

Segmentation of male abdominal fat using MRI

Peter Stanley Jørgensen

Kongens Lyngby 2006
IMM-Master Thesis-2006-

Technical University of Denmark
Informatics and Mathematical Modelling
Building 321, DK-2800 Kongens Lyngby, Denmark
Phone +45 45253351, Fax +45 45882673
reception@imm.dtu.dk
www.imm.dtu.dk

IMM-Master Thesis: ISSN

Abstract

This thesis describes the methods used to construct a pipeline for the automatic and robust segmentation of adipose tissue in the abdominal region of human men. The segmentation is done into 3 classes: subcutaneous adipose tissue, visceral adipose tissue and other tissue.

The MRI data is preprocessed to remove the field of non-uniformity in intensity levels that are present on MR images. A novel way of sampling the field is introduced and the field is estimated using *Thin Plate Splines*.

The initial clustering of the data is done on the preprocessed data using *Fuzzy c-mean clustering*. The results of the clustering are accurate partly due to a successful preprocessing.

The segmentation of adipose tissue into the subcutaneous adipose tissue and visceral adipose tissue classes is done using a combination of *Active Shape Models* and *Dynamic Programming*. This hybrid approach of combining the two methods makes for a both robust and accurate segmentation.

No ground truth is available to verify the accuracy of the results against. The results have however been found accurate by visual inspection of the results on a large number of patients.

Resumé

Denne rapport beskriver de metoder der er blevet brugt for at udvikle en pipeline, der automatisk og robust kan segmentere fedtholdigt væv i maveregionen på mænd. Segmenteringen inddeler vævet i 3 klasser: subkutant fedtvæv, visceralt fedtvæv og andet væv.

MRI dataen bliver preprocessed for at fjerne det felt af ikke uniformitet i intensiteter der er tilstede på MR billeder. En ny metode til at sample dette felt er introduceret og feltet bliver estimeret ved hjælp af *thin plate splines*.

Den første klassificering af dataen bliver udført på den preprocessed data ved at bruge *Fuzzy c-mean clustering*. Resultaterne af klassificeringen er gode, til dels på grund af en succesrig preprocessing.

Segmenteringen af fedtholdigt væv ind i de 3 klasser, bliver gjort ved at anvende en kombination af *Active Shape Models* og *Dynamisk Programmering*. Denne hybride fremgangsmåde giver en både robust og nøjagtig segmentering.

Der er ingen reference resultater at sammenligne resultaterne fra den automatiske metode med. Resultaterne er dog fundet nøjagtige ved visuel inspektion, af de endelige resultater, på en lang række patienter.

Preface

This thesis was prepared at the institute for Informatics and Mathematical Modelling, the Technical University of Denmark in partial fulfillment of the requirements for acquiring the Master of Science degree in engineering (M.Sc.Eng). The work was carried out over a period of 5 months corresponding to 30 ECTS units.

Lyngby, January 2006

Peter Stanley Jørgensen

Contents

Abstract	i
Resumé	iii
Preface	v
1 Introduction	1
1.1 Background	1
1.2 Motivation	2
1.3 Thesis overview	2
2 The MRI data	5
2.1 The two Series	5
2.2 The segmentation tasks	6
2.3 Image variation and quality	7

3	Preprocessing	11
3.1	Bias field correction	11
3.2	Finding adipose tissue Voxels	14
3.3	Estimating the bias field	24
4	Distinguishing adipose tissue from other tissue	35
4.1	Fuzzy c-mean clustering	36
4.2	Results	46
5	Finding image structures	49
5.1	Active Shape Models - (ASM)	50
5.2	Aligning the training set	52
5.3	Modelling shape variation	55
5.4	The two shapes	57
5.5	Finding shapes in unknown images	60
5.6	Obtaining a good initialization	62
5.7	Applying ASM to the MRI data	64
5.8	Results	67
6	Outlining image structures	69
6.1	Dynamic Programming - DP	69
6.2	Transforming the MR image data	71
6.3	The external SAT border	74
6.4	The internal SAT border	76

CONTENTS	ix
6.5 The VAT border	77
6.6 Constraints on other slices	81
6.7 Results	83
7 Set operations and connectivity	85
7.1 Set operations	85
7.2 Connectivity	89
7.3 Results	94
8 Final results	97
9 Conclusion	99
A Program overview	101
B The training set	103
C Volume results	107
D Segmentation results	111

Introduction

This project has been done in collaboration with Kristian Wraae. Kristian Wraae is a PhD student at "Sygehus Fyn". A short description of the background of the medical research project he is working on will first be given to form a basis for understanding the motivation for this project. The projects motivation in the context of the research project will then be described and finally an overview of the thesis as a whole will be given.

1.1 Background

There is growing evidence that obesity is related to several metabolic disturbances such as insulin resistance, impaired insulin secretion, non-insulin-dependent diabetes mellitus (NIDDM), hypertension, dyslipidemia, cardiovascular disease and in the case of the medical study, levels of free testosterone.

It is also becoming clearer that these disturbances are more closely correlated with central (abdominal), rather than peripheral (gluteo-femoral) fat pattern. It is therefore of clinical importance to be able to accurately measure abdominal fat tissues and to be able to distinguish visceral adipose tissue from subcutaneous adipose tissue.

Different techniques are currently available including anthropometry (waist-hip ratio, BMI), ultrasound, dual-energy-x-ray absorptiometry (DEXA), computer tomography (CT) and magnetic resonance imaging (MRI).

These methods differ in terms of cost, reproducibility, safety and accuracy. The anthropometric measures are easy and inexpensive to obtain but do not allow direct quantification of visceral fat. Other techniques like CT will allow for this distinction in an accurate and reproducible way but are not safe to use due to the ionizing radiation used. MRI on the other hand does not have this problem and will also allow a visualization of the adipose tissue.

The potential problems with MRI measures are linked to the technique by which pictures are obtained. MRI does not have the advantage of CT in terms of direct classification of tissues based on Hounsfield units and will therefore usually require an experienced professional to visually mark and measure the different tissues on each image making it a time consuming and expensive technique.

1.2 Motivation

The goal of this project is to develop a pipeline for automatic assessment of Subcutaneous Adipose Tissue (SAT) and Visceral Adipose Tissue (VAT) using MRI data. An accurate and robust method could potentially be an inexpensive and fast way of assessing abdominal fat. The short term goal is to assist PhD student Kristian Wraae at "Sygehus Fyn" in assessing the distribution of adipose tissue in 300 male patients.

1.3 Thesis overview

The thesis is structured into a number of chapters and appendices. The structure is as follows.

- **Chapter 2** gives a description of the MRI data used in the project. It covers what parts of the data are used and gives a sample of the variation of image structures that can be expected from patient to patient.
- **Chapter 3** describes the method used to first sample and then remove the bias field from the MRI data.
- **Chapter 4** covers the initial classification of the preprocessed data.

- **Chapter 5** describes how the desired image structures of each patient are found.
- **Chapter 6** describes how an accurate outline of the image structures is obtained.
- **Chapter 7** covers the steps that are performed to get a final segmentation from the image structure outlines.
- **Chapter 8** gives a short evaluation of the final segmentation result.
- **Chapter 9** is the conclusion.
- **Appendix A** contains an overview of the structure of the programs that have been developed.
- **Appendix B** contains images of the data in the training set that are used in Chapter 5
- **Appendix C** contains tables of the distribution of adipose tissue on a large number of patients.
- **Appendix D** shows the end result of the segmentation on a large number of patients.

At the end of each chapter the results of the described step will be evaluated and the result of the step will be shown on a small number of patients.

The MRI data

The data set consists of MR image data from 300 patients. Only the T1 modality is provided. Adipose tissue shows as high intensity on T1 weighted MR scans and this modality is thus well suited for identifying regions of adipose tissue. The resolution of each scan is 256 times 256. The images were delivered in the DICOM format, which along with the MRI data contains a header with information about the settings used in the image acquisition process.

2.1 The two Series

Two series of MR images have been recorded on each patient. The first series consists of 15 MR scans beginning at the second lumbar vertebra. From the second lumbar vertebra the scans are acquired at a thickness and spacing of 1 cm. That is, true volume data is available.

The second series consist of 5 scans of which only the middle scan is of interest in the context of this project. The middle slice of this series is placed at the lower limit of the fifth lumbar vertebra. For the remainder of the report each scan will be denoted as a *slice* and a numbering will be given to each slice starting at 1 for the slice at the second lumbar vertebra in the example of series 1.

The individual pixels of each MR image will be denoted voxels since each pixel of the image represent a 3-dimensional area in the patient. Note that in the context of explaining methods that work on two dimensional data, not specific to MR data, the term pixel will still be used to denote the smallest component of an image.

The medical research project seeks to determine the volume and distribution of adipose tissue in an anatomically bounded unit. The unit is bounded by the second lumbar vertebra at the top and the fifth lumbar vertebra at the bottom. Since this region has a different extend from human to human, a different number of slices from series 1 will be needed to cover the anatomically bounded unit for each patient. The number of slices to be used from each patient is decided by adding slices from series 1 until the level of the middle slice from series 2 has been reached. The relative location of all slices are available in the DICOM header. All slices outside the anatomically bounded unit are discarded. An Illustration of the 2 series and the anatomically bounded unit can be seen on Figure 2.1.




Figure 2.1: The principle of the anatomically bounded unit. Red lines represent slices from series 2 and blue lines represent slices from series 1.

2.2 The segmentation tasks

The medical research project deals with two types of adipose tissue. The two types are Subcutaneous Adipose Tissue denoted SAT and Visceral Adipose Tissue denoted VAT. Figure 2.2 shows a slice of a patient that has been segmented by hand by an expert. Note that the data is from another research project. The segmentation is done into 5 main classes: subcutaneous, retroperitoneal,

intraperitoneal, bone and muscle. The retroperitoneal and the intraperitoneal areas are combined to form the region that will be denoted as VAT in this report. The subcutaneous class corresponds to the area of the SAT. The task at hand is to construct an automatic pipeline that can perform a classification of all voxels in the anatomically bounded unit into the 3 classes: SAT, VAT and other tissue.




Figure 2.2: A segmentation of a slice in the abdomen done by hand. Light blue is the subcutaneous region, pink is the intraperitoneal region, yellow is the retroperitoneal region, blue is bone and red is muscle tissue.

2.3 Image variation and quality

Figure 2.3 shows a sample of the variation and quality that can be expected from the data. As can be seen the structure and quality of the images varies greatly. The main challenge in creating an automatic pipeline for segmenting the MRI data, is to create a solution that is robust across all possible variations of the external and internal anatomy of the patients. To test the performance of each step in the segmentation process a test set of 80 patients has been collected. This set is chosen at random from the 300 patients in the data set. After each segmentation step the performance on the patients in the test set has been evaluated to ensure that all solutions are robust to the variation of the

patients anatomy and the variation in image quality.




Figure 2.3: An example of the variation in the data. The 6 images are slice 1 from patient: 5198, 5251, 5321, 5367, 5415 and 5480 from top to bottom left to right.

Preprocessing

3.1 Bias field correction

A common problem when dealing with MR images, is the non uniformity in the intensity of same tissue voxels across a slice. This field of biased intensity values, denoted a bias field, usually varies slowly across an image and is caused by poor radio frequency coil uniformity and patient anatomy both inside and outside the field of view. The magnitude of this intensity variation is in the 10-20% range according to Sled and Zijdenbos [1].

The bias field has little impact on visual inspection and can even be hard to notice with the human eye, but it is critical that same tissue voxels have similar intensities across the image if intensity based classification techniques are to be successful.

The fact that the bias field is attributed by the patients geometry and electrical properties makes the field unique for each patient and thus impossible to predict. Hence, a way of estimating the bias field is needed in order to be able to remove it.

The MRI data used in this project have another source of non uniformity in same tissue voxel intensities. This source appears as large peaks in the intensity

of adipose tissue. These peaks are generally fast varying and appear most often as a large wide peak at the bottom of the image or as smaller steeper peaks to the left or right in the image. The peaks always appear over voxels with already high intensities. Figure 3.1(a) shows an example of an image with the large peak at the bottom, and Figure 3.1(b) shows an image with both a bottom and a right peak. Note that the color map used for these images will be used throughout the chapter when plotting images. The intensities in images will always be scaled when plotting images. Thus, Red will always indicate the highest intensity in the image and blue the lowest intensity.




Figure 3.1: (a) An example of an image with a large intensity peak at the bottom. The image is slice 1 from patient 5047. (b) An example of an image with both a bottom intensity peak and a right intensity peak. The image is slice 7 from patient 4950.

After speaking with the medical staff at "Sygehus Fyn" a possible source for this non uniformity was given. The problem lies mainly in the fact that the study includes some very obese patients and they have problems fitting correctly in the scanner. These peaks are however present on all images in some form and hence also on skinny patients. A more complete explanation might be that the equipment is not properly calibrated and or faulty in some way or form. The cause of the peak problem has not been further investigated and emphasis has instead been put on removing the peaks.

Figure 3.2 shows histograms of the intensities in the image on Figure 3.1(a). By inspecting the histogram on Figure 3.2(a) it is seen that the bias peak values have much greater intensity values than the rest of the voxels. It is further observed on Figure 3.2(b) that the bias peak intensity values are spread out over a large range from about 1400 to 3000 in this case. Most of the voxels on Figure 3.2(a) are located near or close to zero, these represent the large amount of background present in the image. Figure 3.2(c) shows a closer inspection of

the histogram in the range 200-1400. This range is where most of the desired information is located, namely the fat and muscle tissue. It is desirable to be able to identify two distinct peaks in this range, corresponding to the adipose tissue and muscle tissue. A clear distinction of these two tissue peaks will greatly improve the results of the intensity based classifier described in Chapter 4. As can be seen this distinction is, at best, hard to do on the uncorrected image. This is due to the slowly varying bias field across the image. It should be noted that the value ranges for the histograms have been chosen by visual inspection of Figure 3.2(a).




Figure 3.2: (a) The full histogram of voxel intensities on slice 1 from patient 5047. (b) The same histogram in the value range 1400-3000. (c) The histogram in the range from 200-1300

There are several problems that need to be addressed for the removal of the bias field. The problems can be divided into two main areas.

- Finding voxels that should have the same intensity.

- Estimating the field from the varying intensities of these voxels.

The next two sections will address each of the problems in turn.

3.2 Finding adipose tissue Voxels

Adipose tissue voxels are used for the bias field estimation since these are the target for the segmentation task and they have the nice property of having the highest general intensity values on T1 weighted MR images. A technique which has been used for similar bias field correction tasks by Engholm et al. [2] involves finding the outer edge of the patients body in the image and then expect adipose tissue voxels to be the voxels on the immediate inside of this edge. This technique works well on images where a significant amount of subcutaneous adipose tissue is present along the entire outer edge of the patients body. This technique fails however on the images used for this study. There are two reasons for this. First, there are some very skinny men included in the study which have next to no subcutaneous fat. This will cause the before mentioned method to base the bias correction on voxels that will have a general lower intensity value than true adipose tissue voxels, thus giving erroneous bias estimation. Second, the bias peaks mentioned above will not be sampled sufficiently since the 2-dimensional extent of the peaks will not be covered.

A new method for finding adipose tissue voxels is needed. This method should be able to sample the intensity values of the large bias peaks, but also give evenly distributed samples of adipose tissue voxel intensities across the entire image. The method used for finding adipose tissue voxels that tries to fulfill these goals is described below.

3.2.1 Methodology

The images contain a large amount of background voxels that contain no useful information. The first task is thus to find a smaller region of interest (ROI) where the search for adipose tissue voxels will be performed. The outline of the patient is represented by low intensity background voxels on the outside and high intensity skin/adipose tissue voxels on the inside. This means that the gradient image will have high values on the outline of the patient. The outline is easily found using simple dynamic programming, driven by high gradient values. Dynamic programming is covered in detail in Chapter 6 and will therefore not be covered further here. Figure 3.3(a) shows the outline found using dynamic

programming. The inside of the outline is then filled to form the Region Of Interest (ROI) mask showed on Figure 3.3(b).




Figure 3.3: (a) The outline found using gradient driven dynamic programming. (b) The ROI mask that has been constructed from the outline.

Inside the ROI all local intensity maxima are located. As can be seen on Figure 3.4 this yields points which clearly represent adipose tissue voxels but also points which represent non adipose tissue voxels. These points thus needs to be trimmed in a way that will only leave us with true fat voxels.

To trim out the non adipose tissue voxels two properties of the bias field and adipose tissue voxels will be used. First, the bias field varies slowly across the image (except for the bias peaks), this means that within a small subregion in the image, adipose tissue voxels tend to have very similar intensities. Second, adipose tissue voxels have the highest intensity values. This leads to the solution to the problem: Divide the image into smaller regions and find the voxels in each region which have high intensity values relative to the other voxels in the region.

In practice this is done by first creating the smallest box that contains the entire ROI. This box is then further divided into smaller boxes. n_r box rows are created vertically and n_c box columns are created horizontally. The boxes overlap by o_r voxel rows and o_c voxel columns. The principle is illustrated on Figure 3.5. The n_r and n_c parameters need to be attuned to a size where it is impossible or at least very unlikely, to find a location anywhere inside the ROI where such a small box can be placed without it overlapping at least one high intensity voxel. If this is not the case a small box could be placed in a region with only low intensity non fat voxels, which would then be passed on as being fat voxels for the purpose of the field estimation. For the MR image material




Figure 3.4: All local intensity maxima in the image showed as red dots

in the study, parameter values of $n_r = 12$ and $n_c = 8$ give good results. Both o_r and o_c is set to 5 voxels.

For each small box the local maxima voxels that lie inside this box are retrieved. Of these, the value of the voxel with the maximum intensity is stored in the variable I_{max} . An intensity percentage threshold is defined p_t . All retrieved voxels that does not satisfy

$$I_j > I_{max} \cdot p_t$$




Figure 3.5: (a) The principle behind the box subdivision of the ROI, only 3 boxes are shown. (b) Up close version of (a), with the parameters shown. o_r and o_c are the overlap vertically and horizontally. n_r and n_c are the number of boxes vertically and horizontally.

where I_j is the intensity of the j 'th voxel inside the box, are discarded. $p_t = 0.85$ was found to give good results. Figure 3.6 shows the points that are left after the trimming. Some of these might look dubious in their location, but by inspecting the zoomed in view on Figure 3.7 it is seen that these locations indeed correspond to high intensity voxels.

Since the large bottom bias peak (see Figure 3.1(a)) is present on almost all images, and always in the bottom subcutaneous adipose tissue layer, a dense sampling of this area is desirable. The method that was first described and discarded in Section 3.2 above is now used, but only in the lower part of the image. In this lower part a thick subcutaneous adipose tissue layer is always present. For each point on the outline of the ROI, the position of the maximum intensity value between the outline and 10 pixels towards the center of the ROI is sampled. The principle is shown on Figure 3.8. The result of adding these points can be seen on Figure 3.9. These points are then added to the points found using the local maxima method.

There is a few problems with the gathered points. They are not equally distributed, rather they tend to cluster together. Furthermore there are quite a few points, which might give unnecessarily large computational load during the field estimation. The points are further trimmed again using boxes. The box containing the ROI is divided into 10 times 10 boxes, and all but the highest




Figure 3.6: The points left after the points shown on Figure 3.4 have been trimmed

intensity sample inside each box is trimmed away. This has the effect of putting an upper bound of 100 on the number of points. It removes many points where the points are tightly clustered and removes few or no points where the points are sparse. The final result can be seen on Figure 3.10.




Figure 3.7: Zoomed in view of the area below the spine on Figure 3.6




Figure 3.8: (a) The principle of finding adipose tissue locations. (b) A zoomed in view of the same image. The green star is the outline point of the ROI, the blue dots are the 10 search points from the outline point towards the ROI center. The red star is the maximum intensity voxel found.




Figure 3.9: (a) The extra fat pixel locations found using the outline inwards search method. (b) A zoomed in view of the right part of (a). Note how points are only added in the lower part of the image where a thick subcutaneous layer of fat is always present.




Figure 3.10: The points left after adding the points from 3.6 and 3.9(a) and subsequently spatially trimming these

3.2.2 Results

Figure 3.11 and 3.12 show the end result on slice 1 and 7 from 5 patients. Note that these patients are not chosen to have either a particular good or bad result, they are simply the first 5 patients in the test set following patient 4950.

Overall the method gives a good sample of evenly distributed high intensity voxels. The bottom bias peak is generally densely sampled but the side bias peaks tend to only have a single sample at their top. This leads to some minor problems later as will be described in the following section. A change that would make the algorithm not trim away points in the vicinity of a peak could be a solution to this. This however raises new problems of determining peak locations and the algorithm might give bad results on future images without any bias peaks. The algorithm described above will work equally well on future data without the bias peaks.

Since the algorithm finds high intensity voxels in general, it will not only find voxels that actually represent adipose tissue but also other tissue types that have high intensity on T1 weighted MR scans. This will generally not influence the bias estimation much since these voxels have intensity values close to the adipose tissue voxel intensities. These tissue types will have to be segmented using non threshold techniques and it is thus of little consequence that they have the same intensity as true adipose tissue.




Figure 3.11: The end results of the high intensity voxel finding algorithm. (a) through (e) are slice 1 from patients 4951, 4952, 4953, 4954 and 4955.




Figure 3.12: The end results of the high intensity voxel finding algorithm. (a) through (e) are slice 7 from patients 4951, 4952, 4953, 4954 and 4955.

3.3 Estimating the bias field

An evenly distributed sample of high intensity voxels have been extracted from the original image. These voxels should all have the same intensity values if the bias field was not present. In order to remove the bias field it must first be known how it varies across the image. To model the field we fit an interpolated smooth surface to the sampled high intensity voxels using thin plate splines.

3.3.1 Thin Plate Splines (TPS)

Thin plate splines were first introduced by Duchon [3] in 1976. An in depth coverage of TPS on finding smooth interpolations of sparse data has been done by Green and Silverman [4]. The idea behind TPS is to model the behavior of infinitely thin plates of metal when forced through points in 3D space. Metal plates forced through specific points in this way will exhibit minimum bending energy. A mathematical formulation of the modelling of metal plates will thus give us the interpolation between points with minimum bending energy. The formulation of TPS carries over to N dimensional space. The theory will first be formulated in 2D and then later be extended to 3D.

Assume N observations in \mathbb{R}^2 , with each observation \mathbf{x} having coordinates $[x_1 \ x_2]^T$ and values z . We seek to find a function f , that describes a surface that passes through these points with minimal bending energy. The problem is formulated as

$$\min_f \sum_{i=1}^N \{z_i - f(\mathbf{x}_i)\}^2 + \lambda J(f) \quad (3.1)$$

where $J(f)$ is a function for the curvature of f :

$$J(f) = \int \int_{\mathbb{R}^2} \left(\frac{\partial^2 f}{\partial x_1^2} \right)^2 + 2 \left(\frac{\partial^2 f}{\partial x_1 \partial x_2} \right)^2 + \left(\frac{\partial^2 f}{\partial x_2^2} \right)^2 dx_1 dx_2 \quad (3.2)$$

λ is a parameter that penalizes for curvature. With $\lambda = 0$ there is no penalty for curvature, this corresponds to an interpolating surface function where the function passes through each observation point. At higher λ values the surface becomes more and more smooth since curvature is penalized. For λ going to

wards infinity the surface will go towards the plane with the least squares fit, since no curvature is allowed.

According to Green and Silverman [4] f is of the form

$$f(\mathbf{x}) = \beta_0 + \beta_1^T \mathbf{x} + \sum_j^n \alpha_j \eta(\|\mathbf{x} - \mathbf{x}_j\|) \quad (3.3)$$

where the η function is defined as

$$\eta(r) = \begin{cases} r^2 \log(r^2) & , r > 0 \\ 0 & , r = 0 \end{cases} \quad (3.4)$$

We now have N equations, one for each observation, but we need to estimate $N + 3$ variables. N variables α_j , one variable β_0 and two variables β_1 . The last 3 equations we get from the 3 linear constraints

$$\sum_{j=1}^N \alpha_j = \sum_{j=1}^N \alpha_j x_{j1} = \sum_{j=1}^N \alpha_j x_{j2} = 0 \quad (3.5)$$

that ensures that the $J(f)$ function is finite. To solve the system of equations we write the system on matrix form. First the matrices

$$\mathbf{P} = \begin{bmatrix} 1 & \cdots & 1 \\ \mathbf{x}_1 & \cdots & \mathbf{x}_N \end{bmatrix} \quad (3.6)$$

and

$$\mathbf{E}_{ij} = \eta(\|\mathbf{x}_i - \mathbf{x}_j\|) \quad (3.7)$$

are defined. The system can then be written as

$$\begin{bmatrix} \mathbf{E} + \lambda \mathbf{I} & \mathbf{P}^T \\ \mathbf{P} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} \mathbf{Z} \\ \mathbf{0} \end{bmatrix} \quad (3.8)$$

where $\mathbf{Z} = [z_1 \cdots z_N]^T$, $\alpha = [\alpha_1 \cdots \alpha_N]^T$ and $\beta = [\beta_0; \beta_1]$. The first line in the matrix equation is the interpolation and smoothing equations and the second line is the constraints. This matrix equation is solved with respect to α and β . An estimate of the TPS at the location \mathbf{x} can now be calculated using Equation 3.3.

To extend the formulation above to 3D only slight changes are needed. We have N observations in \mathbb{R}^3 , with each observation \mathbf{x} having coordinates $[x_1 \ x_2 \ x_3]^T$ and values z . The only major change to the formulation above is that 3.4 becomes

$$\eta(r) = \begin{cases} r^3 & , r > 0 \\ 0 & , r = 0 \end{cases} \quad (3.9)$$

everything else extends trivially to 3D.

3.3.2 Removing the bias field

The bias field is estimated at each voxel location in the patient volume. Having obtained an accurate estimation of the bias field present on the data, the effect of the bias field can now be removed. This is done by dividing the original voxel intensity values with the value of the field estimate at the corresponding locations.

$$I_{cor_j} = \frac{I_{org_j}}{I_{est_j}} \quad (3.10)$$

where I_{cor_j} is the corrected intensity value at location j , I_{org_j} is the intensity of the voxel at location j and I_{est_j} is the bias field estimate at location j . If the bias field estimate is accurate this will yield values close to the 0 to 1 range. Due to noise and inaccurate bias field estimates the actual range will differ slightly. an inaccurate bias field estimate might cause for instance a high original intensity value to be corrected using a too low bias estimate, resulting in a value higher than 1.

3.3.3 Effective degrees of freedom (df_λ)

Determining a good value for λ is an important task in order to estimate an accurate bias field. On one hand the field should be slowly varying and rigid to describe the general bias field and to lessen the effect of random noise on the observations. On the other hand the field should be able to form surfaces that can form a tight fit to the large bias peaks. There is however a problem in determining a set value for λ for all bias fields across all images. Since a certain value of λ will have a different smoothing effect on different observation sets. A method independent of the overall intensity level of the sampled voxels is thus needed to determine λ .

The notion of *effective degrees of freedom* (df_λ) is introduced. Hastie et al. [5] describes a correspondence between λ and a measure similar to degrees of freedom. This measure called effective degrees of freedom gives a more intuitive description of the amount of curvature penalized. For instance $df_\lambda = 3$ would be a plane for a 2D TPS, corresponding to 3 degrees of freedom. If df_λ equals the number of observations the estimated field would be an interpolation passing through all observation points.

Having estimated α and β , the estimate of the target function can be written as

$$\hat{Z} = \begin{bmatrix} \mathbf{E} & \mathbf{P}^T \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} \mathbf{E} & \mathbf{P}^T \end{bmatrix} \begin{bmatrix} \mathbf{E} + \lambda \mathbf{I} & \mathbf{P}^T \\ \mathbf{P} & \mathbf{0} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{Z} \\ \mathbf{0} \end{bmatrix} = H_\lambda \begin{bmatrix} \mathbf{Z} \\ \mathbf{0} \end{bmatrix} \quad (3.11)$$

We call H_λ the hat matrix since its the matrix that puts the "hat" on Z . Hastie et. al defines the *effective degrees of freedom* df_λ as the trace of the hat matrix

$$df_\lambda = tr(H_{\lambda sq}) \quad (3.12)$$

where $H_{\lambda sq}$ is the square part of H_λ that corresponds to the Z values. By specifying df_λ instead of λ directly the rigidity of the field can now be specified. df_λ can only be determined from a set λ value, thus numerical methods has to be used to allow for the df_λ value to be specified as an argument. In practice this has been done by a simple bisection algorithm that calculates df_λ from a given start value of λ and then modifies λ depending on the value of df_λ this results in.

3.3.4 Retrieving intensity values

After inspecting histograms and images of MRI data from 80 patients in the test set an interesting discovery was made. Voxels not located on bias peaks always have intensity values below 1500 as was also hinted on Figure 3.2. Since the peaks only cover voxels that already have high intensity, the locations of all voxels with intensities above 1500 can simply be saved for further use in the tissue classification. For use in the bias estimation all intensities from the locations found above in section 3.2 are cut off at the 1500 level. Thus intensities higher than the cutoff will have their intensity fixed at 1500 instead. This greatly reduces the field estimation errors that are caused by too sparse sampling near the base of the bias peaks or too rigid a field estimation.

The effect of the threshold cutoff on the sparse sampling problem is illustrated on Figure 3.13. The figure shows the effect of applying the cutoff to the bias peaks. This is a purely synthetic one dimensional example using an interpolating spline. The corrected intensities are obtained by dividing the original intensities with the estimated bias field, thus with a perfect field estimation all the corrected intensities would have the value 1. By comparing 3.13(c) and 3.13(d) it is seen that applying the threshold on the bias peaks gives corrected values closer to the desired target value of 1. The problem is not completely eliminated by making the cutoff, since there is still a fast change in curvature where the intensities were cutoff, but the problem is greatly reduced.

To determine an optimal value of df_λ a parametric investigation of df_λ is performed. Figure 3.14 shows part of the parametric investigation of the df_λ parameter on slice 1 from patient 4950. It is seen from Figures 3.14(a)-(c) that $df_\lambda = 5$ gives too rigid a field to be able to give a good bias estimate over the cutoff peaks. The two peaks in the histogram are not clearly defined or well separated. Furthermore the bias peaks are still clearly visible on the corrected image (3.14(b)). As df_λ is increased the bias field estimate becomes faster varying and with greater curvature. Increasing the df_λ value from 80 to 160 gives only minor improvements. After investigation of the full parametric investigation $df_\lambda = 80$ is chosen as a good weighting between having a fast varying field that can fold around the bias peaks and still having some degree of curvature penalization to get a good estimation of the slow varying part of the field.




Figure 3.13: The principle behind the bias field estimation on a bias peak in one dimension with and without a threshold applied at the value 1.2. On (a) and (b) the red line is the voxel intensities, the green crosses are the sample points and the blue line is the interpolating spline through these samples. (c) shows the corrected intensities from the non threshold bias peak on (a). (d) shows the corrected intensities from the threshold bias peak on (b).




Figure 3.14: A parametric investigation of the df_λ parameter. The first column shows the estimated bias field in the ROI. The second column shows the corrected intensities and the final column is the histogram of intensities in row 2. The four rows has df_λ values of 5, 20, 80 and 160 from top to bottom.

3.3.5 Methodology

This section will recap all the methods described above and provide an overview of what is being done in the bias correction pipeline.

- local high intensity voxels are found on all slices from one patient, using the techniques described in Section 3.2.
- The intensity level is cut off at 1500, to remove the top of the bias peaks.
- The bias field is estimated over all voxels in the entire volume using thin plate splines extended to 3 dimensions, as described in Section 3.3.1.
- The corrected volume is obtained by dividing the original volume data with the estimated bias field. Described in Section 3.3.2

The estimation of the bias field could have been done slice by slice instead using 2-dimensional TPS. However, it seems intuitive to perform the bias correction on the whole volume at once, since volume data is available. Furthermore the bias field is expected to vary slowly between neighboring slices, just as it varies slowly across voxels within a slice. This makes whole volume bias field estimation a sensible choice.

3.3.6 Results

Examples of the final result of the bias correction on 3 different patients on slice 1 and 7 can be seen on Figure 3.15 and Figure 3.16. It can be seen from the before and after histograms that the highest valued peak, corresponding to the adipose tissue voxels, is clearly identified on the histograms from all the results. The peak representing muscle tissue is easily identified on the slice 1 results, however it is hard to distinguish on the slice 7 results. This is due to the presence of more non adipose tissue with non uniform intensity values. However, as the overall goal is to separate the adipose tissue from all other tissue, the lack of a clear muscle peak in the histogram is of less importance. The important result of the bias correction is that adipose tissue voxel intensities are clearly distinguishable from the low intensity tissue. By visual inspection of the before and after bias correction images it can be seen that the intensities in areas where adipose tissue is expected have become much more uniform. Note that the scale on the before and after bias correction images are different. The color of the images can not be compared, only the uniformity of the intensities.

It should be noted that the method for removing the bias field presented here will work equally well on data with or without the large bias peaks present on the data at hand. The method is thus robust with regards to future unseen data that might exhibit new forms of bias field variation.

One of the major advantages the described sampling method has over similar methods is that it also samples high intensity voxels from the interior of the patient, thus allowing for a more accurate bias field estimation across the entire ROI. One thing that could cause the method to fail, or at least give worse results, would be the presence of noise with large variation. Since all voxels sampled are local maxima, the sampled voxels will often be from a local peak in the noise contribution. With too high variation in how large this noise contribution is the bias field estimation will give bad results, since the sampled voxel does not represent the local level of the bias field. This could possibly be fixed by applying a smoothing filter to the data before sampling the voxel intensities, but will introduce new problems by lowering the value of high intensity voxels close to low intensity voxels.

Overall the method is robust and performs well. the bias estimates found are not perfect but they are good enough to make the corrected data highly suitable for an intensity based classification, as will be seen in Chapter 4.

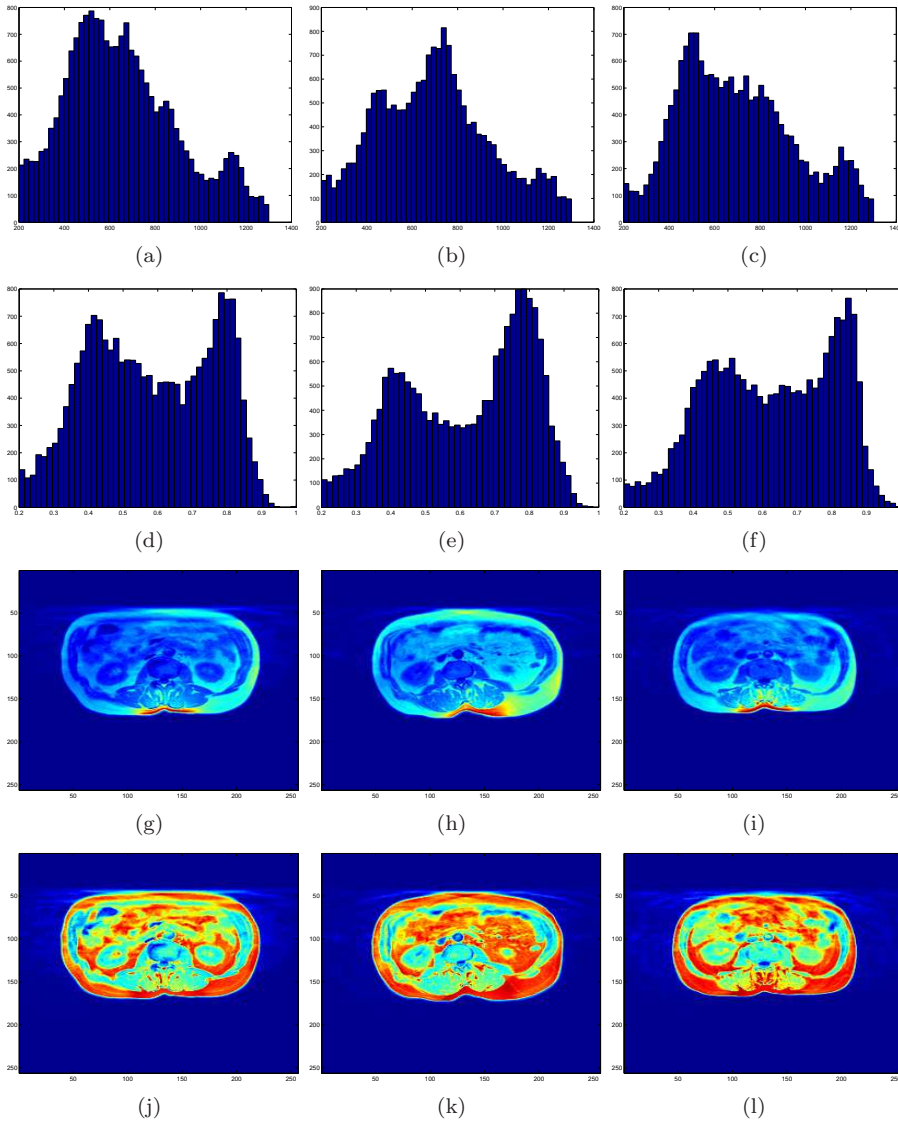


Figure 3.15: The results of the bias correction on 3 patients. The 3 columns are slice 1 from patients 4951, 4952 and 4953. The top row is the histogram of the original image, these histograms are cropped in the same way as Figure 3.2(c) to show the interesting range more clearly. The second row is the histograms after the bias correction has been done. The lower range containing all the background has been cropped away. The third row is the original biased image and the last row is the image after the bias correction has been performed.




Figure 3.16: The results of the bias correction on 3 patients. The 3 columns are slice 7 from patients 4951, 4952 and 4953. The top row is the histogram of the original image, these histograms are cropped in the same way as Figure 3.2(c) to show the interesting range more clearly. The second row is the histograms after the bias correction has been done. The lower range containing all the background has been cropped away. The third row is the original biased image and the last row is the image after the bias correction has been performed.

Distinguishing adipose tissue from other tissue

After having corrected for the bias field in Chapter 3 the data is ready to be classified using an intensity based classifier. This chapter deals with determining which voxels correspond to adipose tissue. The location and thus type of adipose tissue will not be dealt with in this chapter.

Since only the T1-weighted images of the data has been provided for the project multi modality classification techniques can not be used. The absence of ground truth further limits the methods that can be used. A commonly used technique such as *artificial neural networks* could have been a good technique if ground truth had been available. The network could have been trained with inputs from not only the current slice, but also the neighboring slices, thus incorporating volume information in the classification.

Good results have often been achieved by using simple thresholding techniques, and this is also the method that is used here. This method assigns labels to voxels by comparing their intensity values to one or more intensity thresholds. A single threshold segments the image into two classes, multiple thresholds can be used to segment into more classes. These thresholds can be either static or spatially varying. Since the spatial variation was handled with the bias field correction non varying thresholds will be used. The thresholding technique used will be point based and no information from its neighborhood will be used. The

idea is that the spatial context will be applied using other techniques in later chapters.

4.1 Fuzzy c-mean clustering

In order to determine the optimal threshold that will give the best segmentation the technique of *fuzzy c-mean clustering* (FCM) is used. This technique is described in [6] and was used by Positano et al. [7] to segment adipose tissue on MRI data with good results.

The FCM algorithm does not directly determine a threshold that segments the voxels. Instead it performs a fuzzy segmentation where each voxel has a fuzzy membership function constrained to be between 0 and 1. This function reflects the similarity between a given voxel and the typical data value of its class. For instance, a membership value close to 1 means that the voxels intensity is close to the centroid of that class.

The FCM algorithm is formulated as the minimization of the following objective function.

$$J_{FCM} = \sum_{j \in \Omega} \sum_{k=1}^C u_{jk}^q \|y_j - v_k\|^2 \quad (4.1)$$

where j is a location in the image domain Ω , k is the class number, C is the number of classes and q is a parameter greater than 1 that determines the amount of fuzziness of the classification. u_{jk} is the membership value at location j for class k , y_j is the intensity value at location j and v_k is the centroid of class k .

The minimization of J_{FCM} is done by suitably selecting u and v using an iterative process of evaluating the following equations:

$$v_k = \frac{\sum_{j \in \Omega} u_{jk}^q y_j}{\sum_{j \in \Omega} u_{jk}^q} \quad (4.2)$$

$$u_{jk} = \frac{\|y_j - v_k\|^{\frac{-2}{q-1}}}{\sum_{k=1}^C \left[\|y_j - v_k\|^{\frac{-2}{q-1}} \right]} \quad (4.3)$$

u_{jk} is initialized with random values, but under the constraint that the sum of the membership functions for each class for a given location is 1. That is, for each $j \in \Omega$

$$\sum_{k=1}^C u_{jk} = 1 \quad (4.4)$$

The algorithm alternates evaluating equation 4.2 and 4.3 until the change, ΔJ , in J_{FCM} is suitably small.

Figure 4.1(a) shows the similarity measure as a function of voxel intensity for each class for slice 1, patient 4950. This plot is superimposed on the histogram of voxel intensities in the slice. It can be seen that the peak of the similarity measure for each class is well separated and generally follows the 3 peaks in the histogram. Note that the measuring on the y-axis for Figure 4.1(a) is number of observations. The similarity measure has values between 0 and 1. Figure 4.1(b) shows the convergence of the objective function J_{fcm} . The objective function converges to a steady level after 8 iterations. After 18 iterations the change in J_{fcm} is suitably small, meaning that the change is smaller than ΔJ , and the computation stops.

Figure 4.2 shows how the similarity measure behaves during the convergence. At 0 iterations corresponding to the initialization all the similarity measures are chosen at random. After two iterations the peaks of the 3 classes are still very close together. After 4 iterations some separation of the 3 peaks start to become clear, and after 8 operations it can be seen that the graphs start to resemble the graphs on Figure 4.1(a). The algorithm will generally converge to a state where the peaks of the similarity measure function are near peaks in the histogram of the data intensities, and with maximum separation between the similarity measure peaks.

Figure 4.3 shows an investigation of the effect of the q parameter. It is seen that while the q parameter determines the shape of the similarity measure curves, the intersection between the curves are unaffected. $q = 2$ is chosen for further computations. The curves for $q = 2$ resembles gaussian distributed probability curves, which can be exploited as described later. For $q = 2$ Equation 4.3






Figure 4.1: (a) The similarity measure as a function of voxel intensity for each class superimposed on the histogram of the image. Note that the similarity function has values between 0 and 1. Data is slice 1 from patient 4950, $C = 3$ and $q = 2$. (b) The convergence of J_{fcm} .

becomes a calculation of squared distance measures, which is a commonly used measure.


Three distinct classes are identified on the MR images. These classes are denoted **background** (low intensity), **adipose tissue** (high intensity) and **other tissue** (medium intensity). Three classes are chosen because the histograms of intensities on the images after the bias correction has 3 distinct peaks. This makes the FCM algorithm with $C = 3$ a well suited classification method. The similarity measure that was investigated above allows us to use two classification schemes. Discrete classification and fuzzy classification.




(a)



(b)



(c)



(d)

Figure 4.2: The similarity measure as a function of voxel intensity for each class at different stages of convergence




Figure 4.3: The Effect of varying the q parameter. From left to right, top to bottom, the q values are 1.1, 1.5, 2 and 3. Data is slice 1 from patient 4950 and $C = 3$

4.1.1 Discrete classification

By using this classification scheme each voxel will be assigned the label that it has the highest similarity with. This corresponds to setting a threshold at the intersection between the similarity measure curves for each class. All the voxels classified as adipose tissue in this way is simply counted to get a measure of the ratio of fat in the specific slice. Let M_{da} be the measure of adipose tissue using discrete classification, and n be the number of voxels in the the slice. The total measure of adipose tissue is then:

$$M_{da} = \sum_{j=1}^n f(j) \quad (4.5)$$

where

$$f(j) = \begin{cases} 1 & \text{if } (u_{ja} \geq u_{jb}) \wedge (u_{ja} \geq u_{jo}) \\ 0 & \text{if } (u_{ja} < u_{jb}) \vee (u_{ja} < u_{jo}) \end{cases}$$

u_{ja} , u_{jo} and u_{jb} are the similarity measure at location j for the adipose tissue class, other tissue class and background class respectively.

Figure 4.4 shows a color coded classification of the first slice from patient 4950, both before and after the bias correction. The true value of the bias correction can be seen here. The classification of the uncorrected image fails completely while it gives a good result on the corrected image.

At times it might be more costly to classify a certain voxel to a wrong class than to not have it classified at all. Instead of always assigning the class with the highest similarity measure to a given voxel a threshold for the similarity measure level needed can be used instead. using a similarity measure threshold higher than 0.5 will make the voxels classified more probable of being the correct class. Let ε_{sm} denote the similarity measure level threshold, then the new formulation of $f(j)$ becomes:

$$f(j) = \begin{cases} 1 & \text{if } u_{ja} \geq \varepsilon_{sm} \\ 0 & \text{if } u_{ja} < \varepsilon_{sm} \end{cases}$$

Note that this method can only be used to segment into two classes, since the




Figure 4.4: (a) and (b) are the MR images before and after the bias correction on slice 1 from patient 4950. (c) is the classification of the uncorrected image using FCM. $C = 3$ classes have been used and $q = 2$. (d) is the classification of the corrected image using the same parameter values.

scheme only distinguishes between being a certain class and not being that class. Figure 4.5 shows a classification of the adipose tissue class for different values of ε_{sm} . As can be seen from the figure the changes in the classification of adipose tissue are only minor for ε_{sm} in the 0.5-0.7 range. Only for $\varepsilon_{sm} = 0.9$ do large changes start to become evident. This is reassuring since it shows that the classification of adipose tissue is robust to small changes in the intersection location of the similarity measure curves. This way of doing a more certain classification of the adipose tissue class will prove useful in a later chapter.




Figure 4.5: The effect of varying the similarity measure value needed to classify as adipose tissue. Red areas are adipose tissue classification and blue is non adipose tissue. Data is from slice 1 from patient 4950, $C = 3$ and $q = 2$

4.1.2 Fuzzy classification

By using fuzzy classification the similarity measure is used as a measure for the degree of the partial volume effect. The partial volume effect is the effect on the intensity of a voxel that different tissue types within a voxel volume gives. This scheme will give a high similarity measure the interpretation that the voxel contains almost pure fat, and a low similarity measure the interpretation that the voxel contains a lot of non fat tissue. The total ratio of adipose tissue is then calculated as the sum of the similarity measure for all voxels. Let M_{fa} denote the measure of adipose tissue using fuzzy classification. The total measure of adipose tissue is then:

$$M_{fa} = \sum_{j=1}^n u_{ja} \quad (4.6)$$

It only makes sense to use this measuring scheme in areas that are known to contain adipose tissue, since the similarity measure is not strictly zero even at low intensity values. Figure 4.6(a) shows a fuzzy classification. The fuzzy classification generally resembles the discrete classification in its results. The fuzzy classification gives high values where the discrete classification has classified the voxel as adipose tissue. There is however a few problems where intensity artifacts from the original image or created by the bias correction are present. The artifacts make the fuzzy classification give a voxel a smaller value where the discrete classification classifies the same voxel as adipose tissue.

This method will not be used for this project. Since no ground truth is available the accuracy of the results gotten from the fuzzy classification will be much harder to estimate than the results from the discrete classification. Given that the accuracy of the method could be verified, a comparison between the accuracy of the discrete and the fuzzy classifications would be interesting.




Figure 4.6: (a) The fuzzy segmentation of slice 1 from patient 4950. The intensities in the image reflects the degree of the partial volume effect. (b) The segmentation of the same image done using discrete segmentation.

4.2 Results

Figure 4.7 and 4.8 show the result of the FCM classification on slice 1 and slice 7 from patients 4951, 4952 and 4953. The discreet classification scheme was used with highest similarity (Equation 4.5). Without being an expert it is still possible to get a good idea of the performance of the method by comparing the classification to the bias corrected image. There is a good correspondence between what voxels one would expect to be adipose tissue on the bias corrected image and what voxels that gets classified as adipose tissue using the FCM algorithm. An amount of voxels that clearly are neither Subcutaneous Adipose Tissue (SAT) or Visceral Adipose Tissue (VAT) are classified as adipose tissue by the FCM. The adipose tissue class will be further segmented into a SAT, a VAT and a neither class in subsequent chapters. After showing the results to Kristian Wraae he has assured that the classification generally performs very well.

The method is completely automatic and very robust. Further more it requires no data dependent parameters to be set. Thus the method should work equally well on future data without any parameter adjustments. Overall the FCM algorithm performs satisfactory.




Figure 4.7: From top to bottom the data is slice 1 from patient 4951, 4952, 4953. The right column is the bias corrected image, the left column is the result of the FCM classification. Blue is background, red is adipose tissue and green is other tissue.




Figure 4.8: From top to bottom the data is slice 7 from patient 4951, 4952, 4953. The right column is the bias corrected image, the left column is the result of the FCM classification. Blue is background, red is adipose tissue and green is other tissue.

Finding image structures

In the previous chapter the segmentation of adipose tissue was done. However, the distinction between Subcutaneous Adipose Tissue (**SAT**), Visceral Adipose Tissue (**VAT**) and neither has yet to be made. To segment the VAT and SAT, 3 main separations must be made. Figure 5.1 shows a rough hand drawing of the 3 outlines that are wanted. The first contour (red) separates the SAT area from the background. The second contour (green) outlines the internal limit of SAT. Finally the blue contour outlines the VAT area. The separation of the SAT from the background was already done when finding the area of interest in Section 3.2.1. The separation will however be performed again with a greater level of robustness in the following chapters.

The segmentation approach that will be used can be described by 3 steps. The first step is to obtain a rough estimation of points that mark the outline of the 3 contours. Secondly these points will be used to guide more precise contours resembling the contours on Figure 5.1 (described in Chapter 6). Finally the segmentation will be performed by using set operations and connectivity on the areas enclosed by the 3 contours (described in Chapter 7). This chapter describes the first step.




Figure 5.1: The 3 desired separations. The Red outline is the external contour of SAT, the green outline is the internal contour of SAT and blue is the contour of the VAT area border.

5.1 Active Shape Models - (ASM)

Active shape models were introduced in the early 1990's by Tim Cootes. It has since then found wide applications in image analysis and computer vision. The theory behind the implementation used for this project is based on an overview article, Cootes [8].

The human brain can easily recognize known shapes in an image, just as it was very trivial to draw the 3 contours on Figure 5.1 by hand. Active Shape Models tries to mimic this ability. The approach is to build a model of the structure that the computer should be able to recognize, by providing a training set of typical images. By means of this training set a statistical model of the shape and variation of the structure can be constructed. This model can then be used to locate similar structures on unknown images.

5.1.1 Building the training set

The training set is a set of images that are representative for the variation of structures across all images. These images are annotated with landmarks that outline the structures that we wish to identify.

Good choices for landmarks are points that consistently can be located across

images. These landmarks are placed by hand at locations where the structure changes direction, or evenly distributed between such locations if the structure consists of too few such locations. The structures that are sought for this project are the outlines of the 3 curves from Figure 5.1. The outline of the external SAT boundary and the internal SAT boundary is generally easy to identify, since a major shift in intensity always is present at these boundaries. The outline of the VAT area however has poor contrast in the upper part of its boundary and the intensities in the border of this region generally show a high degree of variation. This results in a choice to only landmark the lower well defined part of the VAT area. Figure 5.2 shows the first annotated image included in the training set. Note how the VAT area only has landmarks in the lower well defined part. The complete annotated training set is presented in Appendix B.




Figure 5.2: One of the images that make up the training set annotated with landmark points. Red points are the external SAT outline, blue points are internal SAT outline and green points describe the outline of the well defined part of the VAT area.

The training set used for this project consists of 11 images. The images are chosen to represent the typical variation of the structures on all images. The outline of the external SAT border is annotated with 32 landmark points. The outline of the internal SAT border is annotated with 40 landmark points and the VAT area with 27 landmark points.

5.2 Aligning the training set

Before statistical analysis can be performed on the landmark points the points must be in the same co-ordinate frame, and variation due to global transformations such as rotation and scale must be removed.

Let a shape be defined as the n landmark points that define one or more structures. If the n points are d -dimensional we have nd elements. In 2D we thus have n points (x_1, y_1) that we organize into a $2n$ vector:

$$\mathbf{x} = (x_1, \dots, x_n, y_1, \dots, y_n)^T$$

For a training set consisting of s images we have s such vectors. To align the landmark points of the training set a simple iterative method is used.

1. Translate each example so that the center of gravity of its landmarks is at the origin
2. Chose one example as a reference and use this as the initial estimate of the mean shape. Scale the example landmark points so that $|\bar{x}| = 1$.
3. Save this first mean shape estimate as \bar{x}_0 and let it define the default reference frame.
4. Align all shapes with the current estimate of the mean shape.
5. Re-estimate the mean shape \bar{x}_0 from the aligned shapes.
6. Align the current estimate of the mean shape with $|\bar{x}_0|$ and scale so that $|\bar{x}| = 1$.
7. If not converged return to 4.

Convergence is declared when the change to the mean shape does not change considerably from one iteration to the other.

To align two 2D shapes \mathbf{q} and \mathbf{r} each centered around the origin we find a rotation θ and a scale s , that minimizes the squared distance between \mathbf{q} and the rotated and scaled version of \mathbf{r} .

$$\min(|T_{s,\theta}(\mathbf{q}) - \mathbf{r}|^2) \tag{5.1}$$

The transformation that best aligns the shape \mathbf{q} with the shape \mathbf{r} is given by:

$$T \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} a & -b \\ b & a \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} t_x \\ t_y \end{bmatrix} \quad (5.2)$$

This is the transformation applied to each point (x, y) in the shape \mathbf{q} . a and b are defined as

$$a = \frac{\mathbf{q}^T \mathbf{r}}{|\mathbf{q}|^2} \quad (5.3)$$

$$b = \frac{\sum_{i=1}^n (x_{qi}y_{ri} - x_{ri}y_{qi})}{|\mathbf{q}|^2} \quad (5.4)$$

x_{qi} denotes the i^{th} x-coordinate from shape \mathbf{q} . t_x and t_y are the means of the x and y -coordinates of \mathbf{r} :

$$t_x = \bar{\mathbf{r}}_x \quad (5.5)$$

$$t_y = \bar{\mathbf{r}}_y \quad (5.6)$$

Figure 5.3 shows the mean shape of the training set. The shape contains the landmarks for the external SAT border and the internal SAT border. The shape is centered around the origin.




Figure 5.3: The mean shape of the training set including landmarks from the external SAT border and the internal SAT border.

5.3 Modelling shape variation

We now have s sets of points \mathbf{x}_i that are aligned into a common co-ordinate frame. These vectors of points form a distribution in nd dimensional space. We seek to model this distribution so that new examples similar to those in the training set can be generated and it can be decided if a given new shape is a plausible one. That is, we seek a parameterized model of the form $\mathbf{x} = M(\mathbf{b})$ that can be used to generate new vectors \mathbf{x} by means of the parameters \mathbf{b} .

In order to reduce the dimensionality of the problem from nd to something more manageable, Principal Component Analysis (PCA) is applied to the data. The data form a cloud of points in nd dimensional space. PCA computes the main axis of the point cloud allowing for approximation of the original points by means of these new axis. Thus reducing the dimensionality. The methodology is as follows.

1. Compute the mean of the data:

$$\bar{\mathbf{x}} = \frac{1}{s} \sum_{i=1}^s \mathbf{x}_i \quad (5.7)$$

2. Compute the dispersion matrix of the data:

$$D(\mathbf{x}) = \frac{1}{s-1} \sum_{i=1}^s (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T \quad (5.8)$$

3. Compute the eigenvectors, ϕ_i and corresponding eigenvalues λ_i of $D(\mathbf{x})$. These are then sorted in descending order $\lambda_i \geq \lambda_{i+1}$.

Let Φ contain the t eigenvectors corresponding to the t largest eigenvalues. The data in the training set, \mathbf{x} can then be approximated by

$$\mathbf{x} \approx \bar{\mathbf{x}} + \Phi \mathbf{b} \quad (5.9)$$

where $\Phi = (\phi_1 | \phi_2 | \dots | \phi_t)$ and \mathbf{b} is a t dimensional vector given by

$$\mathbf{b} = \Phi^T (\mathbf{x} - \bar{\mathbf{x}}) \quad (5.10)$$

\mathbf{b} now defines a set of parameters of a deformable model. The shape can then be varied by varying the elements of \mathbf{b} using Equation 5.9. λ_i corresponds to the variance of the i^{th} parameter, b_i , across the training set. By constraining the parameter, b_i , to be within 3 standard deviations ($\pm 3\sqrt{\lambda_i}$) it is assured that the shape generated is similar to those in the training set and thus a plausible shape.

Since the total variance in the data equals the sum of the eigenvalues, we can select the number of eigenvectors to use, t , so that a certain percentage of the total variance is explained by the model. This proportion is set to 98% for any computations in this report.

The i^{th} principal axis is defined as the direction corresponding to the i^{th} highest eigenvector. The i^{th} principal component is the projection of \mathbf{x} on the i^{th} principal axis.

Figure 5.4 shows PCA applied on a 2D example. Figure 5.4(a) shows how the points in 2D can be approximated using a single principal component axis, \mathbf{p} . Figure 5.4(b) shows how \mathbf{x} can be approximated by the nearest point \mathbf{x}' on the principal axis. The approximation is computed as $\mathbf{x}' \approx \bar{\mathbf{x}} + b\mathbf{p}$, where b is the distance along the principal axis from the mean to the point on the principal axis closest to \mathbf{x} .




Figure 5.4: The principle of the principal component analysis. See text above for explanation.

5.4 The two shapes

The landmarks annotated on the training set images are used to build two shape models consisting of two different shapes.

1. **Shape 1** are the points from the external SAT outline and the internal SAT outline. The red and blue points on Figure 5.2.
2. **Shape 2** are the points from the external SAT outline and the VAT area outline. The red and green points on Figure 5.2

If the VAT area points were included in **shape 1**, to form one shape including all points, the variation in the location of individual points would be relatively low in comparison to the total variation of the shape. This would make the points from the internal SAT outline govern the placement of the VAT area points when applied to unknown images, thus making for a less accurate placement of the VAT area points. The external SAT outline points are included in both shapes to make the search for new points more robust. The outer SAT outline is generally quite trivial to find and it thus provides a degree of constraint on the placement of the VAT area points and the internal SAT points. By making two shapes in this way robustness is weighted versus accurate placement of the internal SAT points and the VAT area points.

Figure 5.5 shows the effect of varying the model parameters on the model made from **shape 2**. This corresponds to changing the b_i parameters away from 0 one by one in Equation 5.9. The first few modes are often possible to give a description of what kind of variation they govern. One might say that the first mode governs the width of the area containing the spine. The second mode seems to govern the overall shape of the external SAT points and also the vertical relative placement of the VAT area points. Already at the 3rd mode the variation starts to have a large degree of what seems like noise. It however still seems to govern what could be interpreted as the height of the area containing the spine.

Figure 5.6 shows a plot of the value of the first principal component versus the value of the second principal component for **shape 2** on each of the 11 images in the training set. Training image number 5, Figure 5.6(a), has a low first principal component value and a very high second principal component value. This falls well in line with the definitions of the first two modes that was presented above, since image 5 has a wide spine area and the VAT area points are relatively far away from the bottom of the external SAT points. Training image 11, Figure 5.6(b), has a narrow spine area and the VAT area points are




Figure 5.5: The effect of varying the parameter, b_i , $\pm 3\sqrt{\lambda_i}$, for the first 3 modes of the principal component analysis. The first row shows the first mode, the second row shows the second mode and the third row shows the third mode. The shape is **shape 2**, made up of the external SAT landmarks and the VAT outline landmarks.

close to the bottom of the external SAT points. This falls well in line with the description since training image 11 has a high first principal component value and a low second principal component value.




Figure 5.6: (a) and (b) are training image number 5 and 11. (c) The first principal component as a function of second principal component for each of the 11 images in the training set. The shape used is made up of the external SAT landmarks and the VAT outline landmarks.

5.5 Finding shapes in unknown images

Given a rough starting approximation new model points can be found in an unknown image. This is done by searching the area around each approximated starting point for the pattern of intensity variation that best match the variation of the model intensities around the same point.

In practice this is done by searching along profiles normal to the model boundary at each model point. Figure 5.7 gives a stylistic presentation of the sampling method.




Figure 5.7: The principle of the boundary sampling. A profile is sampled at each model point normal to the model boundary. The blue quarter circle is the wanted structure in the image

If the boundary of the model corresponds to an edge (fast intensity shift) new points can simply be placed at the strongest edge. This is however not always the case. The correct placement could be on a secondary edge or other image structure. The correct profile structure to look for can be learned from the training set. By constructing a statistical model of how the intensities varies across profiles on the training set images, the profile locations that best match the model can be identified. The following describes the approach used.

Suppose for a given point that samples along a profile k pixels to either side of the point are taken. This gives $2k + 1$ samples. This is done for each image in the training set and the samples for the i^{th} image are stored in the vector \mathbf{g}_i . To reduce the effect of global intensity changes in the images the derivative along the profile is computed and stored instead of the intensity value. The samples are normalized by dividing through by the sum of absolute element values. Let g_{ij} be the sample value at profile location j for image i .

$$\mathbf{g}_i \longrightarrow \frac{1}{\sum_j |g_{ij}|} \mathbf{g}_i \quad (5.11)$$

This is done for each image. The assumption that these normalized samples have a multivariate gaussian distribution is made, and their mean $\bar{\mathbf{g}}$ and covariance \mathbf{S}_g are estimated. A statistical model of the intensity profile along each point in the shape is constructed in this way. The result is a model for the intensity profile around each point. The quality of fit for a new sample, \mathbf{g}_s , to the model is given by

$$f(\mathbf{g}_s) = (\mathbf{g}_s - \bar{\mathbf{g}})^T \mathbf{S}_g^{-1} (\mathbf{g}_s - \bar{\mathbf{g}}) \quad (5.12)$$

This is the Mahalanobis distance of the sample from the model mean. The Mahalanobis distance is linearly related to the log of the probability that \mathbf{g}_s is drawn from the distribution. The minimization of $f(\mathbf{g}_s)$ is thus equivalent to the maximization of the probability that \mathbf{g}_s comes from the distribution.

To find the best location for new points along the profiles m pixels on either side of the current profile location is sampled ($m < k$). The quality of fit is then tested (Equation 5.12) at each at the possible $2(k - m) + 1$ locations along the sample profile. The location that gives the best match, corresponding to the lowest value of $f(\mathbf{g}_s)$, is chosen. This is done for all model points thus giving a suggested location for the placement of new points in the unknown image.

The new points are aligned to the mean shape of the model as described in Section 5.2 and are projected into model space using Equation 5.10. The variation is truncated to only allow plausible shapes by limiting \mathbf{b} to be within plus minus 3 standard deviations. Finally the points are converted back into normal coordinate space using Equation 5.9. The Process is illustrated on Figure 5.8.




Figure 5.8: The effect of projecting into model space. The left illustration is the shape found using the profile search method. These points are projected into model space on the middle illustration. The last illustration shows a nearby plausible shape after the truncation of \mathbf{b} is made.

The method described above will need to be performed iteratively to get a good result. The complete search algorithm for new points can be formulated as

follows:

1. Initialize the iterations with a start guess of where the shape is believed to be located.
2. Search the profiles about each start point to find the best nearby match for new shape points in the image.
3. Align the new shape points with the mean shape.
4. Project the aligned points into model space.
5. Truncate the variation to only get plausible shapes.
6. Convert the points back to normal coordinate space.
7. Repeat from 2 with the new shape points as the start guess until a satisfying result is obtained.

5.6 Obtaining a good initialization

In order to be able to find a given structure in an unknown image a good start guess must be made. This start guess must be close enough to the wanted structure in the image that the profiles used to search for new points will overlap the wanted structure. It is usually enough that the majority of the points will have a profile that overlap the image structures. The projection into model space will often ensure that the start points in the following iteration are closer to the wanted structure in the image. To get a good start guess for the point search in the MR image data the following approach is used.

The row and column limits of the external SAT outline is found by searching through each row and column of the unknown image for a shift in intensity. This intensity shift indicates a move away from low intensity background voxels. Let C_j denote the column sum of intensities I at column j , and R_i the row sum of intensities at row i . The limit for the intensity shifts C_{limit} and R_{limit} can be formulated as

$$C_{limit} = p \cdot \max(\mathbf{C})$$

$$R_{limit} = p \cdot \max(\mathbf{R})$$

where

$$C_j = \sum_i I_{ij}$$

$$R_i = \sum_j I_{ij}$$

p is the percentage of the maximum sum of intensities that is considered to indicate a major intensity shift. I_{ij} is the intensity value at column j and row i . The limits of the SAT outline is found at the first and last column where $C_j > C_{limit}$ and the first and last row where $R_j > R_{limit}$. For the MR image data $p = 0.3$ has been found to give good results.

Remember that the mean shape of the model is normalized and centered around the origin. We seek a transformation that will scale and translate the mean shape to match the limits on the unknown image found above. The scale parameters for rows S_r and S_c are found by comparing the range of the unknown image limits and the limits of the mean shape points. Let $C_{e,u}$ be the end column limit for the unknown image and $R_{s,m}$ be the start row limit for the mean shape, then

$$S_c = \frac{(C_{e,u} - C_{s,u})}{C_{e,m} - C_{s,m}}$$

$$S_r = \frac{R_{e,u} - R_{s,u}}{R_{e,m} - R_{s,m}}$$

The correct size of the mean shape is found by multiplying the row coordinates of the mean shape by S_c and multiplying the column coordinates by S_r . The row translation T_r is computed by finding the center of the unknown image using the row limits. This gives

$$T_r = \frac{R_{e,u} - R_{s,u}}{2}$$

Since the center of the mean shape is based on a center of gravity calculation the column center will not be halfway between the column limits. This is due to the shapes only being symmetric for row coordinates. The column center is

found by aligning the start limit of the mean shape with the start limit of the unknown image.

$$T_c = R_{s,u} - S_c R_{s,m}$$

By applying the row and column translations T_r and T_c to the scaled mean shape the transformation is done. Figure 5.9 shows an example of the initialization.




Figure 5.9: The automatic initialization of start guess points. The image is slice 4 from patient 4953

5.7 Applying ASM to the MRI data

Active shape models are used to drive a rough localization of the SAT and VAT area borders. Only one model is made for each patient. Thus only one slice will be used in the finding of structures. The idea is that the slices will be segmented one by one and the results from one slice will be used to set up constraints for the segmentation of the next. This approach exploits the fact that the shape of the structures only change slowly from one neighboring slice to another. To reduce the number of slices that any given segmentation will be away from an ASM driven segmentation, the ASM is performed on the middle slice of each patient. Recall that not all 15 slices are used on all patients to span the anatomically bounded unit, thus the middle slice number will change from one patient to another.

The models of **Shape 1** and **Shape 2** are build with a profile sample width of $k = 20$. The sample match width is set to $m = 2$ making for a $2m + 1 = 5$ wide search window across the $2k + 1 = 41$ wide profile sample. Figure 5.10 shows an example of the ASM search iterations on slice 4 from patient 4955. Note that since patient 4950 4951 and 4952 were used to train the model they will no longer be used to show results. The start guess has been initialized badly on purpose to give a more exciting example. Even though the initialization is quite bad the majority of the points found with the profile search are still found close to where they should be on Figure 5.10(a). There are a few points that are clearly placed wrongly, namely the blue points in the center of the VAT area. Figure 5.10(b) shows how these points are projected into a plausible model shape (the green points). The green points from Figure 5.10(b) become the start guess for the next iteration and are thus the same as the red points on Figure 5.10(c). On Figure 5.10(f) it is seen that after 20 iterations the green points provide a good fit to the structures in the image that was sought. This example gives a good sense of the degree of robustness the method has to a bad start guess. As long as there is some amount of overlap between the start guess and the wanted structures in the image, the algorithm will give a good fit to the structures after a suitable number of iterations.

To determine when the ASM search algorithm has converged the following approach is used. For each profile search that is performed the distance of the best matching location from the center of the profile is recorded. For instance if the best matching profile search point is found directly on top of the start point the recorded distance is 0. The iterations are stopped when p_{close} percent of the points are found within a distance of d_{max} from the profile center. If this criteria is never met the iterations are stopped after N_{max} iterations have been applied. For the MR image data the iterations are stopped if 90% of the points are found within 2 pixels of the profile center, or if more than 20 iterations have already been run.




Figure 5.10: An example of an ASM search. The red points are the start guess of the iteration and the blue points are the points found searching the profiles. The green points are the end results after the projection into model space and back. The figure shows results from iteration 1, 2 and 20. Note that the blue points on the left and right column are the same.

5.8 Results

Figure 5.11 shows the final result of the ASM search on the models build from **Shape 1** and **Shape 2** on 3 patients. The external SAT outline has generally been found quite well across all 3 images. The position of the points are generally close to the edge of the external SAT outline. Figure 5.11(a) and (e) shows how the internal SAT outline also has been found quite accurately with points that are close to the edge of the structure in the image. On patient 4954, Figure 5.11(c), the result is less accurate as the blue points do not follow the wanted image edge on the middle right part of the image. This is mainly due to one yellow point being matched to the an edge on the VAT area outline and thus affecting the shape when the projection into model space is made. By looking at the ASM search on the VAT area outline (the second column) it is seen that it generally suffers from the same lack of accuracy that we saw on Figure 5.11(c). It can be seen on Figure 5.11(f) that even though the yellow points mostly are located on the edge of the wanted image structure the shape that is obtained after the projection into model space does not fit well will the image structure. This is due to the large variation in the shape of this image structure across all images. Even though images with highly varying VAT outline structures are included in the training set the VAT outline structure of the unknown image is simply too different from the structures in the training images to form a good fit.

With all the inaccuracies in mind the results are still well suited for the purpose they are to serve. The shape points are not meant to be used as an exact outline of the image structures. They are created to identify the general location of the image structures and provide a guidance for a more accurate outlining of the structures in the next chapter. While the ASM results for the VAT area outline are not very accurate, they clearly identify the overall form and location of the VAT area structure in the image.

The main strength of using ASM for this task is its degree of robustness. The method is very robust to a bad start guess, mainly because the external SAT outline is included in both models. The identification of the image structures is never (did no occur in any of the over 80 patients used for testing) inaccurate enough to be unsuitable for the further steps in the segmentation. The robustness comes mainly from the fact that even though a number of points are found at wrong locations during the profile search, they only affect the overall shape output to a small degree.




Figure 5.11: The results of the ASM search on 3 patients. The first column is the model build from **Shape 1**. The second column is the model build from **Shape 2**. Magenta dots are the profile search points from the external SAT outline and yellowe dots are search points for the internal SAT and VAT area. The final results of the 3 outlines are red, blue and green for the external SAT, internal SAT and VAT respectively. The data are slice 4, 5 and 4 from patients 4953, 4954 and 4955 from top to bottom.

Outlining image structures

In the previous chapter a rough outline of the SAT and VAT areas were found. This chapter will describe the process used to get an accurate and robust outline of the SAT and VAT areas using these rough outlines.

We seek a method that can follow the path described by these rough point locations, but staying on the edge of the image structure accurately. The method used will be that of Dynamic Programming (DP). An introduction to the method will be given first, followed by a description of the transformations that are needed to apply the method to the MR image data. Finally the segmentation of the SAT and VAT outlines will be described in detail.

6.1 Dynamic Programming - DP

Dynamic Programming is a technique used to compute optimal paths in graphs. A graph is a set of points called *nodes* connected by a set of *links* that connect the nodes. A path through a graph is the set of points that connects the start node with the end node. The graphs used for DP are characterized by the links being one directional (can only be traversed in one direction) and that the graph contains no loops. That is you can only enter a given node once.

Let there exist a matrix of values with i rows and j columns. Each location in the matrix corresponds to a node, and links are constructed using the rule that a node corresponding to row i and column j , denoted $N(i, j)$, can only connect to a node in column $j + 1$ at rows $i - 1$, i and $i + 1$. Figure 6.1 shows the principle.




Figure 6.1: (a) A matrix with values. (b) The DP graph corresponding to the matrix.

Let further the value at each matrix location represent the static cost of the corresponding node in the graph. The optimal path is then defined as the path from any graph node at the first column in the matrix to any graph node at the last column that has the lowest total cost. The total cost is the sum of all node costs in the optimal path. Before the optimal path can be found the cumulative cost matrix is computed. The cumulative cost matrix is computed as follows.

Start at the second column, the cumulative cost of node $N(i, j)$ is computed for each column as

$$N_c(i, j) = \min(N_c(i - 1, j - 1), N_c(i, j - 1), N_c(i + 1, j - 1)) + N_s(i, j) \quad (6.1)$$

where $N_s(i, j)$ is the static cost of node $N(i, j)$. That is, any cumulative cost at a given node is the lowest cumulative cost of the 3 nodes that connects to it plus the static cost of the node itself. Figure 6.2(a) shows the cumulative cost of the matrix from Figure 6.1(a). For each node a pointer to the predecessor node with the lowest cost is saved. The optimal path is found by tracing back through these pointers starting at the last row. Figure 6.2(b) shows the optimal path of the example matrix.

To apply dynamic programming to find an optimal path in an image the values

in the test matrix are simply replaced by the pixel values of the image. In the programs used for this project dynamic programming has been used to find paths with maximum cost instead of minimum cost. This has been more intuitive to work with and all following results will be based on maximum cost optimal paths.

The dynamic programming as presented here only finds optimal paths through columns, thus in effect only a 1D search. The paths we seek are however the outlines of the SAT and VAT areas and these are closed contours. To make the data suitable for dynamic programming a spatial transformation is applied.




Figure 6.2: (a) the accumulated cost matrix computed from the matrix on Figure 6.1(a). (b) The optimal path through the matrix found by tracing back through the cumulative cost matrix.

6.2 Transforming the MR image data

To transform the data to a form well suited for dynamic programming a polar transformation is applied. The transformation is done by sampling along a number of spokes evenly distributed on a circle. During the transformation the samples along a spoke becomes a column in the transformed image matrix, one column for each spoke. The principle is illustrated on Figure 6.3. The transformation is defined by 3 parameters. The center of the circle \mathbf{x}_c , the radius of the circle d_c and the angular resolution of the spokes N_s . Note the coordinate system used, shown on Figure 6.3(b). The first column in the transformed matrix is thus from the spoke at 0, while the center column is from the spoke at π . Samples are taken from the outside and in, meaning that the first row in the transformed matrix corresponds to samples on the circle and the last row corresponds to samples at the center of the original image. Since the sampling location will not generally be in the center of a pixel, bilinear interpolation is

used when extracting the intensities.




Figure 6.3: (a) The principle of the sampling along spokes in a circle. 20 spokes are shown (in red) and the circle (in blue) has a radius of $d_c = 95$ pixels. (b) The polar coordinate system. (c) The result of the transformation of the image using $N_s = 600$ spokes. (d) The result of applying a difference filter to each column.

The form of the image is now well suited for dynamic programming since the wanted outlines can be traced as going from the first column to the last. The values in the matrix are however not suitable. By applying a difference filter across each column in the matrix a suitable image emerges. The difference filter simply computes the difference between two neighboring rows in a column: $X(2) - X(1)$ $X(3) - X(2) \dots X(d_c) - X(d_c - 1)$ where $X(i)$ is the value at row i of a given column. The result can be seen on Figure 6.3(d). In this filtered image, edges that represent a move from high to low intensity will have a low

value, and edges representing a move from low to high intensities will have a high value. The vertical location of each column that best matches the SAT and VAT borders can now be easily identified as ridges or valleys running from left to right in the filtered image.

Because a connected circular contour in the image is sought it is important that the optimal path found starts and ends on the same matrix row. To achieve this, an amount of repetition is added at each end of the matrix. An example can be seen on Figure 6.4. A repetition of 100 columns at each end of the matrix ensures that the optimal path found will pass through the same row at each end of the original matrix. All dynamic programming done for this project incorporates a repetition of 100 columns at each end of the matrix. All results shown will be of paths where the repetitions have been used but removed again to show only the resulting path on the unique part of the matrix.




Figure 6.4: The transformed image with 100 columns of repetition added at each end of the matrix. The blue lines mark the extend of the original matrix.

The use of dynamic programming has usually been coupled with the design of domain-specific cost functions. This approach attempts to guide the dynamic programming optimal path using a complex cost function that is the sum of many individual terms. These terms include gradient terms, terms used to control contour stiffness, terms used to attract the contour to certain intensity levels as well as a number of imaginative domain specific terms. For instance Aboutanos [9] uses a six term cost function. The problem with these many terms are that they need to be carefully balanced against each other to give good results, and might fail if for instance the global intensity level of the images changes. The approach that will be used here will be to only use the gradient as the cost function. The guidance of the dynamic programming will instead be done by applying constraints to the path. This will happen through the use of the shape points found in the previous chapter but also by other means as will be described below. This will make the solution much more robust to global intensity changes and images that generally are extraordinarily different from the norm. In the following the approach used to find the outline of each of the 3 wanted image structure borders will be described individually. The external SAT border the internal SAT border and the VAT area border.

6.3 The external SAT border

The external SAT border is found by dynamic programming. The shape points from **shape 1**, found using ASM, are used to define the transformation of the original image. Note that only the first 32 points from **shape 1** are used since they mark the outline of the desired region. The center of the transformation circle is set halfway between the extremities of the points that make up the shape. Let C_s, C_e, R_s and R_e be the start and end extremities of columns and rows respectively for the shape points. The center of the transformation is then found as

$$\mathbf{x}_c = \begin{bmatrix} \frac{C_e + C_s}{2} \\ \frac{R_e + R_s}{2} \end{bmatrix} \quad (6.2)$$

The radius of the transformation is set at half the norm of the range of the shape points:

$$d_c = 0.5 \cdot \text{norm} \begin{pmatrix} C_e - C_s \\ R_e - R_s \end{pmatrix} \quad (6.3)$$

The angular resolution of the spokes is set at $N_s = 720$. The steps involved in the constrained dynamic programming search can be seen on Figure 6.5. The shape points are transformed with the image. The transformed points are used as soft constraints, meaning that the path is not forced through each shape point but rather highly *encouraged* to do so. This is done by adding a high value to the cost of matrix indices that are close to or coinciding with the shape points. The value added could for instance be 10^6 , where the typical cost of a matrix value is in the range $[-1 : 1]$. A range of row indices are used to form a constraint range r_c pixels on either side of the shape point instead of a single constraint point. The magnitude of the added cost should be several orders of magnitudes higher than the typical cost of a matrix index. This will in effect ensure that the optimal path always will pass through all constraints if not limited by the connectivity in the matrix. Thus if two neighboring constraints are too far apart that the path can not pass through them both, the optimal path will pass through the constraint that gives the overall highest cost. An example of the constraint ranges is shown on Figure 6.5(e). For all shape constraints $r_c = 7$ has been used.

The external SAT border will generally be the most significant ridge in the difference filtered image. The constraints are thus most often not needed to




Figure 6.5: (a) Slice 5 from Patient 4954. (b) The result of the dynamic programming on the image. (c) The transformed image. (d) The transformed image after the difference filter has been applied. (e) The constraints introduced by the ASM. The red dots are the shape points, and the blue dots indicate the range that the path is constrained to. (f) The optimal path found through the transformed image using the constraints.

ensure that the correct ridge is found. This was used to identify the region of interest during the bias correction stage where no constraints were available.

The reason why this same outline is found again here using the constraints is simply to add another level of robustness to the segmentation.

6.4 The internal SAT border

The internal SAT border is found in much the same way as the external SAT border. The transformation is done using the 40 points from **Shape 1** that outline the internal SAT border, while assigning \mathbf{x}_c and d_c in the same way as for the external SAT border. The angular resolution of the spokes is set at $N_s = 720$. Since the outline of the internal SAT area is a valley on the difference filtered transformed images a simple sign change is done on the filtered images before maximum cost dynamic programming is used.

To ensure that the path follows the wanted outline tightly a new set of constraints are introduced. The constraints utilize the classification that was already obtained from Chapter 4 (FCM). The idea is that the optimal path of the internal SAT outline always will be at the immediate end of an extensive area with adipose tissue. The FCM classification will be used to find these locations. The constraint points are found in the following way.

1. Transform the classification map along with the image. The results can be seen on Figure 6.6.
2. For each column in the transformed classification matrix find the first row that corresponds to an adipose tissue classification. From this row location search downwards in the column until the last uninterrupted adipose pixel is found. Save this row number as r_{la} .
3. If the next two underlying row indices (in the same column) following r_{la} are not classified as adipose tissue a hit is found and a constraint point is added at r_{la} and $r_{la} + 1$. If one of the next two row indices in the column are classified as fat no constraint is added.

The area close to the bellybutton has very bad image quality and thus gives a possibly incorrect FCM classification. This results in classifications where the SAT area and the VAT area are connected with adipose tissue classifications, see Figure 4.7(e) for instance. Since the method for obtaining constraints described above will give undesired results when used on these kinds of images the search for constraints are not performed in this area. This corresponds to excluding spokes from $3\pi/4$ to $5\pi/4$ (see Figure 6.3(b)). The constraints added in this way

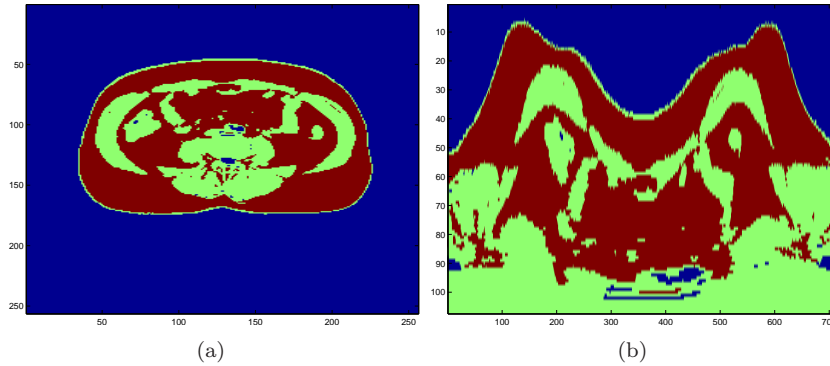


Figure 6.6: (a) The result of the FCM classification. Red is adipose tissue, green is other tissue and blue is void. (b) The transformed classification image. The data is slice 5 from 4954.

can be seen on Figure 6.7. The new constraints are used as soft constraints together with the constraints from the shape points. Meaning that a high cost (10^3 for instance) is added at the constraint locations. The cost added at the two constraint types are balanced so that the shape point constraints have a much higher value added than the new constraints. This in effect gives a priority system where the path will follow the shape point constraints primarily, but given that the new constraints are within the shape constraints the path will follow the new constraints. The optimal path found on slice 5 from patient 4954 can be seen on Figure 6.7(b).

6.5 The VAT border

The dynamic programming is done on two different image transformations. One with the center in the middle of the VAT area (the VAT_{dp}) and another smaller one with its center in the spine area (the SP_{dp}). The points used as constraints are the last 27 points from **Shape 2**. Figure 6.8 shows the constraint points used for the VAT_{dp} and the SP_{dp} . Note that all points from the shape are not used, since some of these will cause an overlap in constraints when the transformation is applied.

The transformation parameters \mathbf{x}_c and d_c for the SP_{dp} are assigned in the same way as for the SAT borders. For the VAT_{dp} R_s (the start row index) is taken from the internal SAT shape points. The computations of the parameters are

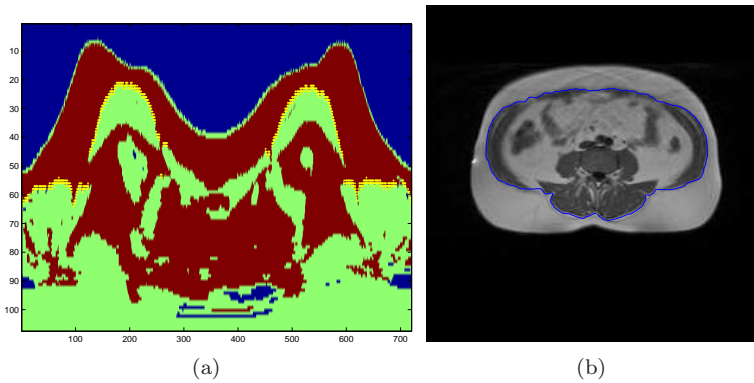


Figure 6.7: (a) The new constraints that are added. The constraint points are yellow dots. (b) The optimal path found using both the shape point constraints and the new constraints. The data is slice 5 from patient 4954

otherwise the same. The angular resolution of the spokes is set at $N_s = 600$ for the VAT_{dp} and at $N_s = 360$ for the SP_{dp} .




Figure 6.8: (a) The 27 shape points that make out the VAT area outline. (b) The points used for the two rounds of dynamic programming. Green points are the constraint points for the VAT_{dp} and the blue dot (located just over the top most yellow point) is the center of the spokes transformation. Yellow points are the constraint points for the SP_{dp} and the red dot is the center of the spokes transformation.

Since only shape constraint points are available for the lower part of the VAT area, a way of guiding the VAT_{dp} path at the upper part is needed. By observing the MR image data it is observed that the border of the VAT area always is near the internal SAT border near the bellybutton. This is utilized to apply a constraint to the optimal path. Within a given spokes range a valid path location is defined as being a maximum of d_{max} pixels away from the internal SAT border found above. A mask is created where valid path locations are given the value 1 and invalid path locations are given the value 0. The valid path locations mask is multiplied with the difference filtered image to give all invalid locations the value 0. The method is shown on Figure 6.9. $d_{max} = 8$ is used and the spokes range are those corresponding to the angular range of $3\pi/4$ to $5\pi/4$.




Figure 6.9: (a) The transformed image. (b) The transformed image with the difference filter applied. (c) The mask with valid path locations shown as white. (d) The difference filtered image multiplied with valid path locations mask.

To control the general direction of the optimal path in the lower area between the shape constraint points one extra constraint point is added. This point is the center of the SP_{dp} transformation (the red point on Figure 6.8(b)). The bottom yellow point on the same figure is the extra point added to the SP_{dp}

constraint points. This point is placed halfway between the center of the SP_{dp} transformation and the highest valued row index of the internal SAT path. These last two constraint points are placed to make the path of the SP_{dp} and the VAT_{dp} overlap. This will prove beneficial in the last part of the segmentation process described in the next chapter. The results of the VAT_{dp} and the SP_{dp} are shown on Figure 6.10.




Figure 6.10: The result of the SP_{dp} and the VAT_{dp} . yellow is the VAT outline and green is the spine area outline. The image is slice 5 from patient 4954.

6.6 Constraints on other slices

The methods described in the last 3 sections have all concerned finding border outlines in a specific slice of the patient, where the shape points were available. In the rest of the slices of the patient there are no shape points available though. Instead the fact that neighboring slices are acquired only 10 millimeters apart is used. The borders outlines of the image structures thus change little from one neighboring slice to the next. This is utilized for constraining the paths in all other slices of a patient. The approach is as follows:

1. Find the outlines of structures in the middle slice (S_m) of the anatomically bounded volume using the shape point constraints as described above.
2. Use all the points that make up the outlines of S_m to form constraints for finding paths in slice $S_m - 1$ and $S_m + 1$.
3. Keep propagating outwards from the middle slice while using constraint points acquired from the neighboring slice until the start and end slice is reached. That is, for a slice number S_i where $S_i > S_m$, constraints are acquired from slice $S_i - 1$. For a slice number S_j where $S_j < S_m$, constraints are gotten from slice $S_j + 1$.

All the transformation parameters are computed from the constraint points in the same way as for the middle slice. Instead of having just 32 constraint points for the external SAT outline 720 are now available corresponding to the spokes resolution of the external SAT outline in the neighboring slice. The extra constraints that are applied for the middle slice are applied in exactly the same way for the other slices and with the same settings. Figure 6.11 shows the constraints and the resulting optimal paths on slice 6 from patient 4954.

It should be noted that another method for adding constraints on other slices was planned at first, but had to be discarded due to time limitations at a late stage. This method was based on building multiple shape models at 5 different depth levels of the anatomically bounded unit. The constraint points for each slice would then be calculated using the ASM model at the depth level that had the closest distance to the depth level of the slice. The annotation of landmarks on 11 images at 4 new depth levels is a big task and since the method used in the project gives good results the time was better spent focusing on other parts of the project. This does not mean that the method of using constraints from the neighboring slices gives bad results. It would however have been the more elegant solution to use the 5 levels of shape models and provide for another level of robustness.

Patient: 4954, Slice: 6




Figure 6.11: The outlines found on slice 6 from patient 4954 using the outlines from slice 5 as constraints. Red is the external SAT outline, blue is the internal SAT outline, yellow is the VAT outline, green is the spine area outline and magenta are the constraint points.

6.7 Results

The main advantage of using constraints to guide the dynamic programming is that it adds some high level structural information about the shape of wanted structures to an inherently low level method. Without the constraints the dynamic programming becomes easily breakable by strong edges in unexpected locations. This is especially true given the overall structure of the MR image data. It is not uncommon to see the desired outline have blurred edges while a random void hole in the visceral area have a very strong edge. This will lead to the optimal path passing along the strong edge of the random void hole instead of the blurred outline of the desired edge. The shape points found using ASM further makes for an easy and robust way of automatically setting the parameters for the polar transformation of the images.

Figure 6.12 shows the result of the dynamic programming on patient 4953, 4955 and 4958. Note that these patients are not chosen because they give either particularly good or bad results. They are the next numbered patients in the test set following the patients that were used to train the shape models. The external SAT outline (red) and the spine area outline (green) gives results that closely match what would be expected by visual inspection. The outlines of the internal SAT border is quite good overall except the area near the bellybutton. The image quality in this area is very poor and suffers from noise, blurriness and artifacts introduced by the image acquisition process. The poor quality results in a general lack of clearly identifiable edges that the path can follow, or in some cases strong edges introduced by artifacts. The result is that the outline of the internal SAT border will be somewhat random in this area on some images, but still constrained by the shape points. This effect translates over to the VAT outline as well since it is constrained by the internal SAT outline in this area.

The results of the dynamic programming are generally good across the patients in the test set and the method is robust. The main inaccuracies are tied to the path of the internal SAT border near the bellybutton area.




Figure 6.12: The outlines found on the middle slice of 3 patients. The first column is the bias corrected images and the second column is the same images with the outlines superimposed. Red is the external SAT outline, blue is the internal SAT outline, green is the VAT outline and yellow is the spine area outline.

Set operations and connectivity

The final segmentation of each slice into the 3 desired classes, SAT, VAT and Other Tissue (OT from now on) will be described in this chapter. The approach used combines the results from the FCM (Chapter 4) and the dynamic programming (Chapter 6). First the set operations that are used will be described and afterwards a description of how connectivity is utilized to improve the results will be given.

7.1 Set operations

The 4 outlines found in the previous chapter, the external SAT outline, the internal SAT outline, the VAT border outline and the spine area outline, are used to construct 4 binary masks. Each mask is constructed by defining a polygon that goes through each point that makes up an outline. All pixels inside or on this polygon are given the value 1 and all pixels outside are given the value 0. The result can be seen on Figure 7.1. The 4 masks will be denoted M_{ROI} , M_{inSAT} , M_{VAT} and M_{SP} for the external SAT outline, the internal SAT outline, the VAT outline and the spine area outline respectively. The SAT area and the VAT area masks can be extracted from these 4 masks using simple

boolean operators.




Figure 7.1: The construction of the 4 sets. (a) The outlines computed by dynamic programming. (b) The region of interest (ROI) mask, M_{ROI} . (c) The internal SAT mask, M_{inSAT} . (d) The VAT mask, M_{VAT} . (e) The spine area mask, M_{SP}

SAT

The SAT area is computed by using the exclusive or operator, **XOR**:

$$M_S = \mathbf{XOR}(M_{ROI}, M_{inSAT})$$

VAT

The VAT area is computed in three steps. First, the common area of M_{VAT} and M_{SP} is computed:

$$M_{com} = \mathbf{XOR}(M_{VAT}, M_{SP})$$

Second, the area that is exclusive to M_{SP} is computed:

$$M_{ex} = \mathbf{AND}(M_{com}, M_{SP})$$

Finally the VAT area is computed:

$$M_V = \mathbf{XOR}(M_{com}, M_{ex})$$

The resulting two masks are shown on Figure 7.2




Figure 7.2: (a) The bias corrected image. (b) The two masks created by the set operations. The green area is the M_S mask and the orange area is the M_V mask.

To get the initial segmentation of the slice the FCM classification (see Figure 7.3(a)) of the slice is combined with the two masks on Figure 7.2(b). Since the outlines for the external and internal SAT were found to be both robust and accurate all voxels that coincide with the M_S mask are classified as being SAT. Thus no consideration of FCM classifications is done in this area. The

FCM classifications are however used indirectly since they were used to form the constraints on the path of the internal SAT border. The FCM classification will sometimes classify a small number of single voxels within the SAT area as being a non adipose tissue class. This classification is generally wrong and due to noise in the original image or shortcomings of the bias correction. By assigning all voxels that coincides with the M_S mask as SAT these erroneous voxels are correctly classified.

To determining the classification of adipose tissue in the VAT area, the boolean **AND** operator is used with the adipose tissue class of the FCM classification as the first argument and the M_V mask as the second argument. The result is that all voxels classified as adipose tissue by the FCM classification that coincide with the M_V mask are classified as VAT. In contrast to what was done with the SAT classification the FCM classification is heavily utilized here. This is necessary since the M_V mask only defines the outline of the VAT area. The VAT area itself might contain only small amounts of adipose tissue. The result of this initial segmentation is shown on Figure 7.3(b) for slice 5 of patient 4954.

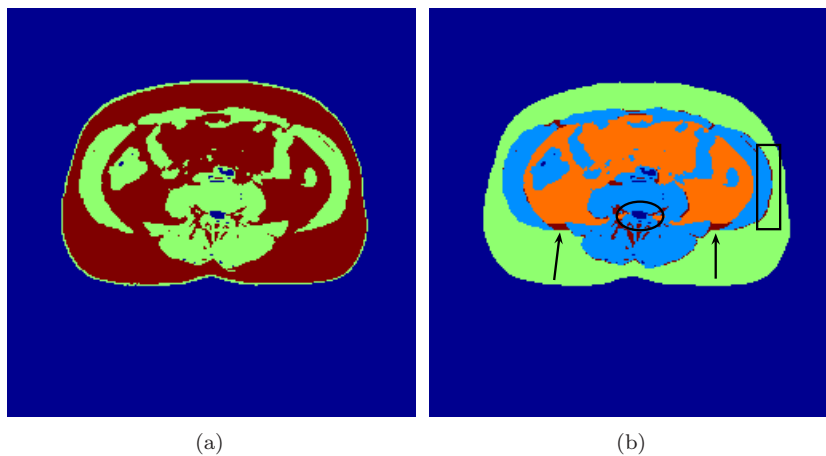


Figure 7.3: (a) The FCM classification of slice 5 from patient 4954. Blue is void, red is adipose tissue and green is other tissue. (b) The initial segmentation. Dark blue is void, light blue is other tissue classified by FCM, green is SAT, orange is VAT and red is adipose tissue classified by FCM.

The result looks promising, but there are a few areas that could be improved. The first troublesome area is indicated by the two arrows on Figure 7.3(b). These side lobes of adipose tissue should be classified as VAT. This problem originates from the dynamic programming where the optimal path is unable to follow the outline here because no edge is present. The second troublesome area

is indicated by the ellipse. Here two small areas have been wrongly classified as VAT. The problem is due to the M_{VAT} and M_{SP} masks not having enough overlap. This results in the M_V mask being too large and thus some non VAT voxels are included in the segmentation. The third troublesome area is indicated by a box on the figure. This problem is less severe and consists of a number of voxels near the internal SAT border that are classified as adipose tissue by the FCM classifier. These should have been included in the SAT class. This problem exist all along the internal SAT border but is most severe in the boxed area.

7.2 Connectivity

By utilizing the structure of the FCM classified adipose tissue voxels a way to rectify the three troublesome areas is sought. The structure property that is utilized is that of connectivity.

Let two voxels be neighbors if they are located in the same image plane and have either the same row number and neighboring column numbers or the same column number and neighboring row numbers. Let two voxels of the same class be connected if it is possible to reach voxel one from voxel two by only traversing voxels of the same class that are all neighbors to an already traversed voxel. This corresponds to normal 4-neighbor connectivity. A connectivity labelling of a boolean class image is done by assigning each body of connected voxels a unique class. A connected body is the collection of voxels that are all connected.

The boolean image is constructed from the FCM classification of the image using a high ε_{sm} value. Recall that ε_{sm} was the parameter used to specify the amount of similarity needed to classify a voxel as a certain class by the FCM. By assigning ε_{sm} a value higher than 0.5 only voxels that are very similar to the class centroid will be classified as that class. To form the boolean image needed for the connectivity labelling FCM is performed on the image with $\varepsilon_{sm} = 0.7$ for the adipose tissue class. The boolean image can be seen on Figure 7.4(a) The reason the normal FCM classification is not used is that since the connectivity labelling will be used to add voxels to the SAT and VAT classes, only voxels that have a strong connectivity with these classes will be added. Furthermore, all voxels added in this way are more likely to be true adipose tissue voxels. A complete connectivity labelling (CL_c) is constructed from this boolean image, See Figure 7.4(b). Another reduced connectivity labelling (CL_r) is constructed as well where the boolean labels that do not coincide with the M_{inVAT} mask are removed. The reduced connectivity labelling can be seen on Figure 7.4(c). Note how each body of connected voxels have a unique labelling corresponding

to a unique collaring on the images. The 3 troublesome areas are corrected in turn using this connectivity labelling.




Figure 7.4: (a) The labelling obtained by assigning all voxels with a similarity measure higher than $\varepsilon_{sm} = 0.7$ to the adipose tissue class. (b) The connectivity labelling of (a) denoted CL_c . (c) the connectivity labelling of (a) after all voxels not coinciding with the M_{inVAT} mask have been removed, denoted CL_r . The data is slice 5 from patient 4954.

Adding the side lobes

First a mask that defines the area where the operation is valid is created, M_{lobes} . The mask is needed since undesired bodies of connected voxels might otherwise be added. The mask is constructed by excluding all voxels with a row number lower than the minimum row number of the path of the spine area outline that was computed previously (Section 6.5). Further voxels are excluded that have column numbers between the column extremities of the spine area outline. The mask is shown on Figure 7.5(a). The confines of the mask are chosen to excluded all the small bodies of voxels below the spine area and the area near the bellybutton that often has bad image quality. The operation that is performed to add the side lobes is as follows:

1. For each body of voxels in CL_r test if any voxels from this body are included in the initial SAT classification.
2. If the test is positive add the voxels to the SAT class if they coincide with the M_{lobes} mask defined above.

In effect this identifies all bodies of voxels that have been cut away from the SAT area by the M_V mask and adds them to the SAT class if they match the mask for the operation. The result of this operation can be seen on Figure 7.5(b).

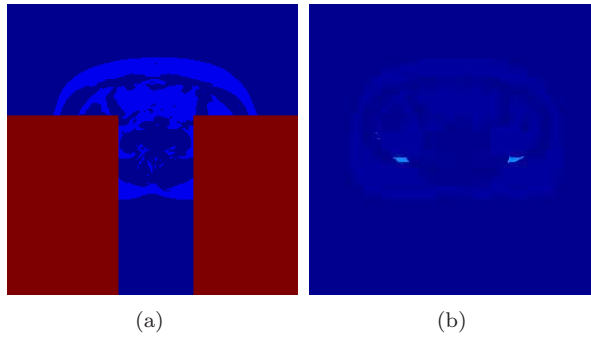


Figure 7.5: (a) The mask of the add side lobes operation. The mask is shown in red. (b) The voxels added by the add side lobes operation. The voxels added are shown in light blue

Removing erroneous VAT voxels

A similar approach is used to remove the unwanted SAT classifications under the spine area. A remove mask is defined (M_{remove}) by a polygon spanned by four points. The upper two points are at the row location of the center of the polar transformation for the spine area (See Section 6.5). The column locations are the minimum and maximum column location of the spine area outline respectively. The lower two points are the lower right and lower left corner of the image respectively, see Figure 7.6(a). The confines of the mask are chosen to only include the area below the spine area where small bodies of voxels that are FCM classified as adipose tissue often are present. The operation that is performed is very simple.

All bodies of small connected voxels in CL_r that have already been classified as VAT are removed if they coincide with the M_{remove} mask. A small body is defined as a body consisting of less than 30 voxels. This works because the adipose tissue surrounding the area under the spine is very homogeneous and usually consists of bodies several hundred voxels in size. Thus any small body in this area must be an unwanted VAT classification and can safely be removed. The result can be seen on Figure 7.6(b).

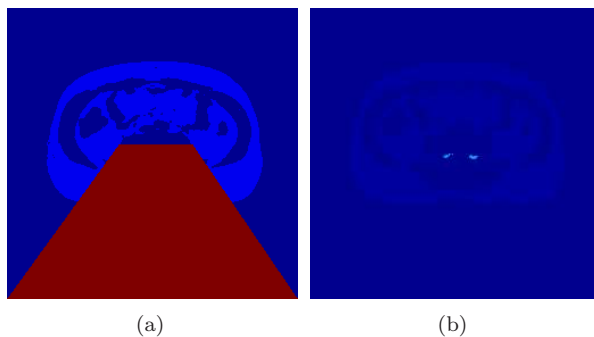


Figure 7.6: (a) The mask of the remove erroneous VAT operation. The mask is shown in red. (b) The voxels removed by the operation. The voxels removed are shown in light blue

Adding small bodies of SAT voxels

The operation used to add the small bodies of adipose tissue near the internal SAT border to the SAT class is very simple. It operates on CL_r , but the bodies are recalculated taking into account the changes that were made when the side lobes were added. The operation adds small bodies from CL_r to the SAT class if these bodies coincide with a body in CL_c where at least one of the voxels has a SAT classification. In other words, it adds small bodies adipose tissue voxels that got separated from the subcutaneous area by the M_S mask to the SAT class. A small area is defined as being a body of voxels consisting of less than 10 voxels. This will ensure that all the small bodies of voxels adjacent to the internal SAT border are added to the SAT class. At the same time no big bodies will be added. This is especially critical near the belly button area where large bodies of voxels that are neither classified as SAT or VAT often are present.




Figure 7.7: (a) The voxels added by the add small bodies operation. The voxels added are shown in light blue. (b) The final segmentation after the operations on all 3 troublesome areas have been performed. Dark blue is void, light blue is muscle classified by the FCM, red is adipose tissue classified by the FCM, green is SAT and orange is VAT.

After the 3 operations are carried out the final segmentation is the result. Figure 7.7(b) shows the end result on slice 5 on patient 4954. Note that the 3 final segmentation classes: SAT, VAT and OT (Other Tissue) are extracted from the final result in the following way. The SAT and VAT classes corresponds to the SAT and VAT classes on Figure 7.7(b). The OT class is the combination of all other classes inside the region of interest defined by the outline of the SAT area. That is, the OT class is the combination of the dark blue, light blue and red areas on Figure 7.7(b) inside the region of interest.

7.3 Results

The removal and adding of voxels to the SAT and VAT class does generally give good results. The side lobes that can not be classified to VAT right away by the set operations are added using the connectivity instead. This operation has the potential to include large possibly unwanted bodies of voxels into the VAT class. There are a number of factors that greatly lessen the likeliness of this happening. First, the operation is only performed on a reduced area due to the introduction of the M_{lobes} mask. Secondly, a FCM classification with a high required similarity measure is used. This combined with the fact that only 4 neighborhood connectivity is used (and not 8) means that only voxel bodies that are strongly connected with the main VAT area will be added. The side lobes operation gives good results on the current data. Further constraints on when the operation should be allowed might be needed if the quality of images in future data should worsen significantly. The operations for the removal of erroneous voxels from the VAT class and the adding of small bodies of voxels to the SAT class are safe to use and give good results.

The final segmentation result on the middle slice of patient 4953, 4955, 4958 can be seen on Figure 7.8. The final segmentation result will be evaluated in the following chapter.




Figure 7.8: The results of the final segmentation on the middle slice of 3 patients. The first column is the segmented image and the second column is the unaltered original MR image.

Final results

Since no ground truth is available it is not possible to give a simple measure for the overall accuracy of the results on all slices of all patients. Instead the final segmentation result of the middle slice from all 80 patients in the test set is included in Appendix D. Only the middle slice from each patient is presented since presenting all slices of 80 patients would take up too much space. The middle slice is the most important slice since it is the basis of the segmentation of all other patient slices. Each segmentation result is presented next to the biased original MR image. This allows for the robustness and accuracy of the segmentation to be assessed as a whole, including the bias correction step.

By looking at the results as a whole, the results seem to be quite close to the segmentation one might do if the task was to be performed manually. The writer of this report is not trained in the manual segmentation of MR images, but showing the final results to Kristian Wraae resulted in an enthusiastic response.

The segmentation has one main troublesome area. The area around the belly-button has proven to be very hard to segment properly. This is due to the very bad image quality that characterize this area. The area will often be one large cluster of voxels classified as adipose tissue after the FCM has been done. This area often extends well into both the subcutaneous and visceral area. This makes it hard to separate these two regions even by visual inspection. Many different methods have been tried to segment this area but none were found that worked

better than letting the dynamic programming find the optimal path through the area. Another problem is the adding of the side lobes with the connectivity operation. This operation sometimes gives questionable results. As an example the side lobes that are added on slice 5 from patient 4958 (See Figure 7.8(e)) can be discussed if they should be included as SAT tissue VAT tissue or neither.

The main strength of the segmentation lies in the robustness of the segmentation. Even though the accuracy varies from image to image the overall correlation between the result of the automatic segmentation and the segmentation one could imagine being done by hand seems high across all images in the test set. This is partly attributed to a well working bias correction resulting in a good FCM classification. Partly to the introduction of ASM acquired constraints to the dynamic programming.

Appendix C contains the combined results for the distribution of tissue of all slices in the anatomically bounded unit for each patient. The tables contain information about the total volume of the anatomically bounded unit and the distribution of tissue on the 3 classes: SAT, VAT and OT. These correspond to the percentage of subcutaneous adipose tissue, visceral adipose tissue and other tissue.

Conclusion

This report has described an approach for implementing a pipeline for the automatic segmentation of adipose tissue in MRI data. The Segmentation was done into the 3 classes: subcutaneous adipose tissue, visceral adipose tissue and other tissue. The segmentation was done using a range of different image analysis techniques.

A novel way of sampling high intensity voxels was developed. This method samples high intensity voxels evenly distributed across the patients anatomy. The method worked well in providing a good sampling of the bias field across the entire volume.

The estimation of the bias field was done over the entire volume using *Thin Plate Splines* extended to 3D. The method allows for a smoothing interpolating surface. The degree of smoothing was controlled using the notion of *effective degrees of freedom*. The bias correction was overall very successful in making the intensities of adipose tissue across the image uniform.

The initial classification of each voxel was done using *Fuzzy c-mean clustering*. This method proved to be very accurate in determining the threshold between intensities of adipose tissue and intensities of other tissue when used on the unbiased data. The method is completely automatic and finds the optimal intensity threshold that best segments the data without any parameter tuning.

The borders of the subcutaneous area and the visceral area were identified using *Dynamic Programming* constrained by points acquired from an *Active Shape Model* (ASM). The shape points acquired from the ASM were very robust in outlining the general location of image structures but the points did not follow the borders accurately. The Dynamic Programming is generally not very robust but is quite accurate when it works. By combining the two methods a both accurate and robust result was obtained where the weakness of one method is covered by the strength of the other.

A last step was performed to improve a number of minor misclassifications inherited from the Dynamic Programming result. The concept of *connectivity* was used to add and remove voxels near the border of the class regions. The technique introduced some potential inaccuracies but did overall improve the result.

The complete pipeline is overall robust to the great variations of anatomy of the patients. This is to a large degree attributed to the use of methods that require a minimum of data specific parameters to be set.

Since no ground truth is available an exact measure for the accuracy of the final results can not be given. The overall accuracy has however been assessed by visual inspection of the results on over 80 randomly picked patients presented in Appendix D. The results were found to be good and a positive response was received from Kristian Wraae (The PhD student that will be using the results to assess adipose tissue in his project).

Program overview

The program that performs the segmentation has been implemented using MATLAB. This appendix will give an overview of the program structure. The most important files and their functionality will be listed. These files are available on the CD that accompanies the report. An indent indicates that the indented file name is a subfunction of a file above.

mainPreprocess - sorts and orders all the DICOM files into a separate data structure for each patient. The script calls **correctbias3d** when all data structures have been build.

correctbias3d - performs all the preprocessing for each patient by calling *findbiasfat* and *tpsmooth*.

findbiasfat - finds the samples for each slice that are to be used to estimate the bias field.

tpsmooth - calculates an estimate of the bias field by using the samples from **findbiasfat**.

annotate - an interactive script that allows for the annotation of landmarks on the training set. The building of the shape and intensity models can be initiated from this script.

buildgradientmodel - builds the model of the gradients along the profiles.

buildshapemodel - builds the model of the shape variation.

batchrun - performs the segmentation on all slices for each patient by calling a number of subfunctions.

fuzzycmean - performs the fuzzy c-mean clustering on a slice.

fitshape - uses the shape model to find a number of points that outline the desired image structures.

dpfromshape - performs the transformation of the image and finds the optimal path using dynamic programming.

setoperations - performs the set operations and the connectivity operations.

APPENDIX B

The training set

Below are the 11 images that make up the training set. They are used to train the model used by the Active Shape Model in [Chapter 5](#).




Figure B.1: images annotated with landmarks that make up the training set.




Figure B.2: images annotated with landmarks that make up the training set.

APPENDIX C

Volume results

Below are the results of the distribution of adipose tissue types in each patient volume. The values are obtained by counting the total number of voxels of each class for each slice in the anatomically bounded volume in each patient. These numbers are then multiplied by the dimensions of each voxel to obtain the volume results. The number of voxels in each class are computed from images with the same labelling as the image in Appendix D. The distributions for SAT and VAT are computed by counting the number of green and orange labels in the segmentation results. The distribution of OT (Other Tissue) is computed by counting all dark blue, light blue and red labels within the confines of the patient boundary.

Patient	Total Volume (<i>Liter</i>)	SAT	VAT	OT
4950	6.43	0.359	0.248	0.394
4951	5.85	0.301	0.222	0.476
4952	5.84	0.358	0.224	0.418
4953	4.60	0.308	0.202	0.490
4954	7.30	0.401	0.244	0.355
4955	4.76	0.266	0.201	0.533
4958	5.49	0.310	0.187	0.503
4996	11.57	0.207	0.322	0.470
4998	7.13	0.275	0.240	0.484
4999	5.38	0.237	0.172	0.591
5000	8.04	0.260	0.377	0.363
5003	4.98	0.214	0.385	0.401
5004	7.06	0.307	0.288	0.404
5045	4.61	0.220	0.251	0.529
5046	5.92	0.229	0.324	0.447
5049	5.31	0.403	0.199	0.397
5050	4.77	0.268	0.234	0.497
5051	5.31	0.217	0.360	0.423
5052	4.24	0.404	0.141	0.455
5094	7.26	0.333	0.280	0.388
5096	5.18	0.252	0.270	0.478
5097	4.43	0.209	0.223	0.567
5099	5.78	0.230	0.294	0.477
5100	6.55	0.236	0.291	0.473
5101	4.02	0.296	0.149	0.555
5145	5.20	0.187	0.349	0.464
5147	4.34	0.272	0.235	0.494
5149	5.77	0.301	0.316	0.382
5150	4.19	0.297	0.178	0.525
5197	7.00	0.333	0.219	0.448

Table C.1: The volume results of the segmentation for each patient. The first column is the patient number. The second column is the total volume of the anatomically bounded unit. The last three columns are the distribution of the volume on the three classes SAT, VAT and OT (Other Tissue).

Patient	Total Volume (<i>Liters</i>)	SAT	VAT	OT
5198	5.45	0.265	0.286	0.448
5247	7.07	0.411	0.232	0.357
5249	6.06	0.292	0.318	0.390
5251	4.83	0.334	0.234	0.432
5252	6.83	0.304	0.350	0.346
5254	5.11	0.255	0.212	0.533
5319	5.96	0.323	0.247	0.430
5321	4.85	0.252	0.276	0.472
5362	4.54	0.373	0.203	0.425
5365	3.65	0.250	0.257	0.493
5366	7.16	0.232	0.363	0.405
5367	6.62	0.289	0.295	0.417
5368	6.20	0.244	0.369	0.387
5371	6.80	0.383	0.235	0.381
5412	4.79	0.252	0.231	0.517
5415	5.70	0.264	0.343	0.393
5480	4.20	0.281	0.123	0.596
5482	5.39	0.232	0.314	0.453
5483	4.93	0.312	0.191	0.497
5484	4.55	0.287	0.228	0.485
5485	5.97	0.209	0.350	0.441
5486	4.87	0.380	0.190	0.430
5531	7.44	0.321	0.321	0.358
5533	5.75	0.327	0.295	0.378
5560	4.24	0.219	0.259	0.522
5565	7.79	0.354	0.248	0.399
5594	3.46	0.187	0.261	0.551
5598	5.97	0.299	0.232	0.469
5628	3.31	0.308	0.153	0.538
5631	2.94	0.179	0.253	0.568

Table C.2: The volume results of the segmentation for each patient. The first column is the patient number. The second column is the total volume of the anatomically bounded unit. The last three columns are the distribution of the volume on the three classes SAT, VAT and OT (Other Tissue).

Patient	Total Volume (<i>Liter</i>)	SAT	VAT	OT
5633	5.54	0.289	0.297	0.414
5718	4.20	0.295	0.186	0.519
5720	5.93	0.253	0.326	0.422
5721	4.93	0.313	0.229	0.458
5723	4.46	0.266	0.134	0.600
5725	6.55	0.219	0.386	0.394
5761	4.33	0.293	0.137	0.570
5763	4.77	0.266	0.221	0.513
5764	4.19	0.234	0.225	0.541
5765	5.58	0.322	0.218	0.460
5766	4.19	0.301	0.201	0.498
5768	4.69	0.273	0.176	0.550
5807	7.54	0.315	0.276	0.409
5846	6.40	0.257	0.303	0.440
5847	3.93	0.331	0.230	0.439
5849	4.69	0.348	0.176	0.476
5850	6.21	0.328	0.251	0.420
5885	4.62	0.270	0.241	0.490
5922	4.75	0.288	0.254	0.457
5924	5.47	0.334	0.266	0.400

Table C.3: The volume results of the segmentation for each patient. The first column is the patient number. The second column is the total volume of the anatomically bounded unit. The last three columns are the distribution of the volume on the three classes SAT, VAT and OT (Other Tissue).

APPENDIX D

Segmentation results

This appendix contains images of the results of the segmentation method described in this report. The slices shown are the middle slice of all the 80 patients in the test set. Each row of images consist of the segmentation result and the original unaltered biased MR image from one patient. The color coding of the segmentation is: green - SAT, orange - VAT, dark blue - void, light blue - other tissue classified by the FCM method, red - adipose tissue classified by the FCM method.




Figure D.1: The final segmentation.




Figure D.2: The final segmentation.




Figure D.3: The final segmentation.




Figure D.4: The final segmentation.




Figure D.5: The final segmentation.




Figure D.6: The final segmentation.




Figure D.7: The final segmentation.




Figure D.8: The final segmentation.




Figure D.9: The final segmentation.




Figure D.10: The final segmentation.




Figure D.11: The final segmentation.




Figure D.12: The final segmentation.




Figure D.13: The final segmentation.




Figure D.14: The final segmentation.




Figure D.15: The final segmentation.




Figure D.16: The final segmentation.




Figure D.17: The final segmentation.




Figure D.18: The final segmentation.




Figure D.19: The final segmentation.




Figure D.20: The final segmentation.




Figure D.21: The final segmentation.




Figure D.22: The final segmentation.




Figure D.23: The final segmentation.




Figure D.24: The final segmentation.




Figure D.25: The final segmentation.




Figure D.26: The final segmentation.




Figure D.27: The final segmentation.

Bibliography

- [1] John G. Sled and Alex P. Zijdenbos. A nonparametric method for automatic correction of intensity nonuniformity in mri data. *IEEE Transactions On Medical Imaging*, 17(1):87–97, 1998.
- [2] R. Engholm, A. Dubinskiy, R. Larsen, L. G. Hanson, and B. Ø. Christoffersen. An adipose segmentation and quantification scheme for the abdominal region in minipigs. In *International Symposium on Medical Imaging 2006, San Diego, CA, USA*. The International Society for Optical Engineering (SPIE), feb 2006.
- [3] J. Duchon. Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces. *IRAIRO Analyze Numérique*, (10):5–12, 1976.
- [4] P. J. Green and B. W. Silverman. *Nonparametric regression and generalized linear models, a roughness penalty approach*. Chapman & Hall, 1994.
- [5] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The element of statistical learning*. Springer, 2001.
- [6] Milan Sonka and J. Michael Fitzpatrick. *Handbook of Medical Imaging: Volume 2. Medical Image Processing and Analysis*. SPIE PRESS, 2000.
- [7] Vincenzo Positano, Amalia Gastaldelli, Anna maria Sironi, Maria Filomena Santarelli, Massimo Lombardi, and Luigi Landini. An accurate and robust method for unsupervised assesment of abdominal fat by mri. *Journal Of Magnetic Resonance Imaging*, (20):684–689, 2004.

- [8] T. F. Cootes and C. J. Taylor. Statistical models of appearance for medical image analysis and computer vision. *appears in Proc. SPIE Medical Imaging*, 2001.
- [9] G. B. Aboutanos, J. Niskanen, N. Watkins, and M. Dawant. Model creation and deformation for the automatic segmentation of the brain in mr images. *IEEE Transactions on Biomedical Engineering*, 46(1):1346–1356, 1999.