

JPEG compression and scientific imaging

John C. Russ

Materials Science and Engineering Department, North Carolina State University, Raleigh, NC

Images are large. Even a modestly priced consumer pocket camera records images with 5-10 million pixels. Although memory prices have dropped significantly in recent years, there is still a desire to reduce the cost of in-camera storage by compressing the images. As the internet connections most of us enjoy have gotten faster, sending large files has become easier, but still, sending or receiving a large image file takes time. There is obvious interest in reducing the file size. Because it became widely established when higher memory costs and slower transmission speeds had a large impact, compression using the JPEG (joint photographers expert group) algorithm has remained with us and is still used. It is built into the firmware of most cameras and is supported by virtually all current imaging software. However, the use of JPEG compression is very unfortunate, and can create real problems for scientific imaging.

JPEG compression was designed to reduce image file sizes by discarding details and modifying pixel color and brightness values in ways that would not significantly affect people's ability to recognize familiar objects in everyday photographs. Human vision is not very sensitive to variations in color (we more-or-less automatically compensate for differences in appearance, for example, under sunlight, fluorescent lighting, or incandescent bulbs). Furthermore, the exact dimensions of objects are generally not important for their recognition, and in any case the object boundaries are defined visually more in terms of brightness than color. Consequently, the JPEG algorithm discards much more color information than brightness. It also shifts boundaries, and removes fine detail that would not be noticed in a casual inspection. Furthermore, since each tiny portion (8x8 pixel blocks) of the image is processed separately, different amounts of detail will be lost in different parts of the image.



Figure 1. Market scene photograph.

As an example, consider the photograph of a market scene in Figure 1 (a portion of an original image that has been cropped to 1000x750 pixels). When viewing the overall image, context and familiarity make it possible to recognize the vegetables and even some of the people. JPEG compression on this image reduced the size of the file from 1.8 MB to 74 kB, a 24:1 reduction. Obviously that facilitates storage and transmission, but the effects on the image contents are high.

Figure 2 shows side-by-side comparisons of enlarged details from two regions in the image. The fine lines (lettering on the sign), location of boundaries and edges, sizes and numbers of features, and the color values of pixels have been altered. In addition, some new lines (e.g., on the sign) have been added. Analysis of the compressed image for scientific or forensic purposes has been seriously compromised, and even more troublesome, it is not possible to determine or predict what information has been altered, added or removed.



Figure 2. Comparison of enlarged portions of the original (left) and jpeg-compressed (center) images. The right hand panel shows the differences between the original and compressed pixel values. Notice the change in colors and the location of feature boundaries.

As unsatisfactory as jpeg-compressed images are for saving grey scale or color images, they are even worse when used for black-and-white (thresholded binary) images. The process of compression only approximates the actual brightness values, and furthermore introduces pixel-to-pixel changes along lines and feature boundaries that produce very significant distortions of the original. And for black and white images, jpeg compression does not even produce the reduction in file size that is the usual justification for its use. The example image shown in Figure 3 consists of a circular region and lines of varying width. Saved in a lossless format such as LZW-compressed TIF, the 400x400 pixel image occupies 72 kB. Saved as a jpeg with the maximum quality setting (12) in Photoshop, the image file is 104 kB, and even with a much lower setting (3) the file size is 84 kB.




Figure 3. Example binary image.

The original image consists of pixels that are either black or white. But when jpeg compression is used, the pixel values are approximated, producing values that are visually close to white or black, so that a simple side-by-side visual comparison of the images does not reveal the differences. This is shown in Figure 4, which compares the histograms for the jpeg images saved with quality settings of 12 and 3. With the highest setting, 2.4% of the pixels are no longer exactly black or white.




Figure 4. Histograms of the image after jpeg compression showing the presence of pixels with values that are not exactly black or white (indicated by red arrows).

The location of these modified pixels is along the edges and boundaries of the features. Since measurement and analysis software treats pixels that are exactly white as background and any darker value as feature, the result is that the shape, dimensions and even number of objects present is remarkably altered, as shown in Figure 5.

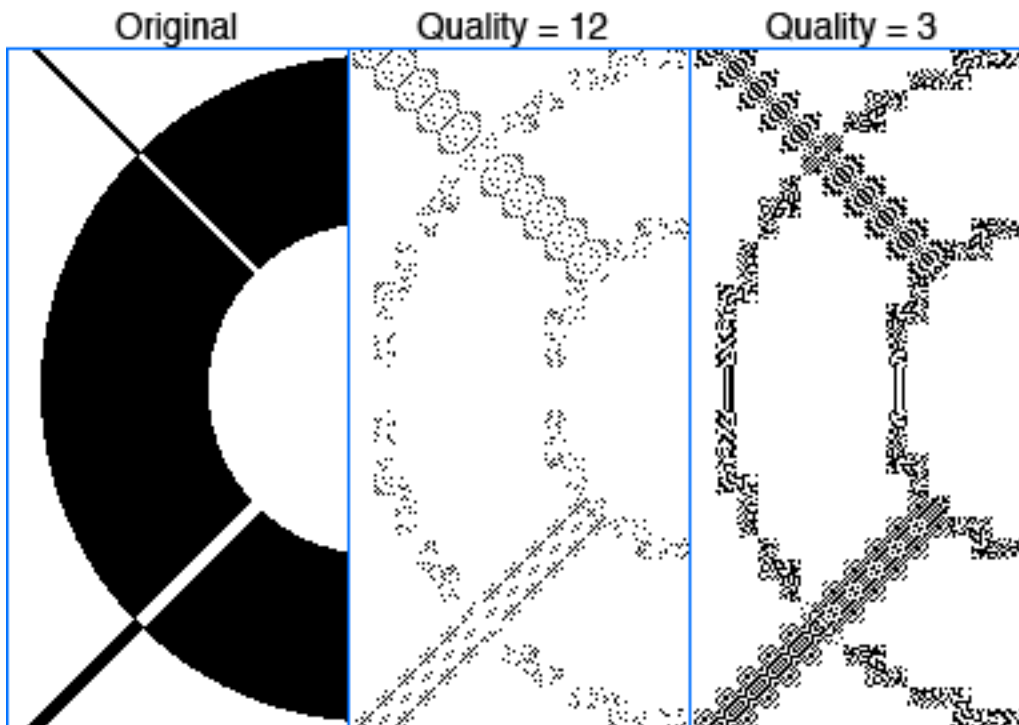


Figure 5. Enlarged detail of a portion of the original image, and the locations of pixels that are no longer black or white after jpeg compression with quality levels 12 and 3.

When higher levels of compression (lower quality levels) are used, the effect is even greater. For a quality level of 3, 30.4% of the pixels in the image are no longer exactly black or white, and as shown in Figure 5 the distortions in the features are greatly increased.

It is also important to understand that the loss of information produced by jpeg compression is unrecoverable, and varies from one part of the image to another. In the simple example shown here, morphological operations such as opening and closing can be applied to the image and will approximately recover the shape of the large circle, but the measurement of its perimeter (and any shape measurements which depend on that) will be significantly altered. The width, integrity and even the presence of the black and white lines will be completely lost. This is illustrated in Figure 6. Notice that the changes in the shape and size of features vary depending on their orientation and position.

JPEG compression is fundamentally incompatible with images that are used for scientific or forensic purposes. With the current availability of relatively inexpensive memory in cameras, large and fast permanent storage using CD or DVD drives, and fast internet connections, there is no reason to even consider using images that have been subject to lossy compression.




Figure 6. Features of varying shape and orientation in a test image: top = original, center = non-white pixels after jpeg compression, bottom = altered features after applying a morphological opening in an attempt to remove the extraneous pixels.