THE UNIVERSITY OF HULL

Optimisation of Computed Radiography Chest Imaging Utilising a Novel Simulation Technique Derived from Real Patient Computed Tomography Data Sets

being a Thesis submitted for the Degree of

PhD

in the University of Hull

by

Craig Steven Moore, BSc (Hons), MSc

September 2011

Acknowledgements

Firstly, my thanks go to my supervisors Prof Andy Beavis and Mr John Saunderson, and my previous manager Mr Viv Whitton for letting me do this research in-service; it would not have been possible otherwise. I would also like to say thanks for their expertise, advice, encouragement and direction.

Many thanks go to the Humberside Radiology Trust Fund, who provided funding for the work.

I would also like to thank Dr Ged Avery, Dr Andy Swift, Mr Steve Balcam and Mr Liam Needler for taking time out from their busy schedules and evaluating simulated chest images.

Most loving thanks go to my wife Shannon for always being there, putting up with me and giving me the inspiration and motivation to finish this thesis.

Publications

The following publications and presentations have arisen from this, and associated work:

- C.S. Moore, G.P. Liney, A.W. Beavis and J.R. Saunderson, “A method to optimise the processing algorithm of a Computed Radiography system for chest radiography.” Br J Radiol. 80:724-730 (2007).
- C.S. Moore, A.W. Beavis and J.R. Saunderson, “Investigation of optimum X ray beam tube voltage and filtration for chest radiography with a Computed Radiography system,” Br J Radiol. 81:771-777 (2008).
- C.S. Moore, A.W. Beavis and J.R. Saunderson, “Investigating the exposure (speed) class of a Computed Radiography system as a means to optimising physical image quality for chest radiography.” Br J Radiol. 82:705-710 (2009).
- C.S. Moore, G.P. Liney, A.W. Beavis and J.R. Saunderson, “A method to produce and validate a digitally reconstructed radiograph-based computer simulation for optimisation of chest radiographs acquired with a computed radiography imaging system.” Br J Radiol. 84:890-902 (2011).
- C.S. Moore, G Avery, A Swift, L Needler, S Balcam, A.W. Beavis and J.R. Saunderson, “Use of a digitally reconstructed radiograph (DRR) based computer simulation for the optimisation of chest radiographic techniques for Computed Radiography (CR) imaging systems.” Accepted for publication in Br J Radiol. May 2011.
- C.S. Moore. "A method to produce and validate a digitally reconstructed radiograph (DRR) based computer simulation for optimization of chest radiographs acquired with a computed radiography (CR) imaging system". 5th Meeting of the IPEM’s Digital Radiology User Group, Glasgow, April 2010.
- C.S. Moore. "Use of a digitally reconstructed radiograph (DRR) based computer simulation for the optimisation of chest radiographic techniques for Computed Radiography (CR) imaging systems ". IPEM RPA Update, Manchester, June 2011.

Abstract

To optimise any medical digital imaging system for chest radiography, it is vital that the images used for optimisation contain projected anatomy, or in other words, anatomical noise. In this thesis, a method to produce and validate a digitally reconstructed radiograph (DRR) computer algorithm that utilises real patient computed tomography (CT) data sets is presented. The algorithm uses a ray casting DRR calculation method to project X-ray pencil beams through CT data and derive the photon energy absorbed in a virtual computed radiography (CR) phosphor. Radiation scatter and CR system noise are added post DRR calculation.

Quantitative and qualitative validation has shown the algorithm simulates chest CR images of average and obese patients with realistic anatomical and system noise. This has allowed images to be generated using various X-ray exposure parameters, i.e. tube potential, scatter rejection and receptor dose, which can then be used in the optimisation exercise. However, the algorithm is not without limitations; the impact of these on the resulting images is discussed.

Simulated images reconstructed at the various X-ray exposure parameters and techniques were scored by experienced image evaluators; optimum tube potential, scatter rejection technique and receptor doses for clinical CR chest radiography have been derived. At the outset of this work, CR chest exposure factors across the Hull & East Yorkshire Hospitals NHS Trust (HEY) were not standardised, and therefore not optimised; this thesis concludes with recommendations to the HEY Radiology Department for optimum exposure factors and technique for chest radiography. These were implemented across the Trust as a result of this work.

In summary, a DRR computer algorithm has been produced (and validated) that adequately simulates anatomical and system noise; image evaluators are able to grade simulated chest images presented at different X-ray exposure parameters in order to optimise radiographic technique for clinical CR chest radiography, without the need for repeat patient exposures.

Table of Contents

List of Figures

List of Tables

Chapter 1: Introduction

This chapter provides an introduction to the work that follows. It comprises of a discussion of the relevant digital imaging modalities used in medicine, and a description of the physics underpinning each modality. Optimisation in medical X-ray imaging and its motivation is summarised. Various optimisation techniques found in the literature are discussed as well as the relevance of the work presented in this thesis.

1.1 Digital imaging in medicine

Digital imaging was introduced into healthcare with the development of nuclear medicine imaging in the 1950's, and later expanded with computerised tomography (CT) in the 1970's and magnetic resonance imaging (MRI) in the 1990's. However, until the last 10 to 15 years, film has remained the prime imaging device for 95% of radiological procedures in the UK.

With the development of ever larger digital storage facilities and the availability of commercial picture archiving and communication systems (PACS), a move away from hard copy film towards digitally stored and viewed images has become a practical proposition. In the last decade in particular, driven by the UK Department of Health's National Programme for Information Technology (NPfIT), there has been a rapid move in UK healthcare away from film to digital image acquisition systems such as Computed Radiography (CR). Almost overnight the accumulated experience of over a century of work with radiographic film has become obsolete, and radiology departments are faced

with the need to adapt to cope with the challenges of, and to utilise the opportunities presented by, the new digital technologies.

1.2 Computed Radiography

1.2.1 Basic physics of Computed Radiography

Computed radiography (CR), scientifically known as photo-stimulable phosphor radiography, is a digital technology for the acquisition of radiographic images [1, 2]. CR is the most common digital radiography modality in the radiology department today, with an estimated 7000 systems in use worldwide in 2001 [3]. At the time of writing, CR is the only method of image generation for general radiography in the Hull and East Yorkshire Hospitals NHS Trust. The technology uses conventional radiographic acquisition geometries to deposit X-ray energy in a photo-stimulable phosphor screen with delayed luminescence properties (phosphorescence). The phosphors most frequently used are those of the barium fluorohalide family [4], in powder form deposited onto a substrate to form the imaging plate. The elemental composition of powder phosphor plates used in modern CR systems is of the form BaSrFBrI:Eu. Barium (Ba) and Fluorine (F) are the principal absorbing elements, and Strontium (Sr), Bromine (Br) and Iodine (I) are halogens used in such a combination that the useful signal is not derived from light emitted immediately after the absorption of X-rays, but rather from the delayed emission of energy from trapped electrons that are optically stimulated. The spectrum of useful light emitted by a powder phosphor is controlled by the ‘activator’, which is an impurity added to the base matrix; in this case Europium (Eu).

A key concept of phosphors is the exciton, which is essentially a bound electron-hole pair (a hole is a mathematical description of the absence of an electron in an electron shell). The exciton is free to move within the crystal lattice and is usually neutral, with the electron in the conduction band and the hole in the valence band (as shown in Figure 1.1).

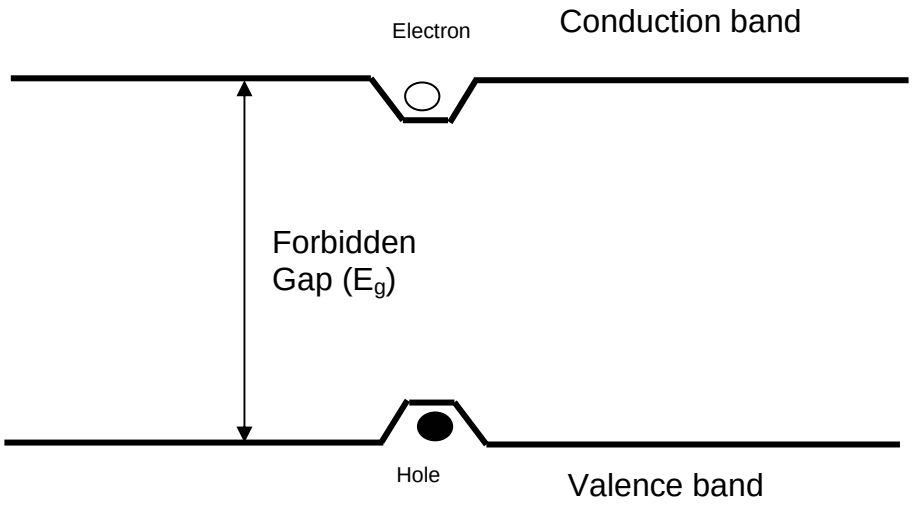

Figure 1.1: Bound electron-hole pair (exciton) in a band structure representation. The electron and hole is not allowed to exist in the forbidden energy gap (Eg) that separates the conduction and valence bands, unless they become 'trapped'.

Trapping of an exciton must occur to produce a latent image in the phosphor, and for the subsequent release of stored energy. The trap comes in the form of the activator (Eu), as demonstrated in Figure 1.2.

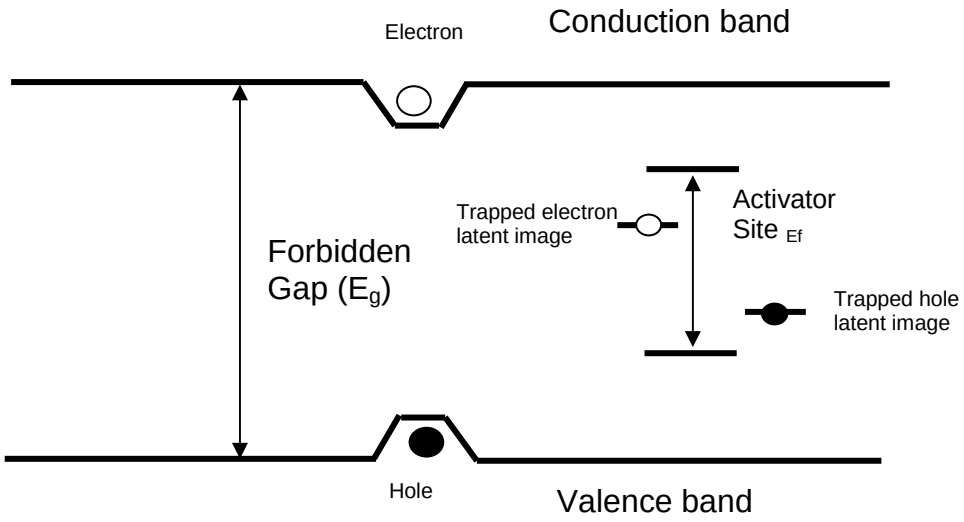

Figure 1.2: The energy levels at which an electron-hole pair can be trapped due to the presence of an activator (Eu) in the powder phosphor.

The first step in the creation of electron-hole pairs, in which an electron is excited from the valence band to the conduction band leaving behind a hole, is through absorption of X-ray energy by the photoelectric effect. The energy of the incoming X-ray photons must be at least that of the forbidden gap energy (E_g) for this process to occur. Some of the excited electrons become trapped at the activator sites (Figure 1.2) in a spatial distribution similar to that of the incoming X-ray photons, producing the latent image. An activated photostimulated luminescent site is thought to be an arrangement of three correlated components: an electron trap, a hole trap, and the activation centre.

Photo-stimulable phosphors have different absorption (K) edge characteristics to the rare-earth screens used in traditional film radiography, owing to their different elemental compositions. Whenever the incoming photon energy is slightly higher than the energy required to remove an electron from a shell in an atom, there is an appreciable increase in photoelectric absorption; this is known as an absorption edge, and in the diagnostic energy range this is dominated by

K shell absorption. The effect this has on the attenuation properties of a typical CR phosphor is shown in Figure 1.3.

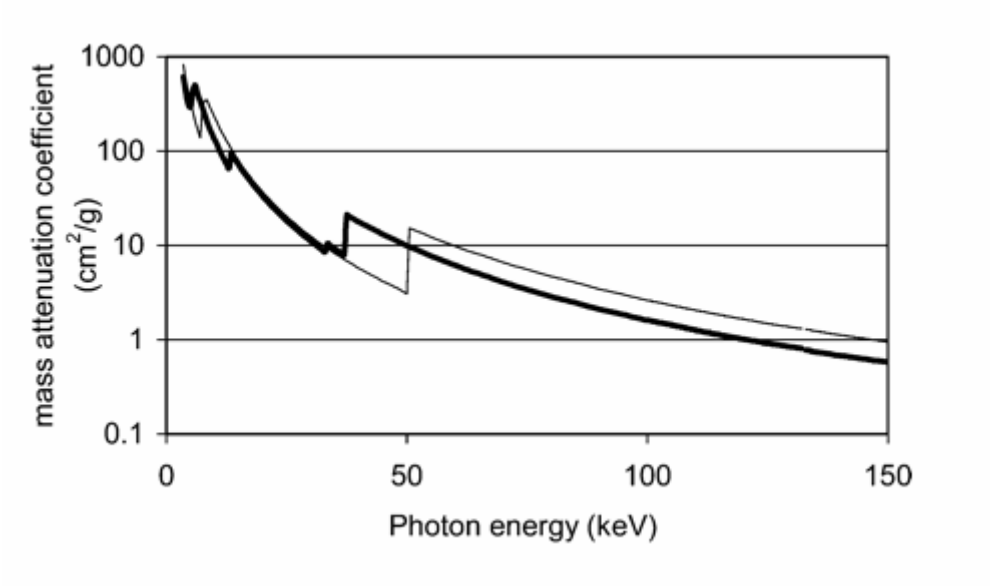

Figure 1.3: Attenuation characteristics of a CR phosphor (thick line). There is a sharp increase in attenuation at the K edge of Barium (37 keV). For illustration, the attenuation characteristics of a rare earth phosphor used in traditional film-screen imaging (thin line) is also shown.

The CR phosphor elememtal composition, BaSrFBrI:Eu, has K-edges at approximately 33 keV for Iodine and 37 keV for Barium, whilst Gd_2O_2S (used for rare-earth screens) has a K-edge at 50 keV, all of which are in the diagnostic energy range. Between energies 37 to 50 keV, the sensitivity of the CR phosphor is increased relative to the rare-earth screen, suggesting X-ray tube potentials lower than that recommended for film-screen would provide improved X-ray absorption, as there would be a greater proportion of photons in the 37 to 50 keV energy range. The properties of photo-stimulable phosphors are therefore inherently different to that of film-screen systems, meaning optimised radiographic technique for film-screen may not be suitable for CR. There is guidance on optimum X-ray radiographic technique with film-screen systems for chest radiography [5] (namely a tube potential of 125 kVp using an anti-scatter grid), but none for CR. Therefore, research must be carried out to determine

optimum radiographic techniques (e.g. tube voltage (kVp), tube current-time product (mAs), filtration, scatter rejection methods) for CR imaging modalities.

1.2.2 Phosphor plate read-out and image production

After X-ray irradiation, the phosphor is stimulated by a scanning laser beam that emits light in the red end of the visible spectrum (typically a solid state diode laser controlled electrically with a wavelength of 680 nm) to release the deposited energy in the form of visible blue light (i.e. the stimulation spectrum). The laser beam is scanned across the CR phosphor plate (denoted the scan direction), while the phosphor plate is continuously translated in a direction perpendicular to the motion of the laser beam (the sub-scan direction) resulting in the phosphor plate being scanned in a raster fashion. The stimulation spectrum results from the process of excitation of the trapped electron; the energy difference between trapped electrons and the conduction band is approximately equal to the energy of the red laser photons contained in the scanning laser beam. This process is called photostimulated luminescence and results in the release of blue light photons in an amount proportional to the intensity of the incident X-ray irradiation. The released photo-stimulated light is captured by a light guide (blue light is optically separated from red light by a filter), which is located as close as possible to the phosphor plate in order to increase its light collection efficiency. Light is channelled down the light guide by total internal reflection to a photomultiplier tube (PMT), and then converted to digital signals in a series of processing steps, the first being that of logarithmic amplification. This reduces the dynamic range of the analogue signal prior to digitization to approximately match the dynamic range of the acquisition monitor. The second step is to filter the signal to prevent aliasing of noise and

eliminate fixed pattern noise from such things as X-ray anti-scatter grids. The third step is digitization of the analogue signal by the analogue to digital converter. Finally, a shading correction is applied due to the varying light collection efficiency of the light guide along its width. This process is shown in Figure 1.4.

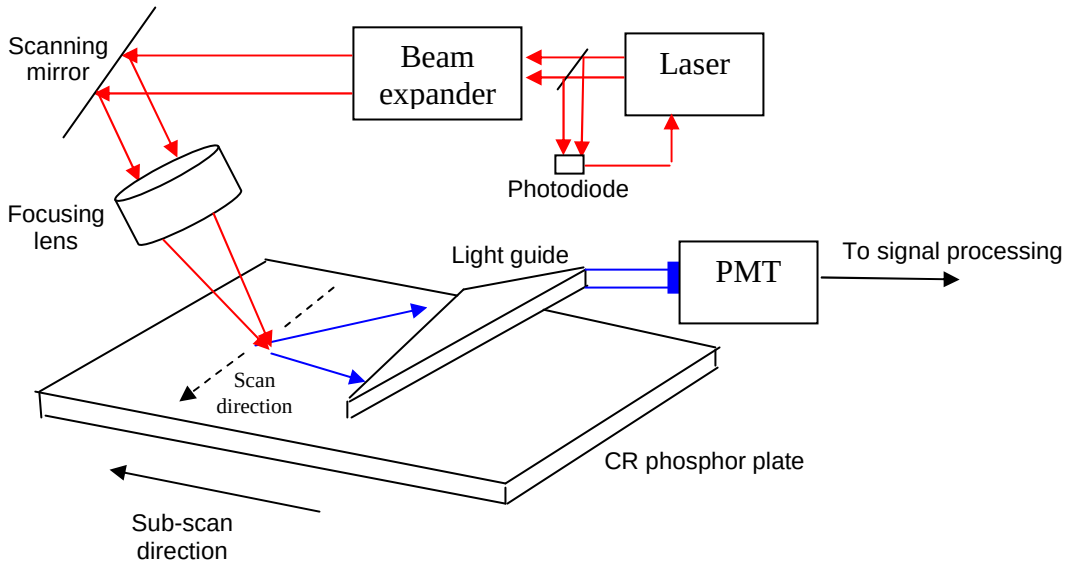

Figure 1.4: Flying spot CR readout scanner. Individual scanner components are labelled. The CR phosphor plate travels in the sub-scan direction, while the focused laser spot travels in the scan direction.

The resultant digital pixel value is calculated and registered with the location on the phosphor from which it has been released. The digital data may then be post-processed for appropriate presentation and sent to a hard copy printer or soft-copy display monitor for medical evaluation.

1.3 Computed Tomography

Computed Tomography (CT) became popular for diagnostic purposes in the UK in the early 1980s, and over the last twenty years its use has become widespread. The objective of CT scanning is to take many one-dimensional sample projections of two-dimensional transverse axial slices (usually within a patient) at many different angles. The data is reconstructed to give detail within the slice. At each angle (or projection), X-rays emanating from the X-ray tube are attenuated in the axial slice and the amount of attenuation is measured by a bank of detectors (that is, the total linear attenuation coefficient in the axial slice between the X-ray tube and detector is measured by each detector). A basic example of one projection is shown in Figure 1.5. The slice through the patient has been ‘voxelated’ to show the corresponding pixel positions in the resulting image.

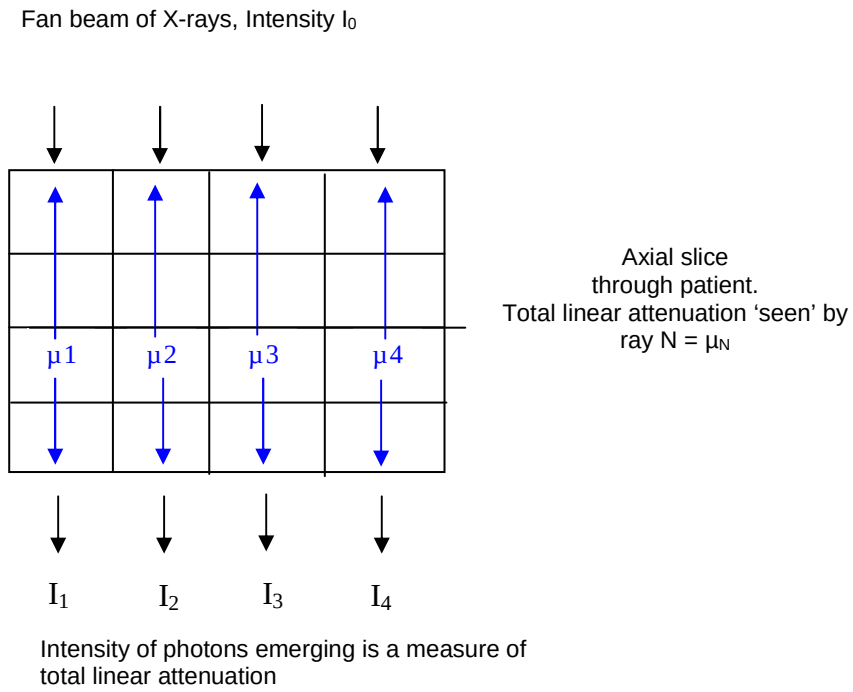

Figure 1.5: Basic depiction of X-rays traversing through a slice in the patient. Total linear attenuation 'encountered' by ray N is given by μN.

Most modern CT scanners use the third generation design, in which the X-ray tube and detector bank rotate continuously whilst maintaining a fixed geometrical relationship, as shown in Figure 1.6.

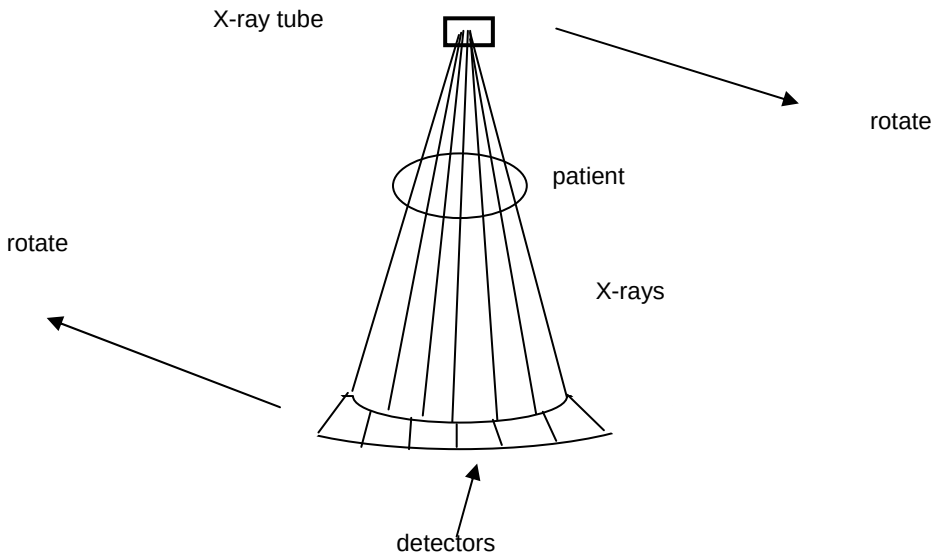

Figure 1.6: Third generation CT scanner design showing positions of X-ray tube, patient and detectors.

Data reconstruction is usually done by a method called filtered back-projection. Back-projection is a process by which total attenuation data measured by each individual detector is ‘smeared’ back across the image field of view (FOV). A highly simplified version of this is demonstrated in Figure 1.7.

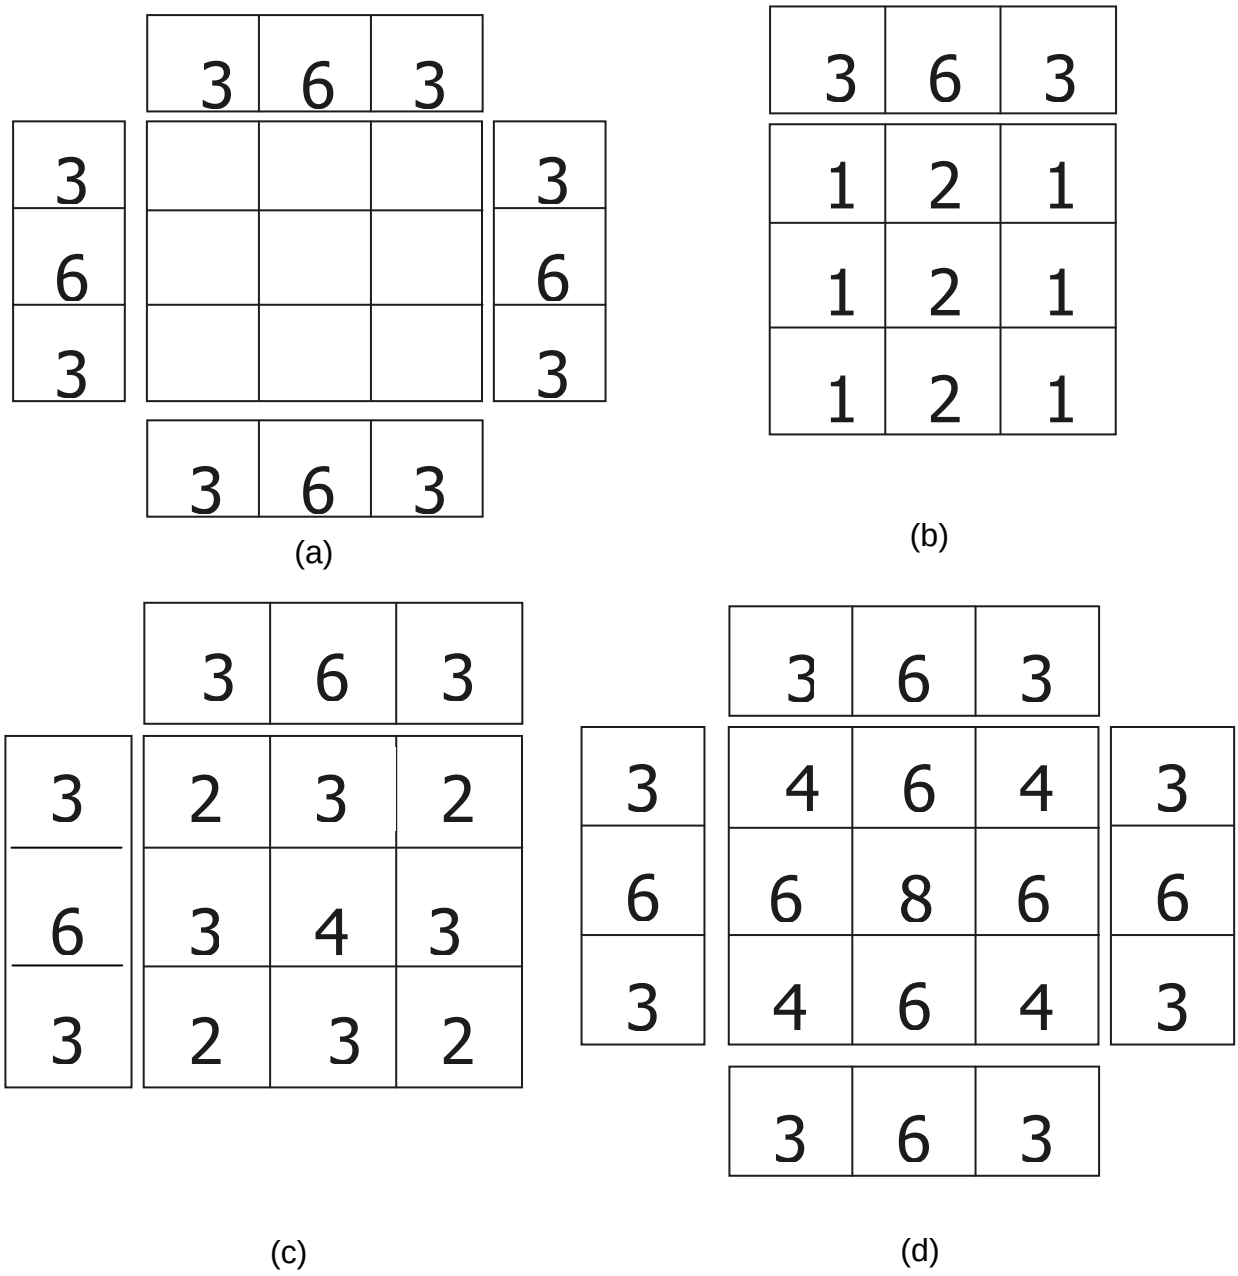

Figure 1.7: (a) Total attenuation measured along four projections, (b) resulting back-projected data from one projection, (c) back-projected data from a second projection summed with the first, and (d), summed back-projected data from all four projections.

Figure 1.7 is a highly simplified demonstration of back-projection. An array of nine pixels is scanned from four successive directions (Figure 1.7(a)). In each case the total attenuation data is recorded by the detectors. Back-projected data from one direction (Figure 1.7(b)) is then summed with another (Figure 1.7(c)) and this is done in turn until all projections are summed together (Figure 1.7(d)) to provide the final image pixel values. The total attenuation data is then stored in the memory locations (i.e. pixels in the resulting image) according to each voxel encountered by each ray. In reality there are many hundreds of projection angles that contribute to the final image, and CT manufacturer specific algorithms modify the pixel values even further.

This method of image reconstruction produces images that are inherently blurred. To overcome this, back-projected data is software filtered, i.e. the data is 2D Fourier transformed into frequency space which is then multiplied by a 2D ramp. The inverse 2D Fourier transform is then performed to give a non-blurred image in real space.

The final image is displayed as an array of pixels called CT numbers, and each CT number is represented by a certain shade of grey. The CT number is an index that compares the linear attenuation of the displayed tissue with that of water, as shown in the following equation;

$$\mathrm{CTnumber} = 1000 \times \left(\frac{\mu_T - \mu_W}{\mu_W} \right) \tag{1.1}$$

where μ_T and μ_W are the linear attenuation coefficients of tissue and water respectively. Water is used as a reference material because its attenuation coefficient is close to that of soft tissue and it is reproducible for scanner calibrations. The multiplier 1000 is used to generate meaningful whole numbers.

1.4 Optimisation in Medical Imaging

The three important principles which are the foundation of radiation protection are justification, optimisation and limitation. The principles of justification and optimisation apply to all exposed individuals, while dose limitation is applicable to radiation workers and the public, but not to patients. The concepts were first proposed by the International Commission on Radiological Protection (ICRP) in 1977 and developed in later recommendations [6]. It is also a requirement in UK law under the Ionising Radiation (Medical Exposure) Regulations 2000 (IR(ME)R2000) to justify and optimise all human medical exposures to ionising radiation [7].

A brief synopsis of each principle is outlined below:

Justification of a practice: No practice involving exposures to radiation should be adopted unless it produces sufficient benefit to the exposed individuals or society to offset the radiation detriment it causes.

In diagnostic radiology, the individual patient's X-ray has to be justified. If the result of an X-ray procedure will not alter the management of the patient then that X-ray is not justified. Justification depends on a number of factors; examples include the patient's age, whether or not they are pregnant, and whether other techniques involving non-ionising radiation can be used.

Optimisation: In relation to any particular source within a practice, the magnitude of individual doses, the number of people exposed, and the likelihood of incurring exposures should all be kept as low as reasonably achievable, economic and social factors being taken into account. This

procedure should be constrained by restrictions on the doses to individuals (dose constraints).

Once a practice has been justified efforts must be made to reduce the radiation risks to the individual. This means reducing doses and minimizing the possibility of accidents and incidents where doses might be raised. This dose reduction should be to a level which is ‘as low as reasonably practicable’ (ALARP). ALARP is central to the concept of optimisation for both staff and patients. For staff, ALARP is usually achieved with the following hierarchy of control measures:

1. Engineering controls (e.g. X-ray beam collimation and local shielding to reduce radiation emitted from a source)

2. Systems of work (e.g. local rules and contingency arrangements)

3. Personal protective equipment (e.g. lead rubber aprons and gloves)

For patients, minimising patient radiation dose reduces the chances of the patient developing stochastic effects such as cancer. ALARP is achieved by optimising the radiographic technique, which usually requires minimising radiation dose whilst maintaining adequate image quality. Currently, there is much effort in the medical field concentrating on reducing X-ray doses to patients while not compromising digital image quality; but ultimately the two are intimately linked, as a reduction in dose usually results in noisier (but possibly adequate) images.

1.4.1 Noise sources with CR imaging systems

Noise in a digital image is defined as fluctuations in the signal level (pixel values) and affects the visibility of low contrast objects. Noise present in clinical CR images can be broken down into two parts: 1) system noise and 2) anatomical noise. System noise is a combination of [8];

- primary quantum noise,
- secondary quantum noise,
- Poisson excess noise,
- structure noise,
- additive electronic noise.

Primary quantum noise is related to the random nature of X-ray photon absorption in the phosphor and is Poisson distributed, i.e. noise is proportional to the square root of the number of photons detected. Secondary quantum noise arises at each stage of the imaging chain where quanta are converted from one form to another. Poisson excess noise occurs when a given amount of primary quanta are absorbed, but differences in the amount of secondary quanta produced is observed. Structure noise occurs due to differences in sensitivity across the CR phosphor (resulting in non-uniformities in the image), and electronic noise is due to the operation of electronic components in the CR system (e.g. dark currents). The noise power spectra (NPS) is a measure of the total noise (i.e. a combined measure of the sources of noise discussed above) in an image at each frequency (of noise), and is the most common metric for describing noise in a digital image [9]. For example, the anode heel effect

manifests itself in a digital image as a slowly changing gross signal, and therefore low frequency noise. Conversely, quantum noise manifests itself as random values between adjacent pixels, and therefore high frequency noise.

Anatomical noise is essentially the influence of the patient's projected anatomy on the detectability of normal chest structures and abnormalities. Many publications [10-17] have shown this projected anatomy is the limiting factor in diagnosis and detection of lesions in the chest, so the term anatomical noise was derived from this collective work. Recently, a European wide study (the RADIUS chest trial) examined various aspects of nodule detection in digital chest radiology, such as the effects of nodule location, system noise, anatomical noise and part of the image background acting as pure noise [18 - 21]. In the summary paper Hakansson *et al.* [22] describe how projected anatomy in chest radiographs affects pathological detection to a much larger extent than system noise, which reinforces the conclusions of earlier published work. Therefore, to optimise digital X-ray systems for chest radiography, it is vital all images used and evaluated by observers contain the essential projected anatomy, otherwise the results obtained would not have been derived from images 'cluttered' with anatomical noise and therefore be potentially misleading.

1.4.2 Optimisation of Chest CR Imaging

There are essentially two types of optimisation techniques used in the literature: 1) methods that utilise physical phantoms containing little clinical detail (anatomical noise) from which optimum X-ray exposure parameters have been assessed by measuring indices such as signal-to-noise (SNR) and contrast-to-noise (CNR) ratios, and 2) simulation of the digital X-ray image at variable exposure parameters using Monte Carlo computing methods.

1.4.2.1 Optimisation using physical phantoms

Early studies typically focused on comparing the performance of CR with rare-earth film-screens. Cook *et al.* [23] used threshold contrast detail detectability phantoms and concluded there was no measurable difference between the two modalities. Similarly, Schaefer-Prokop *et al.* [24] compared CR with film-screen and selenium digital systems. An anthropomorphic chest phantom superimposed with simulated micro-nodules and normal nodules was imaged on each system and the performance compared. No significant difference was found between the modalities for normal sized nodules, but the selenium system outperformed both CR and film-screen for micro-nodule detection.

Hufton *et al.* [25] investigated optimisation of X-ray beam parameters for paediatric chest imaging and showed that doses could be lowered using CR image receptors and high tube potentials. However, other early published work [26-29] contradicted these findings. For example, Dobbins *et al.* [27] studied images of a phantom obtained with CR and standard screen-film and compared them to evaluate observer threshold perception performance with a modified contrast-detail technique. It was found that CR necessitates about 75%-100% more exposure than screen-film radiography to optimally match performance for detection of objects 0.05-2.0 cm in diameter. They also found that performance with CR images was better at lower kilovoltages. Chotas *et al.* [29] measured SNRs in the lung, spine and diaphragm regions of chest phantom CR images and found low tube potentials increased the SNR by about 15%.

Honey *et al.* [30] have recently reported an optimum X-ray tube voltage range for chest radiography with CR of 75-90 kVp; however the phantom used in their study did not mimic the different anatomical regions seen in chest radiography.

Generally, clinical chest post-processing algorithms need to identify these regions to function correctly. Thus their work focused on X-ray tube voltage optimum for the detector (CR phosphor), rather than the complete CR system for chest radiography.

Doyle *et al.* [31] measured the CNR in various areas of a chest phantom developed by Chotas *et al.* [32], and measured a figure of merit (FOM) that factored this in along with effective dose (ED), (FOM = CNR^2/ED). This was used to assess which tube potential, filtration and scatter rejection method provided the highest FOM. The CNR was highest in the lung and heart regions at 60-80 kVp, but highest at 90-110 kVp in the abdomen, and an air gap technique was the preferred scatter rejection method. It must be borne in mind that the phantom used here had non-chest like materials (copper), and no anatomical data. This may have influenced the results obtained.

Other work by Doyle *et al.* [33] used the same chest phantom to optimize a direct digital system. SNR indices in the heart, lung and diaphragm regions of the chest were examined with tube potential for three radiographic techniques, using a grid, air gap technique, and no scatter rejection. Results showed that lower tube potentials and the use of a scatter rejection method were superior for a matched ED. However, questions over the credibility of the phantom due to lack of anatomical detail remain.

Our group [34] examined optimum tube voltage and amount of added copper filtration for processed chest radiographs obtained with an Agfa 75.0 Computed Radiography (CR) system. A CNR ratio was measured in the lung, heart/spine and diaphragm compartments of a chest phantom [35] using various tube voltages and amounts of Cu filtration. CNR was derived as a function of air

kerma at the CR plate and effective dose, and a tissue-to-rib ratio (TRR) was measured to investigate which tube voltages suppress the contrast of this overlying structure. Although processing algorithms affect the signal and noise in a way that is hard to predict, it was found that for a given set of processing parameters, CNR was related to the plate air kerma and effective dose in a logarithmic manner. For imaging of the lung region, a low voltage (60 kVp) produced highest CNR whilst a high voltage (125 kVp) produced highest TRR. In the heart/spine region, 80 - 125 kVp produced highest CNR, while in the diaphragm region 60 - 90 kVp produced highest CNR. Hence, for chest radiography with this CR system, it was concluded that the optimal tube voltage was dependent upon the region of clinical interest. Of the filters tested, 0.1 mm Cu thickness was found to provide a statistically significant increase in CNR in the diaphragm region with tube potentials of 60 and 80 kVp, without affecting CNR in the other anatomical compartments. The phantom used in this study contained very little anatomical detail, therefore the limitations in representing a clinical chest radiograph must be considered.

1.4.2.2 Optimisation using computer methods

Various Monte Carlo studies have been performed in order to find optimum exposure techniques with images simulating projected patient anatomy. Monte Carlo methods are essentially a class of computer algorithms that rely on repeated random sampling to compute their results. This makes it useful for modelling random events such as photon absorption and scatter.

Sandborg *et al.* [36] used a Monte Carlo model of a voxelized phantom and CR system to assess appropriate tube voltage and concluded lower kilovoltages than that recommended for film-screen radiography (125 kVp) were optimum for

a given stochastic radiation risk. Similarly, Ullman *et al.* [37] used differently sized voxelized phantoms in a Monte Carlo model to investigate optimum X-ray beam parameters, filtration and anti-scatter grids for chest CR. They measured the SNR in different regions of the chest to derive optimum conditions. They concluded that the choice of tube voltage depends on whether SNR or contrast ratios (ratio of tissue contrast to that of bone: C(tissue)/C(bone)) are the most relevant for the diagnostic task. They found that SNR increased with decreasing tube voltage, but C(tissue)/C(bone) increased with increasing voltage. However, if SNR and contrast indices are to be taken as the gold standard for determining optimum X-ray beam conditions, there must be a positive correlation between clinical image quality and these indices.

There is a lack of studies investigating optimum conditions for chest CR using clinical observers to score and grade images using images containing the essential anatomical noise. Work by Sandborg *et al.* [36] attempted to answer this question, but they used a voxelized phantom containing very little clinical anatomical data and only 10 patients were graded by radiologists using a Visual Grading Analysis (VGA). Grading was assessed according to European criteria [5]. The resulting Monte Carlo chest images were of low resolution, but the work concluded that lower tube voltages were optimum for a given effective dose. The VGA method has been analysed and validated as an appropriate method to assess the image quality of chest radiographs independently by Mansson *et al.* [38,39], Tingberg *et al.* [40] and Sund *et al.* [41].

Other Monte Carlo studies [42-45] have also investigated optimum tube voltages for chest imaging, but (and in common with those described above) they calculate energy imparted to the CR storage phosphor only; none attempt to model the frequency dependent noise of the CR system in any way. Also, the

organs of the voxel phantom [46] used in most of the Monte Carlo studies are identified with only one of four tissue types, namely soft tissue, bone, bone marrow, lung tissue and air, limiting the contribution of anatomical noise. The resolution of this voxel phantom is also very coarse (voxels are approx 4mm long x 3 mm wide x 3 mm thick). It will therefore produce images of much lower spatial resolution than a real CR image (typical pixel pitch 0.17 x 0.17 mm).

There is very little literature that uses CT data sets to produce the computerised voxelated phantom that can be utilised to generate images for optimisation studies. Fanti *et al.* [47] used human CT data to simulate film-screen images of patient skulls using pencil X-ray beams, but they did not include radiation scatter or system noise.

1.4.2.3 Optimisation with ALARP in mind

Projected anatomy is the limiting factor for chest radiographic evaluation, and its presence in images used for optimisation studies is essential. However, the radiation dose/image quality relationship must not be ignored from a governance perspective. Therefore, CR system noise, as well as anatomical noise, must be present in any image used for optimisation. For every examination, doses to the patient must be kept ALARP whilst obtaining the necessary level of image quality. Bath *et al.* [48] have suggested that clinical images should be used to obtain the highest validity for optimisation, and as such describes a method to simulate lower exposures by adding frequency dependent noise to an original image [49]. However, this methodology only works for a given X-ray beam quality so full optimisation would still necessitate repeat exposures of the same patient, thus increasing the stochastic risk of inducing cancers.

1.5 Thesis outline

It is vital that anatomical noise is present in images used to derive optimum radiographic technique. System noise must also be included if dose reduction studies are to be performed. This thesis proposes the development and validation of a digitally reconstructed radiograph (DRR) computer simulation that models conventional 2-dimensional X-ray CR images created from CT data. It is proposed that clinical chest CT data sets (i.e. the 'virtual patient' or 'computerised phantom') will provide realistic anatomical structures that feature as projected anatomy (anatomical noise) in the final DRR. Methods of DRR generation in the literature will be reviewed and the most appropriate for this work discussed. Each DRR will essentially be a pixelated image of attenuation; that is, a 2-dimensional map of X-ray attenuation projected through the virtual patient at any tube potential desired; free of radiation scatter and system noise. This is vastly more flexible than adapting and utilising existing clinical CR chest images that inherently contain scatter and noise acquired at a given beam quality, receptor kerma and scatter rejection technique; to simulate different exposure factors from the original image would be very difficult. It is proposed that clinical values of CR frequency dependent system noise (including quantum noise) and radiation scatter can be added to the raw DRR; both of which will be examined and discussed. The whole DRR simulation system will then be tested and validated using phantom and real patient CR images.

As there are currently no recommended optimum X-ray exposure parameters for chest imaging with CR systems, this research will also focus on using DRR generated images to derive optimum parameters by utilising experienced image evaluators to grade and score DRRs presented at different X-ray tube

potentials, doses and scatter rejection techniques. As the resulting images will be computer generated using retrospective CT data sets, no extra exposure risk to patients exists.

Based on the results, radiographic technique for CR chest radiography can then be optimised and standardised in the Hull and East Yorkshire Hospitals NHS Trust.

A synopsis of each thesis chapter is as follows:

Chapter 2 – DRR Simulation Methods

Introduces the concept of a DRR simulated image, and their current uses. The two main methods of generation, ray casting and splatting are discussed and critically analysed. The chapter concludes with a justification of the method used in this research.

Chapter 3 – Methods to add scattered X-radiation to digitally reconstructed radiographs

This chapter critically analyses methods described in the literature for measuring and adding scattered X-radiation to digitally reconstructed radiographs. The chapter finishes with a justification for the method used in this research.

Chapter 4 – DRR Generation

Discusses all practical methodologies used in the research, such as X-ray spectra generation, scatter and noise addition. The majority of this chapter is taken up by development of the DRR simulation software, together with the results and discussion.

Chapter 5 – Validation of DRR Generated Images

This chapter discusses how the DRR simulation model generated in Chapter 4 was validated. Figure of merits (FOMs), such as histogram, signal to noise ratio (SNR) and tissue to rib ratio (TRR) analysis were performed on DRR images. These were compared with FOMs derived from real CR images of phantoms and patients and statistical tests carried out, showing that quantitative validation is possible. Expert image evaluator opinion was sought, and it has been shown that DRR images contain realistic projected anatomy. Limitations of the computer model are also discussed. The chapter concludes that the DRR simulator is adequate for optimisation studies.

Chapter 6 – Optimisation study using average and obese patient DRRs

This chapter discusses the use of DRR generated images to optimise CR chest radiographic technique. Expert image evaluators scored and graded the images; optimum tube potential, receptor dose and scatter rejection techniques are subsequently proposed. Change in clinical practice in the Hull and East Yorkshire Hospitals NHS Trust due to the results of this study are discussed.

Chapter 7 – Conclusion

This chapter gives a summary of the conclusions of work presented in this thesis. Future work and further experiments are also discussed.

Chapter 2: Digitally Reconstructed Radiograph Computation Methods

2.1 Introduction

A digitally reconstructed radiograph (DRR) is a simulation of a conventional 2D X-ray image, created from CT data. A radiograph, or conventional X-ray image, is a single 2D view of total X-ray absorption through the body along a given axis. DRRs are created by summing CT attenuation data along a ray from each pixel to the simulated x-ray source; in general, DRR computation is a volume rendering process and a variety of techniques are available to compute a DRR.

A relatively intuitive method of DRR computation is that of ray casting, which is an image-order volume rendering method that tries to find the intersections of a ray with objects in the scene (e.g. voxels in a CT data set). Alternatively, object-order volume rendering techniques are also used for DRR computation; these include Fourier volume rendering [50, 51], splatting and cylindrical harmonics [52]. However, by far the most common types of volume rendering are ray casting and splatting. This chapter contains a description of current DRR applications and a discussion of the two most common DRR computation techniques, and their uses in the literature. It concludes with a justification for the volume rendering method used in this research.

2.2 Current applications of DRRs

2.2.1 Radiation Therapy

In radiation therapy, high energy ionising radiation (primarily photons and electrons), is used to treat cancer by sterilising tumour cells and precluding them from reproducing. It is vital that as little healthy tissue is exposed as possible in order to avoid unessessary damage. As such, patient alignment is very important during treatment. DRRs play a prominent role in ensuring correct patient positioning [53-55], and are computed from CT data of the patient (acquired prior to treatment in order to calculate each patient treatment plan). The DRR can then be compared with images acquired during treatment [56-59], such as portal images captured from the mega-voltage (MV) treatment beam, or kilo-voltage (kV) images taken just before treatment, to ensure the correct part of the anatomy is being treated.

2.2.2 Two dimensional – three dimensional registration

Recently, 2D/3D registration has been gaining in popularity because of the increased use of computer-aided surgery (CAS). CAS systems match pre-operative images to anatomy during the operation. This allows the relative position of surgical tools and anatomy to be determined.

Essentially, the initial position of the patient is taken from a CT scan, and DRRs are computed. During the operation, fluoroscopic images are acquired and compared with the DRR. The patient can then be moved so that there is minimal difference between the DRR and fluoroscopic images.

2.3 Ray casting

The simplest way to project the image is to cast rays through the volume using ray casting. In this technique, a ray is generated for each desired image pixel. Using a simple camera model, the ray starts at the centre of the projection of the camera (usually the eye point) and passes through the image pixel on the imaginary image plane floating in between the camera and the volume to be rendered. The volumetric data is transformed by an attenuation function so that all voxels are converted to an attenuation coefficient. The ray is sampled at regular intervals throughout the volume, interpolated at each sample point, and the process repeated until the ray exits the volume. The output of each ray is deposited in the corresponding image pixel. This process is shown in Figure 2.1.

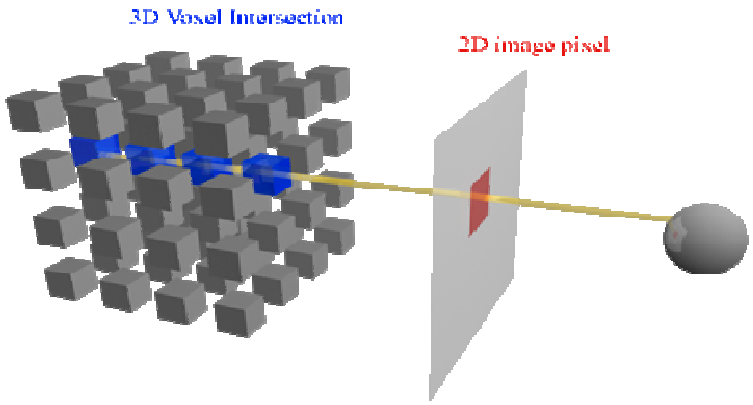

Figure 2.1: Graphical depiction of the ray casting technique. The ray is sampled at each voxel and the resulting output of the ray is deposited on the ray's designated pixel in the image.

Siddon [60] described an exact and efficient algorithm for calculating the radiological path through a three dimensional CT array. Voxels were depicted as the intersection volumes of orthogonal sets of equally spaced, parallel planes, rather than independent elements. To illustrate this in the two

dimensional case, pixels were considered as the intersection areas of orthogonal sets of equally spaced, parallel lines, as shown in Figure 2.2.

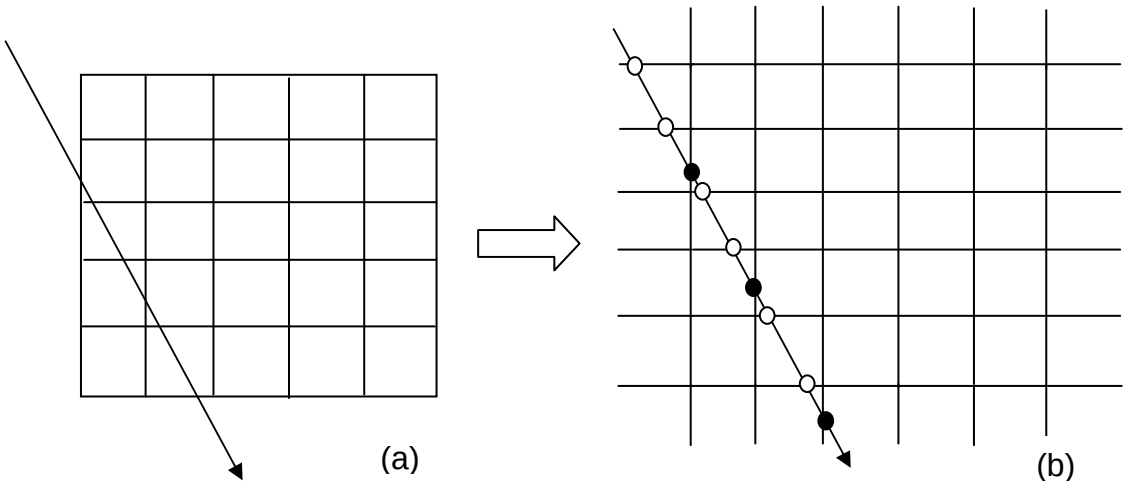

Figure 2.2: (a) Pixels of a CT array (in reality this would be a 3D voxelated array), and (b) pixels as intersection areas of orthogonal sets of equally spaced lines. The interactions of the ray with the pixels can be depicted as interactions with the lines. The intersections of the rays with the lines are given by two equally spaced sets: horizontal (open circles) and vertical (filled circles).

The intersections of the ray with the lines were calculated, rather than intersections of the ray with the individual pixels. As the lines were equally spaced, it was only necessary to determine the first intersection and generate all others by recursion. The intersections of the ray with the pixels was a subset of the intersections with the lines. Identifying that subset allows the radiological path to be determined. For an array of N^3 voxels, considering the planes rather than the voxels allows the algorithm to scale with the number of planes (proportional to N), rather than the number of voxels (proportional to N^3).

Sherouse *et al.* [61] discussed the computation of DRRs for use in radiotherapy treatment design. The DRR algorithm was designed with an emphasis on image quality, rather than speed. The DRR software had three major distinguishing features:

- When casting rays through the CT volume to form an image, tri-linear interpolation was used between slices, and the contribution of every intersected voxel was taken into account.

- A heuristic approach was used to sort the attenuation along the ray path into photoelectric and Compton components.

- The desired resolution and pixel count of the output image can be specified in terms of readily understood parameters such as film size.

The basic ray casting algorithm used by Sherouse *et al.* was a slightly modified version of that developed by Siddon [60]. The algorithm traced rays from the source through a 3D model of the patient made up of voxels determined from CT scans. Each voxel was characterised by its position in the patient model and its dimensions. In contrast to Siddon, the CT numbers were conceptualised as point samples in a continuous space (as opposed to voxelised values) whose physical positions were taken to be at the eight corners of the voxels rather than their centres, and the attenuation coefficient (derived from effective CT number) was determined from these corners using tri-linear interpolation, as shown in Figure 2.3.

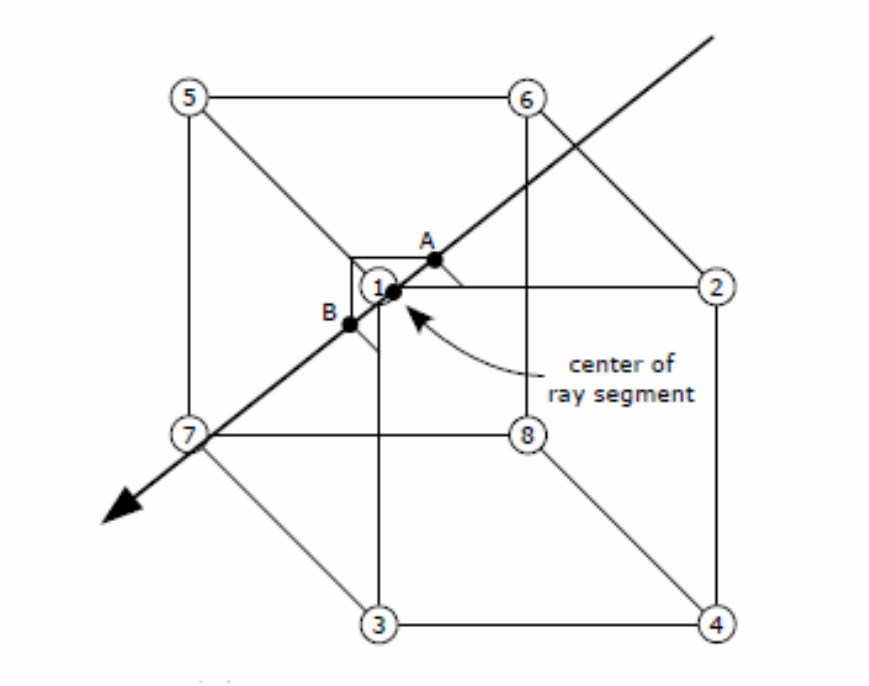

Figure 2.3: Tri-linear interpolation. The ray is interpolated along each of the three axes to the midpoint of the ray segment and involves all eight corners.

Tri-linear interpolation involves interpolating along each of the three axes to the midpoint of the ray segment and involves all eight corners. Other interpolation methods were examined, such as nearest neighbour interpolation where the CT number is taken to be the value of the nearest voxel corner (Figure 2.4(a)), straight average in which the CT number is taken to be the average of the values of all voxel corners (Figure 2.4(b)), and linear interpolation where the nearest neighbours are chosen depending on the intersection of the ray and corners (Figure 2.4(c)). The resulting CT number is found from linear interpolation between the two corners. Figure 2.4 illustrates these interpolation methods.

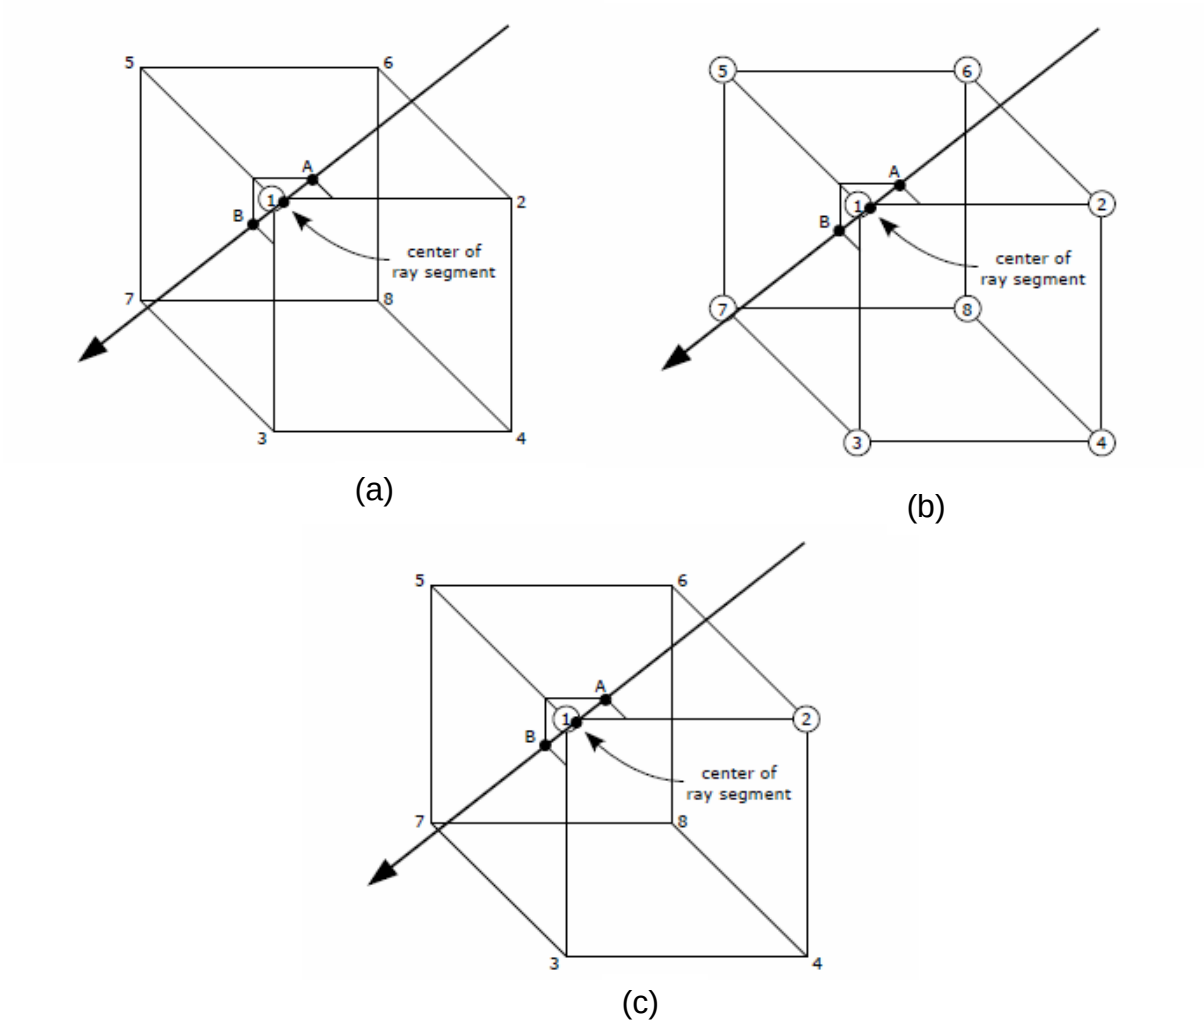

Figure 2.4: Nearest neighbour interpolation where (a) the CT number is taken to be the value of the nearest voxel corner, (b) straight average in which the CT number is taken to be the average of the values of all voxel corners, and (c) linear interpolation where the nearest neighbours depending on the intersection of the ray and corners are chosen and CT number is found from linear interpolation between the two corners.

Tri-linear interpolation was found to provide superior image quality compared to the other methods of sampling, but the final image contained aliasing artefacts owing to ray divergence and the inaccuracy of this method of interpolation (using simple rays is not a realistic situation as real X-rays are not rays). Finally, each ray was traced to the centre of each pixel in the DRR. A piecewise approximation to the line integral of effective linear attenuation coefficient was accumulated along the ray. The length of each segment, computed from the coordinates of the entrance and exit points, was multiplied by the voxel

attenuation coefficient. The sum of the products (ray segment x attenuation coefficient), for each ray was used in an exponential attenuation expression to estimate the beam attenuation. It was concluded that the tri-linear algorithm preserved as much spatial resolution as possible, while explicitly including contributions from every intersected voxel.

Dong *et al.* [62, 63] studied a procedure that used megavoltage (MV) DRRs calculated from patient 3D CT data as a reference image for correlation with on-line electronic portal images to detect patient set-up errors. The DRR calculation method used a ray casting algorithm that calculated the primary transmission only. Voxel values were not conceptualised at the corner of each voxel as per tri-linear interpolation [61], but rather the voxel was considered to be a cubic entity with a uniform value throughout. The primary transmission fluence was calculated using the total linear attenuation coefficient encountered by the ray. The line integration was evaluated along the ray from the virtual source, casting through the patient to the position of the imaging panel. However, only primary transmission was calculated at megavoltage energies by deriving the relationship between CT number and linear attenuation coefficient. CT numbers for air, lung, fat, water, muscle and bone were obtained from a CT calibration phantom of known composition, scanned at a given CT tube potential. Linear attenuation coefficients were then calculated for the tissues using data obtained from National Institute of Standards and Technology (NIST) publications [64] (for other CT numbers linear interpolation and extrapolation was used).

Generation of DRRs using ray casting has been used by Milickovic *et al.* [53] for applications in Brachytherapy. The methodology involved simulation of individual X-rays which were generated from a virtual X-ray source and

projected through the patient CT data set. The attenuation data for an individual X-ray was accumulated while the ray travelled voxel-by-voxel through the data set. The intensity of the outgoing ray was calculated by summing the products of the linear attenuation coefficient of each voxel, and the length of the ray segment in each voxel (i.e. μ x d), as shown in Figure 2.5.

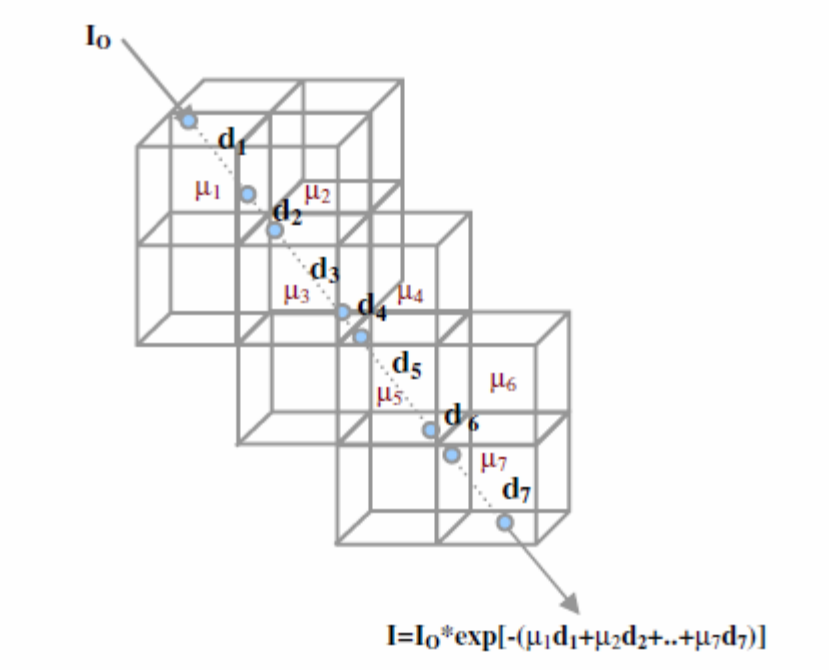

Figure 2.5: Ray casting method used by Milickovic et al. Each ray is segmented in terms of its length traversed through each voxel. This length is then multiplied by the value of the voxel attenuation coefficient. The sum of these products is calculated and used in the exponential function.

As with Dong *et al.* [62], tri-linear interpolation was not used. Together with the knowledge of the input intensity, the exiting intensity was calculated. Linear attenuation coefficients of the various tissue components (i.e. at a given CT number) contained within an appropriate phantom were derived using an XCOM program [65]. The linear attenuation coefficients of the other CT numbers were calculated from linear interpolation and extrapolation. Milickovic *et al.* also generated images using another common DRR volume rendering technique – splatting (see next section of this chapter), but this method was not used as the

resulting images exhibited inferior image quality; brachytherapy catheters were not reconstructed to the relevant degree of resolution.

Killoran *et al.* [66] presented work on optimising DRR reconstruction of the chest for virtual radiotherapy simulation. One conclusion was that spatial resolution in the final DRR image is limited by the underlying CT data set, namely the slice thickness, and suggest that if possible, sub millimeter slice thicknesses should be used to maximise image quality. The DRR calculation method used a virtual simulation by which the 'model' was a 3D data set derived from the CT scan, which was sampled via interpolation to create a contiguous data set of uniform spatial dimensions. A major influence to DRR image quality was CT slice thickness, and in their study 3 mm thicknesses were used. However, this virtual model had in-built interpolation, so DRR calculations were relatively fast (even though 3mm slice thicknesses is relatively low resolution). In a sister paper by Giraud *et al.* [67], it was shown that DRR image quality only gradually decreased as the pitch factor of the CT scanner was rapidly increased, and that image quality was superior with sub-millimeter slices. In the context of this thesis, Giraud *et al.*'s work in general is encouraging, as it is possible to reconstruct slice thicknesses down to 0.8 mm with scanners used in Hull and East Yorkshire NHS Trust.

Bifulco *et al.* [68] looked at estimating out of plane vertebra rotations on radiographic projections using CT data. DRRs from CT data were computed using the ray casting method and tri-linear interpolation to estimate the absorption coefficients of the voxels between slices within the CT volume. Importantly, it was suggested that ray casting techniques produce better quality results than other render-based techniques (such as splatting).

2.3.1 Sheared object space

To speed up the ray casting process, the data volume can be sheared [112] so that rays can pass through object space in a parallel fashion, rather than a computationally expensive diverging manner (i.e. rays entering the volume at various angular displacements). Sheared object space shifts the data slices so that the rays are perpendicular to the slices. A simplified process using parallel rays entering a volume at a given anglular displacement is shown in Figure 2.6.

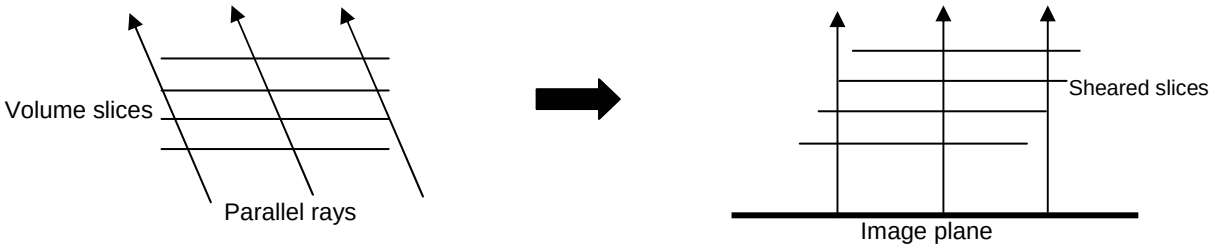

Figure 2.6: Shearing of volume slices to allow the parallel projection of rays. In reality (not shown here), each ray would enter the volume with a different angular displacement (due to divergence) so volume slices would require a non-linear shear.

This technique is relatively fast if the volume is sheared prior to ray casting. It also creates the same level of image quality as diverging rays. However, there is a requirement for storing multiple copies of the volume, which leads to a memory overhead.

2.4 Splatting

Splatting is a technique that trades quality for speed. It has been used to directly render volumes of various grid structures [69,70] for both scalar [71,72] and vector fields [73]. The basic algorithm, first described by Westover [71], projects each voxel to the screen and composites it into an accumulating image.

It visits the voxels in either a back-to-front or front-to-back order, with closer voxels overwriting farther voxels. This process is shown in Figure 2.7.

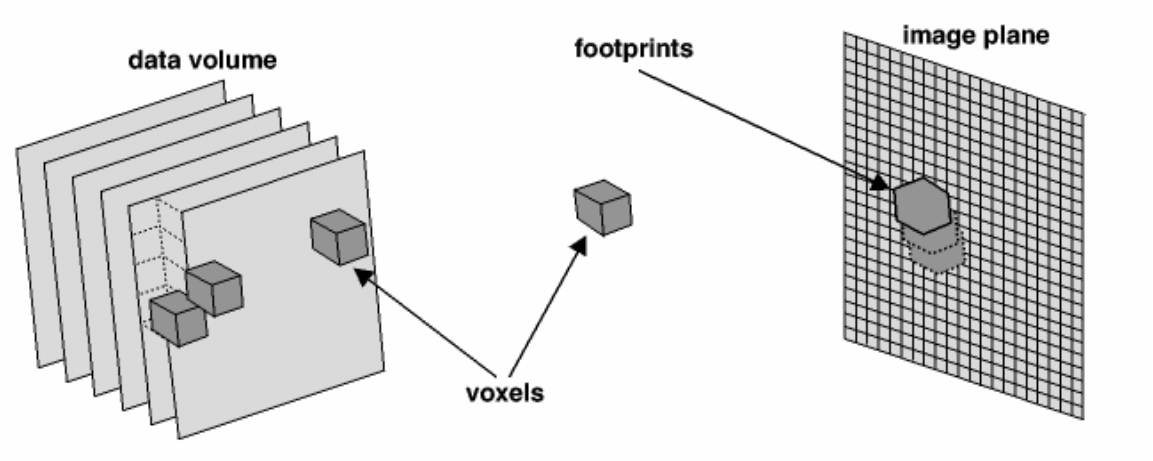

Figure 2.7: Voxels from the data volume 'splatted' onto the image plane.

Splatting is an *object-order* algorithm: the resulting image is built up voxel by voxel. This is in contrast to ray-casting, which is an *image-order* algorithm that builds up the resulting image pixel by pixel. As each voxel is projected onto the image plane, the voxel's energy is spread over the image raster using a reconstruction kernel centered at the voxel's projection point. The reconstruction kernel (or kernel footprint) is used for all voxels, and the resulting image is generated by the following process: (1) the coordinate of the projected voxel in the resulting image is calculated; (2) the kernel footprint is centered around the voxel; (3) the footprint is projected to the image screen. This projection is known as *kernel splatting*, or simply, *splatting*. Individual splats are stored on a screen buffer, and composited with splats from other slices within the data volume on a composite buffer, as shown in Figure 2.8, to build a composite image.

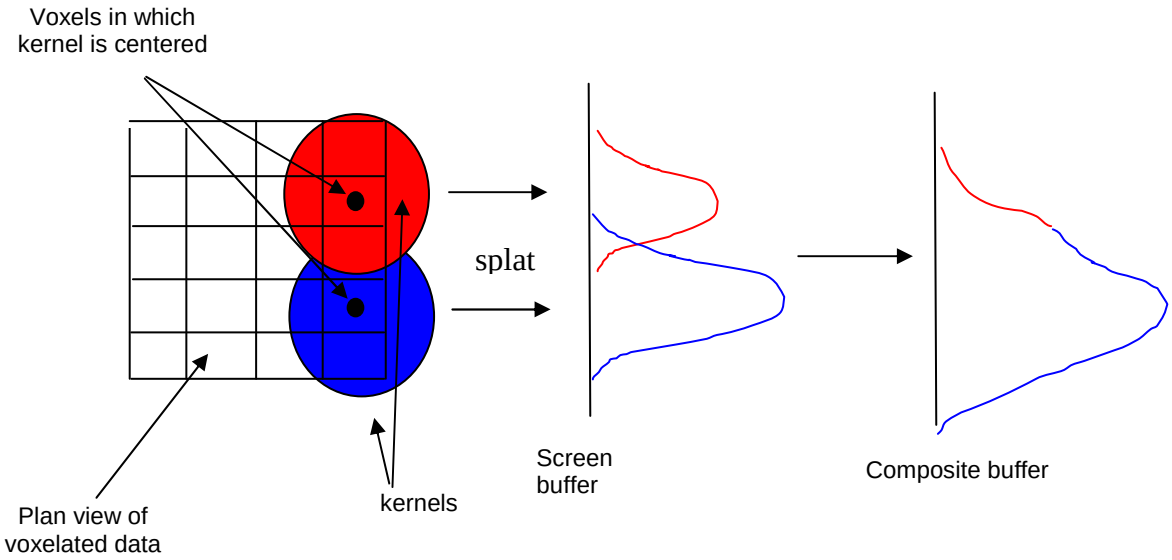

Figure 2.8: Object-order volume rendering with splatting Guassian kernels.

In Figure 2.8, voxels (object points) are traversed in one of two ways; either front-to-back, or back-to-front order. Each point is splatted onto the screen buffer, and then composited with the image that has already been built. However, with this approach, kernels overlap with other slices, which reduces image quality, such as loss of depth resolution, and is known as ‘bleeding’.

An attempt to solve this problem [74] led to what has become known as ‘sheet-buffered splatting’. In this approach, each slice within the voxelated data is separated from its two adjacent ones, and all kernels are sliced and then projected. The process is shown in Figure 2.9.

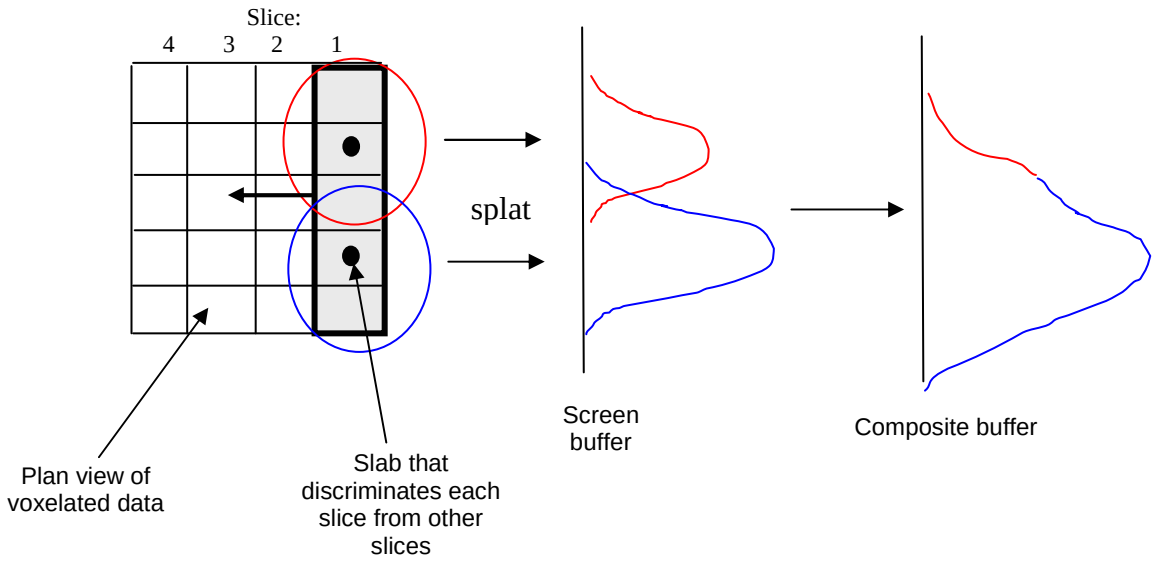

Figure 2.9: Sheet buffered splatting. The slab moves across the voxelated data separating each slice. All kernels are effectively sliced, so that kernels in slice 1 do not 'talk' to kernels in slice 2, and so on.

Sheet-buffered splatting eliminates bleeding between slices, but not blurring between resulting pixels in the composited image (this problem common to all splatting techniques). Footprint (kernel) projection is usually precomputed and represented as a 2D lookup table. The 2D table is centered at the projection point and sampled by the pixels which lie within its extent. Each pixel composites the value it already contains with the new value from the footprint table. Under most conditions the footprint table can be computed once and used unmodified for all voxels.

Recently, Birkfellner *et al.* [75] have modified the conventional splatting algorithm to a so-called 'wobbled splatting' algorithm. This reduced aliasing by distortion of the voxel positions in the 3D data, without the need to apply anti-aliasing kernels.

2.5 Ray casting and splatting: a comparison

There are inherent problems with splatting not encountered with ray casting. Splatting by its very nature requires each kernel centered over a specific voxel to be projected at the image plane. In doing so, image composition is performed on a per-splat basis resulting in incorrect image pixel values where they overlap. However, to ensure the final image is smooth, splats must overlap, resulting in blurred images as well as in-correct pixel values.

Another disadvantage with splatting is sampling voxelated data. With ray casting, each ray is sampled (e.g. tri-liner interpolation) as it projects through the data. However, with splatting, this must be done via reconstruction kernels, which can be a problem because image quality is affected by the size, shape and type of the kernel used. Work has been carried out by various authors to derive optimum kernel features, but no standard has been found. Laur and Hanrahan [76] changed the size of the splatting kernel based upon the voxel in which it is centered, and Mao [69] analysed spherical and ellipsoidal kernels with varying sizes. To date most splatting implementations have used Gaussian reconstruction kernels, but others can be used that influence final image quality.

Splatting is advantageous to ray casting in terms of computational speed. In ray casting, each ray is sampled along its projection which is very computationally expensive. Splatting on the other hand pre-computes each kernel footprint, saving a large amount of time in comparison. Splatting is also an object-order algorithm, allowing static parallel composition. This is not permitted with ray casting, as the ray usually needs to sample many parts of the voxelated data. Splatting can also cull voxels that contribute nothing to the final image (such as those outside the volume of interest), but ray casting has to take into account all

voxels the ray samples, worthless or not. Alakuijala *et al.* [77] directly compared splatting with ray casting and concluded that ray casting performs better in terms of resulting image quality for divergent beams.

In summary, although ray casting is much more computationally expensive, it is not as susceptible as splatting to blurring or incorrect pixel value artefacts. Splatting also requires kernel reconstruction (the foot-print) of which there is no optimum method. This source of error does not exist for ray casting.

2.6 Aliasing in volume rendering

Volume ray casting algorithms project rays from a source through the data. This process can cause aliased signals in the resulting image [78]. This is illustrated in Figure 2.10, where the data volume is shown as a lattice of points.

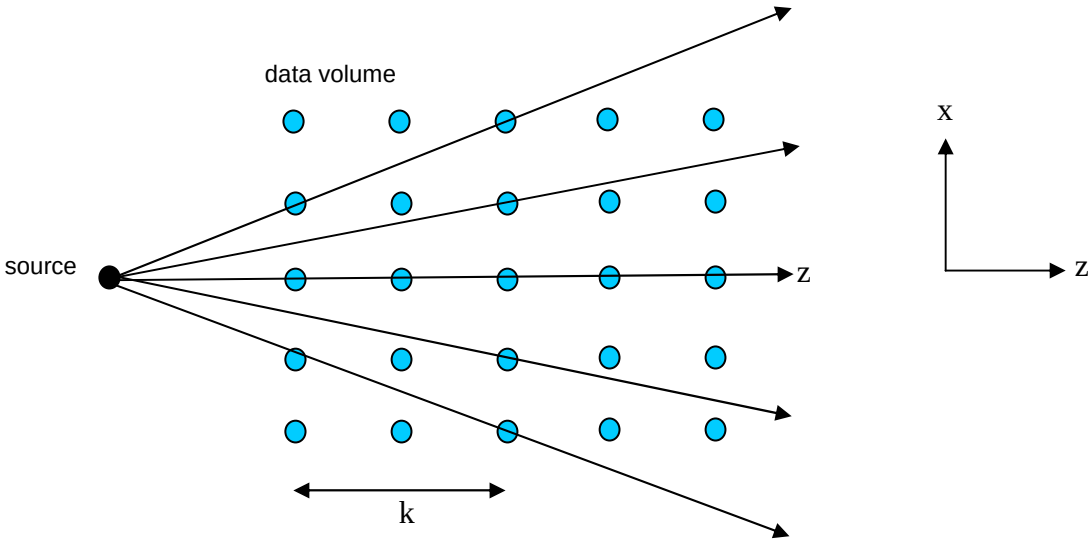

Figure 2.10: Rays projecting through lattice.

Figure 2.10 shows the rays diverging through the lattice (voxels) with the source looking down the z axis. Aliasing in the reconstructed image results from insufficient sampling of the lattice. Notice in Figure 2.10 there is a distance along the z axis denoted k, in which lattice resolution and the sampling rates of

the lattice by the projected rays are equal. At distances less than k, each ray can sample more than one voxel, and so no aliasing artefacts occur. However, at distances greater than k, the distance between adjacent rays (in the x-direction) is greater than one voxel, the data is under-sampled and therefore aliasing artefacts can occur. Some ray casting algorithms minimise aliasing by using low-pass filtering that reduces the frequency content of the volume by employing reconstruction kernels that become larger as the rays diverge [79,80]. However, this filtering results in the loss of information from the data volume, artificially blurs the reconstructed images and deteriorates image quality. All ray casting methods discussed in this chapter are prone to aliasing, and therefore must use low-pass filtering.

Anti-aliasing in splatting is typically accomplished by increasing the area of each splat in proportion to the distance the splat originates from the source, so that the data volume is not sub-sampled. However, increasing the splat size results in increased blurring and less than optimal image quality [75].

2.7 Conclusions – justification of the volume rendering method used in this research

Most, if not all of the literature is heavily biased towards speed when calculating DRRs. This comes at the expense of image quality. Of the main types of DRR calculation methods, ray casting and splatting are by far the most popular. Splatting has been developed to increase the speed of calculation, but is prone to aliasing, blurring and incorrect pixel values. Anti-aliasing algorithms have been introduced but these inherently increase blurring further. Ray casting

produces superior images but usually take much longer to compute than splatting algorithms. All ray casting techniques in the literature use interpolation (mainly tri-linear, but treating each voxel as a uniform cubic structure is also well used) between voxels to increase speed, but this introduces aliasing which must be reduced with low pass filters that cause blurring.

For this research, it is postulated that to produce optimum image quality, rather than using simple rays to render 3D CT volumes, realistic pencil beams projected through the data should be used; this will avoid any confusion of which sampling method to use that is inherent with rays. In this work, as each pencil beam passes through each CT slice in the posterior-anterior (PA) direction, the proportion (area) of each pencil beam (in the centre of a given PA slice) in each neighbouring voxel will be calculated, and a resulting linear attenuation coefficient will be derived based on the weight of each area. This method uses ‘sampling with areas’ rather than ‘sampling with volumes’; as ray casting is a perspective projection, this assumption is not entirely accurate, since this is a method of mapping three-dimensional (volumetric) points to a two-dimensional plane. However, for a given change in beam area in a voxel, the volume would approximately increase or decrease proportionally. Also, pencil beam area sampling of voxels applies to calculation of weighted linear attenuation coefficents only, and as this applies to the whole 3D pencil beam, one is not moving away from a volumetric calculation.

The intention is that image quality will be optimum and there will be very little artefacts; using pencil beams will ensure the CT data is sampled in a realistic manner with respect to rays. However, because each pencil beam is diverging away from one another, aliasing is still expected in the resulting DRR. Low pass filters will not be applied to minimise this artefact, unless the algorithm proves

difficult to validate. The computationally expensive calculation times can be justified by the fact that the research is aimed at producing a computer algorithm which produces chest X-rays that contain sufficient projected anatomy to allow the optimisation of radiographic technique. Therefore, image evaluators only need to see the end point (i.e. the resulting images) rather than have to operate the system and reconstruct images themselves. Nevertheless, the sheared object space technique will be investigated in order to reduce computing time.

In conclusion, this research will look at validating a ray casting technique using realistic pencil beam projection rather than non-realistic rays, as ray casting has been shown to produce superior image quality compared to that of splatting. However, radiation scattered by Compton processes is not included by any of the DRR methods discussed. The next chapter discusses the implications of this as well as methods of scatter measurement and addition to DRRs.

Chapter 3: Methods to add scattered X-radiation to digitally reconstructed radiographs

3.1 Introduction

This chapter introduces the basic physics of scattered radiation at diagnostic X-ray energies, and analyses methods of scatter measurement in chest imaging. Scattered radiation is not added to DRR images produced by ray casting techniques, so the chapter concludes with a justification for a method used to add scatter to DRR images.

3.2 Scattered X-rays in diagnostic radiology

An important effect in diagnostic radiology involves inelastic scattering of X-ray photons with unbound electrons; a phenomenon called the Compton effect. This process occurs when a photon collides with a loosely bound atomic electron in a snooker ball fashion, with both energy and momentum conserved, as shown in Figure 3.1.

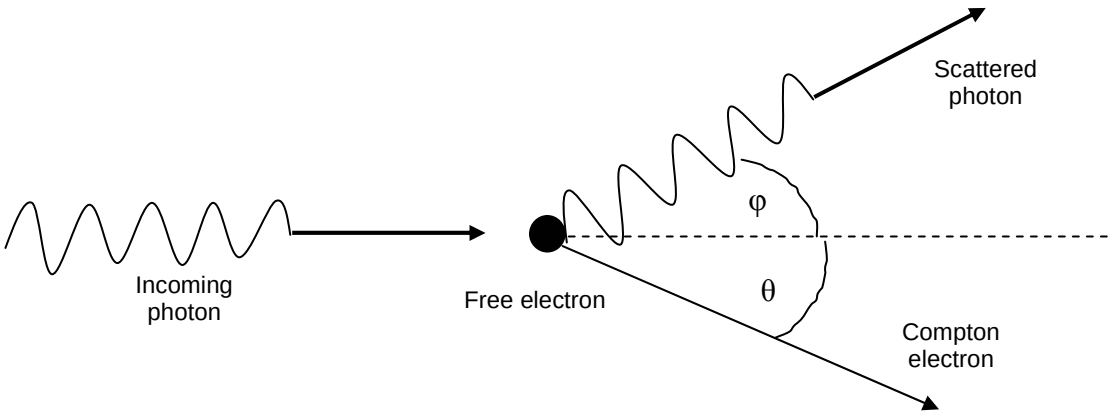

Figure 3.1: Schematic representation of the Compton effect.

Figure 3.1 shows an incoming photon colliding with an electron and being deflected at angles φ and θ, respectively. The proportions of energy and momentum transferred from the photon to the electron are determined by these angles. The electron rapidly loses its kinetic energy through ionisation and excitation within the medium, but the scattered photon, assuming no further interaction, emerges from the medium with lower energy than it entered. This photon may be absorbed by an image receptor (e.g. a CR phosphor) or released into the surrounding environment.

Scattered photons that reach the image receptor do not contain clinically relevant information and degrade the image quality of chest radiographs by creating a non-uniform background that reduces image contrast. The scatter fraction SF, defined as the ratio of the intensity of scattered radiation to that of total (scattered plus primary) radiation recorded on the image, increases with increasing field size and thickness of scattering material. Thus, for large-field examinations of thick body parts, such as the abdomen, scattered radiation can be a serious problem.

In chest radiography, although the equivalent tissue thickness of the lungs is approximately one-half that of the abdomen, scatter is still produced, and is even more appreciable in the spine and diaphragm regions. Also, scatter is present in real CR images and therefore must be present in images produced artificially for optimisation purposes. The DRRs produced from CT data alone will have no scatter present and so it must be added post calculation.

3.3 Measurement of scatter and scatter fractions in chest radiography

Various studies in the literature have investigated measurement of scatter and scatter fractions in chest radiography. Niklason *et al.* [81] measured scatter fractions (SF) in patients and phantoms in various regions of the chest. A single lead beam stop (6 mm thick by 6 mm in diameter) was placed at the surface of the phantom/patient to measure SFs in the resulting chest radiograph. This thickness of the beam stop was used as it is approximately 20 half value layers (HVL) at 150 keV, and therefore sufficient to absorb most of the primary radiation incident upon it. As such, the radiation striking the film in the shadow of the stop was virtually all scattered radiation, as shown in Figure 3.2.

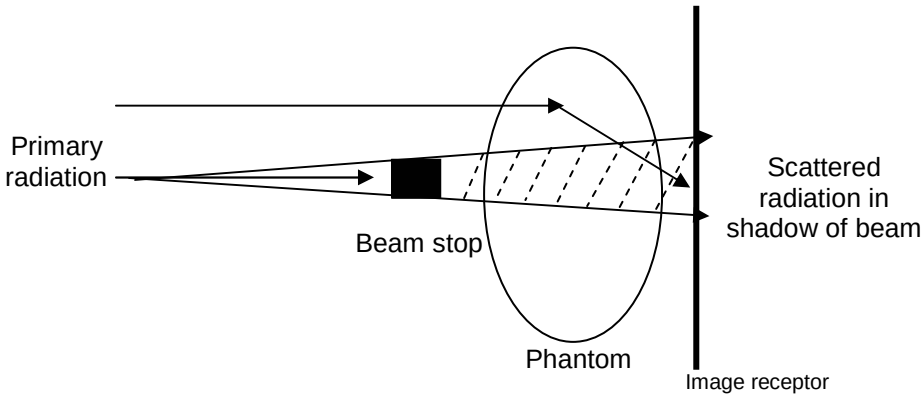

Figure 3.2: Basic depiction of primary radiation being absorbed by the beam stop. Only the scattered radiation is recorded in the shadow of the beam stop, depicted by the dashed area (not to scale).

Radiation recorded on the film surrounding the shadow included both scattered and primary radiation. The optical density (OD) of the film under the beam stop and the surrounding area were measured and converted to exposure using the sensitometric curve for the particular film. SFs were subsequently calculated. This technique for SF measurement has also been used by Bowenkamp *et al.* [82].

Boone [83] describes a method to remove the scatter component from a digital image. SFs were measured in the image using 5 lead disks ranging from 10 to 50 mm diameter. Scatter measurements were extrapolated to 0 mm disk diameter to assess scatter in an open field. The results showed that as the disks get larger, less scattered radiation reaches the image receptor corresponding to the centre of the disk.

In a second paper by Boone *et al.* [84], a method to calculate the point spread function (PSF), and hence the Modulation Transfer Function (MTF) of scattered radiation in a digital imaging system was discussed. Scatter in images was measured with uniform Lucite by firstly positioning the Lucite close to the focal spot, so detector scatter is negligible resulting in the 'primary image', and secondly positioning the Lucite near the detector so that the scatter is detected. The scatter component was measured by subtracting the primary image from the [primary + scatter] image. The PSF and MTF were then derived to analyse the spatial distribution of the scatter for various thicknesses of Lucite. The scatter MTF changed with thickness of Lucite (and therefore patient thickness).

Floyd *et al.* [85] investigated the measurement of quantitative scatter in photo-stimulable phosphor imaging systems. SFs were measured using the beam stop method with five polystyrene phantoms of varying thickness. Similar to Niklason's method, SFs were calculated by measuring pixel values in the shadow of the beam stop and adjacent to it. This was then converted to exposure. Beam stops of various diameters (10, 5 and 3 mm) were also used to give an estimate of SFs at zero diameter (using linear extrapolation). ROI analysis was carried out in the beam stop and around the beam stop. The 'total' value was obtained from the image containing the beam stop. The results show that the thicker the phantom the higher the scatter fraction.

Floyd *et al.* [86] later measured scatter fractions in clinical bedside radiography. A posterior beam stop (PBS) configuration was used that allows measurement of SFs simultaneously with the patient exposure. The PBS is a series of 224 beam stops positioned in an array so that the SF can be measured simultaneously across the entire chest radiograph. It was reported this methodology worked sufficiently, as SFs of different regions of the chest could be measured simultaneously. Of interest, their study found that SFs in the chest were not found to correlate with age, weight or sex. The work described by Floyd *et al.* may be of use in this research as it may allow a 'scatter image' to be generated, which could then be added to the DRR of the virtual patient.

Petrone *et al.* [87] also used a beam stop technique to measure the scatter fractions under three regions of a humanoid chest phantom utilising LaOBr and GdO_2S intensifying screens. Similar to other investigators, measurements in the shadow of the beam stop and the surrounding area were made to evaluate the SFs. The dependence of the point scatter fraction on beam stop size was also estimated by making measurements with 20, 10 and 6 mm diameter beam stops and extrapolating to zero. This allowed for an estimate of error associated with the potential lack of scatter generated in the shadow of the beam stop. SFs decreased only slowly with beam stop size, the largest error being 3.4% in the lung area.

Baydush *et al.* [88] compared the scatter properties of a geometric phantom (consisting of aluminium and copper sheets) versus an anthropomorphic phantom. As per other investigators, an array of lead beam stops 6 mm thick (14 by 16) was placed in front of each phantom to calculate scatter fractions at tube voltages that spanned the diagnostic energy range. Images of 224 equally spaced scatter measurements (25 mm) were obtained and converted from pixel

value to incident exposure using the linear system transfer property of the system previously measured. SFs were determined at each visible beam stop on the chest radiograph using ROI analysis. SFs in each anatomical region (abdomen, mediastinum, lung, retrocardiac and the region behind both the heart and spine (heart/spine)) were determined by averaging the scatter fraction values that were determined for that region.

3.4 Conclusion – justification for the method of scatter measurement

As discussed, scattered radiation will have to be added into each DRR image retrospectively, as ray casting techniques do not compute any component of scatter. To do this an array of beam stops, extensively described in the literature, should be used to calculate scatter fractions across the entire chest radiograph for a range of diagnostic beam energies. It is important to have a realistic chest phantom to carry out this task; that is, either a geometric or anthropomorphic phantom. Intuitively, geometric phantoms due to their simplistic nature are probably not adequate enough to simulate scatter produced by the chest. Therefore, it is proposed that the Alderson RANDO anthropomorphic phantom (described in detail in Chapter 4), can be used together with a beam stop array to measure scatter and SFs in a chest radiographs acquired with different tube qualities. The resulting scatter and SF images can then be added to raw DRRs. Although primarily used for radiation therapy, RANDO muscle has been shown to attenuate diagnostic energy radiation similar to that of water (Shrimpton *et al.* [89]), which in turn is has very similar properties to human muscle. However, this phantom is of fixed size, and therefore does not represent obese patients. It is therefore also proposed that

the phantom be modified by adding a fat equivalent material to represent the obese adult.

Chapter 4: Digitally Reconstructed Radiograph software production, experimental techniques and practical data collection

4.1 Introduction

This chapter discusses the production and development of software used to generate digitally reconstructed radiographs (DRRs), as well as the corresponding practical methodologies, measurements and data collection required, such that, when integrated into the software, adequate DRR images are produced. For purposes of clarity throughout this chapter, it should be noted that all software was written in MATLAB^{TM} Version 8a (The MathWorks Inc, Natick, MA).

The chapter begins with a discussion of phantom and patient image acquisition, X-ray spectra generation and CT data preparation (such as application of appropriate filters and transforming CT voxel values into linear attenuation coefficients). The ray casting algorithm is then addressed, namely accessing and transforming CT data, projecting and sampling of pencil beams, and estimating the energy absorbed in a virtual CR phosphor. Practical methodologies are then discussed, such as the measurement of X-ray scatter and CR system noise. Simulation of lung nodules and fat are addressed, followed by an assessment of how CT noise, DRR resizing and resolution affects final DRR image quality. Finally, the chapter ends with an overview of the entire process.

4.2 Methodology

4.2.1 General practical methodology & phantom/patient image acquisition

The DRR algorithm requires data derived from clinical CR, CT and X-ray systems to enable adequate functionality. The physical characteristics of the X-ray and CR systems were measured in a general purpose X-ray room equipped with a Philips Optimus Diagnost TH ceiling suspended X-ray system (Philips Medical Systems, Surrey, UK) with total inherent filtration equivalent to 3.1 mm of aluminium (see Figure 4.1), and an Agfa CR85-X reader (Agfa, Peissenberg, Germany) with MD4.0 plates (35 cm x 43 cm, effective pixel pitch of 0.17 mm – see Figure 4.2).

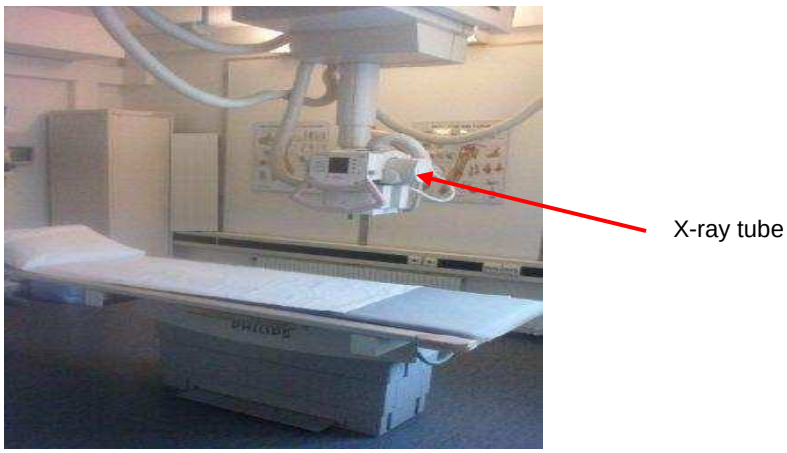

Figure 4.1: Typical Philips X-ray room. The X-ray tube is shown above the bed.

Figure 4.2: Agfa CR85-X digital reader with CR image cassettes shown on the cassette buffer.

4.2.1.1 Phantom used for image acquisition

All phantom acquisitions utilised the chest portion of the Alderson RANDO anthropomorphic phantom. The phantom consists of a natural human skeleton embedded in a synthetic isocyanate rubber with lung substitute and air cavities, simulating the average male, approximately 70 kg. It is constructed from slices of rubber allowing different sections to be used. An image of the phantom is shown in Figure 4.3.

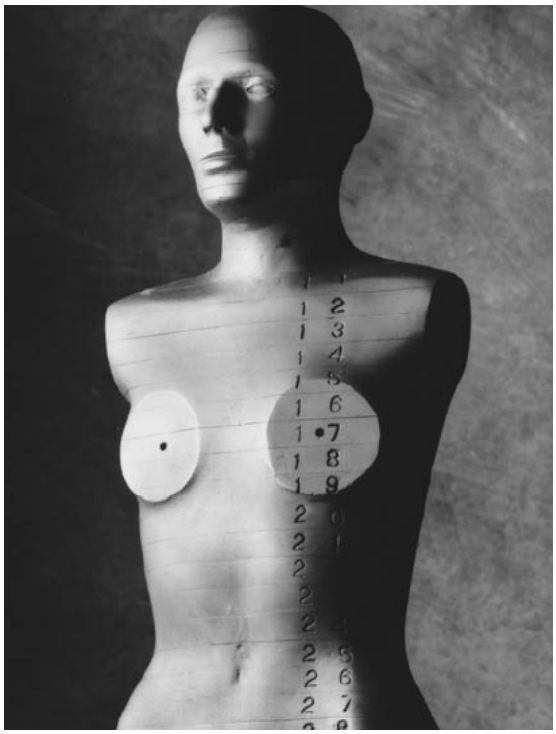

Figure 4.3: Image of the Alderson RANDO phantom. Slices of synthetic rubber are visible.

Although primarily used for radiation therapy, RANDO has been shown to attenuate diagnostic energy radiation in a similar way to water, which in turn has very similar properties to human muscle [89].

Fat was also added (see Figure 4.4(b)) to the phantom periphery to simulate an obese patient. A thickness of 4 cm of fat was added to RANDO on the advice of the superintendent radiographer (private communication, Andrew Stephens, 2009); obese patients typically have approximately this amount of fat around the chest area (as determined from CT images). Grocery store lard was used to simulate the fat. It should be noted that phantom images described below were acquired primarily for validation of the DRR algorithm.

4.2.1.2 Phantom imaging with the CR system

Phantom images collected on the CR system were acquired with a focus to receptor distance of 180 cm and with the CR receptor placed 5 cm behind the

phantom in the cassette holder. A sufficient tube current-time product (mAs) was used to produce a lgM value of 2.00 ± 0.05. The lgM value is a receptor dose indicator displayed on the CR system for every image acquisition; it is recommended by the manufacturer a value of approximately 2.00 should be sought for a correctly exposed chest radiograph (although not necessarily optimum). The X-ray field was collimated to the edges of the phantom and a single CR cassette (digitized in the CR reader with a fixed sensitivity of 400) was used throughout the study. The latter restriction was observed recognizing that different individual receptors do not have exactly matching sensitivities. The cassette chosen demonstrated a median sensitivity (lgM value) of those available. No clinical post-processing was applied.

CR images of the phantom were acquired at tube potentials 50 – 150 kVp in approximate steps of 10 kVp (exact steps of 10 kVp were not possible on the Philips X-ray system). Half and double mAs values were also used at each tube potential to assess the effects of dose reduction and escalation in the resulting DRR images. The phantom set up was chosen to simulate the clinical situation, although only the chest portion of the phantom was used, as shown in Figure 4.4.

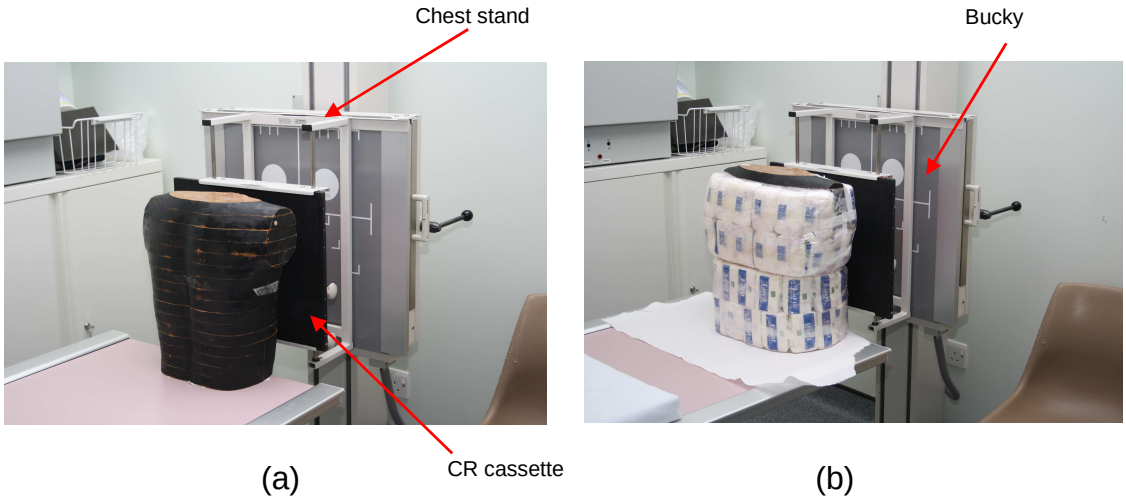

Figure 4.4: The RANDO phantom positioned in front of a CR cassette (a) without added fat, and (b) with added fat. The X-ray tube is out of picture.

4.2.1.3 Phantom imaging with the CT system

To collect CT phantom and patient data, a Philips Brilliance 16 slice multi-detector CT scanner was used, an image of which his shown in Figure 4.5.

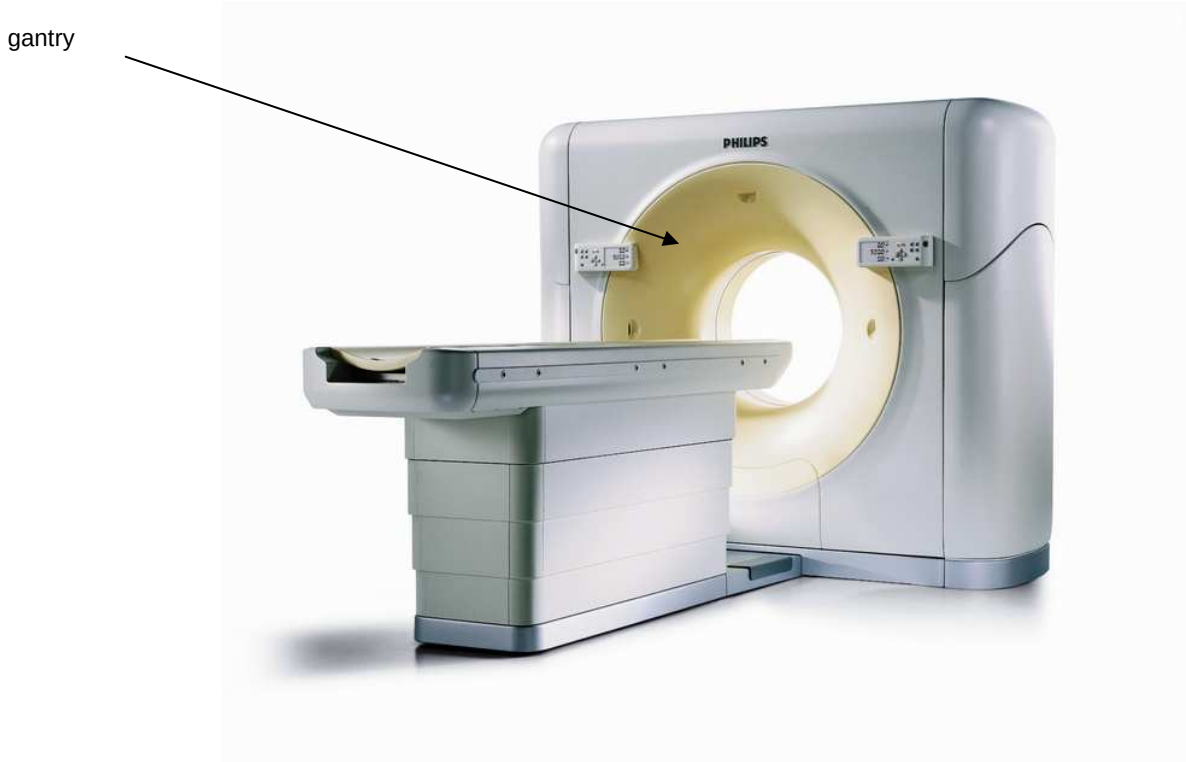

Figure 4.5: Philips Brilliance 16 slice CT scanner.

The chest portion of the RANDO phantom, without and with fat was scanned on the CT scanner and reconstructed using the scan protocol and reconstruction parameters discussed in Section 4.2.4. This experimental set up is shown in Figure 4.6.

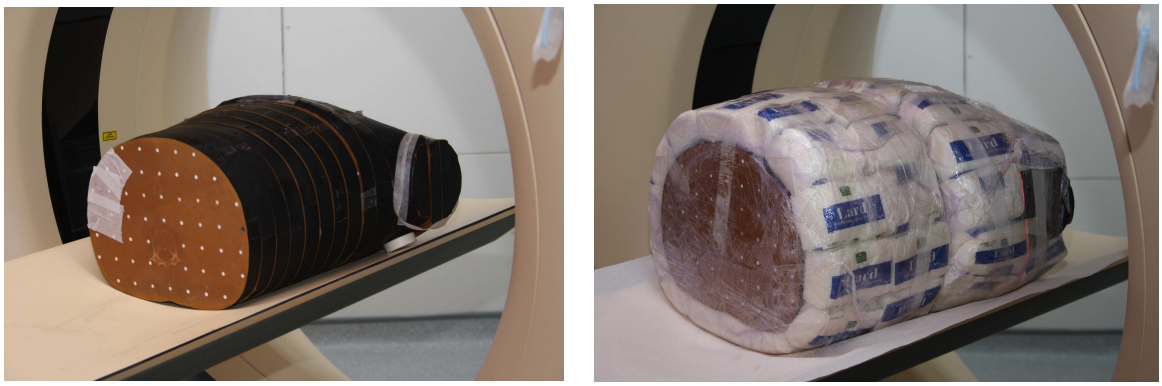

(a) (b)

Figure 4.6: Chest portion of RANDO on the CT scanner (a) without any added fat, (b) with fat added.

4.2.1.4 Clinical image acquisition

Further validation was performed by accessing clinical images (local research ethics committee approval was obtained) to allow the retrospective use of clinical CT and CR images. Image data of suitable patients (average and obese patients identified by the expertise of the examining Radiographer) were identified on the CR system, and beam quality, mAs and focus to detector distance used for their exposures were recorded. Images were transferred to a separate computer after Agfa specific post-processing (MUSICA) was removed.

Average sized patients only were identified on the CT scanner, to compute DRRs that simulated average patients. Fat was also added to the images artificially prior to DRR calculation to simulate obese patients (see Section 4.2.13). The CT data were reconstructed as per the parameters discussed in Section 4.2.4 prior to transfer to a separate computer.

4.2.2 Generating X-ray Spectra

To simulate diagnostic X-ray images using quantitative methods, it is important to ensure that the virtual source of X-rays produces spectra that accurately reproduce X-ray beam qualities used in practice. The X-ray spectra used in this research were generated using the techniques of Birch and Marshall, as described in the Institute of Physics and Engineering in Medicine (IPEM) Report 78 [90-91]. This software has been independently validated by Ay *et al.* [92].

The X-ray spectra can be generated at 0.5 keV intervals from 0.5 keV up to the X-ray tube accelerating potential chosen by the user. The spectral intensity is also determined by user specific variables, namely the target material (Tungsten (W) for this Philips X-ray system), the target angle (13°) and the amount of voltage ripple (assumed to be zero for high frequency generators). The spectral data is specified on a central axis 750 mm from the source. A typical continuous (Bremsstrahlung) spectrum with characteristic peaks as calculated using the methods detailed in IPEM 78 is shown in Figure 4.7.

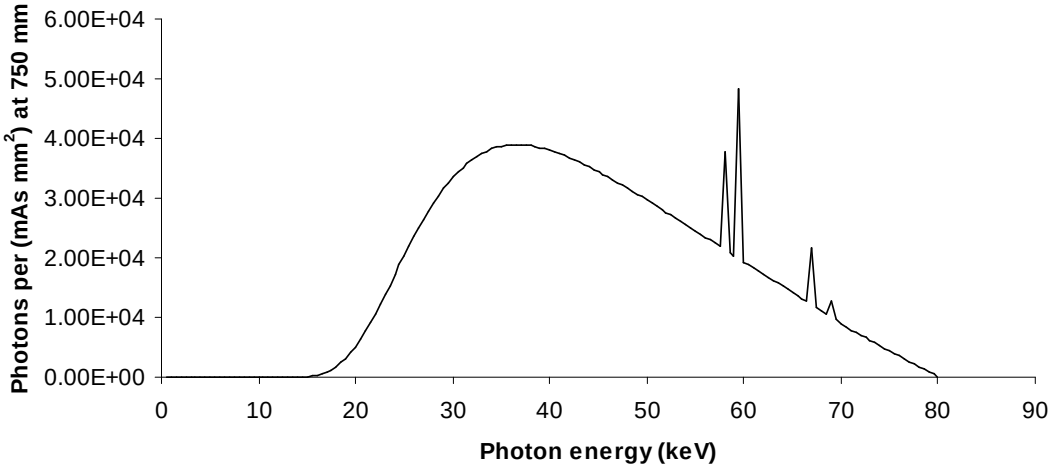

Figure 4.7: Spectrum processed by IPEM 78 for an 80 kVp beam, Tungsten target (angle of 13°) and total filtration of 3.1 mm Al. Intensity is given at the central axis 750 mm from the source.

As the beam originates from a point source, the intensity is inversely proportional to the square of the distance from the source (the inverse square law ISL), assuming insignificant attenuation by air. The intensity used in the simulation was corrected for distance from the source to the virtual patient, using the following equation:

$$I(p_v) = \left(\frac{750}{p_v}\right)^2 \times I_{750} \quad (4.1)$$

where $I(p_v)$ is the intensity at the surface of a given voxel on the virtual patient surface, p_v is the distance from the source to the voxel in millimeters, and I_{750} is the intensity at 750 mm calculated by IPEM 78.

The calculated X-ray spectra are assumed to be from tubes that are 100% efficient. However, this is not the case for real tubes and as such produce spectra with outputs lower than that calculated. The intensity used in this model was therefore corrected by the efficiency of the real X-ray tube used in this study by measuring the air kerma with a calibrated 6 cc ionisation chamber (Radcal Corporation, Monrovia, USA) at 750 mm from the source and comparing this with the air kerma calculated by IPEM Report 78. The tube was found to be approximately 86% efficient. X-ray intensity at the virtual patient surface can then be determined using equation 4.2:

$$I(p_v) = \left(\frac{750}{p_v}\right)^2 \times I_{750} \times 0.86 \quad (4.2)$$

Spectra for X-ray tube voltages across the diagnostic energy range, i.e. 50 to 150 kVp in approximate steps of 10 kVp were calculated (exact steps of 10 kVp were not available on the clinical X-ray system to which the software was

configured and so were not used in the simulation – see Chapter 5, Table 5.1 for the actual kVps used). Central axis intensity (750 mm from the virtual source) at each tube potential, generated with parameters correct for the Philips X-ray system (13° Tungsten anode with zero ripple and 3.1 mm Al total filtration) was converted into a Microsoft Excel file in a format ready for use with the DRR software. The raw Excel data was then accessed by the DRR software and corrected for inverse square law and tube efficiency (in accordance with equation 4.2). Figure 4.8 illustrates the effect inverse square correction has on the X-ray beam intensity at the virtual patient surface, from the 'beams eye view' (BEV) perspective.

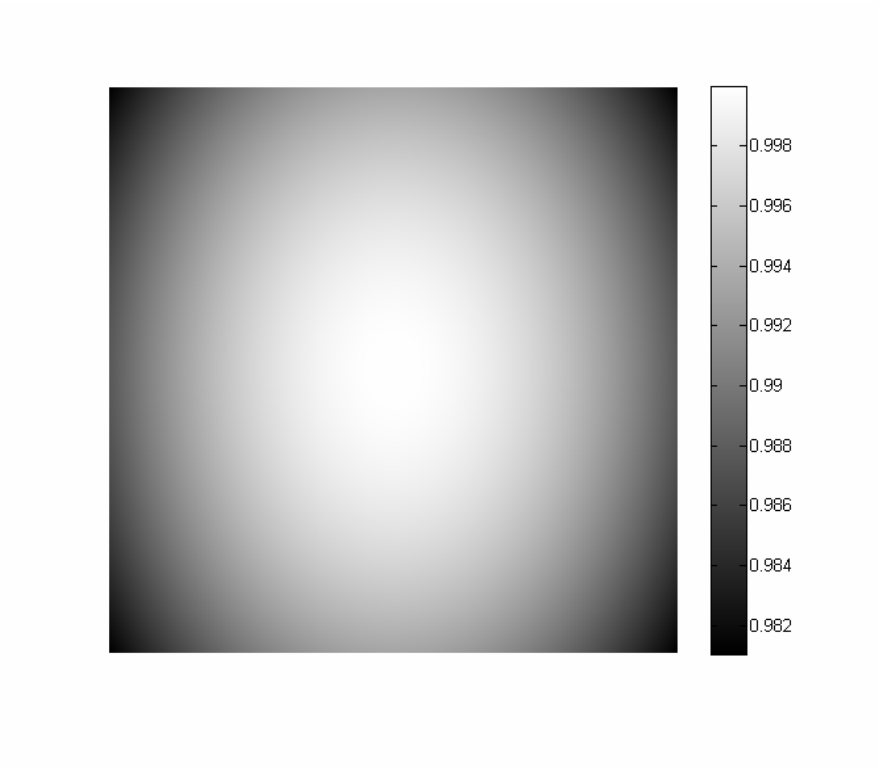

Figure 4.8: Beams eye view of X-ray spectra at the virtual patient surface with inverse square correction. The scale is normalised to 1.

As can be seen from Figure 4.8, the intensity impinging the virtual patient decreases by approximately 2% in the corners. This is expected, as X-rays have to travel further to the corners relative to the centre, assuming the centre

of the beam is normal to the virtual source. It should be noted that the influence of the anode heel effect is not included.

4.2.3 CT Data Preparation

For DRR calculation, it is important that CT images have as little processing applied as possible, so that CT voxel values correspond to the X-ray attenuation properties of the particular tissue. However, as back-projected CT data is inherently blurred, filtering must be applied. The Philips scanner used in this study has a variety of filters, some of which artificially sharpen the image (emphasise high spatial frequencies), and others that smooth the image (emphasise low frequencies). Typical functions associated with filtered back-projection are shown in Figure 4.9.

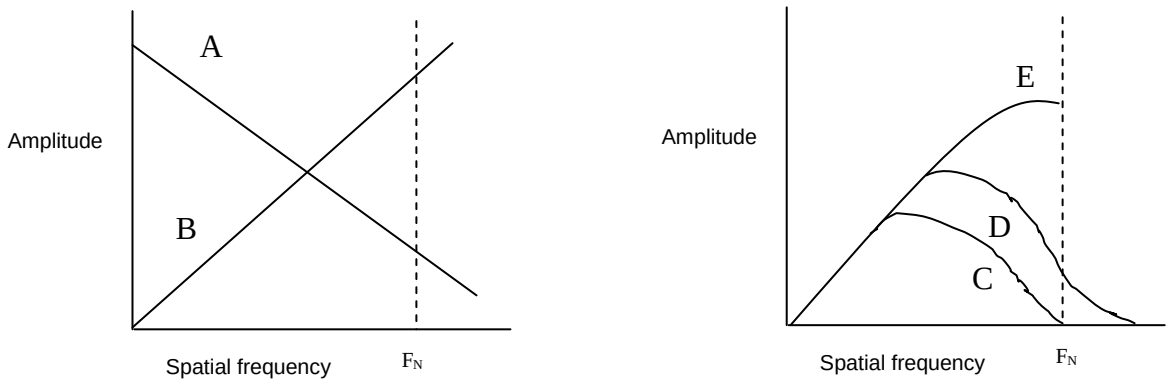

Figure 4.9: (a) Effect filtered back-projection has on amplitudes of spatial frequencies in an image, and (b), those filters used practically. F_N is the Nyquist frequency.

Figure 4.9(a) shows the effect back-projection has on the amplitudes of spatial frequencies present in a CT image (curve A). Curve B shows the theoretical ideal correction function (but this still amplifies high frequency noise and aliasing). Figure 4.9(b) shows typical filters used practically. Filter C will cause

loss of resolution due to the drop off at high frequencies, whilst filter D will give better resolution than filter C but information will still be lost due to high frequency cut off. Filter E gives best resolution but appreciable high frequency noise, and can be described as a ‘ramp’ filter [93, 94].

After discussion with Philips (Private Communication, Paul Klahr, Philips Medical Systems, 2008), filter E (available on our CT scanner) was identified as providing minimum processing, and as discussed above, this is the basic ramp filter that corrects the effect back-projection has on the amplitudes of different spatial frequencies in the image. On the scanner, CT data can be reconstructed with filter E by reprocessing the raw non-filtered data. The field of view (FOV) and reconstruction matrix can also be altered, although the FOV is somewhat dependent on patient size.

CT images also contain noise which may influence final characteristics of a DRR image. However, the impact of this noise can only be investigated once DRR images can be produced. The potential issues around this are discussed in Section 4.2.16.

4.2.4 CT Number to linear attenuation coefficient conversion

The ray casting DRR method requires each CT voxel value to be converted to its linear attenuation coefficient (LAC) in order that one may calculate the exiting photon intensity from a knowledge of the incident photon intensity (see Section 4.2.6). It should be remembered here that CT voxel values are not in Hounsfield Units (HU), but ‘raw values’. This is due to the fact that images transferred from the CT scanner did not have the manufacturer specific conversion function (contained within the DICOM header of each image) applied that converts raw voxel values to HUs (i.e. this processing step was ‘bypassed’). This poses a

potential problem in that these raw voxel values may ‘drift’ because manufacturer calibrations (performed annually) are carried out in terms of HUs; a modification to the conversion function is made by the service engineer if the HU values are out of the accepted range. However, upon comparing the conversion function in the DICOM header of images acquired at the beginning of the study to that at the end (spanning approximately three years), there was no change in conversion function. This demonstrates the service engineer, through the course of this study, never had to recalibrate the HU values, suggesting the raw voxel values did not ‘drift’ significantly.

The Gammex RMI tissue equivalent phantom was used to convert CT voxel values to LACs. The phantom (Gammex-RMI, Broadway Business Centre, Nottingham, UK) is a tissue equivalent phantom (model no. 467), consisting of a Solid Water® cylinder that contains 17 inserts, the attenuation properties of which mimic the different attenuation properties of the various tissues found in vivo, as shown in Figure 4.10. The inserts can be placed anywhere in the phantom.

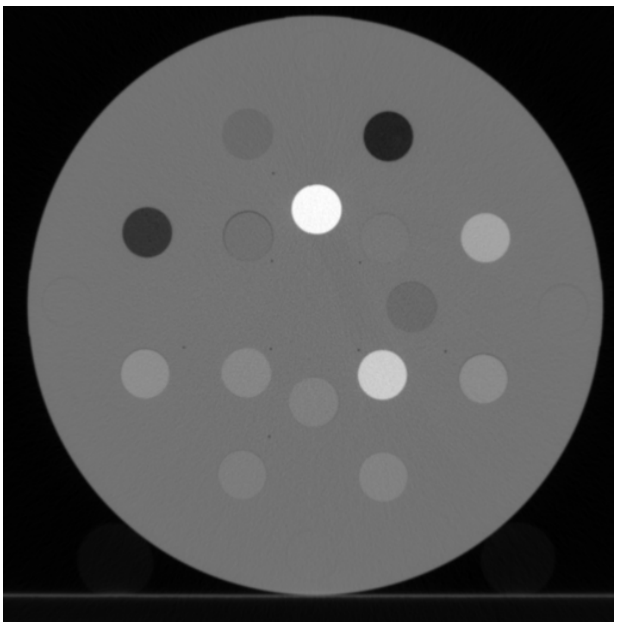

Figure 4.10: CT image of Gammex RMI tissue substitute phantom with each insert visible.

The phantom was scanned with the scanning protocol (see Table 4.1) in use for chest examinations at our institute.

CT Scan Parameter	Value Used
Resolution	Standard
Collimation	16 x 0.75 mm
Pitch	1
FOV (mm)	350
Tube Potential (kVp)	120
Tube Current-time product (mAs)	175

Table 4.1: Parameters used to scan Gammex RMI phantom. These are standard for chest imaging in our Radiology Department.

Each tissue insert was placed at the centre of the phantom and scanned in turn. All other tissue inserts were removed to minimise the effects of scatter. After the scan, raw data were reconstructed with filter E, matrix 1024 and slice thickness 0.8 mm (the same reconstruction parameters to be used to prepare

clinical images for DRR calculation). The mean of the CT voxel values contained within a region of interest (ROI) of each tissue substitute was measured, as shown in Figure 4.11

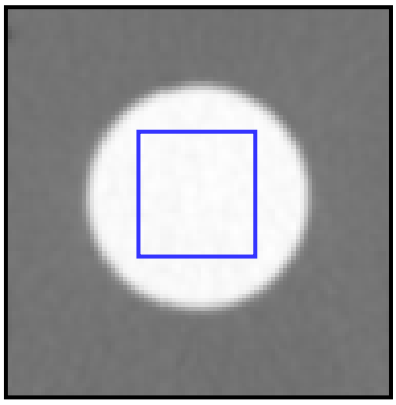

Figure 4.11: Image of an insert against the uniform background. The ROI is shown in blue.

It has been assumed that beam hardening has minimal impact on the results, as the CT scanner utilizes a 'bow-tie' filter to conform the intensity of the beam to compensate for body shape (the RMI phantom has a circular diameter similar in size and shape to the phantom used by the service engineer when testing this artefact), and scanner software corrects for any further artefact. However, to test this assumption, the conversion function (as discussed in the previous section) was used to convert the raw voxel value of solid water to its HU counterpart, and compare this to the service engineer's allowed range of HUs for water. The conversion function has a gradient = 1, and an intercept = -1024, i.e. the conversion function simply takes 1024 voxel value units from the raw voxel values to calculate the corresponding HUs. If one applies this conversion to solid water, which has a raw value of 1016 (see Table 4.3), one derives a HU values for water of -8. This is within the manufacturer tolerance of 0 ± 10, suggesting that beam hardening is indeed minimal on this scanner, and the raw voxel values are 'correct'.

Information from the user manual of the RMI phantom provides the elemental composition of each tissue substitute by weight (e.g. composition of lung by percentage weight is: H=8.33, C=60.32, N=1.67, O=17.38, Cl=0.15, Si=0.61 and Mg=11.54). This data was entered into the XCOM photon cross sections database (developed by the National Institute of Standards and Technology – NIST) [64, 65] together with the X-ray spectrum produced by IPEM Report 78. The XCOM database calculates photon cross sections for scattering, photoelectric absorption and pair production, as well as total attenuation coefficient for any element, compound or mixture at energies from 1 keV to 100 GeV. The database can generate attenuations for user defined photon energies, such as those generated by IPEM Report 78.

The *total* mass attenuation coefficient (cm^2/g) of each tissue substitute, for all photon energies used to generate a given DRR image were derived with the XCOM database. As the density (g/cm^3) of each substitute is readily available, it was simple to convert from total mass attenuations to LAC (cm^{-1}). To illustrate this point, Table 4.2 shows the first five energies of a 50 kVp spectrum together with mass attenuation coefficient and LACs for the lung (LN300) substitute (it should be noted that the XCOM database gives a maximum 10% error on the attenuation values).

Photon Energy (keV)	Total Mass Attenuation Coefficient (cm^2/g) of LN300	Linear Attenuation Coefficient (cm^{-1}) of LN300
13.0	2.49	0.747
13.5	2.24	0.672
14.0	2.03	0.609
14.5	1.84	0.552
15.0	1.68	0.504

Table 4.2: First five photon energies in a virtual 50 kVp spectrum together with their respective total mass attenuation and linear attenuation coefficient for lung substitute (LN300).

This process was repeated for all tissue substitutes with all virtual tube potentials used in this study (50 kVp to 150 kVp in approximate steps of ten). Table 4.3 shows the first three energies of the 50 kVp spectrum with LACs for each substitute.

Tissue Substitute	Mean CT Voxel Value	13.0 keV LAC (cm^{-1})	13.5 keV LAC (cm^{-1})	14.0 keV LAC (cm^{-1})
Air	21	0.00029	0.00026	0.00024
LN300	282	0.07470	0.06720	0.06090
LN450	443	0.11205	0.10080	0.09135
AP6, adipose	909	0.13708	0.12420	0.11316
Polyethylene	911	0.09568	0.08749	0.08022
Breast	970	0.18216	0.16434	0.14949
Solid Water	1016	0.15300	0.13872	0.12648
CB3 resin	1027	0.23094	0.20796	0.18810
Brain	1029	0.15936	0.14413	0.13094
Liver	1113	0.17135	0.15410	0.14030
IB1 inner bone	1114	0.29268	0.26244	0.23760
B200 bone mineral	1143	0.18880	0.17110	0.15576
CB4 resin	1228	0.37440	0.33813	0.30537
CB2-10%	1301	0.55440	0.49840	0.45024
Acrylic	1329	1.06485	0.95837	0.86448
CB2-30%	1571	0.90852	0.81740	0.73834
CB2-50%	2034	1.62240	1.45704	1.31508
SB3 cortical bone	2516	2.48400	2.24480	2.02400

Table 4.3: First three energies of a virtual 50 kVp spectrum together with the LACs and mean voxel value of each tissue substitute. Note that the voxel values are not in terms of Hounsfield Units.

The data from Table 4.3 were then used to derive linear relationships between mean CT voxel values and LACs for each photon energy (Table 4.3 would in reality contain all energies of the 50 kVp spectrum, but for conciseness these are not shown here).

In Table 4.3, for each photon energy, there is a linear relationship (R^2 = 0.9972) between CT voxel value and LACs from air (CT = 21) to lung substitute LN450 (CT = 443). The resulting linear equation was used to convert those voxels in the CT data that range from 21 to 443 (all CT voxel values lower than 21 were converted to 21) to their respective LACs. Moving down the table the next linear relationship that exists with a high degree of correlation (R^2 = 1) is LN450 (443) to adipose (909). As such, all voxel values that lie in the range 443 to 909 were converted to their relevant LACs depending on the linear relationship. Two resulting linear equations of two photon energies are shown graphically in Figure 4.12.

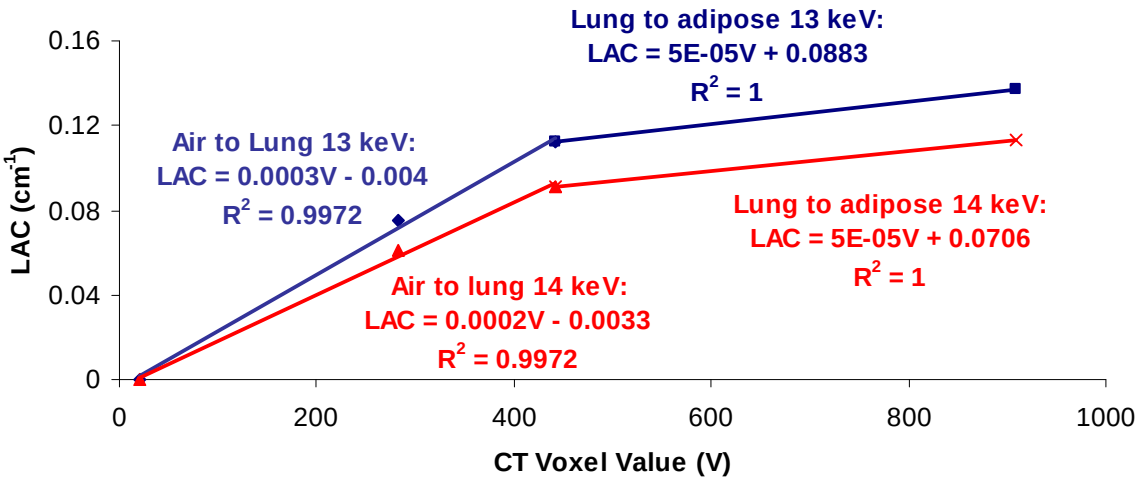

Figure 4.12: Linear equations derived to convert CT voxel values into LACs. The blue plot is for photon energy 13 keV, the red relationship is for photon energy 14 keV. For clarity, other linear relationships are not shown.

The plots in Figure 4.12 clearly show differences in LACs for different photon energies. This process was continued until all CT voxel values were transformed to their respective LAC. *This entailed the derivation of 15 linear equations for each photon energy*. Very few CT voxel values existed above 2516 (cortical bone), but those that did were converted using the linear

relationship for CB2-30% to cortical bone (i.e. the final linear relationship derived from Table 4.3).

It was not possible to derive a single function to convert CT voxel values to LACs because it was not possible to fit a function to the whole data with any degree of accuracy, as shown by the example in Figure 4.13.

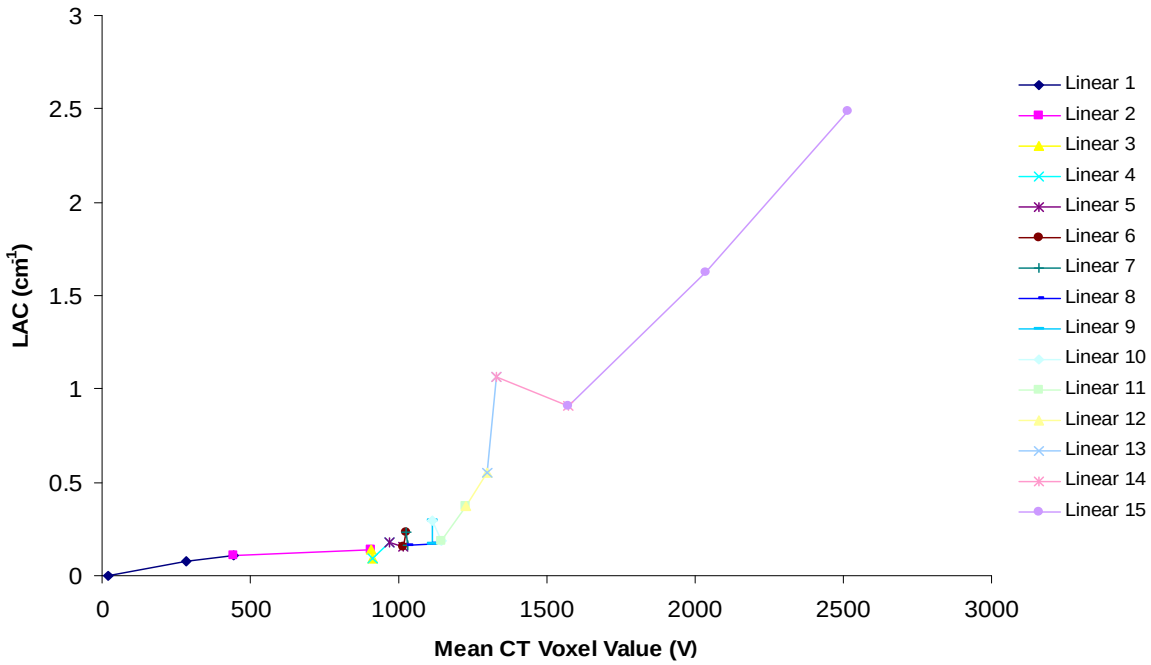

Figure 4.13: All 15 linear equations derived to convert CT voxel values into LACs for photon energy 13 keV.

Figure 4.13 is an extension of Figure 4.12 in that *all* 15 linear relationships are shown. It is clear that one continuous (accurate) function would be very difficult to derive, especially in the voxel value region 900 - 1600. The method described here also has the advantage of using all of the RMI phantom data (i.e. lots of different tissue surrogates), which is a vast improvement over three to four tissue types used in the literature (see Chapter 1).

An example CT slice with voxel values and corresponding LACs correct for a monoenergetic 13 keV beam is shown in Figure 4.14.

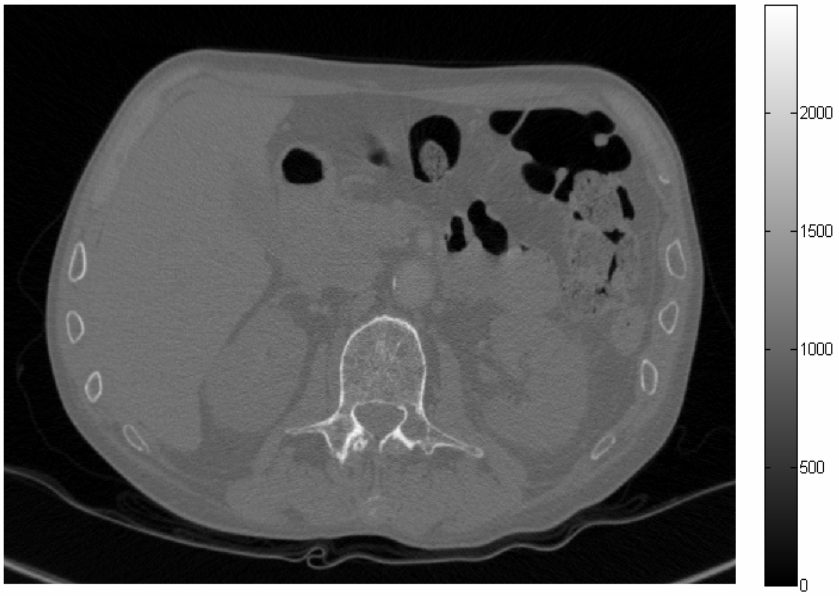

(a)

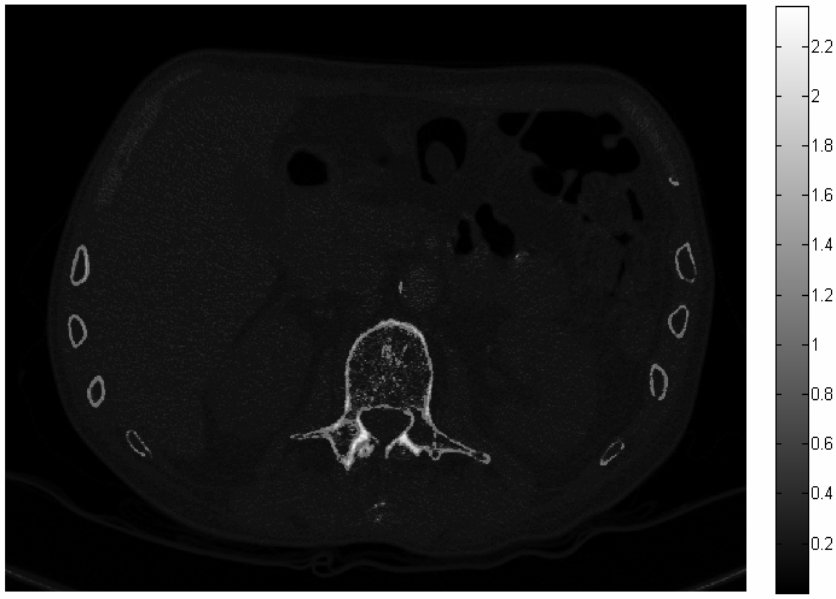

(b)

Figure 4.14: CT image slice with (a) raw pixel values, and (b) the same slice with pixel values converted into LACs (mm^{-1}), i.e. a LAC map.

As can be seen in Figure 4.14, the scale on the right hand side of each image clearly shows the difference between CT voxel values and LACs. Only bone is visible in the LAC map (Figure 4.14(b)) as this has a greater LAC than soft tissue, and therefore stands out in the LAC image.

4.2.5 Photon energy bins – a method to speed up DRR simulation

During image reconstruction, each CT slice was not transformed to a linear attenuation coefficient map (described in Section 4.2.4) for all photon energies present in each virtual X-ray spectrum. To do so would result in impossibly long computation times. As such, 'reference photon energies' in each X-ray spectrum were chosen, the DRR calculation executed for these, and the number of photons emerging from the CT data derived. However, the number of photons of *other* energies emerging from the CT data in each pencil beam (see Section 4.2.6) was estimated purely from one knowing the intensity of photons emerging from the CT data of the reference photon energies. Figure 4.15 and equations 4.3 to 4.8 illustrate the argument.

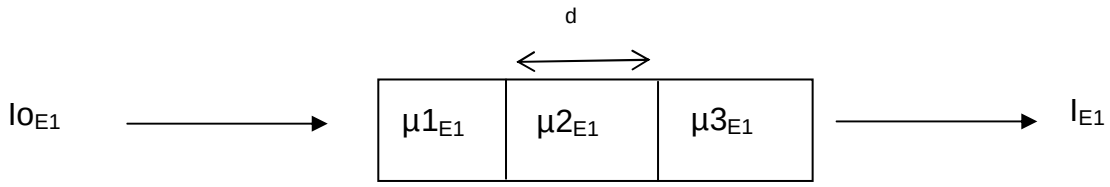

Figure 4.15: Monoenergetic photons of energy E1 incident on voxels with attenuation coefficients μN_{E1}. Each voxel is length d. I_{E1} is the exiting intensity.

Figure 4.15 shows a monoenergetic intensity of photons (energy E1), I_{oE1}, incident on a voxelated structure. The intensity of exiting photons would be:

$$I_{E1} = Io_{E1} \exp[-d \times (\mu 1_{E1} + \mu 2_{E1} + \mu 3_{E1})] \quad (4.3)$$

where d is the length of each voxel. It follows that:

$$I_{E1} = Io_{E1} \exp[-d \times (\mu T_{E1})] \quad (4.4)$$

where μT_{E1} is the total LAC (i.e. sum of LACs) for energy E1. This is also true for incident monoenergetic photon energy E2:

$$\mathrm{I}_{\mathrm{E2}} = \mathrm{Io}_{\mathrm{E2}} \exp\left[-d \times \left(\mu \mathrm{T}_{\mathrm{E2}}\right)\right] \tag{4.5}$$

If the ratio of the two is taken, the following equation is produced:

$$\mathrm{I}_{\mathrm{E2}} = \mathrm{I}_{\mathrm{E1}} \left(\frac{\mathrm{Io}_{\mathrm{E2}}}{\mathrm{Io}_{\mathrm{E1}}} \right) \exp\left[d\left(\mu \mathrm{T}_{\mathrm{E1}} - \mu \mathrm{T}_{\mathrm{E2}}\right)\right] \tag{4.6}$$

So if one knows the total LAC, the incident and exiting intensities of photons for a given energy (e.g. E1), it is relatively simple to calculate the exiting intensity for photons of a different energy assuming the total LAC and incident intensity is known for that energy. This relationship was used in this work. If there is a constant ratio between the total LACs, such that:

$$\frac{\mu \mathrm{T}_{\mathrm{E2}}}{\mu \mathrm{T}_{\mathrm{E1}}} = R \tag{4.7}$$

then equation 4.6 becomes:

$$\mathrm{I}_{\mathrm{E2}} = \mathrm{I}_{\mathrm{E1}} \left(\frac{\mathrm{Io}_{\mathrm{E2}}}{\mathrm{Io}_{\mathrm{E1}}} \right) \exp\left[d\left(\mu \mathrm{T}_{\mathrm{E1}} - R\mu \mathrm{T}_{\mathrm{E1}}\right)\right] \tag{4.8}$$

As discussed in Section 4.2.4, LACs were derived using the XCOM database for each of the tissues present in the Gammex RMI phantom, and for all photon energies in a given X-ray tube potential. For 150 kVp, this entailed 275 energies, as the lowest photon energy with more than one photon present in the simulated 150 kVp spectrum was 13 keV (energies with less than one photon were deleted from the spectrum). One can define a tissue LAC Ratio (TLR – equation 4.7) as the ratio of the (Gammex phantom) tissue LAC of one energy to that of another. Table 4.4 shows TLRs of 13.5 and 20.0 keV photon energies for each of the Gammex tissues compared to that of 13.0 keV.

RMI Tissue	LAC for 13 keV (mm^{-1})	LAC for 13.5 keV (mm^{-1})	Tissue LAC Ratio (TLR) 13.5/13.0	LAC for 20.0 keV (mm^{-1})	Tissue LAC Ratio (TLR) 20.0/13.0
Air	0.00029	0.00026	0.897	0.0001	0.345
LN-300	0.0747	0.0672	0.900	0.0245	0.328
LN-450	0.1121	0.1008	0.899	0.0367	0.327
AP6	0.1371	0.1242	0.906	0.0499	0.364
Poly	0.0957	0.0875	0.914	0.0397	0.415
Breast	0.1822	0.1643	0.902	0.0639	0.351
CB3 resin mix	0.1530	0.1387	0.907	0.0564	0.369
Brain	0.2309	0.2080	0.901	0.0767	0.332
Solid water	0.1593	0.1441	0.905	0.0574	0.360
CB4 resin mix	0.1713	0.1541	0.900	0.0622	0.363
Liver	0.2927	0.2624	0.896	0.0943	0.322
Acrylic	0.1888	0.1711	0.906	0.0675	0.358
CB2 - 10% Bone	0.3744	0.3381	0.903	0.1217	0.325
IB1	0.5544	0.4984	0.899	0.1702	0.307
B200	1.0649	0.9584	0.900	0.3206	0.301
CB2-30% Bone	0.9085	0.8174	0.900	0.2774	0.305
CP2-50% Bone	1.6224	1.4570	0.898	0.4852	0.299
SB3 - Cortical Bone	2.4840	2.2448	0.904	0.7379	0.297
			Mean ratio (MR) = 0.902 St Dev = 0.004 CoV= 0.471%		Mean ratio (MR) = 0.337 St Dev = 0.031 CoV= 9.361%

Table 4.4: Photon Energy Dependent LACs of the RMI elements (13, 13.5 and 20 keV). Tissue LAC Ratios (TLR) 13.5/13.0 and 20.0/13.0 are given as well as the mean ratio (MR), standard deviation of the ratios and coefficient of variance (%) for each. The columns of ratios are highlighted in blue.

As can be seen in Table 4.4, TLRs 13.5/13.0 are very similar to one another, but those for 20.0/13.0 are not, which of course is expected. For equation 4.8 to

work, it is important TLRs are not too different from one another, so as a first approximation, only photon energies with a coefficient of variance (CoV) ≤ 10% compared with the ‘reference photon energy’ were ‘binned’ together. For example, if photon energy 13 keV acts as a reference photon energy, energies 13.5 to 20.0 keV may be placed in the same bin. When the DRR is being generated, CT voxel values are converted into LACs correct for 13 keV, the incident intensity of 13 keV photons is known (Section 4.2.2), and the exiting photon intensity is calculated. Rather than doing this again for energies 13.5 to 20.0 keV (in total 14 separate energies and thus 14 separate DRR calculations), equation 4.8 is used by replacing R with the mean tissue LAC ratio (MR in Table 4.4) of the particular photon energy. This allows the calculation of exiting photon intensities 13.5 to 20.0 keV at the same time as the reference energy 13.0 keV, without having to run the DRR simulation another 14 times. In other words, energies 13.5 to 20.0 keV were binned with the reference energy. As the CoV of photon energy 20.5 keV was 10.176% (larger than the 10% limit), this energy was deemed the reference energy for the second bin, and the process was repeated until enough energy bins were derived for the particular X-ray spectrum. For example, to produce a DRR simulation of a 150 kVp image, the following reference photon energies were used:

1. 13.0 keV
2. 20.5 keV
3. 26.0 keV
4. 31.5 keV
5. 38.0 keV
6. 46.0 keV
7. 59.0 keV
8. 88.5 keV
9. 119.0 keV

This methodology necessitates running the simulation 9 times, rather than 275 as would have been the case without photon energy binning.

4.2.6 Ray casting with pencil beams

As discussed in Chapter 2, rather than using simple rays for DRR generation, this work will use realistic pencil beams emanating from the source to the surface of the virtual patient. A basic example is depicted in Figure 4.16.

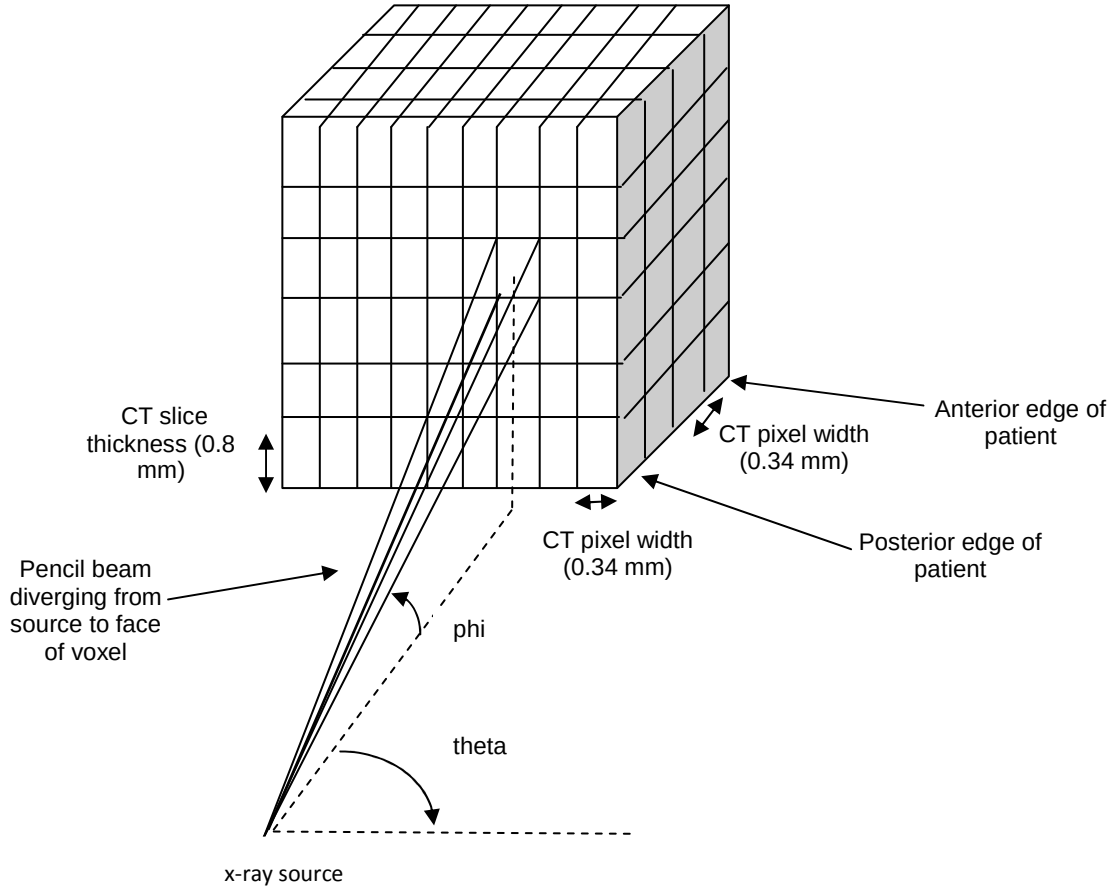

Figure 4.16: Simple depiction of a pencil beam emanating from the source impinging a voxel of the virtual patient. This example shows only one pencil beam, but in reality the software projects pencil beams to each voxel at the patient surface plane.

The virtual patient here was produced with clinical CT slices accessed in a contiguous fashion. However, due to lack of computer memory (MATLAB^{TM} generates an 'out of memory' message whenever it requests a segment of memory from the operating system that is larger than what is currently available; the amount of memory therefore available is dependent on the specification of the computer being used), it was not possible to read in all clinical CT slices at the same time and project pencil beams through the entire dataset, so they were read in batches of twenty. Also, the clinical CT data was only available in an axial orientation, as shown in Figure 4.17.

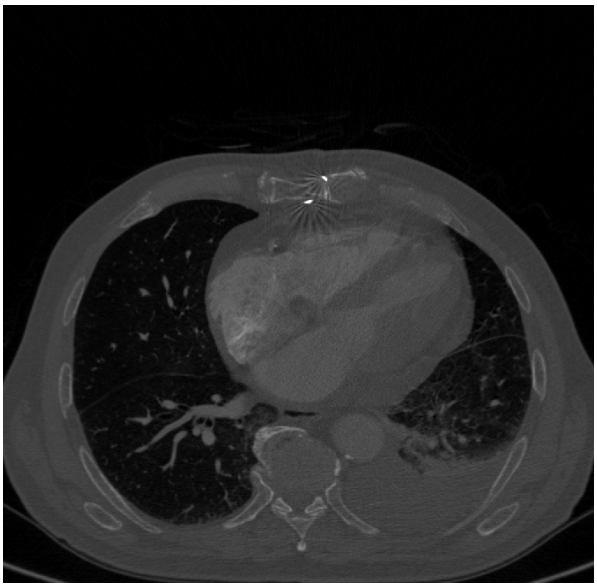

Figure 4.17: Axial CT slice of the human thorax.

Clinical CT data in this axial configuration is not in the orientation required for ray casting, as rays must project through the data in a posterior-anterior (PA) or AP direction (in the configuration shown in Figure 4.17, rays would be projected through the patient axis, superior to inferior (SI), rather than PA or AP). As such, axial CT data orientation was changed to ensure each patient data set was in the correct configuration for pencil beam projection (i.e. in a PA rather than a SI orientation). Also, to avoid even more computer memory problems, some of the coronal slices (i.e. in the PA direction) of the newly orientated CT data were deleted prior to pencil beam projection. The process is described in the flow chart below and is also shown diagrammatically in Figure 4.18:

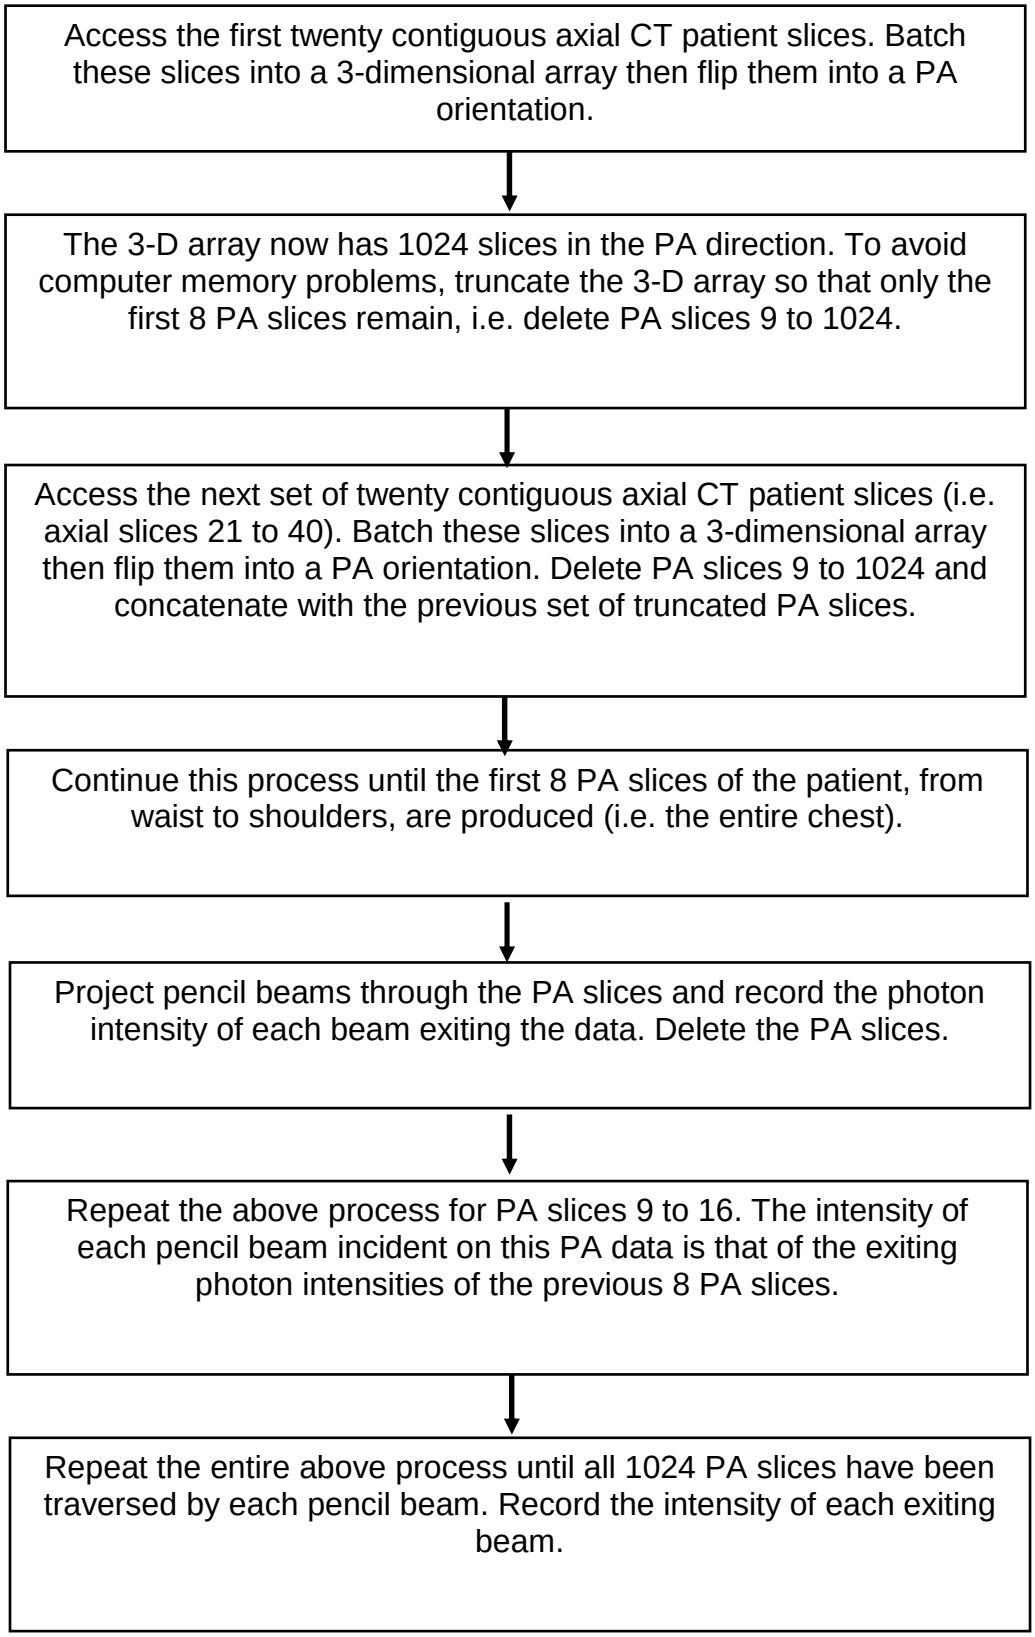
Access the first twenty contiguous axial CT patient slices. Batch these slices into a 3-dimensional array then flip them into a PA orientation.
The 3-D array now has 1024 slices in the PA direction. To avoid computer memory problems, truncate the 3-D array so that only the first 8 PA slices remain, i.e. delete PA slices 9 to 1024.
Access the next set of twenty contiguous axial CT patient slices (i.e. axial slices 21 to 40). Batch these slices into a 3-dimensional array then flip them into a PA orientation. Delete PA slices 9 to 1024 and concatenate with the previous set of truncated PA slices.
Continue this process until the first 8 PA slices of the patient, from waist to shoulders, are produced (i.e. the entire chest).
Project pencil beams through the PA slices and record the photon intensity of each beam exiting the data. Delete the PA slices.
Repeat the above process for PA slices 9 to 16. The intensity of each pencil beam incident on this PA data is that of the exiting photon intensities of the previous 8 PA slices.
Repeat the entire above process until all 1024 PA slices have been traversed by each pencil beam. Record the intensity of each exiting beam.

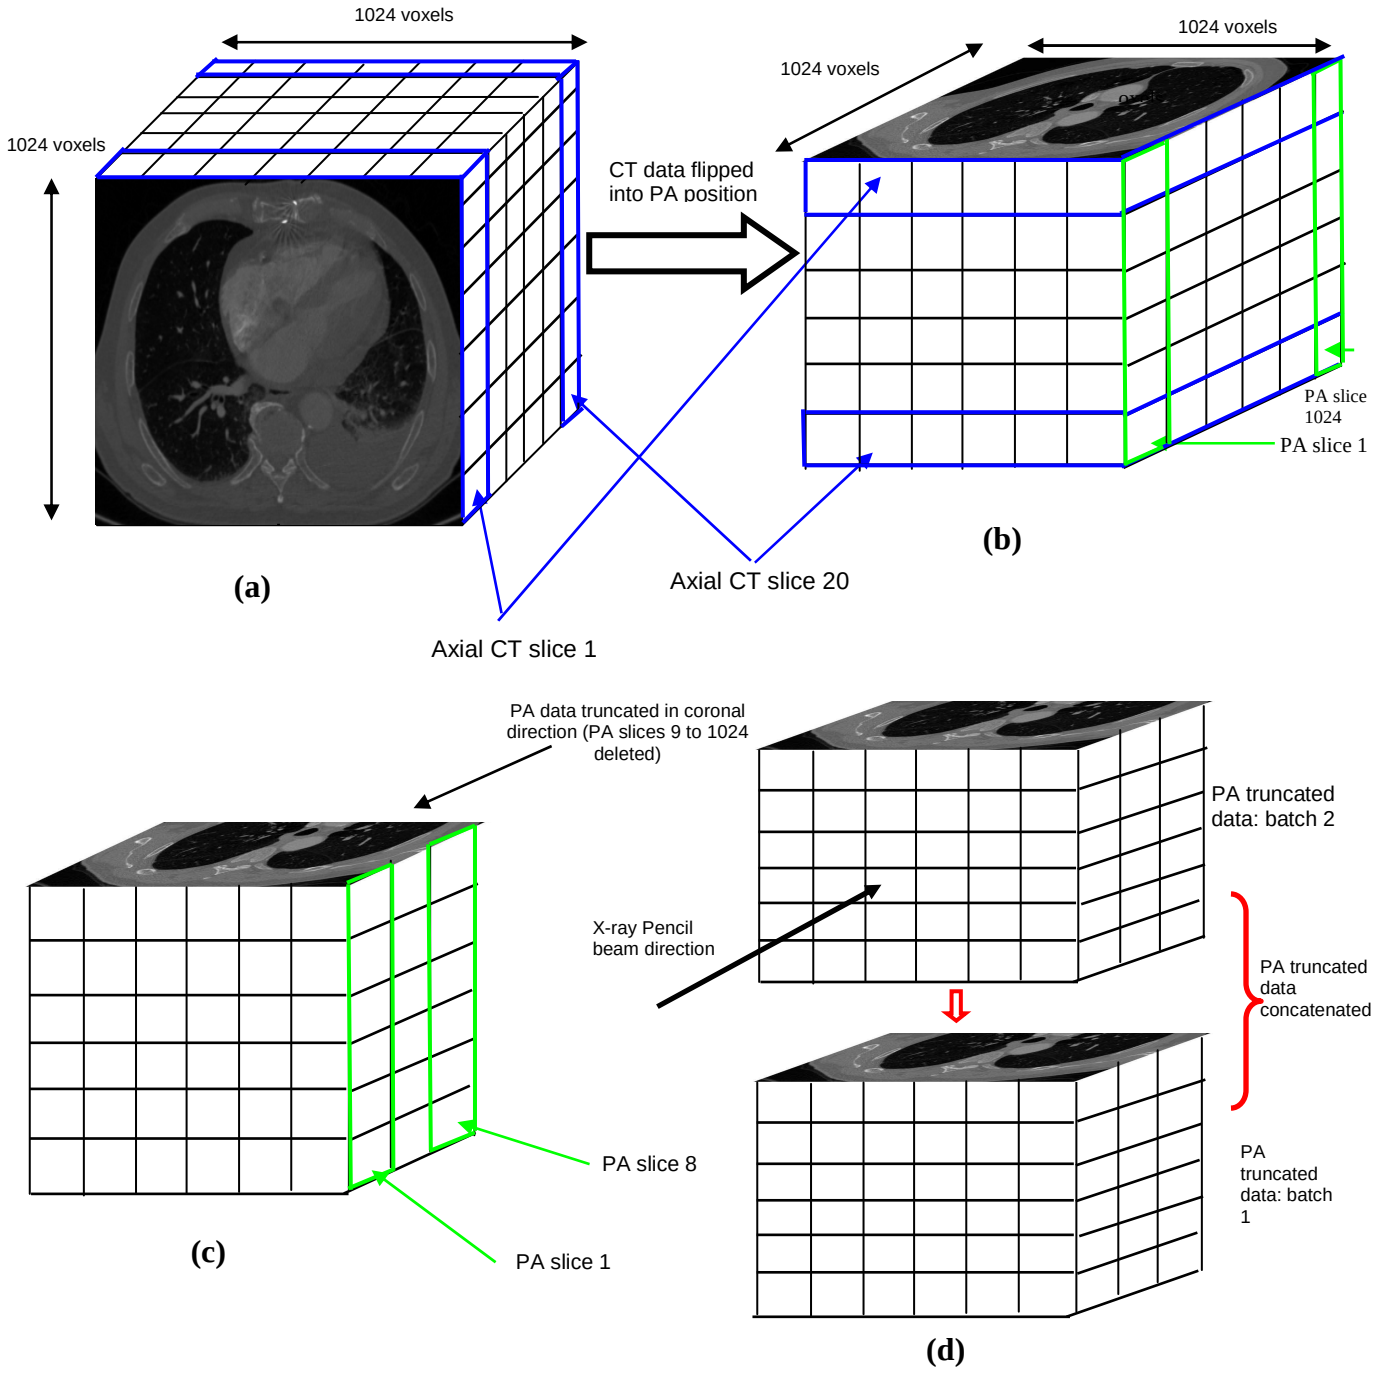

Figure 4.18: Reading in batched CT slices and truncating to minimise memory implications. (a) 3-D array of axial CT data, (b) 3-D array of axial CT data flipped into a PA orientation (there are 1024 PA slices as the axial CT data is 1024 by 1024 voxels in size), (c) PA CT data truncated in coronal direction, (d) second batched CT data flipped in PA orientation concatenated to the first batch. X-ray pencil beam direction is shown by an arrow entering the CT data.

The process shown in Figure 4.18 was repeated from the first *axial* CT batch to the final *axial* batch. That is, if there were 700 axial CT slices in the patient scan, CT batch 1 would include slices 1 to 20, CT batch 2 would include slices

21 to 40 and so on until the last batch which would contain slices 681 to 700. When flipped into PA orientation, the resulting concatenated and truncated array allows pencil beams to be projected through the CT data in a PA direction (i.e. normal to the coronal direction) without any restrictions on memory. Figure 4.19 shows the first 8 PA slices through the patient, after the process described in the above flow chart and Figure 4.18 have been performed.

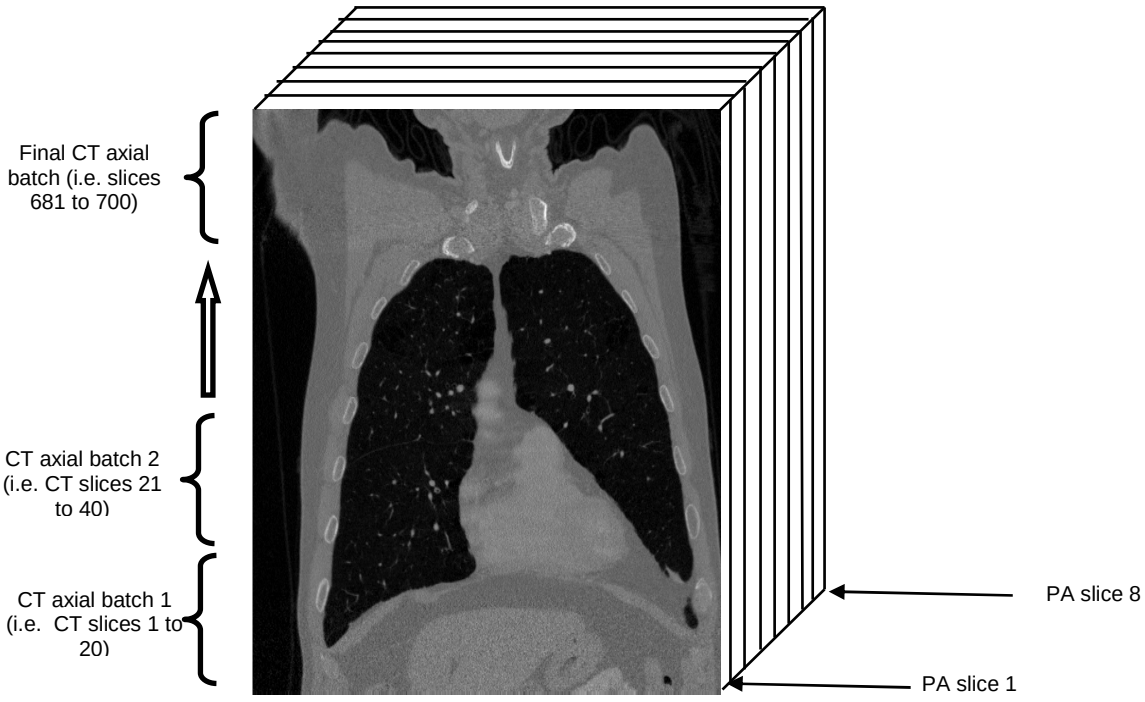

Figure 4.19: First eight PA slices through the virtual patient. Axial CT slices have been batched, flipped PA, truncated and concatenated.

Figure 4.19 shows the first 8 PA slices of a patient in the correct orientation for ray casting. If one takes a normal from the centre of the first PA slice, this would intersect the X-ray source. Using this 8 PA slice configuration, pencil beams can be cast through the data (specific description of how this is done is described below) and the intensity of photons per pencil beam exiting the data can be saved to disc, and the PA CT data deleted. The next set of 8 PA slices may then be accessed (derived using the process described in the flow chart and Figure 4.18 for PA slices 9 to 16), and the process repeated using the photon

intensities per pencil beam saved previously as the input to the new PA data. This process was repeated through the entire PA axis of the patient.

Prior to the above process, pencil beams were projected from the X-ray source (shown in Figure 4.16) to the first PA slice. Enough pencil beams were projected such that one was incident on each voxel of the first PA slice of the patient (i.e. the whole X-ray beam was split into smaller pencil beams), as shown in Figure 4.20.

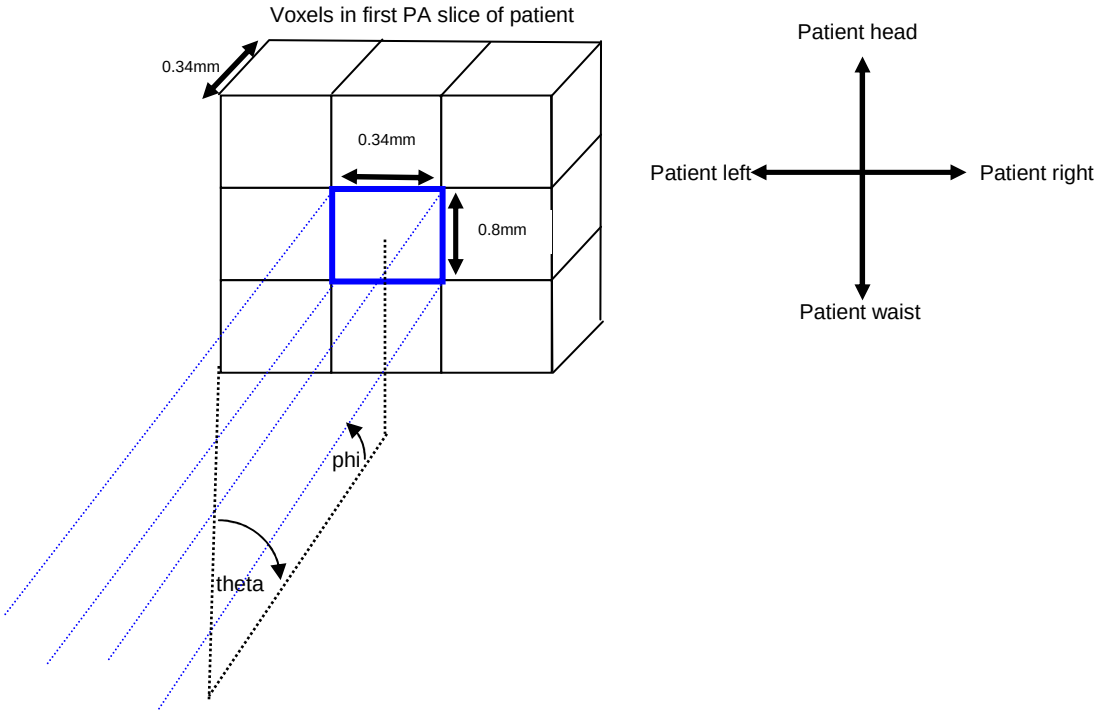

Figure 4.20: Pencil beam impinging a given voxel (dimensions shown) in first PA slice of patient. Pencil beams also impinge neighbouring voxels but for clarity these are not shown here. Pencil beam entrance angles are also shown: phi is the angle which the pencil beam has travelled in the up/down direction from the source, theta is the angle traversed in the left/right direction. Directions are shown for clarity.

Each CT voxel is 0.8 mm x 0.34 mm x 0.34 mm in size (height x width x depth), as each voxel height, width and depth is the same dimension as the CT slice thickness, and CT pixel resolution in x and y respectively (due to flipping the

data). Each pencil beam was set to the same dimensions as a flipped voxel with no increase in size as it projects through the data. In reality, each pencil beam would increase in size ('fan out'), but by very little; for a virtual patient who is 24 cm thick, each pencil beam would increase in size by approximately 7% in each dimension as it moves through the CT dataset. Therefore no correction was applied for this. The projection function uses a 1:1 mapping between pixels on the DRR and voxels on the front face (first PA slice) of the dataset, resulting in a simulated image with a typical pixel density 700 rows x 1024 columns (i.e. no. of CT slices x CT resolution), and resolution of 0.8 mm x 0.34 mm. This results in poorer resolution than a CR image (2800 x 2300, pixel pitch 0.17 mm x 0.17 mm), but is discussed in Section 4.2.14.

The intensity per photon energy of X-ray photons in each pencil beam impinging the face of each voxel in the first PA slice was calculated, as well as the entrance angles of elevation (phi) and azimuth (theta). CT voxel values were converted to the LAC correct for the reference photon energy. It was assumed the central axis of each pencil beam impinges the centre of each voxel in the first PA slice (i.e. no portion of the beam enters neighbouring voxels). However, this is not the case for subsequent PA slices; for example, Figure 4.21 shows the front face of nine voxels together with a pencil beam impinging the central voxel.

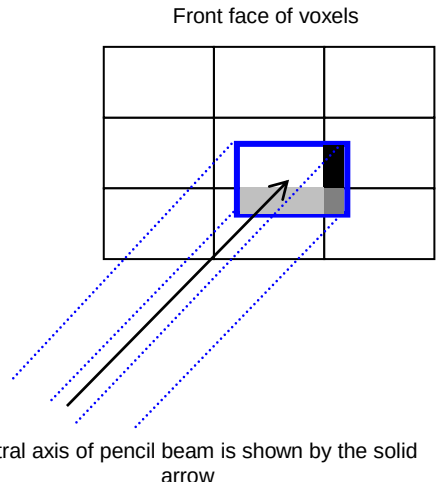

Figure 4.21: Pencil beam impinging on neighbouring voxels in a PA slice subsequent to the first. Area of beam in neighbouring voxels are highlighted by various shades of grey.

Figure 4.21 demonstrates that the central axis of the beam does not impinge the centre of the central voxel. The dark blue square represents the face of the pencil beam impinging not only the voxel in which the central axis impinges, but also the neighbouring voxels.

The ray casting technique requires all pencil beams to sample each voxel they encounter as they project through each PA slice. Therefore, for each pencil beam projected through the PA CT dataset, it is important to know how far its central axis has travelled through the data. In other words, having entered a given voxel in the first PA slice, one must know which voxels the pencil beam samples in all subsequent PA slices. Therefore: (1) the area of pencil beam in the centre and neighbouring voxels can be calculated and, (2) the correct voxels are sampled in each PA slice. To do this, one must know the physical length of the central axis of each pencil beam from the entrance point of PA slice 1 to the central plane of a subsequent given PA slice (as sampling areas were calculated in the central plane), and how far it has travelled in the left/right (and up/down) direction. An example of this process is shown in a plan view

(i.e. looking down on the patient) of the PA data in Figure 4.22, and magnified front view (i.e. looking through the patient) in Figure 4.23.

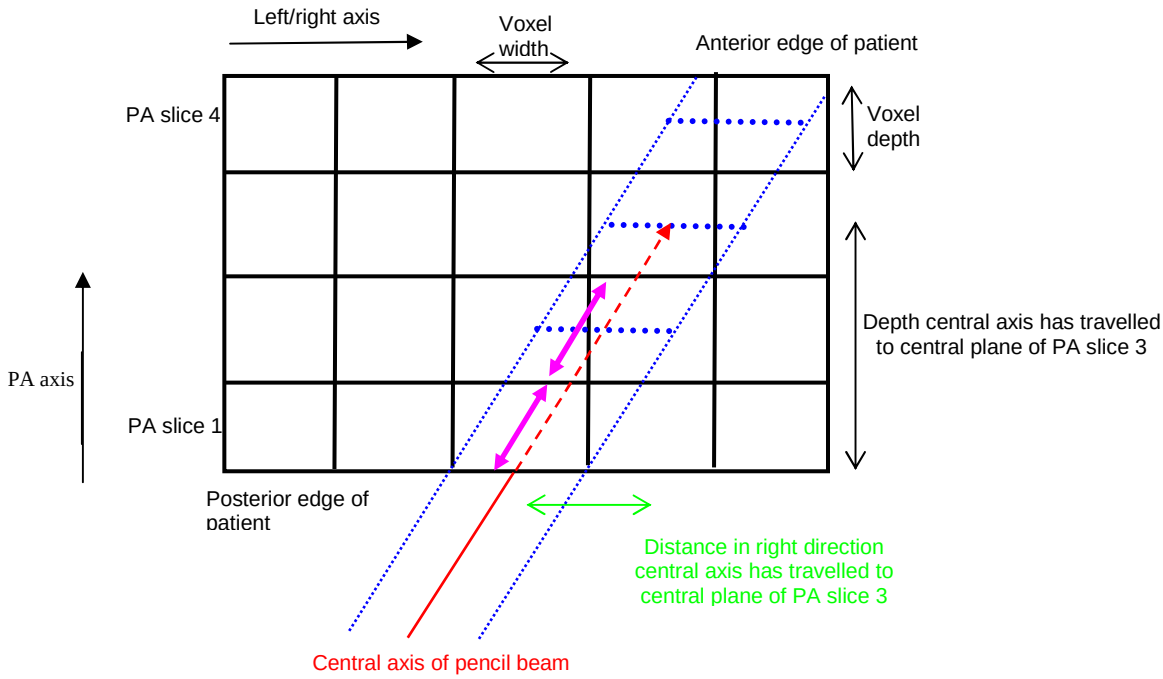

Figure 4.22: Plan view of a pencil beam (thin dashed blue lines). The central axis of the pencil beam is shown by the red arrow (the length of the central axis within the PA dataset up to central plane of PA slice 3 is shown by the dashed red arrow). The thick blue dashed lines are the voxel area sampling points through the CT data. The sampling point of the pencil beam in slice 3 is depicted 'face on' in Figure 4.23. The pathlength of the central axis travelled in each slice is shown by the pink double headed arrows (NB: this is the same for all PA slices).

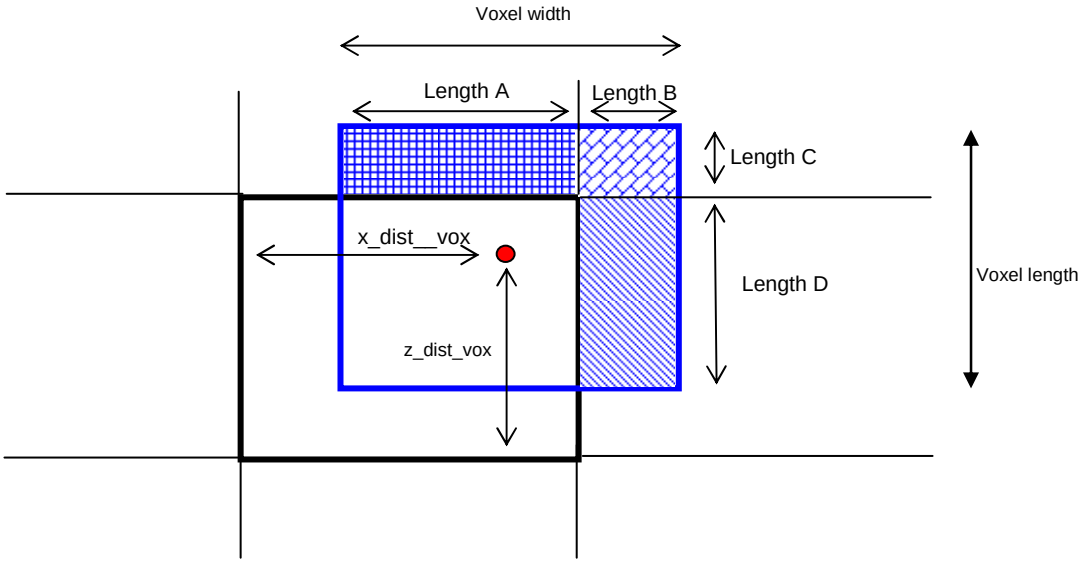

Figure 4.23: 'Face-on' view of a sampling point in a pencil beam (shown by the dark blue dashed lines in Figure 4.22. The solid black lined square is the entrance voxel, the solid blue square is the face of the pencil beam, and the red solid circle is the centre of the beam. The beam has been split into those areas that impinge on neighbouring voxels, as depicted by the dashed markings. Lengths of each side of each area are shown. z_dist_vox is the distance travelled by the central axis in the 'up' direction. x_dist_vox is the distance travelled in the 'right' direction.

The plan view of the PA dataset shown in Figure 4.22 depicts the central axis (red arrow) entering the CT data. The beam samples each voxel (based on area of beam in each voxel) that it encounters along its whole length until it leaves the data, these sampling points are shown in Figure 4.22 as the thick dashed blue lines, and 'face on' in Figure 4.23.

Pencil beam lengths and distances traversed through the PA data were calculated using equations that transform spherical coordinates to Cartesian and vice versa, as shown in Figure 4.24.

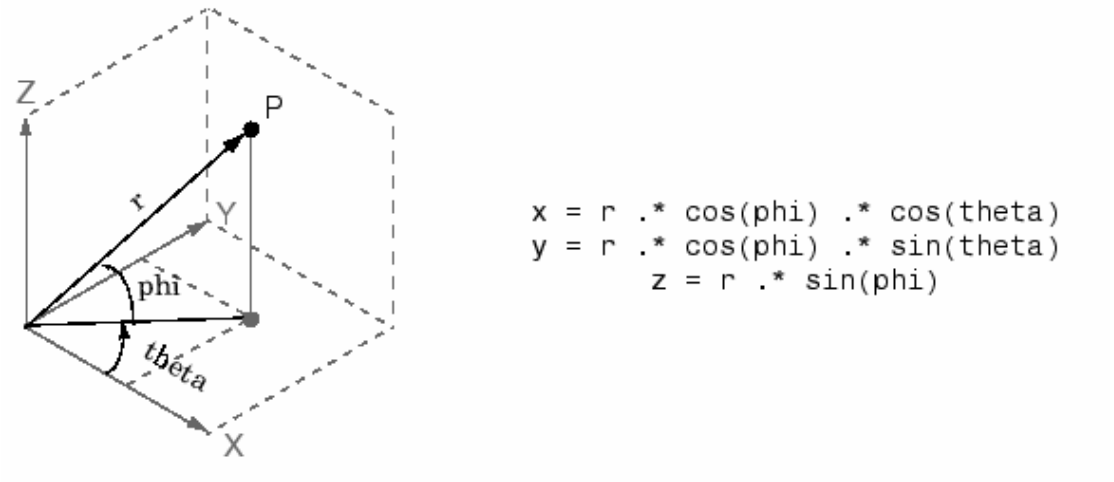

Figure 4.24: The mapping from spherical coordinates to three dimensional Cartesian coordinates.

Using the equations shown in Figure 4.24, the total length (TL) of the pencil beam if it were to reach the end (the whole PA-axis) of the PA dataset is given by:

$$\mathrm{TL} = \frac{(\mathrm{VD} \times 1024)}{\cos(\mathrm{phi}) \times \sin(\mathrm{pi}/2 - \mathrm{theta})} \tag{4.9}$$

where VD is the voxel depth (mm), 1024 is the number of voxels in the PA direction, and phi and theta are the elevation and azimuth angles in which the beam impinges the front surface of the dataset (Figure 4.20). The length of the pencil beam (LPB) inside the PA dataset up to the central plane of a given slice (e.g. up to slice 3 shown with red dashed arrow in Figure 4.22) is then calculated using the following:

$$\mathrm{LPB} = \left(\frac{(\mathrm{slice} \times \mathrm{VD})}{350} \times \mathrm{TL} \right) - \left(\frac{(0.5 \times \mathrm{VD})}{350} \times \mathrm{TL} \right) \tag{4.10}$$

where slice is the number of slices traversed by the pencil beam in the PA-direction, and 350 is the total PA-axis length (mm). The total distance traversed by the pencil beam central axis in the right direction (D_r) after entering the first PA slice (shown with a green arrow in Figure 4.22) is given by:

$$D_r = LPB \times \cos(phi) \times \cos(pi/2 - theta) \quad (4.11)$$

and the number of voxels in the right direction traversed (VTR) is:

$$VTR = \frac{D_r}{VW} \quad (4.12)$$

where VW is the voxel width.

It is now possible to calculate how far the central axis of the pencil beam has travelled within the given voxel (left to right) of the given PA slice (i.e. exactly where the end of the red dashed arrow is in Figure 4.22 in relation to the left edge of the voxel) using the following:

$$DTV_r = (VTR - VTR_{round}) \times VW \quad (4.13)$$

where VTR_{round} is the total number of voxels travelled by the central axis of the pencil beam rounded down to the nearest whole number such that VTR - VTR_{round} = fraction of voxel in right direction the central axis has travelled within the voxel. Note: the nomenclature used here is for right motion of each pencil beam, but it is easily adapted to left and up/down motion so that one can calculate exactly where the central axis is, as shown in Figure 4.23.

The respective sampling areas incident on neighbouring voxels were then computed. In the example in Figure 4.23, the centre of the beam is in the top right of the voxel. As such, the pencil beam impinges on the above, right and above-right voxels (demonstrated with different blue markings). The lengths of the sides of each area were calculated by the following equations:

$$length\_B = DTV_r - \left(\frac{VW}{2}\right) \quad (4.14)$$

$$\text{length\_A} = \text{VW} - \text{length\_B} \tag{4.15}$$

$$\text{length\_C} = \text{z\_dist\_vox} - \left(\frac{\text{VL}}{2}\right) \tag{4.16}$$

where VL is the voxel length,

$$\text{length\_D} = \text{VL} - \text{length\_C} \tag{4.17}$$

It was then a relatively simple matter of using these lengths to calculate the area of pencil beam which samples each voxel. Each area was then found as a ratio of the total beam area (i.e. the area of the blue square in Figure 4.23), and called ratio area (RA).

All path length and RA indices calculated here are universal and irrespective of what patient data is being used. As such, all were pre-calculated and saved to disc, and read back in when required.

During DRR generation, all pencil beams sample each voxel they encounter as they project through the data. The effective LAC of each sampling point (thick blue dashed horizontal lines in Figure 4.22) was found by weighting the LAC of the relevant voxel to the RA in the voxel, and summing them together:

$$LAC_{eff} = (RA_{incident} \times LAC_{incident}) + (RA_{right} \times LAC_{right}) + (RA_{above} \times LAC_{above}) + (RA_{corner} \times LAC_{corner}) \tag{4.18}$$

The intensity of X-ray photons exiting is calculated with the following formula:

$$I_E = I_0 \exp(-\text{pathlength} \times [LAC_{eff\_PA\_slice1} + LAC_{eff\_PA\_slice2} \ldots + LAC_{eff\_PA\_sliceN}]) \tag{4.19}$$

where I_0 is the intensity of the X-ray photons impinging on the surface voxel of the first slice of the PA CT data set, pathlength is the length travelled by the

central axis of the beam in each PA slice as shown by double headed pink arrows in Figure 4.22 (pathlength is a constant value in each slice for a given pencil beam and is irrespective of the number of voxels traversed), and $LAC_{eff\_PA\_sliceN}$ is the effective linear attenuation coefficient for PA slice N. Equation 4.19 was applied for each pencil beam impinging the first PA CT data array. Subsequent PA arrays were subject to pencil beam intensities (i.e. their respective I_0) calculated by equation 4.19 for the previous PA array. In other words, the calculated I_E for PA array 1 was the I_0 for array 2. This process was carried out for all reference photon energies (Section 4.2.5). However, there were still considerable computational timing issues due to vast amounts of sampling by each pencil beam. Therefore, a simplified method of shear transformation (discussed in Chapter 2) was investigated, as described in the next section.

4.2.7 CT voxel shifting prior to pencil beam projection

The DRR calculation process is very computationally intensive as it takes approximately 10 hours to compute a single DRR generated image on a modern PC (processor: Intel® Core™2 Quad 2.5 GHz, memory: 6 GB of RAM). It takes this long because each pencil beam enters a surface voxel with a given angular displacement (theta and phi), and as such the software has to perform a calculation to ‘inform’ each beam which voxels to sample in a given PA slice. In practice this necessitates voxels being shifted into the path of each pencil beam on a DRR by DRR basis.

As already discussed, axial CT data was accessed and re-orientated, so the patient was in the PA position. As shown by the plan view in Figure 4.25, each

pencil beam enters a specific voxel and travels with an angular displacement, sampling voxels it encounters as it projects through the data.

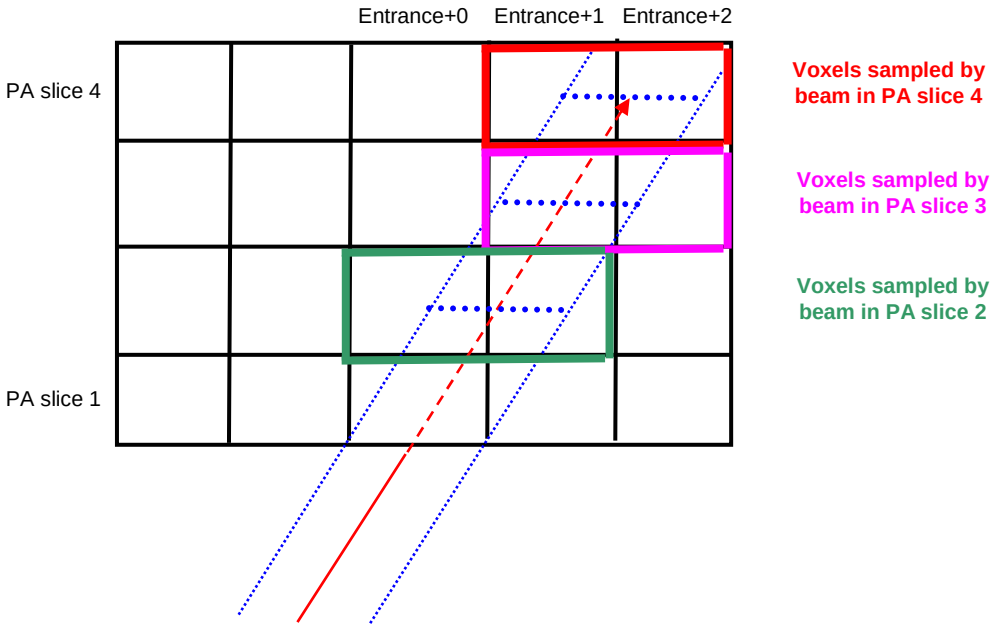

Figure 4.25: Plan view of pencil beam (thin blue dashed lines) sampling voxels at the central plane of each PA slice (thick blue horizontal dashed lines) as it moves through the data. In this example, the beam samples two voxels in each PA slice, all of which have been colour coded in the Figure. Note: for simplicity, only the entrance voxel is sampled in PA slice 1.

Figure 4.25 shows the voxels sampled by the beam as it travels through the data (for right hand movement only), and all sampled voxels are highlighted. For example, in PA slice 2 the beam samples the voxel directly behind the entrance voxel (one may call this 'column' [entrance + 0]) as well as the one to its right [entrance + 1]. However, the central axis is incident on [entrance +1] at the central plane (horizontal blue dashed line) of the PA slice. In PA slice 3, the beam samples [entrance + 1] and [entrance + 2], but its central axis is incident on [entrance + 1] at the central plane. In PA slice 4, the beam samples [entrance + 1] and [entrance + 2] and its central axis is incident on [entrance + 2] at the central plane. For the beam to sample the voxel in which the central axis is indicent at the central plane (in each PA slice), the data can be shifted

such that the beam may travel in a parallel manner through the data, as depicted in Figure 4.26.

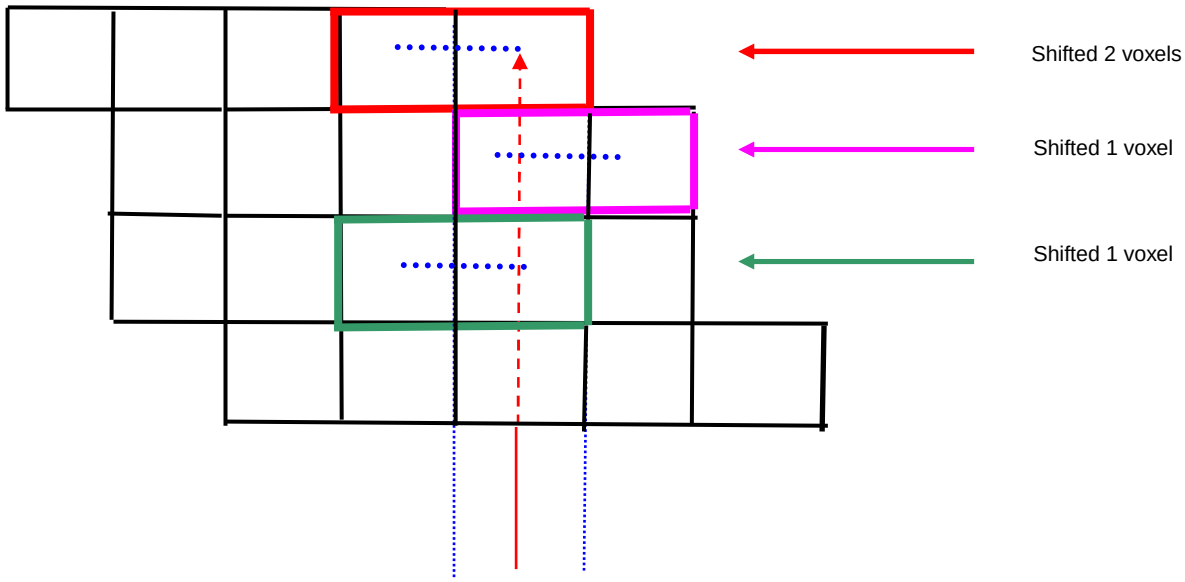

Figure 4.26: Voxelated data shifted such that each pencil beam may pass through in a parallel manner. Note that the beam samples the same voxel area as that in Figure 4.25 (shown by the thick dashed blue lines).

Using this method, the central axis of the beam samples [entrance + 1] in PA slice 2, [entrance + 1] in PA slice 3 and [entrance + 2] in PA slice 4 in a single parallel projection. It should be noted that this is a simplifed version of sheared object space (Section 2.3.1), because the data was sheared linearly, rather than in a non-linear fashion, i.e. voxels were shifted into the path of the pencil beam and were not in themselves 'warped' (it should be remembered here that each pencil beam *does not* itself fan out (for simplification), but stays the same size as it projects through the data; each beam does however move away from one another). In practice, for each pencil beam (using equations 4.9 to 4.12), the number of voxels the central axis (i.e. the red dot in Figure 4.23) had traversed by the time it had reached the central plane of a given PA slice was calculated. For example, if a pencil beam entered the data and by the time it reached PA slice 100 had travelled 10 and 5 voxels in the 'right' and 'up' directions respectively, this voxel was shifted 10 and 5 voxels 'left' and 'down'. However,

rather than do this on a voxel by voxel basis, the software was configured such that within each PA data set (Figure 4.19) the minimum and maximum amount of voxels traversed was calculated for all pencil beams incident on the PA data. The central axis of some incident pencil beams did not move out of their incident voxel columns (i.e. they stayed in [entrance + 0]), or in other words they did not traverse ‘left/right’, ‘up/down’ at all; this was certainly true for beams projected from the X-ray source near to the centre of the PA data, i.e. those with very small values of theta and phi. Voxels where this *was not* true were given a voxel value of zero (voxels where this *was true* were held at their voxel value), the entire data was then shifted by zero and the data saved to disc; in this way more than one voxel at a time with the correct voxel values were shifted. However, some pencil beams did traverse voxels in the ‘right/left’ and ‘up/down’ direction as they travelled through the PA data, e.g. some beams travelled one voxel in the ‘right’ direction within the PA data, and none in the ‘up’ direction. Voxels where this *was not* true were given a voxel value of zero (voxels where this *was true* were held at their voxel value), the entire data was shifted ‘left’ by one voxel and none in the ‘down’ direction. This shifted data was added to the original shifted data (i.e. that for zero shift discussed above). This process was repeated for *all* possible combinations of the above, saved to disc *prior* to DRR calculation and accessed when required (it was possible to do this because the paths of each pencil beam were pre-calculated – see equation 4.9 to 4.17). In carrying out this process, the shifts depicted in Figure 4.26 was possible. This only needed doing once per patient and took approximately 10 hours. If the data was not shifted prior to pencil beam projection and read in when required, every DRR image generated would take this long to compute. Therefore, if there are 11 images per patient (for example, images reconstructed with tube potentials

50 kVp to 150 kVp in steps of 10 kVp) the entire process per patient would take approximately 110 hours to complete. However, using the method described here, shifted data is only accessed when required for parallel DRR computation (see below), and each DRR image took approximately 45 - 90 minutes to compute.

When the shifted data was accessed, each pencil beam was projected through the data in a parallel manner. This shifted data was only correct for area sampling in the incident voxel, so $RA_{incident}$ (also accessed from disc) was multipled with the incident voxel LAC, i.e. $RA_{incident}$ x $LAC_{incident}$ was calculated. For PA slice 2, $RA_{incident}$ x $LAC_{incident}$ was calculated, and the array was then shifted one voxel to the right (Figure 4.25 shows the voxel to the left of the incident one requires sampling, i.e. LAC_{left}), LAC_{left} was accessed and RA_{left} x LAC_{left} performed. This was repeated for the relevant adjacent voxels and equation 4.18 used to calculate LAC_{eff} for that PA slice. For PA slice 3 $RA_{incident}$ x $LAC_{incident}$ was calculated then the array was shifted one voxel to the left to access LAC_{right} (see in Figure 4.25 it is the voxel to the right of the incident one that requires sampling). The relevant adjacent calculations were performed and LAC_{eff} derived. Similarly, $RA_{incident}$ x $LAC_{incident}$ was calculated for PA slice 4, the data shifted one voxel to the right (see Figure 4.35) and RA_{left} x LAC_{left} performed. This was repeated for all parallel pencil beams incident on the data, and effective LACs were calculated (equation 4.18) in each PA slice. The process described in this section was repeated for each reference energy until the photon intensity (per photon energy; I(E)) emerging from the patient in each pencil beam was derived.

4.2.8 Absorption of X-ray energy by the CR phosphor

Having determined the photon intensity exiting the patient, the X-ray energy absorbed by the layer of phosphor was calculated [95] with the following equation:

$$A = \int_{0}^{Emax} E\left[I(E)\{1 - \exp\left(-\left(\frac{\mu_{en}(E)}{\rho} \right)_{CR} \rho x \right)\} \right] dE \qquad (4.20)$$

where $(\mu_{en}(E)/\rho)_{CR}$ is the photon energy dependent mass energy-absorption coefficient of the CR phosphor; ρx the mass loading of the phosphor (mass per unit area of phosphor; g cm^{-2}); I(E) the photon intensity (per photon energy) in each pencil beam incident on the CR phosphor and E the photon energy. The value of ρx used was 0.08 g cm^{-2} (Private Communication, Mark O'Herlihy, Agfa, 2009). The respective atomic weight (g/mol) of each element of the CR phosphor (BaSrFBrI:Eu) was found as a percentage of the total weight, and each value was multiplied with its respective μ_{en}/ρ (taken from the National Institute of Standards and Technology (NIST) database [64]) to derive the total CR phosphor mass energy absorption coefficient at each photon energy, $(\mu_{en}(E)/\rho)_{CR}$. The results are shown in Table 4.5.

Photon Energy (keV)	CR Phosphor Mass Energy Absorption Coefficient (cm^2/g)
10	178.6
15	46.1
20	24.1
30	20.6
40	7.5
50	4.9
60	3.3
80	1.7
100	0.9
150	0.3

Table 4.5: Mass energy absorption coefficients for the CR phosphor used in this work (K edges not shown).

Only the energies shown in Table 4.5 were available from the NIST database (K edges not shown). However, as descibed previously, each X-ray spectra used in this work contains energies 0.5 keV to keV_{max} in steps of 0.5 keV. Therefore, to calculate the energy absorbed in the CR phosphor, mass energy-absorption coefficients matched to the energies present in each X-ray spectra were derived. This was done using cubic interpolation (in-built Matlab function) of the values in Table 4.5. Figure 4.27 shows the CR phosphor mass energy-absorption coefficient in graphical form for all photon energies contained within a 50 kVp spectrum.

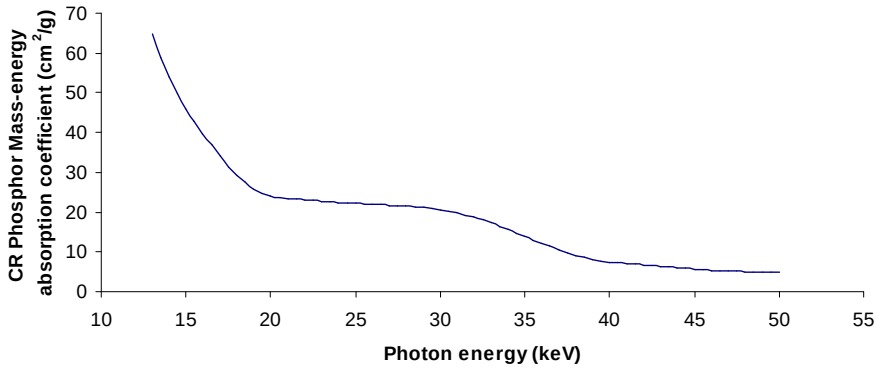

Figure 4.27: Continuous CR mass energy-absorption spectrum used in equation 20 for a 50 kVp DRR reconstruction. K absorption edges are not shown.

Equation 4.20 was used for each photon energy incident on the CR phosphor together with the relevant $(\mu_{en}(E)/\rho)_{CR}$ to calculate the total absorbed energy per pencil beam. Although K-edge information are provided by NIST, increased photon absorption due to the K-edges of Iodine and Barium (33.2 and 37.4 keV respectively) were included in a separate calculation because they were not available in steps of 0.5 keV (i.e. as discussed above, the algorithm uses spectral data at every 0.5 keV). All subsequent DRR pixel values were displayed as energy absorbed by the CR phosphor and were linear with X-ray beam air kerma (and tube mAs) at the DRR image plane. As discussed in Section 4.2.6 each DRR was projected with a 1:1 mapping between pixels in the DRR and voxels on the front face (first PA slice) of the CT dataset. As such, there was one DRR pixel for each pencil beam.

4.2.9 Linearisation of real CR pixel value data

To measure scatter and scatter fractions (SFs), and add frequency dependent noise to a DRR (see Sections 4.2.10 and 4.2.11 respectively), images must be collected from the clinical CR system. DRR images are displayed in terms of energy absorbed by the CR phosphor, and are linear (i.e. DRR pixel values

increase linearly with mAs). Due to this, images collected from the CR system must also be linear. However, CR systems are seldom linear, as is the case with the Agfa system used in this work. Pixel values of images acquired on the CR system were linearised by measuring the system transfer function (detector response). This was done by obtaining a series of six uniform open field exposures at each tube potential (50 kVp – 150 kVp in approximate steps of 10 kVp). The exposures covered a range of approximately 1 – 15 µGy air kerma at the cassette measured with the cassette removed and replaced with a calibrated 6 cc ionization chamber (Radcal Corporation, Monrovia, USA). Patient attenuation was approximated using 20 cm of polymethyl methacrylate (PMMA) at the tube port with a focus to cassette distance of 180 cm (PMMA used due to its similar X-ray absorption properties of water). This experimental set-up is shown in Figure 4.28.

Figure 4.28: X-ray tube with 20 cm PMMA attenuating material at the port. The CR cassette is not shown here but is 180 cm from the tube focus.

Mean CR pixel vales were obtained for each uniform image from a central 12.5 cm square, as shown in Figure 4.29.

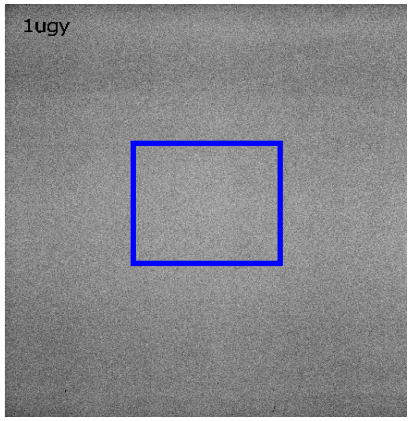

Figure 4.29: Uniform open field image collected from the CR system. The mean pixel value is calculated from central ROI (blue square).

For each tube potential used to collect the open field images, the energy absorbed in the CR phosphor was calculated using the method described in Section 4.2.8 but with X-ray spectra corrected for the extra 20 cm PMMA attenuation (measurement of air kerma here was not used for calculation of energy absorbed in the phosphor, but to ensure exposure factors used were in the correct diagnostic energy range, i.e. 1 – 15 µGy). The relationships between mean CR pixel value and energy absorbed in the phosphor for each tube potential were used to linearise all images acquired on the CR system (i.e. clinical CR pixel values were converted to X-ray energy absorbed by the CR phosphor).

The relationship between energy absorbed by the CR phosphor and CR pixel value were all found to be logarithmic ($\log_e$, all $r^2 > 0.9995$) and demonstrated slight dependence on tube potential, as shown in Figure 4.30.

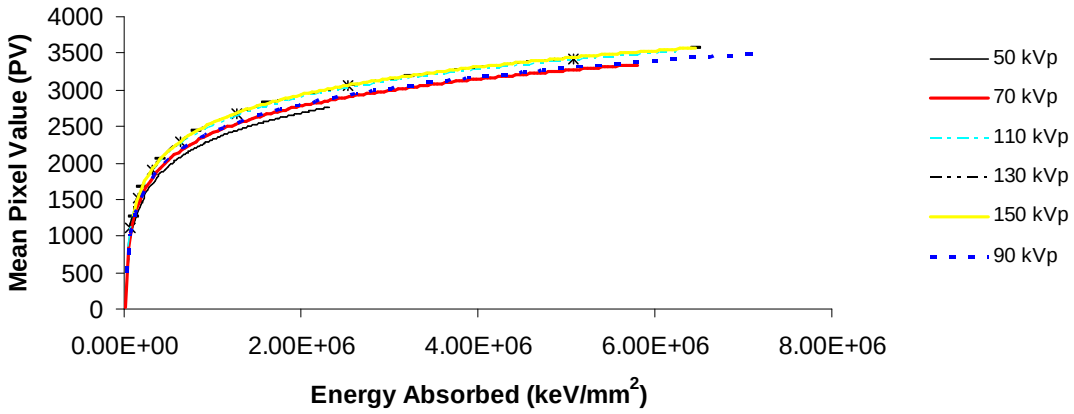

Figure 4.30: Relationship between mean CR pixel value and energy absorbed in CR phosphor. For clarity, plots for other tube potentials are not shown.

The dependence on tube potential of absorbed energy and pixel values is probably due to k-edge absorption. To account for the slight dependence, CR image data were linearised using the correct tube potential dependent logarithmic equations.

4.2.10 Radiation scatter measurement and addition to DRR

As discussed in Chapter 3, scattered photons that reach the image receptor do not contain clinically useful information and degrade the image quality of chest radiographs by creating a non-uniform background that reduces image contrast. The scatter fraction (SF) is defined as the ratio of the intensity of scattered radiation to that of total (scattered plus primary) radiation recorded in the image.

The algorithm that produces the DRR images do not model any scatter contributions, and therefore it must be added post calculation. To do this, the method discussed in Chapter 3 was utilized. Measurements of scatter (and SFs) representative of the whole chest were made using the most common method described in the literature; an array of 224 lead beam stops. The lead stops were each of 6 mm in thickness and 3 mm diameter, 25 mm apart,

suspended on a 1 mm thick PMMA sheet. The array was positioned in front of the chest portion of RANDO, as shown in Figure 4.31.

(a) (b)

Figure 4.31: RANDO phantom with array of lead stops positioned just in front. (a) RANDO without added fat and (b) with added fat.

Figures 4.31(a) and 4.31(b) show the lead stop array positioned in front of the phantom. Images were acquired on the Agfa CR system for a range of diagnostic tube potentials (50 – 150 kVp in approximate steps of 10 kVp) initially with no scatter rejection (as per local Radiology Department procedure). Images across the same range of tube potentials were then acquired with two scatter rejection techniques; firstly with an anti-scatter grid (strips per mm = 4, grid ratio = 12) focused at 140 cm focus to detector distance (useful range 115 to 180 cm), then with an air gap between RANDO and the cassette. The anti-scatter grid set-up was identical to that described above, except the grid was energized, the chest stand removed and the CR cassette placed in the Bucky. The air gap set-up differed from local protocol in that RANDO was positioned 20 cm further away from the cassette (i.e. 20cm closer to the tube) on the advice of the expert Radiographer (private communication, Jo Cook, 2010).

As primary X-radiation is almost entirely absorbed by the lead stops, the resulting shadows in the radiograph provide an estimate of scatter (this

assumes system noise in each shadow is consistent throughout the image, which is verified by annual quality assurance tests). An image of the beam stops with and without RANDO was acquired at each tube potential with a sufficient tube current-time product (mAs), to provide a lgM of 2.00 ± 0.05, as depicted in Figure 4.32.

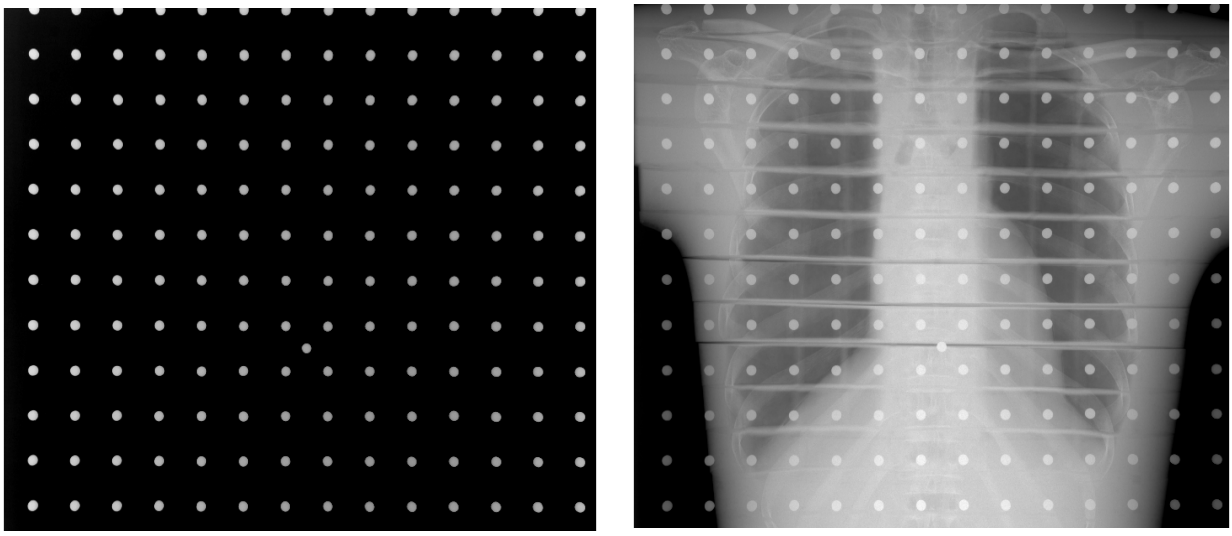

(a) (b)

Figure 4.32: (a) An image of the lead stops without RANDO, and (b) an image of the lead stops positioned in front of RANDO. The shadows of the lead stops are easily visible in both images.

The reason for acquiring an image of the lead stops without RANDO (Figure 4.32(a)) was because it provided a 'lead stop mask', which aided the software's application of a boundary tracing algorithm (Matlab in-built function) to find the coordinates of the perimeter of each shadow. This was much easier to do in an image without the background detail in Figure 4.32(b). The perimeter coordinates were then transferred to the image with RANDO, and, because the images were taken one after the other without the lead stop array being moved (and the CR cassette being carefully repositioned in the same place each time), the coordinates approximately fit over the perimeter of each lead stop shadow in the RANDO image (Figure 4.32(b)). When any mismatch occurred, the mask

was shifted up, down, left or right in a trial and error fashion until an adequate fit over the shadows was achieved. An image of RANDO and lead stops overlaid with the mask is shown in Figure 4.33.

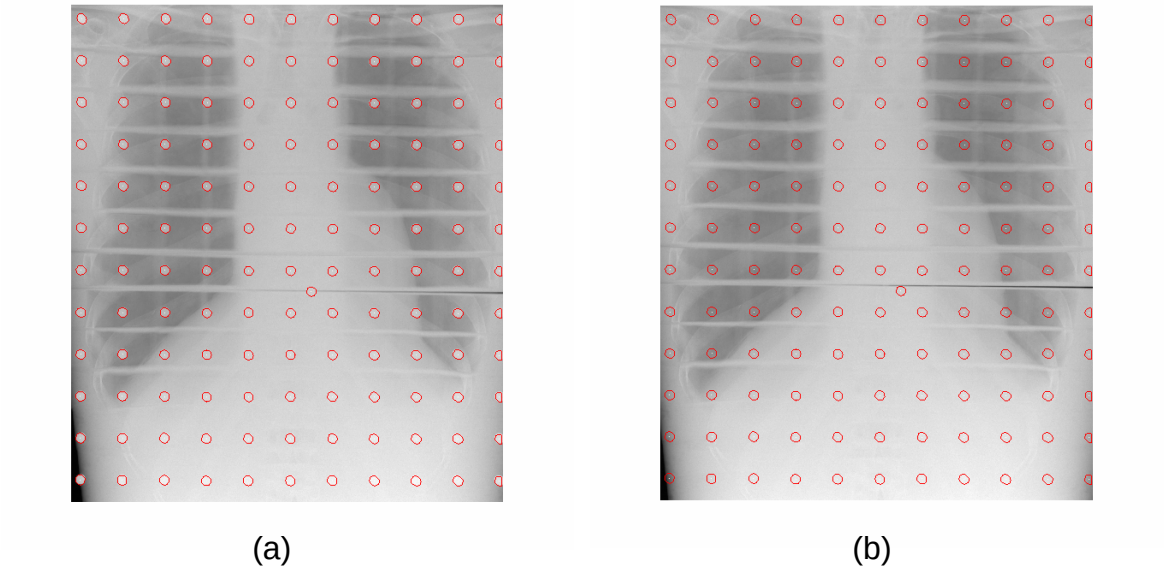

(a) (b)

Figure 4.33: (a) RANDO with lead stops overlaid by the mask. The red circles enclose the lead stop shadows. (b) RANDO overlaid by mask without lead stops present. Notice the pixel values enclosed by the red circles are much different to those shown in Figure 4.33(a).

Each image was linearised in terms of energy absorbed by the CR phosphor, and the mean of the pixel values enclosed within each shadow were calculated. Each mean pixel value is a measure of scatter in terms of the energy absorbed in the phosphor. The total energy absorbed (primary + scatter) was measured in exactly the same manner, but with RANDO images acquired without the lead stops present (as shown in Figure 4.33(b)).

Scatter and SFs (all linear in terms of energy absorbed by CR phosphor) were measured at the position of each lead stop, and a 2-D interpolation program (bi-cubic interpolation that fits a bi-cubic surface through existing data points) was used to calculate the values of scatter and SFs across the entire image. An

example of scatter acquired with two tube potentials without scatter rejection is illustrated in Figures 4.34(a) and 4.34(b).

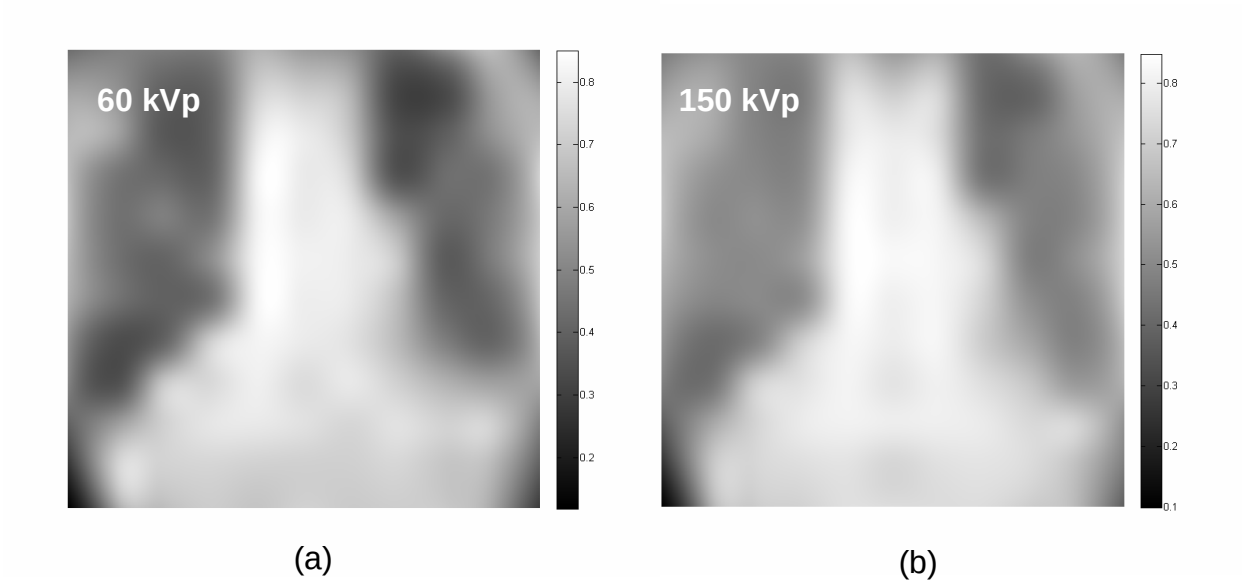

Figure 4.34: Images of scatter factors with RANDO acquired at (a) 60 kVp, and (b) 150 kVp. Scatter factor values are shown on the grey scale bar down the right hand side of each Figure.

Figure 4.34(a) was acquired with a 60 kVp tube potential. Scatter factors in the lung and spine/diaphragm regions range from 0.33 to 0.47, and 0.66 to 0.85 respectively. Figure 4.34(b) was acquired with a tube potential of 150 kVp, and scatter factors range from 0.39 to 0.53 in the lung, and 0.69 to 0.88 in the spine/diaphragm. These values are in general agreement with SFs measured by Floyd *et al.* [86] in humans. There is little change in SF in the spine/diaphragm regions with change in tube potential, but the effect is slightly more pronounced in the lung. This is similar to that reported by Bowenkamp *et al.* [82].

Figure 4.35 shows the difference between scatter factors measured in an image without scatter rejection and an image with scatter rejection (i.e. grid), both acquired with a tube potential of 60 kVp. It is clear that scatter factors are lower in the image derived with a scatter grid (Figure 4.35(b)), which would be

expected. The SF range in lung is 0.15 to 0.33 and spine/diaphragm is 0.26 to 0.54. SFs are therefore, on average, approximately 40% lower in the lung, and 48% lower in the spine/diaphragm regions, when compared with the image acquired with no scatter rejection (Figure 4.35(a)).

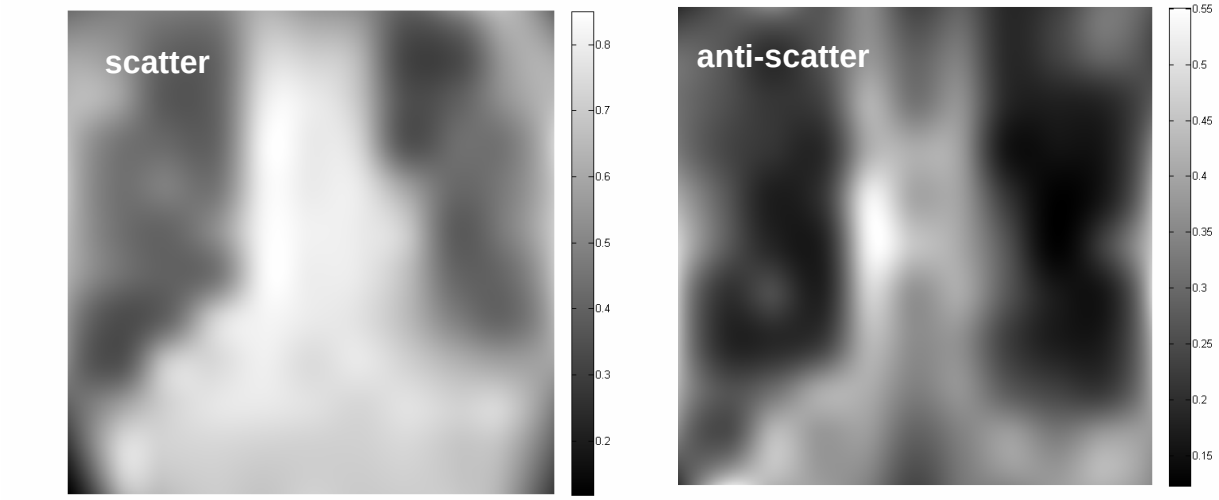

(a) (b)

Figure 4.35: Scatter factor image of RANDO acquired at (a) 60 kVp with no scatter rejection, and (b) image of RANDO acquired with the use of an anti-scatter grid.

Similarly, using an air gap scatter rejection technique, SFs ranged from 0.29 to 0.39, and 0.65 to 0.68 in the lung and spine/diaphragm regions, respectively.

Since each DRR is already linear one can define the total energy absorbed by the CR phosphor as DRR_{total} (as scatter does not exist in the simulation, energy absorbed in the CR phosphor from all pencil beams is not simply primary absorption, but total). The amount of primary absorption (DRR_P) can then be calculated by removing a portion of the signal from DRR_{total} by applying linear SFs measured experimentally using the following equation:

$$\mathrm{DRR_P} = \mathrm{DRR_{total}} \times (1 - \mathrm{SF}) \tag{4.21}$$

Linear scatter (measured experimentally) can then be added to DRR_p. The following equation was used:

$$DRR_{P+S} = DRR_P + \text{scatter} \quad (4.22)$$

where DRR_{P+S} is the primary DRR with linear scatter added. Scatter and SF masks here have been derived with the RANDO phantom and so easily fit over raw DRR images of RANDO. However, although only average and obese sized males have been identified in this study (upon which RANDO has modelled), anatomy differs slightly from patient to patient, and as such scatter and SF masks will not always fit exactly over the raw patient DRR. However, simply adding the scatter and SFs from RANDO over raw DRRs from patient CT data, as a first approximation, was deemed acceptable because any gross errors encountered when validating the model (Chapter 5) could be investigated if required.

It should also be noted that throughout this chapter, SFs were not corrected back to 'zero lead stop diameter' from 3 mm as it has been shown [81, 82, 85, 86] this makes less than 3% difference to the resulting SFs.

4.2.11 Addition of frequency dependent system noise to DRR images

As discussed in Chapter 1, it is important to simulate CR system frequency dependent system noise in the DRR images to allow dose optimisation studies to be carried out.

Frequency dependent system noise was added to each DRR_{P+S} using a slightly different method to that described by Bath *et al.* [49]. Their work demonstrated it

was possible to create or collect an image containing only noise, which, when added to the original clinical image scaled to a lower dose level, resulted in the same noise properties (i.e. same noise power spectrum (NPS)) as an image acquired on the clinical CR system at the lower dose level. They argue that simulating a clinical image at a lower dose (dose of the simulated image, D_{sim} < original image), Im_{sim}, is given by $Im(x,y)_{sim} = Im(x,y)_{orig\ scaled\ down} + Im(x,y)_{noise}$ where $Im(x,y)_{orig\ scaled\ down}$ is the original clinical linear image scaled down to the simulated dose (D_{sim}), and $Im(x,y)_{noise}$ is an noise image acquired on the CR system or produced artificially at the lower simulated dose D_{sim}, resulting in the equality $NPS(u,v)_{Imsim} = NPS(u,v)_{Dsim}$. However, because Bath *et al's* original clinical image is not system noise free, the 2D NPS of their noise image is given by:

$$NPS(u,v)_{Imnoise} = NPS(u,v)_{Dsim} - NPS(u,v)_{Dorig}\left(\frac{D_{sim}}{D_{orig}}\right)^2 \qquad (4.23)$$

The second term in equation 4.23 acts to correct for the noise in the scaled clinical image (pixel values are correct in the scaled clinical image, but the noise properties are not).

The work in *this thesis* differs in that the raw DRR image contains no CR system noise (the raw DRR does not need scaling because it can be produced at the necessary level of dose), as such the second term in equation 4.23 becomes zero. This means that $NPS(u,v)_{Imnoise} = NPS(u,v)_{Dsim}$ suggesting a noise image acquired on the CR system at a given dose level (effectively D_{sim}) will contain the correct frequency dependence required. Also, because the uniform noise image will be acquired directly from the CR system itself, *all* CR system noise sources will be present (see Section 1.4.1). This method assumes

that the NPS is a sufficient descriptor of noise and that detective quantum efficiency (DQE) is constant over the dose range within the image. However, CR systems tend to have a decreasing DQE with dose, but Bath *et al.* argue their method is sufficient for doses used clinically (limitations of the DRR simulation model are discussed in Chapter 5).

A series of uniform noise images were collected from the CR system using the same experimental set-up described in Section 4.2.10. Images were acquired at tube potentials 50 – 150 kVp in approximate steps of 10 kVp across a range of clinically relevant mAs values. Each noise image was linearized and the DC signal (mean value) set to zero (the addition of the noise image must not alter the mean pixel value of the simulated image). As each DRR_{P+S} is an inhomogeneous image, corrections to the uniform noise image must be made since the absolute noise in the low dose areas of a CR image would in reality be lower than that in the high dose areas. To take these local dose variations into account, the following correction is applied to the uniform noise image:

$$\text{Im\_noise}(x,y)_{corr} = \text{Im\_noise}(x,y)_{unif}\sqrt{\frac{DRR_{P+S}}{PV_{mean}}} \tag{4.24}$$

where $\text{Im\_noise}(x,y)_{corr}$ is the corrected noise image, $\text{Im\_noise}(x,y)_{unif}$ is the uniform noise image, DRR_{P+S} is the primary DRR with scatter added and PV_{mean} is the mean pixel value of $\text{Im\_noise}(x,y)_{unif}$. It is very important to remember here that equation 4.24 only holds for images that are dominated by primary and secondary quantum noise (i.e. those noise sources governed by Poisson statistics – see Section 1.4.1). As images acquired at clinical doses are indeed dominated by these noise sources, a simple square-root relationship is an appropriate way of correcting local dose variations.

A uniform noise image at each tube potential was acquired corresponding to the same level of air kerma incident (7.0 ± 0.4 μGy at each tube potential for a lgM = 2.00) at the CR cassette through the lung region of each $DRR_{P+S.}$ This air kerma value was established using a pixel value to air kerma relationship derived previously for this CR system (i.e. the mean pixel value in the lung region was measured and converted to air kerma). Each uniform noise image was corrected according to equation 4.24 and added to the DRR_{P+S}.

Figure 4.36 shows an example of an uncorrected and a corrected noise image derived using the method described above.

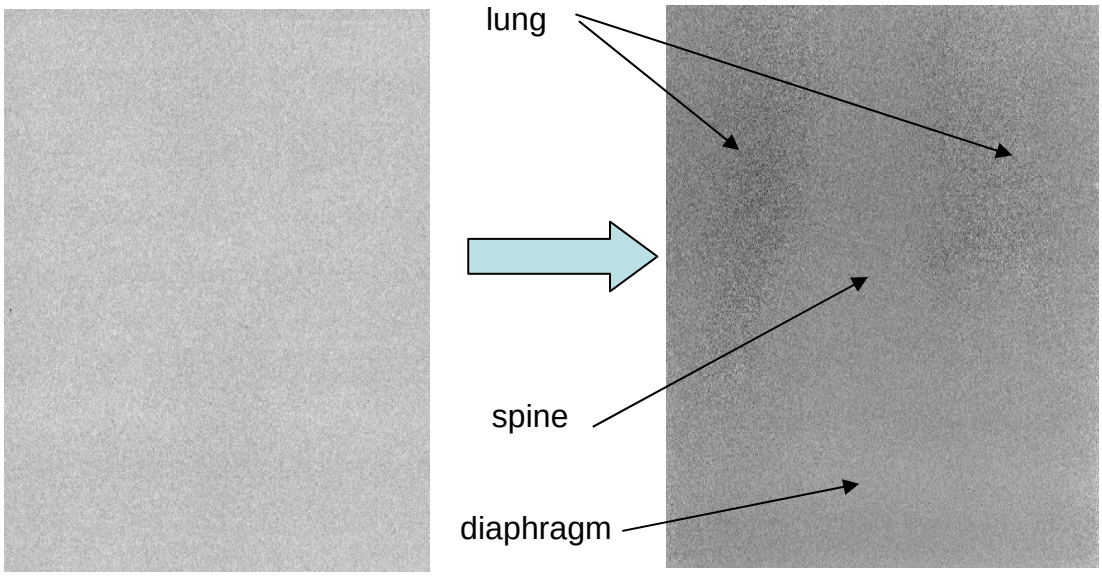

Figure 4.36: (a) Un-corrected and (b) corrected noise images.

The noise image in Figure 4.36(b) shows the effect of local dose variations. The added noise in the diaphragm and spine regions is lower than that in the lung regions. The noise in the lower dose regions would be overestimated (and underestimated in the higher dose regions) if no correction was applied. This method of noise addition is likely to be one of the largest sources of error, and is discussed in more detail in Chapter 5.

4.2.12 Lung nodule simulation

As well as simulating projected normal anatomy, it would be beneficial to simulate nodules in the chest radiograph so that optimisation can also be carried out for these abnormal structures. After discussion with an expert in our Radiology department (Private Communication, Dr Ged Avery, Consultant Radiologist, 2009) it was decided to simulate soft tissue nodules only in the lung. Lung nodules were chosen as they are indicative of common malignant disease such as cancer, and non-malignant diseases such as tuberculosis, pneumonia and sarcoidosis.

The report ICRU 70 [96] also recommends soft tissue lesions in the lung as an indication for chest radiography. Lung nodules were simulated based on the work described and validated by Li *et al.* [97]. They simulated 3D lung abnormalities with realistic characteristics by modelling multiple 2D masks on sequential CT slices. For a given nodule, the peak CT number on contiguous CT slices varies depending on its distance from the centre, and on each CT slice the 2D mask is defined by a contrast-profile equation proposed by Samei *et al.* [98] and reformulated by Burgess *et al.* [99]. This equation is as follows:

$$c(r) = C\left(1 - \left(\frac{r}{R}\right)^2\right)^n \tag{4.25}$$

where c(r) is the contrast profile dependent on distance from the centre of the 2D mask, r is the distance from the centre of the 2D mask, R is the radius of the 2D mask, C is the peak CT number of the 2D mask, and the exponent n is a positive number inversely related to the steepness of the contrast profile, reflecting edge characteristics. For this study, n = 2.416 was used [97]. The

mean CT number (± 1SD) of real lung nodules acquired on the Philips scanner used in this study was 1112 ± 6. Therefore, this value was used for C in equation 4.25. Recently, Gohagen *et al.* [100] carried out a lung cancer study and found that lung nodule sizes typically range from 4 mm to 16 mm. Therefore, in this study we chose to use a diameter of 10 mm (R = 5 mm). Hakansson *et al.* [18] have shown that detectability of lung nodules is dependent on location in the lung so it was decided to simulate nodules in the lateral pulmonary and hilar regions.

A magnified portion of lung in a CT slice with and without a simulated nodule is shown in Figure 4.37.

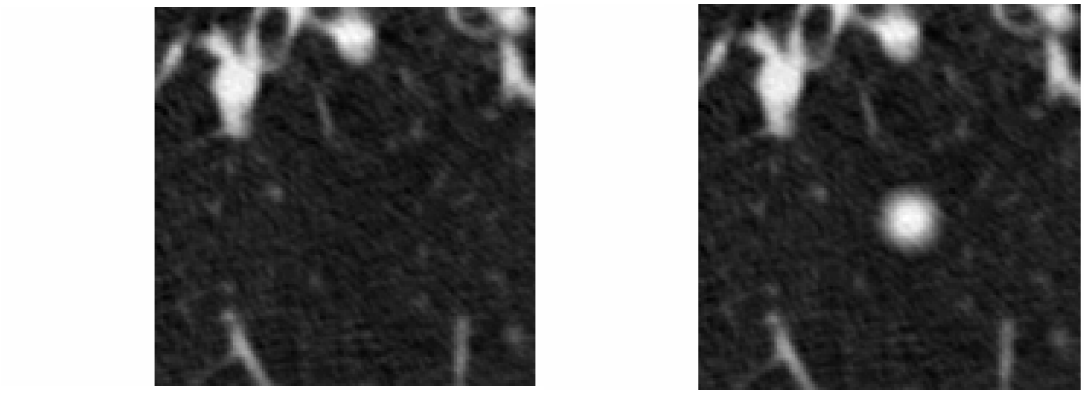

(a) (b)

Figure 4.37: (a) CT slice without a simulated nodule, and (b) with a simulated nodule.

The lung nodule is clearly displayed in 4.37(b). A comparison of resulting DRRs with and without the simulated nodule present is shown in Figure 4.38.

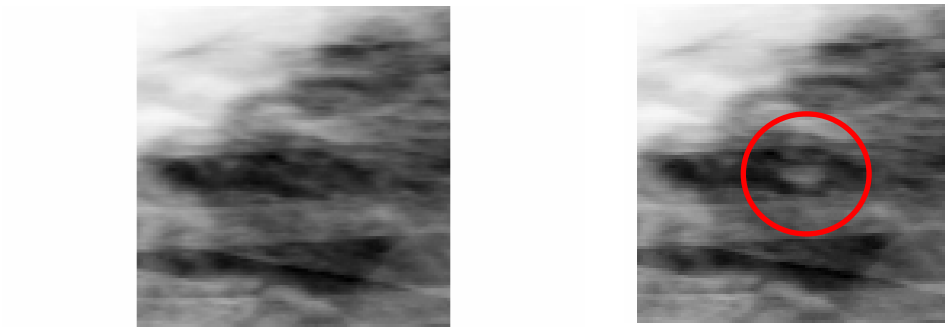

(a) (b)

Figure 4.38: Magnified DRR image of the lung (a) without nodule, (b) with nodule (centre of red circle).

Figure 4.38(b) clearly shows the nodule in the resulting DRR, and as such can be used in dose reduction studies for the purpose of optimisation, assuming qualitative validation provided by image evaluators proves favourable (see Chapter 5). Three dimensional lung nodules were added to CT data prior to X-ray projection rather than in two dimensions to the final DRR image because it would have been extremely difficult to generate the correct X-ray contrast of the nodule by simply adding the nodule post DRR calculation.

4.2.13 Fat simulation

It was necessary to add fat to average patient CT data to simulate obese patients as, at the time of this research, there were too few obese patients available on the scanner. This potentially proves an advantage, as artificially simulating fat would add to the flexibility of the DRR algorithm.

To add fat to average patient CT slices, the mean (± 1 SD) CT voxel value of lard displayed in images of the RANDO phantom with lard added (Section 4.2.1) was measured and found to be 940 ± 10. The thickness of the RANDO fat was 110 voxels (i.e. approximately 4 cm). Secondly, CT voxel values within patient CT images were all converted to one, and the voxel values outside of the

patient were converted to zero (i.e. a binary image was created). Images of an average patient CT slice and the corresponding binary image are shown in Figure 4.39.

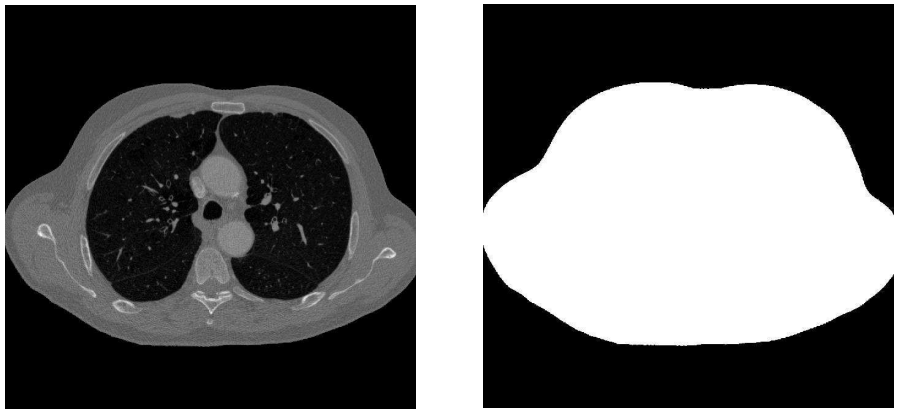

(a) (b)

Figure 4.39: CT slice of (a) an average patient and (b) its binary counterpart.

The binary image (Figure 4.39(b)) was then shifted up by 110 voxels (i.e. thickness of fat) and the original non-shifted binary was taken away from this to leave voxel values of ones where the fat is to be simulated (and zeros everywhere else). This process is shown in Figure 4.40.

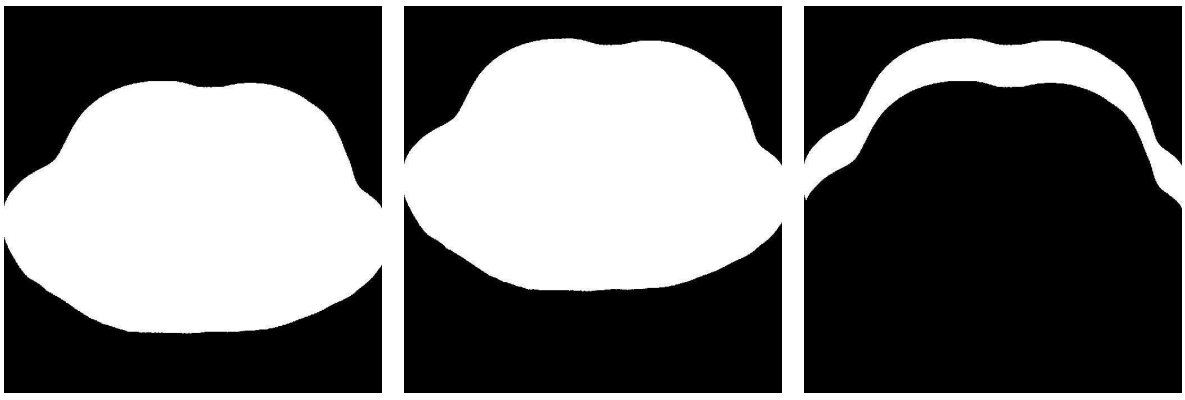

(a) (b) (c)

Figure 4.40: (a) Binary image of average patient, (b) shifted binary image (c) and the difference image. The difference binary image is the layer of fat.

Figure 4.40(c) shows the difference between the non-shifted and the shifted binary images. Each voxel in the difference image was multiplied by 940 (i.e. real CT voxel value for fat) and then added to the original patient image, as shown in Figure 4.41.

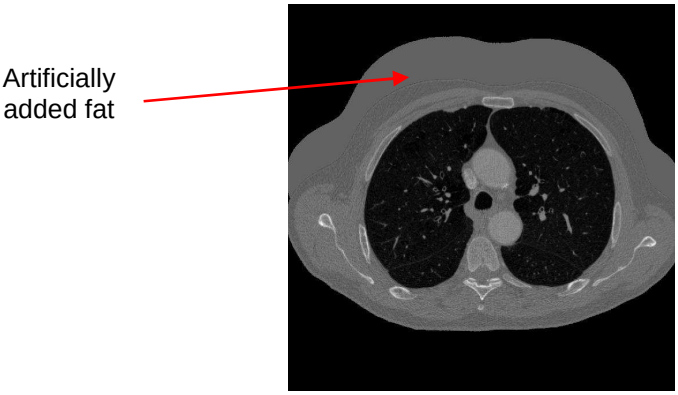

Figure 4.41: Fat added to the top of average patient.

The process described above was repeated for the bottom half of the patient and applied to each CT image slice in the data set. A comparison of an average and obese patient is shown in Figure 4.42.

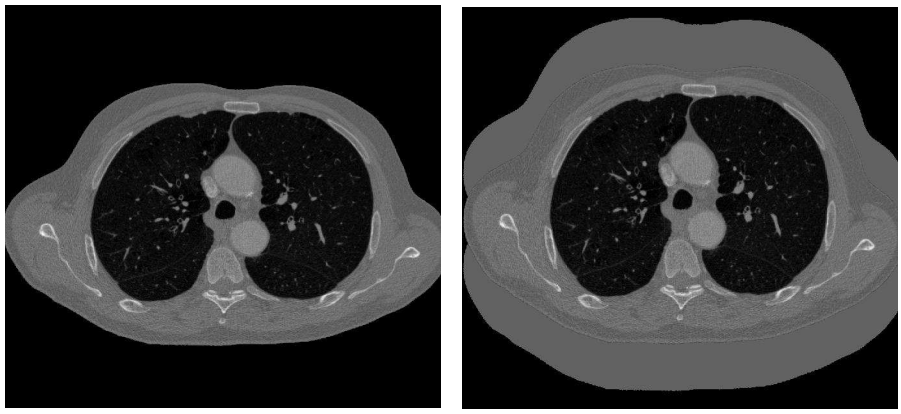

Figure 4.42: (a) Average patient and (b) patient with simulated fat.

Figure 4.42(b) clearly shows the layer of fat around the patient not present in the original image. Fat added artificially here is uniform (i.e. all voxel values are 940) but real fat would not have a single voxel value, and would appear ‘noisy’.

However, noise in CT images has little to no effect on the final DRR image characteristics (see Section 4.2.16) so the uniformity described here is expected to have negligible consequences on the final DRR.

4.2.14 Assessment of DRR image spatial resolution

CT images are inherently of lower resolution than CR images. The highest resolution possible with a DRR is CT slice thickness x CT voxel size, i.e. 0.8 mm x 0.34 mm. This is worse than that of a CR image, which has a pixel pitch of 0.17 mm. Hence, it was deemed necessary to assess the difference in resolution by measuring the modulation transfer function (MTF) of a DRR and CR image respectively using the technique reported by Samei *et al.* [101]. The MTF is widely used as the metric of choice for measurement of the resolution properties of radiographic systems [102].

A 50 mm square, 1 mm thick Tungsten edge test tool was placed on a CR cassette at a 3° angle with respect to the pixel ma trix, and was exposed with a tube potential of 70 kVp and sufficient tube current-time product (mAs) to ensure air kerma at the cassette was approximately 100 µGy. The image was processed as discussed in Section 4.2.1.2 and the MTF was measured along a section of the edge contained within a 60 x 60 pixel ROI, as shown in Figure 4.43 (blue square).

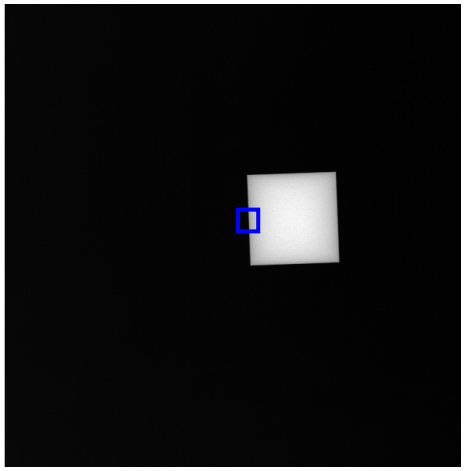

Figure 4.43: CR image of Tungsten edge tool.

The MTF of a DRR image was then measured in a similar fashion as described above, but with a Teflon edge phantom, originally designed to measure MTF of CT scanners, as shown in Figure 4.44.

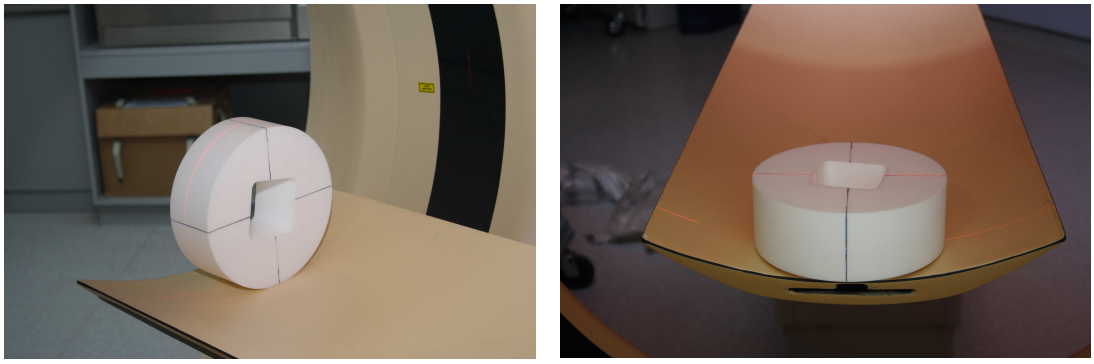

(a) (b)

Figure 4.44: (a) Teflon edge phantom to measure the DRR modulation transfer function, and (b) in the required orientation for scanning.

The CT MTF phantom was scanned in the orientation shown in Figure 4.44(b) using the scan parameters discussed in Section 4.2.4. A DRR was produced of the phantom, as shown in Figure 4.45, and the MTF measured from the reconstructed edge.

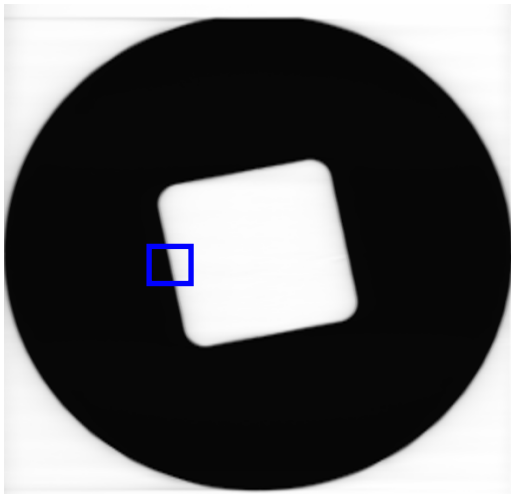

Figure 4.45: DRR of Teflon phantom. Area of MTF calculation is shown by the blue ROI.

The resulting CR and DRR MTFs were plotted on the same axes in Figure 4.46.

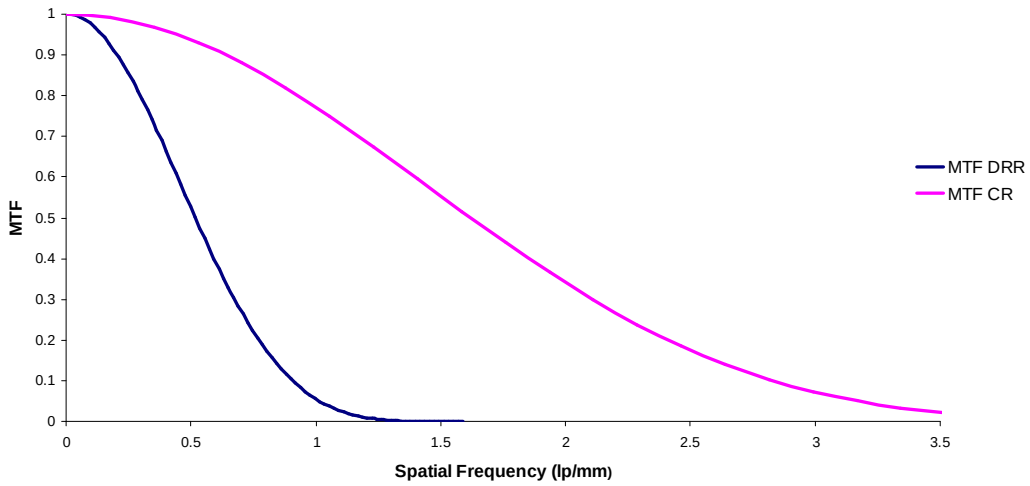

Figure 4.46: MTFs of a DRR image (pink line) and a CR image (blue line).

As can be seen quite clearly in Figure 4.46, the spatial resolution properties of a CR system are superior to that of a DRR. Therefore reconstructed object information will always be presented at a lower resolution, and a DRR image will never look exactly the same as a CR image of the same object. This may be a limiting factor (see Chapter 5).

4.2.15 Accuracy of DRR image resizing

As discussed in Section 4.2.6, each DRR is smaller than a real CR image. As such, the final DRR was resized to match the size of a CR image (2800 x 2300 pixels) using bicubic interpolation (the output pixel value is a weighted average of pixels in the nearest 4-by-4 neighborhood).

The accuracy of image resizing was tested by comparing signal to noise ratios (SNRs) and dynamic ranges of 5 random original and resized patient reconstructed DRRs. All SNRs measured in the lung, spine and diaphragm regions agreed to within 2%, minimum DRR pixel values were always the same, and maximum DRR pixel values agreed to within 3%. This is not surprising since the bicubic interpolation resizing method is designed to minimise differences in output from input. Although this adds a systematic error, it was felt that this was small enough to continue and see if it affected the subsequent validation results.

4.2.16 Assessment of noise in DRR images due to noise inherent within CT data

It is well understood that CT images contain noise due to a number of sources, primarily electronic, quantum and reconstruction filter. Quantum noise is a result of counting a finite number of random events (photons interacting with the detector) and is Poisson distributed [103]. However, noise in a CT image is rarely characterised by a Poisson distribution, as the reconstruction filter has a more significant influence on the final characteristics of the image [105], and it is usually characterized through measurement [104]. It was important to assess whether noise present in the CT data used for the computerised phantom in this

research had any effect on the final characteristics of the DRR image. As such, the type of noise in the CT data was measured using the data from the ROI analysis of the Gammex RMI phantom described in Section 4.2.4. A histogram of pixel values was obtained from each ROI to determine the noise probability density functions (PDF). Analysis of their form illustrates the type of noise in the image e.g. uniform, Guassian, or Poisson [105]. ‘Goodness of fit’ to certain distributions were analysed, and it was demonstrated that a Gaussian noise relationship was dominant in the CT images, as shown in Figure 4.47.

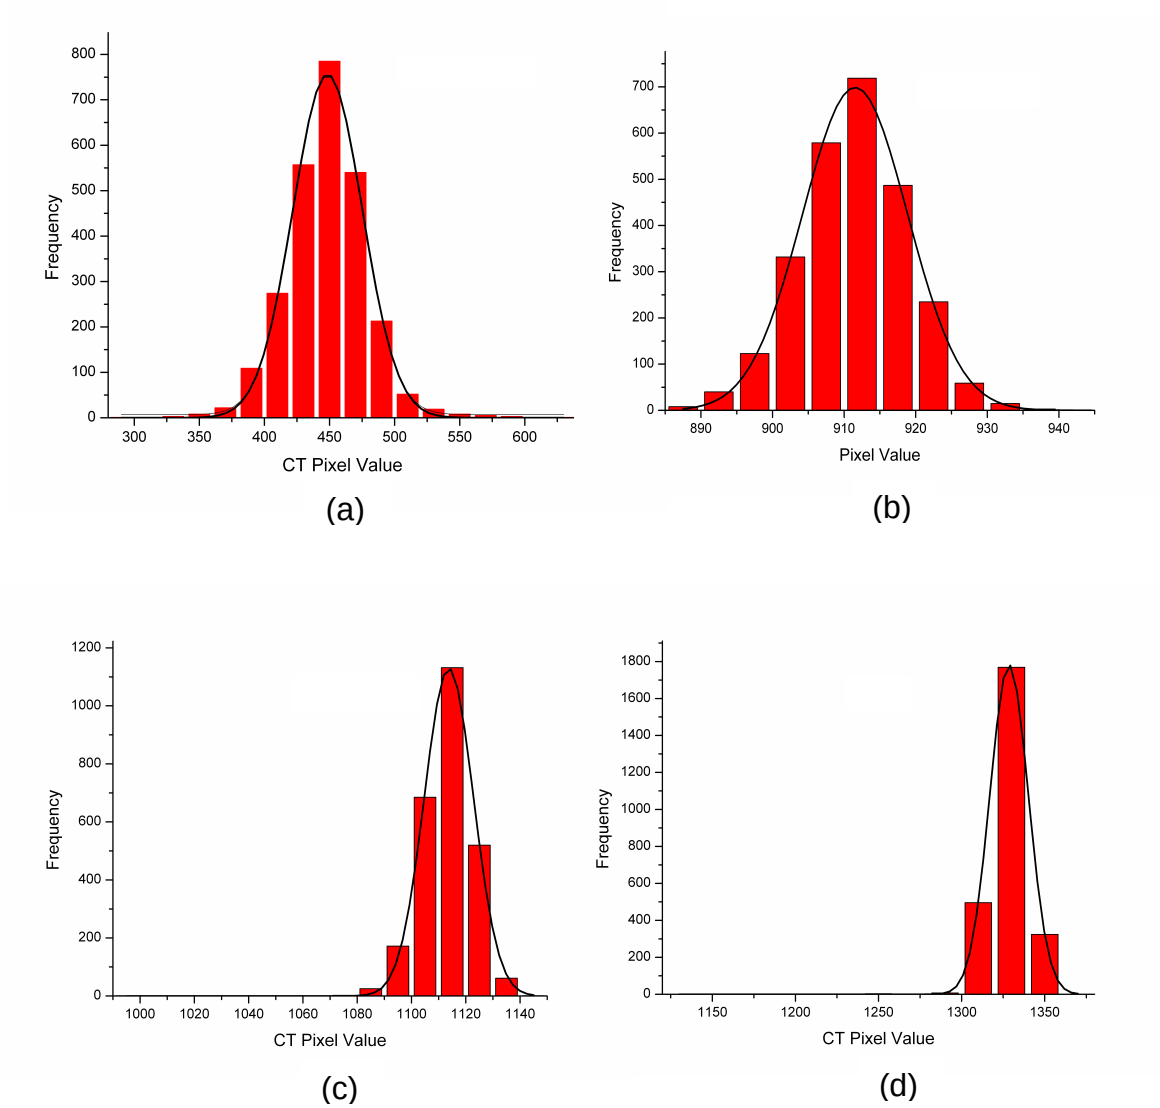

Figure 4.47: (a) PDF of lung insert fit with a Gaussian distribution, (b) PDF of adipose insert, (c) PDF of liver insert and (d), PDF of bone mineral insert. The goodness of fit are all $r^2 > 0.99$.

The PDFs of each insert clearly demonstrate a Gaussian distribution (all $r^2 \geq 0.99$). Although not shown here, all other tissue inserts were also found to be well described by a Gaussian noise distribution (all $r^2 \geq 0.97$), and as such indicates CT image noise is independent of voxel value. This is similar to that reported by Hilts *et al.* [104]. To assess whether this noise has any effect on the final characteristics of a DRR image, a Gaussian noise removing filter (mean adaptive filter) with varying kernel sizes of 0x0 (i.e. no filter applied), 7x7, 14x14

and 21x21 was applied to the CT data prior to DRR calculation (for ten random DRRs). Frequency dependent noise was added to the raw DRRs, and signal to noise (SNR) and dynamic range comparisons were made. CT images also exhibit noise due to scattered radiation to some extent, but the scanner used in this study utilises post-patient collimation which minimises scatter detected (nominally only 1-2% of the signal is due to scatter), so no correction for this was deemed necessary.

SNRs in the lung, spine and diaphragm regions of DRRs reconstructed without a noise removal filter, and with filters of size 7x7, 14x14 and 21x21 demonstrated less than 0.5% change from one another. There was no difference in the dynamic range of the images. This demonstrates that system (including quantum) noise added to the resulting DRR images dominates over any noise that manifests itself in the DRRs due to the presence of Gaussian noise in the CT data. This is probably due to the averaging and summing process of X-ray pencil beam ray casting (averaging and summing causes all the voxels intersected by each pencil beam to tend to their true value). Due to the minimal effect on image quality in the resulting DRRs, it was felt not necessary to apply a noise removal filter to the CT data prior to DRR calculation.

4.2.17 Phantom and patient DRR images

To ensure the DRR system was capable of producing images that represent the chest, CT data of RANDO and a random patient were obtained (Section 4.2.1.3). The processes described in this chapter were used to produce raw DRRs (i.e. images with no noise or scatter) reconstructed with tube potentials of 50 and 150kVp respectively, as shown in Figure 4.48.

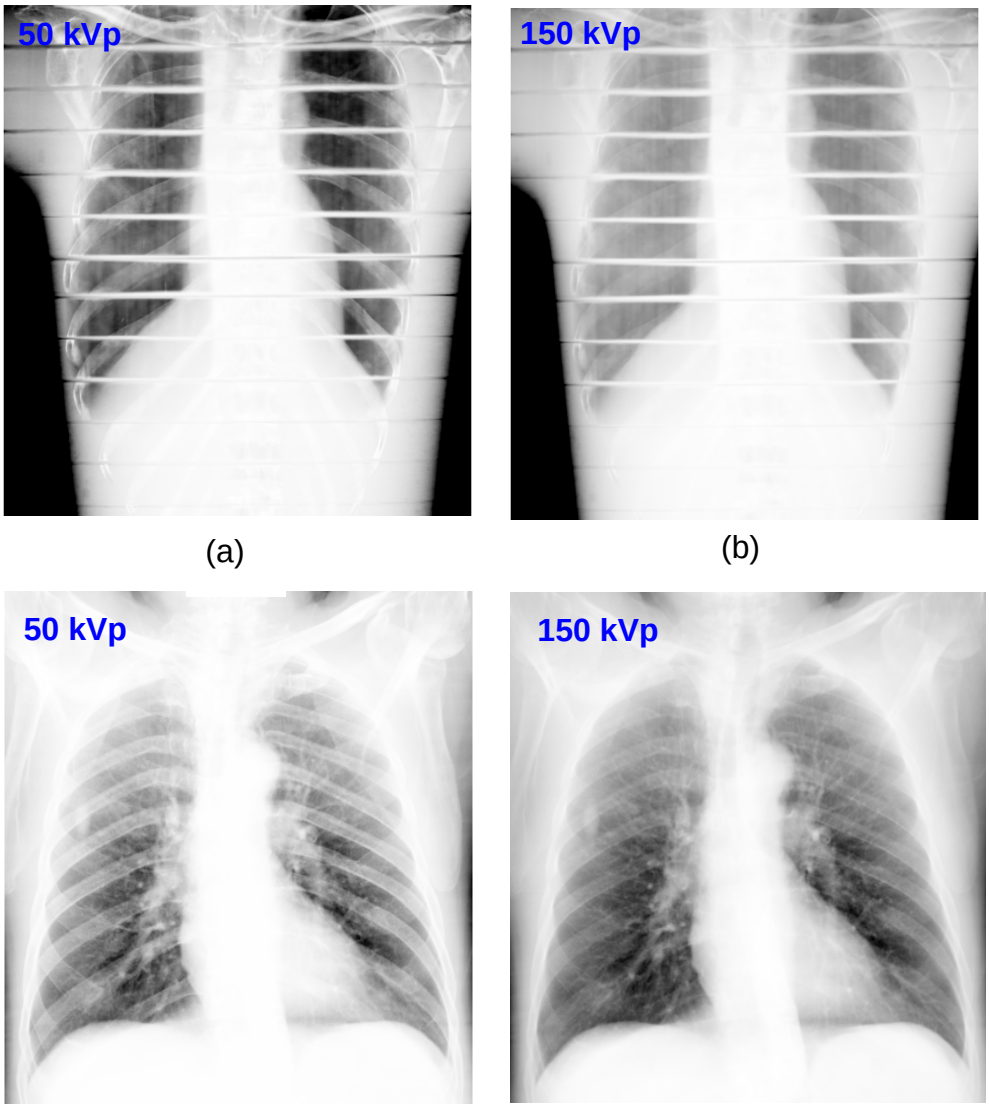

Figure 4.48: (a and b) DRR of RANDO and (c and d) DRR of a patient reconstructed with tube potentials of 50 kVp and 150 kVp. No scatter or noise has been added to the images.

The images in Figure 4.48 visually correlate with actual radiographs of RANDO and patients, which proves the DRR software is capable of reconstructing images that, in the very least, look as they should. However, thorough quantitative and qualitative validation of the images will be required, and this is discussed in Chapter 5. It is apparent from Figures 4.48 that contrast (especially in the lung and ribs) decreases as the tube potential increases. This is because there is a decrease in the differences between the linear attenuation coefficients of different human tissues with an increase in energy, because the photoelectric

cross-section varies with energy as approximately E^{-3}; this is why the photoelectric effect plays a dominant role in producing subject contrast at diagnostic energies. This demonstrates that the DRR software is capable of producing images that reflect the physics of X-ray attenuation.

Scatter and noise (Sections 4.2.10 and 4.2.11) were added to the raw DRRs shown in Figure 4.48 to ensure these processes worked adequately. The resulting images are displayed in Figure 4.49.

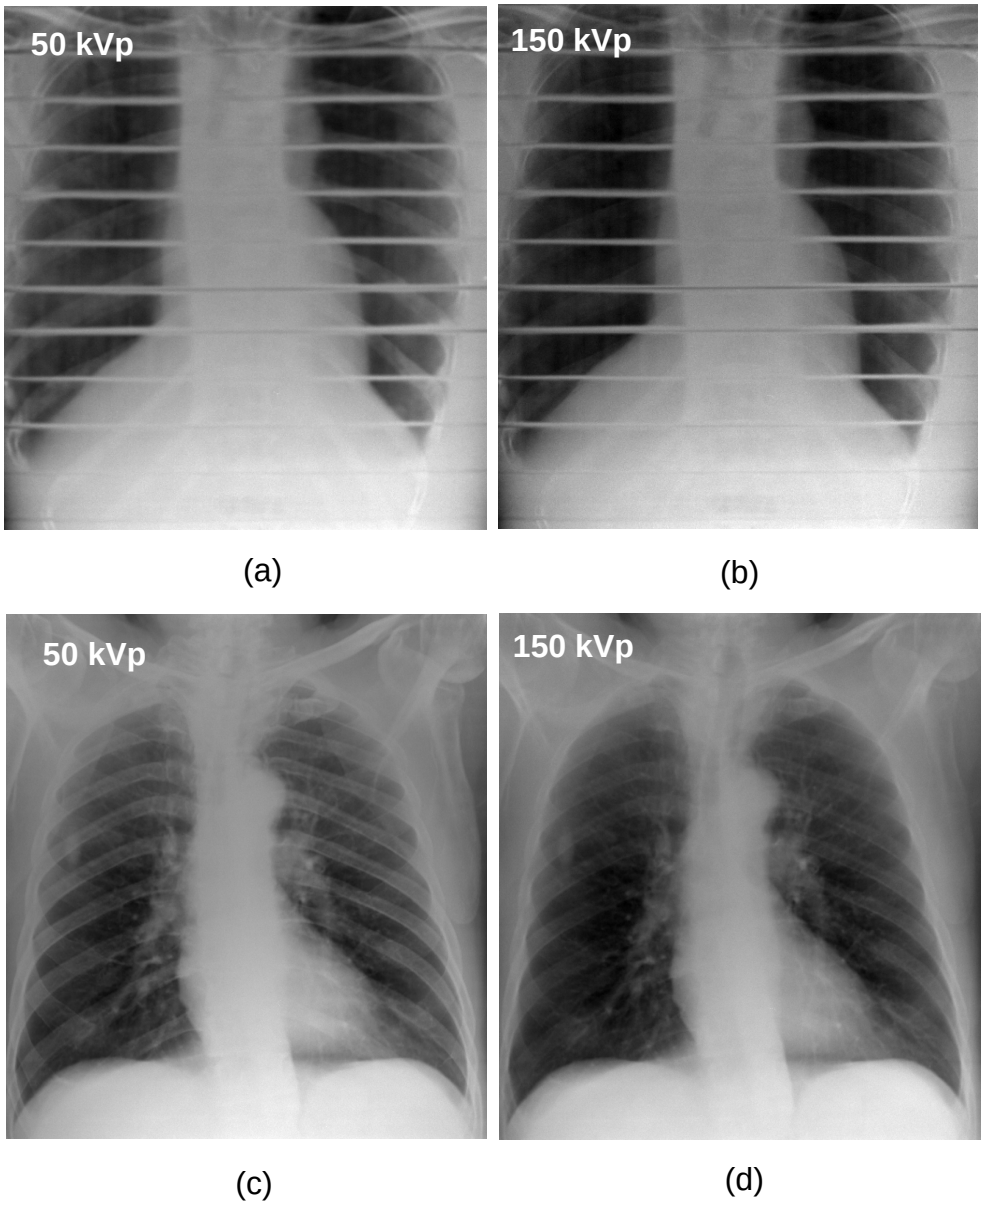

Figure 4.49: (a and b) DRR of RANDO and (c and d) DRR of a patient reconstructed with tube potentials of 50 kVp and 150 kVp. Scatter and noise has been added to the images.

Scatter and noise has been added successfully, as demonstrated in Figure 4.49 and as per Figure 4.48, the images visually correlate with real radiographs.

4.3 Overview of DRR calculation methodology

The following flow charts breaks down the methodology described in this chapter:

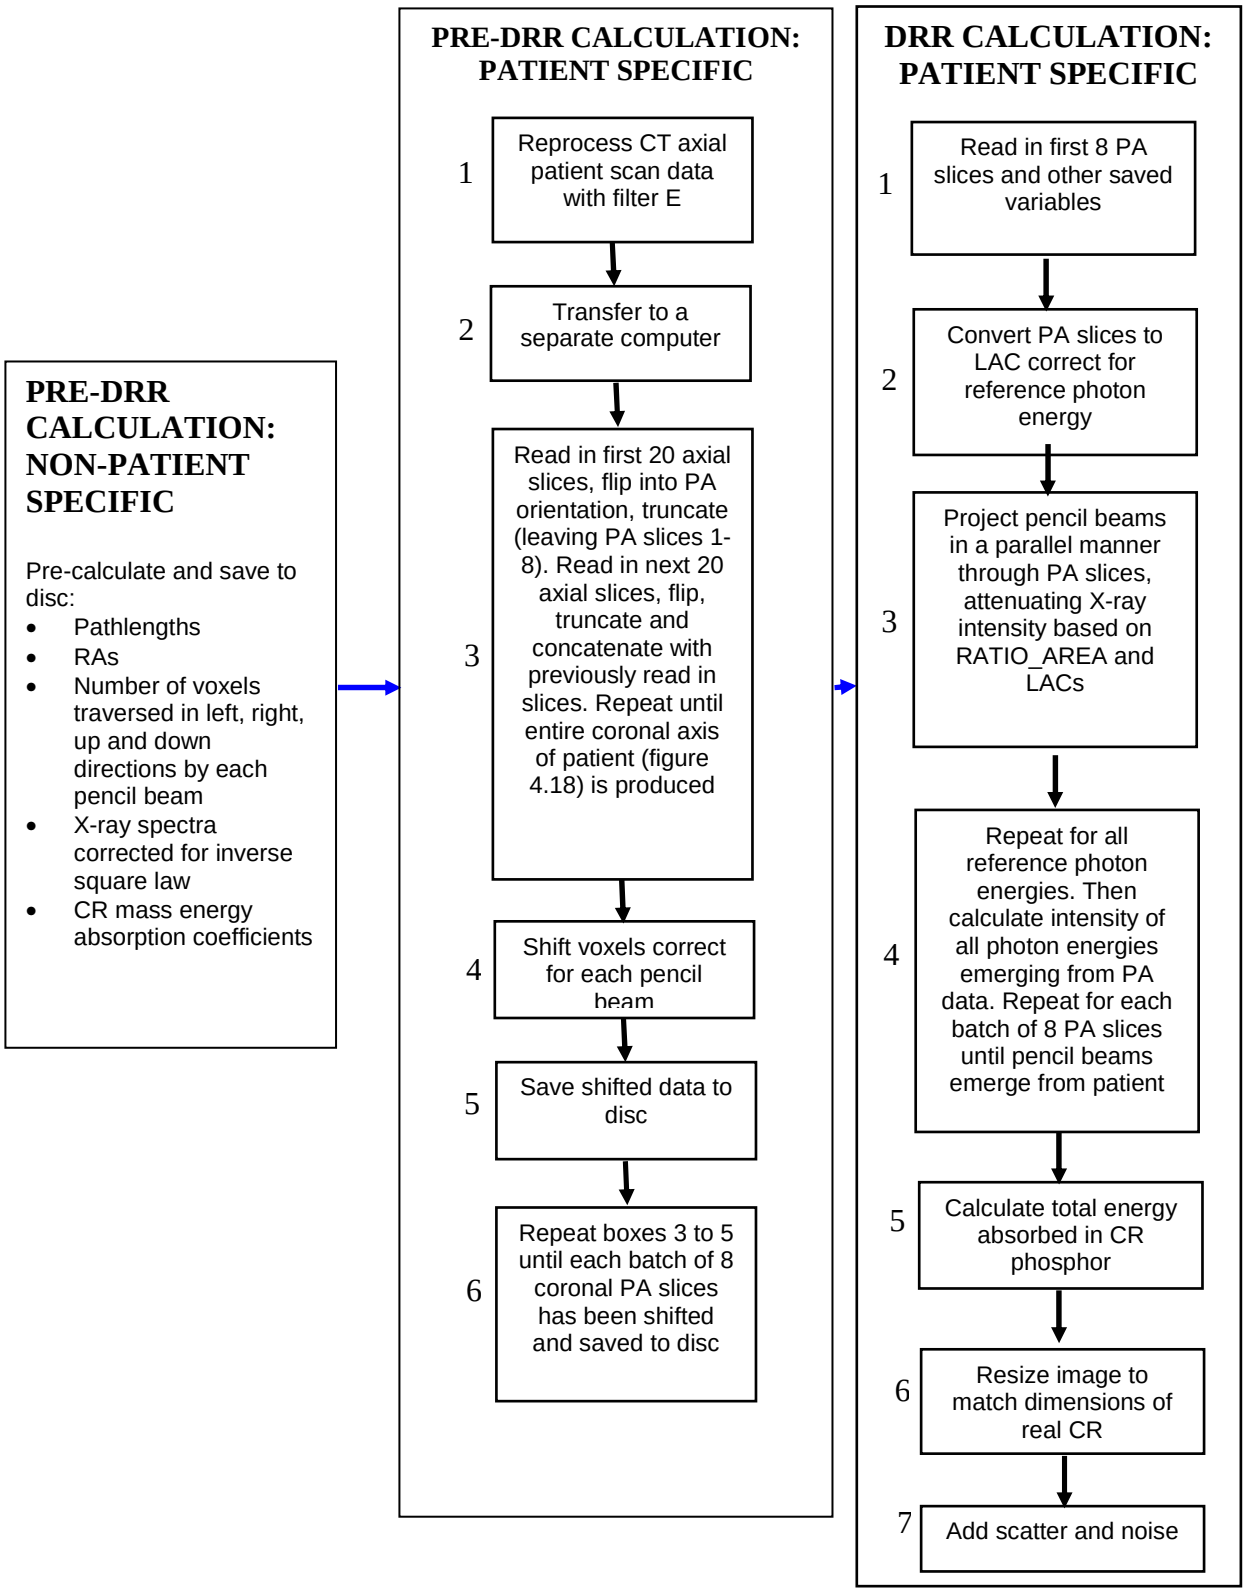

4.4 Conclusions

DRR simulated chest images with real clinical information will be a useful tool in diagnostic radiology for optimising chest radiographs without the need for repeat patient exposure. For this technique to become readily used, there must be a robust methodology that is reliable and repeatable. The work presented in this chapter describes such a method.

It is important that real X-ray, CR and CT systems are used to derive indices that can be utilised by the DRR algorithm to ensure reconstructed images mimic real radiographs as far as possible. A pencil beam ray casting method has been chosen to do this, and has been shown that images can indeed be produced that at the very least visually correlate with real chest radiographs of phantoms and patients. However, the spatial resolution properties of the reconstructed DRRs are inferior to real CR images due to the limitations of current CT systems, which means that some object information will never be faithfully reproduced. This may be a limiting factor.

CT noise is overwhelmingly of a Gaussian nature, but does not require removing as it has very little influence in the final image SNR characteristics (< 0.5% change). This is probably due to the averaging process of voxel sampling and the fact that CR frequency dependent noise dominates, once added to the raw DRR. Resizing DRR images to match the pixel density of a real CR image has little to no effect on DRR image characteristics.

Significant numbers of human thorax tissues have been modelled using the Gammex RMI phantom to enable the dynamic range of the resulting DRR

images to be as close as possible to that of real images, but this will require numerical validation.

Somewhat disappointing is the length of time taken for the software to reconstruct a single image. The process is very computationally intensive. Various methods have been described in this chapter in an attempt to overcome this problem, such as pre-calculating ray lengths and pencil beam areas, binning photon energies and voxel shifting the CT data prior to DRR calculation. These methods speed up the compute time from tens of hours to approximately 45-90 minutes per image. However, this is still not real time image production and one must be cautious about final DRR images due to these simplifications. In future, it may be possible to compile the Matlab code which would speed up compute time, and this would allow the use of the software on any computer. Also, the code may be re-written to allow execution on modern graphics cards which would be beneficial and lead to significantly shorter compute times. Nevertheless, even without code compilation or the use of graphics cards, it is likely computing power and processing speed will increase, and therefore it may be possible to generate DRR images in real time using the methods described in this chapter.

Scatter and frequency dependent system noise must not be left out of the images, as these influence real CR image characteristics. The methodology described in this chapter has measured scatter and noise on a real CR system, and has successfully added these to the raw DRR.

Lung nodules have also been successfully added to the CT data, and these are visible in the final DRR. Therefore, assuming adequate validation, these can be used for optimising CR chest imaging when nodules are indicated. Fat has been

artificially added to average patient CT data to allow DRR reconstruction of obese patients. This adds to the flexibility and usefulness of the software.

Validation of the methodology outlined in this chapter will be vital to ensure this DRR computing model adequately simulates real CR chest radiographs. This is discussed at length in the next chapter.

Chapter 5: Quantitative and qualitative validation of the digitally reconstructed radiograph algorithm

5.1 Introduction

This chapter describes the quantitative and qualitative validation of the algorithm developed to produce digitally reconstructed radiographs (DRRs). Quantitative assessment with the average and obese RANDO phantom is addressed initially, followed by validation with average and obese patient images acquired on a clinical CR system. It should be noted that most of the quantitative validation was done by comparing the arithmetic means of the relevant data; although not stated in the results, the geometric means differed by ≤ 1% c.f. arithmetic means. Qualitative assessment by expert image evaluators is then discussed, as well as the limitations of the DRR algorithm.

5.2 Validation with the average RANDO phantom

5.2.1 RANDO images with no scatter rejection

Initial validation was carried out with real CR and simulated DRR images of the RANDO phantom representing the average sized patient. Phantom images were acquired on the clinical CR system as described in Chapter 4 Section 4.2.1.2. Figure 5.1 shows images acquired at tube potentials of 60, 90 and 150 kVp, as well as their DRR counterparts.

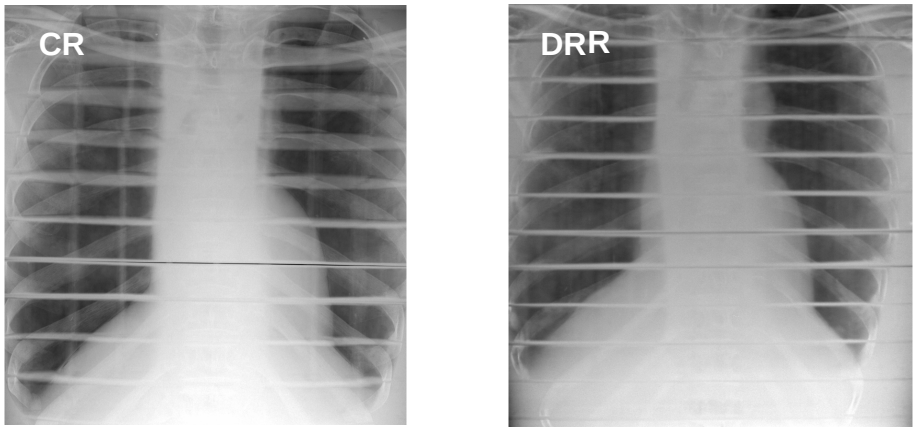

(a)

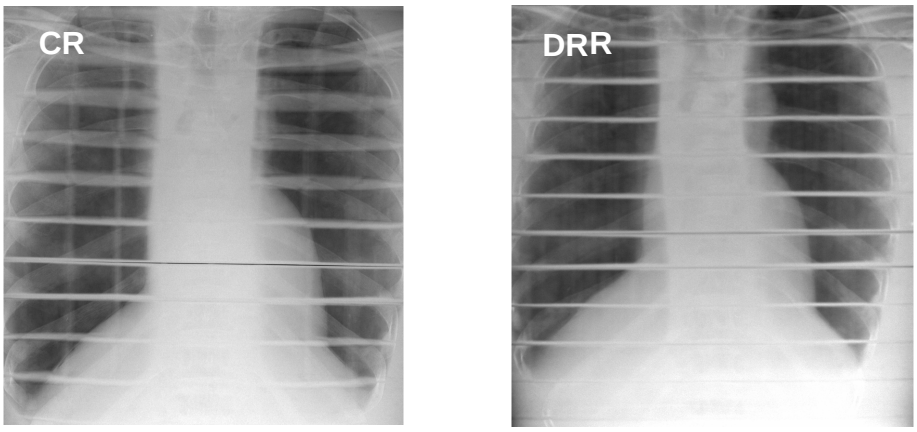

(b)

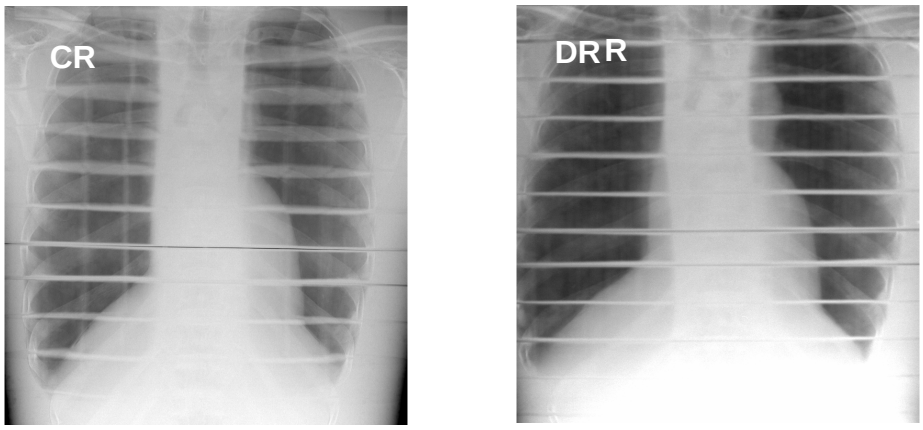

(c)

Figure 5.1: (a) CR & DRR image of the RANDO phantom, 60 kVp 10 mAs, (b), CR & DRR of the RANDO phantom, 90 kVp 2mAs, (c) CR & DRR image of the RANDO phantom, 150 kVp 0.5 mAs.

As can be seen in Figure 5.1, all DRR images correlate visually with the clinically acquired ones. Quantitative validation was carried out by plotting histograms of pixel values, and calculation of signal-to-noise ratios (SNR) and

tissue-to-rib ratios (TRR). It is important to compare SNRs, given the level of both signal and noise in a digital image affect the visualization of normal structures and pathology. The TRR is a metric that compares the mean region of interest (ROI) pixel value of soft tissue to that of rib. Ribs cover a large area of a chest radiograph and can interfere with the detection of soft tissue lesions by distracting the reporting Radiologist. It is therefore important that a good agreement should exist between calculated DRR and acquired CR images. SNRs were measured in the lung, spine and diaphragm regions of each image (both CR and DRR), whilst TRRs were measured in the lateral pulmonary region, as shown in Figure 5.2.

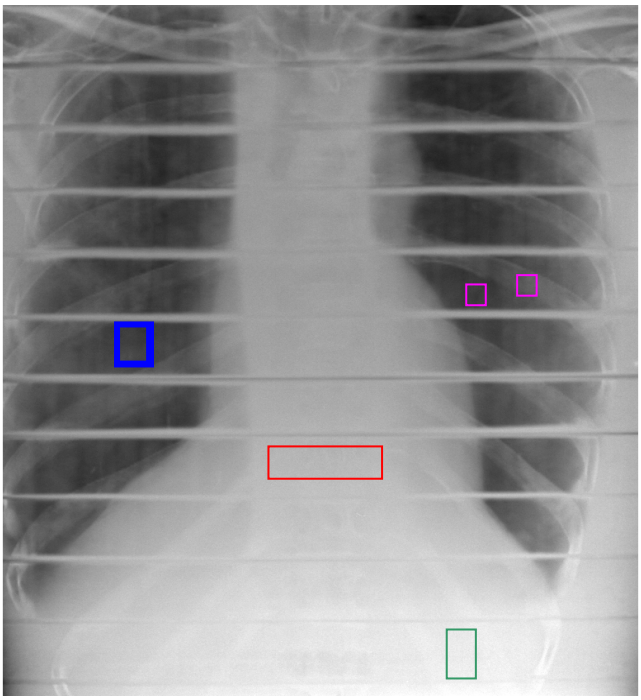

Figure 5.2: DRR of RANDO showing ROI positions used for SNR and TRR measurements. Lung, spine and diaphragm ROIs are depicted as blue, red and green respectively. The two pink ROIs depict the position for TRR measurement.

The size of the ROIs used in Figure 5.2 were dictated by the anatomical region over which they were positioned. Figures 5.3 and 5.4 compare clinical (CR) and simulated (DRR) histograms of pixel values at 60 kVp, 10 mAs and 150 kVp, 0.5 mAs respectively.

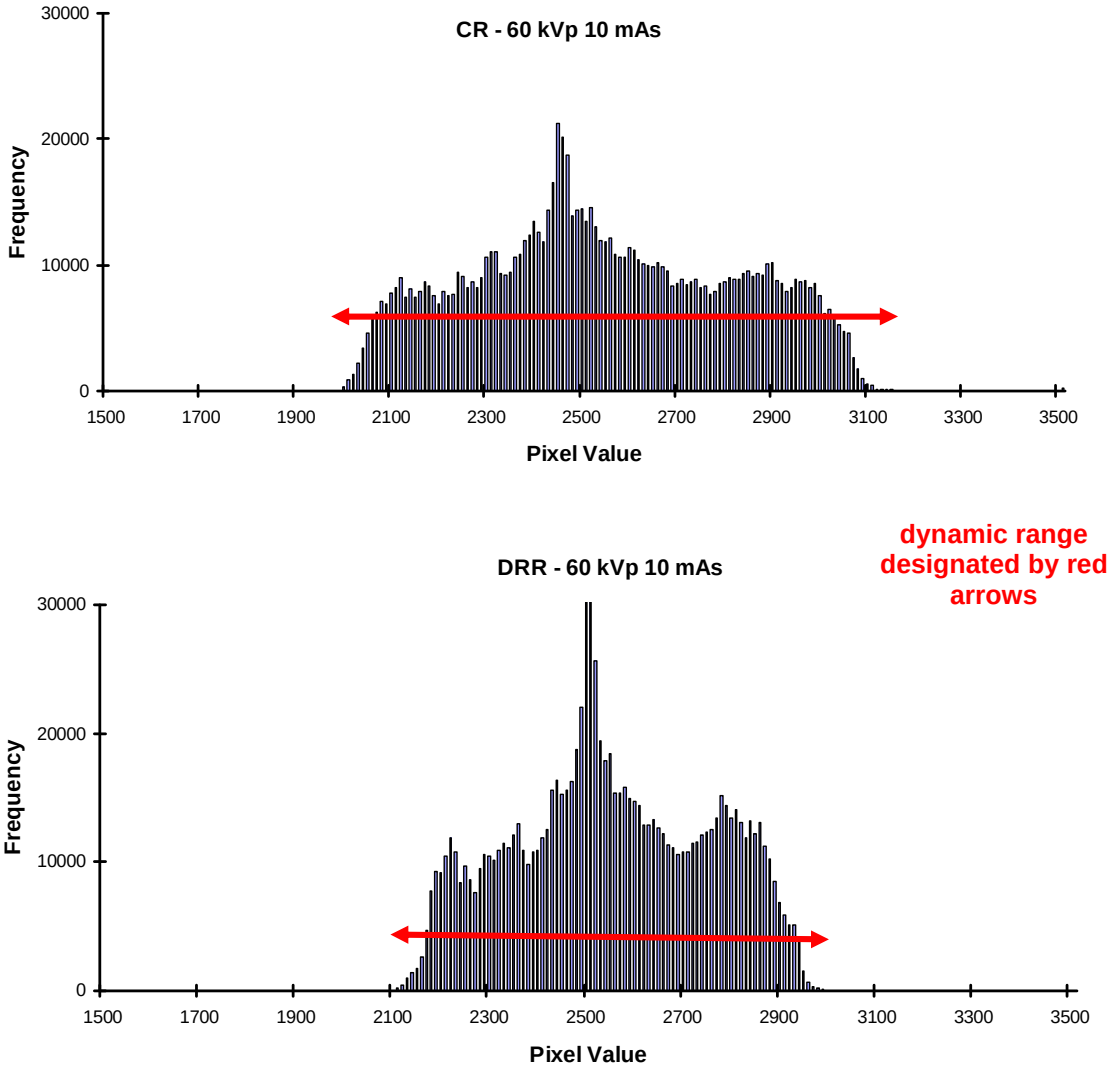

Figure 5.3: CR histogram of pixel values of the average RANDO phantom, 60 kVp & 10 mAs and the corresponding DRR histogram of pixel values of the average RANDO phantom, 60 kVp & 10 mAs. The red arrows indicate the dynamic range of each image.

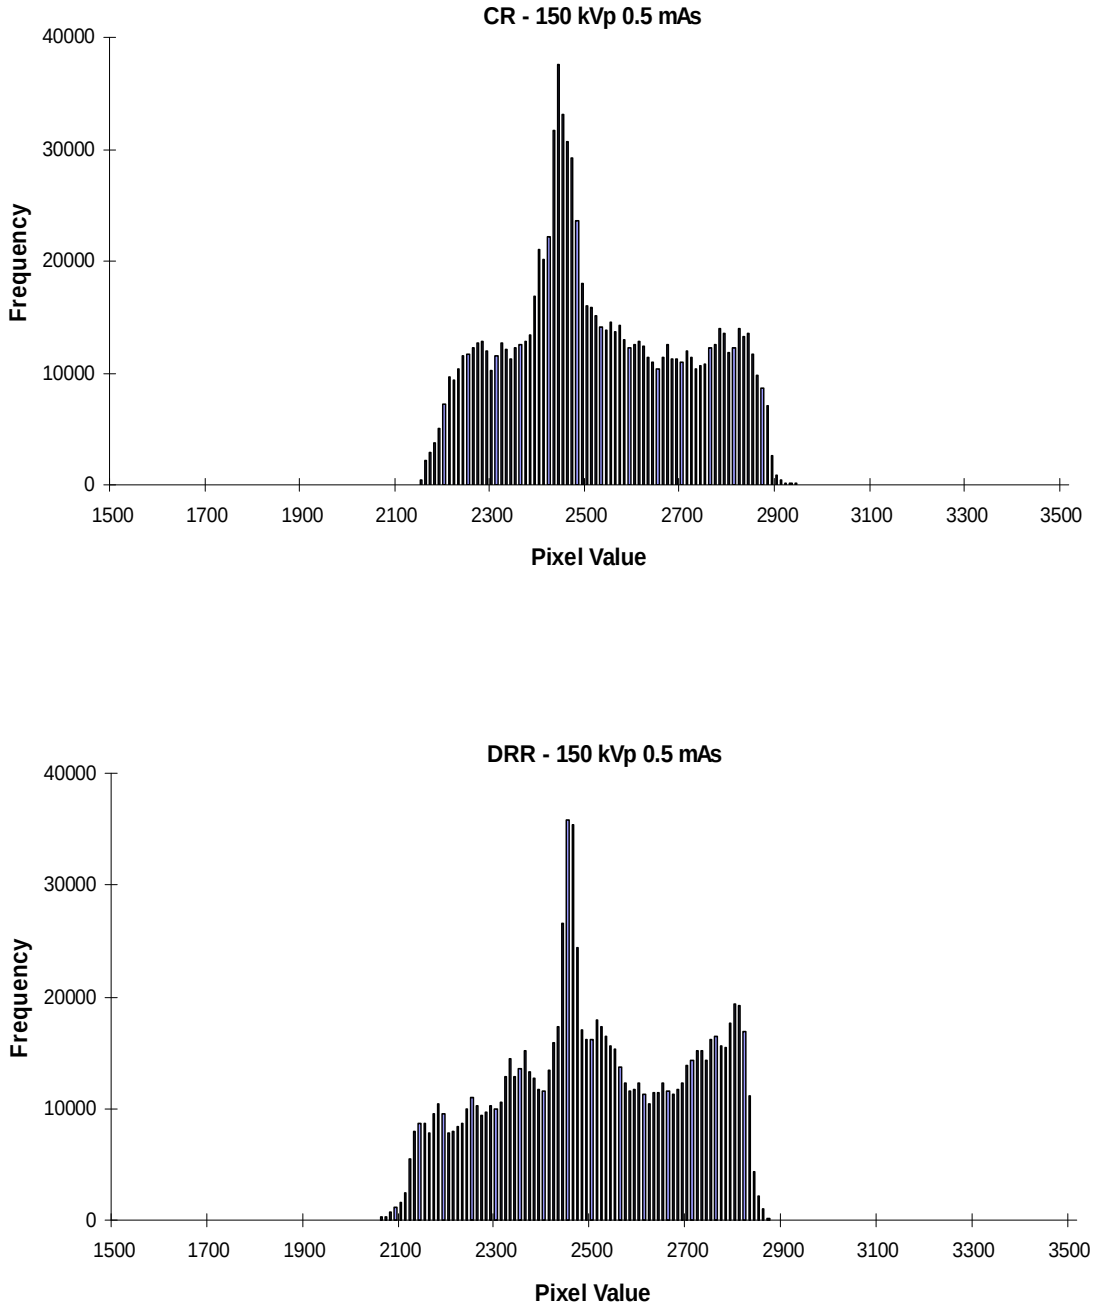

Figure 5.4: CR histogram of pixel values of the average RANDO phantom, 150 kVp & 0.5 mAs and the corresponding DRR histogram of pixel values of the average RANDO phantom, 150 kVp & 0.5 mAs.

It can be seen in Figures 5.3 and 5.4 that histograms produced from simulated DRR images are similar in shape to those produced from real CR images. However, both DRR histograms have a reduced dynamic range (for clarity the dynamic ranges of the images in Figure 5.3 are shown by the red arrows)

relative to the CR histograms. One can argue this to be expected because the Gammex phantom was used to derive equations to convert CT voxel values to linear attenuation coefficient (LAC) as discussed in Chapter 4; whilst this phantom contains many tissue substitutes, there will be fewer than those encountered in the human chest, which results in a coarser LAC transformation than desired. The virtual patient is also a voxelated computerised phantom with inherent loss of data (Sandborg *et al.* [43] reported similar findings with their Monte Carlo computer model). Finally, the voxel size of the virtual patient (0.34 x 0.34 x 0.8 mm) is probably larger than some of the smaller structures within the body resulting in tissues not being present in the virtual patient that would be present in a real one.

The mean (± 2 SD) of the minimum and maximum histogram values (a measure of the dynamic range of an image) for the CR images were 2075 ± 162 and 3000 ± 164 respectively. For the DRR images, these values were 2133 ± 86 and 2926 ± 120. This is probably acceptable, given the large fluctuation in minimum and maximum pixel value between real patients (see Section 5.5). It should be noted that histograms of all other tube potentials tested are not shown (for conciseness), but they all follow the same features as Figures 5.3 and 5.4. Figure 5.5 shows real CR and simulated DRR images together with their histograms acquired and reconstructed at 90 kVp using 1 and 4 mAs respectively.

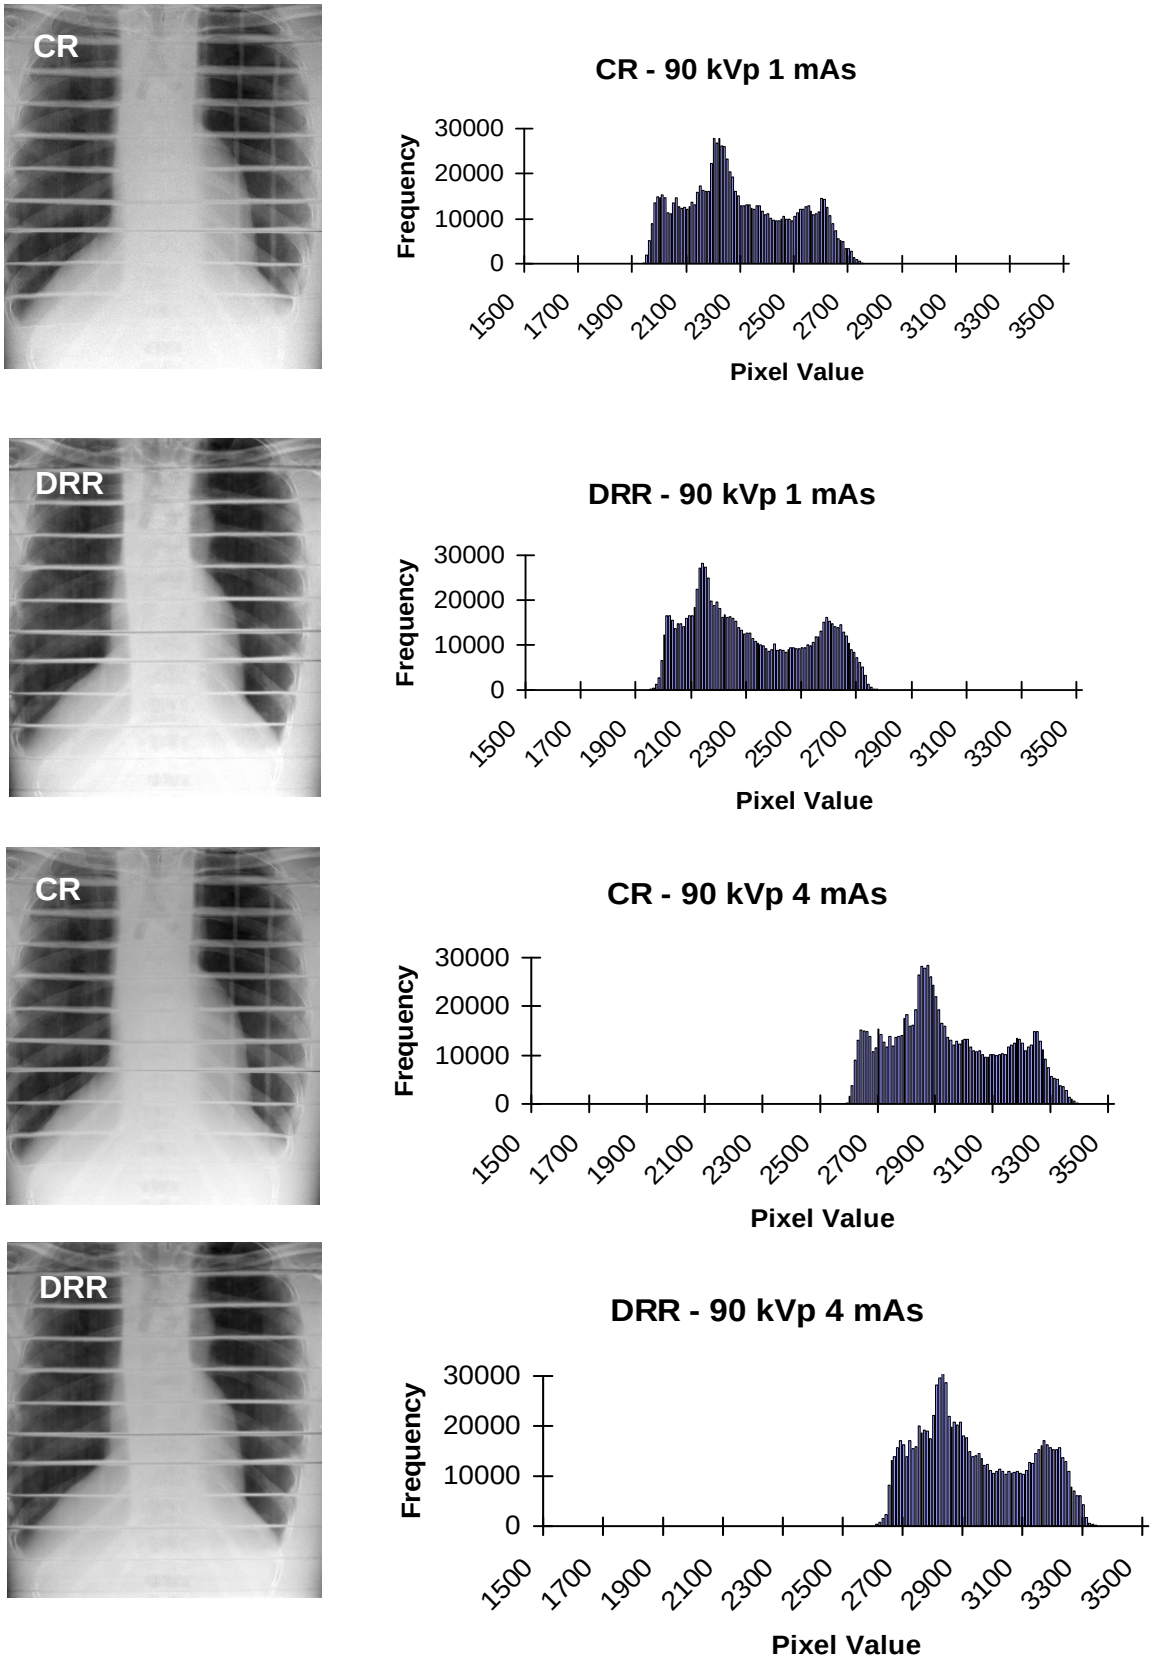

Figure 5.5: Real CR and simulated DRR images and histograms acquired and simulated with exposure factors 90 kVp, 1 and 4 mAs respectively.

Figure 5.5 demonstrates that histograms produced from simulated images are similar in shape to those produced from real CR images, and although their

dynamic range is slightly smaller (as discussed), they are shifted to the correct positions on the pixel value axis. This is encouraging, as it demonstrates the algorithm correctly increases and decreases pixel values according to increased/decreased exposure to the detector; it can therefore be used for dose escalation and reduction studies, assuming levels of noise are also correct.

Table 5.1 compares signal-to-noise ratios (SNRs) of real CR images to that of DRR images in the lung, spine and diaphragm. CR images were acquired at each tube potential on the clinical X-ray system with sufficient mAs to produce lgM values of 2.00 ± 0.05. DRR images were reconstructed with the same exposure parameters.

SNR (all ± 30)						
Tube Potential (kVp)	**DRR - Lung**	**CR - Lung**	**DRR - Spine**	**CR - Spine**	**DRR - Diap**	**CR - Diap**
50	133.7	146.7	95.4	88.7	60.5	54.4
60	155.7	145.9	67.2	68.9	66.8	72.6
70	156.3	154.6	89.8	86.5	63.4	65.4
81	137.3	153.0	118.3	111.5	58.0	56.5
90	155.8	151.5	87.7	83.8	60.7	65.8
102	163.8	159.1	87.1	87.8	60.2	63.9
109	171.1	164.7	91.7	85.2	64.6	67.8
125	172.4	165.5	91.7	88.8	62.7	66.4
133	166.6	158.6	85.3	85.3	57.7	60.9
141	180.8	177.6	94.8	102.6	65.0	71.2
150	186.9	197.8	81.0	89.9	59.0	63.6

Table 5.1: Comparison of SNRs measured in the lung, spine and diaphragm regions in DRR and real CR images.

It is clear from Table 5.1 that all DRR calculated and CR measured SNR values are in good agreement (within error), the maximum deviation being 11% (mean = 5.4%). It is likely the main source of error during DRR reconstruction is the addition of frequency dependant noise, as the method used has its limitations. As discussed in Chapter 4, Section 4.2.11, the Noise Power Spectrum (NPS) does not give a complete description of the noise properties of the system. However, as described by Bath *et al.* [49] this is not significant at dose levels used clinically. Nevertheless, limitations are discussed in depth in Section 5.7.

The SNR error value of ± 30 shown in Table 5.1 was derived by shifting the ROI used to calculate each SNR left/right and up/down a distance equal to half their size (at least 30 pixels). In real CR images, there is a movement of ROIs from one image to the next, as the phosphor is typically returned to the cassette in a slightly different position as previous. This results in a movement of approximately 20 pixels in the subsequent image, so the error here will be an overestimate. The shifted SNR was compared to the 'normal' SNR and the maximum difference encountered was 30.

The effect of dose variation on SNR in the lung, spine and diaphragm at each tube potential was investigated by reconstructing and acquiring images at half and double the tube current-time product (mAs) of that required for a lgM of 2.00. The maximum percentage difference in SNR measurements for half mAs images (DRR vs CR) was 12.1% in the lung, 12.9% in the spine and 15.4% in the diaphragm, and for double mAs images 9.9% in the lung, 13.4% in the spine and 10.2% in the diaphragm. The mean SNR (averaged over all tube potentials) for half and double mAs images are shown in Table 5.2.

Region	Mean SNR (DRR)	Mean SNR (CR)	Mean SNR (DRR)	Mean SNR (CR)
	Half mAs images		Double mAs images	
Lung	138.8 ± 15.9	135.8 ± 13.2	199.3 ± 12.9	207.5 ± 10.4
Spine	75.2 ± 9.0	80.0 ± 8.5	114.7 ± 12.9	117.1 ± 17.0
Diaphragm	57.8 ± 4.1	61.9 ± 7.7	83.9 ± 10.0	86.3 ± 9.4

Table 5.2: Mean SNR values in each chest region for half and double mAs reconstructed (DRR and acquired (CR) images). All errors are 1 standard deviation.

The results in Table 5.2 are in good agreement and demonstrate the algorithm can reproduce the level of noise seen in real CR images at clinical dose levels.

Figure 5.6 demonstrates a visual difference in relative noise levels in the heart region for DRR images reconstructed with 60 kVp, 5 and 20 mAs respectively (i.e. half and double typical clinical doses); the former (Figure 5.6(a)) clearly displays poorer SNR in the heart (area within the red ROI) compared to the latter (Figure 5.6(b)).

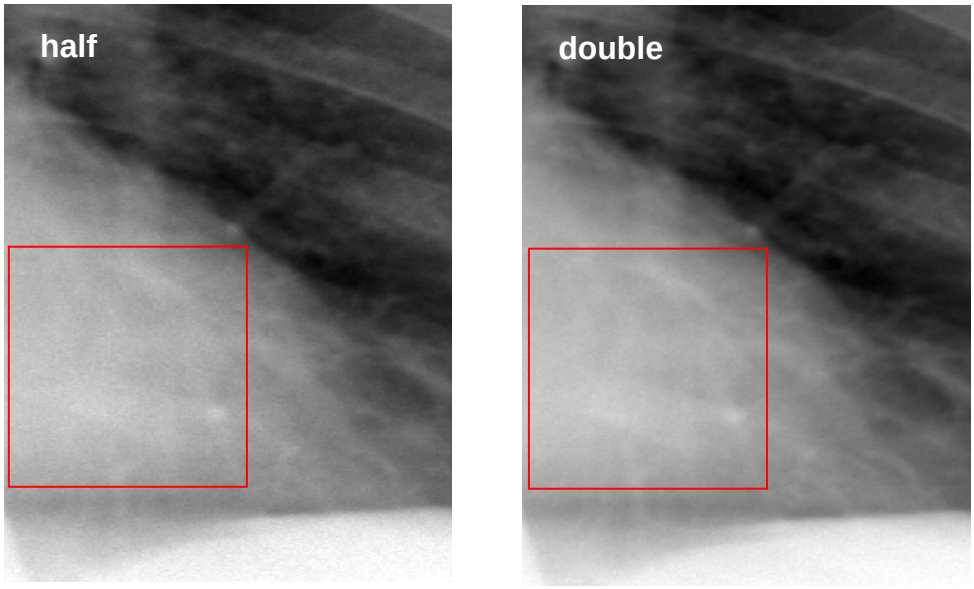

(a) (b)

Figure 5.6: DRR images (heart region) reconstructed with (a) half a typical mAs and (b) double typical mAs. The SNR is poorer in the red ROI of (a) compared to that of (b).

Table 5.3 shows TRRs calculated and measured in DRR and CR images respectively, and clearly demonstrates a good level of agreement between the two. The maximum difference in measured (CR) and calculated (DRR) is 0.6%. It can also be seen that TRR decreases with increasing tube potential; this is probably due to the ribs attenuating a higher percentage of incident photons at lower potentials than soft tissue, thus increasing the TRR.

Tissue to Rib Ratio (all ± 0.005)		
Tube Potential (kVp)	**TRR - DRR**	**TRR - CR**
50	1.041	1.042
60	1.039	1.036
70	1.039	1.034
81	1.029	1.033
90	1.035	1.032
102	1.033	1.027
109	1.029	1.029
125	1.028	1.026
133	1.029	1.025
141	1.026	1.026
150	1.026	1.025

Table 5.3: Comparison of Tissue to Rib Ratios measured in DRR and real CR images.

5.2.2 RANDO images reconstructed with an anti-scatter grid

Validation of RANDO images that were reconstructed with an anti-scatter grid modelled in the algorithm was carried out using the same methods as those described above. Figure 5.7(a) shows a DRR image of RANDO without any scatter rejection compared with a DRR image of RANDO with scatter rejection (Figure 5.7(b)), and with a CR image acquired with the anti-scatter grid in operation (Figure 5.7(c)). Figure 5.7(a) was reconstructed at 60 kVp/10 mAs, and Figures 5.7(b) and 5.7(c) were reconstructed and acquired with exposure factors 60 kVp/40 mAs, respectively. The increase in mAs was required to

overcome the effect of X-ray absorption by the grid, and to achieve a lgM value of 2.00 ± 0.05.

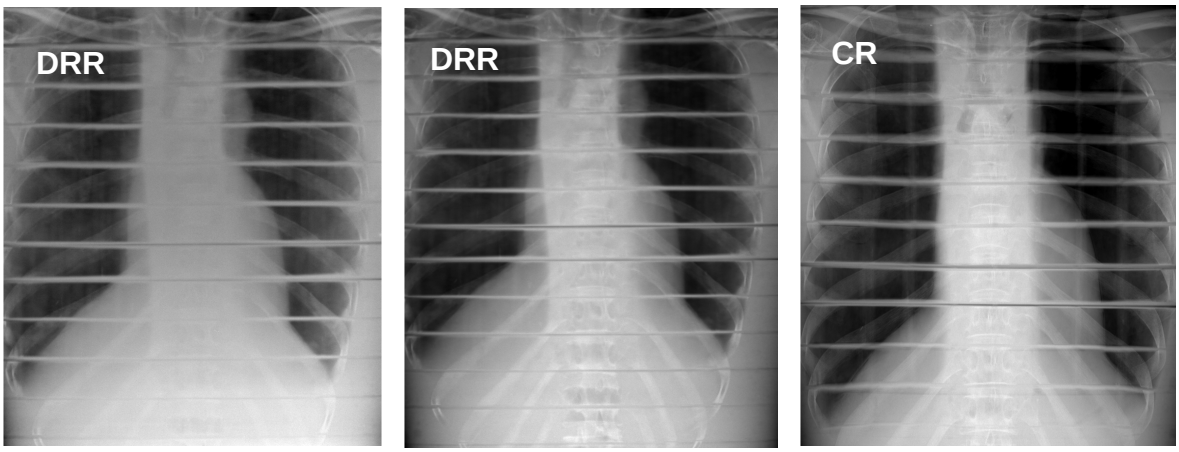

(a) (b) (c)

Figure 5.7: (a) DRR image of RANDO reconstructed without an anti-scatter grid, (b) DRR image of RANDO reconstructed with a scatter grid, and (c) CR image of RANDO acquired with the scatter grid energised (DRR and CR images reconstructed and acquired with tube potential 60 kVp respectively.

The DRR image of RANDO reconstructed with an anti-scatter grid (Figure 5.7(b)) clearly demonstrates improved detail in the spine and diaphragm regions when compared to the DRR without scatter rejection (Figure 5.7(a)). This is of course expected and matches the appearance of the CR image (Figure 5.7(c)).

Figure 5.8 compares clinical (CR) and simulated (DRR) histograms of pixel values at tube potential 125 kVp. This demonstrates that histogram shapes are broadly similar, but the dynamic range of DRR histogram is smaller than that of the CR. For all energies (tube potentials), the mean (± 2 SD) of the minimum and maximum histogram values for the CR images were 1773 ± 203 and 3230 ± 216 respectively. For the DRR images, these values were 1807 ± 260 and 3236 ± 252. As per the average RANDO results, although the dynamic range of the DRR images is smaller than the CR images, there is little difference, given the standard deviations of the means. It should be noted that histograms of all

other tube potentials tested are not shown (for conciseness), but they all have the same features as Figure 5.8.

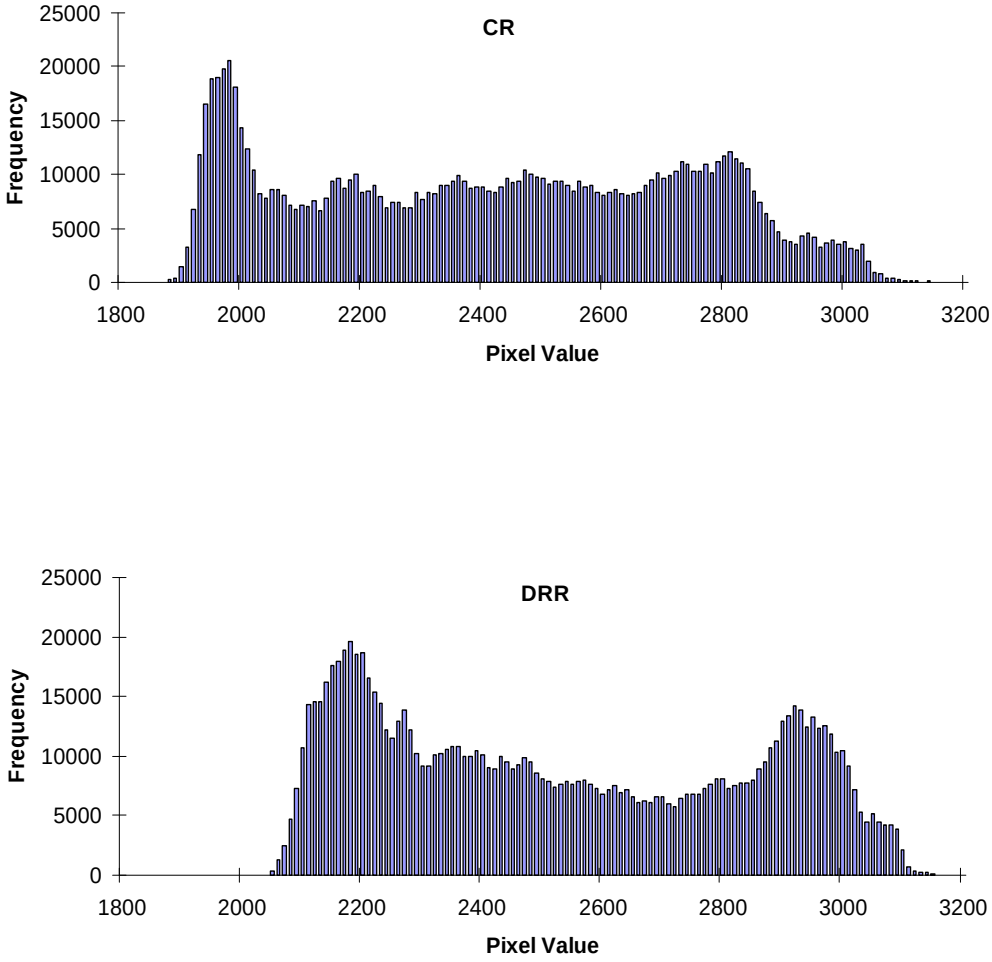

Figure 5.8: Histogram of pixel values of a DRR image, and the corresponding histogram for the CR image.

The maximum percentage difference in SNR measurements (DRR vs CR) was 5.7% in the lung, 7.6% in the spine and 8.6% in the diaphragm. The mean SNR (averaged over all tube potentials) are shown in Table 5.4. These show excellent agreement, demonstrating the DRR algorithm can adequately simulate levels of signal and noise across all diagnostic tube potentials. Furthermore, the maximum difference in TRRs was 0.50%, with a mean difference of 0.24%; again, the results show very good agreement.

Region	Mean SNR (DRR)	Mean SNR (CR)
Lung	185.6 ± 22.8	187.9 ± 23.3
Spine	69.5 ± 25.0	67.9 ± 24.7
Diaphragm	77.4 ± 13.5	77.1 ± 16.0

Table 5.4: Mean SNR values in each chest region of RANDO. All errors are 1 standard deviation of the mean.

All results obtained in this section demonstrate the model can produce images reconstructed with scatter rejection using an anti-scatter grid with adequate confidence.

5.2.3 RANDO images reconstructed with an air gap technique

Validation of RANDO images that were reconstructed with an air gap technique modelled in the algorithm was carried out using the same methods as those described throughout this chapter. Figure 5.9 shows DRR and CR images of RANDO reconstructed and acquired with the air gap technique respectively (both 50 kVp).

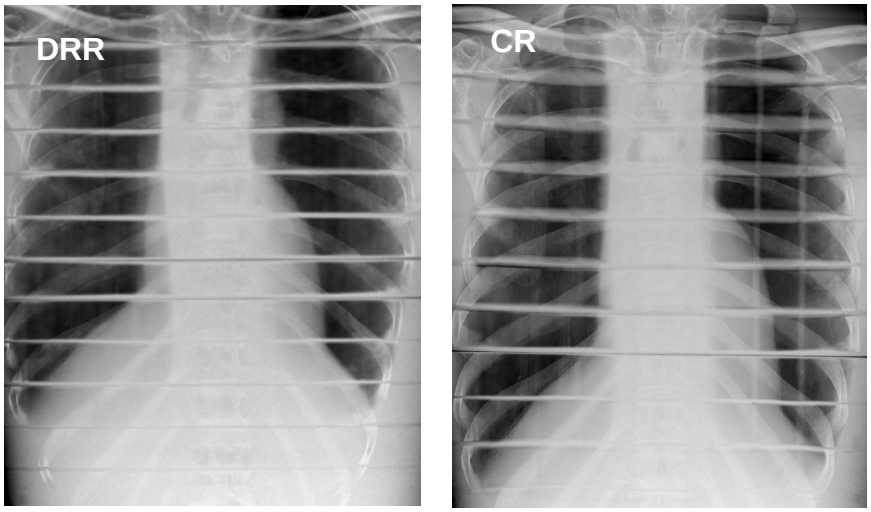

(a) (b)

Figure 5.9: (a) DRR image of RANDO reconstructed with an air gap technique, and (b) a CR image of RANDO acquired with the air gap.

The DRR image of RANDO reconstructed with an air gap technique (Figure 5.9(a)) demonstrates improved detail in the spine and diaphragm regions (although not as detailed as Figure 5.7(b)). This is expected and matches the appearance of the CR image (Figure 5.9(b)).

Figures 5.10 compares clinical (CR) and simulated (DRR) histograms of pixel values at tube potential 50 kVp. This demonstrates that the histogram shapes are very similar, and the dynamic ranges are almost identical. The mean (± 2 SD) of the minimum and maximum histogram values for the CR images were 2131 ± 150 and 3075 ± 206 respectively. For the DRR images, these values were 2115 ± 151 and 3037 ± 165. It should be noted that histograms of all other tube potentials tested are not shown (for conciseness), but they all follow the same features as Figure 5.10.

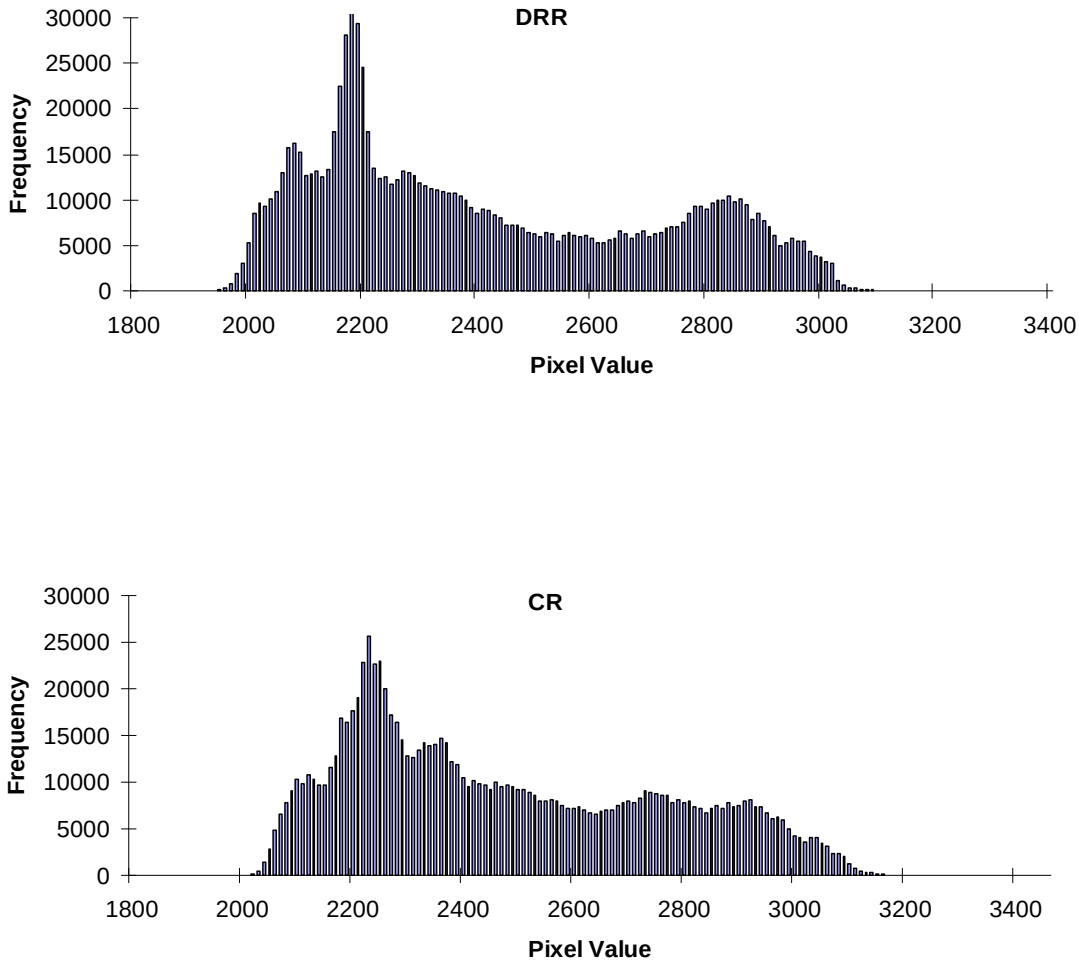

Figure 5.10: Histogram of pixel values of a DRR image, and the corresponding histogram for the CR image.

The maximum percentage difference in SNR measurements (DRR vs CR) was 8.6% in the lung, 10.2% in the spine and 13.5% in the diaphragm. The mean SNR (averaged over all tube potentials) are shown in Table 5.5. These demonstrate very good agreement when one considers the standard deviations are the mean values. The maximum difference in TRRs was 0.3%, with a mean difference of 0.14%. Again, the results show very good agreement.

Region	Mean SNR (DRR)	Mean SNR (CR)
Lung	201.0 ± 18.4	211.3 ± 21.1
Spine	119.8 ± 9.9	124.0 ± 13.9
Diaphragm	69.9 ± 11.8	67.6 ± 11.5

Table 5.5: Mean SNR values in each chest region of RANDO. All errors are ± 1 standard deviation of the mean.

All results obtained in this section demonstrate the DRR algorithm can produce images reconstructed without and with scatter rejection techniques.

5.3 Validation with the obese RANDO phantom

5.3.1 Obese RANDO images with no scatter rejection

Validation was carried out by comparing real CR and simulated DRR images of the obese RANDO phantom. Phantom images were acquired on the clinical CR system as described in Chapter 4 Section 4.2.1.2. Figure 5.11 shows a CR and DRR image of obese RANDO acquired and reconstructed with a tube potential of 60 kVp.

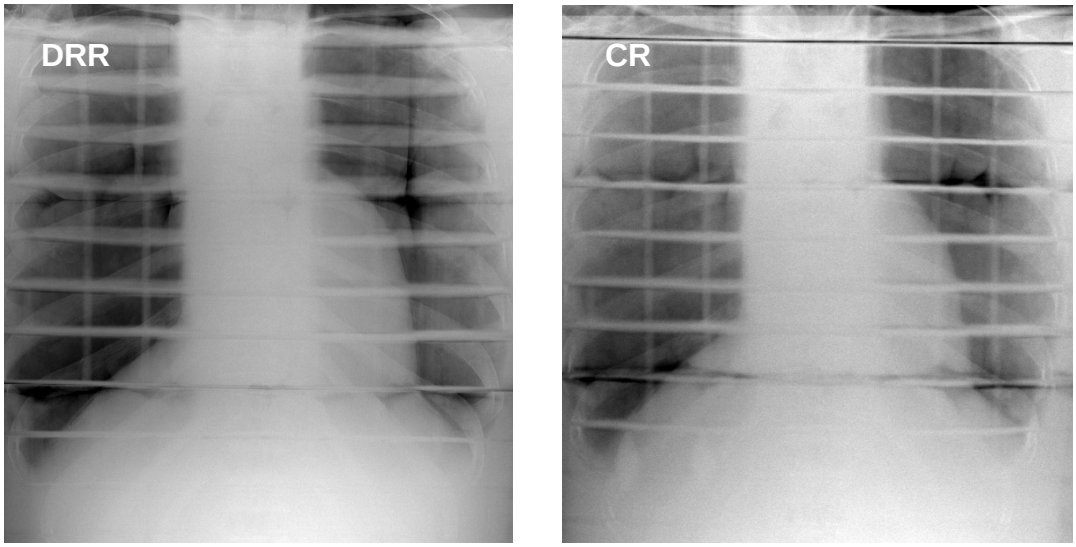

(a) (b)

Figure 5.11:(a) DRR acquired image of RANDO and (b), CR reconstructed image of RANDO.

As illustrated in Figure 5.11 the images correlate visually, and are of lower contrast than those of average RANDO (Figure 5.1). This is expected due to the increased amount of scattered radiation produced by the excess fat.

Figure 5.12 compares clinical (CR) and simulated (DRR) histograms of pixel values at tube potential 60 kVp respectively. The reduction in dynamic range of the DRR image is much more pronounced than in average RANDO. This is probably due to the presence of excess fat; it is thought that in converting CT number of fat to its corresponding linear attenuation coefficient (LAC) through the use of the Gammex tissue equivalent phantom, the chemical composition of the grocery store lard used is not exactly the same as that of the adipose tissue substitute. Any errors would lead to a difference in virtual X-ray attenuation compared with the real situation. It appears here that lard has a larger LAC than that modelled due to the observed reduced dynamic range. This will only prove a problem if validation becomes difficult with patient images.

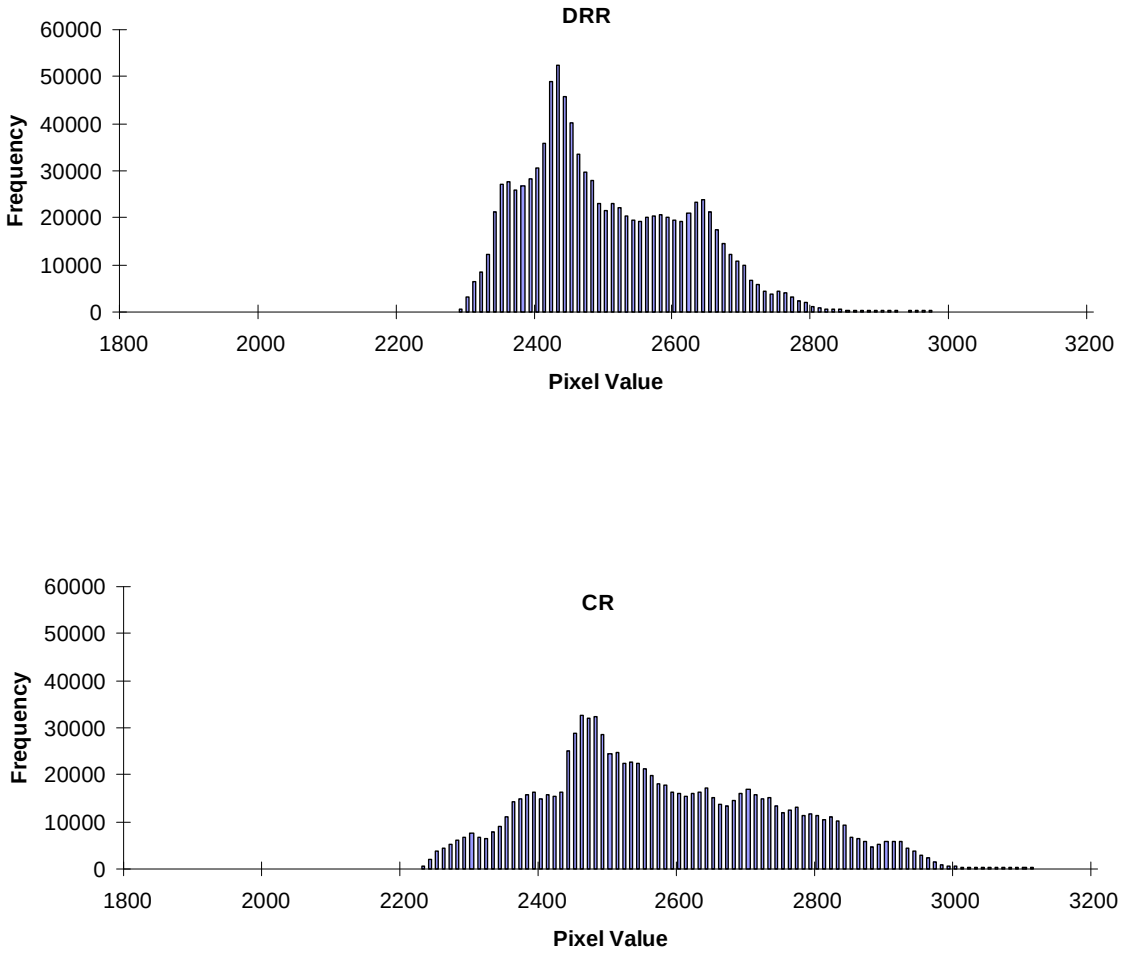

Figure 5.12: DRR and CR histograms of obese RANDO.

Nevertheless, the mean (± 2 SD) of the minimum and maximum histogram values for the CR images were 2213 ± 144 and 2970 ± 64 respectively. For the DRR images, these values were 2250 ± 180 and 2860 ± 62. It should be noted that histograms of all other tube potentials tested are not shown (for conciseness), but they all follow the same features as Figure 5.12.

The maximum percentage difference in SNR measurements (DRR vs CR) was 9.1% in the lung, 10.8% in the spine and 13.9% in the diaphragm. The mean SNR (averaged over all tube potentials) are shown in Table 5.6. These show very good agreement demonstrating signal and noise are similar in DRR and CR images.

Region	Mean SNR (DRR)	Mean SNR (CR)
Lung	193.1 ± 15.0	199.2 ± 13.8
Spine	122.5 ± 26.6	125.1 ± 27.1
Diaphragm	76.8 ± 7.4	75.5 ± 10.6

Table 5.6: Mean SNR values in each chest region of RANDO. All errors are 1 standard deviation of the mean.

Table 5.7 demonstrates TRRs calculated and measured in DRR and CR images respectively. These show good agreement and are typically lower than those of average RANDO. This is to be expected, as the increased amount of soft tissue (fat) increases the scattered radiation absorbed by the phosphor plate, which in turn reduces the contrast of various structures in the chest. With smaller differences between rib and the background the TRR will be lower (i.e. the TRR approaches unity).

Tissue to Rib Ratio (all ± 0.005)		
Tube Potential (kVp)	**TRR - DRR**	**TRR - CR**
50	1.017	1.017
60	1.014	1.014
70	1.011	1.013
81	1.008	1.008
90	1.005	1.005
102	1.005	1.005
109	1.005	1.005
125	1.004	1.006
133	1.005	1.005
141	1.005	1.005
150	1.005	1.006

Table 5.7: Comparison of Tissue to Rib Ratios measured in DRR and real CR images.

The results obtained here clearly demonstrate DRR images exhibit smaller dynamic range with respect to CR images. However, SNR and TRR results are very similar, and therefore the impact of DRR reduced dynamic ranges is likely to be limited, as the visibility of normal and pathological structures are primarily dictated by the level of signal and noise.

5.3.2 Obese RANDO images reconstructed with an anti-scatter grid

Images of obese RANDO acquired (CR) and reconstructed (DRR) with an anti-scatter grid used in the algorithm are shown in Figure 5.13. There is more detail in the spine and diaphragm regions of the chest, which, as discussed in Section 5.2.2, is expected.

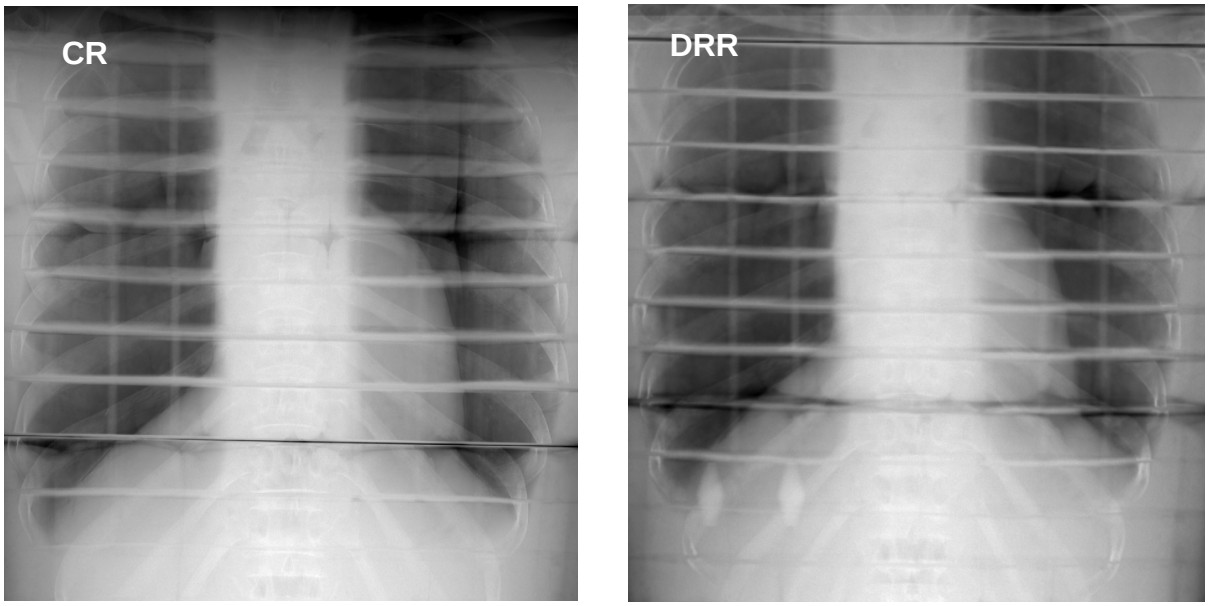

(a) (b)

Figure 5.13 (a) CR acquired image of obese RANDO and (b), a DRR reconstructed image.

Figure 5.14 compares clinical (CR) and simulated (DRR) histograms of pixel values at tube potential 60 kVp respectively. This clearly demonstrates that the histogram shapes are similar, but as expected, the dynamic range of the DRR is smaller than for the CR. The mean (± 2 SD) of the minimum and maximum

histogram values for the CR images were 1827 ± 212 and 3160 ± 82 respectively. For the DRR images, these values were 1885 ± 225 and 3105 ± 85.

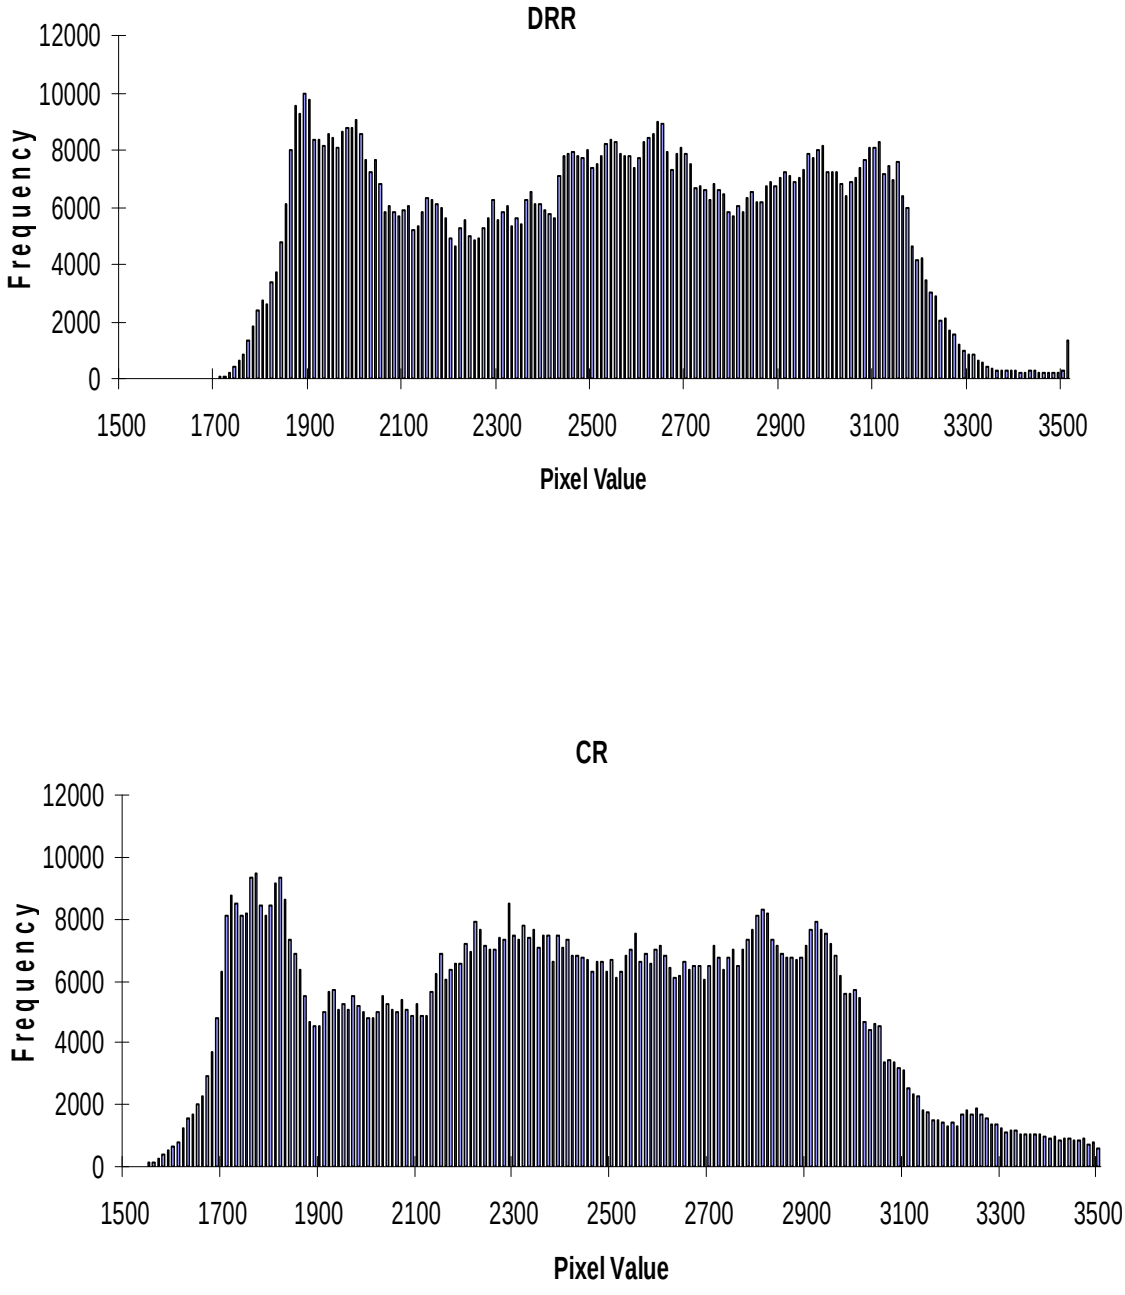

Figure 5.14: Histograms of obese RANDO DRR and CR images respectively.

The maximum percentage difference in SNR measurements (DRR vs CR) was 7.5% in the lung, 10.7% in the spine and 9.6% in the diaphragm. The mean

SNR (averaged over all tube potentials) are shown in Table 5.8; there is very good agreement in the SNR values.

Region	Mean SNR (DRR)	Mean SNR (CR)
Lung	147.6± 23.6	153.6 ± 23.6
Spine	61.1 ± 15.4	61.1 ± 16.4
Diaphragm	42.7 ± 9.2	43.1 ± 9.8

Table 5.8: Mean SNR values in each chest region of RANDO. All errors are 1 standard deviation of the mean.

Finally, the maximum difference in TRRs was 0.50%, with a mean difference of 0.16%. This demonstrates very good agreement.

5.3.3 Obese RANDO images reconstructed with an air gap technique

Images of obese RANDO acquired (CR) and reconstructed (DRR) with an air gap technique are shown in Figure 5.15. The images correlate very well.

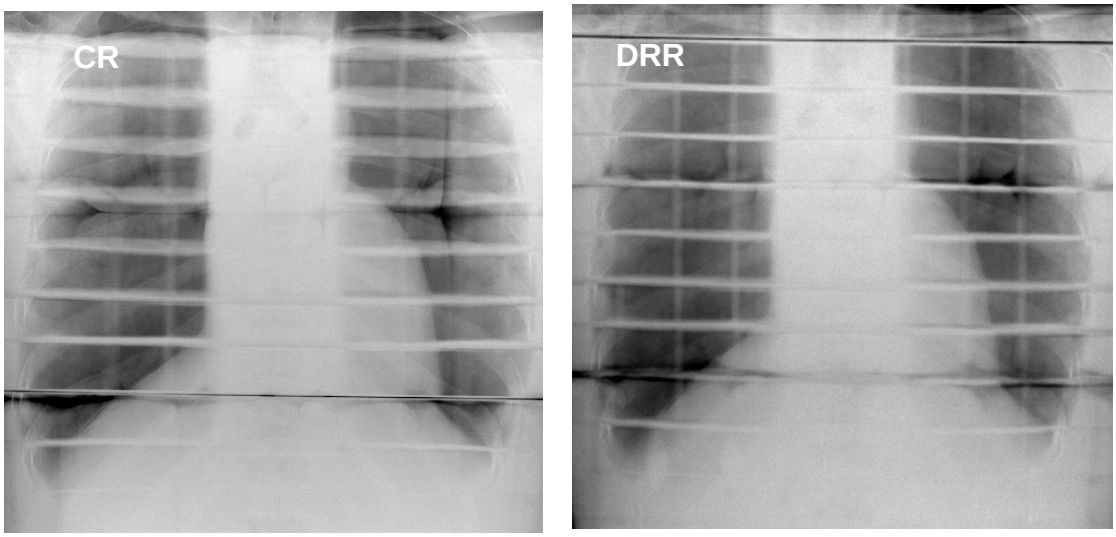

(a) (b)

Figure 5.15: Obese RANDO (a) acquired and (b) reconstructed with an air gap technique.

Figure 5.16 compares histograms of DRR reconstructed and CR acquired obese RANDO. As demonstrated, the shapes are very similar. Maximum percentage difference in SNR measurements (DRR vs CR) was 7.2% in the lung, 5.6% in the spine and 11.1% in the diaphragm. The mean SNR (averaged over all tube potentials) are shown in Table 5.9; there is very good agreement in the SNR values. Finally, the maximum difference in TRRs was 0.40%, with a mean difference of 0.16%. This demonstrates very good agreement.

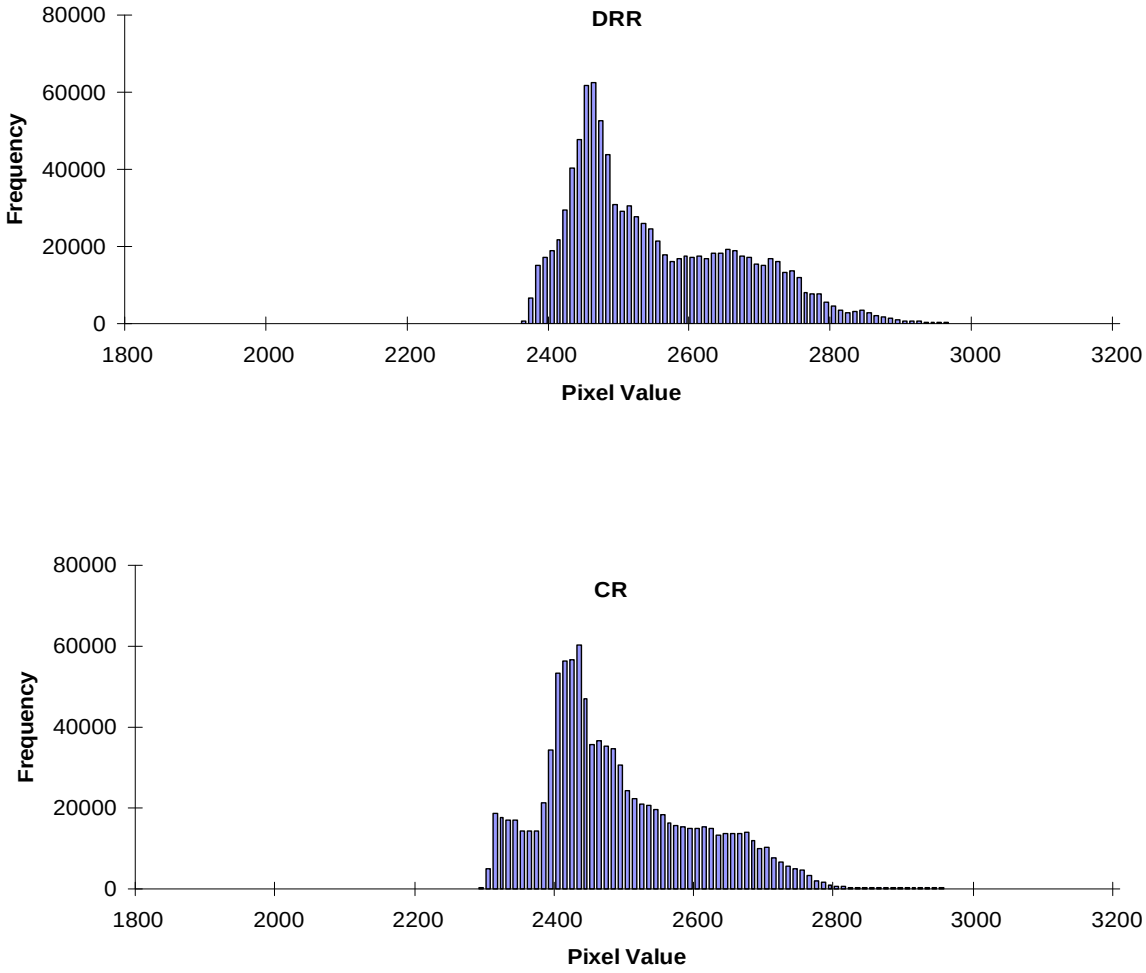

Figure 5.16: DRR and CR histograms of obese RANDO imaged with an air gap technique.

Region	Mean SNR (DRR)	Mean SNR (CR)
Lung	238.3 ± 23.3	245.4 ± 26.4
Spine	145.8 ± 21.6	146.7 ± 22.5
Diaphragm	116.1 ± 10.9	119.9 ± 10.8

Table 5.9: Mean SNR values in each chest region of RANDO. All errors are 1 standard deviation of the mean.

5.4 RANDO phantom validation – conclusions

All the results obtained with the RANDO phantom demonstrate very good agreement visually and quantitatively. The only concern is with reduced dynamic range of DRR images with respect to CR, with this affecting pixel values mainly in the diaphragm region; however this is of little significance in chest radiography (private communication, Ged Avery, Consultant Radiologist, 2009). However, the levels of SNR in each region of the chest are in excellent agreement, and therefore the visibility of normal and pathological structures is expected to be acceptable.

5.5 Validation with real patient data

As has been discussed in this chapter, validation of the DRR computer model has been performed by comparing various indices derived from DRR and CR images of the RANDO phantom. However, to ensure the model was capable of simulating real patient images to an acceptable level, DRR images were compared to real patient CR images, and quantitative analysis was performed in the same manner as for RANDO.

5.5.1 Validation with average patient CR images

As described in Chapter 4, the DRR algorithm has been configured to a specific X-ray and CR system. To assess the 'transferability' of the model, as well as validate it, patient images from different hospital sites in the Hull & East Yorkshire Hospitals NHS Trust (HEY) were used. There are three main hospitals in HEY Trust, each of which has traditionally used different exposure factors for chest radiography. Validation with average patient data was therefore carried out with ten average male CR chest images acquired at each site. These were acquired with exposure factors 60 kVp/10 mAs, 70 kVp/5 mAs, and 80 kVp/5 mAs respectively.

Figure 5.17 compares an image of one of these males to that of a DRR simulated image (both typical of their cohort). It is impossible to determine which is the real patient and which is the reconstructed one. The DRR model is therefore capable of producing images that visually correlate with real ones.

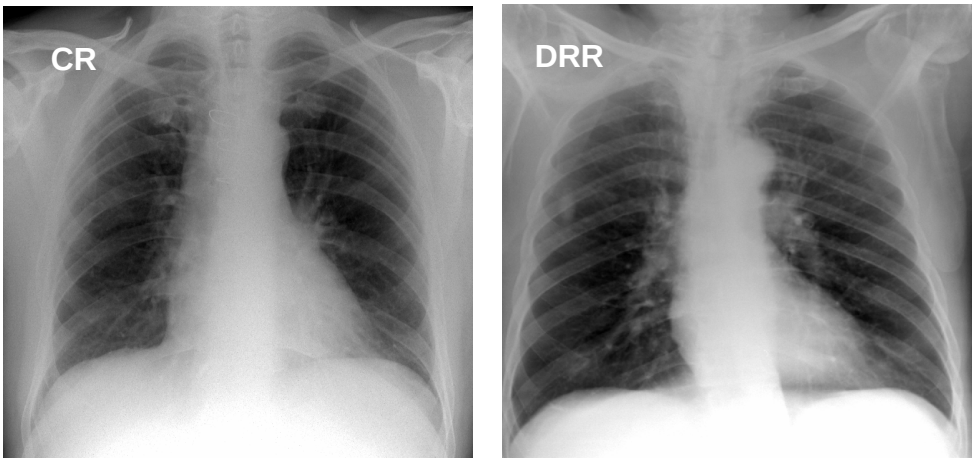

(a) (b)

Figure 5.17: (a) Real chest CR image of an average male, (b) DRR simulated image of an average male. Images (a) and (b) were acquired/reconstructed with exposure factors 70 kVp, 5 mAs. *It should be noted they are not the same patient.*

Figure 5.18 show histograms of pixel values for the images (typical for their cohort) illustrated in Figure 5.17. The histogram shapes are very similar. There is a shift of the histogram peak to lower pixel values for CR images. This is probably due to slightly more scatter recorded in the images acquired with real patients because of the presence of more fat compared to the RANDO phantom (scatter added to the DRR images is based on scatter factors derived from RANDO alone). There is also a reduction in the dynamic range of the DRR image compared to the CR. This is probably due to those factors discussed in Section 5.2.1. Histograms of pixel values were of a similar shape for all other patients selected, but are not shown here for conciseness.

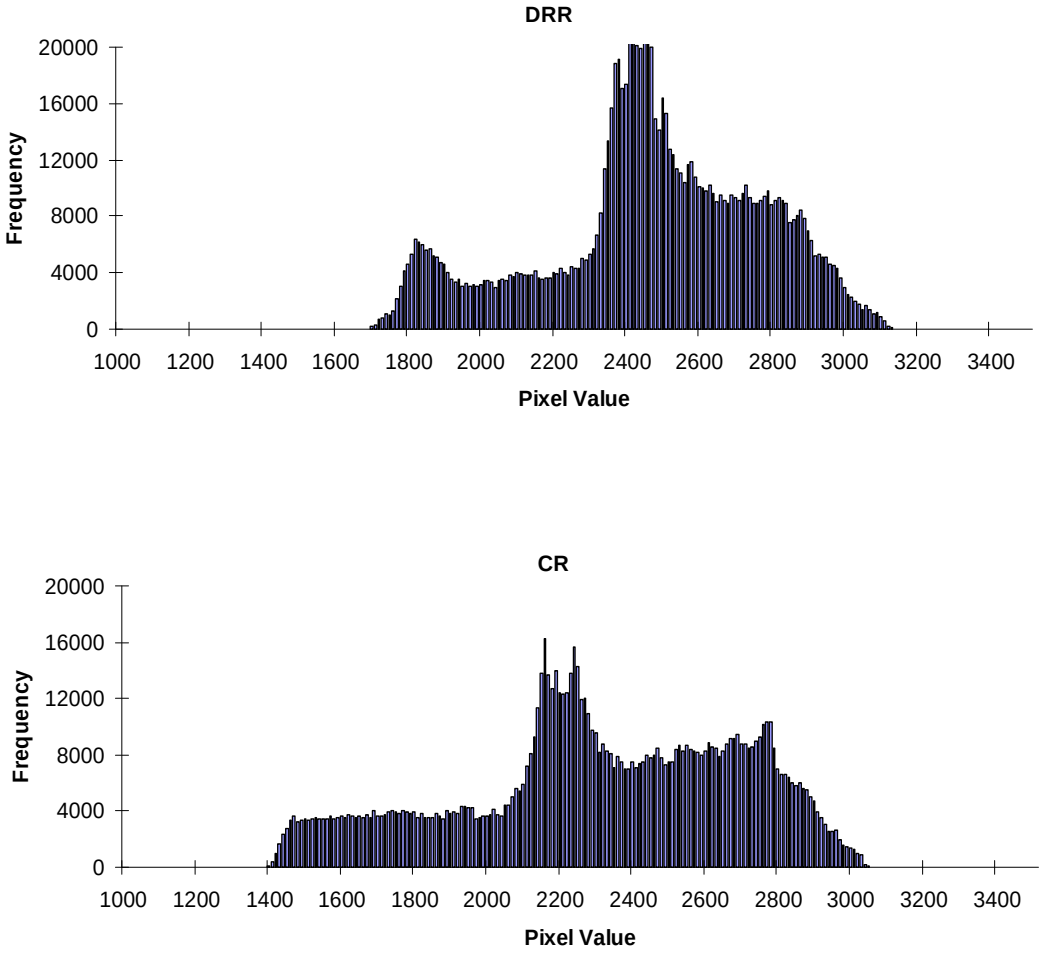

Figure 5.18: Real chest CR histogram of an average male and a DRR simulated histogram of an average male.

The mean (± 2 SD) of the minimum and maximum histogram values for the ten 60 kVp real CR images were 1503 ± 320 and 3120 ± 220 respectively. For ten 60 kVp DRR images, these values were 1700 ± 300 and 3110 ± 300.

Secondly, the mean (± 2 SD) of the minimum and maximum histogram values for the ten 70 kVp real CR images were 1500 ± 320 and 3260 ± 80 respectively. For ten 70 kVp DRR images, these values were 1800 ± 300 and 3060 ± 310.

Finally, the mean (± 2 SD) of the minimum and maximum histogram values for the ten 80 kVp real CR images were 1770 ± 280 and 3265 ± 160 respectively. For ten 80 kVp DRR images, these values were 1860 ± 180 and 3220 ± 150.

The calculated (DRR) measure of dynamic range is smaller than that measured (CR) for all tube potentials and more pronounced at minimum pixel values. As discussed previously, these pixel values typically reside in the diaphragm region which is of little interest clinically and therefore is probably not a limiting factor.

Table 5.10 shows SNRs in the lung area of thirty real patient images, compared with that of thirty simulated DRR images, acquired and reconstructed at 60, 70 and 80 kVp respectively. All images were chosen at random, since it was not possible to simulate a DRR image of a given patient and obtain a CR image of the same patient. The SNRs of the thirty randomly chosen patients (real and simulated) correlate very well. Mean values are the same at each tube potential given the standard deviations. The differences in SNRs are most likely due to differences in patient size (although they are ‘average’, a weight range of 70 ± 10 kg was chosen) leading to changes in X-ray absorption and scatter.

60 kVp		70 kVp		80 kVp	
SNR Real Patient	**SNR Simulated Patient**	**SNR Real Patient**	**SNR Simulated Patient**	**SNR Real Patient**	**SNR Simulated Patient**
77.4	78.4	102.9	105.4	76.8	76.7
109.4	96.5	127.0	107.3	82.2	72.9
155.8	158.4	144.0	127.4	81.9	86.3
104.2	120.2	99.0	105.7	108.8	100.1
135.5	130.2	141.6	142.3	119.1	122.8
85.6	90.2	124.4	145.1	108.7	107.8
123.7	109.9	106.6	100.2	107.3	107.9
109.9	126.0	116.6	113.3	109.2	116.1
126.2	112.9	152.5	149.7	129.9	130.9
95.7	92.6	126.6	134.2	116.3	94.5
Mean = 112.2	Mean = 111.5	Mean = 124.1	Mean = 123.0	Mean = 104.0	Mean = 101.6
St Dev = 23.9	St Dev = 23.5	St Dev = 18.2	St Dev = 18.8	St Dev = 17.7	St Dev = 19.2

Table 5.10: Comparison of SNRs measured in the lung region in thirty random DRR and thirty random CR images.

Table 5.11 shows SNRs in the spine area of thirty real patient images, compared with that of thirty simulated DRR images, acquired and reconstructed at 60, 70 and 80 kVp respectively. All images were chosen at random. The SNRs of the thirty randomly chosen patients correlate very well. Mean values are the same at each tube potential given the standard deviations.

60 kVp		70 kVp		80 kVp	
SNR Real Patient	**SNR Simulated Patient**	**SNR Real Patient**	**SNR Simulated Patient**	**SNR Real Patient**	**SNR Simulated Patient**
44.7	40.0	68.5	56.8	47.3	45.3
69.8	63.6	71	68	38.3	48
75.8	64.9	59.1	57.3	45	48.6
54.4	49.8	52.3	60.7	39.8	54.8
76	85	58.9	58.6	32.9	43.8
59.4	45.3	65.6	60.9	32.3	38.3
90.7	84.2	55.8	55.1	40.3	35.2
68.3	61.9	70	59.2	39.7	44.8
73	79.9	63.2	68.9	42.7	54
83.1	87.9	57.9	50.3	38.9	34.4
Mean = 69.5	Mean = 66.3	Mean = 62.6	Mean = 59.6	Mean = 39.7	Mean = 44.7
St Dev = 13.6	St Dev = 17.5	St Dev = 6.4	St Dev = 5.6	St Dev = 4.7	St Dev = 7.1

Table 5.11: Comparison of SNRs measured in the spine region in thirty random DRR and thirty random CR images.

Table 5.12 shows SNRs in the diaphragm area of thirty real patient images, compared with that of thirty simulated DRR images, acquired and reconstructed at 60, 70 and 80 kVp respectively. All images were chosen at random. The mean SNR value for the diaphragm is typically lower in the real CR images than that of the simulated DRR images. This is probably due to more fat surrounding the abdomen of the patients relative to that of RANDO. This will increase the amount of scatter reaching the CR phosphor and hence noise, thus forcing down the SNR. However, as suggested previously in this chapter, image quality

in the diaphragm is of little importance in chest radiography, so this difference is likely to be of no significance.

60 kVp		70 kVp		80 kVp	
SNR Real Patient	**SNR Simulated Patient**	**SNR Real Patient**	**SNR Simulated Patient**	**SNR Real Patient**	**SNR Simulated Patient**
58.0	31.2	33	21.2	26.6	25.9
22.8	20.9	33.6	21.2	27.1	27.3
14.9	9.2	29.7	27.1	22.5	17.6
9.3	10.8	31.5	26.6	22.1	23.5
22.2	17.9	29.1	27.7	28.3	19.7
30.3	29.2	25.2	19.9	24.3	23.9
26.8	23.1	30.7	22.1	25.1	19.8
31.2	24.9	28.9	20.5	26.8	24
20.5	27.9	31.9	27	21.7	20.3
39.0	29.0	26.7	25.8	23.1	25.7
Mean = 27.5	Mean = 22.4	Mean = 30.0	Mean = 23.9	Mean = 24.7	Mean = 22.8
St Dev = 13.6	St Dev = 7.7	St Dev = 2.7	St Dev = 3.2	St Dev = 2.4	St Dev = 3.2

Table 5.12: Comparison of SNRs measured in the diaphragm region in thirty random DRR and thirty random CR images.

As per the SNR measurements, thirty random real and simulated patients were used to measure the mean Tissue to Rib Ratio (TRR) at each respective tube potential (ten per kVp). Mean (± SD) DRR and CR TRR values were 1.043 ± 0.016 and 1.047 ± 0.018 for 60 kVp, 1.035 ± 0.006 and 1.036 ± 0.005 for 70 kVp, and 1.034 ± 0.08 and 1.033 ± 0.011 for 80 kVp. Differences are probably due to varying patient rib thickness. Nevertheless, there is a satisfactory

agreement, and the same trend of higher TRR at lower tube potential as found with the RANDO phantom exists.

All the results shown here demonstrate the DRR algorithm is quantitatively capable of simulating real average patients and therefore can be used for optimisation studies. DRR dynamic range is typically lower, but SNRs in each region compare very well, especially in the lung and spine. DRR diaphragm dynamic range and SNR are typically lower than that found in patient images, but given the large fluctuation of these indices in real patients (see Tables 5.10 – 5.12) and limited clinical significance of this chest region, the impact is likely to be negligible.

The results also demonstrate that the DRR algorithm is transferable and can be used to simulate CR readers and X-ray systems that it was not originally configured to, despite being produced to simulate a specific system. This adds to the versatility and usefulness of the model.

5.5.2 Validation with obese patient CR images

Validation was carried out using the same methods as that described above, except that only one hospital site was used. Ten obese patients selected at random were chosen to compare with DRR simulated obese patients. All real patients were X-rayed with exposure factors 80 kVp and 8 mAs. Figure 5.19 compares an image of one of these males to that of a DRR simulated image (both typical of their cohort).

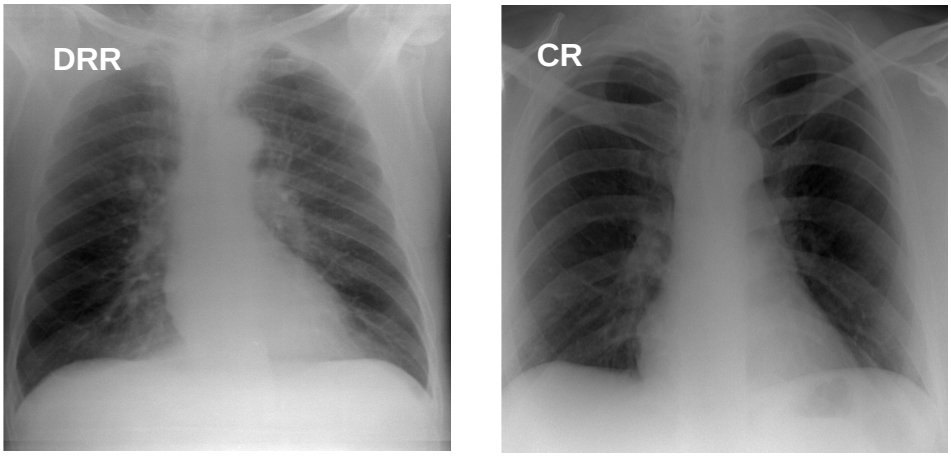

(a) (b)

Figure 5.19:(a) Obese patient DRR reconstructed image and (b), obese patient CR acquired image. *They are not of the same patient.*

The images in Figure 5.19 represent typical reconstructed (DRR) and acquired (CR) obese patients. This Figure shows it is not possible to distinguish between the real patient image and the DRR. They are both of lower contrast than the images shown in Figure 5.17, which is expected.

Figure 5.20 shows the histogram of pixel values for both of the above images. The histogram shapes are broadly similar (both have two peaks), but the dynamic range of the DRR histogram is smaller than the CR acquired one. As discussed throughout this chapter, this is a known limitation of the algorithm, and is discussed in Section 5.7. Other histograms are not shown, but all follow the same trends as those illustrated here.

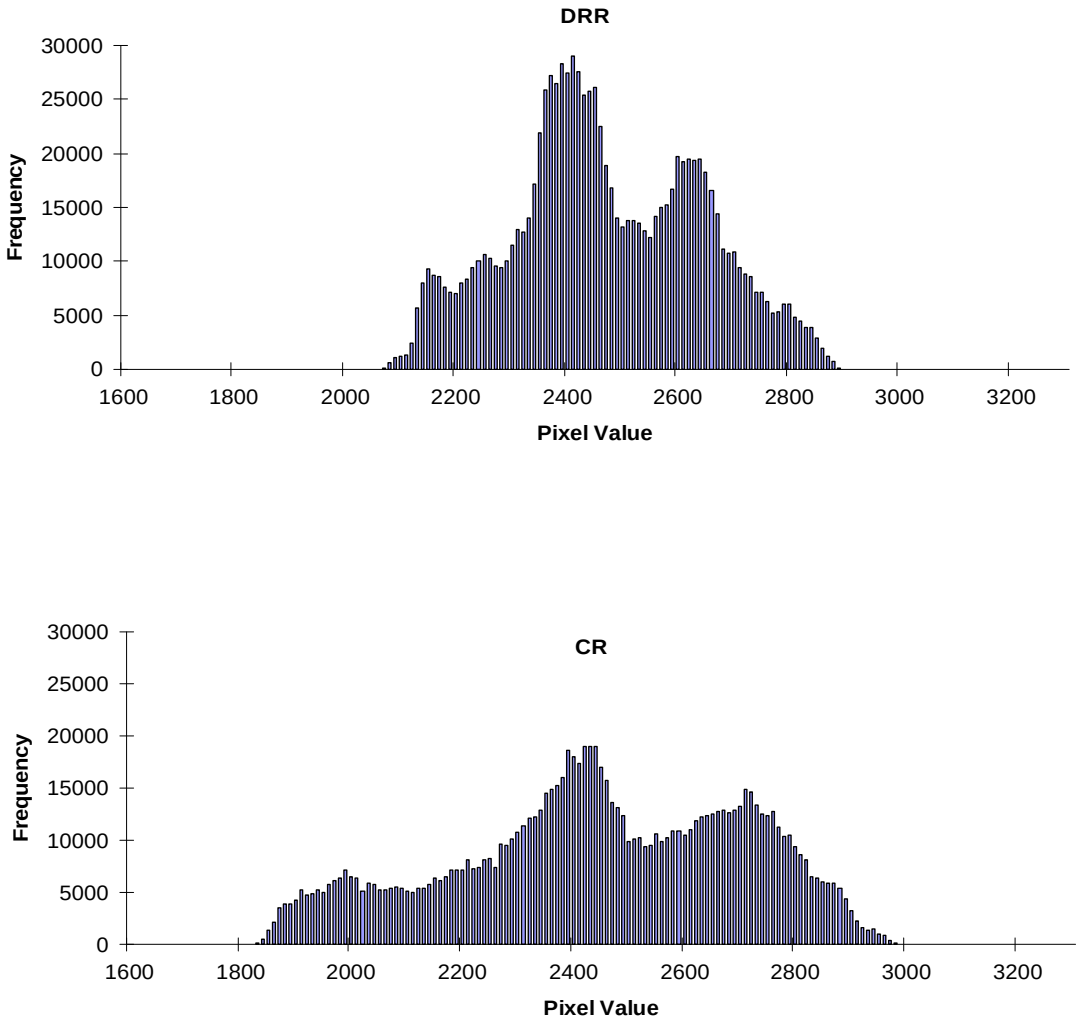

Figure 5.20: Histograms of DRR reconstructed and CR acquired obese patients.

The mean (± 2 SD) of the minimum and maximum histogram values for the ten real CR images were 1790 ± 220 and 3075 ± 280 respectively. For ten DRR images, these values were 2000 ± 100 and 3070 ± 200. Minimum pixel values are lower in the CR acquired images than in the DRR reconstructed ones (maximum pixel values are very similar). These pixels exist in the diaphragm region of the images and therefore surrounded by fat; the reasons discussed in Section 5.5.1 are therefore relevant here as well.

Table 5.13 shows SNRs in each chest region of ten real patient images, compared with that of ten simulated DRR images, acquired and reconstructed

at 81 kVp respectively. All images were chosen at random in the same manner as discussed in Section 5.5.1. There is good agreement of SNRs in real and simulated images.

SNR Real Patient - lung	SNR Simulated Patient - lung	SNR Real Patient - spine	SNR Simulated Patient - spine	SNR Real Patient - diaphragm	SNR Simulated Patient - diaphragm
128.5	191.1	128.1	97.9	69.4	36.1
145.3	121.9	109.6	80.2	55.8	75.4
107.8	79.5	52.3	71.3	54.4	32.6
100.1	102.3	81.9	94.2	54.1	51.7
127.2	109.9	65.2	82.6	47.9	36.4
133.9	125.6	26.4	64.7	38.6	25.6
138.6	121.2	34.1	68.8	39.9	42.3
96.5	115.1	29.7	63.9	42.8	14.7
117.7	130.1	72.5	42.4	60.2	22.7
145.8	114.6	51.4	37.5	34.3	28.8
Mean = 123.6	Mean = 120.6	Mean = 65.1	Mean = 70.4	Mean = 49.7	Mean = 36.6
St Dev = 17.9	St Dev = 28.6	St Dev = 33.9	St Dev = 19.8	St Dev = 11.0	St Dev = 17.1

Table 5.13: SNRs of real and simulated obese patients in each region of the chest.

Ten random real and simulated patients were used to measure the mean Tissue to Rib Ratio (TRR). Mean (± SD) DRR and CR TRR values were 1.023 ± 0.011 and 1.026 ± 0.011 respectively. As discussed in Section 5.4.1, any differences are probably due to varying rib thickness in real patients. Nevertheless, there is a satisfactory agreement.

These results suggest that the DRR algorithm is capable of simulating obese patients to a satisfactory level. SNR and TRR indices are very similar,

suggesting normal and pathological features in the chest will be of the same appearance. DRR dynamic range is smaller than CR, with pixel values in the diaphragm region typically higher. However, as discussed, image quality is of limited importance in the diaphragm region in chest radiographs, so this limitation is of little consequence.

5.6 Image Evaluator Interpretation of DRR Images (Qualitative Validation)

As discussed already in this chapter, the DRR algorithm has been quantitatively validated using histograms, signal-to-noise and tissue-to-rib ratios, and dynamic range measurements. However, all of this becomes redundant if expert image evaluators deem the images inadequate for optimisation studies. To investigate this, 50 reconstructed patients (each containing 11 images) were presented to four independent, experienced image evaluators (2 Radiologists and 2 reporting Radiographers) on calibrated diagnostic reporting monitors (Barco Ltd, Brussels, Belgium). The monitors were calibrated to national standards [106] and were kept in dedicated viewing rooms with lighting levels maintained at an acceptable level. Each image evaluator was asked:

“Do the images contain sufficient clinical detail, are they definitely representative of chest anatomy and are they suitable for optimisation studies requiring anatomical noise?”

Even though all image evaluators knew each image was a DRR and not a real clinical CR image, the answer in all cases was yes which was a very satisfactory response. As introduced in Chapter 1, this research set out to produce a computer tool that adequately simulates CR chest radiographs, and

to use these images for optimisation of radiographic techniques. The excellent and positive comments from the evaluators, as well as the favourable quantitative analysis, has clearly demonstrated the main research objective of this thesis has been fulfilled.

Nevertheless, this qualitative validation was taken a step further by asking each evaluator:

“Do the images adequately mimic real CR images (1 = definitely not, 10 = definitely)?”

The mean (± 1SD) score was 7.5 ± 2.0 demonstrating the algorithm’s capability of producing chest radiographs. However, each evaluator did comment that DRR image spatial resolution was poorer than CR (as discussed and measured in Chapter 4), but they agreed that this was not a limitation, and that optimisation studies would be possible using these images.

The same scoring criterion was applied to the simulated lung nodules (discussed in Chapter 4). Although the method of adding lesions has been validated by Li *et al.* [97], it was felt necessary to qualitatively validate the appearance of these in the resulting DRRs. Each image evaluator was asked to score 50 reconstructed patients, each with lung lesions visible. They were asked:

“Is the nodule realistic in terms of appearance and density (1 = not realistic, 10 = definitely realistic)?”

The mean (± 1 SD) was 7.8 ± 1.2, which demonstrates the appearance of simulated lung lesions are sufficiently realistic and can be used in the images for optimisation.

5.7 Limitations of the DRR computer model

There are numerous limitations with the DRR computer model and they are discussed in turn below.

5.7.1 X-ray spectra

The accuracy of X-ray spectra produced by IPEM Report 78 is dependent upon the variables chosen to produce each spectrum, such as tube potential, filtration, target angle and voltage ripple. The maximum quoted error for photon intensity is 10%. This will influence the results obtained with the DRR model, although they have been minimised as the output of the clinical X-ray system was measured and the spectra produced by IPEM 78 were corrected accordingly. Also, given that the output of real X-ray tubes can easily fluctuate by 10% or more from one annual survey to the next, any error associated with this are unlikely to be observed.

5.7.2 CT number to linear attenuation coefficient (LAC) conversion

The Gammex RMI phantom was used to measure the CT number of various tissue substitutes, and these were converted to LACs using the XCOM program. There is an inherent tolerance of approximately 10% in LAC values provided by this program, and this would affect the LAC values, and therefore X-ray attenuation and dynamic range of the DRR model.

5.7.3 DRR dynamic range

The size of the voxels available in the CT data (0.34 mm width x 0.34 mm depth x 0.8 mm height) are typically larger than the smallest structures within the human chest; this will reduce the number of tissues available in the virtual

patient for X-ray attenuation compared to a real patient (the very small structures will be ‘lost’). This would have the effect of reducing the dynamic range of the attenuated simulated X-rays. Also, the Gammex phantom was used to derive attenuation coefficients for seventeen tissue types; even though seventeen is an improvement on the four used by the computerised Monte Carlo phantom [46], it is still a finite number of tissue substitutes, subsequently affecting the amount of tissues modelled in the virtual patient. The virtual patient will therefore have a different ‘attenuation cross section’ compared to a real patient, reducing the dynamic range of resulting DRR images. However, given the fluctuation of minimum and maximum pixel values in real patient images, the effect of reduced dynamic range is not likely to be observable.

5.7.4 Ray casting voxel interpolation

The ray casting DRR method used in this research has used pencil beams projected through the CT data (virtual patient). In doing so each pencil beam impinges many different voxels requiring interpolation to calculate effective LACs. This interpolation was done by weighting the LAC of a voxel by the amount of area of the front face of the pencil beam impinging the voxel (i.e. sampling by area), and summing the respective contributions. This is not entirely accurate as the three dimensional pencil beam has been reduced to a two dimensional area. A more realistic method would have been to use volumetric interpolation (i.e. calculate the volume of pencil beam in each voxel). However, the area of pencil beam is broadly in proportion to its volume so the effect of this is likely to be minimal.

5.7.5 Photon energy binning

To speed up DRR compute time, reference photon energies were projected through LAC converted CT data correct for the reference energies. The intensity of exiting photons was calculated from a knowledge of the intensity of the incident photons. The subsequent number of photons exiting the virtual patient of all other energies was determined using the formulae derived in Chapter 4. This will inevitably introduce discrepancies in the intensity of photons exiting the virtual patient and those absorbed by the CR phosphor. However, it is anticipated any noise introduced in the DRR image due to this will be swamped by the addition of scattered radiation and frequency dependent noise.

5.7.6 Addition of scattered radiation

The addition of scattered radiation was based on the RANDO phantom alone. Although human chests are very similar, there are slight differences, and therefore using RANDO to model scatter will introduce a systematic error. Also, each scatter map was simply added over the raw DRR without any warping or registration, on the assumption average adult chests are so similar this was not required. This would again lead to systematic errors. However, the results suggest using RANDO to model scatter and adding this over the raw DRR is adequate. The limitations of adding scattered radiation in this way is therefore minimal.

5.7.7 Addition of frequency dependent noise

The method of noise addition used in this thesis assumes that the noise power spectrum (NPS) is a sufficient descriptor of the noise and that detective quantum efficiency (DQE) is constant over the dose variations that exist within

the image. CR systems tend to have a decreasing DQE with dose; however, this method (Bath *et al.* [49]) has been shown to be sufficient for doses used clinically, i.e. it is assumed that quantum noise must be the dominant source of noise in each image. This method also assumes that noise is ergodic (constant over time) but in real CR systems this is not the case. However, good agreement with SNR measurements illustrated throughout this chapter demonstrate there are no issues at clinical dose levels that result from the addition of noise in this way.

5.7.8 Spatial resolution

It is a fact that CT data is of poorer inherent resolution compared with that of CR images (although still much better than voxelated phantoms used in Monte Carlo studies). This has led to DRR images exhibiting lower resolution than desired. This problem will not be an issue in future if CT scan resolution improves and approaches that of CR. However, the prime objective of this research was to develop a simulation system that produces chest images that adequately simulate anatomical noise and pathology. Radiologist comments and scores are excellent, so it appears the spatial resolution of these DRR images is not a limiting factor.

5.7.9 DRR compute time

In spite of the methods used to speed up the DRR compute time, image reconstruction still takes hours. However, this is not a limitation for the end user (i.e. image evaluators), and optimisation of the code is beyond the scope of this work.

5.7.10 Conclusions: Limitations

Despite the numerous limitations described here, the DRR computer model has stood up to thorough quantitative and qualitative validation. The issues described are therefore not limiting factors.

5.8 Conclusions

This chapter has addressed how the DRR algorithm has been quantitatively and qualitatively validated. The images reconstructed by the method outlined in Chapter 4 have been analysed quantitatively in terms of signal to noise ratio (SNR), tissue to rib ratio (TRR), dynamic range and histograms of pixel values.

All DRR simulated images correlate visually with the real CR images. However, this is not sufficient alone for validation. Signal and noise must be evaluated, as the visibility of normal structures and pathology depends strongly upon these. DRR simulated SNR measurements were within 15% of those measured in real CR images using the average and obese RANDO phantom, across all tube voltages and typical tube currents encountered clinically.

Histogram analysis with images acquired and simulated with the RANDO phantom also demonstrated strong agreement, both in terms of shape and pixel value dependence on receptor dose. However, this analysis also showed DRR simulated images tend to have reduced dynamic ranges, probably due to voxel size and tissue variability relative to the Gammex phantom. In future it would be worth addressing this issue, but it is somewhat dependent on CT scan resolution (both in spatial and contrast terms) and the availability of phantoms that contain more tissue substitutes than the Gammex RMI model. It is, however, not an issue for this work as minimum and maximum pixel values

were the same as those encountered in CR images, given the standard deviations of the means.

TRR measurements derived from the average and obese RANDO phantom also demonstrated good agreement across all tube potentials. This is important, as clinical diagnosis of pathology can be disrupted by high contrast rib structures.

SNR, TRR and histogram analysis was also carried out with average and obese patient images. Strong agreement was found with SNR and TRR indices (given the standard deviations of the means), but the dynamic range of the DRR images were smaller. However, this is of little significance since the dynamic ranges were smaller due to differences in the minimum pixel values; as discussed throughout the chapter, these pixel values reside in the diaphragm region of the chest which is of little importance in chest radiography. Also, minimum and maximum pixel values tend to fluctuate quite considerably from one patient image to the next if one examines the standard deviations of the means (probably due to differences in patient sizes, even though they were either all average or obese), so the differences between DRR and CR are not likely to be observed and are no worse than those seen clinically between patients.

Analysis was carried out with CR images acquired at different hospital sites and different tube potentials/currents. The good agreement between indices proves the DRR algorithm is transferable and not restricted to modelling the X-ray and CR systems it is configured to.

Qualitative validation was then carried out with expert image evaluators Their comments and scores were excellent, indicating all images contain the

necessary projected anatomy required for optimisation. They were also satisfied simulated lung lesions were adequate.

Finally, the limitations of the computer model have been discussed and although they are numerous, the only one recognised by image evaluators was the reduced spatial resolution. However, this artefact was not a limiting factor. The main research objective of this work was to develop a computer algorithm that simulates real chest CR images with adequate anatomical noise. The results illustrated in this chapter, both quantitative and qualitative, demonstrate this objective has been achieved.

Chapter 6: Use of the digitally reconstructed radiograph algorithm for the optimisation of chest radiographic techniques for the Agfa Computed Radiography imaging system

6.1 Introduction

This chapter discusses the use of the DRR algorithm presented in this thesis, to optimise radiographic techniques for chest CR imaging. A summary of scores from expert image evaluators of reconstructed images produced at various tube potentials, receptor air kerma and scatter rejection methods is given.

At the outset of this work, chest exposure factors across the Radiology Department in the Hull & East Yorkshire Hospitals NHS Trust (HEY) were not standardised, and therefore not optimised; this chapter concludes with recommendations to the HEY Radiology Department for optimum exposure factors and technique for chest radiography. These were implemented across the Trust as a result of this work.

6.2 Chest image reconstruction

The DRR simulator was used to reconstruct images according to the following seven steps:

1. CT data sets from fifty average sized patients (70 ± 10 kg) were used to generate simulated CR images without any scatter rejection and at 11 different tube potentials; 50 to 150 kVp in approximate steps of 10 kVp. Each image was reconstructed with a matched effective dose of 0.013 (± 1%) mSv (effective dose of 0.013 mSv was chosen as this was the mean value for chest exposures at HEY Trust at the time of the study); this was achieved by deriving, in a trial and error manner, the necessary dose area product (DAP) value at each tube potential (to provide 0.013 mSv) using the effective dose calculation software PCXMC [107]. The required tube current-time product (mAs) at each tube potential was subsequently measured on a clinical X-ray system for each DAP value derived. The DRR computer model uses the mAs values to calculate the intensity of X-ray photons incident on the virtual patient. DAP and corresponding mAs values are shown in Table 6.1. Although chest imaging exposure factors were not standardised in this Trust at the time of the study, scatter rejection was not routinely used. The technique described here was therefore consistent with the chest imaging protocol used in our Radiology department.

Tube Potential (kVp)	DAP ($mGycm^2$)	Tube current-time product (mAs)
50	179	25
60	134	10
70	109	6.4
81	91	4
90	81	2.5
102	72	2
109	68	1.6
125	61	1.3
133	59	1
141	57	0.8
150	55	0.7

Table 6.1: The exposure settings used to reconstruct each image for a matched effective dose of 0.013 mSv. The mAs value is required for image reconstruction.

2. CT data sets from twenty average sized patients were used to generate simulated CR images with an oscillating focused anti-scatter grid (strips per mm = 4, grid ratio = 12) focused at 140 cm focus to detector distance (useful range 115 to 180 cm) incorporated into the algorithm. Each patient DRR set contained 5 images each of a different tube voltage; 60, 81, 102, 125, 141 kVp. Fewer images per patient were used on the advice of the image evaluators, following their evaluations of criteria 1; it became time consuming to score numerous patients and difficult to differentiate differences in image quality in steps of 10 kVp. Each image was reconstructed with a matched effective dose of approximately 0.03 mSv (the increase in dose was required to establish a lgM of 2.00 at the reference tube potential of 102 kVp with the use of the grid).

It should be noted here that the increase in dose to achieve a lgM of 2.00 was done for pragmatic reasons. For film-screen systems the image quality improvement through grid use is often described by a contrast improvement factor (K_c); this requires a fixed total exposure (primary plus scatter) to ensure proper film darkening. Using a constant lgM of 2.00 is akin to maintaining constant film darkening. However, digital X-ray imaging devices are not restricted to total exposure, and so the image quality improvement factors associated with the use of a grid are probably best described by the signal-to-noise ratio (SNR) [113]. The method used in this work may possibly restrict the optimisation, but until a recognised metric for 'image quality standardisation' is established for grid use with digital imaging, it was felt prudent to use the constant lgM method.

3. CT data sets from twenty average sized patients were used to generate simulated CR images with an air gap anti-scatter technique incorporated into the algorithm. Each patient DRR set contained 5 images each of a different tube voltage – 60, 81, 102, 125, 141 kVp. Each image was reconstructed with a matched effective dose of approximately 0.017 mSv (increase in dose due to the patient being closer to the X-ray source).

4. Steps 1 to 3 were repeated but for obese patients, with matched effective doses 0.022 mSv, 0.10 mSv and 0.03 mSv respectively; the increase in doses being required for imaging of obese patients.

5. Ten average sized patients each containing images produced with varying receptor air kerma values through the lung region; 7.0, 4.0, 3.0, 2.0, 1.5, 1.0, and 0.5 µGy.

6. Ten average sized patients each containing three images; an image reconstructed without any scatter rejection, an image reconstructed with an anti-scatter grid, and an image reconstructed with the air gap technique. The three images within each patient were reconstructed with the same tube potential, but each patient was reconstructed with a different potential (patient 1 = 50 kVp, patient 10 = 150 kVp). Each image was reconstructed with the effective dose correct for the method of scatter rejection utilised (as discussed above).

7. As step 6 with ten obese patients.

Images of a given patient were attached in series to a single study and given the name CR_Patient_N, were N was a sequential image number. For example, patient 1 had 11 series (images 50 to 150 kVp – see Table 6.1) and this study was called CR_Patient_1. The study was sent to the Picture Archiving and Communications System (PACS) network for observer evaluation. All other patients were subsequently reconstructed, named and sent to PACS in the same manner. Not all studies were sent with the images in numerical order (i.e. 50, 60, 70…..150 kVp) but in a random manner.

6.3 Evaluation of clinical image quality

The Council of European Communities (CEC) Quality Criteria [5] define important anatomical and image details for various diagnostic examinations, including chest radiography. The image criteria described in the CEC document, slightly revised to reflect modern diagnostic requirements and previous experiences of other groups [40, 41, 44, 108, 109], were used to define anatomical features in each chest image for evaluation. As well as general

chest structures, lung nodules were simulated in each image and evaluated, and are shown in Table 6.2.

	CEC Guidance
1	Vessels seen 3 cm from the pleural margin
2	Thoracic vertebrae behind the heart
3	Retrocardiac vessels
4	Pleural margin
5	Vessels seen en face in the central lung region
6	Hilar region
	Other Criteria
7	Abnormality in the lateral pulmonary region
8	Abnormality in the hilar region
9	Are the ribs a distraction? Y/N

Table 6.2: The chest structures used for visual grading analysis. Structures 1 to 6 are mentioned in the CEC document.

Four experienced expert image evaluators (two Radiologists and two reporting Radiographers) evaluated and graded the images on a diagnostic reporting workstation with a dual monitor configuration (Barco Ltd, Brussels, Belgium). The monitors were calibrated to national standards [106] and were kept in dedicated viewing rooms with lighting levels maintained at an acceptable level. The evaluators were asked to keep the final image in the series of each patient on the right hand screen, as this was the reference image for grading (for example, assuming 11 images per patient, image 11 was kept on the right hand

screen). All other images (the ‘test images’) were displayed in turn on the left hand screen and graded against the reference image. For the tube potential optimisation studies (i.e. points 1 to 4 in Section 6.2), the reference image was reconstructed with a tube potential of 102 kVp; this value was chosen because, as exposure factors were not standardised, a neutral tube potential was deemed appropriate so not to bias the results (the median diagnostic tube potential was therefore used). The evaluators did not have any knowledge of what tube potential the test or reference images represented. Test images were presented in a random order. Evaluators were allowed to change the window and level settings of each image prior to grading to optimise the appearance of each, as per clinical practice. For the scatter rejection study (points 6 and 7 in Section 6.2) the image derived without any scatter rejection was used as the reference image.

The evaluators graded the images using the visual grading analysis (VGA) system used by Tingberg and Sjostrom [44] in a similar study with a physical chest phantom. The VGA method has been analysed and validated as an appropriate method to assess the image quality of chest radiographs by Tingberg *et al.* [40], Mansson *et al.* [38,39], and Sund *et al.* [41]. The image quality for six structures mentioned in the CEC document, as well as lung nodules in the hilar and lateral lung regions in each of the test images (see Table 6.2), was compared with the reference image on a scale incorporating 7 points, as shown in Table 6.3.

Grading	Visibility of Structure
-3	Definitely inferior to the reference
-2	Reasonably inferior to the reference
-1	Slightly inferior to the reference
0	Equal to the reference
+1	Slightly superior than the reference
+2	Reasonably superior than the reference
+3	Definitely superior than the reference

Table 6.3: The grading system for VGA.

This work enhances previous VGA studies in that we have also included lung nodules as well as general chest structures.

Ribs have been shown to interfere with the diagnostic interpretation of chest images [110] therefore the evaluators were also asked to state 'yes/no' whether the ribs were an interference.

For each image, a VGA score (VGAS) was calculated using the equation described by Tingberg and Sjostrom [44]:

$$VGAS = \frac{\sum_{i=1}^{I} \sum_{s=1}^{S} \sum_{o=1}^{O} G_{I,S,O}}{I \times S \times O} \tag{6.1}$$

where $G_{I,\,S,\,O}$ is the grading (-3, -2, -1, 0, +1, +2, +3) given by observer o for image i and structure s, I is the number of images per tube potential (dependent on number of patients), S is the number of structures (eight in total – six general and two abnormal) and O is the number of evaluators (four in this study).

Negative and positive scores indicate inferior and superior image quality respectively in the test image compared to the reference image (see Table 6.3).

Scoring of the dose optimisation study (point 5 in Section 6.2) differed from the above as evaluators were simply asked to give ‘yes/no’ answers to whether the image quality in the lung, spine and diaphragm regions were diagnostically acceptable, and whether lesions in the lateral and hilar regions of the lung were visible.

Also, during the study to determine whether scatter rejection would be beneficial (points 6 and 7 in Section 6.2), as well as VGA scores, image evaluators were asked to answer ‘yes/no’ to the following question: “Image 1 (test image) is at least 1.5 times (up to 4 times) the dose of image 3 (reference image). If image quality of image 1 is ‘better’ than image 3, does it still justify the large increase in dose?” Only the relative dose increase was presented to the evaluators, rather than an absolute measure.

6.4 Statistical analysis

The results in this chapter were tested for significance, firstly using the ANOVA test (analysis of variance) to examine inter-observer variability, and then using the Student’s t-test to examine differences between mean image quality scores (VGAS). A p-value of $< 5\%$ was considered as a statistically significant difference between data sets. Error bars shown in the graphical data of the results were calculated by averaging the standard deviations of each VGAS for each chest structure (Table 6.2) over all patients.

6.5 Phantom experiment to determine minimum possible X-ray exposure times

In the event that the results demonstrated superior image quality with low tube potentials and anti-scatter techniques, it was felt necessary to determine whether a modern X-ray generator (discussed in Chapter 4) could deliver the required dose in exposure times of less than 20 ms (as recommended in the CEC guidance). To investigate this, average then obese RANDO was set up on a Philips X–ray system following the local clinical chest protocol. The phantom was exposed using the tube potentials shown in Table 6.1, using mA and corresponding ms settings sufficient to produce a lgM =2.00 ± 0.05. After each exposure the CR phosphor plate (the same one was used for each exposure) was read through the CR reader. It should be noted that exposure times were not measured with a survey meter; those displayed on the X-ray system were used.

6.6 Results and Discussion

For all images evaluated in this study, the ANOVA test demonstrated a p-value of ≥ 0.08 and therefore it was deemed there was no significant difference between the scores of each image evaluator.

6.6.1 Image quality of 50 average patients reconstructed without scatter rejection

The results of the optimisation study for 50 average patients reconstructed without scatter rejection, (i.e. in the same manner as the chest radiographic

technique currently used in our Radiology department for average adults), are shown graphically in Figure 6.1.

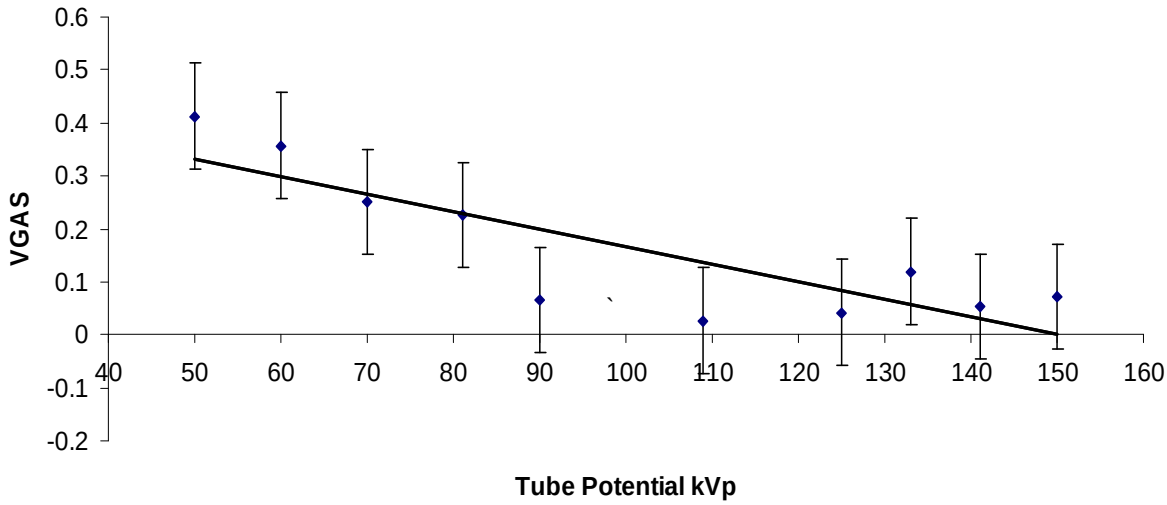

Figure 6.1: Image quality (VGAS) results for average patients reconstructed without any scatter rejection.

It is clear from Figure 6.1 that the VGAS is higher for lower tube potential (kVp) settings, demonstrating that image quality improves with lower potentials. VGAS ranged from 0.41 for 50 kVp, to 0.03 for 109 kVp, but there is very little difference between image quality at tube potentials greater than 102 kVp, as the VGAS ranged from 0.03 at 109 kVp, to 0.07 at 150 kVp. Although there is a trend to increased image quality at low tube potentials ($r^2 > 0.71$), it was only possible to statistically distinguish between 50 and 90 kVp (p-value = 0.033), 80 and 90 kVp (p-value = 0.033), and 80 and 109 kVp (p-value = 0.034). All other combinations demonstrated p-value ≥ 0.102. Nevertheless, these results show that for matched effective dose, image quality improves with lower tube potentials, which is similar to that reported in the literature [27, 29, 36]. However, in lowering tube potential, an increase in mAs would be needed to compensate the reduced intensity of X-ray photons. This is likely to necessitate increased exposure times. This is discussed in Section 6.7. There are a few

reasons why low tube potentials provide superior image quality for CR radiography. Firstly, as discussed in Chapter 1, the efficiency of the CR phosphor is higher for lower photon energies. A reduction in tube potential leads to the production of more low energy photons; one would therefore assume more will be detected. Secondly, there will be less Compton interactions and more photoelectric absorption in the patient, leading to increased radiation contrast and less scatter, both of which result in improved *image* contrast.

6.6.2 Image quality of 20 average patients reconstructed with an anti-scatter grid

The results are shown in Figure 6.2. For average patients, reconstructed with an anti-scatter grid, the results demonstrate there is a trend ($r^2 > 0.58$) of increased image quality to low tube potentials, as per the results in Section 6.6.1. VGAS results ranged from +0.039 (60 kVp) to -0.279 (141 kVp). However, there was no statistically significant differences between any of the observations (p-values ≥ 0.07).

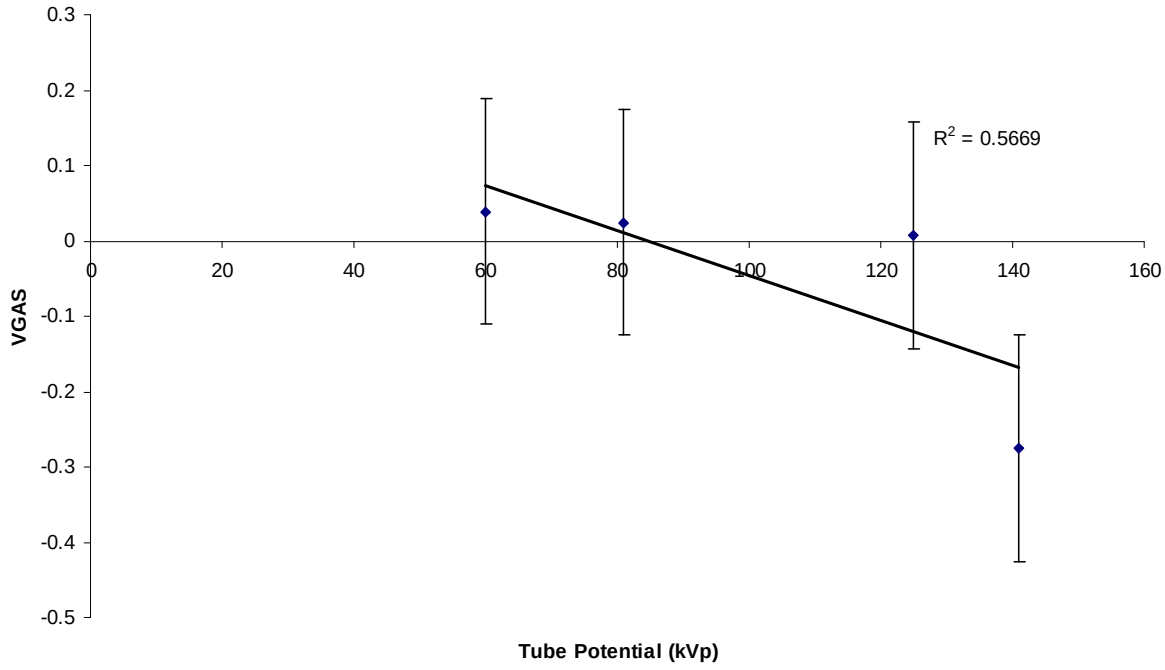

Figure 6.2: Image quality (VGAS) results for average patients reconstructed with an anti-scatter grid.

6.6.3 Image quality of 20 average patients reconstructed with an air gap technique

The results are illustrated in Figure 6.3. For average patients, reconstructed with an air gap, the results demonstrate a trend ($r^2 > 0.78$) of increased image quality to higher tube potentials. VGAS results ranged from -0.297 (60 kVp) to +0.004 (125 kVp) and were all negative (except 125 kVp), suggesting all of the tube potentials were slightly inferior to 102 kVp. However, there was no statistically significant differences between any of the observations (p-values ≥ 0.08). This result is the opposite of that found previously. This may be due to the fact that the air gap is just large enough so that at higher tube potentials the increased scatter misses the CR phosphor leading to relative improved image quality at these higher potentials; as half of the image evaluators were used to

working with 80 kVp clinically, perhaps they were giving higher scores to those images they were used to.

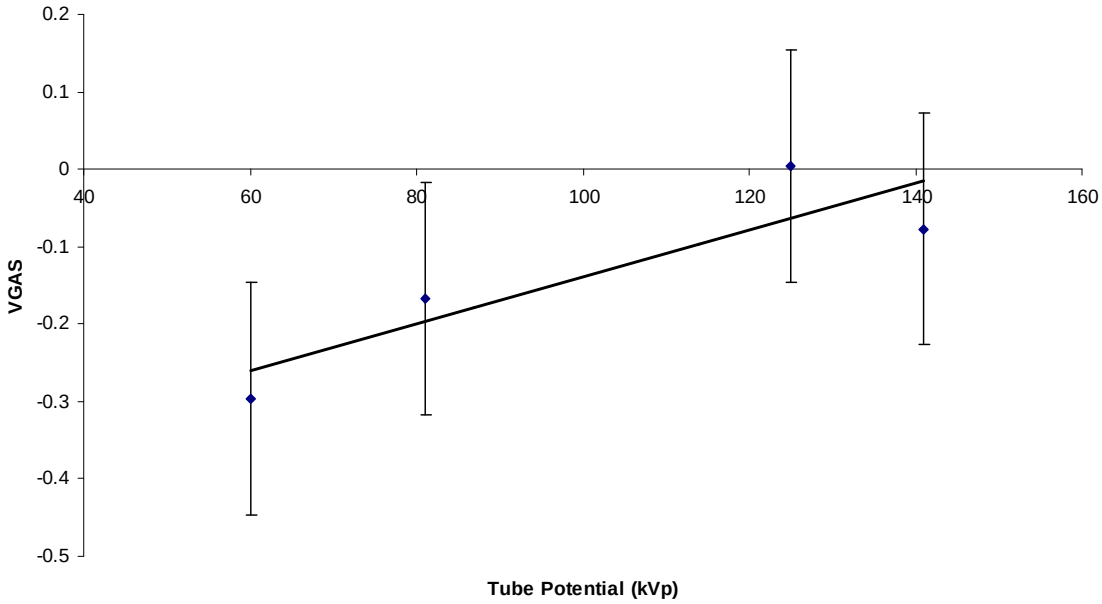

Figure 6.3: Image quality (VGAS) results for average patients reconstructed with an air gap scatter rejection technique.

6.6.4 Image quality of 20 obese patients reconstructed without scatter rejection methods

Figure 6.4 shows the results for obese patients reconstructed without any scatter rejection. There is a trend ($r^2 > 0.73$) of increased image quality to higher tube potentials (min VGAS = -0.088, max VGAS = +0.02) which is the reverse of that found with average patients. There was no statistically significant differences between any of the observations (p-values ≥ 0.06). However, it is likely that the poor penetration at low tube potentials and increased scatter from the obese patients outweighs the benefits of inherent image contrast due to photoelectric absorption. It should be remembered that the mean scatter

fraction (i.e. the ratio of scatter to total radiation incident at the image receptor) in chest radiography changes very little with tube potential [82, 86] and so the poorer image quality seen here at lower tube potentials depends more on patient size than beam quality.

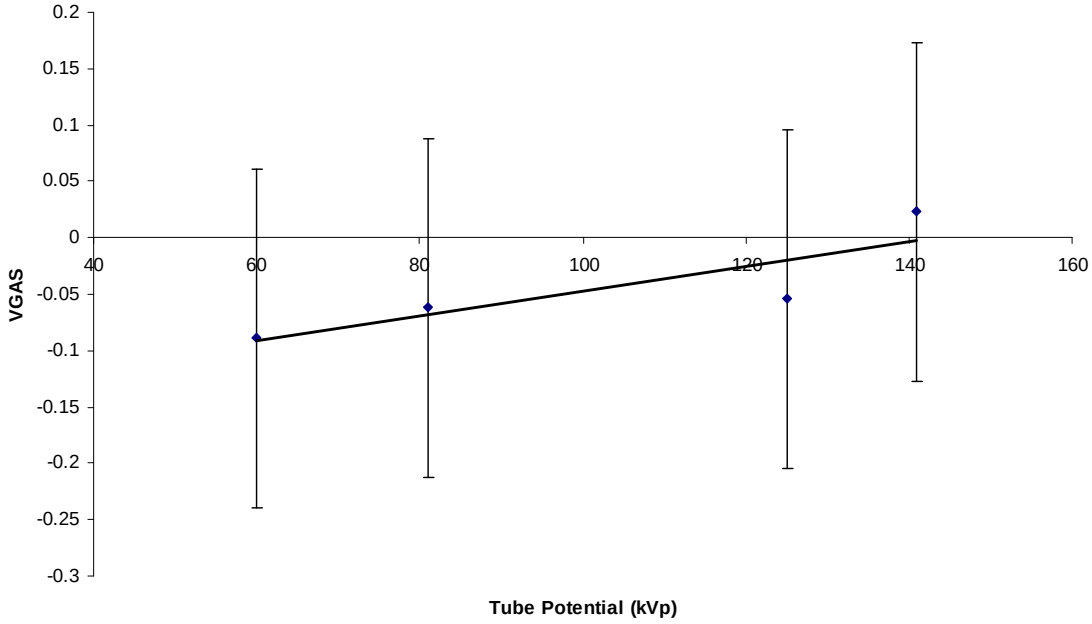

Figure 6.4: Image quality (VGAS) results for obese patients reconstructed without any scatter rejection.

6.6.5 Image quality of 20 obese patients reconstructed with an anti-scatter grid

The results are illustrated in Figure 6.5. They demonstrate that for obese patients reconstructed with an anti-scatter grid there is a strong trend ($r^2 > 0.95$) of increased image quality to low tube potentials. VGAS results ranged from +0.61 (60 kVp) to -0.330 (141 kVp) and there were clinically significant differences between tube potentials 60 and 141 kVp ($p < 0.029$), 81 and 141 kVp ($p < 0.012$), and 125 and 141 kVp ($p < 0.017$). These results suggest that if there is any advantage in using anti-scatter grids for obese patients, the lower

the tube potential the better, and those lower than the reference (102 kVp) would provide superior image quality. However, low tube potentials will increase exposure times, but this is addressed in Section 6.7.

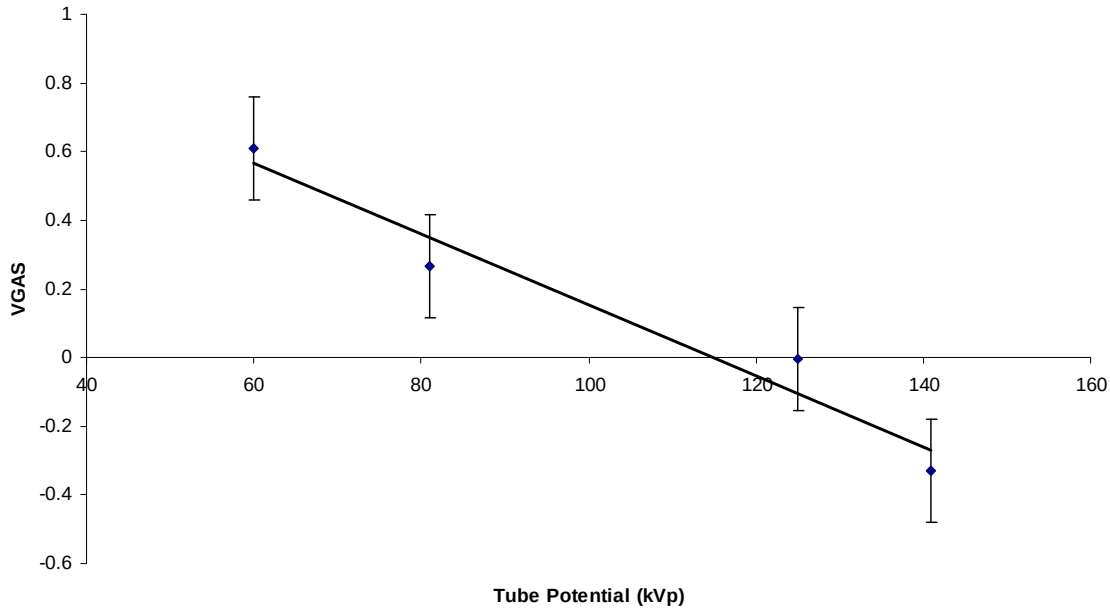

Figure 6.5: Image quality (VGAS) results for obese patients reconstructed with an anti-scatter grid.

6.6.6 Image quality of 20 obese patients reconstructed with an air gap technique

The results are shown in Figure 6.6. They demonstrate that for obese patients reconstructed with an air gap anti-scatter rejection method there is a weak trend ($r^2 > 0.55$) of increased image quality to higher tube potentials, but all average VGAS scores are negative, ranging from -1.42 (60 kVp) to -0.075 (141 kVp). These results suggest image quality is not improved irrespective of the tube potential used. Tube potential 60 kVp was significantly different from all the others ($p \leq 0.004$) but there was no significant difference between any others (p

≥ 0.201). Poorer image quality at low kVps is probably due to the factors mentioned in Section 6.6.4.

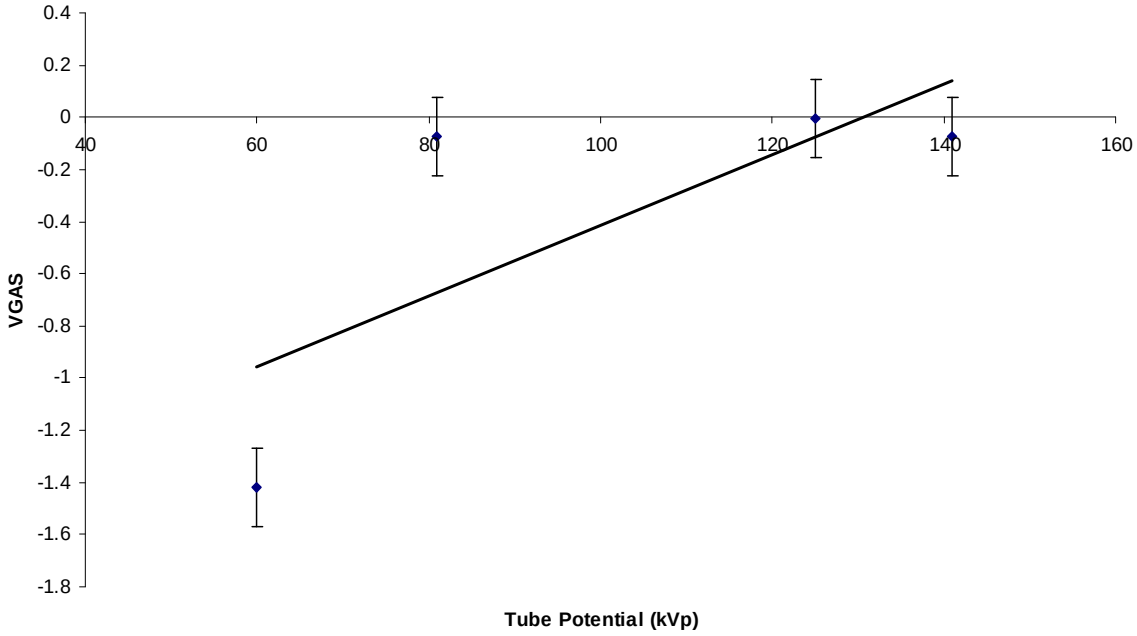

Figure 6.6: Image quality (VGAS) results for obese patients reconstructed with an air gap technique.

6.6.7 Anti-scatter grid v air gap – average patients

Figure 6.7 shows a bar graph of the average VGAS for images reconstructed with an anti-scatter grid and air gap technique (point 6 from Section 6.2).

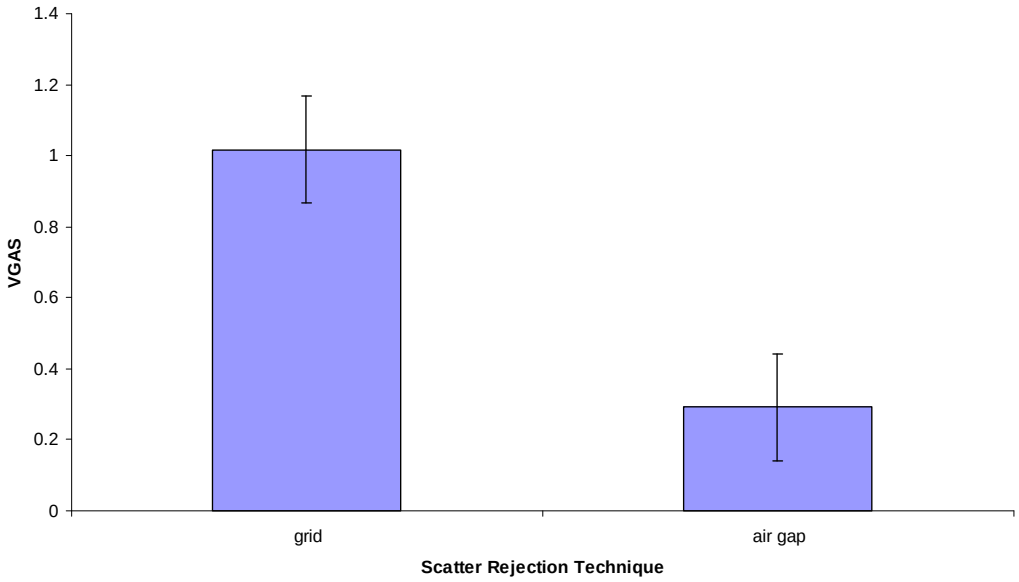

Figure 6.7: Image quality (VGAS) results for average patients for scatter rejection v non-scatter rejection (no scatter rejection VGAS = 0).

It is clear from Figure 6.7 that image quality is superior using scatter rejection techniques, and a grid outperforms the air gap method (note that as the image reconstructed without any scatter rejection was the reference image, it would have a VGAS of zero in Figure 6.7). However, only the use of an anti-scatter grid yielded significant improvements (air gap (p-value < 0.006) and no scatter rejection (p-value <0.0002)); the air gap method was not significantly different from no scatter rejection (p = 0.06). However, image evaluators gave a 'no' response in 100% of the cases when asked if the increase in dose justified using an anti-scatter technique. *Therefore, scatter rejection methods are not indicated for average patients.* As discussed in Section 6.2, the results shown here may be restricted in the fact that a contant lgM of 2.00 was used to derive the increase in dose necessary. If a constant SNR method is used is may be possible to use a grid with lower doses – this is addressed in Chapter 7.

6.6.8 Anti-scatter grid v air gap – obese patients

Figure 6.8 shows the same as Figure 6.7 but for obese patients.

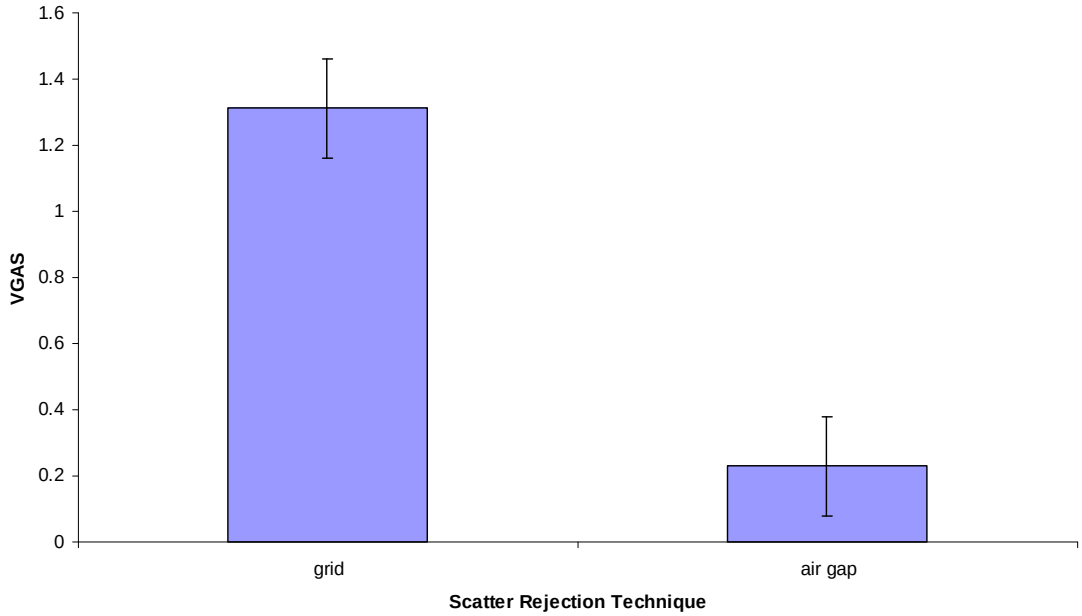

Figure 6.8: Image quality (VGAS) results for obese patients for scatter rejection v non-scatter rejection (no scatter rejection VGAS = 0).

The results are similar to that for average patients, in that image quality is superior using scatter rejection techniques, and a grid outperforms the air gap technique (note that as the image reconstructed without any scatter rejection was the reference image, it would have a VGAS of zero in Figure 6.8). The grid method was significantly different from the air gap method (p-value < 0.0004) and no scatter rejection (p-value <0.0002), and the air gap method was significantly different from no scatter rejection (p < 0.006). Interestingly, image evaluators gave a 'yes' response to 100% of the patients reconstructed with an anti-scatter grid (presented blindly) when asked if the increase in dose was justified, but gave a 'no' response to 100% of the air gap patients. Therefore, use of an anti-scatter grid is indicated for obese patients, but an air gap technique is not, and so the question of exposure times with lower tube

potentials (Section 6.6.5 deemed low tube potentials superior for obese patients imaged with an anti-scatter grid) becomes important.

6.7 Phantom experiment to determine minimum X-ray exposure times possible

The results demonstrated that when average RANDO was exposed without using a grid, it was possible to use a tube current of 630 mA and exposure times < 20 ms to obtain a lgM of 2.00 ± 0.05 with all tube potentials. Therefore it is possible to use low tube potentials as indicated in this work for average patients. However, for obese RANDO imaged with a grid, it was only possible to achieve the 20 ms limit with 109 kVp at 180 cm FDD or 90 kVp at 115 cm FDD (the shortest distance permitted for this focused grid).

6.8 Image quality of 10 patients reconstructed with different receptor air kerma values

The response of each image evaluator was quite surprising in that all deemed image quality acceptable in the lung region down to 1.5 µGy (i.e. all answers were ‘yes’). This is approximately one quarter of the air kerma of that used for patients exposed with our Radiology department’s current standard exposure protocol. However, image quality was not deemed acceptable in the spine and diaphragm regions for an air kerma (through the lung region) of 2 µGy for any of the structures and nodules mentioned in Table 6.2. This suggests chest imaging can be carried out with at least a 50% decrease in receptor dose (and therefore patient dose for a given X-ray beam quality), whilst maintaining image quality that is diagnostically acceptable. A similar result has recently been reported by Veldkamp *et al.* [111].

6.9 Interference of rib contrast

Ribs interfered in approximately 5% of images reconstructed at 50 kVp, but in no other. It is therefore not a limiting factor in chest radiography.

6.10 Conclusions

A DRR based computer simulation of CR chest radiographs that contain clinically realistic anatomical noise has been used to identify optimum radiographic techniques for CR chest radiography with an Agfa CR 85 imaging system and MD-4.0 phosphor plates. Simulated images scored by four experienced image evaluators have shown that for average adult patients scatter rejection is *not* indicated because the increase in dose is not justified; low tube potentials (< 102 kVp) therefore provide superior image quality.

It has been shown that for obese adult patients an oscillating focused anti-scatter grid *is* indicated, and should be used in conjunction with the lowest tube potential possible. Measurements with a chest phantom on a clinical X-ray system demonstrate any tube potential can be used without a grid for average patients to satisfy the 20 ms exposure time limit, but at least 90 kVp must be used for obese patients (with a grid). It has also been shown that receptor air kerma through the lung region can be reduced to 2 µGy whilst maintaining an adequate level of image quality, and the rib contrast interfering with image evaluation is minimal.

6.11 Recommendations for HEY Radiology Department

Clinical chest exposure factors within our Trust were not standardised at the time of this study; the following factors were typically used at each hospital site (it should be noted that the same model of CR system is used at each hospital):

Hull Royal Infirmary:	60 kVp/10 mAs
Princess Royal Hospital:	70 kVp/5 mAs
Beverley Westwood Hospital:	70 kVp/5 mAs
Castle Hill Hospital:	80 kVp/5 mAs

Optimisation of medical exposures given to patients to produce clinical images is a legal requirement under the relevant legislation (IR(ME)R2000); this includes standardising exposure factors. The results of this work recommend the following exposure factors:

<u>Average Patients:</u>

- Scatter rejection *is not* indicated
- Standardise exposure factors across the Trust using 60 kVp and 10 mAs
- After a 'settling in period' of standardisation, drop from 10 to 8 mAs (it should be remembered here that more than 90% of the X-ray tubes in HEY Radiology are of the Philips model (see Chapter 4) and are all of similar output; on average 25.2 ± 2.0 µGy/mAs at 1m).

<u>Obese Patients:</u>

- Standardise exposure factors across the Trust:

- Use focused oscillating anti-scatter grid with an FDD of 115 cm (lowest permissible for focused grid), 90 kVp, 550 mA and exposure time 20ms for lgM = 2.00

- If magnification is not acceptable with the above factors use focused oscillating anti-scatter grid with an FDD of 140 cm (recommended distance for focused grid), 102 kVp, 490 mA and exposure time 20ms for lgM = 2.00

- If magnification is still not acceptable with the above factors use focused oscillating anti-scatter grid with an FDD of 180 cm (maximum allowed distance for grid), 109 kVp, 450 mA and exposure time 20ms for lgM = 2.00

6.12 Change in Radiology practice

The results of this study were presented at the HEY Radiologist Operational Group meeting in March 2011. It was agreed with the Radiologists to change current exposure protocol to reflect the recommendations illustrated above. Chest imaging exposure factors and radiographic technique have now been standardised across the HEY Trust.

6.13 Implications of patient dose due to change in practice

The recommendations suggested in this thesis have led to the Radiology department in HEY Trust standardising exposure factors. It was therefore felt important to quantify the change in patient effective dose due to this change (although image quality has been deemed superior at the lower tube potential,

one could argue that it would not be ethical to change exposure factors if patient dose had increased significantly).

To assess any change in patient dose, the software package PCXMC [107] was used. Tube potentials of 60 and 80 kVp, total filtration of 3.1 mm Al and measured (in-house) entrance air kerma values of 70 and 67 μGy for 60 and 80 kVp respectively were entered into the relevant fields of the software (entrance air kerma values are correct for 60 kVp & 10 mAs, and 80 kVp & 5 mAs respectively). The software calculated the results as shown in Table 6.4.

Tube Potential (kVp)	Tube Current-Time Product (mAs)	Effective Dose (mSv)
60	10	0.00755 ± 0.00007
80	5	0.01053 ± 0.00008

Table 6.4: Effective dose values for exposures with tube potentials 60 and 80 kVp respectively.

As can be seen in Table 6.4, the effective dose and therefore risk decreases with decreasing tube potential. The effective dose actually decreases by approximately 28%. This is probably due to a more pronounced fall off of depth dose in the patient PA direction with 60 kVp with respect to 80 kVp. This result is excellent, as at lower tube potentials, not only does image quality increase, but patient effective dose decreases. This strengthens the argument for the change in HEY Radiology practice at Castle Hill Hospital. Also, once exposure factors have been decreased from 60 kVp & 10 mAs to 60 kVp & 8 mAs, as per the above recommendations, the effective dose will decrease by a further 20%, leading to an overall effective dose reduction of approximately 43%, i.e. the dose reduction resulting from changing exposure factors 80 kVp & 5 mAs to 60 kVp & 8 mAs.

Chapter 7: Conclusions and future work

7.1 Conclusions

The work presented in this thesis has demonstrated that digitally reconstructed radiographs (DRRs) of the chest can be produced, validated and used to optimise radiographic techniques for chest radiography with Computed Radiography (CR) systems. The use of DRRs for optimisation studies has the benefit of eliminating the need for multiple X-ray exposures of the same patient, thus minimising the risk of developing radiation induced cancers. The methodology outlined in this thesis can also be used to produce DRR images of digital imaging modalities other than CR, such as direct digital systems (DR), assuming the physical characteristics of those systems are measured.

At the time of writing, this work has led to the acceptance for publication of two full papers in the British Journal of Radiology (BJR), and two peer reviewed presentations at UK meetings. Work related to this thesis performed in the workup towards this study has also produced three full papers and two presentations.

Initially the work focused on identifying which DRR computation method would be most appropriate. Ray casting and splatting are both popular ways of reconstructing DRRs, but it was decided to use the ray casting method with

pencil X-ray beams as this produces images with superior quality, namely less blurring and aliasing. One disadvantage is that of long computation time, but this is not anticipated to be a concern for the end user. The use of Matlab, which is a developmental prototyping environment, if replaced by computer code written in an appropriate language, would reduce this overhead.

Ray casting DRR methods do not compute any component of radiation scatter at the image receptor, so the most appropriate method to add scatter was investigated. A method widely used in the literature was adopted to calculate scatter and scatter fractions, allowing the addition of this component to the DRR images. This was performed for tube potentials across the diagnostic energy range with an average and obese phantom.

A robust and reliable methodology to produce DRRs of the chest has been presented in this thesis. CT data of RANDO and real patients have been used as computerised voxelated phantoms, and virtual X-ray spectra (produced with commercial software) were successfully projected through the data to produce raw DRR images of the chest. The raw DRR images were presented in terms of energy absorbed (keV) by the CR phosphor material (DRR pixel values linear with absorbed energy), by deriving CR photon energy dependent mass energy absorption coefficients, available from the NIST database. To speed up this calculation, CT pixel value data were converted to linear attenuation coefficient (LAC) for reference photon energies only, and the CT voxels were shifted to allow parallel X-ray pencil beam projection.

Experimental methods were used to derive scatter and scatter fractions of the chest with the RANDO phantom on the clinical X-ray and CR systems for which the DRR algorithm is currently configured. Measured scatter fractions

adequately simulated those found in the human thorax across the diagnostic tube potential range. Frequency dependent noise was also added post raw DRR production based on a slightly modified method described in a recent publication [49].

Lung nodules were successfully simulated and added to the CT data prior to X-ray pencil beam projection, and these were visible on the resulting DRR images. The visualisation of lung lesions is dependent on location, so two lesions were simulated, one in the lateral pulmonary region, and one in the hilar region.

An investigation of noise in the CT data found that it took a Gaussian form. However, a Guassian noise smoothing filter was not required prior to pencil beam projection as this made < 2% difference to the signal and noise characteristics of the DRR image. It is assumed that averaging and sampling of CT voxels (during pencil beam projection) and added frequency dependent noise almost completely mask the effect of CT noise in the final image.

Validation of the DRR model was very successful. It was initially carried out by visually inspecting DRR images of RANDO and patients, and comparing them with real CR images. In all cases on *low* resolution workstations, it was impossible to distinguish the DRR and CR images (although a difference was apparent on reporting monitors).

Measures of signal to noise ratio (SNR), tissue to rib ratio (TRR) and dynamic range were then used to quantitatively validate the DRR algorithm. For average and obese RANDO, measurements of SNR with RANDO differed from the CR image by a maximum of 15%, but the mean values in each chest region were always within good agreement. DRR TRR values were within 0.6% of CR

measurements. It was found that pixel value histograms were very similar in shape, especially for average RANDO, and dose escalation/reduction resulted in histograms that shifted correctly along the pixel value axis. It was observed for all DRR images that the dynamic range was smaller than CR, probably due to limited tissue modelling and voxel size. This effect was more pronounced in obese RANDO which is likely to be due to inadequate modelling of the grocery store lard (i.e. conversion to its correct LAC). Nevertheless, SNR, TRR, and in general, histograms, were very well matched.

The same methods were used to validate the algorithm by comparing CR and DRR images of average and obese patients. Images from other hospital sites were also used. Results were similar to that of RANDO, in that SNRs and TRRs were very well matched and histograms were of the same shape but reduced dynamic range. Nevertheless, minimum and maximum DRR pixel values were very similar to CR minima and maxima (all within two standard deviations of the mean). As images from other hospital sites were used and shown to be very well matched to DRR images, the robustness and transferability of the model has been demonstrated. In other words, the DRR model can be used to simulate clinical systems other than the one that it was originally configured to, without any further work.

Qualitative validation by expert image evaluators was carried out using simple scoring methods. All images were deemed good enough to carry out optimisation studies with. This was an excellent outcome, as it demonstrates the main research objective of this work has been fulfilled.

Limitations of the DRR computer model were discussed, the most important ones being the addition of frequency dependent noise and the reduced spatial

resolution. The addition of noise assumes that DQE is constant with dose (i.e. is constant in all regions of the image) and ergodic. Neither of these conditions is true. However, it has been demonstrated in the literature that these effects are negligible at clinical doses, and the results presented here certainly seem to support this, as SNRs are very well matched. The reduction in spatial resolution means object information (and noise) will never be as faithfully reproduced in a DRR image when compared to CR. This was noticeable when DRR images were presented on high resolution diagnostic reporting monitors. However, expert image evaluators did not feel this was a limiting factor. A further limitation is the computation time. Although methods were introduced to speed up DRR image production, each image takes approximately one hour to calculate. However, this is not a limiting factor for the end user (image evaluators), and it is likely that computing power and speed will increase in future, eventually resolving this issue.

CR chest optimisation studies were subsequently carried out using images produced with the validated DRR software. Expert image evaluators scored images presented at different tube potentials, receptor air kermas and scatter rejection techniques. Scoring criteria were based on slightly modified European criteria. For average patients, low tube potential (< 102 kVp) without scatter rejection was indicated. Similarly, for obese patients, low tube potentials were indicated but with the use of an anti-scatter grid. However, due to the use of a grid, the lowest tube potential possible that did not necessitate exposure times > 20 ms was 90 kVp.

These conclusions were presented to the HEY Radiology Operational Group and it was agreed that chest exposures across the Trust would be standardised, using the exposure factors recommended and discussed in Chapter 6. This

work has therefore changed clinical practice within the HEY NHS Trust, optimised image quality and lowered the effective dose given to our patients.

In summary, an environment has been developed to create virtual patient CR images from clinical CT data which can be used to perform virtual trials which might well be deemed unethical on real patients. It is anticipated that the method presented in this thesis can also be expanded to non-CR imaging modalities, such as direct digital flat panel detectors.

7.2 Future work

The work presented in this thesis has focussed on optimising X-ray beam qualities currently used in the HEY Trust, i.e. tubes containing aluminium filtration only. However, an interesting area of development would be to produce DRR images with added copper filtration to see if it would be possible to lower patient dose whilst maintaining image quality.

It is anticipated that the methodology presented in this thesis can be used to produce simulated images of any anatomical region, assuming adequate access to patient CT data and scatter maps can be derived. Future work would include producing DRRs of the pelvis/abdomen region and subsequently optimising radiographic technique for this area. This specific anatomical region is worth investigating as it is one of the largest patient dose examinations for 2-D imaging (approximate effective dose = 1 mSv).

Other potential areas of research would be to upload DRR images onto a CR system and subject the images to clinical post-processing. This would enable the optimisation of the entire imaging chain as well as X-ray radiographic techniques.

Another interesting area of work would to redo the image evaluation with average patient images reconstructed with an anti-scatter grid at much *lower* doses. It has been shown by Fetterly and Schueler [113] that approximately 76% of the primary radiation is transmitted by the grid. Therefore, for digital imaging modalities where signal-to-noise (not optical density as per film-screen) is the limiting factor one could argue exposures with a grid only need increasing by the reciprocal of this amount (i.e. 1.3), assuming only small amounts of scatter are recorded by the receptor. The dose increases described in Section 6.2 followed a ‘constant lgM’ approach, akin to using a constant optical density for film-screen.

References

1. Schaetzing R, Whiting BR, Lubinsky AR & Owen JF (1990). Digital radiography using storage phosphors, in Digital Imaging in Diagnostic Radiology, edited by J.D. Newall and C. A. Kelsey (Churchill Living Stone pp 107-138)
2. Sonada M, Takano M, Miyahara J & Kato H. Computed radiography utilizing scanning laser stimulated luminescence, Radiology 1983;2:833-48.
3. Samei E, Seibert JA, Willis CE, Flynn MJ, Mah E & Junck KL. Performance evaluation of computed radiography systems. Med. Phys. 2001;28(3):361-71.
4. Barnes GT. Digital x-ray image capture with image intensifier and storage phosphor plates: imaging principles, performance and limitations. Digital Imaging (AAPM Monograph no 22), edited by W.R. Hendee and J.H. Trueblood (Madison: Medical Physics) 1993;23-48.
5. Council of European Communities. European guidelines on quality criteria for diagnostic radiographic images, Rep EUR 16260 EN, 1996.
6. ICRP. The 2007 Recommendations of the International Commission on Radiological Protection. 2007 ICRP Publication 103 (Oxford:Pergamon)
7. Ionising Radiation (Medical Exposure) Regulations 2000. UK: The Stationary Office Limited, 2000.
8. Mackenzie A and Honey ID. Characterization of noise sources for two generations of computed radiography systems using powder and crystalline photostimulable phosphors. Med. Phys. 2007;34:3345-57.
9. Dobbins JT III, Samei E, Ranger NT and Chen Y. Intercomparison of methods for image quality characterization. II. Noise power spectrum. Med. Phys. 2006;33(5):1466-75.
10. Bochud F O, Valley J F, Verdun F R, Hessler C and Schnyder P. Estimation of the noisy component of anatomical backgrounds. Med. Phys. 1999;26:1365-70.

11. Samei E, Flynn M J and Eyler W R. Detection of subtle lung nodules: relative influence of quantum and anatomical noise on chest radiographs. Radiology 1999;213:727-34.
12. Samei E, Eyler W and Baron L. Effects of anatomical structure on signal detection. In: J, Kundel H L and Van Metter R L, editors. Handbook of Medical Imaging. Vol. 1. Physics and Psychophysics. Beutel: 2000:655-82.
13. Burgess A E, Jacobson F L and Judy P F. Human observer detection experiments with mammograms and power-law noise. Med. Phys 2001;28:419-37.
14. Huda W, Ogden K M, Scalzetti E M, Dudley E F and Dance D R. How do radiographic techniques affect mass lesion detection performance in digital mammography. Proc. SPIE 2004;5372:372-82.
15. Keelan B W, Topfer K, Yorkston J, Sehnert W J and Ellinwood J S. Relative impact of detector noise and anatomical structure on lung nodule detection. Proc. SPIE 2004;5372:230-41.
16. Sund P, Bath M, Kheddache S and Mansson L G. Comparison of visual grading analysis and determination of detective quantum efficiency for evaluation system performance in digital chest radiography. Eur. Radiol. 2004;14:48-58.
17. Tingberg A, Bath M, Hakansson M, Medin J, Sandborg M, Alm-Carlsson G, Mattsson S and Mansson L G. Comparison of two methods for evaluation of image quality of lumber spine radiographs Proc. SPIE 2004;5372:251-62.
18. Hakansson M, Bath M, Borjesson S, Kheddache S, Allansdotter Johnsson and Mansson L G. Nodule detection in digital chest radiography: effect of system noise. Radiat. Prot. Dosim. 2005;114(1-3):97-101.
19. Bath M, Hakansson M, Borjesson S, Kheddache S, Grahn A, Bochud F O, Verdun F R and Mansson L G. Nodule detection in digital chest radiography: introduction to the RADIUS chest trial. Radiat. Prot. Dosim. 2005;114(1-3):85-91.
20. Bath M, Hakansson M, Borjesson S, Kheddache S, Grahn A, Bochud F O, Verdun F R and Mansson L G. Nodule detection in digital chest radiography: part of image background acting as pure noise. Radiat. Prot. Dosim. 2005;114(1-3):102-8.

21. Bath M, Hakansson M, Borjesson S, Hoeschen C, Tischenko O, Kheddache S, Vikgren J, and Mansson L G Nodule detection in digital chest radiography: effect of anatomical noise. Radiat. Prot. Dosim. 2005;114(1-3):109-13.
22. Hakansson M, Bath M, Borjesson S, Kheddache S, Grahn A, Ruschin M, Tingberg A, Mattsson S and Mansson L G. Nodule detection in digital chest radiography: summary of the RADIUS chest trial. Radiat. Prot. Dosim. 2005;114(1-3):114-20.
23. Cook LT, Insana MF, McFadden MA, Hall TJ, Cox GG. Comparison of low contrast detectability of a screen film system and third generation computed radiography. Med.Phys. 1994;21:691-5.
24. Schaefer-Prokop CM, Prokpop M, Schmidt A, Neitzel U, Galanski M. Selenium radiography versus storage phosphor and conventional radiography in the detection of simulation of chest lesions. Radiology. 1996;201:45-50.
25. Hufton AP, Doyle SM, Carty HML. Digital radiography in paediatrics: radiation dose considerations and magnitude of possible dose reduction. Br. J. Radiol. 1998;71:186-99.
26. Nobuhiro O, Hajime N, Seiichi m, Kunihiro T, Katsumi N, Akira Y (1996). Optimal beam quality for chest radiography. Invest Radiol;31:126-31.
27. Dobbins JT, Rice JJ, Beam CA, Ravin CE. Threshold perception performance with computed radiography: implications for chest radiography. Radiology. 1992;183:179-87.
28. Tingberg A, Sjostrom D. Search for optimal tube voltage for image plate radiography. Invest. Radiol. 1996;31:126-31.
29. Chotas HG, Floyd CE, Dobbins JT, Ravin CE. Digital chest radiography with photostimulable phosphors: signal-to-noise ratio as a function if kilovoltage with matched exposure risk. Radiology. 1993;186:395-8.
30. Honey I D, MacKenzie A and Evans D S. Investigation of optimum energies for chest imaging using film-screen and computed radiography. Br. J. Radiol. 2005;78:422-27.
31. Doyle P, Martin CJ, and Gentle D. Application of contrast-to-noise ratio in optimizing beam quality for digital chest radiography: comparison of experimental measurements and theoretical simulations. Phys. Med. Biol. 2006;51:2953-70.

32. Chotas HG, Floyd CE, Johnson GA and Ravin CE. A quality control phantom for chest radiography. Radiology. 1997;202:111-16.
33. Doyle P, Martin CJ, and Gentle D. Dose-image quality optimisation in digital chest radiography. Radiat. Prot. Dosim. 2005;114: 269-272.
34. C.S. Moore, A.W. Beavis and J.R. Saunderson. Investigation of optimum X ray beam tube voltage and filtration for chest radiography with a Computed Radiography system. Br. J. Radiol. 2008;81:771-77.
35. Conway BJ, Butler BF, Duff JE, Fewell TR, Gross RE, Jenning RJ, et al. Beam quality independent attenuation phantom for estimating patient exposure from x ray automatic exposure controlled chest examinations. Med. Phys. 1984;11:827-32.
36. Sandborg M, Tingberg A, Ullman G, Dance D R, Alm Carlsson G. Comparison of clinical and physical measures of image quality in chest and pelvis computed radiography at different tube voltages. Med. Phys. 2006;334:169-75.
37. Ullman G, Sandborg S, Dance D R, Hunt R A and Carlsson G A. Towards optimization in digital chest radiography using Monte Carlo modelling. Phys. Med. Biol. 2006;51:2729-43.
38. Mansson LG. Evaluation of radiographic procedures (1994). Investigations related to chest imaging. Thesis, Goteburg University, Goteburg
39. Mansson LG. Methods for the evaluation of image quality: a review. Radiat. Prot. Dosim. 2000;90:89-99.
40. Tingberg A. Quantifying the quality of medical X-ray images (2000). An evaluation based on normal anatomy for lumber spine and chest images. Thesis, Lund University, Luna.
41. Sund P, Herrmann C, Almen A and Mattsson S. Comparison of two methods for evaluating image quality in chest radiographs. Proc. SPIE 2000;3977:437-43.
42. Sandborg M, McVey G, Dance D R, Persliden J and Alm Carlsson G. A voxel phantom based monte carlo computer program for optimisation of chest and lumber spine x ray imaging systems. Radiat. Prot. Dosim. 2000;90(1-2):105-8.
43. Sandborg M, McVey G, Dance D R and Alm Carlsson G. Schemes for the optimization of chest radiography using a computer model of the patient and x-ray system. Med. Phys. 2001;28:2007-19.

44. Tingberg A and Sjostrom. Optimistaion of image plate radiography with respect to tube voltage. Radiat. Prot. Dosim. 2005;114:286-93.
45. Launders J H, Cowen A R, Bury R F and Hawkridge P. Towards image quality, beam energy and effective dose optimisation in digital thoracic radiography. Eur. Radiol. 2001;11:870-75.
46. Zubal G, Harrell C R, Smith E O, Rattner Z, Gindi G and Hoffer P B. Computerized three-dimensional segmented human anatomy. Med. Phys. 1994;21:299-302.
47. Fanti V, Marzeddu R, Massazza G, Randaccio P, Brunetti A and Golosio B. A simulator for X-ray images. Radiat. Prot. Dosim. 2005;114(1-3):350-54.
48. Bath M, Hakansson M, Hansson J and Mansson L G. A conceptual optimization strategy for radiography in a digital environment. Radiat. Prot. Dosim. 2005;114(1-3):230-35.
49. Bath M, Hakansson M, Tingberg A and Mansson L G. Method of simulating dose reduction for digital radiographic systems. Radiat. Prot. Dosim. 2005 ;114(1-3):253-59.
50. Malzbender T. Fourier volume rendering. ACM Trans. Graph 1993; 12(3):233-50.
51. Ntasis E, Maniatis TA and Nikita KS. Fourier volume rendering for real time preview of digital reconstructed radiographs: a web based implementation. Computerized Medical Imaging and Graphics 2002; 26(1):1-8.
52. Wang F, Davis TE and Vemuri BC. Real-time DRR generation using cylindrical harmonics. In Proceedings of the 5th International Conference on Medical Image Computing and Computer-Assisted Intervention-Part II, pages 671-8. Springer-Verlag 2002.
53. Milickovic N, Baltas D, Giannouli S, Lahanas M and Zamboglou N. CT imaging based digitally reconstructed radiographs and their application in brachytherapy. Phys. Med. Biol. 2000;45:2787-800.
54. Ntasis *et al.* Fourier volume rendering for real time preview of digital reconstructed radiographs a web based implementation. Comput. Med. Imaging Graph. 2002;26:1-8.

55. Bollett *et al.* Can digitally reconstructed radiographs replace simulation films in prostate cancer conformal therapy? Int. J. Radiat. Oncol. Biol. Phys. 2003;57:1122-30.
56. Lemieux L, Jagoe R, Fish DR, Kitchen ND, and Thomas DG. A patient to CT image registration method based on digitally reconstructed radiographs. Med. Phys. 1994;21:1749-60.
57. Penny GP, Weese J, Little JA, Desmedt P, Hill DL and Hawkes DJ. A comparison of similarity measures for use in 2D-3D medical image registration IEEE Trans Med. Imaging 1998;17:586-96.
58. Bullitt E, Liu A, Aylward SR, Coffey C, Stone J, Mukherji SK Muller SK and Pizer SM. Registration of 3D cerebral vessels with 2D digital angiograms: clinical evaluation. Acad. Radiol. 1998;6:539-46.
59. Birkfellner *et al.* A faster method for 3D/2D medical image registration – a simulation study. Phys. Med. Biol. 2003;48:2665-79.
60. Siddon RL. Fast calculation of the exact radiological path for a three-dimensional CT array. Med. Phys. 1985;12(2):252-55.
61. Sherouse GW, Novins K, Chaney EL. Computation of digitally reconstructed radiographs for use in radiotherapy treatment design. Int. J. Rad. Onc. Biol. Phys.1991;18(3):651-58.
62. Dong L, Kachilla D, Boyer AL. A ray tracing program based on Bresenham algorithm for generating digitally reconstructed radiographs. Med. Phys. 1994;21:886.
63. Dong L, Boyer AL. An image correlation procedure for digitally reconstructed radiographs and electronic portal images. Int. J. Rad. Onc. Biol. Phys. 1994;33:1053-60.
64. Higgins P D, Attix F H, Hubbell J H, Seltzer S M, Berger M J and Sibata C H. Mass energy transfer and mass energy absorption coefficients, including in flight positron annihilation for photon energies 1 keV to 100 MeV. National Institute of Standards and Technology, NISTIR 4680. Gaithersburg, MD: US Department of Commerce, 1992.
65. Berger MJ and Hubbel JH 1987 XCOM: photon cross sections on a personal computer NBSIR 87-3597.
66. Killoran JH *et al*. A technique for optimization of digitally reconstructed radiographs of the chest in virtual simulation. Int. J. Rad. Onc. Biol Phys. 2001;49:231-39.

67. Giraud J, Killoran J, Svensson G. Quality control, and optimization of DRR. A phantom study (Abstr). Med. Phys. 1998;25(7):A163.
68. Bifulco P, Sansone M, Cesarelli M, Allen R and Bracale M. Estimation of out-of-plane vertebra rotations on radiographic projections using CT data: a simulation study. Medical Engineering & Physics 2002;24:295-300.
69. Mao X. Splatting on non-rectilinear volumes through stochastic resampling. IEEE Transactions on Visualization and Computer Graphics, 1996;2(2):156-70.
70. Westover L A. Interactive Volume Rendering, In: proceedings of Volume Visualisation Workshop (Chapel Hill, NC, May 18-19), Dept of Computer Science, University of North Carolina, Chapel Hill, NC, 9-16, 1989.
71. Westover L A. Footprint evaluation for volume rendering. Computer Graphics (Proceedings of SIGGRAPH), 1990;24(4):367-76.
72. Laur D and Hanrahan P. Hierarchial Splatting: A Progressive Refinement Algorithm for Volume Rendering. Computer Graphics (Proceedings of SIGGRAPH), 1991;25(4):285-88.
73. Crawfis R and Max N. Texture splats for 3D scalar and vector field visualisation. In: proceedings of Visualisation (San Jose, CA, October 25-29), IEEE Computer Society Press 261-266, 1993.
74. K. Mueller, and R. Yagel, "Fast perspective volume rendering with splatting by using a ray-driven approach," *Proc. of IEEE Visualization'96*, pp. 65-72, 1996.
75. Birkfellner W *et al*. Wobbled splatting – fast perspective volume rendering method for simulation of x-ray images from CT. Phys. Med. Biol. 2005;50:N73-N84.
76. Laur D and Hanrahan P. Hierarchial Splatting: A Progressive Refinement Algorithm for Volume Rendering. Computer Graphics (Proceeding of SIGGRAPH), 1991;25(4):285-88.
77. Alakuijala J *et al.* (1996). Reconstruction of digital radiographs by texture mapping, ray casting and splatting. 18th Annual International Conference of the IEEE Engineering in medicine and Biology Society, Amsterdam. 3.1.1: Visualisation and Image-guided Surgery.
78. Swan E II, K.Mueller, T.Möller, N.Shareef, R.A.Crawfis, R.Yagel, "An Anti-Aliasing Technique for Splatting" Proceedings of IEEE Conference on Visualization 1997;197-204.

79. Novins KL, Sillion FX and Greenberg DP. An efficient method for volume rendering using perspective projection. Computer graphics (Proceedings of San Diego Workshop on Volume Visualization), 1999;24(5):95-102.
80. Pfister H and Kaufman A. Cube-4 – A scalable architecture for real-time volume rendering. In proceedings of Visualization 96 (San Fransisco, CA, October 27 – November 1), IEEE Computer Society Press. 1996:47-54.
81. Niklason L T, Sorenson J A and Nelson J A. Scattered radiation in chest radiography. Med. Phys. 1981;8(5):677-81.
82. Quality and choice of Potter Bucky grids: I. A new method for the unambiguous determination of the quality of a grid Bowenkamp JG and Boldingh. Acta. Radiol. 1959;51:479.
83. Boone JM. Scatter correction algorithm for digitally acquired radiographs: Theory and results. Med. Phys. 1986;13(3):319-28.
84. Boone JM, Arnold BA and Seibert JA. Characterisation of the point spread function and modulation transfer function of scattered radiation using a digital imaging system. Med. Phys. 1986;13(6):967.
85. Floyd C E, Lo J Y, Chotas H G and Ravin C E. Quantitative scatter measurement in digital radiography using a photostimulable phosphor imaging system. Med. Phys. 1991;18(3):408-13.
86. Floyd C E, Baker J A, Lo J Y and Ravin C E. Posterior beam stop method for scatter fraction measurement in digital radiography. Invest. Radiol. 1992;27:119-23
87. Petrone TJ and Steidley KD. Rare-earth scatter fractions in chest radiography. Med. Phys. 1993;20(2):475-79.
88. Baydush A H, Ghem W C and Floyd C E. Anthropomorphic versus geometric chest phantoms: A comparison of scatter properties. Med. Phys. 2000;27(5):894-7.
89. Shrimpton P C, Wall B F and Fisher E S. The tissue-equivalence of the Alderson-Rando anthropomorphic phantom for x-rays of diagnostic qualities. Phys. Med. Biol. 1981;26(1):133-39.
90. Sutton DG, Reilly AJ. Report 78 Spectrum Processor. York: IPEM, 1997.
91. Cranley B J, Gilmore G W A, Fogarty A, Desponds L and Sutton D G. Catalogue of diagnostic X-ray spectra & other data. IPEM Report 78 (CD-ROM). York: IPEM, 1997.

92. Ay RA *et al.* Assessment of different computational models for generation of x-ray spectra in diagnostic radiology and mammography. Med. Phys. 2005;32(6):1660-75.
93. Gullberg G T. Reconstruction of fan beam data by filtering the back-projection. Comput. Graph. Image Process. 1979 ;10:30-47.
94. Suzuki S and Yamaguchi S. Comparison between an image reconstruction method of filtering back projection and the filtered back projection method. Appl. Opt. 1988;27:2867-70.
95. Nakano Y, Gido T, Honda S, Maezawa A, Wakamatsu H and Yanagita T. Improved computed radiography image quality from BaFI:Eu photostimulable phosphor. Med. Phys. 2002;29:592-7.
96. ICRU. Image quality in chest radiography ICRU Report 70 (Bethesda, MD:ICRU), 2003.
97. Li X, Samei E, Delong D M, Jones, R P, Gaca A M, Hollingsworth C L, Maxfield C M, Carrico C W T and Frush D P. Three-dimensional simulation of lung nodules for paediatric multidetector array. CT Br. J. Radiol. 2009;82:401-11.
98. Samei E, Flynn M J amd Eyler W R. Simulation of subtle lung nodules in projection radiography. Radiology. 1997;202:117-24.
99. Burgess A E, Li X, Abbey C K. Nodule detection in two component noise: toward patient structure. In: Kundel H L, editor. Proceedings of SPIE: medical imaging – image perception; 1997, February 26; Newport Beach, CA, USA. Bellingham, WA: International Society for Optical Engineering 3036 1997;2-13.
100. Gohagan JK, Marcus PM, Fagerstrom RM Pinsky PF, Kramer BS, Prorok PC, *et al*. Final results of the Lung Screening Study, a randomized feasibility study of spiral CT versus chest X-ray screening for lung cancer. Lung Cancer 2005;47 (1):9-15.
101. Samei E, Flynn M J and Reimann D A. A method for measuring the presampled MTF of digital radiography systems using an edge test device. Med. Phys. 1998;25:102–13.
102. Samei E, Ranger N T, Dobbins J T and Chen Y. Intercomparison of methods for image quality characterization 1. Modulation transfer function. Med. Phys. 2006;33(5):1454–65.

103. Hsieh J. Computed Tomography: Principles, Design, Artefacts and Recent Advances SPIE, Bellingham, WA, 1993.
104. Hilts M and Duzenli C. Image filtering for improved dose resolution in CT polymer gel dosimetry. Med. Phys. 2004;31(1):39-49.
105. Gonzalez R C and Woods R E. Digital Image Processing, 2nd Edition (Prentice-Hall, Upper Saddle River, NJ), 2001.
106. Institute of Physics and Engineering in Medicine Report 91: Recommended standards for the routine performance testing of diagnostic X-ray imaging systems. York, UK, 2005.
107. PCXMC Dose Calculations, Version 2.0.1, Radiation and Nuclear Safety Authority, Helsinki, Finland, 2008.
108. Lanhede B et al. The influence of different technique factors on image quality of chest radiographs as evaluated by modified CEC image quality criteria. Br. J. Radiol. 2002;75:38-49.
109. Tingberg A et al. The influence of the characteristic curve on the image quality of clinical radiographs. Br. J. Radiol. 2004;77:204-15.
110. Hoeschen C, Tischenko O, Buhr E, Illers H. Comparison of technical and anatomical noise in digital thorax X-ray images. Radiat. Prot. Dosimetry. 2005;114:75–80.
111. Veldkamp J H, Kroft L J M and Geleijns J. Dose and perceived image quality in chest radiography. Eu. J. Radiol. 2009;72:209-17.
112. Lacroute P and Levoy M. Fast volume rendering using a shear-warp factorization of the viewing transformation. Computer Graphic Proceedings. 1994;451-8.
113. Fetterly K A and Schueler B A. Experimental evaluation of fibre-interspaced antiscatter girds for large patient imaging with digital x-ray systems. Phys. Med. Biol. 2007;52:4863-80.