

The impact of skull bone intensity on the quality of compressed CT neuro images

Ilona Kowalik-Urbaniak^a, Edward R. Vrscay^a, Zhou Wang^b, Christine Cavarro-Menard^c, David Koff^d, Bill Wallace^e and Boguslaw Obara^f

^aDept. of Applied Mathematics, University of Waterloo, Waterloo, ON, Canada

^bDept. of Elect. and Comp. Eng, University of Waterloo, Waterloo, ON, Canada

^cLaboratory of the Automated systems Engineering (LISA), University of Angers, France

^dDept. of Radiology, McMaster University, Hamilton, ON, Canada

^eAgfa Healthcare Inc., Waterloo, ON, Canada

^fOxford e-Research Centre and Oxford Centre for Integrative Systems Biology,
University of Oxford, Oxford, UK

ABSTRACT

The increasing use of technologies such as CT and MRI, along with a continuing improvement in their resolution, has contributed to the explosive growth of digital image data being generated. Medical communities around the world have recognized the need for efficient storage, transmission and display of medical images. For example, the Canadian Association of Radiologists (CAR) has recommended compression ratios for various modalities and anatomical regions to be employed by lossy JPEG and JPEG2000 compression in order to preserve diagnostic quality.

Here we investigate the effects of the sharp skull edges present in CT neuro images on JPEG and JPEG2000 lossy compression. We conjecture that this atypical effect is caused by the sharp edges between the skull bone and the background regions as well as between the skull bone and the interior regions. These strong edges create large wavelet coefficients that consume an unnecessarily large number of bits in JPEG2000 compression because of its bitplane coding scheme, and thus result in reduced quality at the interior region, which contains most diagnostic information in the image. To validate the conjecture, we investigate a segmentation based compression algorithm based on simple thresholding and morphological operators. As expected, quality is improved in terms of PSNR as well as the structural similarity (SSIM) image quality measure, and its multiscale (MS-SSIM) and information-weighted (IW-SSIM) versions. This study not only supports our conjecture, but also provides a solution to improve the performance of JPEG and JPEG2000 compression for specific types of CT images.

Keywords: JPEG, JPEG2000, medical images, medical image compression, image compression, image segmentation, SSIM, compression ratio, image quality.

1. INTRODUCTION

Due to the recent advances of medical digital imaging technologies such as computed tomography (CT) and magnetic resonance imaging (MRI), a vast amount of medical image data is being acquired and stored in a computer each day. For this reason, the use of lossy image compression techniques for medical images is inevitable. The Canadian Association of Radiologists (CAR) has adopted lossy compression for medical images.¹ Recommended compression ratios for various modalities and well as anatomical areas have been published by Koff and Shulman² for JPEG and JPEG2000. Their study was based on a subjective quality assessment of compressed images (ROC analysis) which, as is well known, is extremely time consuming (and therefore costly). In this regard, objective assessments would be desirable. Unfortunately there is no generally accepted objective quality assessment method for medical images. As such, the mean squared error (MSE) and its close relative, PSNR, remain the most commonly used objective quality measures, even though they are recognized as being deficient in terms of visual quality.³

In an effort to address the limitations of MSE/PSNR, a number of so-called image quality measures have been developed in the image processing literature. Here, we are primarily concerned with the application of the




Figure 1. Compressed CT neuro image at 12:1 comp. ratio (a) JPEG, (d) JPEG2000. SSIM quality maps: (b) JPEG, (e) JPEG2000, local MSE quality maps (c) JPEG, (f) JPEG2000.

structural similarity (SSIM) measure,⁴ along with two variations, namely, multiscale (MS-SSIM)⁵ and information weighted (IW-SSIM).⁶ These measures provide global and local image quality scores. MSE and SSIM global scores are averages of the local MSE and SSIM scores where all regions are treated with equal importance, i.e. no weighting of background and foreground. MS-SSIM has varying weights to different scales, and IW-SSIM has varying weights across both scale and space.

Global scores, however, might lead to wrong conclusions. For instance, the entire image may pass a validity test although some regions are unacceptable. A better indicator of quality is a local quality map, obtained from local quality scores displayed as an image. Figure 1 (a,b) shows JPEG and JPEG2000 compressed CT neuro images with the corresponding MSE (e,f) and SSIM (b,c) local quality maps (where darker regions represent worse quality). SSIM quality maps indicate that degradations are less uniform inside the skull when compared to the MSE quality maps. Our preliminary investigation showed that regions identified as troublesome by SSIM agreed, to a significant extent, with regions that were independently identified as troublesome by a radiologist. This indicates that the SSIM index and the SSIM quality map provide more promising approach to predict subjective quality assessment of compressed brain CT images. We also observe the presence of compression degradations outside the skull in the JPEG2000 compressed CT neuro image in Figure 1 due to the wavelet nature of the compression algorithm.

2. OBJECTIVE QUALITY MEASURES

L^2 -based measures

In the following discussion, we let f denote an $M \times N$ digital image and g its compressed counterpart. The standard measure of error between f and g is the Mean Squared Error (MSE), defined as follows,

$$\text{MSE}(f, g) = \frac{1}{MN} \sum_{i=1}^M \sum_{j=1}^N (f(i, j) - g(i, j))^2.$$

Image quality is often expressed in terms of the Peak Signal-to-Noise ratio (PSNR), derived from MSE, as follows,

$$\text{PSNR}(f, g) = 10 \log_{10} \left(\frac{R^2}{\text{MSE}(f, g)} \right)$$

Here R is the dynamic range, i.e. $0 \leq f(i, j) \leq R$. (For an N bit-per-pixel image, $R = 2^N - 1$. The Root Mean Square Error (RMSE) is simply the square root of the MSE, i.e.

$$RMSE(f, g) = \sqrt{MSE(f, g)}.$$

From a mathematical point of view, RMSE is more relevant since it behaves as a metric.

SSIM

The **Structural Similarity Index (SSIM)**⁴ was defined on the basis that images are highly structured and there exist strong neighbouring dependencies among the pixels. The human visual system is highly sensitive to structural information/distortions in an image and SSIM automatically adjusts to discount the non-structural ones. The SSIM index measures the difference/similarity between two images by combining three components of the human visual system,

Luminance $l(f, g)$, estimated by the mean:

$$\mu_f = \frac{1}{NM} \sum_{i=1}^N \sum_{j=1}^M f(i, j)$$

Contrast $c(f, g)$, measured by variance:

$$\sigma_f^2 = \frac{1}{(NM - 1)} \sum_{i=1}^N \sum_{j=1}^M (f(i, j) - \mu_f)^2$$

Structure $s(f, g)$, measured by covariance:

$$\sigma_{fg} = \frac{1}{(NM - 1)} \sum_{i=1}^N \sum_{j=1}^M (f(i, j) - \mu_f)(g(i, j) - \mu_g)$$

These three components are multiplied together to form the local SSIM index.

$$\text{(local) } SSIM(f, g) = \left(\frac{2\mu_f\mu_g + C_1}{\mu_f^2 + \mu_g^2 + C_1} \right) \left(\frac{2\sigma_f\sigma_g + C_2}{\sigma_f^2 + \sigma_g^2 + C_2} \right) \left(\frac{\sigma_{fg} + C_3}{\sigma_f\sigma_g + C_3} \right) \text{ over } m \times n \text{ pixel neighborhood.}$$

Theoretically,

$$-1 \leq \text{(local) } SSIM(f, g) \leq 1.$$

The “closer” f and g are to each other, the closer $SSIM(f, g)$ is to the value 1.

The local SSIM index is computed within a sliding window of an $m \times n$ pixel neighbourhood. This results in a “SSIM quality map”, which reveals local image quality. The total SSIM score is computed by averaging the local SSIM values. The Structural Similarity Index (SSIM) measures the distance between two images f, g by combining three components of the human visual system (HVS):

SSIM as a variance weighted L^2 distance

Given that $SSIM(f, g)$ approaches 1 as g approaches f , one might conjecture that $1 - SSIM(f, g)$ is a measure of the error between f and g . This is not quite true. However, in the case of zero-mean images, i.e. if we remove the mean of the images f and g , then a metric function can be defined.⁷

Let $f, g \in R^n$ and define:

$$f_0 = f - \mu_f, g_0 = g - \mu_g \implies \mu_{f_0} = \mu_{g_0} = 0$$

Then

$$DSSIM = \sqrt{1 - SSIM(f_0, g_0)} = \frac{\|f_0 - g_0\|}{\sqrt{\|\sigma_{f_0}\|^2 + \|\sigma_{g_0}\|^2}}. \quad (1)$$

It has been shown that $\sqrt{1 - SSIM(f_0, g_0)}$ is a valid distance metric that satisfies the identity and symmetry axioms as well as triangle inequality.⁷

3. JPEG2000 VERSUS JPEG

On the basis of extensive tests employing natural images, it is generally accepted that JPEG2000 provides better rate-distortion performance than JPEG at higher compression ratios. Koff *et al.*,² however, found that JPEG can outperform JPEG2000 in cases of brain and abdominal CT images. It is natural to inquire why this is so. Perhaps one quick answer could be that CT images, especially brain CT images, are not “natural images” - the presence of bone (in the case of brain images, the skull) produces regions of extremely high image intensity and therefore, very sharp edges. But this analysis can be carried a bit further. Speckle patterns together with features such as irregular, small textures (for example, white matter in a CT brain image) are represented by numerous high frequency coefficients of low amplitude. These features, along with noise, are discarded first during quantization. The bitplane coding scheme in JPEG2000 always gives priority to high energy coefficients. When this is applied to an image with sharp edges (such as CT neuro images), the bits are allocated to the sharp edge regions rather than to the tissue regions with lower energy coefficients resulting in local blurring and ringing. JPEG performs better since the 8×8 block DCT is local - as such, it will not allow excessive bits to be allocated to sharp edges.

In this study, a database of 105 CT neuro images (7 studies with 15 images in a study) of slice thickness 0.1 mm, at various compression ratios) obtained from Medical Informatics Research Centre at McMaster (MIRCAM), Hamilton, Canada, has been analyzed. JPEG2000 outperformed JPEG in terms of L^2 (MSE - mean squared error). In terms of SSIM, however, JPEG is seen to perform better for some compression ratios. Figure 2 shows plots of compression ratios vs. RMSE (root mean squared error) and DSSIM. The plots of compression ratio vs. quality measure change their shape when the image under consideration is cropped to a rectangular region inside the skull. Figure 3 reveals that in this case JPEG2000 always performs better than JPEG and that the RMSE and DSSIM curves (obtained by joining points) have very similar shapes. This is expected since the function DSSIM is an inverse variance-weighted L^2 distance (see Equation 1), and the cropped region of the skull interior has many similar details and thus an approximately constant variance throughout the image. This is an indication that the skull edge affects the compressibility of CT neuro images when JPEG2000 is used to compress the image, resulting in worse performance than JPEG. Our study suggests that the SSIM measure and the SSIM quality map provide the most promising approach to predict subjective quality assessment of compressed brain CT images.

4. SEGMENTATION AND COMPRESSION

In order to further investigate the effects of the sharp skull edges in neuro CT images on JPEG and JPEG2000 lossy compression, a straightforward segmentation has been applied to a CT neuro image. A typical CT brain image has a greyscale distribution concentrated mostly in the lower intensities with a peak at a greyscale intensity of approximately 252 representing the skull bone (Figure 4). A CT neuro image is segmented into three parts (background, skull bone and the inside of the skull) using simple thresholding and morphological operators. From each of these separate pieces, a new image is created by assigning the average value of the extracted mask to the remaining pixels. These three separate images are then compressed by JPEG and JPEG2000 at the same compression ratio. After decompression, these three images are merged back into one image.

We emphasize here that the technique described above should not be used to compress images since it cannot be generalized to a larger set of images. We have employed it in this study in order to illustrate the effects of sharp edges on compression.

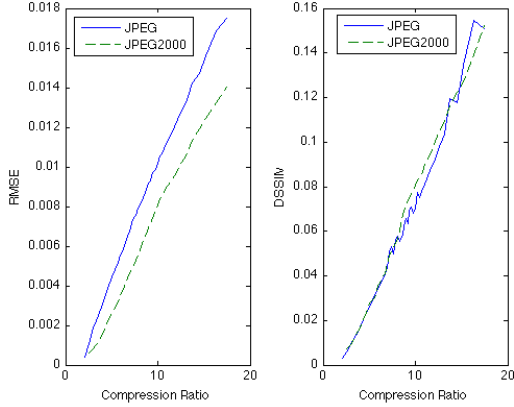


Figure 2. Plots of compression ratio versus RMSE and DSSIM of a CT neuro image

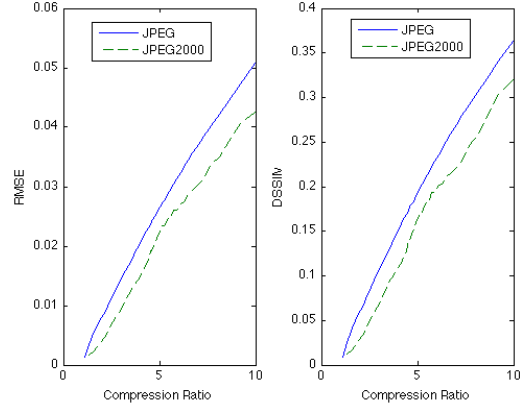


Figure 3. Plots of compression ratio versus RMSE and DSSIM of a cropped CT neuro image




Figure 4. Histogram of a CT neuro image.

CT neuro Image 1	JPEG compression	Segmented JPEG compression	JPEG2000 compression	Segmented JPEG2000 compression
PSNR	32.080	34.528	34.185	36.512
SSIM	0.991	0.996	0.988	0.995
MS-SSIM	0.995	0.997	0.995	0.997
IW-SSIM	0.992	0.996	0.993	0.996
CT neuro Image 2	JPEG compression	Segmented JPEG compression	JPEG2000 compression	Segmented JPEG2000 compression
PSNR	31.572	33.918	33.917	36.883
SSIM	0.989	0.995	0.986	0.995
MS-SSIM	0.994	0.997	0.994	0.997
IW-SSIM	0.991	0.995	0.993	0.996

Table 1. Quality scores using PSNR, SSIM, MS-SSIM and IW-SSIM for two JPEG and JPEG2000 compressed CT neuro images

The objective quality measures that were used include: PSNR, SSIM, MS-SSIM and IW-SSIM. As expected, the quality is improved according to all objective quality measures used (see Table 1). The new method produces less artifacts according to the SSIM local quality map for all the 105 images tested. Figures 5, 6, 7 and 8 show the SSIM and MSE quality maps for two compressed CT images using JPEG and JPEG2000 compression with and without the use of segmentation. These objective quality tests support the hypothesis that pre-compression segmentation of CT neuro images improves the quality of JPEG and JPEG2000 compressed images.

5. 16-BIT VERSUS 8-BIT JPEG2000 COMPRESSION

Baseline lossy JPEG allows only 8- and 12-bit greyscale compression, whereas JPEG2000 allows up to 16-bit of greyscale compression. In this work we have also investigated JPEG2000 compression of 16 bpp (bits per pixel) and 8 bpp CT neuro images. Most viewable images are 8 bits and displays very seldom support a raster depth more than 8 bits (per colour). The typical “12 bit” monitor typically means that the lookup table to drive the monitor is 12 bits, not that the incoming raster depth is 12 bits. For any system which compresses the viewable image, the compression is performed on the viewable image, not necessarily on the source image. This is required




Figure 5. Quality maps of a JPEG compressed CT neuro image, compression ratio: 12:1 (a) compressed image (no segmentation), (b) SSIM quality map, $SSIM = 0.9909$, no segmentation (c) MSE quality map, $PSNR = 32.1$, no segmentation, (d) compressed image (segmentation) (e) SSIM quality map, $SSIM = 0.9959$, segmentation, (f) MSE quality map, $PSNR = 34.5$, segmentation.




Figure 6. Quality maps of a JPEG2000 compressed CT neuro image, compression ratio: 12:1 (a) compressed image (no segmentation), (b) SSIM quality map, $SSIM = 0.9898$, no segmentation (c) MSE quality map, $PSNR = 34.2$, no segmentation, (d) compressed image (segmentation) (e) SSIM quality map, $SSIM = 0.9949$, segmentation, (f) MSE quality map, $PSNR = 36.5$, segmentation.




Figure 7. Quality maps of a JPEG compressed CT neuro image, compression ratio: 12:1 (a) compressed image (no segmentation), (b) SSIM quality map, $SSIM = 0.9887$, no segmentation (c) MSE quality map, $PSNR = 31.6$, no segmentation, (d) compressed image (segmentation) (e) SSIM quality map, $SSIM = 0.9952$, segmentation, (f) MSE quality map, $PSNR = 33.9$, segmentation.




Figure 8. Quality maps of a JPEG2000 compressed CT neuro image, compression ratio: 12:1 (a) compressed image (no segmentation), (b) SSIM quality map, $SSIM = 0.9863$, no segmentation (c) MSE quality map, $PSNR = 33.9$, no segmentation, (d) compressed image (segmentation) (e) SSIM quality map, $SSIM = 0.9945$, segmentation, (f) MSE quality map, $PSNR = 36.9$, segmentation.

in order to use the available imaging libraries which do not typically support 12/16 bpp images (none of the 5 major browsers support more than 8 bpp.)

CT neuro images are usually lossy compressed after their bit-depth has been reduced to 8 bits by means of window levelling. (The window levelling is generally accomplished by using a piecewise linear function). In this case, the compression is 8 bpp, and it is applied to an already altered image. Another option is to apply lossy JPEG2000 compression on the original 16 bpp image, followed by window levelling. There are advantages and disadvantages to both approaches. Figure 9 shows SSIM and MSE quality maps of a compressed CT image using 8- and 16-bit JPEG2000 compressions. Quality scores were computed for each of the compressed images and are presented in Table 2. The result agrees with intuition and we conclude that 16-bit compression does provide better quality than 8-bit compression when the same compression ratio is used. However, 16 bpp compressed images take up more storage space than 8 bpp compressed images! For the image presented in Figure 9, the 8-bit stored image uses 25KB, whereas the 16 bpp compressed version of this image has a size of 45KB. The advantage of using 16-bit compression is that any window levelling (i.e. bone) can be still obtained from the compressed image. With 8 bpp compressed images, that option is no longer available. Furthermore, it is not clear how to compare 8-bit and 16-bit compressions.

The comparison of 8- and 16-bit compression could be carried in several ways. One option is to use the same compression ratio (as performed in our study). Another option is to match one of the quality measures (PSNR, SSIM, etc.). For example, in order to obtain the same image size (25KB) after compression, we would have different compression ratios for 8- and 16-bit compressions (12:1 for 8-bit compression, and 20:1 for 16-bit compression). However, this is not a desirable result as 20:1 compression ratio produces lower quality scores (Table 3) and to the best of our knowledge there were no radiological studies on recommended compression ratios for 16-bit compression. Furthermore, Figure 9 reveals that there are more “edge artefacts” in the 8-bit compressed image than in the 16-bit compressed JPEG2000 image. Thus, we draw the conclusion that the skull edge has an effect on the compressibility of CT neuro images.




Figure 9. Quality maps of a JPEG2000 compressed CT neuro image, compression ratio: 12:1 (a) CT neuro image after 16-bit compression, (b) SSIM quality map (16-bit compression), (c) MSE quality map (16-bit compression), (d) CT neuro image after 8-bit compression, (e) SSIM quality map (8-bit compression), (f) MSE quality map (8-bit compression)

Quality Measure	JPEG2000 16-bit	JPEG2000 8-bit
PSNR	37.124	33.917
SSIM	0.995	0.986
MS-SSIM	0.998	0.993
IW-SSIM	0.997	0.993

Table 2. Quality scores of a JPEG2000 compressed CT neuro image (8- and 16-bit compressions), compression ratio: 12:1

Quality Measure	JPEG2000 16-bit	JPEG2000 8-bit
PSNR	31.339	33.917
SSIM	0.978	0.986
MS-SSIM	0.989	0.993
IW-SSIM	0.987	0.993

Table 3. Quality scores of a JPEG2000 compressed CT neuro image (8-bit with compression ratio 12:1 and 16-bit with compression ratio: 20:1)

CONCLUSION

In summary, this paper answers the question why JPEG performs better than JPEG2000 on some CT images. We conclude that compressing a CT neuro image requires a special treatment due to a specific distribution of its greyscale intensities (Figure 4). We have employed several image quality measures including PSNR, SSIM, MS-SSIM and IW-SSIM. Using the SSIM and its variations, we were able to objectively confirm the subjective radiological results that JPEG sometimes performs better than JPEG2000 in the case of CT neuro images. This paper also investigates the result of 8- and 16-bit compression of CT neuro images. It has been observed that for CT brain images, 8-bit compression results in more edge artefacts compared to 16-bit compression, as measured by the SSIM index.

The effect that sharp edge has on compression may be encountered not only in CT neuro images but also in other images that contain both extremely strong edges and diagnostically important textures with moderate energy. Although the segmentation technique used in this study is not recommended to use in practice, we are working on developing better ways to deal with the troublesome sharp edges in order to improve the performance of both JPEG and JPEG2000 algorithms.

REFERENCES

- [1] Koff, D., Bak, P., Brownrig, P., Hosseinzadeh, D., Khademi, A., Kiss, A., Lepanto, L., Michalak, T., Shulman, H., and Volkening, A., “Pan-Canadian evaluation of irreversible compression ratios (“Lossy Compression”) for development of national guidelines,” *Journal of Digital Imaging* **22**(6), 569–578 (2009).
- [2] Koff, D. and Shulman, H., “An overview of digital compression of medical images: can we use lossy image compression in radiology?,” *Canadian Association of Radiologists Journal* **57**(4), 211–217 (2006).
- [3] Wang, Z. and Bovik, C., “Mean squared error: love it or leave it? - A new look at signal fidelity measures,” *IEEE Transactions on Signal Processing* **26**(1), 98 – 117 (2009).
- [4] Wang, Z., Bovik, A., Sheikh, H., and Simoncelli, E., “Image quality assessment: from error visibility to structural similarity,” *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004).
- [5] Wang, Z. and Bovik, A., “Multi-scale structural similarity for image quality assessment,” in [*IEEE Asilomar Conference on Signals, Systems and Computers*], **2**, 1398–1402 (9-12 November 2003).
- [6] Wang, Z. and Li, Q., “Information content weighting for perceptual image quality assessment,” *IEEE Transactions on Image Processing* **20**(5), 1185–1198 (2011).
- [7] Brunet, D., Vrscay, E. R., and Wang, Z., “A class of image metrics based on the structural similarity quality index,” in [*International Conference on Image Analysis and Recognition (ICIAR ’11)*], (22-24 June 2011).