

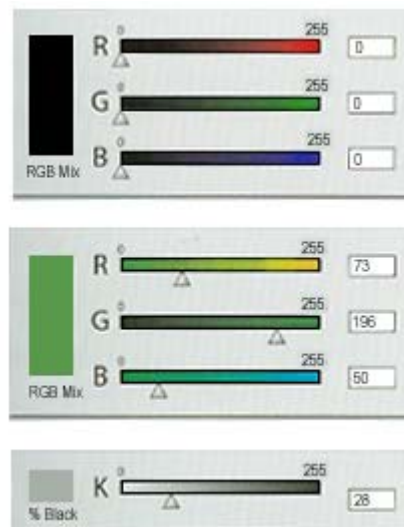
The Image

Image Reproduction and Color Management

The human eye can distinguish millions of different colors, all of which arise from two types of light mixtures: additive or subtractive. The former involves adding together different parts of the light spectrum, while the latter involves the subtraction or absorption of parts of the spectrum, allowing the transmission or reflection of the remaining portions. Computer monitors exploit an additive system, while print color creation is subtractive. This fundamental difference can complicate both accurate reproduction on a computer monitor of the colors of an original work and accurate printing of a digital image.

On a typical video monitor, as of this writing, color is formed by the emission of light from pixels, each of which is subdivided into three discrete **subpixels**, which are in turn responsible for emitting one of the three primary colors: red, green, or blue. Color creation occurs when beams of light from each color **channel** are combined; by varying the voltage applied to each subpixel individually, thus controlling the intensity of light emitted, a full range of colors can be reproduced, from black (all subpixels off) to white (all subpixels emitting at full power). This is known as the **RGB color model** (fig. 1).

Fig. 1. RGB color is a mix of red, green, and blue. Gray scale is a percentage of black.



In print, however, color is created by the reflection or transmission of light from a substrate (such as paper) and layers of colored dyes or pigments, called inks, formulated in the three primary subtractive colors—cyan, magenta, and yellow (CMY). Black ink (K) may be additionally used to aid in the reproduction of darker tones, including black. This system is known as the **CMYK** color model. Printed images are not usually composed of rigid matrices of pixels but instead are created by overprinting some or all of these four colors in patterns that simulate varying color intensities by altering the size of the dots that are printed, in contrast with the substrate, through a process called halftoning. (There are digital printers that combine colors from the CMYK and RGB color models or add gray ink in order to make up for deficiencies in printer inks in representing a wide range of colors.)

Admittedly, this is a highly simplified overview of color. There are many different color models and variations thereof—**HSB/HLS**, which describes colors according to hue, saturation, and brightness/lightness; and **gray scale** (fig. 1), which mixes black and white to produce various shades of gray, are two common systems—and the various devices that an image encounters over its life cycle may use different ones. Variation among

different display or rendering devices, such as monitors, projectors, and printers, is a particularly serious issue: a particular shade of red on one monitor will not necessarily look the same on another, for example. Brightness and contrast may also vary. The International Color Consortium (**ICC**) has defined a standardized method of describing the unique characteristics of display, output, and working environments—the ICC Profile Format—to facilitate the exchange of color data between devices and mediums and ensure color fidelity and consistency, or **color management**. An ICC **color profile** acts as a translator between the **color space** of individual devices and a device-independent color space (**CIE LAB**) that is capable of defining colors absolutely. This allows all devices in an image-processing workflow to be **calibrated** to a common standard that is then used to map colors from one device to another. Color management systems (**CMS**), which are designed for this purpose, should be selected on the basis of their support for the ICC Profile Format rather than competing proprietary systems.

ICC profiling ensures that a color is correctly mapped from the input to the output color space by attaching a profile for the input color space to the digital image. However, it is not always possible or desirable to do this. For instance, some file formats do not allow color profiles to be embedded. If no instructions in the form of tags or embedded profiles in the images themselves are available to a user's **Web browser**, the browser will display images using a default color profile. This can result in variation in the appearance of images based on the operating system and color space configuration of the particular monitor. In an attempt to address this problem, and the related problem of there being many different RGB color spaces, Hewlett-Packard and Microsoft jointly developed **sRGB**, a calibrated, standard RGB color space wherein RGB values are redefined in terms of a device-independent color specification that can be embedded during the creation or derivation of certain image files. Monitors can be configured to use sRGB as their default color space, and sRGB has been proposed as a default color space for images delivered over the **World Wide Web**. A mixed sRGB/ICC environment would use an ICC profile if offered, but in the absence of such a profile or any other color information, such as an alternative platform or application default space, sRGB would be assumed. Such a standard could dramatically improve color consistency in the desktop environment.

Bit Depth/Dynamic Range

The **dynamic range** of an image is determined by the potential range of color and luminosity values that each pixel can represent in an image, which in turn determines the maximum possible range of colors that can be represented within an image's color space or **palette**. This may also be referred to as the **bit depth** or **sample depth**, because digital color values are internally represented by a binary value, each component of which is called a **bit** (from *binary* digit). The number of bits used to represent each pixel, or the number of bits used to record the value of each sample, determines how many colors can appear in a digital image.

Dynamic range is sometimes more narrowly understood as the ratio between the brightest and darkest parts of an image or scene. For instance, a scene that ranges from bright sunlight to deep shadows is said to have a high dynamic range, while an indoor scene with less contrast has a low dynamic range. The dynamic range of a capture or display device dictates its ability to describe the details in both the very dark and very light sections of the scene.

Early monochrome screens used a single bit per pixel to represent color. Since a bit has two possible values, 1 or 0, each pixel could be in one of two states, equivalent to being on or off. If the pixel was "on," it would glow, usually green or amber, and show up against the screen's background. The next development was 4-bit color, which allows 16 possible colors per pixel (because 2 to the 4th power equals 16). Next came 8-bit color, or 2 to the 8th power, allowing 256 colors (compare figs. 2 and 3). These color ranges allow simple graphics to be rendered—most icons, for example, use either 16 or 256 colors—but are generally inadequate for representing photographic-quality images.

Fig. 2. The full color spectrum



Fig 3. The spectrum in 256 colors



The limitations of 256-color palettes prompted some users to develop **adaptive palettes**. Rather than accepting the generic system palette, which specified 256 fixed colors from across the whole range of possible colors, optimal sets of 256 colors particularly suited or adapted to the rendering of a given image were chosen. So, for example, instead of a fixed palette of 256 colors divided roughly equally across the color spectrum (leaving perhaps eight shades of green), the 256 colors might be primarily devoted to greens and blues in an image of a park during summer, or to shades of yellow and gold for an image depicting a beach on a sunny day. While they may enhance the fidelity of any given digital image, adaptive palettes can cause problems. For instance, when multiple images using different palettes are displayed at one time on a system that can only display 256 colors, the system is forced to choose a single palette and apply it to all the images. The so-called **browser-safe palette** was developed to make color predictable on these now largely obsolete 256-color systems. This palette contains the 216 colors whose appearance is predictable in all browsers and on Macintosh machines and IBM-compatible or **Wintel** personal computers (the remaining 40 of the 256 colors are rendered differently by the two systems), so the browser-safe selection is optimized for **cross-platform** performance. While this palette is still useful for Web page design, it is too limited to be of much relevance when it comes to high-quality photographic reproduction.

Sixteen-bit color offers 65,000 color combinations. In the past this was sometimes called "high color," or "thousands of colors" on Macintosh systems, and is still used for certain graphics. Twenty-four-bit color allows every pixel within an image to be represented by three 8-bit values ($3 \times 8 = 24$), one for each of the three primary color components (channels) in the image: red, green, and blue. Eight bits (which equal one **byte**) per primary color can describe 256 shades of that color. Because a pixel consists of three primary color channels, this allows the description of approximately 16 million colors ($256 \times 256 \times 256 = 16,777,216$). This gamut of colors is commonly referred to as "**true color**," or "millions of colors" on Macintosh systems.

As of this writing, 24-bit color display is the highest bit depth obtainable by affordable monitors; although many monitors now offer what is called 32-bit display, this is actually 24 bits of color data and 8 bits of "alpha" or transparency data. It is in fact debatable whether many monitors can even display the full range of 24-bit color, but most do accept 24-bit video signals from their system's **video card**, the circuit board that enables a computer to display information. Experimental monitors that can display 30-bit color (10 bits per color channel) have been demonstrated, and it is possible that such monitors will become more generally available in the future. (The ability of most printers to accurately represent higher bit depths is also limited.)

Given the limitations of computer monitor display, the advantages of capturing any image at greater than 24-bit color may not be obvious, but many institutions are moving toward 48-bit-color **image capture** for archival purposes. This extends the total number of expressible colors by a factor of roughly 16 million, resulting in a color model capable of describing 280 trillion colors. Such "high-bit" or high dynamic range imaging (**HDRI**)—that is, imaging that exploits bit depths of 48, 96, or even higher—actually uses the "extra" bits less to capture ever more colors than to render differences in light and shade (luminance) more accurately. The primary purpose of doing so is to preserve as much original data as possible: since many **scanners** and **digital cameras** capture more than 24 bits of color per pixel, using a color model that can retain the additional precision makes sense for image archivists who wish to preserve the greatest possible level of detail. Additionally, using a high-bit color space presents imaging staff with a smoother palette to work with, resulting in less color banding and cleaner editing and **color correction**.

The following set of images shows the effect of differing levels of sample depth both on the appearance of a digital image and on the full size of the image file. The examples, all of which were captured from a 4-by-5-inch photographic transparency at a resolution of 300 samples per inch (see Resolution), are shown magnified, for comparison.



Fig 4. (above and following page). Francesco Panini (Italian, 1745-1812), from *Il Prospetto del Castel' S'Angiolo con lo sparo della Griandola* (detail) ca. 1780-85. Etching with watercolor and gouache, 58.3 x 88.4 cm (22 15/16 x 34 13/16 in.). Research Library, The Getty Research Institute, 970025.



4-bit (16 colors), 2.55 megabytes (RGB)



8-bit (256 colors), 5.1 megabytes (RGB)



24-bit (16 million colors), 15.3 megabytes (RGB)



48-bit (280 trillion colors), 30.6 megabytes (RGB)

In each example, take note of the high-contrast areas and the color of the clockface. Higher bit images show greater detail and color range.

Resolution

Resolution—usually expressed as the density of elements, such as pixels, within a specific area—is a term that many find confusing. This is partly because the term can refer to several different things: **screen resolution, monitor resolution, printer resolution, capture resolution, optical resolution, interpolated resolution, output resolution**, and so on. The confusion is exacerbated by the general adoption of the dpi (dots per inch) unit (which originated as a printing term) as a catchall measurement for all forms of resolution. The most important point regarding resolution is that it is a relative rather than an absolute value, and therefore it is meaningless unless its context is defined. Raster or bitmapped images are made up of a fixed grid of pixels; unlike scalable vector images, they are resolution-dependent, which means the scale at which they are shown will affect their appearance. (For example, an image that appears to contain smoothly graduated colors and lines when displayed at 100% scale will appear to be made up of discontinuous, jagged blocks of color when displayed at 200%.)

Screen resolution refers to the number of pixels shown on the entire screen of a computer monitor and may be more precisely described in pixels per inch (ppi) than dots per inch. The number of pixels displayed per inch of a screen depends on the combination of the monitor size (15 inch, 17 inch, 20 inch, etc.) and display resolution setting (800 x 600 pixels, 1024 x 768 pixels, etc.). Monitor size figures usually refer to the diagonal measurement of the screen, although its actual usable area will typically be less. An 800-by-600-pixel screen will display 800 pixels on each of 600 lines, or 480,000 pixels in total, while a screen set to 1024 x 768 will display 1,024 pixels on each of 768 lines, or 786,432 pixels in total, and these pixels will be spread across whatever size of monitor is employed. An image displayed at full size on a high-resolution screen will look smaller than the same image displayed at full size on a lower-resolution screen.

It is often stated that screen resolution is 72 dpi (ppi) for Macintosh systems, or 96 dpi (ppi) for Windows systems: this is not in fact the case. These figures more properly refer to monitor resolution, though the two terms are often used interchangeably. Monitor resolution refers to the maximum possible resolution of given monitors. Higher monitor resolution indicates that a monitor is capable of displaying finer and sharper detail, or smaller pixels. Monitor detail capacity can also be indicated by **dot pitch**—the size of the distance between the smallest physical components (phosphor dots) of a monitor's display. This is usually given in measurements such as 0.31, 0.27, or 0.25 millimeters (or approximately 1/72nd or 1/96th of an inch) rather than as a per inch value.

Printer resolution indicates the number of dots per inch that a printer is capable of printing: a 600-dpi printer can print 600 distinct dots on a one-inch line. Capture resolution refers to the number of samples per inch (spi) that a scanner or digital camera is capable of capturing, or the number of samples per inch captured when a particular image is digitized. Note the difference between optical resolution, which describes the values of actual samples taken, and interpolated resolution, which describes the values that the capture device can add between actual samples captured, derived by inserting values between those recorded; essentially the scanner "guesses" what these values would be. Optical resolution is the true measure of the quality of a scanner. Pushing a capture device beyond its optical resolution capacity by **interpolation** generally results in the introduction of

"dirty" or unreliable data and the creation of larger, more unwieldy files. Moreover, generally speaking, when interpolation is required, image-processing software can do it more effectively than can capture devices.

Effective resolution is a term that is used in various contexts to mean rather different things. Generally it refers to "real" resolution under given circumstances, though users should beware of it being used as a substitute term for interpolated resolution in advertisements for scanners. The effective resolution of a digital camera refers to the possible resolution of the photosensitive capture device, as constrained by the area actually exposed by the camera lens. The term is also used to describe the effect of scaling or **resizing** on a file. For instance, a 4-by-6-inch image may be scanned at 400 spi at a scale of 100%—but if the resultant image file is reduced to half size (in a page layout, for instance), its effective resolution will become 800 dpi, while if it is doubled in size, its effective resolution will become 200 dpi. Effective resolution may also be used when accounting for the size of the original object or image when deciding upon capture resolution, when scanning from an intermediary. For example, a 35mm (1.5-inch) negative of a 4-by-6-inch original work would have to be scanned at 2400 spi to end up with what is effectively a 600-spi scan of the original. This number is arrived at through the formula: (longest side of the original x the desired spi) / longest side of the intermediary.

The density of pixels at a given output size is referred to as the output resolution: each type of output device and medium, from monitors to laser printers to billboards, makes specific resolution demands. For instance, one can have an image composed of 3600 pixels horizontally and 2400 pixels vertically, created by scanning a 4-by-6-inch image at 600 spi. However, knowing this gives no hints about the size at which this image will be displayed or printed until one knows the output device or method and the settings used. On a monitor set to 800 x 600 screen resolution, this image would need some four-and-a-half screen lengths to scroll through if viewed at full size (actual size as measured in inches would vary according to the size of the monitor), while a 300-dpi printer would render the image—without modification—as 8 by 12 inches. During digitization, the output potential for an image should be assessed so that enough samples are captured to allow the image to be useful for all relevant mediums but not so much that the cost of **storage** and handling of the image data is unnecessarily high. Many **digitizing** guidelines specify **image resolution** via horizontal and vertical axis pixel counts, rather than a per inch measurement, because these are easier to apply meaningfully in different circumstances.

As discussed in earlier sections (See *Image Reproduction and Color Management* and *Bit Depth/Dynamic Range*), output devices are currently the weakest link in the image-quality chain. While images can be scanned and stored at high dynamic range and high resolution, affordable monitors or projectors are not available at present to display the full resolution of such high-quality images. However, improved output devices are likely to become available in the coming years.

The following set of images shows the effect of differing levels of capture resolution both on the appearance of a digital image and on the full size of the image file. The examples, all of which were captured from a 4-by-5-inch photographic transparency at a resolution of 300 spi and a bit depth of 24, are shown magnified, for comparison.



Fig. 5 (above and following pages). Maria Sibylla Merian (German, 1647-1717), *Pomegranate* . Color engraving, 53.3 X 37.8 cm (21 X 14 7/8 in.), from: *Metamorphosis insectorum Surinamensium* . (Amsterdam: Joannes Oosterwyk, 1719). Research Library, The Getty Research Institute, 89-10750



800 spi, 56 megabytes (RGB)



400 spi, 14 megabytes (RGB)



200 spi, 3 megabytes (RGB)



100 spi, .875 megabyte (RGB)

Compare the antenna in these five examples. Images captured at a higher resolution allow finer detail.



50 spi, .22 megabyte (RGB)

Compression

Image **compression** is the process of shrinking the size of digital image files by methods such as storing redundant data (e.g., pixels with identical color information) more efficiently or eliminating information that is difficult for the human eye to see. Compression **algorithms**, or **codecs** (compressors/decompressors), can be evaluated on a number of points, but two factors should be considered most carefully: compression ratios and generational **integrity**. Compression ratios are simple comparisons of the capability of schemes, expressed as a ratio of compressed image size to uncompressed size; so, a ratio of 4:1 means that an image is compressed to one-fourth its original size. Generational integrity refers to the ability of a compression scheme to prevent or mitigate loss of data—and therefore image quality—through multiple cycles of compression and **decompression**. In the analog world, generational loss, such as that incurred when duplicating an audiocassette, is a fact of life, but the digital realm holds out at least the theoretical possibility of perfect duplication, with no deterioration in quality or loss of information over many generations. Any form of compression is likely to make long-term generational integrity more difficult; for this reason it is recommended that **archival master** files, for which no intentional or unavoidable degradation is acceptable, be stored uncompressed if possible.

Lossless compression ensures that the image data is retained, even through multiple compression and decompression cycles, at least in the short term. This type of compression typically yields a 40% to 60% reduction in the total data required to store an image, while not sacrificing the precision of a single pixel of data when the image is decompressed for viewing or editing. Lossless schemes are therefore highly desirable for archival digital images if the resources are not available to store uncompressed images. Common lossless schemes include **CCITT** (a standard used to compress fax documents during transmission) and **LZW** (Lempel-Ziv-Welch, named for its creators and widely used for image compression). However, even lossless compression is likely to complicate decoding the file in the long term, especially if a proprietary method is used, and it is wise to beware of vendors promising "lossless compression," which may be a rhetorical, rather than a scientific,

description. The technical metadata accompanying a compressed file should always include the compression scheme and level of compression to facilitate future decompression.

Lossy compression is technically much more complex because it involves intentionally sacrificing the quality of stored images by selectively discarding pieces of data. Such compression schemes, which can be used to derive access files from uncompressed (or losslessly compressed) master files, offer a potentially massive reduction in storage and **bandwidth** requirements and have a clear and important role in allowing access to digital images. Nearly all images viewed over the Web, for instance, have been created through lossy compression, because, as of this writing, bandwidth limitations make the distribution of large uncompressed or losslessly compressed images impractical. Often, lossy compression makes little perceptible difference in image quality. Many types of images contain significant natural noise patterns that do not require precise reproduction. Additionally, certain regions of images that would otherwise consume enormous amounts of data to describe in their totality may contain little important detail.

Lossy compression schemes attempt to strike a balance between acceptable loss of detail and the reduction in storage and bandwidth requirements that are possible with these technologies. Most lossy schemes have variable compression, meaning that the person performing compression can choose, on a sliding scale, between image quality and compression ratios, to optimize the results for each situation. While a lossless image may result in 2:1 compression ratios on average, a lossy scheme may be able to produce excellent, but not perfect, results while delivering an 8:1 or even much greater ratio, depending on the type and level of compression chosen. This could mean reducing a 10-megabyte image to 1.25 megabytes or less, while maintaining more-than-acceptable image quality for all but the most critical needs.

Not all images respond to lossy compression in the same manner. As an image is compressed, particular kinds of visual characteristics, such as subtle tonal variations, may produce **artifacts** or unintended visual effects, though these may go largely unnoticed due to the random or continuously variable nature of photographic images. Other kinds of images, such as pages of text or line illustrations, will show the artifacts of lossy compression much more clearly, as the brain is able to separate expected details, such as straight edges and clean curves, from obvious artifacts like halos on high-contrast edges and color noise. Through testing and experience, an image manager will be able to make educated decisions about the most appropriate compression schemes for a given image or set of images and their intended users. It is important to be aware that artifacts may accumulate over generations-especially if different compression schemes are used, perhaps as one becomes obsolete and is replaced by another-such that artifacts that were imperceptible in one generation may become ruinous over many. This is why, ideally, uncompressed archival master files should be maintained, from which compressed **derivative files** can be generated for access or other purposes. This is also why it is crucial to have a metadata capture and update strategy in place to document changes made to digital image files over time.

File Formats

Once an image is scanned, the data captured is converted to a particular file format for storage. File formats abound, but many digital imaging projects have settled on the formula of **TIFF** master files, **JPEG** derivative or access files, and perhaps **GIF thumbnail** files. Image files automatically include a certain amount of technical information (technical metadata), such as pixel dimensions and bit depth. This data is stored in an area of the file (defined by the file format) called the **header**, but much of the information should also be stored externally.

TIFF, or Tagged Image File Format, has many desirable properties for preservation purposes. "Tagged" refers to the internal structure of the format, which allows for arbitrary additions, such as custom metadata fields, without affecting general compatibility. TIFF also supports several types of image data compression, allowing an organization to select the most appropriate codec for their needs, and many users of TIFF opt for a lossless compression scheme such as LZW to avoid any degradation of image quality during compression. Archival users often choose to avoid any compression at all, an option TIFF readily accommodates, to ensure that image data will be simple to decode. However, industry-promoted *de facto* standards, like TIFF, are often implemented inconsistently or come in a variety of forms. There are so many different implementations of TIFF that many applications canread certain types of TIFF images but not others. If an institution chooses such an industry-promoted standard, it must select a particular version of the standard, create clear and consistent rules as to how the institution will implement the standard (i.e., create a data dictionary defining rules for the contents of each field), and make sure that all user applications support it. Without clear consensus on a particular standard implementation, both interoperability and information exchange may be at risk.

The JPEG (Joint Photographers Experts Group) format is generally used for online presentation because its compression is extremely efficient while still giving acceptable image quality. It was developed specifically for

high-quality compression of photographic images where minor perturbations in detail are acceptable as long as overall aesthetics and important elements are maintained. However, JPEG compression is lossy, so information is irretrievable once discarded, and JPEG compression above about 25% often creates visible artifacts. The format that most people know as JPEG is in fact **JFIF** (JPEG File Interchange Format), a public domain storage format for JPEG compressed images. JFIF is a very simple format that does not allow for the storage of associated metadata, a failing that has led to the development of **SPIFF** (Still Picture Interchange File Format), which can be read by JPEG-compliant readers while providing storage for more robust metadata. GIF (Graphics Interchange Format) uses LZW lossless compression technology but is limited to a 256-color (adaptive) palette.

It is possible that the status of TIFF as the de facto standard format for archival digital image files will be challenged by another format in the near future that will be able to serve both master and access functions. Two possible candidates are **PNG** (Portable Network Graphics) and **JPEG2000**. PNG was designed to replace GIF. It supports 24- and 48-bit color and a lossless compression format and is an ISO/IEC standard. Application support for PNG is strong and growing. By contrast, JPEG2000 uses **wavelet compression**, which offers improved compression with greater image quality. It also allows for lossless compression and for the end user to specify resolution to accommodate various bandwidths, monitors, and browsers. The JPEG2000 standard defines two file formats, both of which support embedded XML metadata: JP2, which supports simple XML; and JPX, which has a more robust XML system based on an embedded metadata initiative of the International Imaging Industry Association: the DIG35 specification. However, as of this writing, commercial implementations for JPEG2000 are just beginning to appear.

The following set of images demonstrates the quality and full size of an image file uncompressed and under various compression schemes. The examples are shown magnified, for comparison. The original image was captured from a 4-by-5-inch photographic transparency at a resolution of 400 spi using 24-bit color.



Fig. 6 (above and following pages). Jean-Etienne Liotard (Swiss, 1702-1789), *Maria Frederike van Reede-Athlone at Seven Years of Age*, 1755-56. Pastel on vellum, 57.2 x 47 cm (22 ½ x 18 ½ in.). Los Angeles, J. Paul Getty Museum, 83.PC.273



TIFF file, 18.4 megabytes (uncompressed)
TIFF is an excellent format for master files, which may be up to four gigabytes in size and whose quality is generally limited only by the capabilities of the capture device. Neither uncompressed nor compressed TIFFs are ideal for Web delivery, as the files tend to be large and not supported by all browser.



JPEG file, 3.3 megabytes (lossy compression, 5.5:1 ratio)



JPEG file, 488 kilobytes (lossy compression 38.6:1 ratio)
JPEGs provide high-quality lossy compression of richly-colored continuous-tone images lacking definite lines or sharp contrasts between colors. JPEG is not ideal for images with areas of solid color and clear edges, and at higher compression ratios can create artifacts such as halos, "clockiness," and color banding.



GIF file, 4.6 megabytes (lossy compression, 4:1 ratio) A lossy-compression format best suited for images that contain clean lines and large areas of solid color, GIFs are often used to create thumbnail images and Web graphics. GIF is not recommended for continuous-tone images because it removes colors and replaces them with the nearest approximation in its limited palette.



JPEG2000 file, 332 kilobytes (lossy compression, 56.8:1 ratio)
JPEG2000, though a relatively new format that is not yet widely supported, offers both lossless compression and greater lossy compression at higher quality than JPEG. JPEG2000 is best used for continuous-tone images that do not contain fine detail or "noisy" textures because at higher compression ratios it can produce artifacts such as blurriness and ringing.