The quest for "diagnostically lossless" medical image compression using objective image quality measures

by

Ilona Anna Kowalik-Urbaniak

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Applied Mathematics

Waterloo, Ontario, Canada, 2014

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Given the explosive growth of digital image data being generated, medical communities worldwide have recognized the need for increasingly efficient methods of storage, display and transmission of medical images. There is also a general acknowledgement that lossless compression techniques, with low compression ratios, are no longer adequate and that it is necessary to consider higher compression rates. Since the result, lossy compression, involves a loss of information and possibly visual quality, it is absolutely essential to be able to determine the degree to which a medical image can be compressed before its "diagnostic quality" is compromised. The main goal of this work was to achieve "diagnostically lossless compression", i.e., compression with no loss in visual quality nor the diagnostic accuracy. The degradation produced by "diagnostically lossless compression" is sufficiently minor and will not affect radiological diagnosis.

The quality of a compressed image can be characterized objectively in several ways. Radiologists most often employ the Mean Squared Error (MSE) and its close relative, Peak Signal-to-Noise Ratio (PSNR), even though they are known to correspond poorly to visual quality. The failure of MSE/PSNR is partially due to the fact that spatial relationships are ignored by the L^2 metric on which they are based. A more recent image fidelity measure, the SSIM index, measures the difference/similarity between two images by combining three components of the human visual system – luminance, contrast and structure estimated using the mean, variance and covariance, respectively. The result is a much improved assessment of visual quality.

Recent research by Koff et al. has shown that at higher compression levels lossy JPEG is more effective than JPEG2000 in some cases of brain and abdominal CT images. This is an interesting and somewhat surprising observation as JPEG2000 is generally believed to provide better rate-distortion performance than JPEG. We have investigated the effects of the sharp skull edges in CT neuro images on JPEG and JPEG2000 lossy compression. A CT neuro image is segmented into three parts (background, skull bone and the interior region 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 images are then separately compressed using JPEG and JPEG2000 algorithms. As expected, the quality is improved according to MSE, SSIM and its variations MS-SSIM and IW-SSIM. We provide an explanation why JPEG outperforms JPEG2000 for some CT images.

Another aspect of this study, involving a collaboration with radiologists as well as a leading international developer of medical imaging software (AGFA), is primarily concerned

with improved methods of assessing the diagnostic quality of compressed medical images. The radiological community has not yet accepted a single objective assessment method for the quality of medical images. Recommended compression ratios for various modalities and anatomical regions have been published. To date, these recommendations have been based on experiments in which radiologists subjectively assess the diagnostic quality of compressed images. In this work, we have examined whether compression ratio and MSE actually serve as reliable indicators of diagnostic quality. By this we mean to model the perception of trained radiologists in a satisfactory way. We have also investigated the quality factor, the sole input parameter in the JPEG compression algorithm, since it has also been employed as a reference for quality assessment. The performances of the above indicators are compared to that of SSIM index, based on the data collected in subjective experiments involving radiologists. An ROC and Kolmogorov-Smirnov analyses indicate that compression ratio is not always a good indicator of visual quality. Moreover, SSIM demonstrates the best performance, i.e., it provides the closest match to the radiologists' assessments. We also show that a weighted Youden index can provide SSIM and MSE thresholds for acceptable compression.

We also examine two approaches of modifying L^2-based approximations so that they conform to Weber's model of perception, i.e., higher/lower tolerance of deviation for higher/lower intensity levels. The first approach involves the idea of intensity-weighted L^2 distances. We arrive at a natural weighting function that is shown to conform to Weber's model. The resulting "Weberized L^2 distance" involves a ratio of functions. The importance of ratios in such distance functions leads to a consideration of the well-known logarithmic L^2 distance which is also shown to conform to Weber's model. In fact, we show that the imposition of a condition of perceptual invariance in greyscale space according to Weber's model leads to the unique (unnormalized) measure with density function $\rho(t) = 1/t$. This result implies that the logarithmic L^1 distance is the most natural "Weberized" image metric. From this result, all other logarithmic L^p distances may be viewed as generalizations. We provide numerical implementations of the intensity-weighted approximation methods for natural and medical images.

Acknowledgements

I would like to express my gratitude to my supervisors, Prof. Edward Vrscay and Prof. Zhou Wang, whose understanding and patience helped me considerably to get through my graduate experience. I appreciate their vast knowledge and expertise in many areas. I am particularly grateful to Prof. Vrscay, for inspiration, motivation and all the help he has provided during my time at the University of Waterloo. Prof. Vrscay has consistently demonstrated the model qualities of a graduate student supervisor.

I also would like to express my thanks to Drs. Nadine and David Koff who have provided their valuable knowledge, time and helped with all the subjective experiments involved in this research. I also would like to thank the other members of my committee, Dr. Jeff Orchard, Dr. Justin Wan for their assistance on the research project.

Finally, I would like to thank Prof. Dietmar Saupe for taking time out of his busy schedule to serve as my external examiner and for his most helpful comments and suggestions which improved the quality of my thesis.

I would also like to extend my thanks to Jiheng Wang for creating the image viewer for the subjective experiments. Also, thanks to Faerlin Pulido, who has written an image viewer for the preliminary set of experiments.

A very special thanks goes out to Jane Castelli, who has organized and helped run the subjective experiments with radiologists. I wish to acknowledge the help provided by Bill Wallace, who has been involved in this project from the beginning and has provided valuable and greatly appreciated assistance during our regular research meetings at AGFA Healthcare Inc.

I am particularly grateful for the assistance given by Prof. Paul Marriott on the design of the subjective experiments.

I would also like to thank my family for the support they provided me through my entire life and in particular, I must acknowledge my husband and best friend, Sebastian for his love.

In conclusion, I gratefully recognize the financial assistance of this research provided by the Department of Applied Mathematics at the University of Waterloo (Teaching Assistantships, Graduate Research Scholarships), NSERC/MITACS, AGFA Healthcare Inc. and the University of Waterloo Graduate Studies.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Prelude

An image is a representation of the likeness of some subject. Images are generated by recording the radiation from a source of energy passing through a medium. Medical images are created by the combination of illumination and reflection or absorption of energy from a source by elements of the human body. Some examples of medical images, based on the electromagnetic energy spectrum (Figure 1.1), include X-ray, magnetic resonance (MR) and Gamma-ray images. The various types of medical images are generated using medical modalities having their own specific features corresponding to the physiological phenomena studied [57].

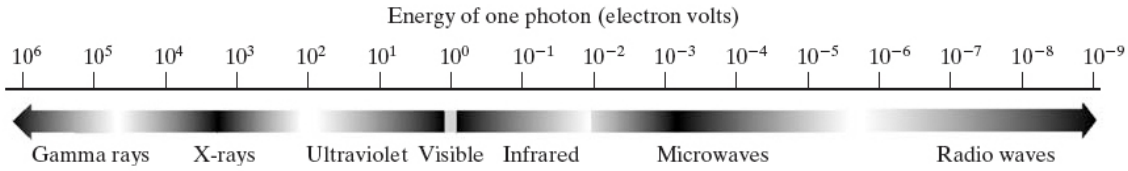

Figure 1.1: The electromagnetic spectrum arranged according to energy per photon. Figure adopted from [31].

Mathematically, a greyscale image can be represented by a real-valued function, $f(x, y)$, of two continuous variables x and y. The value of f at any pair of coordinates (x, y) is called the intensity or grey level of the image at that point. Most medical images are greyscale images. To create a digital image, we need to convert the image function of the

continuous spatial variables into finite, discrete quantities by sampling and quantization. A digital image is represented as a two-dimensional array.

Digital medical images are usually greyscale with 16 bits per pixel (i.e. 65536 greyscale levels). However, the human eye is unable to distinguish that many greyscale levels. It is generally known that the human eye is able to distinguish no more than 10-bits of grey (1024 grey shades). The grey level reduction for medical images is accomplished through the means of window levelling filter. This filter operates on 16-bit images and reduces their greyscale range to 256 shades of grey (8-bit). This operation automatically enhances the contrast. Regular monitors are capable of displaying 256 (8-bit) grey shades whereas medical displays are now capable of displaying 4096 (12-bits) [37] and the latest technology allows for 14-bit display. Although specialized medical displays are able to display more than 8-bit greyscale images, the display is not a "true" 10-bit (or more) display since the window levelling operation transforms the original (raw) 16-bit image into an 8-bit image before it is displayed. It should be noted that displaying images in browsers is limited to 8-bit.

In general, image arrays exhibit coding, spatial and perceptual data redundancies. Image compression exploits these redundancies in order to represent the image using the smallest number of bits. Compression can be defined as "the art and science of representing information in a compact form" [75]. Image compression can be lossless or lossy. In lossless, or reversible compression, the reconstructed image is identical to the original image. In lossy compression, however, which is an irreversible process, we encounter loss of image information/quality. In general, lossy compression techniques achieve much greater compression ratios than lossless algorithms. The most common compression formats used in medical imaging are JPEG and JPEG2000, developed by the Joint Photographic Experts Group.

Lossy compression alters the quality of the original image. Subjective and objective quality assessment methods are used to evaluate the degradations that result from compression. One of the subjective methods used in medical image compression studies is the Receiver Operating Characteristics (ROC) analysis, where a group of experts rates the quality of images. Objective methods, where no human intervention is required, are based on signal, texture analysis and models of the Human Visual System (HVS). The most widely used objective quality measure for images is the Mean Squared Error (MSE). Traditionally, much of the image processing literature has relied on MSE or other "L^2-based" methods. The advantage of these methods is that they are easy to compute. The disadvantage is that they do not necessarily reflect visual quality. It is easy to modify an image to produce another image which is "close" in L^2-distance, yet visually "far" (Figure 2.6). In recent years, there has been some progress in the development of image quality

measures that better reflect visual similarities/dissimilarities. An example is the structural similarity (SSIM) index [88].

1.2 Motivation

A patient's medical record consists of clinical data, images and other physiological signals. Medical information is stored and/or transmitted between healthcare centers. Although, thousands of medical records are still recorded using a radiological film, the vast majority of data acquired are digital. With the rapid advances of medical digital imaging technologies such as computed tomography (CT) and magnetic resonance imaging (MRI) scanners, the amount of image data being acquired each day has been increasing exponentially [57], [62]. Unfortunately, the decreasing cost of computer storage is being largely surpassed by the increasing volume.

CT images take nearly 50% of a Picture Archiving and Communication System (PACS) storage space with CT studies reaching 5000 images [38]. A typical neuro CT image (512×512) pixels requires about 0.5MB of storage space if stored is a raw pixel format. The storage needs are predicted to further increase with 3D volume rendering [57]. It is obvious that physicians, regardless of their location, would like to be able to access medical images efficiently and quickly. Access to imaging data is becoming more widespread with Electronic Health Record (EHR) solutions. Canada Health Infoway is hosting the development of imaging repositories that host data online for legal retention allowing physicians to access medical information dating back many years, in nearly real-time [61]. In order to achieve this goal, new and/or improved technology is needed for efficient storage, transmission between sites, retrieval and display of medical image information in large scale databases. For this reason the use of lossy image compression techniques is inevitable.

Speed limitations of existing networks along with the explosive growth of image modalities with extremely high volume outputs have combined to make the issue of irreversible medical image coding one of the key considerations in the design of future PACS systems [82]. Existing lossy image compression techniques are well suited for images where the only concern is visual quality. Medical images are special; each imaging modality produces images that differ in spatial resolution, contrast and type of noise. Therefore, improved or specially designed compression algorithms are desired.

As expected, increasing the degree of compression of an image leads to decreasing fidelity. The extent of allowable irreversible compression is dependent on the imaging modality and the nature of the image pathology and anatomy. Image compression often

results in distortions of the images, and therefore creates the risk of losing or altering relevant diagnostic information.

The concern that a misdiagnosis could occur as a result of relevant information being lost due to lossy compression has encouraged researchers to systematically review existing research on lossy compression. In order to facilitate the adoption of lossy compression for medical images researchers have investigated the use of JPEG and JPEG2000 compression methods for various modalities and anatomical regions. As a result, recommendations for using lossy compression for medical images have been adapted and published by several radiological societies [63]:

- Royal College of Radiologists (RCR, UK) [62]
- Canadian Association of Radiologists (CAR) [61]
- German Röentgen Society (DRG, Germany) [47]
- The Royal Australian and New Zealand College of Radiologists (RANZCR) [5].

The recommendations all suggest that the use of irreversible image compression is possible without the loss of relevant clinical features required for diagnosis. Legal issues are based on national law.

The guide prepared by RCR in UK outlines the current regulations and legal implications of using lossy compression for medical images: "The main regulatory bodies in the UK, EU, USA, Canada and Australia neither prevent nor endorse lossy compression for medical images. The risk of using lossy compression is no more than in conventional practice provided that the diagnostic quality of the image is not reduced and that the decision to use lossy compression was made at an institutional level as a matter of public policy and resource allocation and there was support from at least a respectable minority of radiologists."

In Germany, irreversible compression is also allowed if "there is no loss in medical information" [63], [47].

In Canada there is a governmental supported/driven process in evaluating and implementing irreversible compression for the national Canada Health Infoway Project.

Two independent literature reviews, sponsored by Canada Health Infoway, were done. These studies reviewed over 120 journal articles and were performed by INSITE Consultancy Inc. [36], and Finnis [25].

From these two independent studies, the following conclusions were drawn: [61]:

- Irreversible compression is a clinically acceptable option for compresing medical images
- The extent of allowable irreversible compression depends on the modality of the image and the nature of the pathology and body part.

Furthermore, two unrelated comprehensive legal reviews including a regulatory assessment, sponsored by Canada Health Infoway, were conducted with regards to the use of irreversible compression for medical images in Canada [11] [52]. From both studies it was concluded that the appropriate use of lossy compression does not increase the legal liability of medical practitioners [61].

One of the main studies considered by CAR, based on which the recommended compression ratios were published, is the pan- Canadian clinical evaluation on JPEG and JPEG2000 irreversible compression in medical imaging by Koff *et al.* [38]. The use of JPEG and JPEG2000 compression methods was investigated for various modalities and anatomical regions. The study was broad and involved 5 modalities (CR/DR, CT, US, MR, NM), 7 body parts, 100 radiologists from across Canada and up to 80 images per session. As a result, they published, for the first time, a standard for the use of compression in medical imaging, supported by CAR. In their study, they found limitations linked to the properties of the compression algorithms. For example, JPEG, based on the block discrete cosine transform, often leads to blurring within blocks and blockiness at block boundaries. JPEG2000, based on the wavelet transform, often produces blurring in image detail structures and ringing around edges.

In Australia, "RANZCR recommends the use of lossless compression where possible. However, where necessitated by infrastructure or cost barriers, conservative levels of data compression may be used, as per the RCR guidelines; if more aggressive lossy compression is used, compression levels should generally be within the guidelines adopted by CAR" [5].

A summary of compression ratio recommendations based on the studies reported in CAR, DRG and RCR is presented in Table 1.2.

	UK	Canada JPEG	Canada J2K	Germany
Radiography - chest	10	30	30	10
Radiography - skeletal	10	30	20	10
Radiography - body	10	30	30	10
Radiography - pediatric		30	30	
Radiography - mammo	20	25	25	15
CT - head	5	12	8	5
CT – skeleton/chest/lung	5	15	15	8
CT - body		15	10	10
CT - angio		15	15	
CT - pediatric		15	15	
MR	5	24	24	7
NM		11	11	
US	10	12	12	
XA	10			6
XRF				6

Figure 1.2: Recommended compression ratios for medical images in UK, Canada and Germany. Table adopted from [63].

Needless to say, if not implemented properly, distortions resulting from lossy compression could impede the ability of radiologists to make confident diagnoses from compressed medical images. However, defining the amount of accepted distortion is a complex task. For this reason, reliable image quality assessment methods are needed in order to achieve what we shall call diagnostically lossless compression, defined as follows:

> **Diagnostically lossless compression:** Irreversible compression of a medical image with little or no loss in visual quality and no effect on diagnostic accuracy. The degradations produced by such compression are sufficiently minor so as not to affect the diagnostic assessments of radiologists.

The above definition was written by the author, however, similar definitions exist in

the relevant literature. For example, "diagnostically acceptable irreversible compression (DAIC)" was introduced in [63] and refers to an irreversible compression that has no effect on diagnostic task.

1.3 Objective

The main goal of this thesis is to achieve diagnostically lossless compression for medical images with a quantitative evaluation of compression-caused degradations and through a subjective experiment with radiologists. Furthermore, we work toward developing intensity-weighted image approximation techniques that incorporate a Weber-type model of perception.

1.4 Statement of Contributions

Below is a summary of the main contributions of this thesis:

- A quantitative analysis of compression degradations is presented. We provide an explanation why JPEG performs better than JPEG2000 for certain types of CT images. According to Koff *et al.* [38], [39], at higher compression levels lossy JPEG is more effective than JPEG2000 in some cases of brain and abdominal CT images. We have investigated how JPEG and JPEG2000 lossy compressions are affected by the sharp skull edges in the case of CT neuro images. Our technique involves segmentation of a brain CT image into background, skull bone and the interior region of the skull, using thresholding and morphological operators. The three image regions are then compressed separately using JPEG and JPEG2000 algorithms. This technique improves image quality as measured by MSE, SSIM and its variations MS-SSIM and IW-SSIM.
- Another aspect of this study, involving a collaboration with radiologists as well as a leading international developer of medical imaging software (AGFA), is primarily concerned with improved methods of assessing the diagnostic quality of compressed medical images. In this study, we have compared the performances of the structural similarity quality measure (SSIM), MSE/PSNR, compression ratio and JPEG quality factor, based on the data collected in a subjective experiment involving radiologists. An ROC and Kolmogorov-Smirnov analyses indicate that compression ratio is not

always a good indicator of visual quality. Moreover, SSIM demonstrates the best performance, i.e., it provides the closest match to the radiologists' assessments. We have also shown that a weighted Youden index can provide SSIM and MSE thresholds for acceptable compression.

- We have proposed two approaches of modifying L^2-based approximations incorporating Weber's model of perception. The first method involves the idea of intensity-weighted L^2 distances. We arrive at a natural weighting function that conforms to Weber's model. The resulting "Weberized L^2 distance" involves a ratio of functions, which leads to a consideration of the well-known logarithmic L^2 distance. The logarithmic L^2 distance also conforms to Weber's model. Furthermore, we have shown that the imposition of a condition of perceptual invariance in greyscale space according to Weber's model leads to the unique (unnormalized) measure with density function $\rho(t) = 1/t$. This result implies that the logarithmic L^1 distance is a natural "Weberized" image metric. As a consequence, all other logarithmic L^p distances may be viewed as generalizations. We provide numerical implementations of the intensity-weighted approximation methods for natural and medical images.

1.5 Thesis organization

The next chapter of this thesis, Chapter 2, covers background information and literature review. It includes a discussion on the creation of medical images, DICOM format, JPEG and JPEG2000 compression algorithms and image quality assessment measures including MSE/PSNR and SSIM index. In Chapter 3, we present an investigation of the effects of irreversible compression on brain CT and body CT images. A quantitative analysis including a discussion on the effects of medical image compression on visual/diagnostic quality is presented. The quest for finding features in medical images that would serve as an indicator of compressibility by employing local analysis of compression artifacts is also one of the topics of Chapter 3. Validation of objective quality assessment models for compressed medical images including results of the main subjective experiment are provided in Chapter 4. Chapter 5 covers a mathematical treatment of intensity weighted approximations of images: "Weberized" L^2 and Logarithmic L^2 methods. We provide numerical implementations of the Weber-model based approximation methods with applications to natural and medical images in Chapter 6. Finally, Chapter 7 contains concluding remarks and possible future directions based on the results provided in this thesis.

Chapter 2

Background and Related work in the Literature

2.1 Image Compression

2.1.1 The Projection Theorem and Frequency Domain Transforms

The Projection Theorem is a well known result that makes transformations from time or space to frequency domain possible.

Theorem 2.1.1. *Projection Theorem:*
Let H be a Hilbert space with dimension M, M could be infinite. Let $\{e_k\}_{k=1}^{M}$ be a complete orthonormal basis of H, i.e., $\langle e_i, e_j \rangle = 1$ if $i = j$ and 0 otherwise. Let us define the subspaces $S_N \subset H$, where $1 \leq N \leq M$. Then the best approximation of f in S_N, to be denoted as $f_N \in S_N$ is given by

$$f \approx f_N = \sum_{k=1}^{N} c_k e_k \tag{2.1}$$

where $c_k = \langle f, e_k \rangle, 1 \leq k \leq N$. c_k are the Fourier coefficients of f.

The approximation f_N is the element in S_N closest to f with respect to the distance/-metric defined by the inner product in H.

Images are finite dimensional, discrete signals. For this reason, it is sufficient to consider N to be finite. Signals or images may thus be represented by N-vectors, which are to be considered as elements of the space C^N of complex N-vectors. Each such N-vector f is indexed as follows,

$$f = (f[0], f[1], ..., f[N-1])). \tag{2.2}$$

The set of functions

$$e_k[n] = \frac{1}{\sqrt{N}} \exp\Big(\frac{i2\pi kn}{N}\Big), n = 0, 1, ..., N-1; k = 0, 1, ..., N-1. \tag{2.3}$$

forms an orthonormal basis on C^N.

Using the basis $\{e_k\}_{k=1}^N$, defined in 2.3, we now present the discrete version of the Fourier transform and its inverse.

Definition 2.1.2. *Discrete Fourier transform (DFT):*

$$c_k = \langle f, e_k \rangle = \sum_{n=0}^{N-1} f[n]\overline{e_k[n]} = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} f[n] \exp\Big(\frac{-i2\pi kn}{N}\Big), \quad k = 0, 1, \cdots, N-1.$$

Definition 2.1.3. *Inverse discrete Fourier transform (IDFT):*

$$f[n] = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} c_k \exp\Big(\frac{i2\pi kn}{N}\Big), \quad n = 0, 1, \cdots, N-1\,.$$

The original vector f can be reconstructed from the DFT coefficients c_k.

Fourier series representation defines the 2π-periodic extension of $f \in L^2[-\pi, \pi]$. This can cause problems at $x = \pm\pi$ if $f(-\pi) \neq f(\pi)$ which is known as the Gibbs phenomenon. One way to avoid the Gibbs phenomenon at the endpoints is to define f on $[0, \pi]$ and construct an even extension $f(x) = f(-x)$ on $[-\pi, 0]$. The 2π-periodic extension of f is continuous at the endpoints. This produces the so-called Fourier cosine series expansion of f, i.e., the orthonormal basis consists entirely of cosines. As a result, ringing artifacts at the endpoints can be avoided. In the discrete case this idea leads to the special case of the Fourier transform and it is known as the Discrete Cosine Transform (DCT). Due to excellent energy compaction of the coefficients, the DCT is used in the JPEG compression algorithm. The discrete cosine transform (1-D and 2-D) and its inverse are given by:

Definition 2.1.4. *Discrete cosine transform (DCT 1-D)*

$$F[k] = \lambda_k \sum_{n=0}^{N-1} f[n] \cos\Big(\frac{\pi}{N}\Big(n + \frac{1}{2}\Big)k\Big), k = 0, 1, .., N-1 \tag{2.4}$$

Definition 2.1.5. *Inverse discrete cosine transform (IDCT 1-D)*

$$f[n] = \sqrt{\frac{2}{N}} \sum_{k=0}^{N-1} F[k]\lambda_k \cos\Big(\frac{\pi}{N}\Big(n + \frac{1}{2}\Big)k\Big), \lambda_k = \begin{cases} \frac{1}{\sqrt{2}} & \text{if } k = 0 \\ 1 & \text{otherwise.} \end{cases} \tag{2.5}$$

The DCT easily translates into two dimensions:

Definition 2.1.6. *DCT 2-D*

$$F[u, v] = \frac{1}{\sqrt{MN}} C[u]C[v] \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} f[x, y] \cos\Big(\frac{(2x+1)u\pi}{2N}\Big) \cos\Big(\frac{(2y+1)v\pi}{2M}\Big) \tag{2.6}$$

Definition 2.1.7. *IDCT 2-D*

$$f[x, y] = \frac{1}{\sqrt{MN}} \sum_{u=0}^{N-1} \sum_{v=0}^{M-1} C[u]C[v]F[u, v] \cos(\frac{(2x+1)u\pi}{2N}) \cos\Big(\frac{(2y+1)v\pi}{2M}\Big) \tag{2.7}$$

$$u = 0, 1, ..., N-1, v = 0, 1, ..., M-1, x = 0, 1, ..., N-1, y = 0, 1, ..., M-1, \tag{2.8}$$

$$C(z) = \begin{cases} \frac{1}{\sqrt{2}} & \text{if } z = 0 \\ 1 & \text{otherwise.} \end{cases} \tag{2.9}$$

2.1.2 Wavelet Transform and Multiresolution Analysis

Let $L^2(\mathbb{R})$ denote the Hilbert space of square integrable functions. We now provide a short summary of multi-resolution analysis.

We begin with the properties of multi-resolution analysis [85, 19], which is defined on the space $L^2(\mathbb{R})$ and consists of a sequence of closed subspaces V_j, $j \in \mathbb{Z}$, such that,

- V_j, $j \in \mathbb{Z}$ are nested, i.e. $\{0\} \subset \cdots \subset V_j \subset V_{j+1} \cdots \subset L^2(\mathbb{R})$
- V_j's are scaled versions of each other with the scaling dilation factor 2^{j-k} for $j > k$, i.e. $v(x) \in V_k \iff v(2^{j-k}x) \in V_j$. (Self-similarity in space)

- V_j's are invariant under shifts by $p \cdot 2^k$, $p \in \mathbb{Z}$, i.e. $v(x) \in V_j \iff v(x + p \cdot 2^k) \in V_j$. (Self-similarity in time)
- V_j's fill the whole space $L^2(\mathbb{R})$, i.e. The set $\cup_{j=-\infty}^{\infty} V_j$ is dense in $L^2(\mathbb{R})$. (Completeness).
- V_j's are not redundant, i.e. $\cap_{j=-\infty}^{\infty} V_j = \{0\}$.
- $\exists$ a generating function $\phi \in V_0$ of span $\{\phi(x-k)\}, k \in \mathbb{Z}$ that forms a Riesz basis (i.e. there exists $A, B > 0$ such that $A \leq \sum_{k \in Z} |\hat{\phi}(\omega + 2\pi k)|^2 \leq B$ where $\hat{\phi}$ is the Fourier Transform of ϕ) of L^2. ϕ should be piecewise continuous and have compact support. ϕ is also known as a scaling function or father wavelet.

For each $j \in \mathbb{Z}$, the space V_{j+1} can be written as

$$V_{j+1} = V_j \oplus W_j.$$

The space W_j is spanned by the set of functions

$$\{\psi(2^j x - k)\}, k \in \mathbb{Z},$$

which are called wavelets. These wavelet functions exist for any multi-resolution analysis [19]. Furthermore, the set of scaling functions

$$\{\phi(2^j x - k)\}, k \in \mathbb{Z}$$

spans V_j. Moreover,

$$\phi(x) = \sum_{k=-\infty}^{\infty} p_k \phi(2x - k) \tag{2.10}$$

and

$$\psi(x) = \sum_{k=-\infty}^{\infty} q_k \phi(2x - k). \tag{2.11}$$

for some sequences of numbers $\{p_k\}$ and $\{q_k\}$.

We denote the scaling and wavelet functions according to their dilations and translates:

$$\phi_{j,k}(x) = \phi(2^j x - k) \tag{2.12}$$

$$\psi_{j,k}(x) = \psi(2^j x - k). \tag{2.13}$$

The dual wavelet

For any multi-resolution analysis, $\exists\ \tilde{\phi}$, satisfying the biorthogonality condition, i.e.

$$\langle \phi(x-k)\tilde{\phi}(x-l)\rangle = \delta_{k,l} \tag{2.14}$$

The dual wavelet may be derived from the dual scaling function by use of the scaling equation, Eq. (2.11). This type of wavelet is called biorthogonal.

Functions of $L^2(\mathbb{R})$ can be represented in the dual wavelet basis:

$$f(x) = \sum_k \langle f, \phi_{j,k}\rangle \tilde{\phi}_{j,k}(x) + \sum_{j,k} \langle f, \psi_{j,k}\rangle \tilde{\psi}_{j,k}(x), \tag{2.15}$$

The inner product $\langle \cdot, \cdot \rangle$ denotes the usual inner product in L^2.
$\langle f, \tilde{\psi}_{j,k}\rangle$ is the wavelet coefficient corresponding to the wavelet function $\psi_{j,k}$.

Image approximation using wavelets

It is well known that a given number N of basis functions will generally approximate a continuous function to more accuracy than a discontinuous one, i.e., a function with discontinuities (in the case of images, edges). Consequently, the rate of convergence of the Fourier series of a discontinuous function is lower than that of a continuous one. Wavelets, on the other hand, have finite support which changes depending on the resolution and provide multiresolution analysis of a signal keeping track of time and frequency. Wavelet decomposition is obtained through the Discrete Wavelet Transform (DWT), which decomposes the signal into its frequency components that lie in ranges bounded by powers of 2. Ingrid Daubechies was the first to construct a family of compactly supported orthogonal wavelets [20]. Further development of multiresolution analysis and the fast wavelet transform by Mallat led to using wavelets in image processing [49].

Biorthogonal transforms are invertible but not orthogonal [84]. The Daubechies wavelet functions that are widely used in image compression (e.g. JPEG2000 compression) do not have a closed form and the decomposition and reconstruction of signals is obtained using high-pass and low-pass digital filters corresponding to parameters (polynomial coefficients) that characterize each wavelet function. The high-pass and low-pass filters that are employed in the DWT use coefficients that do not exist in the data set; this is called the "overspill problem". In order to overcome this issue, signal extensions are performed at

the boundaries. The most common are symmetric and periodic extensions. Periodic extensions introduce boundary discontinuities which cause large number of high frequency coefficients in the DWT. The advantages of symmetric extension over periodic extension have been studied in [81]. It is known that symmetric extension cannot be applied to orthogonal wavelets (for non-expansive DWT, i.e. the number of input samples is the same as the number of wavelet coefficients at any point during the decomposition process) [77]. Symmetric extensions require symmetric filters, which can be employed with biorthogonal wavelets. The filters corresponding to orthogonal wavelets are energy preserving (Parseval's identity). However, the Daubechies biorthogonal filters correspond to a wavelet basis that is "close" to orthonormal and thus preserve energy quite well under the DWT [73].

Let $f(x, y) \in L^2((R))$ represent an image function. The image function f assigns to each point (x, y) an intensity value. These values are the pixel intensities. In the case of a digital image, f is represented as a $N \times N$ matrix with entries (n, m) and integrals are replaced with finite summations.

The "projection theorem" justifies the use of orthogonal or biorthogonal basis functions to best approximate a function in the space L^2. The DWT converts an image into a space where image information is represented in a more compact form, which in turn reduces redundancies. These transforms separate the relevant and irrelevant information that are present in an image, i.e., into low and high frequency coefficients. The goal in image compression is to preserve the most relevant coefficients (low frequency coefficients), whereas the less important coefficients (higher frequency coefficients) are often rounded to zero. The most relevant coefficients are those that have large magnitudes.

2.2 Baseline JPEG Compression

The JPEG Baseline technique was developed by the Joint Photographic Experts Group and became an international standard in 1993 [57]. Figure 2.1 shows the steps involved in JPEG compression of a greyscale image. The algorithm starts with dividing an image into 8×8 pixel blocks. The rest of the algorithm processes each block independently, which is less computationally complex. The DCT is computed for each block. The DCT takes the pixel values of the image and transforms them into a matrix of frequency coefficients. The advantage of this operation is that the coefficients are now decorrelated (to some degree) and most of the image information is contained in a small number of these spectral DCT coefficients. The next step is scalar quantization, where each coefficient is divided by the corresponding quantization number. Depending on the degree of compression, quantization can be more or less strong. The quantized coefficients are then rounded to

the nearest integer. Next, the coefficients are reordered according to increasing spatial frequency using the 1-D zigzag pattern and run-length encoding (RLE) is applied. The RLE takes advantage of the long runs of zeros that usually result from reordering. This operation is performed by grouping similar frequencies together, storing single data value and count instead of the original string of values. Finally, Huffman coding is applied on the remaining data. The idea is that the most frequently occurring characters (numbers in this case) will be represented by shortest code words. Decompression is done by performing the inverse of the same operations in the reverse order. JPEG compression is widely used on the Internet and in digital cameras. It works best for images with smooth variations of tone. [57, 75, 23, 67]

JPEG is not optimized for L^2. A quantization table optimized for L^2 would have the full table with constant entries. However, in the case of JPEG, where non-uniform quantization is performed, finer quantization is applied to mid-low frequency DCT coefficients. The default JPEG DCT quantization table for luminance is given in 2.16. The JPEG quantization tables were claimed to be obtained through psychovisual testing. Interestingly, this does not necessarily mean that a uniform quantization table will lead to better L^2 performance. The real gain of JPEG relies on the number of zero DCT coefficients. The scan of DCT coefficients follows a zig-zag pattern starting from the top-left corner, and it stops when all remaining coefficients are zero. Thus, having larger quantization entries at high frequency coefficients is beneficial because it makes it a lot more likely to create more zeros at the tail of the zig-zag scan. Having more zeros at the tail of the scan makes the scan stop earlier. This is very important in JPEG compression in terms of saving bits. Furthermore, the JPEG default quantization table and zig-zag scan is indeed a coherent joint design.

$$Q = \begin{Bmatrix} 16 & 11 & 10 & 16 & 24 & 40 & 51 & 61 \\ 12 & 12 & 14 & 19 & 26 & 58 & 60 & 55 \\ 14 & 13 & 16 & 24 & 40 & 57 & 69 & 56 \\ 14 & 17 & 22 & 29 & 51 & 87 & 80 & 62 \\ 18 & 22 & 37 & 56 & 68 & 109 & 103 & 77 \\ 24 & 36 & 55 & 64 & 81 & 104 & 113 & 92 \\ 49 & 64 & 78 & 87 & 103 & 121 & 120 & 101 \\ 72 & 92 & 95 & 98 & 112 & 100 & 103 & 99 \end{Bmatrix} \tag{2.16}$$

The default way of using JPEG quantization table is to apply a scale factor on top of the above table, round the resulting table to integers, and use the rounded integers to quantize DCT coefficients.

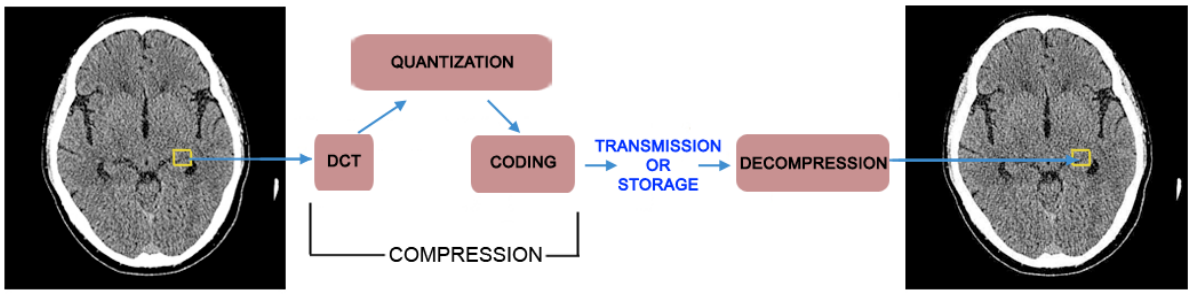

Figure 2.1: Steps in lossy JPEG compression. Each 8×8 block is first transformed using the DCT. These coefficients are then scalar quantized, and reordered using the zigzag pattern. Finally, the run-length encoding (RLE) and Huffman coding is applied. Decompression is done by performing the inverse of the same operations in the reverse order. (Brain CT image: MIIRC@M)

2.3 JPEG2000 Compression Format

JPEG2000 still image compression is based on the DWT and is the newest addition to the family of international standards that were developed by the Joint Photographic Experts Group. JPEG2000 compression involves several steps. These steps are shown in Figure 2.2. First, the greyscale image can be divided into tiles of equal size (this step is unnecessary). Then the DWT is applied (to each tile), which decomposes the image into subbands revealing the collection of details at different resolutions. The details represent the differences between two consecutive resolution levels and correspond to characteristics in the horizontal, vertical and diagonal directions. This subimage pyramid, which contains the "approximations" and "details" of the original image is obtained via an iterative filtering scheme and subsampling operations. The filters used in JPEG2000 compression are the biorthogonal Daubechies (irreversible 9-7 floating point filters and reversible 5-3 integer filters) wavelet filters [31]. Figure 2.3 shows 2-level wavelet decomposition applied to a neuro CT image. Most wavelet compression algorithms compute a four or five level decomposition. Next, each subband's wavelet coefficients are scalar quantized using a quantization step, which depends on the level of compression and on the characteristic of the details that the subband represents. Then, the quantized coefficients are put in codeblocks of sizes 64×64 or 32×32. Individual codeblocks are then divided into bit planes ranging from most significant (MSB) to least significant (LSB). Starting with the MSB plane with a nonzero element, each bit plane undergoes three passes: significance propagation, magnitude refinement and cleanup. Each bit in each bit plane is coded using only one of these passes. The resulting bits are then arithmetically coded and grouped with bits that result from similar passes from other code blocks. This forms layers, which are then divided into

pockets. Pockets allow the extraction of a spatial Region of Interest (ROI) from the total image data. Decompression is done by inverting the above operations. The number of bit planes used for reconstruction can be specified by the user. The compressed image can be recovered at progressive quality from the codeblocks. Due to the nature of the wavelet functions, which satisfy the properties of a multiresolution analysis, lossless and lossy compression can be obtained from one file stream [57, 31, 84, 23].

The goal of JPEG2000 was to provide: better performance at high compression ratios than JPEG, lossless and lossy compression in one file stream, support for 16-bit medical images [23, 39]. The disadvantages of JPEG2000 are that the format is not supported by Web browsers and that it requires complex encoders and decoders.

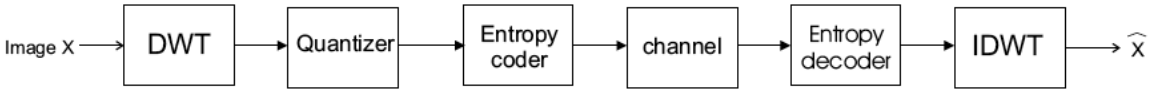

Figure 2.2: Block diagram of a wavelet based lossy compression system (JPEG2000). Figure adopted from [31].

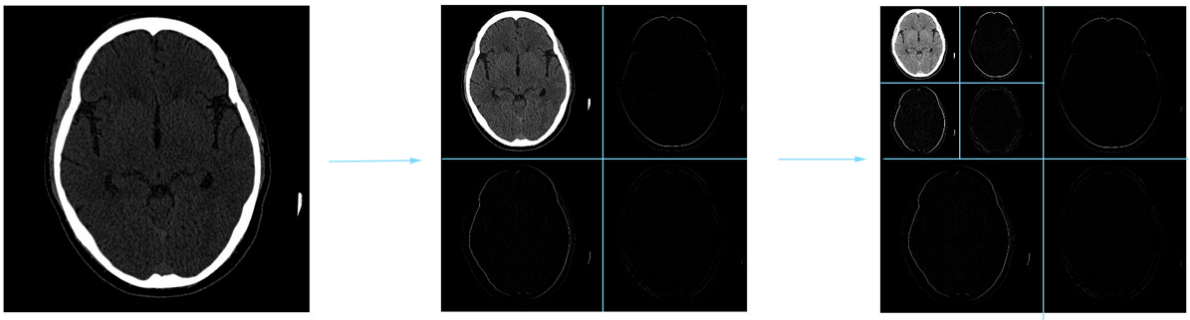

Figure 2.3: 2-level wavelet decomposition. (Brain CT image: MIIRC@M)

2.4 Medical Imaging

The process used to create images of the human body or parts of it for diagnostic purposes or for medical science is called medical imaging. Various physical phenomena are used to create medical images. The most common are: X-rays, γ-rays, ultrasound waves, nuclear

magnetic resonance. The use of computers in the acquisition process (real time treatment of a large amount of information) and for image reconstruction (tomography) has significantly increased over the years.

Medical imaging is a noninvasive process. This means that imaging modalities do not penetrate through the skin physically. However, the electromagnetic radiation that passes through the body alters the physical and chemical reactions of the body in order to obtain data. In some cases, the electromagnetic radiation can be ionizing, i.e. the electromagnetic waves have enough energy to detach electrons from atoms or molecules. Some procedures, however, involve the introduction - either by ingestion or injection - of some kind of contrast medium (e.g. iodine and barium in the case of CT) in order to enhance the contrast of structures or fluids within the body and can be detected quite conspicuously by the imaging modality.

The visual presentation of the human body depends on the image formation process and the features of various medical modalities correspond to the physical and physiological phenomena observed. Consequently, each imaging modality produces images that differ in spatial resolution, contrast and type of noise [57]. Sagittal slices of the brain using four different medical image modalities are shown in Figure 2.4 and include X-ray CT, MRI and ultrasound imaging.

Medical images are generally low contrast images. They contain a complex type of noise, which is introduced due to the various acquisition techniques, transmission, storage, display devices as well as resulting from quantization, reconstruction and enhancement algorithms. Unfortunately, all medical images contain noise, which has grainy, textured or snowy appearance. In CT, MR and ultrasound images noise is very significant and may reduce the visibility of some diagnostically relevant features, especially for low contrast objects [80].

Of particular interest are CT images, which are discussed in the next section.

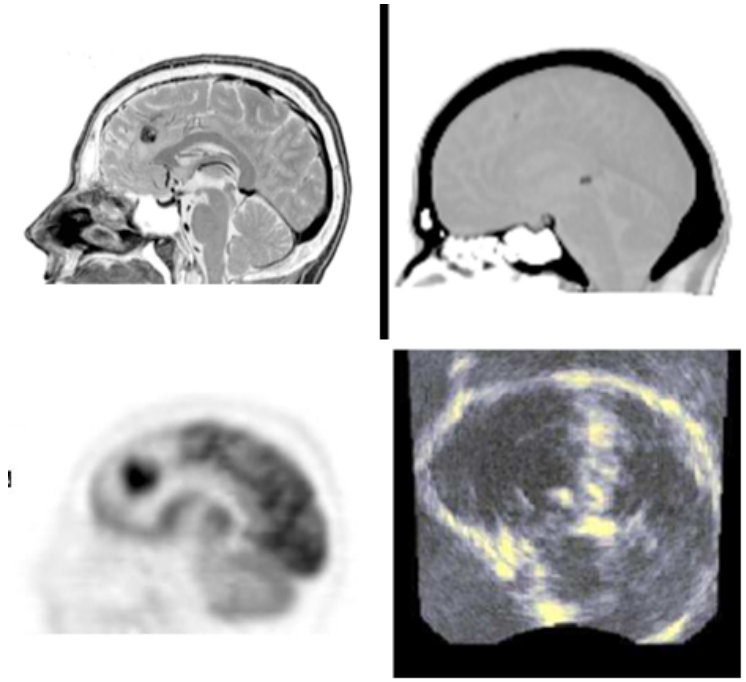

Figure 2.4: Sagittal slices of the brain obtained using four different medical image modalities. a) magnetic resonance imaging (MRI), b) computed tomography (CT), c) positron emission tomography (PET), d) ultrasound. The figure was adopted from [57].

2.4.1 X-ray and Computed Tomography (CT)

X-rays are one of the oldest sources of electromagnetic radiation used in medical imaging. The X-rays are generated by using an X-ray vacuum tube with a cathode and an anode. When the cathode is heated, free electrons are released and then flow to the anode of positive charge. The electrons hit the metal target in the anode and the energy is released in the form of X-rays. The strength of the emitted X-rays is controlled by the voltage on the anode as well as the current in the cathode. The patient is placed between the source of the radiating X-rays and a film that is sensitive to X-rays. The X-rays penetrate the patient's body and images are created using radioluminescent storage phosphor plates or flat matrix panels.

In the former case, the intensities of X-rays are recorded when the phosphor plates are exposed. They are then put in a special reader where a laser beam scans over the panel and converts the X-rays into light, which is then detected by a photomultiplier tube and the signal is digitized.

The flat matrix detectors have receptors with the ability of converting X-rays into an electric charge. For each pixel, an electric charge is received by a low noise electronic system and then digitized. These detectors produce a single projection plane of the region being imaged.

X-rays are used in CT imaging. Each CT image represents a "slice" of a part of the patient's body. CT scanners are moved around the patient in order to obtain a variety of projections from many different angles. The sum of the attenuation coefficients of all tissues that the X-ray beam has passed through is the total X-ray transmission measured by each detector. Although most tissues in the body are composed mainly of water, the variations are sufficient to record differences in the attenuation coefficients. The collection of the ray sums for all the detectors at a given position is a projection (Radon transform); about 800 data points corresponding to the 800 individual detectors is a single matrix array. About a thousand of these projections are needed to produce one CT slice. Sinograms are the graphic plots of projections (amplitude) vs. the X-ray tube angle (phase). Generating a CT image out of these projections requires determining the relative attenuation coefficients of individual pixels in order to discriminate small differences, which are expressed as Hounsfield units (HU). Figure 2.5 shows a plot of Hounsfield units corresponding to different tissues (densities). A cross sectional CT image is reconstructed by performing filtered back-projection on the projected data set. Three dimensional filtered back-projection can be used for 3-D reconstruction of the object being imaged. [57, 35, 6].

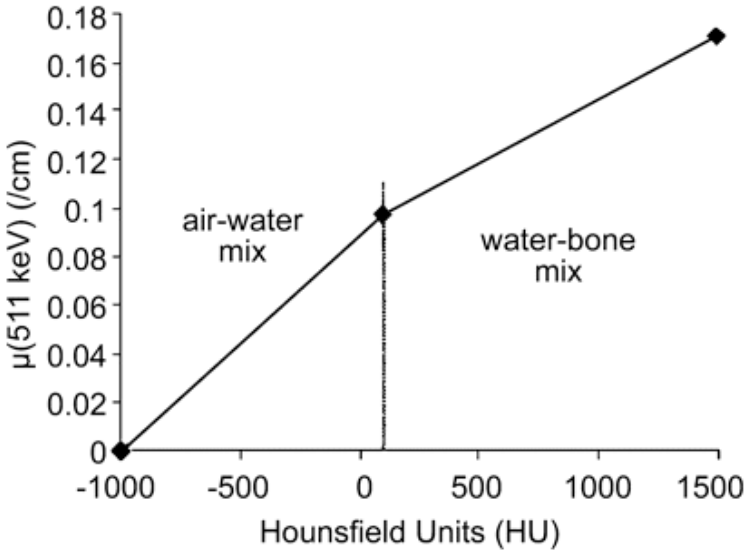

Figure 2.5: The plot shows a function that converts CT numbers to linear attenuation values at 511 keV. Linear attenuation coefficient at 511 keV is a function of corresponding CT value (in Hounsfield units). Figure adopted from [7].

CT provides anatomic information and is used to examine bones, breasts (mammography) or lungs. CT scanners are very useful for the visualization of complicated fractures, examination of organs in the neck or in the abdomen. Angiograms are X-ray images obtained by injecting X-ray contrasting agent in the coronaries through a catheterization process. The resulting images have better contrast of the blood vessels with the surrounding tissue. This technique is used in cardiology for coronary angiography.

A single cardiac (angiography) examination consists of 2000 to 3000 images (512×512×8 bits), which requires about 600 Megabytes of memory. CT imaging is very expensive in terms of storage; for example Angers hospital (France) accumulates 4.5 TB of CT data a year [57], (which comprises 35% of all of its imaging data)!

A single slice CT scanner generates one tomographic image (slice) for each full rotation (360 degrees). The recent developments in CT technology provide multi-slice scanners, which allow for 320 slices (with slice thickness of 0.5mm) to be acquired in one rotation of the X-ray tube. There are more detecting channels, which results in better temporal and spatial resolution of the images. Therefore, CT images obtained using the new technology are of better quality as compared to the older generations CT imaging systems.

Of particular concern is the amount of ionizing radiation that the patient receives during CT examination. Recent developments in CT imaging led to less intrusive and highly efficient imaging. For example, the 3D images enable us to view the aorta and the arteries of the lower limbs through a single acquisition and thus allowing to visualize a total length of more than one meter, made up of sub-millimeter slices [57, 35].

In a CT image, a variety of image artifacts can be observed. Aliasing, observed as dark lines radiating away from edges, results from taking not enough projections or from insufficient penetration of the X-rays. Partial volume effect appears as blurring, which results from the slices being too thick. There are also other artifacts such as ringing and motion artifact [34].

2.5 Medical Image Standards and Formats used in the Healthcare Industry

2.5.1 Picture Archiving and Communication systems (PACS) and Digital Imaging and Communications in Medicine (DICOM)

PACS is a solution to manual filing, retrieval and transfer of digital medical images. PACS allows electronic retrieval and storage and simultaneous access (at more than one site) of medical images, with support of multiple imaging modalities. This communication is usually carried by a server. The universal file format that is used in PACS is DICOM. DICOM is a data transfer, storage and display protocol that covers the functional aspects of digital medical imaging. The DICOM standard differentiates among various medical modalities. It is organized into independent sections, which specify the DICOM file format and rules for printing, communication of related images and data over a network or using a physical media, the security of data exchange, monitor display, and other tasks. Although DICOM has been mainly used in hospitals, small medical practices have started using it.

DICOM is a file format in which the information (such as medical images, patient information, reports, interpretations and other related data) is grouped into data blocks. The image data cannot be separated from the patient data (nor from other DICOM data blocks). The DICOM standard supports several lossless and lossy compression techniques including JPEG (8, 12-bit images) and JPEG2000 (8,12, 16-bit images). PACS allows quick and easy access to images and other related patient data from past examinations, which might in turn lead to a more reliable diagnosis.

Internet browsers play an important role in viewing patient data. Browsers do not support DICOM format however through web based server-client applications, patient's data including the related medical image can be viewed. It is very common that doctors view medical images using an Internet browser such as Chrome. Due to browser limitations as well as transmission speed, only compressed versions of images are displayed. The compression formats that are supported in DICOM and most Internet browsers are 8-bit images compressed with JPEG and JPEG2000 algorithms. This is the reason why we base our studies on these two image formats. As discussed previously, the amount of compression raises several concerns as it is still questionable which measure is most accurate in assessing the quality of compressed images. In the next chapters we are hoping to at least partially answer that question, however, we admit that there is still a lot of research to be done in order to find answers to these important problems in the radiological industry.

2.6 Quality Assessment of Medical Images

Lossy compression alters the original image. When working with medical images, we are not just concerned with the visual quality, we are also concerned with diagnostically relevant features. The goal of many researchers working in the field of image compression is to achieve diagnostically lossless compression as defined in Section 1.2. In order to study and evaluate lossy compression techniques, it is necessary to measure the degradations they produce. The amount of accepted distortion that would qualify a particular lossy compression as diagnostically lossless compression is a complex task and an open discussion among researchers. It depends on the image features, anatomical region and acquisition technique, compression method and compression ratio. As mentioned previously, recommended compression ratios for JPEG and JPEG2000 lossy methods have been proposed by Koff *et al.* [38] for various medical image modalities and anatomical regions. The degradation produced by diagnostically lossless compression should be sufficiently minor and not affecting radiological diagnosis. Clearly, it is absolutely necessary to correctly identify and measure degradations caused by lossy compression applied to medical images. In order to do so, we need reliable quality assessment methods.

Many objective image quality metrics have been proposed in the last decade. Due to the wide variety of image types and applications, image quality assessment is not (yet) fully automatic and subjective approaches are still predominant [57]. How do we measure diagnostic quality? It is the pathological condition that determines the information that must be retained in any given medical data.

There is no standard method to measure the quality of compressed medical images, however, three approaches are usually considered [12]:

1. Subjective image quality rating using psychovisual tests or questionnaires with numerical ratings
2. Diagnostic accuracy measured by simulating a clinical environment with the use of statistical analysis (e.g. ROC)
3. Objective quality measures such as the MSE.

Although subjective experiments are complicated and difficult to conduct [57], they are the most accepted way for measuring diagnosis reliability. In order to overcome problems with subjective quality assessment and to automate the process of assessing degradations, there is a need for reliable objective quality assessment of medical images. Objective medical image quality assessment, where no human intervention is required, must be compliant with the human visual system (HVS) and provide accurate diagnostic quality measures.

2.6.1 Subjective Methods for Image Quality Assessment

Subjective image quality rating experiment is conducted by presenting a randomized set of images to the observer, who will assign a rating to each image (typically from 1 to 5). Depending on the purpose of the test, the observer will be a typical user or a specialist (radiologists, in the case of medical images, for example). This is usually done by comparison (side-by-side or by flickering) of compressed and uncompressed images by readers.

For the purpose of evaluating speech and video quality, the ITU standard that describes several approaches to perform subjective experiments has been developed. These approaches include absolute category rating (ACR), single stimulus continuous quality evaluation (SSCQE), double stimulus continuous quality scale (DSCQS) [1] [70] and Overall quality rating scale (OQRS) [2]. There is no standard way for rating of medical images [12]. Although some of the above rating scales have been used in subjective tests for medical images [72], these standards may not fit well with medical applications, where the ultimate purpose is determining diagnostic quality. For example, in a typical ACR scales, there are five categories: Excellent, Good, Fair, Poor and Bad. Based on these categories alone, it is not clear how to define a threshold for "diagnostically lossless" compression. Specialized rating scales are often used in medical imaging research, where the description of the scale categories is more meaningful to radiologists. In [30], for example, the following modified

scale has been used: "A score of 5 is no distortion (Excellent), score of 4 represents a little distortion which can be ignored (Good), score of 3 shows distortion which can be seen evidently but it can be accepted (Fair), score 2 shows a lot of distortion, which can not be accepted (Bad), and finally score of 1 shows too much distortion, therefore can not be tolerated (Very Bad)." Another specialized scale was used in [46]. The following categories were defined:

1. Better than the reference image
2. Same as the reference image
3. Poorer than the reference image, but acceptable
4. Unacceptable.

A subjective experiment with binary image quality evaluation by comparison of compressed and the original uncompressed images has been conducted in [71].

Some studies report on the results of subjective assessment, with the assumption that there is no tolerance to any visually detectable difference between the compressed image and its uncompressed counterpart. Statistical analysis is usually performed on the data including computation of averages, variability and other statistical qualities. Plots of compression ratio versus PSNR and compression ratio versus the mean subjective scores are often compared and objective and subjective scores are presented in tables. More involved analysis is performed by modelling subjective assessments with an objective quality measure. For medical images, the purpose is to predict the radiologists subjective assessment using an objective criteria. This is typically accomplished by plotting the mean values of the objective measure against the mean values of subjective responses and fitting a curve to the obtained points [12].

The clinical performance of a method can be measured with the use of diagnostic accuracy, i.e. the correct grouping of subjects into clinically relevant subgroups. "Diagnostic accuracy refers to the quality of the information provided by the classification device and should be distinguished from the usefulness, or actual practical value, of the information" [48].

Diagnostic accuracy is assessed using statistical analysis. The most common tool to measure diagnostic accuracy is the Receiver Operating Characteristic (ROC) analysis [55] [83]. The result of an ROC analysis is a threshold, which depends on the trade-off between

the sensitivity and specificity. "For radiological applications, this involves asking radiologists to provide a subjective confidence rating of their diagnoses (typically on a scale of 1-5) which is then used as if it were a threshold to adjust for detection accuracy" [12].

ROC analysis has some pitfalls in the case of diagnostic accuracy measurement. ROC analysis is more appropriate to use in the case of a binary task. Firstly, radiologists are not used to assign values to the confidence of their diagnostic decision. Secondly, in the case of the non-binary scale, it is not clear how to detect the number of abnormalities that are absent/present. In a binary detection task, however, such as detection of a tumour in the brain, if the image is diagnosed as normal then only one abnormality is absent. There exist extensions and variations to the conventional ROC analysis that allow for multiple abnormalities and localization of abnormality:

- Localization ROC (LROC). Use of observers to specify location of pathology. [54] [94]
- Free-Response Operating Characteristic (FROC). Generalized LROC. It allows for multiple abnormalities and multiple locations. [14]
- Differential ROC (DROC) [13].

Furthermore, the ROC curve is a useful visual tool to compare two or more tests on a common set of scales. The performance of a classifier is determined using the area under the curve [60].

Studies on diagnostic accuracy generally involve detecting pathologies in medical images. Some examples include: the radiologists task involving the detection of enlarged lymph nodes in a CT chest scan [12], specifying the quadrant in which the pathology exists [38], and measurement accuracy of MR chest scans [68].

For the purpose of measuring diagnostic accuracy, it is necessary to determine a "ground truth" or a "gold standard", which represents the true diagnosis on which the ROC analysis is based. There are several ways a "ground truth" can be defined. For example, it could consists of subjective opinions of experienced experts, or represent results of an autopsy, surgical biopsy, reading of images from a different imaging modality, etc. Another one is called a "personal gold standard" where each reader's judgement on the original image is considered to be the "ground truth" for the readings of the same reader on the compressed images [12].

In our study, we propose a method that involves ROC analysis with the use of a binary task on detecting unacceptable compression degradations. In order to measure the degree

of allowable compression degradations, in our experiments we consider the radiologists to define the "ground truth". The result of our study is a threshold for diagnostically lossless compression corresponding to an objective quality measure. This method will be discussed in detail in Section 4.

2.6.2 Objective Methods for Image Quality Assessment

There are many full reference image quality assessment algorithms proposed in literature [29]. Full reference methods are based on comparison between the original image and its distorted version. The most popular ones include MSE/PSNR, Structural Similarity Index (SSIM) and its variations [88], Most Apparent Distortion (MAD)[45], Normalized Perceptual Information Distance (NPID) [59], Practical Im age Quality Index [95] and a wavelet-based measures: Visual Signal-to-Noise Ratio (VSNR) [15] and Visual Information Fidelity (VIF) [32]. The MSE/PSNR is related to the L^2 distance between image functions. "The SSIM image quality assessment approach is based on the assumption that the human visual system is highly adapted for extracting structural information from the scene, and therefore a measure of structural similarity can provide a good approximation to perceived image quality"[88], [91]; it is the method that this thesis is based on. The MAD method utilizes two strategies to determine image quality: "Local luminance and contrast masking are used to estimate detection-based perceived distortion in high-quality images, whereas changes in the local statistics of spatial-frequency components are used to estimate appearance-based perceived distortion in low-quality images" [44]. The NPID method is a wavelet based algorithm with laplacian pyramid subband decomposition based on Kolmogorov complexity [29], [59]. The Practical Image Quality Index is relatively simple to implement and it is based on the texture masking effect and contrast sensitivity function [95]. The VSNR is a low complexity method that considers near-threshold and suprathreshold properties of the HVS [15]. VIF is based on visual information fidelity [32].

The objective image quality measures are not necessarily reliable measures of diagnostic quality of medical images. Moreover, compression ratios, generally used as pre-compression quality predictors, indicate poor correlation with image quality and are image dependent [24]. According to Marmolin [51]: "MSE is not very valid as a quality criterion for pictures reproduced for human viewing and the improved measures could be derived by weighting the error in accordance with assumed properties of the visual system." It cannot be assumed that an objective quality metric that performs well for natural images will ensure a superior diagnostic quality for medical images. In spite of these pitfalls, MSE/PSRN and other objective methods have been used in medical image quality assessment.

How do we determine which objective quality assessment measure is best suited for medical images? The common way of comparing performances of lossy compression methods is to plot distortion rate (e.g. MSE, SSIM) versus bit rate curves (or compression ratio). [74] [28]. For medical images, the goal is to find an objective image quality that can best predict the subjective radiologists' assessments. The performance of an objective quality measure to predict subjective scores is usually measured by means of a curve-fitting model. First MSE values are plotted against mean scores of subjective assessments, then curve (e.g. polynomial spline, quadratic, exponential, logistic) is fitted to the resulting points [12]. Another tool used to compare image quality measures is the correlation coefficient. The correlation coefficient relates objective and subjective quality evaluations using statistical correlations. Best performance of an objective measure corresponds to the highest correlation with the subjective method. A review of comparison studies involving computation of the correlation is presented in [72].

More specialized objective quality measures that focus on perceptually significant features incorporating the human visual system have been studied in [10] [50]. Some researchers have focused on designing objective quality assessment methods that were specialized for medical images. An attempt to characterize the class and amount of distortion and to evaluate diagnostic accuracy and acceptable compression level has been proposed by Przelaskowski *et al.* in [71]. Subjective results have been tested against a Chi-square quality measure and recommended compression ratios have been proposed for MR and CT images using a compression based on the block DCT transform. In another study, statistical metrics have been used for quality evaluation of ultrasound compressed images [21]. In a paper by Oh *et al.* [64], a perceptual quality rating (PQR) is introduced. PQR is based on "a multistage perceptual quality assessment for compressed digital angiogram images (MPQA)", which provides distortion maps and contrast error measures with the consideration of the human visual system (HVS). The PQR measure shows satisfactory correlation with radiologists' assessment.

In this thesis, we study and compare the most popular objective quality assessment methods used in literature, namely MSE/PSNR and SSIM index [72]. Although MSE/PSNR is shown to poorly correlate with visual quality, it should not be taken for granted that any perceptual object quality measure must be better. SSIM and other objective measures show better performance for natural image/video content for consumer electronics applications based on subjective tests [89] [90] [3]. Moreover, no objective model has not been yet "established" for medical images, especially when using radiologists as subjects. We aim to validate the SSIM index as a reliable measure of diagnostic quality. The SSIM index provides a meaningful quality map, that allows localized analysis and reveals image quality variation over space closely. The SSIM map may be helpful to explain the differ-

ences between compression algorithms. It may be also very useful in suggesting ways to develop better image compression algorithms for medical images. In what follows, the two objective quality assessment methods, MSE and SSIM, are discussed in detail.

The Mean Squared Error (MSE)

The squared L^2 distance between the compressed image g and the original image f, is defined as

$$\Delta^2 = \sum_{i=1}^{M} \sum_{j=1}^{N} \left(f(i,j) - g(i,j)\right)^2 .$$

The Mean Squared Error (MSE) between the compressed image g and the original image f is given by

$$\mathrm{MSE}(f,g) = \frac{\Delta^2}{NM} = \frac{1}{MN} \sum_{i=1}^{M} \sum_{j=1}^{N} \left(f(i,j) - g(i,j)\right)^2 .$$

In image processing, one also sees the Root Mean Squared Error (RMSE):

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

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

$$\mathrm{PSNR}(f,g) = 10 \log_{10} \left(\frac{L^2}{\mathrm{MSE}(f,g)} \right).$$

Here L is the dynamic range.

Why is MSE so popular? Some of the reasons are: it is simple and fast to implement, it is a metric (non negativity, identity, symmetry and triangle inequality hold), energy preserving for orthogonal transforms (i.e. Parseval's identity), it is convex, differentiable

and symmetric, it is memoryless (i.e. independent of other samples) and it is a convention [88]. According to Marmolin [51], "MSE is not very valid as a quality criterion for pictures reproduced for human viewing and the improved measures could be derived by weighting the error in accordance with assumed properties of the visual system." It is well known that MSE does not correspond well with the Human Visual System (HVS) and it is not a good measure of image diagnostic nor visual quality. For example, it is possible that two images f and g both have the same MSE distance to a reference image and one of them is much "closer" visually than the other. This will be also shown later in this report. The failure of the MSE is partially due to the following: any spatial relationships between the samples of the signal, and any relationships between the original and the distorted image are ignored by the MSE, where image pixels are treated equally. The MSE can also be visualized as a point in an N-dimensional vector space (signal space), in which N represents the number of pixels in the image, and each point is an image. Adding a distortion to an image can be viewed as adding a vector from the original image. The set of all distorted images such that the values of the MSE between the original image and the distorted ones are equal, forms an "equal-MSE hypersphere" in the signal space. Various distortions produce images that differ visually from the original image as well as from each other. The MSE ignores the directions of these distortion vectors, which are also of importance. In the next section, we discuss the Structural Similarity Index (SSIM), a more recent image quality assessment measure which can differentiate the various distortion vectors.

Structural Similarity (SSIM) Index

There has been much effort in developing image quality measures that give more accurate information about visual differences. In a study done by Chandler *et al.*[16] on "Effects of spatial correlations and global precedence on the visual fidelity of distorted images", it was concluded that structural distortions are one of the most perceived distortions by the human eye.

The SSIM index, introduced by Wang and Bovik [89], assumes that the HVS is highly sensitive to structural information/distortions (e.g. JPEG blockiness, "salt-and-pepper" noise, ringing effect, blurring) in an image and automatically adjusts to the non-structural (e.g. luminance or spatial shift, contrast change) ones [88]. Another assumption of the SSIM index is that images are highly structured and there exist strong neighbouring dependencies among the pixels, which the MSE totally ignores. The SSIM Index measures the difference/similarity between two images by combining three components of the human visual system- luminance, $l(f, g)$, contrast, $c(f, g)$ and structure, $s(f, g)$.

The general form of the (local) SSIM is given by

$$SSIM(f,g) = (l(f,g))^{\alpha} \cdot (c(f,g))^{\beta} \cdot (s(f,g))^{\gamma} \tag{2.17}$$

$$= \left(\frac{2\mu_f\mu_g + C_1}{\mu_f^2 + \mu_g^2 + C_1}\right)^{\alpha} \cdot \left(\frac{2\sigma_f\sigma_g + C_2}{\sigma_f^2 + \sigma_g^2 + C_2}\right)^{\beta} \cdot \left(\frac{\sigma_{fg} + C_3}{\sigma_f\sigma_g + C_3}\right)^{\gamma}, \tag{2.18}$$

where μ is the mean:

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

σ_f^2 is the variance:

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

σ_{fg} is the covariance:

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

and α, β and γ are parameters to define the relative importance of the three components. In most applications, we set $\alpha = \beta = \gamma = 1$ and the local (basic) SSIM Index is defined as follows:

$$SSIM(f,g) = \left(\frac{2\mu_f\mu_g + C_1}{\mu_f^2 + \mu_g^2 + C_1}\right) \cdot \left(\frac{2\sigma_f\sigma_g + C_2}{\sigma_f^2 + \sigma_g^2 + C_2}\right) \cdot \left(\frac{\sigma_{fg} + C_3}{\sigma_f\sigma_g + C_3}\right). \tag{2.19}$$

It is normally computed over $M \times N$ pixel neighbourhoods.

The SSIM behaves as follows,

$$-1 \leq SSIM(f,g) \leq 1, \tag{2.20}$$

$$\text{and } SSIM(f,g) = 1 \iff f = g. \tag{2.21}$$

As its name suggests, SSIM measures the similarity between f and g. The closer f and g are to each other, the closer $SSIM(f,g)$ is to the value 1. If f and g are identical, then $SSIM(f,g) = 1$.

The (non-negative) parameters C_1, C_2 and C_3 are stability constants of relatively small magnitude, which are designed to avoid numerical "blowups" which could occur in the case of small denominators. For natural images, there are some recommended default values for these parameters [88]. On the other hand, the question of optimal values for these stability constants for medical images is still an open one. The smaller the values of these constants, the more sensitive the SSIM index is to small image textures such as noise. In our study below, we shall examine a range of values for the stability constant (only one will be used, as explained below) in order to determine the value(s) which are optimal for the assessment of the diagnostic quality of medical images.

Note that in the special case $C_3 = C_2/2$, the following simplified, two-term version of the SSIM index is obtained:

$$SSIM(f,g) = \left(\frac{2\mu_f\mu_g + C_1}{\mu_f^2 + \mu_g^2 + C_1}\right)\left(\frac{2\sigma_{fg} + C_2}{\sigma_f^2 + \sigma_g^2 + C_2}\right). \tag{2.22}$$

The Local SSIM index is computed first within a sliding window that is moved across the image. This results in a "SSIM quality map", which reveals local image quality and is a useful tool in image analysis. The total SSIM score is computed by averaging the local SSIM values [89]. The formulation of the SSIM index in terms of ratios was motivated by the Weber's model of perception (discussed later). The SSIM index gives more accurate measures than the MSE. Figure 2.6 shows brain CT images with similar MSE values for very high and very low quality images. The SSIM values, however, differ relative to human perception. For example, high JPEG compression with clearly visible blocks produces structural distortion (very low quality image, the structure of the brain is lost!) and change in luminance, a non-structural and barely noticeable distortion (very high quality image), have similar MSE whereas the SSIM index is very low for JPEG compressed image and very high for the luminance shifted image. This example illustrates the problem with using MSE as a quality measure for images.

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.

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

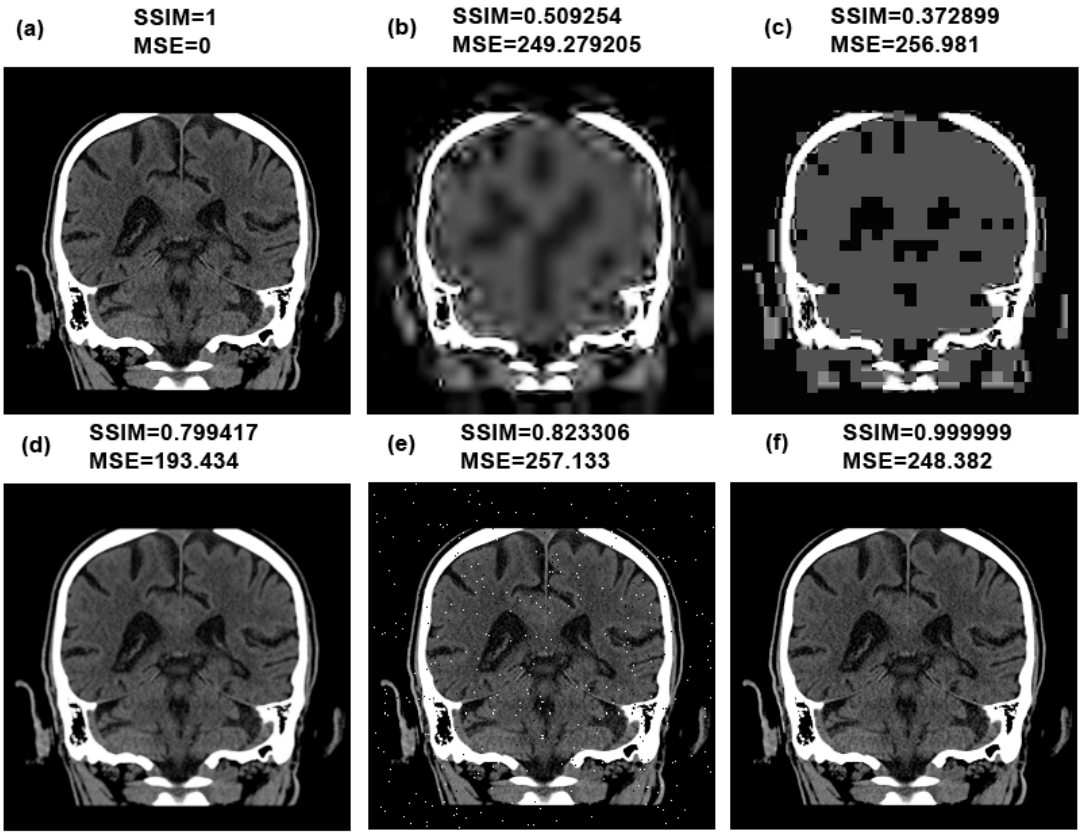

Figure 2.6: MSE and SSIM comparison. (Brain CT image: MIIRC@M) (a) the original brain CT image, (b) JPEG2000 highly compressed image, (c) JPEG highly compressed image, (d) luminance shifted image, (e) salt-and-pepper noise added (f) contrast enhanced image.

$$f_0 = f - \mu_f \, \mathbf{1}_{NM}, g_0 = g - \mu_g \, \mathbf{1}_{NM} \implies \mu_{f_0} = \mu_{g_0} = 0, \tag{2.23}$$

where $\mathbf{1}_{NM}$ denotes an $N \times M$ matrix of ones. Then,

$$DSSIM = \sqrt{1 - SSIM(f_0, g_0)} = \frac{||f_0 - g_0||}{\sqrt{||\sigma_{f_0}||^2 + ||\sigma_{g_0}||^2}}, \text{ for } C_2 = 0. \tag{2.24}$$

It has been shown that in [9] $\sqrt{1 - SSIM(f_0, g_0)}$ is a valid distance metric that satisfies the identity and symmetry axioms as well as triangle inequality.

There are several versions of the SSIM index including the multi-scale SSIM (MS-SSIM)

[91] and complex wavelet SSIM (CW-SSIM)[89].

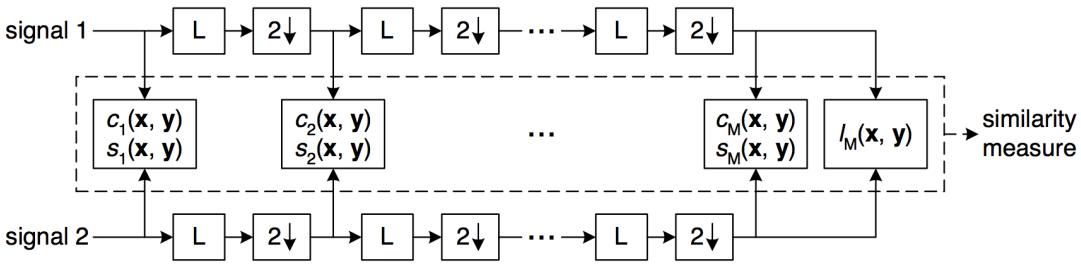

Figure 2.7: MS-SSIM system diagram. L: low-pass filtering; 2 ↓: downsampling by 2. Figure adopted from [91].

MS-SSIM gives more flexibility than the basic single scale SSIM. It incorporates variations of viewing conditions. A low-pass filter is applied iteratively and the filtered image is downsampled by a factor of 2. Figure 2.7 shows system diagram of MS-SSIM. For the original image and its downsampled versions, the MS-SSIM is computed as follows:

$$MS\text{-}SSIM(f,g) = [l_M(f,g)]^{\alpha_M} \cdot \prod_{j=1}^{M} [c_j(f,g)]^{\beta_j} \cdot [s_j(f,g)]^{\gamma_j} \tag{2.25}$$

where α_M, β_j and γ_j represent the relative importance of the three components. M represents the highest scale that is obtained after $M-1$ iterations.

The basic SSIM does not correspond well to non-structural distortions such as relative translations, scalings and rotations. The complex wavelet SSIM (CW-SSIM) [89], solves this issue by comparing the wavelet coefficients of the same spatial location in the same wavelet subband. CW-SSIM makes use of the fact that consistent phase shifts correspond to translations of the image. The CW-SSIM results in high scores for images with luminance and contrast changes, translations and rotations and low scores for images with structural distortions [88].

Another variation of the SSIM index is the Information Content Weighting Structural Similarity Index (IW-SSIM) [90]. The IW-SSIM is a two stage algorithm; Figure 2.8 shows the two steps involved in this approach. The IW-SSIM employs the same localized measurement approach as MS-SSIM, which results in a local quality map, however, the conversion of local quality map into a single quality score is achieved by a pooling algorithm in the

second stage of the algorithm. The content weighting approach involves assigning weights according to the information content of a given region of an image. The pooling method used in the IW-SSIM approach is saliency-based, with the resulting weighting function having connections with quality/distortion-based pooling method [90]. "The success of the IW-SSIM approach may be understood as a natural consequence of an effective combination of several proven useful approaches in image quality assessment research. These include multiscale image decomposition followed by scale-variant weighting, SSIM-based local quality measurement, and information theoretic analysis of visual information content and fidelity" [90] [78],[79].

Figure 2.8: Two stage structure of image quality assessment systems. Figure obtained from [90].

In this thesis we use several variations of the SSIM index. The MS-SSIM is used most often and we will refer to it, in general, as the SSIM-index.

Chapter 3

The Effects of Irreversible Compression on brain CT and body CT Images

3.1 The Effects of Irreversible Compression on Medical Images

Lossy compression does not result in perfect reconstruction. The quality reduction results in the appearance of artifacts in compressed images when the compression ratio is increased. For medical images, the goal is to obtain diagnostically lossless compression, where the degradation produced by compression is sufficiently minor and does not affect the assessment of radiologists. It is essential that the diagnostic interpretation and the visual quality are not altered by compression. The change of image features resulting from compression depends heavily on the characteristics of the image, the compression algorithm and the compression ratio used. At low compression ratios, the quantization step removes high frequency coefficients. These high frequency coefficients are low in magnitude and often correspond to high frequency noise. At low compression ratios, there are no visible degradations and the image is referred to as "visually lossless". When we increase the compression ratio, the first noticeable change is the removal of "salt-and-pepper" noise (it is generally true for medical images). These denoised images are often perceived as having better quality over the original images; they may improve diagnostic accuracy [23]. The differences can also be seen by looking at a quality map obtained from applying an image quality measure.

The greatest disadvantage of block transforms in image compression algorithms at higher compression ratios is the appearance of disjoint blocks, which generally will result in visible discontinuities at the block boundaries ("blocking effect"). The human visual system (HVS) is particularly sensitive to the presence of strong gradients that show up in different directions [84]. These artifacts are the main sources of visual disturbance. JPEG is based on the 8×8 DCT, whereas JPEG2000 is based on a dyadic multi-resolution subband discrete wavelet transform. These two methods also use different quantization and coding techniques; they are not optimized for viewer's present viewing conditions [84]. Blocking artifacts do not appear on images compressed with wavelet techniques, since the transformation is computed usually for the entire image (one tile). At high compression ratios, JPEG2000 produces "rice-shaped" artifacts that result from quantization of high frequency coefficients. Ringing artifacts are present in both JPEG and JPEG2000 algorithms; they can be seen as distortion around edges. Figure 3.1 shows blocking and ringing artifacts caused by JPEG and JPEG2000 at the compression ratio 40:1.

The biggest concern in using lossy compression for medical images is that important features of the image will be lost. The block effect, for example, could result in bad visibility of nodular details. Change of texture might alter diagnosis of a fracture or tumor in a mesh-like bone structures (trabecular bones). The fade in contrast or blurring in JPEG2000 in the high spatial frequency zones of the image reduces the clarity of the image's contours or its details (fibrous aspects). This sometimes causes the image texture to change, which could represent lesions (the "salt and pepper" pattern of glomic tumor, or radial and granular pattern of tumorous tissue) [57]. Dr. Franchetto, a neuroradiologist from Hamilton General Hospital (Hamilton, Ontario), comments that the most subtle brain pathologies are early stroke and trauma (small blood collections). Early strokes are very common, they occur at least twice a day at Hamilton General Hospital. She also states that fractures in the masteroid bone are not always visible and would be more difficult to compress; skull base fractures, however, are better suited for compression since there is more contrast with the surrounding tissue [58].

A large body of research has been dedicated to studying the differences between JPEG based on the conventional DCT block approximation and JPEG2000 based on the wavelet transform [74] [40] [65] [39] [4]. In most cases, JPEG2000 shows better performance.

Koff *et al.* [39] in their pan-Canadian evaluation study of lossy compression, found that: "At low levels of lossy compression, there was no significant difference between the performance of lossy JPEG and lossy JPEG2000. (...) they are both appropriate to use for reporting on medical images. At higher levels, lossy JPEG proved to be more effective than JPEG2000 in some cases of MSK CR, brain and abdominal CT images." A natural question to ask is why does JPEG perform better than JPEG2000 for some images? Brain

CT images, compressed with JPEG2000 show the main source of degradation in regions with speckle noise (acquisition noise). Features with energy spread over a number of small coefficients will be most affected by compression. Speckle patterns together with features such as irregular, small textures (for example, white matter in a brain CT image) are represented by numerous high frequency coefficients of low amplitude. These features, along with noise, are discarded during quantization. This results in local blurring and ringing with orientation and spatial extension corresponding to subbands where the coefficients were largely quantized. This local blurring occurs because wavelets are spatially localized; each DWT coefficient corresponds to a certain frequency and location. The DCT coefficients in JPEG compression, however, are global in nature (over each 8×8 block) and speckle patterns are a combination of various frequencies, which generally are larger and therefore are less susceptible to quantization [39]. The amount of energy that is distributed into low and high frequency coefficients is also referred to as the compaction property of the transform. Thus, it is safer to leave the high frequency coefficients that correspond to noise rather than reduce the texture of fine, irregular tissues [57].

Another paper that studied the differences between JPEG and JPEG200 algorithms and their effect on medical images concluded that subtle low contrast pathologies which might be difficult to see with the human eye, are well preserved by JPEG and JPEG2000 compressions. However, high-frequency patterns with irregular features are more degraded by compression. Moreover, pathologies vary in susceptibility to compression based on the distribution of their energies in the spectral domain [69].

Lossless compression methods applied to medical images yield low compression rates. For example, the JPEG-LS (lossless method) achieves, on average, 4:1 compression ratio for CT images [17]. Thus, the existing lossless compression schemes do not solve the problem of storage and transmission of medical images. For this reason, we focus our attention mainly on lossy compression schemes. The drastically increasing quantity of images generated by modalities such as CT scanners and MR imaging justifies the use of lossy compression for medical images, in order to decrease the storage costs and improve network transmission of data, provided there is no loss of clinically relevant information. The DICOM format, for example, supports lossy JPEG and JPEG2000. In Canada, for example, it is safe to compress neuro CT images using JPEG baseline algorithm with a compression ratio up to 12:1 and using JPEG2000 up to 8:1 [61] [38]. MR images can be compressed up to 24:1 using both JPEG and JPEG2000 algorithms. The goal, however, is to achieve diagnostically lossless compression beyond these recommendations.

Figure 3.1: (Left) original part of brain CT image. Examples of the types of artifacts caused by JPEG (middle) and JPEG2000 (right) at high compression levels. Observe clear blocks due to the division into 8×8 blocks by JPEG and the "rice-shaped" artifacts that result from quantization of high frequency coefficients in JPEG2000. Brain CT image: MIIRC@M

3.2 Investigating the effects of medical image compression on visual/diagnostic quality

One of the initial goals of this project was to predict potentially "bad" subblocks of a (medical) image with respect to a particular method of image compression (JPEG, JPEG2000), that is, those subblocks B_i of an image I with higher rates of degradation with respect to the degree of compression. We consider subblocks B_j with lower rates of degradation to be "good" subblocks. Of course, the term "degradation" of a subblock will have to be specified. As we all know, the usual procedure is to employ the L^2 distance: Let I be a reference image and suppose that a particular compression method is applied to it to produce the compressed image $\tilde{I}$. Let $u(B_k)$ and $\tilde{u}(B_k)$ denote the image functions associated with I and $\tilde{I}$ supported on a given subblock B_k. Then the degradation may be defined by the RMSE as defined in Section 2.6.2.

Given our work with the structural similarity (SSIM) image quality measure as defined in Eq. (2.22), we may wish to consider the following SSIM-based measure of degradation,

$$DSSIM = \sqrt{1 - SSIM(u(B_k), \tilde{u}(B_k))}, \tag{3.1}$$

It is quite obvious that flat regions compress better than those containing structure. However, for some medical images, for example brain CT images, the region inside the skull consists mostly of small structures. The changes of error (DSSIM and RMSE) of 8×8 blocks inside the skull of a brain CT image (Figure 3.2) that is later compressed versus the quality factor of JPEG are presented in Figures 3.3, 3.4, 3.5 and 3.6. These "fan-shaped" plots are also presented for JPEG2000 compressed brain CT image, with compression

ratio in the horizontal axis, in Figures 3.7, 3.8, 3.9 and 3.10. For a given 8×8 block of the image, we have plotted and connected points that correspond to the pair: (quality factor, error (DSSIM or RMSE)) for JPEG and (compression ratio, error (DSSIM or RMSE)) for JPEG2000. It is observed that some regions degrade more than others. Most regions inside the skull (i.e. soft tissue), have structures that appear visually close to each other and have approximately the same size. Furthermore, we expect that compression would have a similar effect on these regions, however, as depicted in the plots, there is certainly a wide variation in rates of degradation. However the variation in rate of degradation seems to be a rather continuous one - there does not seem to be a clear separation between "good" blocks and "bad" blocks. It seems, however, natural to try to classify blocks as "good" or "bad". According to the plots in Figures 3.3 and 3.4, there is a sudden decrease in visual quality according to DSSIM and RSME when quality factor is less than 20. A "close up" plot of quality factors below 20 is shown in Figures 3.5 and 3.6. Another thought provoking observation is the increase in error (DSSIM and RMSE) with the increase in quality for some blocks! What is so special about these blocks? The remainder of this section is devoted to a quantitative investigation of the above observations.

Figure 3.2: Brain CT image (Source: MIIRC@M, McMaster University, Hamilton). Size of the original image is 512×512 pixels.

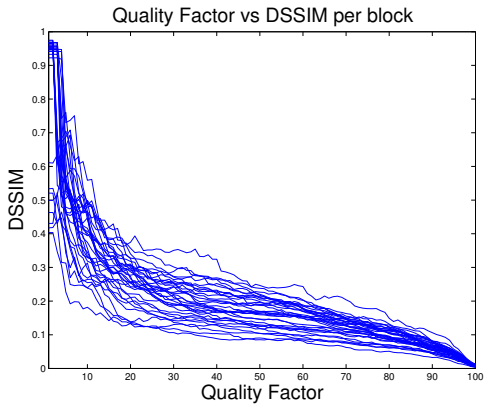

Figure 3.3: (JPEG) Plot of quality factor versus DSSIM corresponding to 32 8×8 blocks of the interior region of the skull of the brain CT image in Figure 3.2.

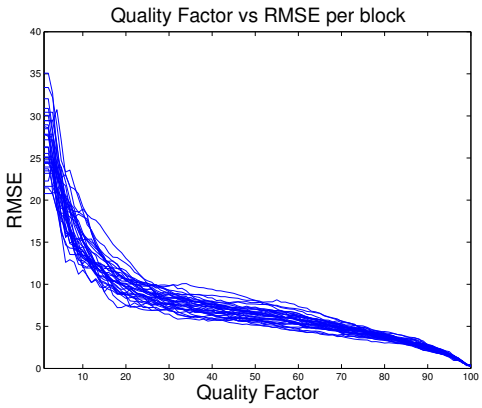

Figure 3.4: (JPEG) Plot of quality factor versus RMSE corresponding to 32 8×8 blocks of the interior region of the skull of the brain CT image in Figure 3.2.

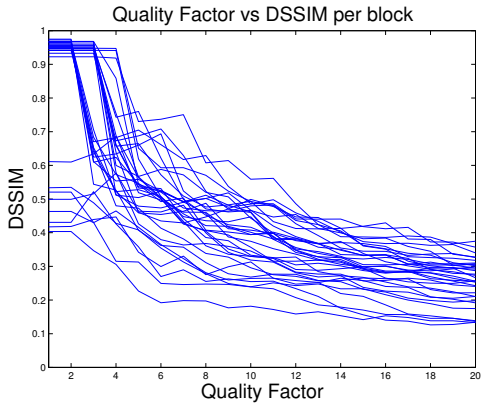

Figure 3.5: (JPEG) Plot of quality factor (< 20) versus DSSIM corresponding to 32 8×8 blocks of the interior region of the skull of the brain CT image in Figure 3.2.

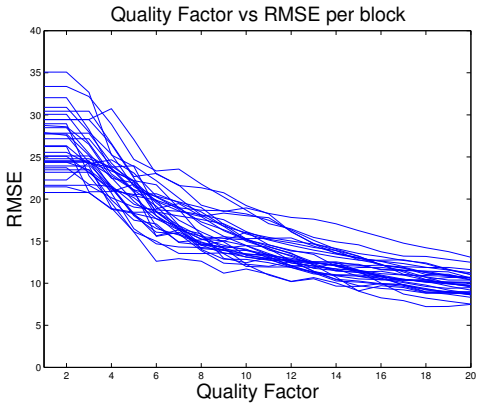

Figure 3.6: (JPEG) Plot of quality factor (< 20) versus RMSE corresponding to 32 8×8 blocks of the interior region of the skull of the brain CT image in Figure 3.2.

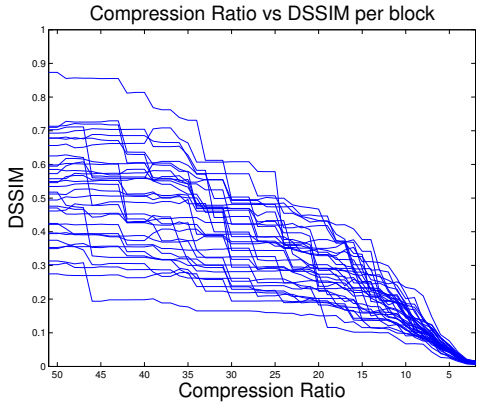

Figure 3.7: (JPEG2000) Plot of compression ratio (< 51) versus DSSIM corresponding to 32 8×8 blocks of the interior region of the skull of the brain CT image in Figure 3.2.

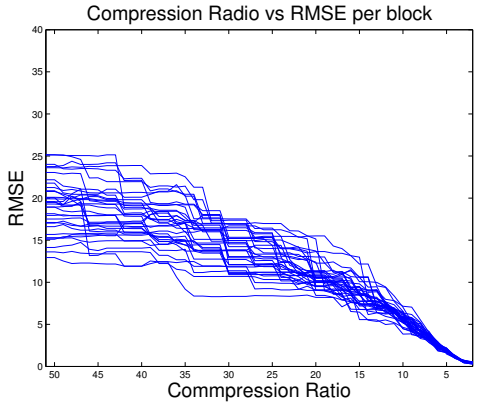

Figure 3.8: (JPEG2000) Plot of compression ratio (< 51) versus RMSE corresponding to 32 8×8 blocks of the interior region of the skull of the brain CT image in Figure 3.2.

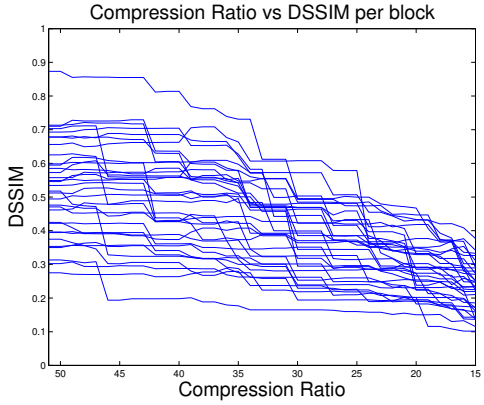

Figure 3.9: (JPEG2000) Plot of compression ratio ($15 - 51$) versus DSSIM corresponding to 32 8×8 blocks of the interior region of the skull of the brain CT image in Figure 3.2.

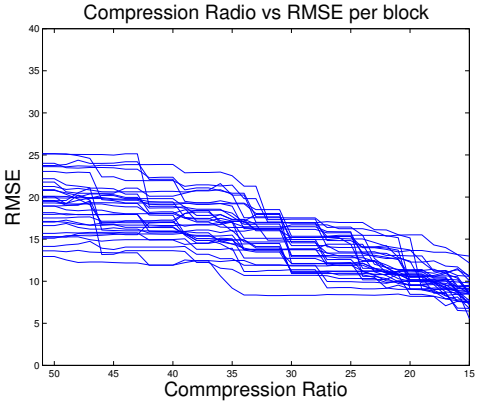

Figure 3.10: (JPEG2000) Plot of compression ratio ($15 - 51$) versus RMSE corresponding to 32 8×8 blocks of the interior region of the skull of the brain CT image in Figure 3.2.

We also present two other medical images, lung CT in Figure 3.12 and body CT in Figure 3.11. The corresponding plots of changes of error score versus quality factor (for JPEG) and compression ratio (for JPEG2000) are shown in Figures 3.13, 3.14, 3.15, 3.16, 3.17, 3.18, 3.19 and 3.20. These plots look similar to the ones corresponding to the presented brain CT image. We therefore conclude that it is not clear that there exists a separation between "good" blocks and "bad" blocks for these types of medical images.

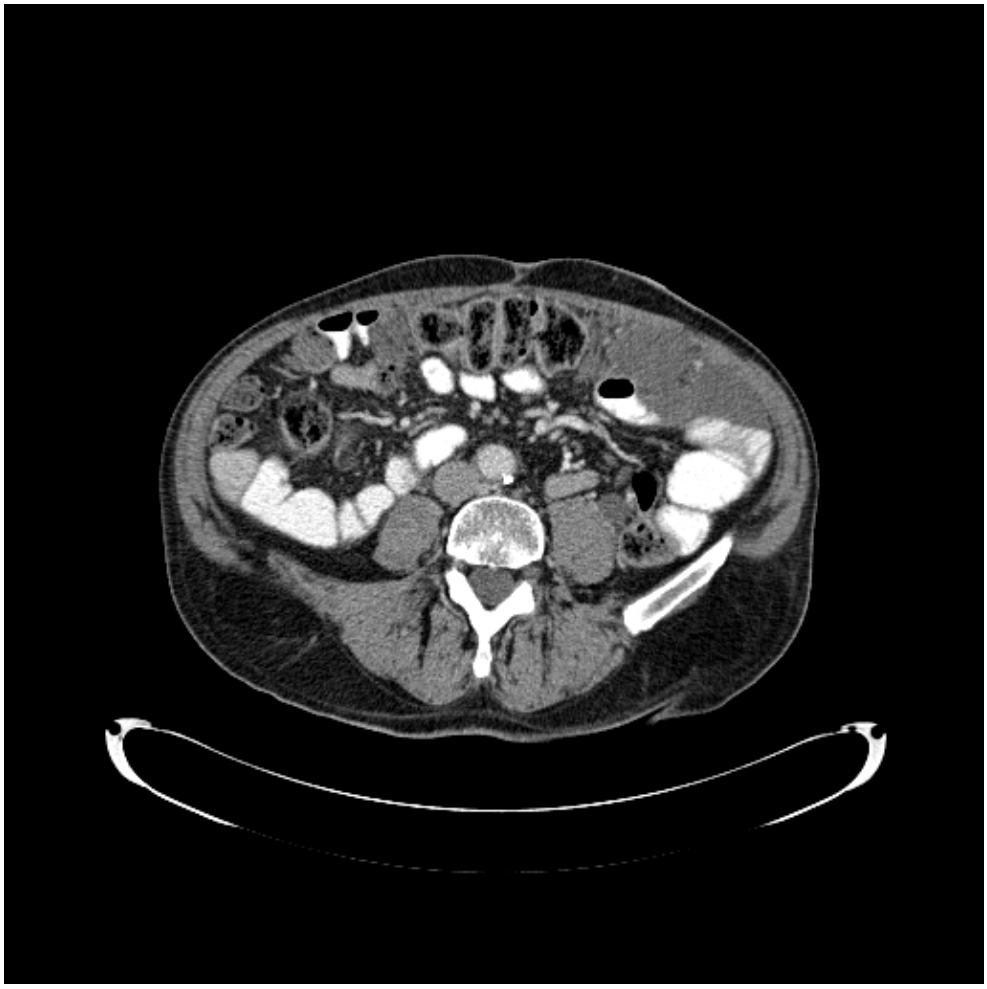

Figure 3.11: Body CT image (Source: MIIRC@M, McMaster University, Hamilton). Size of the original image is 512×512 pixels.

Figure 3.12: Lung CT image (Source: MIIRC@M, McMaster University, Hamilton). Size of the original image is 512×512 pixels.

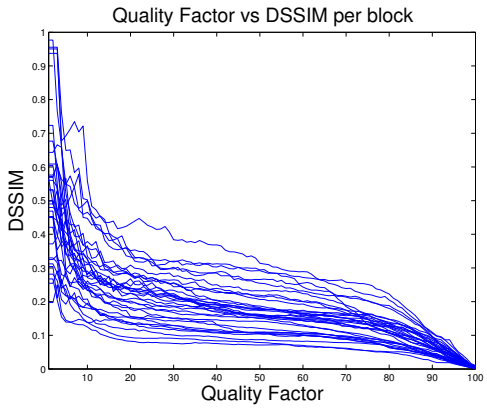

Figure 3.13: Plot of quality factor versus DSSIM corresponding to 32 8×8 blocks of a subregion of the body CT image in Figure 3.11, compressed using JPEG.

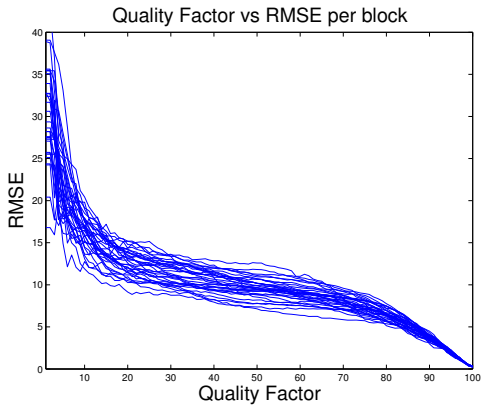

Figure 3.14: Plot of quality factor versus RMSE corresponding to 32 8×8 blocks of a subregion of the body CT image in Figure 3.11, compressed using JPEG.

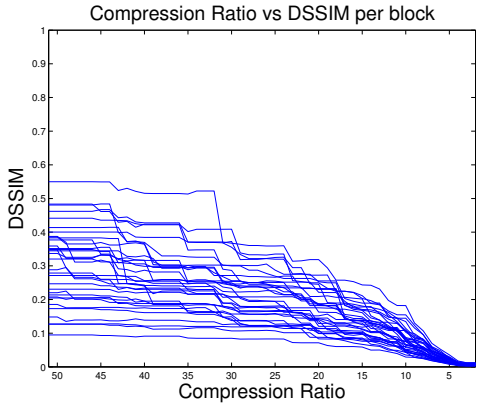

Figure 3.15: Plot of compression ratio (< 51) versus DSSIM corresponding to 32 8×8 blocks of a subregion of the body CT image in Figure 3.11, compressed using JPEG2000.

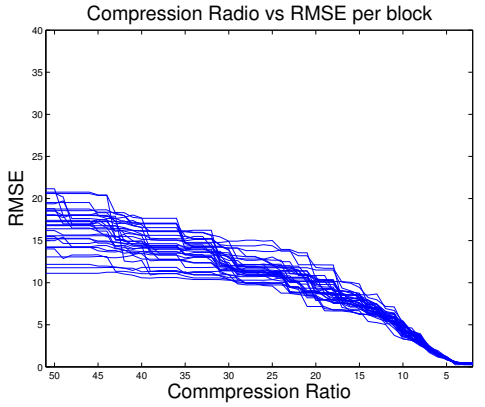

Figure 3.16: Plot of compression ratio (< 51) versus RMSE corresponding to 32 8×8 blocks of a subregion of the body CT image in Figure 3.11, compressed using JPEG2000.

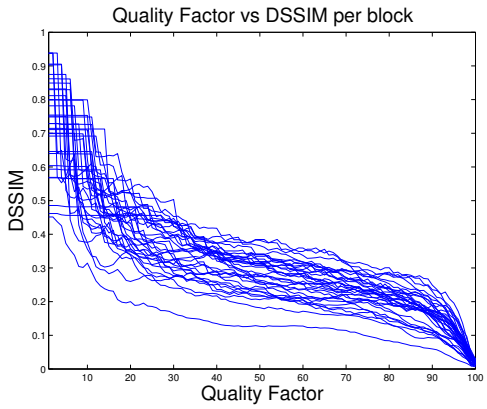

Figure 3.17: Plot of quality factor versus DSSIM corresponding to 32 8×8 blocks of a subregion of the lung CT image in Figure 3.12, compressed using JPEG.

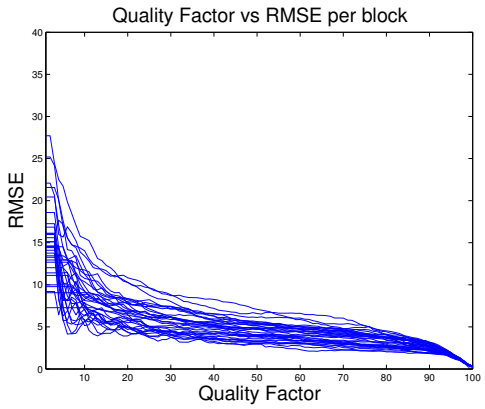

Figure 3.18: Plot of quality factor versus RMSE corresponding to 32 8×8 blocks of a subregion of the lung CT image in Figure 3.12, compressed using JPEG.

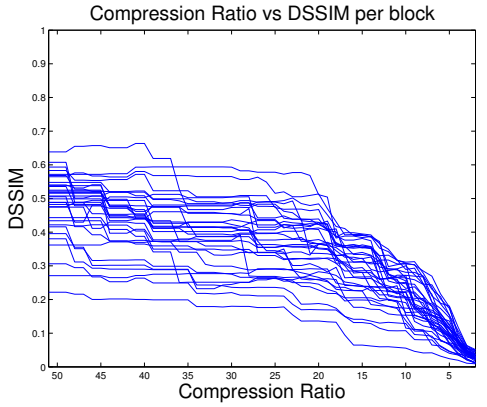

Figure 3.19: Plot of compression ratio (< 51) versus DSSIM corresponding to 32 8×8 blocks of a subregion of the lung CT image in Figure 3.12, compressed using JPEG2000.

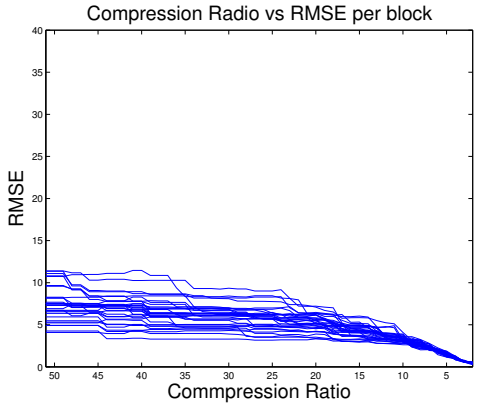

Figure 3.20: Plot of compression ratio (< 51) versus RMSE corresponding to 32 8×8 blocks of of a subregion of the lung CT image in Figure 3.12, compressed using JPEG2000.

3.2.1 Preliminary subjective experiments

We conducted two subjective experiments with radiologists (Dr. David Koff and Dr. Nadine Smolarski-Koff, McMaster University) with the hope of being able to divide a given compressed medical image into "bad" and "good" blocks. In these experiments, only JPEG compression was employed. In the first experiment, a particular brain CT image was selected from the database obtained from Medical Imaging Informatics Research Centre at McMaster University in Hamilton, Ontario (MIIRC@M) and it was displayed using a specially designed for this experiment image viewer (co-op project written by Faerlin Pulido, University of Waterloo). At each step, the uncompressed image was shown on the left and a compressed image on the right. Starting at quality factor Q = 100, and proceeding downwards, compressed images of decreasing quality factor were displayed (on the right) and examined by the radiologists until they noted some degradation. For a typical brain CT image, the first degradations of image features that were important in diagnosis, were noticed by radiologists at Q = 35. Using the viewer, the radiologists selected the pixels in the brain image where such significant degradation was noticed. The image was then subjected to further compression, i.e., lower Q values, until more degradation was noticed. Once again, the pixel locations were recorded. The process was continued.

To illustrate, here is a summary of results obtained for the first brain CT image studied.

- $Q = 35$: 4 pixels recorded.
- $Q = 20$: 4 pixels recorded.
- $Q = 15$: 5 pixels recorded.
- $Q = 10$: 9 pixels recorded.
- $Q = 5$: 19 pixels recorded.

For each particular compression, i.e., Q-value, the following was done:

- The 8×8-pixel block containing that pixel was determined,
- Some characteristic properties/potential indicators of each uncompressed 8×8-pixel block was computed, e.g., standard deviation, total variation, high frequency content, low frequency content.

- The RMSE and DSSIM degradations for all blocks (in other words, the RMSE and DSSIM distances between compressed and uncompressed blocks) were computed. The idea was to look for potential "indicators" or "descriptors" of good/bad blocks, including the following features:

- standard deviation (std),
- L^2 norm of high-frequency DCT components,
- L^2 norm of low-frequency DCT components,
- entropy,
- total variation (tv).

The degradations were plotted with respect to the above characteristic properties, along with the mean value of the degradations. For a given $N > 0$, and a given compression ratio, the N "best" and N "worst" blocks of a given image were selected on the basis of (i) L^2 error and (ii) SSIM error. In each case, two histogram plots of the error vs. each of the above quantities – one for the best blocks and one for the worst blocks – were constructed for each of the above features. If the best-block and worst-block histogram plots would show significant differences, then that feature would be considered as a good indicator. Before proceeding, let us qualify that we focus only on the effects of compression. As such, we do not consider any effects of spatial transformations, as produced by zooming, or greyscale transformations, as produced by windowing.

It is expected that as the JPEG quality factor is reduced, the errors generally increase. Figure 3.21 shows plots of the RMSE error versus DSSIM index for several JPEG quality factors corresponding to a brain CT image. As the quality factor gets larger, the difference between the two quality measures gets smaller. There is a kind of proportionality between DSSIM and RMSE, although there is quite a wide variability as well.

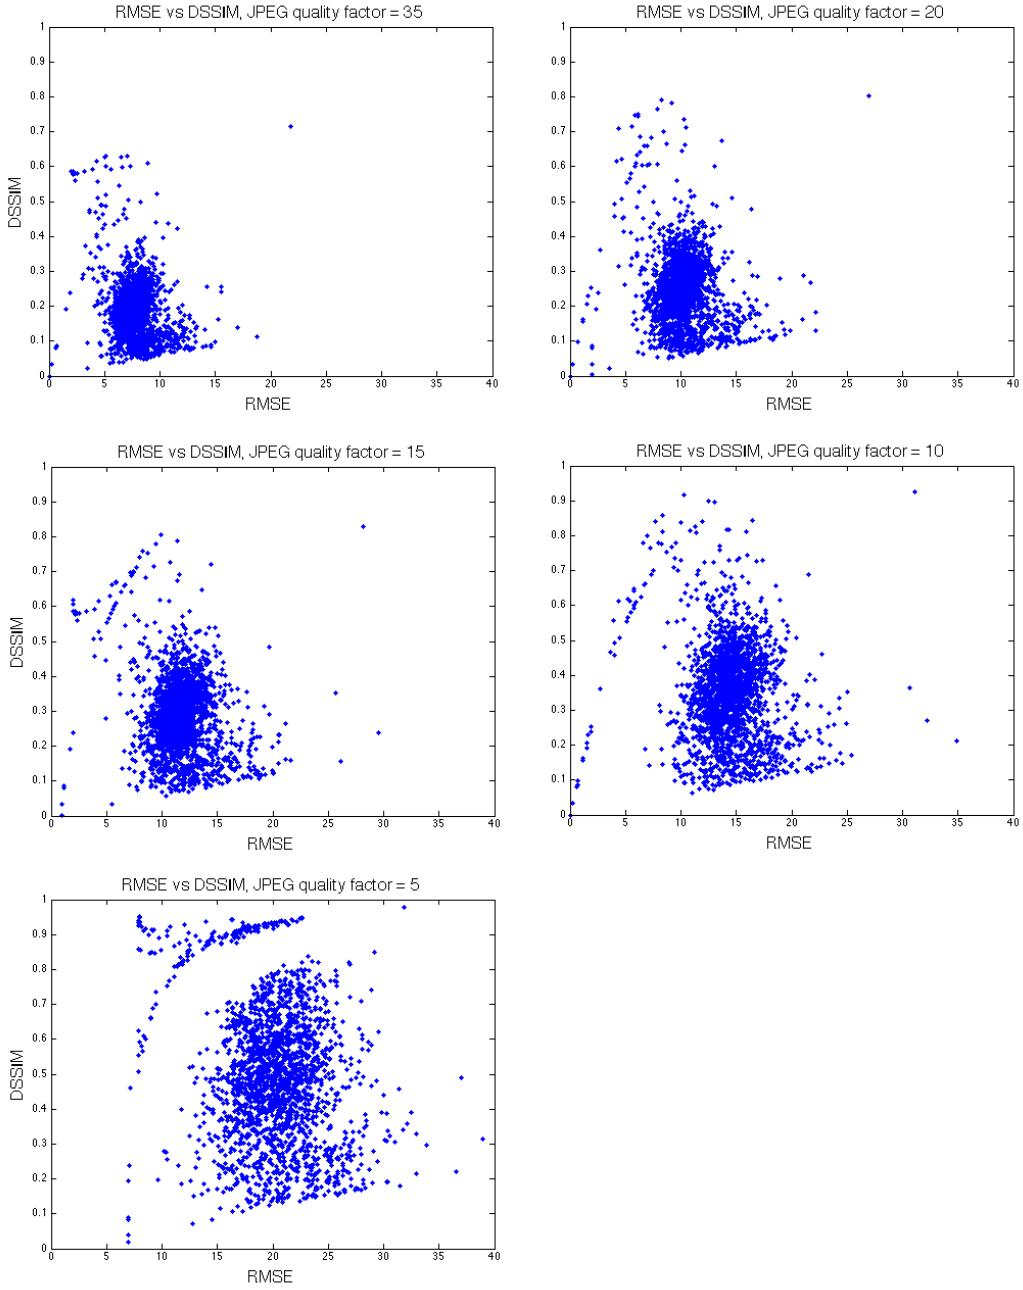

Figure 3.21: Plots of RMSE versus DSSIM corresponding to 8×8 blocks of a CT brain Image in Figure 3.2 compressed using JPEG with quality factors $Q = 35, 20, 15, 10$ and 5.

In the next experiment we asked the radiologists to select regions instead of single pixels. In order to simulate such results in this experiment we also consider neighbouring blocks of those identified by the radiologists in the following set of plots. The size of the neighbourhood is characterized by the parameter $nbhd \geq 0$: For a given block (i, j), where $1 \leq i, j \leq 64$, we consider all $(2 \cdot nbhd + 1)^2$ blocks $(i + k, j + l)$, $-nbhd \leq k, l \leq nbhd$.

Figures 3.22 and 3.23 show plots of statistical properties of the selected ("bad") blocks (shown as red dots) versus other blocks for a brain CT image at Q=10 with $nbhd = 0$ and $nbhd = 2$, respectively. Note that for the case $nbhd = 0$, i.e., only the blocks identified by the radiologists are used, the RMSE and DSSIM degradations of all of these blocks are observed to lie above the mean value of the respective degradation over the image.

In none of the figures shown above is there a clear, or even partial, distinction or separation to be found between good (blue) and bad (red) blocks. As such, we concluded at the time that such separation was not possible. Furthermore, the experiments presented in this section served as the first elements in the sequence of experiments on subjective assessment of compression artifacts that we have conducted. This preliminary study has helped to design a better experiment, presented in Chapter 4, where we have carefully selected the quality levels and asked radiologists to assess the images as "acceptable" or "unacceptable".

An interesting observation has been noted about the DSSIM scores and the respective DSSIM mean values. Here, we consider the blocks that were identified over all quality factors, $Q = 35, 20, 15, 10, 5$, for a total of 41 pixels identified by the radiologists with $nbhd = 0$. The image was compressed with quality $Q = 10$. We note that the DSSIM degradations of the blocks identified by the radiologists lie above the respective mean values over the image. The corresponding plots are shown in Figure 3.24. This is not always true for the RMSE degradations, which may be an encouraging result. This observation seems to be an indication that DSSIM is modelling the subjective assessment of radiologists in a better way than RMSE. Moreover, at first sight it may seem that there is a larger concentration of points above the mean. However, there is a very high unnoticed peak at $(0, 0)$. This peak corresponds to the flat (black) background blocks, which were not affected by compression at the given quality factors.

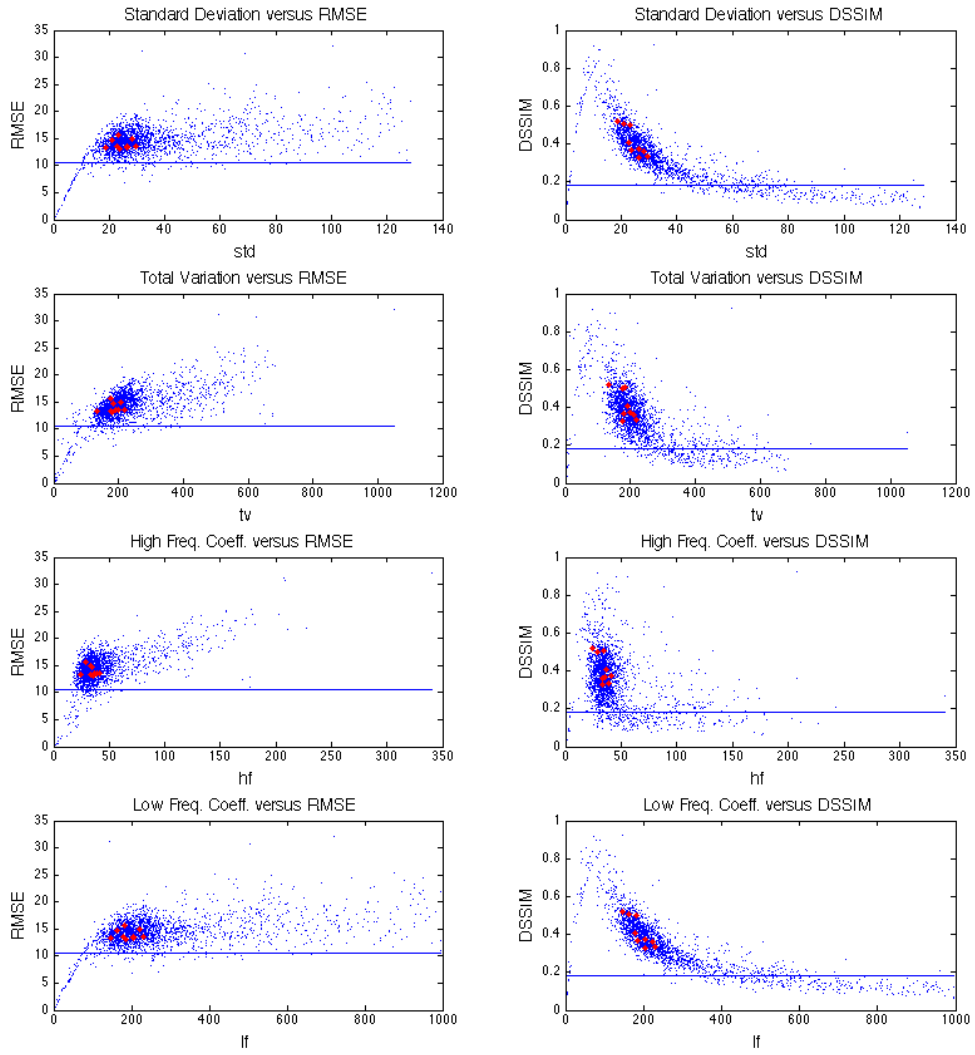

Figure 3.22: Plots of the RMSE and DSSIM errors for all 4096 8×8 blocks vs. (i) standard deviation (std), (ii) total variation (tv), (iii) high-frequency norm (hf) and (iv) low-frequency norm (lf). In each plot, the average value of the appropriate degradation is also plotted for comparison. The red points correspond to the blocks isolated by the radiologists, i.e. "bad" blocks. $nbhd = 0$ and quality factor $Q = 10$.

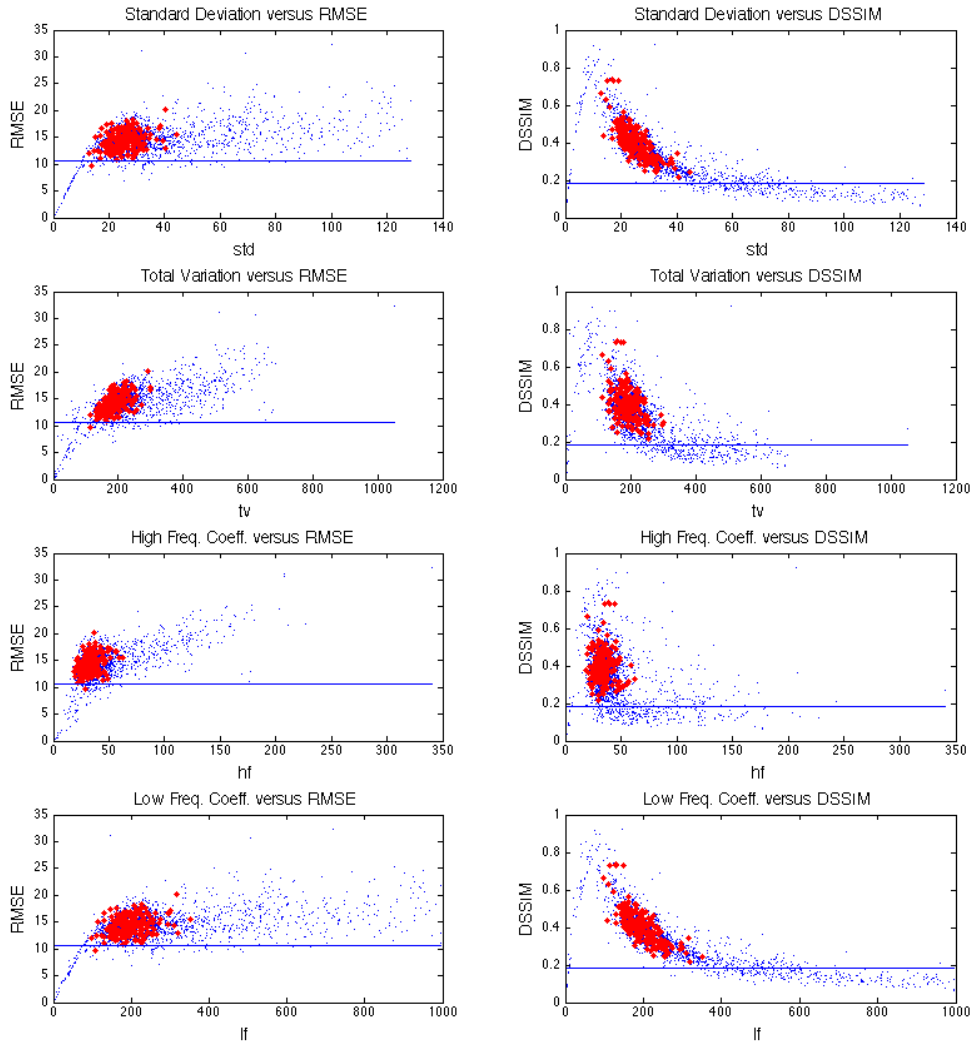

Figure 3.23: Plots of the RMSE and DSSIM errors of a JPEG compressed brain CT image in Figure 3.2 with quality factor $Q = 10$, for all 4096 8×8 blocks vs. (i) standard deviation (std), (ii) total variation (tv), (iii) high-frequency norm (hf) and (iv) low-frequency norm (lf). In each plot, the average value of the appropriate degradation is also plotted for comparison. The red points correspond to the blocks isolated by the radiologists, i.e. "bad" blocks. $nbhd = 2$ and quality factor $Q = 10$.

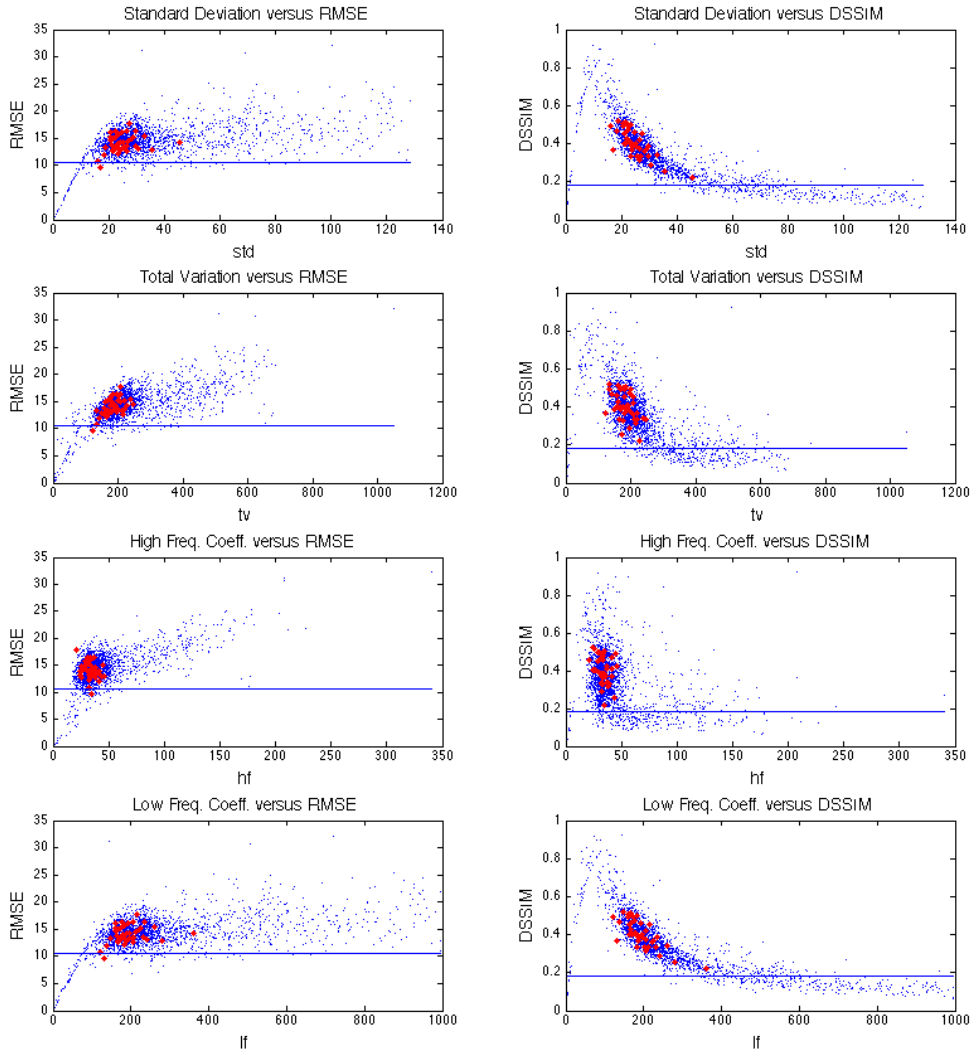

Figure 3.24: Plots of the RMSE and DSSIM errors for all 4096 8×8 blocks vs. (i) standard deviation (std), (ii) total variation (tv), (iii) high-frequency norm (hf) and (iv) low-frequency norm (lf). In each plot, the average value of the appropriate degradation is also plotted for comparison. The red points correspond to the blocks isolated by the radiologists, i.e. "bad" blocks. $nbhd = 0$.

3.3 Quantitative Evaluation of Degradations Resulting from JPEG and JPEG2000 Compression of Neuro CT images

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.* [39], 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 brain CT 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 neuro CT 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 neuro CT images (7 studies with 15 images in a study) of slice thickness 0.1 mm, obtained from Medical Informatics Research Centre at McMaster (MIIRC@M), Hamilton, Canada, was analyzed. JPEG2000 outperformed JPEG in terms of MSE. In terms of the DSSIM index, however, JPEG is seen to perform better for some compression ratios. Figure 3.25 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 3.26 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 2.24), 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 neuro CT 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

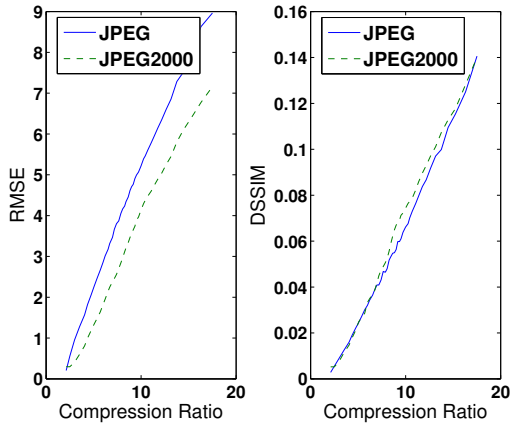

Figure 3.25: Plots of compression ratio versus RMSE and DSSIM of a neuro CT image

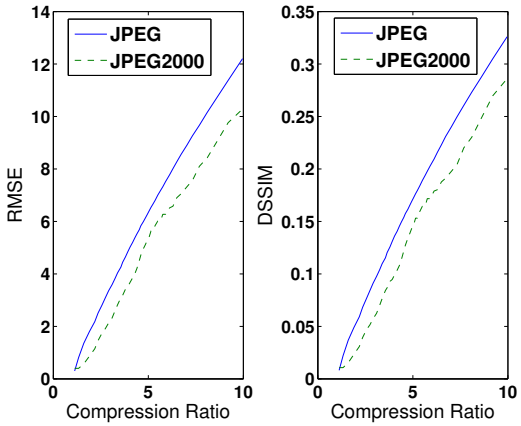

Figure 3.26: Plots of compression ratio versus RMSE and DSSIM of a cropped neuro CT image

images.

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

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 neuro CT image. A typical 8-bit brain CT 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 3.27). A neuro CT 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.

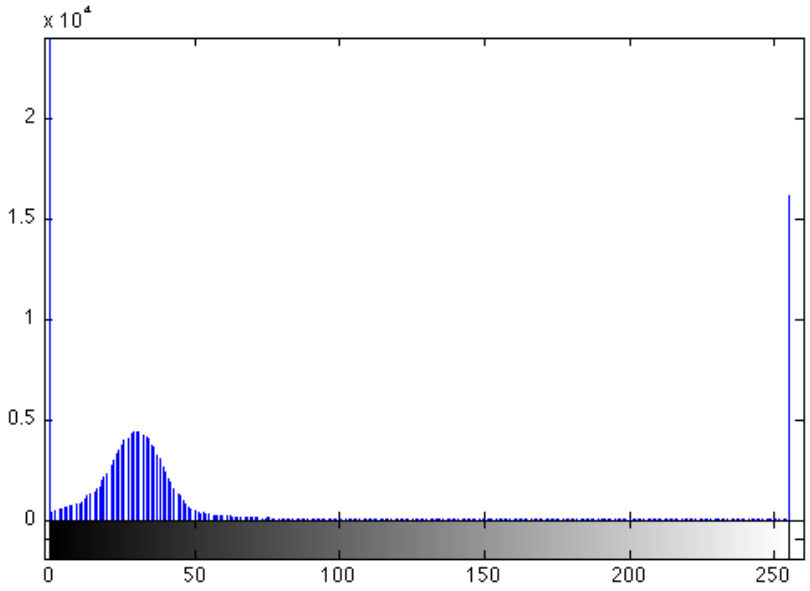

Figure 3.27: Histogram of a neuro CT image.

Neuro CT Image 1	**JPEG**	**Segmented JPEG**	**JPEG2000**	**Segmented JPEG2000**
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
Neuro CT Image 2	**JPEG**	**Segmented JPEG**	**JPEG2000**	**Segmented JPEG2000**
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 3.1: Quality scores using PSNR, SSIM, MS-SSIM and IW-SSIM for two JPEG and JPEG2000 compressed neuro CT 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 3.1). The new method produces less artifacts according to the SSIM local quality map for all the 105 images tested. Figures 3.28, 3.29, 3.30 and 3.31 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 neuro CT images improves the quality of JPEG and JPEG2000 compressed images.

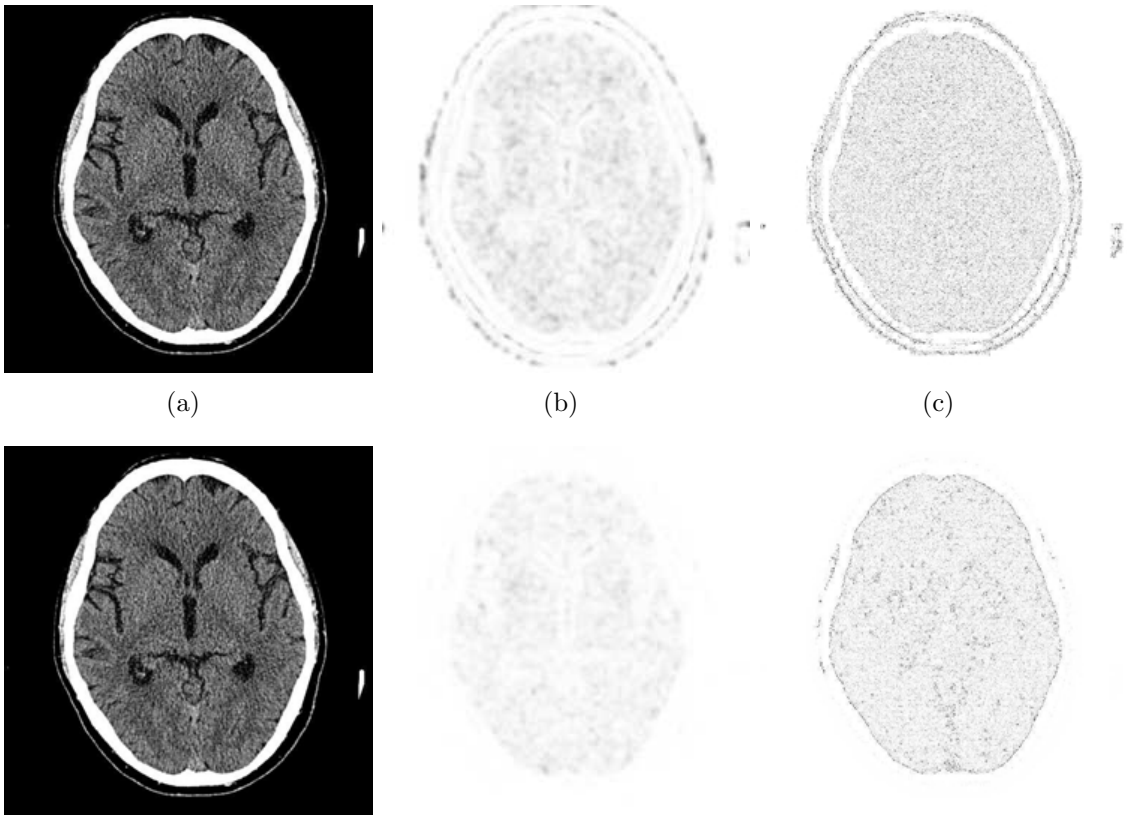

(a) (b) (c)

(d) (e) (f)

Figure 3.28: Quality maps of a JPEG compressed neuro CT 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.

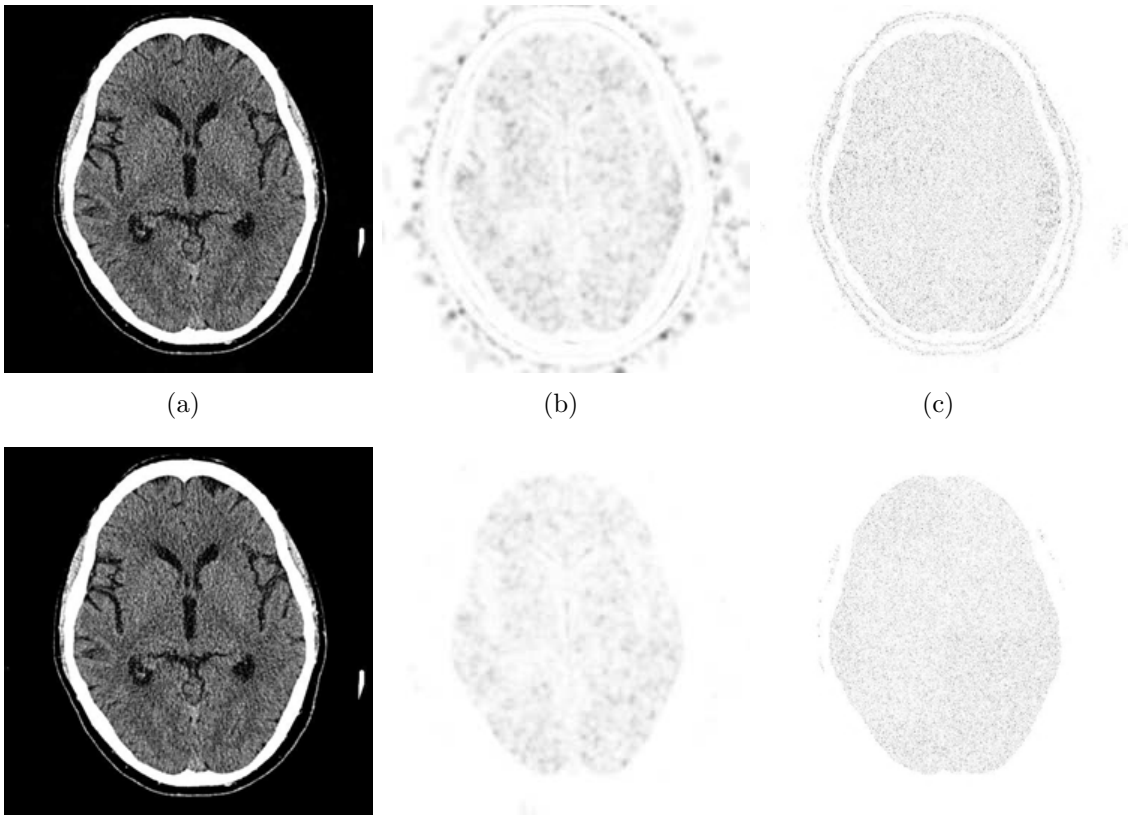

Figure 3.29: Quality maps of a JPEG2000 compressed neuro CT image, compression ratio: 12:1 (a) compressed image (no segmentation), (b) SSIM quality map, $SSIM = 0.9898$, no segmentation (c) MSE quality map, $PSNR = 34.2$, no segmentation, (d) compressed image (segmentation) (e) SSIM quality map, $SSIM = 0.9949$, segmentation, (f) MSE quality map, $PSNR = 36.5$, segmentation.

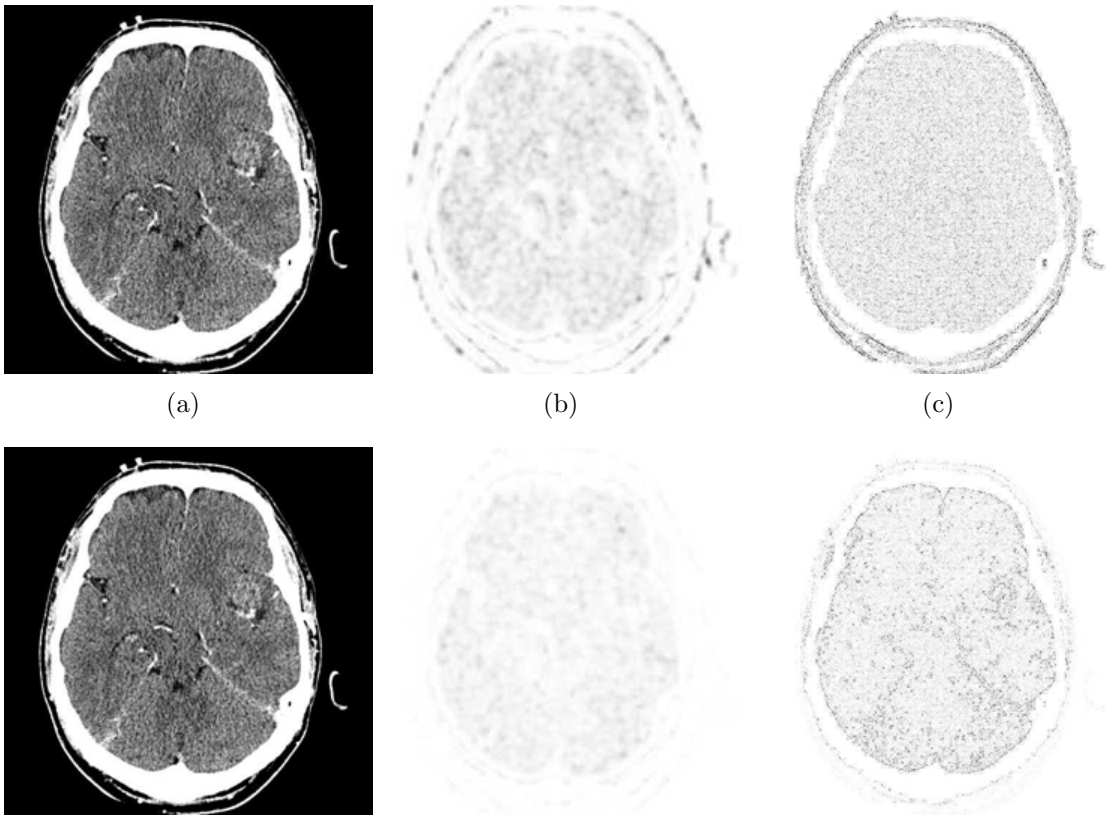

(a) (b) (c)

(d) (e) (f)

Figure 3.30: Quality maps of a JPEG compressed neuro CT 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.

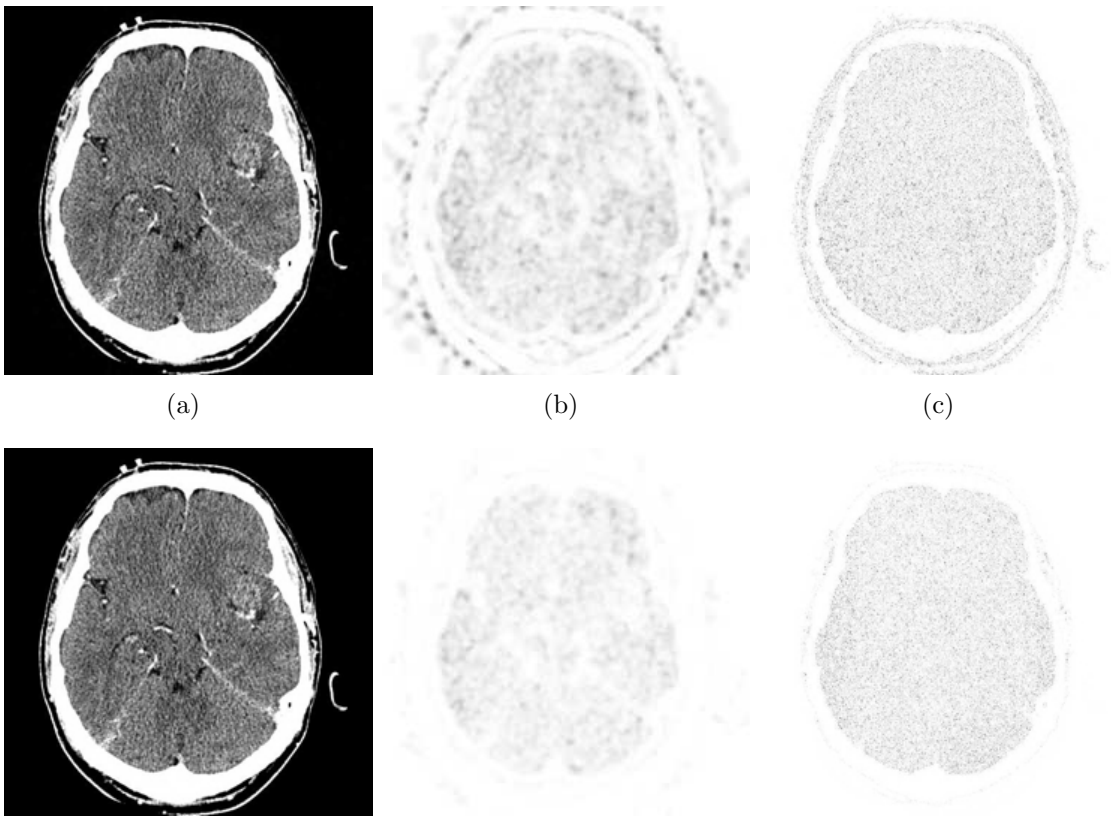

(a) (b) (c)

(d) (e) (f)

Figure 3.31: Quality maps of a JPEG2000 compressed neuro CT 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 summary, we have answered the question why JPEG performs better than JPEG2000 on some CT images. We conclude that compressing a neuro CT image requires a special treatment due to a specific distribution of its greyscale intensities (Figure 3.27). 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 neuro CT images. The result of 8- and 16-bit compression of neuro CT images has also been studied here.

The effect that sharp edge has on compression may be encountered not only in neuro CT 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.

The JPEG2000 algorithm shifts a significant number of bits in order to better encode the edges. As a result, fewer bits are left for the textural brain tissue regions, which are more relevant in medical diagnosis. Applying image segmentation before compression could improve the JPEG2000 compression results because it avoids shifting too many bits from the textural regions.

The procedure presented in this section is not a complete image compression algorithm that is competitive against the existing ones. Instead, the purpose of the segmentation based compression presented herein is to further validate the hypothesis why JPEG sometimes performs better than JPEG2000. The ideas presented in this section might serve as a demonstration of a potential direction that might be useful in the future development of better compression algorithms for medical images.

3.4.1 8-bit versus 16-bit 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 neuro CT 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 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.)

Neuro CT 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 3.32 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 3.2. 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! 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 3.3) and to the best of our knowledge there were no radiological studies on recommended compression ratios for 16-bit compression. Furthermore, Figure 3.32 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 neuro CT images.

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 3.2: Quality scores of a JPEG2000 compressed neuro CT 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 3.3: Quality scores of a JPEG2000 compressed neuro CT image (8-bit with compression ratio 12:1 and 16-bit with compression ratio: 20:1)

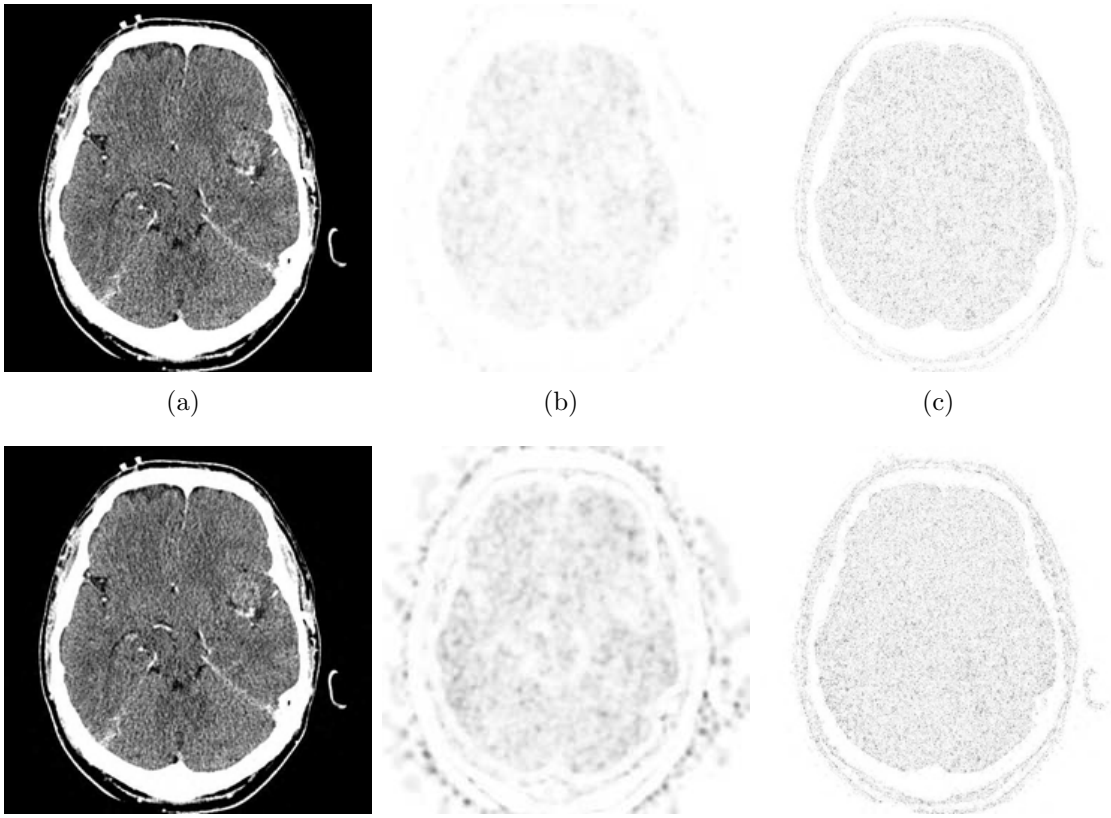

(a) (b) (c)

(d) (e) (f)

Figure 3.32: Quality maps of a JPEG2000 compressed neuro CT image, compression ratio: 12:1 (a) Neuro CT image after 16-bit compression, (b) SSIM quality map (16-bit compression), (c) MSE quality map (16-bit compression), (d) Neuro CT image after 8-bit compression, (e) SSIM quality map (8-bit compression), (f) MSE quality map (8-bit compression)

Chapter 4

Validation of Objective Quality Assessment Models for Compressed Medical Images

4.1 Subjective Experiment

In this study, we compare the performances of the SSIM index, MSE/PSNR, compression ratio CR and JPEG quality factor Q, based on experimental data collected in two experiments involving radiologists. This result was presented at the Medical Imaging, SPIE 2014 conference in San Diego, CA [41]. An ROC and Kolmogorov-Smirnov analysis indicates that compression ratio is not always a good indicator of visual quality. Moreover, SSIM demonstrates the best performance, i.e., it provides the closest match to the radiologists' assessments. We also show that a weighted Youden index [93] and curve fitting method can provide SSIM and MSE thresholds for acceptable compression ratios.

We examine whether compression ratio CR and MSE actually serve as reliable indicators of diagnostic quality. By this we mean "model the perception of trained radiologists in a satisfactory way". We also investigate the quality factor Q, the sole input parameter in the JPEG compression algorithm, since it has also been employed as a reference for quality assessment. The performances of the above indicators are compared to that of the SSIM index, based on experimental data collected in two experiments involving radiologists. The second goal of this work is to determine a method of acceptable compression thresholds using the data from our subjective experiments involving radiologists' image quality assessment.

4.2 Methods: Subjective Experiment Design

Two subjective experiments were designed in order to assess both the local and global prediction of the image quality assessments being examined. The first experiment, designed for a global analysis, employed ten CT slices - five neurological and five upper body images - extracted from volumes stored in the Cancer Imaging Archive [86]. These images were first windowed according to their default viewing parameters (window width and window centre) in order to reduce their bit-depth from 16 to 8 bits per pixel (bpp). Each of the resulting 512×512 pixel, 8 bpp, images were compressed at five different compression ratios using both the JPEG and JPEG2000 compression algorithms. (Since JPEG employs only the quality factor as input, it was adjusted in order to produce compression ratios as close as possible to those used for the JPEG2000 compression.) Preliminary visual observations were used to select the compression ratios employed in the experiment. The range of compression ratios was intended to represent a wide variety of visual qualities, from barely noticeable to fairly noticeable distortion.

An image viewer, was developed specifically for this study by Jiheng Wang (Ph.D. student in the Department of Electrical and Computer Engineering UW), in order to provide an easy-to-use graphical interface for the radiologists. The viewer displayed a compressed image beside its uncompressed counterpart without zoom. (No zooming was permitted in this experiment.) The ten compressed images were presented randomly and independently to each subject. During the course of the experiment, each compressed image was presented twice to each radiologist, but without the radiologists' knowledge. The subjects were not made aware of the compression ratios or quality factors of the compressed images. Two buttons were placed at the bottom of the user interface: acceptable and unacceptable. A confirmation was requested before passing to the next stimulus. In the second experiment, designed for a local analysis, six CT slices - four brain images and two body images - were compressed with JPEG at five different levels. Ten regions (35×35 pixel blocks) were manually selected from each of these images. The regions were chosen in order to obtain both a variety of features as well as local image quality. Ten pairs of buttons in the bottom of the interface allowed the subjects to rate each region as either acceptable or unacceptable. Once again, the experiment was repeated with the same medical images in a random order.

Two radiologist subjects (Dr. David Koff and Dr. Nadine Smolarski-Koff, McMaster University) participated in each of the experiments. (They were not specifically specialists for the sites and types of images presented.) The subjects were instructed to flag an image as unacceptable in the case of any noticeable distortion. The first experiment (global analysis) was held in an afternoon session while the second experiment (local/regional

analysis) took place during the morning. Each experiment lasted about one hour.

In the following sections, we outline the methods used to analyze the data obtained in these experiments. An analysis of the data is then presented in Section 4.4.

4.2.1 Global and Local experiments

Objective Quality Metrics

Since some regions of images are of much less interest than others in terms of diagnosis, an automatic segmentation of the images employed in this study (that is, the uncompressed and all compressed images) was performed in order to remove the background and bony anatomy. In the 16-to-8 bit tone-mapping of the images, background pixels always assumed a zero value, while those corresponding to bony anatomy generally assumed values near or equal to 255. The following simple segmentation operation was therefore sufficient: First threshold the images to separate the foreground from the background and then perform a fill operation in order to include the black pixels within the body part being imaged. A similar operation of thresholding followed by filling was performed in order to remove the skull. The removal of objects such as the couch or the scalp was also automatically performed by selecting only the largest region according to an 8-neighbour connectivity. The couch and the skull are successfully removed, while the bony anatomy inside the body is preserved. Some examples of segmentation are presented in Figure 4.1. All computations reported below were performed on thresholded images.

In this work, we employ a variation of the above SSIM index by considering only the second term, i.e., the structure term. The first term, i.e., the luminance term, will not be considered since it does not make an impact on the quality score. This is because the luminance does not change visibly for the images and compression ratios encountered in this study. In summary, we consider the following SSIM index, which involves only a single stability parameter,

$$SSIM(f,g) = \left(\frac{2\sigma_{fg} + C}{\sigma_f^2 + \sigma_g^2 + C} \right). \tag{4.1}$$

The reader will note that there is a further complication because of the "apples-oranges" nature of MSE and SSIM, i.e., "error" vs. "similarity": Recall that if f and g are "close", then $MSE(f,g)$ is near 0 but $SSIM(f,g)$ is near 1. In order to be able to compare the quality assessments of both indices more conveniently, we define the following quantity,

$$SMSE(f,g) = 1 - MSE(f,g)/D \tag{4.2}$$

where D is a constant. In this paper D was chosen to be 255.

Using this definition, we now have that if f and g are "close", then $SSIM(f,g)$ and $SMSE(f,g)$ are near 1. The reason for choosing $D = 255$ was to have SMSE close to SSIM for the same image quality. Other values of D and other scaling of MSE could also be used.

Finally, we mention another variation in the computation of the SSIM, that is, the computation of the local SSIM. The above discussion of SSIM involved a computation of the similarity of two entire images f and g, in other words, the global similarity of the images. It is often useful to measure the local similarity of images, i.e., the similarity of corresponding regions or pixel subblocks between images. For this reason, one can employ any or all of the above formulas to compute SSIM values between corresponding $m \times n$-pixel subblocks of two images.

One can proceed further and compute a SSIM quality map between two images f and g on a pixel-by-pixel basis as follows: At pixel (i,j) of each image, one constructs an $m \times n$-pixel neighbourhood, or "window", centered at (i,j) and then computes the SSIM index between the two neighbourhoods. This SSIM value is then assigned to location (i,j). The result is a SSIM quality map which reveals local image similarities/differences between images f and g. A total SSIM score may then computed by averaging over all the local SSIM values.

4.3 Classification Performance Metrics

The Receiver Operating Characteristic (ROC) curve [55] [83] is a common tool for visually assessing the performance of a classifier in medical decision making. ROC curves illustrate the trade-off of benefit (true positives, TP) versus cost (false positives, FP) as the discriminating threshold is varied. For convenience, the contingency table is shown in Figure 4.2. At this point, we must qualify that due to the nature of the problem we are investigating, our definitions of FP and TP differ from those normally applied for the purposes of medical diagnosis. In this study, we wish to examine how well different "image quality indicators", e.g., compression ratio, MSE, quality factor, SSIM, compare to the subjective assessments of image quality by radiologists. As such, we must assume that the "ground truth" for a particular experiment, i.e., whether or not a compressed image is acceptable

or unacceptable, is defined by the radiologist(s). From this ground truth, we measure the effectiveness of each image quality indicator in terms of FP, TP, etc... This leads to the following definitions of P, N, TP, FP, etc.:

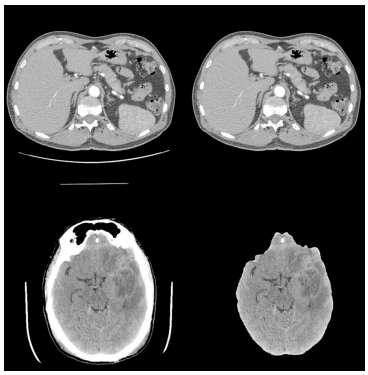

Figure 4.1: Examples of automatic image segmentation and object removal by thresholding and region growing.

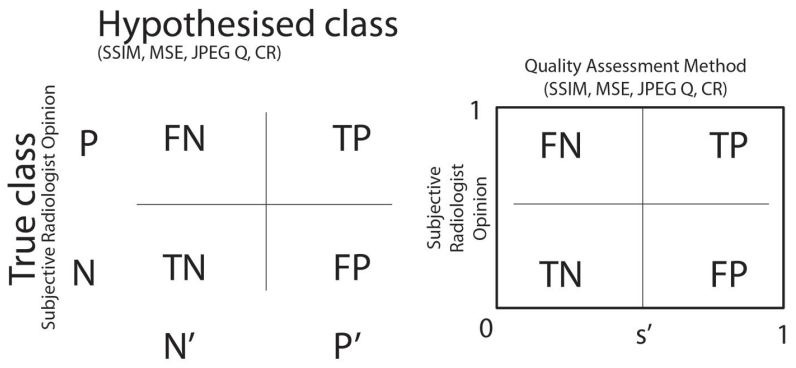

Figure 4.2: Contingency tables for Experiment 1.

1. P (or "1") $= FP + TP$ total positives (acceptable) and N (or "0")$= TN + FN$ total negatives (unacceptable): These refer to radiologists' subjective opinions, which represent the True class. On the other hand, P' and N' belong to the Hypothesis class which, in our experiment, corresponds to a given quality assessment method, i.e., SSIM, MSE, quality factor, and compression ratio. With reference to Figure 4.3:

2. TP (true positives): points that lie to the right of the threshold s' and have a value of 1, i.e., images that are acceptable to both radiologists and a given quality assessment method.

3. TN (true negatives): points that lie to the left of the threshold s' and have a value of 0, i.e., images that are unacceptable to both radiologists and a given quality assessment method.

4. FN (false negatives): points that lie to the left of the threshold s' and have a value of 1, i.e., images that are acceptable to radiologists but unacceptable to a given quality assessment method.

5. FP (false positives): points that lie to the right of the threshold s' and have a value of 0, i.e., images that are unacceptable to radiologists but acceptable to a given quality assessment method.

Figure 4.3 shows an illustration of how the counts for FP, FN, TP, TN are obtained for a given threshold s' corresponding to a discriminant (SSIM, MSE, compression ratio, quality factor).

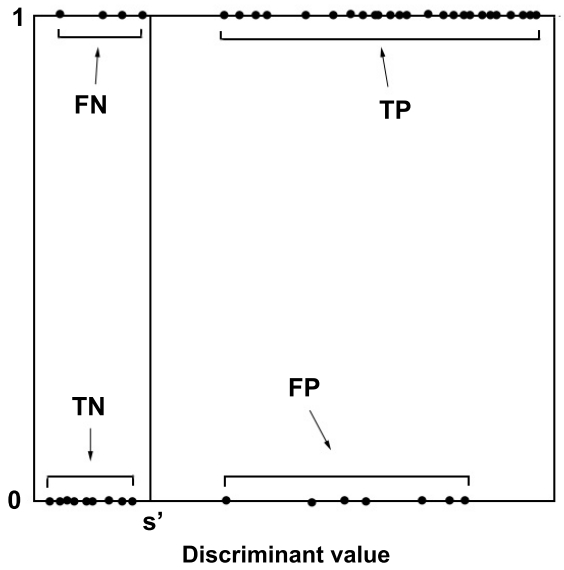

Figure 4.3: Computation of FP, FN, TP and TN for the threshold s'.

This, of course, leads to the question, "What constitutes acceptability/unacceptability to a given quality assessment method?" This is defined with respect to the discrimination threshold s' associated with the method, where $0 \leq s' \leq 1$. Given an objective image quality measure (e.g. SMSE, SSIM), we select a discriminating threshold s', $0 \leq s' \leq 1$ with reference to the figure on the right of 4.2. Using this threshold value s', FPR and TPR are computed. Each threshold value s' generates a point on the ROC curve which corresponds to the pair of values $(FPR, TPR) = (1-SP, SE)$, where SP denotes specificity and SE denotes sensitivity, i.e.,

FPR (false positive rate) $= FP/N = 1 - SP$ (specificity)

TPR (true positive rate) $= TP/P = SE$ (sensitivity)

FNR (false negative rage) $= FN/N = 1 - SE$

TNR (true negative rate) $= TN/N = SP$.

Figure 4.4 shows an example of an ROC curve. The point $(1, 1)$ on the ROC curve corresponds to $FN = 0$, i.e., no false negatives, and $TN = 0$, i.e., and no true negatives.

The opposite scenario occurs at the point (0,0) which corresponds to $FP = 0$, i.e., no false positives, and $TP = 0$, i.e., no true positives. The point $(0, 1)$ corresponds to perfect classification, since $FP = 0$, i.e., no false positives, and $FN = 0$, i.e., no false negatives.

Once again, we mention that the definitions and labels in the contingency table associated with our experiments differ from those associated with a general detection/diagnosis experiment. Here, by false negative we mean that an image with a low objective quality score (hence unacceptable according to the quality method) has actually received a positive subjective score (acceptable, according to the radiologists).

In ROC analysis, as is well known, the performance of a discriminant (here, the SSIM, MSE, JPEG quality factor and compression ratio quality assessment methods) is often characterized by (1) the area under the curve (AUC) and/or (2) the Kolmogorov-Smirnov (KS) statistic or test.

AUC method: The AUC can be computed by numerical integration using the trapezoidal rule. Larger AUC values correspond to better performance of the classifier. It is possible that two ROC curves cross. In this special situation a given classifier might demonstrate better performance for some threshold values whereas another classifier behaves better for other threshold values. In this case, a single AUC may not be the best predictor of the performance of a classifier.

Kolmogorov-Smirnov (KS) test: [18] Given two cumulative probability distributions $P1(x)$ and $P2(x)$, their KS statistic is defined as follows,

$$KS(P1, P2) = \sup_x |P1(x) - P2(x)|. \tag{4.3}$$

In our study, $P1$ and $P2$ are the cumulative distributions of "acceptable" (1's) and "unacceptable" (0's) radiologists responses (respectively) corresponding to a given objective quality measure (i.e. SSIM, SMSE, compression ratio, JPEG quality factor). The larger the difference between the two distributions, the better the performance of a given model. A generic situation is illustrated in Figure 4.5. Examples of cumulative probability distributions with respect to a discriminant are shown in Figures 4.6 and 4.6.

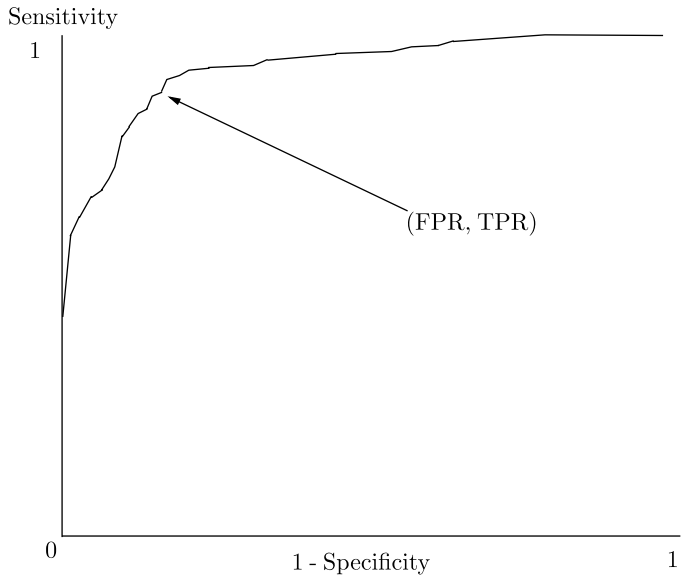

Figure 4.4: Example of an ROC curve.

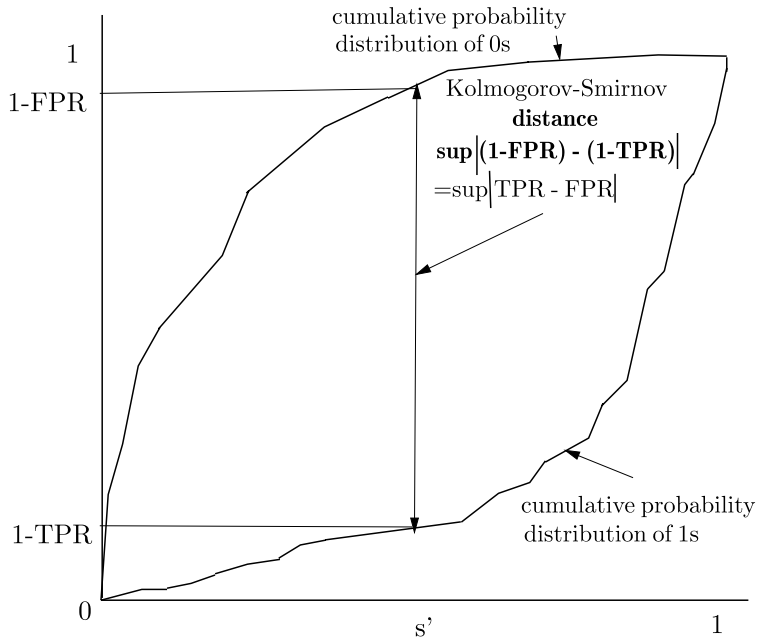

Figure 4.5: Kolmogorov-Smirnov distance.

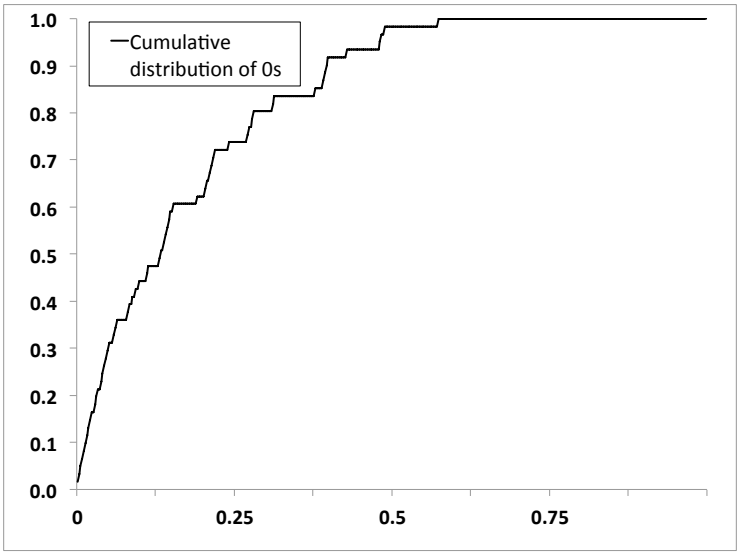

Figure 4.6: Example of a cumulative distribution of 0s with respect to a discriminant.

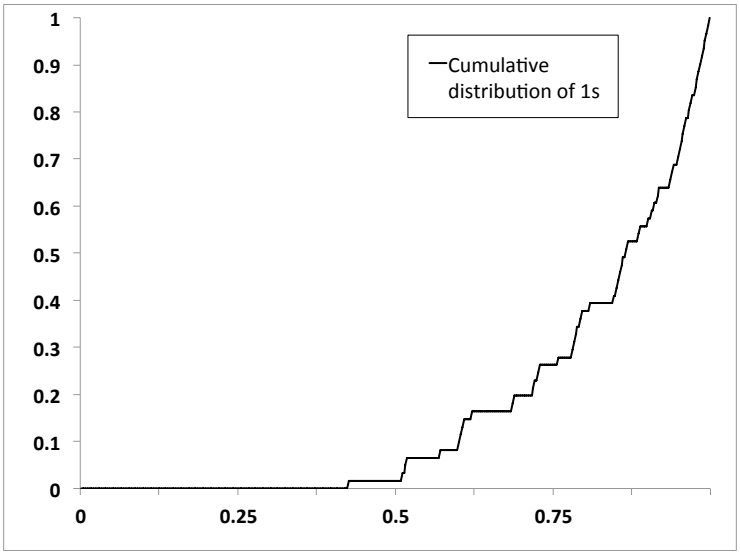

Figure 4.7: Example of a cumulative distribution of 1s with respect to a discriminant.

With reference to Figure 4.2, for a given threshold s' in $[0,1]$, we have the following relations,

Cumulative Probability Distribution of $0's = TN/(TN+FP) = 1-FPR$.

Cumulative Probability Distribution of $1's = FN/(FN+TP) = 1-TPR$.

Thus, the KS statistic translates to

$$KS = \sup_x |TPR(x) - FPR(x)|. \tag{4.4}$$

Youden index: Let us now recall the idea of the Youden index [93]. With reference to Figure 4.5, in which the cumulative probability distribution curves, $1-FPR$ and $1-TPR$, are plotted as a function of the threshold s, the Youden index $Y(s)$ associated with the threshold value s is simply the difference,

$$Y(s) = TPR(s) - FPR(s). \tag{4.5}$$

Now suppose that the maximum value of Y(s), the so-called maximum Youden index, occurs at the threshold value $s = s'$. This implies that the points $TPR(s')$ and $FPR(s')$ lie farthest away from each other which, in turn, implies that the maximum Youden index is the Kolmogorov-Smirnov (KS) value associated with the cumulative probability distributions of 0's and 1's discussed in the previous section.

There is also a connection between the maximum Youden index and ROC curves. Since the values $FPR(s')$ and $TPR(s')$ lie farthest from each other, the corresponding point $(FPR(s'), TPR(s'))$ on the ROC curve (which contains all points $(FPR(s), TPR(s))$ lies farthest away from the diagonal line joining $(0,0)$ and $(1,1)$. The threshold value $s = s'$ for which the Youden index is maximized is considered to be the optimal threshold value.

It is possible to define a weighted version of the Youden index in the following way. First of all, note that the Youden index $Y(s)$ introduced above may be expressed as follows,

$$Y(s) = TPR(s) - FPR(s) = SP(s) + SE(s) - 1, \tag{4.6}$$

where $SE(s)$ and $SP(s)$ denote, respectively, the sensitivity and specificity associated with the threshold value s. The Youden index $Y(s)$ may be viewed as employing an equal weighting of false positives (FP) and false negatives (FN). It may be desirable to employ a non-equal weighting of these statistics in order to alter their relative importance. This

can be accomplished by means of a parameter $\lambda, 0 \leq \lambda \leq 1$ so that the weighted Youden index is given by

$$WY(s) = \lambda SP(s) + (1-\lambda)SE(s) - 1 = -(\lambda FPR(s) + (1-\lambda)FNR(s)). \tag{4.7}$$

4.3.1 Threshold selection

The selection of a threshold is accomplished in two ways: using the Weighted Youden Index and by means of a curve-fitting model.

In order to take into account the variability in the subjective quality assessment of compressed medical images, a logistic cumulative probability distribution is assumed to model the decision of a radiologist to either accept or not accept an image at a given objective score. A robust curve-fitting is performed on plot of the average subjective score over all the radiologists and all the repetitions in function of the objective score. The threshold is selected so that the cumulative probability distribution model represents the desired level of confidence that the quality of the compressed image is diagnostically acceptable. For example, if one requires a 99% confidence, the recommended threshold has to be selected at the value for which the fitted logistic curve is at 0.99.

The recommended threshold for compression was obtained by means of the logistic curve-fitting model and the Youden Index with the use of k-fold cross validation procedure as outlined in [33] on pages 241-242: "K-fold cross- validation uses part of the available data to fit the model, and a different part to test it. We split the data into K roughly equal-sized parts. ... For the kth part, we fit the model to the other $K-1$ parts of the data, and calculate the prediction error of the fitted model when predicting the kth part of the data. We do this for $k = 1, 2, ..., K$ and combine the K estimates of prediction error."

Steps of our 10-fold cross-validation procedure based on the above description using the curve-fitting model:

1. The data is randomly split (random permutation) into K sets. K-1 sets are used for training, kth set is used for testing. The confidence level is chosen by the user.

2. Training step: The threshold of SSIM (and SMSE) is computed on the training set as follows:

- Given x_1 the predicted value (SSIM or SMSE) and y_1 the average subjective score, we determine the parameters a and b of the logistic function $y = 1/(1 + \exp(-a * x + b))$ that produce the least weighted square error, with weighting according to the Bisquare Method (this is a standard robust fitting technique). Alternatively, one could minimize the Absolute Error or the Square Error (standard least square).
- Once the coefficients a and b are found, the recommended threshold for compression is computed from the inverse logistic function $x = (\log(y/(1-y)) + b)/a$ where $y = 1 - \alpha$ is the desired significance level.

3. Testing step: Prediction accuracy (number of FPs and FNs) is computed on the testing set.

4. Steps (2) and (3) are repeated for $k = 1, 2, ...K$ times ($K = 10$).

5. Optimal threshold and prediction accuracy are obtained by averaging the results over all repetitions of the procedure.

Steps of our 10-fold cross-validation procedure based on the above description using the Youden index:

1. The data is randomly split (random permutation) into K sets. K-1 sets are used for training, kth set is used for testing. User defined parameter λ is selected.

2. Training step: Threshold of SSIM (and SMSE) corresponding to the max Youden Index (calculated using the user defined parameter λ) is determined on the training set.

3. Testing step: Prediction accuracy (number of FPs and Fns) is computed on the testing set.

4. Steps (2) and (3) are repeated for $k = 1, 2, ...K$ times ($K = 10$).

5. Optimal threshold and prediction accuracy are computed by averaging the results over all repetitions of the procedure.

4.4 Results

4.4.1 First Experiment: global quality

The first part of our analysis involves all data points accumulated in Experiment 1. These data points include the two image types (brain CT and body CT) and the two compression methods (JPEG and JPEG2000). Figure 4.8 shows the ROC curves that correspond to the two quality measures SSIM and SMSE. Figures 4.9 and 4.10 show the ROC curves corresponding to JPEG and JPEG2000 compressed images and the four quality measures SSIM, SMSE, JPEG quality factor Q and compression ratio CR. We observe that the ROC curve corresponding to CR demonstrates the worst performance, i.e., the lowest area under curve (AUC). Figures 4.11 and 4.12 show ROC curves associated with each of the two image types, i.e., brain CT and body CT. Such an analysis in terms of image types is particularly important since these two classes of images possess different characteristics (e.g. texture) which may yield different compression artifacts. Note that the ROC curves with the largest AUC correspond to the SSIM index quality measure.

In Figures 4.13, 4.14, 4.15 and 4.16 are presented the individual ROC curves corresponding to JPEG and JPEG2000 compression methods, the four quality measures and the two image types. Once again, the ROC curves associated with the SSIM index quality measure yield the largest AUCs. This suggests that of the four quality measures, SSIM performs the best in modeling the radiologists' subjective assessments of compressed images when the AUC is used as a performance indicator.

From Figures 4.13, 4.15 and 4.18, the AUCs associated with JPEG-compressed brain images are seen to be significantly lower than those associated with JPEG-compressed body images. This indicates that assessments of compressed brain CT images agree the least with the subjective assessments of radiologists. A closer examination of the data provides an explanation of this disparity. The radiologists perceived JPEG-compressed brain images as almost always acceptable, even in the cases when the quality of these images was deemed unsatisfactory in terms of SSIM or SMSE. As a result, the ratios FP/TN, which are the false positive rates (FPR) plotted along the horizontal axis of ROC space, assume values of only 0 or 1. This explains why the two JPEG-compressed ROC curves are not only linear but almost horizontal.

Furthermore, Figures 4.13 and 4.15 show that in the case of compressed brain images, JPEG2000 demonstrates better agreement with the radiologists' opinions than JPEG. However, at the same compression ratios, more JPEG images were judged as acceptable by the radiologists. How do we explain this oddity? It is generally accepted that JPEG2000

"performs better" than JPEG for most classes of images - in other words, at a given compression ratio, the error (both visual as well as quantitative) between uncompressed and compressed images is less for JPEG2000 than JPEG. That being said, Koff *et al.* noted that the opposite was often observed in the case of brain images, i.e., JPEG performed better than JPEG2000. This anomaly is due to the bony skull in brain images [43]. The sharp edges between the skull bone and neighbouring regions (both background and interior regions) are difficult to approximate by any method. This is particularly the case for JPEG2000 since a larger number of wavelet coefficients are required to approximate these edges. JPEG can better accommodate these strong edges because of the 8×8 pixel block structure employed in the algorithm. As a result, the visual quality in the interior regions, which contain the most diagnostic information, can be lower for JPEG2000 than for JPEG.

We now summarize the results of applying the Kolmogorov-Smirnov (KS) test to the collected data. Figures 4.19 and 4.20 present the cumulative distributions of subjective radiologist scores (0's and 1's) corresponding to, respectively, SSIM and SMSE quality measures for all data points. Figure 4.21 is a plot of the subjective radiologist scores for JPEG-compressed images using the JPEG quality factor Q as the quality measure. Finally, the distributions of 0's and 1's corresponding to compression ratio are shown in Figure 4.22. As expected, the KS statistic, i.e. the separation between the two cumulative distributions is largest for SSIM (81%), intermediate for JPEG quality factor Q (78%) and smallest for compression ratio CR (60%).

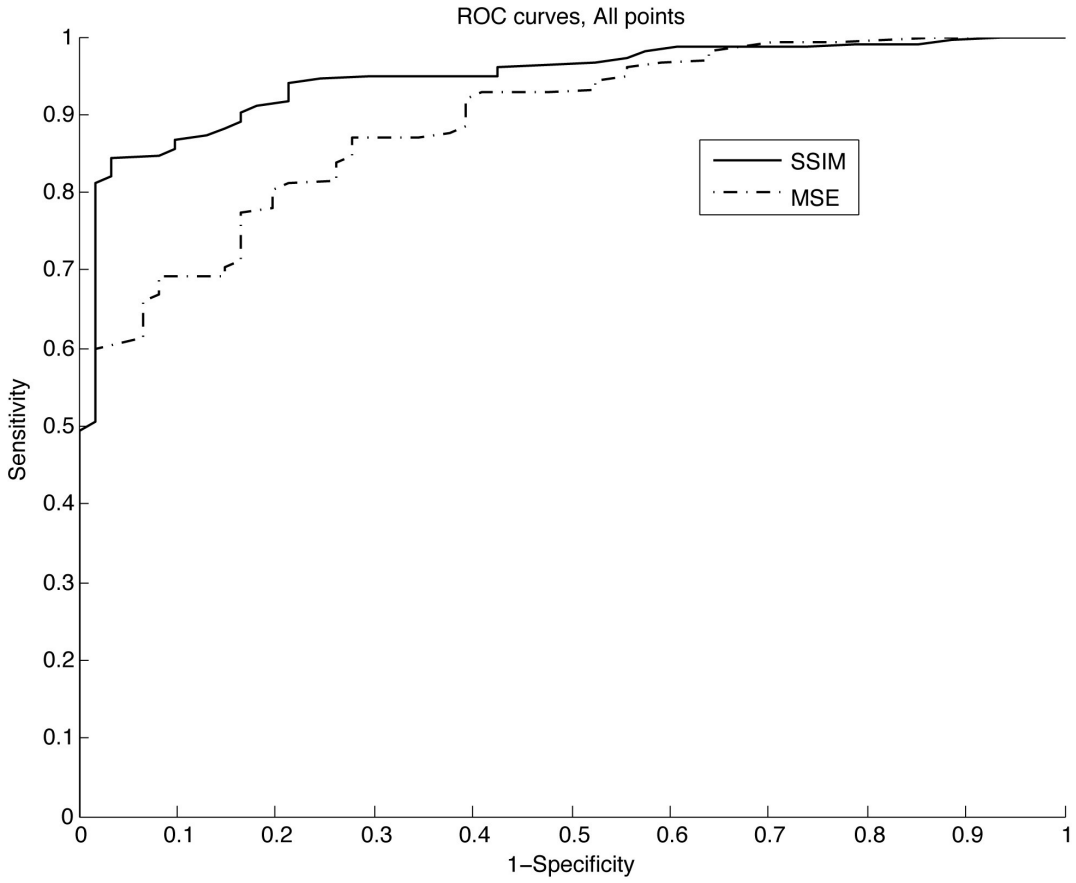

Figure 4.8: ROC curves corresponding to all data points. AUC SSIM = 0.9471 AUC SMSE = 0.8900.

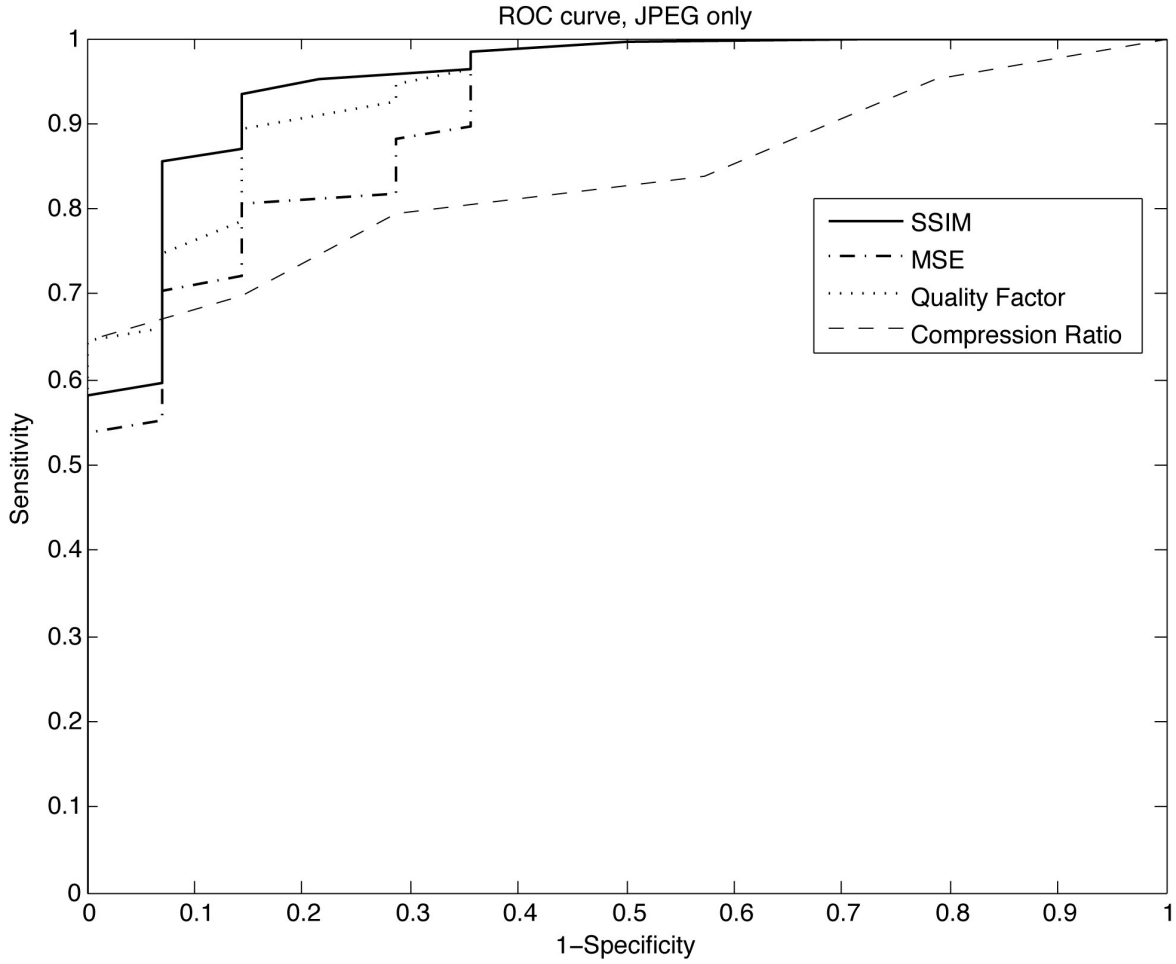

Figure 4.9: ROC curves corresponding to JPEG images and SSIM, SMSE, quality factor, compression ratio. AUC JPEG SSIM = 0.9485 AUC JPEG SMSE = 0.9101 AUC JPEG quality factor = 0.9401 AUC JPEG compression ratio =0.8372.

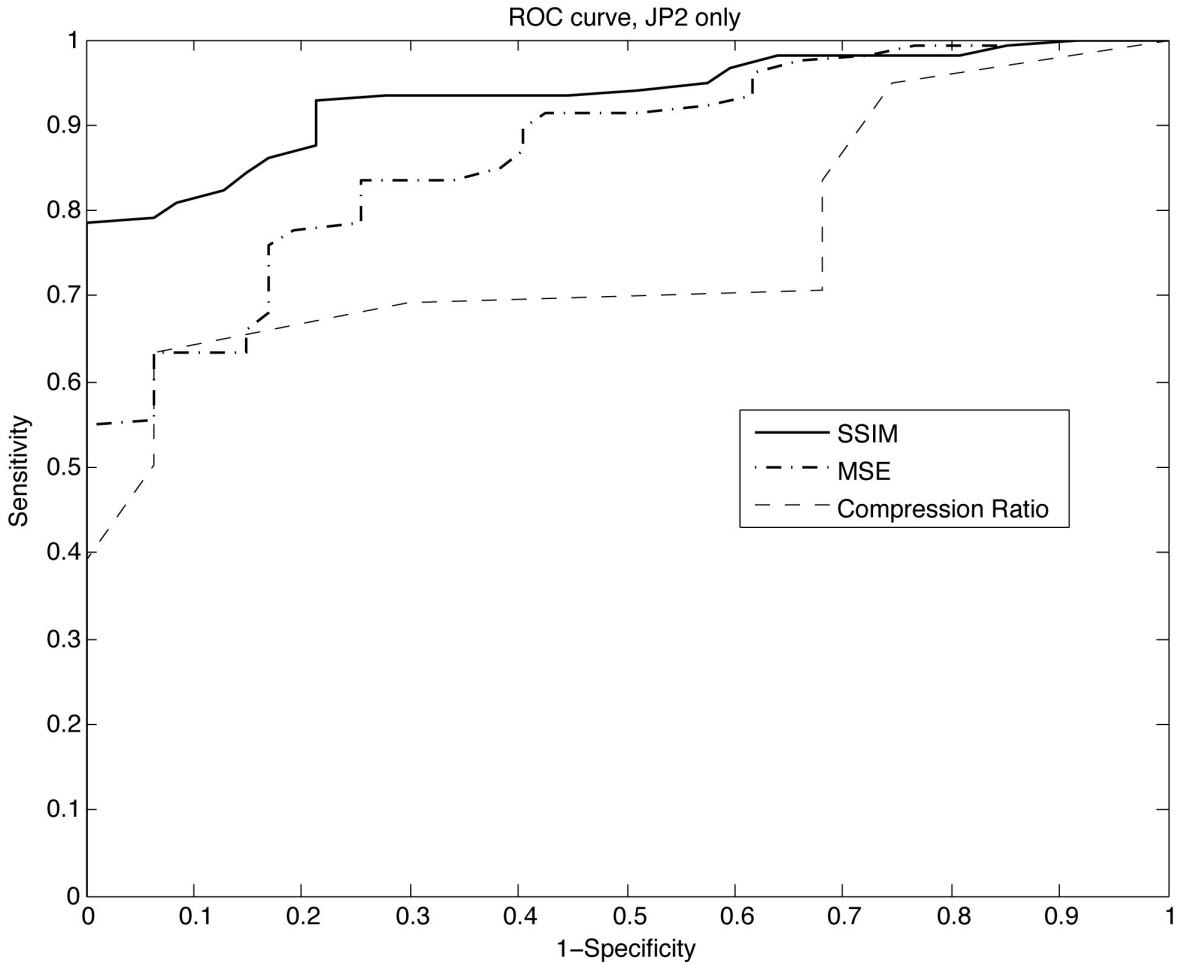

Figure 4.10: ROC curves corresponding to JPEG2000 images and SSIM, SMSE, compression ratio. AUC JPEG2000 SSIM = 0.9330 AUC JPEG2000 SMSE = 0.8691 AUC JPEG2000 compression ratio =0.7573.

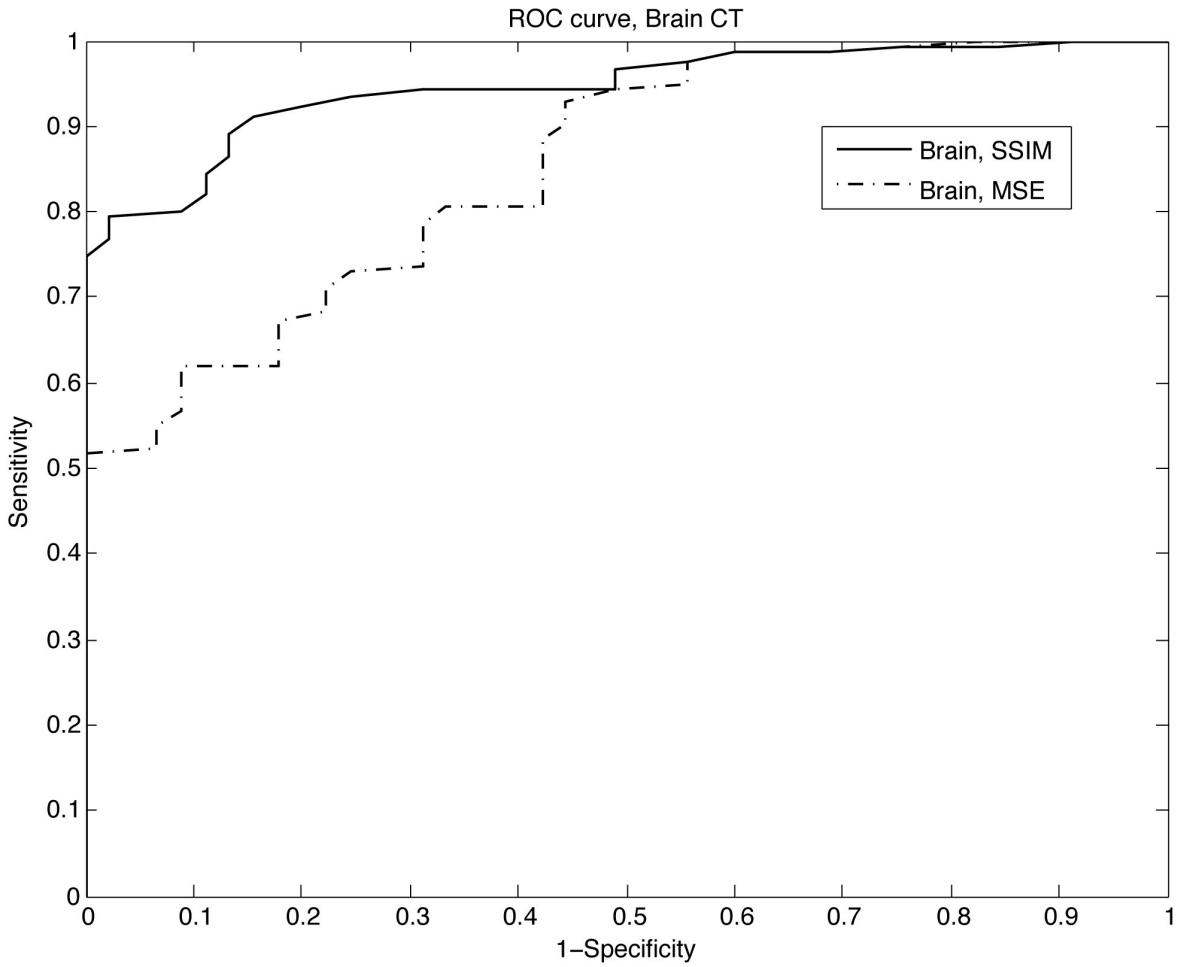

Figure 4.11: ROC curves corresponding to brain CT. The area under the curve for each of the types is: Brain SSIM: AUC = 0.9447 Brain SMSE: AUC = 0.8524.

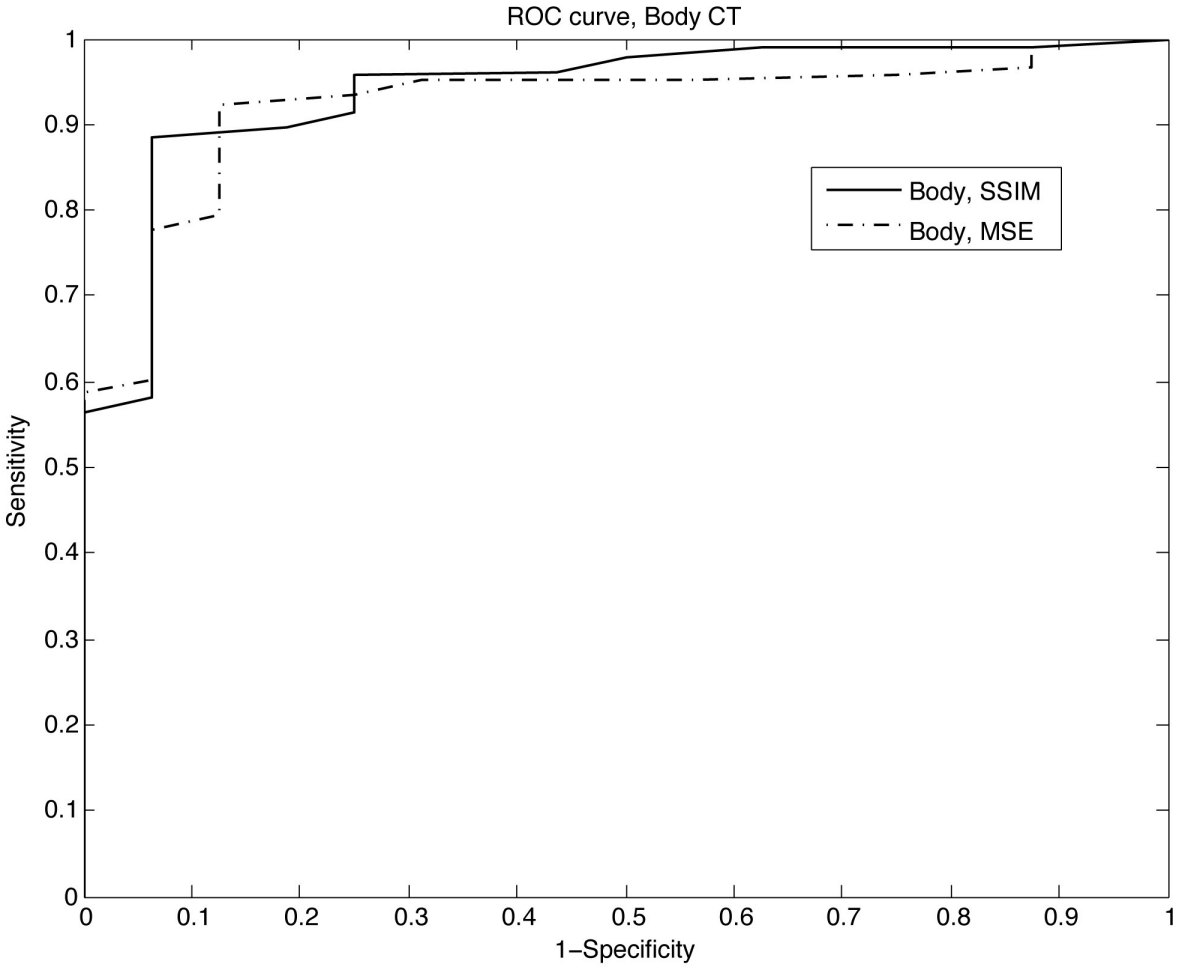

Figure 4.12: ROC curves corresponding to body CT. The area under the curve for each of the types is: Body SSIM: AUC = 0.9389 Body SMSE: AUC = 0.9226.

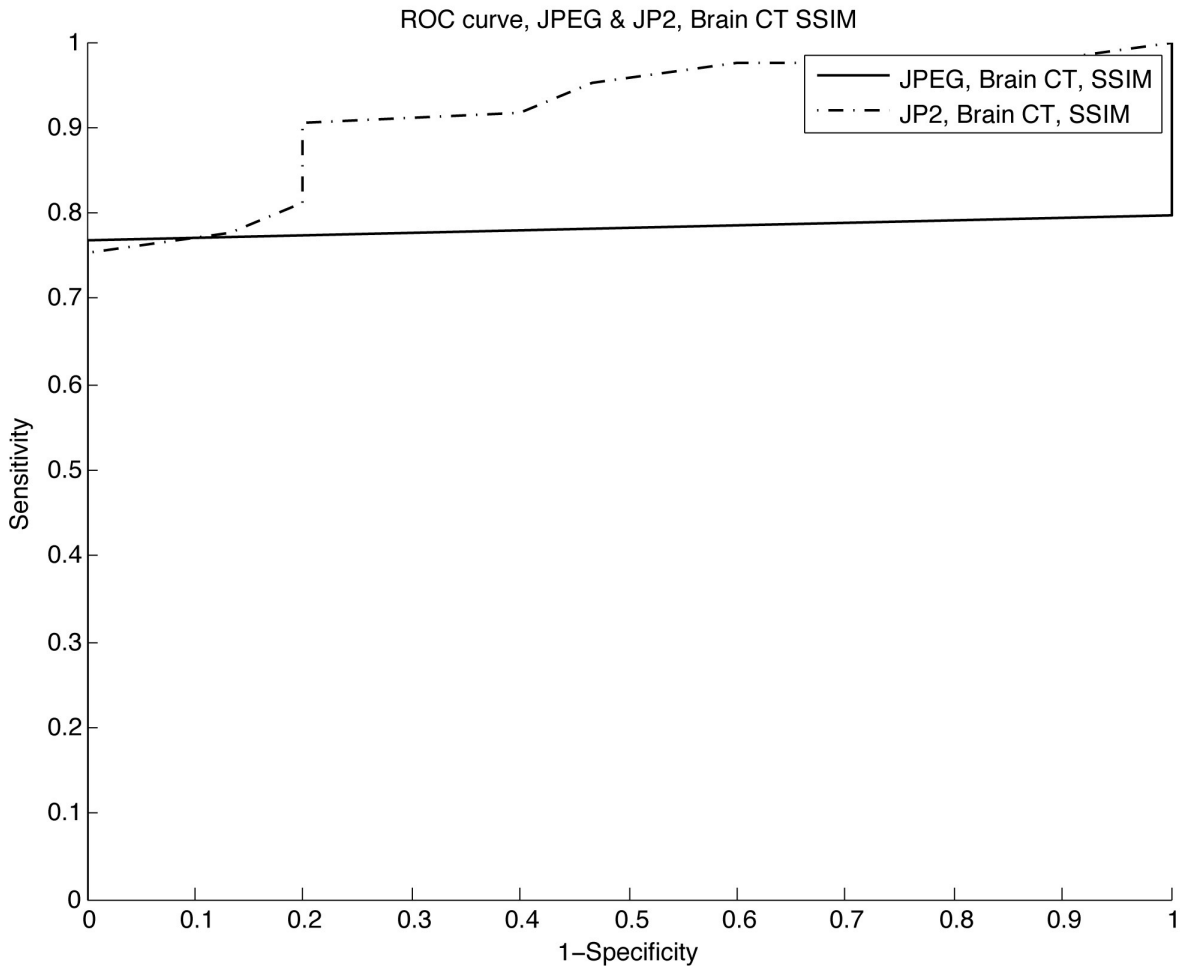

Figure 4.13: ROC curves corresponding to JPEG and JPEG2000 compressed brain CT images and SSIM. JPEG body SSIM: AUC = 0.7828 JP2 body SSIM: AUC = 0.9204.

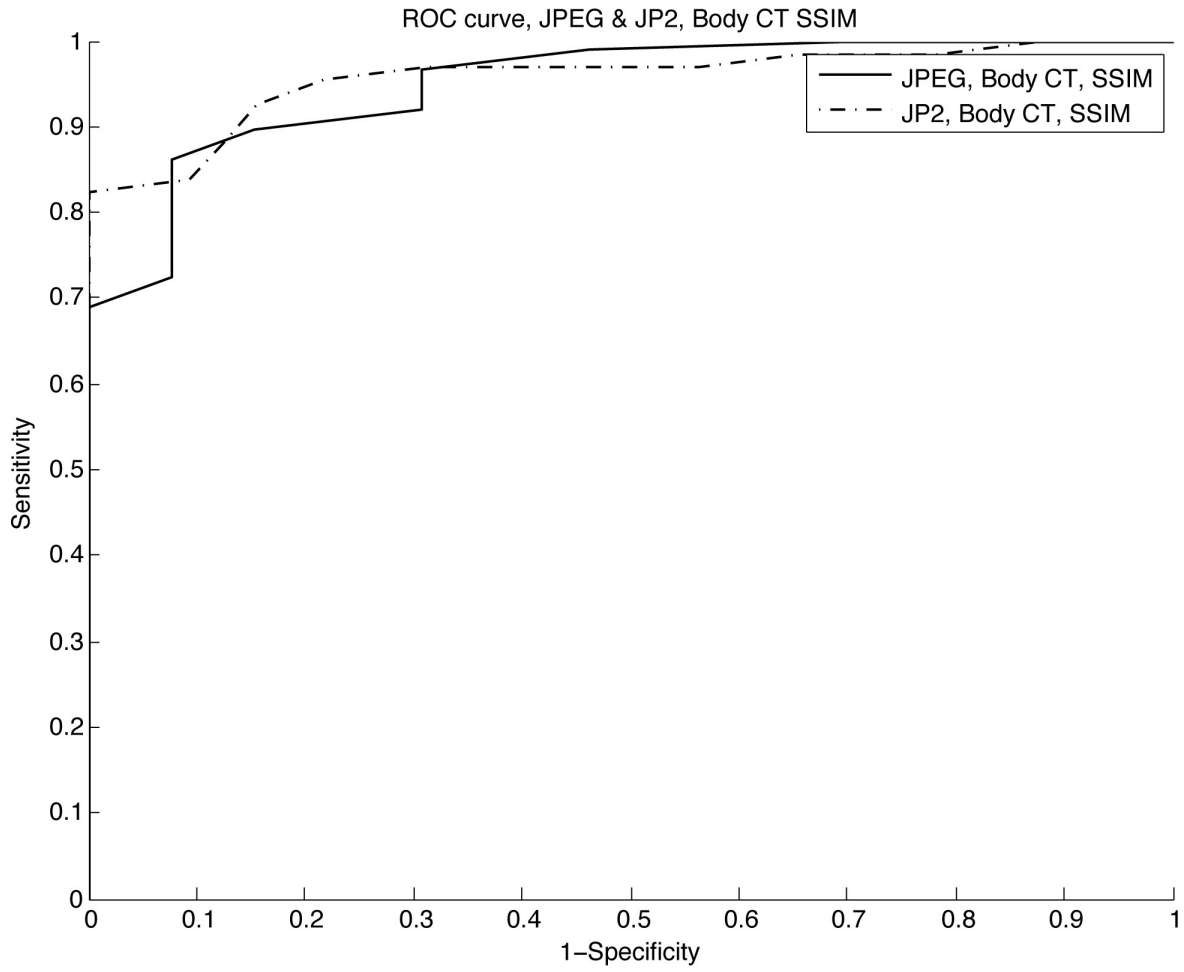

Figure 4.14: ROC curves corresponding to JPEG and JPEG2000 compressed body CT images and SSIM. JPEG brain SSIM: AUC = 0.9492 JP2 brain SSIM: AUC = 0.9577.

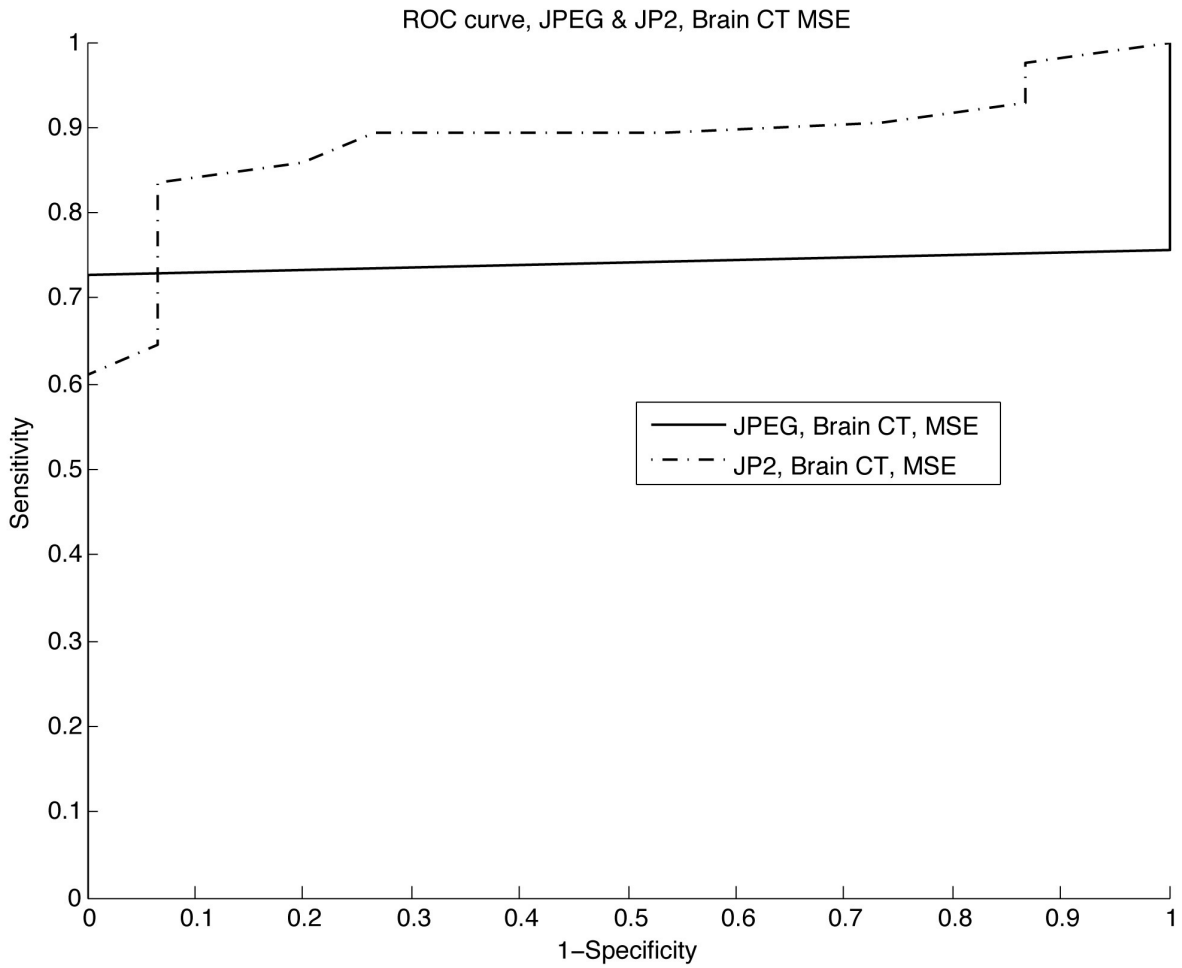

Figure 4.15: ROC curves corresponding to JPEG and JPEG2000 compressed brain CT images and SMSE. JPEG Body SMSE: AUC = 0.7424 JP2 Body SMSE: AUC = 0.8859.

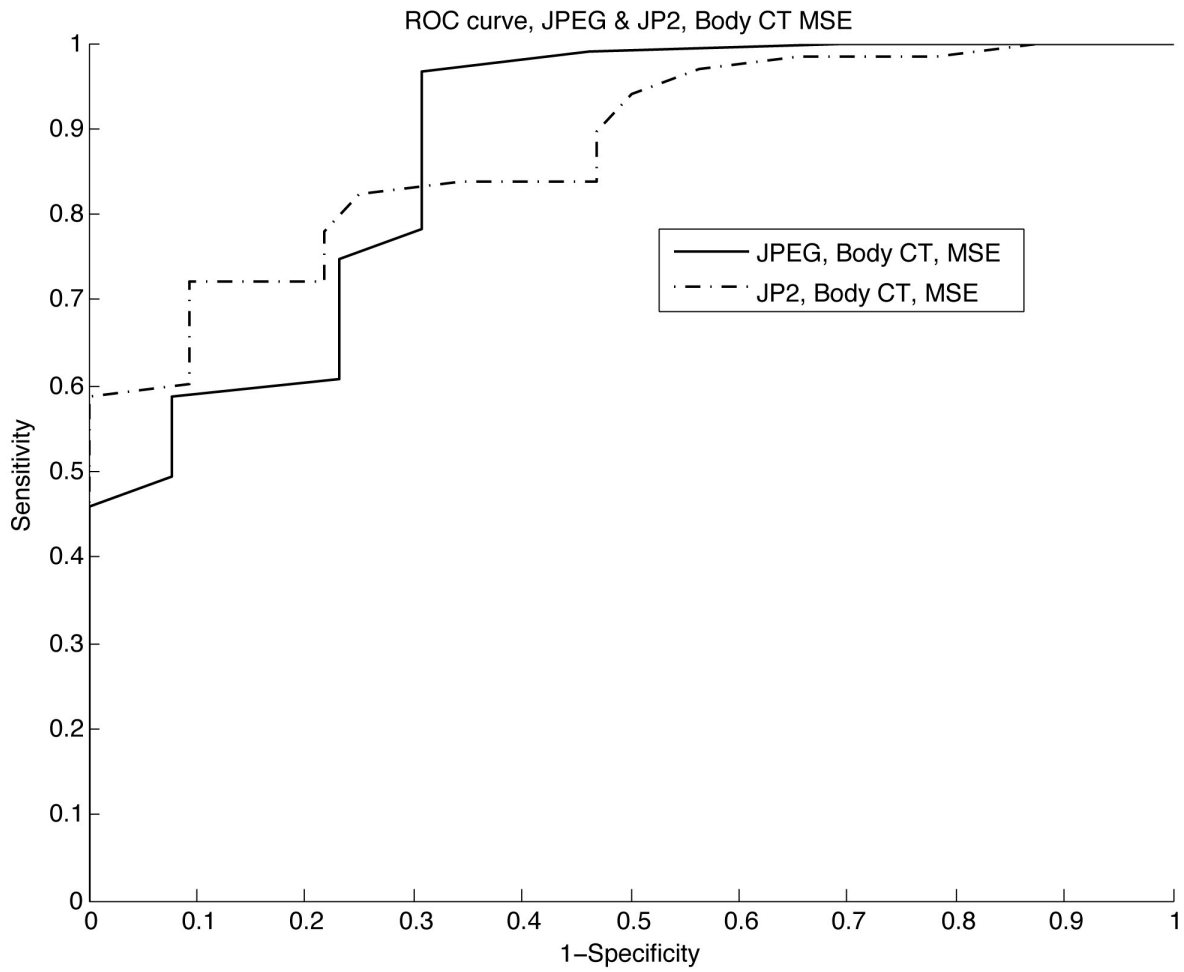

Figure 4.16: ROC curves corresponding to JPEG and JPEG2000 compressed body JPEG brain SMSE: AUC = 0.8749 JP2 brain SMSE: AUC = 0.8750.

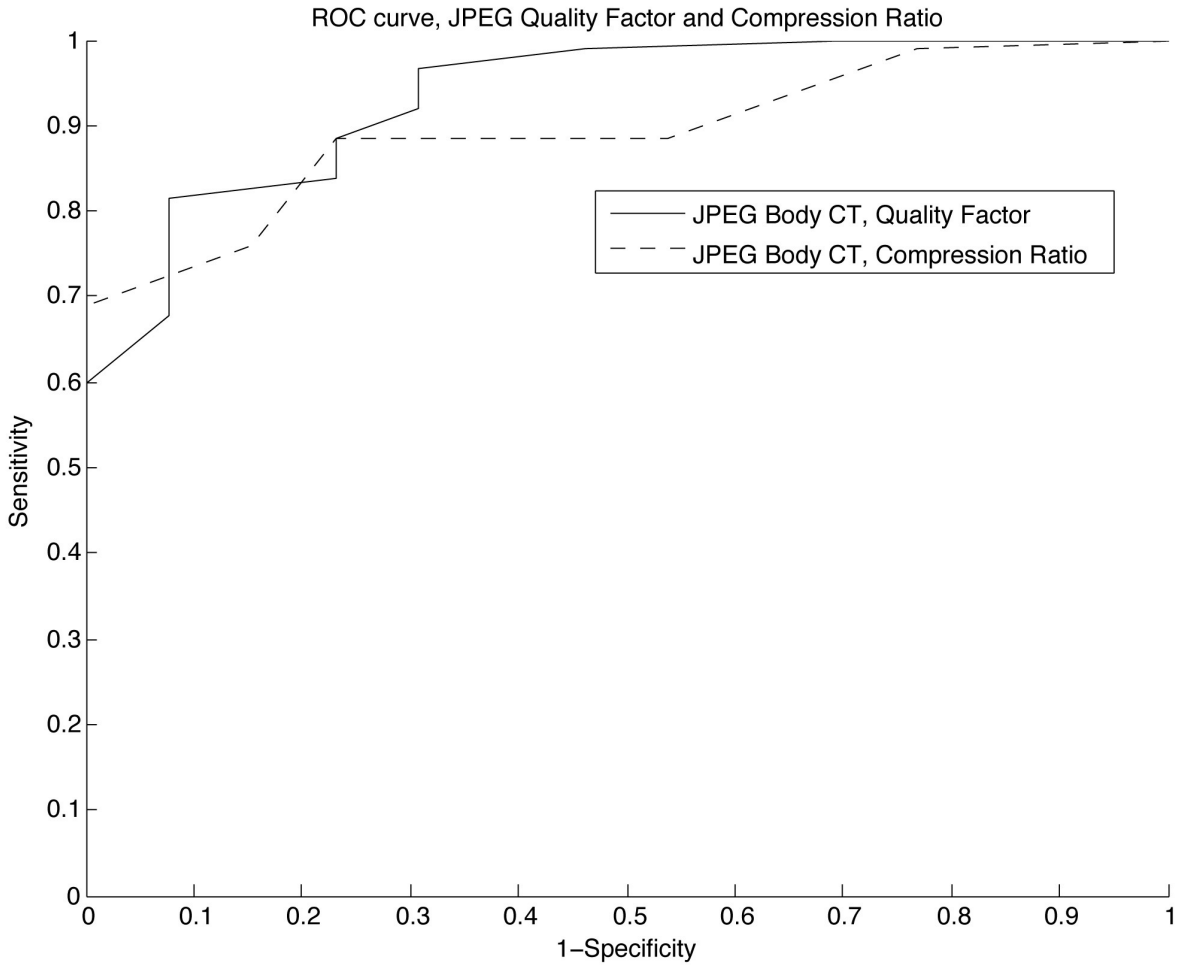

Figure 4.17: ROC curves corresponding to JPEG compressed body images with respect to quality factor and compression ratio. JPEG body, quality factor: AUC = 0.9332 JPEG body, compression ratio AUC = 0.8926.

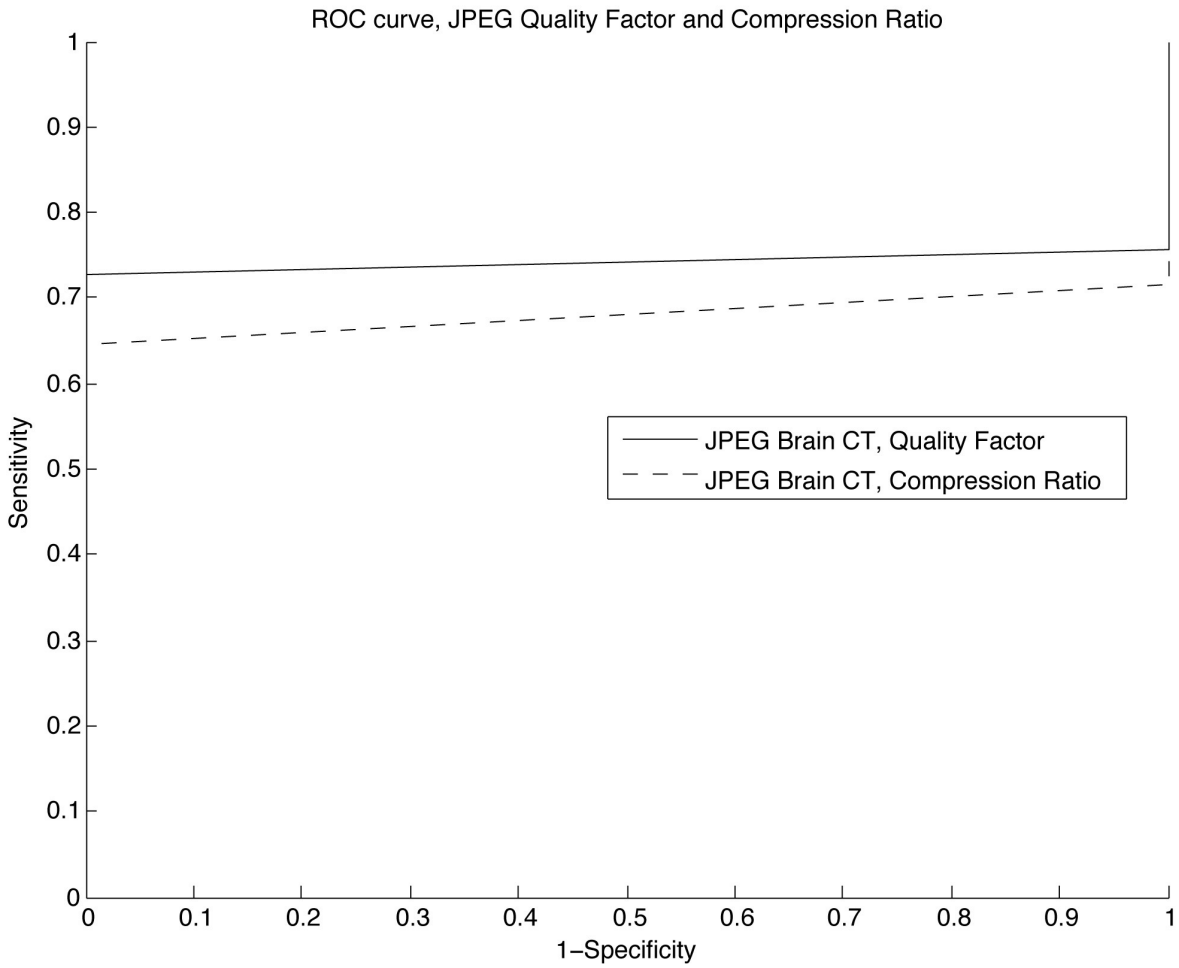

Figure 4.18: ROC curves corresponding to JPEG compressed brain images with respect to quality factor and compression ratio. JPEG body, quality factor: AUC = 0.7424 JPEG body, compression ratio AUC = 0.6818.

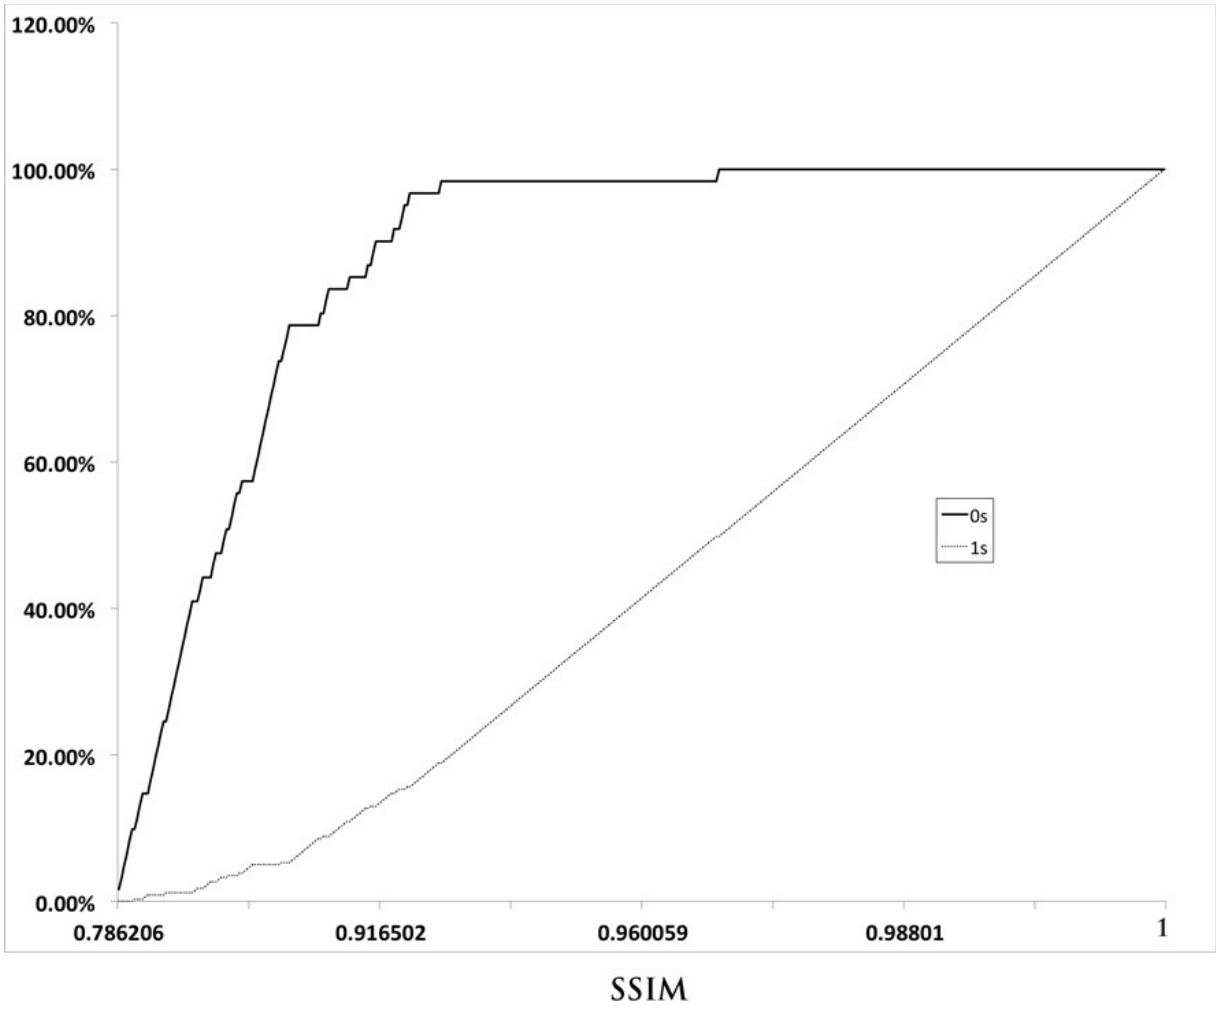

Figure 4.19: Cumulative distributions of subjective radiologist scores corresponding to SSIM. K-S = 81.09%.

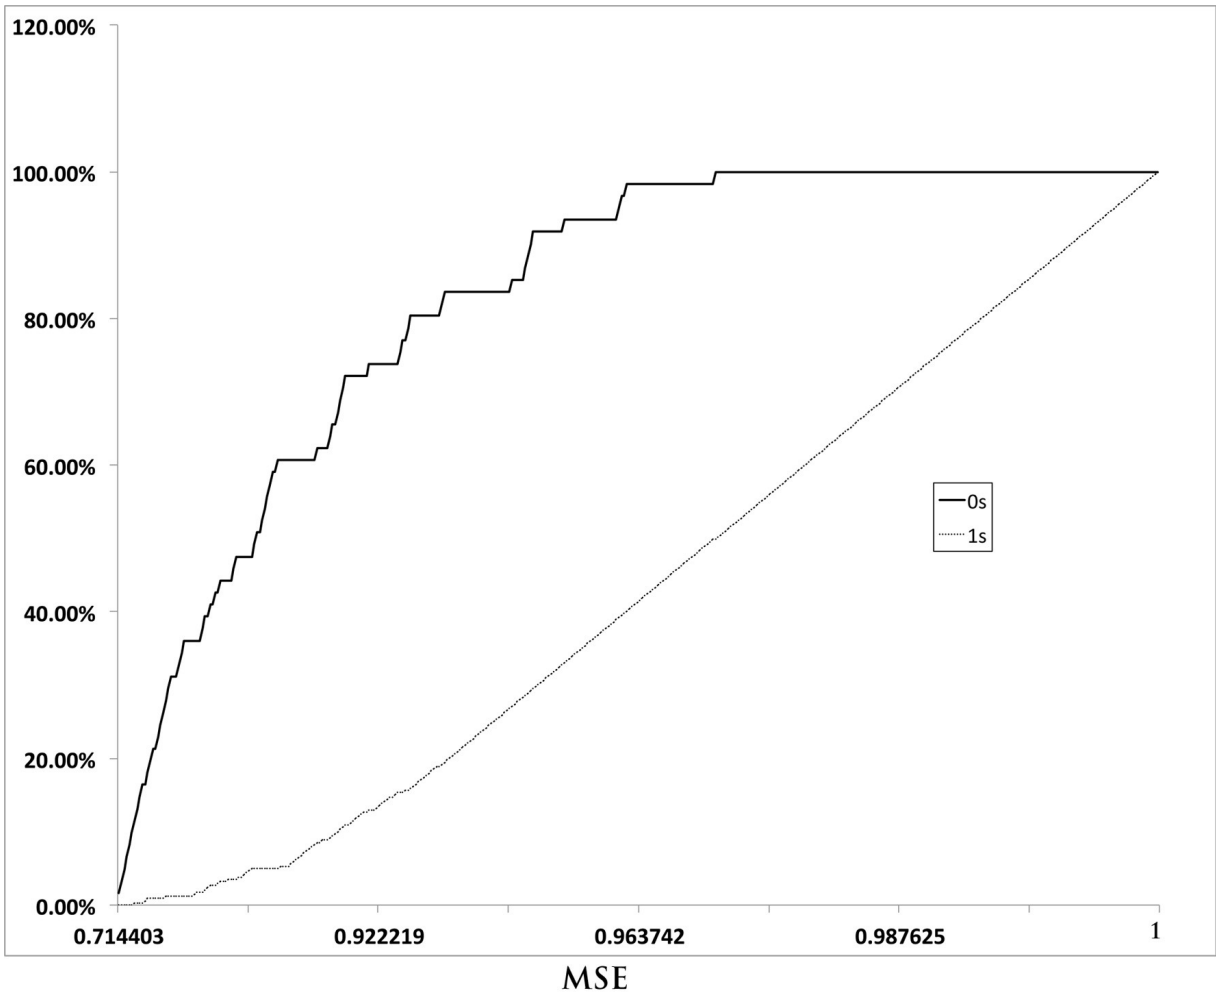

Figure 4.20: Cumulative distributions of subjective radiologist scores corresponding to SMSE. K-S = 64.40%.

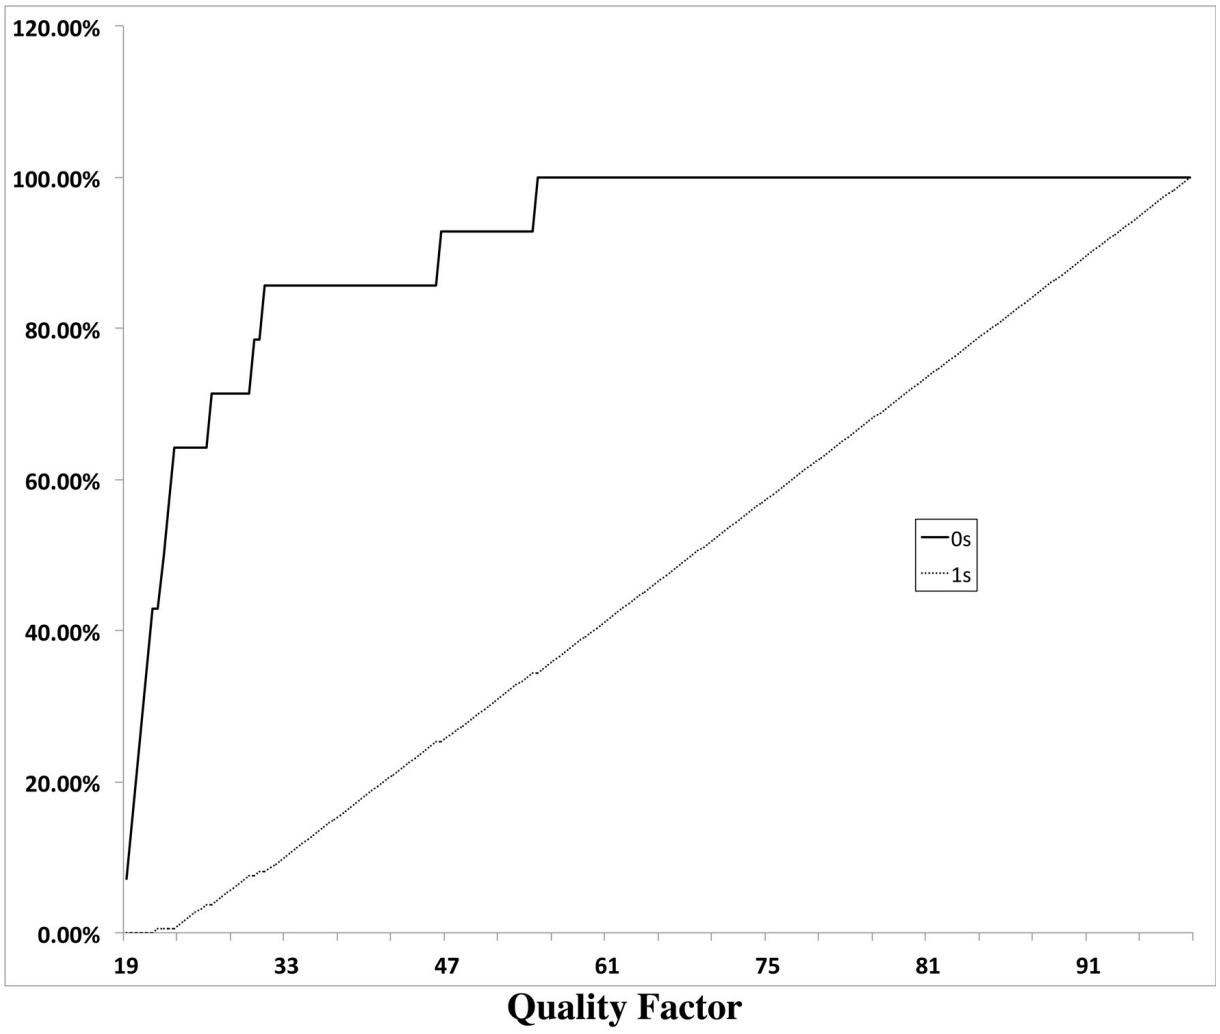

Figure 4.21: Cumulative distributions of subjective radiologist scores corresponding to JPEG quality factor. K-S = 77.65%.

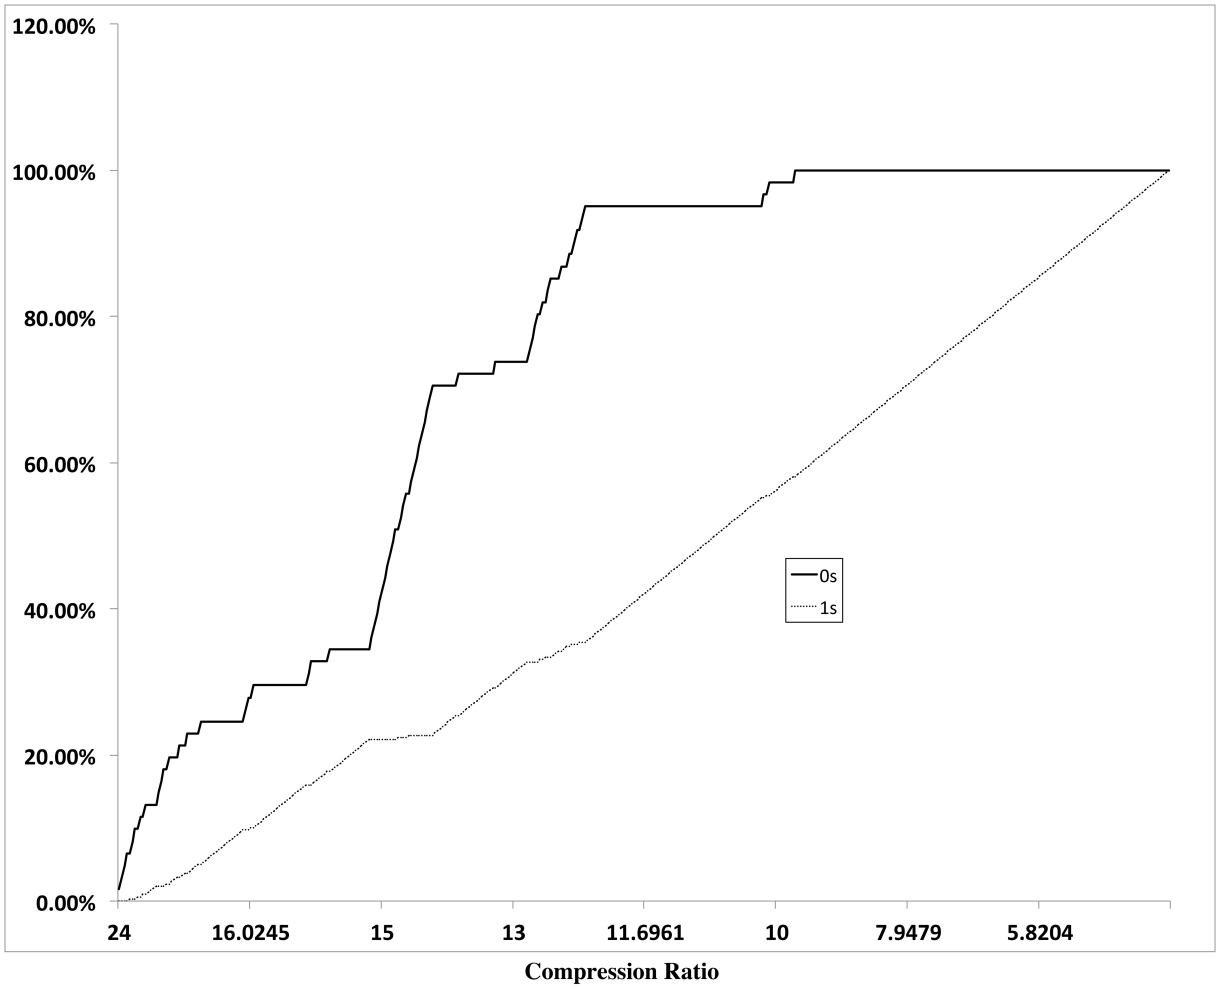

Figure 4.22: Cumulative distributions of subjective radiologist scores corresponding to compression ratio for JPEG and JPEG2000 compressed images. K-S = 59.68%.

The Optimal Threshold of Diagnostically Lossless Compression

The result of the logistic curve-fitting model (described in the Section 4.3.1) applied to all data points is presented in Figure 4.23. For example, for 99.99% confidence, the recommended SSIM threshold is 0.955.

K-fold cross validation was applied to the logistic curve-fitting model and weighted Youden index $WY(s)$.

For the logistic curve-fitting method, the desired level of confidence is provided by the user. For example, if the level of confidence is 99.99%, then the recommended threshold for compression corresponding to the SSIM index is 0.95 with 10% of FPs and 42% of FNs. The threshold results using the curve-fitting model with 10-fold cross validation procedure are presented in Table 4.1.

In the case of the Weighted Youden index, the user specifies the parameter λ, which corresponds to the weighting of FPs versus FNs. Larger values of λ correspond to higher thresholds of SSIM and SMSE and it means that having FPs is very costly. If λ equals 1, there are no TPs and too many FNs. Thus, λ equal or very close to 1 should not be considered. To demonstrate the results, we chose λ to be 0.95. Figure 4.2 shows the thresholds for SSIM and SMSE. The use of the weighted Youden index yields threshold values of 0.955 and 0.96 for the SSIM index and SMSE, respectively. This results in almost a small number of FPs, i.e. there is no risk that an image with higher quality score will be marked as unacceptable.

The rather high percentage of FNs resulting from applying the two methods is caused by the nature of the data. The JPEG compressed Brain CT images were almost always perceived as acceptable by radiologists, even for low quality scores.

The recommended thresholds for diagnostically lossless compression for SSIM and SMSE are similar according to the two methods used (logistic curve-fitting model and Youden index), as described above. However, it is understood that the optimal thresholds for SSIM and SMSE found in this work are dependent upon the collected data samples and may well be different for other data. In order to find optimal thresholds for a given quality measure, the experiment should employ a larger set of images and a larger number of radiologists as subjects. Moreover, the thresholds may vary according to image modality.

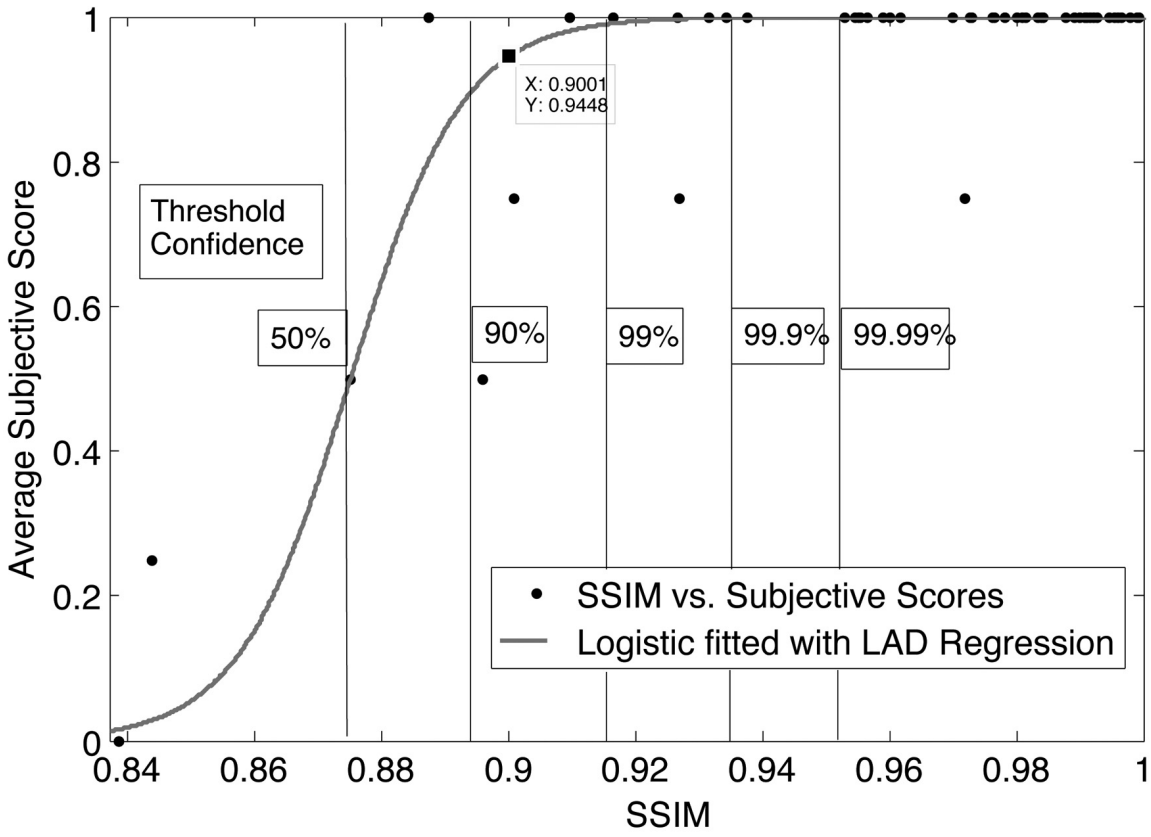

Figure 4.23: Logistic curve fitted with LAD Regression with threshold confidence corresponding to SSIM.

Objective Quality Measure	**Threshold with confidence level** = 99.99%	**Prediction Accuracy**
SSIM	.952	0.03% FPs 27% FNs
SMSE	.96	0.05% FPs 32% FNs

Table 4.1: SSIM index and SMSE thresholds obtained by means of the **logistic curve-fitting model** with 10-fold cross validation procedure using a fixed confidence level of 99.99%. The resulting percentage of FPs, FNs are given for each case.

Objective Quality Measure	Threshold with $\lambda = 0.95$	Prediction Accuracy
SSIM	.955	0.03% FPs 29% FNs
SMSE	.964	0.04% FPs 34% FNs

Table 4.2: SSIM index and SMSE thresholds obtained by means of the **Youden index** with 10-fold cross validation procedure using a fixed $\lambda = 0.95$. The resulting percentage of FPs, FNs are given for each case.

4.4.2 Second experiment: Local quality

In this local analysis only JPEG images were considered. The ROC curves corresponding to the SSIM index, SMSE and quality factor are shown in Figures 4.24, 4.25, 4.26 and 4.27.

Our examination of the data has shown that brain and body CT images should be analyzed separately. These two image types posses different characteristics (e.g. textures, distribution of intensities, variances of subblocks). We therefore expect, and observe, different types of degradations that become noticeable at different compression levels.

In the local analysis, the ROC curves for body CT images corresponding to SSIM, SMSE and JPEG quality factor Q have very similar AUC. However, the ROC curves corresponding to brain CT images show poorer performance. We already observed a similar result for brain images in the first experiment. Again, for the brain, most radiologists' responses were positive, meaning that the images (the regions in Experiment 2) are acceptable. Due to too many "acceptables" for brain images in experiment 2, no proper local analysis is possible for the tested brain CT images.

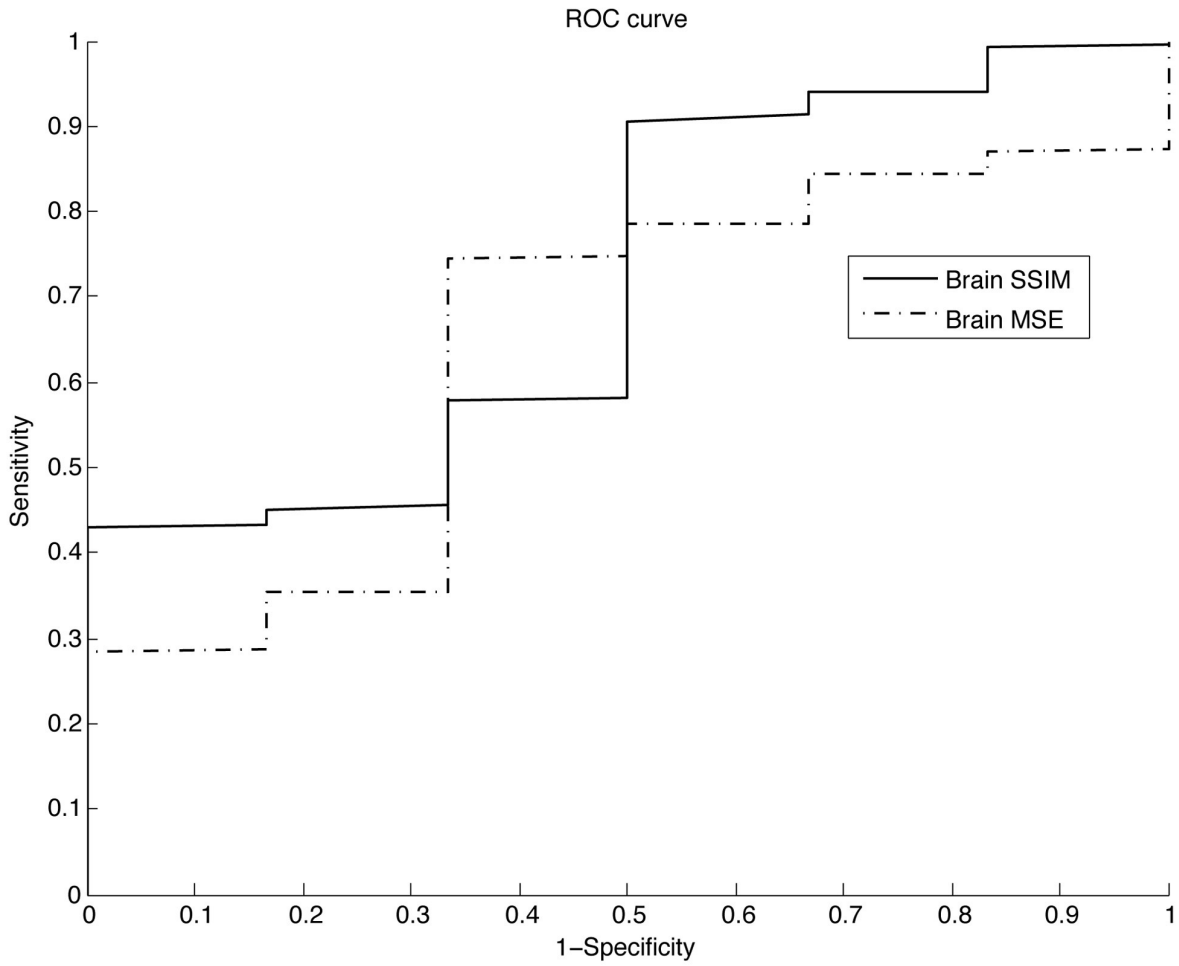

Figure 4.24: ROC curves corresponding to SSIM and SMSE for brain CT images. AUC SSIM = 0.7183 AUC SMSE = 0.6481.

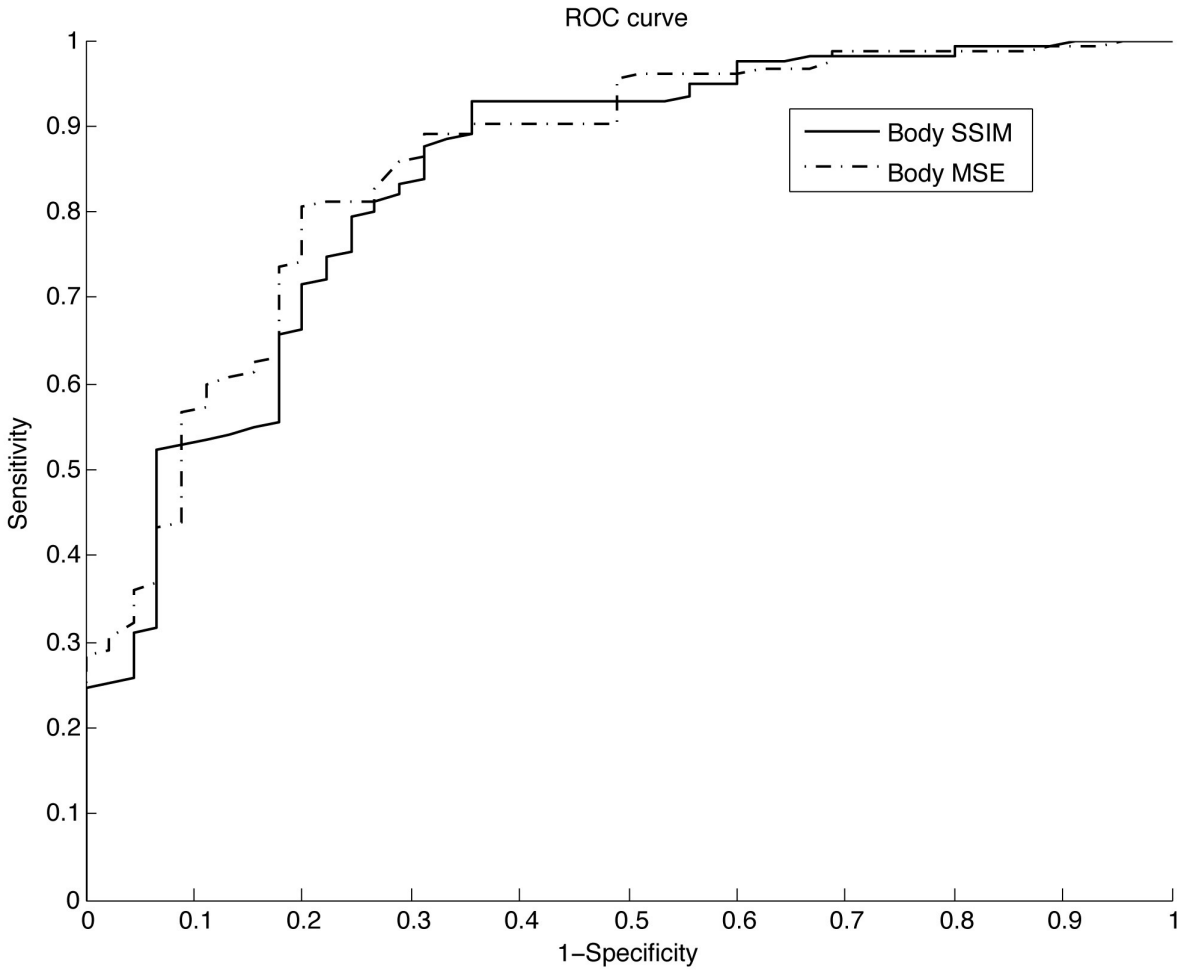

Figure 4.25: ROC curves corresponding to SSIM and SMSE for body CT images. AUC SSIM = 0.8422 AUC SMSE = 0.8541.

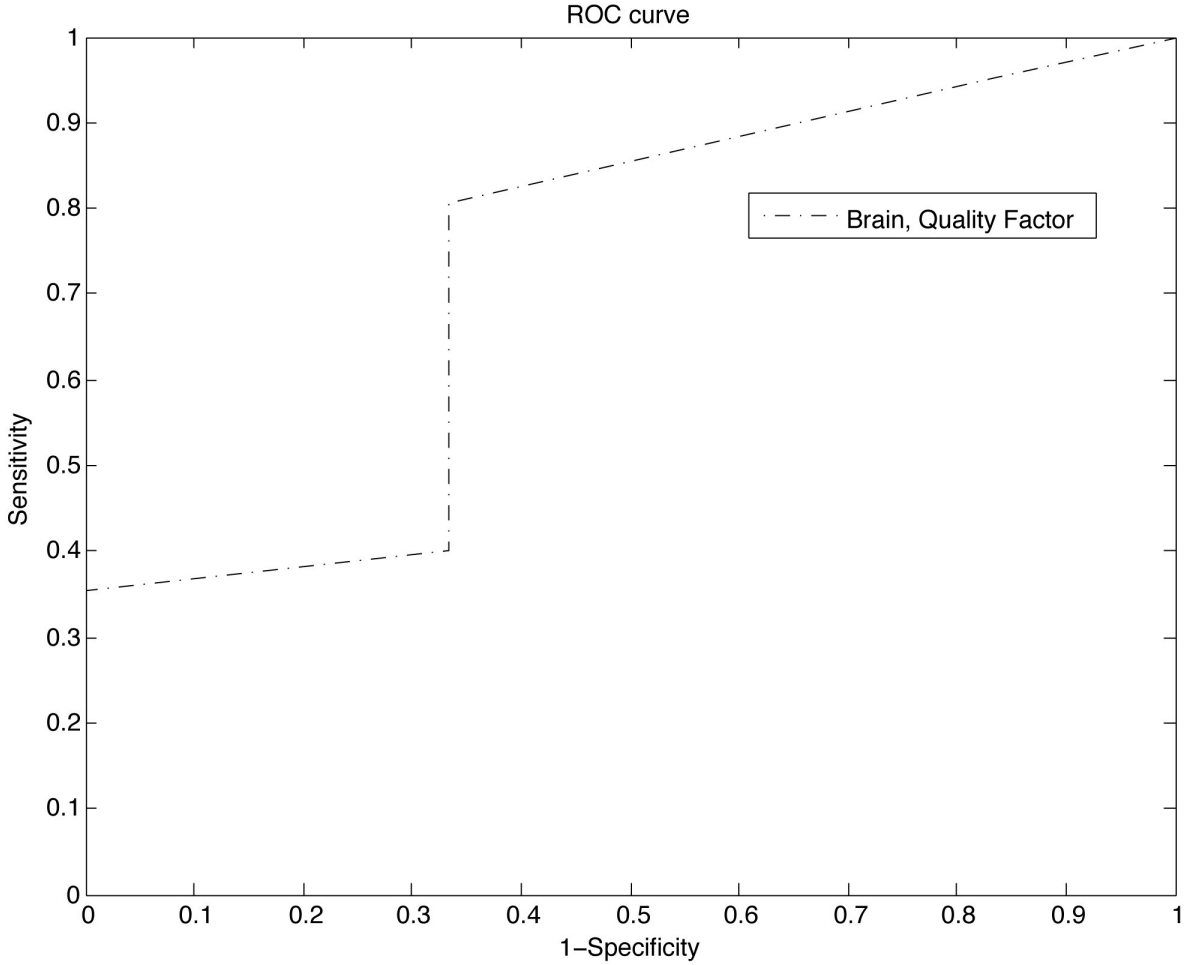

Figure 4.26: ROC curves corresponding to JPEG quality factor for brain CT images AUC QF = 0.7284.

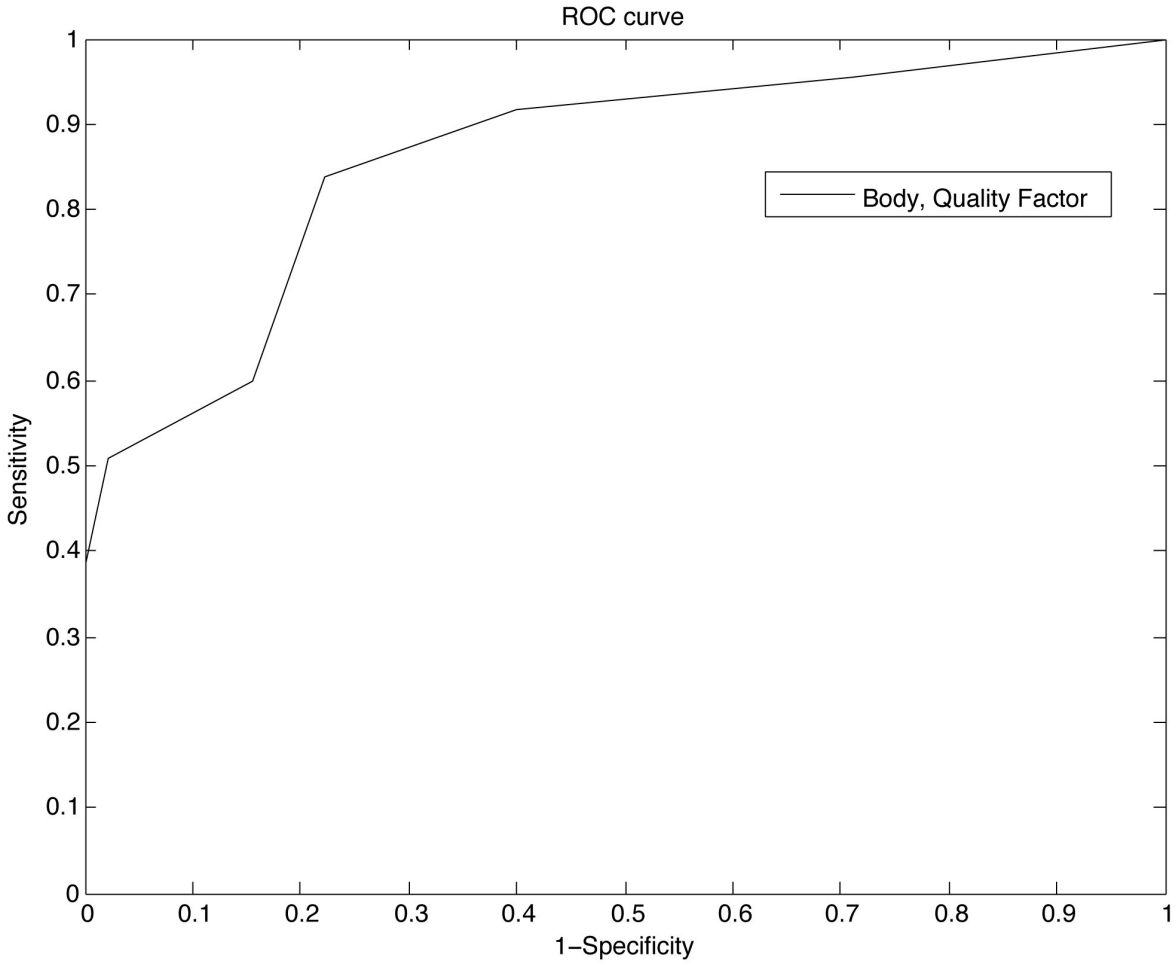

Figure 4.27: ROC curves corresponding to JPEG quality factor for body CT images AUC QF = 0.8613.

4.4.3 Concordance of subjective responses

Here we discuss the consistency of the radiologists' responses. The concordances of the subjective responses corresponding to global and local experiments are shown in Figures 4.28 and 4.29. The term "consistency" in the table refers to the same choice for each repetition for each radiologist (intra-class concordance), whereas "agreement" corresponds to the same choice between radiologists (inter-class concordance). Moreover, we have also computed the coherence between compression level and acceptance, labeled as 'monotonicity' in the figure. According to this summary, during the local experiment the radiologists were more consistent in both intra- and inter-class concordance. However, the consistency statistics percentages are on average close to 100% and are overall satisfying, therefore we conclude that the experiment data are reliable.

GLOBAL EXPERIMENT	Radiologist 1	Radiologist 2
Acceptable Responses	77	83
Unacceptable	15	6
Repetition inconsistency	8	11
Consistency (intra class)	92%	89%
Agreement (inter class)	84%	84%
Monotonicity	99%	99%

Figure 4.28: Consistency of subjective responses of the global experiment.

LOCAL EXPERIMENT	Radiologist 1	Radiologist 2
Acceptable Responses	263	286
Unacceptable	37	14
Consistency (intra class)	100%	100%
Agreement (inter class)	92%	92%
Monotonicity	97%	99%

Figure 4.29: Consistency of subjective responses of the local experiment.

4.5 Discussion

Our AUC and KS analyses of the results of Experiment 1 yield the following important conclusions:

- Compression ratio demonstrates the poorest performance of the four quality measures examined.
- MSE performs inconsistently as indicators of visual/diagnostic quality.
- SSIM demonstrates the best performance, i.e., it provides the closest match to the subjective assessments of the radiologists.

- JPEG quality factor performs nearly as well as SSIM
- Our examination of the data indicates that brain and body CT images should be analyzed separately

Furthermore, we have utilized a weighted Youden index and a curve-fitting model to suggest thresholds for SSIM and SMSE as indicators of acceptable compression levels for brain CT and body CT compressed images. The thresholds reported in this work correspond to the specific set of sample data collected in our experiments. In order to obtain more reliable and statistically significant thresholds, the experiment should employ a larger set of images - separately for the various anatomical regions as well as modalities - and involve more radiologists as subjects.

The local analysis in Experiment 2 was performed separately on brain and body CT images. As expected, the performance of SSIM was quite satisfactory in terms of AUC (area under ROC curves). SMSE and JPEG quality factor also performed well.

Although JPEG quality factor shows a good performance, it is not reliable since it is image-dependent. Please refer to Section 2.2 of Chapter 2 for detailed discussion on JPEG quantization.

How can the SSIM and SMSE thresholds be used in practice? For example, in order to use the recommended threshold, the first compression attempt has to be performed for the purpose of determining the SSIM or SMSE quality score. The corresponding input parameters (e.g. compression ratio and quality factor) can then be adjusted accordingly.

Finally, the recommended threshold results presented here cannot be directly compared with the recommendations based on compression ratio for medical images (as presented in Section 1.2). Based on our findings compression ratio is not a reliable quality predictor. In a study done by Flynn *et al.*, it has been shown that regardless of the type of image, the threshold at which the human eye is not able to determine compression degradations is on average at SSIM= 0.95 [26]. Nevertheless, it would be difficult to find a common threshold on SSIM for diagnostically lossless compression. In practice, the threshold is likely to be application dependent. For example, the same image content viewed on different display devices at different viewing distance would result in different visual quality, however, according to the existing objective measures such as MSE/PSNR and SSIM a fixed quality prediction is given.

4.5.1 Training of the stability constant for the SSIM index

In Figure 4.30 are plotted ROC curves that correspond to the SSIM quality index (actually, the structure term of SSIM) for various values of the stability constant C. The AUC (area under ROC curve) values associated with these values of C are plotted in Figure 4.31. The monotonically decreasing behavior of the curve for values of C near 0 is quite interesting, and suggests that best ROC-performance (i.e., maximum AUC) is obtained. That being said, a nonzero stability constant of very low magnitude will be employed in order to avoid any possible division by zero in the computation of the SSIM index. Such "zero denominators" could arise in the case of "flat" image blocks, i.e., pixel blocks with constant greyscale value, hence zero variance, in both uncompressed and compressed images.

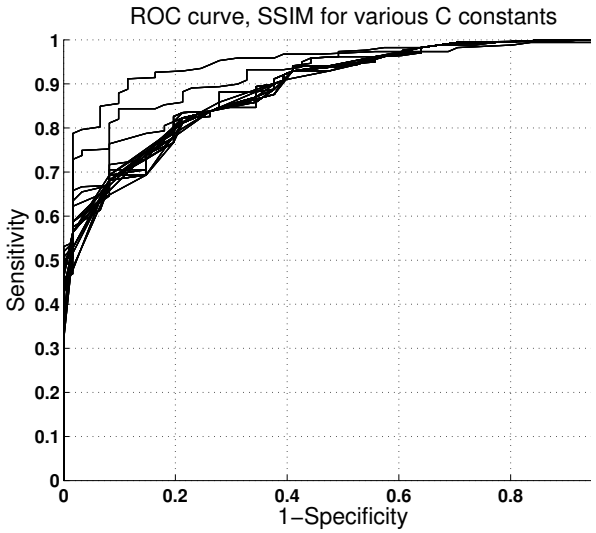

Figure 4.30: ROC curves corresponding to various values of the constant C in the structure term of SSIM.

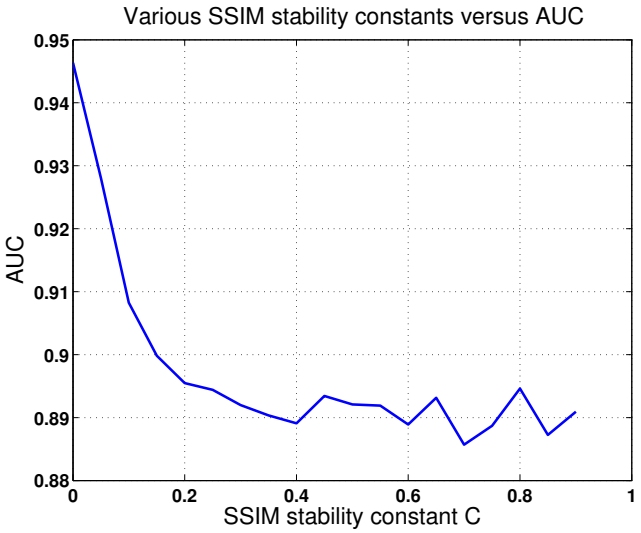

Figure 4.31: Plot of AUC and the various SSIM stability constants.

4.6 Conclusion

In this work, we compared the performances of SSIM, MSE/PSNR, compression ratio CR and JPEG quality factor Q, based on experimental data collected in two experiments

involving radiologists. The first xperiment involved a global quality assessment of 100 brain and body CT images at various compression ratios. The radiologists evaluated compressed images as acceptable or unacceptable as compared to their uncompressed counterparts. An ROC and Kolmogorov-Smirnov analysis indicates that compression ratio is not always a good indicator of visual quality. Moreover, SSIM demonstrated the best performance, i.e., it provides the closest match to the radiologists' assessments. We also show that a weighted Youden index and curve fitting method can provide SSIM and MSE thresholds for acceptable compression ratios. The second experiment involved a local/regional image quality analysis of these images by the radiologists. An ROC analysis once again shows that SSIM provides a closer match to subjective assessments.

4.7 Improved experiment design

The above experiment has several flaws and therefore we are unable to draw all desired conclusions. First of all, the sample sizes (i.e., the number of images and radiologists involved in the subjective assessment) are too small. Of course, at the time of the experiment those were the only available resources. Second of all, there were too few negative responses to make proper ROC analysis for brain CT images. We are currently working on conducting a new experiment- an improved version of the experiment described above. We are in the process of conducting a larger scale experiment where at least ten radiologists are involved in a subjective assessment of diagnostic quality of compressed neuro CT and body CT normal and pathological images. The experience gained during the course of the previous experiments and consultations with Prof. Marriott (Statistics and Actuarial Science Department, University of Waterloo) brought new ideas for designing a better, more statistically valid experiment. Also, the suggestions provided by Nasim Themmati, a radiologist from McMaster University, have helped with the choice of images for the experiment.

Full results are not yet available. In what follows, is the experiment design.

- The group of (al least 10) radiologists subjects includes experienced radiologists as well as residents (McMaster University).
- Types of pathologies: Based on previous findings by Koff *et al.* [38], [39], the pathologies include subtle lesions in the liver for the body, and brain parenchyma and posterior fossa for the heads. By subtle lesions, it is meant two main types:

  - Very small lesions, limit in size, of less than 2 mm, but high contrast (calcifications) or low density (tiny cysts);
  - Subtle parenchymal alterations translating into subtle differences in density such as cerebral infarcts.

- The experiment consists of two parts:
  - Task 1: Subjective classification of compressed normal and pathological brain CT and body CT images into one of three groups: 1) not noticeable and acceptable, (2) noticeable and acceptable and (3) noticeable and not acceptable by radiologists. The compressed images are displayed together, side by side, with their uncompressed versions.
  - Task 2: Subjective classification of images into three groups: (1) Normal, (2) Pathological, (3) Unable to assess.

- The brain CT and body CT images used in the experiment were carefully chosen with the help of radiologists and contain pathological and normal cases (about 1/3 of normal cases).

- Working environment: MIIRC@M office; Eizo Radiforce monitor, 54 cm (21,3") display, with a 1200x1600 native resolution (3:4 aspect ratio) and a viewing size of 324.0 x 432.0 mm. It displays 10-bit colors.

- Images are displayed randomly.

- Number of images
  - Trial experiment includes 6 images from each brain CT and body CT sets. These images are repetitions of images that are included in the main part of the experiment.
  - Main experiment, Part 1

    Number of images: 306 brain CT, 306 body images (30 different images compressed at five compression ratios using JPEG and JPEG2000 algorithms including 6 repetitions added at the end of the sequence).
  - Main experiment, Part 2

    The set of images includes 100 brain CT and 100 body CT randomly chosen images from the set of images displayed in the first part of the main experiment

- Repetitions are included in the trial experiment as well as at the end of the first part of the main experiment to check consistency of the radiologists.
- The second experiment serves as a measure of diagnostic accuracy and concordance of the radiologists.
- Duration of the experiment: The number of images has been adjusted to the time limitation of the experiment. Expected timeline of the experiment: (Total 90 minutes)
  - Trial experiment including explanation of the task: 10 minutes
  - Main experiment, Part 1:
    - Brain CT: 25 minutes
    - Body CT: 25 minutes
  - Break: 10 minutes
  - Main experiment, Part2:
    - Brain CT: 10 minutes
    - Body CT: 10 minutes

All of the above items included in the design of the experiment have been restricted to several limitations. It is important to find the balance between the trade-offs that are involved. First of all, the time of radiologists who are willing to participate in such research study is up to 1.5 hrs. We have chosen 30 images for each brain CT and body CT at five compression levels. This choice is optimal for the given time frame.

With this new experiment we are hoping to answer several questions:

- Determine a more reliable threshold for compression and tuning of the SSIM index.
- Examine the difference in diagnosis between experienced and novice radiologists.
- Understand the difference between quality and diagnostic quality of images by grouping the subjective assessments into two or three groups of responses. This may be possible to accomplish with the use of the third button in the first part of the experiment: Acceptable, noticeable distortions. This third button serves to differentiate between the diagnostic quality and simply visual quality of compressed images. According to radiologists' comments, we believe that some compressed images with very minor distortions are still of "diagnostic quality".

- Measure diagnostic accuracy with a task of detecting a pathology.
- Is the difference among the quality measures clinically important?
- Do pathological and non-pathological images compress in the same way?

It has been mentioned by several radiologists that normal images compress better as compared to pathological images. Also, for brain CT images, compression does not affect diagnosis as much as in the case of body CT images. We are hoping to address these and the related problems using the data from the experiment. Moreover, in further future experiments we will attempt to determine thresholds for SSIM and MSE on the basis of larger sample size data, for several modalities and a variety of anatomical regions.

Chapter 5

"Weberized" approximation of images: An attempt to construct a "Weberized" L^2-based image approximation and similarity measure

In this chapter we consider a modified method of L^2 approximation of functions that takes into consideration a Weber-type model of perception. Weber's law (sometimes called Weber-Fechner law) comes from psychophysics and relates human perception to change in the magnitude of stimuli [87].

5.1 Weber's model of perception

Consider a signal u, for example a grey-scale image where $u(x)$ denotes the intensity at a point x, the minimum perceived sensitivity Δu is related to the intensity u as follows,

$$\frac{\Delta u}{u} = c, \tag{5.1}$$

where c is roughly constant. In other words, the eye/brain/mind is less sensitive to a given change in intensity Δu in regions where the image intensity is high.

In Chapter 2.6.2 we discussed some image quality measures, namely MSE and the SSIM index. The MSE has been shown to behave poorly as perceptual quality assessor, however, it is known for its simplicity and is still being widely employed. A newer measure, the SSIM index, is also simple to compute and it shows better performance (than MSE) to measure the quality of images.

SSIM is also formulated to accommodate Weber's model of perception. In order to see this connection, we take a closer look at one of the components of the SSIM function, the luminance function

$$l(f,g) = \frac{2\mu_f\mu_g + C_1}{\mu_f^2 + \mu_g^2 + C_1}, \text{ (note: } -1 \leq l(f,g) \leq 1). \tag{5.2}$$

The factor $l(x,y)$ measures the discrepancy between the mean values of x and y. Assume that f is fixed and g is an approximation to f. Then $l(f,g)$ can be written as follows,

$$l(f,g) = \frac{2(\frac{\mu_g}{\mu_f}) + C_1'}{1 + (\frac{\mu_g}{\mu_f})^2 + C_1'}. \tag{5.3}$$

This expression of $l(f,g)$ shows a dependence on ratio $\frac{\mu_g}{\mu_f}$, in accordance with the Weber model.

5.2 Approximation of signals using intensity weighting functions

"Weberizing" the L^2 distance function

The L^2 distance function, which in our image context is the MSE, does not incorporate any weighting of the pixels and therefore has no connection with Weber's Law. In fact, all L^p- based metrics for $p \geq 1$, do not conform to the Weber's model of perception since they involve integrations over appropriate powers of intensity differences, $|u(x) - v(x)|$, with no consideration of the magnitudes of $u(x)$ or $v(x)$. One idea to "Weberize" the L^2 distance function is to incorporate some weighting of the pixel intensities inside the integral.

In what follows, we let $u \in L^2[a,b]$ be the signal we wish to approximate and let $\{\phi_k\}_{k=1}^{\infty}$ denote a complete orthonormal basis for the Hilbert space of functions $L^2[a,b]$. In the usual

L^2-based approach, the approximation,

$$u \approx u_N = \sum_{k=1}^{N} c_k \phi_k, \tag{5.4}$$

is obtained by minimizing the squared L^2 error $\Delta_N^2 = \|u - u_N\|^2$ given by

$$\Delta_N^2 = \int_a^b \left[u(x) - \sum_{k=1}^{N} c_k \phi_k(x) \right]^2 dx. \tag{5.5}$$

It is a known result that Δ_N^2 is minimized when

$$c_k = \langle u, \phi_k \rangle, \quad 1 \leq k \leq N. \tag{5.6}$$

The "Weberization" of the L^2 error will be accomplished by dividing the integrand by the intensity at each pixel. The "Weberized" L^2 approximation is then accomplished by minimizing the modified squared L^2 error,

$$\begin{aligned} W_N^2 &= \int_a^b \left[\frac{u(x) - \sum_{k=1}^{N} c_k \phi_k(x)}{u(x)} \right]^2 dx \\ &= \int_a^b \frac{1}{[u(x)]^2} \left[u(x) - \sum_{k=1}^{N} c_k \phi_k(x) \right]^2 dx. \end{aligned} \tag{5.7}$$

Eq. (5.7) is an intensity weighted L^2 approximation. This definition of an error conforms to the Weber Law: Regions of the interval $[a, b]$ where the signal intensity $|u(x)|$ is larger/smaller are weighted to a lesser/greater extent in the integral.

The image function $u(x)$ may contain the value of 0, representing a black pixel. In order to avoid such problems, a "stability parameter" $\epsilon > 0$ is introduced, i.e.,

$$W_N^2 = \int_a^b \frac{1}{[u(x) + \epsilon]^2} \left[u(x) - \sum_{k=1}^{N} c_k \phi_k(x) \right]^2 dx, \tag{5.8}$$

The weighting function $[u(x) + \epsilon]^2$ in Eq. (5.8) can be generalized. This leads to the

idea of a more general form of intensity-weighted approximation, i.e.,

$$W_N^2 = \int_a^b g(u(x)) \left[u(x) - \sum_{k=1}^{N} c_k \phi_k(x) \right]^2 dx, \tag{5.9}$$

where $g(u(x)) > 0$ will accommodate some kind of inverse intensity-weighting, e.g.,

$$g(u(x)) \propto \frac{1}{u(x)^\alpha}, \quad \alpha > 0. \tag{5.10}$$

"Weberized" L^2 approximation

As in the case of the L^2 approximation, the "Weberized" L^2 distance in Eq. (5.7) can be minimized in order to find coefficients to approximate a given function. The coefficients c_k that minimize W_N^2 in Eq. (5.9) may be determined as follows: We impose the stationarity conditions

$$\frac{\partial W_N^2}{\partial c_p} = -2 \int_a^b g(u(x)) \left[u(x) - \sum_{k=1}^{N} c_k \phi_k(x) \right] \phi_p(x)\, dx = 0, \quad 1 \le p \le N. \tag{5.11}$$

After rearranging, we obtain

$$\sum_{k=1}^{N} c_k \int_a^b g(u(x)) \phi_p(x) \phi_k(x)\, dx = \int_a^b g(u(x)) u(x) \phi_p(x)\, dx. \quad 1 \le p \le N, \tag{5.12}$$

Eq. (5.12) represents a system of linear equations in the coefficients c_p having the form $\mathbf{Ac} = \mathbf{b}$, or

$$\sum_{k=1}^{N} a_{pk} c_k = b_p, \quad 1 \le p \le N, \tag{5.13}$$

where

$$a_{pk} = \langle \phi_p, \phi_k \rangle_W, \qquad b_p = \langle u, \phi_p \rangle_W. \tag{5.14}$$

$\langle \cdot, \cdot \rangle_W$ denotes a weighted inner product with the weight function $g(u(x))$, i.e.:

$$\langle u, v \rangle_W = \int_a^b g(u(x)) u(x) v(x)\, dx. \tag{5.15}$$

Consider the special case where the weight function $g(u(x)) = C$ is a constant. We may think of this weighting as "global", i.e., the same at each point. In this case the matrix $\mathbf{A}$ is diagonal, and we obtain

$$c_p = b_p = \int_a^b u(x)\phi_p(x)\, dx, \tag{5.16}$$

which are the Fourier coefficients of Eq. (5.6).

Note that in the case where the weighting function is not a constant, each pixel difference in Eq. (5.12) is weighted according to the value of the intensity at a given pixel. We may regard this kind of weighting procedure as "local".

5.2.1 Another look at the "Weberized" L^2 approximation method

We now consider a particular case of the "Weberized approximation" with the weighting function given by

$$g(u(x)) = \frac{1}{u(x)^2}. \tag{5.17}$$

The choice of the weighting function in Eq. (5.17) is a natural consideration since it involves division of the integrand in Eq. (5.7) by a single power of pixel intensities.

Eq. (5.7) can be written as follows,

$$W_N^2 = \int_a^b \left[1 - \sum_{k=1}^{N} c_k \psi_k(x)\right]^2 dx, \quad \text{where} \quad \psi_k(x) = \frac{\phi_k(x)}{u(x)}, \quad 1 \le k \le N. \tag{5.18}$$

Thus minimizing W_N^2 amounts to finding the best L^2 approximation of the constant function $h(x) = 1$ in terms of the set of functions $\psi_k(x)$. In other words, if $\psi_k(x)$ basis functions are used then the "Weberized" error measures how far the approximated signal is from the unity. This is just another way of looking at the approximation procedure. There is no new mathematics since the form of $g(u(x))$ in Eq. (5.17) implies that the system of linear equations in (5.12) may be rewritten as follows,

$$\sum_{k=1}^{N} c_k \int_a^b \psi_p(x)\psi_k(x)\, dx = \int_a^b 1\ \phi_p\, dx, \quad 1 \le p \le N, \tag{5.19}$$

which is precisely the set of equations corresponding to the approximation of $h(x) = 1$ in the ψ_k basis.

A natural question arises from this viewpoint: "How linearly independent, or even orthogonal, are the functions ψ_k?" One way to characterize their orthogonality is to look at the structure of the matrix $\mathbf{A}$ in Eq. (5.12) – the more concentrated along the diagonal it is, the more orthogonal the basis. If, in our applications, we consider subintervals or subblocks of an image, then one expects that in places where the image function $u(x)$ is relatively constant, the functions ψ_k will behave much like their orthogonal ϕ_k counterparts (Fourier approximation). The interesting cases will occur where the image function $u(x)$ exhibits more "activity," i.e., higher variance and/or edges.

5.3 The "Weberized" L^2 as a distance function

The constructed "Weberized" L^2 error in Eq. (5.7) will be used to measure the error between the original and distorted images, i.e. the distance between the two. From a mathematical point of view, we would like such measure to be a valid metric. We now return to Eq. (5.7), in which the function $u \in L^2[a, b]$ is considered to be the "reference function" that is to be approximated. Instead of employing a linear combination of orthogonal basis functions, we simply consider another function $v \in L^2[a, b]$ as an approximation to u. The error of approximation, in the "Weberized" L^2 sense, and with u as reference, will be denoted as

$$\Delta(u, v) = \left[\int_a^b \frac{1}{u(x)^2} \left[u(x) - v(x) \right]^2 dx \right]^{1/2} . \tag{5.20}$$

Upon interchanging u and v, we obtain the "Weberized" L^2 error of approximating v with u, i.e.,

$$\Delta(v, u) = \left[\int_a^b \frac{1}{v(x)^2} \left[u(x) - v(x) \right]^2 dx \right]^{1/2} . \tag{5.21}$$

Note that both $\Delta(u, v)$ and $\Delta(v, u)$ are examples of weighted L^2 distances having the form,

$$d_2(u, v\,; g(u(x))) = \left[\int_a^b g(u(x)) \left[u(x) - v(x) \right]^2 dx \right]^{1/2} \tag{5.22}$$

and

$$d_2(v,u\,;g(v(x))) = \left[\int_a^b g(v(x))\,[v(x)-u(x)]^2\,dx\right]^{1/2}, \tag{5.23}$$

where $g(u(x)) > 0, g(v(x)) > 0$ for $x \in [a,b]$. Note that we may rewrite the above as follows,

$$d_2(u,v\,;g(u(x))) = \left[\int_a^b [\sqrt{g(u(x))}\,u(x) - \sqrt{g(u(x))}\,v(x)]^2\,dx\right]^{1/2}, \tag{5.24}$$

$$d_2(v,u\,;g(v(x))) = \left[\int_a^b [\sqrt{g(v(x))}\,v(x) - \sqrt{g(v(x))}\,u(x)]^2\,dx\right]^{1/2}. \tag{5.25}$$

In other words,

$$d_2(u,v\,;g(u)) = d_2(\sqrt{g(u)}u, \sqrt{g(u)}v\,;1) = d_2(\sqrt{g(u)}\,u, \sqrt{g(u)}\,v), \tag{5.26}$$

$$d_2(v,u\,;g(v)) = d_2(\sqrt{g(v)}v, \sqrt{g(v)}u\,;1) = d_2(\sqrt{g(v)}\,v, \sqrt{g(v)}\,u), \tag{5.27}$$

where $d_2(u,v)$ denotes the usual (unweighted) L^2 distance between u and v, i.e.,

$$d_2(u,v) = \left[\int_a^b [u(x)-v(x)]^2\,dx\right]^{1/2}. \tag{5.28}$$

This implies that the weighted distance function $d_2(u,v;g(u))$ satisfies all the properties of a metric, including the triangle inequality. Let us confirm the latter statement: For any $u,v,w \in L^2[a,b]$,

$$\begin{aligned} d_2(u,v\,;g(u)) &= d_2(\sqrt{g(u)}u, \sqrt{g(u)}v) \\ &\le d_2(\sqrt{g(u)}u, \sqrt{g(u)}w) + d_2(\sqrt{g(u)}w, \sqrt{g(u)}v) \quad \text{(since } d_2 \text{ is a metric)} \\ &\le d_2(u,w\,;g(u)) + d_2(w,v\,;g(u)). \end{aligned} \tag{5.29}$$

However, $d_2(u,v\,;g(u)) \neq d_2(v,u\,;g(v))$ since, in general, the functions $g(u(x))$ and $g(v(x))$ are not the same. Thus, $\Delta(u,v)$ is not symmetric w.r.t. u and v and therefore is not a metric. Furthermore, it is not necessary to contract a metric out of the "Weberized" measure. Given that our primary applications will be in signal/image processing, let us assume that we are working with square integrable functions that are bounded from above and away from zero. Consider a fixed lower bound $L > 0$ and a fixed upper bound $M > L$

and define the following space of functions,

$$\mathcal{F} = \{u : [a, b] \to [L, M]\}. \tag{5.30}$$

Theorem 5.3.1.

$$\frac{2L - M}{L}\Delta(u, v) \leq \Delta(v, u) \leq \frac{M}{L}\Delta(u, v). \tag{5.31}$$

Proof. $u \in \mathcal{F} \Rightarrow u \in L^2[a, b]$ are square integrable functions. It then follows that

$$\frac{1}{M} \leq \frac{1}{u(x)} \leq \frac{1}{L}, \quad x \in [a, b] \tag{5.32}$$

From (5.20),

$$\frac{1}{M}d_2(u, v) \leq \left\{ \begin{array}{c} \Delta(u, v) \\ \Delta(v, u) \end{array} \right\} \leq \frac{1}{L}d_2(u, v). \tag{5.33}$$

Since our usual convention is to consider the function u as the "reference" or "target" function being approximated by the function v, here we consider $\Delta(u, v)$ and express $\Delta(v, u)$ in terms of it. From Eq. (5.20), we see that both $\Delta(u, v)$ and $\Delta(v, u)$ lie inside the real interval

$$I = \left[\frac{1}{M}d_2(u, v), \frac{1}{L}d_2(u, v)\right] \tag{5.34}$$

where, $d_2(u, v)$ denotes the L^2 distance between u and v, i.e.,

$$d_2(u, v) = \left[\int_a^b [u(x) - v(x)]^2 \, dx\right]^{1/2}. \tag{5.35}$$

This implies that $\Delta(u, v)$ and $\Delta(v, u)$ cannot differ by more than the length of this interval, i.e.,

$$|\Delta(u, v) - \Delta(v, u)| \leq \left(\frac{1}{L} - \frac{1}{M}\right) d_2(u, v). \tag{5.36}$$

But from Eq. (5.20), the above inequality yields

$$\begin{aligned} |\Delta(u,v) - \Delta(v,u)| &\leq M\left(\frac{1}{L} - \frac{1}{M}\right)\Delta(u,v) \\ &= M\frac{(M-L)}{LM}\Delta(u,v) \\ &= \frac{M-L}{L}\Delta(u,v). \end{aligned} \tag{5.37}$$

This implies that

$$\left(1 - \frac{M-L}{L}\right)\Delta(u,v) \leq \Delta(v,u) \leq \left(1 + \frac{M-L}{L}\right)\Delta(u,v), \tag{5.38}$$

which may be simplified to

$$\frac{2L-M}{L}\Delta(u,v) \leq \Delta(v,u) \leq \frac{M}{L}\Delta(u,v). \tag{5.39}$$

□

In Theorem 5.3.1, the coefficient of the LHS may be negative, in which case no new information is yielded since $\Delta(v,u)$ must be nonnegative. As a result, it follows that

$$\Delta(u,v) \to 0 \Rightarrow \Delta(v,u) \to 0. \tag{5.40}$$

The above result implies that it is sufficient to consider only $\Delta(u,v)$ as a measure of closeness between u and v since

$$\frac{1}{M}d_2(u,v) \leq \left\{ \begin{array}{c} \Delta(u,v) \\ \Delta(v,u) \end{array} \right\} \leq \frac{1}{L}d_2(u,v). \tag{5.41}$$

Clearly, if $v \to u$ in d_2 metric, then both $\Delta(u,v)$ and $\Delta(v,u)$ go to zero.

We now address an important property that we require from the weighting function. We could employ a more general form of the weighting function that depends on $u(x)$ and $v(x)$. The dependence on both intensity functions would lead to complicated mathematics for which finding approximated coefficients would be difficult. In our consideration the weighting function depends on one intensity function, $u(x)$ or $v(x)$. The fundamental

property that we require of the weighting function is that it conforms to the Weber's model of perception. Recall Eq. (5.42) in which we declared the weighting function belong to the symmetric, decreasing family of functions:

$$g(u(x)) = \frac{1}{u(x)^\alpha}, \quad \alpha > 0, \tag{5.42}$$

This choice of $g(u(x))$ accommodates Weber's model. As presented in Eq. (5.17), a natural choice of the parameter α is 2. With this choice of α, the approximation errors $d_{2W}(u, v)$ and $d_{2W}(v, u)$ are:

$$d_{2W}(u, v) = \left[\int_a^b \left[1 - \frac{v(x)}{u(x)}\right]^2 dx\right]^{1/2} = \Delta(u, v), \tag{5.43}$$

and

$$d_{2W}(v, u) = \left[\int_a^b \left[1 - \frac{u(x)}{v(x)}\right]^2 dx\right]^{1/2} = \Delta(v, u), \tag{5.44}$$

Note also that from the definitions of $\Delta(u, v)$ and $\Delta(v, u)$ in (5.20) and (5.21) and Eq. (5.27),

$$\begin{aligned} \Delta(u, v) &= d_2(u^{-1}u, u^{-1}v) = d_2(1, v/u) \\ \Delta(v, u) &= d_2(v^{-1}u, v^{-1}v) = d_2(u/v, 1). \end{aligned} \tag{5.45}$$

These results, of course, could have been derived directly from (5.20) and (5.21). From (5.20),

$$\begin{aligned} \Delta(u, v) &= \left[\int_a^b \frac{1}{u(x)^2} [u(x) - v(x)]^2 \, dx\right]^{1/2} \\ &= \left[\int_a^b \left[1 - \frac{v(x)}{u(x)}\right]^2 dx\right]^{1/2} \end{aligned} \tag{5.46}$$

Likewise,

$$\begin{aligned}\Delta(v,u) &= \left[\int_a^b \frac{1}{v(x)^2}\left[u(x)-v(x)\right]^2 dx\right]^{1/2} \\ &= \left[\int_a^b \left[\frac{u(x)}{v(x)}-1\right]^2 dx\right]^{1/2} \end{aligned} \tag{5.47}$$

The Δ distances between two functions $u(x)$ and $v(x)$ are measures of how well their ratios approximate the constant function 1 over $[a,b]$. This may be viewed as an implementation of Weber's Law into the L^2 distance measure. From (5.43) we see that for $\Delta(u,v)$ to be small, the ratio $v(x)/u(x)$ must be close to 1 for all $x \in X$. This already suggests that Weber's model of perception is being followed: Larger values of $u(x)$ will tolerate larger deviations between $v(x)$ and $u(x)$ so that the ratio $v(x)/u(x)$ is kept within a specified distance from 1. To see the differences between the L^2 and the "Weberized" L^2 distances, consider a reference image $u(x) = I$, $I \in \mathbb{R}_g = [A,B] \subset (0,\infty)$, $x \in X \subset \mathbb{R}$. Let $v(x) = I + \Delta I$ such that $\Delta I > 0$ be the constant approximation to $u(x)$. According to the Weber's model of perception, Eq. (5.1), $\Delta I = CI$ is the minimum perceived change in intensity corresponding to I. The L^2 distance between u and v is

$$d_2(u,v) = \left[\int_X \left[(I+\Delta I)-I\right]^2 dx\right]^{1/2} = \Delta I \left[\int_X dx\right]^{1/2} = KCI \tag{5.48}$$

where $K = \left[\int_X dx\right]^{1/2}$

Now, we compute the "Weberized" L^2 distance $\Delta(u,v)$:

$$\Delta(u,v) = K\frac{\Delta I}{I} = KC \tag{5.49}$$

From Eq. (5.48) we see that the higher the intensity level I the larger the L^2. This is expected since ΔI increases with I. The "Weberized" L^2 distance in Eq. (5.49), however, remains constant according to the Weber's model of perception Eq. (5.1). Thus we claim that $\Delta(u,v)$ better accommodates Weber's model of perception: The distance measure remains constant regardless of the perturbations ΔI of image intensities I.

5.3.1 A variation on the "Weberized" L^2 method: Matching of means

The "Weberized" L^2 method was introduced in an effort to accommodate Weber's model of perception. Here we consider a variation of the "Weberized" L^2 approximation where the mean value of the function being approximated is assumed to be the same as the original function. As such, one would hope that optimal "Weberized" approximations would perform better than their optimal L^2 counterparts in terms of image quality as measured, for example, by structural similarity (SSIM). However, this is not generally the case - numerical examples will be presented in the next chapter. It turns out that the SSIM measure of an approximation is affected if the mean value of the function being approximated is not matched very well. Let us just recall the definition for the SSIM measure $S(x, y)$ between two signals, $x, y \in \mathbf{R}^n$:

$$S(x,y) = S_1(x,y)S_2(x,y) = \frac{2\bar{x}\bar{y}}{\bar{x}^2 + \bar{y}^2} \cdot \frac{2s_{xy}}{s_x^2 + s_y^2}, \tag{5.50}$$

where

$$\bar{x} = \frac{1}{N}\sum_{i=1}^{N} x_i, \quad \bar{y} = \frac{1}{N}\sum_{i=1}^{N} y_i, \tag{5.51}$$

and

$$s_{xy} = \frac{1}{N-1}\sum_{i=1}^{N}(x_i - \bar{x})(y_i - \bar{y}), \quad s_x^2 = \frac{1}{N-1}\sum_{i=1}^{N}(x_i - \bar{x})^2, \quad s_y^2 = \frac{1}{N-1}\sum_{i=1}^{N}(y_i - \bar{y})^2. \tag{5.52}$$

Note that $S_1(x, y) = 1$ if $\bar{x} = \bar{y}$. If we consider y as an approximation to x, then if $\bar{x} \neq \bar{y}$, $S_1(x, y) < 1$. Since $S_1(x, y)$ is a multiplicative factor, it can influence the value of $S(x, y)$ significantly.

We can try to alleviate this difficulty by imposing the condition that the mean $\bar{y}$ of the approximation y matches the mean $\bar{x}$ of the signal being approximated. To do this, we shall proceed in a manner similar to that employed in [8], a study of optimal SSIM-based approximation. The following discussion starts with a formulation over a continuous domain, then we will transform to the discrete domain.

We now return to the formulation of the problem in terms of a complete orthonormal basis, but with a minor alteration. The orthonormal basis elements will now be indexed as $\{\phi_k\}_{k=0}^{\infty}$. In this way, the first element of the basis, $\phi_0(x)$, will have a rather special status, as it does in the Fourier and wavelet basis, as well as in our SSIM-based approximations.

The mean value of a function $u(x)$ over the interval $[a, b]$ is defined as follows,

$$\overline{u} = \frac{1}{b-a} \int_a^b u(x)\, dx. \tag{5.53}$$

It would be much easier if we simply formulated the problem over the interval $[0, 1]$, but we will continue with this general formulation. It will be assumed that only the function ϕ_0 has nonzero mean, i.e.,

$$\overline{\phi}_0 \neq 0, \quad \overline{\phi}_k = 0, \quad k = 1, 2, \cdots . \tag{5.54}$$

This holds in the Fourier as well as wavelet cases.

Once again, we consider the following N-function approximation to a function $u(x)$,

$$u \approx u_N = \sum_{k=0}^{N-1} c_k \phi_k, \tag{5.55}$$

but with the condition that

$$\overline{u} = \overline{u}_N, \tag{5.56}$$

which implies that

$$c_0 = (\overline{\phi}_0)^{-1}\, \overline{u}. \tag{5.57}$$

Now define

$$v(x) = u(x) - c_0 \phi_0, \tag{5.58}$$

so that the approximation in (5.55) becomes

$$v \approx v_N = \sum_{k=1}^{N-1} c_k \phi_k. \tag{5.59}$$

The condition on the ϕ_k implies that

$$\overline{v}_N = 0. \tag{5.60}$$

Before discussing the resulting "mean matching" "Weberized" L^2 method, let us note one important result of the above matching method: If the function $\phi_0(x)$ is constant over $[a, b]$ (as is the case for Fourier as well as Haar wavelet bases), then the coefficient c_0 is also

the Fourier coefficient of u in the ϕ_k basis, i.e.,

$$c_0 = \langle u, \phi_0 \rangle = \int_a^b u(x)\, \phi_0(x)\, dx. \tag{5.61}$$

In other words, it is also the "best L^2" coefficient for u, at least for the subspace spanned by ϕ_0.

To see this, note that if the orthonormal basis function ϕ_0 is a constant on $[a, b]$, it follows that

$$\phi_0(x) = (b-a)^{-1/2}, \quad a \le x \le b. \tag{5.62}$$

It is easily verified that this constant value must also be the mean of ϕ_0, i.e.,

$$\overline{\phi}_0 = (b-a)^{-1/2}. \tag{5.63}$$

The Fourier coefficient of a function u with respect to ϕ_0 is

$$\begin{aligned} c_0 &= \langle u, \phi_0 \rangle \\ &= \int_a^b u(x)\, \phi_0(x)\, dx \\ &= (b-a)^{-1/2} \int_a^b u(x)\, dx \\ &= (b-a)^{1/2}\ \overline{u} \\ &= (\overline{\phi}_0)^{-1}\, \overline{u}, \end{aligned} \tag{5.64}$$

in agreement with Eq. (5.57).

We now consider a "Weberized" L^2 approximation to v which minimizes the squared-distance function,

$$\begin{aligned} W_N^2 &= \int_a^b \frac{1}{u(x)^2} \left[v(x) - v_N(x) \right]^2 dx \\ &= \int_a^b \frac{1}{u(x)^2} \left[v(x) - \sum_{k=1}^{N-1} c_k \phi_k(x) \right]^2 dx. \end{aligned} \tag{5.65}$$

The weight function still employs the non-negative function $u(x)$, as opposed to the function $v(x)$, since the latter will have zero mean, i.e., $\overline{v} = 0$, in which case we might encounter problems with zeros.

If we now impose the stationarity constraints,

$$\frac{\partial W_N^2}{\partial c_p} = 0, \quad 1 \leq p \leq N-1, \tag{5.66}$$

we obtain a set of $N-1$ linear equations in the unknowns c_p having the form **Ac=b**, or

$$\sum_{k=1}^{N-1} a_{pk} c_k = b_p, \quad 1 \leq p \leq N-1, \tag{5.67}$$

where

$$a_{pk} = \langle \phi_p, \phi_k \rangle_W, \qquad b_p = \langle v, \phi_p \rangle_W. \tag{5.68}$$

There are two differences between this system and the system in Eq. (5.13):

1. There are only $N-1$ unknowns, c_p, $1 \leq p \leq N-1$.
2. The integrals defining the coefficients b_p involve the function v and not u.

An implementation of the "Weberized" L^2 method with the mean matching procedure is presented in the next chapter. We will see that the difference between the "Weberized" L^2 and its matching of the means variation with small number of basis functions will be negligible for natural images. In the case of some medical images, however, the mean matching procedure shows to produce better visual results than the original "Weberized" L^2 approximation method for low quality approximations.

5.4 Logarithmic L^2

The "Weberized" L^2 approximation accommodates the Weber model of perception, in which we weight the error accordingly to the intensity at a given pixel. The "Weberized" L^2 method is more expensive computationally than the original L^2 method. This is the price we pay for incorporating the Weber model of perception into the L^2 approximation. Furthermore, the "Weberized" L^2 approximation is based on the ratio of intensities, which can also be represented using logarithms. The following discussion introduces another "Weberized" L^2 approximation method, using logarithms of functions, which is simpler to implement than the previous "Weberized" L^2 approach.

Let us first emphasize that taking the ratio of functions is the basis for our "Weberized" L^2 approach. As before, we assume that our signals of interest are represented by positive,

bounded functions with support $[a, b] \subset \mathbb{R}$. The extension to 2D signals, i.e., images, is straightforward. Let us define the function space,

$$\mathcal{F} = \{u : [a, b] \to [L, M]\}. \tag{5.69}$$

where L and M are positive numbers, with $L < M$. Clearly, $\mathcal{F} \subset L^2[a, b]$.

If we now consider $u \in \mathcal{F}$ to be a reference function and $v \in \mathcal{F}$ to be an approximation to u, then a "Weberized" L^2 error associated with the approximation $u \approx v$ may be defined as

$$\Delta(u, v) = \left[\int_a^b \frac{1}{u(x)^2}[u(x) - v(x)]^2\, dx\right]^{1/2}. \tag{5.70}$$

As mentioned earlier, $\Delta(u, v)$ is not a metric since it is not symmetric with respect to u and v. We may also consider

$$\Delta(v, u) = \left[\int_a^b \frac{1}{v(x)^2}[u(x) - v(x)]^2\, dx\right]^{1/2} \tag{5.71}$$

and try to construct a metric from these two components. However, by theorem (5.3.1), it is sufficient to consider only one of these components, for example $\Delta(u, v)$, as a measure of closeness between u and v.

Recall equations (5.43) and (5.44),

$$\begin{aligned} \Delta(u, v) &= \left[\int_a^b \left[1 - \frac{v(x)}{u(x)}\right]^2 dx\right]^{1/2} \\ \Delta(v, u) &= \left[\int_a^b \left[1 - \frac{u(x)}{v(x)}\right]^2 dx\right]^{1/2}. \end{aligned} \tag{5.72}$$

In both cases, a ratio involving u and v is involved. According to Eq. (5.1), such ratios allows for Weber's Law to be accommodated. In order to see this, suppose that $u(x) = A$ on $[a, b]$ and $v(x) = A + \epsilon$, $\epsilon > 0$, is an approximation to u. We assume that A and $A + \epsilon$ lie in the greyscale range $[L, M]$. Then

$$\begin{aligned} \Delta(u, v) &= \left[\int_a^b \left[1 - \frac{A + \epsilon}{A}\right]^2 dx\right]^{1/2} \\ &= \sqrt{b - a}\left(\frac{\epsilon}{A}\right). \end{aligned} \tag{5.73}$$

We conclude that if the ratio ϵ/A is kept constant, the error $\Delta(u, v)$ is constant. This is consistent with the Weber's model of perception, Eq. (5.1).

5.4.1 Logarithms as a natural representation for Weber's model

Before we present the mathematical formalism for a logarithmic L^2-based approximation method, we will attempt to justify why logarithms provide a natural representation for Weber's model of perception. We start this discussion with the reference to an article written by B. Forte and E. Vrscay [27] in which the authors described a method of moving from a strict metric defined over so-called fuzzy set functions to a weaker metric that involves integration. The fuzzy set metric employs the d_∞ metric involving Hausdorff distances between α-level sets of functions. The motivation for a new metric was to measure distances between functions taking into consideration the aspects of visual perception.

The details of the transition from fuzzy sets to L^1 is presented in [27]. Below is a summary of the results. As before, assume that image functions are non-negative functions, i.e. $u : X \mapsto \mathbb{R}_g$ where X is the "base space" on which the image functions are defined, and $\mathbb{R}_g \subseteq \mathbb{R}_+$ denotes the greyscale range. Let ν be a measure on $\mathcal{B}(\mathbb{R}_g)$, the σ-algebra of Borel measurable subsets of $\mathbb{R}_g$.

The authors of [27] consider a more general case of a measure $\mu \in \mathcal{B}(\mathbb{R}_g)$. Since we will be integrating over the space X, we consider uniform Lebesgue measure on X, which we denote by m. As before our two image functions are represented by u and v. We now provide some definitions that will be needed to formally present the ideas and the connection to the logarithmic L^2- based approximation.

Definition 5.4.1. *The α-level set of a function u is defined as follows,*

$$[u]^\alpha = \{x \in X, u(x) \geq \alpha\}. \tag{5.74}$$

Definition 5.4.2. *The symmetric difference of two sets A and B is given by,*

$$A \Delta B = A \cup B - A \cap B. \tag{5.75}$$

The set $A\Delta B$ is composed of points on A and B that are not common to A and B. If A and B are subsets of the plane, then the "greater" the overlap between A and B, the smaller the area of their symmetric difference. However, the area of their symmetric difference is the R^2 Lebesgue measure of $A\Delta B$, i.e., $m(A\Delta B)$.

Here, we construct a distance function that is based on the α-level sets of two functions. The distance function will be defined in terms of the symmetric difference between the α-level sets of u and v, i.e.:

$$[u]^\alpha \Delta [v]^\alpha, \alpha \in \mathbb{R}_g. \tag{5.76}$$

Given a measure μ in the pixel space X (which will typically be Lebesgue measure), define the symmetric difference of the $\alpha-$level sets as

$$G(u, v; \alpha) = \mu([u]^\alpha \Delta [v]^\alpha). \tag{5.77}$$

Now let ν be a measure over the greyscale range $\mathbb{R}_g$ and define

$$g(u, v; \nu) = \int_{\mathbb{R}_g} G(u, v; \alpha) d\nu(\alpha) \tag{5.78}$$

In other words, integrate the ν-measures of the symmetric difference of the α-level sets over $\mathbb{R}_g$. An application of Fubini's theorem (details in [27]) allows $g(u, v; \nu)$ to be expressed as an integration over subsets of the pixel space X. In the special case that the μ-measure is Lebesgue (uniform) measure m_g on X,

$$g(u, v; \nu) = \int_{X_u} \nu(u(x), v(x)) dx + \int_{X_u} \nu(v(x), u(x)) dx \tag{5.79}$$

where

$$X_u = \{x \in X, u(x) < v(x)\}, X_v = \{x \in X, v(x) < u(x)\}. \tag{5.80}$$

Eq. (5.79) is simpler than the general result presented in Eq. (20) in [27] since we are using Lebesgue measure. $g(u, v; m_g)$ is a pseudometric on the space of nonnegative L^1 functions on X. In the case that $\nu = m_g$, the uniform Lebesgue measure on $\mathbb{R}_g$,

$$m_g(u(x), v(x)] = v(x) - u(x), \ m_g(v(x), u(x)] = u(x) - v(x), \tag{5.81}$$

and we obtain another form of Eq. (5.79),

$$g(u, v; m_g) = \int_X |u(x) - v(x)| dx = \|u - v\|_1, \text{‘} \tag{5.82}$$

which is the L^1 distance between u and v.

At the time of writing the article, the authors of [27] did not know about the Weber's model of perception but they wrote:

"In principle, the measure ν may be used to define various types of grey scales, e.g., (i) quantized grey levels, where ν consists of a finite set of Dirac measures, (ii) nonuniform distributions which model the varying sensitivities of the human due to different regions of the grey level spectrum."

In this work, we are interested in finding a measure ν that will accommodate the Weber's model of perception. It seems reasonable to assume that this measure will behave continuously with respect to the greyscale (or intensity) level. As such, we shall assume that it can be defined in terms of a density function $\rho(x)$, i.e., $dv = \rho(x)$ such that the ν-measure of the interval $[a, b] \subset \mathbb{R}_g$ is given by

$$\nu([a, b]) = \int_a^b d\nu = \int_a^b \rho(t)dt \tag{5.83}$$

In order to accommodate the Weber's model of perception, such a density function should decrease with t, i.e., lesser weight is assigned to higher intensities.

Suppose that t_1 and t_2, both positive, are two (background) greyscale intensities. Using the Weber's model of perception, the minimum perceived change in intensities, Δt_1 and Δt_2, obey the relationship,

$$\frac{\Delta t_1}{t_1} = \frac{\Delta t_2}{t_2} = C \implies \Delta t_1 = Ct_1 \text{ and } \Delta t_2 = Ct_2. \tag{5.84}$$

Figure 5.4.1 shows this relationship graphically in terms of the density function $\rho(t)$. The area of region A is

$$\int_{t_1}^{t_1+\Delta t_1} \rho(t)dt = \nu([t_1, t_1 + \Delta t_1]), \tag{5.85}$$

the ν-measure of the interval $[t_1, t_1 + \Delta t_1]$. Likewise, the area of region B is

$$\int_{t_2}^{t_2+\Delta t_2} \rho(t)dt = \nu([t_2, t_2 + \Delta t_2]), \tag{5.86}$$

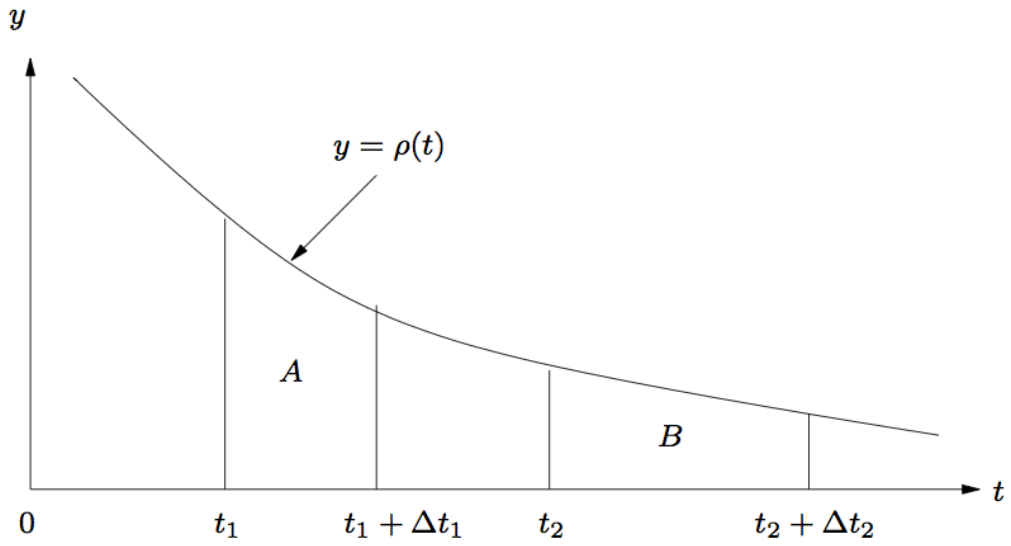

Figure 5.1: Graphical representation of minimum perceived changes in intensities, $\Delta t_1, \Delta t_2$ obeying the Weber's model of perception

We now assume that the area of region A represents a kind of "visual accumulation" of intensity level that contributes to perception. Our goal is to incorporate Weber's law of perception: the area of region A represents the maximum accumulation of intensity levels before a change from level t_1 is perceived. Furthermore, if Weber's law is accommodated, then we propose that

$$\text{area of A} \;=\; \text{area of B} \tag{5.87}$$

Theorem 5.4.3. *The unique density function that satisfies Eq. (5.87) is given by* $\rho(t) = 1/t$.

Proof. For Δt_1 and Δt_2 given in Eq. (5.84), we have

$$\int_{t_1}^{t_1+Ct_1} \rho(t)dt = \int_{t_2}^{t_2+Ct_2} \rho(t)dt. \tag{5.88}$$

Letting $\rho(t) = 1/t$ in Eq. (5.88) gives

$$\text{LHS: } \int_{t_1}^{t_1+Ct_1} \frac{1}{t} dt = \ln(t_1 + Ct_1) - \ln t_1 = \ln(1 + C) \tag{5.89}$$

$$\text{RHS: } \int_{t_2}^{t_2+Ct_2} \frac{1}{t} dt = \ln(t_2 + Ct_2) - \ln t_2 = \ln(1 + C). \tag{5.90}$$

Thus, $\rho t = \frac{1}{t}$ satisfies Eq. (5.87).

We now show that the density function given by $\rho(t) = \frac{1}{t}$ is the unique solution of Eq. (5.87).

Let $F(t)$ be the antiderivative of $\rho(t)$. Using Eq. (5.88), we have:

$$F(t_1 + Ct_1) - F(t_1) = F(t_2 + Ct_2) - F(t_2) \tag{5.91}$$

$$\implies F(t_1 + Ct_1) - F(t_2 + Ct_2) = F(t_2) - F(t_1). \tag{5.92}$$

Upon dividing by $t_1 - t_2$, multiplying by $\frac{1}{1+C}$ and taking the limit as $t_2 \mapsto t_1$, we obtain:

$$F'((1 + C)t_1) = \frac{F'(t_1)}{1 + C}. \tag{5.93}$$

But $F'(t) = \rho(t)$,

$$\implies \rho((1 + C)t_1) = \frac{\rho(t_1)}{1 + C} \tag{5.94}$$

Let $K = 1 + C$.

$$\implies \rho(Kt) = \frac{\rho(t)}{K} \tag{5.95}$$

Assume there exists another solution of the form

$$\rho(t) = \frac{1}{t} \cdot v(t) \tag{5.96}$$

Substituting into Eq. (5.96) gives

$$\rho(Kt) = \frac{1}{K}\frac{v(Kt)}{t} = \frac{1}{K}\frac{v(t)}{t}, t \neq 0 \tag{5.97}$$

$$\implies v(Kt) = v(t). \tag{5.98}$$

This means that $v(t)$ is a homogeneous function of degree 0. Thus Euler's homogeneous function (see Appendix A) theorem applies and

$$pv(t) = v'(t)t, \tag{5.99}$$

where p is the homogeneity degree. Since $p = 0$, we obtain

$$\implies v'(t)t = 0, \forall t. \tag{5.100}$$

$$\implies v'(t) = 0 \text{ since } t > 0. \tag{5.101}$$

Thus $v(t)$ must be a constant.

We conclude that $\rho(t) = \frac{1}{t}$ is the unique solution to Eq. (5.88).

□

Using the measure ν with $\rho = \frac{1}{t}$ as the density function, the distance between u and v in Eq. (5.79) becomes

$$D(u, v; \nu) = \int_{X_u} \left[\int_{u(x)}^{v(x)} \frac{1}{t}\right] dx + \int_{X_v} \left[\int_{v(x)}^{u(x)} \frac{1}{t}\right] dx = \int_X |\ln u(x) - \ln v(x)| dx, \tag{5.102}$$

the logarithmic L^1 distance between u and v. The above discussion hopefully convinces the reader that logarithms provide a natural representation for Weber's model. From the above result, all other logarithmic L^p distances may be viewed as generalizations. The above result was published in [42].

5.4.2 Logarithmic L^2 distance

Looking back at the weighted L^2 metrics in Eqs. (5.20) and (5.21), we see that their accommodation of Weber's model of perception comes from the fact that their integrands involve ratios of the signals/images u and v. A ratio between functions can also be achieved if we consider their logarithms. Taking the logarithm of a signal has been used in image enhancement and is the basis of homomorphic filtering [66]. In our work, however, we are considering logarithms of images for the purpose of image approximation as opposed to image enhancement.

Given two signals $u, v \in \mathcal{F}$, define $U = \log u$ and $V = \log v$, i.e.,

$$U(x) = \log u(x), \quad V(x) = \log v(x), \quad x \in [a, b]. \tag{5.103}$$

Note that *log* represents any base for the logarithm.

Recall from the definition of the space $\mathcal{F}$ that

$$\log L \leq U(x), V(x) \leq \log M\,. \tag{5.104}$$

The usual L^2 distance between U and V is

$$d_2(U, V) = \left[\int_a^b [U(x) - V(x)]^2\, dx\right]^{1/2} < \infty\,. \tag{5.105}$$

Now we define a "Logarithmic L^2 distance" between u and v, i.e.,

$$d_{\log}(u, v) = d_2(U, V) = d_2(\log u, \log v). \tag{5.106}$$

It is easy to see that $d_{\log}$ is a metric on $\mathcal{F}$. First, $d_{\log}(u, v)$ is clearly ≥ 0. Also, $d_{\log}(u, v) = 0 \iff u = v$ we are using the L^2 metric on the logarithms functions, which is continuous and $1-1$. It is also easy to see that $d_{\log}(u, v)$ is symmetric. Let us return to the Logarithmic

L^2 distance between u and v in Eq. (5.106):

$$\begin{aligned} d_{\log}(u,v) &= d_2(\log u, \log v) \\ &= \left[\int_a^b [\log u(x) - \log v(x)]^2 \, dx\right]^{1/2} \\ &= \left[\int_a^b \left[\log \frac{u(x)}{v(x)}\right]^2 dx\right]^{1/2} \\ &= \left[\int_a^b \left[\log \frac{v(x)}{u(x)}\right]^2 dx\right]^{1/2} . \end{aligned} \tag{5.107}$$

In order to verify the triangle inequality, we first define the space

$$\mathcal{G} = \{ f : [a,b] \to [\log L, \log M] \}. \tag{5.108}$$

Then for $u \in \mathcal{F}$ it follows that $U = \log u \in \mathcal{G}$. Furthermore, for any $U \in \mathcal{G}$, the function

$$u(x) = e^{U(x)} \quad \forall x \in [a,b], \tag{5.109}$$

which we shall denote as $u = e^U$, is an element of $\mathcal{F}$.

For any $u, v, w \in \mathcal{F}$, we let $U = \log u$, $V = \log v$ and $W = \log w$. Then

$$\begin{aligned} d_{\log}(u,v) &= d_2(U,V) \\ &\leq d_2(U,W) + d_2(W,V) \\ &= d_{\log}(u,w) + d_{\log}(w,v) . \end{aligned} \tag{5.110}$$

As before, let $u(x) = A$ be the reference signal and $v(x) = A + \epsilon$ an approximation to it. Then

$$\begin{aligned} d_{\log}(u,v) &= \left[\int_a^b \left[\log \frac{A+\epsilon}{A}\right]^2 dx\right]^{1/2} \\ &= \left[\int_a^b \left[\log\left(1 + \frac{\epsilon}{A}\right)\right]^2 dx\right]^{1/2} \\ &= \sqrt{b-a}\, \log\left(1 + \frac{\epsilon}{A}\right) . \end{aligned} \tag{5.111}$$

As before, if the ratio ϵ/A is kept constant, then the distance $d_{\log}(u, v)$ is constant, in accordance with the Weber model of perception, Eq. (5.1).

In the case that the ratio ϵ/A is very small, then

$$\log\left(1 + \frac{\epsilon}{A}\right) \approx \frac{\epsilon}{A}\,, \tag{5.112}$$

so that

$$d_{\log}(u, v) \approx \Delta(u, v)\,, \tag{5.113}$$

the "Weberized" L^2 distance in Eq. (5.73).

Finally, it is also worthwhile to consider the effect of taking the logarithm of a signal. If $U = \log u$, then high values of u are "dampened." It is probably convenient to assume that the lower bound L of our signals is greater or equal to 1. Otherwise, the logarithm produces a "stretching" of the graph of $u(x)$ in the negative direction for values of $u(x)$ between 0 and 1.

5.4.3 Approximation in Logarithmic L^2 distance

We now consider the use of the Logarithmic L^2 distance to approximate a $u \in \mathcal{F}$.

As before, we let $\{\phi_k\}_{k=1}^{\infty}$ denote a complete orthonormal basis for the Hilbert space of functions $L^2[a, b]$ which will therefore serve as a basis for our signals $\mathcal{F} \subset L^2[a, b]$. If we consider the following approximation to u,

$$u \approx u_N = \sum_{k=1}^{N} c_k \phi_k, \tag{5.114}$$

then the squared approximation error in the $d_{\log}$ metric is given by

$$d_{\log}(u, u_N)^2 = \int_a^b \left[\log u(x) - \log \sum_{k=1}^{N} c_k \phi_k(x)\right]^2 dx. \tag{5.115}$$

Setting aside possible complications (e.g., the sum in the above equation assuming negative values at some $x \in [a, b]$), the equations resulting from the stationarity conditions,

$$\frac{\partial d_{\log}(u, u_N)^2}{\partial c_p} = 0, \quad 1 \le p \le N\,, \tag{5.116}$$

are extremely complicated in form, i.e.,

$$\int_a^b \left[\log u(x) - \log \sum_{k=1}^N c_k \phi_k(x) \right] \frac{\phi_p(x)}{\sum_{k=1}^N c_k \phi_k(x)} \, dx = 0. \tag{5.117}$$

These equations are certainly not linear in the unknown coefficients c_k and the problem becomes more complex than solving the simple system of equations as in the "Weberized" L^2 case. An enormous simplification occurs if we consider the approximation of the Logarithmic function $U = \log u$, i.e., removing the logarithms from the approximated function. We then seek to approximate $U \in \mathcal{G} \subset L^2[a, b]$ as follows,

$$U \approx U_N = \sum_{k=1}^N a_k \phi_k, \tag{5.118}$$

by minimizing the squared L^2 distance

$$d_2(U, U_N) = \int_a^b \left[U(x) - \sum_{k=1}^N a_k \phi_k(x) \right]^2 dx. \tag{5.119}$$

The solution to this problem is well known. It is the usual squared L^2 distance and it is minimized by the standard Fourier coefficients of U in the ϕ_k basis, i.e.,

$$a_k = \langle U, \phi_k \rangle = \int_a^b U(x) \phi_k(x) \, dx. \tag{5.120}$$

The sequence of functions U_N will converge to U in L^2 metric, i.e.,

$$\lim_{N \to \infty} d_2(U, U_N) = 0, \tag{5.121}$$

which we shall write as

$$\lim_{N \to \infty} U_N = U. \tag{5.122}$$

The functions U_N are then used to provide approximations to $u = e^U$ as follows,

$$u(x) \approx u_N := \exp(U_N) = \exp\left(\sum_{k=1}^N a_k \phi_k \right). \tag{5.123}$$

Because of the continuity of the exponential function, it follows that

$$\begin{aligned}\lim_{N\to\infty} u_N &= \lim_{N\to\infty} \exp\left(U_N\right) \\ &= \exp\left(\lim_{N\to\infty} U_N\right) \\ &= \exp\left(U\right) \\ &= u. \end{aligned} \tag{5.124}$$

It is interesting to consider the expressions for the approximations $u_N(x)$ in "x-space":

$$\begin{aligned} u_N(x) &= \exp\left(\sum_{k=1}^{N} a_k \phi_k(x)\right) \\ &= e^{a_1\phi_1(x)} e^{a_2\phi_2(x)} \cdots e^{a_N\phi_N(x)} \\ &= \left[e^{\phi_1(x)}\right]^{a_1} \left[e^{\phi_2(x)}\right]^{a_2} \cdots \left[e^{\phi_N(x)}\right]^{a_N}. \end{aligned} \tag{5.125}$$

The last line may also be expressed as follows,

$$u_N(x) = \left[e^{a_1}\right]^{\phi_1(x)} \left[e^{a_2}\right]^{\phi_2(x)} \cdots \left[e^{a_N}\right]^{\phi_N(x)}. \tag{5.126}$$

This provides several ways of computing these approximations.

In summary, the Logarithmic L^2 approximation method is much less computationally expensive that the "Weberized" L^2 approach. First, the Fourier coefficients of the logarithm of the signal are computed, which is the usual procedure in the L^2 approximation. Second the obtained approximation needs to be exponentiated in order to obtain approximation u_N of the function. Moreover, there are no systems of equations to be solved. In the next Chapter, we will see that approximations yielded by the Logarithmic L^2 method are generally slightly better than their "Weberized" L^2 counterparts.

5.5 Best approximation of functions

Let us now return to the scenario of Section 5.1, where $\{\phi_k\}_{k=1}^{\infty}$ denoted a complete orthonormal basis for $L^2[a,b]$. Given a target signal/image $u \in \mathcal{F}$, we consider approximations of the form

$$u \approx u_N = \sum_{k=1} c_k \phi_k \tag{5.127}$$

Once again, the best L^2-approximation is obtained when

$$c_k = \langle u, \phi_k \rangle, \quad 1 \leq k \leq N. \tag{5.128}$$

From the completeness of the ϕ_k-basis, it follows that for this choice of coefficients c_k,

$$d_2(u, u_N) \to 0 \quad \text{as} \quad N \to \infty. \tag{5.129}$$

From (5.34), this implies that the distance between u_N and u in terms of the "Weberized" L^2 metric D_W also goes to zero, i.e.,

$$D_W(u, u_N) \to 0 \quad \text{as} \quad N \to \infty. \tag{5.130}$$

Recall from Eq. (5.121) that in the Logarithmic L^2 approximation

$$D_2(U, U_N) \to 0 \quad \text{as} \quad N \to \infty. \tag{5.131}$$

where $U = \log u$ and $U_N \approx \sum_{k=1}^{N} c_k \phi_k$.

In other words, in the limit $N \to \infty$, the best L^2 approximation to u converges to the best "Weberized" L^2 and the best log L^2 approximation to u. One may then well ask, "So why even consider the "Weberized" L^2 and Log L^2 approaches?" The answer is that we are primarily interested in the finite-dimensional approximations $u \approx u_N$.

This will also be the situation in the finite-dimensional case, i.e., when u is a digitized signal, i.e., $u \in \mathbf{R}^M$. In this case, the orthonormal basis is composed of M functions $\{\phi_k\}_{k=1}^{M}$, and perfect reconstruction is attained when $N = M$, i.e., $u = u_M$. The best L^2, the best "Weberized" L^2 and the best logarithmic L^2 approximations coincide. Again, one may ask, "So why even consider the other approaches?" And once again, the answer is that we are primarily interested in the lower-dimensional approximations to u, i.e., $u \approx u_N$ for $N < M$.

5.5.1 Best constant approximation using the "Weberized" L^2 method

The first is to consider the "best constant approximation" of a function $u(x)$ on $[a, b]$, i.e.,

$$u(x) \approx C, \quad x \in [a, b]. \tag{5.132}$$

It is well known that the best L^2 approximation to $u(x)$, i.e., the value of C which minimizes the distance $d_2(u, C)$ is given by the mean value of $u(x)$ over $[a, b]$, i.e.,

$$C_{L^2} = \bar{u} = \frac{1}{b-a} \int_a^b u(x)\, dx. \tag{5.133}$$

Our goal is to find the best "Weberized" L^2 constant approximation.

In this case, the "Weberized" L^2 error $D_W(u, C)$ will be determined by the two (unsymmetric and squared) distance functions

$$\begin{aligned} D_1(C) &= (\Delta(u, C))^2 = \int_a^b \left[\frac{C}{u(x)} - 1\right]^2 dx, \\ D_2(C) &= (\Delta(C, u))^2 = \int_a^b \left[\frac{u(x)}{C} - 1\right]^2 dx. \end{aligned} \tag{5.134}$$

Proceeding in the usual way, we will impose the following stationarity condition,

$$\frac{\partial D_W}{\partial C} = 0. \tag{5.135}$$

For the moment, we do not assume a particular form for $D_W(u, C)$ in terms of D_1 and D_2 and simply compute the necessary partial derivatives of D_1 and D_2,

$$\begin{aligned} \frac{\partial D_1}{\partial C} &= 2 \int_a^b \left[\frac{C}{u(x)} - 1\right] \frac{1}{u(x)}\, dx \\ &= 2C \int_a^b \frac{1}{u(x)^2}\, dx - 2 \int_a^b \frac{1}{u(x)}\, dx \end{aligned} \tag{5.136}$$

and

$$\begin{aligned}\frac{\partial D_2}{\partial C} &= -2\int_a^b \left[\frac{u(x)}{C} - 1\right]\frac{u(x)}{C^2}\,dx \\ &= -\frac{2}{C^3}\int_a^b u(x)^2\,dx + \frac{2}{C^2}\int_a^b u(x)\,dx. \end{aligned} \tag{5.137}$$

Let us now proceed by computing the critical points of D_1 and D_2:

1. From (5.136), D_1 has a critical point at

$$C = C_1 = \frac{\int_a^b \frac{1}{u(x)}\,dx}{\int_a^b \frac{1}{u(x)^2}\,dx}. \tag{5.138}$$

2. From (5.137), D_2 has a critical point at

$$C = C_2 = \frac{\int_a^b u(x)^2\,dx}{\int_a^b u(x)\,dx}. \tag{5.139}$$

In order to determine the natures of these critical points, we examine the second derivatives of D_1 and D_2. First, note that

$$\frac{\partial^2 D_1}{\partial C^2} = 2\int_a^b \frac{1}{u(x)^2}\,dx > 0, \tag{5.140}$$

for all feasible values of C implying that $C = C_1$ is a global minimum of D_1 on $[L, M]$. (D_1 is a positive quadratic form on $[L, M]$.)

Regarding D_2, we find that

$$\frac{\partial^2 D_2}{\partial C^2}(C) = \frac{6}{C^4}\int_a^b u(x)^2\,dx - \frac{4}{C^3}\int_a^b u(x)\,dx. \tag{5.141}$$

A little algebra shows that

$$\frac{\partial^2 D_2}{\partial C^2}(C_2) = \frac{2}{C_2^3}\int_a^b u(x)\,dx > 0, \tag{5.142}$$

which implies that $C = C_2$ is a local minimum of D_2 on $[L, M]$.

It remains to evaluate D_1 and D_2 at their respective critical points:

$$\begin{aligned} D_1(C_1) &= C_1^2 \int_a^b \frac{1}{u(x)^2}\,dx - 2C_1 \int_a^b \frac{1}{u(x)}\,dx + \int_a^b dx \\ &= (b-a) - \frac{(\int_a^b \frac{1}{u(x)}\,dx)^2}{\int_a^b \frac{1}{u(x)^2} dx} \end{aligned} \tag{5.143}$$

and

$$\begin{aligned} D_2(C_2) &= \frac{1}{C_2^2} \int_a^b u(x)^2\,dx - \frac{2}{C_2} \int_a^b u(x)dx + \int_a^b dx \\ &= (b-a) - \frac{(\int_a^b u(x)dx)^2}{\int_a^b u(x)^2 dx}. \end{aligned} \tag{5.144}$$

If we now wish to consider only $D_1(C) = \Delta(u, C)$ as the error in approximating $u(x)$ with a constant C, then our problem is solved – the best "Weberized" L^2 approximation is

$$u(x) = C_1 = \frac{\int_a^b \frac{1}{u(x)}\,dx}{\int_a^b \frac{1}{u(x)^2}\,dx}. \tag{5.145}$$

The error associated with this approximation is

$$\Delta(u, C_1) = [D_1(C_1)]^{1/2} = \left[(b-a) - \frac{(\int_a^b \frac{1}{u(x)}\,dx)^2}{\int_a^b \frac{1}{u(x)^2} dx}\right]^{1/2}. \tag{5.146}$$

It would be interesting to compare the best L^2 and best "Weberized" L^2 constant approximations for various functions, recalling that the latter would try to accomodate better the portions of the function $u(x)$ with lower magnitudes. In fact, consider the step function,

$$u(x) = \begin{cases} 1, & 0 \le x \le 0.5 \\ 3, & 0.5 < x \le 1. \end{cases} \tag{5.147}$$

First, we prove a claim involving mean values and the best "Weberized" L^2 approximation of the step function $u(x)$.

Claim 5.5.1. *The best "Weberized" L^2 constant approximation is lower than the mean value of $u(x)$ over $[0, 1]$, namely, $\bar{u} = 2$.*

Proof. From Eq. (5.145), the best "Weberized" L^2 constant approximation of $u(x)$ is given by

$$\begin{aligned} u(x) = C_1 &= \frac{\int_{-1}^{0} 1\,dx + \int_{0}^{1} \frac{1}{3}\,dx}{\int_{-1}^{0} 1\,dx + \int_{0}^{1} \frac{1}{9}\,dx} \\ &= \frac{1 + \frac{1}{3}}{1 + \frac{1}{9}} \\ &= \frac{3}{2} \\ &< 2, \end{aligned} \tag{5.148}$$

□

In fact, we can show that this is true in general:

Theorem 5.5.2. *For C_1 in Eq. (5.145),*

$$C_1 \leq \bar{u} = \frac{1}{b-a}\int_a^b u(x)\,dx. \tag{5.149}$$

Proof. We employ the Cauchy-Schwartz inequality,

$$\left|\int_a^b f(x)g(x)\,dx\right| \leq \left[\int_a^b f(x)^2\,dx\right]^{1/2}\left[\int_a^b g(x)^2\,dx\right]^{1/2}, \tag{5.150}$$

where all of the integrals are assumed to exist. In our case, $f(x), g(x) > 0$ so that the absolute values on the left can be removed. For $f(x) = 1$ and $g(x) = 1/u(x)$, we have

$$\begin{aligned} \int_a^b \frac{1}{u(x)}\,dx &\leq \left[\int_a^b 1\,dx\right]^{1/2}\left[\int_a^b \frac{1}{u(x)^2}\,dx\right]^{1/2} \\ &= \sqrt{b-a}\left[\int_a^b \frac{1}{u(x)^2}\,dx\right]^{1/2}. \end{aligned} \tag{5.151}$$

For $f(x) = g(x) = \sqrt{u(x)}$, we have

$$b - a = \int_a^b 1\,dx \leq \left[\int_a^b u(x)\,dx\right]^{1/2}\left[\int_a^b \frac{1}{u(x)}\,dx\right]^{1/2}. \tag{5.152}$$

From Eq. (5.151),

$$\left[\int_a^b \frac{1}{u(x)}\,dx\right]^2 \left[\int_a^b \frac{1}{u(x)^2}\,dx\right]^{-1} \leq b-a. \tag{5.153}$$

From Eq. (5.152),

$$(b-a)^2 \leq \left[\int_a^b u(x)\,dx\right]\left[\int_a^b \frac{1}{u(x)}\,dx\right]. \tag{5.154}$$

From (5.153) and (5.154), along with the fact that $0 < a \leq b$ and $0 < c \leq d$ implies $ac \leq bd$, we have

$$(b-a)^2 \left[\int_a^b \frac{1}{u(x)}\,dx\right]^2 \left[\int_a^b \frac{1}{u(x)^2}\,dx\right]^{-1} \leq (b-a)\left[\int_a^b u(x)\,dx\right]\left[\int_a^b \frac{1}{u(x)}\,dx\right]. \tag{5.155}$$

Therefore,

$$\left[\int_a^b \frac{1}{u(x)}\,dx\right]\left[\int_a^b \frac{1}{u(x)^2}\,dx\right]^{-1} \leq \frac{1}{b-a}\left[\int_a^b u(x)\,dx\right], \tag{5.156}$$

which proves the desired result.

□

5.5.2 Best affine approximation using the "Weberized" L^2 method

We now consider the "best affine approximation" of a reference function $u(x)$ on $[a,b]$, i.e.,

$$u(x) \approx u_{ab}(x) = cx + d, \quad x \in [a,b]. \tag{5.157}$$

As is well known, the minimization of the L^2 distance $d_2(u, u_{ab})$ yields a system of linear equations in the unknowns a and b - the so-called method of *least squares*.

To find the best "Weberized" L^2 affine approximation, we shall continue with the approach of the previous section and minimize the (squared) distance function

$$D_1(a,b) = \Delta(u, u_{ab}) = \int_a^b \left[\frac{cx+d}{u(x)} - 1\right]^2 dx. \tag{5.158}$$

Instead of deriving the linear system of equations in a and b for the above problem, we

will consider the following general squared weighted L^2 distance given in Section 1,

$$W^2 = \int_a^b g(x) \left[u(x) - cx - d\right]^2 \, dx. \tag{5.159}$$

In this way, we can quickly construct the systems of equations for both "Weberized" L^2, where $g(x) = 1/u(x)^2$, and standard L^2, where $g(x) = 1$. We impose the stationarity conditions,

$$\begin{aligned} \frac{\partial W^2}{\partial a} &= -2 \int_a^b g(x)[u(x) - cx - d]x \, dx = 0 \\ \frac{\partial W^2}{\partial a} &= -2 \int_a^b g(x)[u(x) - cx - d] \, dx = 0, \end{aligned} \tag{5.160}$$

to arrive at the following linear system of equations in a and b,

$$\begin{aligned} \left[\int_a^b x^2 g(x) \, dx\right] c + \left[\int_a^b x g(x) \, dx\right] d &= \int_a^b x g(x) u(x) \, dx \\ \left[\int_a^b x g(x) \, dx\right] c + \left[\int_a^b g(x) \, dx\right] d &= \int_a^b g(x) u(x) \, dx. \end{aligned} \tag{5.161}$$

Case 1: Standard L^2 ("least squares") best affine approximation, $g(x) = 1$

$$\begin{aligned} \left[\int_a^b x^2 \, dx\right] c + \left[\int_a^b x \, dx\right] d &= \int_a^b x u(x) \, dx \\ \left[\int_a^b x \, dx\right] c + \left[\int_a^b dx\right] d &= \int_a^b u(x) \, dx. \end{aligned} \tag{5.162}$$

Case 2: "Weberized" L^2 ("weighted least squares") best affine approximation, $g(x) = 1/u(x)^2$

$$\begin{aligned} \left[\int_a^b \frac{x^2}{u(x)^2} \, dx\right] c + \left[\int_a^b \frac{x}{u(x)^2} \, dx\right] d &= \int_a^b \frac{x}{u(x)} \, dx \\ \left[\int_a^b \frac{x}{u(x)^2} \, dx\right] c + \left[\int_a^b \frac{1}{u(x)^2} \, dx\right] d &= \int_a^b \frac{1}{u(x)} \, dx. \end{aligned} \tag{5.163}$$

Some examples:

1. Let $[a,b] = [0,1]$ and $u(x) = 5x + 1$. Since $u(x)$ is affine, we should obtain it as its best approximation in both cases.

   (a) **Case 1:** Standard L^2 "least squares" approximation. After some computation, the linear system of equations in (5.162) is found to be

   $$\begin{aligned} \frac{1}{3}c + \frac{1}{2}d &= \frac{13}{6} \\ \frac{1}{2}c + d &= \frac{7}{2} \end{aligned} \tag{5.164}$$

   The determinant of this system is $D = \frac{1}{12}$ and the solution is $c = 5$ and $d = 1$, as expected.

   (b) **Case 2:** "Weberized" L^2 approximation. After some computation (with the help of MAPLE), the linear system of equations in (5.162) is found to be

   $$\begin{aligned} a_{11}c + a_{12}d &= b_1 \\ a_{21}c + a_{22}d &= b_2, \end{aligned} \tag{5.165}$$

   where

   $$a_{11} = \frac{7}{150} - \frac{2}{125}\ln 2 - \frac{2}{125}\ln 3, \quad a_{12} = a_{12} = -\frac{1}{30} + \frac{1}{25}\ln 2 + \frac{1}{25}\ln 3, \quad a_{22} = \frac{1}{6}, \tag{5.166}$$

   and

   $$b_1 = \frac{1}{5} - \frac{1}{25}\ln 2 - \frac{1}{25}\ln 3, \qquad b_2 = \frac{1}{5}\ln 2 + \frac{1}{5}\ln 3. \tag{5.167}$$

   Computation with MAPLE yields the apparently miraculous result that $c = 5$ and $d = 1$.

2. Let $[a,b] = [0,1]$ and $u(x) = x^2 + 1$.

   (a) **Case 1:** Standard L^2 "least squares" approximation. The coefficients a_{ij} of the linear system are identical to those of Case 1 in the previous example. The coefficients b_j are found to be

   $$b_1 = \frac{3}{4}, \quad b_2 = \frac{4}{3}. \tag{5.168}$$

The solution to the linear system is $c = 1$ and $d = \frac{5}{6}$. Thus the best L^2 affine approximation to $u(x) = x^2 + 1$ is

$$v(x) = x^2 + \frac{5}{6} \approx x + 0.833. \tag{5.169}$$

(b) **Case 2:** "Weberized" L^2 approximation. After some computation (with the help of MAPLE), the coefficients of the the linear system of equations in (5.162) are found to be

$$a_{11} = -\frac{1}{4} + \frac{1}{8}\pi, \quad a_{12} = \frac{1}{4}, \quad a_{21} = a_{12}, \quad a_{22} = \frac{1}{4} + \frac{1}{8}\pi, \tag{5.170}$$

and

$$b_1 = \frac{1}{2}\ln 2, \quad b_2 = \frac{1}{4}\pi. \tag{5.171}$$

The determinant of this system is $D = -\frac{1}{8} + \frac{1}{64}\pi^2$. The solution to the linear system is

$$\begin{aligned} b_1 &= \frac{4(\pi \ln 2 + 2\ln 2 - \pi)}{-8 + \pi^2} \approx 0.904 \\ b_2 &= -\frac{2(4\ln 2 + 2\pi - \pi^2)}{-8 + \pi^2} \approx 0.871. \end{aligned} \tag{5.172}$$

Thus the best "Weberized" L^2 affine approximation to $u(x) = x^2 + 1$ is

$$w(x) \approx 0.904x + 0.871. \tag{5.173}$$

In Figure 5.5.2 are presented plots of (i) the reference function $u(x) = x^2 + 1$ (solid), (ii) the best L^2 affine approximation $v(x) = x + 0.833$ (dot) and (iii) the best "Weberized" L^2 affine approximation $w(x) = 0.904x + 0.871$ (dash). Note that the graph of the "Weberized" L^2 approximation $w(x)$ has a slightly lesser slope than that of the best L^2 approximation $v(x)$. As such, $w(x)$ provides a better approximation to $u(x)$ than $v(x)$ for values of x near 0, where $u(x)$ assumes its lowest values. And $v(x)$ provides a better approximation to $u(x)$ for values of x near 1, where $u(x)$ assumes its highest values. This relative behaviour of $w(x)$ and $v(x)$ is consistent with the "Weberized" L^2 approach.

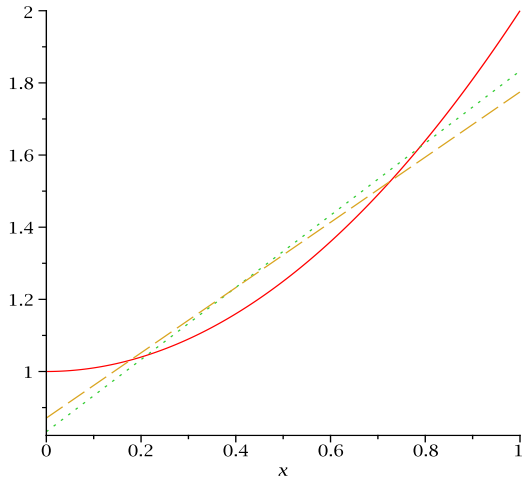

Figure 5.2: Plots of the reference function $u(x) = x^2 + 1$ (solid), the best L^2 affine approximation $v(x) = x + 0.833$ (dot) and the best "Weberized" L^2 affine approximation $w(x) = 0.904x + 0.871$ (dash).

5.5.3 Best constant approximation using the Logarithmic L^2 method

We wish to find the best Logarithmic L^2 constant approximation. i.e.:

$$D_{\log}(C) = \int_a^b (\log u(x) - C)^2 dx = \int_a^b (\log^2 u(x) - 2C \log u(x) + C^2) dx. \tag{5.174}$$

We set

$$\frac{\partial D_{\log}}{\partial C} = 0. \tag{5.175}$$

$$\frac{\partial D_{\log}}{\partial C} = 2 \int_a^b \log u(x) dx + C(b-a) = 0. \tag{5.176}$$

We obtain

$$C = \frac{1}{b-a}\int_a^b \log u(x)dx. \tag{5.177}$$

The second derivatives reveal the nature of this critical point.

$$\frac{\partial^2 D_{\log}}{\partial C^2} = 2(b-a) > 0, \forall C. \tag{5.178}$$

Thus, C is a local minimum of $D_{\log}$ on $[L, M]$.

The value of $D_{\log}$ at the critical point C is:

$$D_{\log}(C) = \int_a^b \Big(\log u(x) - \Big[\frac{1}{b-a}\int_a^b \log u(x)dx\Big] \Big)^2 dx. \tag{5.179}$$

Once again, we return the the step function:

$$u(x) = \begin{cases} 1, & 0 \le x \le 0.5 \\ 3, & 0.5 < x \le 1. \end{cases} \tag{5.180}$$

As in the "Weberized" L^2 constant approximation case, we prove a theorem involving the mean value of Logarithmic L^2 approximation of the step function.

Claim 5.5.3. *The best Logarithmic L^2 approximation is lower than the mean value of $u(x)$ over $[0, 1]$, namely, $\bar{u} = 2$.*

Proof. From Eq. (5.123) , the best Logarithmic L^2 constant approximation of $u(x)$ is given by

$$\begin{aligned} u(x) = e^C &= e^{\frac{1}{b-a}\int_a^b \log u(x)dx} \\ &= e^{\frac{log3}{2}} \\ &= \sqrt{3} \\ &= \approx 1.732 \\ &< 2. \end{aligned} \tag{5.181}$$

□

Theorem 5.5.4. *For C in Eq. (5.177),*

$$e^C \leq \bar{u} = \frac{1}{b-a}\int_a^b u(x)\,dx. \tag{5.182}$$

Proof. We will employ the Jensen's inequality for concave functions,

$$\phi\left(\int_a^b f(x)\,dx\right) \geq \frac{1}{b-a}\int_a^b \phi((b-a)f(x))\,dx \tag{5.183}$$

where $f(x)$ is a non-negative real valued function, ϕ is a concave function on the real line and all of the integrals are assumed to exist.

Using the Jensen's inequality, we obtain:

$$\begin{aligned} e^C &= e^{\frac{1}{b-a}\int_a^b \log u(x)dx} \\ &\leq e^{\frac{1}{b-a}\log\int_a^b u(x)dx} \\ &= \frac{1}{b-a}\int_a^b u(x)dx \end{aligned}$$

which finishes the proof. □

5.5.4 Best affine approximation using the Logarithmic L^2 method

Here, we consider the "best affine approximation" using the Logarithmic L^2 method. Given a function $u(x)$ on $[a,b]$

$$u(x) \approx u_{cd}(x) = cx + d, \quad x \in [a,b]. \tag{5.184}$$

In order to find the "best affine approximation" using the Logarithmic L^2 method, we will proceed as in the previous sections, with minimizing the squared distance function:

$$D^2_{log}(c,d) = \int_a^b \left[\log(u(x)) - (cx+d)\right]^2 dx \tag{5.185}$$

After differentiation, we obtain a system of two equations in two unknowns c and d:

$$\begin{aligned}\frac{\partial D^2}{\partial c} &= 2\int_a^b (\log(u(x)) - (cx+d))x\,dx = 0\\ \frac{\partial D^2}{\partial d} &= 2\int_a^b (\log(u(x)) - (cx+d))\,dx = 0,\end{aligned} \tag{5.186}$$

Then,

$$\begin{aligned}\int_a^b x\log(u(x))dx &= \int_a^b (cx+d)x\,dx\\ \int_a^b \log(u(x))dx &= \int_a^b (cx+d)\,dx.\end{aligned} \tag{5.187}$$

Example: Let $[a,b] = [0,1]$, $u(x) = 3x+1$ and the base of log equal to e. The "best affine approximation" to $u(x)$ using the Logarithmic L^2 method can be obtained by using Eq. (5.187).

$$\begin{aligned}\int_0^1 x\ln(u(x))dx &= \int_0^1 (cx+d)x\,dx\\ \int_0^1 \ln(u(x))dx &= \int_0^1 (cx+d)\,dx.\end{aligned} \tag{5.188}$$

The integrals involving the ln functions have been solved using common formulas from tables of integrals:

$$\begin{aligned}\int x\ln(sx+t)dx &= \frac{tx}{2s} - \frac{1}{4}x^2 + \frac{1}{2}(x^2 - \frac{t^2}{s^2})\ln(sx+t)\\ \int \ln(sx+t)dx &= (x+\frac{t}{s})\ln(sx+t) - x, s \neq 0.\end{aligned} \tag{5.189}$$

After a few steps, we arrive at the following system of equations:

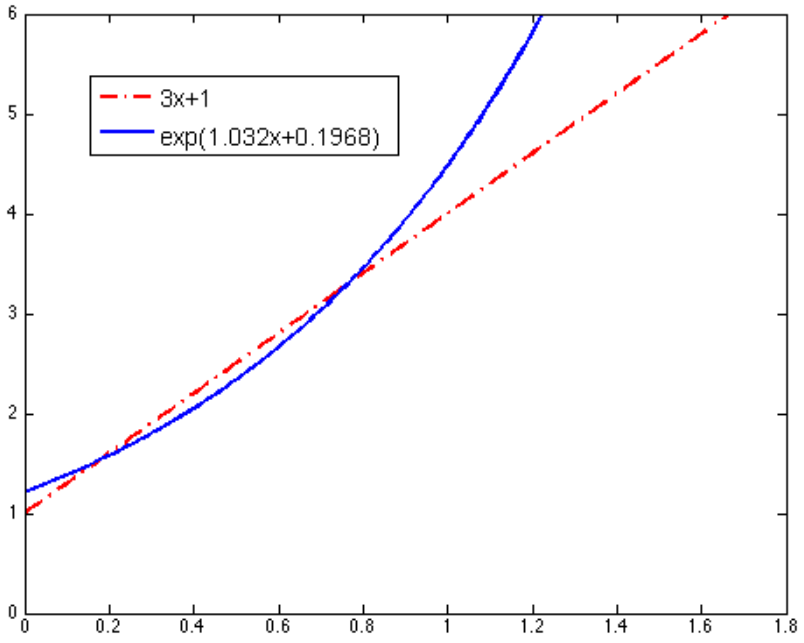

Figure 5.3: Plots of the function $3x + 1$ (dash) and best affine Logarithmic L^2 approximation to $3x + 1$ (solid).

$$\begin{aligned} \frac{1}{3}c + \frac{1}{2}d &= -\frac{1}{12} + \frac{4}{9}\ln 4 \\ \frac{1}{2}c + d &= -1 + \frac{4}{3}\ln 4 \end{aligned} \tag{5.190}$$

Using MATLAB, we obtain: $c = 1.3032$ and $d = 0.1968$. Thus, $\exp(1.3032x + 0.1968)$ is the best affine Logarithmic L^2 fit to the function $3x + 1$.

Perhaps it would be more meaningful to compute the "proper" best Logarithmic L^2 approximation by minimizing the squared logarithmic L^2 distance from Eq. (5.106), i.e.

$$d_{\log}(u, cx + d) = d_2(\log u, \log(cx + d)). \tag{5.191}$$

We impose the stationarity conditions,

$$\begin{aligned} \frac{\partial d^2_{\log}(u, cx+d)}{\partial c} &= 0, \\ \frac{\partial d^2_{\log}(u, cx+d)}{\partial d} &= 0, \end{aligned} \tag{5.192}$$

and arrive at the following set of equations in the unknowns c and d:

$$\begin{aligned} \int_a^b \frac{x \log(u(x))}{cx+d} dx &= \int_a^b \frac{x \log(cx+d)}{cx+d}\, dx \\ \int_a^b \frac{\log(u(x))}{cx+d} dx &= \int_a^b \frac{\log(cx+d)}{cx+d}\, dx \end{aligned} \tag{5.193}$$

The above equations seem a lot more complicated than Eq. (5.187). However, in this case we are approximating the function $u(x)$ with the affine function $cx + d$.

Finding the best affine approximation is closely related to our previous discussions. In this case, the orthogonal basis that is discussed throughout the document, ϕ_k, has been replaced by two functions "1" and "x" which, of course, are not orthogonal but only linearly independent.

Chapter 6

Implementations of the "Weberized" L^2 and Logarithmic L^2 approximation methods

This chapter presents implementations of the intensity weighted function approximation methods presented in the previous chapter. We first illustrate the effect of the "Weberized" L^2 and Logarithmic L^2 by approximating the step function and standard test images. MATLAB code used to compute the approximations is included in Appendix B.

6.1 Approximating the step function using "Weberized" L^2 and Logarithmic L^2

We wish to approximate the step function using three methods discussed in the previous chapter: The usual L^2, "Weberized" L^2 and Logarithmic L^2 approaches. Let

$$u(x) = \begin{cases} 1, & 0 \leq x \leq 0.5 \\ 3, & 0.5 < x \leq 1, \end{cases} \tag{6.1}$$

and the orthonormal basis on $L^2[0,1]$, given by

$$\phi_0(x) = 1, \quad \phi_k(x) = \sqrt{2}\cos(k\pi x), \quad k \geq 1. \tag{6.2}$$

For a given $N > 0$, we compute the best L^2 and best "Weberized" L^2 approximations obtained by minimizing Δ_N^2 and W_2^N in Eqs. (5.5) and (5.7), respectively, using the above cosine basis functions ϕ_k, $0 \le k \le N-1$. We expect that the "Weberized" approximations will provide better fits to $u(x)$ on its left piece. We analyze the approximations using three different number of basis functions for each method.

Case 1: Number of basis functions: $N = 1$

This corresponds to the best constant fit of $u(x)$ over $[0, 1]$. The best L^2 approximation is straightforward and can be easily computed:

$$c_1 = \langle u, \phi_0 \rangle = 2, \quad \Rightarrow \quad u_1(x) = 2, \tag{6.3}$$

the mean of $u(x)$ over $[0, 1]$.

The best "Weberized" L^2 approximation can be found analytically from Eq. (5.12):

$$c_1 \left[\int_0^{1/2} 1 \, dx + \int_{1/2}^1 \frac{1}{9} \, dx \right] = \left[\int_0^{1/2} 1 \, dx + \int_{1/2}^1 \frac{1}{9} \, 3 \, dx \right], \tag{6.4}$$

which implies that

$$c_1 = \frac{6}{5} = 1.2, \quad \Rightarrow \quad u_{1,W}(x) = 1.2. \tag{6.5}$$

The "Weberized" best constant fit is lower than the mean value of $u(x)$ in order to better approximate the lower value $u(x) = 1$ over $[0, 1/2]$.

The best Logarithmic L^2 approximation is computed analytically:

$$c_1 = \langle \log u(x), \phi_0 \rangle = \frac{\log 1}{2} + \frac{\log 3}{2} = 0.55 \tag{6.6}$$

$$u_{1,log_{L^2}}(x) = \exp\left(c_1 \phi_0\right) = 1.73. \tag{6.7}$$

The Logarithmic L^2 best constant fit is lower than the mean value of $u(x)$. As in the "Weberized" L^2 case (although the "Weberized" L^2 methods gives even lower best constant fit), the Logarithmic L^2 method better approximates the lower value $u(x) = 1$ over $[0, 1/2]$.

Cases 2 and 3: Number of basis functions: $N = 5, 20$

Figure 6.1 shows the best L^2, the best "Weberized" L^2 and the best Logarithmic L^2 approximations to $u(x)$ using N-function orthonormal basis sets $\{\phi_0, \cdots, \phi_{N-1}\}$ for the two cases $N = 5$ (left) and $N = 20$ (right). All computations were performed using MATLAB with a grid of 1000 points on $[0, 1]$ to approximate the integrations. In each

case, as expected, the best "Weberized" L^2 and the best Logarithmic L^2 approximations perform a slightly better fit of $u(x)$ on the subinterval $0 \leq x \leq 0.5$ as compared to the best L^2 fit. On the other hand, the "Weberized" L^2 and Logarithmic L^2 approximations will not perform as well as the best L^2 fit on the subinterval $0.5 \leq x \leq 1$.

We now examine the L^2, "Weberized" L^2 error and SSIM index of these approximations. Tables 6.1 and 6.2 present these numerical measures.

N=5	**MSE**	**"Weberized" L^2 error**	**SSIM**
L^2 approx.	0.099	0.055	0.948
"Weberized" L^2 approx.	0.122	0.027	0.937
Log L^2 approx.	0.107	–	0.944

Table 6.1: Quality assessment using MSE, "Weberized" L^2 error and SSIM (with N=5 basis functions)

N=20	**MSE**	**"Weberized" L^2 error**	**SSIM**
L^2 approx.	0.020	0.0112	0.990
"Weberized" L^2 approx.	0.024	0.007	0.988.
Log L^2 approx.	0.22	–	0.989

Table 6.2: Quality assessment using MSE, "Weberized" L^2 error and SSIM (with N=20 basis functions)

As expected, in both cases, the best L^2 approximation yields the lower L^2 error. However, what may be somewhat surprising is that the SSIM values for the best L^2 approximations are better than those for the "Weberized" L^2 and Logarithmic L^2 approximations. After all, SSIM is also formulated to accomodate Weber's model of perception. In order to see this connection, we take a closer look at one of the components of the SSIM function, the luminance function

$$l(f,g) = \frac{2\mu_f\mu_g + C_1}{\mu_f^2 + \mu_g^2 + C_1}, \text{ (note: } -1 \leq l(f,g) \leq 1). \tag{6.8}$$

Assume that f is fixed and g is an approximation to f. Then $l(f,g)$ can be written as follows,

$$l(f,g) = \frac{2(\frac{\mu_g}{\mu_f}) + C_1'}{1 + (\frac{\mu_g}{\mu_f})^2 + C_1'}.$$

This expression of $l(f,g)$ shows a dependence on ratio $\frac{\mu_g}{\mu_f}$, in accordance with the Weber model. However, the factor $l(x,y)$ measures the discrepancy between the mean values of

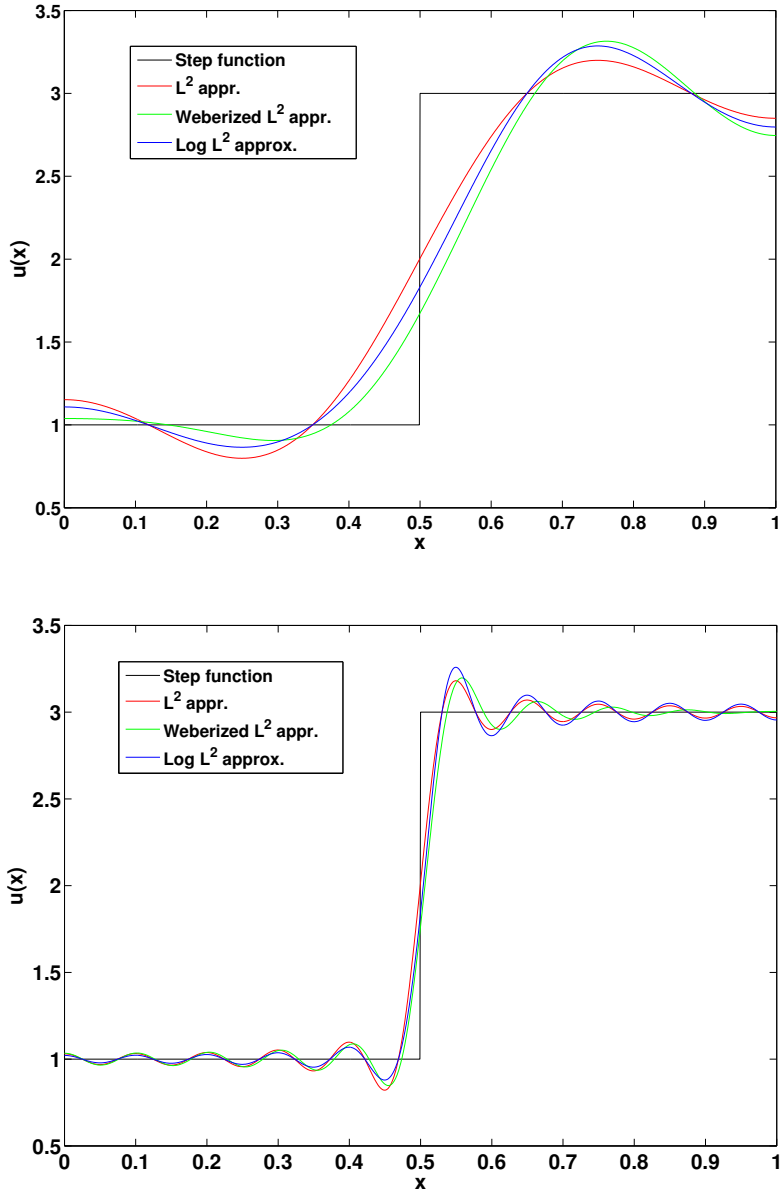

Figure 6.1: L^2, "Weberized" L^2 and Logarithmic L^2 approximations of the step function using $N = 5$ (top) and $N = 20$ (bottom) basis functions.

x and y. With this in mind, we can explain the poorer behaviour of the "Weberized" L^2 method:

1. In the case of the best L^2 approximation u_N always has the same mean value as the function u it is approximating.

2. The above is not necessarily true for the "Weberized" L^2 and the Logarithmic L^2 approximations.

This motivates a reformulation of the "Weberized" L^2 method, in which the mean of the approximation is guaranteed to be the mean of the function u being approximated. This reformulation was discussed in the previous chapter, Section 5.3.1.

The "Weberized" L^2 error was computed for the best L^2 and Weberized L^2 approximations. As expected, the "Weberized" L^2 approximation has the smallest "Weberized" L^2. We omit the computation of the "Weberized L^2 for the Log L^2 approximation due to the latter obtained from a completely different different approach. The use of the "Weberized" L^2 error to compare the Logarithmic L^2 approximation with the other methods studied herein is not well understood.

6.1.1 Comparisons of image approximations using L^2, "Weberized" L^2 and Logarithmic L^2 methods for natural images

This section presents a discussion on the approximation of image functions using the methods described in Chapter 5. For each image (Lena and Mandrill), 8×8 block approximation was performed using 2D DCT, the "Weberized" L^2 approximation (obtained by minimizing the "Weberized" L^2 error) and the log L^2 approximation (minimizing the L^2 error of the logarithm of the signal). Several levels of approximations were used by varying the number of basis functions, N. The approximated images are shown in Figures 6.1.1, 6.1.1 for Lena, and in Figures 6.1.1, 6.1.1 for Mandril. Quality assessment of each approximation was performed and it is presented in Tables 6.3, 6.4, 6.5 and 6.6. As expected, the log L^2 method performs better than the "Weberized" L^2 method according to MSE and the SSIM index. We also show the quality of the "Weberized" L^2 approximation with the matching of the means method implemented.

N=5	**MSE**	**"Weberized" L^2 error**	**SSIM**
L^2 approx.	91.5667	0.0159	0.9231
Web. L^2 approx.	101.3257	0.0115	0.9194
Web. L^2 (m. mean)	96.4424	0.0139	0.9214
Log L^2 approx.	98.9585	–	0.921

Table 6.3: Quality assessment of approximations of Lena using MSE, "Weberized" L^2 error and SSIM (with 8×8 size blocks and $N = 5$ basis functions).

N=20	**MSE**	**"Weberized" L^2 error**	**SSIM**
L^2 approx.	16.3766	0.0027	0.9938
Web. L^2 approx.	17.2275	0.0023	0.9936
Web. L^2 (m. mean)	17.0895	0.0023	0.9936
Log L^2 approx.	17.1163	–	0.9937

Table 6.4: Quality assessment of approximations of Lena using MSE, "Weberized" L^2 error and SSIM (with 8×8 size blocks and $N = 20$ basis functions).

Figure 6.2: Approximations of Lena with $N = 5$ basis functions. Original size of images: 512×512 pixels.

Figure 6.3: Approximations of Lena with $N = 20$ basis functions. Original size of images: 512×512 pixels.

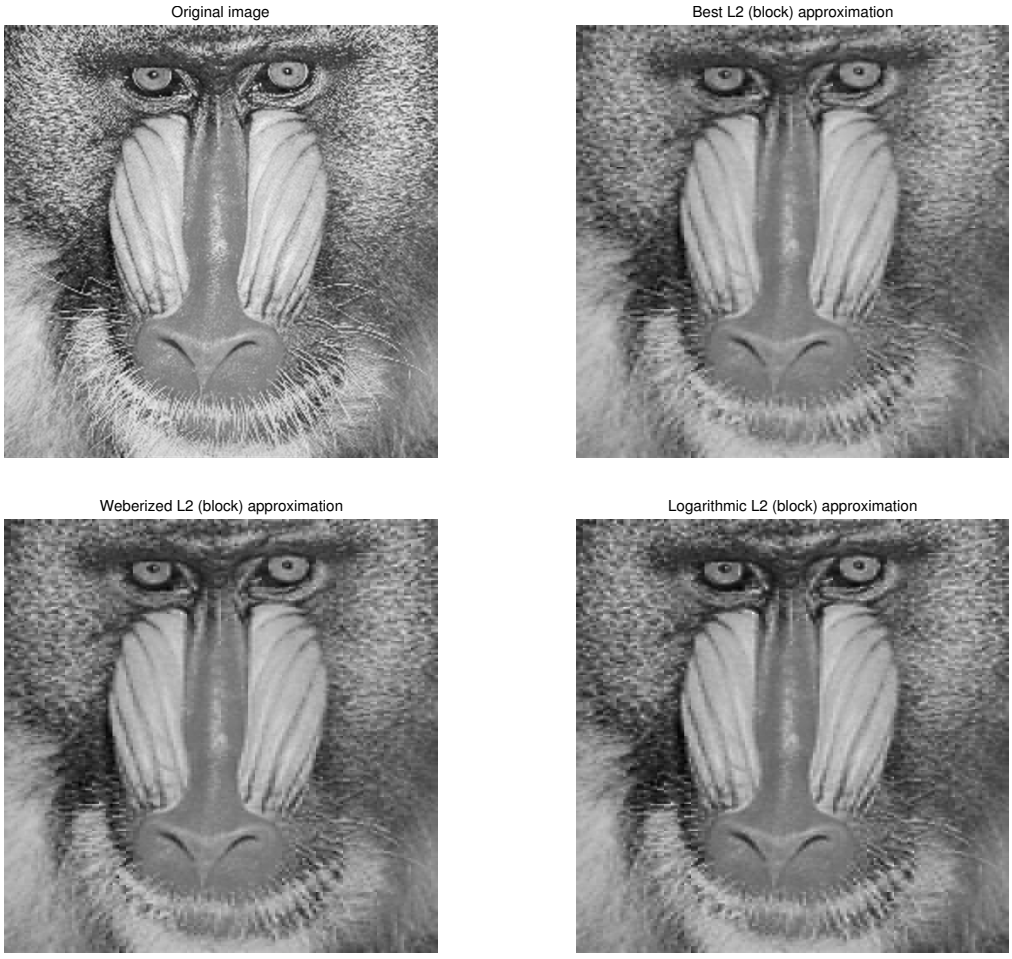

Figure 6.4: Approximations of Mandrill with $N = 5$ basis functions. Original size of images: 512×512 pixels.

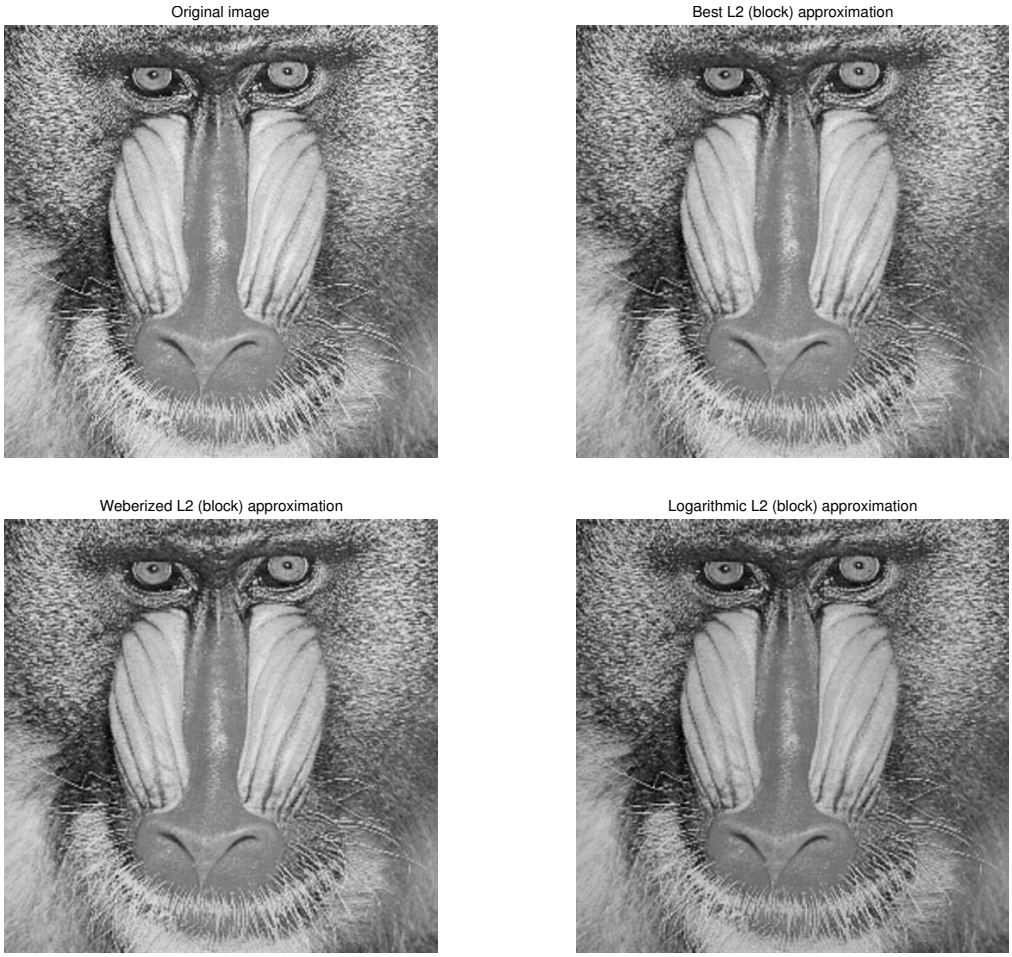

Figure 6.5: Approximations of Mandrill with $N = 20$ basis functions. Original size of images: 512×512 pixels.

N=5	MSE	"Weberized" L^2 error	SSIM
L^2 approx.	254.3766	0.0473	0.8083
Web. L^2 approx.	271.8218	0.0252	0.8038
Web. L^2 (m. mean)	262.7264	0.0303	0.8081
Log L^2 approx.	264.5155	–	0.8050

Table 6.5: Quality assessment of approximations of Mandrill using MSE, "Weberized" L^2 error and SSIM (with 8×8 size blocks and N=5 basis functions).

N=20	MSE	"Weberized" L^2 error	SSIM
L^2 approx.	33.3056	0.0038	0.9935
Web. L^2 approx.	34.6701	0.0031	0.9931
Web. L^2 (m. mean)	34.4869	0.0032	0.9932
Log L^2 approx.	37.1986	–	0.9930

Table 6.6: Quality assessment of approximations of Mandrill using MSE, "Weberized" L^2 error and SSIM (with 8×8 size blocks and N=20 basis functions).

The above examples illustrate the effect of the "Weberized" approximations. The differences are subtle; we do not want to see large differences between the L^2 and the other two approximations as the usual L^2 approximation is already very good. We observe that in darker areas of the images, the "Weberized" L^2 and the Logarithmic L^2 provide better results. This is exactly what is expected. We also note from the above quality assessment tables that the errors do not differ significantly. In the 1 dimensional example of approximating the step function in the previous section, we have seen more significant difference of the error among the three methods. Of course, we expect that if we assign more importance to the low intensity areas then the price that we have to pay is worse approximation of the high intensity areas. The errors according to the quality assessment presented in Tables 6.3, 6.4, 6.5, 6.6 remain similar, although we were hoping that "Weberizing" the L^2 would result in better SSIM estimates.

In order to visually emphasize the effect of "Weberizing" the L^2 approximation, we used 32×32 sized blocks with $N = 78$ with 2 dimensional DCT basis functions, the result is shown in Figure 6.6. Having larger blocks provides more intensity values to be approximated and therefore the weighting will have more effect on the approximation. We used 78 basis functions in this approximation to emphasize the difference of the approximations. It is generally accepted that in the case of L^2 approximation the optimal quality is obtained with the use of 8×8 block sizes. Figure 6.7 shows the result of approximating the Lena image and illustrates the differences between each approximation and the original image

using $N = 43$ blocks. In the next section we will illustrate the differences in approximations using plots of intensities. It is clearly noticeable that in low intensity regions both "Weberized" methods provide better approximation than the L^2 methods. In higher intensity areas, L^2 gives better results. This is again what we expect from weighing the intensities according to the Weber's model of perception. The approximations afforded by the Logarithmic L^2 method are virtually identical to their "Weberized" L^2 counterparts. As such, they display the same kind of "Weberized ringing" over regions with edges separating high and low greyscale intensities, with lesser ringing error over the latter regions.

The matching of means variation of the "Weberized" L^2 approximation is "pushing" the approximation towards the L^2 approximation. Overall, the differences among the errors as presented in the above tables are small. Here, the visual effect that we expect should be somewhere between the L^2 and the "Weberized" L^2 approximations. However, the differences are very subtle; difficult to detect. Please refer to 6.8 for an example of approximation of Lena using this method. A mathematical treatment of the matching of means procedure for the "Weberized" L^2 method was explained in section 5.3.1. Unless otherwise stated, we will be using the "Weberized" L^2 approximation without the matching of means procedure.

If we consider the "Weberized" L^2 error as a measure of quality assessment then we obviously expect to see the lowest errors for the "Weberized" L^2 approximations. The same analogy applies to the L^2 approximation. As discussed above, we would expect that SSIM provides an objective quality assessment for our approximation since it is conceptually consistent with the Weber's model of perception. However, it is strongly related to the mean value of the function, which is used in the usual L^2 approximation. We conclude, although only based on our empirical observation, that SSIM shows stronger dependence on the mean value than to weighting the error according to the Weber's model of perception.

It is somewhat obvious that we should also expect that the more basis functions are used the smaller the difference between these approximations. This can be seen in the above tables 6.3, 6.4, 6.5, 6.6.

Figure 6.6: Best L^2 (top-right), "Weberized" L^2 (bottom-left) and Logarithmic L^2 (bottom-right) approximations to Lena image using $N = 78$ 2D DCT basis functions over 32×32-pixel blocks comprising the shoulder region of Lena image (magnified). Original size of images: 512×512 pixels.

Figure 6.7: Best L^2 (top left), "Weberized" L^2 (middle left) and Logarithmic L^2 (bottom left) approximations to Lena image using $N = 43$ 2D DCT basis functions over 32×32-pixel blocks comprising the shoulder region of Lena image (magnified). Beside each approximation, difference-maps (between each approximation and the original image) are shown. Original size of images: 512×512 pixels.

Figure 6.8: "Weberized" L^2 (left) and "Weberized" L^2 with matching of means (right) approximations to Lena image using $N = 43$ 2D DCT basis functions over 32×32-pixel blocks comprising the shoulder region of Lena image (magnified). Original size of images: 512×512 pixels.

Furthermore, the distribution of variances in Lena is more concentrated toward zero value than that of Mandrill (see Figure 6.9, 6.10). Thus, Lena image has less activity than the Mandrill image and we should expect the L^2, "Weberized" L^2 and Logarithmic L^2 approximations to be further away from each other for the Mandrill image. To illustrate the difference between the approximations, plots of pixel intensities (original versus approximated) are shown in Figures 6.11 and 6.12. As we increase the number of basis functions, the "Weberized" L^2 and Logarithmic L^2 intensities look more like the best L^2 intensities . Of course, in case of better approximations (i.e. more basis functions used) the plots look more like straight lines.

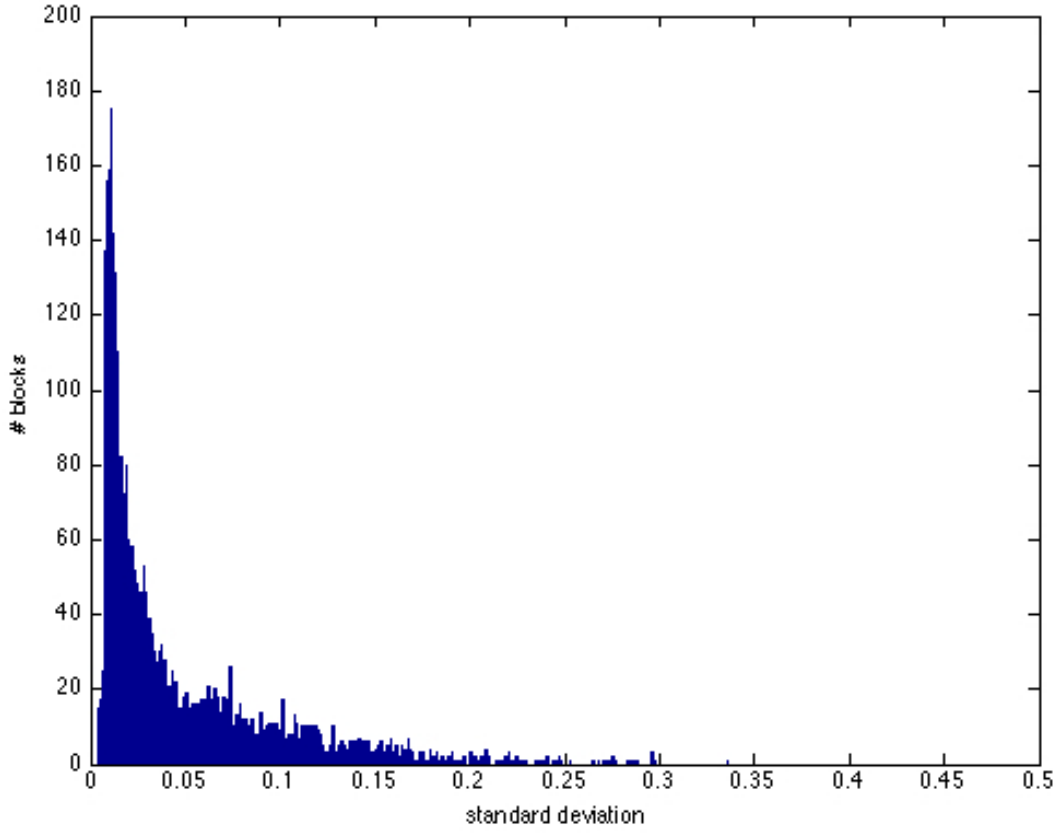

Figure 6.9: Histogram of standard deviations of 8×8 blocks of Lena

.

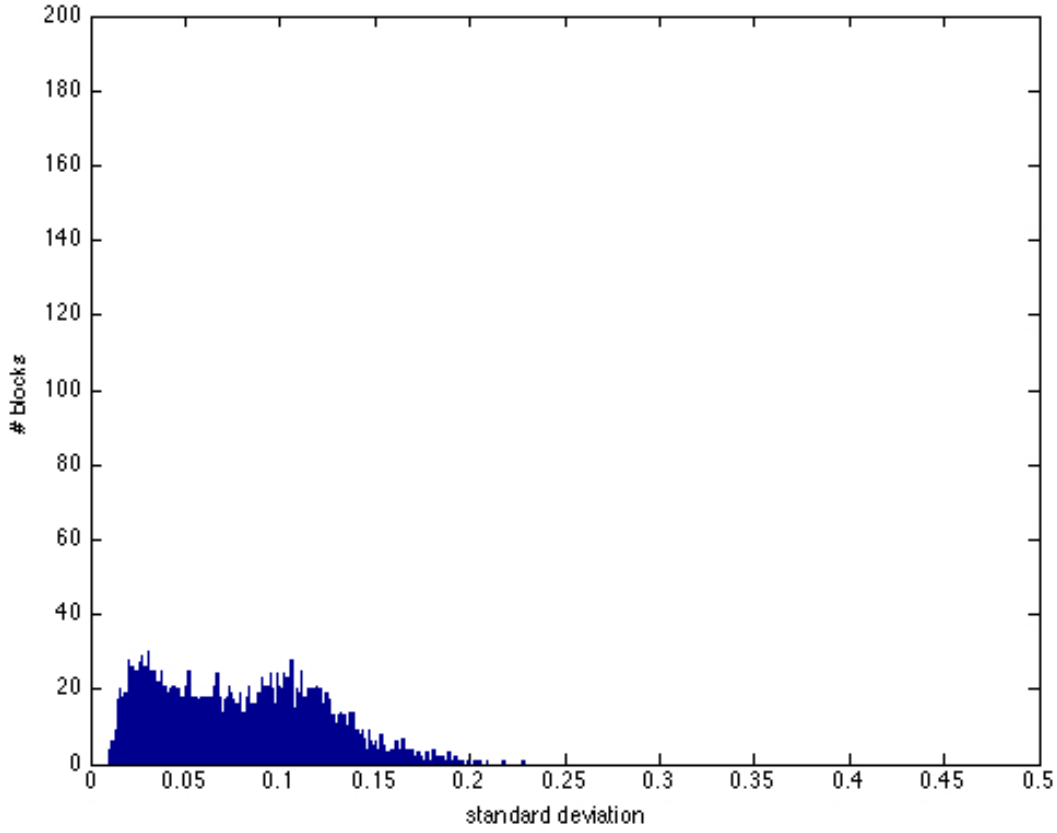

Figure 6.10: Histogram of standard deviations of 8×8 blocks of Mandrill

.

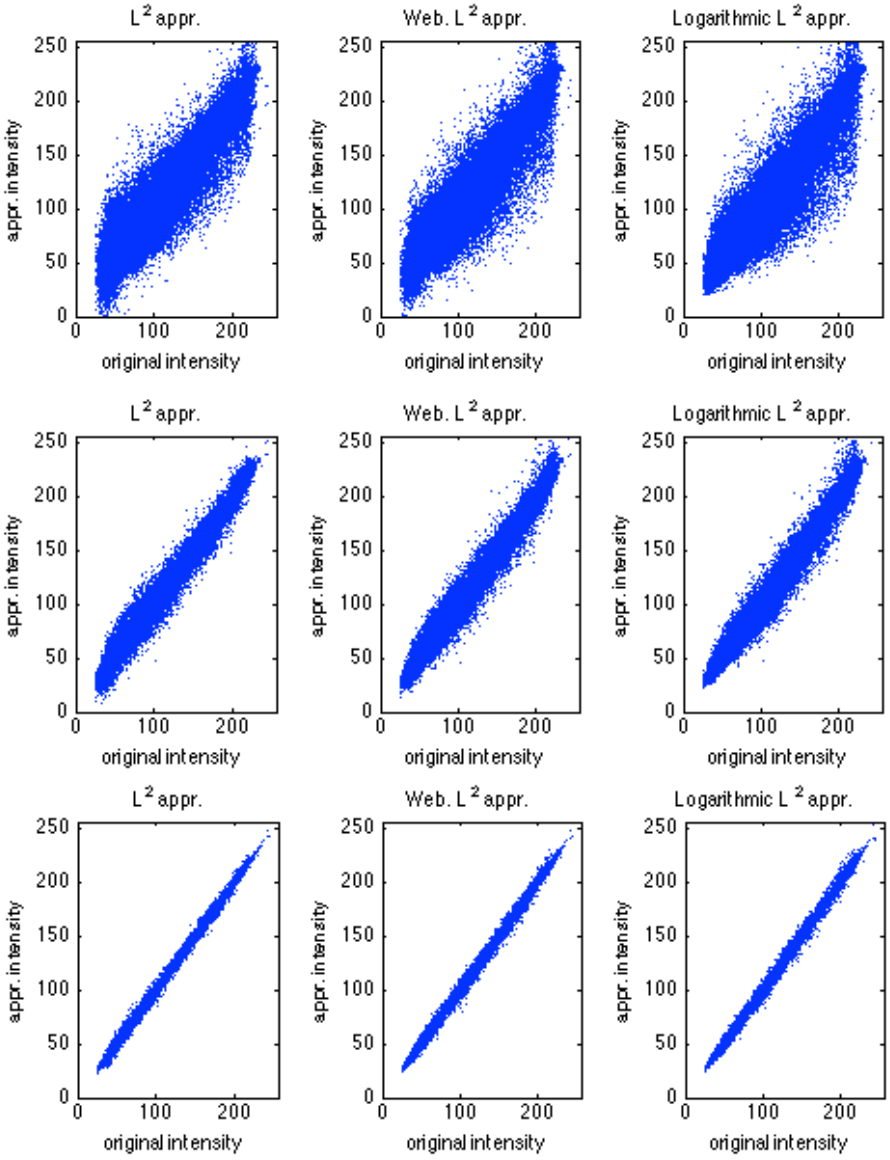

Figure 6.11: Plots of differences between the original intensity and the approximated intensities for Lena with (a) $N = 5$, (b) $N = 20$ and (c) $N = 50$ basis functions.

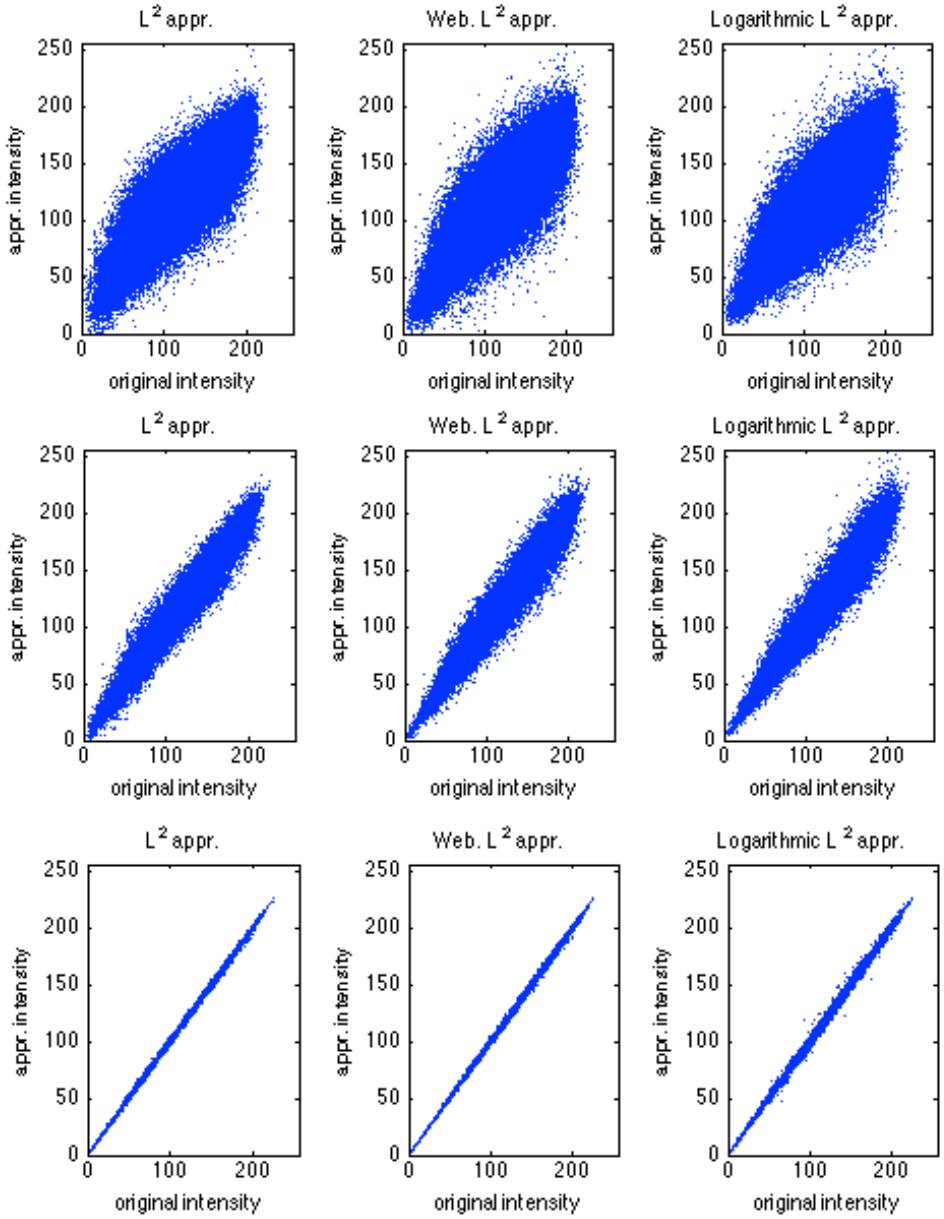

Figure 6.12: Plots of differences between the original intensity and the approximated intensities for Mandrill with (a) $N = 5$, (b) $N = 20$ and (c) $N = 50$ basis functions.

In conclusion, we expect to detect differences in approximation of images having various levels of activities; this is how the intensity weighted methods were constructed (recall the orthogonality of matrix A from Section 5.2.1). The empirical evidence presented in this section confirms the theory that there should be a stronger effect of the intensity weighted methods for images with higher variance.

6.1.2 Comparisons of image approximations using L^2, "Weberized" L^2 and Logarithmic L^2 methods for medical images

In this section we present implementations of the intensity weighted methods on some medical images. We have encountered some problems with approximating the bone-background edges in brain and body CT images using the "Weberized" L^2 method with small number of basis functions. Specifically for extremely high contrast blocks, the "Weberized" L^2 method produces more error than the best L^2 and Log L^2 methods for a small number of basis functions. This effect is shown for the brain CT image in Figures 6.13 and 6.14 together with the corresponding absolute difference maps (between the original and the approximated pixels). This is also reflected in the quality scores, which are presented in Table 6.7. We observe that the "Weberized" L^2 and Log L^2 approximations put less emphasis on higher intensities areas according to the Weber's model of perception. In the case of very poor approximation (i.e. small number of basis functions), the "Weberized" L^2 and Log L^2 approximations approximate the high contrast blocks 8×8 blocks by lower intensity blocks. The Weber's effect can also be seen in the corresponding difference maps in Figures 6.13 and 6.14. The best L^2 approximation is more "uniform" (i.e. the absolute difference map has very similar intensities), whereas in the other two methods, the approximated intensities are weighted according to their magnitudes. According to the SSIM index, the intensity weighted approximations are not as good as the best L^2 approximation. What is causing the poor performance of the Weber-model based algorithms? Could it be that the high contrast blocks (edges of the bone with black background areas) are causing the problems?

Recall that the matching of the mean procedure for the "Weberized" L^2 method was not a significant improvement from the "Weberized" L^2 method for natural images. However, it seems to make a difference in the case of brain CT images. The main difference between natural and medical images is the continuity of grayscale tones. In natural images this change is more continuous whereas medical images may have sudden jumps from white to black pixels. This might be one of the reasons for quite different performance of the intensity weighted approximation methods for some images as the high contrast 8×8

blocks are approximated with lower intensities (i.e. often dark pixels) according to the Weber's model of perception. A similar effect we would observe in the case of body CT images; which also contains blocks of high intensity contrast. Furthermore, there are medical images for which the Weber-model based approximations perform well according to the SSIM index. For example, lung CT images have more continuous tones and are more suitable for the weighted intensity approximations. Approximations of a lung CT image are shown in Figures 6.15 and 6.16 and quality scores are presented in Table 6.8. The contrast of intensities at the edges is not as strong as for brain or body CT images. Table 6.8 shows quality scores for the approximations of the lung CT image in Figure 3.12. We conclude that the intensity based approximation methods studied herein are more suitable for images of continuous tone, i.e. natural images, lung CT, etc. However, brain CT and body CT images are very special having extreme intensity contrast, i.e., almost "black-and-white" blocks and, are more difficult to approximate.

In conclusion, the Weber-model based methods for approximating medical images do not provide satisfactory visual results, which is confirmed by the SSIM scores. However, for certain types of images i.e. lung CT, these methods tend to perform better. These empirical observations are not general enough and require further testing of the weighted intensity approximation methods for brain CT or body CT images, or images of extremely high contrast areas. In this section, we have confirmed that certain types of medical images are more difficult to approximate. These special images are brain CT and body CT images. Recall, that according to Dr. Koff *et al.* [38], brain CT and body CT images behave differently under compression as compared to other medical images (please refer to Chapter 3 for more details on the effects of compression on medical images).

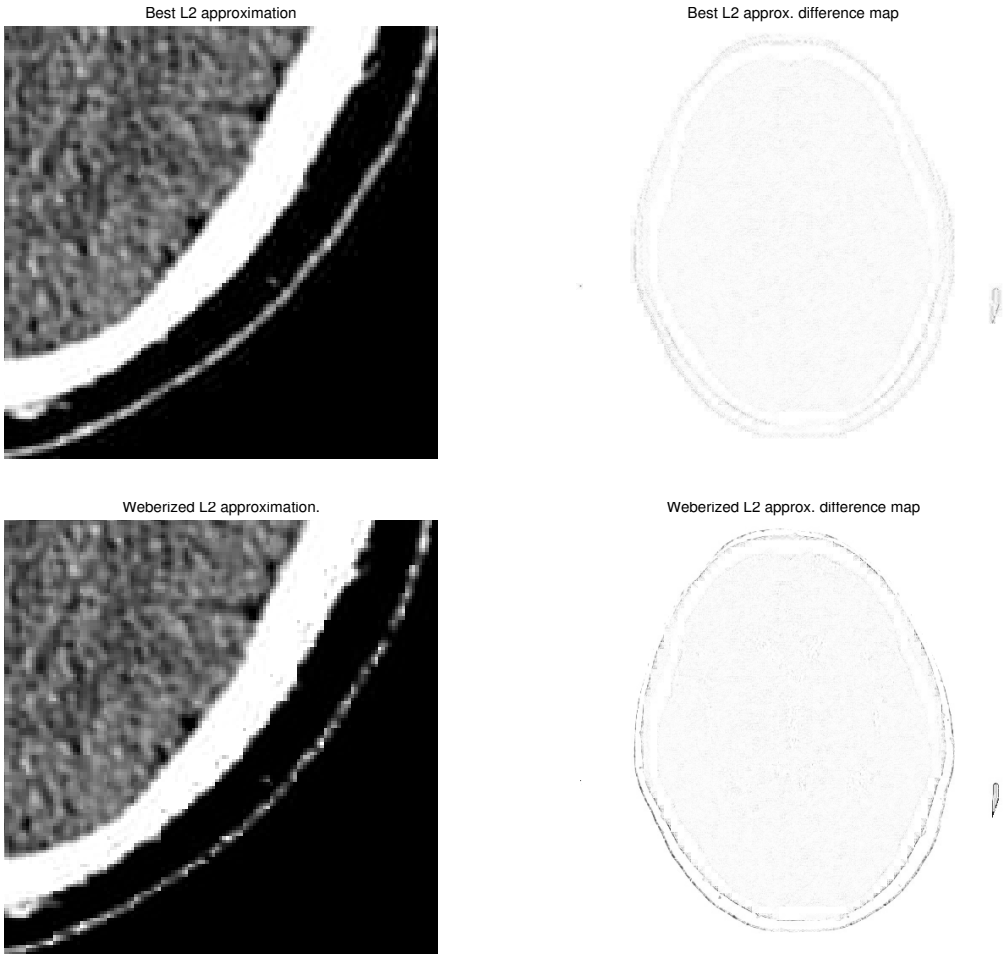

Figure 6.13: Left column: (magnified) approximations of part of the brain CT (Figure 3.2) image using the best L^2, "Weberized" L^2 with 8×8 blocks and $N = 25$ basis functions. Right column: corresponding difference maps. Original size of images: 512×512 pixels.

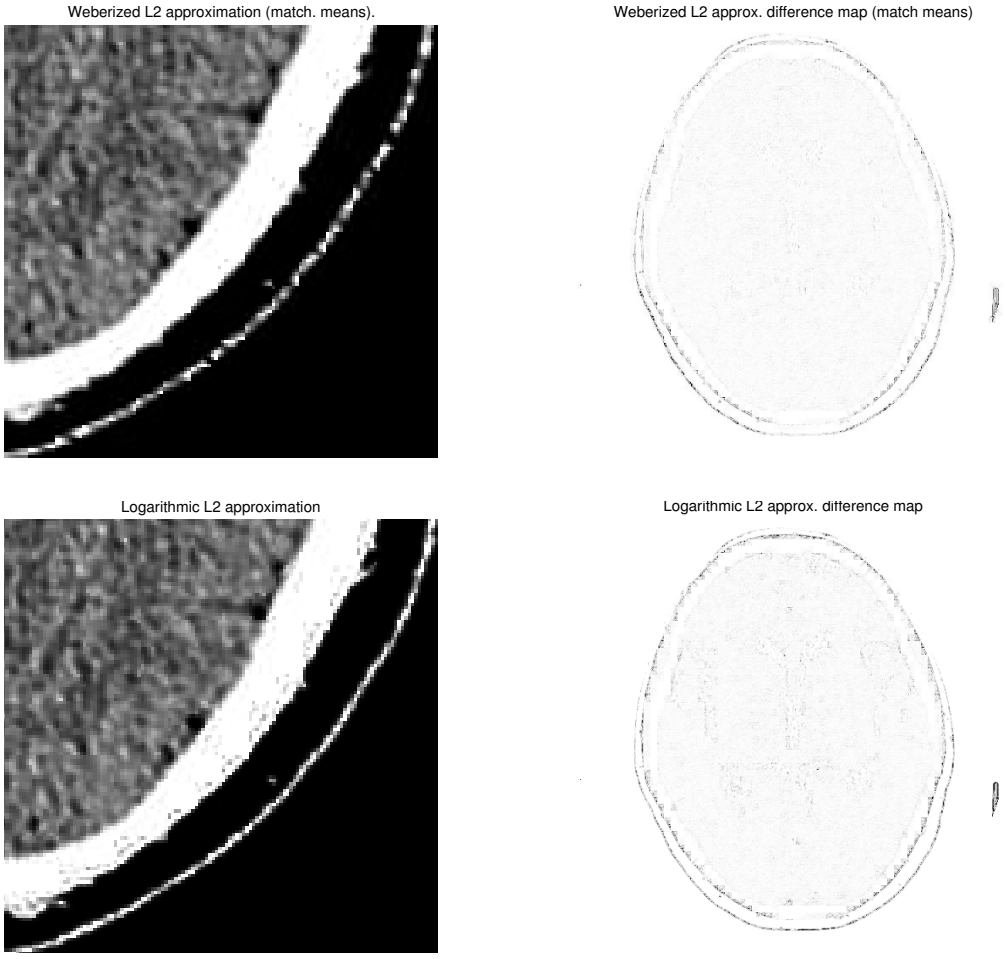

Figure 6.14: Left column: (magnified) approximations of part of the brain CT (Figure 3.2) image using "Weberized" L^2 with matching of the mean procedure and Logarithmic L^2 methods with with 8×8 blocks and $N = 25$ basis functions. Right column: corresponding difference maps. Original size of images: 512×512 pixels.

N=25	MSE	"Weberized" L^2 error	SSIM
L^2 approx.	31.9001	5.6142	0.9963
Web. L^2 approx.	106.8707	0.1318	0.9773
Web. L^2 (m. mean)	102.4284	0.9572	0.9895
Log L^2 approx.	201.6349	–	0.9905

Table 6.7: Quality assessment of approximations of the brain CT image in Figure 3.2 using MSE, "Weberized" L^2 (including the matching of the mean variation) error and SSIM (with 8×8 size blocks and $N = 25$ basis functions).

N=25	MSE	"Weberized" L^2 error	SSIM
L^2 approx.	9.8141	0.0364	0.9964
Web. L^2 approx.	10.3349	0.0020	0.9961
Web. L^2 (m. mean)	10.3261	0.0020	0.9961
Log L^2 approx.	13.4066	–	0.9955

Table 6.8: Quality assessment of approximations of the lung CT image in Figure 3.12 using MSE, "Weberized" L^2 (including the matching of the means variation) error and SSIM (with 8×8 size blocks and $N = 25$ basis functions).

6.2 Conclusion

We have examined two approaches of modifying L^2-based approximations that conform to Weber's model of perception, viz. higher/lower tolerance of deviation for higher/lower intensity levels. The first approach involves the idea of intensity-weighted L^2 distances. We arrive at a natural weighting function that is shown to conform to Weber's model. The resulting "Weberized" L^2 distance involves a ratio of the functions being compared. The importance of ratios in such distance functions leads to a consideration of the well-known logarithmic L^2 distance which also accommodates Weber's model. In fact, it can be shown that logarithms of image functions are the most natural representations with regard to Weber's model.

Moreover, the reason that SSIM measure performs better than L^2-based methods is that it measures the correlation between images, which the L^2-based methods do not do, or at most do indirectly. Nevertheless, it was necessary that the intensity based methods were investigated.

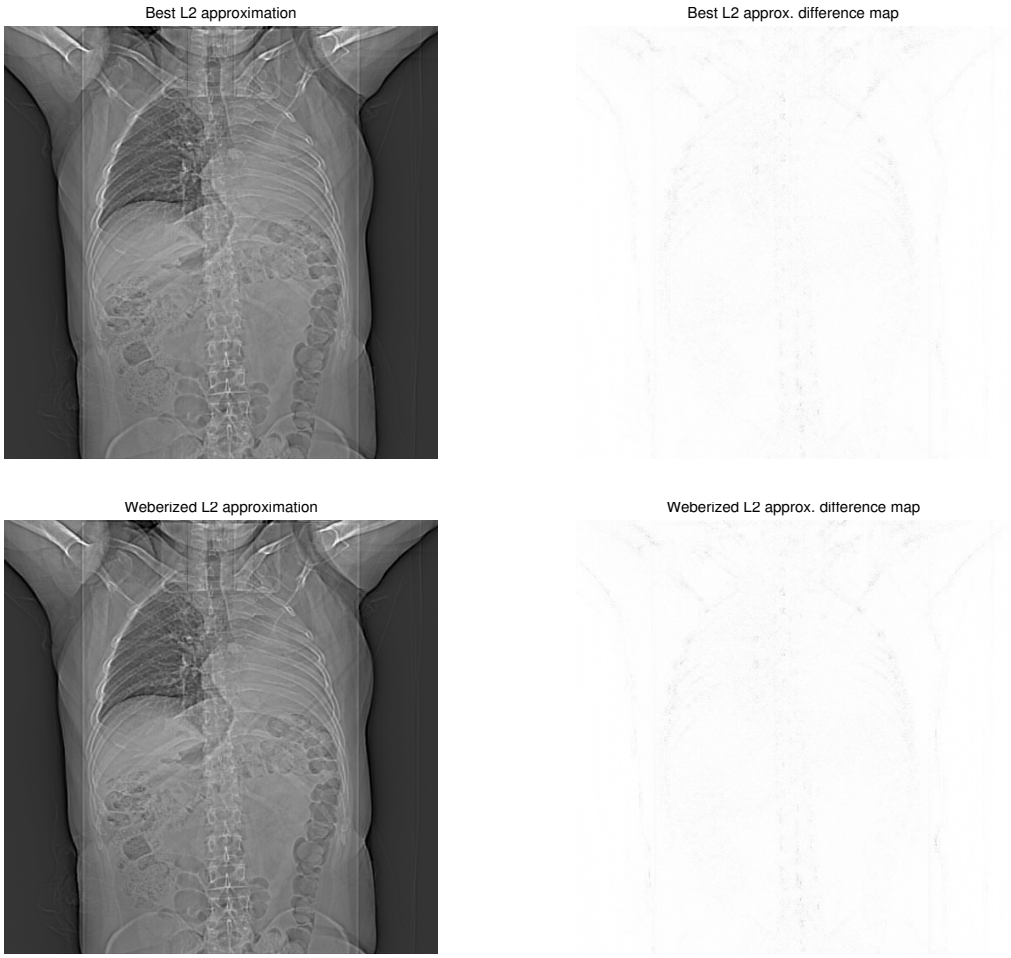

Figure 6.15: Left column: Approximations of part of lung CT image using the best L^2, "Weberized" L^2 methods with with 8×8 blocks and $N = 25$ basis functions. Right column: corresponding difference maps. Original size of images: 512×512 pixels.

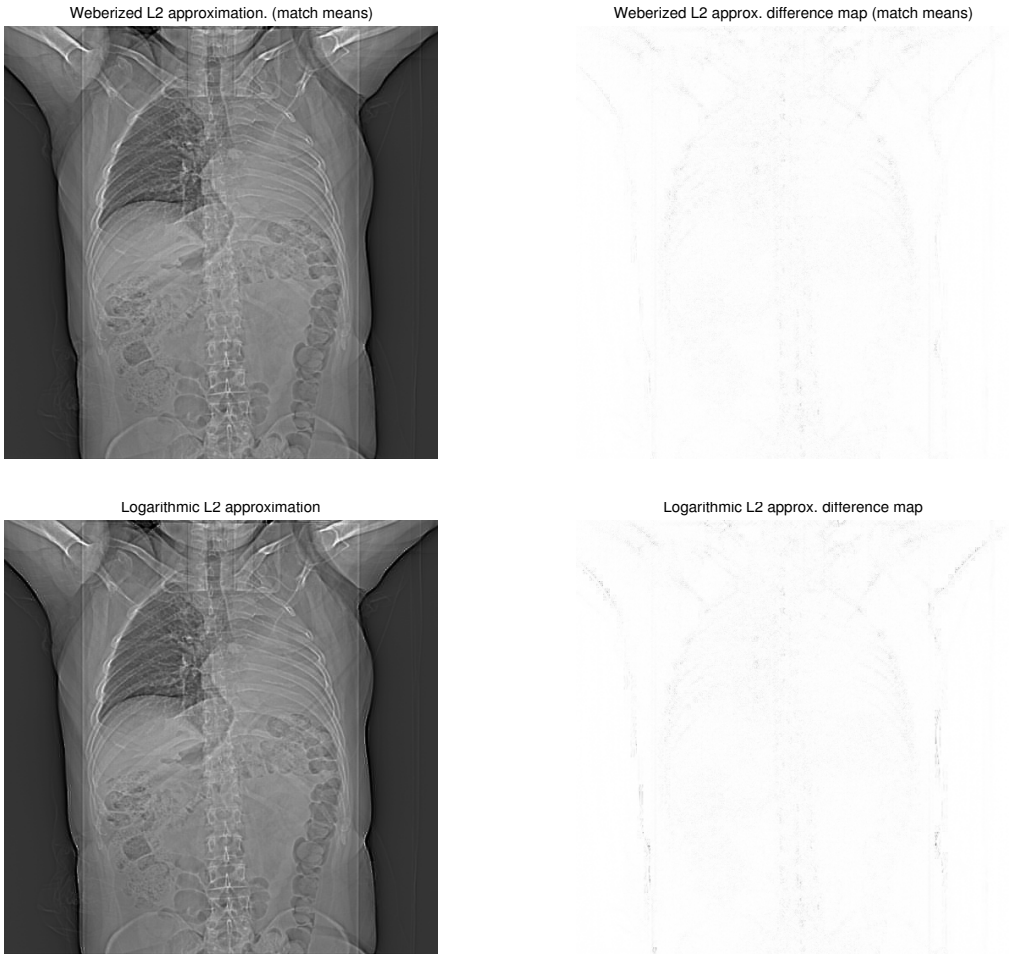

Figure 6.16: Left column: Approximations of part of lung CT image using "Weberized" L^2 with the matching of the means procedure and Logarithmic L^2 methods with with 8×8 blocks and $N = 25$ basis functions. Right column: corresponding difference maps. Original size of images: 512×512 pixels.

Chapter 7

Conclusions

7.1 Quality Assessment of Medical Images

The task of achieving diagnostically lossless compression of medical images is a complex one. It involves tuning technology with radiological subjective responses/preferences. We have used two compression algorithms JPEG and JPEG2000 and several, most commonly used quality assessment algorithms and conducted subjective experiments with radiologists as subjects to assess quality of body CT and brain CT images. The choice of studying JPEG and JPEG2000 compression algorithms comes from the limitation of the DICOM format for medical images.

In particular, we have investigated how the MSE and the SSIM index image quality measures perform in the case of some medical images. By performance of a given quality measure we mean: how closely they match the subjective opinion of a radiologist. Medical images are not just evaluated for their visual quality. By performing a subjective experiment of image quality assessment with radiologists we found that SSIM index best predicts quality of compressed image as compared to MSE, quality factor and compression ratio. Another important result is that the poorest quality predictor for brain CT and body CT images is compression ratio according to subjective radiological assessments. Nevertheless, compression ratio seems to be used as an indicator of image quality for medical images [38] [39]. We are aware that even though SSIM index shows best performance as image fidelity predictor, it will not be feasible to use in certain situations. It is understood that in order to compress images according to the SSIM index, a first pass of compression has to be performed for each individual image and the compression input parameter adjusted accordingly (quality factor for JPEG and compression ratio for JPEG2000). As a further

matter, we are hoping that compression ratio will be at least a satisfactory indicator of quality for specific image types. Perhaps, even within an image group (e.g. brain CT), there are also subgroups that can be categorized where the use of compression ratio will be more appropriate. MSE/PSNR is commonly used to assess the quality of medical images. The results presented in this thesis also indicate that MSE/PSNR is not the best measure to predict image quality.

Medical images posses different properties and should be analyzed separately. Among all types of medical images (body parts as well as various image modalities), body CT and brain CT are the only two types that have shown to behave better under JPEG compression as compared to JPEG2000 algorithm. This oddity has been observed in an experiment conducted by Koff *et al.* [38]. This peculiarity might be caused by the sharp edges in brain CT images between the skull bone and the background. Applying segmentation followed by a morphological operation on brain images before compression improves its JPEG2000 compressibility. A detailed treatment of this topic has been presented in Chapter 3.3.

We are currently in the process of conducting another subjective experiment involving radiological assessment of medical images. This new experiment is an improved version of the one described in Chapter 4. It involves more radiologists and a larger, carefully chosen set of pathological and normal brain and body CT images.

A challenging element of this work, a question that we have posed and still have not answered, involves finding features in medical images that would serve as an indicator of compressibility by employing local analysis of compression artifacts. The details of this investigation have been presented in Section 3.2. Local analysis of compression degradation through SSIM quality maps indicates that there is a variation of quality in a compressed image, i.e., some regions compress better than others. Despite the fact that we see nonuniform compressibility of images according to SSIM quality maps, there does not seem to exist a separator between "good" blocks and "bad" blocks.

7.2 Intensity Weighted Approximations of Images

This thesis also proposes a modification of the L^2-based approach for approximating images that conforms to the Weber's model of perception. In Chapter 5, we have introduced two intensity weighted approximations of images: "Weberized" L^2 and Logarithmic L^2 methods. These methods behave similarly to the best L^2 approximation for some natural images according to MSE and SSIM quality measures. An involved mathematical treatment of this topic is presented in Chapter 5. The "Weberized" L^2 distance involves a ratio of

functions. The well-known logarithmic L^2 distance also shows to conform to Weber's model. Our result shows that the logarithmic L^1 distance is the most natural "Weberized" image metric. From this result, all other logarithmic L^p distances may be viewed as generalizations.

7.3 Final remarks

We are hoping that this study can be useful in defining what diagnostically lossless compression means and that our techniques can be helpful in determining appropriate levels of compression. We are aware that having more available resources in the experiment would provide more statistically significant results. However, the experiment presented in this thesis can be thought of as a "pilot study", which has led to its improved version. In the currently conducted subjective experiment with radiologist we were able to avoid some previous constraints so that the statistical significance is maximized.

The intensity weighted image approximations - the modified L^2 approaches presented in this thesis may serve as a basis for new compression algorithms.

7.4 Future directions

There are several natural extensions of this work that can be considered. One of them is to conduct a subjective experiment with radiologists to compare other quality assessment methods for several modalities and body parts as well as investigation of other lossy compression algorithms. This consideration would lead to statistically valid results including advanced techniques of determining threshold values with regards to a given quality measure that could serve as a reliable quality predictor for compressed medical images.

Another possible direction could be tuning of the SSIM index to correspond better with the radiological opinions. There are several ways to achieve this. One is to determine the best constants for the algorithm. Another one is to consider the IW-SSIM index for medical images and modify the second stage of the algorithm, the pooling procedure in order to achieve the best correspondence with subjective radiological responses.

Our weighted intensity approximation based on the Weber model of perception could possibly be incorporated into an already developed compression algorithm. Consideration of the "Weberized" L^2 approach for high dynamic range (16-bit) images is also a possible extension of this work.

Another future path that could be considered is to study interpolation algorithms and quality assessment of interpolated medical images. It involves producing a larger version of an image, a high-resolution image, by adding new pixels to the existing image. This is particularly important in the radiological community because images are often magnified during diagnostic evaluations. The interpolated pixels are obtained by convolution of a linear interpolation filter. This involves linearly combining the known pixels with some weighted functions that satisfy certain properties and are known as convolution kernels. Even for the same image, different interpolation techniques could produce images that differ significantly. Interpolation is only an approximation and therefore an image will always undergo some loss-of-quality when interpolation is performed. Convolution with an interpolation filter can be thought of as a weighted averaging and therefore produces blurring. Other artifacts include ringing and aliasing. Most interpolation techniques were designed for general sets and therefore they do not necessarily correspond well to human visual perception. Objective quality assessment of such images is a difficult problem since there there is no one-to-one mapping between the original image and the interpolated image. Several workarounds have been suggested in literature [53] [22] [92]; also there has been some advances in the development of no reference quality assessment algorithms [76] [56]. In summary, performing an objective comparison of interpolation algorithms is a difficult task and it is still unknown how to make full-reference objective quality assessment most reliable. A subjective study of interpolation algorithms involving radiologists' responses may be helpful in determining the best interpolation algorithm for given types of medical images.

Another direction is to extend the idea of intensity weighted image approximation to transform domains (e.g. DCT, Fourier and wavelet domains) in light of the fact that the general principle of the Weber's law is applicable to other perceptual quantities such as image contrast. This may lead to perceptually meaningful cost functions that can be directly employed to improve transform-domain image compression algorithms.

APPENDICES

APPENDIX A

7.4.1 Euler's Homogeneous Function Theorem

Assume that $f(x)$, a homogeneous function of order n satisfies

$$f(tx) = t^n f(x) \tag{7.1}$$

Let $x^{'} = xt$. Then

$$\begin{aligned} nt^{(n-1)} f(x) &= \frac{df}{dx^{'}} \frac{dx^{'}}{dt} \\ &= x \frac{df}{dx^{'}} \\ &= x \frac{df}{d(xt)} \end{aligned} \tag{7.2}$$

Let $t = 1$, then

$$x \frac{df}{dx} = nf(x). \tag{7.3}$$

Appendix B

```
%%%%%%%%%%%%%%%%%
% Intensity weighted approximations of images using
% Best L^2, Weberized L^2, Ln L^2 methods.
%
%%%%
close all;
clear all;

%% Reading in images

%natural images
%%%%%%%%%%%
filename_lena='lena_gray_512.tif';
filename_mandrill='mandril_gray.tif';
filename_goldhill='goldhill.tif';

%imjpgorig = double(imread(filename_goldhill, 'tif'));
%imjpgorig = double(imread(filename_lena, 'tif'));
%imjpgorig = double(imread(filename_mandrill, 'tif'));


%read in DICOM medical image
dicom_object = dicominfo('dcm/brainCT.dcm');
%dicom_object = dicominfo('dcm/1.dcm');

%get default window center and window parameters from the dicom
    image
window_center = dicom_object.WindowCenter(1);
window_width = dicom_object.WindowWidth(1);
matrix_16 = dicomread(dicom_object);

%Get the 8-bit matrix representation of the image file
Intercept = dicom_object.RescaleIntercept;
Slope = dicom_object.RescaleSlope;

```

```
35 matrix_16 = double(int16(matrix_16).*Slope + Intercept);
36 matrix_8 = sixteen2eight(matrix_16, window_center, window_width);
     %Medical image
37 imjpgorig = double(matrix_8);
38
39
40 %Select the iterior of the brainCT image
41 %imjpgorig = imjpgorig(110:333, 140:363);
42 %imshow(imjpgorig)
43
44
45 %% set up variables
46 char flag;
47 origsize = size(imjpgorig);
48 NP = origsize(1) * origsize(2);
49 siz = 8;
50
51 imblock = zeros(siz);
52 M = siz^2; %number of points
53
54 %number of basis functions used in approximation
55 N=20;
56
57 eps = 0.9;
58 a = zeros(N,N);
59 b = zeros(N, 1);
60 cw = zeros(N,1);
61 g= zeros(M,1);
62 ssim_k = [0.01 0.03];
63
64 %count for indexing the display of the matrix Aa
65 count = 1;
66
67 %Allocate space for the final image
68 imfin = zeros(siz);
69 imfinal = zeros(floor(sqrt(origsize)));
70
71 %% initialize the DCT basis
```

```
% 2D DCT basis
% All entries are stored in 1D vector
% M = siz^2 (siz=8, M=64 for 8X8 DCT blocks)
% N is the number of basis functions we wish to use

phi = zeros(M,N);

% construction of upper left diagonal of DCT basis set
% (0,0), then (0,1),(1,0), .... (0,8),(1,7), ...(8,0)

      kc=0;
      for kk=1:siz
      for k=1:kk
      l=kk+1-k;
      kc=kc+1;
        if kc <= N
      ic=0;
      for mm=1:siz
        fact1=1;
        if k==1
          fact1=1/sqrt(2);
        end
      term1 = fact1*cos((pi/siz)*((mm-1)+0.5)*(k-1));
      for nn=1:siz
        fact2=1;
        if l==1
          fact2=1/sqrt(2);
        end
      term2 = fact2*cos((pi/siz)*((nn-1)+0.5)*(l-1));
      ic=ic+1;
      phi(ic,kc)= 2*term1*term2/siz;
      end
      end
        end

      end
      end

```

```
% construction (if necessary) of lower diagonal DCT basis

    for kk=1:siz-1
    for l=kk+1:siz
    k=siz+1+kk-l;
    kc=kc+1;
      if kc <= N
    ic=0;
    for mm=1:siz
      fact1=1;
      if k==1
        fact1=1/sqrt(2);
      end
    term1 = fact1*cos((pi/siz)*((mm-1)+0.5)*(k-1));
    for nn=1:siz
      fact2=1;
      if l==1
        fact2=1/sqrt(2);
      end
    term2 = fact2*cos((pi/siz)*((nn-1)+0.5)*(l-1));
    ic=ic+1;
    phi(ic,kc)= 2*term1*term2/siz;
    end
    end
      end

    end
    end

%Compute log of the signal
lnimjpgorig = log(imjpgorig+1);

%% L^2 approximation
%compute approximation for each 8x8 block
  for ind1=1:siz:origsize
count_index = 0;
      for ind2 = 1:siz:origsize

```

```
148 imblock = imjpgorig(ind1:ind1+siz-1, ind2:ind2+siz-1);
149 lnimblock = lnimjpgorig(ind1:ind1+siz-1, ind2:ind2+siz-1);
150 count_index = count_index + 1;
151
152 %compute coefficients
153 for k=1:N
154 yf1(:,k) = imblock(:).*phi(:,k);
155 c = sum(yf1);
156 end
157
158 %approximate the function
159 for xind=1:M
160 ys(xind,:) =c(:)'.*phi(xind,:);
161 end
162 yf = sum(ys,2);
163 yf = reshape(yf,siz,siz);
164 %%
165 if (ind2 == 1)
166     yl2 = yf;
167 else
168     yl2 = [yl2 yf];
169 end
170
171
172 nimjpg = imblock(:);
173
174 [var_l2(count_index) l2_ssim_block(count_index) l2_rmse(
      count_index)] = stsim(double(imblock), double(yf), siz);
175
176
177 %% Ln L^2 approximation
178 %compute coefficients 1D
179 for k=1:N
180 lnyf1(:,k) = lnimblock(:).*phi(:,k);
181 lnc = sum(lnyf1);
182 end
183
184 %approximate the function
```

```
for xind=1:M
ys(xind,:) =lnc(:)'.*phi(xind,:);
end
lnyf = sum(ys,2);
lnyf = reshape(lnyf,siz,siz);

if (ind2 == 1)
    lnyl2 = lnyf;
else
    lnyl2 = [lnyl2 lnyf];
end

[var_lnl2(count_index) lnl2_ssim_block(count_index) lnl2_rmse(
    count_index)] = stsim(double(imblock), double(exp(double(lnyf
    ))+1), siz);

lnnimjpg = lnimblock(:);


%% Weberized L^2 approximation
ap = zeros(M,N,N);
a = zeros(N,N);
b = zeros(N, 1);
cw = ones(N,1);
g= zeros(M,1);
ysw = zeros(M, N);

for xind1=1:M
        if (nimjpg(xind1) == 0)
            g(xind1) = 1/(nimjpg(xind1)+ eps); %.^(2);
         else g(xind1) = 1/nimjpg(xind1); %.^(2);
        end
end


for p=1:N
    for k=1:N
        ap(:,k,p)=g(:).*phi(:,p).*phi(:,k);
```

```
221             a(k,p) = sum(ap(:,k,p));
222         end
223
224     end
225
226     for p=1:N
227         b(p) = sum(g(:).*nimjpg(:).*phi(:,p));
228     end
229
230
231     %% Matching the mean for the Weberized L^2 approximation
232     %Compute the mean
233     imMean = c(1);
234     sum(nimjpg)/64;
235     aMean = a(1:length(a),2:length(a));
236     bMean = b - a(:,1)*imMean;
237
238     %least square
239     cw1 = aMean\bMean;
240     cw(1) = imMean;
241     cw(2:length(a)) = cw1;
242
243     %Conjugate gradient method to solve a system of linear equations
244     [cw2, flag] = pcg(a,b);
245
246     cw';
247
248     %%%%%%%%%%
249     %No matching the mean
250     %Conjugate gradient method to solve a system of linear equations
251     %[cw, flag] = pcg(a,b);
252
253     for xind=1:M
254     ysw(xind,:) =cw(:)'.*phi(xind,:);
255     end
256
257     yfinw = sum(ysw,2);
258
```

```
259   if (ind2 == 1)
260       yfinw = reshape(yfinw,siz,siz);
261       imfin = yfinw;
262
263
264   else
265        yfinw = reshape(yfinw,siz,siz);
266        imfin = [imfin yfinw];
267   end
268
269
270  [var_webl2(count_index) webl2_ssim_block(count_index) webl2_rmse(
       count_index)]  = stsim(double(imblock), double(double(yfinw)),
        siz);
271
272  %% Reshape the original image
273   if (ind2 == 1)
274       yorig = reshape(imblock,siz,siz);
275       yo = yorig;
276
277
278   else
279        yorig = reshape(imblock,siz,siz);
280        yo = [yo yorig];
281   end
282  end
283
284  %% Weber block combining
285    if (ind1 == 1)
286         imfinal = imfin;
287      else
288         imfinal = [imfinal; imfin];
289    end
290
291  %% L2 block combining
292  if (ind1 == 1)
293       yl2final = yl2;
294  else
```

```
295      yl2final = [yl2final; yl2];
296  end
297
298  %% Ln L2 block combining
299  if (ind1 == 1)
300      lnyl2final = lnyl2;
301  else
302      lnyl2final = [lnyl2final; lnyl2];
303  end
304
305  %% Orig block combining
306  if (ind1 == 1)
307      yorigfin = yo;
308  else
309      yorigfin = [yorigfin; yo];
310  end
311
312  %sum of the coefficients of a
313  Aa(count) = sum(sum(abs(a)));
314  count = count +1;
315
316  end
317  %end of computing approximation of each 8x8 block
318
319
320  %% DISPLAY results
321  display('# of Basis functions N=');
322  N
323
324  display('--------MSE---------');
325  msecL2 = mean2(   (double(imjpgorig) - double(yl2final)).^2);
326  msecW = mean2((double(imjpgorig) - double(imfinal)).^2);
327  mseclnL2 = mean2((double(imjpgorig) - double(exp(lnyl2final)-1))
         .^2);
328
329  msecL2
330  msecW
331  mseclnL2
```

```
332
333 display('----------SSIM----------');
334 [mssimL2 ssim_mapL2] = ssim(uint8(imjpgorig), uint8(yl2final));
335 [mssimW ssim_mapWL2] = ssim(uint8(imjpgorig), uint8(imfinal)); %,
      K, window, L);
336 [lnmssimL2 lnssim_mapL2] = ssim(uint8(imjpgorig), uint8(exp(
      lnyl2final)-1));
337 mssimL2
338 mssimW
339 lnmssimL2
340
341 display('----------SSIM non-overlapping----------');
342
343 no_l2_ssim = mean(l2_ssim_block(:))
344 no_webl2_ssiml2 = mean(webl2_ssim_block(:))
345 no_l2ln_ssiml2 = mean(lnl2_ssim_block(:))
346
347 display('----------Web L^2 error----------');
348 msecL2 = mean2( ( (double(imjpgorig) - double(yl2final))./(
      imjpgorig+eps) ).^2 );
349 msecW = mean2(((double(imjpgorig) - double(imfinal))./(imjpgorig+
      eps) ).^2 );
350 nmsecL2 = mean2( ( (double(imjpgorig) - double(exp(lnyl2final)-1)
      )./(imjpgorig+eps) ).^2 );
351 msecL2
352 msecW
353 nmsecL2
354 display('----------');
355
356 resc_lnyl2final = exp(lnyl2final)-1;
357
358 figure(1);
359 sr = 2;
360 sc = 2;
361 subplot(sr,sc,1);imshow(imjpgorig,[]);title(sprintf('Original
      image'));
362 subplot(sr,sc,2);imshow(uint8(yl2final),[]); title(sprintf('DCT (
      block) approximation')); %L^2 appr
```

```
363 subplot(sr,sc,3);imshow(uint8(imfinal),[]); title(sprintf('
    Weberized (block) approximation')); %weberized appr
364 subplot(sr,sc,4);imshow(uint8(resc_lnyl2final),[]); title(sprintf
    ('Ln DCT (block) approximation')); %L^2 appr
365
366 figure(2);
367 sr = 2;
368 sc = 3;
369 subplot(sr,sc,1);imshow(uint8(yl2final),[]); title(sprintf('DCT (
    block) approximation')); %L^2 appr
370 subplot(sr,sc,2);imshow(uint8(imfinal),[]); title(sprintf('
    Weberized (block) approximation')); %weberized appr
371 subplot(sr,sc,3);imshow(uint8(resc_lnyl2final),[]); title(sprintf
    ('Ln DCT (block) approximation')); %L^2 appr
372 subplot(sr,sc,4);imshow(uint8(255 - abs(imjpgorig -yl2final)));
    title(sprintf('DCT (block) approximation')); %L^2 appr
373 subplot(sr,sc,5);imshow(uint8(255 - abs(imjpgorig - imfinal)));
    title(sprintf('Weberized (block) approximation')); %weberized
    appr
374 subplot(sr,sc,6);imshow(uint8(255 - abs(imjpgorig -
    resc_lnyl2final))); title(sprintf('Ln DCT (block)
    approximation')); %L^2 appr
```

Academic Contributions

Published Conference Proceedings

I. A. Kowalik-Urbaniak, Davide La Torre, E. R. Vrscay, Z. Wang. "Some 'Weberized' L2-based methods of signal/image approximation and processing", in Image Analyis and Recognition: Processing, Computer Vision, Pattern Recognition, and Graphics, Proceedings of ICIAR Vol. 8814, 2014.

I. A. Kowalik-Urbaniak, D. Brunet, J. Wang, E. Vrscay , Z.Wang, D. Koff, N. Smolarski-Koff, B. Wallace. "The quest for 'diagnostically lossless' medical image compression: a comparative study of objective quality metrics for compressed medical images", in Medical Imaging 2014: Image Perception, Observer Performance, and Technology Assessment, Proceedings of SPIE Vol. 9037, 903717, 2014.

I.A. Kowalik-Urbaniak, E. Vrscay, Z. Wang. C. Cavaro-Menard, D. Koff B. Wallace and B. Obara. "The impact of skull bone intensity on the quality of compressed CT neuro images", in Medical Imaging 2012: Advanced PACS-based Imaging Informatics and Therapeutic Applications, Proceedings of SPIE Vol. 8319, 83190L, 2012.

Conference Presentations

I. A. Kowalik-Urbaniak, D. Brunet, J. Wang, E. Vrscay , Z.Wang, D. Koff, N. Smolarski-Koff, B. Wallace. "The quest for 'diagnostically lossless' medical image compression: a comparative study of objective quality metrics for compressed medical images". Poster presented as part of the Medical Imaging Conference (SPIE), San Diego, February, 15-20, 2014 (Honourable Mention Poster Award).

A. Kowalik-Urbaniak, E. Vrscay, Z.Wang, D. Brunet, B. Wallace, D. Koff. "Structural Similarity Index for Quality Assessment of Medical Images". Talk as part of the Interdisciplinary Conference Series (AMMCS): Applied Mathematics, Modeling, and Computer Science (AMMCS-2013), Waterloo, Ontario, August 26-30, 2013.

I.A. Kowalik-Urbaniak, E. Vrscay, Z. Wang. C. Cavaro-Menard, D. Koff B. Wallace and B. Obara, "The impact of skull bone intensity on the quality of compressed CT neuro images". Talk as part of the Medical Imaging Conference (SPIE), San Diego, USA, February 14-19, 2012.

I.A. Kowalik-Urbaniak, E. Vrscay, Z. Wang. D. Koff and B. Wallace. "Quantitative Evaluation of Degradations Resulting from JPEG and JPEG2000 Compression of Neuro CT images". Poster presented as part of the 3rd European Workshop on Visual Information Processing - EUVIP, (IEEE), Paris, France, July 4-6, 2011. (Best Student Poster Award).

I.A. Kowalik-Urbaniak, E. Vrscay, Z. Wang and D. Koff. "Objective quality assessment of JPEG- and JPEG2000- compressed CT neuro images". Talk as part of the European Conference on Mathematical and Theoretical Biology - ECMTB (IEEE), Krakow, Poland, June 28 - July 02, 2011.

I.A. Kowalik-Urbaniak, "Objective Quality Assessment of Compressed Medical Images using MSE and SSIM". Talk as part of the 7th Annual Practical Course in Medical Imaging & Teleradiology - MIIT, Toronto, ON, May 13-14, 2011.

I.A. Kowalik-Urbaniak, "Diagnostically Lossless Compression of Medical Images", Poster presented as part of the MITACS - IBM 2nd Annual Research & Innovation Summit, Toronto, ON, September 30, 2010.

Invited Talks

I.A. Kowalik-Urbaniak, "The quest for 'diagnostically lossless' medical image compression: a comparative study of objective quality metrics for compressed medical images", AGFA Healthcare Inc. Convention, Waterloo, ON, Sep 2014. (Keynote Speaker).

References

[1] ITU-T page with recommendations on subjective video quality assessments. http://www.itu.int/rec/T-REC-P.910-200804-I.

[2] Methods for objective and subjective assessment of quality, telecommunication standardization sector of itu, series p: Telephone transmission quality, telephone installations, local line networks, November 2003.

[3] H. Al-Nuaim and N. Abukhodair. A user perceived quality assessment of lossy compressed images. *International Journal of Computer Graphics*, 2(2), November 2011.

[4] H. M. Al-Otum. Qualitative and quantitative image quality assessment of vector quantization, jpeg, and jpeg2000 compressed images. In *Journal of Electronic Imaging*, volume 12, pages 511–521, January 2003.

[5] The Royal Australian and New Zealand College of Radiologists. A guideline for the use of image compression in diagnostic imaging. http://www.ranzcr.edu.au/documents-download/doc_download/574-a-guideline-for-the-use-of-image-compression-in-diagnostic-imaging.

[6] H. Barrett and W. Swindell. *Radiological Imaging, The theory of image formation, detection and processing*, volume 2. Academic Press, 1981.

[7] T. M. Blodgett, C. C. Meltzer, and D. W. Townsend. Pet/ct: Form and function. *Radiology*, 242(2):360–385, February 2007.

[8] D. Brunet, E. R. Vrscay, and Z. Wang. Structural similarity-based approximation of signals and images using orthogonal bases. In *Image Analysis and Recognition*, volume 6111, pages 11–22. Springer Berlin Heidelberg, 2010.

[9] D. Brunet, E. R. Vrscay, and Z. Wang. A class of image metrics based on the structural similarity quality index. In *International Conference on Image Analysis and Recognition (ICIAR '11)*, Burnaby, BC, Canada, June 2011.

[10] Z. Budrikus. Visual fidelitv criteria and modeling. In *Proc. IEEE*, volume 60, pages 771–779, July 1972.

[11] Housser Bull and Tupper LLP. Lossy compression study: Legal assessment. Technical report, December 2004.

[12] P. C., Cosman, R. M. Gray, and R. A. Olshen. Evaluating quality of compressed medical images: Snr, subjective rating, and diagnostic accuracy. In 82, editor, *Proceedings of the IEEE*, volume 6, pages 919–932, June 1994.

[13] D. P. Chakraborty, H. L. Kundel, C. F. Nodine, T. K. Narayan, and V. Devaraju. Differential receiver operating characteristic (droc) method. In *Proc. SPIE 3338, Medical Imaging 1998: Image Processing*, volume 234, June 1998.

[14] D. P. Chakraborty and L. H. Winter. Free-response methodology: alternate analysis and a new observer-performance experiment. *Radiology*, 174(3):873–881, 1990. PMID: 2305073.

[15] D. M. Chandler and S. S. Hemami. Vsnr: A wavelet-based visual signal-to-noise ratio for natural images. *IEEE Transactions on Image Processing*, 16(9), September 2007.

[16] D. M. Chandler, K. H. Lim, and S. S. Hemami. Effects of spatial correlations and global precedence on the visual fidelity of distorted images. *Human Vision and Electronic Imaging XI*, 6057, February 2006.

[17] D. Cluni. Lossless compression of greyscale medical images - effectiveness of traditional and state of the art approaches. In *Proceedings of SPIE*, volume 3980, pages 74–84, 2006.

[18] G. W. Corder and D. I. Foreman. *Nonparametric Statistics: A Step-by-Step Approach*. John Wiley and Sons, Inc, May 2014.

[19] Oxford University Computing Laboratory D. Handscomb. Multiresolution and wavelets. http://eprints.maths.ox.ac.uk/1334/1/NA-95-23.pdf.

[20] I. Daubechies. *Ten lectures on wavelets*. SIAM, Philadelphia, 1992.

[21] C. Delgorge, C. Rosenberger, G. Poisson, and P. Vieyres. Towards a new tool for the evaluation of the quality of ultrasound compressed images. *in IEEE Transactions on Medical Imaging*, 2006.

[22] Max A. Viergever E. H. W. Meijering, Wiro J. Niessen. Quantitative evaluation of convolution-based methods for medical image interpolation. *Medical Image Analysis*, 5(2):111–126, 2001.

[23] N. J. Erickson. Irreversible compression of medical images. *Journal of Digital Imaging*, 15(1):5–14, March 2002.

[24] A. Fidler and B. Likar. What is wrong with compression ratio in lossy image compression? *Radiology*, 245(1):299, 2007.

[25] K. Finnis. A review and comparison of medical image compression. Technical report, Atamai Interactive Visualization Inc., 2004.

[26] J. R. Flynn, S. Ward, J. Abich IV, and D. Poole. Image quality assessment using the ssim and the just noticeable difference paradigm. *Engineering Psychology and Cognitive Ergonomics. Understanding Human Cognition Lecture Notes in Computer Science*, 8019:23–30, 2013.

[27] B. Forte and E.R. Vrscay. Solving the inverse problem for funciton and image approximation using iterated function systems, dynamics of continuous. *Discrete and Impulsive Systems 1*, (1):177–231, 1995.

[28] K. Bindu A. Ganpati and A. K. Sharm. A comparative study of image compression algorithms. *International Journal of Research in Computer Science*, 2(5):37–42, 2012.

[29] A. George and S. J. Livingston. A survey on full reference image quality assessment algorithms. *IJRET: International Journal of Research in Engineering and Technology*, 2(12), December 2013.

[30] S. E. Ghrare, M. A. M. Ali, and M. Ismail. Diagnostic quality of compressed medical images: Objective and subjective evaluation. *European Journal of Scientific Research*, 23(1):6–12, 2008.

[31] R.C. Gonzalez and R.Woods. *Digital image processing*. Prentice Hall, 2008.

[32] Y. Han, Y. Cai, Y. Cao, and X. Xu. A new image fusion performance metric based on visual information fidelity. *Information Fusion*, 14(2):127–135, s 2013.

[33] T. Hastie, R. Tibshirani, and J. Friedman. *Elements of Statistical Learning, Data Mining, Inference, and Prediction*. Springer-Verlag, 2009.

[34] G. T. Herman. *Fundamentals of computerized tomography: Image reconstruciton form projection*. Springer, 2 edition, 2009.

[35] W. Huda. *Review of radiologic physics*. Wolters Kluwer, Lippincott Williams & Wilkins, 2010.

[36] INSITE Consultancy Inc. Irreversible compression in digital radiology – a literature review. Technical report, August 2004.

[37] I. K. Indrajit and B. S. Verma. Monitor displays in radiology: Part 2. *Indian J Radiol Imaging*, 19(2):94–98, May 2009.

[38] D. Koff, P. Bak, P. Brownrigg, D. Hosseinzadeh, A. Khademi, A. Kiss, L. Lepanto, T. Michalak, H. Shulman, and A. Volkening. Pan-canadian evaluation of irreversible compression ratios (lossy compression) for development of national guidelines. *J Digit Imaging*, 22(6):569–578, December 2009.

[39] D. Koff and H. Shulman. An overview of digital compression of medical images: Can we use lossy image compression in radiology?, an overview of digital compression of medical images: Can we use lossy image compression in radiology?, an overview of digital compression of medical images: Can we use lossy image compression in radiology. *CARJ*, 57(4), October 2006.

[40] E. Kofidis, N. Kolokotronis, A. Vassilarakou, S Theodoridis, and D. Cavouras. Wavelet-based medical image compression. *Future Generation Computer Systems*, March 1999.

[41] I. A. Kowalik-Urbaniak, D. Brunet, J. Wang, E. Vrscay, Z.Wang, D. Koff, N. Koff, and B. Wallace. The quest for 'diagnostically lossless' medical image compression: a comparative study of objective quality metrics for compressed medical images. In *Medical Imaging 2014: Image Perception, Observer Performance, and Technology Assessment*, volume 9037, 2014.

[42] I. A. Kowalik-Urbaniak, D. La Torre, E. R. Vrscay, and Z. Wang. Some weberized l2-based methods of signal/image approximation and processing. In *Image Analyis and Recognition: Processing, Computer Vision, Pattern Recognition, and Graphics, Proceedings of ICIAR*, volume 8814, pages 20–29, 2014.

[43] I.A. Kowalik-Urbaniak, E. Vrscay, Z. Wang, C. Cavaro-Menard, D. Koff, B. Wallace, and B. Obara. The impact of skull bone intensity on the quality of compressed ct neuro images. In *Medical Imaging 2012: Advanced PACS-based Imaging Informatics and Therapeutic Applications, Proceedings of SPIE*, volume 8319, pages 83190L–1:7, 2012.

[44] E. Larson and D. M. Chandler. Most apparent distortion: A dual strategy for full-reference image quality assessment. In *Image Quality and System Performance VI*, volume 7242. SPIE IS&T, SPIE Proceedings, 2009.

[45] E. C. Larson and D. M. Chandler. Most apparent distortion: full reference image quality assessment and the role of strategy. *J. Electron. Imaging*, 19(011006):1–21, 2010.

[46] H. Lee, D. Haynor, and Y. Kim. Subjective evaluation of compressed image quality. In *Proceedings of SPIE, Image Capture, Formatting and Display*, volume 1653, pages 241–245, 1992.

[47] R. Loose, R. Braunschweig, E. Kotter, P. Mildenberger, R. Simmler, and M. Wucherer. Compression of digital images in radiology—results of a consensus conference. *Rofo*, 181:32–37, 2009.

[48] G. Campbell M. H. Zweig. Receiver-operating characteristic (roc) plots: a fundamental evaluation tool in clinical medicine. *Clin Chem*, 39:561–577, 1993.

[49] S. Mallat. *A wavelet tour of signal processing*. Academic Press, 1998.

[50] J. L. Mannos and D. J. Sakrison. The effects of a visual fidelity criterion of the encoding of images. In *IEEE Trans. Informat. Theory*, volume IT-20, pages 525–536, July 1974.

[51] H. Marmolin. Subjective mse measures. *IEEE Transactions on Systems Man and Cybernetics*, SMC-16(3):486–489, 1986.

[52] M. McInnes. Digital compression of medical images: an assessment of legal risk. Technical report, Faculty of Law, University of Western Ontario.

[53] E. H. W. Meijering. Spline interpolation in medical imaging: Comparison with other convolution-based approaches. In *European Signal Processing Conference (EUSIPCO)*, volume 4, pages 1989–1996, September 2000.

[54] C. Metz. Receiver operating characteristic analysis: A tool for the quantitative evaluation of observer performance and imaging systems. *Journal of the American College of Radiology*, 3(6):413–422, 2006.

[55] C. E. Metz. Basic principles of ROC analysis. *Seminars in Nuclear Medicine*, 8:282–298, October 1978.

[56] A. Mittal, R. Soundararajan, and A. C. Bovik. Making a completely blind image quality analyzer. *IEEE Signal processing Letters*, 22(3):209–212, March 2013.

[57] A. Nait-Ali and C. Cavaro-Menard. *Compression of Biomedical Images and Signals*. John Wiley and Sons, 2008.

[58] A. Franchetto Neuroradiologist. Discussion session. Hamilton General Hospital, Hamilton Ontario, 2011.

[59] N. Nikvand and Z. Wang. Image distortion analysis based on normalized perceptual information distance. *SIViP*, 7:403–410, 2013.

[60] N. A. Obuchowski. Receiver operating characteristic curves and their use in radiology. In *Radiology*, volume 299, pages 3–8, 2003.

[61] Canadian Association of Radiologists. Car standards for irreversible compression in digital diagnostic imaging within radiology. www.car.ca.

[62] Royal College of Radiologists (RCR). The adoption of lossy data compression for the purpose of clinical interpretation. https://www.rcr.ac.uk/docs/radiology/pdf/IT_guidance_LossyApr08.pdf, 2008.

[63] European Society of Radiology (ESR). Usability of irreversible image compression in radiological imaging. *Insights into Imaging*, 2(2):103–115, 2011.

[64] J. Oh, S. I. Woolley, T. N. Arvanitis, and J. N. Townend. A multistage perceptual quality assessment for compressed digital angiogram images. *IEEE Trans Med Imaging*, 20(12):1352–1561, December 2001.

[65] T. H. Oh, H. S. Lim, and S. Y. Pang. Medical image processing: From lossless to lossy compression medical image processing: From lossless to lossy compression. In *TENCON IEEE Region 10*, 2005.

[66] A.V. Oppenheim, R.W. Schafer, and T.G. Stockham Jr. Nonlinear filtering of multiplied and convolved signals,. In *Proc.*, volume 56 (8), pages 1264–1291. IEEE, 1968.

[67] W. B. Pennebaker and J. L. Mitchell. *JPEG, Still image data compression standard*. Kluwer Academic Publishers, 2004.

[68] S. M. Perlmutter, C. Tseng, P. C. Cosman, K. C. P. Li, R. A. Olshen, and R. M. Gray. Measurement accuracy as a measure of image quality in compressed MR chest scans. In IEEE Computer Society Press, editor, *ICIP*, volume 1, pages 861–865, November 1994.

[69] K. Persons, P. Palisson, A. Manduca, B. J. Erickson, and V. Savcenko. An analytical look at the effects of compression on medical images. *Journal of Digital Imaging*, 10:60–66, 1997.

[70] M. Pinson and S. Wolf. Comparing subjective video quality testing methodologies. In *SPIE Video Communications and Image Processing Conference*. Lugano, Switzerland, July 2003.

[71] A. Przelaskowski, M. Kazubek, and T. Jamrógiewicz. Methods for determining diagnostic accuracy of lossy compressed medical images. In *Proceedings of the 13-th Biennial International Conference BIOSIGNAL*, 1996.

[72] M. Razaak and M. G. Martini. Medical image and video quality assessment in e-health applications and services. In *e-Health Networking, Applications Services (Healthcom), 2013 IEEE 15th International Conference o*, pages 6–10, October 2013.

[73] S. Rout. Orthogonal vs. biorthogonal wavelets for image compression, m.sc. thesis. Technical report, Virginia Polytechnic Institute and State University, Blacksburg, Virginia, 2003.

[74] A. Saffor, A. R. Ramli, and Kwan-Hoong. A comparative study of image compression between jpeg and wavelet. *Malaysian Journal of Computer Science*, 14(1):39–45, June 2001.

[75] K. Sayood. *Introduction to data compression*. Morgan Kaufmann Publishers, 1996.

[76] M. Shahid, A. Rossholm, B. Lövström, and H. Jürgen Zepernick. No-reference image and video quality assessment: a classification and review of recent approaches. *EURASIP Journal on Image and Video Processing*, 40, 2014.

[77] C.E. Shannon. *A mathematical theory of communication*. University of Illinois Press, 1949.

[78] H. R. Sheikh and A. C. Bovik. Image information and visual quality. *IEEE Transactions on Image Processing*, 15(2):pp. 430–444, February 2006.

[79] H. R. Sheikh, A. C. Bovik, and G. de Veciana. An information fidelity criterion for image quality assessment using natural scene statistics. *IEEE Transactions on Image Processing*, 14(12):2117–2128, December 2005.

[80] R. Sivakumar. Denoising of computer tomography images using curvlet transform. *ARPN Journal of Engineering and Applied Sciences*, 2(1), 2007.

[81] M. Smith and S. L. Eddins. Analysis/synthesis techniques for subband image coding. *IEEE Trans. on Acoustics, Speech, and Signal Proc*, 38(8):1446–1456, August 1990.

[82] M. G. Strintzis. A review of compression methods for medical images in pacs. *International Journal of Medical Informatics*, 52:159–162, 1998.

[83] J. A. Swets. ROC analysis applied to the evaluation of medical imaging technique. *Investigative Radiology*, 14:109–121, March-April 1979.

[84] D. S. Taubman and M. W. Marcellin. *JPEG2000, Image compression fundamentals, standards and practice*. Kluwer Academic Publishers, 2002.

[85] G. V. Tcheslavski. Wavelets fundamentals, lamar university. http://ee.lamar.edu/gleb/dip/index.htm.

[86] DCTD/NCI/NI The Cancer Imaging Archive Sponsored by the Cancer Imaging Program. http://cancerimagingarchive.net.

[87] B. A. Wandell. *Foundations of Vision*. Sinauer Associates Inc, 1995.

[88] Z. Wang and A. C. Bovik. Mean squared error: love it or leave it? - a new look at signal fidelity measures. *IEEE Signal Proc. Mag.*, 26(1):98–117, January 2009.

[89] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, April 2004.

[90] Z. Wang and Q. Li. Information content weighting for perceptual image quality assessment. *IEEE Transactions on Image Processing*, 20(5):1185–1198, May 2011.

[91] Z. Wang, E. P. Simoncelli, and A. C. Bovik. Multi scale structural similarity for image quality assessment. In *Proceedings of IEEE Asilomar Conference on Signals, Systems, and Computers*, pages 1398–1402, Pacific Grove, CA, s 2003.

[92] J. Yang, J. Wright, T. Huang, and Y. Ma. Image super-resolution as sparse representation of raw image patches. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.

[93] J. Youden. Index for rating diagnostic tests. *Cancer*, 1(5):32–35, 1950.

[94] F. Zarb, L. Rainford, and M. F. McEntee. Image quality assessment tools for optimization of ct images. *Radiography*, 16(2):147–153, 2010.

[95] F. Zhang, L. Ma, S. Li, and K. N. Ngan. Practical image quality metric applied to image coding. *IEEE Trans. On Multimedia*, 13(4):615–624, 2011.