

# Intensity non-uniformity correction in multi-section whole body MRI

Kevin Robinson<sup>†</sup>, Ovidiu Ghita, Paul F. Whelan

Vision Systems Laboratory, School of Electronic Engineering,  
Dublin City University, Glasnevin, Dublin 9, Ireland

## ABSTRACT

As whole body MRI (WB-MRI) gains currency, the data this class of technique generates presents new challenges for the imaging community. One acquisition protocol currently being applied with considerable success entails imaging the subject in a number of successive coronal sections, resulting in a high resolution, gap free, full body acquisition. However this technique often results in considerable greylevel offsets between adjacent coronal sections. To make the images suitable for the application of automated image analysis procedures these discontinuities in the grey data must be alleviated. We examine the issues related to this problem, and present a solution based on histogram rescaling, which is designed to correct for the non-uniformities while preserving the integrity of the data histogram so that it can be used robustly in later processing steps. The final datasets reconstructed from the resampled coronal sections exhibit superior greyscale homogeneity, visually and in statistical measures, and the image segmentation results achieved using this corrected data are consistently more robust and more accurate than those arrived at using the original raw data. The approach has been tested and successfully validated on a database of sixty two WB-MRI datasets.

**Keywords:** Greyscale matching, Inhomogeneity correction, Whole body MRI

## 1. INTRODUCTION

Histogram based techniques are widely used in two distinct contexts. Firstly, a histogram can be altered, globally or locally, in order to change the data's greylevel distribution in some way. This is often done so as to improve the visual appearance of an image by increasing its contrast or focusing in on a particular grey range in the data. These tasks are performed using such techniques as histogram stretching, windowing, and global and local area equalisation.<sup>1,2</sup> Similar histogram based techniques are also useful for intensity non-uniformity correction in MR images,<sup>3</sup> which is the focus of this paper.


The second area of application of histogram based techniques is in the context of data segmentation and classification tasks, where they are useful for such operations as threshold selection<sup>4-6</sup> and segmentation<sup>7</sup> procedures. In this second class of applications the shape of the histogram is examined in order to guide the processing being applied. If techniques of type one have been used which alter the characteristics of the histogram introducing new minima and maxima, prior to the application of techniques from this second class, then these subsequent histogram analysis based procedures are likely to encounter difficulties. In this paper we describe a method of the first class designed to accommodate the later application of this second type of procedure, by preserving the overall shape of the data histogram during the remapping process.

The histogram matching procedure which we have developed operates in the following fashion. Initially each coronal section of thirty two images is assessed independently, analysing the data histogram in order to identify the two characteristic peaks representing air voxels and lean tissue. The high intensity voxels in the histogram tail represent fatty tissue. The second peak can be wide and low, sometimes barely rising above the descending histogram tail, and for robustness is detected using an adaptive scanning technique. The accurate identification of this peak is crucial as this is the primary marker against which we classify each section's optimal intensity correction scheme. The aim is to transform each section's greylevel distribution such that after resampling, the lean tissue peaks are aligned across all the sections.

---

<sup>†</sup>Corresponding author. *E-mail address:* kevin.robinson@eeng.dcu.ie

The question of identifying the appropriate histogram peaks, and troughs in a robust and consistent fashion is addressed in the analysis which follows. This is in turn followed by an examination of how we scale the individual grey maps so as to preserve the integrity of the rescaled histograms after matching, avoiding the introduction of artificial spikes and voids, which can be caused by histogram bins merging and separating respectively (see Fig. 1). This effect is observed with the more straightforward histogram scaling schemes typically employed in applications such as greyscale windowing for visualisation purposes. Such spike and void features do not adversely affect the visual characteristics of the data as they merely cause certain sets of pixels to change their greyvalues by a single greylevel relative to their intensity neighbours (i.e. pixels close to them in intensity rather than space). As such the visual consequences of such histogram spikes and voids are minimal. However in our work their presence would constitute a significant difficulty as the data histogram is much used, and by addressing this issue here we insure that subsequent histogram based calculations, especially in the automatic selection of threshold bands through histogram analysis, can be performed simply and robustly in the succeeding phases of the processing pipeline.




**Fig. 1:** A histogram compromised with spurious spikes and voids caused by remapping of the data into a new grey range. The original histogram is shown in (a), while (b) illustrates the effect of compressing the gray range thus introducing spikes, and (c) shows the effect of expanding the grey range, introducing voids.

In Fig. 1 the first histogram is that of an unaltered image while the second is from the same slice after all its voxels have been multiplied by a value of 0.96, and in the third the voxels were multiplied by a value of 1.04. Since the voxel values are stored as integers rounding occurs and as a result spikes and voids are created in the histogram, where two adjacent values in the original data are either mapped to the same value in the result, or are mapped to two more widely separated values. Where this issue arises it is most commonly addressed by smoothing the actual data histogram through a process of bin averaging, before any histogram based calculations such as peak detection are performed. While this histogram smoothing approach seeks to overcome the problem as best it can it proves difficult to eliminate completely the spurious features which have been introduced into the histograms in the scaling process, and thus it can result in the misinterpretation of such spurious local minima or maxima found in the compromised data histogram. We aim to avoid these difficulties arising in the first place by performing a more elaborate redistribution of the voxels into the available bins in the rescaled histogram, thus preventing the introduction of artificial spikes and voids.

## 2. DATA CHARACTERISATION

Initially we examined a large number of datasets in order to characterise the nature of the greylevel intensity distributions typically found in WB-MRI data. We observed four characteristic peaks as illustrated in Fig. 2, broadly representing unsampled regions, background air signal, lean tissues, and fat deposits respectively (though the fat peak is often absent as can be seen in a number of the histograms in Fig. 4). The first peak, representing unsampled regions, exists solely at intensity level zero and corresponds to the blank regions observable at the four corners of each sectional image (see Fig. 3). These regions are not sampled during the mri scan and simple square off the rounded regions representing the actually sampled space in the data. We can safely ignore this peak in our processing, and leave the corresponding zero intensity voxels unaltered.



**Fig. 2:** Typical data histogram for a single coronal section taken from a seven section, coronally acquired WB-MRI dataset, illustrating the characteristic features found in the greylevel distributions observed in such data. In this case the unsampled peak actually rises to over four times the height illustrated here.

The second peak represents air spaces within the scanned volume. It is characterised by an extremely compact distribution usually centred around intensity levels six or seven. Its height is extremely variable but is always much greater than that of the lean tissue peak. This variability reflects the variation seen in the amount of airspace within each scanned section. The first and last sections for instance, around the head and feet, contain far more airspace than do the sections around the torso where the body itself occupies most of the scanned volume. The third peak corresponds to lean tissues within the scanned volume. It's location is far more variable than that of the air peak, both from scan to scan and from section to section within a single scan. It is this inter section variability within any given scan which we are concerned with in this paper. The fourth peak is not always present and where it is, it is consistently the smallest of the four. It and the long tail beyond it represent voxels in the higher intensity range within the data, broadly corresponding to fatty tissues.



**Fig. 3:** Three sectional images from a whole body MRI study. The image intensities have been inverted for clarity, such that darker regions represent higher intensity signal. The four zones from Fig. 2 can be identified: unsampled regions at the corners, background noise outside the body, and lean and fat tissues within the body. A complete WB-MRI dataset typically consists of seven sections, each comprised of thirty two such sectional images running from the subjects front to their back.

We note also that the shift observed in the locations of the maxima from one sectional histogram to another does not manifest uniformly between the various characteristic peaks. In general the low intensity peak representing the air signal remains very stable in its location and spread while the lean tissue peak changes significantly in both these respects. Referring to the set of seven unprocessed histograms shown on the left in Fig. 4 we can see that the lean tissue peak changes its location significantly between sections, with the three sections around the chest, torso, and hips demonstrating peaks centred at much lower intensity levels than in the others in this case. This can be recognised in the corresponding image sections as brighter (in the inverted

images) lean tissue regions. Note that in all these histograms only the end of the background peak is visible descending from above, with the unsampled peak not appearing at all. In order to reach the tops of all the background peaks we would need to graph to about seven times the illustrated height, and higher again in order to reach the level of the unsampled peaks at intensity level zero. Note also how small the lean tissue peak is in the seventh section. This illustrates the fact that there is lots of background and very little tissue in this region, the ankles and feet occupying only a small proportion of the total volume in the scanned region.


Finally as the histogram tail extends farther into the higher intensity range it becomes increasingly sparse. The histograms illustrated in this paper have all had their tails uniformly truncated for easier review and comparison. In fact most span more than twice the depicted range and include many empty bins in this upper region. Before processing commences we compress the histograms, squeezing out all the empty bins. This does not adversely affect the data's characteristics and makes the processing easier to manage.

### 3. HISTOGRAM MATCHING


The histogram matching process can be outlined as follows. It commences with the locating of the histograms airspace peak. This will be the largest peak, ignoring the histogram value at intensity level zero. Next the intervening trough and second peak are found in tandem, searching out from the airspace maximum. Once the characteristic peaks representing air and lean tissue have been identified these are used in the matching process which follows. This fixed point location procedure is repeated for each coronal section in the dataset. The scaling necessary in each case is then determined based on the relative locations of these fixed points, and thus each section has its intensities remapped in order to bring the fixed points into alignment across all the sections in such a fashion as to preserve the original histogram shape, avoiding the introduction of spurious spikes and voids.

#### 3.1. Fixed Point Detection


The first step, that of locating the airspace maximum, is a trivial matter. Once the data histogram has been constructed for a section the largest bin (ignoring bin zero) is found and its location in the histogram is recorded (A). This maximum will be located in the air peak. Because the intensity distribution in air voxels is much tighter than it is in tissue voxels the air peak is generally tall and narrow as compared to that of the soft tissue signal. As a result a relatively small region of airspace in the image will result in a peak which is significantly taller than the lean tissue peak, even though the area under the peak (i.e the number of voxels in that intensity range) may be less than the corresponding area under the lean tissue peak.



Next we need to identify the best candidate locations for the trough (B) and second peak (C). Since we are operating in the presence of noise we must be careful that we are locating the correct features in the histogram. After examining the potential difficulties and experimenting with a number of detection strategies we arrived at a scheme which performs in a robust fashion even when presented with a weakly defined trough and second peak in the presence of significant noise, as is the case with the final image section in Fig. 4 for instance.




Most approaches to the question of histogram mode detection use arbitrarily chosen threshold levels or tuning parameters to overcome the problems which the presence of noise introduce. It might be decided that, on moving away from the global maximum, the level must drop to at least 50% of the maximum value before the next peak is sought, or that any peak must be wider than 100 bins, to some definition of width, or must demonstrate a consistent slope direction when averaging over three or five bins for some given distance, or some other such arbitrarily or experimentally selected measure. Our method utilises a robust and meaningfully defined metric which does not require any tuning, initialisation, or parameterisation. By performing the directed search described below we identify the pair of levels which demonstrate the largest trough to peak range in the data, in which the identified local minimum position lies between the histograms airspace maximum and the identified



**Fig. 4:** Slices from the seven sections of a whole body mri study along with histograms for each section, before and after matching. After matching the seven lean tissue peaks are well aligned. The greyscale intensities in the images have been inverted for clarity, such that darker regions represent higher intensity signal.

local maximum. This represents the most reliable and also the most predictable method of identifying the two most significant modes in a multi-modal histogram in the presence of noise.

Rather than attempting to identify the two features one at a time we use a scheme which searches for the maximum differentiating measure in the data. Using the fact that the trough lies between the two peaks we move out from the airspace maximum and every time we identify a local minimum we note its location and level and then continue out looking for all the subsequent local maxima. For every one of these ordered minimum/maximum pairs which we locate we calculate the peak to peak distance ( $\delta$ ). The min/max pair which yield the largest  $\delta$  gives us the pair of bins we seek.



### 3.1.1. Modelling the histograms as a continuous function

In order to perform the desired rescaling we wish to model each histogram as a continuous function so that we can resample that function at a new set of intervals in order to arrive at the required relative levels for the new set of histogram bins. So as to present a consistent treatment throughout and so as to maintain the maximum levels

#### *Locating Local Maxima to Sub-integer Precision*

Consider the three successive sample values  $v_1$ ,  $v_2$ , and  $v_3$ , where  $v_2$  represents a local signal maximum ( $v_1 < v_2 > v_3$ ). If we fit a quadratic ( $y = ax^2 + bx + c$ ) through these three points then the location of the quadratics maximum ( $\frac{dy}{dx} = 0$ ) represents the best estimate of the true local maximum's location.

In order to calculate the position as an offset from  $v_2$  we consider the three points as having coordinates  $(-1, v_1)$ ,  $(0, v_2)$ , and  $(1, v_3)$ . This yields the following three simultaneous equations:

$$v_1 = a - b + c \tag{1}$$

$$v_2 = c \tag{2}$$

$$v_3 = a + b + c \tag{3}$$

We can thus solve for  $a$ ,  $b$ , and  $c$  in the sought quadratic:

$$\begin{aligned} c &= v_2 && 2 \\ b &= \frac{v_3 - v_1}{2} && 3 - 1 \\ a &= \frac{v_3 + v_1 - 2v_2}{2} && 3 + 1 \end{aligned}$$

Now differentiating and setting equal to zero we get:

$$\begin{aligned} \frac{d}{dx} ax^2 + bx + c &= 2ax + b = 0 \\ \Rightarrow x &= -\frac{b}{2a} = \frac{v_3 - v_1}{4v_2 - 2(v_3 + v_1)} \\ y &= -\frac{b^2}{4a} + c = \frac{(v_3 - v_1)^2}{16v_2 - 8(v_3 + v_1)} + v_2 \end{aligned}$$

Hence we know  $x$ , the offset, and  $y$ , the level. Note that if  $v_1$  equals  $v_3$  then the offset is zero, the symmetrical case. Further, since  $v_1$  and  $v_3$  are always smaller than  $v_2$  we can see that the denominator must be positive and hence the offset is negative if  $v_1$  is larger than  $v_3$  and positive if  $v_1$  is smaller than  $v_3$ . If  $v_1 = v_2$  then  $x = -\frac{1}{2}$  and if  $v_2 = v_3$  then  $x = \frac{1}{2}$ . The situation where  $v_1 = v_2 = v_3$  is a degenerate case (the quadratic becomes a straight line). These observations conform to the expected range and motion of the calculated maximum.

of precision in the process, we perform all calculations in continuous space, treating the histogram bin values as a series of sampled data points along this notional continuous function. Thus in the final implementation we use a quadratic model to locate the sub-integer coordinates of the two maxima A and C in each case, (see the text box ‘*Locating Local Maxima to Sub-integer Precision*’ for the mathematical details). We then use linear interpolation (see the text box ‘*Linear Interpolation*’) and cubic interpolation (see the text box ‘*Cubic Interpolation*’) in order to calculate the set of new bin levels, which will be positioned at locations intermediate to the existing bins.

### *Linear Interpolation*

Linear interpolation provides a quick and easy method to calculate intermediate values between pairs of sample points in a data series. Given two points  $(x_1, y_1)$  and  $(x_2, y_2)$  we can fit a line  $y = mx + c$ , and so define  $m$  and  $c$ :

$$m = \frac{y_1 - y_2}{x_1 - x_2}$$

$$c = \frac{x_1 y_2 - x_2 y_1}{x_1 - x_2}$$

Thus for any given  $x$  we can calculate a value for  $y$ :

$$\begin{aligned} y &= mx + c \\ &= \frac{y_1 - y_2}{x_1 - x_2}x + \frac{x_1 y_2 - x_2 y_1}{x_1 - x_2} \end{aligned}$$

And for any given  $y$  we can calculate a value for  $x$ :

$$\begin{aligned} x &= \frac{y - c}{m} \\ &= \frac{y - \frac{x_1 y_2 - x_2 y_1}{x_1 - x_2}}{\frac{y_1 - y_2}{x_1 - x_2}} \\ &= \frac{y(x_1 - x_2) - (x_1 y_2 - x_2 y_1)}{y_1 - y_2} \end{aligned}$$

In order to formalise our presentation we will at this point declare a set of useful functions which are needed to perform the calculations necessary to implement the procedure outlined above. The underlying mathematics required in order to implement all of these functions is provided in the three previously mentioned highlighted text boxes. The first pair of functions use linear interpolation to calculate either the  $x$  or  $y$  coordinate of a point on a straight line ( $y = mx + c$ ), given the other coordinate plus two additional points on the line. Linear interpolation is used in the piecewise linear mapping of each bin in the rescaled histogram back to a corresponding position in the original as illustrated in Fig. 5.

$$x = \text{LinearX}(y, x_1, y_1, x_2, y_2)$$

$$y = \text{LinearY}(x, x_1, y_1, x_2, y_2)$$

The next pair of functions is used to calculate the  $x$  and  $y$  coordinates respectively of the maximum/minimum point on a quadratic ( $y = ax^2 + bx + c$ ) which passes through the three points  $(-1, y_1)$ ,  $(0, y_2)$ , and  $(1, y_3)$ . We use these functions to calculate the best estimate for the true location and height of the maxima corresponding to the positions A and C identified previously. As discussed already we only actually need the locations of the maxima in this case, which are used as the fixed points in the piecewise linear matching to follow, but derivations

for both are provided in the text box.

$$x = QuadraticMaxMinX(y_1, y_2, y_3)$$

$$y = QuadraticMaxMinY(y_1, y_2, y_3)$$

The final function uses a cubic model ( $y = ax^3 + bx^2 + cx + d$ ) to calculate the  $y$  value at any given  $x$ , for a point which lies on a cubic curve passing through the four points  $(-1, y_1)$ ,  $(0, y_2)$ ,  $(1, y_3)$ , and  $(2, y_4)$ . This is used in interpolating the inter-bin levels from the original histogram, used in the construction of the final rescaled version.

$$y = CubicY(x, y_1, y_2, y_3, y_4)$$

### *Cubic Interpolation*

Cubic interpolation is used in a number of places throughout this work in order to resample data at intermediate locations between existing sample points. In general it provides superior performance to linear interpolation, but requires four sample points, to the linear schemes two. The following treatment shows how to estimate  $y$  for any given  $x$  on the basis of four uniformly spaced samples, typically two either side of the desired location, based on a cubic model,  $y = ax^3 + bx^2 + cx + d$ .

Consider four sample points  $(-1, y_1)$ ,  $(0, y_2)$ ,  $(1, y_3)$ , and  $(2, y_4)$ . Inserting these into the cubic leads to four simultaneous equations in the four unknowns  $a$ ,  $b$ ,  $c$ , and  $d$ :

$$y_1 = -a + b - c + d \tag{4}$$

$$y_2 = d \tag{5}$$

$$y_3 = a + b + c + d \tag{6}$$

$$y_4 = 8a + 4b + 2c + d \tag{7}$$

We can combine and rearrange these equations in order to solve for the four unknowns in the standard fashion:

$$\begin{aligned} d &= y_2 && 5 \\ b &= \frac{y_1 - 2y_2 + y_3}{2} && 4 + 6 \\ a &= \frac{-y_1 + 3y_2 - 3y_3 + y_4}{6} && 4 + 4 + 7 \\ c &= \frac{-2y_1 - 3y_2 + 6y_3 - y_4}{6} && 4 \end{aligned}$$

We substitute these values into the cubic to calculate a value for  $y$  at any given intermediate position of  $x$ . Typically  $x$  will be in the range zero to one, spanning the gap between the inner two samples. Moving away from the sampled points reduces the likelihood of a close fit to the real data, although boundary cases can be covered by allowing  $x$  to range from minus one to two in order to facilitate interpolation right up to the edges of the sampled dataspace. We perform interpolation in two or more dimensions by composing a series of one dimensional interpolations in orthogonal directions.



### 3.2. Histogram Resampling

Having defined the requisite functions we can now proceed as follows. Once the values for A and C have been determined for each section in the volume the set of values with the largest C is selected and used to define the target histogram configuration for all the sections. We impose the two additional constraints that the histogram origins are fixed, as are their initial sizes. The rescaled histograms can now be calculated using the identified  $A_{Max}$ ,  $C_{Max}$  pair, along with each sectional histograms own set of calculated fixed points in order to define the required histogram scaling function, mapping the original histograms into their final rescaled versions (Fig. 5). Since initially the number of bins in the histogram is to be held constant the next task is to determine how many voxels should be allocated to each bin in the rescaled version.

Having first calculated values for A and C, and applying the two additional constraints, that the origin stays fixed (bin zero maps to bin zero) and the number of bins in the histogram also initially remains fixed, we can now use the piecewise linear model described (Fig. 5) in order to map each bin in the new histogram back to a real valued position in the original and then, using cubic interpolation, we can calculate the size of each new bin in the final histogram (Fig. 6). These relative sizes must next be scaled uniformly so that the sum of all the histogram bins equals the number of voxels in the section. Some additional manipulation is then required so as to arrive at integer values for the size of each bin, (since only a whole number of voxels can be assigned to a bin). Step one in this process is to perform the piecewise linear histogram re-binning procedure in order to arrive at the real valued resampling of the notional continuous function as illustrated in Fig. 6. These values are normalised and rounded to the nearest whole number, since we can only assign an integer number of voxels to each bin, and the total number of voxels assigned must also be equal to the number of voxels in the section. Consider an image with N voxels, and a series of real valued histogram bin levels stemming from the above process,  $H_1$  to  $H_{Size}$  which sum to the real value  $R$ . The final bin size for each bin is then calculated as  $H_i = round(\frac{H_i \times N}{R})$  and the remainder due to rounding is accommodated at the tail end of the histogram. If the remainder is less than zero then we have allocated too many spaces in the histogram and we remove bins from the tail until we arrive at the required sum. If the remainder is greater than zero then we add to the tail new bins equal in size to the original last bin until we make up the shortfall. Due to the relatively small number of voxels involved and the natural variability of the tail length in the original data this procedure does not unduely alter the characteristics of the data after matching.

It then becomes a question of assigning each voxel to it's new bin, and therefore by definition assigning it it's new intensity value. In order to achieve this we need to define a strict ordering on all the voxels in the section. We already have a partial ordering given by their initial intensities, so that what we now need is to be able to define an ordering within each bin of the original histogram, thus providing a total ordering for all the voxels, allowing us to assign the required number of voxels to each bin in the new histogram.

We use an ordering function based on the offset weighted intensity mean measured in a local neighbourhood around each voxel in order to specify the strict ordering required:

$$\mu_{wt} = \frac{\Sigma(val/\delta)}{\Sigma(1/\delta)}$$

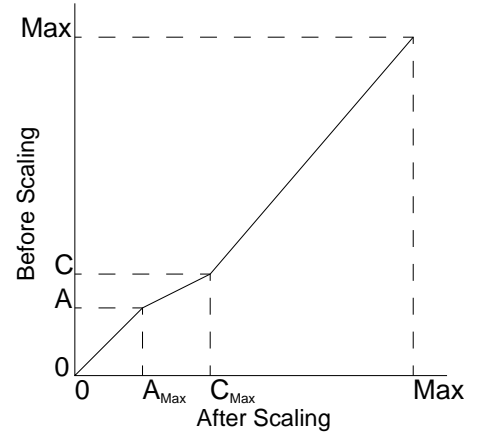


Fig. 5: Piecewise linear scaling

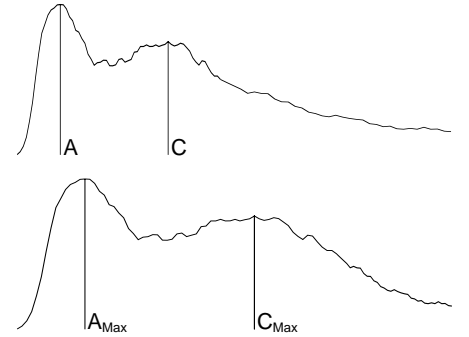


Fig. 6: Before and after scaling

Where  $val$  is the neighbours value and  $\delta$  is the offset from the current voxel to the neighbour. This scheme preserves the underlying greylevel distribution and the relative intensities for individual voxels within each section. No voxels intensity will leapfrog any other's. We are only providing a means of splitting the voxels in a given bin so that we can distribute them in any required proportion between two or more neighbouring bins. The ordering function works on the principle that the brighter a voxels neighbours are, the more likely it is that that voxel should take on a brighter value itself. Thus the intra-bin ordering is performed such that voxels within each bin are given a weighting which specifies a relative tendency towards brighter or darker, based on the local neighbourhood. Essentially each voxel is being assigned a fractional part to their intensity values, allowing for the required strict ordering to be specified.

#### 4. NON-UNIFORMITY CORRECTION RESULTS

Once this procedure has been applied and the set of volume sections have been homogenised such that the characteristic points on their histograms are matched and their greylevel distributions have been brought into alignment, corresponding tissue regions in each section will have been updated to occupy the same grey range from section to section. Fig. 4 shows a set of sectional histograms before and after matching. Regional intensity variations visible from section to section in the original images have been virtually eliminated in the processed version resulting in a far more homogenous set of data regions within the volume as a whole. The two sets of histograms shown illustrate the section by section greyscale distributions before and after intensity non-uniformity correction, and demonstrate the improved levels of histogram feature alignment achieved after the homogenisation procedure has been applied. We can see that the soft tissue peaks in the seven sections have been successfully aligned, bringing the intensity ranges for each imaged tissue region into close agreement.


The voxel ordering function used in this procedure examines a five by five by three neighbourhood so as to determine a voxels weighting, and does not guarantee an absolute ordering, as it is possible for two voxels of the same base intensity to also manifest the same weighting value. As illustrated in Table 1 this situation arises most infrequently in the real world data examined (on average about once in sixty five thousand voxels), and as such is not considered to be a problem.

Volumes processed	24
Total sections	360
Voxels per section	65025
<b>Indistinguishable voxel pairs per section</b>	
Minimum	0
Maximum	7
Sample mean	1.094444
Sample standard deviation	1.348650

**Table 1:** Voxel ordering discriminant function effectiveness.

In all but seriously pathological cases (e.g. a completely uniform sub-volume where no discrimination is possible) this issue could be addressed by an adaptive weighting scheme which grows the neighbourhood used, to whatever level is needed in order to achieve differentiation in each individual case. We consider this to be an unnecessary complication of the situation based on the analysis outlined in Table 1. Instead we simply allow such rare voxel pairs to be ordered based on the order in which they appear in the volume.

In developing this greyscale compensation procedure we have enabled a more robust and more consistent analysis of WB-MRI data to be carried out, while avoiding the introduction of additional difficulties at later stages due to the creation of compromised data histograms. In so doing we maximise the levels of confidence with which we can extract information from the data, and provide the best possible starting point for the processing to follow. The images in Fig. 7 illustrate a successfully processed dataset, indicating the image regions which have been classified as being fat tissue.



**Fig. 7:** Segmentation results on a typical WB-MRI dataset after histogram matching, demonstrating two and three dimensional review of the classified body fat region.

## REFERENCES

1. R. Gonzalez and R. Woods, *Digital Image Processing*, Addison-Wesley, 1992.
2. P. Whelan and D. Molloy, *Machine Vision Algorithms in Java: Techniques and Implementation*, Springer-Verlag, 2000.
3. K. Robinson, O. Ghita, and P. Whelan, "Volumetric reconstruction: Matching and merging in dicom data," in *Computer Assisted Radiology and Surgery*, p. 1236, (Chicago, USA), 2004.
4. N. Otsu, "A threshold selection method for gray-level histograms," *IEEE Transactions on Systems, Man, and Cybernetics* **9**(1), pp. 62–66, 1979.
5. P. Dulyakarn, Y. Rangsanseri, and P. Thitimajshima, "Histogram transformation based threshold selection for image segmentation," in *Proc. Asian Conference on Remote Sensing*, (Hong Kong, China), 1999.
6. K. Robinson, P. Whelan, O. Ghita, and D. Brennan, "Measurement and localisation of body fat in whole body mri," in *IEI Biomedical Engineering Research Award*, (Dublin, Ireland), 2005b.
7. T. Newman, N. Tang, C. Dong, and P. Choyke, "Slice-adaptive histogram for improvement of anatomical structure extraction in volume data," *Pattern Recognition Letters* **23**(1-3), pp. 25–38, 2002.