic image. Nevertheless, about 300 printer pixels per inch at 5 dots x 5 dots per printer pixel will reasonably reproduce the full visual range of sharpness, color, and contrast available in a typical photographic print under normal viewing conditions.

As noted, there are a number of other printing technologies available. Some of these, such as the dye sublimation printers, are capable of providing very high resolution, but typically at a much higher cost than ink-jet printers, both in the printer itself as well as in its consumables. Also, some of the newer ink-jet printers are using more than four inks (typical inks are cyan, magenta, yellow, and black) to provide some tone variation within each printed dot. This reduces the number of dots needed in a half-tone pixel to provide the desired density resolution (~4 percent) in the printed image.

Before you purchase a photo-quality printer, you need to decide what quality you need (want), and investigate the various printer technologies that will work the best for you. Be sure to check the price of the consumables, such as ink, toner, and dye-transfer film, before you settle on a printer. The cost of these consumables can quickly and easily exceed the price of the entire printer.

Storage memory requirements

One of the more serious limiting parameters of digital image processing is the memory requirement for storing the images. There are a number of compression algorithms that can very impressively compress digital image files, but generally with some loss of image information. For our purposes here, we will assume that no compression is applied since it is our purpose to exactly reproduce the full digital image which we capture.

You will remember that a 24-bit scanner provides three 8-bit bytes of pixel intensity information for each pixel. Even the 30 and 36-bit scanners typically generate three 8-bit bytes for each pixel through a software algorithm. So, for this analysis, we will use three 8-bit bytes to represent the full color intensity information of each pixel. The memory required for storing various image formats is computed below.

Consider film digitizing first. About the most common film format is the 35-mm format, and about the largest is the 4 x 5-inch format used primarily by professionals. There are of course a number of format sizes in between, but the two considered here will provide an upper and lower limit of the memory requirements. A 35-mm negative is approximately 24 mm x 35 mm and the 4 x 5 negative about 102 mm x 127 mm. From above, the MTF of a typical color film is about 50 percent at a spatial resolution of about 50 lp/mm. If the negative is a professional image, we need to digitize at about a factor of five higher resolution than the actual film spatial resolution. As a result, for a 24 mm x 35 mm negative with a resolution of 50 lp/mm, we must scan at about 6000 x 8750 pixels. The total number of pixels is then 52.5 M pixels. And because each pixel contains 3 bytes of intensity information, a total of about 158 MB of storage is needed for a single professional 35-mm image. For a 4 x 5 negative, about 2.4 GB of storage is required.

Although these memory requirements may seem totally unmanageable, they are well within the current state of the art in storage media. For example, a single CD-ROM would store about four of these 35-mm images, and a DVD could store about seven 4 x 5 images. So, even though these are rather large files, it is well within the capability of presently available equipment to easily store them. Processing them however, may not be too convenient given the present state of the art in equipment and software.

From the contrast-reversal arguments explored earlier, we need a scanning resolution that is only about a factor of three higher than the image spatial resolution to reach this contrast-reversal reference point. Then about 57 MB is required to store a 35 mm image and about 871 MB for a 4 x 5 image. A CD-ROM can then hold about eleven 35-mm images and a DVD about nineteen 4x5 images.

Finally, if the images are typical amateur photographs, a scanning resolution that is about a factor of two greater than the film spatial resolution is more or less adequate. About 25 MB is required for a 35-mm negative and about 387 MB for a 4 x 5-inch negative. This is roughly the scanning resolution provided by presently available negative scanners like the HP and Nikon units referenced earlier. At this scanning resolution, a CD-ROM can hold about twenty-five 35-mm images and can actually hold one 4 x 5 image. A DVD could hold about forty-four 4 x 5 images and about 675 35-mm images.

If our purpose is to digitally archive our images to accurately preserve them and prevent degradation of the image information over time, the number of images that can be stored on a single medium is not too important. For example, even the limitation of storing only four 35-mm images on a CD-ROM at the maximum scanning resolution of a factor of five greater than the film MTF resolution is acceptable considering the current price of recordable CD-ROM media. At the 2X scanning resolution

Table 1. Memory requirements for various photographic resolutions (without compression).

Image scanned	Memory required
35 mm negative (50% MTF @ ~50 lp/mm)	
5X (6350 dpi)	158 MB
3X (3810 dpi)	57 MB
2X (2540 dpi)	25 MB
4 x 5 negative (50% MTF @ ~50 lp/mm)	
5X (6350 dpi)	2.4 GB
3X (3810 dpi)	871 MB
2X (2540 dpi)	387 MB
8 x 10 print (50% MTF @ ~10 lp/mm)	
5X (1270 dpi)	387 MB
3X (762 dpi)	139 MB
2X (508 dpi)	62 MB
11 x 14 print (50% MTF @ ~10 lp/mm)	
5X (1270 dpi)	745 MB
3X (762 dpi)	266 MB
2X (508 dpi)	119 MB

presently available from several film scanners, the capability of storing about 25 full 35-mm images on a CD-ROM is quite acceptable.

Now let's look at memory requirements for print digitizing. Although the typical resolution available in a print is much lower than film, about 10 lp/mm for prints as compared to 50lp/mm for film, prints are typically very much larger. Most larger scanners provide a scanning area of 8.5 x 14 inch and the smaller units 8.5 x 11 inches. The closest standard photographic print size is 8 x 10 inches.

Consider an 8 x 10 print (203 mm x 254 mm) having a 10 lp/mm resolution. At a 5X scanning resolution, the scanner resolution must be 1270 dpi x 1270 dpi. Then, for a nominal photographic resolution of 10 lp/mm, a storage memory of 387 MB is required at 5X scanning resolution, 139 MB at 3X and 62 MB at 2X. Now consider an 8.5 x 14 scanner (216 mm x 356 mm) and a full 8.5 x 14 photograph. The memory requirement at 5X is 577 MB, at 3X about 208 MB and at 2X about 92 MB. Finally, for an 8.5 x 11 scanner (216 mm x 279 mm) and a full 8.5 x 11 print providing a 10-lp/mm spatial resolution, 452 MB of storage memory is required at a 5X scanning resolution, 163 MB at 3X, and 72 MB at 2X. An 11 x 14 print (which may be scanned in a 11 x 17 "tabloid" scanner) will require 745 MB at the maximum scanning resolution.

Tables 1 and **2** present the memory requirements for several photographic formats and scanners. No interpolation or compression is considered in these tables. If interpolation is used, the memory requirements will increase substantially, and with compression memory requirements may be reduced, but perhaps at the expense of some loss of image information.

Digital cameras

Finally, a brief comment or two on digital cameras is in order. A high-performance consumer digital camera presently provides an image resolution of about 1200 pixels x 1000 pixels (mega-pixel cameras). Note that this is not pixels per millimeter or per inch but rather total pixels in each direction. Recall that the negative scanners referenced above provide a resolution of at least 2400 dpi, or about 3300 pixels x 2300 pixels for a 35-mm negative, which is just below the limit of the scanning resolution required to capture most of the image information of the negative. So the high-end consumer digital camera provides an image area that is about one sixth the image area of a scanned 35-mm image (1200/3300 x 1000/2300).

In other words, the digital camera image is effectively a smaller image format than 35 mm. It is roughly equivalent to a 16-mm format, perhaps slightly larger. Therefore, the printed image from this digital camera must be printed about one-half to one-third as large (one third the length and width) as a print from a digitized 35-mm negative to provide the same spatial pixel density in the print.

So, just like in "analog" photography when using a "small-format" negative, when using a digital camera we must be much more careful to capture only the subject in which we are interested and fully fill the frame with that subject. And because the image is smaller, even if we fully fill the frame with the desired image, we cannot print it as large as a larger-format negative.

Digital cameras are still quite a way from being capable of reproducing the performance of typical 35-mm film cameras. Above we noted that a typical negative film can resolve about 50\~lp/mm and that we need a spatial digitizing pixel resolution of about a factor of three higher than the image resolution. So, in a 24 mm x 35 mm negative, there are 1200 x 1750 resolvable lines. Then, to adequately capture the same spatial information as the film is capable of capturing, the digital-camera resolution must be about 3600 x 5250 pixels (a digitizing resolution that is a factor of three higher than the film resolution). So, the digital cameras are about a factor of four away from 35-mm film cameras.

Basically two breakthroughs are needed to make the digital cameras equivalent to film cameras. One is higher-resolution CCD arrays and the other is high-density, portable digital storage media. As noted in **Table 1**, a memory capacity of about 60 MB is needed to store the digital information from each digitized 35-mm image if digitized at about 3X to capture all the information that 35-mm film is capable of providing. The Super Disk technology presently available provides 100 MB storage on a 3.5-inch floppy-format disk. There are also ZIP drives, optical floppy drives (floptical drives), and perhaps even other technologies that could be used as well to provide this capacity. With the continual advances in processing of solid-state components such as CCD arrays, it is very likely that a nominal 4000 x 5000 pixel device will be commercially feasible in the near future. So, a true digital 35-mm camera may not be too far away.

Conclusions

From all this, we can see that the digital scanning technology available to the consumer at very competitive prices is just at the edge of the minimum resolution required to capture virtually all the image information in a negative (slide) or print. To really capture everything, perhaps an improvement of another factor of two or three in resolution is needed. But the resolution presently available is a reasonable compromise between capturing adequate image information and being able to conveniently store the resulting digital file with available storage technologies. Also, the printing technology is just at the limit of that needed to reproduce all the image information that can be visually discerned. So, in scanning and printing, the present technology is just at the limit of what is available in analog photography. There is still a little room for improvement, but the present technology is certainly capable of providing very impressive results.

But perhaps the most important feature of digital image storage is that the image never degrades. The contrast never diminishes, the colors never fade, and the image is always as pristine as when originally digitized. This is a very important feature in archiving images. Also, anyone may "print" digital images without a lot of costly equipment such as enlargers and lenses or peculiar chemicals, such as developers and fixers. All that is needed is a "photo-quality" color printer. Such a printer is likely a factor of five or 10 lower in price than a reasonable-quality color enlarger, without a lens. And, the cost of printing consumables will be much less than the cost of color-printing chemicals and color photographic paper. To complete the digital photography process, there is a wide range of digital image processing software available to provide anyone the ability to do very elaborate post processing of the digital image.

Digital cameras have a little farther to go to be truly equivalent to analog cameras. The standard analog camera format is the 35 mm. But even the high-end digital cameras are barely equivalent to a 16-mm format analog camera. The digital cameras are limited both by the size of available imaging electronics (CCD chips) and the memory limitations of the storage medium where the digitized image is stored. So, the digital-camera image size must be "improved" by about a factor of three in each dimension to be minimally equivalent to a 35-mm analog camera. We can be almost certain that it will happen, and very likely quite soon. ■

REFERENCES

1. Eastman Kodak, EXR 200T 5293 Film Technical Data Sheet.
2. Eastman Kodak, Kodak Royal Gold 100 Film Technical Data Sheet.
3. AGFA-GEVAERT AG, AGFAColor Optima 100 Negative Film Technical Data Sheet.
4. John Russ, *The Image Processing Handbook*, 2nd Edition, CRC Press, Boca Raton, 1995.
5. J. Bernd, *Practical Handbook on Image Processing for Scientific Applications*, CRC Press, Boca Raton, 1997.
6. Mischa Schwartz, *Information Transmission, Modulation and Noise*, McGraw Hill, New York, 1970.
7. Eastman Kodak, *PMI Student Filmmakers Handbook*, 1997.
8. Hewlett Packard, Hewlett Packard PhotoSmart Scanner Product Data Sheet.
9. Nikon, Nikon CoolScan II Scanner Product Data Sheet.
10. Maxtor, Maxtor 5120 DiamondMax Plus High-Performance Disk Drive Product Data Sheet.

Table 2. Memory requirements for various scanner systems (excluding interpolation and without compression).

Scanner	Memory Required
HP Photo Smart	
2400 x 2400 dpi	???
Nikon CoolScan—35mm	
2700 x 2700 dpi	???
8.5 x 11 scanner	
1200 x 1200 dpi	404 MB
600 x 600 dpi	101 MB
300 x 300 dpi	25 MB
8.5 x 14 scanner	
1200 x 1200 dpi	514 MB
600 x 600 dpi	129 MB
300 x 300 dpi	32 MB
11 x 17 scanner	
1200 x 1200 dpi	808 MB
600 x 600 dpi	203 MB
300 x 300 dpi	50 MB

TECHNICAL CONVERSATIONS

Checking the Smith Chart

Dear Editor:

In "Tech Notes" in the Summer 1999 issue, Steve Sparks, N5SV, describes a Smith Chart solution to his tower matching problem. I believe there is an error in his procedure.

On pages 102 and 105, he states:

> Starting at the antenna end, Z_L, draw a circle of constant impedance using a compass whose center is "0" on the far left of the horizontal resistance axis. The length...is from the worst case...to where it intersects the 1.0 constant resistance circle....

The objective of this step is to determine the shunt capacitance needed to place the transformed impedance (of the load and the shunt capacitance) on the unit resistance circle. But to do this requires travel along a circle of constant conductance. The set circles of constant conductance are all tangential to the "0" on the far left (not centered there) and a radius drawn from "0" does not define either a circle of constant conductance or a circle of constant impedance, as stated above.

Please observe the enclosed **Figure 1**, representing a skeletal Smith Chart. The circle labeled "R=1" is the unit circle of constant resistance. The circle labeled G, whose center is at X1, is a circle of constant conductance through load impedance Z_L. R0 is the radius described by Mr. Sparks, which defines the arc also passing through Z_L. Note that the constant conductance circle centered on X1 intersects the unit resistance circle at P3, while the arc drawn from the "0" point intersects at P2. Generally, the values of both the shunt capacitance and the series inductance will be different from those obtained by using Mr. Sparks' arc. (In his case, the difference is not large, and the use of a variable capacitor masks the error.)

It is not necessary to determine X1. A left-right flip of the chart produces an admittance chart instead of an impedance chart; the constant resistance circles are then constant-conductance circles. Overlaying an admittance chart over an impedance chart allows these two-element matching problems to be solved with ease; photocopying a Smith chart onto a transparency further simplifies the method.

The Smith Chart is an elegant, powerful tool which has been indispensable to engineers for six decades. If repeated matching exercises are anticipated, the purchase of Smith Chart software (such as the modestly priced *MicroSmith* program by Wes Hayward, W7ZOI, available from the ARRL) will both simplify the task and enlighten the user as to the capabilities of the method. Wideband matching, tank circuit design, tolerance analysis with multiple elements, and more, are readily accomplished with such software.

Garry Shapiro, NI6T
Los Gatos, California

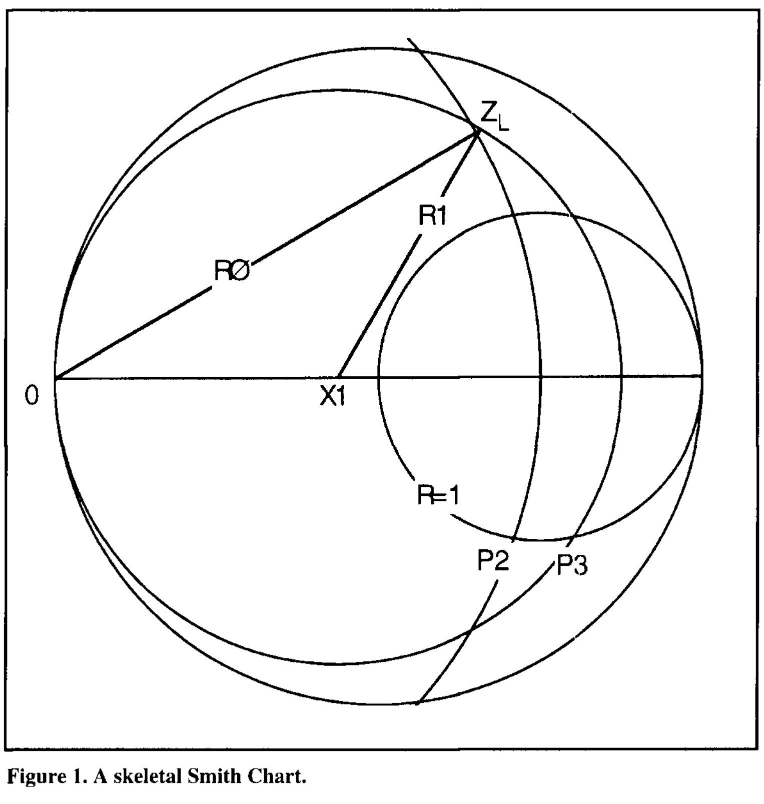

Figure 1. A skeletal Smith Chart.

Clarifying the image

Dear Editor:

I appreciated Mr. Gruchalla's article on image scanning and printing (*Communications Quarterly*, Summer 1999, page 9), especially since I am currently in the market for both a scanner and an ink jet printer, and am waiting for adequate quality digital cameras to become affordable. The numbers which he provided on existing analog processes has helped me to sort out what I really need. I would like to clarify two items which appear to be overlooked in the article, however, and to explore their impact on his conclusions.

First, the test pattern used to determine resolution is "digital" in nature. If it were digitized by a system with infinite bandwidth, the resulting spatial waveform would be a square wave, with infinite harmonics above the frequency determined by the line pair resolution. This is the reason for the artifacts he points out in scanning such an image. In reality, the analog process, like all analog systems, includes some inherent low-pass filtering (although possibly crude) which prevents these sort of artifacts from plaguing it. The resolution limits of the lens is one such filter, the random size and distribution of film grains is another. Were he to photograph the chart and scan the resulting negative or print, the artifacts would be diminished by this low-pass filtering.

When electronic engineers look at digital waveforms we insist on a bandwidth 10 times that of the fundamental if we wish to accurately capture the waveform. It should not be surprising that Mr. Gruchalla wants five times the spatial resolution under similar circumstances. The reason that two times the spatial resolution is adequate for most photographic purposes (besides the low-pass filtering inherent in the analog process which created the image) is that nature rarely provides us with "digital"-type images. Transitions are generally more gradual, requiring less bandwidth.

The second misconception has to do with half-toning and its relation to ink jet printing. The half-tone process used in printing divides each primary color of the image into discrete dots, much like pixels. However, the process relies on spatial limitations and non-linearities of the photographic process to produce dots whose size is a function of intensity, so that each dot carries intensity as well as spatial information. Several of the ink jet printers likewise modulate the dots to produce dots of different sizes depending on the required density (intensity). Such a printer would not need an array of 5 x 5 dots to represent each pixel including intensity. While it is true that the detail structure of any of these processes looks different than an analog photograph under high magnification, the information content can still be identical, and the results indiscernible from the analog counterpart under useful observation conditions. At high magnification, the observer is just comparing one type of artifact (pixels, dithering, etc.) with another (grain structure).

A related point has to do with the interpolated resolution of scanners (or other parts of the digital process). Nyquist's theorem is true in the strictest and most accurate sense only when there exists a reconstruction filter in the decode process. The raw data, while containing all of the information necessary to accurately reproduce the original, does not necessarily represent amplitude levels correctly, and they contain sampling artifacts. A reconstruction filter removes these artifacts and restores the proper amplitude relationships of the frequency components in the data. When properly done, interpolation of the data points in a scanner produces the same effect, and thus renders a more faithful copy of the original. The primary ingredient here is a low-pass filter.

Wilton Helm, WT6C
Franktown, Colorado

Gruchalla responds

Dear Wilton:

Thanks for your comments. My article is simply a brief review of the digital photography process. It was intended to provide a basic insight into the process and limitations. But, admittedly, I did not have the space to thoroughly address the subject (I am pretty sure that Terry considers it more than lengthy enough as it is.)

Your points are well taken; but I believe that I did address them in my piece. Specifically, I believe that I pointed out that the spatial resolving capabilities of most modern lenses far exceed the resolving ability of common films. Therefore, contemporary lenses of average quality do not actually contribute substantially to the analog filtering effect you recite. The image presented to the film, or CCD array, is indeed of quite high spatial frequency content in comparison to the ability to record that content. For film systems, the spatial limitation is primarily in the resolving ability of the film, but the effect is the same. The film provides a smoothing effect of the input image. That is why I based my analyses of scanning of film images on the limitations of the film. But CCD systems are not limited by film limitations, and CCD arrays presently provide much poorer spatial resolution than film. Therefore, if a high-spatial-frequency object is photographed with a CCD camera having a reasonable-quality lens, there will be insufficient spatial filtering to eliminate aliasing.

But you are absolutely correct concerning the smoothing effect. If one were to photograph a test target such as I used in the article (a "square test target" implying abrupt intensity change) and view the result, the sharp edges would be smoothed by the filtering of the resolution limitations of the film. Nevertheless, by using a "perfect" square test target in the article, I could more easily demonstrate scanning limitations by visual example. In reality, the test targets used to measure the spatial frequency response of film (and virtually all other optical elements) have a sinusoidal intensity profile of a specific constant contrast ratio and varying frequency. These are typically analyzed by precision densiometry measurements and mathematical processing, such as Fourier transform, to analytically measure performance. But these are very difficult to analyze visually.

The "square" intensity-function test target which I used allows very easy visualization of how the spatial quantization process affects the quantized image. Also, you may have noticed that I used the 1/2-power point of the film response as an analytic measure of its maximum frequency response. (This use of the 1/2-power point of the film is also a departure from the typical photographic art but common in engineering art.) Film actually provides considerable information beyond its 1/2-power point, so it is actually capable of reproducing fine structure well beyond my limitation of the 1/2-power point. The film will provide substantially sharper edges in the image of the square target than the 1/2-power point would predict. Therefore, my use of the square test target and the 1/2-power point in the analyses are complementary judgment calls—the 1/2-power point under specifies the film frequency capabilities and the square test target over-estimates the film step response. I believe that the aliasing processes that I described using the square test target can actually be reproduced with real film images of such a target object and contemporary scanners. The film provides insufficient smoothing to eliminate the aliasing.

Indeed, if one were to scan such a perfect square-intensity target, for example scan the actual chip of a large memory, the aliasing effects noted in the article would easily be seen. In such an example, the memory chip is the object, and there is no spatial filtering. If we image the chip onto the scanner, say with 1:1 conjugates to allow us to accurately focus through the window of an E-PROM, the lens would provide some filtering; but I argued that if the lens is of even modest quality by modern standards, the lens spatial resolving ability would be high enough that the lens effects would be negligible in comparison to present scanning technology.

One could of course photograph the memory chip and then scan the resulting image, either on film or on a print, in order to "smooth" the image; however, in the case of this example, that photographic image would not be the origi-

(Continued on page 105)

(from page 5)

nal object. But even in such a photograph, if the memory chip were photographed at a magnification providing an image with spatial frequency at the 1/2-power point of the film, the fine structure of the film image would be sufficient to produce aliasing when scanned with a contemporary scanner as reviewed in the article.

The intent of the article was to demonstrate how the scanner, and CCD, spatial quantization affect the results when the object scanned is of sufficient spatial frequency to challenge the limitations of the scanning system. At this point in the technology, I believe that film does indeed challenge contemporary scanning technology. I combined the concept of the square test target and the specified 1/2-power-point spatial resolution of the film to develop a type of maximum requirement for scanning resolution based only on the limitations of the film (or photographic paper) and totally independent of the actual object photographed. I believe this a reasonable estimate of the scanning resolution needed to preserve virtually all the analog intensity information in contemporary film images regardless of the subject.

Similarly, I believe that I addressed briefly the film grain structure. The typical grain size and spacing are well beyond the resolving ability of even the highest-resolution scanners, by perhaps as much as two orders of magnitude. Film resolving ability is limited by other parameters, such as scattering and dispersion of light in the emulsion. Therefore, the actual film or paper grain structure has a negligible effect on the scanned result.

As you point out, and as did I, I believe, most things we may photograph are comparatively smooth analog things (logs, dogs, cats, rats, etc.). But what if we choose to photograph a memory chip, or a high-contrast brick building at a distance? In those cases the objects are not smooth analog subjects, but indeed are very digital in nature. And the aliasing I demonstrated with the square test targets can be reproduced with these real photographic objects. Again, the focus of the article was from the perspective of the limitations of the optical materials and not the subjects, or how the subjects may be filtered to accommodate the scanning technology. The focus of the article was what scanning resolution is needed to preserve the full resolving capability of the analog photographic media independent of the subject being photographed.

Also, I believe that I did comment very briefly on the enhanced printer technologies, such as the HP Photo REt and the improved Photo REt II technology, but specifically limited my discussion to the simple ink-jet "dot matrix" half-tone systems. These are by far the most common. Also, it is much easier to visualize the dot-matrix half-tone process than something such as the Photo REt process. Once the basic process is visualized, one can more easily comprehend the more complex processes.

Also, there is considerable debate as to whether the modulated dot technologies, such as the HP Photo REt II, are visually equal to the micro-dot technology, such as used by Epson and Canon. Personally, I do not like the HP output. To me, it has a very obvious graininess caused by the fewer dots per unit length, regardless of their precision size modulation and spatial dithering. A very obvious example of this is the "banding" seen in images with very slow changing density across the image of comparatively bright subjects, such as clouds for example, and a mottling of continuous-tone features such as skin tones. It is true that over the entire image, the dot modulation does reproduce the average intensity, but the graininess is obvious, at least to me.

The Epson and Canon 1440 DPI systems, and even the 720 DPI systems for most images, also faithfully reproduce the intensity information, but with substantially less grain (assuming the system is working properly, which is not always the case). For example, I see virtually no mottling with these systems, whereas I do with the HP. So, based on my personal preference, I cannot agree with you that the modulated-dot print technologies are equal to the micro-dot technologies. But that is only my subjective personal preference. I use both systems in my work, but I use the micro-dot technology for my personal digital photos. I chose not to make these comparisons in the article both in the interest of space, and because I felt it inappropriate to criticize a particular manufacturer's technology.

But remember, the purpose of my article was to explore what is needed to capture all the information in an analog photographic subject—not what is simply "good enough" if viewed at a great enough distance. In other words, how good must be the scanning and printing technology be to produce a digital print which cannot be distinguished from a true analog photograph, even when the subject (not the medium) is viewed at high magnification?

Finally, the issue of interpolation is quite controversial. The bottom line is that interpolation adds information that is not there. One cannot take the Fourier transform of the image and then "replace" or "improve" the sidebands synthetically to improve the image to make it a more faithful representation of the object unless one knows what the side bands must be. That is a function of the object. But that object information was not recorded, and in general we have no way of recapturing that information. If we could, then we would have a higher-resolution image anyway.

One may guess at some object characteristics; for example, that the object was a generally a smooth analog subject, such as logs, dogs, cats, rats, etc., and synthetically interpolate based on some corresponding estimating function (i.e., a specific reconstruction filter), but that is just a guess. Some a-priori knowledge of the object is needed. What if the object were a memory chip? Since the memory chip is for most practical purposes a true spatially digital object, if we were to interpolate the digitized image to increase the apparent scanning resolution based on a reconstruction filter based on a guess of a smooth subject, the result would be nothing like the original object.

But this does not mean that interpolation is not useful. For example in audio CD players, such interpolation, or over sampling as it is termed, allows much simpler and less expensive filters to be utilized to reconstruct the audio signal. This over sampling, or audio interpolation, provides no improvement to the faithfulness of the reproduced representation of the original audio "image." It simply simplifies (and reduces the cost) of reconstructing the audio image. The quality of the audio image is fundamentally limited by the digitizing precision (number of bits), not the amount of over sampling in reproduction. The over sampling simply makes the hardware easier to build, more stable and robust, and lower in cost, all without degrading the audio quality limited by the digitizing depth.

Nevertheless, here too I cannot agree with you that interpolation of a scanned image, "when properly done," will produce a more faithful copy. I argue it will not. It will produce a smoother image, not a more faithful image—is a smoothed image of a memory chip a more faithful representation than an aliased one? But of course one will look better, regardless of whether it is a faithful representation. That is what the article was about: what do you need to totally faithfully reproduce the analog characteristics of the scanned subject—a film image, for example?

What I think all this means is that each individual must research the various technologies available in the digital photography art and use that most appropriate to their needs (photos of dogs or computer chips). My intent in the article was to provoke the reader to think about the process. I seem to have done that. But of course I welcome any additional comments or suggestions you may have. Perhaps we could actually "do some science" together to explore the limitations of the contemporary scanning art and analog limitations of photographic media, such as film and CCDs?

Mike Gruchalla
Editorial Review Board
Albuquerque, New Mexico

The objective value of interpolation

Dear Mike:

Thanks for the reply. As with the original article, the reply was enlightening, and I appreciate hearing from you. I think we are pretty much in agreement, and you obviously have considerably more direct experience with this. My attempt was to illuminate a couple of areas that other readers might have missed which you had touched on very briefly, if at all.

I still believe that some interpolation has objective value, in the manner you indicated with audio CDs. Digital video systems can sometimes show artifacts of the digitization and presentation processes, which are not in the original, like a 16-bit D/A on a CD that is not followed by a low-pass filter. The over-sampled lower bit converter, as you pointed out, takes less filtering to accurately reconstruct. If a video system lacks the filtering, then interpolation is a useful way of improving the situation. In neither case will the interpolation restore artifacts caused by aliasing because of lack of input filtering, but it can render a more accurate version of what actually is available, especially on an output device lacking in smoothing capacity.

An interesting variation of this is oversampling on the digitization process, where the dot size is larger than the DPI. Obviously some redundant information is being collected, but with proper processing, the results can exceed that predicted by the dot size.

Again, thanks for you reply. I got something I hadn't anticipated—additional insights into the process. I am at present looking for both a scanner and an ink jet printer for a variety of tasks including photographic.

Wilton Helm, WT6C
Franktown, Colorado

The value of interpolation

Dear Wilton:

You are absolutely right that interpolation can have real value—in making prints to be viewed at reasonable distance. By interpolating, the "pixelization" is reduced. No new data is added, but sharp pixel edges are fuzzied up a bit to allow the pixel information to run together. A somewhat fuzzy picture (photographers call that "soft focus") is much more pleasing to view than a sharply pixelized image.

However, I believe that one should apply interpolation only when making a print to be viewed. It should not be used for archiving information. The data for archiving should be as accurate and unmolested as the digitizing and recording technology allows. You can always fiddle with the raw data in the archive later as new tools become available. However, if you fiddle with the raw data and then archive it, you probably cannot un-fiddle it later. Hence the thrust of my article—just how good does scanning (and CCD) technology need to be to provide raw, "un-fiddled with" data that is "just as good as" the real analog source?

As for scanners and printers, I'll give you a couple of ideas. I looked at a bunch of scanners. I decided that I wanted at least a 600 x 600 unit and preferably 1200 x 1200. By the way, make sure the 600 x 600 or whatever is the true "optical" resolution, not "effective" (read interpolated) resolution. If you look at CompUSA and BestBuy, some times you can find a lost-leader sale where the scanner is actually free. A while back, CompUSA had a 300x300 parallel-port unit for sale for less than the rebate. Wasn't much of a scanner, but certainly worth the price. And something like that is a good place to start if you have no idea what you need (or want). But these "low-end" units actually work surprisingly well. I ended up getting an HP6270A with the document feeder (didn't need the feeder, but the unit was on sale for a good price). This is a 1200 x 1200 unit and it was lower in cost (much) than the other 1200 x 1200 units at the time. As scanners go, this one is a bit pricey, about $300, but that is in line with the high-res units. Works great!

As for printers, good luck! I must have looked at every printer made, twice. Some thoughts. I presume you want color. Color lasers are out! They cost a lot, are expensive to run, and produce poor results. Also, stay away from the weird technologies—dye sublimation, wax transfer, etc. Some of these work quiet well, but they cost quite a bit to operate and the prints may not be very durable. That pretty much leaves the ink jet printers. As I mentioned in my previous note, I just do not like the HP images (personal preference). The Epson, Canon, Lexmark, etc. models with the 7-ink system all seem to be about equal as far as resolution. It is pretty easy to get 720 x 720 dpi now and may be even 1440 x 1440. The 720 x 720 seems to be just about all you need, but there is still just a bit of pixelization. Also there are compromise units with 1440 x 720. These are okay too. The trick is finding something that will actually keep working and does not consume an entire set of cartridges on a single print.

In general the printer guys should give the printers away—the consumables eat you alive. Some of the printers have the print head in the machine (Epson) and some in the ink cartridge (HP and Canon). I prefer the head-in-cartridge configuration; it's easy to "replace" the head when (there is no "if" here) it gets clogged. I ended up getting a Canon BJC-5000 1440 x 720 unit (this is now obsolete, replaced by the BJC-5100 I think—the same basic machine). This is the only machine that I could go anywhere in town and find it working where I could print a demo. Virtually all other machines were screwed up in one way or another. I have had virtually no trouble with it at all, I even let it sit for about three months and then it printed perfect pix. I believe that the newer unit (5100 as I recall) is about $150 more or less; that's about the ballpark of all the units. And this unit will do 11 x 17 prints to boot—great for B-size drawings.

You need to go to every place near you that sells printers and do some demos yourself. If possible, take along your own image file to print; everyone here was quite happy to let me use my own demo picture. Then you can compare the resolution and color (accuracy and saturation) of all the different machines (at least the ones you can find that are working!) To do this comparison, you probably should purchase some of the high-resolution paper (not glossy, just high res). Good luck!

Mike Gruchalla
Editorial Review Board
Albuquerque, New Mexico