

DRAFT

Computer Assisted Diagnosis (CAD) for the Rapid Automated Measurement of Body Fat Tissue from Whole Body MRI

Paul F Whelan, Ovidiu Ghita, Kevin Robinson

Vision Systems Group
Dublin City University
Dublin, Ireland

Email: paul.whelan@eeng.dcu.ie, {ghitao, kevin.robinson}@eeng.dcu.ie

Abstract: This paper will examine the technical issues relating to the feasibility of using Computer Assisted Diagnosis (CAD) techniques to automatically identify, localize, and accurately measure body fat tissue from a rapid whole body MRI exam. The aim of this work is the provision of an automated system, which assesses subjects' whole body MRI scans and which provides numerical and visual feedback to illustrate the findings. The system generates real time results allowing for an initial assessment to be performed immediately following the completion of an MRI scan. The paper will focus on the specific issues relating to the formation, volume reconstruction, image processing and analysis of the whole body images. A working system and details of a prospective investigative study of 42 volunteers will be presented.

Key words: Computer Assisted Diagnosis, Volume reconstruction, Whole Body MRI, Body Fat

1. Introduction

The accurate determination of a person's total body fat has become an increasingly important issue in medical analysis as obesity is now recognised to play a significant role in a variety of serious health problems. With the proliferation of cross-sectional imaging modalities over the last twenty years, increased attention has focused on the relative distribution of total body fat, with the recognition that visceral fat, and in particular the ratio of intra-abdominal fat to total fat predispose to altered insulin resistance and subsequent medical problems. However, most available methods for determining body fat content cannot accurately localize the distribution of fat.

Alternative medical techniques have been used to measure body fat utilising tools such as Computed Tomography (CT) and Magnetic Resonance Imaging (MRI). As in the past the quality of MRI images was limited, CT which could provide higher image resolution, was widely used to measure the total body fat [1]. The results proved to be very encouraging but due to the requirement of exposure to ionising radiation this technique is rendered impractical for serial investigations [2]. In addition, whole body imaging using CT would necessitate extensive image interpolation, thus potentially introducing bias. MRI although not as widely available as CT is becoming more widely used and does not entail ionising radiation. While MRI has been previously used for fat analysis, a review of the literature reveals that a wide spectrum of techniques have been applied, often consisting of single or selected slices, with subsequent extrapolation to the remainder of the body [3]. Other early papers on whole body MR imaging for fat analysis used gaps of 1-3 cm between axial slices, which require interpolation and thus inevitably introduce bias. Indeed previous authors have shown that subsampling and limited scanning does introduce bias and increased uncertainty into recorded fat measurements [4]. In our study we achieved whole body coverage without slice gaps, thus removing the potential for bias from this source and ensuring accurate and reproducible results. In addition, developments in hardware,

field homogeneity as well as the use of a gradient-echo sequences have reduced our imaging time to approximately 140 sec.

In this investigation we illustrated how the use of Computer Assisted Diagnostic (CAD) techniques based on the methodologies of advanced image processing and analysis can be used in order to quantify the fat distribution within the body in sequences of full body MR images. The outcome of this research effort is a system which assesses a subject's full body MRI scan, providing numerical and visual feedback to illustrate its findings. This system generates results in a matter of minutes allowing for an initial assessment to be performed immediately following the completion of an MRI scan. By highlighting areas where body fat is concentrated the system allows radiologists to quickly identify and examine regions of interest in the scan.

The systems numerical outputs also provide an accurate measure of body fat as a percentage of whole body mass. This is an important metric, which is difficult and time consuming to arrive at by alternative means such as hydro-density and anthropometrical measurements [5,6,7]. The distribution of actual fat tissue in the body is an important measure of health and overall fitness and is not well quantified by the body mass index (BMI, the measure of body fat based on height and weight = kg/m^2), currently the most generally used metric for quantifying body fat content. But BMI has limitations as it may overestimate body fat in athletes and others who have a muscular build and underestimate body fat in older persons and others who have lost muscle mass.

In this regard the assessment of body fat in athletes involved in programmes of intensive training was an area of particular importance in our study. The ability to localize fat distribution and to show an athlete exactly where it is on their body is of great interest as such information can be used to help shape the training schedule employed by the athlete. This has particular importance for many weight restricted sportsmen, such as rowers, boxers and jockeys. In developing a fast and accurate approach to body fat measurement and localization, the authors believe that it is possible to improve the ease, efficiency, and effectiveness with which the analysis and reporting of these important fitness measures can be conducted on a routine basis.

2. Image Formation

Magnetic resonance imaging is based on the absorption and emission of energy in the radio frequency range. It effectively measures the amount of hydrogen present at each location in the scanned volume, which is in turn used to characterise the various tissues present. In this fashion a detailed picture of the region under examination is built up. The level of detail produced is very high and improving with every new generation of scanners thus enabling a wide range of diagnostic possibilities. A scanner is generally specified in terms of its magnet size. Typically a clinical scanner would have a 1.5T magnet (1 tesla = 10,000 gauss, compared with the earth's magnetic field of 0.5 gauss), although research scanners can produce much higher magnetic fields.

MRI applies a hydrogen specific RF (radio frequency) excitation pulse to the protons that sit in a static magnetic field. The pulse causes the protons within a given area to absorb energy causing the unmatched protons to spin at a given frequency (specific frequency of *resonance*) in a given direction. The RF pulses are generally supplied via a coil in conjunction with a very low strength magnet (18 to 27 millitesla) called a gradient magnet. This is arranged in such a manner inside the main magnet that when they are turned on and off very rapidly in a specific manner, they alter the main magnetic field on a very local level allowing very specific areas to be examined enabling the slicing of any region from any direction. When the RF pulse is turned off, the hydrogen protons return to their natural alignment within the magnetic field and release their excess stored energy. When they do this, they give off a signal that the coil now picks up and sends to the computer system. What the system receives is mathematical data that is converted, through the use of a Fourier transform, into an intensity map, or image.

Measurements are taken at important relaxation times T1 and T2. T1 is the settling time for the atoms to return to equilibrium after being disturbed by the RF pulse while T2 is the decay of the RF signal after it has being created, both these measures are tissue dependent. For example water has a longer T1 time when compared to fat because it does not give up its energy as quickly as fat. In a T1

weighted image the signal from fat is bright, whereas image intensities from areas of muscle and fluid are lower, this is the protocol employed in our study. Similarly water has a longer T2 time when compared with fat, as such a T2 weighted image attenuates the signal from fluid more slowly than the signal from fat, muscle, or normal connective tissues. Using these and other properties a host of different imaging protocols have being devised to optimize image quality..

2.1. Full body MR acquisition

MR images were acquired on a 1.5T imaging unit (Intera-Philips Medical Systems), which was fitted with a table-top extender and which allowed automated table movements. With the table-top extender caudo-cranial coverage of 200 cm is achieved, which allows examination of all but the tallest subjects. Images are acquired in 6-7 fully integrated stacks, with a small overlap. We acquired coronal T1 weighted gradient-echo images with voxels of $2.02 \times 2.02 \times 8.00 \text{mm}^3$. By dividing the body into such small three-dimensional voxels and subsequently using computer aided diagnostic techniques as outlined in this paper, a very definite representation of fat distribution is obtained, and potential inaccuracies due to partial volume effects are kept to a minimum. The entire body is covered, without gaps, in an imaging time of approximately 140 seconds, although actual imaging time does depend on the patient's body size. By this method 32 coronal slices of 8mm thickness are acquired for each of 6-7 stacks. In larger subjects, slice thickness may be increased to 9-10mm to ensure adequate coverage, while in smaller subjects it could be decreased to enable more accurate body mapping.

A degree of controversy exists in the literature as to the exact biochemical consistency of the tissue detected as "Fat" by MRI. However, most authors accept that adipose tissue is composed of 84.67% triglyceride, 12.67% water and 2.66% protein, giving a density of 0.9196 g/cm^3 [8]. In addition, this biochemical consistency appears to be homogenous throughout the human body so that confounding variables that underlie the difficulties with other body composition methods such as age, gender and ethnic origin are removed, and the use of equations apart from the automated total body fat calculation is removed. In our cohort of patients we examined a group with a wide spectrum of body fat levels and distributions without obvious detriment or difficulty, including athletes who are ill served by other available methods.

For the scanning process patients are placed supine, with their hands crossed over the abdomen. This enables the subjects' upper limbs to be imaged. We had initially placed the subjects' hands by their sides, but found that this caused aliasing artefacts due to the geometry of the magnetic field. After imaging, the raw data is transferred, in DICOM (Digital Imaging and Communications in Medicine) format [9,10], to a workstation for analysis. Figure 1 illustrates two images from adjacent coronal sections. Each coronal section contains 32 images running through the body from front to back. When one or more image series in a DICOM study are reconstituted and merged within a volumetric framework the resulting data volume can exhibit characteristics poorly suited to the application of automated image processing techniques. Simplistic reconstruction can result in poorly registered data, while regional greyscale variations between images can cause undesirable intensity shifts within regions in the reconstructed volume. These mismatches in localisation and image intensities can cause serious problems, especially for automated image processing and analysis procedures.




Figure 1. Samples of the individual MRI coronal sectional sub-volumes used to form the whole body image.

We now present a volume reconstruction scheme which addresses the issues of data placement and intensity matching within a three dimensional space using a combination of vector based reconstruction and automatic histogram registration. The resulting data volume represents the combination of the input sub-volumes located and normalised so as to yield a consistent data field across the combined sample space of the sub-volumes. This reconstruction technique yields a matched volumetric dataset specifically designed to produce more accurate and more consistent results under processing and analysis.

3. Volume reconstruction

In the full body MRI protocol used the subject is imaged in a set of overlapping coronal sections. The resulting series of up to 7 coronal sectional sub-volumes must be reconstructed into a volumetric dataset in order to facilitate analysis. There are two issues that require particular attention in this process: spatial registration and greyscale matching. To achieve correct spatial registration we use location and orientation vectors stored in the DICOM headers. This allows us to accurately position each image within a global coordinate system. In this way we can account for the overlap between adjacent coronal slices so as to correctly generate the final volume.




Figure 2. Initial alignment of DICOM sectional sub-volumes. This illustrates the gaps, which can appear between sections, and the grey scale variations across the DICOM sectional images. (a) Six coronal sections composed to form the full image. (b) Close-up of a regional misalignment caused by poor recombination.

3.1 Greyscale Matching

Greyscale matching is necessary because there can be significant intensity offsets between successive coronal sections due to the nature of the MRI acquisition process. It is necessary to minimise these effects in order to optimise the performance of the automated analysis procedure. We achieve this using histogram matching once the incoming images have been grouped into their coronal subsections.

An intensity histogram is calculated for each subsection, Figure 3a, the characteristic peak representing the soft tissues within the data is located algorithmically, and the set of histograms is matched, scaling them linearly so as to align the identified peak locations, while also keeping the histograms lowest intensity levels fixed, Figure 3b. Investigation into the use of a more direct histogram scaling approach based on matching the grey range of the data proved to be insufficient to the task, failing to converge adequately in grey space. The histogram matching process presented here leads to a reformatted intensity histogram for each of the subsections which results in the

reduction of intensity differences introduced during image acquisition, and consequent improvements in the results achieved in subsequent processing steps, (Figure 4).




Figure 3. Histograms of the seven subsections required for a whole body image. The first peak in each plot represents the abundant low intensity background pixels present in the image sections. The second peak represents the soft tissue and the high intensity tail above the second peak is primarily fat tissue. (a) Before normalisation, (b) after normalisation. During the histogram scaling operation some of the bins join hence the noisy appearance of the normalised histograms. Note the precise alignment of the second peak (soft tissue) of these histograms.




Figure 4. Unnormalized and normalized sections.

3.2 3D Reconstruction

At this stage the full set of input images, now intensity matched, can be manipulated uniformly in order to reconstruct the final volume. First a bounding box is calculated by recording the extents of each incoming image within the specified coordinate system and using the maximum and minimum values recorded in x, y, and z in order to calculate the origin and required extents of the final volume. As all the images come from a single DICOM study they share a common real world origin and coordinate system stored in the DICOM header. This information is now used in order to accurately project each image into the merged data space. The location of the first data pixel in each image is recorded in the DICOM header, relative to some real world origin, which remains fixed for all images within a single study. Similarly two vectors parallel to the image rows and columns respectively are also stored. This information allows a transformation matrix to be calculated which yields a direct mapping into the reconstructed volume for every pixel in each of the images. Using this information the final volume is built up by projecting each image into the output volumes data space (Figure 5).




Figure 5. Matching and merging, without (left) and with (right) histogram normalization.

Although the whole body MRI data which we are working with is nominally acquired in a coronal orientation, in reality it is often oriented slightly off the coronal. In order to accurately reconstruct the volume and arrive at a consistent data layout for processing we reorient the incoming data using the axial orientation as our base volumetric coordinate system, as this is the fundamental orientation used in medical imaging applications. In axially sliced data x runs from the subjects right to left (the observers left to right), y runs from the subjects chest to back, and z runs from their feet to head, when the subject is lying head first, supine (face up) in the scanner. Thus in a coronal image, looking at the subject from their chest towards their back, the image x axis maps to the volumetric x axis, but the image y maps to minus z in the axial volumetric coordinate system (the image y runs from head to toe).

In order to achieve the desired normalisation of coordinate systems we apply a matrix transformation to all the data so as to reslice and combine it into a unified volume in the desired axial coordinate system (Figure 6). The information needed to achieve this reorientation is found in the DICOM headers of each image slice. Consider a whole body dataset consisting of 6 sections, each containing 32 slices. This is presented to the system as a set of 192 individual DICOM image files. From each file header we can extract the real world location of the first pixel in that image $O = [O_x, O_y, O_z]^T$, in a fixed global coordinate system, along with the two orientation vectors, $R = [R_x, R_y, R_z]^T$ and $C = [C_x, C_y, C_z]^T$, aligned parallel to the image rows and columns respectively. For an exactly coronal image these vectors would be $R = [1 \ 0 \ 0]$ and $C = [0 \ 0 \ -1]$. But since the nominally coronal acquisition is often slightly off the coronal this may not be the case and as such an accurate reconstruction by straightforward means is hampered by this consideration in addition to the variable overlap which exists between each subsection of 32 images.




Figure 6. Projection of image data into destination volume (coronal to axial).

In order to achieve the desired reconstruction we proceed as follows. First we calculate the spatial extents of each image using the origin and the row and column vectors, along with the pixel width and

height of the image. This yields a three dimensional bounding box for each slice and by extension, for the complete dataset. We then project each slice in turn into the destination volume. We project the four corners of the image into the volume and then using an inverse transformation we project each volume voxel within this subregion back into the source image and interpolate from the neighbouring image pixels in order to arrive at a greyscale value corresponding to this point in the destination volume. The interpolated grey value at position x is calculated using a cubic model

$$v = a + bx + cx^2 + dx^3$$

where x is an offset from v_2 , usually in the range 0 to 1 (spanning the interval v_2 to v_3), v_1 - v_4 are four consecutive, unit spaced grey values along a line and,

$$\begin{aligned} a &= v_2 \\ b &= -v_1/3 - v_2/2 + v_3 - v_4/6 &= (v_3 - (v_1/3.0 + v_2/2.0 + v_4/6.0)) \\ c &= v_1/2 - v_2 + v_3/2 &= ((v_1 + v_3)/2.0 - v_2) \\ d &= -v_1/6 + v_2/2 - v_3/2 + v_4/6 &= ((v_2 - v_3)/2.0 + (v_4 - v_1)/6.0) \end{aligned}$$

This procedure is repeated for each of the 192 slices so that finally the destination volume has been completely populated with all the data from the original source images.

4. Image Segmentation

A visual examination of the images contained in the datasets reveals that the fat tissues tend to have a higher greyscale value than other tissues. But these images also indicate that there is quite a high greyscale variation within the image regions representing fat tissues. Also, even after histogram matching, the greyscale values for fat in some situations overlap those associated with other nominally lower intensity tissues such as those representing liver or brain. Therefore accurate segmentation cannot be achieved by applying simple methods based on thresholding. Thus in order to cope with these problems we have devised a multi-stage segmentation algorithm. The first step involves the segmentation of the full body from MRI data. For this purpose we have to analyse the distribution of the intensity value of the voxels contained in the data set. The histogram of the entire volume dataset is illustrated in Figure 7.

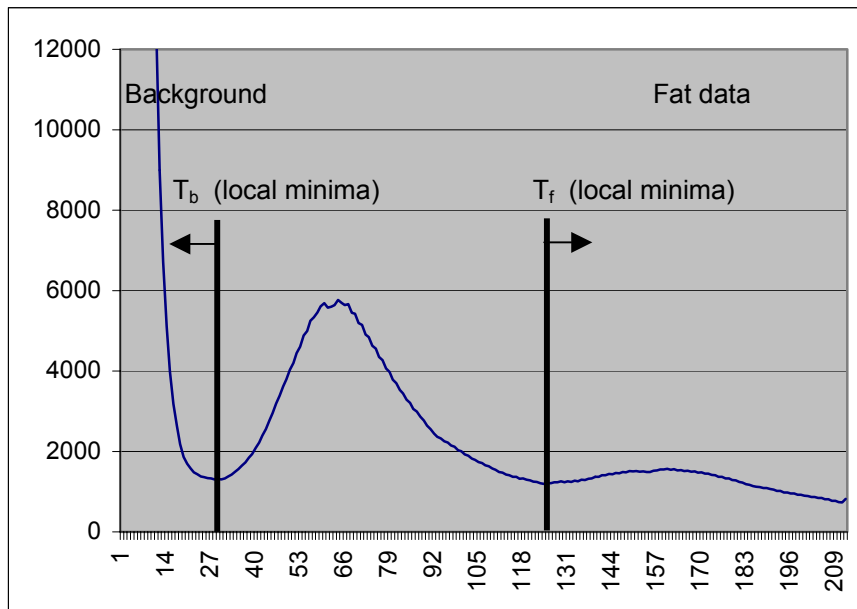


Figure 7. Histogram of the full body volume. The data with intensity values below T_b represents the background while the data with intensity values above T_f represents candidate fat tissue.

Thus, in order to segment the lean and fat body tissues from the background we have to threshold the data where the threshold is set to T_b . The result of this operation represents precise body segmentation except in the region of the lungs whose voxels have intensity values lower than T_b . This first phase segmentation is used to calculate the height and the body mass and to identify the interface surface between the body and background. The extraction of the interface surface (body shell) allows us to apply a boundary enhancement to compensate for signal drop-off in peripheral regions of data. For this purpose we apply an intensity compensation scheme based on a predefined chamfer map that is applied starting from the outer shell of the body region.

4.1 Adaptive Smoothing

Once complete we apply adaptive smoothing (Figure 8) to improve local homogeneity. The aim of this operation is to remove the additive image noise while preserving the image edge information. To this end, we implemented the smoothing algorithm described in [11]. This smoothing algorithm tries to adapt pixel intensities to the local attributes present in the image by evaluating two discontinuity measures (i.e. local and contextual) that should be preserved during the smoothing operation. The local discontinuity is measured using four detectors that approximates the image gradients in four directions:

$$\begin{aligned}
 E_{Hxy} &= |I_{x+1,y} - I_{x-1,y}|, \\
 E_{Vxy} &= |I_{x,y+1} - I_{x,y-1}|, \\
 E_{Cxy} &= |I_{x+1,y+1} - I_{x-1,y-1}|, \\
 E_{Dxy} &= |I_{x+1,y-1} - I_{x-1,y+1}|
 \end{aligned} \tag{1}$$

where $I_{x,y}$ is the pixel intensity at (x,y) . These four detectors respond strongly to local edges and a local discontinuity measure can be defined as follows:

$$E_{x,y} = \frac{E_{Hxy} + E_{Vxy} + E_{Cxy} + E_{Dxy}}{4} \tag{2}$$

As the local discontinuity evaluates the local gradients it is sensitive to image noise. Unfortunately, the MR images reveal a high level of noise and as a result the local discontinuity is not efficient in distinguishing the true local discontinuities from noise. Thus, the local discontinuity has to be augmented with a contextual discontinuity, which evaluates the attributes of the neighbouring pixels. In this implementation the contextual discontinuities are measured by the local variance that is measured in a predefined neighbourhood (see Eq. 3).

$$\sigma_{x,y}^2 = \frac{\sum_{(i,j) \in N_{xy}(R)} (I_{i,j} - \mu_R)^2}{N_{xy}(R)} \tag{3}$$

where σ_{xy}^2 is the measured variance, μ_R is the mean intensity value computed in the square neighbourhood R and N_{xy} represents the total number of pixels in a neighbourhood. The adopted smoothing strategy uses both local and contextual discontinuities and for each pixel its intensity value is iteratively updated with the nonlinear transformation illustrated in Eq. 4.

$$I_{xy}^{t+1} = I_{xy}^t + \eta_{xy} \frac{\sum_{(i,j) \in N_{xy}(1) \setminus \{x,y\}} \eta_{ij} \gamma_{ij}^t (I_{i,j}^t - I_{x,y}^t)}{\sum_{(i,j) \in N_{xy}(1) \setminus \{x,y\}} \eta_{ij} \gamma_{ij}^t} \tag{4}$$

where,

$$\eta_{ij} = \exp(-\alpha\Phi(\sigma_{xy}^2(R), \theta_\sigma)) \quad (5)$$

$$\gamma_{ij}^t = \exp(-E_{ij}^t / S) \quad (6)$$

$$\Phi(\sigma_{xy}^2(R), \theta_\sigma) = \begin{cases} 0 & \sigma_{xy}^2(R) < \theta_\sigma \\ \sigma_{xy}^2(R) & \sigma_{xy}^2(R) \geq \theta_\sigma \end{cases} \quad (7)$$

In Eqs. 4 to 7 the parameter t defines the iteration and the transformation illustrated in Eq. 4 updates the intensity values of each pixel using two weighting parameters, η and γ , which measure non-linearly the contextual and local discontinuities. The variables S and α are important as they determine to what extent the local and contextual discontinuities should be preserved during the smoothing operation. In our implementation we set these parameters to the following values $S = 10$, $\alpha = 10$, $\theta_\sigma = 0.2$, window size $R = 2$, and the algorithm is run for 2 iterations. Figure 8 illustrates the performance of the adaptive smoothing operation. Note that the smoothing operation did not affect the edge localization.




Figure 8. Adaptive smoothing operation. (a) Input image. (b) Filtered image.

An initial threshold level T_f is calculated based on an analysis of the data histogram. The peak representing soft tissue (second peak in Figure 7) is located and voxels whose values fall above the end of this peak are initialised as potential fat voxels. A three-dimensional dilation based (6-connected) region growing procedure [12] is applied to the initial estimate to extract the voxels affected by partial volumes. The adjacent voxels situated in the neighbourhood of the initial estimate with an intensity value above 90% of T_f are reclassified as fat voxels.

4.2 Region Refinement

The final stage in the process involves the application of region refinements to eliminate false positive tissues such as brain and liver [13]. This requires the identification of each disjoint region in the fat estimate. For this purpose a fast 3D labelling procedure has been applied. This process is summarized below:

- The intra-region variability for each disjoint region is evaluated.
- If the region is compact and its intensity mean is high (1.2 x threshold determined from histogram analysis) then the region is classified as fat. The threshold scaling factor was determined experimentally.
- Otherwise each pixel within the region is verified and only the voxels that are above the threshold determined from histogram analysis are retained (hysteresis threshold).
- Isolated voxels are then removed.

Through this process we arrive at a robust segmentation of the signal due to fat tissue within the data volume.

5. Total body fat calculation

Calculation of the total body fat (TBF) is performed using the following formula:

$$\text{TBF} = (N_{\text{FatVoxels}})(\text{Voxel_Dim})(\text{Fat_Density})$$

where $N_{\text{FatVoxels}}$ is the total number of fat voxels contained in the dataset, Voxel_Dim is the voxel dimension (in cm^3), and Fat_Density is the density of the fat tissue (in g/cm^3). The voxel dimensions can be extracted from the DICOM header and the datasets used in our study have mostly had dimensions of $2.02 \times 2.02 \times 8.00$ [mm^3]. The medical literature indicates that fat tissue density can be regarded as constant [1] and is usually assigned a value of 0.9196 [g/cm^3]. The fat content is determined by counting the fat tissue voxels contained in the fat segmented data (marked in yellow in Figure 10). We normalize these values in order to yield the total body fat in kilograms.

Once we know the total body fat (TBF) we can easily calculate the total body mass (TBM) with the following formula:

$$\text{TBM} = \text{TBF} + (N_{\text{FullBody}} - N_{\text{FatVoxels}})(\text{Voxel_Dim})(\text{Lean_Tissue_Density}).$$

The value for $\text{Lean_Tissue_Density}$ used in our calculations is: 1.11 [g/cm^3].

It is worth noting that the protocol used means that our voxels are quite big ($2 \times 2 \times 8 \text{mm}$) so partial volume effects replace the bias as an issue to be aware of and a potential problem and source of inaccuracies (e.g. we have found that a thin layer of fat across the front and back of the subject can be missed). This effect could be greatly reduced with the use of higher resolution data capture or sub-voxel analysis.

6. Basic system functions

Analysis results are presented in a number of ways. Simple orthogonal review allows axial, coronal, and sagittal sections to be examined, highlighting regions, which have been classified as being body fat. Colour mark-up of the images provides effective visual feedback, improving the readability of the data (Figure 9). This form of review also allows detailed examination of the distribution of fat throughout the body and facilitates the easy identification of areas of particular fat concentration.




Figure 9. System display output illustrating the localised regions identified as fat (marked in yellow in the image on the right). The slider bar allows the user to step through the image sequence. The image on the left is the original data after alignment and basic image processing to reduce noise artefacts.

Three-dimensional volume rendered views (Figure 10) provide an excellent overview of the data, and when used with data space clipping can effectively visualise the body fat distribution within a volume of interest, providing a more complete view and comprehensive breakdown of the distribution of fatty tissue within the body.




Figure 10. Volume rendering of a full body MRI dataset. As well as giving an overview of a patient's superficial body-fat distribution, this fully interactive 3D model allows for flexible data clipping to be performed, which enables detailed analysis of regional fat distribution.

In addition to these visual tools, numerical results are presented to the user. Figure 11 shows an example of the typical results generated by the system. Estimates are made of the subjects height and weight from the segmentation of the whole body (see previous section), and measurements are performed to calculate values for the actual and percentage body fat detected, measured by volume

and by weight. All these calculations are performed automatically without any initialisation or subsequent intervention being required on the part of the user. The entire analysis process from raw DICOM data to final results takes less than two minutes.

```
Series #1: 250 x 32 x 937, 2.022059mm x 8.0mm x 2.022059mm
Name:
Weight: 90.0kg
Sex: M
Date of birth: 20/10/1972

Estimated full body volume = 98,422.84cc
Estimated full body height = 1.89m

Calculated Body Fat Table:-

```

	Method 1	Method 2
Body weight DICOM	90kg	90kg
Body weight calculated	103.14kg	102.91kg
Estimated BMI	28.73	28.67
Fat by volume	31,803.91cc	33,016.37cc
Fat by weight	29.2kg	30.31kg
% by volume	32.31%	33.55%
% by DICOM weight	32.44%	33.68%
% by calculated weight	28.31%	29.45%

Figure 11. System output text on a per patient basis. Two slightly different methods for fat calculation are used. The DICOM weight is that entered by the MR technologists and is inaccurate in this patient, whose body weight was 102 Kg. The system provides an accurate assessment of body fat. By applying the clipping tool to the 3D volume rendered image (See Figure 10), similar detailed voxel-by-voxel interrogation of a particular body region can be performed.

A simplistic, threshold based approach was also applied (the results shown as Method 1 in Figure 11) in order to evaluate the advantage gained by using the more involved approach described in this paper (Figure 11, Method 2). This rudimentary approach proved to be extremely unreliable, being highly data dependant. It invariably misclassified significant portions of non-fat tissue as being fat, and occasionally resulted in fat estimates an order of magnitude larger than expected, due to the unsophisticated assessment procedure used.

7. Results

42 patients (Table 1) were recruited through local in-hospital advertising and through a hospital related sports clinic. These included a cohort of international rowers, a group of elite athletes in whom body fat estimation is of particular importance as individuals are weight restricted. Local institutional board approval was obtained. There were 21 male and 21 female volunteers used in the study. Informed consent was acquired from each subject, and their weight and height were recorded. These were used to calculate the BMI in each patient.

The BMI was calculated directly from the patients' mass and height, and indirectly using the segmentation results. The directly measured BMI values of these datasets range from 18 to 35. Comparison of the direct and indirect methods was used to validate the segmentation procedure and a good correlation was found, see Figure 12. Much of the data spread in the graph in Figure 12 is due to incomplete body acquisitions in the MRI scanner. Restricting the data to the 21 most complete datasets, where data loss at the extremities, especially at the feet of tall subjects, is kept to a minimum, results in a far superior straight line fit with much reduced deviations from the fitted line in the restricted data subsample. This is shown in the broken line and its associated subset of data points indicated in Figure 12. The sample points representing incomplete acquisitions all fall above the broken line and pull the full sample trend line (the solid line) in an upward direction. This reflects the relative effect of the missing data on the BMI and % body fat calculations. The missed data represents a relatively small volume, and hence weight (the feet and ankles), but can result in a disproportionately large variation in the estimated height of the subject. With down pointing toes the change can be as much as 30 centimetres in some cases. These two factors result in little difference in the estimated % body fat but can cause a significant elevation in the estimated BMI value since the subject is estimated to be shorter but not much lighter than would be the case were the measurements based on a

complete dataset. Note that in the case of complete datasets the automatic method consistently estimates a BMI less than the measured BMI. This is due to an over estimation in the height caused by down pointing toes. An alteration in the acquisition protocol supporting subject's feet vertically could correct for this. Correspondingly, calculated BMI's based on incomplete datasets tend to be higher than the equivalent measured values, for the reasons given. These observations support the assertion that the automatic measurements achieved represent accurate results for the data under examination.

Table 1 Patient data and associated automatically calculated BMI and percentage body fat values for the 21 male and 21 female subjects assessed in this study. These results are graphed in Figures 12 and 12.

Male Subjects	Age	Measured BMI	Calculated BMI	Calculated % Body Fat	Female Subjects	Age	Measured BMI	Calculated BMI	Calculated % Body Fat
1	31	32.6	28.67	29.45	1	27	19.3	18.11	21.06
2	33	24.1	23.54	17.27	2	29	24.8	23.53	30.54
3	32	25.4	23.73	27.79	3	20	20.3	19.72	19.61
4	51	27.3	24.43	16.85	4	53	26.8	24.07	34.36
5	29	25.0	22.76	17.36	5	56	22.5	20.77	28.66
6	29	24.8	22.59	6.97	6	27	20.8	20.23	25.21
7	29	25.0	23.68	13.06	7	21	24.6	21.73	30.29
8	28	21.6	19.58	4.77	8	43	21.9	20.59	29.02
9	49	31.1	28.42	30.07	9	41	25.6	26.69	37.46
10	52	34.7	32.58	25.98	10	19	18.0	17.68	15.89
11	24	28.7	23.18	6.24	11	26	21.8	22.21	16.39
12	28	28.7	26.91	31.79	12	18	25.3	26.83	28.89
13	23	26.3	22.62	10.32	13	26	30.5	30.96	38.05
14	19	21.5	19.63	11.63	14	33	22.8	24.11	17.99
15	28	26.3	25.58	20.04	15	23	20.5	22.10	23.71
16	25	18.9	22.30	11.57	16	25	20.3	20.89	19.97
17	26	26.9	24.80	12.21	17	25	29.4	30.74	41.67
18	24	20.9	21.12	6.24	18	27	25.7	25.86	31.78
19	26	27.7	27.58	21.19	19	55	25.1	28.52	35.20
20	30	26.0	25.20	14.96	20	53	20.8	24.27	29.66
21	45	29.9	32.70	28.3	21	42	23.7	23.65	25.16




Figure 12. Comparing manually measured BMI and the automatic CAD calculation. This illustrates that the CAD system does a good job in matching the manual BMI measurements. This test was used as an initial validation of the segmentation procedure and demonstrates that the CAD system developed produces reliable measurements. See the text for a full discussion.




Figure 13. BMI mapped against percentage body fat. The data is divided by gender and the two trend lines clearly illustrate the gender difference. Men tend to have a lower percentage body fat than women for the same BMI. The very lean subjects coming from the cohort of athletes examined further lift the trend line for the male subset, as their high muscle mass offsets low body fat and results in relatively high BMI values in this group, situated at the extreme left of the graphed data. The complexity of the relationship between percentage body fat and BMI illustrates the shortcomings of BMI as a reliable indicator of body fat.

The measure which is of greater interest is that of percentage body fat, whose accurate determination was the main goal of the study. The results showed that there was a reasonable correlation between this measurement and the BMI, with the expected gender differences as highlighted in Figure 13. However the complexity of the relationship illustrated in Figure 13 confirms the recognised shortcomings of the BMI as a measure of body fat level. The results for the cohort of international rowers highlights their unusually high muscle mass, demonstrating one of the drawbacks of BMI, its inability to distinguish between body mass stemming from different sources. The results are presented in Table 1 and Figures 12 and 13. In all cases, results were obtained within minutes of receiving the DICOM data. We have implemented an alternative implementation based on image clustering [14], although initial results indicate that this approach is outperformed by the method described in this paper.

8. Conclusion

The focus of this paper has been on the issues relating to the formation, processing and analysis of whole body MRI in the development of an MRI based body fat analyser. The reconstruction method described generates matched and merged volumetric datasets well suited to the application of automatic image processing and analysis techniques. It alleviates a number of problems, which become significant when computerised data analysis procedures are brought to bear on the problem of processing the data. While the human visual system is extremely efficient at compensating for intensity variations and physical misalignments in a scene these effects can cause significant difficulties in an automated approach to image analysis, and their minimisation greatly enhances the accuracy and reliability of such a system.

We have seen that by highlighting areas where body fat is concentrated the system allows radiologists to quickly identify and examine regions of interest in the scan. The main conclusion of this work is that whole body MRI in conjunction with CAD allows a fast, automatic, and accurate approach to body fat measurement and localization and can be a useful alternative to BMI, especially for low fat groups such as athletes. This prospective investigative study contained 42 volunteers (21 male, 21 female) aged 18 to 56. The study group consisted of healthy volunteers with a wide range of body weights, and included a cohort of high performance athletes.

Acknowledgements. This work is done in conjunction with Stephen Eustace (Director of Radiology), Darren Brennan (Radiologist) and Julie O'Brien of the Mater Misericordiae & Cappagh National Orthopaedic Hospital, Dublin, Ireland. We would also like to acknowledge Robert Sadleir, VSG in DCU for his work on the development of our DICOM decoder (NeatMed).

References:

1. RH Nord, RK Payne, Body composition by dual-energy X-ray absorptiometry - a review of the technology, *Asia Pacific Journal of Clinical Nutrition*, 4(1):173-175, 1995.
2. L Sjoström, A computer tomography-based multi-compartment body composition technique and anthropometric predictions of lean body mass, total and subcutaneous adipose tissue, *International Journal of Obesity*, 15:19-30, 1991.
3. EG Kamel, G McNeill, TS Han, et al. Measurement of abdominal fat by magnetic resonance imaging, dual-energy X-ray absorptiometry and anthropometry in non-obese men and women. *Int J Obes Relat Metab Disord* 23:686-692, 1999
4. EL Thomas, N Saeed, JV Hajnal, et al. Magnetic resonance imaging of total body fat. *J Appl Physiol*; 85:1778-1785, 1998.
5. P Deurenberg, JJ Pieters, JG Hautvast, The assessment of the body fat percentage by skinfold thickness measurements in childhood and young adolescence. *Br J Nutr*, 63:293-303, 1990
6. RB Mazess, Accuracy of TOBEC method questioned. *Am J Clin Nutr*, 39:157-159, 1984
7. JC Seidell, A Oosterlee, MA Thijssen et al. Assessment of intra-abdominal and subcutaneous

- abdominal fat: relation between anthropometry and computed tomography. *Am J Clin Nutr*, 45:7-13, 1987
8. LW Thomas, The chemical composition of adipose tissue of man and mice. *Q J Exp Physiol Cogn Med Sci*; 47:179-188, 1962.
 9. National Electrical Manufacturers Association. *Digital Imaging and Communications in Medicine (DICOM)*. Rosslyn, Va: National Electrical Manufacturers Association, 2003; PS 3.1-2003-3.16-2003.
 10. P Mildenerger, M Eichelberg, E Martin. Introduction to the DICOM standard. *Eur Radiol* 2002; 12:920-927.
 11. K Chen, "A feature preserving adaptive smoothing method for early vision", National Laboratory of Machine Perception, Peking University, China, Technical report, 1999.
 12. P.F. Whelan and D. Molloy (2000), *Machine Vision Algorithms in Java: Techniques and Implementation*, Springer (London)
 13. KY Kang, K Engelke, WA Kalender, A new accurate and precise 3-D segmentation method for skeletal structures in volumetric CT data, *IEEE Transactions on Medical Imaging*, 22(5);586-598, 2003
 14. D Ilea, O Ghita, K Robinson, R Sadleir, M Lynch, D Brennan and PF Whelan, Identification of Body Fat Tissues in MRI Data, International Conference on Optimization of Electrical and Electronic Equipment, OPTIM'04, Brasov, Romania, 20-23 May 2004